

REPORT FOR THE  
ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

**ARTIFICIAL INTELLIGENCE  
AND REGULATORY ENFORCEMENT**

Michael Karanicolas

*This report was prepared for the consideration of the Administrative Conference of the United States. It does not necessarily reflect the views of the Conference (including its Council, committees, or members).*

---

**Recommended Citation**

Michael Karanicolas, Artificial Intelligence and Regulatory Enforcement (Dec. 9, 2024) (report to the Admin. Conf. of the U.S.).

# Artificial Intelligence and Regulatory Enforcement

Michael Karanicolas

I. Executive Summary .....	2
II. Introduction.....	5
III. Background: AI, Public Perceptions, and Administrative Agencies .....	9
IV. The Existing Governance Landscape.....	16
V. Key Values Underpinning an Appropriate Framework for AI in Regulatory Enforcement.	26
A. Understanding Risk and Risk Assessments .....	26
B. Public Engagement .....	28
C. Recognizing Failure .....	32
D. Structural Oversight Considerations .....	34
VI. Conclusion: Grappling with the Human-Machine Paradigm.....	39
VII. Recommendations.....	41

## I. Executive Summary

“Artificial intelligence” (AI) comprises a set of technical concepts which may generally be understood through their capacity to automate human capability. In recent years, an increasing number of government agencies have incorporated AI systems into their regulatory enforcement processes. Some examples of these applications include at the Securities and Exchange Commission, which employs machine-learning tools to identify suspicious filings or insider trading, and at the Environmental Protection Agency, which is experimenting with automated facility inspection targeting. The potential efficiency gains from these systems are easy to understand, especially in the context of mounting public demands for the administrative state to do more with less. They may also present opportunities to expand the processing power of government agencies, boosting research and analytical capabilities in novel and unprecedented ways. Such improvements may be necessary in the context of increasing sophistication among targets of regulation.

However, for all of their potential benefits, the incorporation of AI systems into regulatory enforcement functions also raises a number of difficult questions. In particular, there are concerns related to the impact of AI on the fairness of decision-making. In addition to well-documented challenges with bias, drift, and accuracy, there are fundamental questions around transparency, accountability, and procedural fairness. At a time when public trust in government, and in our broader public institutions, has declined to critical levels, the potential for these systems to further erode the relationship between the people and the administrative state should be an extremely serious consideration.

Innovation in the provision of public services is something we should welcome. However, responsible and accountable government necessitates an open and public conversation about the use of AI to perform critical government tasks, particularly related to regulatory enforcement, where due process considerations are paramount. Effective tracking and coordination is a necessary first step to provide a complete picture of the use cases for these technologies. A necessary second step is a robust process for auditing and evaluation of AI systems, where they have been piloted, to honestly assess their successes and failures. Auditing and evaluation is important to guide the next generation of implementation. Ultimately, the goal of administrative agencies, and of our different levels of government, should be to craft an effective, transparent, and relatively consistent decision-making framework around the use of AI across the public sector.

In developing a robust governance framework for the use of AI in regulatory enforcement, there are a number of useful international models to draw from. These include Canada’s *Directive on Automated Decision-Making*, which has been in place since 2019, and which requires a robust risk-based assessment process with a sliding scale of obligations based on the results. The European Union’s *AI Act* is another well-known risk-based model, and a number of states, including California, Maryland, Illinois, and Washington, have proposed their own frameworks to address and mitigate governance and accountability challenges related to these technologies.

There have been a number of executive orders targeting responsible AI use at the federal government, going back to 2016, though the most significant moves to develop concrete and actionable protocols for how the administrative state should approach AI in a regulatory

enforcement context have been the Office of Management and Budget’s AI policy memo, and the National Institute of Standards and Technology’s “Artificial Intelligence Risk Management Framework”. Together, these present a framework for significantly improving transparency around where these systems are being deployed, and to establish an effective framework for identifying and tracking risks, respectively. However, significant governance challenges remain.

In particular, a broad focus on transparency threatens to come at the expense of robust accountability. While the latter depends on the former, it also requires opportunities for the public to meaningfully impact government policy. Likewise, the insertion of a “human in the loop” is unsatisfactory since human review of automated decisions is not always meaningful. The absence of clear performance standards, or a proper operational definition for ubiquitous buzz-words such as “fair”, “safe”, and “explainable”, is another challenge.

Existing governance models’ dependence on risk assessment is another potential weakness. While risk regulation is well-adapted to mitigate certain structural harms, it is less effective at mitigating individualized harms. Experience also suggests that risk regulation is better at addressing predictable and easily quantifiable harms as opposed to the sort of “unknown unknowns” that are prevalent in considerations of the impact of AI. Ultimately, while risk-based processes are an important tool in promoting appropriate guardrails around the use of AI in regulatory enforcement, their efficacy will be heavily dependent on the spirit with which implementing agencies apply these assessments. Above all, agencies should understand the limits of the risk assessment process, and should ensure adequate space to consider worst-case scenarios, even if they are judged to be unlikely. Risk assessments should consider structural impacts on values such as agency trust and public relations. A robust risk mitigation framework should also include appropriate strategies to manage failing systems, to ensure that agencies are willing to retire systems which are not delivering adequate results.

Public engagement is another essential component to a robust governance strategy, including adequate transparency to obtain a complete picture of how AI systems are being implemented and how they are performing. At a minimum, this requires publishing risk assessments and related documentation. At a more advanced level, resources may be required to translate reporting into more accessible formats, or to upskill community partners to engage with more advanced questions. While there may be some sensitivity to disclosing too much material related to regulatory enforcement, this tension is not unique to AI. Agencies should err on the side of transparency, and only withhold information which would clearly harm the effective administration of justice, and where the harm from this disclosure would outweigh the public interest in its release.

Robust consultation may also include establishing participatory forums and the establishment of co-decision-making bodies with leading civil society and other relevant participants. The success of these initiatives will depend on agencies’ ability to offer meaningful opportunities to impact policy. However, robust engagement will be essential to maintaining public trust, generating buy-in from impacted communities, and identifying and quantifying potential or manifested harms.

A broader consideration is how to ensure robust structural oversight over the deployment of AI across the administrative apparatus. Although there are a number of potential models, the Federal

Risk and Authorization Management Program (FedRAMP) would appear to be the most comprehensive, though it remains to be seen whether the structure is capable of delivering robust oversight over the diverse range of regulatory enforcement functions where these technologies are being deployed.

At the end of the day, administrative agencies face a complex set of challenges in navigating a rapidly changing and dynamic political and technological environment. These technologies are likely to fundamentally transform their operations over the coming decade, and the decisions around their development and deployment today will determine whether administrative agencies are able to harness the fruits of AI while maintaining their human and democratic character.

## II. Introduction

Across the federal government, artificial intelligence (AI) has already taken root in a variety of administrative processes and procedures. While the public has been captivated by the possibilities, and dangers, of AI, agencies have been incorporating the technologies into a broad range of institutional functions. In 2020, a report commissioned by the Administrative Conference of the United States (ACUS) identified 157 use cases of AI across 64 federal agencies.<sup>1</sup> Just two years later, a survey by the Government Accountability Office (GAO) identified over 1,200 current and planned use cases, with NASA and the Department of Commerce leading the way.<sup>2</sup> Studies like the ones commissioned by ACUS or by the GAO are likely only presenting a partial picture of where AI is being used across the federal government. What is happening today is likely also just the tip of the iceberg when it comes to AI's longer-term impact on how government operates.

While creative and progressive approaches to the provision of public services are welcome, the expanding use of AI in regulatory enforcement carries natural tradeoffs, including against values like public trust, due process, and the expertise—and the essential human character—that underlies administrative agencies' place in America's constitutional system.<sup>3</sup> There is a need for careful institutional analysis of the pros and cons of incorporating AI into investigations and enforcement activities. These conversations should be public, and should go beyond the narrow, risk-based analysis that dominates current models of assessment and encourage longer term thinking about the agency's character, what it may be giving up through its growing reliance on AI, and the overall impacts on key values like public trust, legitimacy, and fairness. Agency considerations should also involve a careful and critical assessment of where these tools are likely to be effective, as opposed to assuming that technologically-driven solutions will inevitably improve operations.

This paper attempts to provide an initial framework for assessing the role of AI in regulatory enforcement, as well as a set of recommendations for agencies considering introducing or expanding the use of AI for these purposes.

- **Understanding Terms**

Part of the challenge in discussing appropriate impacts and safeguards for “artificial intelligence” lies in pervasive confusion and inconsistencies in how this term is understood and applied. The term is often found alongside references to “machine learning,” and indeed, in many popular contexts, the terms are used interchangeably. But while “machine learning” systems can be defined as algorithms which have the capacity to improve themselves based on training data, “artificial intelligence” belies any such technical definition, since the term is generally used as shorthand for any machine-based system which performs tasks that are traditionally reliant on human

---

<sup>1</sup> DAVID FREEMAN ENGSTROM, DANIEL E. HO, CATHERINE M. SHARKEY, & MARIANO-FLORENTINO CUÉLLAR, GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES (Feb. 2020) (report to the Admin. Conf. of the U.S).

<sup>2</sup> U.S. GOV'T ACCOUNTABILITY OFF., ARTIFICIAL INTELLIGENCE: AGENCIES HAVE BEGUN IMPLEMENTATION BUT NEED TO COMPLETE KEY REQUIREMENTS (Dec. 2023), <https://www.gao.gov/assets/d24105980.pdf>.

<sup>3</sup> See Ryan Calo & Danielle Keats Citron, *The Automated Administrative State: A Crisis of Legitimacy*, 70 EMORY L. J. 798, 802–04 (2021).

intelligence.<sup>4</sup> While these two qualities—the ability to learn from data inputs and the ability to perform tasks generally associated with human intelligence—recur in most legal definitions of AI, there are also substantial differences in how AI is defined across different frameworks.

For example, the 2019 National Defense Authorization Act describes AI as a system that “performs tasks under varying and unpredictable circumstances without significant human oversight.”<sup>5</sup> By contrast, the National Artificial Intelligence Initiative Act of 2020 defines AI as “a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments.”<sup>6</sup> These definitions show considerable overlap, of course, but also differ from one another. Significant variance is further noted throughout the multitude of state-level, agency-level, and international definitions of AI.

When considering these questions from an institutional perspective, consistency is more important than precision. Governance questions are best suited to focus on impacts, which suggest that a relatively open-ended definition is preferable.

The National Institute of Standards Technology (NIST), in its widely cited Artificial Intelligence Risk Management Framework (AI RMF), refers to an AI system as an engineered or machine-based system with varying levels of autonomy that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments.<sup>7</sup> NIST’s definition is largely drawn from the Organization for Economic Cooperation and Development (OECD) Recommendations on AI, which were published in 2019.<sup>8</sup> Because the AI RMF appears to be gaining steam as a decision-making model for agencies considering new AI use cases,<sup>9</sup> this paper will adopt the NIST definition for the sake of consistency, with the general understanding that the term should be understood inclusively.

- **About the Project and Methodology**

This paper was developed through a contract with the Administrative Conference of the United States (ACUS), which called for a study, with accompanying recommendations, to “examine the potential benefits and risks of using algorithmic tools to support agencies’ regulatory enforcement efforts and identify policies, practices, and organizational structures agencies can put in place to ensure they enforce the law fairly, accurately, and efficiently.” Although the project description calls for research into the use of “algorithmic tools”, this paper frames the discussion on AI for two reasons.

---

<sup>4</sup> Michael Karanicolas & Mallory Knodel, *Artificial Intelligence and Bias – An Evaluation*, in AM. ASSOC. FOR THE ADVANCEMENT OF SCI., ARTIFICIAL INTELLIGENCE AND THE COURTS: MATERIALS FOR JUDGES, at 10–11 (2022), [<https://perma.cc/NKE7-2MSG>].

<sup>5</sup> John S. McCain National Defense Authorization Act for Fiscal Year 2019, P.L. 115-232, 132 Stat. 1695 (2018).

<sup>6</sup> H.R. 6216, 116th Cong. § 3 (2020).

<sup>7</sup> NIST, ARTIFICIAL INTELLIGENCE RISK MANAGEMENT FRAMEWORK 1 (2023).

<sup>8</sup> OECD, RECOMMENDATION OF THE COUNCIL ON ARTIFICIAL INTELLIGENCE (2019).

<sup>9</sup> For example, a 2024 AI Policy Memo from the Office of Management and Budget encouraged agencies to incorporate NIST standards, including the Risk Management Framework, into their procurement contracts. See <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> [<https://perma.cc/AD9D-HGTE>].

First, “algorithmic tools” are an extremely broad category of programs, which can include everything from a simple handheld calculator to the operating system used to type this paper. Most of these functions are not subjects of concern, and have been employed across the federal government for decades, attracting little controversy. What is novel is the technology’s sophistication and, more importantly, its function, which has begun to complement, and in some cases even supplant, the judgment of humans within the public service. The issue, in other words, is automation, and the new functions that these tools are beginning to play relative to the regulatory enforcement process, rather than anything inherent in the technology itself.

Second, while, as noted in the previous section, AI is also an imperfect and imprecise term, it has become the lingua franca across the U.S. government, and around the world, to address the economic and social challenges that are coming as a result of the rise of sophisticated machine learning tools which are likely to automate a broad range of work-related functions over the coming decade. This is appropriate when one considers that the core challenge related to these products is their capacity to displace human workers, especially in a decision-making function. AI, as a concept, is fundamentally about displacement of humans by machines, which cuts to the core of the social, legal, and administrative challenges that are the focus of this research project. Because the NIST definition of AI, when interpreted inclusively, already encompasses virtually all of the algorithmic tools which are matters of controversy due to their potential to supplant or replace human workers, this paper defers to using the term AI for the sake of simplicity and clarity.

The focus of this paper is also on regulatory enforcement, which is defined in the project as including detecting, investigating, and prosecuting current and potential noncompliance with the laws that agencies administer.<sup>10</sup> Some examples of these activities include operations at the Security and Exchange Commission, whose mandate includes enforcing federal securities laws,<sup>11</sup> the Internal Revenue Service, which investigates financial crimes such as tax fraud,<sup>12</sup> and the Environmental Protection Agency, which enforces a range of rules related to pollution and waste products.<sup>13</sup> However, while regulatory enforcement is often a particularly high impact use case for AI, it is relatively unusual for accountability structures targeted at AI, in the US or elsewhere, to consider these applications as categorically different from other uses of AI within the decision-making process. As a result, while this paper is specifically focused on regulatory enforcement, most of the recommendations and conclusions are equally applicable to other high impact uses of AI, and are framed more broadly where this is the case.

The paper was developed through a yearlong research process that included consultation with a wide range of experts in AI and administrative law, as well as discussions of initial findings with audiences at Yale, UCLA, and at the University of California Center Sacramento. Although the research process included engagement with public service employees at the state and federal level,

---

<sup>10</sup> ADMIN. CONF. OF THE U.S., REQUEST FOR PROPOSAL— USING ALGORITHMIC TOOLS IN REGULATORY ENFORCEMENT (Aug. 4, 2023), <https://www.acus.gov/sites/default/files/documents/Algorithmic-Tools-in-Enforcement-RFP.pdf> <https://perma.cc/D2Q5-B3DY>].

<sup>11</sup> *Enforcement and Litigation*, U.S. SEC. & EXCH. COMM’N, <https://www.sec.gov/enforcement-litigation> [<https://perma.cc/NS2S-67A8>].

<sup>12</sup> *Office of Fraud Enforcement at a Glance*, U.S. INTERNAL REVENUE SERV., <https://www.irs.gov/about-irs/office-of-fraud-enforcement-at-a-glance> [<https://perma.cc/8EUC-Y4PY>].

<sup>13</sup> *National Enforcement and Compliance Initiatives*, U.S. ENV’T PROT.AGENCY, <https://www.epa.gov/enforcement/national-enforcement-and-compliance-initiatives> [<https://perma.cc/XZ2R-9PH9>].



this paper does not provide an exhaustive list of use cases related to regulatory enforcement, though a few illustrative examples are included in Part II. This is because a general mapping was already carried out through an earlier ACUS report, *Government by Algorithm*, and because a comprehensive use case database is currently under development at AI.gov, including over 700 use cases as of September 2024.<sup>14</sup>

The author offers his sincere thanks to all of the experts who agreed to speak with me, including Kevin De Liban, Andrew Selbst, Jon Michaels, Daniel Ho, Elham Tabassi, Janet Haven, Sanford Williams, Barry Johnson, Reza Rashidi, Phil Lindenmuth, Melodi Dincer, Marc-Etienne Ouimette, Chinmayi Sharma, Margaret Kwoka, Benoit Deshaies, Jith Meganathan, Artur Pericles Lima Monteiro, Alberto De Franceschi, Ji Inn Wong, Jared Koh, Simon Chesterman, Achuta Kadambi, Yuan Tian, Nanyun Peng, and Tom Speaker. Thanks as well to student research assistants who contributed research or background material, including Yuyang (Kate) Hu, Nicholas Wilson, Alyssa Stelmack, and Noah Keith, and to Kazia Nowacki, Adam Cline, Jeremy Graboyes and the rest of the team at the Administrative Conference of the United States, for their helpful feedback and support in preparing this document.

All errors remain the sole responsibility of the author, and opinions expressed do not necessarily reflect those of the Administrative Conference of the United States.

---

<sup>14</sup> *The Government is Using AI to Better Serve the Public*, AI.GOV, <https://ai.gov/ai-use-cases/> [<https://perma.cc/7D8V-B42R>].

### III. Background: AI, Public Perceptions, and Administrative Agencies

Despite the interest surrounding OpenAI’s release of ChatGPT in 2022, excitement around the promise of these new technologies, such as generative AI, has been met by an equal measure of public skepticism, and even fear, about their broader impact on humanity.<sup>15</sup> At the extreme, visions of a “Terminator-style” apocalypse driven by autonomous weapons systems have propagated, alongside highly publicized concerns about the supposed “existential risk” that these technologies pose to humanity.<sup>16</sup> Many experts consider the worst-case fearmongering about AI to be more reflective of excitement around AI than reality, and there has been considerable room for error within the doomsaying. One widely shared story concerning an AI-enabled drone, which had apparently opted to attack its operator when it found the restrictions on its use of force to be too onerous, was ultimately debunked when it turned out the scenario had been a thought-experiment, rather than an actual exercise.<sup>17</sup>

However, not all reports of AI malfeasance are bogus. One commonly cited cautionary tale concerns the Correctional Offender Management Profiling for Alternative Sanctions (COMPAS), a proprietary algorithm built by a private contractor to assess the risk of recidivism for criminal defendants.<sup>18</sup> COMPAS works by developing a “risk score” based on a questionnaire which is meant to predict the likely danger from a person’s release. COMPAS, or systems like it, have been widely incorporated into sentencing or bond hearings, including in Arizona, Colorado, Delaware, Florida, Kentucky, Louisiana, Oklahoma, Virginia, Washington, and Wisconsin. However, an investigation in 2016 found that the software was returning results which were biased against Black subjects. The company’s audits failed to appreciate these discriminatory impacts due to a problematic design choice, which prioritized accuracy of risk scores rather than the assessing the impacts of errors on subjects of different races. As a result, the system developed a tendency to place Black defendants into a higher risk category, and white defendants into a lower risk category. In other words, while there was a mechanism in place to combat racial discrimination, it ultimately proved ineffective due to how the developers framed their notion of what constituted problematic discrimination.

Another high-profile failure concerned the automation of allocation decisions for Medicaid resources, which led to drastic cuts in services for housebound patients in Arkansas, Idaho, and elsewhere.<sup>19</sup> Complaints from individuals subjected to the systems’ decision making went

---

<sup>15</sup> See Ben Schreckinger, *AI vs. Public Opinion*, POLITICO (Sept. 25, 2023), <https://www.politico.com/newsletters/digital-future-daily/2023/09/25/ai-vs-public-opinion-00118002> [<https://perma.cc/644U-WF6J>].

<sup>16</sup> *Pause Giant AI Experiments: An Open Letter*, FUTURE OF LIFE INST. (Mar. 22, 2023), <https://futureoflife.org/open-letter/pause-giant-ai-experiments/> [<https://perma.cc/LT7G-NXE6>].

<sup>17</sup> Zoe Kleinman, *US Air Force Denies AI Drone Attacked Operator in Test*, BBC (June 2, 2023), <https://www.bbc.com/news/technology-65789916> [<https://perma.cc/7YGK-BKUJ>].

<sup>18</sup> Julia Angwin et al., *Machine Bias*, PROPUBLICA (May 23, 2016), <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> [<https://perma.cc/XY2W-YW4Y>].

<sup>19</sup> Colin Lecher, *What Happens When an Algorithm You’re your Health Care*, VERGE (Mar. 21, 2018), <https://www.theverge.com/2018/3/21/17144260/healthcare-medicare-algorithm-arkansas-cerebral-palsy> [<https://perma.cc/3L5A-QUU8>].

unheeded until litigation by the ACLU of Idaho and Legal Aid Arkansas forced the problems into the light.

Every use case for AI is different and should be evaluated based on its specific context and risk profile. There is a world of difference between, for example, an automated tool which manages the allocation of parking spaces for a large agency and one which is setting bombing targets or deciding whether a person should have access to home care visits or bail. However, the preceding cases are a good illustration of why agencies may face an uphill public relations battle in developing and deploying AI for sensitive or high impact purposes, such as regulatory enforcement. There are particularly fraught consequences when considering the importance of trust and perceptions of legitimacy on regulatory enforcement functions. Agencies that wish to capitalize on the potential benefits of AI face a pressing challenge of how to maintain trust and legitimacy while pursuing greater automation.

- **Trust and Legitimacy**

In parallel to the broader expansion of AI across the economy, administrative agencies have begun to pilot these technologies for a range of functions related to regulatory enforcement. At the SEC, for example, a number of AI-based tools have been developed to support investigations into financial crimes. These include the Corporate Issuer Risk Assessment tool (CIRA), which has a machine-learning component that identifies potentially suspicious filings to predict misconduct based on historical datasets.<sup>20</sup> Another SEC tool, the Abnormal Trading and Link Analysis System (ATLAS), attempts to detect insider trading through examining differences in behavior between traders that lost money versus those who profited. At the EPA, a proof-of-concept developed by the Office of Compliance in partnership with the University of Chicago has been used to target facility inspections more accurately and efficiently, resulting in a 47% improvement of detecting violations of the Resource Conservation and Recovery Act.<sup>21</sup> The Animal and Plant Health Inspection Service under the US Department of Agriculture has piloted a similar program aimed at improving the efficiency of methods to detect invasive pest species at ports of entry.<sup>22</sup> Although each of these programs pursues important goals, and the use of AI within their respective enforcement processes offers significant potential in terms of boosting efficacy and efficiency, significant questions remain regarding the broader impact of the expansion of AI in regulatory enforcement on agencies' public trust.

Public trust is a fragile thing which may be built up over generations and destroyed virtually overnight. The importance of cultivating legitimacy through popular participation, regular testing of bureaucratic expertise, and normative reflection of policy choices is baked into the foundations of America's administrative apparatus, particularly through the Administrative Procedure Act.<sup>23</sup> At the core of these procedural protections is the fundamental principle that, in establishing

---

<sup>20</sup> ENGSTROM ET AL., *supra* note 1 at 23.

<sup>21</sup> *EPA Artificial Intelligence Inventory*, U.S. ENV'T PROT.AGENCY, <https://www.epa.gov/data/epa-artificial-intelligence-inventory> [<https://perma.cc/23YE-ZXBT>].

<sup>22</sup> *See 2023 Public Inventory of AI Use Cases*, U.S. DEP'T AGRICULTURE (May 15, 2023), [https://www.usda.gov/data/ai\\_inventory.csv](https://www.usda.gov/data/ai_inventory.csv) <https://perma.cc/V8SN-GGYQ>].

<sup>23</sup> Bruce Ackerman, *The New Separation of Powers*, 113 HARV. L. REV. 633, 697 (2000).

legitimacy, it may often be necessary to accept tradeoffs relative to other institutional values, particularly efficiency of operations.<sup>24</sup>

In addition to legal requirements, there are practical reasons why administrative agencies should prioritize efforts to cultivate public trust and popular perceptions of their legitimacy. Where agencies issue rules that impact the public, and particularly where public compliance is necessary for the success of their mandate, perceptions of legitimacy are of paramount importance to encourage respect for agency authority and voluntary compliance.<sup>25</sup> While the threat of sanctions may loom large over individual decisions to file a tax return, the system, as a whole, depends on voluntary compliance. The sovereign citizen movement is a good illustration of the challenges and frustrations in trying to deal with a group that rejects the legitimacy of government agencies, even if the group's perceptions are based on spurious logic. An agency's sense of legitimacy must be resilient enough to survive even individual unpopular agency decisions or outright errors, which should ideally be seen as exceptions within a fundamentally valid structure even by those who are disadvantaged by them.<sup>26</sup> Public relations are an important operational consideration for modern executive branch structures for practical, as well as political, reasons.<sup>27</sup>

None of this is to argue that agencies should be overly defensive about how they act, or their place in the constitutional order. In *The Procedure Fetish*, Nicholas Bagley argues persuasively that, contrary to concerns about some democratic deficit that is inherent to the administrative state, agencies are natural outgrowths of America's democratic and constitutional structure, which includes balanced roles for the executive, legislative, and judicial branches in maintaining robust public accountability over their operations and decision-making.<sup>28</sup> Nonetheless, in considering their operational future, it would be shortsighted for agencies to ignore the mounting attacks on their legitimacy. While the sovereign citizens are a fringe movement, public skepticism of administrative agencies extends far more broadly.<sup>29</sup> It also appears to be reaching a crescendo in the present political moment, as calls for a significant reduction in the reach of administrative agencies have become firmly entrenched in the political mainstream.<sup>30</sup> While there is no question that politicians have been instrumental in driving this narrative, they are also responding to public sentiment which shows that Americans' trust in government is at historically low levels.<sup>31</sup> It is also worth noting that public support for governments' use of AI tends to correlate with trust in government more generally, leading to the potential for vicious (or virtuous) cycles as adoption

---

<sup>24</sup> See Lisa Shultz Bressman, *Beyond Accountability: Arbitrariness and Legitimacy in the Administrative State*, 78 N.Y.U. L. REV. 461, 546 (2003).

<sup>25</sup> TOM R. TYLER, *WHY PEOPLE OBEY THE LAW* 4 (2006).

<sup>26</sup> See Adrian Vermeule, *Bureaucracy and Distrust: Landis, Jaffe, and Kagan on the Administrative State*, 130 HARV. L. REV. 2463, 2463–65 (2017).

<sup>27</sup> DANIEL CARPENTER, *REPUTATION AND POWER: ORGANIZATIONAL IMAGE AND PHARMACEUTICAL REGULATION AT THE FDA* (2010).

<sup>28</sup> Nicholas Bagley, *The Procedure Fetish*, 118 MICH. L. REV. 345 (2019).

<sup>29</sup> See, e.g., DAVID SCHOENBROD, *POWER WITHOUT RESPONSIBILITY: HOW CONGRESS ABUSES THE PEOPLE THROUGH DELEGATION* 14–18 (1993).

<sup>30</sup> See, e.g., Zachary B. Wolf, *Which Agencies Would GOP Candidates Cut? Try to Remember All of These*, CNN (Aug. 23, 2023), <https://www.cnn.com/2023/08/23/politics/republican-government-cuts-what-matters/index.html> [<https://perma.cc/V9W7-UE5Q>].

<sup>31</sup> See *Public Trust in Government: 1958-2024*, PEW RSCH. CTR. (June 24, 2024), <https://www.pewresearch.org/politics/2024/06/24/public-trust-in-government-1958-2024/> [<https://perma.cc/6WSV-3T5H>].

accelerates.<sup>32</sup> In this moment, it is critically important to scrutinize how automation in general, and the use of AI in particular, are likely to impact public, judicial, and political perceptions of administrative agencies' role in American governance and the legitimacy of their regulatory enforcement functions.

- **Expertise, Discretion, and Regulatory Enforcement**

A traditional justification for the administrative state, generally, is that the technological complexity of the modern world necessitates a level of regulatory expertise which is beyond the capacity of Congress to independently regulate.<sup>33</sup> Today, a similar argument undergirds some of the drive towards incorporating AI into regulatory enforcement, as continued technological progress means that even human experts are outmatched by the volume and complexity of the regulatory challenges they face. Part of this is a matter of scale. The number of daily trades on the New York Stock Exchange is now in the billions, an order of magnitude greater than when the Securities and Exchange Commission was created in 1934, presenting a far more difficult challenge to track suspicious activity.<sup>34</sup> For another example, healthcare devices have become increasingly complex and specialized, with some integrating AI into their features. This presents a unique challenge to the Food and Drug Administration's typical regulatory paradigm due to their tendency to degrade or drift after approval,<sup>35</sup> requiring new and data intensive modes of post-market surveillance.<sup>36</sup> In other words, the spread of AI-enabled devices across the healthcare system creates a need to equip regulators with AI-enabled oversight tools.

The corollary to this increasing focus on automation in regulatory enforcement is an erosion of the human elements in administrative processes, and potentially a hollowing out of the expertise which is a core pillar of legitimacy underlying administrative agencies. Survey data shows a strong correlation between perceptions of administrative expertise and public perceptions of legitimacy.<sup>37</sup> Some of the most trusted federal agencies, such as the Federal Reserve, maintain this status despite being relatively light in terms of the procedural rigor underlying their decision-making.<sup>38</sup>

The spread of AI in regulatory enforcement presents a potential existential challenge to administrative agencies, since it outsources human expertise to a set of rules interpreted and

---

<sup>32</sup> See Miguel Castro et al., *The Citizen's Perspective on the Use of AI in Government*, BOS. CONSULTING GRP. (Mar. 1, 2019), <https://www.bcg.com/publications/2019/citizen-perspective-use-artificial-intelligence-government-digital-benchmarking> [<https://perma.cc/CS4B-DGHC>].

<sup>33</sup> See Kenneth Culp Davis, *A New Approach to Delegation*, 36 U. CHI. L. REV. 713, 720–22 (1969).

<sup>34</sup> *New York Stock Exchange (NYSE) - statistics & facts*, STATISTA, <https://www.statista.com/topics/10015/new-york-stock-exchange-nyse/> [<https://perma.cc/Y422-745H>].

<sup>35</sup> AI models tend to change, and often degrade in their performance, in subtle ways over time, which poses an oversight challenge to systems which are designed to track significant material changes. See Jim Hodsworth, Ivan Belcic & Cole Stryker, *What is Model Drift?*, IBM (July 16, 2024), <https://www.ibm.com/topics/model-drift> [<https://perma.cc/C4G8-M896>].

<sup>36</sup> AKSHAY SREEKUMAR & PETER HORTON, UCLA INST. FOR TECH. L. & POL'Y, *LIABILITY PREEMPTION IN THE NEW REGULATORY FRAMEWORK OF DATA DRIVEN HEALTHCARE* (2022), [https://www.dropbox.com/scl/fi/2gfg6gdx1dd7wsy4fxrvp/LIABILITY\\_AND\\_PREEMPTION.pdf?rlkey=bbe36xhus19qiw2pyumx2xy92&e=1&dl=0](https://www.dropbox.com/scl/fi/2gfg6gdx1dd7wsy4fxrvp/LIABILITY_AND_PREEMPTION.pdf?rlkey=bbe36xhus19qiw2pyumx2xy92&e=1&dl=0) [<https://perma.cc/6YRP-8PCS>].

<sup>37</sup> Jon C. Rogowski, *The Administrative Presidency and Public Trust in Bureaucracy*, 1 J. POL. INSTS. & POL. ECON. 27 (2020).

<sup>38</sup> Bagley, *supra* note 28, at 382.

enforced by machines. This can obviously play into existing concerns about opacity in administrative operations, as well as perceptions of arbitrariness, due to the inscrutability of AI decision-making. AI systems typically deal poorly with edge cases, where an enforcement decision could conceivably go in either direction. Under the “arbitrary and capricious standard”, as spelled out in section 706 of the Administrative Procedures Act, agency decisions on questions of fact or policy may be set aside by a court if found to be “arbitrary, capricious, an abuse of discretion, or otherwise not in accordance with the law.”<sup>39</sup> As a result of the Supreme Court’s ruling in *Loper Bright v. Raimondo*, agency interpretations of law are not entitled to significant deference at all.<sup>40</sup> The fact that administrative agency staff may not be capable of controlling or even explaining the outputs of their tools poses a challenge under these governing standards. They are also problematic insofar as expertise and the exercise of discretion are a key underlying justification supporting the administrative state in the first place.<sup>41</sup>

Where AI systems are being developed by third-party contractors, accountability is further stymied by trade secrecy claims, which can serve as an additional shield against external accountability. Governments have the option to mandate robust public transparency requirements as part of the procurement process, though this is not always done. In the context of regulatory enforcement, secrecy may be motivated by a concern that disclosing too much information will undermine the efficacy or integrity of enforcement processes. The tension between transparency and effective law enforcement is not unique to the AI realm, and it manifests in some form under virtually every freedom of information or right to information framework.<sup>42</sup> Police routinely complain that disclosures of information about their investigative techniques, whether in response to an information request or in a judicial context, will compromise the efficacy of their operations and help bad actors to get away.<sup>43</sup> However, opacity can also undermine efforts to cultivate public trust in favor of the use of AI in regulatory enforcement, frustrating the ability to track and monitor a system’s performance and to obtain buy-in from impacted communities and other key external stakeholders. Some potential solutions to this challenge present their own tradeoffs. Programmers may increase model complexity, reconfigure models periodically, or add randomness in order to make the systems more resilient against attempts to game them, which may help assuage concerns that transparency will undermine the system’s efficacy in regulatory enforcement.<sup>44</sup> But these same measures also make the systems less interpretable, reducing the value of transparency as a mechanism of accountability.<sup>45</sup>

Similarly, while the democratic tensions flowing from the government’s increasing reliance on private sector contractors are not unique to AI, they take on greater salience in the context of AI due to the increasing level of autonomy these systems are capable of exercising, as well as the

---

<sup>39</sup> 5 U.S.C. § 706(2)(A).

<sup>40</sup> See *Loper Bright v. Raimondo*, 144 S. Ct. 2244 (2024).

<sup>41</sup> See Calo & Citron, *supra* note 3, at 818.

<sup>42</sup> See, e.g., MICHAEL KARANICOLAS ET AL., INTERPRETATION OF EXCEPTIONS TO THE RIGHT TO INFORMATION: EXPERIENCES IN INDONESIA AND ELSEWHERE 75–85 (2012), <https://www.law-democracy.org/wp-content/uploads/2010/07/Interpretation-of-Exceptions-To-the-Right-To-Information-Experiences-in-Indonesia-and-Elsewhere.pdf> [<https://perma.cc/5YET-2MQJ>].

<sup>43</sup> Hannah Bloch-Wehba, *Visible Policing: Technology, Transparency, and Democratic Control*, 109 CAL. L. REV. 917, 964-965 (2021).

<sup>44</sup> ENGSTROM ET AL., *supra* note 1 at 87.

<sup>45</sup> *Id.*

challenges that agency staff face in explaining their outputs.<sup>46</sup> In a 2021 article on the subject, Ryan Calo and Danielle Keats Citron pointed out that, if regulatory enforcement is essentially being delegated to AI tools supplied by third-party contractors, there is an argument to be made that administrative agencies are no longer necessary at all, since Congress could just as easily contract directly with the companies providing the enforcement tools in order to achieve their regulatory aims.<sup>47</sup>

While challenges around trust and legitimacy are paramount, since they cut to the core of administrative agencies' functions and mandates, there is a laundry list of other concerns related to AI's integration in regulatory enforcement. These include the fundamentally regressive nature of AI, since the systems necessarily rely on historical data to form their understanding of a particular challenge.<sup>48</sup> Like the proverbial general focused on fighting the last war rather than the next one, AI's dependence on data from earlier enforcement efforts may render it ill-equipped to handle the dynamic and adversarial nature of modern malfeasance. There are also due process concerns, as noted in the 2023 Executive Order (EO) 14110 on "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence."<sup>49</sup> In particular, these technologies present a challenge to both procedural and constitutional due process standards, with the latter including well-understood elements such as an impartial decisionmaker, providing a reason for decisions, notice, and the opportunity to be heard.<sup>50</sup>

There are also concerns that AI-enabled regulatory enforcement tools may be subject to gaming, where targets of enforcement or other interested actors learn how to manipulate the systems to achieve particular outcomes. While this challenge is not specific to machines, Daniel Ho, et al, raise particular concerns around the risk that third-party contractors who design these tools, or employees within the contracting agencies, may sell their knowledge about the inner workings of the machines to enforcement targets or other parties with skin in the game.<sup>51</sup> One can easily imagine that a person with knowledge of what kinds of patterns are likely to be flagged by an AI tool tracking suspected insider trading for the SEC could be tempted to sell this knowledge or even exploit it themselves.<sup>52</sup> The Standards of Ethical Conduct for Employees of the Executive Branch, issued by the Office of Government Ethics for codification at 5 C.F.R. Part 2635, states the general principle that "employees shall not engage in outside employment or activities, including seeking or negotiating for employment, that conflict with the official Government duties and responsibilities."<sup>53</sup> While there is no rule against federal agency employees moving to the private sector, federal agency employees are required to follow recusal rules while seeking employment with "persons whose financial interests would be directly and predictably affect by particular

---

<sup>46</sup> GOVERNMENT BY CONTRACT: OUTSOURCING AND AMERICAN DEMOCRACY (Jody Freeman & Martha Minow eds., 2009).

<sup>47</sup> Calo & Citron, *supra* note 3, at 832–35..

<sup>48</sup> AIURELIEN GERON, HANDS-ON MACHINE LEARNING WITH SCIKIT-LEARN, KERAS, AND TENSORFLOW: CONCEPTS, TOOLS, AND TECHNIQUES TO BUILD INTELLIGENT SYSTEMS 25–26 (Rachel Roumeliotis & Nicole Tache eds., 2nd ed. 2019).

<sup>49</sup> 88 Fed. Reg. 75,191, <https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence>.

<sup>50</sup> Chris Chambers Goodman, *AI, Can You Hear Me? Promoting Procedural Due Process in Government Use of Artificial Intelligence Technologies*, 28 RICH. J.L. TECH 700, 714-732 (2022).

<sup>51</sup> ENGSTROM ET AL., *supra* note 1 at 63.

<sup>52</sup> *Id.* at 87.

<sup>53</sup> 5 C.F.R. § 2635.101(b)(10).

matters in which the employees participate personally and substantially.”<sup>54</sup> Once hired, regulations under 18 U.S.C. 207 prohibit certain acts by former employees, which involve or appear to involve the unfair use of government employment.<sup>55</sup> However, it is unclear whether these rules are fit for purpose in an AI context, given that manipulation can come in subtle varieties, such as knowing what keywords are likely to route a patent application to an examiner with a particularly favorable rate of approval.<sup>56</sup>

The novel nature of these systems, and the under-explored ways in which different components may interact across the administrative apparatus, also gives rise to concerns about cascading failures, wherein the failure of one component triggers subsequent failures in complex interconnected systems.<sup>57</sup> Scholars have also expressed broader worries about the data and energy intensive models that these systems rely on and the implications of their development and expanding use for privacy and for the environment. Kate Crawford has worked extensively to document the massive energy and extractive costs that flow from the rise in popularity of generative AI.<sup>58</sup>

Together, these concerns paint a picture of a need for administrative agencies to tread cautiously in adopting AI, particularly for contentious or sensitive applications such as regulatory enforcement. While it is understandable that agencies are keen to shed public perceptions of bureaucratic inefficiency and portray themselves as being on the leading edge of technological innovation, support for more technologically enhanced government is likely to evaporate if the public comes to believe that AI systems are not trustworthy in performing these functions.<sup>59</sup>

In the next section, the existing legal landscape in the United States and elsewhere is discussed to present a snapshot of the regulatory safeguards that currently exist, and the conceptual gaps in these models.

---

<sup>54</sup> 5 C.F.R. § 2635.601.

<sup>55</sup> 5 C.F.R. § 2641.

<sup>56</sup> ENGSTROM ET AL., *supra* note 1 at 47–49.

<sup>57</sup> See generally Naeem Md Sami & Mia Naeini, *Machine learning applications in cascading failure analysis in power systems: A review* 232 *ELECTRIC POWER SYSTEMS RESEARCH* 110415 (2024).

<sup>58</sup> Kate Crawford, *Generative AI's Environmental Costs are Soaring – and Mostly Secret*, *NATURE* (Feb. 20., 2024), <https://www.nature.com/articles/d41586-024-00478-x> [<https://perma.cc/K7JD-6QLB>]. See also KATE CRAWFORD, *ATLAS OF AI: POWER, POLITICS, AND THE PLANETARY COSTS OF ARTIFICIAL INTELLIGENCE* (2021).

<sup>59</sup> See ENGSTROM ET AL., *supra* note 1 at 6–8.



## IV. The Existing Governance Landscape

Although successive U.S. administrations have taken an increasingly keen interest in the economic and sociological ramifications of AI, there has been relatively little specific focus on the technology's impact on regulatory enforcement, and only sporadic attention paid to its impact on government operations more generally. Legislative attention has been far more concerned with the private sector companies developing these new technologies, and in establishing an appropriate regulatory framework for their work, than on parallel developments across the administrative state.<sup>60</sup> Policy development processes are likely to advance particularly slowly in the regulatory enforcement space, since guiding rules or principles are likely to be tested, refined, and, in some cases, nullified through successive waves of judicial review.

Consideration of the impact of AI on government began under the Obama administration, particularly in a 2016 report which included recommendations on “AI and the Federal Government” and “AI and Regulation”, though the latter was mainly focused on oversight of AI products as opposed to the use of AI for enforcement.<sup>61</sup>

Under the first Trump administration, and as the underlying technologies continued to mature, there was a growing interest in public sector AI applications. Executive Order 13960, Promoting the Use of Trustworthy AI in the Federal Government, advanced the principle that agencies should “seek opportunities for designing, developing, acquiring, and using AI, where the benefits of doing so significantly outweigh the risks, and the risks can be assessed and managed.”<sup>62</sup> It also directed federal agencies to adopt principles that ensure AI use by the federal government is safe, reliable, and accountable, mandated interagency coordination and expertise to support ethical AI practices in government operations, and assigned the Office of Management and Budget to establish policy guidance and mandate public inventories of non-sensitive AI use cases. The AI in Government Act of 2020 created an AI Center of Excellence within the General Services Administration to support federal agencies with AI adoption, improve agency coordination, and publish federal AI initiatives. It also directed the Office of Personnel Management to establish an AI-focused job series and forecast staffing needs for federal AI roles.<sup>63</sup>

Ethical considerations continued to advance under the Biden administration. Executive Order 14110, on the Safe, Secure, and Trustworthy Development and Use of AI, required agencies to implement standardized AI safety evaluations, including for high-risk applications.<sup>64</sup> The order targeted AI-driven discrimination and privacy risks, promoted transparency through public labeling of AI-generated content, and encouraged the creation of global AI standards. The Biden administration also published the Blueprint for an AI Bill of Rights, which is probably the most

---

<sup>60</sup> See, e.g., Safe and Secure Innovation for Frontier Artificial Intelligence Models Act, S.B. 1047, 2023–2024 Reg. Sess. (Ca. 2024).

<sup>61</sup> EXEC. OFF. PRESIDENT, NAT'L SCI. & TECH. COUNCIL COMM. ON TECH., PREPARING FOR THE FUTURE OF ARTIFICIAL INTELLIGENCE 15–18 (2016), [https://obamawhitehouse.archives.gov/sites/default/files/whitehouse\\_files/microsites/ostp/NSTC/preparing\\_for\\_the\\_future\\_of\\_ai.pdf](https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf) [<https://perma.cc/9K2P-XRYF>].

<sup>62</sup> Exec. Order No. 13960, 3 C.F.R. 13960 (Dec. 3, 2020).

<sup>63</sup> AI in Government Act of 2020, Pub. L. No. 116-260, Div. U, tit. I, 134 Stat. 2286, 2286–89 (2020).

<sup>64</sup> Exec. Order No. 14110, 88 Fed. Reg. 75191 (Oct. 30, 2023).

high profile standard-setting document in this space.<sup>65</sup> Its principles and expectations are still relatively high level, though it also includes a number of concrete examples.

The two most significant moves to develop more concrete and actionable protocols for how the administrative state should approach AI in a regulatory enforcement context have come from the National Institute of Standards and Technology (NIST) and the Office of Management and Budget (OMB).

In January 2023, NIST published its “Artificial Intelligence Risk Management Framework” (AI RMF), which provides a model assessment process for agencies to map potential risks, develop tracking mechanisms, and respond appropriately.<sup>66</sup> The AI RMF establishes a taxonomy of potential risks flowing from the use of AI. Challenges defined under the AI RMF include reliability, accuracy, robustness, resilience, security, accountability, explainability, interpretability, privacy, fairness, and bias. The AI RMF provides a model risk assessment process for agencies to map potential risks, develop tracking mechanisms based on this mapping, measure risks as they emerge, and manage and respond appropriately.

Considered in the context of regulatory enforcement, it is particularly important for risks to be assessed institutionally and systematically, rather than purely from the perspective of harms that flow directly to the subjects of the decisions or other direct stakeholders. A robust assessment process should consider the risk that a system, even if it works perfectly, might nonetheless serve to undermine confidence in an agency and perceptions of legitimacy. Likewise, while human performance may provide a useful baseline for comparison, the fact that an AI program may return a lower level of erroneous decisions as compared to a traditional decision-making system, or lower levels of bias, should not be the end of the conversation in assessing whether it is fit for a purpose. If subjects of AI-driven enforcement decisions perceive that they are more unfair or arbitrary, then it is possible that the deployment of these systems will have a net negative impact on an agency’s legitimacy, even if the new system actually improves the accuracy of decision-making. Risk, in other words, should be understood holistically, and considered in the context of an entire organization, as opposed to limiting the assessment to a particular process. Similarly, while the AI RMF specifically mentions the importance of engaging with impacted communities, in the context of regulatory enforcement the perspectives of targets of enforcement, as well as communities that are otherwise impacted by decisions not to prosecute, should be complemented by considerations of the perspectives of the broader public.

Another important characteristic of the AI RMF is that it emphasizes ongoing evaluations throughout the AI lifecycle.<sup>67</sup> In other words, while early-stage assessments to inform decisions on whether to proceed with regulatory enforcement are vital, approval for a system to be developed or deployed should not mark the conclusion of the risk assessment process. Risks, impacts, and trade-offs should be mapped on an ongoing basis and include continuous assessment of whether

---

<sup>65</sup> *Blueprint for an AI Bill of Rights*, WHITE HOUSE OFF. SCI. TECH. POL’Y, <https://www.whitehouse.gov/ostp/ai-bill-of-rights/> [<https://perma.cc/UZF3-D96J>].

<sup>66</sup> U.S. DEP’T OF COM., NAT’L INST. STANDARDS & TECH., ARTIFICIAL INTELLIGENCE RISK MANAGEMENT FRAMEWORK: GENERATIVE ARTIFICIAL INTELLIGENCE PROFILE (2024), <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf> [<https://perma.cc/78BJ-W7UT>] (hereinafter AI RMF).

<sup>67</sup> *Id.*

the technology is delivering as promised or whether it is proving unfit for its purpose. A willingness to retire poor-performing systems, and to avoid falling victim to a sunk-cost fallacy, is vital. In the context of regulatory enforcement, allowing space for systems to be phased out presents a particular challenge insofar as appeals against adverse agency decisions may carry on for years. A post-hoc admission that an AI system—which was instrumental to previous enforcement decisions—is no longer fit for its purpose may undermine the agency over the course of appeals or reconsideration processes. This is an understandable disincentive to critical assessment. However, avoiding such a determination would only cause delay and make the inevitable decision to discontinue a problematic system even more difficult.

While the AI RMF provides a useful starting point, the format leaves significant discretion to implementing agencies. Risks are ultimately contextual determinations. They are resistant to centralized definition since they depend on the particular use case. There can even be significant variance within individual tools. A system may manifest a particular risk profile at the testing phase and introduce completely different problems in implementation.<sup>68</sup> The human element is also a significant factor to consider. Risk may depend on the individuals who are interacting with or using the tools, as well as their expectations and perceptions of its capabilities. The AI RMF relies on a sense of collective responsibility for managing the impacts of AI across the implementing agency, emphasizing the importance of diversity among the team considering potential risks.<sup>69</sup> It also requires a willingness to ask difficult and resource-intensive questions about the tradeoffs flowing from various use cases.

The second major regulatory development which is worth flagging is OMB’s AI policy memo, which was published in March 2024.<sup>70</sup> The memo includes a number of requirements and recommendations for executive branch agencies, including the designation of a Chief AI Officer and the establishment of AI Governance committees at CFO Act agencies to guide and coordinate issues related to AI implementation, including managing risks.

Probably the most noteworthy aspect of the memo is the requirement for agencies to track and publicly report all AI use cases, as well as to identify where AI uses are “rights-impacting” or “safety-impacting.” The memo also includes substantial discussion of risk management and mitigation efforts regarding rights-impacting or safety-impacting uses of AI, including conducting impact assessments and real-world testing, and implementing measures to address discrimination. There is significant overlap between these requirements and the content of the AI RMF. However, the introduction of a centralized framework for collecting these assessments and monitoring how these technologies are being deployed is a vital addition. The lack of such tracking has been a significant impediment to efforts to craft an appropriate response to the use of AI across the federal government. It is difficult to come up with a coherent public policy response if not everyone has a comprehensive understanding of the implications associated with using these technologies.

---

<sup>68</sup> *Id.* at 2.

<sup>69</sup> *Id.* at 23.

<sup>70</sup> Memorandum from Shalanda D. Young, Director, Off. Mgmt. & for the Heads of the Executive Departments and Agencies on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence (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> [<https://perma.cc/AD9D-HGTE>] (hereinafter AI Policy Memo).

OMB’s policy memo serves as an initial attempt to corral federal agencies around a rough and general set of standards by developing a framing of how and where these technologies are being used and by encouraging agencies to construct a model of different types of risk and accompanying mitigation strategies. In response to this prompt, agencies have begun issuing their own internal guidance on the use of AI. One example is the IRS, which in May 2024 issued an interim guidance memorandum which, among other things, establishes a use case inventory and defines an approval and workflow approving new AI applications, as well as establishing minimum practices for safety-impacting or rights-impacting AI.<sup>71</sup> The IRS interim guidance memorandum also designates a cross-functional AI Assurance Team and AI Project Teams to review and execute key governance steps, including steps such as impact assessments and ongoing risk evaluations, which broadly follow the standards spelled out in the AI RMF.

While the OMB policy memo and the AI RMF are the most prominent frameworks, there are other institutional actors which are relevant to potential regulatory efforts. These include ACUS, which in 2020 published a set of standards and considerations for federal agencies using artificial intelligence, including related to transparency, bias, capacity, procurement, privacy, data management, security, oversight, and decisional authority.<sup>72</sup> In 2021, the Government Accountability Office (GAO) published its own “Guidance for Creating Agency Inventories” around federal agency uses of AI. It is worth noting, however, that a follow-up study documented widespread noncompliance: of the 19 agencies to whom GAO offered recommendations, only 10 fully agreed to comply.<sup>73</sup> Compliance challenges are a major issue underlying any attempt to create effective and actionable standards across the executive branch, and there is a long-running debate regarding the relative benefits of binding, sanctions-based systems versus more informal structures built around administrative support and capacity-building.<sup>74</sup> Without commenting on the substance of the GAO recommendations, it is worth noting that there are few things which can undermine the perceived legitimacy of an administrative oversight structure more than issuing a recommendation or requirement which is subsequently ignored.<sup>75</sup> This is not to argue against ambitious standards or the demand for compliance with robust best practices, but it does illustrate the importance of ensuring that the oversight bodies are equipped with the tools and resources to spur meaningful change across administrative agencies.

- **Gaps in the Regulatory Environment**

The emerging regulatory environment, in guiding how AI may be used in administrative enforcement, places a heavy emphasis on transparency. While transparency is an essential ingredient in any effective oversight structure, it does not by itself present a solution, or even a response, to the accountability and other structural challenges posed by AI’s deployment.<sup>76</sup> The

---

<sup>71</sup> Memorandum from Barry Johnson, Chief Data and Analytics Officer, IRS on Interim Guidance for New IRM 10.24.1, Artificial Intelligence (AI) Governance and Principles (May 20, 2024), <https://www.irs.gov/pub/foia/ig/spder/interim-guidance-raas-10-0524-0001-artificial-intelligence-governance-and-principles-redacted.pdf> [<https://perma.cc/VC6G-MZ93>].

<sup>72</sup> Admin. Conf. of the U.S., *Statement #20: Agency Use of Artificial Intelligence*, 86 Fed. Reg. 6616 (Jan. 22, 2021).

<sup>73</sup> See GAO, *supra* note 2.

<sup>74</sup> See Michael Karanicolas & Margaret B. Kwoka, *Overseeing Oversight*, 54 CONN. L. REV. 655, 692–95 (2022).

<sup>75</sup> Bagley, *supra* note 28, at 392.

<sup>76</sup> See *infra* Section IV for a full discussion of transparency and best practices.

AI RMF notes that accountability presupposes transparency.<sup>77</sup> But while the latter is a precondition for the former, transparency is not itself sufficient to provide robust accountability, which in addition to information about these technologies requires opportunities for the public to meaningfully impact government policy related to their use.

Explainability, or the ability to characterize AI decisions in a way which renders their reasoning comprehensible to humans, is often cited as another key value, though it, too, is not an end in itself.<sup>78</sup> Rather, explainability is valuable because of its utility in facilitating meaningful review, supporting human autonomy, facilitating due process, strengthening perceptions of legitimacy, and providing guidance to future decision-makers.<sup>79</sup> The value of explainability, in other words, depends on complementary mechanisms to support these follow-on goals.

In dealing with high-risk applications, such as regulatory enforcement, placing a “human in the loop” is often emphasized as a mitigation tactic. This is unsatisfactory as a solution, in part because creating meaningful review over AI-generated decisions requires more than just human intervention. The oversight must be meaningful, and humans have a tendency to defer to automated recommendations. Meaningful review is made even more challenging due to the fact that, in many instances, there is no single “right answer” to assess AI recommendations against, as they may engage complex and subjective value judgments.<sup>80</sup> As due process rights related to high stakes decisions made by machines escalate, it begs the question as to whether the purported efficiency gains through the use of these systems may, in some instances, be illusory. The intensity of human review required to ensure that enforcement decisions have meaningful oversight may require just as many man-hours as having a human carry out the decision-making process independently. In such circumstances, it is worth asking whether a blanket prohibition against the use of AI as anything other than a research tool for human staff may be preferable in certain highly sensitive or contentious enforcement roles.

A related problem, which pervades much of the AI governance space, relates to challenges in connecting the general principles that usually ground high level guidance to more concrete and operational directions. As Cary Coglianese noted, it is one thing for governments to dictate that AI systems should be “fair”, “safe”, “explainable” and so forth: but determining what that means from an operational perspective is an entirely different matter.<sup>81</sup> The absence of clear performance standards, Coglianese observes, is what gives rise to a reliance on management-based regulation, which relies on process and protocol rather than attempting to achieve particular outcomes.<sup>82</sup> This, in turn, leads to new challenges as regulatory agencies grapple with how to ensure that risk

---

<sup>77</sup> See AI RMF at 12.

<sup>78</sup> See *What is Explainable AI?*, IBM, <https://www.ibm.com/topics/explainable-ai> [<https://perma.cc/A9DD-UWGB>].

<sup>79</sup> See Hofit Wasserman-Rozen, Ran Gilad-Bachrach & Nina Elkin-Koren, *Lost in Translation: The Limits of Explainability in AI*, 42 CARDOZO ARTS & ENT. L.J. 391, 410–11 (2024).

<sup>80</sup> Andrew Selbst, *Negligence and AI's Human Users* 100 B.U. L. REV 1315 (2020).

<sup>81</sup> Cary Goglianese, *How to Regulate Artificial Intelligence*, REGUL. REV. (Jan. 15, 2024), <https://theregreview.org/2024/01/15/coglianese-how-to-regulate-artificial-intelligence/> [<https://perma.cc/J9HX-8CDX>].

<sup>82</sup> *Id.*

assessments and other related processes are actually meaningful, as opposed to becoming mere paperwork exercises.<sup>83</sup>

However, probably the most pervasive structural weakness in existing accountability frameworks is that they assume a continued organic expansion of the use of AI across administrative agencies. Government attitudes to the use of AI frequently emphasize the importance of ensuring that agencies have adequate space to experiment and pilot new applications.<sup>84</sup> The OMB policy memo specifically points to the need to remove barriers to the responsible use of AI and achieve enterprise-wide improvements in AI maturity.<sup>85</sup>

A sense of apprehension at how to structure specific binding limits is understandable, given the novelty of these technologies, their complexity, the rate at which they are evolving, and the incredible range of functions where they are being piloted. At the same time, commenters talk of an “avocado ripeness problem” in finding the opportune time to impose strict regulations: just as an avocado can seemingly transition directly from being underripe to overripe, there is a thin line between when a fast-moving technology is too novel for observers to see clearly and understand its inherent risks, and when its use has become so deeply ingrained in government or the economy as to make effective regulation impossible.<sup>86</sup>

Before moving on, it is useful to examine a few comparable frameworks from the state and international levels, to develop a broader sense of how regulators elsewhere are responding to the emerging challenges posed by AI in regulatory enforcement.

- **AI Regulation at the State Level**

There have been a number of state-level initiatives aimed at regulating AI across the public sector. Though relatively few of them apply specifically to regulatory enforcement, several would likely impact the use of AI by the relevant state governments in different ways.

In California, AB 302 requires the California Department of Technology to coordinate with other interagency bodies to compile a comprehensive inventory of “high-risk automated decision systems” that state agencies are using, developing, or procuring.<sup>87</sup> These systems are defined as those that assist or replace human decision-making and have significant legal impacts, such as around access to housing, education, employment, credit, healthcare, and criminal justice.<sup>88</sup> The inventory must detail the decisions these systems make, the data they use, and any measures to

---

<sup>83</sup> Cary Coglianese & Jennifer Nash, *Compliance Management Systems: Do They Make a Difference?* (2020) (U. Penn. Inst. For L. & Econ. Rsch. Paper No. 20-35).

<sup>84</sup> *See, e.g.*, Exec. Order No. 13960, 3 C.F.R. 13960 (Dec. 3, 2020), which states that agencies should “seek opportunities for designing, developing, acquiring, and using AI, where the benefits of doing so significantly outweigh the risks, and the risks can be assessed and managed.”

<sup>85</sup> AI Policy Memo at 9.

<sup>86</sup> Woodrow Hartzog, *What STS Can (And Can't) Do for Law and Technology*, TECH. L. J. OF THINGS WE LIKE (LOTS) (May 19, 2023), <https://cyber.jotwell.com/what-sts-can-and-cant-do-for-law-and-technology/> [<https://perma.cc/VDS2-GZPK>] (reviewing Ryan Calo, *The Scale and the Reactor* (2022), <https://ssrn.com/abstract=4079851>).

<sup>87</sup> Cal. Gov. Code § 11546.45.5 (2023).

<sup>88</sup> Cal. Gov. Code § 11546.45.5, subd. (a)(4) (2023).

mitigate risks, including cybersecurity and bias.<sup>89</sup> The Department must submit this inventory in a report to the State Legislature by January 1, 2025, and annually thereafter.<sup>90</sup>

Although the AB 302 does not apply specifically to regulatory enforcement uses, these will likely be considered among the “high risk” uses being developed, and the requirements to document and evaluate the performance of these systems, and to implement safeguards against risks such as discrimination or cybersecurity threats, will likely have a substantial impact on how these technologies are deployed.

Another law under consideration in California is SB 896, the Generative Artificial Intelligence Accountability Act. Again, this is not specifically targeted at regulatory enforcement (or even public sector) functions, though it encourages agencies to engage in GenAI-focused rulemaking to clarify if and how existing regulations apply to GenAI or other automated decision-making systems.<sup>91</sup> The bill also mandates that state agencies using GenAI inform members of the public when they are interacting with GenAI regarding government services and benefits.<sup>92</sup>

The Maryland Artificial Intelligence Act, which was passed in May 2024, requires state agencies to develop a publicly available inventory of all systems using high-risk AI, including basic information about the AI systems such as their purpose and intended use.<sup>93</sup> A newly-created AI Subcabinet is tasked with defining “high-risk AI”, though it will also have a broader mandate to support AI innovation across the state government.<sup>94</sup> Earlier versions of the bill contemplated additional responsibilities, like identifying best use cases across state government units and testing proofs of concept, though these were ultimately excluded from the final draft.

In Washington, SB 5356 has been under discussion, in various forms, since at least 2021, and would require public notice and accountability measures for automated decision-making tools used by state agencies to produce legal effects on natural persons.<sup>95</sup> These include requiring each agency to complete an algorithmic accountability report for each automated decision-making tool in use, as well as to require agencies to notify people impacted by the use of automated decision-making tools of the system’s use, how to contest any decision involving an automated decision-making tool, and the degree to which human review resulted in the final decision, among other things.<sup>96</sup> The bill would also make any decision made or informed by an automated decision-making system subject to appeal “if a legal right, duty, or privilege is impacted by the decision”.<sup>97</sup>

In Illinois, there are, as of August 2024, several bills relevant to the use of AI in regulatory enforcement under consideration. HB 5116, known as the Automated Decision Tools Act, applies to deployers, including in administrative agencies, that use an automated decision tool, including those powered by AI, to make consequential decisions that produce significant effects on a

---

<sup>89</sup> Cal. Gov. Code § 11546.45.5, subd. (a)(4) (2023).

<sup>90</sup> Cal. Gov. Code § 11546.45.5, subd. (d)(1) (2023).

<sup>91</sup> S.B. 896 § 2, subd. (h), 2023–2024 Reg. Sess. (Cal. 2024).

<sup>92</sup> *Id.*

<sup>93</sup> S.B. 818, 2024 Reg. Sess. (Md. 2024).

<sup>94</sup> Md. Code, State Fin. & Proc. § 3.5–803, subd. (A)(1), (D)(1) (2024).

<sup>95</sup> S.B. 5356, 68th. Leg., Reg. Sess. (Wash. 2023).

<sup>96</sup> *Id.* § 4, subd. (8)(a), § 5, subds. (1)(f), (4) (2024).

<sup>97</sup> S.B. 5356 § 4, subd. (8)(c) (2024).

person's life and livelihood.<sup>98</sup> HB 5116 imposes several requirements, including that deployers conduct and submit annual impact assessments, and that they inform individuals when an automated decision tool is used to make or influence a consequential decision about them. HB 5116 would also require each deployer to develop and maintain a governance program to mitigate the risks of algorithmic discrimination.

HB 4705, the Artificial Intelligence Reporting Act, would require each state agency to designate a Chief Artificial Intelligence Officer from its existing staff to prepare an annual report detailing its use of covered algorithms for operations including enforcement.<sup>99</sup> These reports are to be published publicly by the Department of Innovation and Technology.

Finally, HB 4836 requires state agencies using AI systems, as well as entities deploying state-funded AI systems, to adhere to NIST standards for trustworthiness, equity, and transparency, and to submit algorithmic impact assessments based on the AI RMF to the Auditor General, and the Department of Innovation and Technology.<sup>100</sup>

- **International Case Studies: European Union**

Probably the best known and most influential international model is the European Union's *AI Act*, which imposes a sliding scale of requirements based on the purported risk of the use case, including obligations related to transparency, auditing, and oversight.<sup>101</sup> However, the *AI Act* also prohibits uses of AI which are deemed unacceptably risky. Although the *AI Act* is not specifically targeted at the public sector, the latter category includes some government applications, particularly the use of AI for predictive policing, to develop social credit scores, or for real-time biometric tracking in public spaces. The *AI Act* also includes some discussion of public sector uses that would be considered high risk, including any use of these systems to determine access to essential public services and benefits (such as healthcare), as well as all uses related to law enforcement, migration, border control, and the administration of justice and democratic processes. High risk systems are also required to be registered in a public database unless their uses are for law enforcement or migration. The *AI Act* also contains blanket exclusions for AI systems that are exclusively designed for military, defense, or national security purposes.

The core of the mitigation practices envisioned by the *AI Act* revolve around a conformity assessment, designed to ensure that the system complies with data quality, traceability, transparency, human oversight, accuracy, cybersecurity, and robustness standards. The assessment is meant to be repeated every time the system or its purposes is substantially modified, though defining a substantial modification may pose a conceptual challenge as a result of the tendency of some AI systems to change in steady but subtle ways after they enter the market.<sup>102</sup> The *AI Act* also requires the development of risk management systems that include testing and assessment at

---

<sup>98</sup> H.B. 5116, 103rd Gen. Assemb. (Ill. 2024).

<sup>99</sup> H.B. 4705, 103rd Gen. Assemb. (Ill. 2024).

<sup>100</sup> H.B. 4836, 103rd Gen. Assemb. (Ill. 2024); AI RMF.

<sup>101</sup> *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts*, COM (2021) 206 final, at 5.2.2 (Apr. 21, 2021), <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52021PC0206> [<https://perma.cc/NLS2-AY53>].

<sup>102</sup> Hodsworth et al., *supra* note 35.



both the piloting and the post-market phases, with accompanying reporting requirements, as well as requirements related to transparency, accuracy, and data quality.

Key oversight functions from the *AI Act* are delegated to technical standard setting organizations, though the main enforcement responsibilities are in the hands of national authorities, as well as the European Commission's AI Office, which provides strategic guidance and governs general-purpose AI models. It is worth noting that the *AI Act* includes substantial sanctions, of up to €35M or 7% of the total worldwide annual turnover of the preceding financial year (whichever is higher) for infringements. The *AI Act* also envisions regular audits and post-market monitoring by these authorities. Together, the rules are meant to create a system which is highly adaptable and iterative, in line with the evolving nature of the underlying technologies.

- **International Case Studies: Canada**

In Canada, government uses of AI are governed by the *Directive on Automated Decision-Making* (the *Directive*), under the auspices of the Treasury Board of Canada Secretariat.<sup>103</sup> Canada was an early mover in the AI governance space, having passed the *Directive* in 2019, though it has been subject to regular updates since then. Although the law does not specifically distinguish between uses for regulatory enforcement and other functions, it is limited in its application to cases where the AI system is processing “client data,” which in practical terms means that it is heavily focused on cases where a person or organization is seeking a government service or benefit, or is a target of enforcement. In other words, use cases like background policy research or personnel management functions are outside the purview of the rules. It is also worth noting that the *Directive* is forward-facing, only applying to applications subsequent to its entry into force, which allows for a gradual ramp up of oversight responsibilities.

The *Directive* relies on a combination of public accountability and risk-based impact assessments to support responsible use of AI. It introduces a requirement to carry out an Algorithmic Impact Assessment (AIA) prior to the production of an automated decision-making system, and to publish the results online.<sup>104</sup> The impact of the decision and the importance of the rights or interests engaged leads to a sliding scale of obligations. At lower levels, these include requirements for data bias testing and the provision of generalized explanations for common decision results. At the higher end, requirements include human intervention in the decision-making process, publication and peer-review, the provision of a “a meaningful explanation” for negative outcomes, and Treasury Board approval for the system to operate.<sup>105</sup>

Regarding public accountability, the *Directive* introduces a robust notification requirement, mandating institutions that utilize automated decision-making systems provide clear, prominent, and plain-language notices to the public of this fact on their website.<sup>106</sup> The *Directive* requires that

---

<sup>103</sup> Treasury Board of Canada Secretariat, *Directive on Automated Decision-Making* (2019), <http://www.tbs-sct.gc.ca/pol/doc-eng.aspx?id=32592&section=html> [<https://perma.cc/TWV6-KE2B>].

<sup>104</sup> *Algorithmic Impact Assessment Tool*, GOV'T OF CANADA, <https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/algorithmic-impact-assessment.html> [<https://perma.cc/22QC-7ETD>].

<sup>105</sup> *Directive on Automated Decision-Making*, *supra* note 103, at Appendix C.

<sup>106</sup> *Id.* at 6.2.

AIAs must be published online, with the intention of spurring public engagement to ensure the process is meaningful. The *Directive* also includes recommendations for consultation and engagement with impacted communities, though cost and logistical concerns mean that these are not currently required.

Like its American counterparts, the *Directive* places few hard and fast restrictions on the use of AI for various applications, though cases using higher levels of risk require approval from a senior political appointee. In practice, administrative agencies have been reticent to hand over direct authority to automated decision-makers, instead incorporating them as research or assessment tools to aid human decision-making. In Canada, as elsewhere, there is a tension between the need for robust centralized oversight and the requirement that impact assessments be delegated to those with the greatest contextual understanding of a particular use case. Although Treasury Board involvement in most assessments is not strictly mandatory, in practical terms agencies have been keen to draw on the expertise that TBS is able to offer in developing a robust assessment process.

- **International Case Studies: Singapore**

Singapore has been another early leader in developing AI governance structures, particularly through the launch of its *Model AI Governance Framework* and, more recently, the development of *A.I. Verify*, a testing framework toolkit to support self-assessment by those developing or deploying AI technologies.<sup>107</sup> These frameworks are framed as voluntary guidance, rather than strictly binding requirements.

The *Framework* leans heavily on the values of explainability, transparency, and fairness, as well as emphasizing that the technologies should be human-centric and focused on supporting human capabilities and the interests of human beings. It also emphasizes the importance of iteration, calling on relevant bodies to institute a documented review process which will “continually identify and review risks relevant to their technology solutions, mitigate those risks, and maintain a response plan should mitigation fail.”<sup>108</sup>

In addition to emphasizing the importance of assessing data sets for inaccuracy or bias, including through maintaining robust records of data provenance and lineage, the *Framework* suggests differentiating the data sets used for training, testing and validation. The *Framework* also suggests expanding human oversight and human involvement in decision-making where risk is particularly heightened, with the latter circumstance being defined as a multiplier of the severity of harm by the probability of harm.

---

<sup>107</sup> *Singapore’s Approach to AI Governance*, PERSONAL DATA PROT. COMM’N (May 2022), <https://www.pdpc.gov.sg/Help-and-Resources/2020/01/Model-AI-Governance-Framework> [<https://perma.cc/G45K-BAAL>].

<sup>108</sup> SINGAPORE PERSONAL DATA COMMISSION, MODEL ARTIFICIAL INTELLIGENCE GOVERNANCE FRAMEWORK 29 (2d. ed. 2020), <https://www.pdpc.gov.sg/-/media/files/pdpc/pdf-files/resource-for-organisation/ai/sgmodelaigovframework2.pdf> [<https://perma.cc/NQ4F-6R28> ].

## V. Key Values Underpinning an Appropriate Framework for AI in Regulatory Enforcement

The use of AI in regulatory enforcement presents opportunities to shift traditional administrative paradigms in novel and valuable ways. Because these systems are more malleable than human decision-makers, they offer new possibilities for achieving regulatory objectives while combating bias. There are also potential gains in efficiency and processing power. As noted in Section II, in many instances the increasing complexity of the modern regulatory landscape has given rise to an arms race of sorts, where regulatory agencies need increasingly advanced technology to keep up with the entities they are overseeing. Nonetheless, the effective use of AI for regulatory enforcement depends on its ability to be deployed in a manner which maintains public trust in the federal government. Trustworthiness is a challenging commodity, especially as it pertains to complicated institutional structures. Moreover, an agency or system is only as trustworthy as its weakest characteristics.<sup>109</sup>

This section considers the earmarks of a strong system of oversight for the use of AI in regulatory enforcement and provides recommendations to safeguard the legitimacy of the federal government in the context of expanding experimentation with AI.

### A. Understanding Risk and Risk Assessments

In the United States and around the world, the dominant governance model for AI focuses on assessing and mitigating risk. This idea is central to NIST's AI RMF and virtually every other major guidance document published across the executive branch, as well as to parallel efforts in the European Union, Canada, and Singapore.

It is easy to understand the appeal of a risk-based framework since it builds on existing models of regulation that are applied to a range of other roughly analogous harms—from environmental pollution to privacy and human rights impacts.<sup>110</sup> A commonality between these categories of harm is that they are all diffuse and difficult to measure or establish strict causality for. Algorithmic impact assessments have emerged as a core component of responsible AI use, as a successor to established models for environmental impact assessments, privacy impact assessments, and human rights impact assessments.<sup>111</sup>

Much of the momentum in favor of risk-assessment models lies in this familiarity, which may be particularly valuable in attempting to build guardrails around a novel and fast-moving technology like AI. But the relatively long track record for this model of governance also demonstrates that, along with its strengths, there are weaknesses and blind spots.

---

<sup>109</sup> See AI RMF at 12.

<sup>110</sup> See A. Michael Froomkin, *Regulating Mass Surveillance as Privacy Pollution: Learning from Environmental Impact Statements*, 5 U. ILL. L. REV. 1713, 1757–58 (2015); Human Rights Council, Rep. of the Special Representative of the Secretary General on the Issue of Human Rights and Transnational Corporations and Other Business Enterprises, John Ruggie, UN Doc A/HRC/17/31 (2011), [https://www.ohchr.org/Documents/Issues/Business/A-HRC-17-31\\_AEV.pdf](https://www.ohchr.org/Documents/Issues/Business/A-HRC-17-31_AEV.pdf) [<https://perma.cc/R9N5-4ZJ2>].

<sup>111</sup> Andrew D. Selbst, *An Institutional View of Algorithmic Impact Assessments*, 35 HARV. J. L. & TECH. 119 (2021).

For example, while risk regulation is well-adapted to mitigate certain structural harms, it is less effective at mitigating individualized harms. This point is key in considering the use of AI in regulatory enforcement, where the consequences of decision-making are particularly sharp for the individual or entity on the other end of the process. Where AI systems are being used for regulatory enforcement, implementing entities should understand the limitations of a risk-based approach. Responsible use of AI for regulatory enforcement may require that risk assessments be complemented by prohibitions on certain particularly sensitive use cases (such as where decisions have a significant impact on fundamental rights),<sup>112</sup> and appropriate mechanisms to compensate individuals for specific harms incurred.<sup>113</sup>

Experience also suggests that risk regulation is better at addressing predictable and easily quantifiable harms as opposed to the sort of “unknown unknowns” that are prevalent in considerations of the impact of AI.<sup>114</sup> In a 2011 article assessing the efficacy of the Nuclear Regulatory Commission’s treatment of uncertain risks, Daniel Farber describes a rulemaking process which assumed that certain wastes would have no impact on the environment since they would be in a sealed repository.<sup>115</sup> Although the agency eventually acknowledged that the risk of a leak was unknown, and was not zero, the perception that the danger was relatively remote led the agency to effectively round-down their assessment.<sup>116</sup> This approach proved misguided when a proposed disposal site was found to have fractures in its bedrock, which would have allowed for water percolation and potential leakage of nuclear materials.<sup>117</sup>

The challenge of “unknown unknowns” is particularly thorny in the context of AI given that many of the harms which are built into existing risk management frameworks, including the AI RMF, are fuzzy at best. There is little agreement on how terms like “fairness” should be applied, in practical terms, and even less consensus on how to understand these principles mathematically.<sup>118</sup> By contrast, audits that are assessing known flaws, particularly data-based ones, such as the impacts of biased or otherwise problematic data sets, are on more familiar ground. The result is that harms like a loss of public trust or legitimacy, which are more difficult to pin down, are likely to be obscured or devalued in a risk-based analysis, which naturally focuses on harms that are easier to measure such as error rates.<sup>119</sup>

As a result, it is important for risk assessment processes to be implemented with what Margot Kaminski calls “epistemic humility,” by acknowledging the tendency of AI systems to surprise us,

---

<sup>112</sup> See Case C-817/198, *Ligue Des Droits Humains v. Conseil Des Ministres*, Official J. of Eur. Union C 34/4 (2022), <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A62019CA0817> [<https://perma.cc/7F5H-QQ89>].

<sup>113</sup> Margot E. Kaminski, *The Developing Law of AI: A Turn to Risk Regulation*, at 15, in *THE DIGITAL SOCIAL CONTRACT: A LAWFARE PAPER SERIES* (2023), [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4692562](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4692562).

<sup>114</sup> *Id.* at 14–15.

<sup>115</sup> Daniel A. Farber, *Uncertainty*, 99 *GEO. L.J.* 901, 910–11 (2011).

<sup>116</sup> *Id.* at 911.

<sup>117</sup> *Id.* at 950.

<sup>118</sup> See Alicia Solow-Niederman, *Information Privacy and the Inference Economy*, 117 *NW. U. L. REV.* 357, 420 (2022).

<sup>119</sup> Jacob Metcalf, Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh, & Madeleine Clare Elish, *Algorithmic Impact Assessments and Accountability: The Co-construction of Impacts*, in *ASSOCIATION FOR COMPUTER MACHINERY CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY* (2021). <https://dl.acm.org/doi/pdf/10.1145/3442188.3445935> [<https://perma.cc/TMB7-R4JB>].

sometimes in harmful or destructive ways, and to incorporate this understanding into the heart of the decision-making process.<sup>120</sup> Similarly, assessments should factor in the natural instinct to perceive AI systems as working more effectively than they do, and as being applicable to settings where they are not fit for purpose.<sup>121</sup> As a consequence of the tendency of new technologies to spread, regulatory assessments of the impact of AI should also incorporate a presumption that systems will expand beyond their initial approved uses. Above all, risk assessments should ensure adequate space to consider worst case scenarios, rather than allowing the outcomes which are perceived as most likely to dominate the calculus.

Different uses of AI can present vastly different risk profiles for administrative agencies. Uses related to regulatory enforcement represent some of the areas of greatest concern, due to the direct proximity of these systems with agency outputs, and the fact that enforcement decisions nearly always engage with thorny procedural and civil rights questions that are core to perceptions of legitimacy in agency decision-making. While it is unlikely that the use of AI in a regulatory enforcement context requires a fundamentally different risk assessment process than other public sector applications, for these use cases a rough corollary may be drawn between the degree of risk and the impacts of the AI on agency outputs.<sup>122</sup> The latter determination may be connected to the level of autonomy that these systems enjoy, though with the caveat that human review does not necessarily mitigate the dominant role AI systems may play in determining outcomes. Nonetheless, the use of AI as a research support tool for human decision-makers is likely to raise fewer concerns than where the AI is directly making enforcement decisions or recommendations. This distinction may be difficult to pin down, given the heavy influence that even early-stage research support can have over the shape of the final decision, and the deference with which humans treat AI-powered recommendations. As a result, in a risk assessment process over the use of AI in regulatory enforcement, the influence of the system over the final decision should be understood as a spectrum, rather than a binary question of whether or not it is producing autonomous outcomes.

Ultimately, while risk-based processes are an important tool in promoting appropriate guardrails around the use of AI in regulatory enforcement, their efficacy will be heavily dependent on the spirit which implementing agencies apply these assessments. Risk regulation is an extremely broad concept, which can mean different things to different groups, guided by a range of legal, sociological, institutional, and historical factors. In order to ensure that assessments are meaningful, and guided by appropriate understandings of the risk landscape, a robust external consultation structure is essential.

## **B. Public Engagement**

It is a core principle of the rule of law that legitimacy rests on adequately justifying laws to the public.<sup>123</sup> Public engagement fulfils several important functions in a robust governance framework,

---

<sup>120</sup> Kaminski, *supra* note 108, at 14.

<sup>121</sup> Theo Araujo, Natali Helberger, Sanne Kruikemeier & Claes H. de Vreese, *In AI we trust? Perceptions about automated decision-making by artificial intelligence*, 35 AI & SOCIETY 611 (2020).

<sup>122</sup> See ENGSTROM ET AL., *supra* note 1, at 10.

<sup>123</sup> See Mireille Hildebrandt, *Privacy as Protection of the Incomputable Self: From Agnostic to Agonistic Machine Learning*, 20 THEORETICAL INQUIRIES L. 83, 113 (2019).

including helping to support democratic accountability and, in the context of regulatory enforcement processes which are targeted at individuals, to support a sense of dignity among those subject to administrative decisions, for better or for worse.

In addition to core democratic reasons why any uses of AI within the public sector should be subject to robust public consultation, external engagement is a valuable tool to boost the quality of risk assessments, and avoid their capture by the private sector or government agencies which have a direct stake in the systems' approval. Representatives from communities which are impacted by algorithmic enforcement, either directly through being a subject of regulation or indirectly through secondary effects of regulatory decisions, may be helpful to quantify and enumerate harms and concerns which may be difficult to isolate as part of an internal risk assessment process.<sup>124</sup> Impacted communities may also be able to spot problems with AI systems' outputs that elude their operators or designers, and which may not even be captured by sophisticated auditing or assessment processes. For example, in the case of Arkansas' infamous experiment with algorithmic healthcare determinations, the patients realized that there were errors in the outcomes even while officials, and the system's designers, insisted that it was working exactly as intended.<sup>125</sup>

Not all public engagement processes are created equal. At the extreme, it is possible to distinguish between meaningful opportunities for the public to influence agency decision-making around whether and how to use AI technologies in regulatory enforcement, and the mere opportunity to provide feedback. The latter may serve as a procedural smokescreen, masking capture or preset decisions behind formal procedural equality.

Although there are aspects of public engagement which are unique to the context of AI in regulatory enforcement, the challenge of providing meaningful avenues for public policy consultation is common to democracies around the world, as well as to intergovernmental organizations, and virtually any other institution which seeks public legitimacy. As a result, there is a robust set of international standards to draw on in attempting to make these processes meaningful.

For example, the Council of Europe's Code of Good Practice for Civil Participation in the Decision-Making Process (Code) sets four levels of participation: information, consultation, dialogue, and partnership.<sup>126</sup> The latter category, which is the highest standard, involves close collaboration, including service provision activities, participatory forums, and the establishment of co-decision-making bodies.<sup>127</sup> The Code envisions an engagement process which spans the entire lifecycle of a governmental decision-making process, beginning at the Agenda Setting phase and extending through implementation to Monitoring and Reformulation processes.

---

<sup>124</sup> See Margot E. Kaminski, *Voices In, Voices Out: Impacted Stakeholders and the Governance of AI*, 71 UCLA L. REV. DISCOURSE 176, 179 (2024).

<sup>125</sup> Lecher, *supra* note 19.

<sup>126</sup> *Code of Good Practice for Civil Participation in the Decision-Making Process*, COUNCIL EUR. CONF. NGOS, at 3, <https://www.coe.int/en/web/ingo/civil-participation> [<https://perma.cc/5WST-YQ24>].

<sup>127</sup> *Id.*

Consultation structures take a number of forms that may include engaging with the public to inform them of rule-making procedures, through stakeholder meetings, the designation of public representatives, structured briefings to air differing views on a controversial public policy question, advertising campaigns to solicit input, negotiated rulemaking, etc.<sup>128</sup>

While transparency is not a panacea to resolve all concerns related to the use of AI in regulatory enforcement, robust transparency is a precondition for effective stakeholder engagement. In a 2024 paper on the subject, Margot Kaminski presents a model of transparency which focuses on the two-way flow of information, including both the “voices in” [to government], through meaningful opportunities to provide feedback, and the “voices out” [from the government], through robust transparency practices which ensure the public is well-informed about the context in which these technologies are being developed and deployed.<sup>129</sup> A good first step to ensuring a robust flow of information to relevant stakeholders, which can inform their participation and responses, is the publication of risk assessment or AI impact assessment results. However, the most relevant community stakeholders may not have adequate subject-matter expertise to process the kinds of technical data developed as part of an AI audit process.<sup>130</sup> Agencies seeking to cultivate relationships with external stakeholders from impacted communities will need to either invest the resources to translate these documents so that they are more generally accessible, or to sponsor training and upskilling for community representatives or organizations to the point where they can engage with advanced questions related to the use of AI in regulatory enforcement.<sup>131</sup>

Agencies interested in boosting public participation in external consultation processes will also need to be mindful of the timing and location of engagement opportunities. A strict adherence to 9-5 business hours may mean that working people are unable to join. Childcare and travel costs may also present an obstacle to in-person participation, which may require resources to help mitigate. These issues are less likely to manifest where a consultation is carried out using remote participation, though an online format may be less satisfying for participants, and less conducive to robust and candid conversations. Where impacted groups include persons with disabilities, or persons who may not speak English, as well as other historically underrepresented communities, there may be a need for additional measures to bridge these challenges. Community organizations can serve as a key liaison to support participation among such historically marginalized groups, though this requires agencies to devote resources to cultivating productive and meaningful relationships with civil society partners.

On the latter point, it is noteworthy that there are a small but growing number of civil society organizations which are specifically focused on AI-related issues. For example, the Algorithmic Justice League, a Cambridge, Massachusetts based non-profit, founded the Algorithmic Vulnerability Bounty Project as a mechanism for outsourcing the identification of AI-driven harms

---

<sup>128</sup> Bagley, *supra* note 28, at 385.

<sup>129</sup> Kaminski, *supra* note 124, at 194–95.

<sup>130</sup> Sasha Costanza-Chock, Inioluwa Deborah Raji, & Joy Buolamwini, *Who Audits the Auditors? Recommendations from a field scan of the algorithmic auditing ecosystem: in* FACCT '22: PROCEEDINGS OF THE 2022 ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY 1571 (2022).

<sup>131</sup> Margot E. Kaminski & Gianclaudio Malgieri, *Algorithmic Impact Assessments Under the GDPR: Producing Multi-Layered Explanations*, 11 INT’L DATA PRIV L. 125, 139 (2021) (encouraging “companies, or regulators, to help fund the involvement” of impacted individuals and to “provide technical expertise or the resources for obtaining technical expertise”).

to the public.<sup>132</sup> This evolved into the Community Reporting of Algorithmic System Harms project, which aims to mobilize an empowered community to report and advocate for the redress of algorithmic harms.<sup>133</sup> While these kinds of initiatives are still relatively thin on the ground, agencies should capitalize on them where they already exist, and should explore options to provide resources to support and expand their work.

- **Policy Engagement vs. Enforcement Engagement**

External consultation related to the use of AI in regulatory enforcement may be divided into two general categories: policy engagement mechanisms, focused on ensuring that members of the public have an ability to impact decisions related to the development and deployment of AI systems; and enforcement engagement mechanisms, which generally revolve around allowing the subjects of an enforcement decision, as well as other interested parties, to opine on a particular enforcement process and its outcomes.

Policy engagement mechanisms, which are meant to address broader systemic concerns, can be particularly valuable in helping administrative agencies adopt an approach to the development and deployment of AI and other algorithmic tools which is in line with public values and expectations. Useful functions of this level of engagement may include helping to define key terms, such as “discrimination” and “fairness” or “less favorable treatment”.<sup>134</sup> External insights may also feed directly into risk assessment processes, both in shaping how assessments are carried out, as well as responding or providing feedback on draft assessments that are being reviewed. Here, stakeholder input may be invaluable to determine whether, for example, a particular harm is being under-appreciated or mischaracterized.

External participation in policymaking may involve an official rulemaking process or may be carried out on a less formal basis as novel issues arise, for example through a standing committee of community participants. The latter model allows for the development of greater subject matter expertise among participants, which in turn can generate more specific and meaningful guidance for the agencies. However, this model is obviously more labor intensive for both the agency and the community participants, which may push the demands on the latter’s time beyond what might be expected on a volunteer basis. It may be worth exploring schemes which compensate community or civil society participants for their engagement, though this needs to be managed carefully in order to avoid creating perverse incentive structures.

By far the most important ingredient in maintaining robust civil society engagement is to ensure that the consultations are meaningful, and that they are perceived as such. Community participants are likely to sour on engagement processes if they feel that their voices are not significantly impacting policy. Where relationships have been forged with trusted and sophisticated community

---

<sup>132</sup> Fast Company, *Meet The Computer Scientist and Activist Who Got Big Tech to Stand Down*, MEDIUM (Aug. 4, 2020), <https://medium.com/fast-company/meet-the-computer-scientist-and-activist-who-got-big-tech-to-stand-down-23d95e0347e7> [<https://perma.cc/KAA8-HXUM>].

<sup>133</sup> *Help Prevent, Report, and Redress Algorithmic Harms*, ALGORITHMIC JUSTICE LEAGUE, <https://www.ajl.org/crash-project> [<https://perma.cc/KJR5-VXNH>] (formerly Algorithmic Vulnerability Bounty Project).

<sup>134</sup> See S.B. 5116, 67th Leg., Reg. Sess. § 4 (Wash. 2021).



partners, and especially where resources have been invested to provide training to these partners, it is critical that agencies manage these connections carefully.

Enforcement engagement mechanisms target a more easily defined group consisting of, first and foremost, the persons who are subject to decisions where AI is a significant part of the assessment process.<sup>135</sup> While this engagement may overlap with the right to due process and explainability, it extends beyond the specific right to challenge outcomes of decisions to broader procedural dissatisfaction. For example, where regulatory enforcement decisions concern pollution impacting a particular geographic region, residents of that region might also be considered as targets for consultation. While these consultations will naturally be more focused on an individual outcome than any broader structural concern, it is possible for individual review or appeals mechanisms to connect back to broader risk assessment processes or other policy determinations, for example by requiring that a successful appeal (or a number of successful appeals above a certain threshold) should trigger a review of the original risk assessment, or that individual appeal outcomes should be factored into regularly scheduled reassessments.

### C. Recognizing Failure

Among the most important characteristics of a robust system of oversight for the use of AI in regulatory enforcement is that there must be adequate consideration of when and whether to phase out these tools. AI is a relatively young technology, whose deployment across the public sector is still in its early stages. And yet, already there are plenty of examples of AI failures. These include not only the well documented issues with bias and discrimination documented in Section II, but also examples where AI has simply failed to deliver on the results that its proponents promised.

In their previous publication for the Administrative Conference of the United States (ACUS), Engstrom et al documented the case of the Sigma system piloted at the USPTO, which was never deployed since it failed to improve efficiency unless its users had a computer science background.<sup>136</sup> A subsequent report on experiences using AI-powered tools to track and clear superfluous regulations, which was authored for ACUS by Catherine Sharkey, noted similar performance challenges in other systems. Staff at the Centers for Medicare and Medicaid Services (CMS) noted numerous false positives, and complained about the labor-intensive process of checking the enormous number of regulations that were flagged.<sup>137</sup> While these failures are not intended to represent the totality of government agencies' experiences in using AI, they demonstrate that, at the very least, the deployment of these technologies has been a mixed bag, though Professor Sharkey's paper suggested that officials were still bullish about the overall potential for AI to improve their agency's operations.<sup>138</sup>

There are a number of factors which may play into a tendency to be unduly optimistic about AI's performance or utility for a given task. Part of this is embedded in human nature, as fascination

---

<sup>135</sup> See Margot E. Kaminski & Jennifer M. Urban, *The Right to Contest AI*, 121 COLUM. L. REV. 1957 (2021).

<sup>136</sup> ENGSTROM ET AL., *supra* note 1 at 48.

<sup>137</sup> CATHERINE M. SHARKEY, ALGORITHMIC RETROSPECTIVE REVIEW OF AGENCY RULES 28 (May 3, 2023) (report to the Admin. Conf. of the U.S.).

<sup>138</sup> See, e.g., *id.* at 34 (quoting a CMS official's opinion that "We absolutely need to add technology into our process"). See also ENGSTROM ET AL., *supra* note 1.

with new technologies can lead to a natural tendency to over-estimate the capability of AI.<sup>139</sup> Research has demonstrated that humans interacting with AI tend to be relatively more tolerant of failures in their performance assessments as a result of the perception that AI has less control over unfavorable outcomes.<sup>140</sup> In circumstances where a considerable amount of time or resources have been invested in developing and deploying a system, sunk cost fallacy may play a role. Biased assessment standards may also manifest as the result of asymmetries in how successes or failures of these systems are tracked and evaluated. For example, a system which under-estimates the threat from a criminal defendant, and recommends their release only to have them reoffend, will receive negative feedback for the mistake.<sup>141</sup> However, if the next defendant's threat level is over-estimated, and they are remanded to custody as a result, there will be no concomitant opportunity to assess whether the system was wrong. A system may therefore learn that it is possible to game its own evaluation, and skew in a biased direction as a result.

- **Assessing Risk and Assessing Failure**

A core challenge to most risk assessment frameworks, including those proposed under the AI RMF, is that they typically adopt a relatively permissive approach to new technologies, with a heavy emphasis on preserving space for agencies to innovate.<sup>142</sup> Risk regulation implies a tradeoff. The assessment process, and the adoption of certain remedial steps, are necessary costs that entities must bear to access the efficiency and processing gains that AI can provide. This “techno-correctionist” tendency is useful for mitigating problematic aspects of widespread adoption, such as through threats to security, but is less useful for answering broader questions about whether a category of technologies are fit for use among regulatory enforcement functions.<sup>143</sup>

Agencies contemplating the use of AI for regulatory enforcement must be mindful of this gap in the risk-based paradigm, and work to complement their risk assessment framework with a careful and critical assessment of the costs and challenges of pursuing an AI-based solution in the first place. This assessment should include difficult questions about whether the operational context presents too great a challenge, or is otherwise unsuited to automation, or whether the potential institutional harm to legitimacy and reputation are too great. It should also consider whether, to put it bluntly, advocates for a particular system are selling snake oil.<sup>144</sup>

While AI, in general, has enormous potential, it is also buoyed by massive amounts of hype, much of which is generated by stakeholders with a direct financial interest in frothing up enthusiasm for new AI applications. Tech-solutionism, and a desire to appear on the cutting edge of innovation, can be powerful drivers in favor of expanding uses of AI. But there are many instances where AI is unfit for a given application, such as where there is insufficient underlying data to power a

---

<sup>139</sup> See AI RMF at 4.

<sup>140</sup> S. Mo Jones-Jang & Yong Jin Park, *How do people react to AI failure? Automation bias, algorithmic aversion, and perceived controllability*, 28 JOURNAL OF COMPUTER-MEDIATED COMMUNICATION 1 (2022).

<sup>141</sup> Noel L. Hillman, *The Use of Artificial Intelligence in Gauging the Risk of Recidivism*, AM. BAR ASS'N: JUDGES' J. (Jan. 1, 2019), [https://www.americanbar.org/groups/judicial/publications/judges\\_journal/2019/winter/the-use-artificial-intelligence-gauging-risk-recidivism/](https://www.americanbar.org/groups/judicial/publications/judges_journal/2019/winter/the-use-artificial-intelligence-gauging-risk-recidivism/) [<https://perma.cc/4ATB-PQ4D>].

<sup>142</sup> See Kaminski, *supra* note 108, at 2.

<sup>143</sup> Jessica M. Eaglin, *When Critical Race Theory Enters the Law & Technology Frame*, 26 MICH. J. RACE & L. 151, 155–58 (2021).

<sup>144</sup> ARVIND NARAYANAN & SAYASH KAPOOR, AI SNAKE OIL (2024).

system.<sup>145</sup> Likewise, as detailed in previous sections, decisions to delegate increasing agency operations to machines may carry significant institutional costs. In the regulatory enforcement context, this may include hollowing out agency expertise and discretion, undermining popular perceptions of legitimacy.

The need to recognize failure is not unique to applications involving regulatory enforcement, though it carries particular salience in a regulatory enforcement context due to the severe and direct impact that enforcement decisions have on their subjects, as well as the long and costly process by which these decisions may be challenged. If an agency waits for a Supreme Court finding that an AI-enabled component of their enforcement process was unreliable or otherwise problematic, the decision may taint years of other enforcement efforts that were based on the same operational paradigm.

As a consequence, any agencies deploying AI technologies, particularly where planned applications include regulatory enforcement, should incorporate retirement strategies into their initial assessment and planning processes. Where an AI system has been significantly incorporated into agency functions, a sudden discontinuation can be highly disruptive. Responsible exit ramps are an essential component of a robust risk mitigation strategy. Where these are in place, they may help to mitigate the impact of sunk cost fallacy on risk assessments, by ensuring that phasing a system out remains a practical response to poor performance.

The challenge in developing a meaningful standard of review which allows for failure is illustrative of a more foundational problem with using risk regulation as the lodestar of the oversight system. Among the main challenges permeating risk-centric frameworks for assessment is that, at their core, their focus is on developing a structure which will allow a proposed use case to move forward. The obvious gap in this approach is that there will inevitably be cases where incorporating AI is inappropriate or ill-advised regardless of how well it performs. Risks to the democratic legitimacy or public trust in an agency may be impossible to mitigate through more rigorous auditing or ensuring that humans remain at key points in the decision-making process. It is worth noting that the AI RMF presents particular challenges on this front insofar as its focus is on mitigating harms to institutions, as opposed to risks to the public or to democracy more broadly.<sup>146</sup>

This is not to suggest that risk-based approaches should be abandoned, though it does demonstrate the importance of ensuring that the assessments are designed in a way that allows for the conclusion that a system under review should be phased out, and that agencies have an effective plan for the system's retirement. It also shows a need to diversify oversight structures through additional mechanisms such as robust engagement with members of the public, including recognizing greater rights to contest decisions and pursue other effective remedies.<sup>147</sup>

#### **D. Structural Oversight Considerations**

---

<sup>145</sup> Sakshi Gupta, *When Should You Not Use Machine Learning?*, SPRINGBOARD (Sept. 25, 2020), <https://www.springboard.com/blog/data-science/when-not-to-use-ml/> [<https://perma.cc/XUG6-6VY2>].

<sup>146</sup> Kaminski, *supra* note 108, at 16.

<sup>147</sup> See Committee of Ministers, Council of Europe, *Recommendation CM/Rec (2020)1 of the Committee of Ministers to Member States on the Human Rights Impacts of Algorithmic Systems*, Doc. No. CM/Rec(2020)1 (April 2020), <https://rm.coe.int/09000016809e1154> [<https://perma.cc/8WU9-Q89P>].

Although virtually every AI governance framework delegates significant responsibilities to the frontline institutions that develop or deploy these systems, there is always a need for a central organization, or multiple organizations, to play a coordinating and oversight role. One may understand these coordinating functions as lying on a spectrum, from a highly devolved system, where the central agency is little more than a clearinghouse for documentation and reporting, to a more rigorous oversight structure which exercises relatively stringent control over how administrative bodies experiment with AI.

As a baseline, it should be relatively uncontroversial to note the value of a central repository which publicly tracks AI use cases across the administrative state and provides public information about where AI has been deployed, and access to relevant background material such as the results of risk assessments or audits. This appears to be the direction that OMB's AI policy memo is pushing agencies toward, in requiring that enhanced reporting and convening take place.

Beyond a limited publicly-facing function, there is additional value in maintaining a central hub for expertise in AI policy. This may include developing and applying best practices for things like risk assessment, transparency, and external consultation processes. There is a clear need for conceptual work to bolster the baseline standards enumerated in documents like the AI RMF and the Blueprint for an AI Bill of Rights. While aspects of the risk assessment process need to be carried out locally, it clearly does not make sense for every agency to be starting from scratch or working from its own independently developed definitions and benchmarks for things like “bias”, “discrimination”, and “fairness”.<sup>148</sup> Similar performance-based metrics are a core feature of environmental regulation, among other fields, and as AI governance matures it is essential to develop similar standards which may be implemented as common measures for performance.<sup>149</sup> Relatedly, the federal government will inevitably need to devote additional resources to providing technical support for agencies that are interested in piloting new AI projects, but which lack the resources to build them. These functions do not all need to be consolidated in a single agency, though there are efficiencies in housing them together.

At the further edge of the spectrum, one may envision an agency that performs a more significant oversight function, and even an enforcement role over how AI is being deployed across administrative agencies. There have been numerous scholarly suggestions for more proactive governance structures, though these typically focus on AI as a whole, as opposed to purely regulating public sector or regulatory enforcement functions. These models include Ryan Calo's 2015 proposal for a federal robotics commission, and Andrew Tutt's proposal for an “FDA for algorithms” that would exercise oversight over all such products before they are marketed.<sup>150</sup> Gianclaudio Malgieri and Frank Pasquale suggested an ex-ante licensing regime targeted at certain highly impactful functions, which would presumably include regulatory enforcement.<sup>151</sup> Imposing pre-market licensing requirements across the private sector may require grappling with certain

---

<sup>148</sup> ENGSTROM ET AL., *supra* note 1 at 6–8.

<sup>149</sup> See Lauren E. Willis, *Performance-Based Consumer Law*, 82 U. CHI. L. REV. 1309 (2015).

<sup>150</sup> Andrew Tutt, *An FDA for Algorithms*, 69 ADMIN. L. REV. 83 (2017); Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 CALIF. L. REV. 513, 556 (2015).

<sup>151</sup> Gianclaudio Malgieri & Frank Pasquale, *From Transparency to Justification: Toward Ex Ante Accountability for AI*, (Brussels Priv. Hub Working Paper No. 33, 2022), [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4099657](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4099657).

constitutional challenges, though the issue is simpler if licensing is considered purely with regard to administrative regulatory enforcement.<sup>152</sup>

Even if licensing goes too far, the potential for high impact failures suggests that a cautious approach which preserves the human-centric nature of government should be adopted. Ryan Calo and Danielle Citron propose that a key guide to assessing the wisdom of incorporating an AI system should be whether the use of AI furthers key substantive commitments and values, such as access, quality, and self-assessment.<sup>153</sup> Whether a particular regulatory enforcement use case is designed to further these values, as opposed to simply reducing costs, is the critical question: will the proposed tools enhance the capabilities of implementing agency, or is it mainly a means to outsource human discretion and expertise to a cheaper and inexhaustible automated decision-maker? Engstrom et al included a similar call for human centered AI in their article, advocating for decision tools which complement, rather than replace, the human element in the process.<sup>154</sup>

The expertise and discretion of administrative agencies, and the broader common interest in maintaining perceptions of trust and legitimacy, both speak to the value of an oversight structure which can work towards the safe and responsible use of AI in regulatory enforcement, rather than acting as merely a clearinghouse of public information.

- **Defining Effective Oversight**

It is possible to identify several features that will be essential for an effective AI oversight body for administrative agency functions, including regulatory enforcement. First, and most importantly, it should possess adequate subject matter expertise related to AI, as well as broader issues within its remit, such as laws and standards around discrimination, procedural fairness, etc. To achieve this, the agency will need to be adequately funded, and staffed with a diversity of experts from different backgrounds, including law, public policy, computer science and engineering, etc. The latter is particularly important if the oversight body is going to play a technical supporting function for agencies which are seeking to develop new AI systems. Though, it is important that subject matter experts from computer science and engineering be complemented by staff with a background in law, public policy, and the humanities, to ensure that there are robust conversations about ethics and public policy which transcend pure performance-based audits.

There are a number of existing agencies which possess robust technical expertise, including the Office of Science and Technology Policy (OSTP) and the General Services Administration AI Center of Excellence, as well as NIST. The latter has been an early mover in the AI space, especially through the AI RMF discussed throughout this report, though there have been concerns raised about their efficacy in performing a direct oversight function.<sup>155</sup> There are also concerns that an expanded role for NIST into a politically contentious realm, like AI in regulatory

---

<sup>152</sup> See, e.g., Cass R. Sunstein, *Artificial Intelligence and the First Amendment* (Apr. 27, 2023) (unpublished manuscript), [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4431251](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4431251).

<sup>153</sup> Calo & Citron, *supra* note 3, at 840.

<sup>154</sup> ENGSTROM ET AL., *supra* note 1 at 83.

<sup>155</sup> See Bryan H. Choi, *NIST's Software Un-Standards*, LAWFARE (Mar. 7, 2024), <https://www.lawfaremedia.org/article/nist-s-software-un-standards> [<https://perma.cc/EGQ5-C5SH>].

enforcement, may jeopardize its other vital roles related to standard setting and technological advancement.<sup>156</sup>

A second important quality is that the agency should be capable of public engagement, both in terms of a robust network of collaborators across academia and civil society, and as a key vector for transparency reporting. Rather than merely providing a clearinghouse of information, a strong coordinating body should have a mandate to promote engagement with the administrative state, while also cultivating awareness of opportunities for civil society and the public to get involved. The Privacy and Civil Liberties Oversight Board (PCLOB) is one model for an agency with a robust mandate to engage with public feedback, and there have been proposals to create a parallel body for AI applications, although several factors have limited PCLOB's record as an effective oversight body.<sup>157</sup>

Effective oversight can be challenging if the responsible agency is not sufficiently independent to enable it to push back against prevailing political priorities to cut costs, and to effectively advocate in support of longer-term institutional and public interest priorities.<sup>158</sup> Ideally, an oversight agency can be empowered with some sort of mechanism to compel compliance with its standards and recommendations. There can be a tension between independence and effective enforcement powers, where arm's length oversight bodies are more likely prominently flag structural problems, but also more likely to encounter institutional resistance in attempting to address them.<sup>159</sup> An agency which is placed outside of the executive branch has the advantage of greater independence, but may be less equipped to push changes directly for executive branch agencies.<sup>160</sup> Mounting skepticism from the judiciary against agency independence from the executive is a further driver of uncertainty. Another important consideration is resilience to regulatory capture, a challenge which manifests in different ways across the different branches of government.

Although independent oversight bodies are relatively common internationally, it is difficult to find parallel examples in the executive branch where an agency maintains sufficient binding oversight powers to be effective in the face of concerted resistance.<sup>161</sup> For example, the Office of Information and Regulatory Affairs (OIRA) plays a role in supporting compliance with the Regulatory Flexibility Act and Congressional Review Act, including through mandated impact analyses of significant regulatory actions. However, it is up to the agencies to determine whether a particular use case for AI is high risk enough to warrant review.

---

<sup>156</sup> Keegan McBride, *Regulating Artificial Intelligence Must Not Undermine NIST's Integrity*, TECH POL'Y PRESS (Aug. 20, 2024), <https://www.techpolicy.press/regulating-artificial-intelligence-must-not-undermine-nist-integrity/> [<https://perma.cc/BN3Y-R5ZJ>].

<sup>157</sup> See Faiza Paten & Patrick C. Toomey, *An Oversight Model for AI in National Security: The Privacy and Civil Liberties Oversight Board*, JUST SEC. (Apr. 26, 2024), <https://www.justsecurity.org/94999/an-oversight-model-for-ai-in-national-security-the-privacy-and-civil-liberties-oversight-board/> [<https://perma.cc/BY9Q-D5DJ>].

<sup>158</sup> See Karanicolas & Kwoka, *supra* note 74, at 689-692, for a discussion of these challenges in the context of the freedom of information system.

<sup>159</sup> *Id.* at 692-695.

<sup>160</sup> *Id.*; see Gordon Unzen, *Artificial Intelligence and the Administrative State: Regulating the Government Use of Decision-Making Technology*, 25 MINN. J.L. SCI. & TECH. 209, 253 (2023).

<sup>161</sup> See Karanicolas & Kwoka, *supra* note 74, at 679-688, for a discussion of the various independent oversight bodies that exist to enforce right to information legislation around the world.

The Federal Risk and Authorization Management Program (FedRAMP) was established in 2011 to facilitate the adoption and use of cloud services by the federal government, presents one interesting model. Congress codified the FedRAMP Authorization Act in December 2022,<sup>162</sup> Per an Office of Management and Budget memorandum, all cloud services that hold federal data must be FedRAMP authorized.<sup>163</sup> In response to EO 14110, FedRAMP established an “Emerging Technology Prioritization Framework,” which, as the name suggests, establishes a framework for prioritizing emerging technologies (ETs) for FedRAMP authorization.<sup>164</sup> The initial categories of ETs include generative AI.<sup>165</sup> There are two types of FedRAMP authorizations: (1) Joint Administrative Board (JAB) or (2) an authorization directly through a federal agency.<sup>166</sup> For both JAB and agency authorization, a requirement is “continuous monitoring,” which includes an annual security assessment and must generally meet the requirements described in the Authorization Playbook.<sup>167</sup>

It is unclear whether the FedRAMP system is equipped to provide comprehensive oversight over agency use of AI, or whether oversight responsibilities should be delegated to another existing agency, or handed over to an entirely new statutory creation.<sup>168</sup> At the end of the day, the likely scale of impact that AI technologies will have across the administrative state, and the importance of ensuring that transitions to greater automation are carefully managed in line with the public interest, suggest that the creation of a standalone oversight body may be justified. However, the question of how to construct effective oversight of the use of AI for regulatory enforcement includes complicated political calculations which go beyond the scope of this paper. Congress has been struggling to pass an update to federal privacy legislation for decades now.<sup>169</sup> Given the pace with which AI is moving, America cannot afford a similarly painful and drawn-out process for effective regulation of government uses of this technology. While a new standalone oversight body may be the best option in the longer term, it is important to focus on solutions which are politically workable in the short term, to avoid a future where AI becomes deeply embedded across sensitive government functions without any meaningful check on its deployment.

---

<sup>162</sup> FedRAMP Authorization Act, H.R. 8956, 117th Congress (2022), <https://www.congress.gov/bill/117th-congress/house-bill/8956/text>.

<sup>163</sup> Shalanda D. Young, Director of OMB, *Memorandum for the Heads of Executive Departments and Agencies*, (July 25, 20224), [https://www.fedramp.gov/assets/resources/documents/FedRAMP\\_Policy\\_Memo.pdf](https://www.fedramp.gov/assets/resources/documents/FedRAMP_Policy_Memo.pdf).

<sup>164</sup> See FedRAMP, *Emerging Technology Prioritization Framework*, <https://www.fedramp.gov/et-framework/>.

<sup>165</sup> *Id.*

<sup>166</sup> FedRAMP, *FedRAMP CSP Authorization Playbook* at 2, [https://www.fedramp.gov/assets/resources/documents/CSP\\_Authorization\\_Playbook.pdf](https://www.fedramp.gov/assets/resources/documents/CSP_Authorization_Playbook.pdf).

<sup>167</sup> *Id.* at 12.

<sup>168</sup> Aram A. Gavor & Raffi Teperdjian, *A Structural Solution to Mitigating Artificial Intelligence Bias in Administrative Agencies*, 89 GEO. WASH. L. REV. ARGUENDO 71, 87 (2021).

<sup>169</sup> See Jessica Rich, *After 20 Years of Debate, It's Time for Congress to Finally Pass a Baseline Privacy Law*, BROOKINGS (Jan. 14, 2021), <https://www.brookings.edu/articles/after-20-years-of-debate-its-time-for-congress-to-finally-pass-a-baseline-privacy-law/> [<https://perma.cc/ZY4Z-G44J>].

## V. Conclusion: Grappling with the Human-Machine Paradigm

Beyond considering the pros and cons of each individual use case, among the most important conceptual questions that agencies need to grapple with over the coming decade concerns the interplay between humans and automated systems, and the appropriate role and limits of automation. The results of this process are likely to shape the administrative state for the next generation.

Addressing this challenge requires a long-term view of agency priorities and values, including balancing the pressure to cut budgets and improve capacity against the need to safeguard perceptions of legitimacy, and the human-centric nature of the administrative state. There are also more subtle impacts that are at play, such as the broader datafication of agency decision-making. AI systems tend to be data intensive, and rely on scale for efficiency, which results in centralizing information flows in order to maximize their utility.<sup>170</sup> This, in turn, is likely to prompt a more concentrated agency structure, which creates its own knock on effects related to how the public interacts with the agency and how decisions are processed. There are no easy answers, and the questions are likely to become more difficult and complicated as technology continues to advance.

As a starting point, it is useful for agencies to isolate and identify areas of governance where “humanity” is of particular importance, such as policing, or other areas where datafication is viewed as fundamentally or intuitively problematic. It is also important to think strategically about the relationship between humans and machines, and the interplay between these two. This can include the need to manage data being collected now in a way that leaves the door open to future automation, but also to figure out impacts that a reduction of staff in favor of AI will have from a human resources perspective, both in terms of likely reactions from the current workforce and the resulting risk to human capacity and expertise within those agencies.

An additional consideration for agencies concerns the need to grapple with apportioning an appropriate role for technical experts in the development and deployment of these technologies, and assessing their performance, without losing sight of the underlying values that are meant to be guiding their work. Given the scientific and mathematical nature of assessments around qualities like fairness and bias, it is easy to forget that both values require a fundamentally human element to ensure that they are meaningful.<sup>171</sup>

From this perspective, one final consideration concerns a persistent skill shortage across the government in terms of AI capacity, and the need for budgets and salaries to reflect the competition for technical expertise in these fields. In their previous publication for ACUS, Engstrom et al relay a story about early government agencies seeking to experiment with AI who, due to hiring rules, were forced to find lawyers with a coding background to develop software for them.<sup>172</sup> While agencies over the coming years will likely enjoy a freer hand to hire technical experts to support their operations, there are likely to be lingering challenges around developing an appropriate pipeline to onboard STEM graduates into public service, how to integrate these new streams

---

<sup>170</sup> Unzen *supra* note 160, at 234–35.

<sup>171</sup> See Deborah Hellman, *Measuring Algorithmic Fairness*, 106 VA. L. REV. 811, 834 (2020).

<sup>172</sup> ENGSTROM ET AL., *supra* note 1 at 44.



alongside traditional policymaking staff, and how to ensure that the federal government is able to compete with the private sector in attracting skilled AI experts.<sup>173</sup>

Administrative agencies face a rapidly changing and dynamic environment, as political, judicial and technological changes combine to fundamentally reshape their role in the constitutional order. Although there are no easy answers, it is possible for agency leadership to cleave to fundamental values, and their core mission, as they attempt to navigate a course which harnesses the benefits of AI while preserving the essential humanity of democratic governance.

---

<sup>173</sup> Jory Heckman, *White House says agencies hired 200 AI experts so far through governmentwide 'talent surge'*, FEDERAL NEWS NETWORK (Jul. 29, 2024), <https://federalnewsnetwork.com/artificial-intelligence/2024/07/white-house-says-agencies-hired-200-ai-experts-so-far-through-governmentwide-talent-surge/> [<https://perma.cc/SVS6-QPJS>].

## VI. Recommendations

### Recommendations for Implementing Agencies:

- While risk assessments are a valuable tool, implementing agencies should understand their limitations in the realm of regulatory enforcement, including by factoring in the tendency of AI systems to produce unpredictable outcomes, and to expand beyond their initial approved uses. They should also ensure that worst case scenarios are considered, even if their likelihood is relatively remote.
- Risk assessment processes should think beyond immediate harms, and include thorough consideration of the potential for an AI system to undermine public confidence in an agency or perceptions of legitimacy.
- In a risk assessment process over the use of AI in regulatory enforcement, the influence of the system over the final decision is a major risk factor. This level of influence should be understood as a spectrum, rather than a binary question that hinges on the presence of human review.
- Risk assessment processes should be complemented by a careful and critical assessment of the costs and benefits of pursuing an AI-based solution, and systems which are failing to perform to an acceptable standard should be shelved. Agencies should be wary of sunk cost fallacies in considering the worthiness or performance of a system and should be cautious of overly ambitious or optimistic assessments of AI capabilities.
- Early stage risk mitigation should include appropriate exit ramps, planning for a system to be retired if it fails to perform. This planning should be updated throughout the life-cycle of the product.
- Implementing agencies should institute robust external consultation processes related to any uses of AI in regulatory enforcement, which span the full lifecycle of the decision-making process, and which include participatory forums and the establishment of co-decision-making bodies with leading civil society and other relevant participants. Key goals for this engagement should include identifying and quantifying potential or manifested harms, and generating buy-in from impacted communities. Strong engagement will depend on agencies' ability to offer meaningful opportunities to impact policy.
- Robust consultation depends on robust transparency, so that external stakeholders can get a complete picture of how AI systems are being implemented and how they are performing. At a minimum, this requires publishing risk assessments and related documentation. At a more advanced level, resources may be required to translate reporting into more accessible formats, or to upskill community partners to engage with more advanced questions.
- In assessing the appropriate scope for disclosure of information, agencies should err on the side of transparency, and only withhold information which would clearly harm the effective administration of justice, and where the harm from this disclosure would outweigh the public interest in its release.

### **Structural Recommendations for the Federal Administrative Apparatus:**

- Responsible use of AI for regulatory enforcement may require that risk assessments be complemented by stronger measures to prevent or mitigate particularly severe harms. These may include prohibitions against particularly problematic use cases, such as where fundamental rights are impacted or particularly sensitive data is involved. Agencies should consider the role of individual remedies in addressing specific harms, as well as their strengths and weaknesses as a vehicle for structural change.
- Effective oversight of AI across the administrative state, and particularly with regards to regulatory enforcement, suggests that there is value in centralized oversight related to AI development and risk assessment. The oversight structure should be staffed with a diverse range of experts across multiple disciplines, and should be equipped with sufficient independence and enforcement powers to play an effective role.