The Automated State: A Realist View

*David Freeman Engstrom**

Abstract

Government use of artificial intelligence ("AI") to make, implement, and enforce law is fueling anxieties among a growing cast of critics. Some are accelerations of concerns raised by other technology adoptions: error, bias, gaming, and the oversight challenges that come with reliance on procurement. Others are more AI-specific: the power of machine learning and big data to draw privacy-violating inferences, or AI's "black box" opacity and inability to engage in democratic reason-giving. The reforms AI's critics demand in response range from the sharp and specific (e.g., outright prohibitions and full-blown, Food and Drug Administration-like licensure) to the gauzy and unproven (e.g., calls to "democratize AI" via participatory design or impact assessments).

This Essay sketches an alternative, "realist" view in three parts. First, policymakers are unlikely to enact the bespoke new regulation of the sort critics are demanding because of AI's twin capacity to cause and cure error, bias, and inequity. For every horror story there will be a success story—a way new tech makes government's work more efficient, accurate, rule-of-law-respecting, and equitable. Second, public law has always been radically limited in reach relative to critics' demands. Indeed, even progressive commentators have long warned that the problem with government may not be too little transparency, but too much, and new procedural burdens, however well-meaning, can stultify a government that already lacks vigor. A final insight follows: the urgent task ahead may be both more pedestrian and more ambitious than AI's critics suggest. Indeed, if algorithmic accountability will be litigated rather than legislated, then efforts should focus more than they have on legal adaptation—that is, the tailoring of existing legal frameworks, particularly ordinary administrative law, to the government's new algorithmic toolkit. Wise adaptation, not the blue-sky regulatory overhauls that occupy much of the scholarly literature, should be the order of the day. This Essay makes a start down that road.

* LSVF Professor in Law and Co-Director, Deborah L. Rhode Center on the Legal Professions, Stanford Law School; dfengstrom@law.stanford.edu. This Essay is an expanded version of a keynote address first given at HEC Paris, France in December 2021.

Table of Contents

Introduction

In 2020, an African American man named Robert Williams became a symbol of artificial intelligence ("AI") run amok when a facial recognition system wrongly matched his face to a criminal suspect.[1] Williams spent thirty hours in a Detroit jail before police realized the mistake.[2] Stories like this did not prevent a company named Clearview Technology from scraping millions of social media pictures and then successfully marketing a facial recognition system to police departments—now reportedly in use in at least 600 jurisdictions.[3] Then there are the families in Arkansas who lost their disability benefits when a crude algorithm wrongly declared them ineligible,[4] or thousands in Michigan who lost their unemployment benefits when an equally crude algorithm called MiDAS wrongly flagged them for benefits fraud with only a few days to respond to the automated accusation.[5] And, from European shores comes the tale of an automated system that wrongly accused 20,000 Dutch families of child support benefits fraud, a grave breach of public trust that reportedly ended a government.[6]

These revelations about the abuses of algorithmic governments are deeply worrying, and they have fueled righteous anger among academics and activists. The critique of a growing cast of critics is long

1 Kashmir Hill, *Wrongfully Accused by an Algorithm*, N.Y. Times (Aug. 3, 2020), https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html [https://perma.cc/G738-B3BX].

2 *Id.*

3 Kashmir Hill, *The Secretive Company That Might End Privacy as We Know It*, N.Y. Times (Nov. 2, 2021), https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html [https://perma.cc/5E5M-L8SM].

4 Colin Lecher, *What Happens When an Algorithm Cuts Your Health Care*, Verge (Mar. 21, 2018, 9:00 AM), https://www.theverge.com/2018/3/21/17144260/healthcare-medicaid-algorithm-arkansas-cerebral-palsy [https://perma.cc/Y3L9-UDZV].

5 *See* Robert N. Charette, *Michigan's MiDAS Unemployment System: Algorithm Alchemy Created Lead, Not Gold*, IEEE Spectrum (Jan. 24, 2018), https://spectrum.ieee.org/michigans-midas-unemployment-system-algorithm-alchemy-that-created-lead-not-gold [https://perma.cc/49LW-C7YA].

6 Jon Henley, *Dutch Government Resigns over Child Benefits Scandal*, Guardian (Jan. 15, 2021, 9:32 AM), https://www.theguardian.com/world/2021/jan/15/dutch-government-resigns-over-child-benefits-scandal [https://perma.cc/SNG4-JRK6].

and withering. AI's agonistes warn of a *secret state* when government uses opaque, "black box," and very often proprietary, algorithmic systems that cannot meet basic reason-giving demands when taking actions that impair rights.[7] They warn of a *regressive* and a *racist state*, because government use of algorithms to surveil and control tends to fall most heavily on the poor and people of color and reproduces or even exacerbates structural inequalities.[8] They warn of a "*myopic*" *state*, its "line of sight" altered by creeping dataism, that uses the power of machine learning and big data to reactively manage the effects of policy problems rather than addressing their root causes.[9] They warn of a *relentless state*, its digitized exercise of coercive power too perfect and uncompromising, leaving too little space for human empathy[10] or democracy's deliberative froth and friction.[11] And, finally, they warn of a *repressive state* that crushes human autonomy by treating citizens not as individuals but as embodiments of their data—a digital violation of the Kantian imperative to treat individuals as ends not means.[12]

The reforms AI's agonistes demand in response range from the specific and slashing to the gauzy and unproven. In the first category are what Professor Pasquale has called "first-wave" algorithmic accountability mechanisms[13]: prohibitions on particular techniques or uses;[14] a

---

7 *See* Frank Pasquale, The Black Box Society: The Secret Algorithms That Control Money and Information 3 (2015); Mike Ananny & Kate Crawford, *Seeing Without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability*, 20 New Media & Soc'y 973, 974 (2018); Jenna Burrell, *How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms*, Big Data & Soc'y, Jan.–June 2016, at 1, 3; Hannah Bloch-Wehba, *Access to Algorithms*, 88 Fordham L. Rev. 1265, 1265 (2020).

8 *See* Virginia Eubanks, Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor 6–7 (2018); Ruha Benjamin, Race After Technology: Abolitionist Tools for the New Jim Code 53 (2019); Cathy O'Neil, Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy 22–23 (2016).

9 *See* Marion Fourcade & Jeffrey Gordon, *Learning Like a State: Statecraft in the Digital Age*, 1 U.C. Davis J.L. & Pol. Econ. 78, 87 (2020) (quoting Louise Amoore, *Lines of Sight: On the Visualization of Unknown Features*, 13 Citizenship Stud. 17, 17 (2009)); Karen Yeung, *Algorithmic Regulation: A Critical Interrogation*, 12 Regul. & Governance 505, 508 (2018); Eubanks, *supra* note 8, at 178.

10 *See* Sofia Ranchordás, *Empathy in the Digital Administrative State*, 71 Duke L.J. 1341 (2022).

11 *See* Julie E. Cohen, Configuring the Networked Self: Law, Code, and the Play of Everyday Practice 34–35 (2012).

12 *See* Colin Koopman, How We Became Our Data: A Genealogy of the Informational Person 8 (2019).

13 Frank Pasquale, *The Second Wave of Algorithmic Accountability*, L. & Pol. Econ. Project (Nov. 25, 2019), https://lpeproject.org/blog/the-second-wave-of-algorithmic-accountability/ [https://perma.cc/N5AK-JQ8D].

14 David Freeman Engstrom & Amit Haim, *Regulating Government AI and the Challenge of Sociotechnical Design*, 19 Ann. Rev. L. & Soc. Sci. 277, 278 (2023).

pre-deployment licensure or certification scheme;[15] rules prescribing how agencies design and deploy algorithmic systems, such as the requirement that a human provide "meaningful oversight";[16] and "mandat[ory] open sourcing of code [or] data"[17] and extensive rights for data subjects to challenge automated decisions[18] to achieve something like "decision-level" transparency[19] whenever government uses an algorithmic tool.[20] The other category of fixes is different. It encompasses sundry calls to "democratize AI" via an often-vague menu of "participatory design,"[21] "impact assessments,"[22] and "algorithmic audits,"[23] despite questions about the efficacy of each.[24]

This Essay sketches an alternative view, or at least a more complicated one. Call it a realist view—though some will label it an apologist's view, or, perhaps worse, a lawyer's view. Whatever label attaches, the argument comes in three parts.

First, we are unlikely to get the robust regulation the agonistes want because government automation will always be fundamentally

---

15 David Freeman Engstrom & Daniel E. Ho, *Algorithmic Accountability in the Administrative State*, 37 Yale J. on Regul. 800, 826 (2020).

16 *See* Ben Green, *The Flaws of Policies Requiring Human Oversight of Government Algorithms*, Comput. L. & Sec. Rev., July 2022, at 1, 8; *see also Commission 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*, at 9, COM (2021) 206 final (Apr. 21, 2021).

17 Engstrom & Haim, *supra* note 14, at 278.

18 Engstrom & Ho, *supra* note 15, at 826.

19 *Id.* at 825.

20 *See generally* Ada Lovelace Inst., AI Now Inst. & Open Gov't P'ship, Algorithmic Accountability for the Public Sector (2021), https://www.opengovpartnership.org/wp-content/uploads/2021/08/algorithmic-accountability-public-sector.pdf [https://perma.cc/YJL8-RABQ] (analyzing algorithmic accountability policies).

21 *See* Min Kyung Lee et al., *WeBuildAI: Participatory Framework for Algorithmic Governance*, Proc. ACM on Hum.-Comput. Interaction, Nov. 2019, at 1, 2, 5; Angela Zhou, David Madras, Inioluwa Deborah Raji, Smitha Milli & Richard Zemel, *Participatory Approaches to Machine Learning*, Int'l Conf. on Mach. Learning (July 17, 2020), https://participatoryml.github.io/ [https://perma.cc/NYP2-UVVA].

22 *See, e.g.*, Dillon Reisman, Jason Schultz, Kate Crawford & Meredith Whittaker, Algorithmic Impact Assessments: A Practical Framework for Public Agency Accountability (2018); Kate Crawford & Ryan Calo, *There Is a Blind Spot in AI Research*, 538 Nature 311, 311 (2016).

23 Inioluwa Deborah Raji, Peggy Xu, Colleen Honigsberg & Daniel Ho, *Outsider Oversight: Designing a Third Party Audit Ecosystem for AI Governance*, 2022 Proc. AAAI/ACM Conf. on A.I., Ethics, & Soc'y 557, 557.

24 *See id.* at 566; Jacob Metcalf, Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh & Madeleine Clare Elish, *Algorithmic Impact Assessments and Accountability: The Co-Construction of Impacts*, 2021 Proc. ACM Conf. on Fairness, Accountability, & Transparency 735, 735; Mona Sloane, Emanuel Moss, Olaitan Awomolo & Laura Forlano, *Participation Is Not a Design Fix for Machine Learning*, 2022 Proc. 2nd ACM Conf. on Equity & Access Algorithms, Mechanisms, & Optimization 1, 1.

ambiguous in at least two senses. To start, it will be ambiguous in terms of what the thing to be regulated is. Beyond well-bounded AI technologies, such as facial recognition technology ("FRT"),[25] or specific uses, such as criminal risk assessment for bail, sentencing, or parole decisions, statute-drafting challenges abound.[26] What is covered and what is not? Narrowly drawn definitions targeting sophisticated machine learning applications miss highly concerning forms of automation, many of them low-tech tools that government has used for decades.[27] But, if drawn too broadly, new regulations encompass virtually all of what modern government does.[28]

More important is the brute fact that government automation will always be fundamentally ambiguous in its effects. For every horror story like Williams's, there will be a success story—a way the government's new digital toolkit makes its work more efficient, more accurate, more rule-of-law-respecting, and more equitable.[29] Indeed, most debates around government AI bottom out at the decades-old question of the merits and demerits of clinical versus actuarial judgment.[30] Unless a sudden parade of outrages opens stubborn policy windows in a politically polarized and paralyzed time, algorithmic accountability for the overwhelming majority of government uses of AI is likely to be litigated under existing laws, not legislated via bespoke new ones.[31]

Second, a host of legal complexities picks up where automation's dual ambiguities leave off. Public law has always been radically limited in reach relative to agonist demands. In the United States, United Kingdom, and elsewhere, law has *never* demanded anything approaching

---

25 *See* David Freeman Engstrom & Alicia Solow-Niederman, Federalism and the Automated State 4 (2023) (unpublished article) (on file with authors).

26 *See infra* notes 79–104 and accompanying text.

27 *Cf.* Herbert A. Simon, Administrative Behavior 91, 119 (4th ed. 1997).

28 *See infra* note 81 and accompanying text; *see also* Karen Levy, Kyla E. Chaslow & Sarah Riley, *Algorithms and Decision-Making in the Public Sector*, 17 Ann. Rev. L. & Soc. Sci. 309, 311 (2021) (explaining that "[d]efinitional decisions are more than semantic squabbles," as they "determine which systems are subject to regulation").

29 *See* Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan & Cass R. Sunstein, *Discrimination in the Age of Algorithms*, 10 J. Legal Analysis 113, 163 (2018); Matthew M. Young, Justin B. Bullock & Jesse D. Lecy, *Artificial Discretion as a Tool of Governance: A Framework for Understanding the Impact of Artificial Intelligence on Public Administration*, 2 Persps. on Pub. Mgmt. & Governance 301, 301 (2019).

30 *See* Robyn M. Dawes, David Faust & Paul E. Meehl, *Clinical Versus Actuarial Judgment*, 243 Sci. 1668 (1989); William M. Grove & Paul E. Meehl, *Comparative Efficiency of Informal (Subjective, Impressionistic) and Formal (Mechanical, Algorithmic) Prediction Procedures: The Clinical-Statistical Controversy*, 2 Psych., Pub. Pol'y, & L. 293 (1996); Jon Kleinberg, Jens Ludwig, Sendhill Mullainathan & Ziad Obermeyer, *Prediction Policy Problems*, 105 Am. Econ. Rev.: Papers & Proc. 491 (2015).

31 *See infra* note 80 and accompanying text (describing the relatively few and anemic state-level enactments regarding government use of AI).

full transparency or accountability from government.[32] Indeed, much of American administrative law—the body of constitutional and statutory law that governs how government agencies operate—offers only shallow review and a thin hedge against arbitrary agency action.[33] Moreover, heightened scrutiny brings steep trade-offs. It was not long ago that progressive voices called for *less* government transparency, not more, because of concern about stultifying an already risk-averse government.[34] More recent voices warn that even well-intentioned transparency measures and procedural burdens can drain the government of vigor and even play into the hands of its neoliberal opponents.[35]

Training a wider-angle lens on the long history of government automation reveals a wider landscape and still sharper trade-offs. Two macrotrends dominate the story of government automation in recent decades. The first is that bigger data and better analytics have steadily moved automation toward the center of the state's coercive and redistributive powers—toward more and more consequential tasks—ranging from the selection of enforcement targets to the adjudication of valuable rights and privileges. Put another way, automation has entered *gray zones*—highly consequential modes of state action not previously thought tractable to automation.[36]

---

32 *See* Wendy Wagner & Martin Murillo, *Is the Administrative State Ready for Big Data?*, Knight First Amend. Inst. (Apr. 30, 2021), https://knightcolumbia.org/content/is-the-administrative-state-ready-for-big-data [https://perma.cc/Y3CE-ARZL]. *See generally* Wendy Wagner with Will Walker, Incomprehensible! A Study of How Our Legal System Encourages Incomprehensibility, Why It Matters, and What We Can Do About It 3 (2019) (describing how society accepts excessive and incomprehensible information as a necessary feature of the contemporary world).

33 *See generally* Adrian Vermeule, Law's Abnegation: From Law's Empire to the Administrative State (2016) (noting the increasing prevalence, among judges and lawyers, of the view that administrators should have broad discretion, and that there is little constraint in administrative law against the exercise of such discretion).

34 *See* William J. Stuntz, *Against Privacy and Transparency*, New Republic (Apr. 17, 2006), https://newrepublic.com/article/65393/against-privacy-and-transparency [https://perma.cc/VH6T-CVUC].

35 *See* David E. Pozen, *Transparency's Ideological Drift*, 128 Yale L.J. 100, 146–48 (2018). *See generally* Paul Sabin, Public Citizens: The Attack on Big Government and the Remaking of American Liberalism (2021) (describing the unintended consequences of the rise of the new liberal movement in the 1960s and 1970s that challenged the federal government and its officials because of the belief that insulated and unaccountable government officials would exploit existing measures of accountability and transparency in ways that would not serve the public interest); Nicholas Bagley, *The Procedure Fetish*, 118 Mich. L. Rev. 345 (2019) [hereinafter Bagley, *Procedural Fetish*] (arguing that procedures designed to cabin government action can prevent even progressive governments from achieving desired ends); Nicholas Bagley, *The Puzzling Presumption of Reviewability*, 127 Harv. L. Rev. 1285, 1287 (2014) [hereinafter Bagley, *Puzzling Presumption*] (same).

36 *See* Alicia Solow-Niederman, *Algorithmic Grey Holes*, 5 J.L. & Innovation 116, 122 (2023); Michael Veale & Irina Brass, *Administration by Algorithm? Public Management Meets*

The other macrotrend is that, as government automation has moved into governance gray zones, it has simultaneously pushed into legal *gray holes*—that is, places where law does not reach or does not reach well.[37] Some gray holes are practical. It is well-known that legal action focused on error correction—for instance, individual challenges to agency benefits decisions—does little to fix systemic problems.[38] Other gray holes, however, are deliberate. As an example, American law has long hived off enforcement decisions from legal review out of concern about generalist judges second-guessing agency priority-setting.[39] In short, a good deal of the opacity that hangs over an increasingly automated state is not technical. It is often legal, and it is an *intentional design choice*.

This all places demands for robust regulation of government AI in critical perspective. On the one hand, many proposals, particularly calls to "democratize AI,"[40] are often, on closer inspection, just gauzier versions of current administrative law, such as the requirement that agencies ventilate new policies via "notice and comment."[41] Other proposals raise the opposite problem. They would be a quantum leap in scrutiny of government action. Of course, for those worried about a state that is secret, regressive, racist, myopic, relentless, or repressive, that is good. For those who, instead, worry about a government that is already risk averse and even sclerotic and that is also easily captured by well-heeled "haves" who exploit accountability measures more readily than "have-nots," that is bad.[42]

Either way, government use of AI might be less a thing to be regulated and more a referendum on *all* of our democratic institutions.[43] That need not be a bad thing. AI might thus hold up a mirror to our democracy and make it better. But AI can also become a stalking horse for remaking much of public law, hamstringing government in the process, but without rigorous reckoning with the trade-offs that have

Public Sector Machine Learning, *in* ALGORITHMIC REGULATION 121 (Karen Yeung & Martin Lodge eds., 2019).

37 *See* Solow-Niederman, *supra* note 36, at 118; Amit Haim, *The Administrative State and Artificial Intelligence: Toward an Internal Law of Administrative Algorithms*, 14 U.C. IRVINE L. REV. 103, 144 (2024).

38 David Ames, Cassandra Handan-Nader, Daniel E. Ho & David Marcus, *Due Process and Mass Adjudication: Crisis and Reform*, 72 STAN. L. REV. 1, 23 (2020); *see* Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249, 1254 (2008).

39 *See, e.g.*, Heckler v. Chaney, 470 U.S. 821, 837–38 (1985) ("presumption that agency decisions not to institute proceedings are unreviewable"); Engstrom & Ho, *supra* note 15, at 832.

40 *See* Johannes Himmelreich, *Against "Democratizing AI,"* 38 A.I. & SOC'Y 1333, 1337 (2023).

41 *See* Wagner & Murillo, *supra* note 32, at 15, 34–35.

42 *See, e.g.*, SABIN, *supra* note 35; Bagley, *Procedural Fetish*, *supra* note 35, at 392; Bagley, *Puzzling Presumption*, *supra* note 35, at 1322–23; Wagner & Murillo, *supra* note 32, at 16–17, 22.

43 *See* Himmelreich, *supra* note 40, at 1333–34, 1337.

long occupied judges, policymakers, and those worried about a state, particularly an administrative state, that lacks vigor.[44]

A third and final insight follows: the most urgent task ahead may be both more pedestrian and more ambitious than AI's agonistes suggest. Indeed, if algorithmic accountability will be litigated, not legislated, then efforts should focus more than they have to this point on *legal adaptation*.

Adaptation will, first and foremost, require unflashy work at a basic descriptive level to understand how AI is reshaping public bureaucracies and how those bureaucracies are adopting, designing, and overseeing new algorithmic tools. More recycling of headline-grabbing, worst-cases—Michigan's decade-old MiDAS fiasco, perennially featured in critical scholarship, is an example—will not advance the ball. Instead, if critiques of the automated state are to land and drive real reform, we must get under the hood of a wider array of tools performing a wider array of governance tasks and then connect our findings to a fast-emerging field of inquiry at the intersection of computer science, social science, and organizational theory. Only once we know far more about how AI changes bureaucracies, and how bureaucracies incorporate AI into their routines, can we craft durable and wise legal constraints.

Legal adaptation will also mean finding creative ways for law to do what it has always struggled to do: induce public bureaucracies to think critically about upstream, systemic design and oversight of algorithmic systems throughout the AI lifecycle, not just prior to deployment or after they cause harm.[45]

The end result may well disappoint agonistes. Rather than hard-edged prohibitions, full transparency, or open-sourced code, the most likely regulatory outcomes will be measured and incremental adaptations to existing law, particularly administrative law, perhaps with efforts to boost agencies' internal technical capacity so they can more responsibly adopt and oversee new tools. But this might be better than the alternatives, and it might also be the best that can be hoped for. Indeed, we need to be realists—neither agonistes nor apologists. If we are not, we risk ceding the field at what is a critical moment in the evolution of the modern liberal-democratic state.

---

44 *See generally* Himmelreich, *supra* note 40. For a classic statement of liberal anxiety about government vigor, see Elena Kagan, *Presidential Administration*, 114 HARV. L. REV. 2245 (2001).

45 *See* Kurt Glaze, Daniel E. Ho, Gerald K. Ray & Christine Tsang, *Artificial Intelligence for Adjudication: The Social Security Administration and AI Governance*, *in* OXFORD HANDBOOK OF AI GOVERNANCE 779 (Justin B. Bullock et al. eds., 2022) ("[P]ublic sector AI innovation . . . requi[res] continuous analysis, evaluation, and iteration."); Evelyn Douek, *Content Moderation as Systems Thinking*, 136 HARV. L. REV. 526 (2022).

That, warts and all, is the beginnings of a realist view. From here, this Essay fleshes out the above claims in three parts. Part I bolsters the claim that the latest round of AI-fired government automation is fundamentally ambiguous. Part II turns to the deep legal challenges that have arisen while potent new forms of automation have pushed into (governance) gray zones and (legal) gray holes. Part III begins to ask what is to be done—and, in particular, what a new field of inquiry might look like in order to arm regulators and, perhaps more importantly, judges with the knowledge they will need to bring meaningful accountability to the automated state.

I. Automation's Ambiguities

A 1963 headline in the *Middletown, Ohio, Journal* blared: "Tax Cheaters Out of Luck: That Robot Will Getcha."[46] Taxpayers' "raw data," the article explained, was "punched on cards" in the Internal Revenue Service's ("IRS") Philadelphia office, "transfer[red] . . . to magnetic tape," flown to the National Computer Center in West Virginia, and then analyzed "against a master file."[47] The result was a list of taxpayers who owed back-taxes or whose data had thrown flags as "suspicious"—information that was recorded on another tape and flown back to Philadelphia.[48] Once there, the journalist breathlessly noted, a machine could "automatically write notices" to taxpayers requesting payment or announcing an audit.[49]

Early IRS automation efforts are a useful reminder that the current round of innovation is not the first time government has adopted new technologies to do its work.[50] Far from it. Nearly seventy-five years before the IRS's "robot," the U.S. Census Bureau used the Hollerith machine, an ingenious electromechanical punch card device, to tally the 1890 census.[51] From there, new technologies came in waves: more advanced tabulating systems in a growing New Deal state;[52] the uptake of "expert systems" beginning in the 1960s as computers made it possible to automate more complex bureaucratic tasks;[53] the spread of data analytics as part of the "performance management" and "reinventing government"

46 Raymond J. Crowley, *Tax Cheaters Out of Luck: That Robot Will Getcha*, Middletown, Ohio, J., Apr. 3, 1963, at 25.

47 *Id.*

48 *Id.*

49 *Id.*

50 *See generally id.*

51 *U.S. Census Bureau History: Herman Hollerith and Mechanical Tabulation*, U.S. Census Bureau (Jan. 5, 2023), https://www.census.gov/history/www/homepage_archive/2016/january_2016.html [https://perma.cc/N3TT-N5AK].

52 *See* Koopman, *supra* note 12, at 6, 16–17.

53 *See* Simon, *supra* note 27, at 134–35.

crusades of the 1980s and 1990s;[54] and the emergence of "e-government" initiatives in the 2000s, powered by rising digital connectivity that both remade state-citizen relations[55] and unlocked data previously housed in separate agencies and courthouses (while also sharpening privacy and cybersecurity concerns).[56] Taking a long view, it is possible to see the AI governance revolution as just the latest stop in a century of formalization, rationalization, datafication, and digitization—the "long computerization" of government.[57]

Nor is the current digital moment the first time the government's embrace of technology has fueled heated claims about its virtues or anxious hand-wringing about its perils. In the years after the IRS created its "robot" auditor, debate steadily mounted about government automation, pitting techno-optimists who imagined "a bureaucracy of almost celestial capacity"[58] against leaders of the "cybernation" scare who pressed now-familiar concerns about the brittleness of machine decisions, bias, labor displacement, and privacy in high-profile hearings before Congress.[59]

As computers flooded government in the decades that followed, concern about growing use of "computer matching" and "computer profiling"—the latter a proto-algorithmic use of prediction to target enforcement resources—produced more heated congressional inquiry and, quickened by the Watergate scandal's abuses of executive power, the Privacy Act of 1974.[60] At least some observers at the time, however,

---

54 *See generally* David E. Osborne & Ted Gaebler, Reinventing Government: How the Entrepreneurial Spirit Is Transforming the Public Sector (1992).

55 *See* Harold C. Relyea & Henry B. Hogue, *A Brief History of the Emergence of Digital Government in the United States*, *in* Digital Government: Principles & Best Practices 16, 30 (Alexei Pavlichev & G. David Garson eds., 2003).

56 *See generally* Tal Z. Zarsky, *Governmental Data Mining and Its Alternatives*, 116 Penn St. L. Rev. 285 (2011) (discussing privacy concerns related to the government's collection and use of personal data acquired through digital means); Fred H. Cate, *Government Data Mining: The Need for a Legal Framework*, 43 Harv. C.R.-C.L. L. Rev. 435 (2008) (discussing the same).

57 *See* David Freeman Engstrom & Connor Hoge, Enforcement by Algorithm 7–15 (2023) (unpublished article) (on file with authors).

58 Robert MacBride, The Automated State: Computer Systems as a New Force in Society 154 (1967).

59 *See* George Terborgh, The Automation Hysteria 51 (1966) (discussing the IRS's auditing automation and Congressional hearings on automation); Daniel Bell, *The Bogey of Automation*, N.Y. Rev. Books (Aug. 26, 1965), https://www.nybooks.com/articles/1965/08/26/the-bogey-of-automation/?printpage=true [https://perma.cc/9C4E-EYKM] (arguing that claims about automation's negative effects on the economy are false or unprovable); Howard Brick, *Optimism of the Mind: Imagining Postindustrial Society in the 1960s and 1970s*, 44 Am. Q. 348 (1992) (discussing the sources of some anti-automation themes in the 1960s); MacBride, *supra* note 58, at 168 (discussing individual freedom in an automated future).

60 5 U.S.C. § 552(a). This history is well told in two reports: U.S. Cong., Off. of Tech. Assessment, OTA-CIT-296, Federal Government Information Technology: Electronic Record Systems and Individual Privacy (1986) [hereinafter OTA Systems Report]; U.S. Cong., Off. of Tech. Assessment, OTA-CIT-297, Federal Government Information Technology: Management,

recognized that computer profiling was both new and decidedly not.[61] Red flags and suspicious behavior had always triggered closer scrutiny.[62] And many saw computerization as a net benefit. Conceding concerns about bias and privacy, they noted that datafication and digitization could empower agencies to detect "organizational offense[s]" that, with scattered victims and complex causation, were not well suited to old school investigative efforts based on tips and informants.[63] Moreover, whereas precomputer profiles could be especially "crude," relying largely on stereotyping, computer profiles could employ "sophisticated modeling" and generate more objective predictions.[64]

One of history's virtues is its capacity to chasten present-day views with a reminder that sometimes, perhaps often, there is nothing new under the sun. But revisiting earlier upheavals illustrates a further, critically important point as we size up the automated state and think about what to do about it: the automated state is fundamentally *ambiguous*, both in its methods and in its effects.

Start with methodological ambiguity—or, more concretely, what we mean by automation, or AI, in the first place. A credible case can be made that the current frontier of machine learning and big data makes the present round of government automation different in kind from past ones. Modern machine learning applications are nothing short of astounding in their power. Current efforts to automate selection of taxpayer audits[65] would be unrecognizable to past generations of agency administrators toting "expert systems."[66] The 1963 article, after all, noted that the IRS's "robot" was "lightning fast" but "pretty dumb."[67]

Security, and Congressional Oversight (1986) [hereinafter OTA Management Report]. For congressional hearings around amendments to the Privacy Act that ultimately became the Computer Matching and Privacy Protection Act of 1988, see *Computer Matching and Privacy Protection Act of 1987: Hearing on S. 496 Before a Subcomm. of the H. Comm. on Gov't Operations*, 100th Cong. (1987). For press accounts reflecting mounting public concern, see David Burnham, *Calculating the Cost of Government by Computer*, N.Y. Times (Apr. 17, 1983), https://www.nytimes.com/1983/04/17/weekinreview/calculating-the-cost-of-government-by-computer.html [https://perma.cc/A9U8-BWFW]; David Burnham, *Senators to Examine Official Use of Computer Data on Individuals*, N.Y. Times (Dec. 13, 1982), https://www.nytimes.com/1982/12/13/us/senators-to-examine-official-use-of-computer-data-on-individuals.html [https://perma.cc/H229-R66G]; Ross Gelbspan, *Computer Matching Stirs Up Criticism*, Bos. Globe, June 9, 1985, at A1.

61 OTA Systems Report, *supra* note 60, at 88.

62 *Id.*

63 Nancy Reichman, *Computer Matching: Toward Computerized Systems of Regulation*, 9 L. & Pol'y 387, 394 (1987).

64 OTA Systems Report, *supra* note 60, at 88.

65 Peter Henderson, Ben Chugg, Brandon Anderson, Kristen Altenburger, Alex Turk, John Guyton, Jacob Goldin & Daniel E. Ho, *Integrating Reward Maximization and Population Estimation: Sequential Decision-Making for Internal Revenue Service Audit Selection*, 37 Proc. AAAI Conf. on A.I. 5087, 5087 (2023).

66 *See supra* note 53 and accompanying text.

67 Crowley, *supra* note 46.

Too often missing in present-day debate, however, is recognition that many of the most concerning forms of government automation, including those that feature in many agonistes critiques, may have more in common with those earlier, low-tech implementations that governments have used for decades than with the current frontier of the "machine-learning state."[68]

Take benefits adjudication. A perennial challenge for government is efficiently and accurately distributing public benefits to qualified applicants. As the scale of this "[a]llocative [s]tate" has grown, agencies have tapped technology to help.[69] Sometimes that process has gone awry. The Michigan MiDAS fiasco is a good example, as are episodes in Colorado, Arkansas, and New Mexico.[70] However, many of these outrages did not involve machine learning. More often, the algorithm was just a "set of logical if-then statements—[a] step-by-step, hard-coded" recipe, sometimes called "logical"[71] or "symbolic" AI, not a "black box" neural net.[72] Its workflow could be printed on paper and pieced together even without programming experience.

Compare these efforts to tools under development at mass adjudicatory agencies such as the Social Security Administration ("SSA").[73] Dogged by revelations about case backlogs and stark inter-judge decisional disparities—some adjudicators grant disability benefits 10% of the time, others 90%[74]—the SSA built a team of lawyers and technologists and developed a trio of AI-based tools.[75] Most striking is the Insight System, which uses natural language processing—the branch of machine learning that performs text analytics, and a forerunner of the current large language models—to analyze judges' draft decisions and catch errors.[76]

---

68 *See generally* Aziz Z. Huq, *Constitutional Rights in the Machine-Learning State*, 105 Cornell L. Rev. 1875 (2020).

69 *Id.* at 1894–95.

70 Citron, *supra* note 38, at 1256 (Colorado); Ryan Calo & Danielle Keats Citron, *The Automated Administrative State: A Crisis of Legitimacy*, 70 Emory L.J. 797, 820–21 (2021) (Arkansas); Engstrom & Solow-Niederman, *supra* note 25, at 14–15, 18–19 (Michigan, New Mexico, and other states).

71 *See* Engstrom & Haim, *supra* note 14, at 281.

72 *See* AI Now Inst., Litigating Algorithms: Challenging Government Use of Algorithmic Decision Systems 7–8 (2018), https://ainowinstitute.org/wp-content/uploads/2023/04/litigatingalgorithms.pdf [https://perma.cc/677F-EYVN].

73 *See* Glaze et al., *supra* note 45; David Freeman Engstrom, Daniel E. Ho, Catherine M. Sharkey & Mariano-Florentino Cuellar, Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies 37–45 (2020), https://www.acus.gov/sites/default/files/documents/Government%20by%20Algorithm.pdf [https://perma.cc/XP9B-JE4Q] (report for the Administrative Conference of the United States).

74 *See* Engstrom et al., *supra* note 73, at 38; *see* Ames et al., *supra* note 38, at 31, 36.

75 *See* Engstrom et al., *supra* note 73, at 39.

76 *See id.* at 40–41.

The stark contrast between Michigan's MiDAS tool and the SSA's sophisticated suite of applications is important, for it suggests that much anxiety about automated government may miss the mark. The problem may not be the technical opacity of the "machine-learning state."[77] Rather, it is a slower burn of accountability concerns, particularly in thousands of smaller government units where low technical capacity and budgetary and political imperatives yield low-tech, underpowered, or poorly designed tools—whether homegrown or procured from contractors—that do not have enough agency wherewithal to meaningfully oversee their use.[78]

AI's methodological ambiguity creates a fundamental challenge for regulators as to where to direct their efforts: The machine-learning state? Or the vast iceberg of automation that sits below its machine-learning tip, decades in the making and deploying a near-infinite variety of computational approaches, from expert systems to number-crunching Excel spreadsheets? New York City learned this the hard way. Its high-profile Automated Decision Systems Task Force, tasked with proposing a regulatory framework, floundered after spending weeks trying to decide whether Excel was within its remit.[79] Framed too narrowly, regulation of public sector AI captures only the iceberg's high-tech and often less worrying tip.[80] Framed too broadly, however, and the effort soon encompasses the entirety of what modern government does.[81]

---

77 Huq, *supra* note 68, at 1890–91.

78 *See* Engstrom & Solow-Niederman, *supra* note 25, at 8; *see also* Levy et al., *supra* note 28, at 327.

79 *See* AI Now Inst., Confronting Black Boxes: A Shadow Report of the New York City Automated Decision System Task Force 13 (Rashida Richardson ed., 2019), https://ainowinstitute.org/ads-shadowreport-2019.html [https://perma.cc/4LEQ-6PHY]; Diana Budds, *New York City's AI Task Force Stalls*, Curbed N.Y. (Apr. 16, 2019, 11:35 AM), https://ny.curbed.com/2019/4/16/18335495/new-york-city-automated-decision-system-task-force-ai [https://perma.cc/6FMU-5BSS].

80 For instance, a recently enacted Connecticut law that mandates impact assessments focused on uncovering bias or disparate impact—already a limited substantive scope—applies only to advanced forms of AI, defined as

> (A) [A]n artificial system that (i) performs tasks under varying and unpredictable circumstances without significant human oversight or can learn from experience and improve such performance when exposed to data sets, (ii) is developed in any context, including, but not limited to, software or physical hardware, and solves tasks requiring human-like perception, cognition, planning, learning, communication or physical action, or (iii) is designed to (I) think or act like a human, including, but not limited to, a cognitive architecture or neural network, or (II) act rationally, including, but not limited to, an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communication, decision-making or action, or (B) a set of techniques, including, but not limited to, machine learning, that is designed to approximate a cognitive task.

Conn. Gen. Stat. § 4-68jj (2024).

81 *See* Engstrom & Haim, *supra* note 14, at 279 (discussing the challenge of "how, but also whether, to fortify existing public law paradigms given AI's power to transform government's

Ambiguity about methods is only half the story. The long history of government automation at the IRS and elsewhere also helps to surface a second of automation's ambiguities: the automated state is deeply ambiguous in its *effects*.

Consider here another key governance task: enforcement. Enforcement is central to government. It is how government gives life to legal mandates, converting "law in books" into "law in action."[82] Too little enforcement means socially costly lawbreaking. But too much, or coercive action targeting the wrong people or certain groups over others, is wasteful and unfair. And yet agencies also face an acute resource dilemma, for potential regulatory targets are infinite and agency resources finite.[83]

New automated tools under development at large U.S. enforcement agencies such as the Securities and Exchange Commission ("SEC"), the Environmental Protection Agency, and the IRS aim to address these dilemmas by arming line-level enforcement staff with data-based predictions as to who is violating the law, thus "shrinking the haystack" of potential targets.[84] As with any algorithmic system, these

---

structure and operations, but without hamstringing useful innovation" because AI is now used to make many significant governmental decisions); Levy et al., *supra* note 28, at 309 (discussing the government's broad use of AI in several consequential contexts such as "criminal justice, education, and benefits provision"). As an example from the federal level, the proposed Stopping Unlawful Negative Machine Impacts through National Evaluation Act applies to "dynamic or static machine learning algorithms or other forms of artificial intelligence, including a data system, software, application, tool, or utility" with the latter language seemingly encompassing virtually any form of automation. S. 5351, 117th Cong. § 2 (2022). Also at the federal level, the proposed Data Protection Act, which imposes disclosure requirements on federal and state agencies (but is mostly trained on private sector entities) applies to any "automated decision system," defined as "a computational process, including one derived from machine learning, statistics, or other data processing or artificial intelligence techniques, that makes a decision, or facilitates human decision making." S. 2134, 117th Cong. § 2 (2021). Another way to limit the scope of regulation is a risk-based approach in which the stringency of a new regulatory mandate is keyed to the gravity of the potential consequences. A law or regulation might, to cite two real-world examples, apply only to "critical," as in the Connecticut law just noted, or "high risk" decisions, as with the E.U. AI Act. *See European Approach to Artificial Intelligence*, EUR. COMM'N (Aug. 1, 2024), https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence [https://perma.cc/GJ6F-78VE]. But this only underscores the challenge: regulatory architects must make hard choices and, as elaborated below, large swathes of government automation are likely to go unregulated except by existing law.

82 Margaret H. Lemos, *Democratic Enforcement? Accountability and Independence for the Litigation State*, 102 CORNELL L. REV. 929, 931 (2017).

83 Engstrom & Ho, *supra* note 15, at 815.

84 *See* Henderson et al., *supra* note 65 (IRS); Elinor Benami, Reid Whitaker, Vincent La, Hongjin Lin, Brandon R. Anderson & Daniel E. Ho, *The Distributive Effects of Risk Prediction in Environmental Compliance: Algorithmic Design, Environmental Justice, and Public Policy*, 2021 PROC. ACM CONF. ON FAIRNESS, ACCOUNTABILITY, & TRANSPARENCY 90 (Environmental Protection Agency); ENGSTROM ET AL., *supra* note 73, at 22 (SEC); Engstrom & Ho, *supra* note 15, at 805 (SEC, IRS, and Environmental Protection Agency).

tools raise error and bias concerns.[85] But compared with dispersed line-level staff working up cases in traditional gumshoe fashion, these new tools might also make the government's exercise of coercive powers more efficient *and* more equitable. Indeed, frontier AI development at the IRS shows how automation, and the digitization and datafication that underpins it, opens up the possibility of bias at scale but also permits useful formalization of trade-offs between equity and efficiency in targeting tax audits.[86] Contrary to pervasive concerns about AI's "black box," government automation also brings transparency not possible in many analog systems.[87]

Predictive policing offers another example. Those tools have rightly raised concerns. Biased datasets can fuel runaway feedback loops: sending police back into the same neighborhoods again and again creates a self-fulfilling prophecy as to where crime is worst.[88] But it is also the case that use of data to predict where crime is likely to occur is as old as policing itself.[89] And, where data is not available or declared off-limits, anecdotal hunches about where to send patrols are the inevitable default.[90] Of course, human hunches might be better than bias-infused, data-based approaches, but they can also be a lot worse—a basic ambiguity that hangs over decades of debate around data mining in law enforcement.[91]

A final example bridges adjudication and enforcement and shows government automation at its most fraught but also its most ambiguous. In 2016, Allegheny County, Pennsylvania became one of the first U.S. jurisdictions to use an algorithm to help agency officials decide whether to "screen in" a household that had drawn a report of abuse and thus subject it to further investigation by a social worker.[92] Investigation can

---

85 Engstrom et al., *supra* note 73, at 35, 45, 79.

86 Emily Black, Hadi Elzayn, Alexandra Chouldechova, Jacob Goldin & Daniel E. Ho, *Algorithmic Fairness and Vertical Equity: Income Fairness with IRS Tax Audit Models*, 2022 Proc. ACM Conf. on Fairness, Accountability, & Transparency 1479, 1480, 1485.

87 Kleinberg et al., *supra* note 29, at 145.

88 Danielle Ensign, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger & Suresh Venkatasubramanian, *Runaway Feedback Loops in Predictive Policing*, 81 Proc. Mach. Learning Rsch. 160, 160 (2018).

89 *See* Sarah Brayne & Angèle Christin, *Technologies of Crime Prediction: The Reception of Algorithms in Policing and Criminal Courts*, 68 Soc. Probs. 608, 611 (2021); Andrew Guthrie Ferguson, *Policing Predictive Policing*, 94 Wash. U. L. Rev. 1109, 1123 (2017); Bernard E. Harcourt, Against Prediction: Profiling, Policing, and Punishing in an Actuarial Age 1–2 (2007); Sarah Brayne, Predict and Surveil: Data, Discretion, and the Future of Policing 5, 13 (2020).

90 *See* Brayne & Christin, *supra* note 89, at 609.

91 *See id.*; Zarsky, *supra* note 56, at 309; Cate, *supra* note 56, at 440.

92 *See* Alexandra Chouldechova, Emily Putnam-Hornstein, Diana Benavides-Prado, Oleksandr Fialko & Rhema Vaithianathan, *A Case Study of Algorithm-Assisted Decision Making in Child Maltreatment Hotline Screening Decisions*, 81 Proc. Mach. Learning Rsch. 134, 144–45 (2018); Marissa Gerchick, Tobi Jegede, Tarak Shah, Ana Gutierrez, Sophie Beiers, Noam Shemtov,

lead to additional services—a seeming win for struggling families—but it can also lead to a child's removal, one of the gravest things government does.[93]

Use of AI in child protection has rightly become ground zero in the debate about how to design and oversee algorithmic systems. The tools have drawn sharp criticism because of faulty forecasts and the risk of reproducing racial bias from historical over-surveilling of Black families.[94] Child protection may also be the clearest place where algorithms are crowding out the clinical judgments of trained professionals. Their introduction has touched off battles between "screen-level" agency managers and "street-level" social workers, who have learned how to game the new systems, whether out of concern for their clients or resentment at being reduced to glorified data entrants.[95] Finally, child welfare is the place where the critique of a myopic state—as noted previously, a government that shortsightedly focuses on managing the effects of policy problems rather than addressing root causes—is on most visible display.[96]

But child welfare is *also* the clearest place where debate bottoms out at longstanding disputes about the merits and demerits of clinical versus actuarial judgment. The purely human approach to child protection was deeply biased.[97] New algorithmic approaches might be, too, but they at least hold the promise of mitigating bias through thoughtful attention to data and modeling.[98] Moreover, child welfare shows

Kath Xu, Anjana Samant & Aaron Horowitz, *How Policy Hidden in an Algorithm Is Threatening Families in This Pennsylvania County*, ACLU (Mar. 14, 2023), https://www.aclu.org/news/womens-rights/how-policy-hidden-in-an-algorithm-is-threatening-families-in-this-pennsylvania-county [https://perma.cc/N38J-L454].

93 *See* Chouldechova et al., *supra* note 92, at 136–37.

94 *See* EUBANKS, *supra* note 8 (discussing how automation and algorithms can increase bias against the poor); DOROTHY ROBERTS, TORN APART: HOW THE CHILD WELFARE SYSTEM DESTROYS BLACK FAMILIES—AND HOW ABOLITION CAN BUILD A SAFER WORLD (2022) (discussing racial bias in the child welfare system).

95 Jennifer Raso, *Displacement as Regulation: New Regulatory Technologies and Front-Line Decision-Making in Ontario Works*, 32 CANADIAN J.L. & SOC'Y 75, 92 (2017).

96 *See* Stephanie K Glaberson, *Coding Over the Cracks: Predictive Analytics and Child Protection*, 46 FORDHAM URB. L.J. 307, 320 (2019).

97 Brett Drake, Melissa Jonson-Reid, Maria Gandarrilla Ocampo, Maria Morrison & Darejan (Daji) Dvalishvili, *A Practical Framework for Considering the Use of Predictive Risk Modeling in Child Welfare*, 692 ANNALS AM. ACAD. POL. & SOC. SCI. 162, 170 (2020); ANGELA WHITE & PETER WALSH, RISK ASSESSMENT IN CHILD WELFARE 4 (2006); *see* Stephanie Cuccaro-Alamin, Regan Foust, Rhema Vaithianathan & Emily Putnam-Hornstein, *Risk Assessment and Decision Making in Child Protective Services: Predictive Risk Modeling in Context*, 79 CHILD. & YOUTH SERVS. REV. 291, 292 (2017).

98 Chouldechova et al., *supra* note 92, at 135; s*ee* Maria De-Arteaga, Riccardo Fogliato & Alexandra Chouldechova, *A Case for Humans-in-the-Loop: Decisions in the Presence of Erroneous Algorithmic Scores*, 2020 PROC. CONF. ON HUM. FACTORS IN COMPUTING SYS. 1, 5–7; RHEMA VAITHIANATHAN, EMILY PUTNAM-HORNSTEIN, NAN JIANG, PARMA NAND & TIM MALONEY, DEVELOPING

how AI can be used to "surveil," but also "spotlight," with great risk of exacerbating inequities but great potential to expose and then address inequities with well-targeted resources.[99] AI's agonistes too often resort to rhetoric here. In one formulation, automation "trade[s] the possibility of human bias for the guarantee of systemic bias."[100] But human biases are also pervasive.[101] And although algorithmic tools can be biased at scale, they are also "*perfectible* at scale."[102] A single, centralized decisionmaker can surely be worse than a dispersed army of human decisionmakers, but it can also be better. In fact, some AI optimists predict that law will soon *require* use of AI where actuarial judgments are demonstrably more accurate and equitable than the clinical judgments of human bureaucrats.[103]

To note automation's twin ambiguities is not to suggest that all is well with the automated state. Far from it. Critics' concerns about a secret, regressive, racist, myopic, relentless, and repressive state retain substantial force. And just because automation comes in good and bad forms does not defeat the case for robust regulation. Even a few bad implementations in a sea of good ones might violate a Rawlsian difference principle—the notion that any proposed distribution of resources or regulatory burdens should, at minimum, benefit the worst-off.[104]

From a realist's perspective, however, automation's ambiguities carry a further, critically important implication: they raise considerable doubt as to whether policy windows will *ever* open for the large mass of government automation outside of well-bounded technologies, such as FRT, or "high-risk" uses, such as criminal risk assessment, that raise especially acute or visceral concerns. For most other implementations, algorithmic accountability will be *litigated*, using existing laws and doctrines, not *legislated*, via bespoke new rules.

Predictive Models to Support Child Maltreatment Hotline Screening Decisions: Allegheny County Methodology and Implementation 4, 11–12 (2019) [hereinafter AFST Report], https://www.alleghenycountyanalytics.us/wp-content/uploads/2019/05/16-ACDHS-26_PredictiveRisk_Package_050119_FINAL-2.pdf [https://perma.cc/54U3-ETUQ].

99 Rebecca A. Johnson & Tanina Rostain, *Tool for Surveillance or Spotlight on Inequality? Big Data and the Law*, 16 Ann. Rev. L. & Soc. Sci. 453, 454–55 (2020).

100 Calo & Citron, *supra* note 70, at 819.

101 Cary Coglianese & Alicia Lai, *Algorithm vs. Algorithm*, 72 Duke L.J. 1281, 1293–99 (2022). *See generally* Daniel Kahneman & Amos Tversky, *Subjective Probability: A Judgment of Representativeness*, 3 Cognitive Psych. 430 (1972) (discussing heuristics and how they affect evaluation by humans).

102 *See* Engstrom & Solow-Niederman, *supra* note 25, at 11.

103 Cary Coglianese & Kat Hefter, *From Negative to Positive Algorithm Rights*, 30 Wm. & Mary Bill Rts. J. 883, 910–11 (2022).

104 Iason Gabriel, *Toward a Theory of Justice for Artificial Intelligence*, 151 Daedalus 218, 226 (2022).

II. Gray Zones, Gray Holes, and Law's Limits

If algorithmic accountability is to be litigated, not legislated, then a critically important task is to look under the hood of modern administrative law—the mix of constitutional and statutory rules that constrain government action—and understand its contours.[105] Here again, a dose of history informs the realist view. For sitting beneath anxieties about a secret, regressive, racist, myopic, relentless, and repressive state lies a pair of deeper, tectonic shifts that too often go missing in current debate. First, as AI has advanced, government automation has steadily pushed into *gray zones*—that is, highly consequential governance tasks previously thought to require human discretion because they were too complex to be tractable to quantification, reduction, or encoding in automated systems.[106] But, at the same time that machines have moved into gray zones, government automation has simultaneously pushed into legal *gray holes*, where legal accountability is, and has long been, either practically or deliberately weak or even nonexistent.[107]

Start with gray zones. Even a cursory glance across the current landscape of government automation leaves the undeniable impression that something is different this time round. For AI agonistes, it is a state that is *more* secret, regressive, racist, myopic, relentless, and repressive. But, as already noted, each of these claims raises empirical questions because of AI's ambiguities.[108] A rich debate is emerging as to each critique and will doubtless continue for years to come.[109]

In the meantime, it is safer empirical ground to say that the current round of government automation is different from past ones on at least two counts. First, potent new analytics can exploit what were previously low value and often unstructured data—the sea of numbers and text in administrative records or audio and video footage collected by growing surveillance cameras and microphones.[110] Machine learning is more

---

105 *Supra* note 33 and accompanying text.

106 *See* Solow-Niederman, *supra* note 36; Veale & Brass, *supra* note 36, at 125.

107 *See* Solow-Neiderman, *supra* note 36; Haim, *supra* note 37.

108 *Supra* notes 7–24 and accompanying text.

109 An interdisciplinary scholarly community organized under the heading "fairness, accountability, and transparency" is spearheading these efforts. *ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT)*, Ass'n for Computing Mach. (Aug. 10, 2024), https://facctconference.org/ [https://perma.cc/9Q2L-QDN8].

110 *See* Yeung, *supra* note 9, at 505; *see also* Sallyann Bergh, Alyssa Davis, Amber Ivey, Dan Kitson & Jennifer Thornton, How States Use Data to Inform Decisions: A National Review of the Use of Administrative Data to Improve State Decision-Making 5–6 (2018), https://www.pewtrusts.org/-/media/assets/2018/02/dasa_how_states_use_data_report_v5.pdf [https://perma.cc/M6ML-DLBP].

than just souped-up analytics; it also taps the latent potential of data that previously seemed unusable or unimportant.[111]

Second, and relatedly, the new algorithmic governance tools are more deeply embedded in the work of government than past waves of automation.[112] Advances in computing, more sophisticated analytics, and bigger data make it possible to automate a wider range of government tasks than ever before.[113] Today's tools go well beyond longstanding uses of technology to make reports to agency overseers, weigh policy choices, or perform back-office paper pushing. Machines touch "decisions as significant as whether to initiate an audit, inspection, or enforcement action; whether an applicant is entitled to public benefits or services; and what punishments should be meted out."[114] Put another way, advances in data and analytics have allowed public sector automation efforts to move into a set of "unstructured" and "'semi-structured' problems"—governance gray zones—where machines were previously thought unable to go.[115] The importance of this latter point cannot be overstated. Much of our anxiety about automation occurs because new governance technologies are being applied in ever more pointed exercises of sovereign power affecting more discrete individuals and interests than possible previously.[116] AI is quickly moving to the center of the coercive and (re)distributive power of the state, and we are rightly anxious.

At the same time, and just as important from a realist perspective, the movement of machines into governance gray zones has simultaneously pushed them into a very different gray area: *legal gray holes*, where law has not historically gone, or at least not gone well.[117] Sometimes, these legal gray holes are practical. Much of American public law is built around retail enforcement—a "liability in tort" model—that arms aggrieved individuals with a right of action against the government after provable harm has occurred.[118] But we have long known that ex-post accountability mechanisms, however potent in theory, skew

111 Alicia Solow-Niederman, *Information Privacy and the Inference Economy*, 117 Nw. L. Rev. 357, 388–90 (2022).

112 Levy et al., *supra* note 28, at 312; *see* Engstrom et al., *supra* note 73, at 11.

113 *See* Engstrom et al., *supra* note 73, at 11.

114 *See* Engstrom & Solow-Niederman, *supra* note 25, at 11.

115 *See* Veale & Brass, *supra* note 36, at 125 (quoting Robert Hoppe, The Governance of Problems: Puzzling, Powering and Participation 30 (2010)); Young et al., *supra* note 29, at 302.

116 *See* Engstrom & Solow-Niederman, *supra* note 25, at 10–12; Engstrom & Hoge, *supra* note 57, at 20.

117 Solow-Niederman, *supra* note 36, at 118. *See generally* David Dyzenhaus, The Constitution of Law: Legality in a Time of Emergency 3, 30 (2006).

118 Huq, *supra* note 68, at 1883 (quoting Steven Shavell, *Liability for Harm Versus Regulation of Safety*, 13 J.L. Stud. 357, 357 (1984)).

toward proving and reversing individual error and do little to shape systemic design choices or practices.[119]

The more important gray holes, however, may be the deliberate ones. To see this, one cannot do better than a pair of iconic cases—a tale of two *Heckler*s, as it turns out. Consider, first, a cornerstone of American administrative law, *Heckler v. Chaney*.[120] The case has heady facts. Death penalty opponents challenged the Food and Drug Administration's refusal to initiate enforcement against prison officials for using drugs to perform executions by lethal injection that were not—and could not be—"safe and effective" under the Food, Drug, and Cosmetic Act.[121] The Supreme Court held the agency's refusal to enforce unreviewable and ended the case.[122] Thus, the presumption of unreviewability was born—a mostly impenetrable shield that has built up around agency enforcement decisions, rendering them largely immune to legal challenge.[123]

Part of what is going on in *Heckler* is the problem of inaction and passivity—what to do when agencies with broad legislative mandates refuse to make policy or enforce parts of it. Later cases, however, have steadily expanded the general principle to cloak nearly all enforcement decisions, including challenges by regulatory targets themselves.[124] In most instances, the best a target can do is wait out a costly enforcement action and then seek to have an adverse decision set aside on the merits.

The reasons for *Heckler*'s protective shield are, at the least, plausible. One is concern about generalist judges reviewing expert bureaucrats, especially around how to spend a scarce budget—a principle of institutional deference.[125] Another reason is epistemic: no decisionmaker

119 Ames et al., *supra* note 38, at 23; *see* Citron, *supra* note 38; Daniel E. Ho, David Marcus & Gerald K. Ray, Quality Assurance Systems in Agency Adjudication: Emerging Practices and Insights 18–19 (2021), https://www.acus.gov/sites/default/files/documents/ACUS_QA_Report_Final_Nov30.pdf [https://perma.cc/3QC3-AHHZ]; Daniel E. Ho & Sam Sherman, *Managing Street-Level Arbitrariness: The Evidence Base for Public Sector Quality Improvement*, 13 Ann. Rev. L. & Soc. Sci. 251, 265 (2017).

120 470 U.S. 821 (1985).

121 21 U.S.C. §§ 301–399; *Heckler*, 470 U.S. at 824 (quoting 21 U.S.C. § 355).

122 *Heckler*, 470 U.S. at 837.

123 *Id.* at 832 ("An agency's decision not to take enforcement action should be presumed immune from judicial review under [the Administrative Procedure Act].").

124 This part of the protective shield built up around enforcement comes from administrative law's timing doctrines, including limits on pre-enforcement review and the "legislative rule" doctrine, as discussed *infra* notes 203–206. But the main component is ripeness: the Court's holding in *FTC v. Standard Oil Co. of California* (*SoCal*), 449 U.S. 232 (1980), that an agency's decision to initiate an investigation is not ripe, and thus not reviewable on a quasi-interlocutory basis. *Id.* at 246. *SoCal* can be thought of as a complement to *Heckler*, limiting even a regulatory *target's* ability to challenge agency enforcement decisions. For more discussion, see Engstrom & Hoge, *supra* note 57, at 39.

125 *Heckler*, 470 U.S. at 831 (noting that an "an agency decision not to enforce often involves a complicated balancing of a number of factors which are peculiarly within its expertise," including

(court or otherwise) can reliably reconstruct specific agency decisions to enforce or not against a regulatory target.[126] When dispersed human prosecutors are poking around in a vast haystack of potential violators, there is no focal point for review.[127] But a third and broader concern hovers over both of these: lawsuits work best when they seek to enjoin *specific* actions causing *concrete* harm to identifiable persons.[128] Law is less good—and, indeed, is strongly disfavored—when its aim is, as a related Court decision put it, "wholesale improvement" of government operations.[129]

Reasonable minds—and, of course, lawyers—can disagree about whether deference, indeterminacy, or something else is driving the bus in *Heckler*, but the legal consequences are not in doubt. Under current doctrine, almost any effort to challenge an agency's use of an algorithmic tool to identify enforcement targets would be a nonstarter.[130]

But that is only the beginning. Even when a litigant gets into court, legal requirements are famously weak. Indeed, modern American administrative law offers only a thin hedge against arbitrary government decisions—often little more than a gut check for overall "rationality."[131] Courts give particular deference to agency scientific and predictive judgments in recognition of the fact that agencies often operate out at the regulatory frontier under conditions of uncertainty.[132] Formulations vary, but most courts require no more than that agencies' analytic models sit "within the range of scientific defensibility" or bear a "reasonable . . . relationship" to data.[133] Moreover, administrative

---

"whether agency resources are best spent on this violation or another"); *id.* at 831–32 ("[An] agency is far better equipped than the courts to deal with the many variables involved in the proper ordering of its priorities."); *see also* Engstrom & Ho, *supra* note 15, at 832 n.120; Engstrom & Hoge, *supra* note 57, at 28. In a separate opinion that same term, the Court offered some further detail on the "many variables" that *Heckler* noted are at issue in enforcement decisions:

> This broad discretion rests largely on the recognition that the decision to prosecute is particularly ill-suited to judicial review. Such factors as the strength of the case, the prosecution's general deterrence value, the Government's enforcement priorities, and the case's relationship to the Government's overall enforcement plan are not readily susceptible to the kind of analysis the courts are competent to undertake.

Wayte v. United States, 470 U.S. 598, 607 (1985).

126 Engstrom & Ho, *supra* note 15, at 833.

127 *See Heckler*, 470 U.S. at 832 ("[W]hen an agency *does* act to enforce, that action itself provides a focus for judicial review, inasmuch as the agency must have exercised its power in some manner."); Engstrom & Ho, *supra* note 15, at 833.

128 *See* Norton v. S. Utah Wilderness All., 542 U.S. 55, 64 (2004).

129 *Id.* (emphasis removed) (quoting Lujan v. Nat'l Wildlife Fed'n, 497 U.S. 871, 891 (1990)).

130 Engstrom & Ho, *supra* note 15, at 834–36; Engstrom & Hoge, *supra* note 57, at 2.

131 Vermeule, *supra* note 33, at 19–20.

132 Kristin E. Hickman & Richard J. Pierce, Jr., Administrative Law Treatise 1116 (6th ed. 2019).

133 *Id.* at 605 (quoting Nat'l Wildlife Fed'n v. EPA, 286 F.3d 554, 565 (D.C. Cir. 2002)); Nat. Res. Def. Council v. EPA, 16 F.3d 1395, 1404 (D.C. Cir. 1993).

law has never demanded anything approaching full transparency when reviewing agency decision-making processes—the "how" versus the "what" of policymaking.[134] Even for major policy decisions, an agency need only ventilate its choices by publishing "notice" of its intent to issue them, soliciting "comment" from the public, and then defending its ultimate policy choices in writing.[135] Subsequent legal review does not typically require anything like open-source code or data or elaborate explications of analytics or otherwise breach the "computational wall."[136] In short, it is not just that American administrative law is often ineffectual for agency actions at scale or that it hives off entire areas from legal review. In addition, current law sharply limits rummaging around in agency policy formation even when review is available.

The other *Heckler*—the Court's 1983 decision in *Heckler v. Campbell*[137]—further illustrates the nonrummaging principle. Criticized for backlogs and inconsistent disability benefits decisions, the SSA promulgated, after notice and comment, an elaborate grid—sometimes called a "plug-and-chug" rule, or, within SSA, "the grid"[138]—to streamline adjudication by administrative judges.[139] Under the SSA's new approach, an adjudicator would take evidence on facts about a claimant's disability and past employment and then, reading across the chart, determine if benefits would flow.[140] The Supreme Court upheld the grid and its application—thus blessing a substantial narrowing of human discretion in favor of what amounts to a coarse algorithmic approach.[141] Where a grid has survived notice and comment, the Court concluded, it does not violate due process, except where an agency denies all opportunity to present individualized information before the plugging and chugging begins.[142]

In an age of algorithms, *Heckler v. Campbell* seems the more redeemable of the two *Heckler* cases. The year was 1983, and the grid's guts were fully transparent in the same way that low-tech, "expert systems"—the "if-then" recipes noted previously—can be printed out and assessed even by laypeople,[143] not like a neural net that makes a benefits/no-benefits decision after being fed a set of medical reports

---

134 *See* Wagner with Walker, *supra* note 32, at 158–63; Wagner & Murillo, *supra* note 32, at 15–16.

135 Wagner & Murillo, *supra* note 32, at 13.

136 *Id.* at 5, 15–17.

137 461 U.S. 458 (1983).

138 Fast v. Barnhart, 397 F.3d 468, 471 (7th Cir. 2005).

139 *Campbell*, 461 U.S. at 460–62.

140 *Id.*

141 *See id.* at 467–68.

142 *Id.* For the classic treatment of the SSA's medical-vocational guidelines, see generally Jerry L. Mashaw, Bureaucratic Justice: Managing Social Security Disability Claims (1983).

143 *See supra* notes 53, 71 and accompanying text.

but provides no explanation for how it got there.[144] Perhaps the current Court would balk at such a system, declaring its "black box" nature inconsistent with due process and forbidding its use. It seems more likely, however, that the current Court would bless the system, so long as its basic attributes—for instance, a "system-level" description of data features and modeling approach[145]—had been ventilated via notice and comment. This more modest, algorithms-as-rules approach would fall well short of agonistes' calls for full, decision-level transparency, but it would be better than nothing.[146]

Might the Court, instead, choose an in-between option and, tracking contemporary calls for robust regulation of algorithmic systems, call for more demanding, "decision-level" scrutiny of an SSA algorithmic decision-making system? Here, it is useful to note that, when the SSA first created the grid in the 1970s, the agency defended the grid's specifics during notice and comment as "an elaboration of longstanding policy" supported by "considerable agency expertise developed over many years."[147] "Organizations and professions," the agency continued, "commonly recognize the value of experience even though it may not always be presented in statistical form."[148] Thus, in blessing the SSA's grid some years later, the Court seemed to accept the agency's reliance on the collective expertise and experience of past and present employees—what some have called "institutional intelligence."[149] In the current algorithmic

---

144 *See supra* text accompanying note 72.

145 *See* Cary Coglianese & David Lehr, *Transparency and Algorithmic Governance*, 71 ADMIN. L. REV. 1, 1, 24–25, 38–39, 42 (2019). *See generally* Engstrom & Ho, *supra* note 15, at 825–26.

146 Note that, as elaborated below, an "algorithms-as-rules" approach would also bring substantial costs. *See infra* notes 207–08 and accompanying text.

147 Rules for Adjudicating Disability Claims in Which Vocational Factors Must Be Considered, 43 Fed. Reg. 55,349, 55,357–62 (Nov. 28, 1978) (to be codified at 20 C.F.R. pts. 404, 416). For another round of debate, focused on whether the grid would discriminate against young people, fail to account for local job conditions or relevant claimant characteristics, or would implement an "average man" approach that was insufficiently attuned to case by case nuance, see Rules for Adjudicating Disability Claims in Which Vocational Factors Must Be Considered, 43 Fed. Reg. 9284 (Mar. 7, 1978) (to be codified at 20 C.F.R. pts. 404, 416).

148 Rules for Adjudicating Disability Claims in Which Vocational Factors Must Be Considered, 43 Fed. Reg. at 55,357; *see also* Rules for Adjudicating Disability Claims in Which Vocational Factors Must Be Considered, 43 Fed. Reg. at 9289 ("Prior experience of the Social Security Administration in determining when age makes a difference in disability determinations has also been considered . . . .").

149 JERRY L. MASHAW, RICHARD A. MERRILL, PETER M. SHANE, M. ELIZABETH MAGILL, MARIANO-FLORENTINO CUELLAR & NICHOLAS R. PARRILLO, ADMINISTRATIVE LAW: THE AMERICAN PUBLIC LAW SYSTEM 478 (8th ed. 2019); *see also* Cameron Averill, *Algorithmic Reason-Giving, Arbitrary-and-Capricious Review, and the Need for a Clear Normative Baseline*, SSRN 9, 21 (Feb. 6, 2024), available at https://ssrn.com/abstract=4640667 [https://perma.cc/LY4N-7TAP] (likening algorithmic systems to policies as "tools of bureaucratic management" to "standardize and centralize decision-making" and noting that both "will encode the expertise of many individuals" and are often built using years of data based on the application of policies).

age, that notion puts the Court in a bind. On one hand, the Court could adopt an algorithm-specific rule in response to the longstanding concern that, as agency rules have grown more technical, agency administrators may not hide controversial policy choices behind scientific or technical claims.[150] However, were an AI-wary Court to subject new algorithmic tools to hefty validation and testing requirements, many an agency would be deterred from using cutting-edge analytics in order to avoid onerous scrutiny that does not apply to "organizational" or "collegial" decisions justified by reference to tacit expertise and professional judgment.[151]

Understanding the two *Heckler*s offers much-needed perspective on current debate about the automated state. For starters, the two cases provide rich context for thinking about agoniste proposals, from bold prohibitions on the more sophisticated AI model types and demanding transparency requirements to lofty calls to "democratize AI." In some instances, those proposals largely duplicate rules already in place. Indeed, calls to "democratize AI" by requiring agencies to seek stakeholder input, publicly disclose details of new algorithmic systems, or perform impact assessments, are not dissimilar from the notice-and-comment process at the heart of American administrative law. Many of them thus violate a commonsense maxim against imposing redundant regulatory burdens where reasonably good statist institutions already exist.[152] In other instances, current proposals are the opposite. They are revolutionary rather than redundant in demanding decision-level transparency, open-source code, or elaborate validation of algorithmic systems prior to deployment. They would represent a significant increase, even a quantum leap, in scrutiny of government action.

For those who see in the automated state a government that is growing more secret, regressive, racist, myopic, relentless, and repressive, increased regulatory stringency would be welcome. But understanding the two *Heckler*s should, at the very least, be chastening. To begin, the *Heckler* cases are a stern reminder that the prohibitions, prescriptions, and transparency measures that dominate first-wave algorithmic accountability proposals bring tradeoffs, often steep ones. Some tradeoffs

---

150 *See* Wendy E. Wagner, *A Place for Agency Expertise: Reconciling Agency Expertise with Presidential Power*, 115 COLUM. L. REV. 2019, 2023, 2025 (2015) (noting the growing "hypertechnicality of agency rules" and "[s]kepticism about the agency-as-expert"); Emily Hammond Meazell, *Super Deference, the Science Obsession, and Judicial Review as Translation of Agency Science*, 109 MICH. L. REV. 733, 735–36 (2011) (noting concern that agency decisions might strategically hide policy decisions behind scientific claims); Wendy E. Wagner, *The Science Charade in Toxic Risk Regulation*, 95 COLUM. L. REV. 1613, 1617 (1995) (same). For general discussion of agency science, see generally WENDY WAGNER, SCIENCE IN REGULATION: A STUDY OF AGENCY DECISIONMAKING APPROACHES (2013).

151 For the notion of "organizational" and "collegial" decisions, see MASHAW ET AL., *supra* note 149, at 898. For a general discussion, see generally Wagner & Murillo, *supra* note 32.

152 *See* Himmelreich, *supra* note 40, at 1339–40.

are specific to particular governance tasks, as with *Heckler*'s focus on enforcement decisions.[153] Indeed, a requirement of decision-level transparency for an algorithmic tool that helps identify enforcement targets would quickly drain the tool of utility by telegraphing government action and facilitating gaming by regulated parties.[154] Transparency might make sense for welfare benefits but less so enforcement.[155]

More generally, we should worry that more stringent review of state action will exacerbate the risk aversion of already-cautious bureaucrats. It was not so long ago that even progressive voices argued that the problem with government was not too little transparency but too much, given government's well-known sclerotic tendencies.[156] More recent commentary has questioned transparency's premises and value,[157] including its weaponization by government's neoliberal critics.[158] Still others warn that heightened scrutiny via new procedural burdens, however well-meaning, can stultify a government that already lacks vigor.[159] As Justice Holmes put it long ago, there must be limits to legal challenges "if government is to go on."[160]

This line of thinking highlights a final challenge for the agonistes. Many of their proposals apply *only* to automated systems. That focus could well make sense. Concerns about a secret, regressive, racist, myopic, relentless, and repressive state may, as the automated state's trajectory comes into clearer focus, prove spot on and more than justify intrusive interventions or bans. But it is not hard to see that the approach could also yield the worst of all worlds. Agencies might shy away from salutary technological innovations if they subject them to more stringent scrutiny.[161] The public-private technology gap will widen further, with easy ways to improve public administration, and thus bolster the legitimacy of an embattled liberal-democratic state, left in a drawer or, more likely, a computer hard drive and never implemented.

The alternative, of course, is to apply greater scrutiny to *all* government action, not just the algorithmic sort—a leveling up of regulatory stringency across the board. Even before AI, one could argue,

---

153 *See* Heckler v. Chaney, 470 U.S. 821, 831–32 (1985).

154 *See* Jane Bambauer & Tal Zarsky, *The Algorithm Game*, 94 Notre Dame L. Rev. 1, 33–38 (2018).

155 *See* Engstrom & Ho, *supra* note 15, at 824–27.

156 *See* Stuntz, *supra* note 34.

157 Ananny & Crawford, *supra* note 7, at 2; Lilian Edwards & Michael Veale, *Slave to the Algorithm? Why a 'Right to an Explanation' Is Probably Not the Remedy You Are Looking For*, 16 Duke L. & Tech. Rev. 18, 21–22 (2017); *see* Joshua A. Kroll, Joanna Huey, Solon Barocas, Edward W. Felten, Joel R. Reidenberg, David G. Robinson & Harlan Yu, *Accountable Algorithms*, 165 U. Pa. L. Rev. 633, 636–38 (2017).

158 Pozen, *supra* note 35, at 147–48.

159 Sabin, *supra* note 35, at xvii; Bagley, *supra* note 35, at 348–49.

160 Bi-Metallic Inv. Co. v. State Bd. of Equalization of Colo., 239 U.S. 441, 445 (1915).

161 Wagner & Murillo, *supra* note 32, at 20–21.

government was plenty secret, regressive, racist, myopic, relentless, and repressive. Why not subject *all* government operations to more scrutiny? AI might thereby hold up a mirror to our democratic system and force us to do better. But here again, there are risks. We should not level up lightly or allow AI to be a stalking horse for a wholesale revision of public law—at least not without more careful thinking.

III. Adapting Law to an Algorithmic Age: A Research Agenda

For AI agonistes, the news so far is not good. AI is ambiguous in its methods and effects, making bespoke new regulatory approaches unlikely for all but a few well-bounded and especially concerning or "high risk" applications. Algorithmic accountability for the rest will be litigated, not legislated. Worse, modern administrative law is anemic compared with a growing menu of proposals for how to bring the automated state to heel. Legal gray holes, many of them deliberately created or at least abided, abound. And, even where judicial review can be had, it is often ineffectual for agency actions at scale, and it rarely requires anything approaching the level of scrutiny that many current proposals contemplate. The result? If it is to be achieved at all, algorithmic accountability will come not via first-best legislative overhauls but a decidedly second-best process of legal adaptation, worked out case by case by judges primed to see tradeoffs all around.

These conclusions are a blow to agonistes, but there is power in their recognition, for they allow us to focus our energies and think constructively about what work needs to be done. Armed with a more realist view, we can begin to glimpse some plausible paths forward for bringing the automated state to heel, even if they look different from the first-wave accountability measures that have so far dominated the AI governance debate.

How to get there? A trio of research tasks can light the way. First, we must map the automated state's many contours. Second, we must take better and more detailed account of how AI is reshaping agencies of all stripes and the myriad ways they are designing, implementing, and overseeing algorithmic systems. Third, we must think creatively about how existing law can be adapted to the world uncovered by the first two tasks. The rest of this Part considers each task in turn.

The first task—mapping the automated state's still-emerging contours—is a rote descriptive exercise, but it will pay large dividends. Key questions include the following: How many of the more concerning use cases within an increasingly automated state are, as with the enforcement decisions of civil regulatory agencies[162] or local agencies not subject

162 Engstrom & Ho, *supra* note 15, at 828; Engstrom & Hoge, *supra* note 57, at 27.

to much administrative law at all,[163] simply "hive[d] off"[164] from existing legal review mechanisms? And what proportion of the more concerning use cases are low-tech and transparent, yet deployed by agencies that cannot meaningfully oversee them? As noted previously, it is plausible that, once the automated state has been fully surveyed, the main problem will not be AI's "black box" but rather cruder implementations by beleaguered and underresourced agencies, particularly state and local ones, without enough technical capacity to responsibly oversee their use.[165]

The second task continues the descriptive project, but it quickly turns prescriptive. We need to know more than we do about how agencies design, implement, and evaluate algorithmic systems. That means less rehashing of Michigan's decade-old MiDAS debacle and more analysis of the multitudinous other places and ways that agencies have begun to develop practices—some exemplary, others dubious—in automating bureaucratic routines.[166]

One of the critical questions to answer before we can adapt current law is how AI will change agencies. Bureaucracies can already be thought of as forms of AI—that is, structures, processes, and decision-making nodes through which inputs are diffused in order to generate outputs.[167] Given this, a pithy version of the question can be asked: What happens when you add AI to AI? There is already substantial evidence that AI is reshaping agencies in fundamental ways. Dreamy futurists say algorithmic governance tools will eventually supplant human discretion entirely.[168] Although that day may or may not come, in the meantime, AI will surely shift discretion around. A useful shorthand here—though one that grossly oversimplifies longstanding research on how "ownership and control of knowledge redistribute power within organizations"[169]—is that AI's absorption into government agencies will shift discretion "up, over, and out."[170]

---

163 *See* Maria Ponomarenko, *Substance and Procedure in Local Administrative Law*, 170 U. PA. L. REV. 1527, 1531 (2022).

164 Engstrom & Hoge, *supra* note 57, at 27.

165 *See* Engstrom & Solow-Niederman, *supra* note 25, at 12, 29–32. *See generally* ENGSTROM ET AL., *supra* note 73, at 71–74 (discussing the importance of technical capacity).

166 *See* Engstrom & Haim, *supra* note 14, at 278.

167 *Compare* Felipe A. Csaszar & Tom Steinberger, *Organizations as Artificial Intelligences: The Use of Artificial Intelligence Analogies in Organization Theory*, 16 ACAD. MGMT. ANNALS 1, 13 (2022), *with* SIMON, *supra* note 27, at 1–17 (discussing decision-making in the administrative organization).

168 *See* Anthony J. Casey & Anthony Niblett, *The Death of Rules and Standards*, 92 IND. L.J. 1401, 1424, 1426–31 (2017).

169 *See* Peter Duchessi, Robert O'Keefe & Daniel O'Leary, *A Research Perspective: Artificial Intelligence, Management and Organizations*, 2 INTELLIGENT SYS. ACCT., FIN. & MGMT. 151, 152 (1993). *See generally* Daniel E. O'Leary & Efraim Turban, *The Organizational Impact of Expert Systems*, 7 HUM. SYS. MGMT. 11 (1987) (discussing the impact of expert systems on the distribution of power within organizations).

170 Engstrom & Haim, *supra* note 14, at 283 (citation omitted).

AI will shift discretion *up* by improving managerial control over the "street-level bureaucrats"—the social workers, administrative judges, and law enforcement officials—who have long served as discretion-wielding foot soldiers in the complex and uncertain policy spaces that pervade government action.[171] AI will shift discretion *over* by empowering technologists and new "algorithmic occupations" over the many other types of actors—policy analysts, economists, lawyers—that populate agencies.[172] And AI will shift discretion *out* by empowering private actors over public officials through the procurement process.[173]

Important questions are everywhere: Which shifts will be most powerful, and what will be the net effects? For instance, on the *up* shift, a long literature suggests we will inevitably move toward "screen-level" and then "systems-level" bureaucracy,[174] with agency managers using a growing tech toolkit to direct, evaluate, and discipline frontline employees.[175] But those employees are simultaneously developing "strategies of resistance"—from simple "foot-dragging" to "data obfuscation"—out of "fear[] of deskilling" and loss of professional autonomy.[176] For the moment, it remains unclear whether AI curtails or enables frontline and street-level discretion.[177] Time will tell. The larger lesson is that AI changes bureaucracy, but bureaucracy also changes AI.[178] We need to better understand how, when, and with what effects.

Descriptive questions about AI and bureaucracy are more than an academic exercise because the answers inform emerging views about how agencies *should* go about designing, implementing, and overseeing algorithmic systems. The standard starting point for that analysis is to note that algorithmic tools are not merely technical implementations

---

171 Young et al., *supra* note 29, at 301; MICHAEL LIPSKY, STREET-LEVEL BUREAUCRACY: DILEMMAS OF THE INDIVIDUAL IN PUBLIC SERVICES 3, 13–16 (1980).

172 *See* Elizabeth Magill & Adrian Vermeule, *Allocating Power Within Agencies*, 120 YALE L.J. 1032, 1035 (2011); Katherine C. Kellogg, Melissa A. Valentine & Angele Christin, *Algorithms at Work: The New Contested Terrain of Control*, 14 ACAD. MGMT. ANNALS 366, 388 (2020).

173 Deirdre K. Mulligan & Kenneth A. Bamberger, *Procurement as Policy: Administrative Process for Machine Learning*, 34 BERKELEY TECH. L.J. 773, 773 (2019); ENGSTROM ET AL., *supra* note 73, at 89; *see* Fourcade & Gordon, *supra* note 9, at 95.

174 Mark Bovens & Stavros Zouridis, *From Street-Level Bureaucracies to System-Level Bureaucracies: How Information and Communication Technology Is Transforming Administrative Discretion and Constitutional Control*, 62 PUB. ADMIN. REV. 174, 177–80 (2002).

175 *See* Kellogg et al., *supra* note 172, at 368.

176 Brayne & Christin, *supra* note 89, at 608; Kellogg et al., *supra* note 172, at 391; Raso, *supra* note 95, at 79.

177 Justin B. Bullock, *Artificial Intelligence, Discretion, and Bureaucracy*, 49 AM. REV. PUB. ADMIN. 751, 755–56 (2019).

178 *See* Duchessi et al., *supra* note 169, at 151–55.

but rather "socio-technical assemblages."[179] Because the design and deployment of an algorithmic system is social, not just technical,[180] meaningful accountability will depend on understanding how humans and machines interact to generate a decision, rather than assessing a system's technical outputs alone.[181]

This literature on the technical, managerial, and organizational dimensions of AI's incorporation into bureaucratic routines is still growing,[182] but a few key lines of inquiry are visible. How can agencies mitigate bias in their deployment of algorithmic systems through data and modeling choices?[183] How can agencies structure tech teams—fully centralized, decentralized, hub-and-spoke, something else—or blend personnel with technical and domain expertise to ensure that a mix of values and ethical perspectives are considered at each development stage?[184] And how best to put "humans in the loop" to leverage the best of human and machine judgment, or a blend of the two, while mitigating their weaknesses?[185]

On the latter, and as noted previously, child protection agencies' use of AI is already the site of rich debate and experimentation around sociotechnical design.[186] For instance, system designers in Allegheny County have been admirably attentive to feedback loop concerns in structuring the "hand-off" of machine outputs to human decisionmakers.[187] Agency staff who field calls reporting possible neglect or abuse use risk scores to prioritize investigations, but those scores are withheld from the street-level caseworkers who visit homes and make removal decisions to avoid the self-fulfilling prophecy of machine prediction that infects predictive policing tools.[188] Moreover, systems in place in Allegheny County and elsewhere feature a growing array of override procedures and behavioral nudges—steps agency staff must take to deviate from machine recommendations—to guard against overreliance ("automation bias") and underreliance ("algorithmic aversion")

---

179 Rob Kitchin, *Thinking Critically About and Researching Algorithms*, 20 INFO., COMMC'N & SOC'Y 14, 20 (2017).

180 *See* Michael Veale, Max Van Kleek & Reuben Binns, *Fairness and Accountability Design Needs for Algorithmic Support in High-Stakes Public Sector Decision-Making*, 2018 PROC. CHI CONF. ON HUM. FACTORS COMPUTING SYS. 1, 2.

181 *See* Engstrom & Haim, *supra* note 14, at 287–88; Levy et al., *supra* note 28, at 323; Veale & Brass, *supra* note 36, at 136.

182 *See* Engstrom & Haim, *supra* note 14, at 279.

183 *See* Kleinberg et al., *supra* note 29, at 115.

184 *See* Glaze et al., *supra* note 45, at 4, 16–21.

185 *See* Rebecca Crootof, Margot E. Kaminski & W. Nicholson Price II, *Humans in the Loop*, 76 VAND. L. REV. 429, 497–502 (2023); Engstrom & Haim, *supra* note 14, at 279, 289.

186 *Supra* notes 94–96 and accompanying text.

187 VAITHIANATHAN ET AL., *supra* note 98, at 26–32.

188 *Id.*

on machine outputs.[189] These design experiments in child protection and beyond are generating a nascent science and social science of the loop.[190]

Now the third task and the hardest of all: How might existing law be adapted to an algorithmic age? It is early days, but one can already glimpse several promising tracks. Consider three possibilities.

One obvious track is to reverse or substantially revise longstanding doctrines in light of the new realities of an algorithmic age. The obvious candidate here is *Heckler*. By centralizing and formalizing decision-making, algorithmic systems create the "focus for judicial review" that the *Heckler* Court found lacking in analog agency refusals to enforce.[191] In particular, digitization and datafication can reveal enforcement patterns, making it possible to test an agency's fidelity to its governing statute or its own regulations, and they can also make error rates explicit, rendering a court's review far more tractable.[192] As noted above, AI-based decision systems might, despite hand-waving references to AI's "black box," bring greater transparency over agency enforcement policies and practices than was possible in an analog world of dispersed line-level enforcers working up individual enforcement cases the old-fashioned way.[193]

But there are clear problems. For starters—and as noted above—the protective shield that has been built up around enforcement discretion is rooted in deference to agency expertise, not just the inscrutability of enforcement decision-making.[194] Algorithms shift some, but not all, of the positive and normative underpinnings of the doctrine. Just as important, *Heckler* is not the only barrier to review of agency enforcement decisions: Article III standing likewise hives off those decisions from review, as the Court recently reminded us in a muscular statement.[195] Still,

---

189 *See* De-Arteaga et al., *supra* note 98, at 2 (describing automation bias and algorithm aversion); Linda J. Skitka, Kathleen L. Mosier & Mark Burdick, *Does Automation Bias Decision-Making?*, 51 INT'L J. HUM.-COMPUT. STUD. 991 (1999) (explaining automation bias); Berkeley J. Dietvorst, Joseph P. Simmons & Cade Massey, *Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them*, 64 MGMT. SCI. 1155 (2016) (discussing algorithm aversion).

190 *See, e.g.*, Crootof et al., *supra* note 185, at 429, 497–503; Engstrom & Haim, *supra* note 14, at 277, 279, 288–89.

191 *See* Heckler v. Chaney, 470 U.S. 821, 832 (1985); Averill, *supra* note 149, at 27 (noting the possibility that algorithmic systems might provide "focal points" for judicial review); *id.* ("For as much as algorithms and policies undermine reason-giving, they can also facilitate more generalized contestation of agency decisions."); *see also* Engstrom & Ho, *supra* note 15, at 833.

192 *See* Engstrom & Ho, *supra* note 15, at 821; Averill, *supra* note 149, at 36.

193 *Supra* note 87 and accompanying text.

194 *See, e.g.*, *supra* note 149 and accompanying text.

195 *See* United States v. Texas, 599 U.S. 670, 674, 682 (2023) (first quoting Town of Castle Rock v. Gonzales, 545 U.S. 748, 761 (2005); then quoting Linda R.S. v. Richard D., 410 U.S. 614, 619 (1973)) (citing the "deep-rooted nature of []enforcement discretion" and the "fundamental Article

another problem returns to a puzzle noted above[196]: Would *Heckler*'s relaxation apply across the board, or just to algorithmic enforcement? Either approach could plausibly connect to older cases suggesting that *Heckler*'s properly limited domain is "single-shot" enforcement decisions as opposed to "expressions of broad enforcement policies [that] are abstracted from the particular combinations of facts the agency would encounter in individual enforcement proceedings."[197] One possible interpretation of this line of cases is that agency enforcement decisions are more likely reviewable as the agency's proffered reasons move outward from case-specific managerial concerns to legal reasons that implicate a broader swath of enforcement activity, as is true with algorithmic systems operating at scale.[198] Subsequent decisions, however, have mostly moved away from this view in favor of a near-categorical *Heckler* approach in which *any* agency invocation of resource constraints triggers immunities.[199] The broader point, which future work might consider, is that lowering the *Heckler* bar would, applied across the board, represent a significant change in current doctrine, opening litigation floodgates to spurned regulatory beneficiaries of all sorts. But limiting *Heckler*'s relaxation to algorithmic systems might be just as bad, chilling agency use of algorithmic tools even where they are likely to yield more equitable and more effective enforcement.

A final challenge in any rethinking of *Heckler* is what happens next once *Heckler*'s reviewability bar comes down. The answer is likely arbitrary and capricious review—American administrative law's default, all-purpose substantive review standard.[200] As currently formulated, that doctrinal test contains a quasi-procedural component (did the

III principle" that a challenger "lacks standing to contest the policies of the prosecuting authority when he himself is neither prosecuted nor threatened with prosecution").

196 *Supra* notes 143–52 and accompanying text.

197 Crowley Caribbean Transp., Inc. v. Peña, 37 F.3d 671, 673, 676–77 (D.C. Cir. 1994) (emphasis omitted) (noting that an agency announcement of nonenforcement in a particular case that also "lay[s] out a general policy delineating the boundary between enforcement and non-enforcement and purport[s] to speak to a broad class of parties" might be reviewable).

198 For an academic synthesis of this position, see Eric Biber, *The Importance of Resource Allocation in Administrative Law*, 60 ADMIN. L. REV. 1, 31 (2008) ("Because courts balance upholding statutory supremacy with deference to agency resource allocation, courts should be more willing to step in when the benefits to enforcing statutory supremacy are higher—i.e., when the judicial action will result in the correction of legal errors that might harm a wide range of private parties or public interests.").

199 A good example is *Citizens for Resp. & Ethics in Wash. v. FEC*, 892 F.3d 434 (D.C. Cir. 2018), in which Senior Circuit Judge Randolph, writing for a 2–1 majority, held that the agency's invocation of resource constraints barred review despite an express statutory provision authorizing review to determine if the agency's dismissal of a citizen-filed complaint was "contrary to law"—a determination that, if unremedied by the agency, authorized a private lawsuit. *Id.* at 436, 438–40.

200 *See* Motor Vehicle Mfrs. Ass'n. v. State Farm Mut. Auto. Ins. Co., 463 U.S. 29, 44–45 (1983).

agency apply the right factors?) and a quasi-substantive component (was there "a rational connection between the facts found and the choice made"?).[201] Both would require substantial updating to capture a growing body of knowledge about the optimal design, implementation, and oversight of algorithmic systems, and scholars have only just begun to consider how to do it. One thoughtful proposal would have judges consider a range of factors that includes (1) the agency's "systemic reasons" for its adoption and design choices, (2) the technical aptitude of agency officials or staff who might incorporate machine outputs into their decisions and, in addition, of those who are acted upon, (3) "the nature of the decision being made," and (4) "the conditions that led the agency to [choose] an algorithm[ic] [approach] in the first place," such as a high degree of inconsistency or inaccuracy in the existing data.[202] Future work could refine or add to these inquiries—but all will reraise the question of whether generalist judges are up to the task and whether similarly searching questions could be asked of the nonalgorithmic decision systems that pervade agencies without grinding the entire enterprise to a halt.

A second and less disruptive track is to find ways existing doctrines are already well-adapted to an algorithmic age and, without significant reworking and leveraging what we have learned about sociotechnical design, could provide principled line-drawings subjecting only some, but not all, algorithmic systems to legal scrutiny. For instance, administrative law already contains a potentially attractive option: the "legislative rule[]" doctrine.[203] Recall that a pillar of American administrative law is the requirement that agencies ventilate major policies by publishing "notice" of the agency's plans, taking "comment" from the public, and then defending ultimate policy choices in light of those comments.[204] But not every statement an agency makes must go through that process. Only policies that bind lower-level agency staff or are seen as binding by regulated parties out in the world trigger notice and comment.[205]

Here, doctrine already supplies a plausible way to decide when an algorithmic system should be subject to more stringent review. When an algorithm makes a decision outright or dictates decisions by lower-level bureaucrats—whether prosecutors making enforcement decisions or adjudicators making benefits decisions—it is binding and must go

201 *Id.* at 43 (quoting Burlington Truck Lines, Inc. v. United States, 371 U.S. 156, 168 (1962)). For theory and evidence on the present-day application of the standard, see generally Jacob Gersen & Adrian Vermeule, *Thin Rationality Review*, 114 MICH. L. REV. 1355 (2016).

202 *See* Averill, *supra* note 149, at 53–54. For another smart take on how to adapt arbitrary and capricious review to algorithmic systems, see Peter Henderson & Mark Krass, *Algorithmic Rulemaking vs. Algorithmic Guidance*, 37 HARV. J.L. & TECH. 105 (2023).

203 Engstrom & Ho, *supra* note 15, at 836–38.

204 *Id.*

205 *Id.* at 837.

through full notice and comment. When, however, an algorithm is a recommender system only—for instance, a machine merely provides a risk score or other piece of information to a human decisionmaker, or there is a robust review process empowering the human decisionmaker to override the machine decision—it is not binding.[206]

Caveats apply to this algorithms-as-rules approach. Growing research suggests that putting humans in loops is hardly an all-purpose or uncomplicated fix,[207] despite calls by some to make human oversight a strong default.[208] There also remains the question of the degree of transparency required during notice and comment: "system-level" transparency, meaning a description of data features and models, or something closer to "white-box" testing,[209] in which code and data are made publicly available to run extensive simulations? Lawyers will see still further challenges. Under current doctrine, a rule can only be changed by another rule, meaning another round of notice and comment.[210] Declaring that algorithms are rules thus introduces significant rigidity into a process of design, implementation, and evaluation that is highly iterative and rarely follows "a set of distinct, ordered steps."[211]

A third and final possible track for thinking about legal adaptation is conceptually attractive and yet may also be hardest to achieve. Some commentators have come to recognize that the key to bringing meaningful accountability to the automated state is to induce agencies to think critically about the design, implementation, and evaluation of automated systems throughout a system's lifecycle, not just prior to deployment or after it causes harm.[212] But if that is the goal, it is not hard to see in existing and proposed regulatory frameworks a basic mismatch with an emerging science of sociotechnical design[213]—what prior work has labeled the "snapshot" and "bookend" problems.[214] In a nutshell, many existing oversight frameworks "are unlikely to achieve meaningful accountability because they are static assessments of [algorithmic] systems at particular [points] in time"—the "snapshot" part of the problem—and they are also often "rigidly ex ante or ex post," focused at the very front- or back-end of the process of designing and implementing an automated tool—the "bookend" part of the

---

206 *See id.* A case that contains hints of this approach, applied to a statistical model, is *McLouth Steel Products Corp. v. Thomas*, 838 F.2d 1317, 1318–22 (D.C. Cir. 1988).

207 Green, *supra* note 16, at 2; Engstrom & Haim, *supra* note 14, at 289.

208 *See, e.g.*, FRANK PASQUALE, NEW LAWS OF ROBOTICS: DEFENDING HUMAN EXPERTISE IN THE AGE OF AI 3–12 (2020).

209 Kroll et al., *supra* note 157, at 650–51.

210 *See* Engstrom & Ho, *supra* note 15, at 836–37.

211 *See* Levy et al., *supra* note 28, at 310.

212 *See* Glaze et al., *supra* note 45, at 18–21; Douek, *supra* note 45, at 528, 568–69, 585.

213 *See* Engstrom & Ho, *supra* note 15, at 826–28, 849; Wagner & Murillo, *supra* note 32, at 12.

214 Engstrom & Haim, *supra* note 14, at 291.

problem.[215] However, a growing science of sociotechnical design and deployment suggests that trustworthy AI is "built in the in-between spaces," where system designers make critical decisions throughout a messy and iterative process of designing, implementing, and testing new algorithmic tools.[216] The core challenge is finding creative ways for law to do something it has not typically done, at least not well: shift accountability "upstream" to institutional designers and overseers[217] and induce those upstream actors to adopt "good" systemic practices, ones centered around careful design followed by continuous evaluation and improvement, and eschew "bad" ones.

That kind of sifting of agency process has long been one of public law's most vexing challenges. Indeed, an influential line of analysis in American administrative law circles holds that agency accountability has more reliably come from the internal norms and practices through which agencies define and implement policy mandates—dubbed "internal administrative law"—than external regulation imposed from without.[218] But it is worse, for external regulation can *reduce* agency incentives to adopt salutary internal rules and constraints by creating attachment points for legal review.[219]

A possible workaround might be an algorithm-specific adjunct to American administrative law's embattled deference regime—a kind of *Chevron*[220] deference for agencies that adhere to then-current guidance about how best to mitigate algorithmic harms.[221] A virtue of a deference approach is its status as a positive inducement for agencies to behave well, not a source of legal vulnerability when they do not. It also fits well with the focus of the traditional deference doctrines on agency legal interpretations. After all, an automated system that predicts likely violators of regulatory mandates, or one that makes a benefits/no-benefits decision, must encode law. As these systems come to pervade government, litigants will perennially raise questions about whether a particular implementation is faithful to congressional command or an agency's own regulation.[222] An agency that has engaged in careful

215 *Id.*

216 *Id.*

217 *See* Green, *supra* note 16, at 15.

218 *See* Mashaw, *supra* note 142, at 1; Gillian E. Metzger & Kevin M. Stack, *Internal Administrative Law*, 115 Mich. L. Rev. 1239, 1239 (2017).

219 *See* Metzger & Stack, *supra* note 218, at 1239, 1281, 1288; *see also* Wagner & Murillo, *supra* note 32, at 16.

220 467 U.S. 837 (1984), *overruled by* Loper Bright Enters. v. Raimondo, 144 S. Ct. 2244 (2024).

221 Deference regimes are embattled because of *Loper Bright*, 144 S. Ct. 2244, which held that courts "must exercise their independent judgment in deciding whether an agency has acted within its statutory authority." *Id.* at 2273. *See* Engstrom & Ho, *supra* note 15, at 853; Wagner & Murillo, *supra* note 32, at 25.

222 An AI-based tool that does not faithfully implement an agency's own regulation could support a challenge under the so-called "*Accardi* doctrine," which holds that agencies must follow

design, oversight, and evaluation of a tool could receive what amounts to interpretive deference to machine-aided decisions, thus empowering courts to give life to then-current "best practices" in the agency's design and deployment of automated tools.

Here again, there are caveats. To begin, after the Supreme Court's *Loper Bright Enterprises v. Raimondo*[223] decision, *Chevron* deference, at least, is no longer good law, leaving less of a doctrinal foothold for a deference-based, algorithm-specific approach.[224] Moreover, a long literature in sociology and law details the risk that internal organizational structures and processes will become symbolic indicators of compliance and that oversight and enforcement efforts will steadily morph into pro forma, "check list" compliance.[225] There is always slippage between announced policies and their implementation, so regulatory interventions will perennially be caught between a need to make organizations more critically reflective about the choices they make in the adoption, design, and deployment of algorithmic systems and the possibility that interventions, whether a deference-style approach or something else, will quickly become empty or toothless.[226]

Conclusion

AI's critics will bridle at much or all of the above. This Essay opened with concerns about a secret, regressive, racist, myopic, relentless, and repressive state and ended with some narrow ways to refashion existing legal doctrines. Those who call for new "imaginaries"[227] will see a whimper where a bang is needed—a spineless reduction of urgent questions about the future of the liberal-democratic state to a lawyer's game.[228]

---

their own rules. *See* United States *ex rel.* Accardi v. Shaughnessy, 347 U.S. 260 (1954).

223 144 S. Ct. 2244 (2024).

224 With *Chevron* deference no longer, an algorithm-specific deference regime would perhaps have to be tethered to the weaker—some would say tautological—*Skidmore* deference, whereby an agency receives deference to the extent its analysis is well-reasoned and careful. *See* Skidmore v. Swift & Co., 323 U.S. 134 (1944).

225 Lauren B. Edelman, Linda H. Krieger, Scott R. Eliason, Catherine R. Albiston & Virginia Mellema, *When Organizations Rule: Judicial Deference to Institutionalized Employment Structures*, 117 Am. J. Socio. 888, 898, 901–05 (2011); *see* Paul J. DiMaggio & Walter W. Powell, *Introduction*, *in* The New Institutionalism in Organizational Analysis 1, 27–30 (Walter W. Powell & Paul J. DiMaggio eds., 1991).

226 *See* Metcalf et al., *supra* note 24, at 735–36; Ari Ezra Waldman, *Privacy Law's False Promise*, 97 Wash. U. L. Rev. 773, 793–97 (2020).

227 Laura Sartori & Giulia Bocca, *Minding the Gap(s): Public Perceptions of AI and Socio-Technical Imaginaries*, 38 A.I. & Soc'y 443, 443–44 (2023).

228 *See* Calo & Citron, *supra* note 70, at 802–04; Annette Zimmermann, Elena Di Rosa & Hochan Kim, *Technology Can't Fix Algorithmic Injustice*, Bos. Rev. (Jan. 9, 2020), https://www.bostonreview.net/articles/annette-zimmermann-algorithmic-political/ [https://perma.cc/H3DA-HLEQ].

One can already hear other critiques leveled at realist perspectives. The so-called "Collingridge dilemma"[229] looms especially large with efforts to regulate technology. Early in the life of new technologies, we worry we lack the information to regulate them well.[230] But soon enough we lack an equally important asset: sufficient power.[231] The technology is entrenched, creating vested interests in its continuation and even a measure of public comfort that comes with familiarity.[232]

From a realist perspective, however, these concerns should not deter us. We need a different frame and a different type of research going forward if we are to achieve meaningful accountability in the automated state. We need to better understand how AI will change bureaucracies for better and worse, and how law and sociotechnical design and embedded technical capacity can steer them toward the former. That work has already begun in a vital new field at the intersection of computer science, social science, organizational theory, and law. But making that knowledge useable also means being a realist—neither agoniste nor apologist—about the uses and abuses of state power and the ways we have tried, however imperfectly, to use law to constrain it. The stark reality is that judges will soon be making key decisions about the future of the automated state, and those with a deep understanding of new technologies cannot stand on the sidelines hoping for a more thoroughgoing overhaul of American public law. If we do not begin to equip decisionmakers with plausible ways to adapt the current rules, we risk being cut out of the conversation entirely at what is plainly a hinge moment in the modern democratic state's evolution.

---

229 Audley Genus & Andy Stirling, *Collingridge and the Dilemma of Control: Towards Responsible and Accountable Innovation*, 47 RSCH. POL'Y 61, 63 (2018).

230 *Id.*

231 *Id.*

232 DAVID COLLINGRIDGE, THE SOCIAL CONTROL OF TECHNOLOGY 11 (1980); THE GROWING GAP BETWEEN EMERGING TECHNOLOGIES AND LEGAL-ETHICAL OVERSIGHT: THE PACING PROBLEM xvi (Gary E. Marchant et al. eds., 2011); Paul J. DiMaggio & Walter W. Powell, *The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields*, 48 AM. SOCIO. REV. 147 (1983) (discussing organizations and bureaucracies' response to the incentives of their external environments, which leads to convergence around behavior due to risk-aversion and familiarity).