

# ADMINISTRATIVE LAW’S FOURTH SETTLEMENT: AI AND THE CAPABILITY-ACCOUNTABILITY TRAP

*Nicholas Caputo\**

## ABSTRACT

Since 1887, administrative law has navigated a “capability-accountability trap”: technological change forces government to become more sophisticated, but sophistication renders agencies opaque to generalist overseers like the courts and Congress. The law’s response—substituting procedural review for substantive oversight—has produced a sedimentary accretion of requirements that ossify capacity without ensuring democratic control. This Article argues that the Supreme Court’s post-*Loper Bright* retrenchment is best understood as an effort to shrink administration back to comprehensible size in response to this complexification. But reducing complexity in this way sacrifices capability precisely when climate change, pandemics, and AI risks demand more sophisticated governance.

AI offers a different path. Unlike many prior administrative technologies that increased opacity alongside capacity, AI can help build “scrutability” in government, translating technical complexity into accessible terms, surfacing the assumptions that matter for oversight, and enabling substantive verification of agency reasoning. This Article proposes three doctrinal innovations within administrative law to realize this potential: a Model and System Dossier (documenting model purpose, evaluation, monitoring, and versioning) extending the administrative record to AI decision-making; a material-model-change trigger specifying when AI updates require new process; and a “deference to audit” standard that rewards agencies for auditable evaluation of their AI tools. The result is a framework for what this Article calls the “Fourth Settlement,” administrative law that escapes the capability-accountability trap by preserving capability while restoring comprehensible oversight of administration.

---

\* Researcher – Legal Alignment on Frontier AI, Oxford Martin AI Governance Initiative, University of Oxford. Thanks to Jack Boeglin, Molly Brady, Glenn Cohen, Noah Feldman, Gillian Hadfield, Noam Kolt, Peter Salib, and Ben Rolsma for their insightful comments. Thanks to Ryan Copus, Guha Krishnamurthy, Alex Platt, and the faculty of the Johns Hopkins University School of Government and Policy for stress testing these arguments.

## Table of Contents

ABSTRACT.....	1
INTRODUCTION .....	2
<b>I. A TECHNOLOGICAL HISTORY OF ADMINISTRATIVE LAW .....</b>	<b>10</b>
<i>A. Railroads, the Telegraph, and the Birth of Expert Administration (1850–1920)</i> .....	11
<i>B. Mass Society and Statistical Governance (1920–1946)</i> .....	15
<i>C. Computers and Complex Science (1946–1985)</i> .....	19
<b>II. A CRISIS OF ACCOUNTABILITY .....</b>	<b>24</b>
<i>A. The Sedimentary State</i> .....	25
<i>B. Technological Stagnation</i> .....	27
<i>C. Judicial Dismantling for Administrative Simplification</i> .....	31
<b>III. THE FOURTH SETTLEMENT: AI BEYOND THE CAPABILITY–ACCOUNTABILITY TRAP .....</b>	<b>35</b>
<i>A. The Cognitive Infrastructure of Administration</i> .....	35
<i>B. AI as Scrutability Infrastructure for Participation and Oversight</i> .....	42
<i>C. AI as Accountability Infrastructure</i> .....	51
<i>D. Beyond Procedure: AI for Substantive Accountability</i> .....	62
<b>CONCLUSION: CRISIS AS OPPORTUNITY .....</b>	<b>67</b>

### INTRODUCTION

In 1870, the United States had roughly 53,000 miles of railroad track in operation<sup>1</sup> and zero federal railroad regulators.<sup>2</sup> By 1887, operated miles of track had almost tripled to 150,000,<sup>3</sup> and the railroads had transformed the structure of markets and corporations, lengthened supply chains outside of local jurisdictions, and introduced new safety risks and pathways for economic shocks.<sup>4</sup> That year, Congress finally responded to this new technology's new problems, creating the Interstate Commerce Commission (ICC) and so laying the foundations for the modern administrative state.<sup>5</sup>

The ICC, with its expert commission and politically insulated staff,<sup>6</sup> represented both institutional innovation and a Faustian bargain that has structured administrative governance ever since. To solve the novel problems presented by novel technologies, government must become more

<sup>1</sup> U.S. Bureau of the Census, *Historical Statistics of the United States: Colonial Times to 1957* 427–28, ser. Q 15–22 (U.S. Dep't of Commerce 1960) (reporting U.S. railroad mileage over time).

<sup>2</sup> THOMAS K. McCRAW, *PROPHETS OF REGULATION*: CHARLES FRANCIS ADAMS, LOUIS D. BRANDEIS, JAMES M. LANDIS, ALFRED E. KAHN 61–62 (1984).

<sup>3</sup> U.S. Bureau of the Census, *supra* note 1.

<sup>4</sup> ALFRED D. CHANDLER, JR., *THE VISIBLE HAND: THE MANAGERIAL REVOLUTION IN AMERICAN BUSINESS* 79–107 (1977); Bruce Wyman, *The Rise of the Interstate Commerce Commission*, 7 *YALE L.J.* 915, 915–16 (1915).

<sup>5</sup> McCraw, *supra* note 2, at 61–62; STEPHEN SKOWRONEK, *BUILDING A NEW AMERICAN STATE: THE EXPANSION OF NATIONAL ADMINISTRATIVE CAPACITIES, 1877-1920* 4–5 (1982) (“Generally speaking, the expansion of national administrative capacities in America around the turn of the century was a response to industrialism.”).

<sup>6</sup> *Id.* at 62.

scientifically and technologically sophisticated. But this sophistication and its attendant complexification threaten to undermine democratic control because democratic bodies cannot easily understand how scientific government works and so cannot ensure that it is acting in the public interest.<sup>7</sup> The ICC required professional economists, telegraph networks, statistical analysis, and modern filing systems as well as rulemaking and adjudicatory authority to solve the railroads' problems.<sup>8</sup> Its expert commissioners and staff were insulated from political interference, including through for-cause firing protections and partisan balancing requirements, to enable regulators to act without worrying about retaliation.<sup>9</sup> Scientific expertise and the uptake of new technologies combined with a novel legal structure to create modern administration.

But this new structure pushed up against democratic government. Courts lacked the railroad economics expertise to substantively review ICC determinations that a rate was not "reasonable and just."<sup>10</sup> Congress could not specify in advance how to set rates across different contexts and in the face of a changing technology, necessitating legislative vagueness and creating delegation problems.<sup>11</sup> Agency independence might enable the ICC to pursue statutory goals rather than political exigencies, but insulated it from democratic control.<sup>12</sup> Ratemaking, the main task of the ICC, seemed to unconstitutionally transfer the legislative and judicial powers into the executive branch. The ICC needed these powers to accomplish its goals, but each threatened accountability<sup>13</sup> and the creation of "a government of a bureaucratic character alien to our system,"<sup>14</sup> leading to legal challenges.

This brief history of the ICC illustrates what I term the "capability-accountability trap," a pattern that has structured American administrative governance since its inception.<sup>15</sup> Here, "capability"

---

<sup>7</sup> See generally MAX WEBER, *ECONOMY AND SOCIETY: A NEW TRANSLATION* 350–54 (Keith Tribe ed., trans., Harv. U. Press 2019) (1921) (describing the growth of technically trained bureaucracies as a functional response to new technology and to complex social and economic problems, and warning that rationalization and expert administration may be necessary but also threaten democratic self-rule by concentrating authority in domains inaccessible to lay understanding); JAMES M. LANDIS, *THE ADMINISTRATIVE PROCESS* 23–30 (1938) (arguing that modern industrial conditions require governance by scientifically trained experts while acknowledging the resulting tension with democratic accountability); HERBERT A. SIMON, *ADMINISTRATIVE BEHAVIOR* 46–47, 65–67, 87–118 (4th ed. 1997) (explaining that bounded rationality necessitates reliance on expert specialization in complex decision environments, with implications for democratic control over value judgments in decision-making); SHEILA JASANOFF, *THE FIFTH BRANCH: SCIENCE ADVISERS AS POLICYMAKERS* 1–32 (1990).

<sup>8</sup> See GABRIEL KOLKO, *RAILROADS AND REGULATION, 1877–1916*, 57–150 (1976).

<sup>9</sup> McCraw, *supra* note 2, at 62.

<sup>10</sup> Interstate Commerce Act of 1887, ch. 104, 24 Stat. 379 § 1 (Feb. 4, 1887) (repealed and recodified as 49 U.S.C. § 10101 et seq. (Supp. III 1980)).

<sup>11</sup> Richard B. Stewart, *The Reformation of American Administrative Law*, 88 HARV. L. REV. 1667, 1675, 1684–85 (1975).

<sup>12</sup> See Jacob Gersen, *Designing Agencies*, in RESEARCH HANDBOOK ON PUBLIC CHOICE AND PUBLIC LAW 333, 337–38 (Daniel A. Farber & Anne Joseph O'Connell eds., 2010).

<sup>13</sup> Wyman, *supra* note 4, at 532–34; I.L. Sharfman, *The Interstate Commerce Commission: An Appraisal*, 46 YALE L.J. 915, 918–20 (1937).

<sup>14</sup> 285 U.S. 22, 57 (1932).

<sup>15</sup> My framing builds on Professor Adrian Vermeule's "deference dilemma," describing how the background conditions of complexity in administration necessitate broad and vague delegations of power to agencies at the same

refers to the government's ability to process information, coordinate action, and implement solutions at the scale and sophistication that the problems facing it demand. "Accountability" encompasses both democratic control, ensuring that government responds to the public will through direction and oversight, and legality, maintaining rule-of-law constraints on arbitrary power. Both are necessary, because without capability, government will not be able to handle the public's hard problems, and without accountability, it may not choose to do so. However, scholars of administrative law have traditionally conceptualized capability and accountability as trading off against each other, as, historically, increases along one dimension have come with diminishment along the other.<sup>16</sup> If government's capabilities too far exceed accountability, it can grow misaligned with the public interest and arbitrary or abusive in its actions. If accountability is excessive, government will be stifled by oversight and unable to exercise its expertise to solve problems.

Administrative law evolved over time to mediate the capability-accountability trap, picking points along the tradeoff curve where both are sufficient to meet the demands government faces. However, the underlying field of problems and solutions across which the law must operate changes over time as new technologies emerge. New technologies create capability crises that overwhelm existing institutions, like railroad rate discrimination overwhelmed the direct regulatory capacity of legislatures, and upset existing doctrinal settlements. New technologies can also be integrated into government to increase its capabilities and potentially to enable better oversight, as punch card tabulation facilitated the mass information processing necessary to administer Social Security during the Great Depression while also creating a record of the information processed that was scrutable to courts.<sup>17</sup> As such, new technologies can be understood to shift the capability-accountability tradeoff curve while the law picks points along or within it.<sup>18</sup>

---

time that the need for legitimating judicial review pushes back against such delegations. Adrian Vermeule, *The Deference Dilemma*, 31 GEO. MASON L. REV. 619, 619–20 (2024). However, whereas Professor Vermeule focuses on the implications of his dilemma for delegation of interpretive authority, I argue more broadly that the tension between administrative capability and accountability structures the whole of administrative law and also draw the connection to how new technologies change this picture.

<sup>16</sup> See, e.g., Nicholas R. Bednar, *Presidential Control and Administrative Capacity*, 77 STAN. L. REV. 823, 825–32 (2025) (discussing the empirical correlation between independence and capability); Elena Kagan, *Presidential Administration*, 114 HARV. L. REV. 2245, 2352–58 (2001) (framing "the apparent tradeoff between politics and expertise" in administrative decisionmaking as a central problem for institutional design and oversight); Matthew C. Stephenson, *Optimal Political Control of the Bureaucracy*, 107 MICH. L. REV. 53, 53–56 (2008) (taking as a starting point the conventional view that greater bureaucratic insulation can improve expert decisionmaking but "reduces the agency's responsiveness to the preferences of political majorities"); Mathew D. McCubbins, Roger G. Noll, & Barry R. Weingast, *Administrative Procedures as Instruments of Political Control*, 3 J.L. ECON. & ORG. 243, 244–45, 259–61 (1987) (arguing that delegation generates principal-agent problems because agencies may become "more expert" than elected principals; describing administrative procedures as devices to reduce informational asymmetries and reassert democratic/political control).

<sup>17</sup> See Carolyn Puckett, *The Story of the Social Security Number*, 69 SOC. SEC. BULL. 55, 61–63 (2009).

<sup>18</sup> Not all new technologies shift the curve out along both dimensions. For example, advanced statistical models enable better governmental information processing but are incomprehensible to non-expert overseers, so they push the capability dimension out while potentially shifting the curve in on accountability, unless oversight can be restored by some other means. Government's overall ability to solve technological problems requires both sufficient capability and sufficient accountability and each problem sits in some position on the grid defined by what it requires from each

A third characteristic mediates between capability and accountability: “scrutability,”<sup>19</sup> the cognitive tractability of the administrative system for those who must oversee and participate in it. Courts, Congress, and the public can only hold government accountable for what they can understand. When the complexity of administration exceeds the comprehension of its overseers, accountability requires substitutes for substantive review that make government scrutable again.<sup>20</sup> This is precisely what happened across the history of American administrative law. As government became more sophisticated to handle increasingly complex problems, from railroad economics to macroeconomic stabilization to probabilistic risk assessment, generalist overseers lost the ability to evaluate whether agencies were reaching substantively correct conclusions. Administrative law compensated by shifting from substantive to procedural oversight: if courts could not assess whether an EPA risk model was scientifically sound, they could at least verify that the agency had compiled a record, considered alternatives, and provided reasons tied to evidence. Procedural requirements, which courts are cognitively well-equipped to handle, became scrutability proxies, preserving a form of accountability even as the underlying substantive decisions became opaque to non-experts.

Since the establishment of the ICC (and thus the modern administrative state<sup>21</sup>) in 1887, three major technological shifts and doctrinal settlements rebalancing capability and accountability have occurred: the initial formation of the administrative state in response to the disruptions of the Industrial Revolution, its expansion into something like its modern form in the face of the crisis of the Great Depression, and the *Chevron*<sup>22</sup> settlement following the advent of computers and complex science. This history presents three main doctrinal tools of resettlement: reallocating interpretive authority, ratcheting standards of review toward reasoned decision-making on the administrative record, and establishing information-forcing baselines that made agency reasoning visible. Together, these tools allow agencies to retain the expertise, technological sophistication,

---

of these dimensions. If government can't understand or act on a problem, it fails to resolve it, and if government is misaligned to the public interest through capture or some other mechanism, it will also fail.

<sup>19</sup> I use “scrutability” as something like the inverse of James C. Scott’s “legibility.” Where “legibility” means the comprehensibility of the world to regulators, I mean the comprehensibility of regulators to the world. See JAMES C. SCOTT, *SEEING LIKE A STATE* 2–3 (1998). See also Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 *FORDHAM L. REV.* 1085, 1094 (2018) (“inscrutability” as “a situation in which the rules that govern decision-making are so complex, numerous, and interdependent that they defy practical inspection and resist comprehension”); Katherine J. Strandburg, *Rulemaking and Inscrutable Automated Decision Tools*, 119 *COLUM. L. REV.* 1851, 1851–54 (2019).

<sup>20</sup> See Herbert A. Simon, *From substantive to procedural rationality*, in *25 YEARS OF ECONOMIC THEORY* 65 (T. J. Kastelein, S. K. Kuipers, W. A. Nijenhuis, & G. R. Wagenaar eds. 1976).

<sup>21</sup> Professor Jerry L. Mashaw has persuasively argued that administrative law predates the ICC and structured earlier forms of American government. See *generally* JERRY L. MASHAW, *CREATING THE ADMINISTRATIVE CONSTITUTION: THE LOST ONE HUNDRED YEARS OF AMERICAN ADMINISTRATIVE LAW* (2012). However, those earlier forms were much narrower than the modern administrative state, and today’s administrative law doctrine looks to the debates over the ICC as its starting point. See Susan E. Dudley, *Milestones in the Evolution of the Administrative State*, 150 *DAEDALUS* 33, 34 (2021) (describing the ICC as the first modern regulatory agency).

<sup>22</sup> Referring to *Chevron U.S.A., Inc. v. Natural Resources Defense Council, Inc.*, 467 U.S. 837 (1984), a landmark case establishing deference to agency interpretations of ambiguous organic statutes.

and independence necessary to achieve their goals while giving courts and legislatures new oversight mechanisms that limit the risk of administrative lawlessness, capture, or abuse. While new technologies like the internet and “narrow AI”<sup>23</sup> have emerged since *Chevron*, government has not become significantly more capable or less accountable through their use,<sup>24</sup> and so, until recently, the *Chevron* settlement has held.

However, with each turn of the technological wheel, government became more complex and inscrutable. Accountability can be preserved either by reviewing the substance or procedure of agency actions or by increasing the extent to which the agency is verifiably aligned to the public interest. Alignment of human interests is difficult<sup>25</sup> and substantive review has been defeated by complexity, so the burden of accountability has increasingly shifted to procedure. Each settlement required new procedural protections to keep government accountable, but over time this process layered new requirements atop existing ones without removing the old. The result is the “scrutability tax:” the accumulated cost of maintaining comprehensibility for external overseers as the underlying complexity of administration grows. Procedure designed to ensure accountability has become a source of ossification,<sup>26</sup> paid in delay, cost, and defensive documentation that channels agency energy toward producing review-proof records rather than solving problems,<sup>27</sup> and turning procedure into a tool of pluralist bargaining among interest groups rather than a mechanism for deploying legitimate technical expertise.<sup>28</sup> The EPA’s 2024 revision of particulate matter standards illustrates the pattern. It took three years to change one number, required the processing of 700,000 comments, and produced a 200-page preamble in what was intended to be a “concise general statement of [...] basis and purpose.”<sup>29</sup> While this shift is a reasonable response to increasingly complex government by boundedly rational humans given our cognitive limitations,<sup>30</sup> the long retreat from substantive oversight has increased the distance between administration and the people.

---

<sup>23</sup> Narrow AI systems are designed and trained for specific, bounded tasks like determining whether a person should be let out on bail. This contrasts with “general” AI that can reason across domains. Frontier AI systems like ChatGPT are increasingly general in this sense. STUART RUSSELL & PETER NORVIG, *ARTIFICIAL INTELLIGENCE: A MODERN APPROACH* 1–5 (4th ed. 2021).

<sup>24</sup> See Jennifer Pahlka, *Testimony on Harnessing AI to Improve Government Services and Customer Experience*, S. COMM. ON HOMELAND SEC. & GOV’T AFFS. (Jan. 10, 2024); U.S. Gov’t Accountability Off., *High-Risk Series: Critical Actions Needed to Urgently Address IT Acquisition and Management Challenges 2–4* (GAO-25-107852, 2025).

<sup>25</sup> Simon, *supra* note 7, at 152–63.

<sup>26</sup> See Thomas O. McGarity, *Some Thoughts on “Deossifying” the Rulemaking Process*, 41 DUKE L.J. 1385, 1385–87 (1992).

<sup>27</sup> See Nicholas Bagley, *The Procedure Fetish*, 118 MICH. L. REV. 345, 352–55, 362–65 (2019).

<sup>28</sup> See Stewart, *supra* note 11, at 1711–12; Pahlka, *supra* note 24; Steven M. Teles, *Kludgeocracy: The American Way of Policy*, NEW AMERICA FOUNDATION (Dec. 2012), [https://static.newamerica.org/attachments/4209-kludgeocracy-the-american-way-of-policy/Teles\\_Sтивен\\_Kludgeocracy\\_NAF\\_Dec2012.d8a805aa40e34bca9e2fecb018a3dcb0.pdf](https://static.newamerica.org/attachments/4209-kludgeocracy-the-american-way-of-policy/Teles_Sтивен_Kludgeocracy_NAF_Dec2012.d8a805aa40e34bca9e2fecb018a3dcb0.pdf).

<sup>29</sup> Reconsideration of the National Ambient Air Quality Standards for Particulate Matter, 89 Fed. Reg. 16,202 (Mar. 6, 2024) (to be codified at 40 C.F.R. pts. 50, 53, 58). For the statutory reference, see 5 U.S.C. § 553(c) (2018).

<sup>30</sup> See Simon, *supra* note 7, at 118–39.

And yet, despite no increase in government capabilities of the kind that traditionally precipitates doctrinal resettlement, the administrative state faces an unprecedented crisis. In a series of recent cases, the Supreme Court has undone foundational elements of modern administrative law. In *Loper Bright v. Raimondo*,<sup>31</sup> the Supreme Court overruled *Chevron*, overturning its rule of deference to expert agencies in interpreting statutory ambiguity. In *West Virginia v. EPA*,<sup>32</sup> the Court expanded the major questions doctrine to require a clear statement from Congress in significant regulatory areas, limiting agency power over key economic and social questions. And in *SEC v. Jarkesy*,<sup>33</sup> it curtailed agency adjudication when the government seeks civil penalties, instead requiring judicial process under the Seventh Amendment. The Court has continued to operationalize this retrenchment through a set of follow-on decisions that expand the timing, availability, and granularity of judicial control over regulatory programs.<sup>34</sup> Together, these decisions significantly curtail agency capabilities.

The Supreme Court's recent doctrinal unsettlement can be understood as a response to the scrutability crisis. Faced with agencies operating in domains that exceed judicial comprehension, deploying methodologies that judges cannot evaluate, and producing records that require expertise to interpret, the Court has opted to shrink the administrative state back to comprehensible size. *Loper Bright* reclaims interpretive authority for courts.<sup>35</sup> The major questions doctrine reserves significant policy choices for Congress. *Jarkesy* channels enforcement into judicial fora with familiar procedural forms. These moves sacrifice capability for scrutability, and so seek a different kind of accountability than that provided by procedure alone. If courts and Congress cannot understand what agencies are doing, agencies must do less.

This doctrinal dismantling is occurring just as a new set of complex problems is emerging that will require more capable government, not less. Federal climate governance rests on statutes passed before the problem was understood.<sup>36</sup> Pandemics may become more frequent.<sup>37</sup> Advanced AI presents new kinds of risks for which government is unequipped,<sup>38</sup> from deepfakes eroding public

---

<sup>31</sup> 603 U.S. 369 (2024).

<sup>32</sup> 597 U.S. 697 (2022).

<sup>33</sup> 603 U.S. 109 (2024).

<sup>34</sup> See, e.g., *Corner Post, Inc. v. Bd. of Governors of the Fed. Reserve Sys.*, No. 22-1008, slip op. at 1 (U.S. July 1, 2024) (holding an APA claim “does not accrue” for § 2401(a) until the plaintiff is injured by final agency action); *Ohio v. EPA*, Nos. 23A349, 23A350, 23A351 & 23A384, slip op. at 2 (U.S. June 27, 2024) (rejecting the argument that applicants had to return to the agency before raising an objection the agency had already confronted); *City & Cnty. of San Francisco v. EPA*, No. 23–753, slip op. at 1 (U.S. Mar. 4, 2025) (holding the Clean Water Act does not authorize “end-result” provisions in NPDES permits); *Garland v. Cargill*, No. 22–976, slip op. at 1 (U.S. June 14, 2024) (holding ATF exceeded statutory authority in classifying bump stocks as “machineguns”).

<sup>35</sup> Whether courts actually have the capacity to interpret complex organic statutes they now claim is unclear. See Vermeule, *supra* note 15, at 624–25.

<sup>36</sup> See Jody Freeman & David B. Spence, *Old Statutes, New Problems*, 163 U. PENN. L. REV. 1, 20–22 (2014).

<sup>37</sup> B. Adam Williams, Charles H. Jones, Verna Welch, & Jane M. True, *Outlook of pandemic preparedness in a post-COVID-19 world*, 8 NPJ VACCINES 1–2 (2023).

<sup>38</sup> See Markus Anderljung et al., *Frontier AI Regulation: Managing Emerging Risks to Public Safety*, CENTRE FOR THE GOVERNANCE OF AI (2023), [https://cdn.governance.ai/Frontier\\_AI\\_Regulation\\_Managing\\_Emerging\\_Risks.pdf](https://cdn.governance.ai/Frontier_AI_Regulation_Managing_Emerging_Risks.pdf).

trust<sup>39</sup> to AI-enabled bioterrorism<sup>40</sup> to complicated multi-agent risks.<sup>41</sup> The Supreme Court's retrenchment<sup>42</sup> may reflect genuine concern about administrative complexity.<sup>43</sup> But reverting to simpler governmental forms will not solve society's growing problems.

Neither is the solution a simple return to *Chevron*-style deference and the accumulation of still more procedural sediment. That path leads only to further ossification and deeper alienation of administration from democratic control. If the capability-accountability trap is to be escaped rather than merely navigated, something must change in the underlying relationship between governmental sophistication and public comprehension.

This Article makes two contributions in response to this crisis. In Parts I and II, it offers a historical and doctrinal account of how administrative law has repeatedly mediated the coevolution of governance and technology to manage the capability-accountability trap. Drawing on the three doctrinal settlements of the twentieth century, it demonstrates that law can work to make government both powerful enough to handle society's new technological problems and accountable enough to stay legitimate and responsive to the needs of the people. Importantly, however, it also shows that with each doctrinal settlement, the complexity of government grew, and administrative law could provide only a partial solution to the alienating effects of this complexification. This framework provides a key for understanding why the Supreme Court has recently attacked the administrative state and suggests that, even as new governance problems from climate change to AI emerge, the complexity crisis that has driven the Court to take its recent actions will persist unless administrative capability and accountability can be reconciled.

Part III argues that AI offers a way out of this fundamental trap, not by adding another layer of complexity but by changing the relationship between administrative sophistication and democratic comprehensibility. Previous technologies have enhanced government's capabilities while also increasing inscrutability because of tradeoffs and resource bottlenecks intrinsic to human

---

<sup>39</sup> See generally Robert Chesney & Danielle K. Citron, *Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security*, 107 CAL. L. REV. 1753 (2019).

<sup>40</sup> See *Report on Reducing the Risks at the Intersection of Artificial Intelligence and Chemical, Biological, Radiological, and Nuclear Threats*, U.S. DEPARTMENT OF HOMELAND SECURITY (Apr. 26, 2024), [https://www.dhs.gov/sites/default/files/2024-06/24\\_0620\\_cwmd-dhs-cbrn-ai-co-report-04262024-public-release.pdf](https://www.dhs.gov/sites/default/files/2024-06/24_0620_cwmd-dhs-cbrn-ai-co-report-04262024-public-release.pdf); David Manheim et al., *Understanding the Theoretical Limits of AI-Enabled Pathogen Design*, RAND (Sep. 24, 2025), [https://www.rand.org/pubs/research\\_reports/RRA4087-1.html](https://www.rand.org/pubs/research_reports/RRA4087-1.html).

<sup>41</sup> Yonadav Shavit et al., *Practices for Governing Agentic AI Systems*, OPENAI (Dec. 14, 2023), <https://cdn.openai.com/papers/practices-for-governing-agentic-ai-systems.pdf>; Lewis Hammond et al., *Multi-Agent Risks from Advanced AI*, COOPERATIVE AI FOUNDATION (Feb. 19, 2025), <https://www.cooperativeai.com/post/new-report-multi-agent-risks-from-advanced-ai>; Alan Chan et al., *Infrastructure for AI Agents*, arXiv:2501.10114 (Jun. 19, 2025).

<sup>42</sup> Gillian E. Metzger, *Foreword: 1930s Redux: The Administrative State Under Siege*, 131 HARV. L. REV. 1, 1–6 (2017).

<sup>43</sup> It may also be politically motivated by a sense that the administrative state is opposed to conservative interests or by more traditional “capture”-style concerns. For more discussion of these possibilities, see *infra* Part II.C.

reasoning about complex problems.<sup>44</sup> AI possesses characteristics that could improve scrutability by shifting complexity to where it is needed, while enabling courts, legislatures, and the public to better understand administration.

Part III develops this possibility. Subpart A draws on Herbert Simon's cognitive model of administration to explain why scrutability has historically declined as administration has complexified, establishing a typology of administrative action that clarifies where AI intervention could be most effective. Subpart B argues that AI can serve as *scrutability infrastructure* for participation and oversight, including by translating complex rulemaking proposals into accessible terms, generating individualized explanations in adjudication, and surfacing the relevant assumptions embedded in technical analysis for congressional oversight. These interventions reduce the cognitive costs of understanding agency action, restoring meaningful participation to those currently excluded by procedural complexity. Subpart C proposes an accountability infrastructure for AI-assisted administration, including an expanded administrative record (the "Model and System Dossier"), doctrine for when AI system changes trigger procedural requirements, and a new review posture based on deference to audit rather than expertise.

Finally, Subpart D argues that AI's distinctive characteristics can enable forms of substantive accountability unavailable for human-staffed administration, reducing rather than increasing the need for procedural safeguards and so helping combat ossification. Technical alignment, including reinforcement learning from human feedback,<sup>45</sup> constitutional AI,<sup>46</sup> and related techniques,<sup>47</sup> can shape AI systems' objectives more directly than the coarse tools of incentives and monitoring that align human employees to organizational goals.<sup>48</sup> Advances in interpretability<sup>49</sup> create possibilities for examining AI reasoning that do not exist for human cognition because we may one day be able to examine the decision pathways of an AI system to verify that its outputs reflect appropriate inputs and reasoning processes. And AI systems may be able to help oversee other AI systems at scale, extending accountability to the millions of individual determinations that currently escape human review entirely.<sup>50</sup>

---

<sup>44</sup> Simon, *supra* note 7, at 92–106.

<sup>45</sup> Long Ouyang et al., *Training Language Models to Follow Instructions with Human Feedback*, arXiv:2203.02155 (Mar. 4, 2022) (introducing RLHF for instruction-following and reporting improvements in human preference, truthfulness, and toxicity).

<sup>46</sup> Yuntao Bai et al., *Constitutional AI: Harmlessness from AI Feedback*, arXiv:2212.08073 (Dec. 15, 2022) (principle-based training with self-critique and "RL from AI feedback" (RLAIF) to scale supervision and improve harmlessness).

<sup>47</sup> See, e.g., Melody Y. Guan et al., *Deliberative Alignment: Reasoning Enables Safer Language Models*, arXiv:2412.16339 (Jan. 8, 2025) (describing alignment techniques in reasoning systems).

<sup>48</sup> Avinash Dixit, *Incentives and Organizations in the Public Sector: An Interpretative Review*, 37 J. HUM. RES. 696, 696–700, 708–12 (2002).

<sup>49</sup> See Finale Doshi-Velez & Been Kim, *Towards A Rigorous Science of Interpretable Machine Learning*, arXiv:1702.08608 (Mar. 2, 2017).

<sup>50</sup> See Samuel R. Bowman et al., *Measuring Progress on Scalable Oversight for Large Language Models*, arXiv:2211.03540 (Nov. 11, 2022).

These capabilities suggest that the trajectory toward ever-more-procedural oversight might be reversed. If AI systems can be verified to be pursuing appropriate objectives through appropriate reasoning, the procedural proxies that substituted for such verification become less necessary. The scrutability problem that drove proceduralization might be addressed not by adding new procedures but by restoring the substantive oversight that procedure was always a second-best replacement for. The result is not more procedure piled on existing procedure, but a rebalancing toward substantive verification, getting stronger capability and accountability with reduced ossification.

Decades ago, Herbert Simon wrote that “[t]he central concern of administrative theory is with the boundary between the rational and the nonrational aspects of human social behavior. Administrative theory is peculiarly the theory of intended and bounded rationality—of the behavior of human beings who satisfice because they have not the wits to maximize.”<sup>51</sup> But these foundational constraints of human rationality are being transformed by the advent of AI,<sup>52</sup> and this Article seeks to provide a framework for integrating AI systems with human minds in ways that can handle the difficult problems of contemporary governance. The argument that follows is not prediction. Whether AI develops along the trajectory that would make this settlement possible, whether political conditions permit its adoption, and whether agencies and courts would implement it well are all uncertain. But the current moment demands more than diagnosis of administrative law’s failures. It demands a constructive vision of what legitimate and effective administration could look like in an era of technological transformation. The previous settlements were not inevitable. They were built by reformers who saw that new technological tools could mediate the tension between capability and accountability in government in new ways. This Article aims to do the same for the settlement that must come next. The question ahead is not whether AI will transform government—transformation is inevitable given today’s problems. The question is whether democratic societies can harness AI to break the capability-accountability trap, creating government that is both sophisticated enough for today’s challenges and comprehensible enough for democratic control.

## I. A TECHNOLOGICAL HISTORY OF ADMINISTRATIVE LAW

The American administrative state has evolved through recurring cycles of technological disruption and legal adaptation. This Part traces three such cycles—the railroads, the New Deal, and computerized risk regulation—to show how technological shocks increased administrative capability while pushing accountability toward procedural instead of substantive review. Across cycles, doctrinal “settlements” recur to manage the capability-accountability trap using the tools

---

<sup>51</sup> Simon, *supra* note 7, at 50.

<sup>52</sup> Simon was also of course a pioneer in the study of AI and attentive to the effects of computerization on administration. *See, e.g.*, Herbert A. Simon, *From Substantive to Procedural Rationality*, in *25 YEARS OF ECONOMIC THEORY* (T. J. Kastelein et al., eds.) (1976) (arguing that human and machine cognition are based on the same fundamental components, though here as symbolic systems rather than neural networks).

of allocation of authority, standards of review, and transparency-forcing. Underlying these tools is a recurring dynamic: as administrative action grew more technically complex, it became less scrutable to generalist overseers, that is less tractable for courts and legislators operating under the constraints of bounded rationality. Oversight compensated by shifting from substance to procedure, substituting questions courts could answer (did the agency follow required steps?) for questions they could not (did the agency reach correct conclusions?). Understanding this pattern is essential for two reasons. First, it reveals that our current administrative framework is not the product of abstract constitutional theory but rather accumulated responses to concrete technological problems. Second, it exposes why each doctrinal settlement has proven temporary: new technologies have presented new problems that force adjustments in the administrative state and add to the complexity of administration and its oversight.

Fully developing each historical case, particularly the political contestation that drives and shapes lawmaking,<sup>53</sup> exceeds the scope of this Article, but the pattern of capability crisis, government complexification, accountability gap, and doctrinal settlement clearly emerges across them. Importantly, each step in the cycle also represents an increase in the overall complexity of administration as scientific problems have demanded ever more sophisticated government, increasing the burdens on the administrative settlements and distancing administration from democratic control by introducing more procedural intermediation.

#### *A. Railroads, the Telegraph, and the Birth of Expert Administration (1850–1920)*

On August 26, 1871, a speeding locomotive plowed into a stationary passenger train at Revere station near Boston, Massachusetts.<sup>54</sup> Twenty-nine passengers died in the wreckage. The disaster emerged from a cascade of failures exemplifying the governance crisis railroads had created. Although the railroad was found liable for negligence, an investigation found that “the corporation had broken no existing law.”<sup>55</sup> Traditional legal mechanisms struggled to address the systemic

---

<sup>53</sup> The ICC emerged not mechanically from railroad complexity but from decades of agrarian mobilization, Granger legislation, and intraparty conflict over the terms of federal regulation. See ELIZABETH SANDERS, *ROOTS OF REFORM: FARMERS, WORKERS, AND THE AMERICAN STATE, 1877–1917* 179–213 (1999). The New Deal agencies reflected coalition politics and constitutional brinkmanship as much as functional demands. See STEPHEN SKOWRONEK, *THE POLITICS PRESIDENTS MAKE: LEADERSHIP FROM JOHN ADAMS TO BILL CLINTON* 287–324 (1997). And the *Chevron* settlement was shaped by deregulatory politics and judicial ideology, not only by the cognitive limits of courts. See Thomas W. Merrill, *The Story of Chevron: The Making of an Accidental Landmark*, 66 ADMIN. L. REV. 253, 275–82 (2014). This Article foregrounds the technological dimension not because politics was secondary but because the relationship between technological change and institutional adaptation has received less systematic attention in administrative law scholarship, aiming for “pattern identification” rather than full explanation. KAREN ORREN & STEPHEN SKOWRONEK, *THE SEARCH FOR AMERICAN POLITICAL DEVELOPMENT* 7–8 (2004).

<sup>54</sup> McCraw, *supra* note 2, at 26–28.

<sup>55</sup> *Id.*

complexity of railroad operations.<sup>56</sup> The tools of pre-industrial governance had become inadequate for rapidly developing industrial problems.<sup>57</sup>

This inadequacy reflected a fundamental capability crisis. Between 1850 and 1890, railroad track in operation expanded from 9,000 to 166,000 miles,<sup>58</sup> while the telegraph enabled instantaneous communication across the country.<sup>59</sup> Together, these technologies knit local markets into national ones and created new problems and forms of economic power that existing institutions could not comprehend or control, as the railroads demonstrate.<sup>60</sup> Passengers, passers-by, and employees were killed and injured in large numbers by speeding trains, and derailment catastrophes like the one at Revere occurred with some frequency.<sup>61</sup> Rate discrimination by the railroads harmed travelers and producers, but state legislatures and courts lacked the sophistication to determine what setting “just and reasonable”<sup>62</sup> rates for a particular railroad might mean. To effectively regulate a railroad, the regulator had to ascertain the value and volume of different commodities, if there was water competition, the distance of hauls, and the railroad’s need for return on capital investments. Rate decisions crossed jurisdictional lines, such that a railroad’s intrastate rates in Illinois could devastate Iowa farmers shipping to Chicago, but state governments lacked the authority to regulate outside their jurisdiction after the Supreme Court’s decision in the *Wabash* case.<sup>63</sup> Discriminatory pricing that favored long-distance shippers over local ones could shift patterns of economic development.<sup>64</sup> Aggressive rate regulation might protect local shippers but drive railroads to recoup losses through higher interstate rates beyond the state’s control, while lenient regulation invited exploitation of captive local markets to subsidize competitive long-haul routes.

This governmental capability crisis eventually precipitated a response, in the form of the expert commission. Massachusetts created the first expert railroad commission in 1869, staffed by specialists who could comprehend railroad operations using the industry’s own tools of statistical analysis, standardized accounting, and telegraph communication.<sup>65</sup> As Professor Thomas McCraw writes, this “agency began life as an institutional adaptation called into being because, after a generation of uninterrupted growth, the railroad system had become too much for the old legal machinery to handle.” The old regulatory apparatus was unable to match the “unprecedented scale

---

<sup>56</sup> Wyman, *supra* note 4, at 531.

<sup>57</sup> The pace of technological change began to increase during the nineteenth century as the technologies of the Industrial Revolution disseminated around the world. See JOEL MOKYR, *THE LEVER OF RICHES* 82–83 (1990).

<sup>58</sup> U.S. Bureau of the Census, *supra* note 1.

<sup>59</sup> Chandler, *supra* note 4, at 79–107.

<sup>60</sup> *Id.*

<sup>61</sup> McCraw, *supra* note 2, at 25–26.

<sup>62</sup> Interstate Commerce Act of 1887, ch. 104, 24 Stat. 379 § 1 (Feb. 4, 1887) (repealed and recodified as 49 U.S.C. § 10101 et seq. (Supp. III 1980)).

<sup>63</sup> 118 U.S. 557, 564–66, 577–79 (1886).

<sup>64</sup> See Felix L. Barton, *Economic Effects of Discriminatory Freight Rates*, 12 L. & CONTEMP. PROBS. 507, 520–31 (1947).

<sup>65</sup> McCraw, *supra* note 2, at 19–21.

of operations the railroads represented, their financial and technological complexity, and the increasingly important role of transportation in the Massachusetts economy.”<sup>66</sup> In subsequent years, state railroad commissions sprang up in other states with varying degrees of legal capability.<sup>67</sup>

The state commissions set an important precedent for the creation of the ICC in 1887, which sought to provide a national solution to the problems of national railroad governance.<sup>68</sup> Strengthened over time through a series of laws expanding its powers and jurisdiction,<sup>69</sup> the ICC was the first federal regulatory agency with legal protections guaranteeing its independence and quasi-legislative and quasi-judicial powers to set rates and adjudicate cases, require standardized accounting practices, and adjudicate disputes across state lines.<sup>70</sup> It compelled testimony, examined books, and mandated uniform reporting systems that allowed for the first truly national picture of railroad operations. Yet the ICC faced immediate challenges that revealed the persistent capability gap that railroad governance still faced. In particular, its initial powers were severely limited by judicial review often favoring the railroads, characteristic of the broader attitude of *laissez-faire* government that the administrative state would struggle with into the 1930s.<sup>71</sup>

The emergent expert commission form, with its specialist staffs and partial insulation from politics, increased government capability but threatened accountability. Should unelected commissioners make rules and adjudicate cases so deeply affecting American life? How would individual rights be protected from harmful or arbitrary agency actions? When the ICC set rates, affected shippers had little recourse if commissioners misunderstood market conditions because judges could not substantively evaluate the agency’s economic analysis, legislators met too infrequently to effectively intervene, and voters could not remove appointed commissioners. Expertise and independence threatened the traditional mechanisms of democratic accountability, from elections to judicial review to legislative oversight.

The legal cases immediately brought against the ICC upon its founding illustrate these worries and demonstrate how courts struggled with balancing the need for expert government to handle the railroads with legitimacy and accountability. Early cases went badly for the ICC, “reduc[ing] it, by the late 1890s, to a mere collector of data.”<sup>72</sup> *ICC v. Cincinnati, N.O. & Texas Pacific Ry. Co.*,<sup>73</sup> decided in 1897, is characteristic: the ICC, set up to resolve the problem of railroad rate

---

<sup>66</sup> *Id.* at 20.

<sup>67</sup> *Id.* at 57.

<sup>68</sup> *Id.* at 61–62.

<sup>69</sup> *Id.* at 62–63.

<sup>70</sup> Interstate Commerce Act of 1887, ch. 104, § 11, 24 Stat. 379, 383–84 (creating a five-member commission; fixed terms; bipartisan limitation; and for-cause removal).

<sup>71</sup> Robert L. Rabin, *Federal Regulation in Historical Perspective*, 38 STAN. L. REV. 1189, 1208–35 (1986) (situating the ICC within Populist and Progressive conflicts, particularly around the Supreme Court).

<sup>72</sup> McCraw, *supra* note 2, at 62.

<sup>73</sup> 167 U.S. 479 (1897).

discrimination, sought to do so in part by setting prospective national railroad rates pursuant to what it understood to be a grant of authority from Congress in the Interstate Commerce Act of 1887 to set “reasonable rates.”<sup>74</sup> However, the Supreme Court held that the power to set rates was a legislative power that Congress could only constitutionally delegate through explicit language, which it had not done.<sup>75</sup> As such, the ICC had acted in excess of its authority in setting rates and enjoining the plaintiff railroad to follow them,<sup>76</sup> despite these actions being taken to address the very regulatory problem that it had been established to resolve.

But the problems of the railroads did not go away, and between 1897 and 1920 a doctrinal settlement emerged that managed, though did not resolve, these tensions.<sup>77</sup> This doctrinal settlement established the three characteristic tools through which administrative law handles the capability-accountability trap. First, *allocation of authority*: Congress responded to *ICC v. Cincinnati* by expressly granting rate-setting power through the Hepburn Act of 1906<sup>78</sup> and Mann-Elkins Act of 1910.<sup>79</sup> Courts, meanwhile, began distinguishing legal questions (their province) from substantive technical analysis (the agency’s domain).<sup>80</sup> Second, *procedural review*: In *ICC v. Louisville & Nashville Railroad*, a 1913 rate discrimination case, the Supreme Court held that while courts could not review an agency’s substantive conclusions in an adjudication, they could review the agency’s record of administrative proceedings to determine whether the agency had acted based on sufficient evidence and with sufficient procedural protections.<sup>81</sup> Taken with *Illinois Central*’s injunction that courts should not question the “wisdom” of agency actions<sup>82</sup> unless “so arbitrary and unreasonable as to render it invalid,”<sup>83</sup> these cases created mechanisms of procedural review that compensated for the declining ability of courts to substantively review the actions of expert agencies. Third, *information-forcing*: In *ICC v. Goodrich Transit Co.*, the Court upheld ICC orders requiring carriers to adopt uniform accounting and submit comprehensive reports, making the railroads more comprehensible to government.<sup>84</sup> Furthermore, under *Louisville & Nashville Railroad*, the ICC was required to provide information to regulated parties, holding hearings and

---

<sup>74</sup> 167 U.S. at 485–88.

<sup>75</sup> 167 U.S. at 505–11. See also Thomas W. Merrill, *Rethinking Article I, Section I: From Nondelegation to Exclusive Delegation*, 104 COLUM. L. REV. 2097, 2111 (2004).

<sup>76</sup> 167 U.S. at 511–12.

<sup>77</sup> Professor Thomas W. Merrill specifically dates this appellate review model to judicial review of ICC orders around 1910. Thomas W. Merrill, *Article III, Agency Adjudication, and the Origins of the Appellate Review Model of Administrative Law*, 111 COLUM. L. REV. 939, 942 (2011).

<sup>78</sup> Ch. 3591, 34 Stat. 584 (1906). See McCraw, *supra* note 2, at 62.

<sup>79</sup> Ch. 309, 36 Stat. 539 (1910). See McCraw, *supra* note 2, at 62–63.

<sup>80</sup> 215 U.S. 452, 464–66 (1910). Unless the agency’s action was so arbitrary or unreasonable that it was invalid. 215 U.S. at 471.

<sup>81</sup> 227 U.S. 88, 92–93 (1913).

<sup>82</sup> 215 U.S. at 470.

<sup>83</sup> 215 U.S. at 452.

<sup>84</sup> 224 U.S. 194, 212–16 (1912).

making decisions based on written records to enable review.<sup>85</sup> Constitutional barriers to federal regulation also began to diminish, enabling a national administrative state.<sup>86</sup>

Taken together, this first settlement, with its doctrinal framework of allocation of authority, procedural judicial review, and information-forcing, laid the foundations for administrative law. Governmental capacity rose and accountability survived through these doctrinal mechanisms, but the rail settlement was imperfect. Statutory delegation could authorize expert government but not ensure that it served public rather than private interests. The substitution of procedure for substance meant that courts could ensure that agencies gave reasons for their decisions but not that those reasons were good. Transparency through information-forcing and statistical analysis still required expertise to interpret it. By 1920, the railroad settlement had created government capable of managing these new industrial technologies, but at the price of increasing distance from democratic control. By the end of the decade, the Great Depression demanded another round of governmental innovation.

### *B. Mass Society and Statistical Governance (1920–1946)*

Since the turn of the twentieth century, the United States had completed a transition from a fragmented, regional economy to a national one.<sup>87</sup> The economy was much more productive, complex, and interlinked,<sup>88</sup> making it sensitive to demand shocks. Some new instrumentalities of national economic regulation, including the Federal Trade Commission and Federal Reserve, emerged in this period.<sup>89</sup> However, especially after World War I, a succession of comparatively non-interventionist presidents mostly left the economy to run on its own, believing that the conditions for “a permanent plateau of prosperity” had been established.<sup>90</sup>

In 1929, their hopes were dashed and the Great Depression began.<sup>91</sup> The technologically-enabled mass society that had birthed the Roaring Twenties now magnified and spread the hurt of the Crash around the country. By 1933, more than 13 million workers were unemployed, around a quarter of the total American workforce.<sup>92</sup> The old remedies of retrenchment and economic self-correction failed. Government lacked the tools to comprehend, much less correct, systemic collapse affecting

---

<sup>85</sup> 227 U.S. at 104–06.

<sup>86</sup> See 230 U.S. 352, 381–82 (1913); 234 U.S. 342, 351–52 (1914).

<sup>87</sup> Chandler, Jr., *supra* note 4, at 5.

<sup>88</sup> DAVID A. HOUNSHELL, FROM THE AMERICAN SYSTEM TO MASS PRODUCTION, 1800–1932: THE DEVELOPMENT OF MANUFACTURING TECHNOLOGY IN THE UNITED STATES 289 (1984); Vaclav Smil, *Nitrogen Cycle and World Food Production*, 2 WORLD AGRICULTURE 9 (2011).

<sup>89</sup> Federal Reserve Act, ch. 6, 38 Stat. 251 (1913); Federal Trade Commission Act, ch. 311, 38 Stat. 717 (1914).

<sup>90</sup> ARTHUR M. SCHLESINGER, JR., THE CRISIS OF THE OLD ORDER, 1919-1933 155 (1957).

<sup>91</sup> *Id.* at 159.

<sup>92</sup> *Id.* at 440.

millions simultaneously. The Great Depression was not merely an economic crisis like the business cycle-driven crashes of the previous century but a new kind of capability crisis.

These new problems required mass government, new kinds of expertise, and novel statistical approaches. Macroeconomic stabilization required population-level measures rather than case-by-case relief, overproduction and underconsumption required sectoral intervention, and industry-wide labor disputes called for standardized fact-finding and program rules.<sup>93</sup> Legacy governance institutions were insufficient to handle these new challenges. Courts could only adjudicate backward-looking disputes between specific parties,<sup>94</sup> Congress lacked the technical expertise to calibrate monetary policy or agricultural production quotas to shifting market conditions, and the expert commissions were designed to handle specific sectors rather than coordinate recovery across interdependent industries.<sup>95</sup> State governments, which had traditionally handled relief for the poor, found their treasuries exhausted and their regulatory reach stopped at borders that economic distress freely crossed.

The modern administrative state coalesced and was vastly expanded in the New Deal to address these systemic challenges. More than 20 “alphabet agencies” were created in 1933 and 1934 alone, dealing with problems from rural electrification to agricultural price adjustment to the banking crisis.<sup>96</sup> The New Deal agencies were mostly staffed by specialists, experts in their fields who were able to grapple with the unique challenges each of them presented, but at a new level of scale and statistical instrumentation.

New technologies played a crucial role in enabling the New Deal institutions to operate effectively, as they had enabled the railroad commissions. The Social Security Administration used punch-card systems to register 26 million workers in its first four months and began distributing benefits to 222,000 elderly Americans by 1940.<sup>97</sup> The Agricultural Adjustment Administration’s statistical surveys facilitated production controls that raised farm income by 50% between 1932 and 1935, rescuing rural America from collapse.<sup>98</sup> The Securities and Exchange Commission’s mandatory disclosure regime and market surveillance restored enough confidence that new securities issues

---

<sup>93</sup> See DAVID M. KENNEDY, *FREEDOM FROM FEAR: THE AMERICAN PEOPLE IN DEPRESSION AND WAR, 1929–1945* 131–60 (1999); JOHN MAYNARD KEYNES, *THE GENERAL THEORY OF EMPLOYMENT, INTEREST AND MONEY* 246–64 (1936).

<sup>94</sup> See Lon L. Fuller, *The Forms and Limits of Adjudication*, 92 HARV. L. REV. 353, 363–70, 394–404 (1978) (explaining how adjudication’s party-presentation form limits its ability to address polycentric, system-wide problems and complex repercussions).

<sup>95</sup> See CONG. RSCH. SERV., *ECONOMICS OF FEDERAL RESERVE INDEPENDENCE* 9–10 (RL31056, Apr. 17, 2007) (explaining that because the legislative branch is not positioned to exercise day-to-day control over monetary policy, delegation structures tend to place operational control in expert agencies).

<sup>96</sup> See Bob Clark, *FDR, Archivist: The Shaping of the National Archives*, 38 PROLOGUE MAG. (Nat’l Archives, 2006).

<sup>97</sup> See Carolyn Puckett, *Administering Social Security: Challenges Yesterday and Today*, 70 SOC. SEC. BULL. 27, 29 (2010); *Historical Background and Development of Social Security*, SOC. SEC. ADMIN. (n.d.), <https://www.ssa.gov/history/briefhistory3.html>.

<sup>98</sup> See E.M. Brooks, *As We Recall: The History of Agricultural Estimates, 1933 to 1961*, U.S. DEP’T OF AGRIC. NAT’L AGRIC. STAT. SERV. (1977).

rose from near zero in 1933 to \$3 billion by 1936.<sup>99</sup> Tabulation machines and mechanical calculators enabled processing of the masses of data collected by the new agencies.<sup>100</sup> Mass communication technologies like the radio, the telephone, and microfilm facilitated coordination within government and between it and society, as well as mass record-keeping.<sup>101</sup> These agencies operated at population scale, using statistical methods and modern technologies to transform abstract economic aggregates into concrete improvements in human welfare.

But, as had occurred in the railroad era, the scale and power of these new agencies led to concerns that they were growing unaccountable. Congress had delegated power to the agencies, often with little detail about how they should use it. The agencies began to be described as a “fourth branch” outside of the constitutional bounds of American government.<sup>102</sup> Substantive oversight by Congress and the courts became harder again. For example, when the Agricultural Adjustment Administration set crop quotas affecting millions of farmers, no individual could easily understand or challenge the sector-wide view and statistical models underlying production targets. Social Security’s punch-card systems processed citizens as data points, transforming democratic subjects into statistical objects who might have little sense of why they received the governmental treatment that they did. The demands of independence, expertise, and technological sophistication again created a gap between accountability and capability in American government.

The initial accountability pushback occurred in the courts. In *Panama Refining*<sup>103</sup> and *Schechter Poultry*,<sup>104</sup> two cases from 1935, the Supreme Court struck down parts of the National Industrial Recovery Act on the basis that it contained unconstitutional delegations of legislative power and also exceeded congressional authority under the Commerce Clause. This “nondelegation” doctrine echoed earlier concerns about the delegation of legislative power in *ICC v. Cincinnati* but went further in its limitations on congressional delegation.<sup>105</sup> In *Carter v. Carter Coal Co.* and *United States v. Butler*, the Court limited Congress’s basic ability to regulate federally under the Commerce Clause and Spending Power.<sup>106</sup>

---

<sup>99</sup> SECURITIES AND EXCHANGE COMMISSION, THIRD ANNUAL REPORT (FISCAL YEAR ENDED JUNE 30, 1937) 12, (reporting “new securities proposed to be offered for cash” of approximately \$2.84 billion during the fiscal year).

<sup>100</sup> Puckett, *supra* note 17, at 61–63.

<sup>101</sup> See, e.g., Franklin D. Roosevelt, *Fireside Chat on Banking*, AMER. PRES. PROJ. (Mar. 12, 1933), <https://www.presidency.ucsb.edu/node/207762>; see generally Michael A. Cronin, *Fifty Years of Operations in the Social Security Administration*, 48 SOC. SEC. BULL. 14, 17–26 (1985) (describing the use of different technologies in social security management).

<sup>102</sup> THE PRESIDENT’S COMM. ON ADMINISTRATIVE MANAGEMENT, ADMINISTRATIVE MANAGEMENT IN THE GOVERNMENT OF THE UNITED STATES 30 (1937); *FTC v. Ruberoid Co.*, 343 U.S. 470, 487 (1952) (Jackson, J., dissenting) (“[A]dministrative agencies have become a veritable fourth branch of the Government”).

<sup>103</sup> 293 U.S. 388 (1935).

<sup>104</sup> 295 U.S. 495 (1935).

<sup>105</sup> See 293 U.S. at 430 (establishing the nondelegation doctrine).

<sup>106</sup> 298 U.S. 238, 303–08 (1936); 297 U.S. 1, 68–69 (1936).

Despite this pushback, the elected branches demanded the administrative state, leading to the threat of court packing and the “switch in time” of 1937.<sup>107</sup> That year, in *West Coast Hotel Co. v. Parrish*, the Court upheld a Washington minimum wage law<sup>108</sup> and in *NLRB v. Jones & Laughlin Steel* it took a relatively expansive view of the Commerce Clause’s limitations on federal regulatory power, coming to terms with the New Deal and its administrative framework.<sup>109</sup> By 1944, the Court was willing in *Yakus v. United States* to authorize even actual price controls based on broad delegations of congressional power that would not have survived the nondelegation test established a decade prior.<sup>110</sup>

But worries of government overreach and abuse persisted, especially once the emergencies of the Great Depression and World War II receded. The courts had retreated from substantive oversight, even as the administrative state pressed against constitutional constraints. Statistical governance made it difficult to ensure due process to the millions now subjected to regulation, and courts lacked capacity to review every breach given the rate and complexity of agency actions.<sup>111</sup> A new doctrinal framework was needed. The solution would be to further substitute procedural regularity for substantive accountability—if courts couldn’t evaluate mass statistical decisions, they could at least ensure proper processes had been used.

The Administrative Procedure Act (APA)<sup>112</sup> was passed in 1946 to provide the foundation for this new settlement. This “super-statute”<sup>113</sup> structures administrative law to this day, applying to all agencies except to the extent that their authorizing statute supersedes it.<sup>114</sup> Its key provisions expand on the allocation, review, and information-forcing dimensions of the railroad era’s administrative settlement. First, the APA provides the basic structure for the roles played by administrative agencies and by courts. Agency actions, now statutorily classified as either rulemakings or adjudications, are subject to a baseline expectation of judicial review with specific standards to govern it.<sup>115</sup> Second, each of these kinds of agency action must follow an enumerated set of formal and informal procedures, including decision-making on a record, which provide the basis for procedural review of agency action.<sup>116</sup> Third, the APA established procedural transparency requirements for agency actions that, when combined with caselaw holding that

---

<sup>107</sup> See 2 BRUCE ACKERMAN, *WE THE PEOPLE: TRANSFORMATIONS* 279–344 (1998).

<sup>108</sup> 300 U.S. 379, 391–97 (1937).

<sup>109</sup> 301 U.S. 1, 37 (1937).

<sup>110</sup> 321 U.S. 414, 423–27 (1944).

<sup>111</sup> See Jerry L. Mashaw, *The Management Side of Due Process: Some Theoretical and Litigation Notes on the Assurance of Accuracy, Fairness, and Timeliness in the Adjudication of Social Welfare Claims*, 59 CORNELL L. REV. 772 (1974) (analyzing due process constraints under conditions of mass administration).

<sup>112</sup> Administrative Procedure Act, Pub. L. No. 79-404, 60 Stat. 237 (1946) (codified as amended at 5 U.S.C. §§ 551–559, 701–706).

<sup>113</sup> See generally William N. Eskridge Jr. & John Ferejohn, *The APA as a Super-Statute: Deep Compromise and Judicial Review of Notice-and-Comment Rulemaking*, 98 NOTRE DAME L. REV. 1893 (2023).

<sup>114</sup> 5 U.S.C. § 559.

<sup>115</sup> 5 USC §§ 701–06.

<sup>116</sup> 5 USC §§ 553–58.

agencies had to make decisions actually based on the evidentiary record in front of them, operate as an information-forcing mechanism to bring administration further into the public and judicial view.<sup>117</sup>

Taken together, these elements of the APA allowed the population-level governance necessary to handle the challenges raised by the Great Depression and a national economy and society while preserving accountability and guarding against the worst abuses of administrative power. Rather than restore complete individual-level accountability, impossible as agencies now governed through statistics and needed to make substantive decisions that were difficult for courts to evaluate, it created procedural protections that aim to protect individual rights and seek to incentivize correct agency decisions through mandating process. The APA could not make mass administration comprehensible to the general public but gave it procedural legitimacy instead.<sup>118</sup> Statistical and expert governance were reconciled with the demands of legality. The APA settlement remains the foundation for the administrative state today.

However, the APA settlement again came with a cost of increased complexification and distancing from the public. Statistical governance became permanent, procedural complexity became mandatory, and democratic accountability became increasingly mediated through technical processes. Government grew significantly, in terms of the number of operating agencies, the people employed working in them, the number of rulemakings and adjudications that they undertook, and the expertise required to understand and engage with the agencies and the growing body of administrative law itself. The government had become more complex again, with a profusion of agencies and a federal workforce that reached into the millions. In the decades to come, environmental and safety crises would again demand a technologically upgraded government, destabilizing the APA settlement and paving the way for another cycle of unsettlement.

### *C. Computers and Complex Science (1946–1985)*

The New Deal had mostly focused on concrete, population-level harms like unemployment or labor abuses, but new and more diffuse problems would soon emerge. The new environmental regulation of the 1960s and 70s responded to problems like water, air, and soil pollution which were being worsened by synthetic pesticides and similar inventions.<sup>119</sup> These problems demanded more sophisticated scientific understanding, measurement, and prediction than earlier forms of

---

<sup>117</sup> See 5 U.S.C. § 552(a)(1)–(2) (publication/availability requirements); *id.* § 706 (record-based review standards); *SEC v. Chenery Corp. (Chenery I)*, 318 U.S. 80, 87 (1943) (agency action must stand or fall on the grounds invoked by the agency in the record).

<sup>118</sup> See JERRY L. MASHAW, *DUE PROCESS IN THE ADMINISTRATIVE STATE* 1–31 (1985) (explaining proceduralized legality as the central currency of legitimacy in mass administration).

<sup>119</sup> See Jack Lewis, *A Historical Perspective on Environmental Regulations*, EPA J. (Mar. 1988), <https://www.epa.gov/archive/epa/aboutepa/looking-backward-historical-perspective-environmental-regulations.html>.

administration could provide. Regulators had to identify issues in complex biological systems, anticipate secondary effects and tradeoffs, and design implementation frameworks that could operationalize scientific judgments at scale.<sup>120</sup>

Beginning in the early 1960s, Congress passed laws obliging agencies to resolve these scientific problems. The Clean Air Act,<sup>121</sup> Clean Water Act,<sup>122</sup> National Environmental Policy Act (NEPA),<sup>123</sup> Occupational Safety and Health Act (OSHA),<sup>124</sup> and others presented new kinds of questions. What concentration of lead endangers public health? What is a “significant” carcinogenic risk in the workplace? What probability of containment failure is acceptable at a nuclear facility? Congress could not legislate specific answers, both because it lacked substantive expertise and because if it tried to legislate specifically, new evidence or shifting technology might render the statutory answer obsolete in practice but mandated in law. Government faced problems as statistical probabilities, environmental harms measurable only through complex instrumentation, and systemic risks comprehensible only through computer simulation, but could not fully turn governance over to agencies because of the normative questions underlying scientific government.<sup>125</sup>

Agencies had to adopt new technologies to respond to these problems. Statistical models of complex systems including atmospheric transport, exposure-response curves, and probabilistic risk assessment were needed to estimate harms and regulatory impacts.<sup>126</sup> Building and analyzing large datasets required new techniques of measurement and the expansion of computing capacity. The Environmental Protection Agency implemented dispersion models and multi-pollutant inventories.<sup>127</sup> The Nuclear Regulatory Commission piloted probabilistic risk assessment at the scale of  $10^{-6}$  chance of harm occurring.<sup>128</sup> The Securities and Exchange Commission began electronic market surveillance that could keep up with computerized trading.<sup>129</sup> These models created regulatory facts through computerized analysis. When EPA’s atmospheric transport models predicted pollution concentrations, their outputs were legally binding despite resting on layers of contested modeling assumptions. Computing both increased and limited the scrutability of administrative action. Runs of models, parameter files, and validation studies could be filed,

---

<sup>120</sup> See Sheila Jasanoff, *Science, Politics, and the Renegotiation of Expertise at EPA*, 7 OSIRIS 194, 197–205 (1992).

<sup>121</sup> Clean Air Act § 109(b)(1), 42 U.S.C. § 7409(b)(1).

<sup>122</sup> Clean Water Act § 304(a)(1), 33 U.S.C. § 1314(a)(1).

<sup>123</sup> NEPA § 102(2)(C), 42 U.S.C. § 4332(2)(C).

<sup>124</sup> Occupational Safety and Health Act § 6(b)(5), 29 U.S.C. § 655(b)(5).

<sup>125</sup> See Alvin M. Weinberg, *Science and Trans-Science*, 10 MINERVA 209, 209–215 (1972) (arguing that many public problems are scientific, political, and also “trans-scientific” in that they deal with questions that scientific experiments cannot answer).

<sup>126</sup> See, e.g., 40 C.F.R. pt. 51, app. W § 1.0(b)–(d) (explaining that air monitoring data alone are often insufficient and impacts of new sources “can only be determined through modeling”).

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

<sup>128</sup> See Weinberg, *supra* note 125, at 210 (discussing such risk modeling).

<sup>129</sup> SECURITIES AND EXCHANGE COMMISSION, ANNUAL REPORT 1981 11–12 (1981).

audited, and contested, but were still more difficult for non-specialists, including courts and legislatures, to understand than simpler earlier models.<sup>130</sup>

These capability gains again created accountability problems. Agencies operated in specialized domains where courts could not even fulfill their traditional function of using statutory interpretation to guarantee the legality of administrative action. Many terms in authorizing statutes, from the Clean Air Act to OSHA, were highly technical and diverged from ordinary usage.<sup>131</sup> The era's debates over statutory interpretation reflected anxiety on the part of judges as to whether they would be able to actually interpret the laws such that they could sufficiently constrain agency action.<sup>132</sup> In the *Chevron* deliberations in 1984, Justice John Paul Stevens reportedly admitted that "when I get so confused, I go with the agency," an admission that the complexity of modern administration had exceeded the capacity of the courts to oversee it.<sup>133</sup>

Procedural protections could substitute for substantive review only in part. The gap between agency action and review often stretched to years, meaning that unless courts enjoined agency interventions (often burdening the parties the government was supposed to protect), remedies were far away.<sup>134</sup> Where courts relied on increased proceduralization to protect disempowered groups, resource asymmetries emerged as sophisticated parties marshaled resources and experts to navigate procedures and influence agency action that under-resourced groups could not.<sup>135</sup> This "capture" dynamic and the flowering of public choice theory would become key influences in attempts to rein in the administrative state.<sup>136</sup> Another resource asymmetry emerged between agencies, with large staffs and high technology, and courts, which still ran as small organizations and often on paper.<sup>137</sup> Thus, once capture had occurred or the agency otherwise grew misaligned, it was difficult for the judiciary to identify and correct it. As such, increasing capabilities in administration gave rise to new risks of loss of accountability.

---

<sup>130</sup> Jasanoff, *supra* note 120, at 203, 209–11 (describing dissonance between public demands for "untainted knowledge" and the reality that policy-relevant science assimilates to politics, and highlighting contestation over models and underlying assumptions).

<sup>131</sup> See *Ethyl Corp. v. EPA*, 541 F.2d 1, 36–38 (D.C. Cir. 1976) (emphasizing that in "technical" cases a reviewing court must work to "understand enough about the problem confronting the agency" to perform its appellate function, while insisting review ultimately asks for a "rational basis" rather than judicial substitution of judgment).

<sup>132</sup> See *Indus. Union Dep't, AFL-CIO v. Am. Petroleum Inst. (The Benzene Case)*, 448 U.S. 607, 652–58 (1980).

<sup>133</sup> Mark Sherman, *High Court Climate Case Looks at EPA's Power*, HUFF POST (Feb. 23, 2014, 3:59 PM), <https://perma.cc/SQ6K-5ANA>.

<sup>134</sup> See Bagley, *supra* note 27, at 350, 360–63 (explaining how procedural demands and litigation risk can delay or derail protective regulation and shift policymaking forms).

<sup>135</sup> See Marc Galanter, *Why the "Haves" Come Out Ahead: Speculations on the Limits of Legal Change*, 9 LAW & SOC'Y REV. 95, 97–104 (1974).

<sup>136</sup> See George J. Stigler, *The Theory of Economic Regulation*, 2 BELL J. ECON. & MGMT. SCI. 3, 3–5 (1971); Daniel A. Farber & Philip P. Frickey, *The Jurisprudence of Public Choice*, 65 TEX. L. REV. 873, 880–82 (1987).

<sup>137</sup> JOHN G. ROBERTS, JR., 2023 YEAR-END REPORT ON THE FEDERAL JUDICIARY 2–5 (charting the history of technology in the judiciary).

Administrative law again sought a new settlement of capability and accountability, relying again on the three tools of allocation of authority, review, and information-forcing. This new settlement built on the previous ones, but a few key innovations stand out. First, courts fine-tuned review of agency actions. The administrative record became the basis for review, and the “hard look” doctrine established in the 1971 case *Citizens to Preserve Overton Park v. Volpe* required that courts consider the whole record to determine whether the agency had acted based on “all relevant factors” without a “clear error in judgment.”<sup>138</sup> *Motor Vehicle Manufacturers Association v. State Farm* held that courts would use the record to assess whether agencies had considered important aspects of the problem, addressed reasonable alternatives and counterarguments, and gave “adequate reasons” for changes.<sup>139</sup> The effect was to proceduralize rationality, expanding the role of the courts into evaluating how agency reasoning occurred rather than if it was correct.

Second, deference doctrine crystallized as a form of allocating authority and oversight to preserve both capabilities and accountability. *Chevron U.S.A. Inc. v. NRDC*, holding that if Congress had not directly spoken to an issue of interpretation, courts should defer to agencies’ reasonable construction of ambiguous statutes, became the canonical statement of interpretive deference. As Justice Stevens remarked, statutory interpretation in the modern administrative state was so confusing for judges that it was better to defer to the agencies.<sup>140</sup> When the Clean Air Act used “stationary source,” courts faced a term whose meaning emerged from engineering practice, atmospheric chemistry, and economic modeling, domains where legal training provided no insight.<sup>141</sup> In 1983, in *Baltimore Gas & Electric Co. v. NRDC*, the Court upheld the Nuclear Regulatory Commission’s assumption that permanent nuclear waste storage would have no significant environmental impact because courts should be deferential to agencies under conditions of uncertainty such as estimating the progress of future technology.<sup>142</sup> Combined with *Chevron*, these cases meant that courts would defer substantive questions, including interpretation of ambiguous statutes, to agencies, and would police procedure using the tools of hard look review. Courts had retreated again from evaluations of substance, again seeking in procedure a partial substitute that could prevent lawless administration.

Finally, a variety of legislative, judicial, and executive actions worked to make the informational underpinnings of the administrative state legible to review and oversight through information-forcing. The Freedom of Information Act of 1966<sup>143</sup> (FOIA) and the Government in the Sunshine

---

<sup>138</sup> 401 U.S. 402, 416, 420 (1971).

<sup>139</sup> 463 U.S. 29, 42–44 (1983).

<sup>140</sup> 467 U.S. 837, 865–66 (1984).

<sup>141</sup> 467 U.S. at 840–41 (describing the “bubble concept” dispute over the meaning of “stationary source”); see Richard J. Pierce, Jr., *What Actually Happened in Chevron*, YALE J. REG. NOTICE & COMMENT (Jan. 23, 2018), <https://www.yalejreg.com/nc/what-actually-happened-in-chevron-by-richard-j-pierce-jr/>.

<sup>142</sup> 462 U.S. 87, 103 (1983).

<sup>143</sup> Freedom of Information Act, Pub. L. No. 89-487, § 1, 80 Stat. 250, 250–52 (1966).

Act<sup>144</sup> enabled public review of agency files across the administrative state. NEPA<sup>145</sup> and the D.C. Circuit case *Calvert Cliffs' Coordinating Committee v. AEC* institutionalized environmental impact statements and alternatives analysis, forcing agencies to publicize and defend the scientific bases of their actions.<sup>146</sup> The Paperwork Reduction Act of 1980 established the Office of Information and Regulatory Affairs (OIRA) within the Office of Management and Budget to review agency actions.<sup>147</sup> Executive Order 12,291 centralized regulatory review and required cost-benefit analysis.<sup>148</sup> These information-forcing mechanisms did not rely upon courts or Congress to develop the technical expertise to govern the substance of agency actions, but rather empowered interested parties to understand and, if necessary, object to the workings of government.

The Supreme Court did seek to limit this mounting proceduralization in *Vermont Yankee Nuclear Power Corp. v. NRDC*<sup>149</sup> and *Heckler v. Chaney*.<sup>150</sup> These limitations proved insufficient.<sup>151</sup> Procedure is an imperfect substitute for substance, and while the courts were doing what they could to extend the doctrinal settlement of allocation, review, and information to deal with the challenges of complex science, these moves came at the cost of procedural intensification and slower government, which hampers administration today.<sup>152</sup>

By 1985, the administrative state had accumulated three layers of complexity, each persisting alongside its successors: sectoral expert regulation in the commission tradition; mass statistical administration for social programs; and model-driven risk governance for environmental and technological hazards. A single major EPA regulatory action could now require deep technical records and expert judgment, population-scale statistical inference and risk assessment, and computerized modeling. It also came with procedural requirements from all three eras that had likewise layered over time: APA notice-and-comment, record-based “hard look” judicial review, and (for executive agencies) centralized OMB review and cost–benefit analysis under E.O. 12,291. The sedimentary layers had compressed into governmental geology that no single actor could fully comprehend, and the balance between capability and accountability grew unsteady.

---

<sup>144</sup> Government in the Sunshine Act, Pub. L. No. 94-409, § 3(a), 90 Stat. 1241, 1241–49 (1976) (codified at 5 U.S.C. § 552b).

<sup>145</sup> National Environmental Policy Act of 1969, Pub. L. No. 91-190, § 102(2)(C), 83 Stat. 852, 852–53 (1970) (codified as amended at 42 U.S.C. § 4332(2)(C)).

<sup>146</sup> 449 F.2d 1109, 1112–14 (D.C. Cir. 1971).

<sup>147</sup> 44 U.S.C. § 3503(a)–(b) (creating OIRA within OMB); id. § 3504(a) (to “oversee the review and approval of information collection requests,” federal statistical activities, records management, and use of information technology).

<sup>148</sup> Exec. Order No. 12,291, §§ 2(b)–(d), 3(a), 3(c)–(e), 46 Fed. Reg. 13,193, 13,193–95 (Feb. 19, 1981).

<sup>149</sup> 435 U.S. 519, 524–25 (1978).

<sup>150</sup> 470 U.S. 821, 831–35 (1985).

<sup>151</sup> Bagley, *supra* note 27, at 350–55, 360–64; McGarity, *supra* note 26, at 1387–93.

<sup>152</sup> McGarity, *supra* note 26, at 1385–87; Richard J. Pierce, Jr., *Rulemaking Ossification Is Real: A Response to “Testing the Ossification Thesis”*, 80 GEO. WASH. L. REV. 1493, 1493–96 (2012).

## II. A CRISIS OF ACCOUNTABILITY

For a hundred years, American administrative law followed a clear pattern: new technologies created capability crises, government complexified in response, accountability gaps emerged, and doctrinal settlements restored balance. Procedural oversight through administrative law created sufficient accountability at each step. Today, that pattern has broken. The Supreme Court is dismantling the administrative state without precipitating technological crisis. New technologies like the internet and narrow AI have emerged since the *Chevron* settlement but did not fundamentally transform government in such a way as to create a capability crisis. A backlash has emerged nonetheless, this novel accountability crisis threatening to push administration back past the New Deal.<sup>153</sup>

This discontinuity demands explanation. This Part argues that administrative law has collapsed not from new technological disruption but from accumulated procedural weight, each settlement adding procedural requirements without removing previous ones. These accumulated procedural requirements transformed the character of administration. As Professor Richard Stewart diagnosed, the New Deal “transmission-belt” model of neutral experts simply enacting legislative directives has yielded to a regime in which expertise and pluralistic interest-representation sit uneasily together.<sup>154</sup> As Congress and the courts lost the ability to substantively understand and oversee administration and came to rely increasingly on the discretion of administrators,<sup>155</sup> pluralistic bargaining became the underpinning for a kind of legitimacy through interest representation.<sup>156</sup> The procedural tools that Part I identified as the law’s way of mediating capability and accountability became the bargaining rules.

This procedural accretion helps explain why government failed to integrate transformative digital technologies despite their revolutionizing everything else. Procedural complexity creates veto points, which empower incumbents. Incumbents resist changes that threaten their bargaining position and seek to strengthen that position by adding more procedure, “play[ing] for rules.”<sup>157</sup> Fundamental changes to the structure of administrative law making use of technologies like the internet and narrow AI are blocked by powerful existing groups that benefit from the status quo.

The result is a double deficit. The federal workforce is smaller *per capita* today than it was in the 1960s even as demands have expanded.<sup>158</sup> Far from alleging that government is doing too much,

---

<sup>153</sup> See Metzger, *supra* note 42, at 1–6 (describing the rising tide of anti-administrative sentiment in the Supreme Court).

<sup>154</sup> Stewart, *supra* note 11, at 1711–13.

<sup>155</sup> *Id.* at 1676–78; 1702–11.

<sup>156</sup> *Id.* at 1712.

<sup>157</sup> See Galanter, *supra* note 135, at 100.

<sup>158</sup> *The evolution of government employment*, FRED BLOG (Dec. 30, 2024), <https://fredblog.stlouisfed.org/2024/12/the-evolution-of-government-employment/>; U.S. GOV'T ACCOUNTABILITY

complaints that government cannot do anything have become a key part of contemporary political discourse.<sup>159</sup> Government has become both less capable and less accountable, throttled by procedure while simultaneously captured by those sophisticated enough to navigate it.

Faced with this crisis, the Supreme Court has chosen a blunt remedy: shrinking administration back to comprehensible size. This response is coherent on its own terms. If oversight requires comprehension, and comprehension requires simplicity, then legitimacy requires cutting back. But it addresses the accountability deficit by deepening the capability deficit, precisely when new problems from climate change to AI demand more sophisticated governance than ever. The capability-accountability trap is biting harder than ever, with government losing both the ability to address complex problems and the legitimacy to act even if it could. This Part examines the causal chain that produced this crisis: first, how accumulated procedural complexity transformed administration into a bargaining arena that advantages repeat players and resists reform; second, how this dynamic explains the failure of internet-era technology to transform government; and finally, how the resulting legitimacy crisis triggered the current judicial retrenchment.

#### *A. The Sedimentary State*

The procedural accumulation described in Part I both made government increasingly complicated and changed the character of administration. The traditional model of administrative action was of neutral, scientific experts simply carrying out the will of Congress through their specialized capabilities.<sup>160</sup> However, this “transmission belt” model relied upon a separation between politics and expertise that did not accurately describe how administrators exercised their discretion, undermining the legitimacy of scientific administration.<sup>161</sup> By the early 1970s, a new justification for administration arose, arguing that agency actions operated as pluralist forums in which organized interests could make their voices heard through legally structured procedures.<sup>162</sup> Groups with strong preferences could invest more in participation, and procedures became rules of these “markets” rather than guarantors of the fairness of administrative action.

In this model, regulatory outputs are legitimated by structured participation and judicially policed reason-giving rather than by claims of technocratic neutrality. Because Congress could no longer translate the will of the people into comprehensive guides for agency action,<sup>163</sup> the popular will would be expressed through pluralist bargaining. Judicially-governed process provided the

---

OFF., INTERNAL REVENUE SERVICE: ABSORBING BUDGET CUTS HAS RESULTED IN SIGNIFICANT STAFFING DECLINES AND UNEVEN PERFORMANCE (GAO-14-534R) (Apr. 2014).

<sup>159</sup> See generally EZRA KLEIN & DEREK THOMPSON, ABUNDANCE (2025).

<sup>160</sup> Stewart, *supra* note 11, at 1671–75.

<sup>161</sup> *Id.* at 1776–81.

<sup>162</sup> *Id.* at 1711–13.

<sup>163</sup> *Id.* at 1671–75.

framework. Notice-and-comment opens the door to informed outsiders, an expanded administrative record makes choices contestable, and “hard look” review enforces responsiveness to salient criticisms. Courts would defer to agency interpretations under *Chevron* because they also could not claim to be able to truly determine or specify what agencies should do regarding scientific questions, and substituting their judgment for that of the agencies would also undermine government’s legitimacy. Administrative agencies must decide questions that, while they involve science, are fundamentally normative and subject to the discretion of administrative officials, like what counts as “fair” or “unhealthy.”<sup>164</sup> Political participation through bargaining where different interests are represented reduces the extent to which these questions are insulated from public input.

However, the agency forums are not truly representative. Repeat players have the resources and familiarity with the regulatory process to exert disproportionate influence over agency decisions.<sup>165</sup> McCubbins, Noll, and Weingast reframe the procedural moves that supported this shift as deck-stacking, institutional design that embeds legislative compromises and distributes leverage across favored constituencies.<sup>166</sup> McCubbins and Schwartz argue that Congress deliberately relies on fire-alarm oversight, deputizing organized groups to monitor agencies through those procedures instead of conducting continuous “police-patrol” oversight itself.<sup>167</sup> In domains with a high degree of complexity of agency action, this kind of expert interest group oversight can compensate for substantive congressional incapacity. Yet the groups best positioned to pull the fire alarm are those with resources to monitor agencies continuously, not the diffuse public that agencies are meant to serve.

Because rules became tools for interest groups, procedures no longer simply maintain rule of law. Instead, they shape bargaining domains, and shifting them advantages or diminishes specific parties. Notice-and-comment rulemaking is formally an information-gathering device but practically operates as an agenda-shaping and litigation-positioning stage.<sup>168</sup> Repeat players file submissions that target the fulcrum assumptions of an agency’s analysis and lay the groundwork for *State Farm* challenges if the final rule does not squarely address reasonable alternatives. Agencies respond strategically, pre-negotiating with likely litigants, over-documenting “consideration of relevant factors,” and drafting response-to-comment matrices with an eye toward remand risk rather than clarity.<sup>169</sup> Each procedural requirement, defensible for external reasons in isolation, becomes another surface on which sophisticated parties can apply pressure. The cumulative effect is ossification, the increasing structural paralysis of government,<sup>170</sup> and increased inscrutability.

---

<sup>164</sup> Jasanoff, *supra* note 120, at 203–06.

<sup>165</sup> Stewart, *supra* note 11, at 1713–15.

<sup>166</sup> McCubbins, Noll, & Weingast, *supra* note 16, at 249–53.

<sup>167</sup> Mathew D. McCubbins & Thomas Schwartz, *Congressional Oversight Overlooked: Police Patrols versus Fire Alarms*, 28 AM. J. POL. SCI. 165, 166 (1984).

<sup>168</sup> See McCubbins, Noll, & Weingast, *supra* note 16, at 258–60.

<sup>169</sup> See Bagley, *supra* note 27, at 351–52, 360–61.

<sup>170</sup> See McGarity, *supra* note 26, at 1386–90.

This dynamic creates a two-sided legitimacy problem. As veto points proliferate and administrative reasons are crafted to withstand litigation rather than to be genuinely transparent to the public, the outputs of administrative action look like lowest common denominator compromises rather than what is genuinely in the public interest. The distributional fights that often are the main focus of conflict in these areas do not seem well-governed by the judicial procedural review that is intended to give administration legitimacy. External legitimacy, in the form of presidential administration, can provide needed energy and a democratic justification for the actions of the administrative state.<sup>171</sup> However, even with OIRA and similar forms of executive review, much of agency action remains outside of presidential control because of its size and complexity.<sup>172</sup>

More fundamentally, pluralist bargaining is not representative. Influence shifts from diffuse citizens to repeat players who can credibly threaten delay, remand, or appropriations retaliation.<sup>173</sup> The construction company that will litigate an EPA rule has a seat at the table that the asthmatic children who will breathe the air do not. Pluralist bargaining allows public input on key normative questions that agencies must decide, but when only certain players can provide input effectively, its utility is undermined. Administrative law, which is intended to be the fundamental democratic check on an increasingly powerful, complicated, and expert government becomes a kind of politics by other means which does not even have the virtue of complete democratic representation.

Attempts to reform the system disadvantage particular interest groups in bargaining, leading to resistance. It seems plausible that the internet has not transformed government for something like this reason—efforts to build simpler tax filing systems have been met with lobbying efforts by established players like TurboTax that threaten their interests.<sup>174</sup> Procedure often only piles up, and so both capacity and accountability tend to erode. This erosion sets the stage for Part II.B's puzzle: why, unlike every prior era, did the transformative new technologies of the internet and narrow AI fail to transform government?

### *B. Technological Stagnation*

The technologies developed since 1985 have transformed modern life, but mostly not modern government. In 1990, less than one percent of Americans used the internet.<sup>175</sup> By 2023, that

---

<sup>171</sup> See Kagan, *supra* note 16, at 2341.

<sup>172</sup> See Bednar, *supra* note 16, at 830–31.

<sup>173</sup> Galanter, *supra* note 135, at 97–104.

<sup>174</sup> Lawrence Zelenak, *Complex Tax Legislation in the TurboTax Era*, 1 COLUM. J. TAX L. 91, 92–95 (2010).

<sup>175</sup> World Bank, *Internet users for the United States* (FRED series ITNETUSERP2USA) (reporting 1990: 0.7850 internet users per 100 people).

number had reached 93%,<sup>176</sup> as the World Wide Web, email, the smartphone, and social media pushed life online. Government investments, both technical and political, underpinned this growth. The internet emerged from research at the Defense Advanced Research Projects Agency,<sup>177</sup> and international protocols that enabled an integrated, global internet required political will.<sup>178</sup>

Yet agencies underwent only superficial digitalization. The E-Government Act of 2002 mandated online services and<sup>179</sup> Regulations.gov digitized comments.<sup>180</sup> Electronic filing systems, including in the U.S. Patent and Trademark Office,<sup>181</sup> the Securities and Exchange Commission's EDGAR,<sup>182</sup> and the federal courts' PACER,<sup>183</sup> moved much information access online. But these incremental improvements did not constitute the kind of institutional transformation that had previously driven administrative evolution but rather a shift in how to access existing tools.<sup>184</sup> Even on its own terms, government uptake of technology has been relatively ineffective. As of 2025, the Government Accountability Office continues to list federal IT acquisition and management as a "high-risk" area, finding that investments "too frequently fail or incur cost overruns and schedule slippages while contributing little to mission-related outcomes," notwithstanding annual federal IT spending exceeding \$100 billion.<sup>185</sup> Crises from the failure of HealthCare.gov<sup>186</sup> to the scramble to find COBOL coders to fix unemployment systems that crashed during the COVID-19 pandemic demonstrate government's underlying incapacity.<sup>187</sup> Instead of reducing the need to rely on procedures developed for a prior technological era, and so making government more scrutable, the internet left government much as it was—often mediated by private actors providing interfaces optimized for their own purposes.

---

<sup>176</sup> *Id.* (reporting 2023: 93.1444 internet users per 100 people).

<sup>177</sup> *Birth of the Commercial Internet*, NAT'L SCI. FOUND. (n.d.), <https://www.nsf.gov/impacts/internet>.

<sup>178</sup> *Management of Internet Names and Addresses*, U.S. DEP'T COMMERCE NAT'L TELECOMM. & INFO. ADMIN. (Jul. 22, 2000) (establishing ICANN).

<sup>179</sup> E-Government Act of 2002, H.R. 2458, 107th Cong. (as enrolled).

<sup>180</sup> *See Presidential Initiatives: E-Rulemaking*, GEORGE W. BUSH WHITE HOUSE (archives.gov) (n.d.), <https://georgewbush-whitehouse.archives.gov/omb/egov/c-3-1-er.html>.

<sup>181</sup> Zandra Smith & Lisa Tran, *Filing for a Patent Using EFS Web*, U.S. PATENT AND TRADEMARK OFFICE (2017), <https://www.uspto.gov/sites/default/files/documents/Website%20PDF%20-%20Invention%20Con%202017%20Filing%20for%20a%20Patent%20-%20OPLA.pdf>.

<sup>182</sup> Mauri L. Osheroff, Mark W. Green, & Ruth Armfield Sanders, *Electronic Filing and the EDGAR System: A Regulatory Overview*, U.S. SECURITIES AND EXCHANGE COMMISSION (Oct. 3, 2006), <https://www.sec.gov/info/edgar/regoverview.htm>.

<sup>183</sup> *About Us*, PUBLIC ACCESS TO COURT ELECTRONIC RECORDS (PACER) (n.d.), <https://pacer.uscourts.gov/about-us>.

<sup>184</sup> As Jane E. Fountain predicted might happen in her pathbreaking book, "reorganizing and restructuring the institutional arrangements in which [government to citizen] transactions are embedded" proved to be much harder than simply taking up new technology. JANE E. FOUNTAIN, *BUILDING THE VIRTUAL STATE: INFORMATION TECHNOLOGY AND INSTITUTIONAL CHANGE* 6–7 (2001).

<sup>185</sup> U.S. Gov't Accountability Off., *High-Risk Series: Critical Actions Needed to Urgently Address IT Acquisition and Management Challenges* 2–4 (GAO-25-107852, 2025).

<sup>186</sup> JENNIFER PAHLKA, *RECODING AMERICA: WHY GOVERNMENT IS FAILING IN THE DIGITAL AGE AND HOW WE CAN DO BETTER* 118–134 (2023).

<sup>187</sup> Michael A. Navarrete, *COBOLing Together UI Benefits*, at 1–2 (Brookings Working Paper 98, 2024).

Why have these technologies not transformed government? Institutional factors, like civil service rules, procurement processes, risk aversion, and inadequate appropriations, present obstacles. Fundamentally, however, proceduralization and the rise of bargaining through the sedimentary state discussed in Part II.A provide a deeper explanation for why transformation has not worked well.

Procedure disincentivizes innovation within government. An agency redesigning operations would need to alter workflows formalized in regulations, change processes organized interests have optimized around, and modify record-keeping that forms the basis for judicial review. Each change would create winners and losers among the constituencies that have learned to navigate existing procedures, many of whom would bring legal challenges arguing that the new process violates statutory requirements or denies procedural rights. Each change would require approval through the same ossified channels that Part II.A described: interagency review, OIRA clearance, notice-and-comment on any rule changes, and anticipation of litigation from those disadvantaged by reform. As such, procedure has not just prevented government from effectively carrying out its duties in many domains but also from upgrading its ability to do so, incapacity reaching to the roots.

TurboTax and the problem of tax reform illustrate the dynamic. Efforts to build simpler tax filing systems, a useful technology that would make government more accessible to ordinary citizens, have been met with lobbying efforts by established players whose business models depend on tax complexity.<sup>188</sup> The ability of private companies to uptake technology and fulfill quasi-governmental functions has created a constituency that is opposed to more effective government which might displace them. The same pattern appears across domains. As such, the procedural sediment described in Part II.A is not merely inert but actually defended by those who have adapted to it.

Internet-era technological participation has amplified the surface dynamics without changing the underlying structure. Mass comments increase volume and noise, but agencies rationally prioritize the credible-litigant subset of those who have commented. A proposed rule might generate a million form comments from an advocacy campaign, but the agency's lawyers will focus on the detailed technical submissions from regulated industries that preview the arguments likely to appear in a petition for review.<sup>189</sup> Online dockets make agency proceedings more transparent, but transparency without capacity for engagement merely lets citizens watch a game they cannot play.

Agencies have failed to undergo the kind of fundamental process redesign that digital technology enables elsewhere. While some scholars have tried to develop visions for such transformation

---

<sup>188</sup> Zelenak, *supra* note 174, at 92–95, 97–100.

<sup>189</sup> See Nina A. Mendelson, *Foreword: Rulemaking, Democracy, and Torrents of E-Mail*, 79 GEO. WASH. L. REV. 1343, 1346–50, 1359–76 (2011).

through the internet or narrow AI systems deployed in administrative domains, their visions of genuine transformation remain unrealized.<sup>190</sup> The use of AI in administration presents serious risks both to legitimacy, because of the black box nature of many systems,<sup>191</sup> and of abuse, such as biased algorithms.<sup>192</sup> However, these problems may be more correctable in AI systems than they are in humans,<sup>193</sup> and government will have to adopt new technology in order to keep up with the rapid advances made in the private sector.

As agencies fall further behind the technological frontier, the gap between governmental and private-sector capabilities widens. Regulated entities operate with tools that regulators cannot match, and the rise of advanced AI threatens to make this problem even worse. The resource asymmetry between agencies and sophisticated parties that Part II.A identified as a feature of pluralist bargaining grows more severe when that asymmetry includes technological capacity. Agencies that cannot understand how regulated entities operate cannot effectively oversee them, and the procedural substitutes for substantive oversight that Part I traced become even more inadequate.

This account helps resolve the apparent puzzle of technological stagnation in government. Unlike prior eras, no capability crisis preceded the Supreme Court's current doctrinal unsettlement because government never gained the capabilities that new technologies could have provided. The relatively limited integration of the internet and narrow AI into government has left it technologically under-prepared to face emerging challenges. The sedimentary state struggles to transform itself. Procedures pile up, but capability does not accumulate alongside them. Government is left with the worst of both worlds: the complexity costs of three eras of procedural accretion without the technological benefits that might justify those costs.

This paralysis helps explain why no new settlement or other significant transformation has emerged in administrative law despite four decades passing since *Chevron*. There was no capability expansion to trigger an accountability gap, no new governmental power requiring new

---

<sup>190</sup> See, e.g., Cary Coglianese & David Lehr, *Regulating by Robot: Administrative Decision Making in the Machine-Learning Era*, 105 GEO. L.J. 1147, 1167–68, 1174–75, 1180–81 (2017) (describing “adjudicating by algorithm” and “rulemaking by robot,” and explaining how ML could be nested within larger models to support “automated rulemaking”); Anthony J. Casey & Anthony Niblett, *The Death of Rules and Standards*, 92 IND. L.J. 1401, 1403–05, 1410–14 (2017) (introducing “micro-directives,” machine translation of legislative goals into “simple commands,” and the claim that micro-directives can “update[] as quickly as conditions change”); David Freeman Engstrom & Daniel E. Ho, *Algorithmic Accountability in the Administrative State*, 37 YALE J. ON REGUL. 800, 800–02, 804–06, 849–54 (2020) (proposing “prospective benchmarking” to “learn about, validate, and correct errors” in algorithmic decisionmaking).

<sup>191</sup> See Dierdre K. Mulligan & Kenneth A. Bamberger, *Procurement as Policy: Administrative Process for Machine Learning*, 34 BERKELEY TECH. L.J. 773, 816–17 (2019) (explaining that ML can displace discernible reasons with models whose internal “reasoning” may be difficult or impossible to recover).

<sup>192</sup> See Julia Angwin et al., *Machine Bias*, PROPUBLICA (May 23, 2016), <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>; Joy Buolamwini & Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*, in PROC. FAT\* 77, 80–83, 86–88 (2018).

<sup>193</sup> Jon Kleinberg et al., *Human Decisions and Machine Predictions*, 133 Q. J. ECON. 237, 251–57 (2018).

constraints. The traditional cycle of change—technological disruption, capability crisis, governmental response, doctrinal settlement—did not occur. But the absence of a capability surge has not meant the absence of crisis. Instead, a different kind of crisis emerged, one driven not by governmental overreach but by governmental dysfunction. Part II.C examines how the Supreme Court has responded to this legitimacy vacuum.

### *C. Judicial Dismantling for Administrative Simplification*

Parts II.A and II.B described a government that is often ossified, controlled by sophisticated repeat players, and unable to upgrade itself to handle the new problems society confronts. This is a crisis, but not the kind that prior cycles of doctrinal disruption and resettlement addressed. That traditional pattern involved capability surging ahead of accountability, as when the mass administration of the New Deal outpaced the ability of courts and legislatures to substantively review decisions. Here, both have eroded together, and over-capability has not been the driver of worries about governmental abuse.

Nevertheless, the United States is again undergoing a serious crisis in administrative law aimed at creating a greater degree of accountability in administration. The Supreme Court, in a series of recent decisions, has struck at the foundations of administrative agencies, creating what Professor Gillian Metzger calls “1930s Redux” as elements of administrative law settled since the New Deal come under attack.<sup>194</sup> The Court is not merely adjusting doctrine at the margins, as it does during the periods of settlement. Instead, it is dismantling parts of the fundamental architecture that resulted from the *Chevron* settlement, making significant changes to the same three tools of allocation, review, and information-forcing that enabled administrative law to mediate capability and accountability across a century of technological change.

It is unclear exactly why the Court has decided at this moment to undo the doctrinal settlement that has structured administrative law since at least the *Chevron* era. It is possible that increased partisanship among the Justices<sup>195</sup> and a growing sense that administrative agencies are captured by Democrats<sup>196</sup> have led to this turn, even though the Supreme Court has had a majority of Republican-appointed Justices since 1970 and *Chevron* itself was championed by Antonin Scalia.<sup>197</sup> But the framework developed in this Article suggests a structural explanation beyond partisanship: the accumulated complexity described in Parts II.A and II.B has made administration

---

<sup>194</sup> Metzger, *supra* note 42, at 1–3.

<sup>195</sup> Richard Re, *Foreword: To a Conservative Warren Court*, 139 HARV. L. REV. 1, 19 (2025); Metzger, *supra* note 42, at 15.

<sup>196</sup> Steven Teles, *It's time for liberals to confront their own anti-majoritarianism*, NISKANEN CENTER (Mar. 21, 2025), <https://www.niskanencenter.org/its-time-for-liberals-to-confront-their-own-anti-majoritarianism/>.

<sup>197</sup> See Antonin Scalia, *Judicial Deference to Administrative Interpretations of Law*, 1989 DUKE L.J. 511, 511–12 (1989).

genuinely incomprehensible to generalist judges, and the Court is responding by trying to shrink government back to a size it can understand and therefore oversee. From this perspective, the Court can be understood to be seeking to restore accountability and democratic legitimacy to an administrative state that has been distanced from the popular will by complexity and pluralist bargaining.

This reading finds support in the Justices' recent decisions. Requiring explicit delegation from Congress on major questions, centralizing executive power in the President, and increasing the role of courts with their politically appointed judges in overseeing administrative actions can each be seen as ways that the Supreme Court is seeking to increase the extent to which administrative agencies must respond to democratic authority, even if they may in doing so undermine the output legitimacy that comes from effective, capable, and independent government. These moves also track the three pillars of doctrinal settlement that Part I identified.

First, the end of *Chevron* deference with the decision in *Loper Bright* transformed the allocation pillar of the prior administrative settlement.<sup>198</sup> Deference regimes like *Chevron* and *Auer*<sup>199</sup> had put much of legal interpretation into the hands of agencies who were deemed to be better equipped to understand how to resolve ambiguity in technical matters. In *Loper Bright*, the Court explicitly overruled *Chevron*, holding that “[c]ourts must exercise their independent judgment in deciding whether an agency has acted within its statutory authority” even where a statute is ambiguous.<sup>200</sup> Some degree of deference to agencies does persist through the background doctrine of *Skidmore* deference.<sup>201</sup> But overall the recent shift toward the use of the traditional tools of statutory interpretation by the judiciary to address these ambiguous problems, as *Loper Bright* and *Kisor v. Wilkie*<sup>202</sup> require, represents a more confident judicial posture towards technical interpretation and a shift in authority away from agencies toward the courts. The allocation that enabled expert agencies to resolve technical ambiguity has been reversed. Courts have reclaimed interpretive authority they ceded to agencies because they lacked the ability to understand what agencies were doing, and requiring that agencies act under judicial interpretations will push them to act in ways that courts understand.<sup>203</sup>

---

<sup>198</sup> *Loper Bright Enters. v. Raimondo*, 603 U.S. \_\_\_, slip op. at 18–20, 34 (June 28, 2024).

<sup>199</sup> 519 U.S. 452 (1997). See Thomas E. Nielsen & Krista A. Stapleford, *What Loper Bright Might Portend for Auer Deference*, HARV. L. REV. BLOG (Jul. 5, 2024), <https://harvardlawreview.org/blog/2024/07/what-loper-bright-might-portend-for-auer-deference/>.

<sup>200</sup> *Loper Bright*, slip op. at 18–20, 27–28.

<sup>201</sup> *Id.* at 5, 16–17 (describing *Skidmore* respect and quoting “body of experience and informed judgment”). *Skidmore v. Swift & Co.*, 323 U.S. 134, 140 (1944).

<sup>202</sup> 588 U.S. 558, 574–76 (2019).

<sup>203</sup> Whether this shift will allow for effective government is unclear. Forcing agencies to follow the interpretations of courts that do not understand the domains within which they operate could simply undermine their capacity to regulate effectively. See Vermeule, *supra* note 15, at 625–26.

Second, recent doctrinal shifts have broadened the set of areas in which courts can review agency authority. The expanding major questions doctrine demonstrates a new front in which judicial review of agency authority has opened up. Since the New Deal settlement, which ended the use of the nondelegation doctrine to police what Congress can and cannot delegate to the administrative state, delegations were doctrinally permissible if some “intelligible principle” guides the delegation.<sup>204</sup> Practically, however, no delegation has been struck down on the basis of this test since the New Deal, rendering it a dead letter. However, the major questions doctrine, which blocks delegations of “major question[s]” of “economic or political significance” to agencies unless there is clear congressional authorization for them to handle such issues, has become an increasingly active tool of limiting agency authority.<sup>205</sup> In *West Virginia v. EPA*, the Court deployed the major questions doctrine to invalidate EPA’s Clean Power Plan, holding that the agency’s interpretation of the Clean Air Act to allow generation-shifting as an emissions reduction strategy required clearer congressional authorization given its economic significance.<sup>206</sup> The effect is to require Congress to speak clearly on precisely the technical and contested questions that Congress delegated to agencies, attempting to force democratic deliberation over these questions even if Congress may struggle to effectively engage in it.

Third, and finally, the Court has begun to restrict agency adjudication, narrowing the domain in which agencies can act as both prosecutor and judge and forcing more adjudications into the judicial branch. In *SEC v. Jarkesy*, the Court recognized a jury-trial right for SEC civil-penalty cases under the Seventh Amendment, curtailing in-house adjudication in that category.<sup>207</sup> This decision strikes at the adjudicatory capacity that has been central to agency enforcement since the New Deal; administrative courts with their own specialized judges and sets of procedures made adjudications more efficient while providing fairly strong procedural protections to defendants.<sup>208</sup> Agencies across domains may face similar challenges to their adjudicatory authority, potentially requiring them to litigate enforcement actions in courts already struggling with docket backlogs, creating more procedural protections for affected parties but potentially adding delay to an already sclerotic system.<sup>209</sup>

Taken together, these decisions represent a coherent judicial project: restoring accountability by making administration scrutable to the courts, even if at the cost of administrative capability. If legitimacy requires judicial review, and judges cannot understand what agencies are doing, then, the Court’s implicit logic runs, agencies must be cut back until judges can oversee them. This

---

<sup>204</sup> *J.W. Hampton, Jr., & Co. v. United States*, 276 U.S. 394, 409 (1928) (articulating “intelligible principle”); *Whitman v. Am. Trucking Ass’ns*, 531 U.S. 457, 472–76 (2001).

<sup>205</sup> *West Virginia v. EPA*, 597 U.S. \_\_\_, slip op. at 19–20 (2022).

<sup>206</sup> 597 U.S. 697 (2022).

<sup>207</sup> *SEC v. Jarkesy*, No. 22-859, 603 U.S. \_\_\_, 2024 WL 3187811, \*8–10, 17 (U.S. June 27, 2024).

<sup>208</sup> See CONG. RSCH. SERV., *INFORMAL ADMINISTRATIVE ADJUDICATION: AN OVERVIEW* (R46930 Oct. 1, 2021).

<sup>209</sup> The year before *Jarkesy*, in *Axon Enterprise, Inc. v. FTC*, the Court held that litigants in administrative proceedings could directly seek review of collateral constitutional claims about agency actions in the courts rather than having to first go through the agency appeals process. *Axon Enters., Inc. v. FTC*, 598 U.S. 175, 180 (2023).

approach is not irrational. The complexity crisis is real, and the Court is correct that accountability has eroded. But the remedy addresses one deficit by deepening another. Shifting interpretive authority from agencies to courts does not make technical questions less technical, it merely relocates them to institutions less equipped to resolve them.<sup>210</sup> Requiring clear congressional statements on major questions does not make Congress more capable of anticipating regulatory challenges, it creates risks that novel problems go unaddressed until political coalitions can form around specific statutory language, especially given Congress's current incapacity.<sup>211</sup> Requiring jury trials for agency enforcement does not make enforcement more effective, but merely adds procedural costs to agencies already struggling with resource constraints and facing lengthy and onerous judicial review.

The Court's retrenchment thus threatens to produce an administrative state too constrained to address emerging challenges yet still too complex for meaningful democratic oversight. The status quo of *Chevron*-style doctrinal settlement setting the terms for pluralist bargaining might have persisted until some new wave of governance problems emerged. But the accumulated weight of procedural complexity and rising sense of administrative illegitimacy has instead created a different kind of crisis. The Court perceives, correctly, that something has gone wrong. However, its diagnosis, that agencies have accumulated too much power and need to be cut back, mistakes symptom for cause. The problem is not that agencies are too capable, but that capability and accountability have both eroded, leaving a system that is neither efficient nor democratic.

The question, then, is whether there is another way out of the trap. Can government become both more capable and more accountable, rather than sacrificing one for the other? The traditional doctrinal tools (allocation, review, and information-forcing) have reached their limits. Each additional procedural safeguard adds complexity without restoring comprehensibility. The cycle that worked from 1887 to 1985 has broken down.

Part III argues that AI offers a path forward by changing the underlying relationship between capability and accountability. Where prior technologies enhanced governmental capacity while increasing opacity, AI could represent a new chance to make government scrutable, increasing both capability and accountability. The capability-accountability trap has structured administrative law for over a century. Breaking it will require more than doctrinal adjustment. It will require technological transformation of the kind that government has so far failed to achieve, but that AI might finally make possible.

---

<sup>210</sup> Vermeule, *supra* note 15, at 624–26.

<sup>211</sup> See Freeman & Spence, *supra* note 36, at 20–22 (arguing that the reliance on outdated statutes to tackle new problems by agencies is rational given congressional incapacity).

### III. THE FOURTH SETTLEMENT: AI BEYOND THE CAPABILITY–ACCOUNTABILITY TRAP

This Part asks whether AI could change the underlying relationship of capability and accountability that has structured administrative law through cycles of technological change. Unlike prior waves of administrative technology that increased opacity as they increased capacity, AI plausibly functions as cognitive infrastructure to increase *scrutability*, the comprehensibility of governance to overseers and the public. If deployed correctly, existing AI systems can translate technical material into accessible forms, surface the assumptions and tradeoffs that matter for oversight, and reduce the costs of participation and comprehension for courts, Congress, and the public. More ambitiously, advances in alignment and interpretability suggest a path toward *auditable* administration, oversight that targets the reasons and constraints that actually drive governmental determinations, rather than relying exclusively on procedural proxies that approximate substantive review.<sup>212</sup> As such, though updating administrative law for AI might require more procedure suited to the AI context, if government can use AI to follow procedures more efficiently or if substance can substitute for procedure at different levels, the overall ossification of government could be reduced even as it becomes more accountable. The argument that follows is not a prediction that transformation *will* happen—important questions of political economy must be answered and political will gathered—but aims to suggest how AI *could* help construct a new settlement that makes government both more capable and more accountable in the future.

The discussion proceeds in four steps. Subpart A makes explicit the cognitive foundations of administrative behavior—why bounded rationality has repeatedly forced the substitution of procedure for substance—and identifies where AI intervention could matter most. Subpart B argues that AI can improve participation and oversight by acting as scrutability infrastructure, producing legally usable inputs for rulemaking, adjudication, and congressional review. Subpart C proposes accountability mechanisms suited to AI-assisted administration: an expanded administrative record that makes algorithmic reasoning auditable, review postures that reward transparency rather than merely policing procedure, and institutional architecture for ongoing oversight of AI systems in government.

#### *A. The Cognitive Infrastructure of Administration*

Administrative organizations are, at their core, systems for coordinating people to make decisions aimed at solving social problems under conditions of uncertainty.<sup>213</sup> They gather information about problems facing their constituencies, identify possible responses, model the consequences of each,

---

<sup>212</sup> I mean auditing in the expansive sense of technically-supported investigation suggested by Kroll et al., though also acknowledge the limits they point out. See Joshua A. Kroll et al., 165 U. PA. L. REV. 633, 660–61 (2017).

<sup>213</sup> Simon, *supra* note 7, at 18–19.

incorporate preferences from the public about which consequences are desirable, and act. These tasks are performed by human minds operating within institutional structures of rules that coordinate behavior, hierarchies that allocate authority,<sup>214</sup> and technologies that extend cognitive capacities beyond what any individual could achieve alone.<sup>215</sup> For example, the Social Security Administration is made up of a set of objectives and rules, both legal and operational, that give it purpose and structure it, a host of managers and employees who make decisions and carry out the tasks necessary to fulfill the goals of the organization, and technologies like punch card tabulation systems that extend and collectivize human cognitive capacities of memory and search to enable the administration of a national welfare program that requires tracking millions of people and their relevant data.

Herbert Simon's model of administrative behavior provides the theoretical foundation for understanding why this "cognitive" infrastructure of administration matters. Administrators work under conditions of bounded rationality, with limited cognitive computational resources,<sup>216</sup> imperfect information,<sup>217</sup> and time pressures that push them to "satisfice" across goals rather than perfectly maximize them.<sup>218</sup> They cannot examine every possible course of action, perfectly predict consequences, or hold in mind the full complexity of the systems they seek to govern.<sup>219</sup> "Rationality implies a complete, and unattainable, knowledge of the exact consequences of each choice. In actuality, the human being never has more than a fragmentary knowledge of the conditions surrounding his action, nor more than a slight insight into the regularities and laws that would permit him to induce future consequences from a knowledge of present circumstances."<sup>220</sup> Fundamentally, as Simon writes, "human rational behavior... is shaped by a scissors whose blades are the structure of task environments and the computational capabilities of the actor."<sup>221</sup>

Organizations help people overcome some of the limitations of bounded rationality by creating stable expectations about the behavior of other members in the organization, establishing channels of communication that facilitate coordination, and allowing specialization of work in expectation

---

<sup>214</sup> See Herbert A. Simon, *Organizations and Markets*, 5 J. ECON. PERSP. 25, 39–40 (1991) (explaining that a major use of authority is to coordinate behavior by promulgating "standards and rules of the road," enabling actors to form "more stable expectations" about others' behavior).

<sup>215</sup> See Herbert A. Simon, *A Behavioral Model of Rational Choice*, 69 Q. J. ECON. 99, 101 (1955); Simon, *supra* note 7, at 25.

<sup>216</sup> See George A. Miller, *The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information*, 63 PSYCHOL. REV. 81, 81–83 (1956) (classic statement of short-term/working-memory capacity limits). See also Alan Baddeley, *Working Memory*, 255 SCIENCE 556, 556–57 (1992) (surveying working memory as a limited-capacity system supporting storage and processing).

<sup>217</sup> See JOHN W. PAYNE, JAMES R. BETTMAN, & ERIC J. JOHNSON, *THE ADAPTIVE DECISION MAKER* 1–5 (1993) (arguing decision strategies are adaptive responses of a limited-capacity information processor to task complexity).

<sup>218</sup> Simon, *supra* note 7, at 110–11.

<sup>219</sup> See Herbert A. Simon, *The Architecture of Complexity*, 106 PROC. AM. PHIL. SOC'Y 467 (1962) (explaining why inferring system-level properties from interacting parts is nontrivial in complex systems).

<sup>220</sup> Simon, *supra* note 7, at 94.

<sup>221</sup> Herbert A. Simon, *Invariants of Human Behavior*, 41 ANN. REV. PSYCHOL. 1, 7 (1991).

of later combination.<sup>222</sup> Each of these operations increases the total computational capabilities of the administrative organization. Technologies like filing, email, and statistical analysis enable people working in organizations to extend their cognitive capacities in different ways and combine with each other to become more effective.<sup>223</sup> An EPA scientist modeling particulate matter dispersion need not also understand the legal constraints on agency action because the lawyers handle that. A Social Security claims examiner need not design eligibility criteria because the rulemakers already did.

But specialization creates its own problems, particularly regarding oversight and control within an organization. The same division of labor that enables expertise also fragments knowledge across the organization, as scientists, economists, lawyers, and managers struggle towards mutual comprehensibility. Coordination mechanisms like meetings, memos, and hierarchical review consume resources and introduce delays. And each person working in an administrative organization has an idiosyncratic set of goals that only imperfectly corresponds to the goals of the organization, which they pursue sometimes at the expense of the organization's goals.<sup>224</sup> The organization has various tools for aligning the goals of the individual to the goals of the organization or of increasing the costs of pursuing the individual's own goals at the expense of those of the organization.<sup>225</sup> But these alignment and oversight mechanisms are themselves costly for the organization to run, as managers must spend resources to monitor subordinates, provide incentives, and cultivate loyalty. Furthermore, as the difficulty of determining whether an official is substantively pursuing the goals of the organization increases because of specialization or the quantity of actions taken, so do oversight costs.<sup>226</sup> Where the administrative official (or the agency itself) is particularly expert or operating in a specialized decision domain, substantive oversight of their decisions can become impossible, making it very difficult for even her managers to determine whether the agency's true goals are in fact being pursued.<sup>227</sup> Because administrative agencies in particular function to improve the capacity of the public to solve complicated problems that require specialization, expertise, and technical capacity to understand and structure the decision space, these risks are high.

These dynamics are internal to the administrative organization. But the capability-accountability trap arises from a further problem: the relationship between the organization and its external overseers necessary to create democratic legitimacy in government. Courts, Congress, and the

---

<sup>222</sup> *Id.* at 19. Simon, *supra* note 7, at 39–40.

<sup>223</sup> See Lee Sproull & Sara Kiesler, *Reducing Social Context Cues: Electronic Mail in Organizational Communication*, 32 MGMT. SCI. 1492–94 (1986) (showing email alters organizational information exchange and coordination, not merely speed).

<sup>224</sup> See Michael C. Jensen & William H. Meckling, *Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure*, 3 J. FIN. ECON. 305, 308–10 (1976) (defining agency relationships and why agents' incentives diverge from principals').

<sup>225</sup> Simon, *supra* note 7, at 9.

<sup>226</sup> Jensen & Meckling, *supra* note 224, at 308–10 (agency costs include monitoring, bonding, and residual loss).

<sup>227</sup> Stephenson, *supra* note 16, at 53–56.

public must evaluate whether agencies are acting lawfully, effectively, and in the public interest. To do so, they need to understand what agencies are doing and why. This understanding is what I call *scrutability*: the cognitive tractability of administrative action for those who did not participate in producing it.

Scrutability is constrained by the same bounded rationality that constrains administration itself. Judges have limited time, small staffs, and no specialized training in the domains they review. Members of Congress face attention scarcity across hundreds of policy areas and rely on skeletal committee staff. Citizens have lives to live and cannot devote significant cognitive resources to monitoring agency behavior. When administrative action is simple, that is when the relevant considerations are few, the causal chains short, and the tradeoffs intuitive, external overseers can evaluate substance directly, like a nineteenth century court assessing whether a municipality had authority to grant a particular license.

However, when administrative action grows complex, scrutability declines. The shift from simple licensing to railroad rate-setting, from local poor relief to national social insurance, and from common law nuisance to probabilistic risk assessment each represented an increase in the sophistication required to assess whether an agency had acted correctly. Over time, government has become inscrutable.

Administrative law's response to declining scrutability has been the substitution of procedural for substantive review discussed in Parts I and II of this Article, because generalist overseers have the cognitive capacity to administer procedural compliance. Courts may not be able to evaluate whether an agency's scientific conclusions are correct, but they can determine whether the agency compiled an adequate record, considered relevant factors, and provided reasoned explanations for its choices. Congress cannot specify every technical determination, but it can require that agencies follow specified procedures before acting and can deputize affected parties to raise alarms when those procedures are violated. Courts are specialists in procedure and legal interpretation, so it is natural that, faced with a declining ability to substantively oversee administration as it became more complex and scientific, they would respond by moving oversight to the grounds where they are best equipped to provide it.

Modern administrative law operationalizes this procedural substitution. The record rule cabins review to what can be compiled and transmitted to a generalist overseer; *Chenery* requires agencies to stand on the reasons they actually gave; and *State Farm's* "hard look" centers oversight on whether the agency confronted relevant material and explained its choices, while information-forcing regimes like FOIA, the Sunshine Act, and NEPA push the same accountability work onto documentation and disclosure. This architecture produces real constraint, but it also induces agencies to invest in review-proof paper trails rather than better decisions.<sup>228</sup> Over time, those

---

<sup>228</sup> Bagley, *supra* note 27, at 351–52.

layers compound into a scrutability tax that crowds out substance, privileges repeat players experienced in handling procedural complexity, and makes engaging with the lengthy record impossible for parties with limited resources.<sup>229</sup> The structure of government becomes inscrutable alongside its content. However, those tradeoffs operate differently across the administrative state. Consider three types of administrative action, distinguished by the nature of the decisions involved and the role of technical expertise:

*Type 1: High-volume, low-discretion adjudication.* Consider benefits determinations, licensing decisions, permit approvals, and other administrative adjudications that involve applying relatively determinate criteria to individual cases. The administrative function in these domains is matching facts to legal categories: Does this applicant's medical condition meet the statutory definition of disability? Does this invention satisfy the non-obviousness requirement?

Procedural review still operates here, though it could be improved through statistical evaluations. Due process requirements<sup>230</sup> can be verified case by case, and error rates can be estimated through sampling and appeals, though measurement and good evaluation practices would have to be applied. The criteria used in these determinations are articulable *ex ante*, so reviewing courts can assess whether the agency applied the right standard even if they cannot evaluate the underlying factual determinations with confidence in every case appealed to them. The main challenges are scale (reviewing agencies make far more decisions than courts can supervise) and resource constraints (both agencies and claimants may lack capacity to document and contest individual determinations adequately). AI could improve this Type of adjudication, for example by generating personalized explanations of why an applicant is denied welfare benefits, but review works decently here where it can actually be applied.

*Type 2: Technical rulemaking with complex records.* Rules like EPA air quality standards and FDA drug approvals require synthesizing scientific evidence, economic analysis, and policy judgments into generally applicable requirements. The record in a major rulemaking is a complex artifact integrating multiple disciplinary perspectives that courts often lack expertise in. The agency must characterize the problem (what harms does the pollutant cause?), evaluate evidence (what does the epidemiology show?), model consequences (how will affected parties respond to different regulatory options?), and make tradeoffs (how much cost is justified to achieve how much risk reduction?). While Congress sometimes guides agencies in making these determinations, often there is little specific statutory guidance in how they should be done.

Procedural review faces difficulties in Type 2 contexts. Courts can check that the agency compiled a record and offered reasons, and sometimes even look to the quality of these reasons, but they

---

<sup>229</sup> See Wendy E. Wagner, *Administrative Law, Filter Failure, and Information Capture*, 59 DUKE L.J. 1321, 1324–28 (2010).

<sup>230</sup> See Henry J. Friendly, *Some Kind of Hearing*, 123 U. PA. L. REV. 1267, 1277–1304 (1975).

cannot evaluate whether a key scientific synthesis was sound on substantive principles. Sophisticated challengers learn to exploit the focus on procedure and box-checking that this dynamic generates, flooding the record with studies and objections that create procedural litigation hooks regardless of their substantive merit. Agencies respond in the “defensive crouch”, over-documenting their rulemaking and pre-answering criticisms at length, with the effect of extending rulemaking timelines and increasing costs without necessarily improving outcomes.<sup>231</sup>

*Type 3: High-salience, value-laden tradeoffs.* Some administrative decisions involve fundamental normative choices about which reasonable people disagree, even where there is scientific consensus. How should we balance tradeoffs between reducing emissions and disrupting economic production? What balance should be struck between drug safety and patient access? How should agencies weigh privacy against security in surveillance contexts? Politicians debate these questions in elections and in Congress, but often it is agencies that make the key determinations in their application in the world. These questions have technical dimensions, but their resolution ultimately depends on value judgments that science cannot dictate.<sup>232</sup>

Procedural review is least adequate for Type 3 decisions. The major questions doctrine reflects the sense among courts that some choices are too important and political for agencies to decide them without clear congressional delegation. But Congress often cannot or will not specify answers to these questions, because it may lack the technical capability to resolve them, because they may emerge after legislation is already passed, or because no new law can get through congressional gridlock. Agencies step in to resolve critical societal problems like climate change without clear authority to make the relevant value judgments, and procedural review cannot supply the fundamental democratic legitimacy that congressional authorization would provide.

This typology clarifies where AI might help and where it cannot. The scrutability problem is most tractable in Type 1 adjudications, where AI systems can match facts to criteria using their natural language capabilities and the main difficulties are scale and documentation. Indeed, AI can likely improve Type 1 adjudications significantly, both for regulated parties and for agencies. AI systems that can process applications, flag edge cases, and generate individualized explanations could dramatically reduce the scrutability tax while maintaining or improving accuracy. The challenge is ensuring that AI-assisted adjudication satisfies due process requirements, most significantly that affected parties can understand and contest the grounds for decisions affecting them, but having AI assistance would likely make these adjudications easier to understand.

Type 2 contexts present more substantial challenges. AI cannot resolve scientific uncertainty or make contested empirical questions disappear, but it can reduce the costs of producing and processing the complex records that technical rulemakings generate. Advanced AI is starting to be

---

<sup>231</sup> See Bagley, *supra* note 27, at 351–52.

<sup>232</sup> See Jasanoff, *supra* note 7, at 1–32.

used in different scientific domains to answer empirical questions, and plugging scientific models into LLM-based frontier AI systems that can model regulatory domains, summarize generated evidence, and translate technical findings into language that is comprehensible to agencies, their overseers, and the public could improve scrutability without pretending to resolve substantive disputes. The goal is not to replace expert judgment but to make that judgment more auditable, agencies benefiting from the capabilities of this advancing technology.

Type 3 contexts present the most serious risks from AI. Value-laden tradeoffs require democratic legitimation that technical systems, no matter how capable, cannot supply. This principle is core to the purpose of judicial oversight of the administrative state. AI systems that purport to resolve such tradeoffs without human participation risk “values laundering,” obscuring normative choices behind a veneer of algorithmic neutrality.<sup>233</sup> In major questions-type cases, AI should not make decisions but rather clarify what decisions are being made. These systems could surface tradeoffs, identify winners and losers from different policy alternatives, and present the choice in terms that democratic institutions can evaluate. AI systems must be designed to prevent AI-assisted framing from manipulating rather than informing deliberation, but there is a chance to significantly improve how administration handles this type of problem here.

More generally, AI might enable oversight to depend less on procedure and more on substance or even true alignment of agencies to the public will, reducing procedure’s attendant ossification. These systems can understand the world and perform research, identifying and using relevant information and demonstrating advanced knowledge of scientific domains critical for modern governance both in and outside the lab.<sup>234</sup> Multimodal systems can take in data of various forms and use coding tools to analyze it statistically, finding important patterns.<sup>235</sup> When deployed in agentic environments, these systems can identify problems and opportunities, and sometimes even solve them. Courts and Congress could ask AI systems about a given agency decision, creating a new means for low-cost provision of expertise that enables the overseeing bodies to at least somewhat substantively evaluate whether an agency acted correctly. Furthermore, advances in technical alignment, including reinforcement learning from human feedback, constitutional AI methods, and related techniques, can bring AI systems’ objectives closer to human ones, while interpretability might allow deeper examinations of AI reasoning than is possible for human

---

<sup>233</sup> See Andrew D. Selbst et al., *Fairness and Abstraction in Sociotechnical Systems*, in FAT\* ‘19: Proceedings of the Conference on Fairness, Accountability, and Transparency 59, 59–68 (2019). Though note that the process of scientific advising for political questions itself has been described as a kind of values laundering through the apparent neutrality of science. See Jasanoff, *supra* note 7, at 1–32.

<sup>234</sup> See Karan Singhal et al., *Large Language Models Encode Clinical Knowledge*, 620 NATURE 172 (2023) (evaluating LLM clinical knowledge); Bernardino Romera-Paredes et al., *Mathematical Discoveries from Program Search with Large Language Models*, 625 NATURE 468 (2024) (FunSearch discovers new algorithms/scientific results); Daniel J. Mankowitz et al., *Faster Sorting Algorithms Discovered Using Deep Reinforcement Learning*, 618 NATURE 257 (2023) (AlphaDev discovers improved sorting routines integrated into LLVM).

<sup>235</sup> See OPENAI, GPT-4 TECHNICAL REPORT 1 (2023) (multimodal inputs).

cognition. Truly aligning AI systems to the public good would mean that oversight of AI decisions is less necessary, reducing its social costs.

The next two Subparts develop these possibilities. Subpart B examines how AI can serve as scrutability infrastructure for participation and oversight—reducing the costs of understanding agency action and enabling more meaningful engagement by courts, Congress, and the public. Subpart C proposes accountability mechanisms suited to AI-assisted administration, including an expanded administrative record, review postures that reward transparency, and institutional architecture for ongoing audit. Together, they sketch a candidate settlement that uses AI to restore substantive accountability rather than merely adding new procedural layers to the accumulated sediment of past reforms, seeking to provide a positive vision for how to transform administration.

### *B. AI as Scrutability Infrastructure for Participation and Oversight*

The procedural substitution described in Subpart A was a reasonable response to declining scrutability resulting from the increased cognitive costs of handling the complexity of government. If overseers cannot evaluate substance, they can at least police process. But procedural proxies have limits. They verify that agencies followed required steps without confirming that those steps produced good outcomes, create compliance burdens that accumulate over time, leading to ossification, and they advantage sophisticated parties who can navigate procedural complexity over diffuse publics who cannot.

AI offers a different approach: rather than substituting procedure for substance, use technology to restore substantive comprehension and create more accountable but also acceptably complex government. If the scrutability problem arises because complex agency action exceeds the cognitive capacity of overseers and participants, then tools that reduce the cognitive costs of understanding or increase that cognitive capacity could ameliorate the problem at its source.<sup>236</sup> Here, AI serves as a kind of scrutability infrastructure that does not replace human judgment, but enables humans to exercise their judgment over material they could not otherwise comprehend without significant investment.

This Subpart examines three contexts where AI-enhanced scrutability could transform administrative governance: rulemaking, where AI can reduce barriers to meaningful public participation; adjudication, where AI can improve both the accuracy of initial determinations and the accessibility of review; and congressional oversight, where AI can help legislators focus scarce attention on consequential agency choices. It then addresses the risks, including manipulation, astroturfing, and inequality in the use of technologies, that AI-assisted participation creates and proposes initial safeguards to mitigate them.

---

<sup>236</sup> See Andy Clark & David Chalmers, *The Extended Mind*, 58 ANALYSIS 7, 7–12 (1998).

## 1. Rulemaking: From Mass Comments to Structured Contestation

Notice-and-comment rulemaking is the paradigmatic form of public participation in administrative governance. The APA requires agencies to publish proposed rules, accept written comments from interested parties, and respond to significant objections before issuing final rules.<sup>237</sup> In theory, this process enables democratic input into technical policymaking because affected parties can inform agencies about real-world impacts, identify errors in analysis, and propose superior alternatives.<sup>238</sup> In practice, notice-and-comment often fails to achieve these goals.

The failures of notice-and-comment are well-documented. Major rulemakings generate thousands or millions of comments, most of which are duplicative form letters.<sup>239</sup> Agencies cannot meaningfully handle this volume of comments, so they triage, prioritizing comments from sophisticated parties who frame objections in legally salient terms.<sup>240</sup> Someone writing in on a form letter seems much less likely to sue than an affected company with technical expertise, and understanding a technical comment requires more agency resources than noting a form letter written in broad terms. At the same time, ordinary people affected by a proposed rule often lack the knowledge that the rule exists and don't have the ability to understand the proposal or how it might affect them. Notice-and-comment thus becomes dominated by repeat players who can invest in technical engagement rather than operating as a true form of democratic input.

AI can help address these problems. Consider a layered AI interface that enables citizen participation while preserving the procedural protections of the APA:

*Layer 1: Record-anchored explanation.* Alongside the Notice of Proposed Rulemaking, agencies would publish AI-generated explanatory materials that translate the proposal into terms that normal citizens could understand along with a chatbot that could help people process what they are reading. These materials would identify what the rule changes, who is affected, what tradeoffs the agency is making, and what alternatives were considered, and the chatbot could answer questions based on the particular circumstances of affected parties. The chatbot's explanations would be anchored to the administrative record, and each claim it made would cite the underlying technical document, data source, or analytical assumption, enabling users to drill down into the basis for agency conclusions and acting as a knowledgeable guide to the rulemaking docket. This layer is achievable with current technology. LLMs can summarize complex documents, answer

---

<sup>237</sup> 5 U.S.C. § 553(b)–(c).

<sup>238</sup> See Stewart, *supra* note 11, at 1711–60.

<sup>239</sup> See Mendelson, *supra* note 189, at 1346–50, 1359–76.

<sup>240</sup> *Id.*

questions about their contents, and maintain citation links to source material.<sup>241</sup> Though there might be some risk of hallucinations, it could be mitigated, and red teaming to validate the accuracy of the chatbot would further reduce any risks here.

This layer should not create significant new procedural costs, and agencies could work with frontier AI providers to develop their chatbot aids on top of powerful existing tools. Furthermore, the chatbot explanations should be explicitly legally non-binding, aids to understanding the record but not themselves authoritative interpretations.<sup>242</sup> The agency's obligation remains to respond to comments on the actual proposed rule as documented in the record, and the AI layer simply reduces the cost of accessing that record for those affected by it.

*Layer 2: Structured commenting.* Beyond explaining what's in the record to affected parties, AI could help commenters articulate their views in forms that agencies would pay attention to and process. Current notice-and-comment accepts free-text submissions, creating an aggregation problem: how should an agency weigh a heartfelt but technically naive letter against a sophisticated critique? How much weight should a mass of form letters get?

A structured commenting interface would help citizens translate their concerns into legally cognizable form. The system could guide users through questions once it has helped them understand the record: What aspect of the proposal concerns you and how are you affected by it? What assumption or conclusion do you dispute? What evidence supports your position? What alternative would you prefer? What tradeoffs are you willing to accept? The AI would help users understand what they are agreeing to or contesting, identify the parts of the record relevant to their concerns, and frame their comments in terms the agency must address. Where the commenter does not really have a strong, personal objection, the system could help them understand that too, and guide them in expressing whatever political or other objection is causing them to want to comment.

This is more ambitious than Layer 1, requiring not just summarization but interactive dialogue that maps user concerns onto the structure of the rulemaking. It also raises design questions: How much should the system shape user input? Is there a risk that structured commenting forces diverse concerns into bureaucratic categories that lose important nuance? And how should agencies handle the reduction of friction that this system would entail, even though such a reduction might be necessary to increase the democratic legitimacy of rulemaking? These are real issues, but the system still represents an improvement on the status quo. The interface could also preserve unstructured input while offering structure for users who want help.

---

<sup>241</sup> See Paul Rust et al., *Evaluation of a large language model to simplify discharge summaries and provide cardiological lifestyle recommendations*, 5 COMMS. MED. 208, 208 (2025).

<sup>242</sup> See 5 U.S.C. § 553(b)(A).

*Layer 3: Comment aggregation and transparency.* Layers 1 and 2 have focused on how AI can improve the experience of affected parties, but agencies can also benefit from advanced AI. Agencies already triage comments because of the massive quantities they deal with, but the triage criteria are opaque. Which comments receive substantive response? How are duplicates handled? What weight attaches to different types of input? AI can make this process more systematic and transparent.

Comment clustering<sup>243</sup> would group submissions by the issues they raise, creating a map of public concerns that agencies must address. The agency's response-to-comments could be organized around clusters rather than individual submissions, showing how each significant objection was considered. Weighting rules would be disclosed. For example, duplicative comments might be counted for prevalence (indicating public salience) without being treated as independent technical evidence. Near-duplicate comments generated by AI systems could be flagged and handled appropriately.

This layer addresses the ossification critique directly. If agencies can process public input more efficiently by using AI to identify genuine substantive objections and document how each was addressed, the burden of notice-and-comment declines. Faster rulemaking becomes possible without sacrificing the democratic input that notice-and-comment is supposed to provide. The interface reduces the extent to which the process is a litigation-positioning exercise and increases the extent to which it is actual information transmission from the public to agencies.

## 2. Adjudication: Procedural Justice at Scale

Administrative adjudication presents different challenges than rulemaking. Where rulemaking produces generally applicable standards, adjudication applies standards to individual cases. The due process stakes are immediate and personal because they affect things like the disability benefits or immigration status of specific people. Procedural justice requires that affected parties understand the grounds for decisions affecting them and have meaningful opportunity to contest adverse determinations.<sup>244</sup>

However, high-volume adjudication has put serious strains on this requirement. The Social Security Administration decides millions of disability claims annually.<sup>245</sup> Immigration courts face crushing backlogs.<sup>246</sup> Benefits agencies across federal and state governments make countless

---

<sup>243</sup> See generally BRIAN S. EVERITT, SABINE LANDAU, & MORVEN LEESE, CLUSTER ANALYSIS (2009).

<sup>244</sup> See Friendly, *supra* note 230, at 1277–1304.

<sup>245</sup> SOC. SEC. ADMIN., FY 2024 ACTUAL PERFORMANCE: WORKLOAD AND OUTCOME MEASURES 1 (2024).

<sup>246</sup> Holly Straut-Eppsteiner, *FY2024 EOIR Immigration Court Data: Caseloads and the Pending Cases Backlog* 1, 4, CONG. RSCH. SERV. (IN12492, Jan. 24, 2025).

determinations that affected individuals may never fully understand.<sup>247</sup> The form letter denying a claim, citing regulatory criteria in boilerplate language, satisfies minimal due process requirements but provides little actual comprehension for the affected party. Claimants often do not know why they were denied, what evidence would support their case, or how to navigate the appeals process.

AI can improve procedural justice in adjudication across three dimensions: the quality of initial determinations, the comprehensibility of decisions, and the accessibility of appeals.

*Initial determinations.* AI-assisted case processing can improve accuracy and consistency in high-volume adjudication. Systems trained on past determinations could flag cases that warrant closer review, identify relevant evidence that adjudicators might miss, and detect inconsistencies across similar cases. This is Type 1 administrative action in the typology introduced in Subpart A: high-volume, relatively low-discretion application of criteria to facts. AI is well-suited to such tasks, not because it replaces human judgment but because it ensures that human adjudicators have relevant information and that similar cases are treated similarly. Centralizing adjudication in one or a few task-specific systems allows rigorous evaluation and correction in ways human-staffed adjudication cannot. When one of thousands of human adjudicators has idiosyncratic bias, correcting a single decision on appeal does little to improve overall system performance or likely even the later performance of that person. But a judicial decision requiring improvement of an algorithm propagates across every case the system handles, making system-level oversight tractable.

Such uses of AI present real risks. Automation bias, error propagation, and opacity all undermine procedural justice in AI administration.<sup>248</sup> Subpart C addresses the accountability mechanisms necessary to manage these risks. Here, the point is simply that AI-assisted initial processing can improve baseline accuracy in systems where current error rates are substantial and inconsistency is endemic. Human review has deep problems, and AI can help us do better.

*Comprehensible decisions.* AI can transform the explanation of adverse determinations. Rather than boilerplate denial letters, agencies could provide individualized explanations at much lower cost than having a human write it. For example, here is a simple explanation to an affected party in a disability hearing: “Your claim was denied because the medical evidence you submitted did not establish that your condition prevents you from performing sedentary work. The vocational expert found that jobs X, Y, and Z remain available to someone with your functional limitations. To succeed on appeal, you would need to provide evidence showing [specific deficiency].” While the relevant agency could write dozens of template form letters and have officials fill in the

---

<sup>247</sup> Donald Moynihan, Pamela Herd, & Hope Harvey, *Administrative Burden: Learning, Psychological, and Compliance Costs in Citizen-State Interactions*, 25 J. PUB. ADMIN. RSCH. & THEORY 43, 43–46 (2015).

<sup>248</sup> See Coglianese & Lehr, *supra* note 190, at 1213–17.

variables customized to each person after each adjudication, doing so would be costly for a human but cheap for an AI.

Such a step moves us toward constitutionally grounded due process. *Mathews v. Eldridge* requires that procedures be calibrated to the stakes involved and the risk of erroneous deprivation.<sup>249</sup> Individualized explanation reduces error by helping claimants identify when the agency has made a mistake and by focusing appeals on actual points of dispute. It would also improve the perceived legitimacy of adverse decisions because people are more likely to accept unfavorable outcomes when they understand the reasons.

Current large language models can generate such explanations.<sup>250</sup> The system would need access to the case record, the criteria applied, and the specific findings that led to denial. It would need to be accurate, which requires measuring and addressing hallucinations.<sup>251</sup> The system would also need safeguards against manipulation from within the agency, because the explanation should reflect the actual reasons for decision, not a *post hoc* justification designed to discourage appeals. However, human-written explanations present the same problem, and doctrinal requirements that they reflect the agency's true reasoning are of limited effect because it is difficult to see into the reasoning processes of human officials.

*Accessible appeals.* The same interface that explains decisions can help claimants with appeals. What additional evidence would strengthen a claimant's case? What are the deadlines? What procedural steps must they take to launch an appeal? The AI helps claimants who cannot afford representation by human lawyers, a partial remedy for the inequality that pervades administrative adjudication.

This application raises concerns about unauthorized practice of law.<sup>252</sup> The system should not provide legal advice in the sense of making rights-affecting recommendations about whether and how to appeal, but instead it should provide information about process and requirements that claimants are entitled to know. The line is blurry, but it is the same line that legal aid organizations and court self-help centers must work across daily. AI could dramatically expand claimants' access to useful information and provide procedural guidance that is currently available only to those with resources to obtain representation, and the benefits of such tools should not be overlooked.

---

<sup>249</sup> 424 U.S. 319 (1976).

<sup>250</sup> See Rust et al., *supra* note 241.

<sup>251</sup> See Yue Zhang et al., *Siren's Song in the AI Ocean: A Survey on Hallucination in Large Language Models*, arXiv:2309.01219 (Sep. 14, 2025).

<sup>252</sup> See *Modernizing Unauthorized Practice of Law Regulations to Embrace AI-Driven Solutions and Improve Access to Justice*, AI POLICY CONSORTIUM FOR LAW & COURTS (Aug. 2025).

### 3. Congressional Oversight: AI as Attention Multiplier

Congress is structurally incapable of overseeing the administrative state. Members have limited time, divided across hundreds of policy areas and the demands of campaigning and constituent service.<sup>253</sup> They have small staffs who might be able to help with certain areas, but cannot dive into more than a few problems their boss cares deeply about. Help in hiring more people and staffing up offices does not seem forthcoming. Committee staff possess expertise but are also stretched thin, and cannot match agency specialists who have spent careers in narrow domains and invested substantial time in the problems at issue.<sup>254</sup> The “information asymmetry” between agencies and their legislative overseers is vast, and agencies control the production of the records that oversight relies upon.<sup>255</sup>

AI cannot solve Congress’s fundamental attention scarcity or its political dysfunctions. But it can reduce the costs of informed oversight, enabling legislators to focus their limited attention on consequential agency choices rather than drowning in information. Consider an AI system that operates as a congressional query interface for major rulemakings. Staff could interrogate the rulemaking record through natural language, asking questions about the assumptions driving a cost estimate, where the agency acknowledges uncertainty, what alternative courses the agency considered or should have, and what statutory terms the agency is interpreting and relying on. These questions would be tailored for the interests of Members of Congress by virtue of being asked by their staff and would reduce the cost of sitting down with a complex record substantially, increasing congressional capabilities.

This capability matters most for allocation of interpretive authority, the first pillar of administrative settlements. Post-*Loper Bright*, courts will exercise independent judgment on statutory interpretation, but Congress retains primary authority to resolve ambiguity through legislation. The problem is that Congress often does not know which interpretive questions will matter until litigation surfaces them years later, and changing circumstances and new technologies can create new ambiguities where none existed before. AI-assisted oversight could identify interpretive pressure points as they arose through agency action or litigation. For example, a system could flag to an interested legislator or her staff that a new rule that has been published in the Federal Register interprets “stationary source” in a way that extends or shrinks the agency’s authority to regulate new pollution categories or that a new guidance interprets “waters of the United States” to include features that prior administrations excluded. Congress could legislate prophylactically, clarifying delegations before they generate litigation rather than after.

---

<sup>253</sup> Russell W. Mills & Jennifer L. Selin, *Don’t Sweat the Details! Enhancing Congressional Committee Expertise Through the Use of Detailees*, 42 LEGIS. STUD. Q. 611, 613 (2017).

<sup>254</sup> *Id.* at 611–13.

<sup>255</sup> McCubbins, Noll, & Weingast, *supra* note 16, at 250–51.

The interface could also surface the value-laden choices that are embedded in technical analysis. For example, cost-benefit analyses require assumptions about discount rates, the value of statistical lives, and the baseline against which impacts are measured.<sup>256</sup> Choices about what values to use encode normative judgments about how to weigh present against future harms, whose lives count, and what world we are comparing the regulatory intervention to.<sup>257</sup> AI systems can clarify the assumptions being made choices and transform them into something that legislators can engage with like a statement that “the agency’s conclusion that benefits exceed costs depends on valuing statistical lives at \$X, discounting future harms at Y%, and assuming baseline Z. Here is how the conclusion changes under alternative assumptions.”

None of this substitutes for congressional will. If legislators do not want to oversee agencies, perhaps because oversight is politically costly, or because they prefer to blame agencies for unpopular outcomes, AI cannot compel them to do so. But if Congress wants to exercise meaningful supervision, AI could reduce its cost. It reduces the extent to which Congress must rely on proxies like fire-alarm oversight<sup>258</sup> toward something more systematic without requiring resources that Congress does not have.

Implementation would require the development of institutional infrastructure. Congressional support agencies, including the Government Accountability Office, the Congressional Research Service, and the Congressional Budget Office could develop and maintain AI tools for legislative oversight. A revived Office of Technology Assessment<sup>259</sup> or the U.S. Center for AI Standards and Innovation<sup>260</sup> could provide independent technical analysis of agency systems and actions. The investments necessary for such oversight are modest compared to the significance of effective congressional oversight of AI administration.

#### 4. Safeguards: Preventing AI-Assisted Participation from Becoming Manipulation

AI-enhanced participation creates new risks that require new safeguards. The same tools that help citizens understand rulemakings could help bad actors flood agencies with synthetic comments. If agencies are not aligned to the public interest, the interfaces that translate complexity into publicly-comprehensible terms could frame issues in ways that manipulate rather than inform. Capture risks

---

<sup>256</sup> Office of Mgmt. & Budget, Exec. Office of the President, Circular A-4 (2023).

<sup>257</sup> See CASS R. SUNSTEIN, *THE COST-BENEFIT STATE: THE FUTURE OF REGULATORY PROTECTION* (2002).

<sup>258</sup> See McCubbins & Schwartz, *supra* note 167, at 166.

<sup>259</sup> See Darrell M. West, *It is time to restore the US Office of Technology Assessment*, BROOKINGS (Feb. 10, 2021), <https://www.brookings.edu/articles/it-is-time-to-restore-the-us-office-of-technology-assessment/>.

<sup>260</sup> See *Statement from U.S. Secretary of Commerce Howard Lutnick on Transforming the U.S. AI Safety Institute into the Pro-Innovation, Pro-Science U.S. Center for AI Standards and Innovation*, U.S. DEPT. COMMERCE (Jun. 3, 2025), <https://www.commerce.gov/news/press-releases/2025/06/statement-us-secretary-commerce-howard-lutnick-transforming-us-ai>.

from sophisticated parties seeking to add prompt engineers to their arsenal alongside lawyers are also real.

Four categories of risk require particular attention:

*Propaganda risk.* Under the model proposed here, agencies would control the AI interfaces that explain their proposals. What prevents an agency from designing explanations that favor its preferred outcome, for example emphasizing benefits, downplaying costs, or framing alternatives uncharitably? The risk is not that agencies will lie outright but that they will deploy the same rhetorical strategies through AI that they currently deploy through selective presentation in preambles and fact sheets, now with the imprimatur of technological neutrality.

Mitigation requires structural constraints on interface design. Agencies should be required to present credible alternatives and their strongest supporting arguments, treating the AI interface as a neutral guide rather than an advocacy tool. Review of interface outputs by OIRA or an independent AI-expert body could verify that explanations are balanced through sampling-based substantive review. Disclosure requirements could include interface design choices, enabling affected parties to identify and challenge bias.

*Astroturf risk.* Rulemakings are already being flooded by AI-generated synthetic comments.<sup>261</sup> The comment process will become even noisier, and agencies would face an impossible task distinguishing genuine public input from manufactured volume.

Mitigation requires authentication and provenance mechanisms. Identity verification (at least to the level of confirming that a comment comes from a real person) could filter out purely synthetic submissions. Provenance signals like watermarking or evaluations of whether the text is human- or AI-written could enable appropriate weighting without excluding AI-assisted comments entirely.<sup>262</sup> These detection techniques are still nascent but represent the beginnings of a way to raise the costs of astroturfing.

*Inequality risk.* AI tools are not equally accessible. Sophisticated parties will fine-tune their use of AI interfaces, optimizing prompts and exploiting system capabilities in ways that unsophisticated users cannot match. The risk is that AI-enhanced participation reproduces existing inequalities rather than ameliorating them as industry lawyers use AI to generate more effective comments and participate in more processes while ordinary citizens receive only marginal assistance.

---

<sup>261</sup> Mark Febrizio, *Will ChatGPT Break Notice and Comment for Regulations?*, GEO. WASH. UNIV. REG. STUD. CTR. (Jan. 30, 2023), <https://regulatorystudies.columbian.gwu.edu/will-chatgpt-break-notice-and-comment-regulations>.

<sup>262</sup> See Aloni Cohen, Alexander Hoover, & Gabe Schoenbach, *Watermarking Language Models for Many Adaptive Users*, arXiv:2405.11109 (Jun. 28, 2024).

This risk is real but should be evaluated against the current domination of participation by sophisticated parties. AI interfaces that provide meaningful assistance to ordinary participants improve on that baseline, even if repeat players benefit more. The goal is to ensure that citizens without resources can participate meaningfully, even if participation remains easier for those with resources. Interfaces should be optimized for accessibility, focusing on plain language, intuitive navigation, and responsive assistance for confused users rather than advanced features that only experts will use. Public investment in user-friendly tools can partially offset private investment in optimization.

*Capture risk.* Who builds and maintains the AI systems that agencies use? If vendors control the infrastructure of administrative participation, they acquire leverage over the agencies that depend on them. Vendor interests may not align with public interests and opacity and lock-in benefit vendors even when they harm governance.

Mitigation requires attention to procurement and governance. Agencies should retain audit rights over vendor systems, including access to model specifications and training data sufficient to evaluate system behavior. Portability requirements can prevent lock-in, enabling agencies to switch vendors or bring systems in-house. Developing evaluative capacity within government could help ensure that capture is identified and challenged where it emerges. For core administrative functions, there is a strong argument for public development of AI infrastructure, government-built systems that serve governmental purposes without private intermediation. This is a significant undertaking, but if AI becomes foundational to administrative participation, the case for public provision is compelling.

Taken together, these safeguards would help legitimate AI-assisted administration. AI-enhanced scrutability depends on design choices that prioritize public comprehension and democratic access. Subpart C turns from participation to accountability: how should administrative law adapt its mechanisms of record, review, and audit to a world where AI systems play significant roles in agency decision-making and ensure that AI-assisted administration remains subject to meaningful oversight?

### *C. AI as Accountability Infrastructure*

Subpart B described how AI can improve scrutability for those who participate in and oversee administrative governance. But participation and oversight require something to oversee: a record of agency action that makes reasoning visible and contestable. The administrative record, developed across the settlements examined in Part I, serves this function for human-driven administration. When agencies act, they compile documentation of the information considered, the alternatives evaluated, and the reasons for their choices. Courts review this record to determine

whether agencies have met threshold requirements of rationality and legality. The record is the interface through which external accountability operates.

AI-assisted administration requires a new kind of record. The “reasons” for an algorithmic determination exist, if anywhere, in model weights, training data, and activation patterns that no human wrote and few humans can interpret. If accountability depends on legible records, and AI reasoning is not legible in traditional forms, then accountability requires new forms of documentation.

This Subpart proposes the accountability infrastructure for AI-assisted administration. It begins with a corrective: AI does not eliminate complexity but relocates it, and the legal goal is auditable complexity rather than simplicity. It then proposes an expanded administrative record—a Model and System Dossier—that makes AI reasoning reviewable through auditing, considering how review should work, when it should occur, and what kind of institutional architecture might work best, and addressing remedies and failure modes.

### 1. The Relocation of Complexity

A threshold clarification: the AI administration proposed here does not make administration simple but rather shifts complexity from one part of government to another. AI can simplify the interface that courts, Congress, and the public use with administration while adding backend complexity in the form of model architectures, training pipelines, evaluation procedures, vendor relationships, security considerations, drift monitoring, and other things.

Relocating complexity is necessary because modern, scientific problems require complex government. However, complexity has ended up in the wrong places. The scrutability problem arose because the complexity of administrative action exceeded the cognitive capacity of overseers. AI addresses this by concentrating complexity in systems that can be audited through technical means, while presenting simplified interfaces to humans who cannot process the underlying detail. A court cannot evaluate whether an atmospheric dispersion model uses appropriate parameters, but it might be able to evaluate whether an agency followed appropriate procedures for validating the model, testing its performance, and monitoring its behavior over time. Courts shifted from substantive to procedural review of administration because procedure was comprehensible to them. The approach proposed here simplifies and strengthens procedural review while updating it for the AI era, using AI’s technical characteristics to make procedural review more efficient.

The legal goal, then, is not to eliminate complexity but to make it *auditable*. Auditability means that the behavior of AI systems can be examined, tested, and verified by parties with appropriate

expertise,<sup>263</sup> even if those parties are not the courts or legislators who ultimately hold agencies accountable. The accountability chain becomes: courts verify that agencies followed appropriate audit procedures; auditors verify that AI systems behave as claimed; the combination provides assurance that AI-assisted decisions are lawful and reasonable. This is not so different from current practice, where courts defer to agency scientific conclusions while verifying that appropriate analytical procedures were followed. The difference is that AI systems require different audit procedures than human experts, and administrative law must specify what those procedures are.

## 2. The AI Administrative Record: A Model and System Dossier

Adapting administrative accountability for AI requires an expanded administrative record that documents how AI systems are designed, deployed, and monitored. This “Model and System Dossier” would be a required addendum to the traditional administrative record that is required whenever an agency adopts a new AI system that it will rely on in rulemaking or adjudication. The Dossier, inspired by existing frameworks like the EU AI Act and its Code of Practice,<sup>264</sup> the Frontier Safety Frameworks of leading AI companies,<sup>265</sup> and model cards<sup>266</sup> and datasheets,<sup>267</sup> but adapted for administrative law, could include the following elements, among others:

*System identity, scope, and decision role.* What is the AI system, and what is it used for? This element defines the boundaries of the system subject to documentation and distinguishes consequential AI applications from trivial ones. A system that recommends denial of disability benefits requires extensive documentation, while a system that sorts incoming mail does not. Relevant questions might include the following: Is the AI system advisory or determinative? Does a human review every output, or only flagged cases, or none? What discretion does the human reviewer have? Do they empirically use their discretion? The description should specify both what the system does and what it does not do, clarifying the human functions that remain even when AI assists.

*Data provenance.* What data was used to train the system, and what data does it process in operation? This element enables evaluation of whether the system’s informational basis is

---

<sup>263</sup> See Edwin A. Farley and Christian R. Lansang, *AI Auditing: First Steps Towards the Effective Regulation of Artificial Intelligence Systems*, 38 HARV. J. L. & TECH. DIGEST (Feb. 19, 2025), <https://jolt.law.harvard.edu/digest/ai-auditing-first-steps-towards-the-effective-regulation-of-artificial-intelligence-systems>.

<sup>264</sup> European Commission, *General-Purpose AI Code of Practice* (July 10, 2025), <https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>.

<sup>265</sup> See Anthropic, *Responsible Scaling Policy* (Sept. 19, 2023, updated May 14, 2025), <https://www.anthropic.com/rsp>; OpenAI, *Preparedness Framework* (Dec. 18, 2023, updated Apr. 15, 2025), <https://openai.com/index/updating-our-preparedness-framework/>; Google DeepMind, *Frontier Safety Framework* (May 17, 2024, updated Sept. 22, 2025), [https://storage.googleapis.com/deepmind-media/DeepMind.com/Blog/strengthening-our-frontier-safety-framework/frontier-safety-framework\\_3.pdf](https://storage.googleapis.com/deepmind-media/DeepMind.com/Blog/strengthening-our-frontier-safety-framework/frontier-safety-framework_3.pdf).

<sup>266</sup> Margaret Mitchell et al., *Model Cards for Model Reporting*, in FAT\* ‘19: PROCEEDINGS OF THE CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY 220, 220–29 (2019).

<sup>267</sup> Timnit Gebru et al., *Datasheets for Datasets*, 64 COMM’NS ACM 86 (2021).

appropriate for its task. A model trained on historical decisions may encode past biases, and a model processing incomplete records may produce unreliable outputs. Complete data documentation may sometimes be infeasible, but agencies should document what they know about where their data is from, if there are gaps or limitations in it, and the steps they have taken to address data quality issues. Data poisoning can also present security risks, and agencies must be careful not to use any potentially-compromised data sources.

*Performance evaluation.* How well does the system perform its intended function? This element requires agencies to specify metrics appropriate to whatever task the AI is aimed at, like accuracy, consistency, error rates, and processing speed, and to report measured performance against those metrics. Testing the system before deployment on relevant metrics would reduce the chance that harm occurs from a non-performant system, though some problems may only become apparent through use. Effective performance evaluation requires disaggregated analysis: does the system perform differently for different populations, case types, or contexts? If the system does perform differently across protected groups, that would create a red flag requiring explanation and, potentially, remediation, under the law.

*Stress testing and red teaming.* How has the system been tested for failure modes? Adversarial testing can reveal vulnerabilities that ordinary performance evaluation misses, like inputs designed to fool the system, edge cases that the system struggles with, or scenarios where the system produces harmful outputs after a user or attacker manipulates it.<sup>268</sup> The science of red teaming advanced AI is improving, and should be integrated into agency evaluations of their systems before deployment. Agencies should document what testing was conducted, what vulnerabilities they identified in their systems, and what mitigation steps were taken. If vulnerabilities were not remediable, the agency should explain why they decided to go ahead with deploying a system anyway, and what processes they have in place to respond to any exploitation of the vulnerabilities.

*Monitoring plan.* How will the system be monitored in operation? AI systems can drift over time as the world changes and the relationship between inputs and appropriate outputs shifts. For example, a model trained on pre-pandemic disability claims may perform poorly on post-pandemic cases involving long COVID, and people can respond to the use of AI by changing their behavior, making it shift off target. Agencies should specify how they will detect drift, what thresholds trigger review, and what retraining or adjustment procedures are in place. Consistent monitoring would be cheaper and less harmful than monitoring through lawsuits, where harm to a person is what causes the agency to check if its system is operating incorrectly.

*Explainability and traceability outputs.* What can be reconstructed about the reasoning behind individual decisions? This element does not require that AI systems be fully interpretable or even

---

<sup>268</sup> See Jessica Ji, *What Does AI Red-Teaming Actually Mean?*, CENTER FOR SECURITY & EMERGING TECHNOLOGY (Oct. 24, 2023), <https://cset.georgetown.edu/article/what-does-ai-red-teaming-actually-mean/>.

explainable, but it requires agencies to document what explanatory outputs are available. Feature importance scores, attention patterns, citation to evidence in the record, natural language rationales generated by the system and the like should be documented, along with known limitations of these explanatory outputs. As interpretability and similar sciences advance, their outputs could be integrated into this process.

*Change log.* How has the system changed over time? Version control for AI systems enables accountability for system evolution. When was the model retrained? What changes were made to preprocessing or postprocessing? Why? The change log creates a historical record that courts and auditors can examine to understand how current system behavior relates to past documentation.

*Governance and accountability.* Who is responsible for the system? This element documents the organizational structure within the agency that is using the AI system, including who owns it, who can authorize changes, who reviews performance, and what escalation pathways exist when problems arise. It enables identification of responsible officials when accountability questions arise.

*Procurement and vendor constraints.* If the system was developed by an external vendor, what constraints govern the relationship? Agencies should document audit rights, data access provisions, portability requirements, and any trade secret limitations on disclosure. Where vendor claims of trade secrecy limit documentation,<sup>269</sup> agencies should explain what information is unavailable and why, enabling courts to evaluate whether the limitation is justified.

Of course, not every element of the Dossier requires exhaustive documentation for every system. Instead, it should be scaled to the stakes involved, requiring more extensive documentation for systems that make or influence high-stakes decisions and lighter documentation for lower-stakes applications. But the framework should be consistent, enabling cumulative learning about what AI documentation practices work and facilitating comparison across agencies and systems.

The Dossier becomes part of the administrative record, subject to the same preservation and disclosure requirements as other record materials. Challengers can contest agency action by identifying deficiencies in the Dossier, for example arguing that performance evaluation was inadequate, that known failure modes were not addressed, or that monitoring was insufficient. Courts can evaluate these challenges without themselves understanding the technical details of model architecture, asking instead whether the agency followed appropriate procedures for ensuring that the AI system works as intended.

---

<sup>269</sup> See Note, *Wisconsin Supreme Court Requires Warning Before Use Of Algorithmic Risk Assessments In Sentencing*, 130 HARV. L. REV. 1530, 1530–37 (2017).

### 3. Material Model Change: When Updates Trigger Process

AI systems are rapidly developing, as the course of frontier AI over the past few years demonstrates clearly. They are retrained as new data becomes available, adjusted as performance issues emerge, and updated as the underlying technology evolves. This constant change creates a problem for administrative law, which assumes relatively discrete agency actions subject to defined procedural requirements. When does an update to an AI system constitute “agency action” requiring notice and comment or other process?

The answer cannot either be “every update” or “no updates.” Requiring full process for every minor model adjustment would freeze AI systems by creating much more procedure, sacrificing the adaptability that makes them valuable and preventing full use of the benefits of the technology. However, a change that significantly alters how the system behaves is substantively equivalent to a change in the underlying rule, though measuring that alternation is a key part of the problem. Allowing such changes without process would circumvent the procedural requirements that legitimate agency action.

As such, the doctrine should track *material* model changes: updates that significantly affect outcomes, reasoning, or affected populations. Materiality is necessarily a standard rather than a rule, but its application can be guided by several substantive factors:

*Outcome effects.* Does the update change the rate at which something like disability benefit applications are approved or denied? Does it alter the distribution of outcomes across protected groups like by race or gender? Changes that move outcome metrics by more than some substantial defined threshold would presumptively require enhanced process.

*Reasoning effects.* Does the update change the factors that drive system outputs? A model that previously relied heavily on work history might, after retraining, rely instead on medical evidence. Even if aggregate outcomes are similar, the change in reasoning may affect how parties should present their cases and what evidence is relevant. For example, technical interpretability tools could indicate if the activations in the system’s neural network have changed when applied to a constant set of test cases before and after retraining to determine if a shift in reasoning has occurred.

*Population effects.* Does the update affect different populations differently? A change that improves overall accuracy while worsening performance for a particular demographic group could raise concerns about bias against that group.

*Architecture effects.* Does the update change the fundamental structure of the system rather than just updating parameters? Moving from one model architecture to another might warrant more scrutiny than retraining an existing architecture on new data.

When a material model change occurs, the agency should be required to update the Model and System Dossier, provide notice, and explain the expected impact of the change. This is less onerous than full notice-and-comment rulemaking but provides more notice than silent updates. The requirements should scale with the significance of the change.

The analogy to existing doctrine is the “logical outgrowth” test for final rules.<sup>270</sup> A final rule must be a logical outgrowth of an agency’s proposed rule at the beginning of rulemaking. If the final rule differs so substantially that affected parties could not have anticipated it, the agency must provide a new opportunity for comment.<sup>271</sup> Material model changes that alter the effective rule being applied trigger similar requirements, not because the formal rule has changed, but because the system implementing the rule operates differently. Congressional oversight of significant AI system changes might also provide a democratic check on *sub silentio* AI changes.

#### 4. Deference to Audit: A New Review Posture

How should courts review AI-assisted agency action? Not by deferring to agency expertise, as they have traditionally done, nor by evaluating algorithmic details directly, but by assessing whether agencies have subjected their AI systems to rigorous technical audits. Post-*Loper Bright*, courts exercise independent judgment on statutory interpretation. But they must also evaluate whether agencies acted arbitrarily or capriciously—and here, traditional approaches falter when algorithms contribute to agency reasoning.

The traditional answer would be to require agencies to explain their AI-assisted decisions in terms that courts can evaluate, to translate algorithmic outputs into *State Farm* reason-giving. This approach has surface appeal: it preserves familiar doctrinal forms and does not require courts to develop new technical competencies. But it has serious problems. It invites *post hoc* rationalization, explanations for AI outputs that do not reflect the actual drivers of the decision (if those drivers are understood at all), and does not actually verify that the AI system is working properly but only that the agency can tell a good story about the system’s outputs.

A better approach shifts from deference to expertise toward deference to audit. The question is not whether the agency’s AI-assisted reasoning is substantively correct, because courts will still struggle to evaluate that, but whether the agency has subjected its AI systems to appropriate scrutiny and documentation. If an agency complies with a robust audit regime, for example by producing the Model and System Dossier, conducting performance evaluation and stress testing, and maintaining monitoring and change logs, courts would treat the technical components of the

---

<sup>270</sup> See Patti Zettler, *The Logical Outgrowth Doctrine and FDA’s Intended Use Revisions*, YALE J. REG. NOTICE & COMMENT (Jul. 18, 2017), <https://www.yalejreg.com/nc/the-logical-outgrowth-doctrine-and-fdas-intended-use-revisions/>.

<sup>271</sup> *Id.*

agency's action as presumptively reasonable. The agency has earned deference not by claiming expertise but by demonstrating verification.

This safe harbor approach channels agency behavior in productive directions. Agencies that want judicial approval for AI-assisted decisions have a clear path: invest in the audit infrastructure that makes AI systems reviewable. Agencies that do not invest in this infrastructure would face more skeptical review. The doctrine does not require courts to become machine learning experts but to use their procedural expertise to evaluate whether agencies have followed appropriate procedures for ensuring that their AI systems are trustworthy.

The safe harbor proposed here is not absolute. Challengers can still prevail by identifying specific, record-supported flaws in agency AI systems, like evidence of bias that the agency's evaluation did not detect, failure modes that stress testing should have revealed, or drift that monitoring should have caught. The Dossier creates the evidentiary basis for such challenges, acting as both a shield for agencies that comply and a sword for challengers who can identify deficiencies. What challengers cannot do is simply assert that AI systems are untrustworthy without engaging the agency's documentation of system behavior.

This approach has antecedents in existing administrative law, expanding procedural review to the new AI context. Doctrinally, it keeps judicial review to something like procedure, where courts are expert, while importing substantive review through technically sophisticated auditing. It also responds to the post-*Loper Bright* environment. The current Court is skeptical of agency claims to deference based on expertise. But deference to audit is different: it is deference to verifiable compliance with procedural requirements that courts themselves specify. This is less a claim that agencies know better than a claim that agencies have done the work necessary to earn judicial respect for their technical judgments. Courts retain ultimate authority to define what audit requires and cede only the technical evaluation that they cannot perform themselves.

## 5. Institutional Architecture: Who Audits?

Deference to audit presupposes auditors. Courts can verify that agencies have produced documentation and followed specified procedures, but they cannot verify that the documentation is accurate or that the procedures are adequate. Technical scrutiny of AI systems requires technical capacity that courts lack and cannot develop. Three institutional mechanisms, operating in combination,<sup>272</sup> can provide the audit capacity that AI accountability requires:

*Internal agency audit units.* Agencies that deploy AI systems should develop internal capacity to evaluate those systems. Agency AI audit units would review system design, conduct performance

---

<sup>272</sup> It might also be worth incentivizing private development of auditing capacity. See Gillian K. Hadfield & Jack Clark, *Regulatory Markets: The Future of AI Governance*, arXiv:2304.04914 (Apr. 25, 2023).

evaluation, oversee stress testing, and monitor deployed systems for drift or failure. They would produce or verify the Model and System Dossier, certifying that documentation is complete and accurate. However, internal audit is necessary but not sufficient. Auditors embedded within agencies may lack independence or share the blind spots of the teams that developed the systems they review.

*External government auditors.* The Government Accountability Office (GAO) and agency Inspectors General could provide external audit capacity that can supplement internal review. GAO already evaluates agency technology programs,<sup>273</sup> so extending this function to AI systems is a natural evolution. Inspectors General can investigate AI-related failures and assess whether internal audit processes are adequate. External government auditors can also develop standards and best practices across agencies. What should a Model and System Dossier contain? What stress testing protocols are appropriate for different types of systems? What monitoring procedures detect drift reliably? However, these bodies would need enhanced technical capacity, requiring investment.<sup>274</sup>

*Centralized review for high-impact systems.* Not all AI systems warrant the same level of scrutiny. A chatbot that answers frequently asked questions poses different risks than a system that recommends benefit denials. Centralized review, analogous to OIRA review of significant regulations, could focus intensive scrutiny on the AI systems that matter most. The threshold for centralized review should track impact: systems that affect large numbers of people, that make or significantly influence high-stakes determinations, or that operate in sensitive domains, like law enforcement, immigration, and public benefits, warrant more scrutiny than routine administrative applications. OIRA, or a dedicated AI review function within OMB, could conduct this review, evaluating whether agency AI documentation is adequate and whether appropriate safeguards are in place. Centralized review also enables cross-agency learning. Agencies deploying AI for similar functions—eligibility determination, fraud detection, risk assessment—can learn from each other's successes and failures. A central review body can identify common problems and promote solutions across the government.

## 6. Remedies: What Courts Do When AI Governance Fails

When AI-assisted agency action is unlawful, what remedies are appropriate? The standard remedial toolkit of vacatur, remand, and injunction<sup>275</sup> applies, but AI systems present distinctive considerations.

---

<sup>273</sup> U.S. Gov't Accountability Office, *Science, Technology Assessment, and Analytics*, <https://www.gao.gov/about/careers/our-teams/STAA>.

<sup>274</sup> See West, *supra* note 259.

<sup>275</sup> 5 U.S.C. § 706(2) (2018).

*Vacatur versus remand.* Vacatur sets aside unlawful agency action while remand returns the matter to the agency for further proceedings. For AI-assisted adjudication at scale, vacatur may be impractical or harmful. Setting aside all determinations made with a flawed AI system could affect millions of beneficiaries, overwhelming agency capacity to redetermine cases and leaving claimants in limbo until their case can be reevaluated. Remand without vacatur may better balance remedial goals with practicability, though if the AI system systematically denied benefits to eligible claimants, vacatur may be necessary. If the flaw is procedural, remand to cure the deficiency may suffice.

*Targeted injunctions.* Injunctive relief can address specific failures in how an AI system has been developed or deployed without halting AI-assisted administration entirely. This is analogous to injunctions requiring specific procedural steps rather than enjoining agency action altogether.

*Disclosure under protective order.* Vendors sometimes claim that model specifications, training data, and performance evaluations are proprietary, and that their disclosure would harm competitive interests. Courts could respond through disclosure under protective order. Confidential submission to the court and, where appropriate, to challengers' experts, with restricted public dissemination, would avoid harm to the interests of companies and prevent trade secrecy from being an absolute bar to review.<sup>276</sup> If a vendor's system is so secret that no one can evaluate it, that system should not be making decisions that affect people's rights and undermines the democratic legitimacy of administration.

*Structural remedies.* When the use of AI in administration creates systematic problems, reflecting inadequate agency capacity, vendor capture, or institutional dysfunction, case-by-case remedies may be insufficient. Courts could order structural relief, for example requiring external audits of agency uses of AI or procurement reforms.

## 7. Failure Modes and Hard Problems

No accountability regime is perfect. AI-assisted administration will fail in ways that even robust oversight cannot prevent, and some failures will cause harm in ways outside the protection of procedural mechanisms. Some such potential problems follow, as well as the beginnings of ways to address them, but more research is needed in this area.

*Hallucinated rationales.* AI systems, particularly ones based on LLMs, can generate plausible-sounding explanations that do not reflect the actual drivers of an output or decision.<sup>277</sup> A system that denies a benefit claim might produce an explanation citing factors that were not, in fact,

---

<sup>276</sup> Fed. R. Civ. P. 26(c)(1).

<sup>277</sup> See Zhang et al., *supra* note 251.

determinative or that cites evidence that does not exist in the record. Agencies and courts may credit these explanations without recognizing their unreliability.

Mitigation requires grounding explanations in verifiable record materials. Explanations should cite specific evidence, citations should be checked against the actual record, and claims about what factors drove the decision should be tested against system behavior in similar cases. Technical interpretability research might also enable the evaluation of neuron activations inside LLMs to provide a substantive guide to their reasoning and what they considered. This process would not guarantee fidelity, but it reduces the risk of hallucination and creates opportunities for detection. Furthermore, hallucination rates have been declining and should be compared to the rate of incorrect statements by humans in the relevant decision-making areas.

*Automation bias.* Humans tend to defer to algorithmic recommendations in many domains.<sup>278</sup> For example, caseworkers reviewing AI-flagged cases might approve the system's recommendation without independently reasoning about the case or verifying if the AI was correct in its determination. Such deference would undermine the extent to which human review is actually present, even when legally required.

Mitigation requires institutional design that forces human officials to actually use their judgment. Random audits of human decisions can detect excessive agreement with AI recommendations. Override rates can be monitored, and if a given official never overrides a system or seems to do so randomly, there is no meaningful human oversight. Training can sensitize reviewers to automation bias. And for the highest-stakes decisions, procedures can require affirmative human justification with reasons given rather than simply affirming or denying AI outputs.

*Vendor capture.* When agencies depend on external vendors for AI systems, those vendors acquire leverage. Vendors may resist audit requirements that reveal system limitations, arguing that they must protect trade secrets, or design for lock-in, making it costly for agencies to switch providers. They may prioritize features that generate revenue over features that improve system performance as experienced by regulated parties. The public interest in accountable AI may conflict with vendor interests in profitable AI.

The prescriptions are familiar from other procurement contexts: audit rights in contracts, portability requirements that prevent lock-in, data access provisions that enable agency oversight, and, for core governmental functions, potentially even public development of AI systems to eliminate vendor dependence entirely. The case for public AI infrastructure is strongest where the stakes for democracy, rights, and public accountability are highest, including in systems that determine eligibility for public benefits, that inform enforcement decisions, or that affect

---

<sup>278</sup> See Jennifer M. Logg, Julia A. Minson & Don A. Moore, *Algorithm Appreciation: People Prefer Algorithmic to Human Judgment*, 151 ORGANIZATIONAL BEHAV. & HUM. DECISION PROCESSES 90 (2019).

constitutional rights. These domains should not depend on vendors whose interests may diverge from the public's.

*Value laundering.* AI systems may provide cover for value choices that should be made through democratic processes by laundering the underlying political question through opaque AI systems. An agency that announces a policy of denying benefits to certain claimants faces political accountability for that choice, while an agency that deploys an AI system that happens to deny benefits to those claimants can disclaim responsibility by blaming the system, if the underlying bias is noticed at all. When used in this way, AI becomes a mechanism for laundering contestable value judgments through the appearance of technical neutrality.

As discussed above, the requirement that agencies identify normative choices is one principal safeguard. When an AI system's behavior reflects value-laden parameters, like how much weight to place on different factors, what error rates to tolerate, or which tradeoffs to make, the agency must own those choices in the administrative record. The system embodies these values because we chose them, not because the algorithm decided. Courts reviewing AI-assisted action should scrutinize claims that "the model determined" an outcome, as the question of what choices humans made in designing, training, and deploying the model is always relevant. This requirement connects to the major questions doctrine. If an AI system's embedded value choices are significant enough to constitute major questions, those choices require clear congressional authorization regardless of their algorithmic implementation. AI cannot launder major questions into technical determinations that escape heightened scrutiny.

#### *D. Beyond Procedure: AI for Substantive Accountability*

The accountability mechanisms proposed thus far are mostly procedural in character, an odd turn in an Article partly about how procedure stifles government. They verify that agencies followed appropriate processes for developing and deploying AI systems without directly evaluating whether those systems reach correct conclusions. As I have argued, this is a familiar posture for administrative law, which has long substituted procedural for substantive review when the latter exceeds judicial competence.<sup>279</sup> And it invites an obvious objection: if AI accountability is just more procedure, will it not simply add to the ossification that already burdens administrative governance and prevent us from effectively adopting AI?

The objection has force against a regime that relies on procedure alone. But there are three responses: First, practically speaking, AI will have to be integrated into administration using the existing frameworks and competencies of administrative agencies and their overseers in the courts and Congress. The framework developed above integrates the particular characteristics of AI with

---

<sup>279</sup> See *supra* Part I.

administration to make near-term adoption by government possible. Second, the fundamental cost of procedure is that it creates a drain on limited human cognitive resources, which must be dedicated to operating and responding to it. If AI increases the cognitive capacity of administrators, their overseers, and the public more than these new procedures burden it, then it will have reduced net ossification in the system. Third, AI possesses characteristics that may enable forms of substantive accountability unavailable for human-staffed administration, potentially reducing the extent to which oversight must rely on ossifying procedure. Three areas deserve particular attention: technical alignment, interpretability, and scalable oversight. Together, they suggest that AI-assisted governance can be subject to verification mechanisms that go beyond process to examine whether systems are actually pursuing appropriate objectives through appropriate means. Though the technology is nascent and its course speculative, to the extent these mechanisms mature and prove reliable, they reduce rather than increase the need for procedural safeguards and so could help de-ossify government.

*Technical alignment.* Human employees of administrative agencies have their own goals, like career advancement, ideological commitments, cognitive ease, and personal relationships, that only partially overlap with their agencies' statutory missions.<sup>280</sup> Organizations attempt to align employee behavior through compensation, promotion, discipline, professional socialization, and hierarchical oversight, but these mechanisms are costly and imperfect.<sup>281</sup> They shape behavior at the margin but cannot ensure that employees genuinely internalize organizational objectives. The principal-agent problems that pervade bureaucracy reflect the fundamental difficulty of aligning human goals to institutional purposes.<sup>282</sup>

AI systems can be aligned through different mechanisms. Reinforcement learning from human feedback fine-tunes model outputs to match human evaluators' preferences.<sup>283</sup> Constitutional AI extends this approach to AI graders, making it more efficient and effective,<sup>284</sup> while deliberative alignment extends it to AI "reasoning."<sup>285</sup> Reward modeling and preference learning shape system behavior toward specified goals. These techniques, and their cousins, are still nascent and face problems like reward hacking,<sup>286</sup> goal misgeneralization,<sup>287</sup> and alignment faking.<sup>288</sup> But they operate on the system's objectives more directly than the external incentives used to align humans. A technically aligned AI system pursues its specified objectives because its training has in some

---

<sup>280</sup> Avinash Dixit, *Incentives and Organizations in the Public Sector: An Interpretative Review*, 37 J. HUM. RES. 696, 696–700, 708–12 (2002).

<sup>281</sup> *Id.*

<sup>282</sup> See Michael C. Jensen & William H. Meckling, *Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure*, 3 J. FIN. ECON. 305, 308–10 (1976).

<sup>283</sup> Ouyang et al., *supra* note 45.

<sup>284</sup> Bai et al., *supra* note 46.

<sup>285</sup> Guan et al., *supra* note 47.

<sup>286</sup> Dario Amodei et al., *Concrete Problems in AI Safety*, arXiv:1606.06565 (Jul. 25, 2016).

<sup>287</sup> Lauro Langosco et al., *Goal Misgeneralization in Deep Reinforcement Learning*, 162 PMLR 12004 (2022).

<sup>288</sup> Ryan Greenblatt et al., *Alignment faking in large language models*, arXiv:2412.14093 (Dec. 18, 2024).

sense made those objectives its own, not because external rewards and punishments make compliance instrumentally rational, representing a real improvement on human alignment.

The implications for administrative accountability are significant. To the extent an AI system is reliably aligned to statutory objectives, the oversight burden shifts. We need not verify through procedural proxies that the system is trying to achieve the right goals or even check its work substantively because we can have reasonable confidence in its objectives and focus oversight on whether it is achieving them effectively. Alignment does not eliminate accountability—systems can be aligned to the wrong objectives or can fail to achieve objectives they are aligned to—but it changes what accountability through oversight must accomplish and lessens the extent to which it is necessary. Procedural requirements designed to detect or deter goal divergence become less necessary when goal alignment can be technically verified. Procedural requirements substitute for the alignment we cannot achieve easily in humans. AI alignment techniques, however imperfect, offer a path toward reducing this substitution.

*Interpretability.* Human reasoning is opaque. We can ask an expert why she reached a conclusion, but her introspective report may be inaccurate, incomplete, or strategically constructed. Cognitive science has demonstrated that humans often cannot accurately report the factors that influenced their judgments and that *post hoc* rationalization is the norm rather than the exception.<sup>289</sup> This has serious implications for administration, which relies upon delegating important decisions to opaque human agents. When a claims examiner denies a disability application, we cannot verify that the stated reasons reflect the actual cognitive process that produced the denial. The examiner herself may not know.

This opacity is a primary driver of procedural substitution. If we cannot examine the substance of expert reasoning, we demand that experts document their reasoning in reviewable form, like the records, explanations, and responses to objections that make up so much administrative paperwork. This documentation requirement creates accountability, but it is accountability for what experts say they did, not for what they actually did. Sophisticated actors learn to produce documentation that survives review regardless of the underlying reasoning process.

AI systems are not introspectively transparent, but they are mechanistically inspectable in ways that human brains are not. Interpretability research aims to understand how neural networks process information and produce outputs, including which features activate in response to which inputs, how information flows through network layers, and what internal representations the system constructs.<sup>290</sup> The field is young, and current techniques provide limited insight into large model behavior. But to the extent that interpretability works, it enables a form of substantive

---

<sup>289</sup> Richard E. Nisbett & Timothy D. Wilson, *Telling More Than We Can Know: Verbal Reports on Mental Processes*, 84 PSYCHOL. REV. 231 (1977).

<sup>290</sup> See Santhosh Kumar Ravindran, *Adversarial Activation Patching: A Framework for Detecting and Mitigating Emergent Deception in Safety-Aligned Transformers*, arXiv:2507.09406 (Jul. 12, 2025).

review that human oversight of human experts cannot achieve. Rather than asking the AI system to explain itself and having to trust that explanation, reviewers could examine the system's actual reasoning process. Did the model attend to legally relevant features of the input? Did its internal representations track the distinctions that the governing legal standard requires? Did the decision pathway reflect appropriate weighting of relevant factors? These questions probe the substance of reasoning, not just its process.

Interpretability-based review would not replace traditional accountability mechanisms entirely. Interpretability techniques require technical expertise to apply, and there is debate about when interpretations of black box models are meaningful.<sup>291</sup> But interpretability outputs could be translated into reviewable artifacts like reports documenting what the system attended to, how it weighted factors, or where its reasoning diverged from expected patterns. These artifacts enable substantive review at one remove: courts could verify that appropriate interpretability analysis was conducted and evaluate its findings, much as they currently verify that appropriate scientific analysis was conducted and evaluate agency reliance on it.

The combination of alignment and interpretability suggests a different accountability posture than pure proceduralism. For AI systems that are both technically aligned to statutory objectives through validated alignment techniques and subject to interpretability analysis verifying that their reasoning processes track legal requirements, substantive accountability becomes tractable in ways it is not for human administrators. The procedural requirements that exist because substantive review was impossible can be relaxed when substantive review becomes possible. As such, the burdens of paperwork could be lifted while actually producing more accountable government.

*Scalable oversight.* A further characteristic of AI enables accountability at scale: AI systems can oversee other AI systems. Human oversight of administrative action faces the fundamental bottleneck of limited human time, attention, and expertise. The millions of individual determinations that agencies make annually cannot all receive human review.<sup>292</sup> Sampling, prioritization, and triage are necessary compromises, but most agency actions escape external scrutiny entirely.

AI oversight does not face the same constraint. A system designed to audit other systems can process every case, not just a sample, especially if the costs of auditing are less than making the initial decision. It can flag anomalies, detect patterns, and identify cases warranting human review, compare decisions across adjudicators to detect inconsistency, monitor outcomes over time to detect drift, and examine reasoning processes to detect deviation from expected patterns. We might

---

<sup>291</sup> See Cynthia Rudin, *Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead*, 1 NATURE MACHINE INTELLIGENCE 206 (2019); Gillian K. Hadfield, *Explanation and justification: AI decision-making, law, and the rights of citizens*, SCHWARTZ REISMAN INSTITUTE FOR TECHNOLOGY AND SOCIETY (May 18, 2021), <https://srinstitute.utoronto.ca/news/hadfield-justifiable-ai>.

<sup>292</sup> Bowman et al., *supra* note 50.

prefer human to AI oversight for normative reasons, but the alternative is often not human review but no review, and such a system would extend accountability to actions that might otherwise escape review entirely. Furthermore, human attention could be concentrated on the hard or normatively important Type 3 cases that AI surfaces, providing broader coverage than human review alone while preserving human authority over consequential judgments.

*Implications for ossification.* If alignment, interpretability, and scalable oversight mature as accountability mechanisms, they change the calculus of procedural design. Procedural requirements exist in significant part because substantive review is impossible. We cannot verify that agencies are pursuing appropriate objectives through appropriate means, so we require them to follow processes that create proxies for such verification. But if AI systems can be aligned to appropriate objectives (verifiably), if their reasoning processes can be examined (through interpretability), and if their outputs can be monitored comprehensively (through scalable oversight), the need for procedural proxies diminishes.

This does not mean procedures become unnecessary. Process requirements serve multiple functions beyond enabling substantive accountability. They provide participation opportunities, create deliberation-forcing mechanisms, and establish legitimacy through public engagement. These functions persist even if substantive accountability becomes more tractable. But the accretive proceduralism that characterizes modern administrative law, each oversight failure generating new procedural requirements layered atop existing ones, can be arrested if substantive verification becomes possible. Procedures that exist solely because substance was unreviewable can be streamlined or eliminated when substance becomes reviewable.

The anti-ossification potential of AI accountability is thus twofold. First, the mechanisms proposed in this Part—the Dossier, audit-based review, centralized oversight—are designed to be standardized and automatable, reducing the case-by-case procedural burden that slows agency action and increasing the cognitive capacity available to handle it. Second, to the extent alignment and interpretability enable substantive accountability, they reduce the need for procedural requirements that exist only as substitutes for substance. The result is not more procedure piled on existing procedure, but a rebalancing toward substantive verification that procedure was always a second-best replacement for.

This rebalancing is contingent on technical progress that remains uncertain. Alignment techniques may plateau short of reliable goal verification. Interpretability may never fully illuminate large model behavior. Scalable oversight may introduce its own errors and biases. The argument here is not that substantive AI accountability is achieved but that it is achievable, that AI possesses characteristics enabling forms of verification that human-staffed administration does not. Administrative law should be designed to take advantage of these characteristics as they mature,

creating incentives for agencies to invest in alignment and interpretability rather than treating procedure as the only path to legitimacy.

#### CONCLUSION: CRISIS AS OPPORTUNITY

Since the Industrial Revolution, American administrative law has followed a characteristic pattern: technological disruption creates capability crises that overwhelm existing institutions, government complexifies in response, accountability gaps emerge, and doctrinal settlements restore equilibrium. The railroads demanded the ICC. The Great Depression demanded the New Deal agencies and the APA. Environmental and safety regulation demanded computerization and *Chevron*. Each settlement was forged in crisis, when the inadequacy of existing institutions became undeniable and the political barriers to transformation temporarily fell.

The current moment is anomalous. The Supreme Court is dismantling administrative law's foundations without a precipitating capability crisis. No new technological disruption has overwhelmed governmental capacity; if anything, government has failed to exploit the transformative technologies of the internet and narrow AI that have revolutionized everything else. The Court's retrenchment responds not to governmental overreach but to governmental inscrutability. Administration has become too complex for generalist overseers to comprehend, and the Court has opted to shrink it back to manageable size. This is a scrutability intervention masquerading as an accountability intervention, sacrificing capability for comprehensibility at precisely the moment when complex problems demand sophisticated responses.

But a capability crisis is coming. Climate change, pandemics, and advanced AI present problems for which existing regulatory frameworks are unprepared, and these problems will not wait for government to resolve its internal contradictions. They will force resolution, one way or another.

History suggests that such crises create openings for institutional transformation that are otherwise unavailable. The procedural sediment described in Part II accumulated precisely because, absent crisis, incumbent interests can block reforms that threaten their bargaining positions. This is why the internet did not transform government as it transformed everything else: the political economy of procedural complexity blocked the fundamental process redesign that digital technology could have enabled.

Crises change this calculus. When they come, the political barriers to transformation temporarily fall. Reforms that were blocked for decades become possible in months, but they may not be enacted. The question is whether, when the next crisis arrives, reformers will be prepared to build institutions that escape the capability-accountability trap rather than merely adding another layer of procedural sediment atop the accumulated weight of past settlements.

This Article has argued that AI offers tools for a different kind of settlement. Unlike previous technologies, which enhanced capabilities while reducing scrutability, AI possesses characteristics that could improve comprehensibility even as it enables more sophisticated administration. AI can translate complex technical material into accessible forms, process massive records to surface the assumptions and tradeoffs that matter for oversight, and help citizens understand and engage with government in ways that procedural formality has long precluded, creating infrastructure for participation and oversight that reduces the cognitive costs of democratic engagement.

But the more fundamental contribution lies in AI's potential to restore substantive accountability—verifying that government is actually pursuing appropriate objectives through appropriate means, not merely that it has followed required procedures. Technical alignment, interpretability, and scalable oversight offer mechanisms for such verification that do not exist for human-staffed administration. If these capabilities mature, and their trajectory, while uncertain, is promising, they could enable a settlement that previous technologies could not. Procedure substituted for substance because courts could evaluate whether agencies followed steps but not whether they reached correct conclusions. AI changes this calculation. To the extent alignment can be verified and reasoning examined, the need for procedural proxies diminishes. The scrutability problem that drove proceduralization might be addressed not by adding more procedures but by restoring the substantive oversight that procedure was always a second-best replacement for.

The institutional architecture sketched in Part III.C provides a framework for this settlement that is implementable with current technology, compatible with existing institutions, and designed to provide safeguards against failure while enabling benefits as AI capabilities mature. Without such a framework, crisis-driven reforms may reproduce past patterns: expanding agency authority to address immediate problems, layering new procedural requirements to address accountability concerns, and leaving the capability-accountability trap intact for the next generation to confront. The sediment will continue to accumulate.

The opportunity before us is to break this cycle. The capability crisis that climate change, pandemics, and AI risks will generate can be the occasion for building administrative institutions that are both more capable and more accountable than any that have come before, dissolving the familiar tradeoff between sophistication and comprehensibility. The historical pattern shows that such transformation is possible: the ICC was at least as radical an innovation in 1887 as anything proposed here, and the New Deal agencies seemed to many contemporaries a fundamental break with constitutional government.

This Article has sought to provide a framework for thinking about governmental transformation before it becomes urgent. The capability-accountability trap has structured American administrative law for over a century. It need not structure the next century. If the reformers who

built previous settlements could match the railroad with the expert commission, the Depression with the administrative procedure, and complex science with deferential review, this generation can match the age of artificial intelligence with institutions worthy of its challenges. The question is not whether crisis will come but whether we will be ready when it does.