10

PART II · THE INTELLIGENT STATE

Capable and Accountable Administration in the Time of AGI

Nicholas Caputo



Federal Radio Commission staff surrounded by accumulated correspondence, 1929. The volume of mail became so overwhelming that the commission sought congressional permission to destroy old letters.

National Photo Company / Library of Congress.

SUMMARY

Since its inception, American administrative government has been caught in a capability-accountability trap. Governing a complex society demands a sophisticated state, but sophistication renders that state opaque to its generalist overseers in courts and Congress. The law's recurring answer has been to substitute procedural for substantive review, but this has ossified government. AGI threatens to deepen the trap from both directions, drowning agencies in automated objections while letting them launder arbitrary choices through machine-generated records. However, used correctly, AGI could break the trap by making government more scrutable, comprehensible to its overseers and to the public, rather than less. This essay sketches administrative law for AGI based on a new administrative record for AI, a "deference to audit" standard of judicial review, and AI tools that let Congress and the public oversee agencies at low cost, reenabling substantive oversight and providing the foundation for accountable administration in the AGI era.

IN 2021, THE ENVIRONMENTAL PROTECTION AGENCY (EPA) DECIDED TO RECONSIDER the national ambient air quality standard for particulate matter. Following the informal rulemaking procedures laid out in the Administrative Procedure Act of 1946 (APA) and the judge-made doctrines of administrative law that have accumulated since, after three years, 700,000 public comments, and the writing of a 200-page preamble legally intended to be "concise general statement of [. . .] basis and purpose," it changed a single number, lowering the primary annual PM2.5 level from 12.0 $\mu g/m^3$ to 9.0 $\mu g/m^3$.[1] Upon publication of the new rule, the EPA was sued immediately, and on March 12, 2025, announced that it would reconsider the rule again, restarting the process.

It is a commonplace to complain that American administration is slow and unresponsive, even when it works. This form of government, built to handle the challenges of the rapid technological progress of the late nineteenth and twentieth centuries, has ossified even as the challenges it confronts have grown more difficult, the legal doctrines used to mediate between expert power and democratic control losing their ability to find a balance.

Administrative law and its procedures seem to be a particular culprit, and AGI will make this situation worse. On the one hand, procedural rules designed to give voice to regulated parties bound by human cognitive constraints are about to come under incredible pressure with AI. Regulated parties already produce huge volumes of comments that agencies must address in rulemakings.[2] It's trivial to imagine them building an automated platform that writes and submits every conceivable objection to a proposed rule, overwhelming the human regulators on the other side in a kind of regulatory Distributed Denial of Service (DDoS) attack. Technical measures against DDoS-by-comment will help, but a smart adversary can likely train against them faster than government can deploy them.

On the other hand, administrative agencies themselves could use AGI to circumvent the congressional and judicial oversight that has long aimed to keep them accountable, finding new ways to exercise arbitrary and unlawful power. Administrative law assumes that humans review data and write documents, and that the existence of a document is proof that a person thought about its contents. But what if an agency simply decides on a course of action and then has an AGI come up with a justification and answer any objections? Pretextual reasoning of this kind is unlawful, but could a court tell, if the administrative record looks as good as it used to?[3]

These dangers point to a deeper problem, one that AGI might actually help solve. As I have argued,[4] since at least the dawn of the modern administrative state, American

1 Reconsideration of the National Ambient Air Quality Standards for Particulate Matter, 89 Fed. Reg. 16,202 (March 6, 2024) (to be codified at 40 C.F.R. pts. 50, 53, 58); see 5 U.S.C. § 553(c) (informal rulemaking).

2 On the duty to respond to significant comments, see *Motor Vehicle Mfrs. Ass'n v. State Farm Mut. Auto. Ins. Co.*, 463 U.S. 29, 43 (1983); 5 U.S.C. § 553(c).

3 See *SEC v. Chenery Corp.* (Chenery I), 318 U.S. 80, 87 (1943) (agency action must stand or fall on the grounds the agency invoked in the record).

government has been structured by a capability-accountability trap. Technological society demands more capable, scientific government. To maintain freedom and basic respect for rights, that government must also be kept accountable, mostly through judicial and congressional oversight. But expert administration is difficult for generalists to understand, so substantive review of agency actions was (mostly) abandoned early on in favor of procedural review. Then procedure piles up, leading to ossification, which reduces administrative capability.[5] American administrative law has moved through a series of such settlements, wherein government adopts new technologies to handle new problems, an accountability backlash follows, and administrative law doctrine rebalances capability against accountability through procedure.

But more procedure doesn't actually lead to greater understanding. As Justice John Paul Stevens reportedly remarked in the *Chevron* conference in 1984, "when I am so confused, I go with the agency."[6] Whether an agency did its procedural homework is a proxy for whether it thought hard and impartially, but only a proxy. And the public, interacting with agencies, now must handle procedures that take lawyers to navigate. Administration has thus become doubly incomprehensible: substantively because of the required scientific expertise and procedurally because of piled-up doctrine.

AGI can help. The capability-accountability trap is at bottom about allocating complexity. A complicated world demands complex government and oversight able to handle it. Managing information complexity is perhaps what AGI will do best. Frontier systems already read and summarize massive corpora, surfacing what matters so people can act on it. The same technology that enables DDoS-by-comment could be turned on the problem, clustering and answering at machine speed, doing science and presenting results.

Every prior settlement bought accountability by piling on procedure, making government inscrutable. AI may be the first administrative technology that can deliver capability and scrutability together, the first that doesn't force a trade and so enables both capable and accountable government.[7] Every administrative task—scanning the environment for problems, modeling interventions, gathering input, and implementing decisions—could be supplemented by AI. Imagine routine welfare adjudication run through a single AI rather than thousands of caseworkers, humans keeping the value judgments and the final appeal. A single system attends to every case alike and can be evaluated and improved directly through statistical tests of accuracy and fairness

4 For a much more thorough discussion of my perspective, see Nicholas Caputo, "Administrative Law's Fourth Settlement: AI and the Scrutable State," SSRN preprint, January 10, 2026.

5 See Thomas O. McGarity, *"Some Thoughts on 'Deossifying' the Rulemaking Process*," *Duke Law Journal* 41 (1992): 1385–1387; Nicholas Bagley, *"The Procedure Fetish*," *Michigan Law Review* 118 (2019): 345, 352–365.

6 *Chevron U.S.A. Inc. v. Natural Resources Defense Council, Inc.*, 467 U.S. 837 (1984). The remark is reported in Mark Sherman, "High Court Climate Case Looks at EPA's Power," *Huffington Post*, February 23, 2014.

7 The term inverts James C. Scott's "legibility": where legibility is the comprehensibility of the world to the state, scrutability is the comprehensibility of the state to the world. See James C. Scott, *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed* (New Haven: Yale University Press, 1998), 2–3; Andrew D. Selbst and Solon Barocas, *"The Intuitive Appeal of Explainable Machines*," *Fordham Law Review* 87 (2018): 1085, 1094.

where humans are often biased, tired, and inattentive.[8] An AI could likewise hand the EPA a menu of modeled interventions for the air-quality rulemaking above.

This capability must come with accountability, and AGI can help there too. But scrutable government means more than a system that writes some text explaining itself. A plausible explanation is exactly what a pretextual agency, or a model rationalizing after the fact, already produces. The test is not whether a decision can be explained but whether an independent auditor can verify that it was made well.[9] Advances in methods like interpretability research may eventually let an auditor read a model's reasoning or look at other technical measures of compliance,[10] but even if it never fully does, behavior can be audited at scale, itself the kind of knowledge work that AGI makes cheap. We need not read the machine's mind to hold it to account.

That's a rosy picture. But even if it pans out technically, courts, Congress, and the public may lack the capacity to use these tools. A new era of administrative law can help. First, AGI administration needs a new administrative record, in two parts: a "Model and System Dossier" documenting information including the AI's provenance, training data, and evaluation and red-teaming scores;[11] and a record of each decision, its inputs, outputs, explanations, and neural activations. When an agency is sued, these records would provide grounds for review.

Second, courts should defer to agency decisions that have undergone thorough auditing and evaluation. Courts already defer to the substantive expertise of agencies, focusing review on procedure. "Deference to audit" would keep courts in that role as guardians of procedure, checking whether an audit was adequate, while shifting deference from the agency to third-party auditors who have an incentive to identify problems in the agency's AI. An expert body like the U.S. Center for AI Standards and Innovation could accredit auditors and set standards.[12] Congress and the public could benefit too, including by having AI agents that investigate agency actions and explain relevant ones in plain language, reducing the cost of oversight.

Each earlier period of administrative law was forged in a crisis that briefly lowered the barriers to reform: the railroads produced the Interstate Commerce Commission, the Depression the Administrative Procedure Act, environmental degradation and the computing revolution the *Chevron* settlement.[13] Each bought new capability at the price of new procedure, and the sediment kept rising, eventually burying government. Today,

8 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*," *Cornell Law Review* 59 (1974): 772.

9 See Cynthia Rudin, *"Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead*," *Nature Machine Intelligence* 1 (2019): 206; Gillian K. Hadfield, *"Explanation and Justification: AI Decision-Making, Law, and the Rights of Citizens*," Schwartz Reisman Institute, May 18, 2021; see also Richard E. Nisbett and Timothy D. Wilson, *"Telling More Than We Can Know: Verbal Reports on Mental Processes*," *Psychological Review* 84 (1977): 231.

10 See Finale Doshi-Velez and Been Kim, *"Towards a Rigorous Science of Interpretable Machine Learning*," arXiv preprint, arXiv:1702.08608, March 2, 2017.

11 On red-teaming, see Jessica Ji, *"What Does AI Red-Teaming Actually Mean?"* Center for Security and Emerging Technology, October 24, 2023.

12 U.S. Department of Commerce, "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," June 3, 2025.

with AGI, we face another crisis and the opportunity that it brings. The technology forcing government to complexify could, for the first time, also render its complexity legible. A century of administration has perfected the craft of making the people legible to the state. A fourth settlement would, at last, help resolve the tension between administrative capability and accountability—making the state both more effective and more legible to the people it is meant to serve.

ABOUT THE AUTHOR

Nicholas Caputo is an assistant professor at the Johns Hopkins School of Government and Policy. He studies how law can govern artificial intelligence, and how AI can in turn renew and improve our legal and administrative institutions.



From *Futurecraft: The Day After AGI*, Foundation for American Innovation, 2026.
Read or listen online at futurecraft.thefai.org/essays/capable-and-accountable-administration-in-the-time-of-agi.

© 2026 by the Foundation for American Innovation. All rights reserved.

13 Interstate Commerce Act of 1887, ch. 104, 24 Stat. 379; Administrative Procedure Act, Pub. L. No. 79-404, 60 Stat. 237 (1946); *Chevron U.S.A. Inc. v. NRDC*, 467 U.S. 837 (1984). On the political contingency of these settlements, see Elizabeth Sanders, *Roots of Reform: Farmers, Workers, and the American State, 1877–1917* (Chicago: University of Chicago Press, 1999); Stephen Skowronek, *The Politics Presidents Make: Leadership from John Adams to Bill Clinton* (Cambridge: Harvard University Press, 1997).