GOVERNING FOR IMPACT

AI in Agency Rulemaking: Legal Guardrails

July 2025

Issue Brief

INTRODUCTION

As artificial intelligence and machine learning systems grow in sophistication, they are likely to become increasingly integrated into many fields—including governance. As a tool for governance, AI would seem to have promise. If used responsibly, it could, among other things, analyze and synthesize data in service of effective policy, increase the technical capacity of regulators and legislators, aid policymakers in satisfying procedural requirements, and facilitate public participation. It also poses risks: It might be used to displace human discretion and decisionmaking, hollow out state capacity, decrease transparency, evade accountability, replicate and reinforce bias and arbitrariness in public policy, and compromise essential privacy and data governance safeguards. But for all the hype—both the utopian and apocalyptic varieties—federal agencies have been exploring AI tools and confronting these concerns for some time.[1] Now, the Trump administration appears set to increase the use of AI in governance, including by bringing AI's power to bear in service of its deregulatory agenda.[2]

This Issue Brief focuses on one aspect of AI's application to governance: the role it might play in the notice-and-comment rulemaking process at federal agencies. First, it lays out the functions AI might be capable of serving in rulemaking. Then it explores how commenters, litigants, and other members of the public might use existing doctrines of administrative law—in particular, rules of disclosure and reasoned decisionmaking—to impose guardrails on agencies' use of AI. Requiring agencies to disclose, explain, and justify their adoption of AI may be one way to steer them away from irresponsible and arbitrary applications of the technology without unduly hampering an agency's responsible use of AI. We emphasize, though, that our assessment of these issues is preliminary. To that end, this Issue Brief closes by identifying some shortcomings of the current doctrine and questions requiring further attention.

[1] *See, e.g.*, David Freeman Engstrom, Daniel E. Ho, Catherine M. Sharkey & Mariano-Florentino Cuéllar, Governance by Algorithm (Feb. 2020) (report to the Admin. Conf. of the U.S.).

[2] *See* Hannah Natanson, Jeff Stein, Dan Diamond & Rachel Siegel, *DOGE Builds AI Tool to Cut 50 Percent of Federal Regulations*, Wash. Post (July 26, 2025); Elon Musk & Vivek Ramaswamy, *The DOGE Plan to Reform Government*, Wall St. J. (Nov. 20, 2024).

CONTENTS

I. BACKGROUND 3

Artificial Intelligence: Basic Concepts 3

AI in Government 4

II. POSSIBLE USES OF AI IN THE RULEMAKING PROCESS 7

Identifying Subjects and Generating Ideas for Rulemaking 8

Determining the Substance of Rules 9

Explaining and Supporting Rules 10

III. HOW EXISTING ADMINISTRATIVE LAW DOCTRINES MIGHT APPLY TO USES OF AI IN RULEMAKING 11

Disclosure 12

Reasoned Decisionmaking 17

IV. DOCTRINAL SHORTCOMINGS AND OPEN QUESTIONS 20

V. CONCLUSION 22

I. BACKGROUND

Artificial Intelligence: Basic Concepts

Congress has defined the term "artificial intelligence" to include, among other things, "[a]ny artificial system that . . . can learn from experience and improve performance when exposed to data sets"; "[a]n artificial system designed to think or act like a human, including cognitive architectures and neural networks"; and "[a] set of techniques, including machine learning, that is designed to approximate a cognitive task."[3] This Issue Brief adopts this definition, which we understand to include a range of systems—including both machine learning tools (which can recognize patterns in and draw inferences from large sets of data) as well as generative AI (which can generate new data, such as text, in the case of so-called large language models).[4] Not every matter discussed below implicates all types of AI tools equally; the Issue Brief will aim to differentiate among them as necessary.

Given their speed, capacity, and the number of variables they can consider, AI systems can produce results and draw inferences probative of trends and correlations both perceptible and imperceptible to human analysts.[5] Trained on vast amounts of historical data, AI tools can

[3] John S. McCain National Defense Authorization Act for Fiscal Year 2019, Pub. L. 115-232, § 238(g), 132 Stat. 1363, 1697–98 (2018). The full statutory definition states:

> (g) Artificial Intelligence Defined.—In this section, the term 'artificial intelligence' includes the following:
>
> (1) Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets.
>
> (2) An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action.
>
> (3) An artificial system designed to think or act like a human, including cognitive architectures and neural networks.
>
> (4) A set of techniques, including machine learning, that is designed to approximate a cognitive task.
>
> (5) An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making, and acting.

Id. This definition of AI has been also incorporated by reference into more recent statutes. *See* Advancing American AI Act, Pub. L. 117-263, § 7223(3), 136 Stat. 3668, 3669 (2022) (codified at 40 U.S.C. § 11301 note).

[4] Stephen J. Bigelow, *Generative AI v. Machine Learning: Key Differences and Use Cases*, TechTarget (Dec. 30, 2024).

[5] *See* Cary Coglianese, A Framework for Governmental Use of Machine Learning 27–30 (Dec. 8, 2020) (report to the Admin. Conf. of the U.S.) [hereinafter Coglianese, Framework] .

pick stocks,[6] detect cancer,[7] forecast the weather,[8] and assess risk in insurance underwriting.[9]

Of course, it is well known that generative AI can also produce arbitrary or incorrect responses to prompts, a phenomenon anthropomorphically referred to as "hallucination."[10] The proliferation of AI-drafted legal briefs with invented citations is a well-reported example.[11] There are even indications that a federal agency has relied on false AI-generated information.[12] AI companies are reportedly moving to improve the reliability of their models,[13] although it is unclear the extent to which those measures have been effective.[14]

Unlike algorithms that simply execute preprogrammed rules, AI tools frequently produce results in ways that categorically cannot be explained.[15] That is, neither the systems themselves nor their designers can genuinely or reliably explain particular results. To the extent that a generative AI tool like a large language model chatbot can generate a plain text "explanation" of its work, that is not really an explanation at all; it is just another statistically likely response to the user's prompt.[16] For this reason, AI is sometimes described as a "black box."[17]

AI in Government

Congress and the executive branch have weighed in on, including to encourage, AI's use in the federal government. Among other things, the AI in Government Act of 2020[18] directed the White House to produce guidance for agencies to "inform the development of policies

[6] *An AI Analyst Made 30 Years of Stock Picks—and Outperformed Human Investors by a "Stunning" Degree*, Stan. Rep. (June 9, 2025).

[7] *Artificial Intelligence (AI) and Cancer*, Nat'l Cancer Inst. (May 30, 2024).

[8] Tim Fernholz, *The Future of Weather Prediction Is Here. Maybe.*, N.Y. Times (July 15, 2025).

[9] AI in Insurance Underwriting: A Complete Guide, Salesforce.

[10] *When AI Gets It Wrong: Addressing AI Hallucinations and Bias*, MIT Sloan Teaching & Learning Techs.

[11] *See* Sara Merken, *AI "Hallucinations" in Court Papers Spell Trouble for Lawyers*, Reuters (Feb. 18, 2025).

[12] Nia Prater, *Did RFK Jr.'s Crew Use AI to Write Error-Filled MAHA Report?*, N.Y. Mag.: Intelligencer (May 30, 2025); Sarah Owermohle, *FDA's Artificial Intelligence is Supposed to Revolutionize Drug Approvals. It's Making Up Studies.*, CNN (July 23, 2025).

[13] Megan Morrone, *Why AI Is Still Making Things Up*, Axios (June 4, 2025).

[14] Cade Metz & Karen Weise, *A.I. Is Getting More Powerful, but Its Hallucinations Are Getting Worse*, N.Y. Times (May 6, 2025).

[15] Boris Babic & I. Glenn Cohen, *The Algorithmic Explainability "Bait and Switch,"* 108 Minn. L. Rev. 857, 866–70 (2023).

[16] *See* Matthew Kosinski, *What is Black Box Artificial Intelligence (AI)?*, IBM (Oct. 29, 2024).

[17] *Id.*

[18] Pub. L. No. 116-260, 134 Stat. 2286 (codified at 40 U.S.C. § 11301 note).

regarding" AI usage; "recommend approaches to remove barriers for use" of AI "while protecting civil liberties, civil rights, and economic and national security"; and "identify best practices for identifying, assessing, and mitigating . . . unintended consequences" of AI usage, like "discriminatory impact."[19] The statute tasked agencies with producing plans to conform to that guidance.[20]

The Advancing American AI Act of 2022[21] renewed the mandate for the executive branch to develop AI guidance and policies for agencies on multiple topics, including procurement, data security, privacy, and civil rights and civil liberties.[22] It also directed each agency to "prepare and maintain an inventory of the artificial intelligence use cases of the agency" to be made "available to the public" by the Office of Management and Budget (OMB).[23]

Consistent with Congress's directive, recent presidential administrations have issued executive orders and memoranda on AI policy. Among other efforts, the Biden administration issued two executive orders focused on AI,[24] a "Blueprint for an AI Bill of Rights,"[25] and OMB memos to implement the statutes referenced above.[26]

For its part, the Trump administration has instructed agencies to "adopt a forward-leaning and pro-innovation approach that takes advantage of [AI] technology to help shape the future of government operations."[27] Agencies are to "lean forward on adopting effective, mission-enabling AI" and "establish clear expectations for their workforce on appropriate AI use" while "ensuring that rapid AI innovation is not achieved at the expense of the American people

[19] § 104(a)(1)–(3), 134 Stat. at 2287.

[20] § 104(c), 134 Stat. at 2289.

[21] Pub. L. 117-263, 136 Stat. 3668 (2022) (codified at 40 U.S.C. § 11301 note).

[22] § 7224(a), 136 Stat. at 3669 (calling for general guidance); § 7224(b), 136 Stat. at 3670 (mandating issuance of procurement policies and procedures); § 7224(d), 136 Stat. at 3671 (requiring development of "an initial means by which to" protect, among other things, data security, "privacy, civil rights, and civil liberties").

[23] § 7225(a)–(b), 136 Stat. at 3671–72. The central inventory of AI use cases is available at *2024 Federal Agency AI Use Case Inventory*, GitHub (Jan. 23, 2025).

[24] Executive Order 14141 of January 14, 2025 (Advancing United States Leadership in Artificial Intelligence Infrastructure), 90 Fed. Reg. 5469, *revoked by* Executive Order 14318 of July 23, 2025, 90 Fed. Reg. 35,385; Executive Order 14110 of October 30, 2023 (Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence), 88 Fed. Reg. 75,191, *revoked by* Executive Order 14148 of January 20, 2025, 90 Fed. Reg. 8237.

[25] Office of Science and Technology Policy, Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People (2022).

[26] *See, e.g.*, OMB, M-24-18, Advancing the Responsible Acquisition of Artificial Intelligence in Government (Sept. 24, 2024), *revoked by* OMB, M-25-22, Driving Efficient Acquisition of Artificial Intelligence in Government (Apr. 3, 2025); OMB, M-24-10, Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence (March 28, 2024), *revoked by* OMB, M-25-21, Accelerating Federal Use of AI through Innovation, Governance, and Public Trust (Apr. 3, 2025).

[27] OMB, M-25-21, *supra* note 26.

or any violations of their trust."[28] Most recently, the Trump administration has released "America's AI Action Plan," which asserts that the government can use AI to "accelerat[e] slow and often manual internal processes" and "streamlin[e] public interactions."[29] It calls for government-wide initiatives to, among other things, train employees on AI usage and share technology.[30]

Trump administration actions have indicated that the federal government's use of AI in administrative decisionmaking will increase—perhaps quickly.[31] This Issue Brief discusses how administrative law doctrines might be applied or adapted to the use of AI tools in one particular part of administrative governance: the rulemaking process.

There are several important topics at the intersection of AI and administrative decisionmaking that this Issue Brief does *not* address. This Issue Brief does not:

- identify principles or best practices for how agencies *ought* to use AI, or the pros and cons of AI more generally;[32]
- catalog how agencies are currently using AI—scholars have undertaken this work in the past,[33] and, as described above, agencies are required by statute to publish AI "use case" "inventor[ies],"[34] which OMB compiles;[35]
- discuss data governance and privacy concerns;
- propose ways for Congress, courts, or agencies themselves to adjust administrative law or practice to account for AI (although the Issue Brief does identify certain gaps in existing doctrines relevant to AI usage and suggest ways the law could be applied or modestly extended in light of AI);
- opine on how agencies ought to procure, design, or fine tune AI products;
- discuss how AI might be used in other domains of agency work, like enforcement, adjudication, and internal management; or

[28] *Id.* at 2–3.

[29] Exec. Off. of the President, America's AI Action Plan 10 (July 2025).

[30] *Id.* at 11.

[31] *See, e.g.*, Natanson *et al.*, *supra* note 2; Sareen Habeshian, *Musk's xAI Announces $200 Million Contract with Pentagon*, Axios (July 14, 2025).

[32] Literature abounds on the broader topic. *See, e.g.*, John Cassidy, *How to Survive the A.I. Revolution*, New Yorker (Apr. 14, 2025).

[33] *See* Engstrom et al., *supra* note 1.

[34] Advancing American AI Act, Pub. L. 117-263, § 7225(a), 136 Stat. 3668, 3672 (2022).

[35] *See 2024 Federal Agency AI Use Case Inventory*, GitHub (Jan. 23, 2025).

- address how Congress or agencies should regulate AI development or use by private entities or individuals—Governing for Impact and the Center for American Progress have produced a report exploring aspects of that topic.[36]

II. POSSIBLE USES OF AI IN THE RULEMAKING PROCESS

Agencies could conceivably use AI at various stages of the rulemaking process, which typically (but not always) involves the following steps: To start, the agency decides to act on a certain subject matter based on, for example, a rulemaking petition or other information from the public, the agency's own retrospective review of its own regulations, an analysis of scientific literature or public reporting, or a political directive. The agency then develops a proposed rule, which includes both proposed regulatory text as well as an initial explanation of the agency's basis for regulating. That proposal will often reflect consultation with other executive branch offices as well as internal executive branch policies like the requirements to analyze benefits and costs. The agency publishes a "notice of proposed rulemaking"—which describes and explains the agency's proposal, supplies the critical factual material on which the agency relied, and invites public comment—in the *Federal Register.*

The agency then reviews the public's comments and develops a final rule. That generally entails drafting final regulatory language and a preamble offering a legal and policy justification for the agency's action, along with responses to comments.[37] As at the proposed rule stage, the agency will frequently submit its draft final rule for interagency review coordinated by OMB. The final rule is then published in the *Federal Register* and generally takes effect no fewer than thirty days thereafter.

As that description should make plain, the rulemaking process can often be lengthy and time-consuming, with many rules taking years to finalize.[38] Administrations of both parties have noted their frustration with the rulemaking process and how it can serve to slow down their ability to effectuate their policy priorities. Therein lies the promise of artificial intelligence: it may provide additional capacity to agencies as they regulate, including to speed their

[36] *See* Will Dobbs-Allsopp et al., Taking Further Agency Action on AI (June 17, 2024).

[37] For a description of the rulemaking process, see Off. of the Fed. Reg., A Guide to the Rulemaking Process (Sept. 2013).

[38] Richard L. Revesz, *New Challenges for Federal Regulations: Executive Branch Responses*, 100 N.Y.U. L. Rev. (forthcoming 2025) (manuscript at 22) (observing that "it has historically taken two years of work to publish the proposed rule and a year-and-a-half of additional work to publish a final rule").

development and finalization of regulatory proposals. But AI may also introduce error, bias, or arbitrariness to administrative decisionmaking.

AI may be capable of facilitating the following steps:

Identifying Subjects and Generating Ideas for Rulemaking

AI tools could be used to perform several tasks that might aid agencies in determining where and how to regulate. An agency might use an AI system to conduct a "retrospective review" of its current regulations in order to identify obsolete provisions, scrivener's errors of various sorts, and areas that could benefit from additional elaboration.[39] In a similar vein, a team from Stanford's Regulation, Evaluation, and Governance Lab has used an AI program to identify every provision of San Francisco's city code requiring a municipal department to create a report, enabling city officials to decide "which reports could be tweaked for efficiency, combined with similar requirements, or slashed altogether."[40]

Agencies might also use AI to digest and synthesize voluminous data and identify trends and patterns relevant to policymaking. That could mean analyzing, among other things, environmental, economic, meteorological, or pharmaceutical data.[41] Or it could mean using a large language model to perform literature reviews, identifying important findings in academic and public policy research.[42]

AI tools might be used to evaluate the substance of existing regulations, including to assess their legality. The U.S. DOGE Service is reportedly planning to use an "AI Deregulation Decision Tool" to identify more than 100,000 regulations—or "50%" of all federal regulations—that it believes can be rescinded.[43] According to a DOGE presentation, this tool has powered "100% of [d]eregulations [w]ritten" at the Consumer Financial Protection Bureau and "[d]ecisions on 1,083 [r]egulatory [s]ections in [t]wo [w]eeks" at the Department of

[39] *See* Catherine M. Sharkey, Algorithmic Tools in Retrospective Review of Agency Rules (May 3, 2023) (report to the Admin. Conf. of the U.S.).

[40] Chase DiFeliciantonio, *San Francisco Wants to Use AI to Save Itself from Bureaucracy*, Politico (June 5, 2025); Nicholas Bagley, *Using Artificial Intelligence to Build State Capacity*, Divided Argument (June 9, 2025); *see also* Faiz Surani et al., What is the Law? A System for Statutory Research (STARA) with Large Language Models (2025).

[41] *See* Coglianese, Framework, *supra* note 5, at 25–26.

[42] *Cf.* Helen Pearson, *Can AI Review the Scientific Literature—and Figure Out What It All Means?*, Nature (Nov. 13, 2024).

[43] Dep't of Gov't Efficiency, DOGE Deregulation Opportunity 2 (July 1, 2025).

Housing and Urban Development.[44] Notably, when used by that agency, DOGE's AI system reportedly "made several errors."[45]

Finally, agencies are required by law to accept petitions for rulemaking from the public.[46] They also no doubt receive less formal correspondence and communications from stakeholders advocating that they take certain actions. Generative AI might synthesize petitions and other communications and flag relevant trends for agency personnel, such as recurring requests or concerns—an application of AI that, as discussed below, could figure into other stages of the rulemaking process as well.

Determining the Substance of Rules

Having identified subjects for regulation, agencies might use AI tools of all sorts to develop specific regulatory proposals. For instance, an agency could use machine learning tools to identify trends and patterns in voluminous data suggestive of where standards should be set and what attributes of, say, a product or entity ought to be regulated.

A 2020 report commissioned by the Administrative Conference of the United States offers some examples of how this might work. In 2016, the Food and Drug Administration began using a natural language processing tool to process reports of drug adverse events and identify "previously undetected relationships" between certain adverse outcomes and particular drugs and drug combinations.[47] It is possible to imagine similar analysis informing standards underlying rulemaking or other regulatory activity. The report also recounts that the SEC used "a machine learning tool that helps identify which filers might be engaged in suspect earnings management."[48] "[T]rained on a historical dataset of past issuer filings," the tool "predict[s] possible misconduct using indicators such as earnings restatements and past enforcement actions."[49] While the most obvious application of this tool might be identifying targets for enforcement actions, it could also be used to inform prospective disclosure requirements. These methods, though less transparent and more powerful,[50] are similar to

[44] *Id.*; *see* Natanson *et al.*, *supra* note 2.

[45] *Id.*

[46] *See* 5 U.S.C. § 553(e) ("Each agency shall give an interested person the right to petition for the issuance, amendment, or repeal of a rule.").

[47] Engstrom et al., *supra* note 1, at 55–56.

[48] *Id.* at 23.

[49] *Id.*

[50] *See* Cary Coglianese & David Lehr, *Transparency and Algorithmic Governance*, 71 Admin. L. Rev. 1, 16 (2019).

the traditional mathematical models agencies have long used insofar as they use algorithms to sort through large datasets.[51]

Large language models might be used in less sophisticated ways to generate policies. This might entail agency staff prompting chatbots to generate ideas for the substance of regulations. Some commentators observed that certain commercially available AI chatbots, upon appropriate prompting, proposed tariff rates strikingly similar to those announced by President Trump in early April 2025.[52]

Explaining and Supporting Rules

Generative AI, which is increasingly capable of producing high-verisimilitude natural language, including in specialized legal contexts,[53] might be used for several rulemaking tasks required by the Administrative Procedure Act and other federal laws.

For example, under the APA, agencies' final rules must include "a concise general statement of their basis and purpose."[54] These preambles, though, are not generally concise. They frequently must address numerous legal and policy issues in order for the underlying rule to survive judicial review.[55] As a result, they are often extremely resource-intensive to produce. Agency personnel might use large language models to draft rule preambles and prepare analyses addressing particular legal or policy concerns implicated by a rule. Generative AI could be put to similar work with respect to supporting analyses agencies often must produce. For instance, for certain actions, an agency may be required to produce one or more analyses under the National Environmental Policy Act.[56] And executive branch policy generally requires agencies promulgating significant regulatory actions to assess their benefits and costs.[57] Large language models could also be prompted to translate dense rule preambles into plain English.

Along similar lines, agencies must respond to important comments submitted in response to proposed rules.[58] That task can be daunting as well. Some rulemakings garner hundreds of

[51] *But see* Note, *Machine Rulemaking: Arbitrary and Capricious Review in the Age of AI*, 138 Harv. L. Rev. 1821, 1831–34 (2025) (distinguishing traditional and machine learning models using the FDA case study).

[52] *See* Dominic Preston, *Trump's New Tariff Math Looks a Lot Like ChatGPT's*, The Verge (Apr. 3, 2025, 11:06 AM).

[53] *See* Adam Unikowsky, *In AI We Trust, Part II*, Adam's Legal Newsl. (June 16, 2024).

[54] 5 U.S.C. § 553(c).

[55] *See generally Motor Vehs. Mfrs. Ass'n v. State Farm Mut. Auto. Ins. Co.*, 463 U.S. 29, 43 (1983); Governing for Impact, Arbitrary-and-Capricious Challenges (May 2025).

[56] *See* 42 U.S.C. § 4336.

[57] *See* Exec. Order No. 12,866, 58 Fed. Reg. 51735, § 6(a)(3)(B)–(D); *see also* Off. of Info. and Regul. Affs., Regulatory Impact Analysis: A Primer 2.

[58] *Perez v. Mortgage Bankers Ass'n*, 575 U.S. 92, 96 (2015).

thousands—even, in rare cases, millions—of comments (the drafting of which can now, of course, be assisted by large language models).[59] Currently, agencies might use software to, among other things, identify form and duplicate comments.[60] Generative AI tools capable of natural language processing might be able to go significantly further: digesting the comment record, provisionally flagging which comments might be significant enough to require a response under governing law, topically grouping comments and summarizing the points they make, and even drafting possible responses. Generative AI could similarly summarize comments agencies receive from other agencies and the White House when a rule undergoes interagency review.

Here, too, DOGE has proposed the aggressive use of its AI tool. Its internal presentation contemplates using AI to "draft[] all submission documents" and "analyze[] 20 to 500,000 citizen comments for Final Rule inclusion."[61] While reporting indicates the AI tool's outputs would be subject to "some staff feedback," DOGE's proposal is to "automate[]" large portions of the regulatory process.[62]

III. HOW EXISTING ADMINISTRATIVE LAW DOCTRINES MIGHT APPLY TO USES OF AI IN RULEMAKING

With the use of AI in the rulemaking process almost certain to increase, members of the public—including commenters and litigants—might consider how existing doctrines of administrative law can be brought to bear. Two principles seem particularly important. First, disclosure: Under established administrative law doctrine, agencies are required to disclose the critical factual material on which they have relied in proposing and promulgating rules. That disclosure obligation may extend to whether, how, and to what extent an agency has relied on an AI to inform a rulemaking. Second, reasoned decisionmaking: The statutory requirement that agency action not be arbitrary and capricious has been understood to

[59] *See* Paul Hitlin, Kenneth Olmstead, and Skye Toor, *Public Comments to the Federal Communications Commission About Net Neutrality Contain Many Inaccuracies and Duplicates*, Pew Rsch. Ctr. (Nov. 29, 2017) (reporting that 21.7 million comments were submitted regarding net neutrality regulations proposed by the FCC).

[60] *See DocketScope's Public Comment Software*, DocketScope.

[61] DOGE Deregulation Opportunity, *supra* note 43, at 4.

[62] Natanson *et al.*, *supra* note 2.

require agencies to take actions that are reasonable and reasonably explained. Rules of reasoned decisionmaking may thus require agencies to justify their use of AI by explaining why they believe it has produced reliable results in a particular case. This section outlines how these requirements might apply to the use of AI in rulemaking.

Disclosure

Agencies are required to disclose information on which they have relied in proposing and promulgating rules. One disclosure requirement is imposed by 5 U.S.C. § 553, the section of the APA establishing procedures for notice-and-comment rulemaking. The statute requires agencies to publish "notice" of a proposed rule's "terms or substance . . . or a description of the subjects and issues involved,"[63] and provide "an opportunity to participate in the rule making."[64] To allow the public to meaningfully participate in rulemakings where an agency has used sophisticated methodologies, Section 553 has long been interpreted to require agencies to disclose the data, systems, and assumptions underlying models sometimes used to develop rules. This is sometimes called the *Portland Cement* doctrine, after a case that articulated it.[65] As the D.C. Circuit aptly put it in another case more than forty years ago:

> The safety valves in the use of . . . sophisticated methodology are the requirement of public exposure of the assumptions and data incorporated into the analysis and the acceptance and consideration of public comment, the admission of uncertainties where they exist, and the insistence that ultimate responsibility for the policy decision remains with the agency rather than the computer.[66]

Agencies have also been required to "reveal[] for public evaluation . . . the technical studies and data upon which [they] rel[y]."[67] Disclosing the "critical factual material" that the agency "has employed in reaching the decisions to propose particular rules" is "especially important"

[63] 5 U.S.C. § 553(b)(3).

[64] *Id.* § 553(c).

[65] *See Portland Cement Ass'n v. Ruckelshaus*, 486 F.2d 375, 393 (D.C. Cir. 1973).

[66] *Sierra Club v. Costle*, 657 F.2d 298, 334 (D.C. Cir. 1981) (Wald, J.) (footnotes omitted).

[67] *Am. Radio Relay League, Inc. v. FCC*, 524 F.3d 227, 236 (D.C. Cir. 2008) (internal quotation marks and alterations omitted) (quoting *Chamber of Com. v. SEC*, 443 F.3d 890, 899 (D.C. Cir. 2006)). *But see id.* at 246 (Kavanaugh, J., concurring in part, concurring in the judgment in part, and dissenting in part) (arguing that "*Portland Cement* stands on a shaky legal foundation" because it "cannot be squared with the text of § 553 of the APA," implicating "the basic principle of *Vermont Yankee* that Congress and the agencies, but not the courts, have the power to decide on proper agency procedures" (internal quotation marks omitted)).

"to allow for useful criticism" by the public.[68] The agency must facilitate "[a] 'genuine interchange' about the accuracy of the data" on which it relied.[69]

The APA's bar on "arbitrary" or "capricious" decisionmaking also entails disclosure.[70] Because an agency must provide a reasoned explanation for why it chose a particular approach, "[w]hen an agency uses a computer model, it must explain the assumptions and methodology used in preparing the model."[71] As discussed below, reasoned decisionmaking may require an agency to offer a substantive explanation of its use of AI. But for present purposes, it also would appear to require a disclosure similar to that mandated by Section 553: An agency must be transparent about the methods, assumptions, and data it used to generate a rule.

Under these requirements, agencies might be obligated to disclose information about their actual or intended use of AI in rulemaking when they issue proposed and final rules. When an agency uses an AI tool to determine the substance of a regulation, information about that system and how it was used might well qualify as "critical factual material" subject to disclosure, as would the workings of a mathematical model used to inform a regulatory standard. But, as explored above, AI might also play roles in rulemaking better characterized as "supporting"—like producing literature reviews to inform regulatory action or drafting responses to public comments.

The further an agency gets from using AI to inform the substance of its action, the less clear *Portland Cement*'s application becomes. At one end of the spectrum, disclosure requirements apply most obviously when an agency uses an AI tool to determine the content of a regulation. The *Portland Cement* line of cases developed in response to agencies' growing reliance on "sophisticated methodology"—complex mathematical models—to inform regulatory decisions.[72] If agency models come to incorporate AI capabilities, or if AI tools otherwise supplement or supplant those models, that should implicate the same disclosure requirements. So, if the FDA uses AI to determine which drug combinations and attributes should be subject to regulation, the EPA uses AI to set an acceptable limit for a certain pollutant, or the SEC uses AI to identify additional requirements for filers, the disclosure case law should apply. In short, if a rule is justified based on a trend or correlation identified by a sophisticated AI tool, then the usage and workings of that tool are subject to disclosure.

---

[68] *Id.* (quoting *Chamber of Com..*, 433 F.3d at 900; *Conn. Light & Power Co. v. Nuclear Regul. Comm'n*, 673 F.2d 525, 530 (D.C. Cir. 1982)).

[69] *Window Covering Mfrs. Ass'n v. Consumer Prod. Safety Comm'n*, 82 F.4th 1273, 1283 (D.C. Cir. 2023) (quoting *Conn. Light & Power*, 673 F.2d at 530).

[70] 5 U.S.C. § 706(2)(A).

[71] *Owner-Operator Indep. Drivers Ass'n, Inc. v. Fed. Motor Carrier Safety Admin.*, 494 F.3d 188, 204 (D.C. Cir. 2007) (internal quotation marks omitted) (quoting *U.S. Air Tour Ass'n v. Fed. Aviation Admin.*, 298 F.3d 997, 1008 (D.C. Cir. 2002)).

[72] *Sierra Club v. Costle*, 657 F.2d 298, 334 (D.C. Cir. 1981) (Wald, J.).

Things might become more complicated if an agency uses AI substantively in developing a rule but justifies its action on grounds independent of the AI tool. In that event, the government might argue in litigation that the characteristics and uses of the AI product do not qualify as "critical factual material" subject to disclosure. Taking that approach might result in other problems for the agency. For one thing, an agency disclaiming AI's role in its action would be prohibited from subsequently relying on an AI product to justify or explain its actions.[73] That might put additional pressure on the agency's stated bases for acting. For instance, if an agency relied on an AI product's output behind the scenes to aid in the formulation of a rule, its stated justifications may prove insufficient to justify the chosen outcome. Along these same lines, if an agency has not been truthful about the basis for its action, it may have opened itself to an argument that its stated rationale is impermissibly pretextual or otherwise arbitrary and capricious.

Further along the continuum of potential AI uses, an agency might use an AI tool to support its decisionmaking. To the extent a particular usage of AI helps the agency come to a decision—for instance, if an agency uses a large language model to produce the equivalent of a "staff report"—*Portland Cement* might well come into play. In that situation, disclosure of critical factual information would likely include the AI tool's output.

The requirement that critical factual material be disclosed, as presently understood, might be less likely to encompass the use of AI to help articulate and justify decisions already made. For instance, generative AI might "draft" legal justifications for rules or responses to adverse comments. We are unaware of cases mandating that agencies disclose the technologies they use in discharging their reasoned decisionmaking obligations, such as software used to identify duplicative public comments. Moreover, agencies may be able to avoid disclosing material that is "predecisional" and "deliberative," that is, "actually related to the process by which policies are formulated," on the grounds that it is privileged.[74] That "deliberative process privilege" generally protects agencies from having to disclose tentative or provisional justifications drafted by subordinate officials, and so may, by analogy, shield draft reasoning generated by AI for review by senior officials.

Nevertheless, litigants might explore arguments that disclosure rules should apply—or be modestly extended, if necessary—to the use of AI to justify rules and even to the outputs an

[73] *Dep't of Com. v. New York*, 588 U.S. 752, 780 (2019) ("[I]n reviewing agency action, a court is ordinarily limited to evaluating the agency's contemporaneous explanation in light of the administrative record."); *Dep't of Homeland Sec. v. Regents of the Univ. of Cal.*, 591 U.S. 1, 21 (2020) (an agency may not rely on a "*post hoc* rationalization" for its action (internal quotation marks omitted)).

[74] *ACLU v. NSA*, 925 F.3d 576, 592 (2d Cir. 2019) (internal quotation marks and alterations omitted); *see also Oceana, Inc. v. Ross*, 920 F.3d 855, 865 (D.C. Cir. 2019) (recognizing in the context of APA review of rulemaking that "predecisional and deliberative documents are not part of the administrative record" (internal quotation marks omitted)). *But see Ksanka Kupaqa Xa'lcin v. United States Fish & Wildlife Serv.*, 2020 WL 4193110, at *1 & nn.1, 2 (D. Mont. Mar. 9, 2020) (surveying cases, including that reach the opposite conclusion).

AI tool has generated. A rule's justification, as distinct from the content of the rule itself, has increasingly become a locus of judicial review.[75] And AI-generated "drafts"—which might well be infected with arbitrariness in the forms of, say, impermissible bias, factual errors, or faulty logic—risk anchoring the agency to erroneous reasoning.[76] As Professor Bridget Dooling put it, "first drafts are sticky" because, in Judge Posner's words, a reviewer may "glide over the smoothly written product of an able assistant without searching out the stresses and weak joints."[77] *Portland Cement*, which broadly requires that the information shaping an agency's decision be opened to public scrutiny, might be understood to cover draft reasoning produced by AI.[78]

And even assuming that the APA would not require disclosure of draft justifications for a rule if generated by a human, an agency's interest in nondisclosure is likely less weighty with respect to AI-generated reasons. The deliberative process privilege is intended to foster frank and candid consideration among agency personnel.[79] Yet even if a large language model's work can qualify as "deliberative" at all, an AI would presumably not be deterred or intimidated by the prospect of public disclosure, nor would the possibility of AI usage being disclosed necessarily lead agency personnel to self-censor. At the very least, the *fact* that an AI tool was used to justify a regulation and information about the product's procurement, design, and validation are arguably closer to "factual information," which generally is not privileged,[80] than to "part of the agency give-and-take . . . by which the decision itself is made," which generally is protected.[81] But, to be clear, we are unaware of any cases that have addressed the applicability of the deliberative process privilege to AI-generated materials.

As an initial step toward disclosure, those submitting public comments on proposed rules might simply ask agencies whether and to what extent they have used AI as part of the rulemaking. Such requests might be grounded in the *Portland Cement* doctrine and also raise policy reasons favoring disclosure. For instance, in apparent recognition of the public's strong interest in transparency around government uses of AI, Congress has required agencies to

[75] *See, e.g.*, *Ohio v. EPA*, 603 U.S. 279, 293 (2024).

[76] *Cf.* Amos Tversky & Daniel Kahneman, *Judgment Under Uncertainty: Heuristics and Biases*, 185 Science 1124, 1128 (1974) (describing the cognitive bias of anchoring).

[77] Bridget C.E. Dooling, *Ghostwriting the Government*, 109 Marq. L. Rev. (forthcoming 2026) (manuscript at 26–27) (quoting Richard A. Posner, The Federal Courts 111 (1985)).

[78] Note that this reasoning would not necessarily require agencies to disclose more ministerial applications of technology, like word processors or de-duping software.

[79] *Rudometkin v. United States*, 140 F.4th 480, 487 (D.C. Cir. 2025) ("[T]he goal of the deliberative-process privilege is to ensure that debate and candid consideration of alternatives within an agency are not subject to public inspection." (internal quotation marks and alterations omitted)).

[80] *Loving v. Dep't of Defense*, 550 F.3d 32, 38 (D.C. Cir. 2008).

[81] *Abtew v. DHS*, 808 F.3d 895, 898–99 (D.C. Cir. 2015) (internal quotation marks omitted).

publish AI "use case" "inventor[ies]."[82] The Administrative Conference of the United States has observed that "[a]gencies' efforts to ensure transparency in connection with their AI systems can serve many valuable goals," and it therefore recommends that "agencies might prioritize transparency in the service of legitimizing its AI systems, facilitating internal or external review of its AI-based decision making, or coordinating its AI-based activities."[83] And for its part, the White House recently emphasized that agencies, in using AI, must "provide improved services to the public, while maintaining strong safeguards for civil rights, civil liberties, and privacy."[84] Given all this, commenters might ask agencies:

- Whether AI was used in the rulemaking and, if so, how;
- What product the agency used and how it was selected;
- How the product was procured;
- Whether it was fine-tuned for a particular agency use;
- What categories of data it was trained on;
- What prompts or inputs the agency used to elicit responses or outputs;
- What responses the AI product produced;
- How agency staff used those outputs;
- What quality control and validation agency staff performed on the outputs;
- What measures the agency took to ensure that its usage of AI complied with applicable data security and privacy requirements;
- Whether and to what extent persons and entities not employed by the agency developed, modified, provided access to, or used AI as part of a rulemaking.[85]

[82] Advancing American AI Act, Pub. L. 117-263, § 7225(a), 136 Stat. 3668, 3672 (2022). This Issue Brief does not address the open question whether the statutory requirement that agencies publish AI uses cases obligates agencies to disclose AI usage in particular proceedings, or whether any such requirement would be judicially enforceable.

[83] Admin. Conf. of the U.S., Statement #20, Agency Use of Artificial Intelligence, 86 Fed. Reg. 6612, 6616 (Jan. 22, 2021).

[84] OMB, M-25-21, *supra* note 26.

[85] Governing for Impact has submitted a comment requesting this information with respect to the Office of Personnel Management's proposal to establish a "Schedule Policy/Career" in the civil service. *See* Governing for Impact, Comment on Improving Performance, Accountability, and Responsiveness in the Civil Service (May 22, 2025).

Commenters should also ask agencies, if they have used AI but did not disclose that usage in the notice of proposed rulemaking, to provide the public with an additional opportunity to comment, relying on the *Portland Cement* doctrine.

Agencies might not reply satisfactorily to these questions. Among other things, agency staff may not have full visibility into all the ways that AI touches their work. For instance, numerous software systems incorporate machine learning capabilities and agencies may not understand the code behind the software they use as part of rulemaking. But agencies are under a legal obligation to provide reasoned responses to important comments.[86] To the extent an agency offers a deficient or false response to commenters' requests, that might provide a basis for litigation.

Disclosure serves at least two useful purposes. First, in a moment of transition in both technology and governance, disclosure can provide interested parties and the broader public with valuable factual information about how agencies are using AI. That transparency can spur efforts to shape and, where necessary, reform AI's role in administrative governance. Second, requiring disclosure may impose some discipline on agencies. Agencies compelled to disclose how and when they use AI may tend toward more prudent uses of the technology than agencies whose deployment of AI is not subject to public scrutiny. These considerations no doubt spurred Congress's repeated moves to require the executive branch to disclose both its general policies regarding AI usage and specific use cases.[87]

Importantly, disclosure does not generally impose any express limit on an agency's discretion or authority, including with respect to AI (though, as noted above, there are some questions about the inner workings of AI systems that cannot be answered; requiring disclosure as to these questions could thus have broader implications). Outright limitations on agencies' use of AI might have benefits and downsides, but they are beyond the scope of this Issue Brief.

Reasoned Decisionmaking

The APA instructs courts to "hold unlawful and set aside agency action . . . found to be" "arbitrary, capricious, an abuse of discretion, or otherwise not in accordance with law."[88] Courts have interpreted that to mean that agency action must be "reasonable and reasonably explained."[89] Among other things, agency action is invalid under the APA if it is based on

[86] *See Ohio v. EPA*, 603 U.S. 279, 293 (2024).

[87] This Issue Brief reserves the question whether agencies could be required to disclose information about AI usage in particular matters under the Freedom of Information Act, 5 U.S.C. § 552.

[88] 5 U.S.C. § 706(2)(A).

[89] *FCC v. Prometheus Radio Proj.*, 592 U.S. 414, 423 (2021).

factual errors, insufficient evidence, internal contradictions, or otherwise faulty reasoning.[90] It could violate the requirements of reasoned decisionmaking for an agency to base its policy decision or its supporting reasons on a system prone to making errors, reproducing biases, or otherwise acting unpredictably and unreliably. Therefore — and, again, we emphasize that our thinking in this regard is preliminary — litigants might argue that when an agency uses an AI tool to determine the content of a rule, the APA requires that the agency reasonably explain, among other things:

- the purposes for and means by which the AI product used was designed and fine-tuned;
- how the product was prompted; and
- whether and how agency staff oversaw, validated, and performed quality control on the product's responses to ensure their reliability and compliance with statutory and constitutional requirements.[91]

This manner of explanation, of course, dovetails with — and builds upon — the disclosures described above. To the extent the agency provides deficient answers to these questions, its decision may be challengeable as well. An agency that announced that it had simply asked ChatGPT how it should deal with a topic would be providing something of an explanation, but certainly not a well-reasoned one.[92]

Courts have long required such explanations with respect to agencies' usage of mathematical models to set regulatory standards. In that context, "the agency must 'explain the assumptions and methodology used in preparing' any model used and 'provide a full analytical defense' of any challenged aspects."[93] Courts "will also look for evidence that the agency is conscious of the limits of the model."[94] Moreover, an agency "retains a duty to

[90] *See* Governing for Impact, Arbitrary-and-Capricious Challenges 12–15 (May 2025).

[91] *See* Coglianese & Lehr, *supra* note 50, at 43 ("What will matter to the courts is that the agency has sufficiently justified its design of and reliance on a particular algorithmic tool. The agency will need to reveal and justify its choice of an outcome variable and objective function. As the selection of an objective function and design of an algorithm necessarily entail making tradeoffs that call for policy judgment, an agency will need to explain its choices about these tradeoffs in terms of factors that are consistent with the agency's statutory authority, and it will need to respond to meaningful public comments submitted during a rulemaking. Agencies will also need to validate that the algorithm performs as intended and that it achieves the justified objectives." (footnotes omitted)).

[92] There may well be other legal rules implicated by an agency's reliance on AI in rulemaking. For instance, the use of a model shown to produce biased results could raise questions under the Due Process and Equal Protection Clauses, as well as other sources of nondiscrimination law.

[93] *Powder River Basin Res. Ctr. v. Dep't of Interior*, 749 F. Supp. 3d 151, 164 (D.D.C. 2024) (citations omitted) (quoting *Owner-Operator Indep. Drivers Ass'n, Inc. v. Fed. Motor Carrier Safety Admin.*, 494 F.3d 188, 204 (D.C. Cir. 2007); *Columbia Falls Aluminum Co. v. EPA*, 139 F.3d 914, 923 (D.C. Cir. 1998)).

[94] *Small Refiner Lead Phase-Down Task Force v. EPA*, 705 F.2d 506, 535 (D.C. Cir. 1983).

examine key assumptions as part of its affirmative 'burden of promulgating and explaining a non-arbitrary, noncapricious rule.'"[95]

These precedents would seem to apply squarely when an agency relies on any form of AI in setting rules, either as part of or in place of a traditional model. It has been suggested that courts, which may already struggle to comprehend and review complex mathematical models, may not be able to assure themselves of the reliability of complicated and opaque AI tools.[96] That remains to be seen. For the moment, though, that possibility presents litigators with an opportunity. It is an agency's burden to reasonably explain its analysis.[97] If it cannot, that increases the likelihood that an arbitrary and capricious challenge will be successful.

As in the disclosure context, an agency may argue that it is not under this same obligation to explain if it used AI in formulating a rule but did not rely on the AI's output as a basis for its action. As discussed above, that move would increase the analytical pressure on the agency's stated reasons.

It is less clear that the reasoned decisionmaking caselaw would require agencies to explain the use of AI for "supporting" tasks, like producing a response to a particular comment. But even if it does not, arbitrary-and-capricious review should still offer an important guardrail against erroneous or faulty AI-generated reasons, just as it does with respect to human decisionmaking. An unsound reason, an unsupported assertion, or a factual error may well be apparent to commenters and litigators whether it was initially produced by human agency staff or a large language model. Thus, even if agencies are able to successfully argue that they need not explain certain uses of AI (and it is far from clear that they will be), their actions will still be subject to legal scrutiny.

Importantly, while arbitrary-and-capricious review is frequently described as deferential, it frequently has real teeth in practice.[98] Litigants can point to a wide range of agency errors to argue for a rule's invalidation. But the Supreme Court has instructed that judicial review of an agency's "predictions, within its area of special expertise, at the frontiers of science" is "at its most deferential."[99] An agency might argue that because of AI's cutting-edge nature, this extra-deferential standard ought to apply to a court's review of its use of AI. Litigants might respond by emphasizing that frontiers-of-science deference applies only when an agency is operating "within its area of special expertise."[100] Agencies are at a comparative advantage,

[95] *Id.* at 534 (quoting *Nat'l Lime Ass'n v. EPA*, 627 F.2d 416, 433 (D.C. Cir. 1980).

[96] Note, *Machine Rulemaking*, *supra* note 51, at 1831–32.

[97] *Small Refiner Lead*, 705 F.2d at 535.

[98] Governing for Impact, Arbitrary-and-Capricious Challenges 1 (May 2025).

[99] *Balt. Gas & Elec. Co. v. Nat. Res. Def. Council, Inc.*, 462 U.S. 87, 103 (1983).

[100] *Id.*

in other words, when the substance of their predictive judgments pushes the envelope, or, at least, courts defer when there is genuine uncertainty about a technical question. But it does not follow that an agency is owed special deference when it uses a technologically advanced tool to make a more ordinary decision. What matters is not "the complexity of the *tools* used, but the complexity of the *question* posed to the agency."[101]

We close by emphasizing that the explanations suggested in this section relate to how agencies deploy, design, prompt, and oversee AI systems.[102] In light of the black box nature of many AI tools, a requirement that an agency explain how an AI system *actually* reached a particular result may prove impossible to satisfy and could severely limit whether agencies can use AI in rulemaking at all.

IV. DOCTRINAL SHORTCOMINGS AND OPEN QUESTIONS

The statutes and doctrines shaping administrative law are generally geared toward values like transparency, participation, and accountability. If the rules governing the use of AI in the rulemaking process are to serve these same purposes, existing administrative law doctrine may be an imperfect fit. Over the last several decades, courts have applied the APA's disclosure and explanation requirements to analogous issues, like agencies' use of mathematical models in rulemaking; to the extent agencies use AI products in a similar fashion to those models, the existing case law might prove useful to litigants. But, as this Issue Brief has suggested, if agencies use generative AI to perform internal functions like drafting justifications for rules and responding to comments—practices that may be of immense importance and public interest—the law may need to evolve, either at the urging of litigants or the direction of Congress, if it is to ensure adequate transparency and oversight.

The use of AI to perform rulemaking tasks generally handled by agency staff also raises a number of other questions, some verging on the metaphysical. Questions administrative law may need to answer include:

[101] Note, *Machine Rulemaking*, *supra* note 51, at 1838; *see also id.* at 1836–39 (discussing the *Baltimore Gas* rule's application to agency use of machine learning).

[102] *See* Coglianese & Lehr, *supra* note 50, at 47–49.

- Must agency decisions ultimately be made by human decisionmakers?[103] If so, to what extent may AI aid in human decisionmaking by, say, generating initial drafts of rule preambles or comment responses?[104] And what manner and degree of human involvement is sufficient for the decision to be attributable to a human? Can an agency discharge its duty to respond to comments by delegating the task to an AI tool in the first instance? What manner of human ratification, review, or validation ought to be required?

- Should the output of a large language model be considered the legal equivalent of, say, a draft rule produced by subordinate personnel? Or are there material differences between the two types of output?[105]

- Should there be legal limits on how agency personnel working on a draft rule may prompt a large language model? May agency staff prompt a model to reach a particular substantive conclusion, by, for instance, instructing it to refute comments? Or does that contravene the requirement that agencies make decisions with an open mind?[106] Would that practice be legally different from human agency staff operating under similar explicit or implicit instructions? Can human review cure any error in this regard? Is it legally relevant that large language models often seem to pick up on cues in prompts and tell humans what they want to hear?[107]

- When a large language model drafts reasons for a rule, what is the focus of judicial review? The AI-generated reasons? When those reasons are adopted or ratified by a human, is it legally relevant that an AI product drafted them? Should human agency decisionmakers be required to produce their own reasoning as to why they have adopted or rejected the large language model's output?

- To the extent the AI tool's own reason-giving is subject to judicial review, should it be subject to review on the same terms as human decisionmaking? Or should the agency have to make some additional showing of cogency and reliability? If so, is a second-order explanation of how the model was designed, fine-tuned, trained, and prompted sufficient? Or is a first-order explanation of the AI's "thinking" required? If so, is it

[103] *Sierra Club v. Costle*, 657 F.2d 298, 334 (D.C. Cir. 1981) (Wald, J.) (describing "the insistence that ultimate responsibility for the policy decision remains with the agency rather than the computer" as a "safety valve[] in the use of . . . sophisticated methodology").

[104] *See* Dooling, *supra* note 77, at 12.

[105] *See id.* at 50–51.

[106] *Cf. Miss. Comm'n on Env't Quality v. EPA*, 790 F.3d 138, 183 (D.C. Cir. 2015) (explaining that "[a]n individual should be disqualified from rulemaking" "when they act with an unalterably closed mind and are unwilling or unable to rationally consider arguments" (internal quotation marks omitted)).

[107] *See* Luis Prada, *AI Chatbots Are Telling People What They Want to Hear. That's a Huge Problem.*, Vice (June 8, 2025); Stephanie Nguyen & Erie Meyer, *Tech Brief: AI Sycophancy & OpenAI*, Geo. Inst. for Tech. L. & Pol'y (2025).

enough for the large language model to generate plausible reasons for its own output?[108] Or must it actually give its "genuine" reasons,[109] which is frequently impossible,[110] and not actually required of human decisionmakers?[111] Does the black box nature of a large language model legally distinguish it from a human mind, which is often similarly incapable of recognizing or accounting for its own true reasons and frequently unreliable as a rational decisionmaking tool?[112]

The law currently does not seem to answer these questions, among many others. Nor does this Issue Brief.

V. CONCLUSION

AI will very likely come to play a role in the notice-and-comment rulemaking process, a prospect that presents both benefits and risks. The stakeholders in the administrative process—agencies, Congress, courts, and those who benefit from and are subject to regulation—must decide how administrative law will adapt. For the moment, litigants have at their disposal tools that ought to shed some light on how agencies are using AI in rulemaking and, ideally, push agencies toward prudent and responsible AI usage.

[108] *Cf. Dep't of Com. v. New York*, 588 U.S. 752 (2019) (an agency may act based on "unstated reasons" and "political considerations," *id.* at 781, but may not give "contrived" reasons, *id.* at 784).

[109] *Cf. SEC v. Chenery Corp.*, 318 U.S. 80, 94 (1943) ("[T]he orderly functioning of the process of review requires that the grounds upon which the administrative agency acted be clearly disclosed and adequately sustained.").

[110] *See* Babic & Cohen, *supra* note 15, at 864.

[111] *See* Coglianese & Lehr, *supra* note 50, at 43 ("[T]he arbitrary and capricious standard has never required a full explanation of the kind that psychologists, historians, or political scientists might demand if they wanted to understand exactly why government officials reached a decision.").

[112] *Cf.* Coglianese, Framework, *supra* note 5, at 10–21 (describing the physical limitations and cognitive biases that affect human decisionmaking).