

Potential Legal Challenges to AI Rubber- Stamping

November 2025
Issue Brief

The information in this document is provided for informational purposes only and does not contain legal advice, legal opinions, or any other form of advice regarding any specific facts or circumstances and does not create or constitute an attorney-client relationship. You should contact an attorney to obtain advice with respect to any particular legal matter and should not act upon any such information without seeking qualified legal counsel on your specific needs.

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	AGENCY USE OF AI IN RULEMAKING	4
III.	ADMINISTRATIVE LAW GUARDRAILS	8
	A. Duty to Give Reasons	10
	B. Duty to Engage in Reasoned Decisionmaking	12
	C. Duty to Address Significant Comments	18
	D. Duty to Keep an Open Mind	20
IV.	CONCLUSION	22

I. INTRODUCTION

The potential for artificial intelligence, machine learning systems, and large language models to transform administrative decisionmaking has attracted interest from politicians, policymakers, and academics alike. In particular, the use of AI to expedite and streamline federal rulemaking—a lengthy, complicated, and difficult process—has drawn significant attention. To that end, a presentation obtained by the *Washington Post* indicates that the Trump administration is considering using an AI tool developed by the U.S. DOGE Service to facilitate the rescission of 50 percent of all federal regulations by January 20, 2026, the first anniversary of President Trump’s second inauguration.¹ That proposal, if adopted, would impose substantial burdens upon advocates and litigators seeking to preserve vital regulatory protections.

In previous publications, Governing for Impact has offered suggestions for how advocates might require federal agencies to disclose and reasonably explain their use of artificial-intelligence tools,² including through comments on proposed regulatory actions.³ GFI has taken as an assumption in this work that the Administrative Procedure Act, the statute that governs agency rulemaking efforts, does not categorically preclude agencies from using sophisticated AI programs to carry out various tasks in the rulemaking process. AI is simply one of many tools an agency may use to perform its functions; it can be used in “better” (more reliable and participatory) or “worse” (more error-prone and exclusionary) ways. Indeed, a future administration might see immense value in the responsible use of AI to improve the quality of administrative decisionmaking and to expand state capacity to address the multitude of serious problems that the country faces.

¹ See Hannah Natanson et al., *DOGE Builds AI Tool to Cut 50 Percent of Federal Regulations*, Wash. Post (July 26, 2025), <https://www.washingtonpost.com/business/2025/07/26/doge-ai-tool-cut-regulations-trump/>.

² See *AI in Agency Rulemaking: Legal Guardrails*, Governing for Impact (July 2025), https://governingforimpact.org/wp-content/uploads/2025/07/AI-in-Agency-Rulemaking_Legal-Guardrails.pdf.

³ See Jordan Ascher, *Seeking Disclosure of AI Usage in Agency Rulemaking*, Yale J. on Reg.: Notice & Comment (July 19, 2025), <https://www.yalejreg.com/nc/seeking-disclosure-of-ai-usage-in-agency-rulemaking-by-jordan-ascher/>.

However, there are other potential uses of AI that arguably would violate the APA’s requirements—and undermine its broader goals of ensuring reasoned agency decisionmaking and meaningful public participation. In this Issue Brief, we identify one such circumstance: where an agency simply “rubber stamps”⁴ the output of an AI tool without engaging in any meaningful consideration of the issues posed by a given rulemaking or the comments submitted by the public. That hypothetical might involve several steps.

- An agency official might begin by asking an AI tool to identify regulations that purportedly are not required by statute, exceed the agency’s statutory authority, or are in conflict with the administration’s policy priorities (or, conversely, to propose new regulations that would advance the administration’s priorities).
- The official would then instruct the AI tool to prepare a notice of proposed rulemaking for publication in the Federal Register, including, as necessary, a description of the rule, an articulation of reasons for adopting it, its legal authority, draft regulatory text, and any other required regulatory reviews.
- Finally, the official would direct the AI tool to review the public comments on the proposed rule—which could number in the hundreds of thousands—and prepare a final rule that restates the agency’s reasons for action and why it chose to reject the issues raised by the comments.

At each of these steps, the official would not substantively engage with the AI’s output or the record before the agency; aside from perhaps a basic proofread, he or she would simply publish the AI’s output as-is—rubber-stamping it.

This description may not be hypothetical. The presentation obtained by the *Washington Post* sketches out a particularly extreme version of how the Trump administration may be planning to use AI—one in which DOGE’s AI tool could ultimately save “93% of Man Hours” by identifying regulations for rescission, “[a]utomatically draft[ing] all submission documents for attorneys to edit,” and

⁴ Cf. Adam Samaha, *Rubber Stamps*, 1 *Indep. L.J.* (forthcoming 2025), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5400421.

“[a]utomatically analyz[ing] 20 to 500,000 citizen comments for Final Rule inclusion.”⁵ According to the presentation, “[r]esearch,” proposed rule “[w]riting,” comment “[a]nalysis,” and final rule “[w]riting” will all be “[a]utomated.”⁶

On the other hand, the presentation states that agencies will “comment [on] and modify” the AI’s list of regulations recommended for rescission and “attorneys” will “edit” the AI’s “draft[] ... submission documents,” such that “[p]olicy and [l]egal [t]eams” will “[m]ake [a]ll [t]he [d]ecisions.”⁷ Whether that review would be meaningful remains to be seen; that would seem impossible if the administration proceeds at the extraordinary pace and scale proposed in the presentation, but if the administration’s deregulatory output is substantially slower, perhaps. It also remains unclear whether the Trump administration will ultimately follow through on the plan in the face of practical and legal concerns about the viability of the DOGE AI tool, among other obstacles.⁸

But the closer that the Trump administration’s use of AI is to the hypothetical sketched out above — the more it involves substantial amounts of rubber-stamping — the more likely it is to encounter significant legal obstacles. In this Issue Brief, we identify four potential administrative-law requirements that might prohibit agencies from simply rubber-stamping the output of an AI tool:

- Agencies must set forth their *own* reasons for acting, not those provided by a machine.
- Agencies must use AI in a reasoned and reasonable manner — that is, with safeguards to ensure reliability and a lack of impermissible bias.
- Agencies must meaningfully consider significant public comments, rather than prompting an AI tool to reject them out-of-hand.

⁵ *DOGE Deregulation Opportunity*, Dep’t of Gov’t Efficiency 3 (July 1, 2025), <https://www.washingtonpost.com/documents/857b6c65-0690-4b3c-b438-e3dc1dc87340.pdf>.

⁶ *Id.* at 8.

⁷ *Id.* at 3, 8; see also *id.* at 9 (referring to attorney “review”).

⁸ See, e.g., Andrew Coats, *A DOGE AI Tool Called SweetREX Is Coming to Slash US Government Regulation*, *Wired* (Aug. 14, 2025), <https://www.wired.com/story/sweetrex-deregulation-ai-us-government-regulation-doge/>; Ian Kullgren, *DOGE’s AI Plan to Kill 100,000 Rules Has a Bark Bigger Than Bite*, *Bloomberg Law* (Aug. 12, 2025), <https://news.bloomberglaw.com/daily-labor-report/doges-ai-plan-to-kill-100-000-rules-has-a-bark-bigger-than-bite>.

- Agencies must keep an open mind throughout the rulemaking process, rather than using an AI tool to rationalize a predetermined outcome.

To the extent an agency incorporates additional human oversight or review in its use of AI, such an approach may pose more complicated questions about whether the level of human involvement is sufficiently meaningful to pass legal muster.

To that end, we acknowledge that the rubber-stamping hypothetical this Issue Brief posits, and that the DOGE presentation could be read to embrace, might not realistically capture how Trump administration agencies end up using AI in the rulemaking process. Nevertheless, we believe the hypothetical is useful as a way to tease out the administrative law principles implicated by the use of AI in rulemaking. As such, we identify potential arguments here as minimum standards that might govern agency uses of AI. We invite readers to consider whether and to what extent these arguments stand up under different facts, such as when agencies temper their reliance on AI with meaningful human review and oversight. We leave that line-drawing question — which will likely be the focus of litigation, regulation, and perhaps legislation — for another day.⁹

II. AGENCY USE OF AI IN RULEMAKING

To contextualize our analysis, we provide a brief description of how agencies have begun to use AI, including the approach proposed by the DOGE AI presentation. For years, federal agencies have worked to incorporate AI and machine learning tools into their work.¹⁰ Congress has encouraged the responsible and transparent use of AI in governance, including by repeatedly directing the White House to issue guidance

⁹ In this piece, we do not address other potential legal constraints upon the use of AI, including other statutory limits (such as the Privacy Act or federal procurement statutes) or constitutional limits (such as equal protection or due process).

¹⁰ See, e.g., David Freeman Engstrom et al., *Government by Algorithm* (2020) (report to the Admin. Conf. of the U.S.), <https://law.stanford.edu/wp-content/uploads/2020/02/ACUS-AI-Report.pdf>.

for the executive branch on effective and prudent AI usage.¹¹ And in 2022, it specifically directed each agency to prepare a public “inventory of artificial intelligence use cases of the agency.”¹²

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 shape the future of government operations,”¹³ and it has released an “AI Action Plan” asserting that the government can use AI to “accelerat[e] slow and often manual internal processes” and “streamlin[e] public interactions.”¹⁴ It has been reported that the administration has been using AI, among other things, to evaluate federal workers’ response to the government-wide “Fork in the Road” email,¹⁵ “munch” agency contracts it viewed as nonessential,¹⁶ assist General Services Administration staff,¹⁷ and to “accelerat[e]” immigration enforcement.¹⁸

The DOGE presentation, however, sets forth a far more extreme approach toward adopting AI as a tool of administrative governance. As described, it promises that “50% of All Federal Regulations Can Be Eliminated.”¹⁹ The DOGE presentation claims that the “DOGE AI Program Saves 93% of Man Hours” by “[c]reat[ing] a regulation/statute database, and determin[ing] which regulations are required” by

¹¹ See AI in Government Act of 2020, Pub. L. No. 116-260, 134 Stat. 2286; Advancing American AI Act of 2022, Pub. L. No. 117-263, 136 Stat. 3668 (both codified at 40 U.S.C. § 11301 note).

¹² *Id.* § 7225(a)–(b), 136 Stat. at 3671–72.

¹³ Off. of Mgmt. & Budget, M-25-21, Accelerating Federal Use of AI through Innovation, Governance, and Public Trust (Apr. 3, 2025), <https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf>.

¹⁴ Exec. Off. of the President, America’s AI Action Plan 10 (July 2025), <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>.

¹⁵ Makena Kelly, *DOGE Used a Meta AI Model to Review Emails from Federal Workers*, Wired (May 22, 2025), <https://www.wired.com/story/doge-used-meta-ai-model-review-fork-emails-from-federal-workers/>.

¹⁶ Brandon Roberts et al., *DOGE Developed Error-Prone AI Tool to “Munch” Veterans Affairs Contracts*, ProPublica (Jun. 6, 2025), <https://www.propublica.org/article/trump-doge-veterans-affairs-ai-contracts-health-care>.

¹⁷ Gen. Servs. Admin., *GSA Launches Early Stage AI Tool for Internal Use, Seeks Staff Feedback* (Mar. 20, 2025), <https://www.gsa.gov/about-us/newsroom/news-releases/gsa-launches-early-stage-ai-tool-for-internal-use-seeks-staff-feedback-03202025>.

¹⁸ Angélica Franganillo Díaz, *‘Like Prime, but with Human Beings’: How the Trump Administration Is Using AI to Ramp Up Immigration Enforcement*, CNN (Sept. 22, 2025), <https://edition.cnn.com/2025/09/22/politics/artificial-intelligence-immigration-enforcement>.

¹⁹ *DOGE Deregulation Opportunity*, *supra* note 5, at 2.

statute, “[a]utomatically draft[ing] all submission documents for attorneys to edit,” and “[a]utomatically analyz[ing] 20 to 500,000 citizen comments for Final Rule inclusion.”²⁰ It cites two “Case Studies with DOGE AI Deregulation Decision Tool,” claiming that “100% of Deregulations [Were] Written Using the Tool” at the CFPB and that the “HUD Policy Team Completed Decisions on 1,083 Regulatory Sections in Two Weeks.”²¹ DOGE advertises that “[r]esearch” on what regulations can be rescinded, “[a]nalysis” of public comments, and “[w]riting” of “regulatory submission documents” and comment responses will all be “[a]utomated.”²²

At the same time, the presentation allows for some level of human involvement. In addition to noting that agency staff will be able to review and edit the AI’s output, as described above, it allocates some time toward “Policy + AI” and “Legal + AI” review.²³ Some of these figures are so low as to indicate a lack of any *meaningful* review; for instance, it allocates 18 minutes for staff review of “100,000+ comments.”²⁴ And yet the presentation suggests that, even without AI, it would only take 12.5 hours for staff to conduct such a review—a laughably low estimate to anyone with even a passing familiarity with the rulemaking process.²⁵ The extent to which agency staff are permitted to review the AI tool’s output may well vary by agency and depend to some degree on how well the AI is performing.

To achieve the goal of deleting half of federal regulations by “Relaunch America on Jan 20, 2026”²⁶—what DOGE describes as “an Audacious Goal that Requires Only the Will to Prioritize it”²⁷—DOGE proposed the following timeline:

²⁰ *Id.* at 3.

²¹ *Id.* at 4.

²² *Id.* at 8.

²³ *Id.*

²⁴ *Id.*

²⁵ *Id.*

²⁶ *Id.* at 4.

²⁷ *Id.* at 9.



Inexplicably, this timeline has agencies and their AI systems drafting final rules while the comment windows on the proposed rule are still open – responding to comments as they come in, seemingly – and does not provide time for OIRA review of draft final rules.

The current status of the DOGE plan is unclear. Aspects of it are hard to take seriously, and it purports to be just that – a plan. The White House commented that “no single plan has been approved or green-lit” in response to the *Washington Post*’s reporting.²⁸ If the DOGE AI tool is in fact being used as proposed, it is unknown how far it has progressed, if it is proceeding in line with DOGE’s proposed schedule, whether it involves meaningful human review of AI-generated content, or whether it has encountered agency resistance or legal or practical obstacles. For instance, while the Environmental Protection Agency has recently disclosed its intention to use AI to “help sort and categorize public comments,” it maintains that “the basis for the agency’s response remains the comments themselves and responses to comments as drafted by humans.”²⁹ Even the former general counsel of DOGE has allowed that “agencies still need human lawyers and human policy experts to do substantial work after capitalizing on all available AI tools.”³⁰ The government shutdown that lasted

²⁸ Natanson et al., *supra* note 1.

²⁹ *Commenting on EPA Dockets*, U.S. Env’tl Prot. Agency, <https://www.epa.gov/dockets/commenting-epa-dockets> (last visited Oct. 27, 2025).

³⁰ James Burnham, *State Farm and Making Deregulation Make Sense*, Yale J. on Reg.: Notice & Comment (Oct. 14, 2025), <https://www.yalejreg.com/nc/state-farm-and-making-deregulation-make-sense-by-james-burnham/>.

from October 1 to November 13 also likely slowed agency rulemaking efforts in general, including those that involve AI.³¹

Nevertheless, it is possible that, in the next few months, federal agencies will release a slew of proposed rules that would roll back vast swaths of the Code of Federal Regulations and give the public only 30 to 60 days to respond.³² It will likely be difficult for those opposing the administration's deregulatory efforts—groups who are overtaxed already—to effectively respond to an onslaught of that scale. Advocates may struggle to find the bandwidth to scrutinize all of the administration's proposed changes, draft effective comments, and ultimately, prepare for litigation as necessary. This Issue Brief aims to provide a starting point for how affected communities might respond to this deregulatory wave, if it materializes.

III. ADMINISTRATIVE LAW GUARDRAILS

Below, we sketch out several administrative-law doctrines that might be implicated by agencies rubber-stamping AI output in the rulemaking context. Depending on how the Trump administration implements the proposed DOGE AI plan, some of these doctrines might provide viable APA claims in a lawsuit challenging the plan.³³ We reiterate, however, that this assessment is preliminary, and any such challenge would depend upon the specific manner in which the DOGE AI tool is trained and utilized. Again, our analysis is directed at a potential hypothetical use of AI in the rulemaking process that may or may not overlap with how agencies intend to use the DOGE AI tool—one in which AI is used with only minimal human oversight to flag rules that ought to be rescinded; draft regulatory text accomplishing that rescission; and

³¹ See, e.g., Ilisa Halpern Paul & Ian Rockwell, *The Federal Regulatory Process During a Government Shutdown*, Venable LLP (Oct. 20, 2025), <https://www.venable.com/insights/publications/2025/10/the-federal-regulatory-process-during-a-government>.

³² See *DOGE Deregulation Opportunity*, *supra* note 5, at 9.

³³ Such arguments could be asserted in the context of challenges to specific deregulatory actions prepared using the DOGE AI Tool, or potentially, in a challenge to any mandate (likely from the Office of Management and Budget) directing agencies to use the DOGE AI Tool in a particular way, although the latter may involve complicated justiciability and other questions.

produce supporting documents justifying the agency’s action and purporting to “consider” legally relevant issues, alternatives, contrary evidence, arguments, and public comments.

Several cross-cutting issues also lurk in the background. The most obvious is that the APA was not drafted with artificial intelligence in mind, and the seminal decisions interpreting and applying the APA were issued long before the advent of the large language model—and, in some cases, modern computing. Any effort to apply the APA to agencies’ use of artificial intelligence is therefore likely to confront numerous issues of first impression.

Another issue is courts’ general unwillingness to police the internal deliberative processes of agencies, including how they use their staff and sophisticated technological systems. “The presumption of regularity supports the official acts of public officers, and, in the absence of clear evidence to the contrary, courts presume that they have properly discharged their official duties.”³⁴ Thus, “inquiry into the mental processes of administrative decisionmakers is usually to be avoided,” except where there has been a “strong showing of bad faith or improper behavior.”³⁵ And yet courts do sometimes grant discovery beyond the scope of the administrative record to probe an agency’s decisionmaking where litigants have made the requisite showing.³⁶ The leak of the DOGE plan, coupled with any available evidence about how agencies are using the DOGE AI tool, might assist litigants in doing so—as might the many cases where courts have found the presumption of regularity to be overcome by the Trump administration’s conduct.³⁷

Finally, courts are prohibited from imposing procedural requirements upon agencies beyond those set forth in the APA or other statutory or constitutional provisions.

³⁴ *United States v. Chem. Found.*, 272 U.S. 1, 14 (1926).

³⁵ *Citizens to Preserve Overton Park v. Volpe*, 401 U.S. 402, 420 (1971).

³⁶ See, e.g., Governing for Impact, *Challenging Agency Action Based on Pretextual Reasons* 6–9 (2025), <https://governingforimpact.org/wp-content/uploads/2025/05/Challenging-Agency-Action-Based-on-Pretextual-Reasons.pdf>; Governing for Impact, *Challenging DOGE* 12–16 (2025), <https://governingforimpact.org/wp-content/uploads/2025/02/Challenging-DOGE-Primer-final.pdf>; *Nat’l TPS Alliance v. Noem*, 2025 WL 2419266, at *2 (N.D. Cal. Aug. 21, 2025) (“The Court concludes that this case falls within the ‘bad faith’ exception to the general practice of prohibiting extra-record discovery in APA cases.”).

³⁷ See Ryan Goodman et al., *The “Presumption of Regularity” in Trump Administration Litigation*, Just Security (Oct. 15, 2025), <https://www.justsecurity.org/120547/presumption-regularity-trump-administration-litigation>.

“Time and again,” the Supreme Court has reiterated that the APA “sets forth the full extent of judicial authority to review executive agency action for procedural correctness.”³⁸ “Beyond the APA’s minimum requirements, courts lack authority ‘to impose upon [an] agency its own notion of which procedures are “best” or most likely to further some vague, undefined public good.’”³⁹ Given that the APA’s text does not expressly impose any procedures or requirements on how agencies use AI (how could it?), any assertion that the APA governs the use of AI will need to find an anchor in established administrative-law doctrines.

Notwithstanding these difficulties, litigants might press four plausible APA-based challenges to AI rubber-stamping, founded on four overlapping yet distinct duties that apply to agencies engaging in rulemaking: the duty to give reasons, the duty to engage in reasoned decisionmaking, the duty to address significant comments, and the duty to keep an open mind.

A. Duty to Give Reasons

By requiring agencies to “incorporate in the rules adopted a concise general statement of their basis and purpose,” the APA directs agencies to formulate and disclose their reasons for adopting a particular rule.⁴⁰ As the Supreme Court put it in *SEC v. Chenery Corp.* (1943), “the orderly functioning of the process of review requires that the grounds upon which the administrative agency acted be clearly disclosed and adequately sustained.”⁴¹ To be sure, agencies are not required to account for every reason motivating their actions, but the explanations they offer cannot be “contrived” or “a distraction.”⁴² In other words, they must accurately reflect the *agency’s own reasons* for promulgating a particular rule.

An agency that does no more than rubber-stamp an AI-generated justification has arguably not satisfied that reason-giving requirement. The APA and cases interpreting it suggest that agencies not only have a duty to offer reasons, but to actually engage in the process of reasoning. The APA’s notice-and-comment

³⁸ *Perez v. Mortg. Bankers Ass’n*, 575 U.S. 92, 101–02 (2015) (quoting *FCC v. Fox Television Stations, Inc.*, 556 U.S. 502, 513 (2009)).

³⁹ *Id.* at 102 (quoting *Vt. Yankee Nuclear Power Corp. v. NRDC*, 435 U.S. 519, 549 (1978)).

⁴⁰ 5 U.S.C. § 553(c).

⁴¹ 318 U.S. 80, 94 (1943).

⁴² *Dep’t of Commerce v. New York*, 588 U.S. 752, 784–85 (2019).

procedures permit agencies to issue rules only “after *consideration* of the relevant matter presented.”⁴³ And Section 706 prohibits arbitrary-and-capricious agency action, which, the Supreme Court explicated in *Motor Vehicles Mfrs. Ass’n v. State Farm Mut. Auto. Ins. Co.*, means that agencies “must *examine* the relevant data,” “*articulate* a satisfactory explanation,” and base decisions “on a *consideration* of the relevant factors.”⁴⁴ That is the “domain which Congress has set aside exclusively for the administrative agency.”⁴⁵

Congress, in enacting these requirements, expected that agencies would reach sound results by actually reasoning through policy problems. This process, of course, generally involves collaboration and dialogue between senior and subordinate agency personnel.⁴⁶ Ultimately, as the D.C. Circuit explained in *Independent U.S. Tanker Owners Committee v. Lewis*, the APA’s reason-giving requirement

play[s] an integral role in the decisionmaking process. It ensures thoughtful consideration of the various issues raised before a decision is published. ... Furthermore, in the process of preparing the statement [of a rule’s basis and purpose], staff members themselves may become aware of unforeseen difficulties with the decision, with the result that individual decisionmakers change their mind.⁴⁷

An agency that simply adopts an AI-generated reason for a particular action has not engaged in this process.

The D.C. Circuit, in discussing agencies’ use of “sophisticated methodology,” has referred to an “insistence that ultimate responsibility for the policy decision remains with the agency rather than the computer.”⁴⁸ In that vein, former Justice of the Supreme Court of California Mariano-Florentino Cuéllar has posited a “distinction between essentially turning a decision over to a computer program, and having the

⁴³ 5 U.S.C. § 553(c) (emphasis added); see also *National Small Shipments Traffic Conf., Inc. v. ICC*, 725 F.2d 1442, 1451 (D.C. Cir. 1984) (describing “the consideration requirement” of § 553).

⁴⁴ 463 U.S. 29, 43 (1983) (emphases added).

⁴⁵ *SEC v. Chenery Corp.*, 332 U.S. 194, 196 (1947).

⁴⁶ See *Small Shipments*, 725 F.2d at 1050 (recognizing that agency leadership delegates tasks to subordinates).

⁴⁷ 690 F.2d 908, 920 (D.C. Cir. 1982).

⁴⁸ *Sierra Club v. Costle*, 657 F.2d 298, 334 (D.C. Cir. 1981).

program merely support a human decision-maker.”⁴⁹ While it is difficult to identify “a cogent standard of how much reliance on a computer is too much,” simply adopting an AI’s work without any human deliberation or reflection would seem to fall on the wrong side of the line, just as it would violate the APA for an agency to reflexively adopt the output of a mathematical model in other contexts. Moreover, an agency that rubber-stamps an AI-generated justification for an action it has already decided to take may well be violating the APA’s requirement that agency reasons not be “contrived” or otherwise pretextual.⁵⁰

Of course, as soon as duly serving agency personnel become involved in revising or iterating on the AI’s output, things become far more complicated. For instance, the DOGE plan sends mixed signals about the degree to which agency staff will be involved in rulemaking. On the one hand, it provides that “Policy and Legal Teams Make All The Decisions.”⁵¹ But it also trumpets that “[a]nalysis” and “[w]riting” are “[a]utomated.”⁵² To that end, it contemplates only “2.4” hours of human work per deregulatory action: 12 minutes for deciding what regulatory action to take, 48 minutes for “[d]raft[ing] the regulatory submission documents,” 18 minutes for “[a]nalyz[ing] & respond[ing] to 100,000+ comments,” and one hour for “drafting the Final Rule including comment responses.”⁵³ Whether agencies carrying out the DOGE plan would satisfy the duty to give reasons likely turns on which of these seemingly contradictory claims turns out to be true — i.e., whether humans meaningfully review the AI’s output or merely rubber-stamp it.

B. Duty to Engage in Reasoned Decisionmaking

Not only must agencies formulate and articulate their own reasons for acting, but the agency action must itself be “reasonable and reasonably explained.”⁵⁴ Under the APA’s arbitrary-and-capricious standard, a court must ensure that the agency “has

⁴⁹ Mariano-Florentino Cuéllar, *Cyberdelegation and the Administrative State* 20 (Oct. 2016) (unpublished manuscript), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2754385.

⁵⁰ *Dep’t of Commerce*, 588 U.S. at 784–85.

⁵¹ *DOGE Deregulation Opportunity*, *supra* note 5, at 8.

⁵² *Id.*

⁵³ *Id.*

⁵⁴ *FCC v. Prometheus Radio Proj.*, 592 U.S. 414, 423 (2021).

reasonably considered the relevant issues and reasonably explained its decision.”⁵⁵ In other words, the APA “requires agencies to engage in reasoned decisionmaking.”⁵⁶ Rubber stamping AI-generated rules risks running afoul of that requirement.

At present, large language models are prone to a number of common errors which, if they manifest in agency action, would likely lead a court to strike the action down as arbitrary and capricious.⁵⁷ These errors include:

- **Hallucination:** LLMs have been known to generate false information— a problem that has persisted even as AI tools have advanced in other respects.⁵⁸ In one high-profile example involving the Trump administration, a report on chronic disease issued by the Department of Health and Human Services appears to have contained AI-generated citations to non-existent studies.⁵⁹
- **Sycophancy:** LLMs exhibit a tendency to “pursue[] human approval” by generating “overly flattering or agreeable responses.”⁶⁰ “For instance, research shows that models will even agree with objectively incorrect mathematical statements if primed by users.”⁶¹
- **Bias:** LLMs frequently replicate biases or errors present in their training data.⁶² One way that can happen is that “LLMs are trained to give high probability to human-generated text, and in so doing they learn latent

⁵⁵ *Id.*

⁵⁶ *Dep’t of Homeland Sec. v. Regents of the Univ. of Calif.*, 591 U.S. 1, 16 (2020) (quotation omitted).

⁵⁷ Some of these errors relate to the broader challenge of AI alignment, which refers to “encoding human values and goals into AI models to make them as helpful, safe and reliable as possible.” Alexandra Jonker & Alice Gomstyn, *What Is AI Alignment?*, IBM, <https://www.ibm.com/think/topics/ai-alignment> (last accessed Nov. 18, 2025).

⁵⁸ See, e.g., Cade Metz & Karen Weise, *A.I. Is Getting More Powerful, but Its Hallucinations Are Getting Worse*, N.Y. Times (May 5, 2025).

⁵⁹ See, e.g., Phie Jacobs, *Trump Officials Downplay Fake Citations in High-Profile Report on Children’s Health*, Science (May 30, 2025), <https://www.science.org/content/article/trump-officials-downplay-fake-citations-high-profile-report-children-s-health>.

⁶⁰ See, e.g., Erie Meyer et al., *Tech Brief: AI Sycophancy & OpenAI*, Georgetown Inst. for Tech L. & Pol’y (2025), <https://www.law.georgetown.edu/tech-institute/insights/tech-brief-ai-sycophancy-openai-2/> (quotations omitted).

⁶¹ *Id.*

⁶² James Holdsworth, *What is AI Bias?*, IBM, <https://www.ibm.com/think/topics/ai-bias> (last accessed Nov. 18, 2025).

structure that underlies that text,” but “can’t distinguish structure connected with harmful biases from any other aspect of latent structure.”⁶³

- **Context window limitations:** LLMs have limited “context windows,” a term that broadly refers to the amount of text or information they can consider at one time when generating a response.⁶⁴ LLMs may thus struggle to accurately review or analyze long documents — a matter of particular concern in the rulemaking process, which often involves lengthy and complicated records.⁶⁵

These systemic issues map onto common categories of APA violations, such as when an agency “relie[s] on factors which Congress has not intended it to consider, entirely fail[s] to consider an important aspect of the problem, [or] offer[s] an explanation for its decision that runs counter to the evidence before the agency,” or when the agency’s decision “is so implausible that it could not be ascribed to a difference in view or the product of agency expertise.”⁶⁶ An LLM incapable of faithfully reviewing complex documents or portions of the administrative record might produce a result that runs contrary to the evidence. Similarly, proposals premised on hallucinations are likely to be substantively unreasonable. Thus, to the extent agencies increasingly rely on AI to generate and explain regulatory proposals, they increase the risk of exposure to an arbitrary and capricious claim.

These reliability concerns only emphasize an APA obligation that would constrain agency use of AI in any event: the requirement that the agency explain and justify its methodology. When agencies use mathematical models to generate substantive regulations, they must justify and validate those models. They “retain[] a duty to

⁶³ Philip Resnik, *Large Language Models Are Biased Because They Are Large Language Models*, 51 *Computational Linguistics* 885 (2025), <https://direct.mit.edu/coli/article/51/3/885/128621/Large-Language-Models-Are-Biased-Because-They-Are>.

⁶⁴ See Dave Bergmann, *What is a Context Window?*, IBM, <https://www.ibm.com/think/topics/context-window> (last accessed Nov. 18, 2025).

⁶⁵ See, e.g., Samuel Estreicher & Lior Polani, *Introduction: The Promise and Reality of Legal AI*, *Justia Verdict* (Aug. 8, 2025) (“Legal work demands precisely what AI systems struggle to provide: comprehensive simultaneous analysis of interconnected documents.”).

⁶⁶ *State Farm*, 463 U.S. at 43. For a discussion of various kinds of errors in reasoning courts find arbitrary and capricious, see Governing for Impact, *Arbitrary-and-Capricious Challenges* (May 2025), <https://governingforimpact.org/wp-content/uploads/2025/05/Arbitrary-and-Capricious-Challenges.pdf>.

examine key assumptions.”⁶⁷ That same obligation likely holds true as technology evolves; agencies must affirmatively explain their usage of AI in similar fashion.⁶⁸ Indeed, agencies’ “ex ante” explanations of how their AI systems are designed, prompted, and validated may prove insufficient to explain the actual basis for the agency’s action.⁶⁹ It follows even more strongly that an agency has not complied with its reasoned decisionmaking obligations if it offers *no* explanation as to the reliability of its AI usage.

Agencies must also be “conscious of the limits of the model.”⁷⁰ This principle derives from a concern that agencies will come to reflexively rely on mathematical models to decide regulatory matters for them. That concern may well be even more acute for AI, to the extent that human decisionmakers place excessive trust or confidence in AI-generated output. Even draft reasons generated by AI might anchor decisionmakers to faulty reasons or otherwise inhibit the reasoning process.⁷¹ The D.C. Circuit has said that agencies may not even rely on *human staff* who “distort the record” or produce “severely skewed” work product.⁷² AI errors might similarly taint an agency’s decisionmaking process.

An agency might be required to consider these issues throughout the process of utilizing AI to formulate a rule and an accompanying justification. The agency might need to provide a reasonable explanation for why it chose to use an AI system at all, or to use a particular large language model, in light of these known deficiencies. The agency might also need to reasonably explain how it used the model, including such questions as how it was trained and prompted and how the agency reviewed its output. And the agency almost certainly would need to ensure that both the action and explanation supplied by the AI system are themselves reasonable, as agencies

⁶⁷ *Columbia Falls Aluminum Co. v. EPA*, 139 F.3d 914, 923 (D.C. Cir. 1998); accord *Powder River Basin Res. Ctr. v. Dep’t of Interior*, 749 F. Supp. 3d 151, 164 (D.D.C. 2024); *Owner-Operator Indep. Drivers Ass’n v. FMCSA*, 494 F.3d 188, 204 (D.C. Cir. 2007)

⁶⁸ See Note, *Machine Rulemaking: Arbitrary and Capricious Review in the Age of AI*, 138 Harv. L. Rev. 1821, 1830 (2025).

⁶⁹ See *id.* at 1832–34 (describing the limits of “ex ante review” in the AI context).

⁷⁰ *Small Refiner Lead Phase-Down Task Force v. EPA*, 705 F.2d 506, 535 (D.C. Cir. 1983).

⁷¹ See Bridget Dooling, *Ghostwriting the Government*, 109 Marq. L. Rev. (forthcoming 2026), manuscript at 26–27 (describing anchoring dynamics); *Independent U.S. Tanker Owners Committee*, 690 F.2d at 920 (same).

⁷² *Small Shipments*, 725 F.2d at 1451.

have been required to do since long before the advent of AI. Simply rubber-stamping the output of an AI system would be arbitrary and capricious.

Such an argument admittedly faces counterarguments and objections.

First, perhaps outsourcing the drafting of regulatory material to AI is no different than outsourcing such work to subordinate agency staff, which occurs as a matter of course. Moreover, within statutory limits, agencies frequently have discretion to structure and define their operations.⁷³ As far as we are aware, agencies have never been required to justify the choice to use human staff in a particular manner, and perhaps they should not be required to justify their use of AI either.

Our tentative view, however, is that an AI tool is not the equivalent to human staff under these circumstances. The indispensable role of subordinates in the administrative process has long been established by policy and practice and, as discussed below, blessed by law.⁷⁴ And agencies are structured—by statute and regulation—to foster exchange between junior and senior agency staff, along with necessary accountability and supervision. Courts have long recognized that the operation of automatic tools cannot, without appropriate explanation and oversight, serve as a substitute for that process.⁷⁵ And that legal result is also good policy: automated tools cannot be held accountable for bad decisions in the same way, and are not subject to the same ethical constraints, as human agency staff.⁷⁶ Of course, once humans get involved in working “collaboratively” with AI, that makes the analysis more complicated. That kind of back-and-forth might reflect that agency staff are using AI as a decision-support tool, not as a wholesale substitute for their own work, a dynamic that could more readily satisfy the APA’s reasoned decisionmaking requirement.

Second, courts may not find it relevant whether an AI tool was used to generate or justify agency action. Generally, agency action and supporting reasoning should be

⁷³ Cf. *FCC v. Pottsville Broadcasting Co.*, 309 U.S. 134, 143 (1940) (agencies “should be free to fashion their own rules of procedure and to pursue methods of inquiry capable of permitting them to discharge their multitudinous duties”).

⁷⁴ See, e.g., *Small Shipments*, 725 F.2d at 1450.

⁷⁵ See *Sierra Club*, 657 F.2d at 334.

⁷⁶ See *Machine Rulemaking*, *supra* note 67, at 1835–36 (“Because humans know that they may have to explain themselves, they are more likely to operate in a way that could be justified and considered reasonable. In contrast, the [machine learning] model is constrained in no such way.”).

judged on their own terms—regardless of who or what generated them, they are either reasonable or they are not.⁷⁷ Moreover, as discussed above, if an agency chooses to rely on an AI system rather than expert staff, it runs the risk of producing an arbitrary and capricious action. If that risk materializes, the court need not pry into the agencies’ decisionmaking procedures to perform the task of setting that rule aside.

But even if this reasoning precluded AI-specific inquiries as to an agency’s rule preambles or supporting reasons, it would not block review of the use of AI in crafting the underlying agency action, at least so long as the agency relied on the AI tool as a basis for its action. Agencies justifying their actions based on the outputs of mathematical models have to justify the assumptions and methods of their models. It follows that agencies basing their actions on AI systems must offer some defense of those systems.

Rubber-stamping could be relevant to judicial review in other ways, too. Arbitrary and capricious review is frequently described as deferential to agency expertise and careful consideration.⁷⁸ But maximalist use of AI minimizes the role of those factors, undercutting the basis for deference. And procedurally, evidence of a significant deviation from lawful agency procedure—like rubber-stamping the DOGE AI tool’s output—could well rebut the presumption of regularity and the rule that review of agency decisions is based on the existing administrative record, justifying discovery into particular rulemakings.

Third, some might point out that because it is an open question whether AI systems can explain their decisions,⁷⁹ a requirement that LLMs tasked with regulatory drafting offer genuine reasons for particular actions could make it essentially impossible to use AI as a regulatory tool. They might also note that such a requirement would go beyond what is required of human agency staff, who need only

⁷⁷ See Kullgren, *supra* note 8 (former DOGE General Counsel James Burnham opining that “courts should judge rules on the quality of the final product and nothing else”).

⁷⁸ *E.g.*, *Dep’t of Commerce*, 588 U.S. at 773.

⁷⁹ Compare Boris Babic & I. Glenn Cohen, *The Algorithmic Explainability “Bait and Switch,”* 108 *Minn. L. Rev.* 857 (2023), with Vrunda Gadesha et al., *What is Chain of Thought (CoT) Prompting?*, IBM, <https://www.ibm.com/think/topics/chain-of-thoughts> (last accessed Nov. 18, 2025); *Reasoning Models Don’t Always Say What They Think*, Anthropic (Apr. 3, 2025), <https://www.anthropic.com/research/reasoning-models-dont-say-think> (“Reasoning models are more capable than previous models. But our research shows that we can’t always rely on what they tell us about their reasoning.”).

produce plausible reasons and are not subject to psychological examination by reviewing courts. But these concerns are misplaced. We are not addressing here whether and to what extent AI systems must be able to authentically explain the basis for their decisions. The question is whether an agency complies with *its* legal obligation by simply rubber-stamping AI-generated reasons.

C. Duty to Address Significant Comments

Rubber-stamping an AI-generated final rule may also violate the APA’s requirement that agencies adequately consider the comments submitted during the rulemaking process. That duty stems from both 5 U.S.C. § 553(c), which requires agencies to “give interested persons an opportunity to participate in the rule making through submission of written data, views, or arguments” and then engage in “consideration of the relevant matter presented,”⁸⁰ as well as from the APA’s general reasoned decisionmaking requirement which, in this context, directs agencies to “demonstrate the rationality of [their] decision-making process[es] by responding to those comments that are relevant and significant.”⁸¹ In that way, requiring agencies to invite and consider public comment serves the APA’s underlying goals of “foster[ing] public participation and facilitat[ing] reasoned decisionmaking.”⁸²

Admittedly, the APA does not specify the precise manner in which agencies must process, evaluate, and address comments. In a pre-APA case, *Morgan v. United States* (1936), the Supreme Court held that a decisionmaker must “address[] himself to the evidence,” and “upon that evidence ... conscientiously reach[] the conclusions which he deems it to justify” — in other words, “[t]he one who decides must hear.”⁸³ But the Court also hedged, clarifying that “[t]his necessary rule does not preclude practicable administrative procedure in obtaining the aid of assistants in the department.”⁸⁴ Today, it is well-established that “an agency decisionmaking body ...

⁸⁰ 5 U.S.C. § 553(c).

⁸¹ *Grand Canyon Air Tour Coal. v. FAA*, 154 F.3d 455, 468 (D.C. Cir. 1998).

⁸² *Humane Soc’y v. U.S. Dep’t of Agric.*, 41 F.4th 564, 568 (D.C. Cir. 2022).

⁸³ *Morgan v. United States (Morgan I)*, 298 U.S. 468, 481 (1936).

⁸⁴ *Id.* *Morgan I* was the first of four decisions issued by the Supreme Court in the *Morgan* case, and over the course of the litigation, the Court increasingly made plain that it did not intend to permit intrusive scrutiny of administrative proceedings; “[j]ust as a judge cannot be subjected to such a scrutiny, so the integrity of the administrative process must be equally respected.” *United States v. Morgan (Morgan IV)*, 313 U.S. 409, 422 (1941); see also *Nat’l Nutritional Foods Ass’n v. FDA*, 491 F.2d 1141, 1144 (2d Cir. 1974) (Friendly, J.) (“[T]he life of this aspect of *Morgan I* was extremely brief.”).

may delegate detailed consideration of the administrative record to its subordinates while retaining the final power of decision for itself.”⁸⁵ It is therefore difficult to argue that the APA requires the official who signs a final rule to personally review and consider every comment submitted by the public.

But that does not mean that the APA has nothing to say about how agencies process comments. To the contrary, the D.C. Circuit suggested in *National Small Shipments Traffic Conference v. ICC* (1984) that an agency must process comments in a manner sufficient to indicate that the agency meaningfully considered and ventilated the issues that the comments raised. Although the “strong presumption of regularity” means that “courts will not normally entertain procedural challenges that [agency decisionmakers] inadequately considered the issues before reaching a final decision,” “[a]t some point, ... staff-prepared synopses may so distort the record that an agency decisionmaking body can no longer rely on them in meeting its obligations under the law.”⁸⁶ “More particularly, in informal rulemaking employing notice-and-comment procedures, dependence on severely skewed staff summaries may breach the decisionmaker’s statutory duty to accord ‘consideration’ to relevant comments submitted for the record by interested parties.”⁸⁷ *National Small Shipments* therefore suggests that defects in an agency’s procedures that prevent agency decisionmakers from adequately considering comments may violate the APA, although the contours of that rule are far from clear.

If any use of AI might violate the rule from *National Small Shipments*, rubber-stamping of AI-generated material likely does. In our hypothetical, an agency decisionmaker would not substantively engage with comments at all, leaving that work to AI. Even more extensive human review may be insufficient if, for example, the AI tool cannot reliably summarize or explain adverse comments in a way that enables a human decisionmaker to consider the issues that they raise. Or a human decisionmaker might prompt the AI tool to explain why the comments are irrelevant or misguided, making it more likely that the AI’s output presents a distorted version

⁸⁵ *Small Shipments*, 725 F.2d at 1450.

⁸⁶ *Id.* at 1450–51.

⁸⁷ *Id.* at 1451 (quoting 5 U.S.C. § 553(c)); see also *United Steelworkers of Am., AFL-CIO-CLC v. Marshall*, 647 F.2d 1189, 1217 (D.C. Cir. 1980) (requiring “specific proof” that the responsible official “failed to confront personally the essential evidence and arguments in setting the final standard”); *S. Garment Mfrs. Ass’n v. Fleming*, 122 F.2d 622, 625–26 (D.C. Cir. 1941) (“[H]ear’ is used in the artistic sense of requiring certain procedural minima to insure an informed judgment by the one who has the responsibility of making the final decision and order.”).

of the record. As noted above, notwithstanding the DOGE presentation’s assertions of human review, the extremely accelerated timeline for human analysis of comments — 18 minutes⁸⁸ — certainly gives the impression that agency staff (and, by extension, the decisionmaker) will not have sufficient time to scrutinize comments and consider the issues that they raise.

That need not be so: a well-designed and trained AI system could be used in ways that facilitate meaningful human consideration of comments. For example, agencies might use AI to group or summarize comments while leaving responses to staff or prompt the AI to explain which comments are relevant and substantial so that staff can focus on those comments and make any necessary changes to the agency’s proposal. Agency staff might also use the AI’s output as a starting point but nonetheless review individual comments on a case-by-case basis. Moreover, AI might even enhance public participation, in the sense that it might help agency staff devote greater time to those comments which merit the most extensive consideration. These alternative approaches only underscore the impermissibility of rubber-stamping AI output in a manner that frustrates, and indeed, nullifies, the APA’s requirement of public input.

D. Duty to Keep an Open Mind

Similarly, relying on AI to “consider” and perfunctorily reject adverse comments and serious flaws in an agency’s proposal might violate an agency’s duty to keep an open mind during the rulemaking process.⁸⁹ The standard for such a claim is demanding, and rarely met: a decisionmaker can be disqualified from an administrative process “only when there has been a clear and convincing showing that the [official] has an unalterably closed mind on matters critical to the disposition of the proceeding.”⁹⁰ These claims are typically advanced in circumstances where an official has publicly opined on the policy issues presented by a rulemaking, and are in that context almost always rejected.⁹¹ After all, such public statements do not necessarily indicate that

⁸⁸ See *DOGE Deregulation Opportunity*, *supra* note 5, at 8.

⁸⁹ Such a claim draws from both the Due Process Clause and general principles of administrative law. See *Cinderella Career & Finishing Sch., Inc. v. FTC*, 425 F.2d 583, 584 (D.C. Cir. 1970).

⁹⁰ *Ass’n of Nat. Advertisers, Inc. v. FTC*, 627 F.2d 1151, 1170 (D.C. Cir. 1979).

⁹¹ See, e.g., *C & W Fish Co. v. Fox, Jr.*, 931 F.2d 1556, 1565 (D.C. Cir. 1991) (“The harm that would result were courts to disqualify agency members whenever they express views in public, as Fox did here, is readily apparent.”); *Ass’n of Nat. Advertisers*, 627 F.3d at 584 (“We would eviscerate the proper

the official has “a mind that has been closed to the evidence in the past or that would disregard any significant new material subsequently introduced.”⁹²

Prompting an AI tool to identify a rule for rescission, draft all rulemaking materials, and automatically refute any adverse comments gives rise to a unique form of an unalterably closed mind problem, however. Depending on how an AI system is prompted and the degree to which it exhibits sycophancy, it may be incapable of reconsidering its views, admitting that an adverse comment raises an important point, or flagging an idea up the chain for further consideration. In contrast, even a human decisionmaker with strong opinions about the matter in question might at least conceivably be swayed by contrary evidence. Put differently, an AI system’s mind can be unalterably closed at the push of a button. There is therefore some reason to think that litigants challenging AI rubber-stamping may find more success in bringing unalterably closed mind claims than those challenging biased human decisionmakers historically have.

Of course, a human decisionmaker who reviews the output of such an AI system may not themselves possess an unalterably closed mind—they remain capable of reviewing the AI system’s output and deciding whether it warrants a different conclusion. But a decisionmaker who simply rubber-stamps the AI system’s output has failed to engage in meaningful consideration of contrary viewpoints. Even a decisionmaker who engages in some amount of review of what the AI has generated may be unable to reopen what biased prompting has closed, to the extent they rely upon the AI’s distorted summaries or are anchored to the AI’s responses. Indeed, the abbreviated and cursory nature in which DOGE proposes to use the DOGE AI tool suggests that the tool will be prompted simply to rationalize a predetermined conclusion — akin, perhaps, to instructing a human subordinate to ignore any contrary facts or ideas.

evolution of policymaking were we to disqualify every administrator who has opinions on the correct course of his agency’s future action.”).

⁹² *Consumers Union of U.S., Inc. v. FTC*, 801 F.2d 417, 427 (D.C. Cir. 1986).

IV. CONCLUSION

By sketching out arguments as to why the APA forbids agency rubber-stamping of AI output, this Issue Brief might provide a starting point for those developing arguments as to why a specific agency’s use of AI in a given rulemaking is unlawful — depending, of course, on the particular facts involved, and subject to the caveats and open questions raised above. AI may have salutary uses as an administrative tool in the hands of a responsible administration. But rubber-stamping threatens to displace reasoned decisionmaking with AI-generated output. Litigants might be able to challenge extreme uses of AI without unduly trenching on a future administration’s legitimate discretion or prudent deployment of AI.

The information in this document is provided for informational purposes only and does not contain legal advice, legal opinions, or any other form of advice regarding any specific facts or circumstances and does not create or constitute an attorney-client relationship. You should contact an attorney to obtain advice with respect to any particular legal matter and should not act upon any such information without seeking qualified legal counsel on your specific needs.