Yale Journal of Law & Technology
Volume 29, Issue 1

Deference in the Age of AI

Juan Caballero[*]

The United States is in the midst of two concurrent revolutions; one is technological, the other precedential. The current technological revolution began in November 2022 when OpenAI released ChatGPT. This release led to a rapid acceleration and proliferation of AI technology, particularly large language models (LLMs), on the global stage. Companies and governments alike were eager to exploit this new technology. In the United States, agencies across the administrative state have deployed AI systems to automate and streamline a wide variety of functions, from determining public benefit eligibility, to screening immigrant visa applicants, to reviewing classified government records for publication. Generative AI technology is ubiquitous.

The second, precedential revolution began two years later when the Supreme Court decided Loper Bright Enters. v. Raimondo. This decision overturned the Chevron deference paradigm that had shaped administrative law for over four decades. In the wake of this decision, lower courts are still in the process of determining the appropriate analysis for deciding when and how to defer to administrative decisions.

As federal agencies increasingly rely on AI systems, the judiciary faces novel questions about when and how much deference these AI-assisted and AI-driven decisions warrant. As legal doctrine and technology continue to develop, courts must adapt deference principles to ensure meaningful oversight in an

---

[*] Associate Professor of Law, Seattle University School of Law. For helpful comments, I am grateful to participants in the Works-in-Progress Panel at the 2025 Clinical Law Review Workshop including Professors Ali Ekber Cinar, Rebecca Delfino, and Michele Gilman.

The author used AI systems to assist with limited tasks, including preliminary research, bluebooking, and editing for grammar, clarity, and flow. The author alone generated the ideas, wrote the initial drafts, and organized the arguments. Any errors are mine.

era of automated governance.

This article proposes an inverse proportionality rule for allocating judicial deference. Under such a rule, judicial deference to an agency's interpretations will be inversely related to the agency's reliance on generative AI technology, to a point; complete outsourcing of reasoning responsibility to AI systems violates the APA and should be struck down. The traditional deference doctrine that the Supreme Court revived in Loper Bright typically rests on the premise of reasoned decision-making by accountable human experts. AI systems, however, operate as opaque "black boxes" that can sever the chain of democratic accountability underlying such delegations. Between full automation and purely human decision-making lies a spectrum requiring a nuanced analytical structure.

Article Contents

Introduction 178
I. The End of an Era: *Loper Bright* and the Future of Administrative Deference 184
  A. Early 20th Century Development of the Administrative State 184
  B. Relentless Attacks: Loper Bright and Chevron's Demise 188
    1. What Remains: "Respect" versus "Deference" 190
    2. An Uncertain Landscape Post-*Loper Bright* 192
II. The AI Revolution in Administrative Governance 196
  A. Artificial Intelligence: Technology, Definitions, and Current State 196
  B. The Proliferation of AI in Federal Agencies 200
  C. The Promise and Peril of AI in Government 206
    1. The Promise: Efficiency, Consistency, and Enhanced Government Performance 206
    2. The Perils: Transparency Deficits and Systematic Risks 209
      a. Public Skepticism 209
      b. The Black Box Problem: Transparency and Interpretability 212
      c. Hallucinations and the Reliability Challenge 215
      d. Bias Amplification and Systematic Discrimination 218
III. Judicial Deference and Its Theoretical Underpinnings 220
  A. The Challenge to Traditional Expertise Models Posed by AI 222
    1. The Fundamental Humanity and Fallibility of Expertise 223
    2. The Human Observation Paradigm in Legal Doctrine 225
    3. False Rationality of AI Systems 231
IV. A Framework for AI-Era Deference: The Sliding

Scale Approach 232
A. The Inverse Proportionality Rule 234
B. The Spectrum of Human-AI Collaboration 236
1. Purely Human Decision-Making 237
a. Rulemaking 238
b. Adjudications 240
2. Partially Automated Decision-Making 242
a. Rulemaking 243
b. Adjudications 247
3. Fully Automated Decision-Making 251
a. Rulemaking 253
b. Adjudications 256
Conclusion 261

Introduction

"[A] country of geniuses in a datacenter"[1]

Two revolutions are reshaping the landscape of American administrative law. The first arrived on November 30, 2022, with the public release of OpenAI's GPT-3.5, a large language model, and the subsequent proliferation of generative artificial intelligence ("AI") software.[2] As a result, global society is experiencing an unprecedented acceleration of generative AI technology that can be adapted to a wide range of tasks and settings.[3] The second revolution took place when the Supreme Court declared the end of four decades of deference precedent in its *Loper Bright Enterprises v. Raimondo* decision.[4] *Loper*

---

[1] Dario Amodei, *Machines of Loving Grace* (Oct. 2024), https://darioamodei.com/essay/machines-of-loving-grace#4-peace-and-governance [https://perma.cc/LQT5-LBNJ].

[2] *See Introducing ChatGPT*, OpenAI (Nov. 30, 2022), https://openai.com/index/chatgpt/ [https://perma.cc/XKE6-2VCQ] (announcing the release of ChatGPT and signaling the beginning of the modern AI arms race).

[3] *See* Sha Sajadieh et al., *Artificial Intelligence Index Report 2026*, STAN. UNIV. INST. FOR HUM.-CENTERED AI 1, 71 (Apr. 2026), https://arxiv.org/pdf/2606.15708 [https://perma.cc/CU5W-ZHFA] ("AI capability is outpacing the benchmarks designed to measure it, and surpassing human-level performance. Frontier models gained 30 percentage points in a single year on Humanity's Last Exam, a benchmark built to be hard for AI and favorable to human experts. Evaluations intended to be challenging for years are saturated in months, compressing the window in which benchmarks remain useful for tracking progress."); *id.* at 171 ("Generative AI reached close to 53% population-level adoption within three years of its mass-market introduction, faster than the personal computer or the internet . . . ."); Daniel E. Ho, *Hallucinating Law: Legal Mistakes with Large Language Models Are Pervasive*, STAN. UNIV. INST. FOR HUM.-CENTERED AI (Jan. 11, 2024), https://hai.stanford.edu/news/hallucinating-law-legal-mistakes-large-language-models-are-pervasive [https://perma.cc/GA7T-K4AN] (discussing AI use cases in finance, education, healthcare, computer science, and the law); *see also* Paul Ford, *The A.I. Disruption We've Been Waiting for Has Arrived*, N.Y. TIMES (Feb. 18, 2026), https://www.nytimes.com/2026/02/18/opinion/ai-software.html [https://perma.cc/HB9R-5LQT] ("The [AI] market keeps convulsing, and I wish we could hit the brakes. But we live in a brakeless era.").

[4] 603 U.S. 369, 412 (2024).

Bright overturned the previous *Chevron* paradigm for allocating judicial deference to agency interpretations of ambiguous statutes. The convergence of these dual revolutions presents courts with an unprecedented challenge: how to review agency actions increasingly produced by nonhuman decision-making.

The challenges posed by these dual revolutions are just the beginning. AI systems continue to develop at breakneck speeds and their transformative potential demands that the government begin planning now for its implications. The meteoric rise of large language models (LLMs) promises to reconstitute entire industries, transform institutions, and fundamentally alter the public's relationship with the government.

This phenomenon has spared no one, not even the federal government. The second Trump administration has zealously embraced AI. In contrast to the Biden administration's approach, which emphasized caution and responsibility in the use of AI by government agencies,[5] the Trump administration has adopted an "AI-first strategy" that values quick adoption of AI products at scale[6] with minimal safeguards.[7] Despite their differences, both presidents agreed on the fundamental promise of AI and adopted policies of incorporating this

---

[5] Exec. Order No. 14110 (2023).

[6] *See generally* Michael J. Kratsios et al., *Winning the Race: America's AI Action Plan* (July 2025), https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf [https://perma.cc/U9W5-XHQN]; Exec. Order No. 14179 (2025) (titled *Removing Barriers to American Leadership in Artificial Intelligence*). *See also* Q. Vera Liao & Jennifer Wortman Vaughan, *AI Transparency in the Age of LLMs: A Human-Centered Research Roadmap*, HARV. DATA SCI. REV., 9 (Feb. 29, 2024), https://doi.org/10.1162/99608f92.8036d03b [https://perma.cc/2HN7-KG88] (discussing the "AI arms race" taking place within the private sector); Makena Kelly, *Elon Musk Ally Tells Staff 'AI-First' Is the Future of Key Government Agency*, WIRED (Feb. 3, 2025), https://www.wired.com/story/elon-musk-lieutenant-gsa-ai-agency/ [https://perma.cc/9EM9-6BFF] (discussing the "AI-first strategy" adopted by the Trump administration's DOGE).

[7] *See, e.g.*, Exec. Order No. 14409 (2026) (titled *Promoting Advanced Artificial Intelligence Innovation and Security*) (outlining a "voluntary framework" for federal review of frontier AI systems).

emerging technology to reshape industries and society. This technology is here to stay.[8]

As a result, agencies across the administrative state are deploying AI systems to automate and streamline a wide variety of functions, from determining benefit eligibility,[9] to screening immigrant visa applicants,[10] to reviewing classified government records for publication.[11] From 2023 to 2024, federal agencies' AI use cases grew from 710 to 1,757,[12] and

---

[8] *See* Nestor Maslej et al., *Artificial Intelligence Index Report 2025*, STAN. UNIV. INST. FOR HUM.-CENTERED AI 1, 406-07 (Apr. 2025), https://hai.stanford.edu/assets/files/hai_ai_index_report_2025.pdf [https://perma.cc/UWT9-EWDL] ("60% of [global] respondents believe AI is likely to change how they do their job in the next five years and 36%, or more than one in three, believe that AI is likely to replace their current job in the next five years . . . in 2024, 67% of Gen Z compared to 49% of boomers agree with the statement that AI will likely affect their current jobs.").

[9] *See* Kevin De Liban, *Inescapable AI: The Ways AI Decides How Low-Income People Work, Live, Learn, and Survive*, TECHTONIC JUSTICE 1, 30-42 (Nov. 2024), https://static1.squarespace.com/static/65a1d3be4690143890f61cec/t/673c7170a0d09777066c6e50/1732014450563/ttj-inescapable-ai.pdf [https://perma.cc/K2N8-BD7C] (explaining agency staff are using AI to administer Medicaid).

[10] Dan Maranci, *Government Use of AI and Technology and the Implications for Immigration Law Practitioners*, AILA Doc. No. 25052900 (May 28, 2025), https://www.aila.org/library/government-use-of-ai-and-technology-and-the-implications-for-immigration-law-practitioners [https://perma.cc/Z4ET-XLWQ] (describing the use of AI by immigration enforcement officials).

[11] Miranda Nazzaro, *Gabbard: AI used to decide continued JFK files' classification*, THE HILL (June 11, 2025), https://thehill.com/policy/technology/5344158-trump-administration-ai-jfk-files/ [https://perma.cc/G5PW-6LBY].

[12] Madison Adler, *Federal Government Discloses More than 1,700 AI Use Cases*, FEDSCOOP (Dec. 18, 2024), https://fedscoop.com/federal-government-discloses-more-than-1700-ai-use-cases/ [https://perma.cc/R3PR-H568]; *see also* Madison Adler, *From Translation to Email Drafting, State Department Turns to AI to Assist Workforce*, FEDSCOOP (Dec. 11, 2024), https://fedscoop.com/state-department-ai-chatbot-email-drafting-northstar-famsearch/ [https://perma.cc/C2GX-AGHV] (reporting that the U.S. State Department launched a campaign of incorporating various AI technologies into its workforce procedures in 2024).

then to 3,611 in 2025.[13] This number will almost surely be higher by the time of this article's publication.

While AI revolutionizes administrative practice, judicial doctrine lags behind. Courts have yet to consider how they should review government decisions made wholly or in part with AI systems. This article addresses that gap by proposing a framework for judicial review of AI-assisted agency actions. Judicial deference should be inversely proportional to AI involvement in agency decision-making.[14] As agencies delegate

---

[13] *2025 Federal Agency AI Use Case Inventory*, https://github.com/ombegov/2025-Federal-Agency-AI-Use-Case-Inventory [https://perma.cc/7P2C-GNQU].

[14] The Court, in *Loper Bright v. Raimondo*, distinguished between "deference," which it equated with *Chevron*, and "respect," which courts have long afforded to agency decisions. Loper Bright Enters. v. Raimondo, 603 U.S. 369, 372-73 (2024); *see also* Chavez v. Bondi, 134 F.4th 207, 221 (4th Cir. 2025) ("The Supreme Court in *Loper Bright* may have instructed us to review agency interpretation with a keener eye, but it didn't write off respect for settled, consistent, and persuasive precedents."). The language of *Loper Bright* mirrors the Court's language in *Skidmore*, which is traditionally viewed as an alternative deference standard. *See, e.g.*, United States v. Mead Corp., 533 U.S. 218, 228 (2001) ("The fair measure of deference to an agency administering its own statute has been understood to vary with circumstances, and courts have looked to the degree of the agency's care, its consistency, formality, and relative expertness, and to the persuasiveness of the agency's position, *see* Skidmore, *supra*, at 139-140."); Bax v. Drs. Med. Ctr. of Modesto, Inc., 52 F.4th 862, 872 (9th Cir. 2022) ("A proposed regulation is entitled to respect under *Skidmore* if it has the power to persuade. The deference given to an agency action may range from great respect to near indifference, depending on the degree of the agency's care, its consistency, formality, and relative expertness, and the persuasiveness of the agency's position." (internal citations and quotation marks omitted)); *see also* Cary Coglianese & Daniel E. Walters, *The Great Unsettling: Administrative Governance After* Loper Bright, 77 ADMIN. L. REV. 1, 15-20 (2025) (discussing legal debates over the implications of *Loper Bright* and whether the case requires de novo review or a different "brand" of deference, more closely associated with *Skidmore*); Sanne Knudsen, *Sidestepping Substance: How Administrative Law Plays an Outsized Role in Shaping Environmental Policy and Why Recalibration Is Necessary*, 76 ADMIN. L. REV. 519, 578 (2024) ("[W]hile the particular consequences of *Loper Bright* will unfold in the coming years, it is probably safe to say that the reversion to a *Skidmore*-like approach will give lower courts wide

more decision-making authority to AI systems, courts should correspondingly reduce the deference afforded to an agency's decisions. Traditional judicial deference rests on the premise of reasoned decision-making by accountable human experts.[15] AI systems, however, operate as opaque "black boxes" whose decision-making processes remain inscrutable even to their designers.[16] When agencies rely heavily on AI, the human

---

latitude in determining how much respect to afford agencies."). *Contra* Adrian Vermeule, *The Old Regime and the* Loper Bright *"Revolution"*, 2024 SUP. CT. REV. 235, 246-47 (2024) (distinguishing between the deference courts owe to federal agencies under *Chevron* and the "respect" owed to persuasive agency interpretations under *Skidmore*).

[15] Chevron U.S.A., Inc. v. Nat'l Res. Def. Council, Inc., 467 U.S. 837, 865-66 (1984); *see also* Nat'l Cable & Telecomms. Ass'n v. Brand X Internet Servs., 545 U.S. 967, 980 (2005) (discussing Chevron justifications); Emily Hammond Meazell, *Presidential Control, Expertise, and the Deference Dilemma*, 61 DUKE L.J. 1763, 1775 (2012) ("Although courts had long emphasized that agencies are more politically accountable than the judiciary, Chevron tied agency accountability to the president himself, which gives these executive agencies a claim to democratic legitimacy."); John F. Manning, *Constitutional Structure and Judicial Deference to Agency Interpretations of Agency Rules*, 96 COLUM. L. REV. 612, 626 (1996) ("Chevron's reasoning devotes far greater emphasis to the broader assumptions underlying our structure of government. Specifically, the Court emphasized that our constitutional system favors relatively more accountable agencies, and not relatively less accountable courts, as repositories of policymaking discretion.").

[16] *See* Rishi Bommasani et al., *Improving Transparency in AI Language Models: A Holistic Evaluation*, HAI POL'Y & SOC'Y, STANFORD UNIV. 1, 2 (Feb. 2023), https://hai.stanford.edu/assets/files/2023-02/HAI%20Policy%20&%20Society%20Issue%20Brief%20-%20Improving%20Transparency%20in%20AI%20Language%20Models.pdf [https://perma.cc/N749-CXY6] ("Black-box decision-making remains a challenge for policymakers, researchers, company executives, and the public seeking to understand why an AI model is generating a particular output."); Liao & Vaughan, *supra* note 6, at 7 ("In essence, then, *such models can only be probed in a black-box manner,* which may not be sufficient to meet the transparency requirements for stakeholders, and poses challenges for the research community to develop transparency approaches."); *see also* Hjalmar Wijk & Ajeya Cotra, *Brief Independent Investigation of Agents' Behavior, Reasoning, and Collaboration in the OpenAI / Hugging Face Hacking Incident*, METR & REDWOOD RSCH., 26,

judgment and democratic accountability that justify deference may be reduced or absent entirely. If the AI systems, rather than federal employees, are ultimately responsible for the agency's proposed interpretation, then courts should strike down the proposed rule as illegal under the APA.[17] In such cases, there may be no meaningful "agency" to defer to beyond the AI system itself.

The framework proposed here does not require abandoning AI in government but instead offers a nuanced approach that allows for beneficial AI applications while maintaining the human oversight essential to democratic governance. By calibrating deference to the degree of meaningful human involvement in agency decision-making, courts can ensure that the administrative state reaps AI's benefits without sacrificing the accountability that legitimizes governmental power.

This article proceeds in four parts. Part I discusses the Court's historical and current treatment of deference to federal agencies. This section explores early 20th century precedent that informs the current Court's approach to "deference" and culminates in the 2024 decision in *Loper Bright v. Raimondo*, which ended the *Chevron* paradigm.[18] Part II discusses recent technological developments in AI software and their implications for federal agency practice. This section discusses both the promise and potential perils of this software if it is deployed widely and at scale. Part III examines the theoretical underpinnings of the current deferential framework. I argue that the policy justifications at the heart of judicial deference to administrative interpretations, even post-*Loper Bright*, do not support granting deference to administrative interpretations that are the products, either in whole or in significant part, of these black box "thinking" machines. Finally, Part IV proposes a standard by which courts should

---

34, 48-49, 60 (Aug. 26, 2026), https://metr.org/hugging-face-incident-report-aug-2026.pdf [https://perma.cc/2TS4-UUK7] (describing investigators' incomplete access to transcripts detailing multiple AI agents' coordinated hack of Hugging Face).

[17] *See* 5 U.S.C. § 706.

[18] *Loper Bright*, 603 U.S. at 369.

review agency actions that incorporate AI software into their decision-making processes.

The convergence of *Loper Bright*'s doctrinal revolution with AI's technological transformation creates both opportunity and urgency. As agencies increasingly automate their core functions, the question is not whether courts will confront AI-driven government decisions, but how they will review them. The framework developed in this Article provides a path forward—one that adapts judicial review for the algorithmic age while preserving the democratic values that underlie administrative law itself.

I. The End of an Era: *Loper Bright* and the Future of Administrative Deference

A. Early 20th Century Development of the Administrative State

The relationship between courts and administrative agencies that define modern government took shape through decades of institutional conflict and accommodation. The foundations of the modern administrative state can be traced back to the first half of the twentieth century, when an explosion of federal agencies during the New Deal fundamentally transformed the relationship between government and society.[19] In response to the proliferation of new agencies created during the New Deal, courts explored different, sometimes competing, schemes for deferring to agency decisions.

While the Court developed several different considerations in this deference jurisprudence, the most relevant for our purposes are those articulated in *Skidmore v. Swift & Co.,*

[19] United States v. Morton Salt Co., 338 U.S. 632, 644 (1950) ("The Administrative Procedure Act was framed against a background of rapid expansion of the administrative process as a check upon administrators whose zeal might otherwise have carried them to excesses not contemplated in legislation creating their offices."); *see also* Juan Caballero, *Allocating Deference in Shared Administrative Space*, 56 LOY. L.A. L. REV. 883, 890-901 (2023) (discussing the development of the administrative state in the United States).

decided two years prior to the adoption of the Administrative Procedure Act.[20] *Skidmore* articulated a deference framework that incorporated factors to weigh the persuasive value of an agency's interpretation.[21] This list of considerations included the validity of the agency's reasoning, its consistency, and "*all those factors* which give it power to persuade."[22]

These early deference standards reflected an idealized vision of administrative expertise. Courts viewed agency interpretations of statutory ambiguity as apolitical applications of solution-oriented expertise designed to "fix the nation's ills." [23] Early twentieth-century courts understood agency expertise as guiding decisions toward objective outcomes, devoid of political considerations.[24] This belief in singular, objective agency interpretation was reflected in judicial preference for longstanding statutory interpretations and those that were contemporaneously adopted by an agency with the passage of the underlying statute. These deference regimes incentivized agencies to lock in their initial interpretations and limited their ability to adopt subsequent ones.[25]

In 1984, the Supreme Court issued its unanimous opinion

---

[20] Administrative Procedure Act, Pub. L. No. 79-404, 60 Stat. 237 (1946).

[21] Skidmore v. Swift & Co., 323 U.S. 134, 140 (1944) ("The weight of such a judgment in a particular case will depend upon the thoroughness evident in its consideration, the validity of its reasoning, its consistency with earlier and later pronouncements, and all those factors which give it power to persuade, if lacking power to control.").

[22] *Id.* at 140 (emphasis added); *see also* United States v. Mead Corp., 533 U.S. 218, 228 (2001) (articulating "the degree of the agency's care," "formality," and "relative expertness" as additional considerations under *Skidmore*).

[23] *See* Lisa Schultz Bressman, *Procedures as Politics in Administrative Law*, 107 COLUM. L. REV. 1749, 1759-61 (2007).

[24] *See* id.

[25] Antonin Scalia, *Judicial Deference to Administrative Interpretations of Law*, 1989 DUKE L.J. 511, 517 (1989) ("[T]here is no longer any justification for giving 'special' deference to 'long-standing and consistent' agency interpretations of law. That venerable principle made a lot of sense when we assumed that both court and agency were searching for the one, permanent, 'correct' meaning of the statute; it makes no sense when we acknowledge that the agency is free to give the statute whichever of several possible meanings it thinks most conducive to accomplishment of the statutory purpose.").

in *Chevron U.S.A., Inc. v. Natural Resources Defense Council, Inc.*, marking an inflection point in administrative law.[26] Authored by Justice Stevens, this decision distilled decades of scattered deference precedent into a generally applicable formula that would define the field for the next four decades.[27] *Chevron* would become a foundational doctrine that was "virtually synonymous" with administrative law itself.[28]

Under *Chevron*, courts reviewed agency interpretations of their governing statutes by engaging in a two-step analysis. First, courts would examine the statutory language at issue to determine whether it was clear or ambiguous.[29] If the statutory language clearly expressed congressional intent, courts had to give effect to that unambiguous directive; if, however, the statute contained ambiguity, courts would proceed to step two and consider whether the agency's proposed interpretation represented a permissible interpretation of the ambiguous language. If the agency's interpretation satisfied both steps, courts would defer to it regardless of whether it matched the court's preferred reading of the statute.[30]

Chevron deference was built on a legal fiction that statutory ambiguity signals congressional intent to delegate interpretive authority to the implementing agency.[31] Rather than requiring

---

[26] 467 U.S. 837 (1984).

[27] *See* Coglianese & Walters, *supra* note 14, at 2-3 (describing Chevron as "super precedent" that was "virtually synonymous" with the field of administrative law for four decades); Peter M. Shane & Christopher J. Walker, Chevron *at 30: Looking Back and Looking Forward*, 83 FORDHAM L. REV. 475, 475 (2014) (describing *Chevron* as "the most-cited administrative law decision of all time"); Lauren Mattiuzzo, *Most-Cited U.S. Supreme Court Cases in HeinOnline: Part III*, HEINONLINE BLOG (Sept. 26, 2018), https://home.heinonline.org/blog/2018/09/most-cited-u-s-supreme-court-cases-in-heinonline-part-iii [https://perma.cc/5RSU-XRGP] (reporting *Chevron* as the seventh most-cited Supreme Court decision in legal scholarship).

[28] *E.g.,* Coglianese & Walters, *supra* note 14, at 2-3 (using the term "super precedent" to describe *Chevron*).

[29] *Chevron*, 467 U.S. at 842-43.

[30] *Id.*

[31] *See* Kisor v. Wilkie, 588 U.S. 558, 576 (2019) ("[W]e give Auer deference

courts to divine specific congressional intent statute-by-statute, *Chevron* established a general presumption that Congress desired agencies to serve as primary interpreters of ambiguous provisions in their governing statutes.[32] *Chevron* marked a departure from the Court's earlier approach to deference, where it understood political accountability and technical expertise to be contradictory variables when weighing deference to federal agencies.[33]

For four decades, *Chevron* remained the dominant framework. This deference regime would shape not only how

---

because we presume, for a set of reasons relating to the comparative attributes of courts and agencies, that Congress would have wanted us to."); Smiley v. Citibank (S.D.), N.A., 517 U.S. 735, 740-41 (1996) ("We accord deference to agencies under *Chevron* . . . because of a presumption that Congress, when it left ambiguity in a statute meant for implementation by an agency, understood that the ambiguity would be resolved, first and foremost, by the agency, and desired the agency (rather than the courts) to possess whatever degree of discretion the ambiguity allows." (citing *Chevron*, 467 U.S. at 843-44)); Cass R. Sunstein, *Beyond Marbury: The Executive's Power to Say What the Law Is*, 115 YALE L.J. 2580, 2589-90 (2006) (describing the legal presumption of "congressional will" to be a legal fiction that provides the foundation of *Chevron* itself); Scalia, *supra* note 25, at 517 ("[A]ny rule adopted in this field represents merely a fictional, presumed intent, and operates principally as a background rule of law against which Congress can legislate.").

[32] *See* Vermeule, *supra* note 14, at 249 ("The resulting regime represents a retail rather than wholesale version of *Chevron*. Whereas *Chevron* was said to rest on a general presumption that gaps and ambiguities represented a delegation to agencies, judges will now have to decide, statute by statute and problem by problem, whether a *Loper Bright* delegation is the best reading of the statute. But it is also true that *Chevron* was always already a retail regime, in the sense that judges always had to decide, statute by statute and problem by problem, what decision Congress had made about the issue at hand.").

[33] *See, e.g.*, Humphrey's Ex'r v. United States, 295 U.S. 602, 624 (1935) ("The [Federal Trade Commission] is to be nonpartisan; and it must, from the very nature of its duties, act with entire impartiality. It is charged with the enforcement of no policy except the policy of the law. Its duties are neither political nor executive, but predominantly quasi-judicial and quasi-legislative."); Elena Kagan, *Presidential Administration*, 114 HARV. L. REV. 2245, 2261 (2001) (discussing the writing of James Landis, principal spokesman for the idea that expertise was the primary justification for federal agencies' "enhanced bureaucratic power" during the New Deal Era).

courts reviewed agency action, but also how agencies approached their own interpretive responsibilities, knowing that reasonable interpretations of ambiguous statutes would likely receive judicial deference.[34]

B. Relentless Attacks: Loper Bright *and* Chevron*'s Demise*

While *Chevron* defined the deference paradigm for four decades, its eventual demise was long foretold. Almost since its inception, *Chevron* received continuous and protracted criticism from scholars and jurists who questioned its political and constitutional legitimacy.[35] For years, the Court was reluctant to cite *Chevron* directly,[36] and when it did discuss the precedent, it was usually to impose new limitations or carve out

---

[34] *See* Loper Bright Enters. v. Raimondo, 603 U.S. 369, 411 (2024) ("Chevron thus allows agencies to change course even when Congress has given them no power to do so. By its sheer breadth, Chevron fosters unwarranted instability in the law, leaving those attempting to plan around agency action in an eternal fog of uncertainty."); *see also* Coglianese & Walters, *supra* note 14, at 25-26 ("Although judicial review is never an inevitability, it undoubtedly looms over much agency decision-making and action. Agencies always undertake action against the potential backdrop of judicial review.").

[35] *See* Richard J. Pierce, Jr., *Democratizing the Administrative State*, 48 WM. & MARY L. REV. 559, 562 n.1 (2006) (cataloging law review articles wherein scholars question the political and constitutional legitimacy of the administrative state); Kristin E. Hickman & Aaron L. Nielson, *Narrowing Chevron's Domain*, 70 DUKE L.J. 931, 954 (2021) (cataloging critiques of the implicit delegation justification of *Chevron* deference); Manning, *supra* note 15; Connor N. Raso & William N. Eskridge, Jr., Chevron *as a Canon, Not a Precedent: An Empirical Study of What Motivates Justices in Agency Deference Cases*, 110 COLUM. L. REV. 1727 (2010) (critiquing jurists' inconsistent application of *Chevron* in different contexts).

[36] *See* Isaiah McKinney, *The Chevron Ball Ended at Midnight, but the Circuits Are Still Two-Stepping by Themselves*, YALE J. ON REG., NOTICE & COMMENT (Dec. 18, 2022), https://www.yalejreg.com/nc/chevron-ended/ [https://perma.cc/E2DN-UPJZ] ("Although the Supreme Court has not overruled *Chevron*, it has not deferred under the doctrine at Step Two in six years...Since *Michigan v. EPA* and *King v. Burwell*, the Supreme Court has analyzed ten cases under *Chevron*. In only one of those, *Cuozzo Speed Techs., LLC v. Lee*, did the Court continue to Step Two. All the other cases were decided under Step One or under an exception, such as *United States v. Mead*.").

exceptions.[37] A change in Supreme Court personnel during the first Trump administration all but assured *Chevron*'s end.[38]

After years of direct and indirect challenges, the Supreme Court finally overturned *Chevron* in *Loper Bright Enterprises v. Raimondo* in 2024.[39] Writing for the majority, Chief Justice Roberts adopted many longstanding critiques of *Chevron*. The opinion began by detailing the legal history surrounding the 20th century development of the administrative state and the passage of the Administrative Procedure Act.[40] Roberts used this framing to contend that *Chevron* constituted an illegal delegation of authority from the courts to the federal agencies.[41] According to Roberts, *Chevron* rested on a flawed assumption that Congress delegates interpretive authority to agencies whenever a statute is ambiguous and thereby violated the judiciary's constitutionally mandated role.[42] Chief Justice Roberts' opinion explicitly overruled *Chevron*, vacated the lower courts' opinions, and remanded the case for decision using the court's "independent judgment."[43]

---

[37] *E.g.*, King v. Burwell, 576 U.S. 473, 485 (2015) (applying the major issues exception to avoid applying *Chevron*); Kisor v. Wilkie, 588 U.S. 558, 580 (2019) (modifying *Auer* deference, the deference afforded to agency interpretations of their own regulations); *see also Loper Bright*, 603 U.S. at 375 (describing "the Court's constant tinkering with and eventual turn away from *Chevron*" as justification for overturning the 40-year precedent).

[38] *See* Shoba Sivaprasad Wadhia & Christopher J. Walker, *The Case Against Chevron Deference in Immigration Adjudication*, 70 DUKE L.J. 1197, 1199-200 (2021) ("[T]he potential demise of *Chevron* deference was a core talking point against Gorsuch's elevation to the Supreme Court."); Ann Graham, *Searching for* Chevron *in Muddy* Watters*: The Roberts Court and Judicial Review of Agency Regulations*, 60 ADMIN. L. REV. 229, 271 (2008) ("Analysis of the Roberts Court cases shows that classic *Chevron* analysis is dead—or at least critically wounded. Unfortunately, it appears that stray bullets, in the form of inconsistent applications of the doctrine, may have done it in.").

[39] *Loper Bright*, 603 U.S. at 369. *Loper Bright* was merged with a similar case arising in the U.S. Court of Appeals for the First Circuit titled *Relentless, Inc. v. Department of Commerce*. *See* Relentless, Inc. v. Dep't of Com., 144 S. Ct. 325 (2023) (granting certiorari).

[40] *Loper Bright*, 603 U.S. at 394-95.

[41] *Id*. at 398-99.

[42] *Id.* at 384-96.

[43] *Id.* at 412.

The Court also rejected *Chevron*'s reasoning that agency policymaking in light of statutory ambiguity is inherently political. Chief Justice Roberts wrote, "[i]t is reasonable to assume that Congress intends to leave policymaking to political actors. But resolution of statutory ambiguities involves legal interpretation. That task does not suddenly become policymaking just because a court has an 'agency to fall back on.'"[44]

Thus, after forty years, *Loper Bright* became the vessel that sank *Chevron*.

1. What Remains: "Respect" versus "Deference"

While *Loper Bright* marked the end of the *Chevron* deference paradigm, it did not eliminate judicial deference to agency interpretations entirely.[45] The opinion acknowledges that Congress retains authority to "specifically" grant agencies broad interpretive authority when it does so explicitly.[46] Thus,

---

[44] *Id.* at 403 (citing Kisor v. Wilkie, 588 U.S. 558, 575 (2019)).

[45] *See, e.g.,* Texas v. U.S. Env't Prot. Agency, 137 F.4th 353, 365 (5th Cir. 2025) ("To be clear, discarding *Chevron* deference does not mean ignoring agency interpretations. The Supreme Court discussed in considerable detail its 1944 precedent that to a large but not complete extent *Chevron* displaced in 1984. The *Loper Bright* Court agreed with *Skidmore* that an agency's expertise has always been one of the factors which may give an Executive Branch interpretation particular 'power to persuade, if lacking power to control.' It is not necessary to give *Chevron* deference in order 'to ensure that the resolution of statutory ambiguities is well informed by [an agency's] subject matter expertise. . . . Congress expects courts to do their ordinary job of interpreting statutes, *with due respect* for the views of the Executive Branch.'" (internal citations and quotations omitted)); United States v. Poore, No. 22-3154, 2025 WL 1201946, at *3 (7th Cir. Apr. 25, 2025) (noting that *Loper Bright* did not purport to overrule other court precedent discussing deference to agency interpretations of their own regulations).

[46] *Loper Bright*, 603 U.S. at 394-95 ("In a case involving an agency, of course, the statute's meaning may well be that the agency is authorized to exercise a degree of discretion. Congress has often enacted such statutes. For example, some statutes 'expressly delegate[ ]' to an agency the authority to give meaning to a particular statutory term . . . Others empower an agency to prescribe rules to 'fill up the details' of a statutory scheme, . . . or to regulate subject to the limits imposed by a term or phrase that 'leaves

explicit delegations are still permitted, but *implicit* congressional delegations in the form of ambiguous statutes no longer justify the level of deference accorded under *Chevron*.

In situations where agencies are not granted explicit congressional delegations of authority, however, *Loper Bright* still recognizes a type of judicial "respect" for agency interpretations.[47] While *Loper Bright* frames this "respect" as an alternative to "deference,"[48] it appears to be a distinction without a difference. To illustrate what the Court meant by "respect," Chief Justice Roberts cites *Skidmore v. Swift & Co.*, a pre-*Chevron* form of deference but a form of deference nonetheless.[49] Thus, the opinion does little more than supplant *Chevron* with a repackaged *Skidmore* deference under the guise of affording "respect" to agency interpretations; as both a practical and precedential matter, courts will continue to defer to agencies. According to Chief Justice Roberts, "respect," unlike *Chevron* deference, functions as an

---

agencies with flexibility,' . . . such as 'appropriate' or 'reasonable.' When the best reading of a statute is that it delegates discretionary authority to an agency, the role of the reviewing court under the APA is, as always, to independently interpret the statute and effectuate the will of Congress subject to constitutional limits. The court fulfills that role by recognizing constitutional delegations").

[47] *Id.* at 412-13 ("Careful attention to the judgment of the Executive Branch may help inform that inquiry . . . . But courts need not and under the APA may not *defer* to an agency interpretation of the law simply because a statute is ambiguous").

[48] *Id.* at 399 (describing "respect" as consistent with courts' exercise of independent judgment, unlike "deference").

[49] *Id.* (citing Skidmore v. Swift & Co., 323 U.S. 134, 140 (1944)). Here, *Skidmore* is used as an example of the "respect" historically afforded to agency interpretations. Despite Chief Justice Roberts' turn of phrase, scholars and jurists have long recognized *Skidmore* as constituting a *type* of deference that is less generous than *Chevron*. *See, e.g.*, *Loper Bright*, 603 U.S. at 476 (Kagan, J., dissenting) ("[T]he majority makes clear that what is usually called Skidmore deference continues to apply"); United States v. Mead Corp., 533 U.S. 218, 239 (2001) (criticizing the majority opinion's decision to apply "so-called Skidmore deference" in select circumstances); *see also* William N. Eskridge, Jr. & Lauren E. Baer, *The Continuum of Deference: Supreme Court Treatment of Agency Statutory Interpretations from Chevron to Hamdan*, 96 GEO. L.J. 1083, 1109-11 (2008) (describing *Skidmore* as a node on a continuum of deference).

"interpretive aid" like others the judiciary deploys.[50] Chief Justice Roberts also noted that this "respect" was especially warranted "when an Executive Branch interpretation was issued roughly contemporaneously with enactment of the statute and remained consistent over time."[51]

2. An Uncertain Landscape Post-*Loper Bright*

The immediate effects of *Loper Bright* have yet to be fully settled. Chief Justice Roberts' retention of *Skidmore* deference in an opinion demanding that courts engage in independent judgment presents a "Rorschach test" for jurists.[52] *Loper Bright* reasons that anything short of a court's "independent judgment" is inconsistent with Section 706 of the APA,[53] but it also retains *Skidmore* deference as the default deferential standard going forward.[54]

Loper Bright's attempt to distinguish deference from the "respect" that must now be accorded to agency opinions is confusing considering the Court's prior precedent. While *Chevron* was the dominant form of deference for four decades, the Court previously revived and incorporated *Skidmore*

---

[50] *Loper Bright*, 603 U.S. at 386.

[51] *Id.* (citing United States v. Ala. Great S. R. Co., 142 U.S. 615, 621 (1892) and Nat'l Lead Co. v. United States, 252 U.S. 140, 145-146 (1920)).

[52] Coglianese & Walters, *supra* note 14, at 5 (describing *Loper Bright* as "a Rorschach test inside a crystal ball"); Knudsen, *supra* note 14, at 583-84 ("After *Loper Bright*, which instructs lower courts to revert to a *Skidmore*-style review, there will undoubtedly be another surge of discussions and cases trying to find the right balance of flexibility and structure."). *Compare* Chavez v. Bondi, 134 F.4th 207, 221 (4th Cir. 2025) ("The Supreme Court in *Loper Bright* may have instructed us to review agency interpretation with a keener eye, but it didn't write off respect for settled, consistent, and persuasive precedents.") *with* Miller Plastic Prods., Inc. v. NLRB, 141 F.4th 492, 503 (3d Cir. 2025) (considering the NLRB's "body of experience and informed judgment" as a persuasive factor in its review of the statutory interpretation).

[53] *Loper Bright*, 603 U.S. at 391 (quoting 5 U.S.C. § 706).

[54] *Id.* at 394 ("In exercising such judgment, though, courts may—as they have from the start—seek aid from the interpretations of those responsible for implementing particular statutes. Such interpretations 'constitute a body of experience and informed judgment to which courts and litigants may properly resort for guidance' consistent with the APA." (quoting *Skidmore*, 323 U.S. at 140)).

deference into this dominant paradigm.[55] In *Mead*, the Court created a threshold step to the *Chevron* analysis wherein courts considered whether the agency acted with the force of law. Acts that carried such weight would undergo the *Chevron* two-step analysis, while agency actions that did not were afforded *Skidmore* deference.[56] Thus, *Mead* modified the *Chevron* analysis by imposing what became known as *Chevron* "step zero"—a threshold inquiry into whether *Chevron* deference applied in the first place.[57] Agency actions that failed this threshold inquiry became subject to the less deferential *Skidmore* standard. While *Skidmore* deference differed substantively from *Chevron* in application, both precedents advanced a similar proposition: "courts were not required or expected to engage in fully independent, or *de novo*, review of agency interpretations."[58]

Prior to *Chevron*, administrative deference was premised primarily on the consistency of an agency's interpretation and application of its governing statutes over time.[59] *Chevron* rejected this foundational premise, though, reasoning instead that interpreting ambiguous statutory provisions required agencies to make policy decisions that were inherently

---

[55] United States v. Mead Corp., 533 U.S. 218, 226-27 (2001).

[56] *Id.*

[57] *E.g.*, Cass R. Sunstein, *Chevron Step Zero*, 92 VA. L. REV. 187, 191 (2006) (describing "initial inquiry into whether the Chevron framework applies at all" as "Step Zero").

[58] Coglianese & Walters, *supra* note 14, at 11-12 ("Although Skidmore's language is in some ways similar to Chevron's, in practice the case came to stand for a conceptually distinct, nonbinding form of deference that could apply in cases where Chevron deference did not apply. Whether agency interpretations went down the *Chevron* track or the *Skidmore* track, one thing was perfectly clear: when deferring to agencies, courts were not required or expected to engage in fully independent, or de novo, review of agency interpretations of their ambiguous statutes."); *see also* Kristin E. Hickman & Matthew D. Krueger, *In Search of the Modern* Skidmore *Standard*, 107 COLUM. L. REV. 1235 (2007) (showing that most courts understood Skidmore as endorsing "sliding scale" deference rather than independent judgment).

[59] *See, e.g.,* Immigr. & Naturalization Servs. v. Cardoza-Fonseca, 480 U.S. 421, 446 n.30 (1987) (stating that pre-*Chevron* deference depended on consistent application of regulations or statutes over time to legitimize the agency interpretation) (quoting Watt v. Alaska, 451 U.S. 259, 273 (1981)).

political. [60] Under this reasoning, agencies were better positioned than courts to make such determinations because they, unlike the judiciary, operated within the political branches.[61] By recognizing the possibility of multiple "right" legal conclusions, courts viewed an agency's flexibility in light of newly discovered facts and shifting policy preferences as a beneficial mechanism of statutory implementation.[62]

The evolution in judicial justification for deference is reflected in *Mead*'s expansion of relevant considerations. [63] While the original *Skidmore* decision afforded deference based on effectiveness of agency reasoning and consistency with other pronouncements, the 2001 *Mead* opinion expanded the list of factors to include "the merit of [the] writer's thoroughness, logic, and expertness, its fit with prior interpretations, and any other sources of weight."[64] Agency expertise, while briefly mentioned in the original *Skidmore* decision, [65] became a central consideration in *Skidmore's* revival post-*Chevron*. While *Loper Bright* rejects the

---

[60] Chevron U.S.A., Inc. v. Nat'l Res. Def. Council, Inc., 467 U.S. 837, 865 (1984) ("Judges are not experts in the field, and are not part of either political branch of the Government. Courts must, in some cases, reconcile competing political interests, but not on the basis of the judges' personal policy preferences.").

[61] *Id.* at 865-66; *see also* Nat'l Cable & Telecomms. Ass'n v. Brand X Internet Servs., 545 U.S. 967, 980 (2005) (discussing Chevron justifications); Meazell, *supra* note 15, at 1775.

[62] *See* Russell L. Weaver, *Deference to Regulatory Interpretations: Inter-Agency Conflicts*, 43 ALA. L. REV. 35, 54-55 (1991) ("As an agency gains experience with a regulatory scheme, its understanding of that scheme increases. This understanding can alter the agency's view of regulatory problems and affect its interpretations. Thus, even though an agency may have interpreted a regulation one way when the regulation was promulgated, it might interpret that same regulation in a different way at a later date."); Scalia, *supra* note 25, at 519.

[63] Skidmore v. Swift & Co., 323 U.S. 134, 139 (1944); *see also* Midship Pipeline Co. v. FERC, 45 F.4th 867, 875 (5th Cir. 2022) ("[P]ersuasiveness is the touchstone in determining how much to defer to an agency interpretation.").

[64] United States v. Mead Corp., 533 U.S. 218, 220 (2001); *Skidmore*, 323 U.S. at 140.

[65] *Skidmore*, 323 U.S. at 137, 139-40 (referencing agency "experience" as justifying deference).

generalized presumption that agency expertise justifies wholesale deference to their interpretations as it had under *Chevron*, Roberts' opinion recognizes that an agency's subject-matter expertise remains one among many factors courts weigh in their deference analysis.[66]

Loper Bright signaled a return to pre-*Chevron* justifications for deference—namely consistency and contemporaneity.[67] Where *Chevron* incentivized agency flexibility and political responsiveness by recognizing that policymaking was informed by political realities, *Loper Bright* incentivizes agencies to maintain consistent interpretations of statutes. This favoring of longstanding interpretations risks ossification of agency interpretations by disincentivizing agencies from revising their understanding as circumstances evolve.[68] Ultimately, *Loper Bright* empowers courts while constraining agencies. Unlike *Chevron*, which provided judges a structured—albeit imperfect—analytical framework for allocating deference, the *Skidmore* paradigm that re-emerged in *Loper Bright* is vague and open-ended, affording courts wide latitude to determine whether and how much "respect" to afford agency

---

[66] Loper Bright Enters. v. Raimondo, 603 U.S. 369, 401-02 (2024) ("[W]e recently noted that interpretive issues arising in connection with a regulatory scheme often 'may fall more naturally into a judge's bailiwick' than an agency's. We thus observed that '[w]hen the agency has no comparative expertise in resolving a regulatory ambiguity, Congress presumably would not grant it that authority.' *Chevron*'s broad rule of deference, though, demands that courts presume just the opposite. Under that rule, ambiguities of all stripes trigger deference." (internal citations omitted)).

[67] *Id.* at 385-86 (describing the history of judicial deference to agency interpretations).

[68] *See* Nat'l Cable & Telecomms. Ass'n v. Brand X Internet Servs., 545 U.S. 967, 983 (2005) (holding that limiting agency discretion to interpret ambiguous statutes based on stare decisis would "lead to the ossification of large portions of our statutory law" (quoting United States v. Mead Corp., 533 U.S. 218, 247 (2001) (Scalia, J., dissenting))); *see also* Scalia, *supra* note 25, at 517 ("One of the major disadvantages of having the courts resolve ambiguities is that they resolve them for ever and ever; only statutory amendment can produce a change.").

interpretations.[69]

II. The AI Revolution in Administrative Governance

A. Artificial Intelligence: Technology, Definitions, and Current State

"AI" has no universally recognized definition.[70] The term "AI" is used as an umbrella concept encompassing a wide range of technologies and applications, from simple automated decision trees to sophisticated LLMs.[71] One technical definition describes AI as "the study and development of intelligent machines and software that can reason, learn, gather knowledge, communicate, manipulate and perceive objects."[72]

---

[69] Knudsen, *supra* note 14, at 578 ("To that end, while the particular consequences of *Loper Bright* will unfold in the coming years, it is probably safe to say that the reversion to a *Skidmore*-like approach will give lower courts wide latitude in determining how much respect to afford agencies; and the unstructured nature of *Skidmore* respect will give rise to concerns about judicial review that is prone to manipulation and marked by inconsistency.").

[70] USPTO, PUBLIC VIEWS ON ARTIFICIAL INTELLIGENCE AND INTELLECTUAL PROPERTY POLICY i, ii (Oct. 2020), https://www.uspto.gov/sites/default/files/documents/USPTO_AI-Report_2020-10-07.pdf [https://perma.cc/3Q3A-GEXP] (noting that one "General Theme" present in the nearly 200 comments it received is that there is no universally recognized definition of "AI").

[71] *Ethics of Artificial Intelligence and Robotics*, STAN. ENCYC. OF PHIL. (Apr. 30, 2020), https://plato.stanford.edu/entries/ethics-ai/ [https://perma.cc/V5BS-EFNV] (defining AI as "any kind of artificial computational system that shows intelligent behaviour, i.e., complex behaviour that is conducive to reaching goals. In particular, we do not wish to restrict 'intelligence' to what would require intelligence if done by *humans*, as Minsky had suggested (1985). This means we incorporate a range of machines, including those in 'technical AI', that show only limited abilities in learning or reasoning but excel at the automation of particular tasks, as well as machines in 'general AI' that aim to create a generally intelligent agent").

[72] Avneet Pannu, *Artificial Intelligence and Its Application in Different Areas*, 4 INT'L J. ENG'G & INNOVATIVE TECH. (IJEIT) 79 (2015); *see also Explanatory Memorandum on the Updated OECD Definition of an AI*

Current AI models demonstrate remarkable versatility, capable of applications ranging from question answering to object recognition. [73] Frequently, numerous disparate technologies are grouped together under the AI moniker, creating definitional challenges for legal analysis.[74] The key

---

System, OECD PUBL'G., 4 (Mar. 2024), https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/03/explanatory-memorandum-on-the-updated-oecd-definition-of-an-ai-system_3c815e51/623da898-en.pdf [https://perma.cc/GB2F-CS2C] (explaining the OECD's updated definition of an AI system).

[73] Pannu, *supra* note 72, at 79 ("AI technologies have matured to the point in offering real practical benefits in many of their applications. Major Artificial Intelligence areas are Expert Systems, Natural Language Processing, Speech Understanding, Robotics and Sensory Systems, Computer Vision and Scene Recognition, Intelligent Computer-Aided Instruction, Neural Computing. From these Expert System is a rapidly growing technology which is having a huge impact on various fields of life. The various techniques applied in artificial intelligence are Neural Network, Fuzzy Logic, Evolutionary Computing, and Hybrid Artificial Intelligence.").

[74] NAT'L SEC. COMM'N ON A.I., FINAL REPORT 32 (2021), https://assets.foleon.com/eu-central-1/de-uploads-7e3kk3/48187/nscai_full_report_digital.04d6b124173c.pdf [https://perma.cc/638B-D5F5] ("Artificial Intelligence (AI) is not a single piece of hardware or software, but rather a constellation of technologies."); Xavier Rodriguez, *Artificial Intelligence (AI) and the Practice of Law*, 24 SEDONA CONF. J. 783, 788-89 (2023) ("AI applications employ algorithmic models that receive and process large amounts of data and are trained to recognize patterns, thus enabling the applications to automate repetitive functions as well as make judgments and predictions. Machine learning is a subset of AI . . . Generative AI is [another] specific subset of AI used to create new content based on training on existing data taken from massive data sources in response to a user's prompt, or to replicate a style used as input." (quotations omitted)); William D. Eggers et al., *AI-Augmented Government: Using Cognitive Technologies to Redesign Public Sector Work*, DELOITTE INSIGHTS (Apr. 26, 2017), https://www.deloitte.com/us/en/insights/topics/emerging-technologies/artificial-intelligence-government.html [https://perma.cc/W73R-P7RT] ("AI-based technologies include machine learning, computer vision, speech recognition, natural language processing, and robotics."); *Ethics of Artificial Intelligence and Robotics*, *supra* note 71, at § 2.3, https://plato.stanford.edu/entries/ethics-ai/

feature across most definitions is automated decision-making or recommendations based on external data.[75]

Much of the technology that is collectively referred to as "AI" is not new. Many of the architectures and training techniques underlying contemporary AI systems have existed in some form for decades.[76] What distinguishes the current generation of AI systems is the unprecedented scale and scope

---

[https://perma.cc/A2SV-R3PX] ("Many AI systems rely on machine learning techniques in (simulated) neural networks that will extract patterns from a given dataset, with or without 'correct' solutions provided; i.e., supervised, semi-supervised or unsupervised. With these techniques, the 'learning' captures patterns in the data and these are labelled in a way that appears useful to the decision the system makes, while the programmer does not really know *which* patterns in the data the system has used."); Cristian Gonzalez, *A New Blue Sky: SEC Considerations in the Regulation of Autonomous AI Misconduct*, 10 ALR ACCORD 33, 38-39 (2025) ("GenAI is a type of AI that allows users to input prompts and data to generate new outputs like text, images, code, and other types of data. GenAI also learns as it is trained on more data. Similarly, machine learning is a subset of AI that involves training algorithms to recognize patterns and make decisions based on data; however, this is different from GenAI because machine learning does not create new, original data.").

[75] De Liban, *supra* note 9, at 18; 15 U.S.C. § 9401(3) (defining "Artificial Intelligence Systems").

[76] Maslej et al., *supra* note 8, at 46-58 (detailing findings of "notable AI models" recorded by Epoch AI, an AI Index data compiler, that has compiled a database of 900 AI models released since the 1950s); NAT'L SEC. COMM'N ON A.I., *supra* note 74, at 32-33 (describing "first wave," "second wave," and current iterations of AI technology spanning decades); Rishi Bommasani et al., *On the Opportunities and Risks of Foundation Models*, CTR. FOR RSCH. ON FOUND. MODELS (CRFM), STAN. UNIV. INST. FOR HUM.-CENTERED AI 1, 3-4, 22 (2021), https://crfm.stanford.edu/assets/report.pdf [https://perma.cc/EWN5-HLAJ] (describing the history of "machine learning," "self-supervised learning," "deep learning," and "[n]atural language processing"); Jens Ludwig & Sendhil Mullainathan, *Fragile Algorithms and Fallible Decision-Makers: Lessons from the Justice System*, NAT'L BUREAU OF ECON. RESEARCH, Working Paper No. 29267 (2021), http://www.nber.org/papers/w29267 [https://perma.cc/2CRK-M35Q] (discussing "algorithms" in both the general sense and in reference to specific technology used by AI and machine learning which the authors argue have been "widely used in the criminal justice system" prior to this article); *Ethics of Artificial Intelligence and Robotics*, *supra* note 71.

of tasks they can perform.[77] For instance, present-day LLMs function by predicting the probability of each token in a sequence within its surrounding context.[78] These tokens – defined as the units LLMs use to process text[79] – might consist of individual characters, words, or phrases.[80] The LLM's "neural networks" make calculations while conforming to billions or trillions of parameters developed through the model's training.[81] While LLMs from a few years prior were trained on moderate-sized datasets, current LLMs utilize massive datasets containing hundreds of billions or even trillions of tokens.[82]

These vast training datasets enable an LLM's heightened sophistication and expressiveness compared to predecessors, though they also present transparency challenges. AI developers typically do not disclose data sources used to train the technology, thus limiting transparency and public insight into their training methodologies and the model's overall functioning.[83] While some base LLMs have broad capabilities, they can be tailored to exhibit specific behaviors or excel at particular tasks like text summarization, question answering, or

---

[77] Bommasani et al., *supra* note 76; *see also* Jens Ludwig et al., *Large Language Models: An Applied Econometric Framework*, Nat'l Bureau of Econ. Research, Working Paper 33344, at 5 (Jan. 2025), http://www.nber.org/papers/w33344 [https://perma.cc/3E8J-44KU] ("The central difference between LLMs and supervised learning algorithms is that LLMs aspire to serve as useful tools not just for a single task, but all tasks — to be a general-purpose technology (GPT). The scope of that ambition is both the most remarkable feature of LLMs and also their Achilles' heel."); Gonzalez, supra note 74, at 38 ("[A] multimodal large language model (MMLLM) is an LLM that can process and generate many different types of data, for example, text, audio, and graphics, creating complex outputs such as articles with images and text, video captioning, audio narrations for videos, and graphical data analysis.").

[78] Liao & Vaughan, *supra* note 6.

[79] *See, e.g.*, *Understanding and Counting Tokens*, OPENAI, https://help.openai.com/en/articles/4936856-understanding-and-counting-tokens [https://perma.cc/PB3J-WTZ6].

[80] Liao & Vaughan, *supra* note 6.

[81] *Id.* at 5.

[82] Liao & Vaughan, *supra* note 6.

[83] *Id.* at 6.

code generation.[84] Previously, each individual task would likely have required bespoke features tailored by software engineering teams for specific applications.[85]

B. The Proliferation of AI in Federal Agencies

Many LLM developers promote techno-utopian visions of AI's potential.[86] These visions predict a future where AI systems serve as multipliers of human ingenuity and creativity, ushering in an age of economic prosperity and scientific discovery.[87] AI developers, like OpenAI, anticipate that such a

---

[84] *Id.* at 5; s*ee also* Bommasani et al., *supra* note 76, at 4 (describing popular fine-tuning approaches).

[85] Bommasani et al., *supra* note 76, at 4.

[86] *See, e.g.*, Mark Zuckerberg, *The Future is for Everyone*, META (Aug. 10, 2026), https://www.meta.com/thefutureisforeveryone/ [https://perma.cc/JRM6-AAZS] (describing AI as yet another "transformative advance" that will create "the abundance of the future" that "can be shared by everyone"); Nick Lichtenberg, *Google's Nobel-winning AI leader sees a 'renaissance' ahead—after a 10- or 15-year shakeout*, FORTUNE (Feb. 11, 2026), https://fortune.com/2026/02/11/demis-hassabis-nobel-google-deepmind-predicts-ai-renaissance-radical-abundance/ [https://perma.cc/E74J-JS5D] (quoting Sir Demis Hassabis, CEO of Google DeepMind, describing an AI powered "new golden era of discovery" where the technology will lead to "radical abundance"); Amodei, *supra* note 1 (articulating "positive applications" for AI in the areas of biology and physical health; neuroscience and mental health; economic development and poverty; peace and governance; and work and employment); *Planning for AGI and Beyond*, OPENAI (Feb. 24, 2023), https://openai.com/index/planning-for-agi-and-beyond/ [https://perma.cc/F77J-ZD8S] ("If AGI is successfully created, this technology could help us elevate humanity by increasing abundance, turbocharging the global economy, and aiding in the discovery of new scientific knowledge that changes the limits of possibility.").

[87] Lichtenberg, *supra* note 86 ("In this near future, [Sir Demis Hassabis] predicted that 'medicine won't look like it does today,' with AI enabling personalized treatments and curing major diseases. Beyond health, he said he foresees AI unlocking new materials to solve the energy crisis through fusion or solar breakthroughs, eventually allowing humanity to 'travel the stars and . . . explore the galaxy.'"); Amodei, *supra* note 1 (articulating "positive applications" for AI in the areas of biology and physical health; neuroscience and mental health; economic development and poverty; peace and governance; and work and employment); *Planning for AGI and*

future will be achieved through iterative cycles of model development, deployment, and revision as society incorporates these systems into various functions.[88] This process is arguably already underway as developers release increasingly powerful models at accelerating rates.[89] These models, in turn, have become increasingly embedded in everyday life.

Industries are rapidly seeking to leverage AI technologies for increased efficiencies across virtually every sector of the economy.[90] This widespread adoption, including within federal

---

Beyond, *supra* note 86 ("AGI has the potential to give everyone incredible new capabilities; we can imagine a world where all of us have access to help with almost any cognitive task, providing a great force multiplier for human ingenuity and creativity.").

[88] Zuckerberg, *supra* note 86 ("Invention, not automation, will be the greatest contribution of superintelligence. Early AI could answer questions and do routine work. Soon it will increasingly help discover new knowledge -- ranging from discovering new drugs to cure a family member's disease to finding new ways to improve your business. While the number of questions a person can ask in a day is limited, the number of valuable things superintelligence can invent to help achieve your goals is unlimited.").

[89] *See* Maslej et al., *supra* note 8, at 27 ("New research finds that the training compute for notable AI models doubles approximately every five months, dataset sizes for training LLMs every eight months, and the power required for training annually."); *see also id.* at 46 (fig. 1.3.2 showing a marked increase in "notable AI models" released between 2003 and 2024).

[90] *See, e.g., Legislative Use of Artificial Intelligence 2026 Survey*, NAT'L CONF. OF STATE LEGISLATURES (Aug. 10, 2026), https://www.ncsl.org/center-for-legislative-strengthening/legislative-use-of-artificial-intelligence-2026-survey [https://perma.cc/KS5U-W7WV] (describing AI adoption amongst legislative staff); Joel Becker, *Measuring the Self-Reported Impact of Early-2026 AI on Technical Worker Productivity*, METR (May 11, 2026), https://metr.org/blog/2026-05-11-ai-usage-survey/ [https://perma.cc/D8RL-PNFS] (describing AI adoption amongst technical workers—software engineers, researchers, academics, and founders/managers); Jeffrey S. Allen, *Monitoring AI Adoption in the US Economy*, BD. OF GOVERNORS OF THE FED. RSRV. SYS., FEDS NOTES (Apr. 3, 2026), https://www.federalreserve.gov/econres/notes/feds-notes/monitoring-ai-adoption-in-the-u-s-economy-20260403.html [https://perma.cc/P26N-5G5C] (surveying AI adoption in "professional services," "finance and insurance," "health care and social assistance," "manufacturing," and "retail trade"); *2026 Physician Survey on Augmented Intelligence*, AM. MED. ASS'N, CTR. FOR DIGIT. HEALTH & A.I., 4-6 (Mar.

agencies, combined with AI's growing sophistication, requires that various stakeholders across the government consider when and how to deploy these systems. One branch, the judiciary, has the added challenge of confronting this technological revolution while it fundamentally alters its relationships with the agencies it oversees.

AI systems have proliferated into our federal government. Federal agencies are rapidly integrating AI across virtually every aspect of their operations, from routine administrative tasks to complex policy decisions affecting millions of Americans. [91] This widespread adoption reflects both the technology's promise for increased efficiency and the Trump administration's aggressive push toward "AI-first" governance. [92] The breadth and depth of current AI deployments almost guarantees that courts will increasingly confront agency actions that have been shaped by algorithmic decision-making.

To properly situate this discussion and illustrate the prominence of AI in current government work, here are a few examples of how federal agencies currently incorporate AI systems into their work.

2026), https://www.ama-assn.org/system/files/physician-ai-sentiment-report.pdf [https://perma.cc/XC45-ZJ3A] (describing AI adoption amongst medical practitioners); *The Total Economic Impact of LexisNexis Lexis+ AI for Corporate Legal Departments*, FORRESTER TOTAL ECONOMIC IMPACT STUDY, LEXISNEXIS, 9-11 (June 2025), https://www.lexisnexis.com/en-us/products/lexis-plus-ai/roi-corporate-legal.page [https://perma.cc/8SZC-E5GL] (describing AI adoption amongst lawyers).

[91] *See, e.g.,* Press Release, U.S. GEN. SERVS. ADMIN., *GSA and ServiceNow Strike Landmark OneGov Deal to Accelerate AI-Driven Government Modernization* (Sept. 3, 2025), https://www.gsa.gov/about-us/newsroom/news-releases/gsa-servicenow-strike-landmark-onegov-deal-09032025 [https://perma.cc/BA8M-8LYL] (announcing a OneGov agreement with AI service provider, ServiceNow, to create an AI licensing model with pre-negotiated discounts for federal agencies); Request for Information on the Development of an Artificial Intelligence (AI) Action Plan, 90 Fed. Reg. 9088 (Feb. 6, 2025) (requesting input from all interested parties on the Development of an AI Action Plan for government in public services, cybersecurity, data privacy, and general governance requirements).

[92] Kelly, *supra* note 6.

- **Revenue Collection and Financial Operations**: In May 2025, Treasury Secretary Scott Bessent testified before the House Appropriations Committee that reduced IRS staffing would not impair agency operations because the "AI boom" would allow the IRS to "enhance" collections.[93] Secretary Bessent imagined that AI would compensate for fewer revenue agents by automating the process of auditing tax returns.[94]
- **Immigration and Border Security**: The Department of Homeland Security's "public use inventory" documents AI applications that track noncitizens. This includes "automated detection systems" that are equipped with the capability to analyze field images, conduct cyber threat analysis, and utilize facial recognition technology to identify individuals of interest.[95] One such platform, created by Palantir, a Colorado-based defense contractor and software company, includes ImmigrationOS.[96] This AI platform aggregates data

---

[93] Brandon Vigliarolo, *IRS hopes to replace fired enforcement workers with AI*, THE REGISTER (May 8, 2025), https://www.theregister.com/2025/05/08/irs_ai_plans/ [https://perma.cc/SJN2-N7WP].

[94] *Id.*

[95] *Artificial Intelligence Use Case Inventory Library*, DEP'T OF HOMELAND SEC., https://www.dhs.gov/publication/ai-use-case-inventory-library [https://perma.cc/M7YB-KXZZ]. *See also* Maranci, *supra* note 10 (surveying developments in federal immigration enforcement's adoption of generative AI technology).

[96] *Investigative Case Management—Additional Capabilities*, U.S. IMMIGRATION & CUSTOMS ENF'T, OFF. OF ACQUISITION MGMT., https://sam.gov/opp/f71acee6010c423db4902446a59a690c/view#description [https://perma.cc/U6ZY-QQEJ] (public post announcing DHS contract award for "ICE's Immigration Lifecycle Operating System (ImmigrationOS)" which will be used for "Targeting and Enforcement Prioritization," "Self-Deportation Tracking," and "Immigration Lifecycle Process" monitoring); Rosemarie Ho, *ICE just ordered $30 million worth of new technology from Palantir to track immigrants*, BUSINESS INSIDER (Apr. 17, 2025), https://www.businessinsider.com/ice-palantir-new-technology-30-million-visa-overstays-self-deportation-2025-4 [https://perma.cc/85VF-WKFF] (describing Palantir software aimed at tracking self-deportations and immigrants who overstayed their visas).

from across government databases to surveil noncitizen populations.[97]

- **Food and Drug Regulations and Reporting:** The FDA announced plans to deploy AI across multiple functions, from reviewing chemicals in food supplies to prioritizing facility inspections and assisting in drug and medical device approvals.[98] More controversially, Health and Human Services Secretary Robert F. Kennedy Jr. has discussed using AI to replace or alter the Vaccine Adverse Event Reporting System (VAERS), potentially automating the identification of vaccine side effects and risks.[99] Such changes could fundamentally alter how the CDC assesses and communicates information related to vaccine safety.[100]
- **Defense and National Security**: Military agencies represent another frontier of AI integration, with the

---

[97] Ho, *supra* note 96; Caroline Haskins, *ICE Is Paying Palantir $30 Million to Build 'ImmigrationOS' Surveillance Platform*, WIRED (Apr. 18, 2025), https://www.wired.com/story/ice-palantir-immigrationos/ [https://perma.cc/GC6K-WCW3] (describing passport records, Social Security files, IRS tax data, and license-plate reader information as part of the aggregated data).

[98] Christina Jewett, *F.D.A. to Use A.I. in Drug Approvals to 'Radically Increase Efficiency'*, N.Y. TIMES (June 10, 2025), https://www.nytimes.com/2025/06/10/health/fda-drug-approvals-artificial-intelligence.html [https://perma.cc/N9VF-CMUU].

[99] Matt Novak, *RFK Jr. Says AI Will Approve New Drugs at FDA 'Very, Very Quickly'*, GIZMODO (July 1, 2025), https://gizmodo.com/rfk-jr-says-ai-will-approve-new-drugs-at-fda-very-very-quickly-2000622778 [https://perma.cc/9H6K-QS2D]. *See also* Daniel Payne, *RFK Jr. plans changes to vaccine side effect reporting system*, STAT NEWS (Apr. 15, 2025), https://www.statnews.com/2025/04/15/rfk-jr-remarks-vaccine-adverse-event-reporting-system-vaers-changes-indiana-maha/ [https://perma.cc/H7XQ-9P56] (reporting on Health Secretary Kennedy's announcement of plans to reform VAERS by automating the system and increasing data sharing arrangements).

[100] *See* Matthew Perrone, *Ex-official says he was forced out of FDA after trying to protect vaccine safety data from RFK Jr.*, AP NEWS (Apr. 7, 2025), https://apnews.com/article/kennedy-rfk-vaccines-measles-fda-injury-marks-5eda3335bae9b8df88795c2d5e09ae69 [https://perma.cc/KE2D-MGR8] (describing former Food and Drug Administration vaccine chief Dr. Peter Marks' apprehension of Kennedy's accessing VAERS).

> Pentagon's "Thunderforge" program exemplifying the scope of current initiatives. The Pentagon awarded Scale AI a contract to incorporate AI agents into military planning and operations, with systems designed to simulate outcomes, plan scenarios, and refine proposed courses of action.[101] This is not the first AI platform used by the armed forces. The Army has previously experimented with "Scylla," AI surveillance software that detects weapons deployment, intruders, and other security threats in real-time and alerts security personnel. [102] The Defense Department emphasizes that these systems operate "always under human oversight," with humans making final decisions based on AI recommendations.[103]

Several patterns emerge from this survey of AI deployment. [104] First, agencies are using AI not merely as

---

[101] *Introducing Thunderforge: AI for American Defense*, SCALE (Mar. 5, 2025), https://scale.com/blog/thunderforge-ai-for-american-defense [https://perma.cc/Q78K-V76J]; *DIU's Thunderforge Project to Integrate Commercial AI-Powered Decision-Making for Operational and Theater-Level Planning*, DEFENSE INNOVATION UNIT: PROJECT SPOTLIGHT (Mar. 5, 2025), https://www.diu.mil/latest/dius-thunderforge-project-to-integrate-commercial-ai-powered-decision-making [https://perma.cc/R2L4-KZZ2].

[102] Greg Kocher, *A New Tool for Intruder Detection: BGAD Tests AI for Enhanced Security*, U.S. ARMY (Oct. 8, 2024), https://www-army-mil.cdn.ampproject.org/c/s/www.army.mil/article-amp/280348/a_new_tool_for_intruder_detection_bgad_tests_ai_for_enhanced_security [https://perma.cc/K65S-HU9G].

[103] *Introducing Thunderforge: AI for American Defense*, supra note 101; *see also* Brandon Vigliarolo, *It Begins: Pentagon to Give AI Agents a Role in Decision Making, Ops Planning*, REGISTER (Mar. 5, 2025), https://www.theregister.com/2025/03/05/dod_taps_scale_to_bring/ [https://perma.cc/8ZAY-N7XV] ("The end goal of the project, the DIU said, is to help military decision-makers pore over and assess more info more quickly and make judgement calls more rapidly based on AI suggestions . . . . We're given the strong impression humans make all the final decisions, albeit ones guided by software.").

[104] These reported uses of AI platforms by federal agencies are limited to factual and policy determinations, which—unlike statutory interpretation—trigger judicial review under the arbitrary and capricious or reasoned

analytical tools but as auxiliary or in some cases *primary* decision-makers in matters affecting national security, individual rights, and access to public benefits.[105] Second, the complexity of these systems often makes meaningful human oversight difficult, even when formally required. Third, the integration of multiple data sources and pattern recognition capabilities enables unprecedented surveillance and assessment capabilities that blur traditional boundaries between enforcement, investigation, and adjudication.

C. The Promise and Peril of AI in Government

The integration of AI systems into federal agencies presents a classic double-edged sword: the same AI systems that offer opportunities for increased efficiency and consistency also have the potential to exacerbate existing issues like procedural opacity and bias. As agencies increasingly turn to AI systems to address operational challenges and meet public expectations, understanding both AI's transformative potential and its inherent limitations becomes essential for evaluating how reviewing courts should respond. This section examines AI's promise alongside its perils.

1. The Promise: Efficiency, Consistency, and Enhanced Government Performance

Democratic governments worldwide face a mounting

---

decision-making standards of review. Nevertheless, agencies' zealous embrace of AI in the context of factual and policy determinations portends future extensions of AI reliance, including into the realm of statutory interpretation, which is subject to the *Loper Bright* deferential framework. While these frameworks sometimes overlap in practice, they apply to different types of agency choices and are triggered under specific procedural and contextual conditions.

[105] *See, e.g.*, *Introducing Thunderforge: AI for American Defense, supra* note 101; *DIU's Thunderforge Project to Integrate Commercial AI-Powered Decision-Making for Operational and Theater-Level Planning, supra* note 101 ("The system's end-state will accelerate decision-making, allowing planners to more rapidly synthesize vast amounts of information, generate multiple courses of action, and conduct AI-powered wargaming to anticipate and respond to evolving threats. . . . Team Thunderforge will deliver a unified planning ecosystem where AI agents simulate wargaming and planning scenarios and refine proposed courses of action.").

credibility crisis with their constituents, driven—in part—by perceptions that public institutions are neither responsive nor transparent.[106] AI proponents view the deployment of this technology as an opportunity to showcase powerful capabilities that could restore public confidence in government competence.[107] Researchers note that, "[a]rtificial intelligence has the advantages over the natural intelligence as it is more permanent, consistent, less expensive, has the ease of duplication and dissemination, can be documented and can perform certain tasks much faster and better than the human."[108]

In addition to efficiency gains, AI has the potential to

---

[106] *Trust in Major Government Service Providers*, PERFORMANCE.GOV, https://www.performance.gov/cx/data/ [https://perma.cc/74AT-WSEV] (listing ineffectiveness, difficulty, inefficiency, lack of transparency, and unfairness as the top five factors cited by respondents where they indicated room for improvement by federal agencies). *See also* Valrie Wirtschafter, *For AI to Make Government Work Better, Reduce Risk and Increase Transparency*, BROOKINGS (Jan. 16, 2025), https://www.brookings.edu/articles/for-ai-to-make-government-work-better-reduce-risk-and-increase-transparency/ [https://perma.cc/UZA4-743R] ("Democracies around the world face a fundamental challenge: their citizens do not believe they can deliver results. Recent surveys of the countries that are members of the Organization for Economic Cooperation and Development show declining trust in government, driven in part by the perception that public institutions are neither responsive nor transparent. And a series of Pew Research Center polls show a clear decline in satisfaction with democracy across 12 advanced economies, including in the United States.").

[107] Bommasani et al., *supra* note 76, at 60 (the other two legal areas identified as susceptible to AI use include "criminal law" and "civil justice"); Wirtschafter, *supra* note 106 ("With a new wave of government advisers, many of whom come from the 'techno-optimist' space, it seems inevitable that technological advances, including the deployment of artificial intelligence (AI), will be part of any proposed solution.").

[108] Pannu, *supra* note 72, at 79; *see also* Erik Brynjolfsson et al., *Generative AI at Work* (Nov. 7, 2024), https://arxiv.org/abs/2304.11771 [https://perma.cc/PV8M-Y7QL] (study of 5,000 customer service sector employees that found, among other things, that "access to AI assistance markedly improves how customers treat agents, as measured by the sentiment of their chat messages. We also find that customers are less likely to question the competence of agents by requesting to speak to a supervisor").

reduce bureaucratic arbitrariness by systematically producing higher-quality and more consistent decisions. Data from the adoption of AI in various workplace settings suggest that the technology can "improve output quality, save workers time, decrease disparities in performance, and raise job satisfaction."[109] These performance enhancements appear readily transferable to government operations, where agencies face mounting pressure to deliver better services with less resources.[110] AI's cost-effectiveness and scalability make it a particularly attractive option for resource-constrained agencies seeking to maintain or expand service delivery.

However, efficiency gains alone may be overvalued as a driver of constituent satisfaction. For instance, research on airport security experiences suggests that improvements in interpersonal communication in areas like professionalism, respect, and understanding matter more for public satisfaction than reduced wait times achieved through technological solutions. [111] Thus, AI's efficiency benefits may not

---

[109] Brynjolfsson et al., *supra* note 108; Wirtschafter, *supra* note 106.

[110] *See* Bommasani et al., *supra* note 76 ("A major promise is that foundation models can improve access to justice and government services by leveling procedural and financial barriers to legal services. The challenges posed by legal applications can, in turn, motivate basic research questions for foundation models."); FY 2023 ANNUAL PERFORMANCE REPORT/FY 2025 ANNUAL PERFORMANCE PLAN, U.S. PAT. & TRADEMARK OFF. 15, https://www.uspto.gov/sites/default/files/documents/USPTO_FY23_FY25_APPR.pdf [https://perma.cc/9WN9-A7N8] ("Emerging technologies using AI, machine learning (ML), and robotic process automation (RPA) bring potential opportunities to maximize both the USPTO's operations and workforce. These tools can serve as catalysts for operating efficiencies and cost savings, enhancing the work of patent examiners and increasing access to prior art."); Jonathan Bright et al., *Generative AI Is Already Widespread in the Public Sector*, ALAN TURING INST., 4 (Jan. 3, 2024), https://arxiv.org/abs/2401.01291 [https://perma.cc/FP95-7U4Q] (reporting that 32.62% of 938 UK public sector employees reported using "Any AI").

[111] Wirtschafter, *supra* note 106 ("[A]lthough perceptions of the airport security experience have improved in the United States, a survey of more than 13,000 TSA customers found that this shift had less to do with perceived wait time—which facial recognition technology has helped to

automatically translate directly into enhanced public trust.

2. The Perils: Transparency Deficits and Systematic Risks

a. Public Skepticism

Despite AI's potential benefits, its deployment brings notable risks including potential misuse, systemic accidents, and societal disruption.[112] First, many communities across the political spectrum harbor broad distrust of AI systems.[113] While this skepticism may partially reflect a lack of familiarity with current AI capabilities,[114] poorly implemented systems risk

---

reduce—and more to do with improved 'interpersonal communication' like 'professionalism, respect, and understanding of security procedures.' In short, friendly TSA agents may make more of a difference to travelers' experiences than a quicker wait in line.").

[112] *Planning for AGI and Beyond*, *supra* note 86.

[113] Michelle Faverio & Emma Kikuchi, *Key Findings About How Americans View Artificial Intelligence*, PEW RSCH. CTR. (Mar. 12, 2026), https://www.pewresearch.org/short-reads/2026/03/12/key-findings-about-how-americans-view-artificial-intelligence/ [https://perma.cc/Q7XM-JTSE] ("Americans continue to be wary of AI's impact on daily life. Half of U.S. adults say the increased use of AI in daily life makes them feel more concerned than excited, according to a June 2025 survey. Just 10% say they are more excited than concerned. Another 38% say they are equally concerned and excited."). *See also* Wirtschafter, *supra* note 106 (cautioning that "harmful deployment of these tools—whether by governments or in the private sector—could alienate citizens, lead to deeper skepticism of their potential benefits, or even slow the development of transformative solutions. It may also further erode public trust in government writ large").

[114] *See* Michelle Faverio & Alec Tyson, *What the Data Says About Americans' Views of Artificial Intelligence*, PEW RSCH. CTR. (Nov. 21, 2023), https://www.pewresearch.org/short-reads/2023/11/21/what-the-data-says-about-americans-views-of-artificial-intelligence/ [https://perma.cc/FK4C-KDPH]; Liao & Vaughan, *supra* note 6, at 8 ("[LLMs'] public perception is still evolving and shaped by complex mechanisms including mass media, marketing campaigns, ongoing events, and design choices of popular LLM-infused applications."). *Cf. Americans Are Split on AI*, AM. CUSTOMER SATISFACTION INDEX 1, 3 (2026), https://theacsi.org/wp-content/uploads/2026/04/AI-Survey-Report_ACSI_2026.pdf [https://perma.cc/B5U8-WHMM] (study of six AI platforms found that

delegitimizing AI solutions across government agencies.[115] Although some American populations have grown cautiously optimistic about AI products, current data suggests that there is a persistent and enduring public anxiety about the technology that is unevenly distributed across communities.[116]

Current technological limitations of AI systems, including hallucinations and bias, may not permanently undermine their adoption in administrative settings. As AI technology advances, accuracy concerns may become less significant impediments to implementation. Global sentiment on AI has begun to shift in favor of the new technology in certain countries.[117] AI developers are increasingly focused on

---

more than half of U.S. adults [56%] have no recent experience using any AI platform compared with 44% who do; of the adults who report using AI platforms, the majority [52%] report regular, daily engagement with these platforms).

[115] Wirtschafter, *supra* note 106.

[116] *See also* Maslej et al., *supra* note 8, at 397 ("Early data suggests growing public anxiety about AI, with some regions expressing significantly more pessimism than others."); Nick Lichtenberg, *Most of You Are Rejecting AI. The Data Shows You're Running Out of Time*, FORTUNE (Apr. 16, 2026), https://fortune.com/2026/04/16/ai-resistance-running-out-of-time-rebellion-quiet-quitting-trust/ [https://perma.cc/HTJ3-F44Q] (reporting on a global survey of 3,750 executives and employees across 14 countries which found "roughly eight in 10 enterprise workers are either avoiding or actively rejecting the technology their employers are spending record sums to deploy"). *Compare* Faverio & Kikuchi, *supra* note 113 ("Americans continue to be wary of AI's impact on daily life. Half of U.S. adults say the increased use of AI in daily life makes them feel more concerned than excited, according to a June 2025 survey. Just 10% say they are more excited than concerned. Another 38% say they are equally concerned and excited.") *with* Faverio & Tyson, *supra* note 114 ("Americans are increasingly cautious about the growing role of AI in their lives generally. Today, 52% of Americans are more concerned than excited about AI in daily life, compared with just 10% who say they are more excited than concerned; 36% feel a mix of excitement and concern.").

[117] Maslej et al., *supra* note 8, at 4 ("Global AI optimism is rising—but deep regional divides remain. In countries like China (83%), Indonesia (80%), and Thailand (77%), strong majorities see AI products and services as more beneficial than harmful. In contrast, optimism remains far lower in places like Canada (40%), the United States (39%), and the Netherlands (36%).

building public confidence by incorporating data privacy and security measures.[118]

Nevertheless, widespread AI adoption may promote additional public skepticism by further deepening the public's disconnect from their government and the federal agencies that comprise it. The public currently holds mixed views about federal agencies; [119] adopting AI systems that arguably "cheapen" the public's relationship with the agency has the

---

Still, sentiment is shifting: Since 2022, optimism has grown significantly in several previously skeptical countries, including Germany (+10%), France (+10%), Canada (+8%), Great Britain (+8%), and the United States (+4%).").

[118] *See, e.g., Path to Astra: Critical Capabilities and Frontier Safeguards*, OPENAI (Sept. 1, 2026), https://openai.com/index/path-to-astra/ [https://perma.cc/4W6P-A5VV] (discussing OpenAI's decision to postpone release of its "astra" model while it "strengthened and tested protections against cyber misuse and unauthorized model actions"); *The Hugging Face Incident and the Road Ahead*, OPENAI (Aug. 26, 2026), https://openai.com/index/hugging-face-incident-and-the-road-ahead/ [https://perma.cc/AD34-8E25] (describing an OpenAI model's hack of Hugging Face and subsequent actions taken by OpenAI to improve security); LEXISNEXIS, *supra* note 90, at 31-32 (discussing Lexis+ AI satisfying "both enterprisewide information security and privacy requirements while also meeting the unique needs of legal departments with commercial-grade private cloud infrastructure, independently trained models with siloed user sessions, customer data encryption, and help meeting regulations like the EU's AI processing center regulations"); *see also* Benedict Vigers & Justin Lall, *Americans Prioritize AI Safety and Data Security*, GALLUP (Sept. 16, 2025), https://news.gallup.com/poll/694685/americans-prioritize-safety-data-security.aspx [https://perma.cc/S7HE-98ZS] (reporting that approximately 80% of Americans prefer "maintaining rules for AI safety and data security, even if it means developing AI capabilities at a slower rate," but are almost evenly split between whether the ultimate responsibility for AI regulation lies with the U.S. government or with the companies developing the AI platforms).

[119] *See* Jocelyn Kiley & Andy Cerda, *Republicans' Views of Justice Department, FBI Rebound as Democrats' Views Shift More Negative*, PEW RSCH. CTR. (Aug. 2025), https://www.pewresearch.org/wp-content/uploads/sites/20/2025/08/PP_2025.8.27_agency-favorability-report.pdf [https://perma.cc/D493-4CKE].

potential to only exacerbate these problems.[120]

b. The Black Box Problem: Transparency and Interpretability

While the flexibility and adaptability of AI models make them a highly desirable technology, they achieve these outputs through unknown and potentially unknowable processes. Modern LLMs are trained on vast, disparate datasets that span multiple modalities, enabling them to "learn" exceptionally wide-ranging behaviors that make them adaptable to diverse downstream tasks.[121] This capability allows AI models to perform unforeseen tasks in unforeseen ways.[122] Researchers still do not fully comprehend how these systems operate in practice.[123] At this point, transparency remains inadequate,[124]

---

[120] Tara Aida, *Ticking the Boxes: AI and the Notice-and-Comment Process,* YALE J. ON REG., NOTICE & COMMENT (Feb. 13, 2026), https://www.yalejreg.com/nc/ticking-the-boxes-ai-and-the-notice-and-comment-process-by-tara-aida/ [https://perma.cc/U5CQ-V595] ("[I]n a world where LLMs have suddenly made reading, summarizing, and drafting quite 'cheap,' § 553 might miss the mark by treating the written word as a proxy for intellectual engagement.").

[121] Bommasani et al., *supra* note 76, at 53-67, 81-84.

[122] *E.g.*, *Our Framework for Reporting Model Misalignment*, OPENAI (Sept. 16, 2026), https://openai.com/index/model-misalignment-reporting-framework/ [https://perma.cc/TH4T-Q25P] (disclosing six "misalignment" incidents wherein OpenAI systems acted without authorization, coordinated with other models, or otherwise attempted to evade oversight during models' training or evaluation); Wijk & Cotra, *supra* note 16, at 5-20 (describing an unprecedented, AI-coordinated hack involving tens of thousands of AI agents across multiple models seeking to trick or tamper with the automated performance evaluation system).

[123] Bommasani et al., *supra* note 76, at 124-25.

[124] Maslej et al., *supra* note 8. In addition to intellectual property concerns, public safety is also a justification for reduced transparency around AI products. *See, e.g.*, Erin Woo & Robert McMillan, *Gemini Hacked Three Companies in First Known Breakout by Google's AI*, WALL ST. J. (Sept. 18, 2026), https://www.wsj.com/tech/ai/gemini-hacked-three-companies-in-first-known-breakout-by-googles-ai-5c0baba2 [https://perma.cc/WK9H-927T] (reporting on journalists' discovery of autonomous hacks by Google's Gemini models during third-party testing, which Google previously failed to disclose); *Our Framework for Reporting Model Misalignment, supra* note

and it fosters further uncertainty about the capabilities and failure modes (i.e., failsafes) of AI models.[125] Base LLMs are often adapted for specific applications by parties other than their original developers, who may only have restricted access to application training interfaces rather than underlying model architectures.[126] When models build upon each other, defects in foundation models—the common basis from which subsequent task-specific models are built—may create unforeseen effects in subsequent iterations.[127] This layered opacity compounds when multiple parties adapt base models by incorporating domain-specific data or enabling in-context learning by the end user.[128] Any transparency issues in pretrained models thus

---

122 ("We do not believe that the AI industry has solved alignment and monitoring to a sufficient degree to continue responsibly scaling at maximum speed for much longer. . . . At the moment, there is no industry-wide framework with explicit standards for how AI developers should disclose examples of misalignment in their models."); Wijk & Cotra, *supra* note 16; *Project Glasswing*, ANTHROPIC, https://www.anthropic.com/glasswing [https://perma.cc/M7T7-GK3Y] (describing "Project Glasswing," a coordinated defensive effort to share Claude Mythos Preview, an unreleased frontier model, with select software partners to address "high-severity vulnerabilities" in *every* major operating system and web browser before the model's release to the public); *see also* Nicholas Carlini et al., *Assessing Claude Mythos Preview's cybersecurity capabilities*, ANTHROPIC (Apr. 7, 2026), https://red.anthropic.com/2026/mythos-preview/ [https://perma.cc/9R9T-N579] ("As we discuss below, we're limited in what we can report here. Over 99% of the vulnerabilities we've found have not yet been patched, so it would be irresponsible for us to disclose details about them (per our coordinated vulnerability disclosure process)."). Shortly after Anthropic announced the existence of Claude Mythos, the Trump administration issued Exec. Order No. 14409, which articulated a voluntary framework by which the federal government would review new AI platforms prior to release to the general public. Exec. Order No. 14409 (June 2, 2026).

125 Bommasani et al., *supra* note 76, at 6.

126 Liao & Vaughan, *supra* note 6, at 5.

127 Bommasani et al., *supra* note 76 ("Though foundation models are based on standard deep learning and transfer learning, their scale results in new emergent capabilities, and their effectiveness across so many tasks incentivizes homogenization. Homogenization provides powerful leverage but demands caution, as the defects of the foundation model are inherited by all the adapted models downstream.").

128 Liao & Vaughan, *supra* note 6.

cascade through to adapted LLM or LLM-infused applications, which may be the only systems end users directly encounter.[129]

In recent years AI developers have sought to address these concerns by building tools to allow users to peer into the AI black box, but that technology remains fundamentally opaque.[130] Current LLMs are uniquely characterized by their emergent properties, where system behavior is "implicitly induced rather than explicitly constructed."[131] These emergent properties create massive, opaque architectures that are so large and complex that providing complete pictures of the "knowledge" reflected in models or their reasoning processes becomes impossible. [132] Even AI experts do not fully

---

[129] *Id.* at 7 ("End-users may not interact with LLMs directly, but rather through LLM-infused applications. Emerging applications include general and specialized chatbots, web search, programming assistants, productivity tools such as for writing support or presentation generation, and text analysis tools such as for customer insights discovery."); *Ethics of Artificial Intelligence and Robotics*, *supra* note 71, at § 2.3 ("[T]he quality of the program depends heavily on the quality of the data provided, following the old slogan 'garbage in, garbage out.' So, if the data already involved a bias (e.g., police data about the skin colour of suspects), then the program will reproduce that bias.").

[130] Maslej et al., *supra* note 8, at 16 ("The updated Foundation Model Transparency Index—a project tracking transparency in the foundation model ecosystem—revealed that the average transparency score among major model developers increased from 37% in Oct. 2023 to 58% in May 2024. While these gains are promising, there is still considerable room for improvement.").

[131] Bommasani et al., *supra* note 76, at 3 ("[Emergence] is both the source of scientific excitement and anxiety about unanticipated consequences.").

[132] Liao & Vaughan, *supra* note 6; *see also Ethics of Artificial Intelligence and Robotics*, *supra* note 71, at § 2.3 ("Many AI systems rely on machine learning techniques in (simulated) neural networks that will extract patterns from a given dataset, with or without 'correct' solutions provided; i.e., supervised, semi-supervised or unsupervised. With these techniques, the 'learning' captures patterns in the data and these are labelled in a way that appears useful to the decision the system makes, while the programmer does not really know *which* patterns in the data the system has used. In fact, the programs are evolving, so when new data comes in, or new feedback is given ('this was correct', 'this was incorrect'), the patterns used by the learning system change. What this means is that the outcome is not transparent to the user or programmers: it is opaque.").

understand how these models arrive at particular outputs.[133]

c. Hallucinations and the Reliability Challenge

Unlike other forms of automation, AI systems present unique challenges for building trustworthy governance systems because LLMs communicate through natural language.[134] Previous studies have found that LLMs exhibit tendencies to produce content deviating from factual accuracy or well-established principles, a phenomenon termed "hallucinations."[135] Current training methods tend to incentivize hallucinations in LLMs, as models are designed to "guess" the next chunk of text in a sequence.[136] And to make matters worse, these hallucinations are not rare occurrences.[137]

---

[133] Liao & Vaughan, *supra* note 6, at 6 ("There are currently no established answers, even in the research community, to questions such as precisely how and why these models work as well as they do, why they can or cannot perform certain tasks, and how characteristics of the training data impact model capabilities.").

[134] Bommasani et al., *supra* note 76, at 49 ("One might argue that we cannot trust a system's linguistic behavior unless it understands the language it is using. Of course, we currently trust engineered systems to do things (e.g., manufacturing auto parts) without the question of understanding even arising, but language might be special in this regard, since it is uniquely human. In addition, language can be used to deceive and misrepresent, so understanding alone clearly does not imply trust.").

[135] Ho, *supra* note 3.

[136] *Tracing the Thoughts of a Large Language Model*, ANTHROPIC (Mar. 27, 2025), https://www.anthropic.com/news/tracing-thoughts-language-model [https://perma.cc/ZD7U-JJC2] ("At a basic level, language model training incentivizes hallucination: models are always supposed to give a guess for the next word. Viewed this way, the major challenge is how to get models to not hallucinate.").

[137] *See, e.g.*, Tripp Mickle et al., *How Accurate Are Google's A.I. Overviews?*, N.Y. TIMES (Apr. 7, 2026), https://www.nytimes.com/2026/04/07/technology/google-ai-overviews-accuracy.html [https://perma.cc/9ER4-F75Z] ("A recent analysis of AI Overviews found that they were accurate approximately nine out of 10 times. But with Google processing more than five trillion searches a year, this means that it provides tens of millions of erroneous answers every hour (or hundreds of thousands of inaccuracies every minute), according to an analysis done by an A.I. start-up called Oumi."); Matthew Dahl et al., *Large*

The complex capabilities of LLMs lead to unpredictable behaviors and ultimately make regulating hallucinations challenging. [138] Unlike previous machine learning systems operating within well-defined input-output structures, LLMs demonstrate greater flexibility. [139] While this adaptability represents a marketable advantage, it limits developers' abilities to fine-tune outputs and guarantee consistency. [140] Consequently, LLM capabilities cannot be fully anticipated, nor is their output always reliable.

Despite technological improvements that have reduced the rate of hallucinations, significant challenges remain. Recent research suggests that LLM hallucinations could be diminished through various techniques, but these methods have notable limitations. [141] Legal hallucinations remain particularly problematic, posing especially difficult challenges for the legal profession. While recent AI models excel at law-related tasks such as first-year law school examinations[142] or the Uniform Bar Exam,[143] legal hallucinations remain pervasive, with one

---

Legal Fictions: Profiling Legal Hallucinations in Large Language Models, J. LEGAL ANALYSIS 1, 21-33, 35-36 (forthcoming Apr. 25, 2024), https://arxiv.org/pdf/2401.01301 [https://perma.cc/9ESQ-YZJC] (detailing study findings that hallucinations were "generally widespread" and persistent across four tested LLMs—GPT 4, GPT 3.5, PaLM 2, Llama 2—and detailing variables that affect the hallucination rate of tested models).

[138] Liao & Vaughan, *supra* note 6, at 6.

[139] *Id.*

[140] *See, e.g.*, *Anthropic's Transparency Hub*, https://www.anthropic.com/transparency [https://perma.cc/ZTL9-4W3L] (last updated Aug. 17, 2026) (noting that during training of Claude Opus 4, the model demonstrated "problematic behaviors," which included "inconsistent responses between conversations, role-playing undesirable personas, and occasional agreement with harmful requests like planning terrorist attacks when given certain instructions").

[141] *See* Dahl et al., *supra* note 137, at 35-36 (discussing the promise and limitations of techniques such as retrieval-augmented generation, advanced prompting, specialized fine-tuning, factuality-focused decoding methods, and external database checks).

[142] *Id.* at 3.

[143] Daniel Martin Katz et al., *GPT-4 Passes the Bar Exam*, 382 PHIL. TRANSACTIONS ROYAL SOC'Y A, 3-4 (2024) (finding that GPT-4 performed a 75.7% accuracy on the "Multistate Bar Exam" and outperformed the student average).

study finding hallucination rates ranging "from 69% to 88% in response to specific legal queries."[144] The likelihood of legal hallucinations increases with the complexity of the research task.[145]

Some of these technical challenges facing legal AI applications could potentially be addressed through proper investment and training. For example, in the legal context, decision-making requires integration of various knowledge bases—historical context, canons of construction, precedent, and balancing competing considerations through deliberation—but current AI technology has several significant weaknesses in this area.[146] For instance, legal document translation requires precision and an understanding of highly technical terminology, which makes data collection for this specific type of task laborious and costly.[147] While this challenge could be offset to an extent by appropriate investment in training datasets, there remain additional limitations that no amount of financial investment could resolve. Only a fraction of legal decisions result in published opinions, for example, which limits the amount of relevant legal text available for training AI models in this function.[148]

---

[144] Ho, *supra* note 3 (reporting on the author's study that found "hallucination rates range from 69% to 88% in response to specific legal queries for state-of-the-art language models."); *see also* Dahl et al., *supra* note 141, at 22-29.

[145] Dahl et al., *supra* note 137, at 22-29.

[146] *Id.* at 34 ("Instead of examining LLMs' ability to engage in legal reasoning, we assess LLMs' capacity to internalize legal knowledge. Ultimately, LLMs will need to excel in both of these respects if they are going to be effectively integrated into the legal profession. So long as they suffer from gaps in their background legal knowledge—as our results suggest—they will be unable to function as reliable sources of legal counsel and advice, no matter how strong their in-context reasoning abilities become.").

[147] Bommasani et al., *supra* note 76, at 63 ("The average U.S. Supreme Court opinion contains around 4,700 words, a brief on the merits to the Supreme Court can have as many as 15,000 words, a law review article often contains 20,000 to 30,000 words, parole transcripts can be hundreds of pages long [Bell et al. 2021], and trial records can be even longer. Current foundation models have struggled with such long contexts and outputs.").

[148] *Id.* at 66.

Similarly, challenges in processing colloquial speech, local dialects, and specialized language could be addressed through improved training datasets and model adaptation to low-resource contexts.[149]

Legal settings in particular require precision in both factual accuracy and in using technical terminology, as imprecise statements can have drastic, unanticipated consequences and potentially even result in sanctions.[150] Since hallucinations result directly from the core structures that power current LLMs, no guarantee exists that they can ever be fully eliminated.[151]

d. Bias Amplification and Systematic Discrimination

AI systems can compound pre-existing inequalities by producing unfair outcomes, entrenching power systems, and disproportionately distributing negative technological consequences to already marginalized communities.[152] LLMs and AI systems more generally have been repeatedly found to

---

[149] *Id.* at 61.

[150] *Id.* at 65; *see also* Jewett, *supra* note 98 (reporting that FDA staff members found AI tools "far from transformative" in part because of their tendency to hallucinate or produce false information).

[151] ANTHROPIC, *supra* note 136 ("[L]anguage model training incentivizes hallucination: models are always supposed to give a guess for the next word. Viewed this way, the major challenge is how to get models to *not* hallucinate. Models like Claude have relatively successful (though imperfect) anti-hallucination training; they will often refuse to answer a question if they don't know the answer, rather than speculate. . . . By intervening in the model and activating the 'known answer' features (or inhibiting the 'unknown name' or 'can't answer' features), we're able to *cause the model to hallucinate* (quite consistently!) that Michael Batkin plays chess."); *What are AI Hallucinations?*, GOOGLE CLOUD, https://cloud.google.com/discover/what-are-ai-hallucinations [https://perma.cc/XL8H-8L38] (defining AI hallucinations and explaining how they occur).

[152] Bommasani et al., *supra* note 76, at 130; *see, e.g.*, Suzanne Lucas, *AI Hiring Tools Are Advising Women and Minorities to Ask for Lower Pay in Salary Negotiations*, INC. (July 23, 2025), https://www.inc.com/suzanne-lucas/ai-hiring-tools-are-advising-women-and-minorities-to-ask-for-lower-pay-in-salary-negotiations/91217775 [https://perma.cc/A2QB-8GJ8].

encode bias.[153] This creates reliance risks, as these technologies may perpetuate harmful discrimination, stereotypes, and exclusions at scale. LLMs are trained on human-produced data that includes intrinsic bias and representational bias, wherein people are misrepresented by pernicious stereotypes or negative attitudes or have their presence distorted through underrepresentation or overrepresentation in the datasets.[154] Since foundation models serve as bases for myriad applications, biases in foundation models propagate across many applications and settings.[155] Subsequent AI model iterations inherit problematic biases from earlier foundation models.[156]

It is true that there is no guarantee that humans perform better than AI systems at addressing inequities and biases, but human decision makers offer a decentralized structure that limits the potential harm of any single decision maker. The scale of AI models, and the centralized nature of their operation, effectively removes such barriers to systemic bias and discrimination.[157]

The ubiquity and scale of these systems can cause compounding injustices by amplifying arbitrary exclusions, perpetuating misclassifications, and producing results that systematically disadvantage individual decision subjects.[158]

---

[153] Liao & Vaughan, *supra* note 6, at 3.

[154] Bommasani et al., *supra* note 76, at 131; *see also Ethics of Artificial Intelligence and Robotics*, *supra* note 71 (describing three forms of biases that are potentially exacerbated by AI systems: (1) implicit "learned bias;" (2) "cognitive bias," which "tend[s] to interpret information as confirming what they already believe"; and (3) "statistical bias," wherein the dataset itself distorts reality).

[155] Bommasani et al., *supra* note 76, at 131.

[156] *Id.*

[157] *Id.* at 152-53 ("To the extent that models train on similar data, they are likely to acquire similar patterns of behavior, biases.").

[158] *Id.* at 152 ("In contrast to the current distributed and varied model of decision making, employing many adaptations of the same foundation model for multiple automated decision-making tasks means that decision subjects may face a more homogeneous set of judgments rooted in the underlying foundation model. This algorithmic monoculture could lead to

Researchers have already proposed solutions for many AI concerns, including requiring transparent documentation and auditing training data to address bias and harmful content.[159] Despite developers' efforts to train explicitly unbiased models, however, these attempts have been unsuccessful to date.[160]

III. Judicial Deference and Its Theoretical Underpinnings

The doctrine of judicial deference to administrative agencies rests on foundational assumptions about human expertise, democratic accountability, and reasoned decision-making that AI adoption upends. As LLMs and AI more broadly function as general-purpose technologies, human actors can increasingly outsource significant portions of their analytical thinking to these platforms, raising profound questions about when and whether traditional justifications of administrative deference remain valid.[161] Understanding these theoretical underpinnings is essential for evaluating how courts

---

consistent and arbitrary rejection, mis-classification, or ill-treatment of individual decision subjects. We will call this homogenization." (internal citations omitted)); NAT'L SEC. COMM'N ON A.I., *supra* note 74, at 32-33 ("We have entered the age of deployed AI. AI is now ubiquitous, embedded in devices we use and interact with on a daily basis—for example, in our smartphones, wireless routers, and cars. We routinely rely on AI-enriched applications, whether searching for a new restaurant, navigating traffic, selecting a movie, or getting customer service over the phone or online.").

[159] Bommasani et al., *supra* note 76, at 101-04.

[160] Maslej et al., *supra* note 8, at 165 ("Many advanced LLMs—including GPT-4 and Claude 3 Sonnet—were designed with measures to curb explicit biases, but they continue to exhibit implicit ones. The models disproportionately associate negative terms with Black individuals, more often associate women with humanities instead of STEM fields, and favor men for leadership roles, reinforcing racial and gender biases in decision making. Although bias metrics have improved on standard benchmarks, AI model bias remains a pervasive issue."); *Tracing the Thoughts of a Large Language Model*, *supra* note 136 (discussing how training Claude 3.5 Haiku to refuse to answer questions by default that it doesn't know the answer to reduces the model's hallucinations).

[161] *See* Rebecca Crootof, *Cyborg Justice and the Risk of Technological-Legal Lock-In*, 119 COLUM. L. REV. F. 233, 244 (2019) (discussing "automation bias" and the related risks of overreliance and "skill fade").

should approach agency actions that are increasingly shaped by algorithmic decision-making.

Loper Bright reasserted early-twentieth-century justifications for administrative deference; the implicit congressional intent reasoning that underpinned *Chevron* did not survive.[162] *Loper Bright* similarly reduced the importance of agencies' subject-matter expertise when weighing whether to grant deference. Instead, by equating the act of policymaking with traditional statutory interpretation, the Court reclaimed expertise as a feature belonging to the judiciary, not agencies.[163] Nevertheless, post-*Loper Bright*, agency expertise still informs the "respect" afforded to agency interpretations, albeit only to a point.[164] The majority opinion also notably omits any meaningful discussion of political accountability in its discussion of deference save for a passing reference in its discussion of *Chevron*.[165]

In dismissing or diminishing the policy justifications that undergirded *Chevron*, *Loper Bright* revives the deference justifications that predate the modern administrative state and which place significant weight upon historical practice. In describing the "respect" afforded to executive interpretations, the Court cites nineteenth century precedent for the

---

[162] Loper Bright Enters. v. Raimondo, 603 U.S. 369, 399-400 (2024) ("In neither case does an ambiguity necessarily reflect a congressional intent that an agency, as opposed to a court, resolve the resulting interpretive question. And many or perhaps most statutory ambiguities may be unintentional."). Explicit congressional delegations of broad interpretive authority to federal agencies, however, remain permissible. *Id.* at 394-95.

[163] *Id.* at 401-02 ("[W]e recently noted that interpretive issues arising in connection with a regulatory scheme often 'may fall more naturally into a judge's bailiwick' than an agency's. . . . *Chevron*'s broad rule of deference, though, demands that courts presume just the opposite. Under that rule, ambiguities of all stripes trigger deference." (quoting Kisor v. Wilkie, 588 U.S. 558, 578 (2019))).

[164] *Id.* at 386 ("The Court also gave 'the most respectful consideration' to Executive Branch interpretations simply because '[t]he officers concerned [were] usually able men, and masters of the subject,' who were '[n]ot unfrequently ... the draftsmen of the laws they [were] afterwards called upon to interpret.'" (quoting United States v. Moore, 95 U.S. 760, 763 (1878))).

[165] *Id.* at 397.

proposition that longstanding, contemporaneous government practices legitimize agency interpretations.[166] Much like the Court applies the First Congress canon of construction that places greater weight on actions of the original Congress interpreting the Constitution,[167] the Court in the pre-*Chevron* era viewed contemporaneous constructions as granting greater legitimacy to agency action.[168] This renewed interest in historical practice as a policy justification inverts the *Chevron* paradigm that favored flexible policymaking's responsiveness to modern needs for one that favors administrative stasis.[169] This bygone flexibility seems more crucial than ever in the age of AI.

A. The Challenge to Traditional Expertise Models Posed by AI

Agency expertise has long served as a central justification for judicial deference to administrative interpretations. Courts

---

[166] *Id.* at 385-86 (citing the example of Edwards' Lessee v. Darby, 25 U.S. 206, 208 (1827)); *see also id.* at 386 ("Such respect was thought especially warranted when an Executive Branch interpretation was issued roughly contemporaneously with enactment of the statute and remained consistent over time.").

[167] Michael Bhargava, Comment, *The First Congress Canon and the Supreme Court's Use of History*, 94 CAL. L. REV. 1745, 1746 (2006) ("The *Eldred* decision is only the most recent example of the Supreme Court's use of a method of constitutional interpretation that looks to the actions of the First Congress as weighty or even dispositive evidence of the meaning of the Constitution. The Court has sporadically trotted out this First Congress canon of constitutional interpretation in more than thirty decisions over the last 200 years.").

[168] *Loper Bright*, 603 U.S. at 386 ("[T]he longstanding 'practice of the government'—like any other interpretive aid—can inform a court's determination of 'what the law is.'" (internal quotation marks omitted) (quoting NLRB v. Noel Canning, 573 U.S. 513, 525 (2014))).

[169] *See* United States v. Ala. Great S. R.R. Co., 142 U.S. 615, 621 (1892) ("It is a settled doctrine of this court that in case of ambiguity the judicial department will lean in favor of a construction given to a statute by the department charged with the execution of such statute, and, if such construction be acted upon for a number of years, will look with disfavor upon any sudden change whereby parties who have contracted with the government upon the faith of such construction may be prejudiced."). *Loper Bright* cites *Alabama Great Southern Railroad Co.* in its discussion of "respect." *Loper Bright*, 603 U.S. at 386.

have consistently recognized that agencies possess technical knowledge, practical experience, and specialized understanding that generalist federal judges typically lack. This expertise rationale underlies *Chevron*, *Skidmore, and now Loper Bright*, though each doctrine operationalizes expertise differently.

While democratic accountability as it was understood in the *Chevron* era is no more, the *Loper Bright* Court still recognized that the supposed expertise informing contemporaneous constructions of ambiguous statutes was often by "masters of the subject," who were frequently "the draftsmen of the laws they [were] afterwards called upon to interpret."[170] To the extent that nineteenth and early-twentieth century precedent continues to inform the Court's view of this expertise, this early precedent valued administrative expertise as a "nonpartisan vehicle of governance."[171]

1. The Fundamental Humanity and Fallibility of Expertise

A crucial corollary to expertise-based deference is that courts have consistently withheld deference from administrative interpretations that fail to demonstrate how their conclusions incorporate an agency's subject-matter expertise. The *Skidmore* framework explicitly imagined deference as a sliding scale where "respect" afforded to administrative decisions was proportional to the thoroughness of an agency's consideration. As the Court explained, "[t]he weight of such a judgment in a particular case will depend upon the thoroughness evident in [the agency's] consideration, the validity of its reasoning, its consistency with earlier and later pronouncements, and all those factors which give it power to persuade, if lacking power to control."[172] This formulation directly connects the degree of deference to the quality and depth of deliberation, which up to that point was understood

---

[170] *Loper Bright*, 603 U.S. at 386 (quoting United States v. Moore, 95 U.S. 760, 763 (1878)).

[171] DANIEL R. ERNST, TOCQUEVILLE'S NIGHTMARE: THE ADMINISTRATIVE STATE EMERGES IN AMERICA 1900-1914, at 6 (2016).

[172] Skidmore v. Swift & Co., 323 U.S. 134, 140 (1944).

to be a wholly human ingredient evident in the agency's work.[173]

This concept can also be observed in cases where the judiciary reviews agencies' attempts to interpret their own regulations. Normally, these interpretations receive what is now referred to as *Auer* deference, wherein courts defer to reasonable agency constructions of their ambiguous rules.[174] For instance, in *Gonzales v. Oregon*, the Supreme Court denied *Auer* deference to an agency regulation that did "little more than restate the terms of the statute itself" because the language at issue came from Congress and not the agency.[175] Similarly, in *Fogo De Chao (Holdings) Inc. v. U.S. Department of Homeland Security*, the D.C. Circuit held that "[n]o deference is due . . . to an agency's interpretation of its own regulation when, instead of using its expertise and experience to formulate a regulation, it has elected merely to paraphrase the statutory language."[176] These opinions exemplify an implicit precedential requirement of both human expertise and effort when determining whether or not to grant deference. Rote regurgitations or reconstitutions of pre-existing requirements lack this necessary "human" ingredient and therefore do not receive deference. To the extent that agency interpretations produced wholly or in part by AI software do the same, they too should not benefit from deference.

Our system is built on the imperfection of human intelligence and decision making. The Court acknowledged

---

[173] *See* Encino Motorcars, LLC v. Navarro, 579 U.S. 211, 222 (2016) (declining to defer to a Department of Labor regulation that "offered barely any explanation" for its policy change); Christopher v. SmithKline Beecham Corp., 567 U.S. 142, 155 (2012) ("[D]eference is likewise unwarranted when there is reason to suspect that the agency's interpretation 'does not reflect the agency's fair and considered judgment on the matter in question.'" (citing Auer v. Robbins, 519 U.S. 452, 462 (1997))).

[174] *See* Kisor v. Wilkie, 588 U.S. 558, 580 (2019) (reifying *Auer* deference).

[175] Gonzales v. Oregon, 546 U.S. 243, 257 (2006) (denying *Auer* deference to a regulation).

[176] 769 F.3d 1127, 1135 (D.C. Cir. 2014) (quoting *In re* Polar Bear Endangered Species Act Listing & Section 4(d) Rule Litig., 709 F.3d 1, 18 (D.C. Cir. 2013)) (internal quotation marks omitted).

that this premise is central to our constitutional system in *Loper Bright* when it noted that the Framers recognized that ambiguities would inevitably follow from "the complexity of objects . . . [and] the imperfection of the human faculties."[177]

2. The Human Observation Paradigm in Legal Doctrine

Legitimacy derives its value not from authority alone, but from the transparency of its foundations. When we can trace the reasoning process and examine the specific information that informed an expert judgment, we can properly evaluate its worth and apply it meaningfully to our own decisions. Opaque expertise, where conclusions are presented without their underlying logic or data, offers little more than sophisticated guesswork, regardless of the expert's credentials. Transparency allows us to distinguish between genuine insight and mere pattern matching, to understand the boundaries of expert knowledge, and to adapt their reasoning to contexts they may not have considered.

The integration of AI into agency decision-making fundamentally challenges the transparency assumptions underlying judicial deference by creating an unprecedented "black box" problem that renders traditional accountability mechanisms ineffective.[178] Unlike human expertise, which

---

[177] Loper Bright Enters. v. Raimondo, 603 U.S. 369, 400 (2024) (quoting THE FEDERALIST NO. 37, at 236); *see also id.* at 385 (noting that the Framers of our Constitution were "[c]ognizant of the limits of human language and foresight").

[178] *See* Bommasani et al., *supra* note 16, at 2 ("Black-box decision-making remains a challenge for policymakers, researchers, company executives, and the public seeking to understand why an AI model is generating a particular output."); Liao & Vaughan, *supra* note 6, at 7 ("In essence, then, *such models can only be probed in a black-box manner*, which may not be sufficient to meet the transparency requirements for stakeholders, and poses challenges for the research community to develop transparency approaches." (emphasis added)). *Contra* J. Nathan Matias, *Using Machine Learning for Precision Nudging: Jens Ludwig at Princeton*, MEDIUM (Feb. 22, 2018), https://natematias.medium.com/using-machine-learning-for-precision-nudging-jens-ludwig-at-princeton-78f4b5442514

operates through reasoning processes that can be examined, questioned, and understood,[179] AI systems produce outputs through opaque computational processes that may be fundamentally unknowable to the humans who deploy them.

This is best observed by judicial emphasis on the point of human observation and interpretation. For instance, in the context of Fourth Amendment jurisprudence, U.S. courts have developed an "exposure-based approach" that treats a "search" as occurring "when information is exposed to possible human observation, rather than when it is copied or processed by a computer."[180] This exposure-based approach presumes that the Fourth Amendment's purpose is protecting "whether and when a person's information will be kept private or exposed and shared with the police."[181] Prior to the emergence

---

[https://perma.cc/Q9E2-3HG5] ("[A]lgorithms aren't black boxes in the same way as a human mind–we have the power to design & adjust the algorithm. We can use that power to create fairer systems.").

[179] Michigan v. EPA, 576 U.S. 743, 760 (2015) (holding that the EPA acted unreasonably by failing to consider costs when determining whether regulation was "appropriate and necessary"); Util. Air Regul. Grp. v. EPA, 573 U.S. 302, 327 (2014) ("[A]n agency may adopt policies to prioritize its expenditures within the bounds established by Congress."); Morton v. Ruiz, 415 U.S. 199, 230 (1974) (stating in dictum that the Secretary of the Interior could "create reasonable classifications and eligibility requirements in order to allocate the limited funds available to him for this purpose" if it observed the necessary administrative procedures); *see also* 5 U.S.C. § 706 (authorizing courts to review agency actions to determine whether they are arbitrary, capricious, an abuse of discretion, or otherwise not in accordance with law).

[180] Ashley Deeks, *An International Legal Framework for Surveillance*, 55 VA. J. INT'L L. 291, 357 (2015); *see also* Orin S. Kerr, *Searches and Seizures in a Digital World*, 119 HARV. L. REV. 531, 550-51 (2005) ("[I]f the general process of accessing information on a computer can constitute a search, at exactly what stage does the search occur--when the hard drive heads read the data from the drive, when the computer collects the data, when the analyst can see the data, or at some other point? . . . The best answer is that a search occurs when information from or about the data is exposed to possible human observation, such as when it appears on a screen, rather than when it is copied by the hard drive or processed by the computer.").

[181] Kerr, *supra* note 180, at 551, 554 ("What made the conduct in Kyllo a search was not the existence of the radiation signal in the air, but the output

of AI, this dichotomy created a clear distinction: digital information was "searched" when exposed to human observation, not merely at the point of automated collection.[182]

AI systems that closely emulate human thought and reasoning challenge this paradigm. If the human observation—or in this case human supervision—element is missing or otherwise limited, then the AI output is akin to an unobserved dataset. As agencies incorporate AI software into their work, they must consider what constitutes meaningful human involvement in agency decision-making processes.

To fully appreciate how large language models review and perform tasks, users and courts require complete access to the models' internal architecture and reasoning processes.[183] Yet current AI models remain essentially uninterpretable, even to their creators and AI experts studying them.[184] However, transparency proves both difficult to define and even more

---

of the thermal image machine and what it exposed to human observation."); *see also* United States v. Karo, 468 U.S. 705, 712 (1984) ("[W]e have never held that potential, as opposed to actual, invasions of privacy constitute searches for purposes of the Fourth Amendment. A holding to that effect would mean that a policeman walking down the street carrying a parabolic microphone capable of picking up conversations in nearby homes would be engaging in a search even if the microphone were not turned on. It is the exploitation of technological advances that implicates the Fourth Amendment, not their mere existence.").

[182] Kerr, *supra* note 180, at 554.

[183] Liao & Vaughan, *supra* note 6, at 6; *see also Ethics of Artificial Intelligence and Robotics*, *supra* note 71, at § 2.3 ("Many AI systems rely on machine learning techniques in (simulated) neural networks that will extract patterns from a given dataset, with or without 'correct' solutions provided; i.e., supervised, semi-supervised or unsupervised. With these techniques, the 'learning' captures patterns in the data and these are labelled in a way that appears useful to the decision the system makes, while the programmer does not really know *which* patterns in the data the system has used. In fact, the programs are evolving, so when new data comes in, or new feedback is given ('this was correct', 'this was incorrect'), the patterns used by the learning system change. What this means is that the outcome is not transparent to the user or programmers: it is opaque.").

[184] Liao & Vaughan, *supra* note 6, at 7 ("There are currently no established answers, even in the research community, to questions such as precisely how and why these models work as well as they do, why they can or cannot perform certain tasks, and how characteristics of the training data impact model capabilities.").

challenging to implement in the LLM context. The proprietary nature of the technology fundamentally inhibits such transparency. This opacity becomes particularly problematic when agencies rely on AI built using proprietary technology from private parties.[185] The vast majority of LLM developers decline to disclose critical information including the "size, makeup, and provenance of training data," "number of parameters," or resources required for model training.[186] This opacity severely limits a third party's ability to inspect the technology or understand its functioning, effectively insulating these systems within impenetrable "black boxes."[187] Some LLMs have begun incorporating tools aimed at facilitating transparency through model reporting, where developers publish information about model functioning and capabilities.[188] Nevertheless, these measures remain inadequate to address the fundamental interpretability challenges.

Researchers have increasingly called for greater transparency into AI models to enable stakeholders to form appropriate understanding of a system's capabilities, limitations, and functioning, as well as users' ability to control outputs.[189] Some AI models attempt to address transparency concerns by incorporating explanatory components that interpret the reasoning behind an AI action and display that

---

[185] *Id.* ("An elephant in the room that will inevitably inhibit attempts at transparency for LLMs is the proprietary nature of the models.").

[186] *Id.*

[187] *See* Bommasani et al., *supra* note 16, at 2 ("Black-box decision-making remains a challenge for policymakers, researchers, company executives, and the public seeking to understand why an AI model is generating a particular output."); Liao & Vaughan, *supra* note 6, at 7 ("In essence, then, such models can only be probed in a black-box manner, which may not be sufficient to meet the transparency requirements for stakeholders, and poses challenges for the research community to develop transparency approaches.").

[188] Liao & Vaughan, *supra* note 6, at 15-16 (describing model-reporting frameworks such as model cards, "a popular framework for model reporting that has been adopted by companies like Google and Hugging Face").

[189] *Id.* at 3-4 (describing transparency as essential for developers, stakeholders, and end-users to have responsible AI usage and for affected parties seeking recourse).

interpretation for users while the model is computing an output.[190] While these explanatory approaches seek to increase transparency, their implementation introduces additional risks since they operate as functionally separate LLMs. This separation creates significant challenges because the provided explanations may not be truly accurate and have even been found to mislead audiences about the actual causes for the system's behavior.[191] More troubling still, these explanatory systems occasionally produce fabricated reasoning to justify factually incorrect answers; one recent study of Claude 3.7 Sonnet revealed that the model can "make up plausible-sounding steps to get where it wants to go," with this "faked" reasoning often appearing "very convincing" to users.[192] Without a complete and continuous auditing of the underlying LLM's functioning, these "think out loud" models prove to be unreliable guides to an AI's actual reasoning processes.[193]

Traditional judicial review assumes that courts can examine agency reasoning, assess the quality of decision-making processes, and hold officials accountable for their analytical

---

[190] Bommasani et al., *supra* note 76, at 125-26 ("Current explanatory approaches can generally be understood as distinct models, which are designed to provide an explanation of particular behaviors of another *black box* model.").

[191] *Id.* at 126 (describing explanations that "can lack faithfulness, by being unreliable and misleading about the causes of a behavior"); *Tracing the Thoughts of a Large Language Model*, *supra* note 136 ("Recently-released models like Claude 3.7 Sonnet can 'think out loud' for extended periods before giving a final answer. Often this extended thinking gives better answers, but sometimes this 'chain of thought' ends up being misleading; Claude sometimes makes up plausible-sounding steps to get where it wants to go. From a reliability perspective, the problem is that Claude's 'faked' reasoning can be very convincing"); *id.* (detailing Claude 3.7 Sonnet's "bullshitting" when it attempted to compute the cosine of a large number).

[192] *Id.* (detailing Claude 3.7 Sonnet's "bullshitting" when it attempted to compute the cosine of a large number).

[193] Bommasani et al., *supra* note 76, at 126 ("It is thus important to be discerning of the difference between the ability of a model to create plausible-sounding explanations and providing true insights into its behavior.").

choices.[194] When AI systems generate recommendations or make decisions through opaque processes, courts lose the ability to conduct any form of meaningful review of the reasoning that traditionally justified deference in the first place. The preservation of meaningful human oversight represents more than a mere procedural safeguard; it addresses a fundamental legitimacy crisis for democratic governance. There exists an inherent contradiction in relying on systems that may be superior to humans in analytical capability but which cannot provide a consistent nor accurate explanation of their decisions to the citizens who are subject to those decisions.[195] Legitimacy, the essential property that enables coercive public decision-making processes to be viewed as rightfully exercising authority, depends on the possibility of

---

[194] *See* 5 U.S.C. § 706; Encino Motorcars, LLC v. Navarro, 579 U.S. 211, 222 (2016) (declining to defer to a Dep't of Lab. regulation that "offered barely any explanation" for its policy change); Christopher v. SmithKline Beecham Corp., 567 U.S. 142, 155 (2012) ("[D]eference is likewise unwarranted when there is reason to suspect that the agency's interpretation 'does not reflect the agency's fair and considered judgment on the matter in question.'" (quoting Auer v. Robbins, 519 U.S. 452, 462 (1997))).

[195] *See* Henry Kissinger, *How the Enlightenment Ends: Philosophically, Intellectually—in Every Way—Human Society Is Unprepared for the Rise of Artificial Intelligence*, ATLANTIC (June 2018), https://www.theatlantic.com/magazine/archive/2018/06/henry-kissinger-ai-could-mean-the-end-of-human-history/559124/ [https://perma.cc/N6XX-53BK] ("The Enlightenment started with essentially philosophical insights spread by a new technology. Our period is moving in the opposite direction. It has generated a potentially dominating technology in search of a guiding philosophy."); *see also* John Danaher, *The Threat of Algocracy: Reality, Resistance and Accommodation*, 29 PHIL. & TECH. 245, 246 (2016) (arguing that "increasing reliance on algorithms gives rise to the threat of algocracy—a situation in which algorithm-based systems structure and constrain the opportunities for human participation in, and comprehension of, public decision-making"); JILL LEPORE, THE RISE AND FALL OF THE ARTIFICIAL STATE 9 (2026) ("Nearly two and a half centuries ago, Americans decided that societies of people are indeed capable of governing themselves by reflection and choice rather than being ruled by accident and force. But the question has returned, slightly revised: Are humans to abandon government by consent for government by machine?").

human comprehension, participation, and accountability.[196]

3. False Rationality of AI Systems

AI systems create a dangerous illusion of intellectual and analytical rigor that can mask fundamentally flawed or harmful algorithmic processes. As one comprehensive study observed, "AI provides a cloak of unwarranted rationality to actions that intensify or perpetuate injustices."[197] This false veneer of objectivity becomes particularly problematic in administrative contexts, where "the alleged rationality and objectivity of the system allow the users of AI to justify harmful actions, like benefit cuts or law enforcement harassment of students, or reinforce status quo power imbalances, such as that between employers and employees."[198]

This pattern reflects a deeper risk identified by scholars: the emergence of bureaucratic systems that become indifferent to human suffering by perceiving the world as an elaborate math problem.[199] Tellingly, AI systems applied to vulnerable populations "almost never improve access to benefits or other opportunities," suggesting that the appearance of analytical sophistication may actually facilitate rather than prevent harmful outcomes.[200]

The scope of AI systems expands the scale of risk beyond the reach of individual decision-makers by centralizing and

---

[196] Danaher, *supra* note 195, at 245-50 (discussing the three approaches to political legitimacy: pure instrumentalists, pure proceduralists, and mixed/pluralist approaches).

[197] De Liban, *supra* note 9, at 14.

[198] *Id.* at 14-15.

[199] *See* Stephen Cave, *To Save Us from a Kafkaesque Future, We Must Democratise AI*, GUARDIAN (Jan. 4, 2019), https://www.theguardian.com/commentisfree/2019/jan/04/future-democratise-ai-artificial-intelligence-power [https://perma.cc/PWC3-D7BV] ("In other words, the 'intelligence' in 'artificial intelligence' is not the intelligence of the human individual – not that of the composer, the care worker or the doctor – it is the systemic intelligence of the bureaucracy, of the machine that processes vast amounts of data about people's lives, then categorises them, pigeonholes them, makes decisions about them, and puts them in their place. The problems of AI resemble those of the Kafkaesque state because they are a product of it.").

[200] De Liban, *supra* note 9, at 15.

expanding the risks of misapplied policies, coding errors, intentional cruelty, or bias to masses of constituents.[201] Many of the mass functions enabled by AI were previously unfeasible.

IV. A Framework for AI-Era Deference: The Sliding Scale Approach

The challenge for post-*Loper Bright* jurisprudence lies in developing workable standards that distinguish between AI assistance that enhances human decision-making and AI delegation that displaces it entirely. This distinction becomes crucial as agencies increasingly rely on AI systems that may exceed human analytical capabilities in specific domains while lacking the broader judgment, political accountability, and transparency that justify judicial deference in the first place.[202]

This Article proposes a framework designed to preserve the benefits of AI-assisted statutory interpretations while maintaining the human oversight and transparency/accountability principles that justify judicial deference in the first place.[203] The goal is not to halt AI

---

[201] *Id.* at 30-46 (discussing the tens of millions of low-income Americans who are potentially subject to AI decision-making through public benefits programs like Medicaid, Medicare, SNAP, etc.); *see also The Federal Government in the Age of Artificial Intelligence: Hearing Before the H. Comm. on Oversight & Gov't Reform*, 119th Cong. (2025) (statement of Linda Miller, founder and chief growth officer of TrackLight, answering whether AI could audit Medicaid and Medicare: "Yes, that is a very, very good use case for AI because… you can look at all that data using AI very quickly and identify patterns that look suspicious").

[202] Ryan Calo & Danielle Keats Citron, *The Automated Administrative State: A Crisis of Legitimacy*, 70 EMORY L.J. 797, 806 (2021) (discussing scholarship warning of bureaucratic overreliance on computers, which "resist accountability" and "rob participants of their dignity" by removing their capacity for understanding the procedures to which they are subjected).

[203] Agencies have already zealously embraced AI models in contexts of law enforcement, fact-finding, and discretionary decision-making. *See supra* Part II.B (discussing AI deployment in tax auditing, immigration enforcement, and autonomous warfare, to name a few). Such use of AI

adoption within the government, but to ensure that the allocation of deference reflects the actual degree of human expertise, judgment, and accountability present in agency decisions. This objective requires courts to adopt review standards to ensure that beneficial AI deployment does not undermine the transparency and human control essential to administrative legitimacy.[204]

The framework proposed in this article establishes a sliding scale for determining judicial deference. The sliding scale operates on a fundamental inverse relationship: the greater an agency's reliance on AI in reaching its final decision, the less deference that decision should receive from reviewing courts. This inverse proportionality reflects the core insight that deference doctrine assumes human expertise and comprehensible reasoning processes that may be diminished or absent when agencies delegate decision-making authority to opaque algorithmic systems.

At one end of the spectrum lie agency actions where AI serves merely as a research tool or analytical aid, similar to legal databases or statistical software. On this end of the continuum, human officials retain full authority to interpret, modify, or reject AI-generated information, and the final decision reflects meaningful human expertise and judgment. Such cases may warrant traditional *Skidmore* "respect," provided that the agency's reasoning process remains transparent and the human contribution substantial. At the opposite end lie fully automated decisions where AI systems determine outcomes with minimal human involvement beyond initial system deployment. This would include, for instance, agencies using AI systems to propose statutory interpretations that satisfy policy objectives or automate the process of

---

models is beyond the scope of this Article, which discusses AI usage in statutory interpretation. These fact-finding and discretionary decision-making actions are generally reviewed by courts under the arbitrary and capricious standard, but nevertheless present analogous concerns regarding AI transparency and accountability. Such issues are set aside for future discussions and are not addressed here.

[204] *Cf.* Bommasani et al., *supra* note 76, at 114 (describing AI safety's "control problem" of reaping an AI system's benefits "while at the same time leaving us with sufficient control").

reviewing and replying to comments in the rulemaking process with minimal human interaction. These use cases warrant no deference and they run afoul of the APA requirement that agencies engage in "reasoned" decision-making.[205]

Between these poles exists a range of AI-human collaboration models requiring nuanced analysis. The framework provides specific factors for courts to consider when determining appropriate deference levels, including the nature and extent of human oversight, the transparency of AI operations, the ability of humans to understand and override AI recommendations, and the procedural safeguards surrounding AI deployment.

A. The Inverse Proportionality Rule

This proposed approach serves multiple critical functions. First, it allows for foreseen and unforeseen AI applications that enhance human decision-making capabilities while maintaining essential accountability mechanisms. Agencies can continue leveraging AI's analytical power and efficiency gains provided they maintain meaningful human control over final decisions. This incentive structure encourages AI applications that augment rather than replace human intelligence. Additionally, it provides courts with practical guidance for navigating AI-assisted administrative actions without requiring judges to become AI technical experts.

To this end, I propose that courts first explicitly recognize AI usage as one of, and potentially the most significant of, the persuasive factors considered under the *Loper Bright/Skidmore* analysis.[206] To allow for this, courts must

---

[205] Motor Vehicle Mfrs. Ass'n v. State Farm Mut. Auto. Ins. Co., 463 U.S. 29, 42-43 (1983) (holding that an agency must articulate a satisfactory explanation for its action to satisfy "arbitrary and capricious" review).

[206] *See* Skidmore v. Swift & Co., 323 U.S. 134, 140 (1944) ("The weight of such a judgment in a particular case will depend upon the thoroughness evident in its consideration, the validity of its reasoning, its consistency with earlier and later pronouncements, *and all those factors which give it power to persuade*, if lacking power to control." (emphasis added)).

impose an AI use disclosure requirement.[207] Such a requirement would be easily administrable as the federal government already discloses thousands of current AI uses.[208] This proposed change would simply extend this disclosure practice to agency decision-making.

Second, when considering AI reliance, courts should afford deference inversely proportional to the agency's reliance on AI technology. Such an approach would not lend itself to a single, bright-line rule but would instead require courts to engage in a balancing test wherein they consider a totality of the circumstances. This would require courts to closely examine the AI systems, their capabilities, their limitations, and what role the AI played in the agency's decision-making. They would consider how the AI platform's deployment comports or fails to comport with its intended and authorized use. Courts should also consider who within the agency used the AI and the frequency with which it was used.

Just as courts decline to defer to agency interpretations offering perfunctory analyses,[209] they should similarly scrutinize whether AI-driven decisions reflect (human)

---

[207] Several jurisdictions have begun to require such disclosures from practicing attorneys. See, e.g., N.D. Tex. Civ. R. 7.2(f) ("A brief prepared using generative artificial intelligence **must disclose** this fact on the first page under the heading 'Use of Generative Artificial Intelligence.' If the presiding judge so directs, the party filing the brief must disclose the specific parts prepared using generative artificial intelligence . . . (3) A party who files a brief that does not contain the disclosure required by subsection (f)(1) of this rule certifies that no part of the brief was prepared using generative artificial intelligence." (emphasis added)).

[208] *See* Adler, *Federal Government Discloses More than 1,700 AI Use Cases*, *supra* note 12; *see also* Adler, *supra* note 12 (reporting that the U.S. State Department launched a campaign of incorporating various AI technologies into its workforce procedures in 2024).

[209] *See, e.g.*, Encino Motorcars, LLC v. Navarro, 579 U.S. 211, 221 (2016) ("Agencies are free to change their existing policies as long as they provide a reasoned explanation for the change."); Nat'l Cable & Telecomms. Ass'n v. Brand X Internet Servs., 545 U.S. 967, 981 (2005) ("Agency inconsistency is not a basis for declining to analyze the agency's interpretation under the *Chevron* framework. *Unexplained* inconsistency is, at most, a reason for holding an interpretation to be an arbitrary and capricious change from agency practice under the Administrative Procedure Act." (emphasis added)).

reasoned decision-making.[210] The process of human deliberation itself serves as a proxy for democratic engagement—ensuring that policy choices reflect considered judgment by accountable officials rather than algorithmic outputs divorced from political oversight.

B. The Spectrum of Human-AI Collaboration

There are potentially infinite administrative uses for generative AI and multiple methods of incorporating them into administrative functioning. Tasks like analyzing public comments in notice-and-comment rulemakings, patent examinations, redacting Freedom of Information Act responses, and conducting removal adjudications are all potentially subject to automation by AI.[211] The adaptability and flexibility of AI models create the potential for their adoption to improve quality, efficiency, utility, and accessibility of government services where resources are constrained.[212] However, their expanded use threatens to undermine the principles that form the foundation of judicial deference precedent.

This section will use immigration law examples to highlight potential applications of the inverse proportionality rule in distinct contexts. As a form of administrative law, immigration law provides ample fodder for considering how such a system would work in practice. Immigration law is also a broad,

---

[210] *See* Calo & Citron, *supra* note 202, at 803 ("In other words, many consequential government decisions were once made by people, attended by accountability mechanisms suited to people. Now that machines make these decisions, law or technology must change to restore the rights and values afforded individuals under the previous arrangement.").

[211] *See, e.g.,* U.S. PAT. & TRADEMARK. OFF., *supra* note 70, at i-ii (reporting results of USPTO requests for comments related to AI); U.S. PAT. & TRADEMARK OFF., FY 2023 ANNUAL PERFORMANCE PLAN, *supra* note 110, at 15 ("Emerging technologies using AI, machine learning (ML), and robotic process automation (RPA) bring potential opportunities to maximize both the USPTO's operations and workforce. These tools can serve as catalysts for operating efficiencies and cost savings, enhancing the work of patent examiners and increasing access to prior art."); Bommasani et al., *supra* note 76, at 62; Eggers et al., *supra* note 74 (discussing AI's potential to reduce the USPTO's patent-application backlog).

[212] Bommasani et al., *supra* note 76, at 62.

intricate, and highly influential corpus of law that has been central to the development of administrative law, generally.[213] Moreover, this area of law is administered by several distinct agencies using a variety of techniques for statutory interpretation and implementation. For instance, two federal agencies tasked with the implementation of immigration laws, the Department of Homeland Security and the Department of Justice, regularly engage in notice-and-comment rulemakings and informal adjudications in their respective roles.[214] They offer two distinct modalities of statutory interpretation across which to explore the potential impacts of AI adoption.

To illustrate how such a rule would function in immigration law, this section will discuss three nodes on an AI reliance continuum. These nodes include purely human, partially automated, and fully automated interpretations.

1. Purely Human Decision-Making

The first node on this AI reliance continuum is the one that predates the generative AI revolution. It most closely resembles the paradigm in which the Court issued *Skidmore*, wherein federal agency decision-making is exercised purely by human intelligence. As such, this model of LLM (non)use is the

---

[213] *See, e.g.*, Adam B. Cox, *The Invention of Immigration Exceptionalism*, 134 YALE L.J. 329, 333-34 (2024) ("At the dawn of the twentieth century, American public law underwent an enormous transformation. The Supreme Court, confronting the explosive growth of the administrative state and myriad other legal and political forces, cast aside nineteenth-century ideas about due process and the separation of powers that had, for so long, organized public law. In place of that old approach, the Court created an entirely new framework for thinking about the role federal courts would play in policing executive-branch officers. Crucially, immigration law during this period evolved apace with the rest of public law; it underwent the same dramatic transformation taking place elsewhere. Indeed, immigration law was often on the leading edge of the transformation, driving the development of the legal regime we now call 'administrative law.'").

[214] *See, e.g.*, 8 U.S.C. § 1103(a)(3) (conferring on the Secretary of Homeland Security the authority to adopt regulations to aid in the enforcement of immigration and naturalization laws); *id.* § 1103(g)(1) (vesting in the Attorney General the authorities and functions exercised by the Executive Office for Immigration Review).

easiest to address. Such decision-making allows for auditability, disclosure of criteria, and opportunities to contest ultimate outcomes in a transparent manner.

a. Rulemaking

Despite aggressive AI deployment, this analog world of purely human intelligence-driven rulemaking continues to play out at DHS in select settings. For example, public disclosures suggest that rulemaking, DHS's preferred method of statutory interpretation, is still done without much reliance on generative AI.[215]

The traditional, human-driven rulemaking process embodies the deliberative, transparent decision-making that the *Skidmore* Court assumed would characterize agency action. When DHS promulgates regulations interpreting immigration statutes, these interpretations, to date, emerge from identifiably human processes. In this current system of rulemaking, career civil servants draft proposed rules, political appointees review and modify them, and the agency solicits public comment before finalizing regulatory text.[216] This process frequently generates significant public engagement.[217]

---

[215] *See The Department of Homeland Security Super Intelligence Use Case Inventory*, DEP'T OF HOMELAND SEC. (Oct. 2, 2026), https://www.dhs.gov/ai/use-case-inventory [https://perma.cc/Z3R5-TA6J] (listing a variety of AI deployment scenarios). However, note that this disclosure page discusses artificial intelligence technology and does not preclude the possibility that the agency uses other algorithms or technologies that it does not classify as "AI" in its review.

[216] 5 U.S.C. § 553 (establishing notice-and-comment procedures for rulemaking).

[217] Collection and Use of Biometrics by USCIS, 90 Fed. Reg. 49062 (proposed Nov. 3, 2025) (to be codified at 8 C.F.R. pts. 1, 103, 204, 207, 208, 209, 210, 212, 214, 215, 216, 235, 236, 240, 244, 245, 245a, 264, 287, 333, 335), https://www.regulations.gov/document/USCIS-2025-0205-0002 [https://perma.cc/H94P-438D] (6,661 comments received by the Jan. 2, 2026 deadline to submit comments); Public Charge Ground of Inadmissibility, 90 Fed. Reg. 52168 (proposed Nov. 19, 2025) (to be codified at 8 C.F.R. pts. 103, 212), https://www.regulations.gov/docket/USCIS-2025-0304 [https://perma.cc/4MFE-7P98] (8,846 comments received by the Dec. 19,

Each step involves human judgment about statutory meaning, policy consequences, and public input, which is precisely the kind of reasoned elaboration that merits judicial respect under *Skidmore*'s framework.

This human-centric process ensures the transparency and accountability that a *Skidmore* analysis demands. Courts can examine the rulemaking record to assess the "thoroughness evident in [the agency's] consideration, the validity of its reasoning, and all those factors which give it power to persuade." [218] When DHS explains why it interprets "public charge" in a manner that renders large classes of noncitizens inadmissible, records of comments and agency responses allow reviewing courts to trace an agency's interpretation through an evidentiary trail detailing how the agency considered statutory text, legislative history, and public comments.[219] Furthermore, the interpretive process remains auditable because it relies on cognitive operations familiar and innate to human judges.

It is precisely because human-generated rulemaking most closely aligns with the *Skidmore* Court's imagining of agency actions that any decisions arrived at through pure human decision-making need not undergo a more searching judicial inquiry than that articulated in *Skidmore*. Human language and intellect may be imperfect machines, but these imperfections are core to the system of government created by the Framers.[220] The Constitution's structure assumes that fallible human

---

2025 deadline to submit comments); Establishing a Fixed Time Period of Admission and an Extension of Stay Procedure, 90 Fed. Reg. 42070 (proposed Aug. 28, 2025) (to be codified at 8 C.F.R. pts. 214, 248, 274a), https://www.regulations.gov/docket/ICEB-2025-0001 [https://perma.cc/Z8DC-9P3S] (21,950 comments received by the Sept. 29, 2025 deadline to submit comments).

[218] Skidmore v. Swift & Co., 323 U.S. 134, 140 (1944).

[219] *See, e.g.*, Inadmissibility on Public Charge Grounds, 84 Fed. Reg. 41292, 41316-19 (Aug. 14, 2019) (to be codified at 8 C.F.R. pts. 103, 212, 213, 214, 245, 248) (responding to public comments challenging statutory interpretation).

[220] *See* Loper Bright Enters. v. Raimondo, 603 U.S. 369, 400 (2024) ("As the Framers recognized, ambiguities will inevitably follow from 'the complexity of objects, . . . the imperfection of the human faculties,' and the simple fact that 'no language is so copious as to supply words and phrases for every complex idea.'" (citing THE FEDERALIST NO. 37, at 236)).

judgment, when subject to procedural constraints and public accountability, can serve as a legitimate basis for administrative action.[221] Courts applying *Loper Bright*/*Skidmore* to purely human rulemaking therefore need not grapple with novel questions about algorithmic transparency, training data bias, or machine-generated reasoning; they can instead focus on the traditional factors of thoroughness, validity, and persuasive power that have guided administrative law for eight decades since *Skidmore* was decided in 1944.

b. Adjudications

Immigration adjudication within the Executive Office for Immigration Review ("EOIR"), the sub-agency within the Department of Justice that houses the Immigration Courts and the Board of Immigration Appeals ("BIA"), represents another straightforward application of traditional *Skidmore* analysis. These proceedings rely fundamentally on uniquely human attributes that cannot be mechanically applied or easily replicated through automated processes, and thus they operate within a framework resembling the human-centric model the *Skidmore* Court envisioned in 1944.

In removal proceedings, immigration judges generally rely on their purely human cognition in the exercise of interpretive authority over complex statutory and regulatory schemes. These adjudications involve real-time analysis of witness credibility, application of legal standards to factual scenarios, and interpretive judgments about statutory meaning.[222] The

---

[221] *See* INS v. Chadha, 462 U.S. 919, 953 n.16 (1983) (explaining that administrative action is constrained by enabling statutes enacted through bicameralism and presentment and remains subject to judicial review).

[222] *See, e.g.*, 8 C.F.R. § 1003.10(b) (2026) (describing immigration judges' authority to conduct removal proceedings); 8 C.F.R. § 1240.1 (2026) (establishing procedures for removal proceedings conducted by immigration judges); 8 C.F.R. § 1246.4 (2026) (tasking immigration judges with the authority to receive and examine evidence in proceedings related to the rescission of adjustment of status grants).

[222] *See, e.g.*, 8 C.F.R. § 1003.10(b) (2026) (describing immigration judges' authority to conduct removal proceedings); 8 C.F.R. § 1240.1 (2026)

immigration judge's decision-making process is intended to be transparent and auditable: hearings are conducted on the record, decisions must be explained in writing, and parties retain full opportunity to challenge both factual findings and legal conclusions through the BIA and federal court review.[223] The BIA, however, recognizes that witness/respondent credibility, a concept of significant importance in removal proceedings, is best assessed by human observers in the room rather than from the documents filed on appeal.[224] Similarly, various forms of immigration relief are discretionary. These discretionary delegations require immigration judges to balance adverse factors against the "social and humane considerations" that favor a noncitizen's petition.[225]

This transparency distinguishes EOIR adjudications from an administrative proceeding where AI systems might obscure decision-making processes. When an immigration judge determines whether an asylum applicant meets the "particular

---

(establishing procedures for removal proceedings conducted by immigration judges); 8 C.F.R. § 1246.4 (2026) (tasking immigration judges with the authority to receive and examine evidence in proceedings related to the rescission of adjustment of status grants).

[223] *See* 8 U.S.C. § 1229a (providing for review procedures in removal proceedings); 8 C.F.R. § 1003.1 (2026) (establishing Board of Immigration Appeals review); 8 C.F.R. § 1240.9 (2026) (describing the contents of the record of proceedings).

[224] *In re* A-S-, 21 I. & N. Dec. 1106, 1111 (B.I.A. 1998) ("[W]e emphasize that the Immigration Judge is in the unique position of witnessing the live testimony of the alien at the hearing. Because an appellate body may not as easily review a demeanor finding from a paper record, a credibility finding which is supported by an adverse inference drawn from an alien's demeanor generally should be accorded a high degree of deference." (citations omitted)); *In re* V-T-S-, 21 I. & N. Dec. 792, 796 (B.I.A. 1997) ("The Board ordinarily will not disturb an Immigration Judge's finding concerning the credibility of a witness, since the Immigration Judge has the advantage of observing the alien as he testifies.").

[225] *In re* Edwards, 20 I. & N. Dec. 191, 195 (B.I.A. 1990); *see also In re* C-V-T-, 22 I. & N. Dec. 7, 11 (B.I.A. 1998) ("The [BIA] has long noted both the undesirability and the difficulty, if not impossibility, of defining any standard in discretionary matters . . . which may be applied in a stereotyped manner. Accordingly, there is no inflexible standard for determining who should be granted discretionary relief, and each case must be judged on its own merits." (internal citations and quotations omitted)).

social group" requirement under 8 U.S.C. § 1101(a)(42)(A), that interpretation flows directly from human analysis of statutory text, precedent, and factual circumstances. Rather than require rote application of elements, the "particular social group analysis" requires immigration judges to consider varied cultural and contextual factors to determine whether amorphous requirements of immutability, "particularity," and "social distinction" are satisfied.[226] The resulting decision merits *Skidmore* deference not because it represents perfect algorithmic reasoning, but because it involves a subjectively informed assessment of rules and discretion.

Moreover, the adversarial structure of removal proceedings ensures that interpretive decisions are scrutinized by advocates before they acquire any claim to administrative legitimacy.[227] Unlike *ex parte* rulemakings or internal agency guidance, statutory interpretations emerge from contested proceedings where affected parties have a full opportunity to present contrary interpretations and challenge the judge's reasoning.[228] This process produces the type of thorough consideration and reasoning privileged in *Skidmore*.[229]

2. Partially Automated Decision-Making

Between the two extremes on this continuum lies a range of human decision-making assisted by generative AI technology, wherein the technology is used to support the ultimate human reasoning but does not replace it. Industry leaders have already begun to articulate an idealized near future where AI systems

---

[226] *See, e.g.*, *In re* M-E-V-G-, 26 I. & N. Dec. 227 (B.I.A. 2014) (establishing the current requirements to establish a "particular social group").
[227] *See, e.g.*, 8 C.F.R. § 1240.2 (2026) (describing the role of DHS counsel in removal proceedings).
[228] *See* 8 C.F.R. § 1240.10(a)(1), (4) (2026) (requiring immigration judges to advise respondents of their right to representation and their opportunity to examine and object to evidence, present evidence, and cross-examine witnesses); *Id.* § 1240.11 (2026) (guaranteeing right to representation and presentation of evidence in removal proceedings).
[229] Skidmore v. Swift & Co., 323 U.S. 134, 140 (1944).

work alongside humans to aid decision-making.[230] Whether this idealized future will come to fruition depends on decisions made today.

While retaining human supervision risks retaining inconsistency and human fallibility in systems designed for efficiency and uniformity,[231] this tension reveals a fundamental challenge: the more effective AI systems become, the more difficult it becomes for humans to meaningfully oversee their operations while preserving the benefits that justify AI adoption.[232] This is illustrated below across both rulemaking and adjudication modalities.

a. Rulemaking

Notice-and-comment rulemaking is a modality ripe for automation. This process may require rulemaking agencies to review thousands, or even millions, of public comments in order to identify recurrent arguments and draft responses

---

[230] Amodei, *supra* note 1 (suggesting that future AI technology may be incorporated more fully into the judicial system and democratic institutions with the "potential to *reduce* bias and be fairer for everyone." (emphasis in original)).

[231] Ludwig & Mullainathan, *supra* note 76, at 4 ("Another reason is that in criminal justice settings the algorithm is not the final 'decider' – a human is. Building good algorithms requires understanding how human decisions respond to algorithmic predictions. Algorithm builders too often fail to address these types of technical challenges because they haven't had to. Existing regulations provide weak incentives for those building or buying algorithms, and little ability to police these choices.").

[232] *See* Isabell Steidel & Benedikt Gieger, *The Oversight Paradox: Why Human Control over AI May Be Eroding the Very Competence It Requires*, WORLD ECON. F. (July 2, 2026), https://www.weforum.org/stories/artificial-intelligence/oversight-paradox-human-control-ai [https://perma.cc/EEG5-B7W6] (describing an oversight paradox wherein human supervisors increasingly outsource AI oversight responsibilities to increasingly complicated AI models and thereby erode their own capacity to oversee the technology); Jiaxin Wen et al., *Automated Weak-to-Strong Researcher*, ANTHROPIC (Apr. 2026), https://alignment.anthropic.com/2026/automated-w2s-researcher [https://perma.cc/6AAZ-DNRX] (describing Anthropic's adoption of automated AI supervision approach, "weak-to-strong supervision," for AI systems that are smarter than the humans overseeing them).

when revising or adopting proposed regulations.[233] The advent of mass comment campaigns has exacerbated the procedural demands made of agencies using this process.[234]

Currently, there is little evidence that generative AI is deployed by DHS in its rulemaking, but it could be in the near future. For instance, DHS has disclosed use of AI technology "Relativity One" to review and categorize documents for litigation and FOIA requests.[235] While there is no representation that DHS currently uses this technology in rulemaking, software that categorizes, organizes, and synthesizes large quantities of documents could readily be deployed in informal rulemaking to review and categorize comments. The Department's proposed rules regularly generate substantial news coverage,[236] which—in turn—

---

[233] *Compare, e.g.*, Establishing a Fixed Time Period of Admission and an Extension of Stay Procedure, 90 Fed. Reg. 42070 (proposed Aug. 28, 2025) (to be codified at 8 C.F.R. pts. 214, 248, 274a), https://www.regulations.gov/docket/ICEB-2025-0001 [https://perma.cc/Z8DC-9P3S] (21,950 comments received by the Sept. 29, 2025 deadline to submit comments) (reporting 21,924 comments received), *with* Jacob Kastrenakes, *FCC Received a Total of 3.7 Million Comments on Net Neutrality*, THE VERGE (Sept. 16, 2014), https://www.theverge.com/2014/9/16/6257887/fcc-net-neutrality-3-7-million-comments-made [https://perma.cc/HF4N-YQEG] (reporting that the FCC received approximately 3.7 million comments responding to its proposed rule on net neutrality, an unprecedented amount), *and* Cassandra Handan-Nader, *Do Fake Online Comments Pose a Threat to Regulatory Policymaking? Evidence from Internet Regulation in the United States*, 15 POL'Y & INTERNET 100, 103-04 (2023) (discussing the nearly 24 million comments submitted to the Federal Communications Commission's 2017 "Restoring Internet Freedom" proposed rule).

[234] *See, e.g.*, Admin. Conf. of the U.S., Recommendation 2021-1, Managing Mass, Computer-Generated, and Falsely Attributed Comments, 86 Fed. Reg. 36075 (July 8, 2021).

[235] *Department of Homeland Security Enterprise*, DEP'T OF HOMELAND SEC. (July 15, 2026), https://www.dhs.gov/ai/use-case-inventory/dhs-enterprise [https://perma.cc/A2L7-TS9Y].

[236] *See, e.g.*, Alex Ashley, *Trump's DNA Dragnet: The Law That Turns Us All Into Suspects*, ROLLING STONE (Dec. 21, 2025), https://www.rollingstone.com/politics/politics-features/trump-dna-dragnet-immigration-policy-biometric-surveillance-1235487564/

generates significant public engagement in the form of thousands of comment submissions.[237]

The deployment of AI assistance in rulemaking presents both compelling advantages and legitimate concerns in a judicial deference analysis. On the positive side, AI tools have the potential to enhance the thoroughness of an agency's consideration of public input that might otherwise overwhelm human reviewers. This thoroughness may, in turn, inform the agency's reasoning in a manner that the *Skidmore* Court observed would be a factor justifying deference. [238] For instance, when DHS received over 20,000 comments on a

---

[https://perma.cc/J45Q-2FVX] (reporting on the Collection and Use of Biometrics by USCIS proposed rule); Johanna Alonso, *Institutions, Faculty Oppose Plan to Limit Foreign Students' Time in U.S.*, INSIDE HIGHER ED (Oct. 3, 2025), https://www.insidehighered.com/news/global/international-students-us/2025/10/03/duration-status-rule-prompts-opposition-commenters [https://perma.cc/2JK4-X2PY] (reporting on the Establishing a Fixed Time Period of Admission and an Extension of Stay Procedure proposed rule); Catherine S. Ramirez, *The New Wealth Test for Immigrants Is Un-American*, N.Y. TIMES (Feb. 24, 2020), https://www.nytimes.com/2020/02/24/opinion/trump-immigration-public-charge.html [https://perma.cc/HQL3-63US] (reporting on a previous version of the Public Charge proposed rule).

[237] *E.g.*, Collection and Use of Biometrics by USCIS, 90 Fed. Reg. 49062 (proposed Nov. 3, 2025) (to be codified at 8 C.F.R. pts. 1, 103, 204, 207, 208, 209, 210, 212, 214, 215, 216, 235, 236, 240, 244, 245, 245a, 264, 287, 333, 335), https://www.regulations.gov/document/USCIS-2025-0205-0002 [https://perma.cc/H94P-438D] (6,661 comments received by the Jan. 2, 2026 deadline to submit comments); Public Charge Ground of Inadmissibility, 90 Fed. Reg. 52168 (proposed Nov. 19, 2025) (to be codified at 8 C.F.R. pts. 103, 212), https://www.regulations.gov/docket/USCIS-2025-0304 [https://perma.cc/4MFE-7P98] (8,846 comments received by the Dec. 19, 2025 deadline to submit comments); Establishing a Fixed Time Period of Admission and an Extension of Stay Procedure, 90 Fed. Reg. 42070, https://www.regulations.gov/docket/ICEB-2025-0001 [https://perma.cc/Z8DC-9P3S] (21,950 comments received by the Sept. 29, 2025 deadline to submit comments).

[238] Skidmore v. Swift & Co., 323 U.S. 134, 140 (1944) ("The weight of such a judgment in a particular case will depend upon the *thoroughness evident in its consideration*, the validity of its reasoning, its consistency with earlier and later pronouncements, and all those factors which give it power to persuade, if lacking power to control." (emphasis added)).

proposed rule governing student visa procedures,[239] AI systems could identify patterns, categorize substantive arguments, and flag novel concerns that human reviewers might miss due to the volume of comments.[240]

For this reason, rulemaking that *partially* delegates interpretive authority to AI systems should receive diminished deference, but not *de novo* review. Such a result, however, should only be rewarded with deference to the extent that the human actors within the agency themselves review, consider, and engage with the comments provided. An agency's expert-informed reasoning—which justified the Court's conclusion in *Skidmore*[241]—is embodied in the human employees of the agency, not in generative AI technology. To this end, if human actors within the agency cannot explain their statutory interpretations without reliance on algorithmic analysis, or if AI systems influence the substance of regulatory choices rather than merely organize data, courts should question whether the resulting rules reflect the reasoned, human judgment that *Skidmore* was designed to privilege.[242]

This approach recognizes that AI assistance can enhance administrative procedures, but at the potential expense of

---

[239] Establishing a Fixed Time Period of Admission and an Extension of Stay Procedure, 90 Fed. Reg. 42070, https://www.regulations.gov/docket/ICEB-2025-0001 [https://perma.cc/Z8DC-9P3S] (21,950 comments received by the Sept. 29, 2025 deadline to submit comments).

[240] *See generally* Aida, *supra* note 120.

[241] *Skidmore*, 323 U.S. at 140 ("We consider that the rulings, interpretations and opinions of the Administrator under this Act, while not controlling upon the courts by reason of their authority, do constitute a body of experience and informed judgment to which courts and litigants may properly resort for guidance.").

[242] *Id.* (listing "thoroughness evident in its consideration" and "the validity of its reasoning" as persuasive factors in favor of administrative deference); *see also* Encino Motorcars, LLC v. Navarro, 579 U.S. 211, 221 (2016) ("One of the basic procedural requirements of administrative rulemaking is that an agency must give adequate reasons for its decisions. The agency 'must examine the relevant data and articulate a satisfactory explanation for its action including a rational connection between the facts found and the choice made.'" (quoting Motor Vehicle Mfrs. Ass'n v. State Farm Mut. Auto. Ins. Co., 463 U.S. 29, 43 (1983))).

agency legitimacy and functionality. [243] Courts have traditionally used adherence to statutorily mandated procedures as a proxy for the thoroughness of an agency's reasoning.[244] Here, however, courts would be an integral part of holding federal agencies accountable to this standard by requiring agency disclosure around AI use in generating rules and reviewing agency reasoning to ensure that human actors understand the outcome. Courts applying *Skidmore* to partially AI-generated rules should focus on the human reasoning that underlies the agency's statutory interpretations as doing so would allow them to prevent pretextual decision-making.[245]

b. Adjudications

Partially automated adjudications present more serious concerns for *Skidmore* deference than AI-assisted rulemaking because they directly compromise the human judgment that underlies administrative legitimacy. AI adoption in this scenario can range from using AI to organize and categorize filings to drafting entire legal opinions. In this latter extreme, where immigration judges rely on AI systems to analyze case facts, assess credibility, or recommend decisions, the resulting adjudications would reflect algorithmic logic rather than any human interpretation of applicable legal standards.[246] As these systems process thousands of cases daily, the practical capacity for meaningful human review of each decision diminishes, regardless of formal oversight requirements. Thus, as AI systems are increasingly deployed, assertions of human supervision may become increasingly aspirational. Adjudications require real-time application of legal principles

---

[243] *Encino Motorcars*, 579 U.S. at 220-21 (noting that judicial deference to administrative interpretations is improper where regulations are "procedurally defective").

[244] *Id.*

[245] SEC v. Chenery Corp., 318 U.S. 80, 94 (1943) ("In either event the orderly functioning of the process of review requires that the grounds upon which the administrative agency acted be clearly disclosed and adequately sustained. 'The administrative process will best be vindicated by clarity in its exercise.'" (quoting Phelps Dodge Corp. v. NLRB, 313 U.S. 177, 197 (1941))).

[246] *Skidmore*, 323 U.S. at 140.

to specific factual scenarios, which is precisely the kind of contextual judgment that current AI systems cannot reliably perform.

The current immigration court backlog presents both opportunity and potential pitfalls for the unbridled adoption of AI technologies. As of August 2026, the Immigration Court backlog included nearly 3.1 million pending cases.[247] This deluge of cases presents real problems for the approximately 700 immigration judges around the country whose average docket consists of over 5,000 cases.[248] While U.S. immigration courts have not fully embraced generative AI technology, other non-U.S. jurisdictions have begun experimenting with AI automation to clear case backlogs.[249]

The EOIR's current approach to AI regulation reveals the challenges courts will face when evaluating partially automated adjudications. As of this writing, EOIR has not articulated a blanket prohibition of generative AI, nor have they imposed a mandatory disclosure policy for AI use by practitioners appearing before immigration courts or the BIA. In a policy memorandum to all EOIR staff on August 8, 2025, Acting Director Sirce Owen eschewed the adoption of an agency-wide policy. The memorandum instead allowed individual immigration judges and the Board of Immigration Appeals to adopt their own generative AI rules if they so choose.[250]

---

[247] *Immigration Court Quick Facts*, TRAC IMMIGRATION, https://tracreports.org/immigration/quickfacts/eoir.html [https://perma.cc/8BWT-T4K8]. *See also* Press Release, Exec. Off. for Immigr. Rev., *EOIR Announces Significant Immigration Court Milestones* (Sept. 4, 2025), https://www.justice.gov/eoir/pr/eoir-announces-significant-immigration-court-milestones [https://perma.cc/FC2Z-KSA5] (reporting the EOIR backlog consisted of 3.75 million pending cases).

[248] *Are immigration judges keeping up with rising caseloads?*, USAFACTS (Nov. 6, 2025), https://usafacts.org/articles/are-immigration-judges-keeping-up-with-rising-caseloads/ [https://perma.cc/H59D-BZU5] (reporting 735 immigration judges around the country and an average caseload of 5,286 cases per judge in fiscal year 2024).

[249] Crootof, *supra* note 161, at 233.

[250] Memorandum from Acting Dir. Sirce Owen, Exec. Off. for Immigr. Rev., *Use of Generative Artificial Intelligence in EOIR Proceedings*, PM 25-40, at

The 2025 Owen Memorandum focuses externally on regulating generative AI use by practitioners appearing before EOIR rather than addressing internal usage of the technology.[251] To the extent the policy memorandum discusses use of AI technology by EOIR employees, it simply encourages them to verify the accuracy of the output themselves.[252] This guidance offers an interesting counterpoint to earlier EOIR internal policy that prohibited employees from using "unauthorized AI services" on government-furnished equipment.[253] The internal policy includes a list of prohibited AI websites and implies that it was published in response to unauthorized use of these technologies observed by an EOIR cybersecurity team. [254] The subsequent 2025 Owen Memorandum can only be reconciled with this earlier internal policy if EOIR employees use generative AI on computers other than their workplace equipment, which further highlights the ad hoc nature of current AI governance.

Immigration laws afford decision-makers discretion as a variable in most forms of immigration relief because Congress saw fit to trust human decision-makers with the tools to exercise their judgment and, relatedly, their mercy, when reviewing relevant claims.[255] Any system that involves partially

---

2 (Aug. 8, 2025), https://www.justice.gov/eoir/media/1410621/dl?inline [https://perma.cc/B2WP-BVJK] ("EOIR has neither a blanket prohibition on the use of generative AI in its proceedings nor a mandatory disclosure requirement regarding its use. Nevertheless, nothing in EOIR's rules prohibit [sic] individual adjudicators or courts from adopting standing orders . . . . The BIA, too, may prescribe rules, with the approval of the Director, related to the use of generative AI in pleadings filed with it." (citations omitted)).

[251] *Id.* at 3.

[252] *Id.* ("[B]e vigilant to ensure that decisions are not based on hallucinated or inaccurate case citations and information.").

[253] Matthew Hoppock, *EOIR's AI Memo: Don't Use It on Government Computers*, HOPPOCK L. FIRM (Aug. 29, 2025), https://www.hoppocklawfirm.com/eoirs-ai-memo-dont-use-it-on-government-computers/#comment-3291 [https://perma.cc/4MLM-9G23].

[254] *Id.*

[255] Asylum, a discretionary form of relief, is one of the most heavily sought-

automated decision-making must involve active human operators who are sufficiently engaged in the process to ensure congressional intent is adhered to and constitutional demands are satisfied.[256]

For this reason, appellate courts should extend diminished or no *Skidmore* deference to statutory interpretations emerging from partially automated adjudications depending on the level of administrative reliance on generative AI technology. Unlike purely human decision-making, where courts can assess the thoroughness and validity of agency reasoning, AI-influenced adjudications obscure the actual basis for interpretive conclusions. When an asylum decision by an immigration judge reflects AI analysis of country conditions, witness credibility, or automated legal research, the interpretation may appear to flow from human judgment while depending on algorithmic processes that resist judicial evaluation and, as explored earlier, reason using potentially hallucinated facts.

Deference to agency adjudications should, therefore, depend on whether a human adjudicator retains decisional authority, provides reasons grounded in statutory and regulatory criteria, and offers parties a meaningful opportunity to contest the AI system's contribution.[257] Absent this, the human actors become little more than figureheads within the agency structure.

If an agency official can understand the basis of the model's recommendation, override it with documented reasoning, and a respondent can inspect and rebut the model's core inputs, assumptions, and outputs, some reduced form of deference

---

after forms of relief in removal proceedings. 8 U.S.C. § 1158 (codifying "asylum" as a discretionary form of relief); *Immigration Court Quick Facts*, *supra* note 247 ("At the end of August 2026, out of the total backlog of 3,092,988 cases, 2,278,978 immigrants have already filed formal asylum applications and are now waiting for asylum hearings or decisions in Immigration Court."). *See also* 8 U.S.C. § 1229b (codifying "cancellation of removal" as a discretionary form of relief); 8 U.S.C. § 1255 (codifying "adjustment of status" as a discretionary form of relief).

[256] *See* Crootof, *supra* note 161, at 240-45 (discussing the perils of "hybrid human-AI judicial systems").

[257] *Id.* at 243-44.

may be warranted. If, however, the model is effectively dispositive, opaque, or insulated by vendor secrecy such that the party cannot challenge it, the rationale for deference diminishes and procedural enhancements—disclosure, an evidentiary hearing, or independent review—become necessary.

3. Fully Automated Decision-Making

The final category on the continuum consists of agency statutory interpretations that are fully automated with no human supervision. This represents the most radical departure from traditional administrative practice and poses the gravest threat to the legitimacy principles underlying judicial deference doctrines.

The federal government has already begun contemplating such deployment of generative AI in "high-impact" cases, or situations where outputs serve as "a principal basis for decisions or actions that have legal, material, binding, or significant effect on rights or safety."[258] Critically, these high-impact determinations may occur in the absence of human oversight.[259] This explicit acknowledgment that consequential government decisions might proceed without human supervision represents a fundamental shift from traditional administrative practice, where human judgment was assumed to drive all consequential government action.

Unlike human decision-makers who can be questioned, disciplined, or replaced, AI systems operate according to training data and algorithms that resist meaningful oversight. Courts cannot assess the "thoroughness evident in [the agency's] consideration" or "the validity of its reasoning" when that reasoning occurs within opaque computational processes

---

[258] Memorandum from Russell T. Vought, Dir., Exec. Off. of the President, on *Accelerating Federal Use of AI through Innovation, Governance, and Public Trust*, M-25-21, at 14 (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 [https://perma.cc/Z52W-US5S]. *See also id.* at 19 (definitional section for "High-Impact AI").

[259] *Id.* at 14.

rather than through identifiable human judgment.[260] The administrative authority vested in the human representatives of federal agencies becomes meaningless if they merely serve as human vessels implementing algorithmic outputs without any form of independent analysis.

Full automation of interpretive decisions renders the mechanisms that traditionally ensure agency accountability ineffective. The notice-and-comment process loses its corrective function when agencies cannot explain how public input influenced their final interpretive decisions or explain why they dismissed documented concerns. Moreover, fully automated interpretations violate long-established principles of administrative law that require agencies to justify their actions when issuing their interpretations and prohibit agencies from offering *post hoc* rationalizations.[261] When AI systems

---

[260] Loper Bright Enters. v. Raimondo, 603 U.S. 369, 388 (2024) (quoting Skidmore v. Swift & Co., 323 U.S. 134, 140 (1944)); *see also* Liao & Vaughan, *supra* note 6, at 7 ("In essence, then, *such models can only be probed in a black-box manner,* which may not be sufficient to meet the transparency requirements for stakeholders, and poses challenges for the research community to develop transparency approaches." (emphasis added)); *Machine Rulemaking: Arbitrary and Capricious Review in the Age of AI*, 138 HARV. L. REV. 1821, 1832 (2025) ("Now compare this to the use of an ML model. Although the agency might describe the training data, testing data, and other design decisions, no one—not the agency, the expert witnesses, or the engineers themselves—can describe how real-world variables or concepts relate to the model's outputs. Thus, instead of evaluating the reasonableness of certain real-world assumptions, judges could be faced with inherently mathematical questions like: Is it appropriate for the model to use stochastic gradient descent or root mean squared propagation as its optimization algorithm? Even with expert opinions on both sides, judges may struggle to understand, let alone review these decisions. Compared to judicial review of traditional models, second-order review in the ML context provides less transparency and a weaker check on irrational decision-making.").

[261] Dep't of Homeland Sec. v. Regents of the Univ. of Cal., 140 S. Ct. 1891, 1909-10 (2020) ("An agency must defend its actions based on the reasons it gave when it acted. This is not the case for cutting corners to allow DHS to rely upon reasons absent from its original decision."); Motor Vehicle Mfrs. Ass'n of U.S. v. State Farm Mut. Auto. Ins. Co., 463 U.S. 29, 57 (1983)

generate statutory interpretations, agencies cannot identify the specific legal reasoning that led to particular conclusions, making honest judicial review impossible.

Courts must reject fully automated agency interpretations. Not only is it inappropriate to afford deference to such automated decision-making, but such decision-making also violates the APA's prohibition on agency actions that are "arbitrary, capricious, an abuse of discretion, or otherwise not in accordance with law." [262] Courts have interpreted this provision to require agencies to engage in "reasoned decision-making" when rulemaking or adjudicating.[263] While courts may not "substitute [their] judgment for that of the agency,"[264] they must review an agency's concurrent justification for its actions to ensure there is a "rational connection between the facts found and the choice made."[265] Taken together, this precedent "requires a certain level of transparency from agencies so that courts may effectively review their decisions."[266] This approach recognizes that, under the APA, such interpretations lie beyond the scope of permissible agency action. Since fully automated interpretations lack statutory legitimacy and resist traditional accountability mechanisms, it falls to the courts to ensure that the APA's statutory constraints are observed.

a. Rulemaking

Fully automated rulemaking would entail some form of AI

(extending the *Chenery* prohibition on *post hoc* rationalizations to all forms of agency action, including rulemaking); Burlington Truck Lines, Inc. v. United States, 371 U.S. 156, 168-69 (1962) (reaffirming the *Chenery* prohibition on *post hoc* rationalizations as essential for preserving the separation of powers doctrine); SEC v. Chenery Corp., 318 U.S. 80, 87 (1943) ("The grounds upon which an administrative order must be judged are those upon which the record discloses that its action was based.").

[262] 5 U.S.C. § 706(2)(A).

[263] *Motor Vehicle Mfrs.*, 463 U.S. at 42-43 (citing SEC v. Chenery Corp., 332 U.S. 194, 196 (1947)) (holding that agencies must articulate a satisfactory explanation for their action, whether rulemaking or adjudication, in order to satisfy "arbitrary and capricious" review).

[264] *Id.* at 43.

[265] *Burlington Truck Lines*, 371 U.S. at 168.

[266] *Machine Rulemaking: Arbitrary and Capricious Review*, *supra* note 260, at 1829.

systems generating regulatory interpretations without meaningful human oversight. Agencies are likely to prioritize AI automation for decisions that require repetitive and rote tasks.[267]

The fundamental problem with fully automated rulemaking lies in its opacity. Judges and agencies alike have little insight into the functioning of the AI systems; they must instead rely on the "'second-order' data about the design, training, and testing of the model." [268] This relationship severs any meaningful tie between the agency and a fully automated rule. Attempts by agencies to explain how their AI systems reached their conclusions would be indistinguishable from *post hoc* rationalizations. [269] Courts would be unable to determine

---

[267] *See* Melissa Mortazavi, *Rulemaking Ex Machina*, 117 COLUM. L. REV. ONLINE 202 (2017) (discussing the advent of "natural language processing software" and the allure of its deployment in rulemaking); Rebecca Heilweil, *ChatGPT, Meet DHSChat: Homeland Security Has a New AI Bot*, FEDSCOOP (Dec. 17, 2024), https://fedscoop.com/chatgpt-meet-chatdhs-homeland-security-ai-bot/ [https://perma.cc/YZS8-3FQ6] (quoting then-Director of the Department of Homeland Security's AI Corps Michael Boyce, "[b]y using DHSChat, employees are able to perform routine work more efficiently. This includes summarizing complex documents and reports, generating computer code, and streamlining repetitive tasks like data entry"); De Liban, *supra* note 9, at 30 (reporting that public benefit program administrators use AI to "determine program eligibility and benefit amounts, enforce income and asset limits, detect fraud, and verify identity"); David Klepper, *Gabbard Says AI is Speeding up Intel Work, Including the Release of the JFK Assassination Files*, ASSOCIATED PRESS (June 10, 2025), https://apnews.com/article/gabbard-trump-ai-amazon-intelligence-beca4c4e25581e52de5343244e995e78 [https://perma.cc/N7CN-8QDS] (quoting then-Director of National Intelligence Tulsi Gabbard: "[w]e have been able to do that through the use of AI tools far more quickly than what was done previously—which was to have humans go through and look at every single one of these pages"); Vigliarolo, *supra* note 93 (reporting on Secretary of the Treasury Scott Bessent's testimony before the House Appropriations Committee, wherein he proclaimed AI would compensate for the reduced number of employees who audit tax returns for accuracy).

[268] *Machine Rulemaking: Arbitrary and Capricious Review*, *supra* note 260, at 1823.

[269] *E.g.*, Dep't of Homeland Sec. v. Regents of the Univ. of Cal., 140 S. Ct.

whether the resulting interpretations truly motivated the agency's action or if it satisfied the basic requirements of reasoned decision-making.

While both human decision-makers and generative AI technologies function as "black boxes," only human intelligence allows for the possibility of transparency into its reasoning. Human officials can be questioned about their interpretive choices, required to explain their statutory analysis, and held accountable for regulatory decisions. AI systems, by contrast, operate through pattern recognition and algorithmic processing that resist explanation even by their designers.[270] The complexity of generative AI systems creates the possibility of these platforms acting in unprecedented and unexpected ways; they also may not generate explanations that accurately reflect the AI's "thought" process.[271]

When AI tools draft and implement rules—by filtering which comments are "seen," summarizing technical input, or proposing draft rationales—the resulting regulations fail the requirement of reasoned explanation imposed by *Motor Vehicle Manufacturers v. State Farm*.[272] If agencies treat AI summaries as dispositive, when they do not or cannot preserve the underlying clustering logic or training data, or cannot explain why certain comments were deemed "insignificant," courts should recognize such rulemaking as fundamentally deficient. The Administrative Procedure Act's notice-and-comment requirements become meaningless when agencies

---

1891, 1909-10 (2020) ("An agency must defend its actions based on the reasons it gave when it acted. This is not the case for cutting corners to allow DHS to rely upon reasons absent from its original decision.").

[270] David Freeman Engstrom & Daniel E. Ho, *Algorithmic Accountability in the Administrative State*, 37 YALE J. ON REG. 800, 824 (2020) ("[E]ven their own engineers cannot necessarily understand how the most advanced models arrived at a given result.").

[271] Crootof, *supra* note 161, at 241 ("Depending on the structure of the algorithm, it may not be possible to plumb the 'reasoning' behind a determination to understand where the error occurred. Some forms of AI are simply too complex to disentangle, while the structure of others render them just as opaque to their designers as a human brain." (footnote call number omitted)).

[272] Motor Vehicle Mfrs. Ass'n v. State Farm Mut. Auto. Ins. Co., 463 U.S. 29, 42 (1983).

cannot demonstrate that human officials actually considered public input rather than algorithmic abstractions of that input.

b. Adjudications

Fully automated adjudication also poses severe threats to legal legitimacy and constitutional governance. While agencies have not yet admitted to adopting such practices, resource constraints relative to the output demands incentivize federal agencies to adopt AI automation to replace human labor where possible.[273] This pressure threatens to undermine fundamental principles that have guided American legal development for centuries.

---

[273] *See, e.g.*, Timothy Gardner, *More than 2,600 US Energy Dept Staffers Accept Second Offer to Resign, Sources Say*, REUTERS (Apr. 10, 2025), https://www.reuters.com/world/us/more-than-2600-us-energy-dept-staffers-accept-second-offer-resign-sources-say-2025-04-10/ [https://perma.cc/725E-SCUU] (reporting that approximately 3,817 of the DOE's 17,000 staffers—22.4% of the agency's total workforce—accepted the first two rounds of resignation requests at the agency); Vivian La & Barbara Moran, *Local Scientists, Fisheries and Weather Forecasters Feeling Impact of NOAA Cuts*, WBUR (May 13, 2025), https://www.wbur.org/news/2025/05/13/noaa-cuts-local-impacts-new-england-massachusetts-climate-science-fisheries-weather-forecasts [https://perma.cc/67EN-8V3M] (reporting that the approximately 20% reduction in the National Oceanic and Atmospheric Administration workforce degraded weather forecasts and created chaos for commercial fisheries); Leah Douglas & Nathan Layne, *Staff Losses and Low Morale Are Derailing FEMA Hurricane Preparations, Internal Document Says*, REUTERS (May 15, 2025), https://www.reuters.com/world/us/fema-chief-says-agency-will-raise-state-burden-disasters-50-2025-05-15/ [https://perma.cc/XQ9N-Z32Z] ("[FEMA] lost 2,000 full-time staff, or roughly one third of its total, to terminations and voluntary incentives."); Valerie Volcovici, *US EPA Plans to Cut Staff to 1980s Levels, Dissolve Research Office*, REUTERS (May 2, 2025), https://www.reuters.com/business/world-at-work/us-epa-plans-reduce-staff-1980s-levels-cut-budget-by-300-million-2025-05-02/ [https://perma.cc/AE9R-LNCH] ("The U.S. Environmental Protection Agency announced plans on Friday to slash its budget by $300 million in fiscal year 2026 . . . ."); Russell Contreras, *Nearly 100 Immigration Court Staff Retiring, Resigning Amid Swelling Backlog*, AXIOS (May 7, 2025), https://www.axios.com/2025/03/07/immigration-court-staff-retire-resign-backlog [https://perma.cc/6J7B-HFWL] ("[S]taff reduction will likely add to the historic backlog of cases . . . . The union estimates that the loss of judges and members of judge teams will add 24,000 cases to the courts' backlog in 2025.").

Human actors and machine systems process information and reach conclusions in fundamentally different ways, with AI being particularly ill-suited for the rule application and value balancing often required of judges.[274] AI produces results by detecting patterns in training data and applying preprogrammed rules and algorithms to process information.[275] Conversely, human intelligence that has shaped precedent is "sensitive to context"; it can consider unanticipated, extenuating circumstances as well as shifting social norms in a flexible approach to rule application.[276] We expect judges and agencies alike to interpret and apply legal rules within a dynamic social context.[277]

The rule of law is premised on the principle of *stare decisis*, which balances stability with adaptability to social shifts and new circumstances.[278] For this reason, precedent evolves

---

[274] *See* Crootof, *supra* note 161, at 237-40 (discussing "fundamental differences in human and machine intelligence").

[275] *Id.* at 238 (quoting Harry Surden, *Artificial Intelligence and Law: An Overview*, 35 GA. ST. U. L. REV. 1305, 1308 (2019)).

[276] *Id.*

[277] *Id.* at 237; *see also* Frank Pasquale, *A Rule of Persons, Not Machines: The Limits of Legal Automation*, 87 GEO. WASH. L. REV. 1, 45 (2019) ("Because '[l]aw as a social institution is multifaceted and embedded in particular political systems and traditions,' a technology that 'reduces a legal relationship to a "clear prescription that exists prior to its application and that determines appropriate conduct or legal outcomes"' will be inherently flawed." (quoting Richard H. Fallon, Jr., *"The Rule of Law" as a Concept in Constitutional Discourse*, 97 COLUM. L. REV. 1, 14 (1997))); Kiel Brennan-Marquez, *"Plausible Cause": Explanatory Standards in the Age of Powerful Machines*, 70 VAND. L. REV. 1249, 1298 (2017) ("To say that a problem is best resolved by prudence rather than principle is to express doubt about the possibility of fashioning second-order rules for navigating the collision between first-order values. . . . When that happens, case-specific judgements—as opposed to generalized principles—must carry the day."); Comment, *The Dead Hand of the Common Law*, 27 YALE L.J. 668, 670-71 (1918) ("The judge will not ordinarily lag much behind the mores of society in its forward march; but he dare not advance much ahead of them either. At least, he dare not lay down a new rule in direct conflict with the prevailing opinion of the community.").

[278] Michigan v. Bay Mills Indian Cmty., 572 U.S. 782, 798 (2014) (discussing

gradually and organically as it develops in response to new applications.[279] Fully automated adjudications risk ending this rich tradition of methodical evolution by effectively freezing precedent in amber.[280]

Unlike human adjudicators, AI systems are unable to consider shifting social mores or contextual developments. Precedential cases gain strength and reliance interests over time. [281] Human intelligence is one of perpetual evolution, development, and growth through exposure to new information and experiences in a manner fundamentally distinct from generative AI technology.

This gradual, cautious approach to the evolution of precedent applies to the EOIR as it does to other adjudicatory bodies. The dynamic nature of immigration law requires agencies to adapt interpretations as circumstances change. The Board of Immigration Appeals (BIA) and the Attorney General have overturned earlier statutory interpretations in asylum and cancellation of removal cases on a number of occasions. The evolution of immigration law can reflect

---

stare decisis as a "foundation stone of the rule of law" that is essential to the development of the law in a "principled and intelligible fashion"); Planned Parenthood of Se. Pa. v. Casey, 505 U.S. 833, 854-55 (1992), *overruled on other grounds by* Dobbs v. Jackson Women's Health Org., 597 U.S. 215 (2022) (describing evolving legal principles and changed facts as two of the four factors the Court considers in its *stare decisis* analysis).

[279] *E.g.*, Lawrence v. Texas, 539 U.S. 558, 573 (2003) (considering the evolving nature of state prohibitions on sodomy when overruling Bowers v. Hardwick, 478 U.S. 186 (1986)); Brown v. Bd. of Educ., 347 U.S. 483, 494-95 (1954) (overruling the Court's prior doctrine of "separate but equal"); Crootof, *supra* note 161, at 246-47 (discussing the organic evolution of common law).

[280] Note that there is tension between the principles that undergird our administrative system and the practical implications of fully automated AI systems. While fully automated AI systems would likely have the effect of ossifying interpretations, which would create greater stability in precedent, they would do so at the expense of the democratic accountability of the federal agency.

[281] Crootof, *supra* note 161, at 247.

evolving social mores,[282] but it also frequently reflects changes in the political preferences of an administration.[283] Judges have commented on the pattern of frequent reversals in immigration court precedent.[284] This pattern of "ping-ponging" policy

---

[282] *See, e.g.*, *In re* M-E-V-G-, 26 I. & N. Dec. 227, 227 (B.I.A. 2014) (altering earlier BIA precedent concerning "particular social group" for the purposes of asylum law); *see also In re* Toboso-Alfonso, 20 I. & N. Dec. 819, 823 (B.I.A. 1990) (recognizing sexual orientation constitutes a "particular social group" for the first time); Pitcherskaia v. INS, 118 F.3d 641, 648 (9th Cir. 1997) (overturning BIA precedent and holding that involuntary psychiatric treatment to "cure" the petitioner of homosexuality constitutes persecution).

[283] *E.g.*, *In re* S-S-F-M-, 29 I. & N. Dec. 207, 208 (A.G. 2025) (Attorney General decision overturning the preceding Attorney General's decision and BIA precedent that recognized "married women in Guatemala who are unable to leave their relationship" as a cognizable particular social group for asylum purposes); *In re* M-S-, 27 I. & N. Dec. 509, 510 (A.G. 2019) (overruling *In re* X-K-, 23 I. & N. Dec. 731 (B.I.A. 2005), which allowed asylum-seekers with a positive credible fear determination for persecution or torture to be eligible for a custody redetermination hearing before an immigration judge); *In re* A-B-, 28 I. & N. Dec. 307, 307 (A.G. 2021) (vacating *In re* A-B-, 28 I. & N. Dec. 199 (A.G. 2021), which restricted asylum claims based on domestic or gang violence, and *In re* A-B-, 27 I. & N. Dec. 316, 316 (A.G. 2018), which itself overruled *In re* A-R-C-G-, 26 I. & N. Dec. 338 (B.I.A. 2014), and recognizing domestic violence as a basis for asylum); *In re* A-C-A-A-, 28 I. & N. Dec. 351, 351 (A.G. 2021) (vacating *In re* A-C-A-A-, 28 I. & N. Dec. 84 (A.G. 2020) and directing immigration judges to return to the "longstanding review processes" that the previous Attorney General prohibited); *In re* L-E-A-, 28 I. & N. Dec. 304, 305 (A.G. 2021) (vacating *In re* L-E-A-, 27 I. & N. Dec. 581 (A.G. 2019), reversing BIA findings, abrogating previous training guidance to immigration officers that a family may constitute a particular social group consistent with existing case law, and instructing immigration judges to revert to the "preexisting state of affairs"); *In re* Cruz-Valdez, 28 I. & N. Dec. 326, 326 (A.G. 2021) (vacating *In re* Castro-Tum, 27 I. & N. Dec. 271 (A.G. 2018)).

[284] Edwards v. U.S. Att'y Gen., 97 F.4th 725, 747-48 (11th Cir. 2024) (Jordan, J., concurring) ("What this has meant in practice over the last two decades is that existing immigration precedent is subject to change every four or so years. Indeed, between the last two administrations alone, there have been at least five vacaturs of prior precedential decisions by the BIA or the Attorney General—many of those decisions themselves vacaturs of even earlier precedential decisions. . . . One need not be a legal savant to recognize that this state of affairs is tantamount to chaos.").

appears to only be accelerating in the immigration context.[285] The cyclical nature of these reversals means that legal strategies and case outcomes can shift dramatically based on changes in political administration rather than legal developments.

Fully automated decision-making risks fostering stability at the expense of evolution.[286] As one scholar observes: "As laws, norms, policies, and decision processes are translated into algorithmic programs, they are cemented . . . . [T]he technological, economic, and political edifices within which code is entrenched may be far more rigid."[287] Generative AI programs cannot "automatically update over time." Pattern-recognition AI systems are trained on past data to predict future outcomes, and while they may excel at identifying traits that were once relevant, AI programs cannot easily adapt what once worked to new situations. Instead, they reflect back decisions biased by prior human judgment as to what was once

---

[285] *See* AM. BAR ASS'N, Resolution 121A & Report 2-3 nn.9-14 (Aug. 13, 2019) (noting that "the [Attorney General] certification process has been used, as opposed to rulemaking (or legislative recommendations), to establish not only procedural and docket management policies, but also substantive questions of law governing immigration proceedings that have resulted in reversing longstanding precedential decisions"); *see also* Alberto R. Gonzales & Patrick Glen, *Advancing Executive Branch Immigration Policy Through the Attorney General's Review Authority*, 101 IOWA L. REV. 841, 857-58 (2016) (describing the historical use of the certification authority, including during the George W. Bush administration); Jennifer S. Breen, *Labor, Law Enforcement, and "Normal Times": The Origins of Immigration's Home within the Department of Justice and the Evolution of Attorney General Control over Immigration Adjudications*, 42 U. HAW. L. REV. 1, 58 (2019) (outlining the increasing use of the self-referral mechanism and noting that "[n]ow, self-referral is the *only* way in which the Attorney General asserts his review power and that power is increasingly used to reshape immigration procedure and settled areas of immigration law"). Note that the practice of the BIA overturning earlier statutory interpretations creates significant practical challenges for immigration practitioners and their clients. Legal uncertainty pervades the field as attorneys must constantly monitor for new precedential decisions that could retroactively affect pending cases. This creates substantial strategic complexity in case preparation and advice to clients.

[286] Crootof, *supra* note 161, at 247.

[287] *Id.*

relevant.[288]

For these reasons, courts should consider reliance on AI as a negative factor affecting the overall persuasiveness of an agency's statutory interpretation. Judicial review of administrative decisions should examine the specific AI systems used, the manner of their deployment, and any available documentation of the models' reasoning processes.[289] Agencies seeking deference for AI-assisted interpretations should bear the burden of demonstrating that human officials retained meaningful control over substantive legal conclusions and that those officials can articulate their reasoning without reliance on algorithmic outputs. In cases involving fully automated decision-making, courts should presume that agencies failed to provide a reasoned basis for their actions.

This approach recognizes that some degree of technological assistance may be inevitable in modern adjudication while insisting that interpretive authority remain with accountable human decision-makers. Courts, however, should distinguish between AI tools that enhance human analysis and AI systems that influence substantive legal judgments about statutory meaning or factual conclusions. Only interpretations that emerge from demonstrably human reasoning processes should receive the respect that *Skidmore* deference accords to agency actors.

Conclusion

As federal courts themselves grapple with AI-generated briefs and motions,[290] they cannot ignore the parallel challenge

---

[288] *Id.* at 248.

[289] *See* Loper Bright Enters. v. Raimondo, 603 U.S. 369, 388 (2024) (quoting Skidmore v. Swift & Co., 323 U.S. 134 (1944)) (instructing reviewing courts to assess all relevant factors in evaluating agency interpretations).

[290] *See AI Hallucination Cases*, ARTIFICIAL AUTH. (Oct. 3, 2026), https://www.damiencharlotin.com/hallucinations/ [https://perma.cc/8Y2Z-XXQ8] (database tracking legal decisions in cases where generative AI produced hallucinated content that listed 2,137 cases as of Oct 3, 2026); Jay Adkisson, *AI-Generated Court Submissions Drawing Increasingly More*

of reviewing agency decisions that are increasingly shaped by the same opaque technologies. The question, then, is not *whether* agencies will deploy AI systems, but how courts will respond *when* an agency inevitably relies on AI systems.

The sliding scale deference framework proposed here offers a principled response that preserves both the benefits of AI and the policy objectives of the current deference framework. Under this framework, agencies remain free to use AI tools to *assist* human decision-makers, but they cannot escape accountability for interpretive choices simply by claiming that those choices emerged from automated systems. Courts should treat fully automated interpretations as though they lack any special claim to judicial respect or legitimacy under the APA. By calibrating deference based on human involvement and supervision, courts can distinguish between AI that enhances human judgment and AI that displaces it.[291]

As AI capabilities advance and agency adoption of AI accelerates, the stakes of this analysis will only increase. The administrative state stands at a crossroads wherein technological efficiency, transparency, and accountability must be reconciled rather than traded against one another.

Severe Sanctions, FORBES (Jan. 9, 2026), https://www.forbes.com/sites/jayadkisson/2026/01/09/ai-generated-court-submissions-drawing-increasingly-more-severe-sanctions/ [https://perma.cc/NM76-N869] (describing hallucinated legal citations in submissions as reaching "contagion level").

[291] Willis v. U.S. Bank Nat'l Ass'n, No. 3:25-CV-516-BN, 2025 WL 1408897, at *4 (N.D. Tex. May 15, 2025) ("The use of artificial intelligence must be accompanied by the application of *actual* intelligence in its execution.").