

Artificial Authorship and Judicial Opinions

Richard M. Re*

92 GEO. WASH. L. REV (forthcoming 2024)

Generative AI is already beginning to alter legal practice. If optimistic forecasts prove warranted, how might this technology transform judicial opinions—a genre often viewed as central to the law? This symposium essay attempts to answer that predictive question, which sheds light on present realities. In brief, the provision of opinions will become cheaper and, relatedly, more widely and evenly supplied. Judicial writings will often be zestier, more diverse, and less deliberative. And as the legal system’s economy of persuasive ability is disrupted, courts will engage in a sort of arms race with the public: judges will use artificially enhanced rhetoric to promote their own legitimacy, and the public will become more cynical to avoid being fooled. Paradoxically, a surfeit of persuasive rhetoric could render legal reasoning itself obsolete. In response to these developments, some courts may disallow AI writing tools so that they can continue to claim the legitimacy that flows from authorship. Potential stakes thus include both the fate of legal reason and the future of human participation in the legal system.

| | |
|--|----|
| I. Framing the Inquiry | 4 |
| A. Motivating Premises | 4 |
| B. Varieties of AI Assistance | 8 |
| II. Initial Effects on Authors | 9 |
| A. Quality and Quantity | 10 |
| B. Uniformity and Diversity | 11 |
| C. Authenticity and Accountability | 12 |
| D. Deliberation and Direction | 13 |
| E. Reason and Rhetoric | 16 |
| F. Canonicity and Customization | 17 |
| III. Reactions by Readers | 18 |
| A. The Perfect Adversarial System | 18 |
| B. Artificially Balkanized Readership | 20 |
| C. Rhetoric’s Rise – and Reason’s Demise | 22 |
| IV. Seeking Equilibrium | 25 |
| A. Transitions and Tradition | 25 |
| B. The Limitations of Wisdom | 27 |
| Conclusion | 29 |

* Professor, UVA School of Law. Thanks to Sam Armstrong, Sridevi Dayanandan, and Mary Grace Triplett for excellent research assistance. Thanks also to Minji Doh, Mike Livermore, Alicia Solow-Niederman, and my co-panelists at a George Washington Law School symposium. [Additional grateful thanks to be added]

In September 2023, a British court of appeals judge attracted global media attention by describing his enthusiastic use of generative AI tools to draft court opinions.¹ Lord Justice Colin Birss provided a specific example from his own work:

“I asked ChatGPT, ‘Can you give me a summary of this area of law,’ and it gave me a paragraph. I know what the answer is because I was about to write a paragraph that said that, but it did it for me and I put it in my judgment. It’s there and it’s jolly useful.”²

Based on these sorts of experiences the Lord Justice reached a conclusion about generative AI that is as intuitive as it is simple: “It is useful, and it will be used.”³

Judicial enthusiasm for generative AI may seem premature, as these tools are currently prone to error, bias, and other serious defects.⁴ Yet generative AI has already affected many aspects of life and work, including legal practice, and optimistic assessments foretell greater advances. In time, generative AI tools may become far more reliable and effective, even as they remain fast, easy, and efficient. Just as important, people are learning how best to use the technology.⁵

If optimistic predictions prove warranted, how would generative AI affect judicial opinions—a genre often viewed as the heart of legal practice, if not of the law itself?⁶ This essay aims to answer that question, which calls for an exercise of speculative imagination. As with many thought experiments, the goal is not just to make heroic predictions, but also to shed light on the present we already inhabit.⁷ For instance, legal practice generally takes for granted that persuasive resources are finite and costly, yielding an economy of persuasiveness—yet AI tools will challenge that premise.⁸ So there is something to learn here, even if technological optimism proves unwarranted.

¹ Hibaq Farah, *Court of Appeals Judge Praises ‘Jolly Useful’ ChatGPT After Asking It For Legal Summary*, GUARDIAN (Sept. 15, 2023). When referring to “generative AI” or “AI,” I generally have in mind artificial intelligence systems based on large language models, or LLMs.

² *Id.* (punctuation altered for readability).

³ *Id.* See also Chief Justice Roberts, *2023 Year-End Report on the Federal Judiciary* at 6. For additional examples of judges using ChatGPT, see Juan David Guierrez, *Judges and Magistrates in Peru and Mexico Have ChatGPT Fever*, TECH POLICY PRESS (April 19, 2023).

⁴ For instance, existing tools can generate “hallucinations.” See, e.g., Molly Bohannon, *Lawyer Used ChatGPT In Court—And Cited Fake Cases*, FORBES (June 8, 2023). For a measured prognosis, see Steve Lohr, *A.I. Is Coming For Lawyers, Again*, N.Y. TIMES (April 10, 2023).

⁵ See Rebecca Crootof, Margot E. Kaminski & W. Nicholson Price II, *Humans in the Loop*, 76 VAND. L. REV. 429 (2023) (discussing pitfalls in human-tech integration); see also *infra* notes.

⁶ For theories that center judicial opinions, see JAMES BOYD WHITE, *HERACLES’ BOW: ESSAYS ON THE RHETORIC AND POETICS OF THE LAW* (1985); DAVID A. STRAUSS, *LIVING CONSTITUTION* (2010).

⁷ See text accompanying *infra* note 26. Such efforts may even be fictional. See L. Bennett Capers, *Afrofuturism, Critical Race Theory, and Policing in the Year 2044*, 94 N.Y.U. L. REV. 1, 3–4 (2019) (discussing “futurist legal scholarship”); see, e.g., Derrick Bell, *The Space Traders*, in *FACES AT THE BOTTOM OF THE WELL: THE PERMANENCE OF RACISM* 158 (1992).

⁸ See *infra* Section III.C.

My argument poses two interrelated paradoxes. The first is that a proliferation of legal reasoning may lead to its abandonment. The legal system is already soaked in texts—not just statutes, regulations, and rules, but also case law, briefs, and commentary. Because of AI authorship, we will soon bear witness to an exponential increase in textuality on at least the same scale as the development of text-searchable databases (and perhaps comparable to the advent of published case reports).⁹ This transformation could conceivably bring about a perfect adversarial system, in which dueling AIs efficiently generate the best possible arguments for opposing views, and the legally right answer (or the lack of such an answer) consequently becomes clear.¹⁰ Or it could yield an ever more balkanized discourse, as readers increasingly exalt in the flawless rhetoric of their favorite side while disparaging the contrary opinions of their opponents.

At least as likely, however, is a quite different outcome. An explosion in the supply of effective, cogent verbiage will render persuasive reasoning better and cheaper, regardless of its legal correctness. Even lay readers will know that AI is mainly responsible, that almost anyone (with AI tools) could produce similarly impressive texts, and that an AI-authored opinion could support virtually any conclusion in a contested case. Judicial opinions, and legal reasoning in general, will become even more demystified than they already are. Would-be readers who have come to see the futility of reading judicial opinions will then put down the case reports and attend to other indicia of merit or correctness, such as the identity of the author, the views of trusted commentators, and the likely consequences of court determinations.¹¹ In a legal system overflowing with persuasive reasoning, there might as well be none at all.

The second paradox is partly, but only partly, derivative of the first one: courts will find use of AI tools almost irresistibly attractive—yet these tools threaten the courts’ institutional interests and so will become an object of concerted judicial resistance. The appeal of these tools is largely self-evident, as they make judicial work easier, faster, and more effective. Legal practice aside, moreover, these tools are already becoming staples of society at large. Yet artificial authorship will tend to dissolve judicial authority by democratizing the ability to both produce and understand sophisticated legal prose. Courts might also fear that regulators will want a say in how courts make use of transformative new technologies. So while generative AI may indeed be “jolly useful” for courts, it is also quite perilous.

⁹ See STUART BANNER, *THE DECLINE OF NATURAL LAW: HOW AMERICAN LAWYERS ONCE USED NATURAL LAW AND WHY THEY STOPPED* (2021) (describing the explosion of case reports).

¹⁰ AI is of course altering many aspects of the adversarial process, including research. LexisNexis, for instance, has rolled out an AI-assisted research tool, called Lexis+ AI. One related but distinct use of AI will be to assess empirical propositions of legal import, some of which may be included in a judicial opinion. See, e.g., Yonathan Arbel & David A. Hoffman, *Generative Interpretation*, 99 N.Y.U. L. REV. (forthcoming 2024).

¹¹ As in the 18th Century, the law “would be what it is, not because it has been so reported, but because it has been so decided.” See M. Tiersma, *The Textualization of Precedent*, 82 NOTRE DAME L. REV. 1187 (2007) (quoting the REPORT OF THE LORD CHANCELLOR’S COMMITTEE (1940)).

Those considerations, in addition to the fact that most current judges have no familiarity with these newfangled gizmos, and much intuitive distrust of them, make it plausible to imagine that exclusively or primarily human authorship could become part of the judicial ethic. By abjuring AI authorship, the judiciary could avoid becoming an object of regulation and preserve the impression that justice is a proper subject for human judgment. As AI becomes ubiquitous elsewhere, legal culture might come to view great legal reasoners and writers as akin to chess grandmasters. Sure, a computer could do their job just as well, or much better. But the impressiveness of the human judges' achievement might still remain. The resulting pride of human authorship might delay or curtail AI authorship in some quarters.

So the stakes in the present inquiry are quite large, touching on both the fate of reasoned legal argument and the shape of human participation in the legal system.

I. Framing the Inquiry

Let me start with how I approach the topic of AI and judicial opinions, including some methodological and empirical assumptions.

A. Motivating Premises

The rise of artificial authorship is a widespread phenomenon, affecting private correspondence, government propaganda, and corporate public relations (among other things). But courts have a distinctive—and in some ways heightened—interest in this topic. Unlike other government actors, courts tend to operate *through* their justificatory public statements.¹² The meaning of a judgment often depends on its accompanying opinion (“The case is remanded for proceedings consistent with this opinion”), and precedential rules—the proverbial “holding” of a court—derive much if not all of their content from their surrounding justification. By comparison, legislators and presidents tend to generate either operative directives or rhetorical soundbites—one or the other. Administrative agencies offer a closer example, insofar as their directives may stand or fall depending on their accompanying explanations.¹³ But even then, the directive and justifications remain separated, both conceptually and practically.

Anglo-American courts' merger of directive and justification has not just functional but ethical implications. And by *ethical* I mean to capture both the idea of moral rightness as well as a professional ethos. That is, a judicial opinion is often thought to convey full authority or legitimacy only because (or if) its author has offered an adequate justification. Similarly, the

¹² See Frederick Schauer, *Opinions as Rules*, 62 U. CHI. L. REV. 1455 (1995).

¹³ See generally DAVID FREEMAN ENGSTROM ET AL., GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES (2020); Danielle K. Citron & Ryan Calo, *The Automated Administrative State: A Crisis of Legitimacy*, 70 EMORY L. J. 797 (2021).

judges who produce judicial opinions are often thought to be fulfilling their roles—to be instantiating the character of their occupation—only if they generate adequate justifications. The issue, in other words, isn't whether the decision is justified *at all* or *by anyone*. The issue, at least in part, is whether the deciding court itself has adduced an acceptable justification. Judicial decisions without any accompanying justification can therefore be unsettling and are usually deemed tentative or otherwise peripheral—even though such decisions could be rationalized by one of the parties or by outside commentators.¹⁴

But if authorship is integral to the judiciary, its role is also complex and qualified. Judicial authors often borrow heavily from one another as well as from the filings of parties, submissions by so-called “friends of the court,” and scholarly articles. These forms of borrowing are not always or fully attributed. Further, judges are generally assisted by officially anonymous clerks who “ghostwrite” the vast majority of judicial opinions.¹⁵ Judicial authorship is therefore something of a collective and specialized undertaking. The production of a judicial “voice,” while important to any judge’s professional reputation, may be more akin to an art studio in which a senior virtuoso manages the apprentices’ work before signing it with his own name.¹⁶ More concretely, judicial authorship already makes extensive use of *ratified rationalizations*, that is, justificatory texts that are authored in the first instance by someone other than the judge but then officially endorsed after the fact.

Against this backdrop, one might think that AI tools will simply replace the law clerk, without otherwise disrupting legal practice. But the advent of the law clerk is itself thought to have changed judicial practice in meaningful ways, as did the subsequent rise of email, searchable legal databases, and word processing.¹⁷ New capabilities engender novel choices, altered incentives, and reshuffled power relationships. At the Supreme Court, for instance, justices armed with law clerks and laptops have tended to create opinions that are longer, more citation-studded, and similar to one another.¹⁸ Meanwhile, lower courts with the same resources may be

¹⁴ Any discomfort with unexplained rulings does not prevent them from being commonplace and sometimes justifiable, *see* Mathilde Cohen, *When Judges Have Reasons Not to Give Reasons: A Comparative Law Approach*, 72 WASH. & LEE L. REV. 483 (2015), though they often attract attention as part of critiques. *See, e.g.*, STEPHEN VLADECK, *THE SHADOW DOCKET: HOW THE SUPREME COURT USES STEALTH RULINGS TO AMASS POWER AND UNDERMINE THE REPUBLIC* (2023) (criticizing “shadow docket” rulings in part for often being unexplained); *see also infra* note 32.

¹⁵ *See, e.g.*, Stephen J. Choi and G. Mitu Gulati, *Which Judges Write Their Opinions (And Should We Care)?*, 32 FLA. ST. U. L. REV. 1077 (2005) (“Law clerks are said to draft the majority of opinions for judges.”).

¹⁶ *See* Peter Friedman, *What Is a Judicial Author?*, 62 MERCER L. REV. 519 (2011); *see, e.g.*, Michiel Franken and Jaap van der Veen, *The Signing of Paintings by Rembrandt and His Contemporaries*, in *THE LEIDEN COLLECTION CATALOGUE* (Arthur K. Wheelock Jr. and Lara Yeager-Crasselt, eds.) (2023).

¹⁷ *See, e.g.*, RICHARD POSNER, *THE FEDERAL COURTS: CRISIS AND REFORM* 102 (1985).

¹⁸ *See* Keith Carlson, Michael A. Livermore & Daniel Rockmore, *A Quantitative Analysis of Writing Style on the U.S. Supreme Court*, 93 WASH. U. L. REV. 1461 (2016).

more adept at either avoiding or, if they choose, creating circuit splits for the Supreme Court to review.¹⁹ And analogous technologies have empowered courtwatchers, granting instant access to court rulings as well as the ability to fact check them in real time.²⁰ In earlier eras, by contrast, judges tended to be more personally responsible for their pronounced or published opinions, lower courts had trouble keeping track of what their colleagues were deciding, and the public had limited, delayed access to most court decisions.

To a great extent, artificial authorship is not just inevitable but actual. Thanks to programs like ChatGPT, millions of people are already taking advantage of artificially intelligent editing and drafting. So one might think that we need only look around to discover how new writing tools affect legal practice. Yet both the technology's development and its uses are still very much in flux. Just as important, the law itself is a dynamic system with many opportunities for re-equilibration and adjustment. So when a new technology generates pressure for institutional change over here, the upshot may simply be a new counter-pressure over there, somewhere else within the system. This systemic perspective can help us recognize the importance of technologically driven change in the law without succumbing to either utopianism or doomsaying.²¹ That is, we can anticipate the relevant changes, as well as responses, countermoves, and interventions.

The present inquiry is more focused and modest than many recent discussions of artificial intelligence and the law. “Robot judges” have understandably attracted a great deal of attention as it has become more plausible to imagine chatbots and other algorithmic tools dictating specific decisions presently left to human judgment.²² Bail decisions and the imposition of punishment—both of which can be determined largely if not entirely based on algorithms—are among the most developed and salient examples.²³ While these applications are certainly important, they have too much eclipsed a set of subtler and now seemingly more imminent scenarios involving language, communication, and reason-giving.

¹⁹ See Allison Orr Larsen and Neal Devins, *The Amicus Machine*, 102 VA. L. REV. 1901, 1949 (2016) (“there is at least some support for the claim that circuit splits are less common in a world in which the lower courts have greater access to one another's opinions”).

²⁰ See, e.g., Dan Farber, *More About EPA's Victory*, LEGAL PLANET (April 29, 2014).

²¹ See Edward A. Parson, Richard M. Re, Alicia Solow-Niederman, & Elana Zeide, *Artificial Intelligence in Strategic Context: An Introduction*, SSRN (2019). Cf. ABDI AIDID & BENJAMIN ALARIE, *THE LEGAL SINGULARITY: HOW ARTIFICIAL INTELLIGENCE CAN MAKE LAW RADICALLY BETTER* 1410–41 (2023) (providing an optimistic long-term perspective).

²² See Eugene Volokh, *Chief Justice Robots*, 68 DUKE L. J. 1135 (2018); Benjamin Minhao Chen, Alexander Stremitzer & Kevin Tobia, *Having Your Day in Robot Court*, 36 HARV. J. L. & TECH. 127 (2022); Richard M. Re & Alicia Solow-Niederman, *Developing Artificially Intelligent Justice*, 22 STAN. TECH. L. REV. 242 (2019); Rebecca Crootof, “Cyborg Justice” and the Risk of Technological-Legal Lock-In, 119 COLUM. L. REV. FORUM 233 (2019); Tim Wu, *Will Artificial Intelligence Eat the Law? The Rise of Hybrid Social-Ordering Systems*, 119 COLUM. L. REV. 2001 (2019).

²³ See Alexis Morin-Martel, *Machine Learning in Bail Decisions and Judges' Trustworthiness*, AI & SOC'Y (2023); Jessica M. Eaglin, *Racializing Algorithms*, 111 CAL. L. REV. 753 (2023); Sandra G. Mayson, *Bias In, Bias Out*, 128 YALE L.J. 2218 (2019).

With a human “in the loop,” moreover, many problems with AI judges diminish or disappear.²⁴ There is a clearer basis for political legitimacy, a possibility of empathy and interpositionality, and an assurance of reasonableness deriving from the judge’s everyday life in a flesh and blood society. In effect, the gradual schedule of technological change is forcing us to grapple with cases of mixed human-and-AI decision-making, in all their complexity, before fully resolving or understanding some of the conceptually “easier” (but technologically harder) cases of near-total automation. Still, the complex and easier inquiries are interlinked, not only because our view of the stark scenarios can inform the blurry ones, but also because the judiciary’s actual use of AI today will shape the technology’s development for years to come.

Throughout, I will take for granted certain basic features regarding the technology behind AI authorship. Perhaps most importantly, I generally assume that highly effective AI tools will be widely available and used, including by many courts. This assumption seems quite plausible given the current use of fairly effective AI tools at no charge by millions of people. Still, a dramatic increase in draconian regulation or an abrupt halt to technological progress in this area could disrupt this premise. At the same time, I assume that the returns to AI authorship are both finite and, at some point, diminishing. That is, even the most effective AI-authored text will not be tantamount to mind control or the song of the sirens. And while different AI tools will doubtless exhibit varying degrees of quality, and are already being specially designed for certain goals and clients,²⁵ I assume that these comparative differences will be relatively small compared with the effects of having an AI tool at all.

The point of exploring this topic has more to do with seeing the various possibilities, pressures, and equilibria than making ironclad forecasts or even placing odds on specific outcomes. Imaginative predictions of the type offered here are partly aimed at better understanding what is happening in the present.²⁶ Our goal is not just to get ahead of a technological wave that is already rising above the legal profession, but also to find a fresh way of thinking about how legal writing has long worked, what role it currently plays, and how it might imminently be improved. AI underlines these questions, and may change how we answer them, but the questions mattered long before anyone had heard of ChatGPT.

²⁴ See Crootof, Kaminski & Price, *supra* note 5; Thomas Julius Buocz, *Artificial Intelligence in Court: Legitimacy Problems of AI Assistance in the Judiciary*, 2 COPENHAGEN J. OF L. STUD. 41 (2018); Aziz Z. Huq, *A Right to a Human Decision*, 105 VA. L. REV. 611 (2020); Kiel Brennan-Marquez & Stephen E. Henderson, *Artificial Intelligence and Role-Reversible Judgment*, 109 J. CRIM. L. & CRIMINOLOGY 137 (2019).

²⁵ Some firms are already touting their creation of bespoke LLMs designed specifically for legal services. See Patrick Smith, *Sullivan & Cromwell’s Investments in AI Lead to Discovery, Deposition ‘Assistants’*, Law.com (August 21, 2023).

²⁶ See *supra* note 7 (collecting sources).

B. Varieties of AI Assistance

Artificial authorship is a complex category potentially involving every aspect of resolving a case. As a rough first cut, we might begin by noting that judicial decisions generally involve the selection of case results, precedential rules, or justificatory opinions. And human judges could act independent of AI, receive guidance from AI tools, or delegate the AI effective decisional authority.²⁷ These two dimensions of AI assistance, along with their simplified categories, are outlined in Figure 1 below:

Figure 1: Dimensions and Varieties of AI Assistance

| | Independence | Guidance | Delegation |
|---------|--|---|--|
| Result | Judge identifies the result, independent of AI | AI suggests or assesses a result | AI identifies the result |
| Rule | Judge crafts the rule for an identified result, independent of AI | AI proposes or assesses a rule reaching an identified result | AI crafts the rule for an identified result |
| Opinion | Judge drafts the opinion for an identified rule and/or result, independent of AI | AI drafts or assesses an opinion for an identified rule and/or result | AI drafts the opinion for an identified rule and/or result |

Some scenarios fall neatly into one or another of the above categories. For example, a judge might simply ask an AI which way to rule and then abide by the resulting recommendation (Result Delegation). Or an AI might suggest edits to the penultimate draft of a nearly finished opinion (Opinion Guidance).

Still, these categories are interlinked, mix-and-matchable, and present to various degrees. For instance, the selection of an outcome or rule will naturally influence what kind of opinion is practically available. A judge might opt for Result Independence, then seek Rule Guidance, and finally engage in Opinion Delegation. Or a judge might independently select both the result and the rule (Result and Rule Independence), only to change her mind on both counts during an AI-assisting writing process (Opinion Guidance). In that last scenario, when the experience of writing an opinion suggests a new outcome or rule, judges might say that the initially envisioned

²⁷ See ARTEMUS WARD & DAVID L WEIDEN, *SORCERERS' APPRENTICES: 100 YEARS OF LAW CLERKS AT THE UNITED STATES SUPREME COURT* (2007) (discussing justices' working relationship with clerks generally shifting from a "retention model" to a "delegation model").

opinion “wouldn’t AI-write.”²⁸ An intended instance of Opinion Guidance would thus become an occasion for both Result Guidance and Rule Guidance.

For present purposes, the most important category is Opinion Guidance. AI now seems most immediately able to supply assistance in that domain, generating significant practical effects in the near future. This conclusion is somewhat surprising, as most discussion of AI adjudication has focused on tools bearing on Result Guidance, or even Result Delegation—that is, use of algorithmic tools to furnish technical information bearing, with various degrees of conclusiveness, on discrete case outcomes.²⁹ By comparison, the idea that AI would so quickly be able to engage in the seemingly ultra-nuanced task of drafting or improving sophisticated human expression seemed improbable just a few years ago.³⁰ The rise of AI rhetoricians, moreover, is virtually the opposite of the time-honored trope that robots are, well, robotic. In fact, it seems that AI will be lucid and charming long before it can accurately or reliably make up its own mind.

That said, Opinion Guidance can come in many forms, including every stage or aspect of a judge’s work on an opinion:

- Argumentation: generating possible arguments and counterarguments
- Composition: drafting or proposing revisions to a draft opinion
- Commentary: analyzing a draft opinion’s style, strengths, and weaknesses
- Recommendation: identifying the best way to compose an opinion

Here, too, conceptual distinctions blur in practice, as existing AI tools invite iterative and reflective use. True, the AI will sometimes do far more than simply act as an editor, in that it will compose sentences, paragraphs, and entire sections for human review and further modification. Yet the AI’s first response will often lie far from the final product that is used or published. Generally, then, there is no clear line between the generation of options, the refining of options already on the table, and the selection among options.

II. Initial Effects on Authors

What are the first-order or immediate effects of artificial authorship?

²⁸ Judges sometimes conclude that “an opinion would not write.” Richard A. Posner, *Judges’ Writing Styles (and Do They Matter?)*, 62 U. CHI. L. REV. 1421, 1448 (1995).

²⁹ See earlier notes.

³⁰ See Ajeya Cotra, *Language Models Surprised Us*, PLANNED OBSOLESCENCE (BLOG) (August 29, 2023) (“ML researchers, superforecasters[,] and most others were all surprised by the progress in large language models in 2022 and 2023.”).

A. Quality and Quantity

A judicial author faces a series of tradeoffs that can be framed in economic terms.³¹ Producing a strong opinion takes time, which means less work for other opinions, for other types of judicial work, and for leisure. There is also a maximum amount of effort that a judge can realistically expend over any given period. So when judges expend effort on opinions, they are optimizing along several variables.

Against that backdrop, the most immediate effect of AI tools is that the cost of producing effective writing will decline. In principle, then, any given judge would be able to maintain the same effort as to other forms of judicial work while improving the quality of her opinions. Or, equivalently, the judge could maintain the same quality of opinion but devote more time to other aspects of her work, such as identifying correct results. Viewed from either of these standpoints, AI tools would be an unmitigated improvement to the legal system.

But the answer isn't that simple. When the cost of a good declines, the optimal response is often to consume more of that good while trading down consumption of other goods. In the context of judicial opinions, that might mean spending additional hours making opinions more readable and entertaining, thereby enhancing the judge's reputation, rather than using the time freed up by AI tools to make sure that cases are decided correctly. Writing more opinions, or more persuasive ones, could distract judges from the task of identifying the best outcomes and rules.

In addition, demand for judicial opinions might change in light of their diminished cost. Today, many forms of adjudication lack any opinion whatsoever or come with only short, technical explanations.³² But AI will soon allow almost every determination—from certiorari denials to routine appellate affirmances to trial-court minute orders—to come with an instantly generated, artificially authored explanation.³³ And those explanations could be tailor-made for the specific, legally unsophisticated individuals involved in many retail-level disputes.³⁴ The result would be an explosion of judicial prose. And increases in the total volume of judicial opinions will tend to mitigate improvements in quality.³⁵

³¹ See RICHARD POSNER, *HOW JUDGES THINK* (2008).

³² For data on summary denials of appeal in the federal system, see Merritt E. McAlister, “Downright Indifference”: *Examining Unpublished Decisions in the Federal Courts of Appeals*, 118 MICH. L. REV. 533 (2020).

³³ See text accompanying *infra* note 58. AI authorship will operate differently in different courts. For instance, front-line adjudicators who resolve disputes at scale will likely veer toward Opinion Delegation, whereas apex appellate tribunals will tend toward Opinion Guidance. See also text accompanying *infra* note 59.

³⁴ Generating internal court memos, such as the initial “pool” memos evaluating petitions for certiorari at the Supreme Court, also seem like apt tasks for AI tools.

³⁵ Cf. RICHARD SUSSKIND, *ONLINE COURTS AND THE FUTURE OF JUSTICE* (2019).

Whether the overall effect is viewed as desirable partly depends on how we imagine artificially authored opinions that are relatively low-quality but that would not have existed at all without AI assistance. These opinions could be viewed as a pure gain for the legal system. Something is better than nothing, one might say. From another standpoint, however, these opinions could pull down the average quality of judicial writing in the system, and they risk degrading the public’s view of judicial opinions as a genre. If these opinions are lacking in quality, or are discounted as cheap robot-talk (which in some sense they would be), then they might tarnish the overall perception of the judiciary and its work product.³⁶

B. Uniformity and Diversity

In this enhanced writing environment, weak or mediocre writers will enjoy the greatest relative improvement.³⁷ Almost anyone will be able to produce AI-assisted prose, after all. The writing quality of a standard AI tool will therefore tend to establish a baseline or floor for all minimally competent users of AI tools, even if those users are neither particularly good at using the tools nor talented writers on their own. So while excellent and especially hard-working writers will likely be able to eke out meaningful improvements over the AI-facilitated baseline, the marginal returns on that effort will be both small and diminishing.

It would be tempting to conclude that prose will become more uniform and bland as AI guides all writers to converge on the same efficient and artificial style.³⁸ Yet AI tools also enable authors to express their personalities or adopt idiosyncratic writing personas. Mediocre writers, whether judges or clerks, may be trapped in a familiar style or simply unable to conceive of a creative way to express themselves. Few writers are poets. And judges and clerks are always selected based on many criteria other than writing virtuosity.

But artificial authorship can already convert prose into poetry with the touch of a button. And it can alter the tone of any text, including by assuming the voice of a desired speaker. The AIs, in

³⁶ See Elise Karinshak, Sunny Liu, Joon Park, & Jeffrey Hancock, *Working with AI To Persuade: Examining A Large Language Model’s Ability to Generate Pro-Vaccination Messages*, 7 PROCEEDINGS OF THE ACM ON HUMAN-COMPUTER INTERACTION 1 (2023) (presenting a study indicating that, while ChatGPT created effective messaging, audiences devalued those messages when they knew that the message was created by an AI). One obvious point of comparison is the therapeutic ELIZA chatbot from the 1960s, which people felt comfortable chatting with even though—or because—they knew it was a machine. See JOSEPH WEIZENBAUM, *COMPUTER POWER AND HUMAN REASON: FROM JUDGMENT TO CALCULATION* (1976). Similarly, some people might prefer to be “judged” only by a machine. “It isn’t personal,” they might think.

³⁷ See Jonathan H. Choi & Daniel Schwarcz, *AI Assistance in Legal Analysis: An Empirical Study* (SSRN) (reporting a study in which weaker scorers showed the greatest improvements from AI use).

³⁸ See Vishakh Padmakumar and He He, *Does Writing with Language Models Reduce Content Diversity?*, arXiv preprint arXiv:2309.05196 (2023) (presenting study evidence that writing with “InstructGPT (but not the GPT3) . . . increases the similarity between the writings of different authors and reduces the overall lexical and content diversity”).

other words, will be much more versatile writers than clerks. In some ways, these tools already are superior. To see this, copy a passage of a judicial opinion into ChatGPT and ask it to convert the text into a haiku, sonnet, or hymn. As these capabilities increase, the result may be an increase in rhetorical personality and diversity.

Paradoxically, AI tools may tend to promote both uniformity and panache. Capabilities are a critical determinant of style in part because they limit what is possible. But writers produce for audiences, so consumer demand is often the most important factor. As AI makes it easier to write both clearly and entertainingly, writers will take advantage of those opportunities. And with artificial authorship enabling almost anyone to write more like the most popular writers around, professional standards will rise. Judicial writing could also become stylistically dynamic, even faddish, as jurists instruct their AIs to match the latest trends.

C. Authenticity and Accountability

Will judicial opinions tend to represent authentic expressions of a judge's actual views and personality, as opposed to rationalizations or an assumed persona?

Artificial authorship might seem incompatible with any kind of human authenticity, and in many instances that assumption will be borne out. Again consider the vast number of decisions that presently come with no explanations, such as one-line denials of appeal.³⁹ A court might set up an AI that reviews the record and generates a plausible explanation for the ruling, thereby affording the parties the dignity of an understandable explanation. Yet the automated explanation might fail to capture, or even be unrelated to, the actual basis for decision.⁴⁰

In many other situations, however, artificial authorship will foster or enhance authenticity. Life coaches sometimes talk about making people “the best versions of themselves.” AI tools, like human editors and law clerks, might similarly make opinion writers realize their best selves, achieving what might be called aspirational authenticity. When judges are motivated by a complex legal argument, or else by an elusive, ineffable moral intuition, writing assistance might help the judges both form and communicate their ideas.

Better writing would thus mean better access to authors and their true thought processes. Several consequences follow. Public engagement with legal decision-making would increase.⁴¹ The

³⁹ For an argument that unexplained or summary denials of appeal are objectionable, see McAlister, *supra* note 32.

⁴⁰ See Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 *FORDHAM L. REV.* 1085 (2018). And, again, parties may be put off by the knowledge that they are getting only a robotic explanation. See *supra* note 36.

⁴¹ In this sense, generative AI will foster “demosprudence.” See Lani Guinier, *Foreword: Demosprudence Through Dissent*, 122 *HARV. L. REV.* 4 (2008).

public's better understanding of the judiciary could facilitate political efforts to hold courts accountable. And readers who can understand what judges are talking about might end up being persuaded, leading them to afford courts greater legitimacy. What was once a technocratic and inaccessible profession might become open and relatable.

But what about a judge who uses writing tools to produce empathetic, entertaining, or relatable prose simply to placate readers? Whereas the authentic judge wants an opinion whose form and substance harmonizes with her own considered thoughts, the cynical judge seeks a particular effect on the audience, even though the result misaligns with his own personality or views. A self-absorbed, egotistical jurist might opt to generate opinions that exhibit manufactured empathy.⁴² Any resulting legitimacy would then flow from inauthenticity.

Yet inauthenticity has its benefits. Imagine an angry judge who, to promote a favorable public reputation, always asks a clerk or AI tool to render his opinions calm and courteous and so becomes a popular symbol of collegiality.⁴³ Though subject to personal criticism for his motives, a cynic may still have engaged in beneficial conduct. The persona, though inauthentic and therefore misleading in some respects, might even be for the best. Far from being mere deception, the persona would have become a beacon of virtue.⁴⁴

So both authenticity and its opposite are appealing in different ways and circumstances, a truism discernible in popular aphorisms like “Do as I say, not as I do,” or “Hypocrisy is the tribute vice pays to virtue.” Whether beneficial inauthenticity is well-motivated—that is, whether it is aspirational or cynical—is thus generally of secondary importance. Nor is the reality/perception distinction particularly critical. The key question, which admits of no simple answer, is whether to prioritize accountability based on the judge's public or private self.

D. Deliberation and Direction

AI can and often will improve judicial deliberation. For example, a judge could call upon an AI to brainstorm arguments and counterarguments or to conduct research that parties overlooked. Or the judge could instruct the AI to point out draft prose that has certain problematic features,

⁴² Thanks to Larry Sager for this phrase.

⁴³ Insincere respectfulness could be viewed as a key life skill, including in the judiciary.

⁴⁴ People might come to view judicial affect as too easily constructed to place any faith in it. Fake empathy, for instance, might be so hard to separate from the genuine article that readers assume that no judge is truly empathetic. In essence, readers might become more cynical to avoid being fooled. *See also infra* Section III.A. Judges who sincerely express virtuous ideas might then fail to get credit for them. *Cf.* Robert Chesney and Danielle Keats Citron, *Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security*, 107 CALIF. L. REV. 1753 (2018) (discussing the “liar’s dividend” resulting from deep fakes).

much as a confident editor or intrepid clerk might “push back” on an errant passage.⁴⁵ AI tools may thus increase both the volume and the quality of internal debate among judges.

To some extent, the AI’s deliberative efforts—like a clerk’s—will supplement the adversarial system of litigation. AI’s ability to find missing arguments or details might prove especially useful when parties are under-resourced or otherwise fail to advance the best litigation positions.⁴⁶ AI-facilitated deliberation could therefore render deliberation more fairly distributed among claimants, mitigating the often inegalitarian distribution of attention that results from inequalities in wealth and representative capacity.

Yet opinion-writing could become less deliberative in some respects. The experience of composing a judicial opinion is thought to improve the final product,⁴⁷ and jurists sometimes assert that their intended opinion “wouldn’t write,” leading them to adopt a different and presumably better conclusion. As noted earlier, AI would sometimes help judges realize that a planned opinion just won’t do.⁴⁸ On balance, however, that sort of experience will become rarer as artificial authorship makes it easier to generate a strong opinion with respect to any given issue or viewpoint. Fewer intended dispositions or draft opinions will seem like dead ends, and writing will be easier overall.

Deliberation also bears on the author’s character over time. For example, simply *choosing* to publish respectful opinions is compatible with having a standing order that all clerks and AIs are to write in a respectful style. But it is quite another thing to *enact* respectfulness. Doing so means sitting down and actually writing one courteous opinion after another, generating considered text while suppressing snark and snideness.⁴⁹ The intellectual labor that goes into that effort can transform authors into their persona. Easy writing makes virtuous writing easier to display, but also circumvents deeper processes of transformation.

A similarly formative intellectual labor currently goes into translating sophisticated legal information to lay audiences. Yet, as we have seen, an AI could help explain technical decisions for mass consumption. An interminable hearing transcript riddled with jargon and ending with “Claim denied” could thus be transformed into a compact, readable essay.⁵⁰ The adjudicator would then be freer to live and think exclusively in terms of stylized, professional reason, rather

⁴⁵ See, e.g., Gil Seinfeld, *The Good, the Bad, and the Ugly: Reflections of a Counterclerk*, 114 MICH. L. REV. FIRST IMPRESSIONS 111 (2016).

⁴⁶ On the related possibility of an “AI *Gideon*,” that is, a right to the assistance of artificially intelligent counsel, see *infra* note 66.

⁴⁷ See, e.g., Hon. Roger Traynor, *Some Open Questions on the Work of State Appellate Courts*, 24 U. CHI. L. REV. 211, 218 (1957) (“I have not found a better test for the solution of a case than its articulation in writing, which is thinking at its hardest.”).

⁴⁸ See *supra* note 28.

⁴⁹ Hence the old trope about “building character through hard work.”

⁵⁰ See text accompanying *supra* note 36.

than imagining a lay audience's priorities or engaging in what sometimes travels under the heading of "common sense." Once again, style would shape substance by shaping the author herself.

Finally, AI will reduce the deliberation that stems from interpersonal friction within the writing process. Judges may personally employ AI tools, circumventing clerks. And even when clerks are used, they will know that an AI tool stands ready to execute any judicial instruction. So there will be less chance of a key clerk having her own stubborn style or unchecked point of view.⁵¹ The opinion-writing variability that comes from rotating human clerks would dwindle, allowing each judge's opinions to become predictable and consistent. With clerk-based friction removed or reduced, the judge herself will become more accustomed to the seamless execution of her initial directives—thereby cutting off further deliberation.

Ultimately, AI's effects on deliberation greatly depend on what judges ask AI to do. If judges ask the AI to generate the strongest versions of competing views, especially before the judges make up their minds, then artificial authorship might foster deliberative virtues. Cases might then seem harder than they initially appeared. But if judges instead make quick, knee-jerk decisions before asking the AI to implement them, then the deliberative costs might be substantial. And the judges will often opt for the easier path, given their biases and desire for leisure. Deliberative problems are only exacerbated when the AI itself favors certain views. Systematic or arbitrary skews in an AI tool's recommendations, for instance, can bias the human adjudicator.⁵²

So regulation might be helpful.⁵³ For example, a law governing AI tools might be designed to require or nudge judges to engage with opposing arguments. Or, to similar effect, judges might not be allowed to engage strong AI writing tools until they have already grappled with the opposing arguments set forth through the adversarial system. Advocates, after all, will be coming to court with their own AI-assisted arguments and writing. Already, some courts impose similar rules on themselves, such as by postponing opinion drafting until after oral argument.⁵⁴ Whether judges would welcome the imposition of such a regime is, of course, another matter.

⁵¹ See *supra* note 45.

⁵² See Maurice Jakesch, Advait Bhat, Daniel Buschek, Lior Zalmanson, and Mor Naaman, *Co-writing With Opinionated Language Models Affects Users' Views*, Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (2023). For discussion of possible corrective steps, see Tamara N. Lewis Arredondo, *Incorporating ChatGPT into Human Rights Pedagogy and Research Practices*, *Opinio Juris* (January 2, 2024).

⁵³ For arguments that courts themselves are or will largely be left to deal with the challenges of AI on their own, without new regulation, see David Freeman Engstrom, *The Automated State: A Realist View*, *G.W. L. REV.* (forthcoming 2024); Solow-Niederman, *Do Cases Generate Bad AI Law?* (forthcoming 2024).

⁵⁴ Some courts, by contrast, have a practice of drafting opinions before oral argument. See Daniel Bussell, *Opinions First—Argument Afterwards*, 61 *UCLA L. REV.* 1194 (2014).

E. Reason and Rhetoric

AI tools may be strongest when it comes to the art of rhetoric, and they are bound to get stronger still. By rhetoric, I mean efforts at persuasion, rather than showing what is actually known based upon facts and reason.⁵⁵ Reason itself is often persuasive, but not always—or not as much as other techniques. If you are trying to get someone to eat a particular cereal, fully explaining its health benefits (even if true) might not get the job done. A celebrity testimonial or catchy jingle might be far more effective, even if it lacks any rational basis whatsoever.⁵⁶

Once reason and rhetoric are teased apart, it quickly becomes clear that they frequently come into conflict—and not just because someone with both reason and rhetoric on her side might have to choose which one to pursue. Reason and rhetoric can point in opposite directions—a distressingly common scenario. In the face of rhetoric, should proponents of reason persist in arguing from reason? Ought they try to censor rhetoric?⁵⁷ Or, perhaps most alarmingly, should they arm themselves with reason-free rhetoric of their own?

AI will facilitate rhetoric of all types. When it does so simply by making existing reasons more understandable, the effects are salutary. Clearer explanations, after all, generally foster debate, refinement, and accountability. But what about implicit appeals to preconceptions, prejudice, stereotypes, and allegiances? Or pages of easy reading that gloss over critical logical flaws and legal vulnerabilities? Human writers of course use these techniques today; but, as we have seen, AI will spread and enhance writing capabilities.

The precise way in which AI generates rhetoric will naturally vary by context and, in some cases, may allow for techniques that are not presently realistic. In retail or low-level adjudications, for instance an AI could enable messages that are targeted at the specific individuals in the dispute. Professor Alicia Solow-Niederman and I have given the following example:

Imagine an AI adjudicator whose “opinions” are leavened with personal touches informed by instantaneous social media research. After discovering that a losing party is a Rolling Stones fan, for instance, the AI might comment that “you can’t always get what you want” and then play the hit song’s refrain. The song’s aphoristic familiarity might be both emotionally comforting and cognitively distracting ...⁵⁸

⁵⁵ See ARISTOTLE, RHETORIC 1357a23–33.

⁵⁶ See, e.g., Richard F. Yalch, *Memory in a Jingle Jungle: Music as a Mnemonic Device in Communicating Advertising Slogans*, 76 J. OF APPLIED PSYCH. 76, 268 (1996).

⁵⁷ See Kenji Yoshino, *The City and the Poet*, 114 YALE L.J. 1835 (2005) (discussing Plato’s proposed banishment of the poet, along with contemporary implications).

⁵⁸ Re & Solow-Niederman, *supra* note 22.

As this passage illustrates, effective rhetoric can come at the cost of at least two forms of accuracy relevant to adjudication. First is accuracy in the sense of what actually brought about the adjudicative result—an issue most relevant to what I have referred to above as authenticity and accountability.⁵⁹ Second is accuracy in the sense of whether the result is in fact justified—an issue I am now associating with reason.

In salient or high-level adjudication, AI tools will facilitate feats of rhetoric, as judges produce grand opinions to impress the public, mollify critics, and increase their supporters' admiration. The target audience here is far larger than in most retail adjudication and requires greater finesse, as the desires of various audiences might be in competition. A great appellate opinion might have to be legalistic, breezy, funny, distinguished, and authoritative—all at the same time.

AI is likely to be well-suited to this task. For example, I just asked ChatGPT to generate an essay in favor of the income tax, to revise the essay to be persuasive to libertarians, and then to liven up the essay with humor. In seconds, the AI accomplished all three of these tasks. Future tools may be able to access supplemental information, like polling data or trending social media memes. With such a broad base of training data to draw on, AI would be especially skillful at playing to different audiences simultaneously. And, in special situations, the AI could tailor its work to particular audiences, such as specific judges or swaths of the public.

F. Canonicity and Customization

Judicial opinions are hardly a set medium, much less a fixed genre. Common law rulings even now are largely oral in some nations—and were almost exclusively so until recent centuries made widespread printing feasible.⁶⁰ Early decisions of the U.S. Supreme Court, for instance, were not formally published by the judiciary.⁶¹ Instead, a private individual became designated as “court reporter,” transcribed materials, and sold print copies for profit.⁶² In recent decades, internet posting has now essentially supplanted print publication.

Still-newer technologies will allow for dynamic, interactive judicial opinions. There is no need for a single canonical judicial opinion, after all. Appellate decisions sometimes feature various quite different rationales put forward in separate opinions by concurring judges.⁶³ In some countries, official copies of judicial opinions are released in different languages, and different

⁵⁹ See *supra* Section II.C.

⁶⁰ See *infra* notes 61 and 99.

⁶¹ See Richard J. Lazarus, *The (Non)finality of Supreme Court Opinions*, 128 HARV. L. REV. 540, 552 (2014).

⁶² See *id.*

⁶³ *E.g.*, *Furman v. Georgia*, 408 U.S. 238 (1972).

readers effectively opt in to one or the other language version.⁶⁴ Many courts already publish relatively concise syllabuses to ease case digestion.⁶⁵

In a similar spirit, courts might “publish” a program that interacts with its reader’s preferences, creating a personalized version or presentation of the relevant judicial opinion. Some readers may want a version without citations, or one that is written in plain English. Others might want a version with jokes and flair, while others prefer a “just the facts” narration. And still others might want a dialogic version, complete with a visual avatar, that explains the decision through questions and answers, as a human would in conversation.

In this regime of customization, one might wonder what version—if any—would be canonical, that is, either legally authoritative or dominant in the public eye? A range of options is available, roughly tracking existing practices. Perhaps a base text would be treated as judicially authoritative, whereas derivative, customized versions would be legally irrelevant—much like a syllabus today. Or different versions could be ranked so that an inferior version is trumped whenever it came in conflict with another. Especially popular customizations might inform the meaning or identity of the canonical version. And so on.

III. Reactions by Readers

How will the judiciary’s readers react to artificial authorship? I suggest three basic answers. While in some sense mutually exclusive, each of these reactions is likely to take place at one time or another.

A. *The Perfect Adversarial System*

If AI tools become sufficiently powerful and accessible, then virtually any party could generate a maximally persuasive brief for any proposition. Those without access to an AI might have one appointed by the court—effectively yielding a right to an artificial attorney.⁶⁶ Or, as we have seen, the court itself might use an AI to generate some of the arguments in favor of one side or another, yielding a practice somewhat resembling an inquisitorial system.

⁶⁴ For example, Canada’s federal courts sometimes publish in both English and French. Marie-Ève Hudon, *Bilingualism in Canada’s Court System: The Role of the Federal Government*, available at https://lop.parl.ca/sites/PublicWebsite/default/en_CA/ResearchPublications/201733E#a3.1.1.

⁶⁵ See *United States v. Detroit Timber & Lumber Co.*, 200 U.S. 321, 337 (1906).

⁶⁶ We could even imagine a constitutional right along these lines, yielding an “AI *Gideon*.” Cf. *Gideon v. Wainwright*, 372 U.S. 33 (1963) (establishing that certain indigent criminal defendants have a right to an appointed attorney). For predictions that AI tools will facilitate access to justice in various ways, see, e.g., Benjamin Alarie, Anthony Niblett, Albert H Yoon, *How Artificial Intelligence Will Affect the Practice of Law*, 68 U. TORONTO L.J. 106 (2018); SUSSKIND, *supra* note 35.

AI-generated arguments could even prove to be more legally accurate than similar work by humans. To be sure, the AI would often fall prey to its own biases, including because it could rely on data that is itself shaped by racism, sexism, classism, and other forms of discrimination.⁶⁷ Yet human work, too, is regularly clouded by bias, as well as by self-interest, fatigue, and other flesh-and-blood limitations. Moreover, courts could test their AIs for bias in ways that would be infeasible with respect to human judges or clerks.⁶⁸ We can even imagine parties using AI to point out the problematic biases in one another's AI-generated briefs.

The result could be viewed as a large step toward the perfection of the adversarial system. With each side always putting its best foot forward, the stronger view would become manifest in court. And a similar dynamic could arise in the public square. Consistent with the famous “marketplace of ideas” metaphor,⁶⁹ the public would be well positioned to assess which advocates and judges have the better position. No longer would asymmetries in talent or resources cloud the pursuit of truth and the exercise of reasoned judgment.

The judicial process would adapt accordingly. We might imagine a rule dictating that a certain type of neutrally validated AI must write a dissent for every judicial opinion, whether or not any of the actual judges wish to compose one. If the dissent were written and circulated before the court's decision, this rule would foster deliberation.⁷⁰ And, if the automatically generated AI dissent were ultimately published, this reform would also help hold judges accountable in the event that they are shading the facts or distorting the law.

The value of automated dissent is most evident when there is only a single adjudicator, such as in federal district court or many administrative proceedings. But it would also be significant for multi-member appellate courts where dissent is already possible. For one thing, not just majority judges but dissenters, too, might want to shade or distort the truth in various ways, given their own biases.⁷¹ Perhaps the best grounds for dissent would be embarrassing for the dissenters, given their own past writings or their desire to align with present-day political trends.

⁶⁷ See SAFIYA NOBLE, *ALGORITHMS OF OPPRESSION* (2018); Dorothy E. Roberts, *Digitizing the Carceral State*, 132 HARV. L. REV. 1695 (2019).

⁶⁸ See David Freeman Engstrom, *The Automated State: A Realist View*, G.W. L. REV. (forthcoming 2024); e.g., Hadi Elzayn, Evelyn Smith, Thomas Hertz, Arun Ramesh, Robin Fisher, Daniel E. Ho & Jacob Goldin, *Measuring and Mitigating Racial Disparities in Tax Audits* (2023).

⁶⁹ See *Abrams v. United States*, 250 U.S. 616, 630 (1919) (Holmes, J., dissenting); JOHN STUART MILL, *ON LIBERTY* (1859).

⁷⁰ See, e.g., Hon. Ruth Bader Ginsburg, *Remarks on Writing Separately*, 65 WASH. L. REV. 133, 143 (1990). Similar logic has underlied practices like the Sanhedrin's rule against unanimous verdicts and the “devil's advocate” during the Catholic Church's canonization deliberations.

⁷¹ See Patricia M. Wald, *The Rhetoric of Results and the Results of Rhetoric: Judicial Writing*, 62 U. CHI. L. REV. 1371, 1374 (1995) (reporting that judges compromise to avoid dissents).

Further, unanimous rulings are often more difficult or problematic than they appear, and AI dissents can reveal that fact.⁷² Legislatures and regulators might then be in a better position to react to, or simply override, harmful judicial precedents. More generally, AI dissents could help the public learn whether the judges are, for lack of a better expression, just making it up. And that knowledge, whether viewed as desirable or not, could inform efforts at court reform, either by legitimating or casting doubt on the judiciary's performance.

Because legal indeterminacy is often viewed as undesirable,⁷³ the possibility that AI-authored dissents could both surface and increase it may likewise be problematic. Judges might not want their role in law-creation to be so evident, and onlookers might not want that role to be so robust. To solve that perceived problem, jurists might begin to rigidify their decisional process or promote determinate legal principles.⁷⁴ A proliferation of AI assistance might therefore foster efforts to alter the law itself, whether by judges, legislators, or others.

I have just sketched a doubly optimistic scenario: the technology not only works well, but it also interacts favorably with social practices. It is worth pausing to note the key premise underlying that happy outcome—namely, the assumption that optimal arguments will surface the true state of the law, whether that be a right answer or the lack of any such answer. But a clash of optimal arguments may leave different readers with equal and opposite confidence regarding the right way to resolve a case. Perhaps readers would be persuaded by whichever opinion they happen to have read first—or last. This kind of worry leads to the scenarios that follow.

B. Artificially Balkanized Readership

The foregoing section generally assumed that maximally persuasive opinions would tend to produce homogeneous reactions among readers. That is, the relevant set of readers would more or less all agree that one side or the other had the better of the exchange, or else that there simply was no clear winner to choose. That prospect of consensus is what makes it possible for a perfect adversarial system to draw out the truth, or at least the best available answer. And, in many situations, that consensus could indeed be achieved.

But the prospect of legal indeterminacy tees up another possibility—namely, that different groups of readers would have highly divergent reactions to the same opinion. There might be

⁷² *See id.*

⁷³ Indeterminacy certainly has its virtues, too. *See generally* HRAFN ASGEIRSSON, *THE NATURE AND VALUE OF VAGUENESS IN THE LAW* (2020).

⁷⁴ *Cf. Re & Solow-Niederman, supra* note 22 (arguing that AI creates certain incentives to increase legal determinacy); David Freeman Engstrom and Nora Freeman Engstrom, *Legal Tech and the Litigation Playing Field*, in *LEGAL TECH AND THE FUTURE OF CIVIL JUSTICE* 150–51 (2023) (Engstrom, ed.) (skeptically assessing the prospects of “recalibrating substantive law” to solve a litigation playing tilted by technological asymmetries).

several maximally persuasive opinions available, each pitched to different constituencies. Or the one uniquely superior opinion might be highly partisan in content and appeal. AI authors might be especially adept at exploiting these sorts of fissures in the legal community or the public at large, even without being specifically asked or designed to do so.

AI tools could therefore reflect or even compound reader prejudices, at the expense of truth. Biases aren't evenly distributed across an opinion's readership. So the most persuasive opinion either expanding or constricting gun rights, for instance, might be designed exclusively for conservative or liberal readers, respectively. And there may not be anything that the other side could do to win over those readers, particularly when their priors are being stoked by writing generated by an AI.

On this view, the rise of AI authorship, far from being a perfect adversarial system or truth finding machine, will be a source of sophistry. This is rhetoric of the deceptive sort that Plato warned against, the kind that brings about many false beliefs.⁷⁵ Even worse, the rhetoric in question would be balkanizing, in that it would increase the confidence of warring factions even as it also encouraged them to adopt new partisan views.⁷⁶ The shared beliefs that are traditionally thought necessary to have a legal system might be put under strain.

The most straightforward solution is simply to require that the AI not play to partisan prejudices or other biases. That is, the AI might be asked or trained to avoid playing to groups and instead to write only for a legalistic audience or someone with middle-of-the-road political views.⁷⁷ Because the judges have an incentive to garner whatever support is available, a requirement of this sort might have to be imposed on the judiciary. In essence, some forms of persuasiveness might be ruled out of bounds, at least when AI is concerned.

But that remedy leads to difficult questions about the proper means of regulating adjudication. Shaping AI tools depends on judgments about the proper goals of judicial writing. And those goals are highly contested and varied across judges. Just think of Justice Scalia's norm-busting transformation of judicial writing.⁷⁸ Moreover, forcing the AI to favor consensus will almost necessarily come at the cost of maximal persuasiveness. Sterile rhetoric and on-the-one-hand arguments might come at the expense of creativity, decisiveness, and zest.

⁷⁵ See PLATO, SOPHIST 233. Self-interest of course plays a key role, as people and institution often have an incentive to pretend to knowledge.

⁷⁶ Cf. CASS R. SUNSTEIN, GOING TO EXTREMES: HOW LIKE MINDS UNITE AND DIVIDE (2011).

⁷⁷ Relatedly, courts might try to strengthen norms against rhetoric—though the difficulty of telling reason from rhetoric may undermine this effort. See Nina Varsava, *Professional Irresponsibility and Judicial Opinions*, 5 HOUSTON L. REV. 101 (2021).

⁷⁸ See, e.g., JUSTICE SCALIA: RHETORIC AND THE RULE OF LAW (eds. Brian G. Slocum and Francis J. Mootz III) (2019).

Perhaps regulating the AI tools available to courts would be legitimate for reasons akin to existing legislation constraining who can serve as a clerk.⁷⁹ Yet existing regulations of clerks tend not to focus on their reasoning qualities or legal views. Regulation confining the uses or nature of AI tools might thus resemble unprecedented rules, like a ban on hiring any clerk with a sharp wit or a record of criticizing the political branches. And a prohibition like that would probably encroach on judicial independence and deliberation.

Moreover, efforts to regulate AI might simply be ineffectual. Various rules currently purport to restrain judges' use of private email or their ability to learn classified information that has been leaked to the press and widely published.⁸⁰ Yet the efficacy of those restrictions is easily called into doubt, because the relevant technology is so pervasive. In a similar vein, AI writing tools may soon be so ubiquitous as to be beyond realistic governmental control.

C. Rhetoric's Rise – and Reason's Demise

There is at least one other potential response to the extraordinary proliferation of reasoning and rhetoric that comes with AI authorship: ignore it. When powerful writing becomes ubiquitous, it might stop being quite so powerful. Unable to tell sound reasoning from persuasive rhetoric, readers might stop caring to read at all, preferring instead to evaluate some or even all legal questions based on other qualities. The case's outcome, rule, or author might matter, as contrasted when the reasoning put forward by any judge.

Several factors might conspire to bring this counterintuitive result into reality. One is the effect of legal indeterminacy. What if it turns out that indeterminacy is nearly everywhere?⁸¹ And necessarily so. Consistent with legal realism, critical legal studies, and the Priest-Klein hypothesis,⁸² contested cases, particularly at the appellate level, may almost always be toss ups. Why else would they be litigated? AI authorship could surface that legal uncertainty, thereby revealing the true, flimsy state of the law to professionals and the public alike.

Even more unsettling is the prospect that AI authorship will generate so much rhetoric that it becomes difficult or impossible to discern which side is correct. One version of this worry would focus on writing quantity. Judges and other readers might be inundated with so much purported

⁷⁹ For example, federal clerkships are limited to U.S. citizens and certain lawful residents.

⁸⁰ See Lauren Aratani, *US Supreme Court Justices Use Personal Emails for Work, Report Says*, *GUARDIAN* (Feb 4, 2023).

⁸¹ As fleshed out in the main text, my basic claim here can be framed in terms of either metaphysical indeterminacy (that is, an actual absence of a legally correct answer) or epistemic indeterminacy (that is, a lack of ascertainable knowledge regarding the correct answer).

⁸² See George L. Priest & Benjamin Klein, *The Selection of Disputes for Litigation*, 13 *LEG. STUD.* 1 (1984); Mark Tushnet, *Following the Rules Laid Down: A Critique of Interpretivism and Neutral Principles*, 96 *HARV. L. REV.* 781 (1983); JEROME FRANK, *COURTS ON TRIAL: MYTH AND REALITY IN AMERICAN JUSTICE* (1949).

reasoning—thousand-page opinions, endless motions or briefs, and so forth—that there is no time to grapple with many important matters.⁸³ More fundamentally, strong AI writing could transform what presently seem like easy answers into head-scratching tossups. In many actual cases, for instance, commonplace intuitions might favor one side of a case, and the opposing side might not have the resources or talent to reveal that there is a strong counterargument, perhaps based on arcane legal sources. At present, that kind of case would likely generate a quick and unanimous outcome. But, with the benefit of AI writing or automated dissents, the underdog could fight to a draw.

And by turning many easy cases into hard ones and right outcomes into wrong ones, AI authorship would increase the effective scope of legal indeterminacy. Again calling sophistry to mind, AI authorship would be corrosive, injecting uncertainty or error where none had existed. We could even imagine that virtually every contested case features two equally plausible opinions. Readers, then, would quickly learn that it is a waste of time to look to AI opinions for guidance. What they would find there is nothing more than persuasive pabulum.

To a great extent, of course, legal and popular culture are already quite skeptical of judicial opinions. Almost everyone these days is a legal realist to some extent, and a multitude of legal commentators stands ready to inveigh against the courts at any moment. Perhaps oracular judges in prior eras could credibly claim to be doing legal science or exhibiting profound sagacity, but those days are long gone.⁸⁴ Some people do keep an open mind about major cases until they “read the opinion,” as Justice Barrett has implored.⁸⁵ But the very fact that she had to make that plea suggests that many people, probably most, do not.

AI authorship could take legal culture several steps farther along this path. Today, judicial dissent and media engagement both focus on a relatively small sample of all cases—especially ones with political salience. That limited focus stems partly from resource constraints. It takes time and talent for a human author to digest and debunk a legalistic argument. And, at first blush, judicial opinions often appear plausible and well-reasoned. Legal culture accordingly operates on the assumption that courts do—or at least could—engage in sound reasoning.⁸⁶

⁸³ AI itself might offer a cure here, insofar as AI reading assistants could boil down voluminous prose. Confronted with stacks of AI-generated briefs, the judge might ask her own AI to pick out the facts and arguments that the judge is likely to view as most important—and even compose a draft opinion based on them. In a rhetorical arms race, then, AI readers might effectively counteract AI writers. This culling process would introduce even more separation between human authors and the texts that humans read.

⁸⁴ See Oliver Wendell Holmes, *The Path of the Law*, 10 HARV. L. REV. 457 (1897).

⁸⁵ See Justice Amy Coney Barrett, Remarks at the Reagan Library (April 4, 2022).

⁸⁶ Of course, most opinions find few if any readers. But the opinions that *are* read—whether by the parties, lawyers arguing the next case, or students perusing their casebooks—sustain the general assumption that courts trade in legal reason. Cf. Douglas H. Ginsburg, *Remarks Upon Receiving the Lifetime Service Award of the Georgetown Federalist Society Chapter*, 10 GEO. J.L. & PUB. POL’Y 1, 9 (April 26, 2011) (“When I was new on the court, my colleague . . . told me, tongue in

AI authorship would pose new challenges to that basic assumption. For one thing, automated dissents could puncture the aura of authority that presently accompanies unanimous, business-like rulings.⁸⁷ Commentators often point out how many appellate rulings, such as at the Supreme Court, are unanimous.⁸⁸ And even divided rulings often feature partial or muted dissent. Critical writings generated or facilitated by AI could render those statistics, and the impression they convey, obsolete.

Unanimous rulings would also lose the authority that comes from inscrutability. Legal jargon and technicalities can impress an audience.⁸⁹ Other times, they make it hard for the reader to engage meaningfully with whatever the adjudicator is asserting. AI will greatly reduce both of those effects. Not only would judicial writing instantly be translated into understandable prose, but an AI tool could also answer follow-up questions from the reader.⁹⁰

And then there is the knowledge that an AI probably revised or shadow-authored what the judges have published. By comparison, the public's awareness that a judge often works with clerks (effectively, mini-judges⁹¹) does not challenge the fundamental idea that court personnel are expert—and special. With the AI's help, however, most anyone might endeavor to write their own faux judicial opinions. And those opinions, though generated by amateurs, could meet almost any desired standard of professional competence.

Skeptics of judicial authority routinely lament that courts are making it up, rather than enforcing predetermined legal norms.⁹² But those sorts of critics often struggle to convince their audiences of how judging “really” works. In a world of AI authorship, by contrast, lay readers might immediately understand that a judge could simply ask an AI to generate persuasive arguments for virtually any conclusion, at least in most contested cases. After all, the lay public would itself often be using AI in much the same way as part of their daily lives. What could more thoroughly demystify the courts? Everyone would know not merely that the judges are making it up, but that the AIs are making it up for the judges.

cheek, of course, ‘Remember, the only people who read these opinions are the winning lawyer, the losing lawyer, and the winning lawyer’s mother.’”).

⁸⁷ See text accompanying *supra* note 70.

⁸⁸ See, e.g., Nora Donnelly and Ethan Leib, *The Supreme Court Is Not as Politicized as You May Think*, N.Y. TIMES (Oct. 8, 2023); Devin Dwyer, *Supreme Court Defies Critics With Wave of Unanimous Decisions*, ABC News (June 29, 2021).

⁸⁹ See Richard A. Posner, *Judges’ Writing Styles (And Do They Matter?)*, 62 U. CHI. L. REV. 1421 (1995).

⁹⁰ See *supra* Section II.F.

⁹¹ See Alvin A. Rubin, *Views From the Lower Court*, 23 UCLA L. REV. 448, 456 (1976) (“para-judges”).

⁹² This sort of accusation has long—perhaps always—been a staple of anti-court rhetoric, whether launched from the political left or the right. See, e.g., ROBERT BORK, *TEMPTING OF AMERICA* (1990); ERWIN CHEMERINSKY, *THE CONSERVATIVE ASSAULT ON THE CONSTITUTION* (2010).

What would remain are non-rhetorical proxies for desirable legal outcomes. Having been disillusioned about the nature of judges' work, the lay public and sophisticates alike would generally evaluate case outcomes based on factors unrelated to the persuasive content of any published legal analysis: who voted for the result, what interest groups applauded it, and what do trusted authorities have to say about the ruling's likely consequences? Rule-based adjudication would largely disappear, leaving instead something quite different: an avowed species of policymaking. A judiciary operating in that environment would resemble a legislature or lay jury, rather than a specialized body populated by bureaucrats.

AI authorship thus draws attention to an important premise of real-world legal systems. The existence of legal norms and elites has always depended on there being a scarcity of persuasive resources and arguments. It takes time for lawyers to be trained, their talents honed, and their arguments crafted for each new case. That basic reality means that it is expensive or infeasible to litigate in the teeth of straightforward and accessible law—and costly to litigate even when the law is unclear.⁹³ By undermining these constraints, a surfeit of persuasiveness threatens the effectiveness of legal norms.⁹⁴ The question then arises: would (or should) legal actors attempt to restore an economy of persuasive ability?

IV. Seeking Equilibrium

We have discussed the initial ways in which AI authorship will alter judicial craft and affect readers. But the complexity of the legal system fosters dynamic change, featuring interactions that are iterative, cross-cutting, or reinforcing. And we have seen that the initial effects of AI authorship will generate certain reactions among readers. Judges will quickly anticipate or experience those reactions and so must re-adjust their own behavior. How, at this third stage of the dialectic, will the judiciary account for reader reactions to AI authorship?

A. Transitions and Tradition

The judiciary has a lot to lose from the long-term trends that AI authorship is setting in motion. For reasons we have seen, judicial authority may be undermined by an endless stream of rhetorically effective challenges to their rulings. In addition, the judiciary may become demystified as lay persons realize that they, too, can understand, criticize, and even author sophisticated legal opinions—all with AI assistance. Finally, AI assistance could invite novel forms of regulation, such as restrictions on how an AI assistant is trained.⁹⁵

⁹³ See Frederick Schauer, *Easy Cases*, 58 S. CAL. REV. 399, 401 (1985).

⁹⁴ For different suggestions that the advent of AI might render law as we know it obsolete, see Wu, *supra* note 22; Anthony J. Casey & Anthony Niblett, *The Death of Rules and Standards*, 92 IND. L.J. 1401 (2016).

⁹⁵ Alicia Solow-Niederman, *Administering Artificial Intelligence*, 93 S. CAL. L. REV. 633 (2020).

Because these effects are reactive, they will take longer to materialize than the initial effects discussed earlier. We must therefore pay special attention here to issues of temporality and to the development path that AI authorship is presently following.⁹⁶ For example, additional technological changes—like the invention of general artificial intelligence—could overwhelm the effects of AI authorship, rendering human judging altogether obsolete.⁹⁷ One implication of that contingency is that we have to downgrade our predictive confidence.

Even more importantly, the existence of a development path creates an opportunity for interested parties to resist and shape these relatively long-term effects. So even if AI authorship is bound to become ubiquitous, for many years judges might eschew, or regulations might successfully block, its use. And, if we weaken the premise of inevitability, a concerted effort to swear off AI tools could generate institutional dynamics that entrench traditional writing methods, even as AI authorship elsewhere prevails.

Imagine this. Motivated largely by the anxiety that often accompanies new technologies,⁹⁸ judges might disavow use of certain AI tools. And, to make their disavowal credible, judges might even return to historical practices of judicial decision-making, such as ruling orally from the bench after hearing public arguments without AI involvement.⁹⁹ This sort of policy would afford human judges the legitimacy that comes from authorship. And it would insulate the courts from regulatory interference. While the public debates how to train, limit, and monitor AI tools, judges can remain detached. Rather than becoming exhibits in public controversy, they can pass judgment on novel regulations created elsewhere.

Once controversies subside and AI tools become uncontroversial and standardized, judges might reconsider their choice to abstain from using them. By then, judges would have spent much or all of their lives with these tools, both lay and sophisticated audiences wouldn't be fazed in the slightest by the knowledge of their use, and the case for intrusive regulations would have subsided. Judicial avoidance of AI authorship would be temporary but consequential, essentially allowing the least powerful branch to empower itself.

Or perhaps not. Instead of being sloughed off, the ideal of human authorship could become a permanent, self-conscious component of judges' professional ethic—perhaps even as recognized

⁹⁶ See Re & Solow-Niederman, *supra* note 22 (discussing development paths).

⁹⁷ See *supra* note 14 (collecting sources).

⁹⁸ See generally CALESTOUS JUMA, *INNOVATION AND ITS ENEMIES: WHY PEOPLE RESIST NEW TECHNOLOGIES* (2016). Impulses toward technological panic are, paradoxically, also counteracted by tendencies toward technological utopianism and trust in “science.”

⁹⁹ This approach might be cast as a new rule of orality. See ROBERT J. MARTINEAU, *APPELLATE JUSTICE IN ENGLAND AND THE UNITED STATES: A COMPARATIVE ANALYSIS* 102–03 (1990) (discussing a traditional rule demanding orality, which fostered public scrutiny of judicial work).

by codes of professional conduct. And that gambit could preserve a sense of respect or awe for the work that human judges continue to perform, even as lay people realize that the same work could be automated.¹⁰⁰ Likewise with other intellectual endeavors that have become dominated by computers. Chess players know that computers can beat any human.¹⁰¹ Yet we still marvel at the human accomplishment of grandmasters. A twenty-first century Cardozo might be regarded similarly.¹⁰²

This sort of reform would recreate a degree of scarcity with respect to persuasive resources and arguments. That is, a social norm (or set of norms) would check the unwanted productive capabilities unleashed by rhetorical technologies.¹⁰³ And that development would halt or slow the social processes associated with both the potential demise of legal reasoning and (relatedly) the decline in judiciary's prestige. In this way, the interests of judges tend to align with the preservation of traditional legal norms and practices.

B. The Limitations of Wisdom

So far, I have focused on just two dimensions of opinion-writing success: rhetorical persuasiveness and legal correctness. But other criteria are available. Moral rectitude, for example. AI tools that aim at persuasiveness or lawfulness might tend toward popular moral views, but morality is often thought to be independent of, and possibly contrary to, popular opinion. Pragmatic virtues are similarly beyond an AI's expertise. The AI tool might be great at making proposals *sound* practicable. But would its prescriptions actually be sound?

These additional criteria for success offer alternative ways of evaluating AI tools, and some are more objectively ascertainable (at least in hindsight) than notions of morality or pragmatism. For instance, courts are often thought to care whether their rulings will promote the judiciary's perceived legitimacy.¹⁰⁴ This criterion involves a factual prediction. Yet there is little reason to think that imminent technologies can offer reliable guidance here. The tools simply are not trained on these sorts of empirical prognoses.

¹⁰⁰ See IAN BANKS, LOOK TO WINDWARD 319 (2001).

¹⁰¹ Students of chess—and many other games—routinely watch computers give devastating real-time assessments of humanity's greatest players.

¹⁰² Cardozo is often viewed as a master of judicial rhetoric, even by his critics, *see generally* RICHARD POSNER, CARDOZO: A STUDY IN REPUTATION (1993). Could a well-tuned AI tool one day recreate Cardozo's style—making it available to one and all?

¹⁰³ *See generally* Ryan Calo, *The Scale and the Reactor*, SSRN (arguing for the social contingency of seemingly inevitable technological change).

¹⁰⁴ Versions of this point are sometimes explicit. *See, e.g.*, *Dobbs v. Jackson Wome's Health Organization*, 142 S. Ct. 2228, 2278 (2022) (“The *Casey* plurality was certainly right that it is important for the public to perceive that our decisions are based on principle, and we should make every effort to achieve that objective by issuing opinions that carefully show how a proper understanding of the law leads to the results we reach.”).

So there is a large and multifaceted category of opinion-writing virtues other than either legal correctness or persuasiveness, and those various ideas might be collected under the heading of “wisdom.”¹⁰⁵ This large, multifarious category represents a limitation on the effectiveness and appeal of AI tools. Though they are already designed to *appear* wise, they are not in fact wise, or at least not reliably so.¹⁰⁶ Any wisdom they exhibit is incidental to their pursuit of what amounts to persuasive writing, or writing of the style that has been requested.

The idea of persuasiveness itself also turns out to be complex. Foreseeable AI tools will be much better at writing what is regarded as a good opinion today, rather than predicting what will be most persuasive or laudable years into the future. When wise judges select outcomes, rules, or rationales, they sometimes aim in part to change popular views, rather than simply appealing to them.¹⁰⁷ They are leading their audience, in addition to appealing to it. AI tools can certainly facilitate that effort, but—unless guided by humans—they will usually neglect it.

Moreover, judges will have a strong incentive to emphasize the asserted importance of wisdom. This essay has been concerned with AI authorship, not the more elaborate and seemingly remote prospect of robot judges.¹⁰⁸ Yet the one might help lead to the other. As low-level adjudications increasingly feature automated opinions, people will become more accepting of adjudication without human judgment—or even human participation. The role of judges would decline, unless there remains demand for what AI assertedly cannot offer.¹⁰⁹

The catch, of course, is that humans, too, are often lacking in wisdom.¹¹⁰ They may be motivated primarily by self-interest, fads, and bias. They may believe they are being wise when, in fact, they are self-deluded. And one might think that using AI to appeal to the considered popular views of today is a more reliable path to wisdom than catering to imagined future readers. For any and all of these reasons, judges will hesitate before sacrificing immediate persuasiveness in favor of appeasing the distant view of history. And so they should.

¹⁰⁵ See Norman W. Spaulding, *Is Human Judgment Necessary? Artificial Intelligence, Algorithmic Governance, and the Law*, in THE OXFORD HANDBOOK OF ETHICS OF AI (2020) (Markus D. Dubber, Frank Pasquale & Sunit Das, eds.); see also John Tasioulas, *Ethics of Artificial Intelligence: What It Is and Why We Need It* (Oct. 2023); Cass R. Sunstein, *The Use of Algorithms in Society*, REV. AUSTRIAN ECON. (2023).

¹⁰⁶ This is almost Plato’s definition of sophistry. See *supra* note 75.

¹⁰⁷ Salient examples include major rulings like *Brown v. Board of Education*, 347 U.S. 483 (1954), and *Obergefell v. Hodges*, 575 U.S. 994 (2015).

¹⁰⁸ See text accompanying *supra* note 22.

¹⁰⁹ See FRANK PASQUALE, NEW LAWS OF ROBOTICS: DEFENDING HUMAN EXPERTISE IN THE AGE OF AI 229 (2020) (“As automation advances, we must now [adopt] a commitment to ‘a rule of persons, not machines.’”).

¹¹⁰ See Volokh, *supra* note 14, at 1139 n.12 (“The question is never whether a proposed computer solution is imperfect; it’s whether it’s good enough compared to the alternative.”).

Wisdom, then, will not only limit the use of AI rhetoric but also be limited by it. And that tension will form part of a still larger matrix of competing influences on judicial behavior. Yes, judges will be drawn toward persuasiveness, self-interest, and bias, but they will also care about legal correctness, prudence, and morality. Judges, moreover, will have to account for the dynamic reactions of audiences, the risk of becoming objects of regulation, and—perhaps most of all—the felt need to justify their own human involvement in the legal system.

Conclusion

AI authorship won't be limited to courts, and many of the tradeoffs and dynamics that arise in connection with judges will find parallels elsewhere. Rhetorical craft will crowd out reason, skill levels will quickly become flattened, and human professionals will struggle to preserve the roles long allotted them. True, courts are special. In other domains, authenticity may be less important, reason may more easily be separated from rhetoric, and incumbent professionals may be less able to assert that their role is essentially human. Still, the present study can be viewed as a model or point of departure when analyzing broader social trends.