2024 Forum for State Appellate Court Judges

Artificial Intelligence and the Courts

Saturday, July 20, 2024, Nashville, TN

Judicial Economy in the Age of AI

By Prof. Yonathan A. Arbel*

Executive Summary

In his introduction, Professor Arbel explains how so many Americans struggle to access justice. Reasons range from the cost of paying a lawyer to how difficult it is for those without legal training to know when a lawsuit is appropriate. He then explains how artificial intelligence ("AI") could be an important lifeline at first glance. This sets up a discussion about how AI will dramatically increase the number of lawsuits the legal system must adjudicate each year and the risks this development poses for access to justice.

Part I explains how AI can already do many legal tasks. Professor Arbel discusses research showing that AI tools can spot legal issues and summarize legal documents as well as some lawyers. As time goes on, these tools will only become more sophisticated. Professor Arbel then shows how these tools could lead to more litigation. For example, a person in a dispute with a landlord who has no idea about whether to sue and which claims to bring could explain the situation and ask AI what to do. The relevant tool would inform the person about the legal duties landlords have and what remedies might be available. The AI could then draft a complaint. A person who might have seen litigation as impossible would receive the help they need to begin the process. Given that millions of Americans live with civil justice problems, one can understand why AI tools making it easier for them to sue could lead to a deluge of lawsuits.

In Part II, Professor Arbel explains the concept of legal thermostats and why it might suggest caution about whether AI can truly improve access to justice. In most houses, there is a heating and air conditioning system that keeps the temperature stable even as the weather fluctuates outside. It might heat the home in some seasons and cool it in others. The legal system operates on a similar principle. When there is a concern about too many lawsuits overwhelming the system, judges and

* Professor of Law, Silver Faculty Scholar, University of Alabama, School of Law. The author would like to thank Matt Tokson, Russell Gold, Benjamin McMichael, Mirit Eyal-Cohen, and Heather Elliott. Justin Heydt provided invaluable research assistance. This essay will appear as an invited essay as part of the National Civil Justice Institute Forum Paper.

policymakers have turned to devices like statutes of limitations, exhaustion requirements, and heightened pleading standards to filter out cases and keep the system stable. There is a risk that a similar development will occur if AI leads to an explosion of lawsuits.

In Part III, Professor Arbel considers possible strategies the legal system might use to deal with increased caseloads. These are (1) refraining from taking action while AI continues to develop, (2) preventing litigants from using AI or requiring them to disclose when they have used it, (3) providing additional funding for the legal system, (4) adjusting the procedural thermostat discussed in part II to make initiating and completing litigation more difficult, and (5) carefully integrating AI into the judicial process. Professor Arbel argues that embracing AI and thoughtfully considering how it could help judges and administrators is the most promising and realistic path. Doing so, he argues, can help keep caseloads manageable without creating additional hurdles for the new plaintiffs AI brings into the legal process.

- - - - - - - - - - - - - - -

The majority of legal claims go unvindicated because of access to justice barriers. This entrenched state of affairs is now facing a disruption. Lawyers and non-lawyers alike are adopting artificial intelligence (AI) tools to perform legal tasks, tools that sharply reduce the costs of generating legal materials. There is finally hope that AI might allow many more to access justice.

Paradoxically, what we gain in access to justice we might lose in the delivery of justice. The problem is not that AI tools are ineffective. Indeed, they are even more effective than most realize, affecting every stage of the naming, blaming, and claiming process. The problem is that this change threatens judicial economy.

Historically, judges and legislatures have often met challenges to judicial economy by adjusting "legal thermostats." That concept encompasses ad-hoc adaptations to procedural rules and even substantive doctrines meant to curb the flow of litigation. But these adaptations necessarily imply the shrinking of substantive rights. We run the risk, then, that litigants who finally gain access to justice will find narrow rights and stringent procedures. To avoid this trajectory, I advocate a proactive integrative approach, where AI is used to enhance and scale up the legal process itself. By thoughtfully and carefully incorporating AI tools, we can ensure that we reap the fruits of greater access to justice, even in the face of a rapidly expanding caseload.

Introduction

Most legal disputes are not filed anywhere. While estimates on access to justice barriers are notoriously unreliable, [1] a recent study suggests that about 75 million legal issues go unresolved every year.[2] Around 75 percent of low-income Americans suffer significant civil legal issues, but 92 percent of these problems receive little to no legal aid.[3] One commentator estimates that 100 million Americans live with "civil justice problems," many of which affect their "basic human needs."[4]

The barriers to justice are legion, but most can be expressed in terms of cost.[5] Lawyers charge an average of $327 per hour,[6] with common disputes costing between $2,754 and $6,370.[7] On the other side of the cost spectrum, commercial actors will spend roughly $2 million in outside legal fees to litigate, from start to finish.[8] The sheer size of the investment required to close the gap bedevils attempts to resolve access to justice problems. Even doubling legal aid budgets has done little to narrow the gap.[9] 1.8 million people are turned down annually due to resource

[1] *See generally* Rebecca L. Sandefur, *Paying Down the Civil Justice Data Deficit: Leveraging Existing National Data Collection*, 68 S.C. L. REV. 295 (2016) ("In the arena of civil justice, we face a severe data deficit").

[2] *Justice Needs and Satisfaction in the United States of America*, IAALS (Sept. 1, 2021), https://iaals.du.edu/sites/default/files/documents/publications/justice-needs-and-satisfaction-us.pdf. [hereinafter *Justice Needs*] The authors also report issues resolved in an unfair manner, which amounts to a total of 120 million per annum. For comparison, one estimate considers that 100 million cases are handled by state courts every year. *State of the State Courts: 2022 Presentation*, NCSC (2022), https://www.ncsc.org/__data/assets/pdf_file/0019/85204/SSC_2022_Presentation.pdf (last accessed June 19, 2024).

[3] *Justice Gap Report 2022*, Legal Services Corporation, https://lsc-live.app.box.com/s/oi1atcgn8xmvofc70aildz3bhg5p0zn5 (last accessed June 19, 2024).

[4] Sandefur, *supra* note 1, at 446.

[5] *See generally* DEBORAH RHODE, ACCESS TO JUSTICE (2004). https://pubs.aeaweb.org/doi/pdfplus/10.1257/jel.20201330 ("The principal reason that so few individuals and small businesses avail themselves of legal services is cost and availability."); *see also* Gillian K. Hadfield, *Higher Demand, Lower Supply? A Comparative Assessment of the Legal Resource Landscape for Ordinary Americans*, 37 FORDHAM URB. L.J. 129 (2010) (noting that access to justice affects not just poorer Americans but also middle America). On sociolegal barriers, *see* discussion *infra* Part I.3.

[6] Ashley Merriman, *What Does Hiring a Lawyer Cost?* U.S. NEWS & WORLD. REP. (Mar. 1, 2024, 9:21 AM), https://law.usnews.com/law-firms/advice/articles/what-does-hiring-a-lawyer-cost#:~:text=The%20majority%20of%20attorneys%20charge,The%20size%20of%20their%20firm.

[7] *See Justice Needs*, *supra* note 2, at 47.

[8] *Litigation Cost Survey of Major Companies*, U.S. Courts, https://www.uscourts.gov/sites/default/files/litigation_cost_survey_of_major_companies_0.pdf p. 14

[9] According to the LSC data, between 2013-2022, total funding for legal aid has increased (inflation adjusted) from $1 billion to $1.76 billion. *See The Justice Gap: The Unmet Civil Legal Needs of Low-Income Americans*, Legal Services Corporation P. 11 (2022), https://lsc-live.app.box.com/s/h2bajpr3gps4s4a1iio6fwiddhmu1nwb; Nora Freeman Engstrom & David Freeman Engstrom, *The Making of the Access-to-Justice Crisis*, 75 STAN. L. REV. ONLINE 1 (2023). ("even a vast increase over current commitments would barely dent the current crisis").

constraints.[10] To put this in perspective, the rate of legal aid lawyers to eligible clients is 1 to 15,625.[11]

Recently, Nora and David Freeman Engstrom have sought to center the problem of access to justice around legal tech.[12] While others have already noted legal tech as a barrier,[13] they draw on the debt collection litigation literature to fashion a somewhat different argument.[14] In their view, the asymmetry in power is due to asymmetry in legal tech adoption patterns. While firms zealously adopt legal tech, they only see "anemic adoption" by individuals, relative to more zealous adoption by firms.[15] In particular, they claim that large firms systemize and automate litigation, whereas individuals are still reliant on "analog tools."[16] While this argument is too strong to be true, it does have a kernel of truth to it.[17] Or at least it used to.

We are now witnessing a sea change in the patterns of technological adoption. Most are by now familiar with the occasional news story of a hapless lawyer using AI to comedically bad outcomes.[18] The narrative involves a work-shy lawyer submitting an AI-generated and hallucination-riddled brief to an exasperated judge, who then admonishes and sanctions the lawyer. Such widespread stories seem to draw their memetic power from commonplace

[10] *Final FY2025 Budget Request*, Legal Services Corporation (2024), https://lsc-live.app.box.com/s/oi1atcgn8xmvofc70aildz3bhg5p0zn5

[11] Hanna Kozlowska, *There's a Devastating Shortage of Lawyers in the US Who Can Help the Poor with Eviction or Child Custody Cases*, QUARTZ (May 12, 2016), https://qz.com/681971/for-every-10000-poor-people-in-the-united-states-theres-less-than-1-lawyer-who-can-help-them-with-an-eviction-or-child-custody-case/

[12] Nora Freeman Engstrom & David Freeman Engstrom, *The Making of the Access-to-Justice Crisis*, 75 STAN. L. REV. ONLINE 1 (2023) [hereinafter *A2J Crisis*]. However, *see* Gillian K. Hadfield, *Legal Markets*, 60 J. ECON. LIT. 1264 (2022), https://doi.org/10.1257/jel.20201330 (arguing that regulation favors traditional lawyering across the board at the expense of legal tech).

[13] Gillian K. Hadfield, *Legal Markets*, 60 J. ECON. LIT. 1264 (2022), https://doi.org/10.1257/jel.20201330

[14] Yonathan A. Arbel, *Adminization*: *Gatekeeping Consumer Contracts,* 71 VAND. L. REV. 121 (2018) (discussing robo-signing and other problematic creditor practices in debt collection cases and offering administrative-technological solutions); Daniel Wilf-Townsend, *Assembly-Line Plaintiffs*, 135 HARV. L. REV. 1704, 1708 (2022). ("Assembly-line plaintiffs show no sign of slowing down. Because of both the increases in consumer debt and the improvements in their litigation technology, they continue to grow").

[15] *See A2J Crisis supra* note 12, at 1. This asymmetry is also discussed in Yonathan A. Arbel & Roy Shapira, *Theory of the Nudnik: The Future of Consumer Activism and What We Can Do to Stop It*, 73 VAND. L. REV. 929 (2020) (focusing on the concern that firms employ advanced tools to defang litigation-prone consumers at very early stages of their claiming process).

[16] *See* Engstrom and Engstrom, *supra* note 12.

[17] Most litigants rely on the internet and other digital tools to amass information, communicate about it, and draft and file litigation. *See e.g.*, Margaret Hagan, *Data on People's Reliance on the Internet for Legal Problems*, A BETTER LEGAL INTERNET (Nov. 2, 2022), http://betterinternet.law.stanford.edu/2022/11/02/data-on-peoples-reliance-on-the-internet-for-legal-problems/; *See also* Benjamin H. Barton, *The Future of American Legal Tech: Regulation, Culture, Markets*, *in* LEGAL TECH AND THE FUTURE OF CIVIL JUSTICE 21, 29 (David Freeman Engstrom ed., 2023) ("Nor has legal aid shied away from using technology to forward its mission).

[18] *See e.g.,* Benjamin Weiser, *Here's What Happens When Your Lawyer Uses ChatGPT*, N.Y. TIMES (May 27, 2023), https://www.nytimes.com; Molly Bohannon, *Breaking: Lawyer Used ChatGPT In Court—And Cited Fake Cases. A Judge Is Considering Sanctions*, FORBES (May 12, 2023), https://www.forbes.com.

Shakespearean perceptions of our profession. Incidentally, they also reify an elitist notion that only artisanal lawyering is real lawyering. And perhaps most alluring, they affirm a comforting thought: that down to brass tacks, AI is but a cold machine that will not be able to usurp our jobs.

Reassuring and entertaining as such surface themes are, they also distract from the broader reality that they unwittingly reveal. These stories reveal how AI is being deployed in practice, with two surprising patterns. First, they are being adopted even by small law firms who, at least traditionally, are rarely early adopters of cutting-edge of technologies. Second, they are being adopted *despite* broad knowledge that these tools are imperfect. The point is that even if these tools are only sometimes reliable, they are always convenient. And this convenience and accessibility seems to drive many end-users.

The expected outcome of democratizing litigation technology is a sharp pruning of the cost of producing legal materials.[19] As such, the technology presents a heavyweight contender to existing recalcitrant barriers to justice, leading to a litigation boom by those currently denied access to justice. And while our first instinct might be to celebrate the dismantling of access to justice barriers, realism about judicial economy cautions great care. What we must ask ourselves is whether a legal system already critiqued for being clogged and dilatory, a system whose judges are overworked and under-resourced, will be capable of handling the impending AI boom in litigation?[20] What changes will be made to our laws, rules, and standards to accommodate such a spike? Would we find ourselves, at the end of the day, with a system with a truly greater degree of access?

The goal of this Essay is to sound the alarm about judicial economy in the age of AI, consider central implications, and then offer some constructive mitigation steps. The Essay is organized around three principal contributions.

First, the Essay argues that AI can bring about a litigation boom. Its size is commensurate with the access to justice gap, if not larger. Existing estimates suggest that there is a considerable volume of unmet demand for legal services.[21] I argue, drawing on legal sociology, that these estimates likely understate the true potential.[22] Beyond visible barriers like court and lawyer fees, sociolegal literature suggests that there are much less visible barriers at very early stages. The naming-blaming-claiming model of litigation suggests a pyramid filtering model that prevents many individuals from even thinking about their accidents as legal matters.[23] AI Assistant can assist with

[19] For cost comparisons between human lawyers and state of the art AI models, *see infra* p. 8 and note 31. The point here is static, but there are important dynamic effects, given that costs will decline across the industry.

[20] *See Justice Delayed*, THE ECONOMIST (July 13, 2023) https://www.economist.com/united-states/2023/07/13/judge-and-staff-shortages-are-leaving-americans-in-limbo

[21] *See* LSC data *supra* note 9.

[22] *See* discussion *infra* Part I.3.

[23] The model was developed by William L. F. Felstiner, Richard L. Abel & Austin Sarat, *The Emergence and Transformation of Disputes: Naming, Blaming, Claiming...*, 15 LAW & SOC'Y REV. 631 (1980). It has since become a

these pent-up claims by articulating matters in legally cognizable terms, thereby unearthing an even larger volume of cases. Moreover, existing estimates predominantly focus on unaddressed meritorious claims. However, the same filtering mechanisms that obstruct access to justice also serve beneficial purposes by excluding abusive litigation aimed at harassing individuals with trumped-up charges.[24]

Second, the Essay draws on control theory to consider the implications of a potential AI litigation boom.[25] The entire equilibrium of judicial economy rests on a balance between litigation patterns and judicial resources. One repeated lesson from legal history is that technological and social shocks that threaten judicial economy are met with adjustments of various procedural and substantive doctrines. Even though these doctrines and ostensibly about substantive and procedural rights, they double as legal "thermostats." This effect can be broad and deep. Orin Kerr famously argued that the entire body of Fourth Amendment law, often seen as erratic and "embarrassing,"[26] can be rationalized as a series of "equilibrium adjustments" the courts make in response to new technologies. Here, I generalize this insight and provide illustrations of legal thermostats used across the justice system.

The upshot is that by trying to achieve homeostasis, judges may feel compelled to adjust the thermostats that are at their disposal. They would reach out, by necessity, to procedural and substantive rights. They would be pressured to require, perhaps, more demanding standards of proof, or may require more exacting evidence, or may expand the scope of de-minimis. The degree of thermostat adjustment may be so large that, from the viewpoint of any individual litigant, there would be no sense of progress. They would overcome initial barriers only to crash on ever more limited rights. If we stay the course, it seems that we might squander the opportunity to make a real dent in the access to justice problem by simply reshuffling it.

The third and most practical contribution lies in considering the menu of reactions judges and judicial administrators can make to take advantage of this specific moment. The proposed course of action involves a proactive approach that works to integrate AI into the judicial process itself. This will allow the system to scale up and meet the challenge, without compromising the substantive rights of litigants. Grounding the case for judicial integration in the problematic nature of the realistic alternatives helps motivate adoption even if AI tools are imperfect. Doing so proactively today will help mitigate the harms and ensure responsible adoption.

mainstay of socio-legal analysis.

[24] Paul Ohm and Brett Frischmann developed a framework for thinking about the positive effects of friction as tools of governance, and many of litigation barriers can be conceived along similar lines. Paul Ohm & Brett Frischmann, *Governance Seams*, 36 HARV. J.L. & TECH. (forthcoming 2024)

[25] *See infra* Part II. Control theory is devoted, loosely speaking, to the study of maintaining desired states in dynamic systems. Home thermostats are a common example of tools used by control theory to maintain temperature equilibrium in light of changing outside temperature.

[26] Orin S. Kerr, *An Equilibrium-Adjustment Theory of the Fourth Amendment*, 125 HARV. L. REV. 476 (2011)

1. The AI Litigation Boom

How big is the access to justice gap and what chances do advanced AI systems have to put a dent in it?

This Part opens by first evaluating the technical skills of current-generation AI systems to establish that they can perform many legal tasks *adequately*. Obviously, "adequately" is the load-bearing part of the sentence, but part of the goal here is to show that it covers a fairly broad range of legal capabilities.

It then considers the uptake patterns among end users and the size of the access to justice gap to provide a qualitative and semi-quantitative sense of its size. The combination of cheap but capable AI systems with this large gap leads to the expectation of an AI litigation boom effect in the coming years.

1.1. How Good are AI Lawyering Skills?

Any sufficiently advanced technology can appear indistinguishable from magic.[27] In practice, much commentary on AI seems to fall into this trap, leading commentators down one of two erroneous paths. Either believing in AI omnipotence (AI can do *everything*) or in AI as a cheap magic trick (AI can't do *anything*). In reality, AI tools are a mix. The goal of this section is to avoid a simplistic view of AI and discuss examples of the current state of the art in legal AI.

In assessing the evidence, it is important to remember that we are writing on ice. The evidence of capabilities shows tentative floors, while limitations are tentative ceilings.[28] We do not know which limitations are here to stay, and which can be resolved with future development. We only know that we are still in early stages of development, and that we are still seeing constant improvements.

The first piece of evidence comes from a recent study that evaluated AI on contract review tasks.[29] The models were presented with a contract and some necessary context, and then asked to locate and determine legal issues. Comparing against the benchmark of senior lawyers, the researchers found that GPT-4 "exhibited a level of accuracy in identifying legal issues that was on

[27] ARTHUR C. CLARKE, PROFILES OF THE FUTURE: AN INQUIRY INTO THE LIMITS OF THE POSSIBLE 36 (1962)

[28] *See* Yonathan A. Arbel & Samuel Becher, *Contracts in the Age of Smart Readers*, 90 GEO. WASH. L. REV. 83 (2022) (discussing the capabilities of smart readers as well as the risks associated and the need to regulate and integrate with caution).

[29] Lauren Martin et al., *Better Call GPT: Comparing Large Language Models Against Lawyers*, arXiv (2023), https://arxiv.org/html/2401.16212v1. There are other claims, less open to scrutiny, about AI/ML systems replacing lawyers in various repetitive tasks. For example, JP Morgan reports of a software that reviews contracts and "reviews approximately 12,000 new wholesale contracts per year and replaced "360,000 hours" of staff time between lawyers and loan officers" Hugh Son, *JPMorgan Software Does in Seconds What Took Lawyers 360,000 Hours*, BLOOMBERG (Feb. 27, 2017), https://perma.cc/J548-GSUB.

par with that of Junior Lawyers."[30] The models have performed their tasks at 8 percent of the time it would take a junior lawyer to perform them. Critically, where the lawyer would charge an average of $74.26 for the task, the model operating cost was a single quarter.[31]

While the models were relatively accurate, they were not perfect, and their failure modes prove interesting. Relative to senior lawyers, models showed "a preference for precision over recall," i.e., they preferred to be accurate rather than comprehensive. This offers greater confidence in the issues identified, but risks overlooking some issues. This type of failure mode, however, is not much different than that exhibited by junior lawyers, who also follow a similar preference. Next, based on the two examples the authors provide, the mistakes appear transient rather than fundamental – when presented to new models (Opus-3, Gemini Pro) it answered them both correctly without any tuning.[32]

A related study evaluated the ability of LLMs to serve as "smart readers" that assist consumers with their contracts, privacy policies, and other legal documents.[33] It found that smart readers reduce the length of contracts by 66.9 percent; reduce reading time by 14:41 minutes; improve text readability by reducing reading levels from college-level to fifth-grade level; and, finally, do so without compromising the essential information in the original documents. There were failures, but at least some are attributable to the length of the documents, which the LLMs examined could only read in parts (this problem has since been mostly mitigated).[34]

A different study evaluated the performance of large language models on tax code questions.[35] These questions involve logical complexity (e.g., exploring taxation of vested reversible, transferable shares, and cost basis following a sale of inherited property), but also tend to have a fairly crisp, unique answer. They find that GPT-4 achieves around 77 percent accuracy on questions related to the CFR (with as much as 100 percent on basic tax problems), and 53 percent

[30] Lauren Martin et al., *Better Call GPT: Comparing Large Language Models Against Lawyers*, arXiv (2023), https://arxiv.org/html/2401.16212v1.

[31] *See id.*

[32] Presenting Claude and Gemini with a contract and some context and asking it them to identify the legal issues, https://claude.ai/chat/77338278-0036-469c-8d22-615c331f8c58; https://gemini.google.com/app/560bd35270464077

[33] Yonathan A. Arbel and Samuel Becher, *How Smart are Smart Readers? LLMs and the Future of the No-Reading Problem*, *in* THE CAMBRIDGE HANDBOOK ON EMERGING ISSUES AT THE INTERSECTION OF COMMERCIAL LAW AND TECHNOLOGY (Nancy Kim & Stacey-Ann Elvy eds., 2024) (hereinafter *How Smart are Smart Readers*), Arbel & Becher, *Contracts in the Age of Smart Readers*, *supra* note 28, 94-106; Noam Kolt, *Predicting Consumer Contracts*, 37 BERKELEY TECH. L.J. 71 (2022)

[34] Arbel and Becher, *How Smart are Smart Readers*; *see also* Kolt, *supra* note 33.

[35] John J. Nay et al., *Large Language Models as Tax Attorneys: A Case Study in Legal Capabilities Emergence*, 382 Phil. Trans. R. Soc'y A 20230159 (2024), https://doi.org/10.1098/rsta.2023.0159. Importantly, the design employs retrieval-augmented generation and prompt-engineering techniques.

on general US Code questions.[36] Critically for the interpretation of these numbers, the questions involve four to 10 possible answers, so chance accuracy would only be between 10-25 percent.[37]

These results are consistent with the previous ones. They show a high but inconsistent level of performance. Unfortunately, this study did not include a human benchmark, so we cannot tell how much better or worse these numbers are relative to a professional. However, given that the questions rely on legal and financial fluency, it is safe to assume that they considerably exceed the accuracy levels of the average lay tax preparer, and possibly even of the average non-tax lawyer. This highlights the margin of substitution point: LLM will replace not your white shoe lawyer, but your neighborhood HR Block representative or estate planner.

A persistent failure mode in these studies is "hallucinations," the invocation of non-existing facts such as precedents and their presentation as facts. One study found that "legal hallucinations are alarmingly prevalent" in LLMs, occurring 69 percent (ChatGPT 3.5) to 88 percent (Meta's Llama 2) of the time when asked specific questions about federal court cases. [38] Two factors ameliorate this concern, however. False sources, while a severe problem, can be checked with relatively little work, often involving a short internet search for verification. Moreover, while our current understanding suggests that *some* degree of model inaccuracy is inevitable, advances in modeling have shown promise in reducing this problem significantly.[39]

Assessed more holistically, two recent papers tried to determine whether models can act as generalist lawyers by comparing the performance of humans to models on the bar exam. A technical report by OpenAI famously reported that GPT-4, at launch and without modifications, has passed the Uniform Bar Exam at the 90th percentile.[40] This puts GPT-4 above the average test taker. Digging more deeply, Eric Martinez argued that these results are confounded by the timing of the specific comparison exam (February), which included many repeat test-takers with lower scores.[41] Applying several corrections, he concludes that when compared to exam passers in the July administration, GPT-4 performance is estimated to be at the median of test takers, and bottom

[36] I focus here on the few-shot experiment. The relative weakness of the US Code is probably associated with the weakness of the retrieval augment generation method, which is degraded on large corpora of text. Data taken directly from the data files reported in https://github.com/JohnNay/llm-tax-attorney/tree/main/data.

[37] Data retrieved from https://github.com/JohnNay/llm-tax-attorney/tree/main/data.

[38] Matthew Dahl et al., *Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models*, arXiv (2023), https://arxiv.org/abs/2401.01301.

[39] Ziwei Xu, Sanjay Jain & Mohan Kankanhalli, *Hallucination is Inevitable: An Innate Limitation of Large Language Models*, arXiv (2024), https://arxiv.org/abs/2401.11817. on mitigation techniques, *see* S.M. Towhidul Islam Tonmoy et al., *A Comprehensive Survey of Hallucination Mitigation Techniques in Large Language Models*, ARXIV (2023), https://arxiv.org/abs/2401.01313.

[40] OpenAI.: GPT 4. Accessed: 2023-04-24. https://openai.com/research/, Daniel Martin Katz, Michael James Bommarito, Shang Gao & Pablo Arredondo, *GPT-4 Passes the Bar Exam*, 382 Phil. Trans. R. Soc'y A 20230254 (2024), https://doi.org/10.1098/rsta.2023.0254.

[41] Eric Martínez, *Re-Evaluating GPT-4's Bar Exam Performance*, LawAI Working Paper No. 1-2023, 24 pages, (Mass. Inst. Tech. 2024), https://.ssrn.com/abstract=4441311.

15th percent on the essay section.[42] This aligns with an earlier study of GPT 3.5 that showed that on law school exams, GPT 3.5 performed at a C+ level.[43]

It is worth bearing in mind that we should be cautious about extrapolating from bar performance and law school exams to real-world performance. At the same time, we also cannot completely discount their relevance, given the critical gatekeeping role they play in our regulatory apparatus.[44] Moreover, bar exams offer one of the sharpest ways of test performance differentials between models and highly motivated quasi-experts.

Finally, and most importantly, are the real-world studies of AI effectiveness. These are early days, so caution is advised. One study asked a trained lawyer and a GPT-4 to draft a complaint letter to the employer. 80 percent of human referees, in a blind test, preferred the model's letter to that of the trained lawyer.[45] Another study recruited legal aid lawyers, and gave them access to GPT-4, with some of them getting access to other AI tools. The lawyers reported a productivity increase, although they remained worried about these tools. It is worth noting that most appreciated GPT-4, but found the other tools fairly unhelpful.[46]

To conclude, if we can provide an estimate of the general level of AI models in 2024, it will be in line with Martinez' ultimate conclusions. Rigorous testing shows that these systems are fast, and cheap, but perform below the level of median lawyers. This conclusion should be made alongside the observation made at the outset, i.e., that what we see today are tentative floors and ceilings. In fact, the tests discussed not only do not account for future developments, but they also do not fully take advantage of *present* developments, such as deep prompt engineering, fine-tuning on specific datasets, or ensembling. They also do not account for the realistic alternatives that people have to using these systems.

1.2. Uptake

How are people reacting to this new technology? The potential seems quite large, with a Goldman Sachs report from 2023 claiming that AI will automate 44 percent of legal tasks within ten years of broad adoption.[47] Various reports show that law firms are experimenting with AI tools

[42] *See id.*

[43] Jonathan H. Choi, Kevin E. Hickman, Amy B. Monahan & Daniel Schwarcz, *ChatGPT Goes to Law School*, 71 J. LEGAL EDUC. 387 (2021).

[44] Kyle Rozema, *Does the Bar Exam Protect the Public?*, SSRN, 2-3 (2021) https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3612481. (showing that the "bar passage requirements have a modest negative effect on public sanctions.")

[45] Lena Wrzesniowska, Can AI Make a Case? AI Vs. Lawyer in the Dutch Legal Context, THE INT'L J. LAW, ETHICS, & TECH. (2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4614381

[46] Colleen V. Chien & Miriam Kim, *Generative AI and Legal Aid: Results from a Field Study and 100 Use Cases to Bridge the Access to Justice Gap*, 25 LOY. L.A. L. REV. (2024), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4189837.

[47] Goldman Sachs, The Potential Large Effects of Artificial Intelligence on Economic Growth 6 (Mar. 26, 2023),

in their practice.[48] For example, Allen & Overy deployed a model called Harvey and found that quickly 25 percent of the firm's practice used the tool daily.[49]

Industry surveys provide a broader picture. A survey in 2023 found that 82 percent of lawyers believed that AI can be applied to legal work, while showing more hesitancy on the appropriateness of doing so, with only 51 percent answering in the affirmative.[50] An ABA survey from 2023 reported usage among 11 percent of lawyers,[51] a Lexis survey reported 16 percent,[52] and a survey of legal aid lawyers found 21 percent usage.[53]

While these surveys suggest only small to moderate adoption, lawyers also see a broad room for integration of AI tools into their practice. Among the most common use cases, lawyers reported creating drafts, brainstorming ideas, summarizing complex documents, and assisting in writing emails.[54] It is quite reasonable to expect that as AI tools develop specifically to meet the needs of law firms, and as more lawyers graduate from law schools after using AI tools, the levels of integration will consistently increase. This is especially true given client pressure to reduce billing through the integration of these tools.[55]

Equally remarkable is the rate of change: slowly, then suddenly. A recent survey on AI adoption in the workplace (not specifically legal) has shown that 75 percent of knowledge workers use AI at work.[56] What is remarkable is that 46 percent of workers started using AI tools less than

https://www.gspublishing.com/content/research/en/reports/2023/03/27/d64e052b-0f6e-45d7-967b-d7be35fabd16.html

[48] Frank Fagan, *A View of How Language Models Will Transform Law*, TENN. L. REV. (forthcoming 2025) (manuscript at 26).

[49] Bob Ambrogi, As Allen & Overy Deploys GPT-based Legal App Harvey Firmwide, Founders Say Other Firms Will Soon Follow, LAWSITES.COM (Feb. 17, 2023), https://www.lawnext.com/2023/02/as-allen-overy-deploys-gpt-based-legal-app-harvey-firmwide-founders-say-other-firms-will-soon-follow.html

[50] *New Report on ChatGPT & Generative AI in Law Firms Shows Opportunities Abound, Even as Concerns Persist*, THOMSON REUTERS (May 3, 2023), https://www.thomsonreuters.com/en-us/posts/technology/chatgpt-generative-ai-law-firms-2023/.

[51] Darla Wynon Kite-Jackson, *2023 Artificial Intelligence (AI) TechReport*, AM. BAR ASS'N (2024), https://www.americanbar.org/groups/law_practice/resources/tech-report/2023/2023-artificial-intelligence-ai-techreport/.

[52] LEXISNEXIS, *International Legal Generative AI Report: Detailed Survey Findings* 6 (2023), https://www.lexisnexis.com/pdf/lexisplus/international-legal-generative-ai-report.pdf.

[53] Colleen V. Chien & Miriam Kim, *Generative AI and Legal Aid: Results from a Field Study and 100 Use Cases to Bridge the Access to Justice Gap*, 25 LOYOLA L.A. L. REV. (Forthcoming 2024), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4189837.

[54] Caroline Hill, *ILTA's Annual Tech Survey 2023 Reveals All on Collaboration Tools Adoption, Governance, and Yes, Lots on Gen AI*, LEGAL IT INSIDER (Sept. 29, 2023), https://legaltechnology.com/2023/09/29/iltas-annual-tech-survey-2023-reveals-all-on-collaboration-tools-adoption-governance-and-yes-lots-on-gen-ai/

[55] Logan Lathrop, *Law Firms Leveraging AI: Maximizing Benefits and Addressing Challenges*, HARV. J.L. & TECH. DIG. (Nov. 20, 2023), https://jolt.law.harvard.edu/digest/law-firms-leveraging-ai-maximizing-benefits-and-addressing-challenges

[56] *AI at Work Is Here. Now Comes the Hard Part*, MICROSOFT WORKLAB (May 8, 2024), https://www.microsoft.com/en-us/worklab/work-trend-index/ai-at-work-is-here-now-comes-the-hard-part/.

6 months ago. This spells a staggering rate of adoption. It is highly unlikely that law firms will lag behind for much longer.

1.3. Breadth of the AI Impact on Access to Justice

Access to justice is a large umbrella term, and barriers exist for several reasons. Still, many of these reasons can be reduced, in some way or another, to a cost-based explanation. AI tools drastically cut down on these costs. The point that should be made emphatically is that AI produces a *holistic* shock to the access to justice problem, going well beyond the cost of lawyering.

Legal sociology teaches the critical importance of upstream filters. "[D]isputes are not things: they are social constructs."[57] For a mischief to be conceived as a legal dispute, it must undergo at least three transformations given by the "naming, blaming, and claiming" (NBC) model.[58] That is, the injured party must perceive that they were injured; that a recognizable actor injured them (rather than an act of Fortuna), and then to be able to conceptualize of that accident in terms of a legal assertion of rights against the violator. While data is scarce, sociologists believe that these filters have a dramatic effect: "we know that most of the attrition occurs at [the NBC] early stages."[59] An important facet is distributional; the NBC filter asymmetrically affects claimants, as the ability to name, blame, and claim is predicated on access to educational, social, and plain-vanilla capital.[60]

Generative AI can be a powerful agent of transformation. To illustrate, I presented a simple query to a model: "my landlord wants me to pay to fix the mold in the basement and I don't know what to do." The model responded with some fairly generic reminders that landlords are responsible for the habitability of their residences, that it is advisable to read the lease, and that it might be appropriate to consult a legal professional. To a lawyer, burdened with the curse of knowledge, this may not seem very informative. But this response *quickly* and *cheaply* takes the user through all three of the NBC stages. If the NBC filter is as powerful as sociologists claim, and if it is as regressive in effect as commonly believed, its removal would have broad implications for both substantive rights and litigation patterns.

Beyond the early stages, AI continues to contribute to every aspect of the litigation journey. We have noted above the research showing how AI can perform legal functions at the level of a low-quality or middling lawyer. For many individuals, this will be much better than any realistic

[57] L. F. Felstiner, Richard L. Abel & Austin Sarat, *The Emergence and Transformation of Disputes: Naming, Blaming, Claiming...*, 15 LAW & SOC'Y REV. 631 (1980)

Id.

[59] *Id*. at 636.

[60] *Id.* at 637.

alternative. Today, people surveyed report that they seek lawyers for legal information in only 29 percent of their cases, often depending on the internet and family or friends for orientation.[61]

The AI system will help them with legal strategy, including matters such as whether to send a demand letter, talk to a lawyer, write to a government agency, and so on. When individuals turn to AI tools, they can use them as powerful smart readers, tools that not only summarize the information but also make it accessible to one's specific sociolinguistic needs.[62] The models can then draft the required communications, demand letters, complaints, and other litigation materials. There are many other smaller frictions along the way, and in all of them, it is straightforward to see the contribution of even today's models.

There is also considerable scope for more traditional machine learning techniques. In a recent overview, Frankenreiter and Nyarko offer a broad exploration of the utility of narrower predictive and classification models.[63] They provide persuasive use cases related to automated review of documents to identify privileged information, or using a model to predict case outcome and then use that prediction to inform one's selection of attorneys and venues.[64] More generally, the extraction of legal data from troves of documents presents a compelling and highly useful use case.[65]

The removal of barriers to access to justice will also have less salutary effects. The economic theory of signaling offers an unexpected lesson here. In the classic telling, a peacock is carrying a large flashy tail not because the tail contributes to its survival, but exactly because it hinders it. A naïve bystander would think that removing such hindrance would promote the welfare of the peacock. But the tail serves a critical function. The peahen can easily see it and then, instinctively infer, that only a fit mate would be able to survive despite such a tail.[66] The lesson is that some forms of friction, constraints, and filters, are sometimes necessary and serve indirect but important social functions.

Back to barriers on the way to justice: because they make litigation harder, they work to filter many cases. These barriers may have an important role in deterring filing by strategic players who seek to use legal enforcement to extract rents or otherwise engage in antisocial behavior.[67]

[61] *Justice Needs and Satisfaction in the United States*, IAALS (Sept. 1, 2021), https://iaals.du.edu/sites/default/files/documents/publications/justice-needs-and-satisfaction-us.pdf. (showing legal aid services account for additional 8 percent and court provided information for additional 7 percent).

[62] *See* Arbel and Becher, *supra* note 28.

[63] Jens Frankenreiter and Julian Nyarko, *Natural Language Processing in Legal Tech*, *in* LEGAL TECH AND THE FUTURE OF CIVIL JUSTICE 21, 29 (David Freeman Engstrom ed., 2023).

[64] *Id.* at 74.

[65] *Id.* at 75.

[66] For an overview of signaling in social life, *see* Yonathan A. Arbel, *A Social Status Theory of Defamation Law,* U. C. IRVINE L. REV. (2023).

[67] To be sure: that barriers to the legal system serve a positive function do not make them net positive. They also filter many truly important cases and their effect is likely regressive. The point here is only that they *also* chill low-quality cases.

To provide one common example, consider debt collection litigation. Despite a common view that these lawsuits are frequently abusive, matters could actually be worse. Professional debt buyers, who buy large debt portfolios, are effectively deterred by access frictions from filing claims for claims that lie below $500, and often not even $1,000.[68]

These filters, then, limit a great variety of cases and the untapped potential for litigation is very large. Third party financing ameliorated the liquidity barrier that prevented litigants with strong cases from filing them, and this had the effect of a litigation spike.[69] If the access to justice literature correctly mapped the barriers and their size, we have a strong reason to expect an AI litigation boom in the coming years. Exactly how large it would be is hard to gauge with any accuracy, but some estimates suggest a doubling to quintupling of current litigation rates. Moreover, it is not just the raw number of cases that matters; AI systems are excellent providers of *verbose* materials, making it effortless to write briefings and other filings that are long-winded.

To be sure, there are some trends that would work to mitigate the litigation boom. It is possible that adoption rates will be lower, or high only among those already prone to litigate their cases. It is also possible that the higher risk of litigation would lead people to adapt their behavior into greater compliance; or that would-be defendants will settle at earlier stages. AI labs, by pressure of regulation or exposure to unauthorized practice of law rulings, might also try to prevent their models from producing effective materials. Such possibilities exist, even if I find them quite unlikely. The economic incentives are simply too strong, and the temptation of convenience too large. Even if the quality is not quite there, the cost of convenience usually takes the upper hand.

2. Legal Thermostats

We have just reviewed arguments suggesting an impending surge in AI-driven litigation. A rapid increase in case volume can have systemic repercussions on substantive justice throughout the legal system. This is partly because justice delayed is justice denied, and partly because judges are ultimately humans, with only so many hours in a day. [70] Bert Huang demonstrated that a rise in administrative cases can lead to "lightened scrutiny" of civil appeals.[71] Not because judges work

[68] Dave T., *Debt Collection Agencies: What Is The Minimum Amount They Would Sue For?*, Man vs. Debt (Sept. 22, 2022), https://manvsdebt.com/debt-collection-agencies-what-is-the-minimum-amount-they-would-sue-for/.

[69] Third party financing is meant to alleviate the liquidity constraints of litigants, and its effect is said to be to "increase[] the volume of litigation in any jurisdiction where it is available." *Third Party Financing*, U.S. Chamber of Commerce Institute for Legal Reform (Oct. 2020), https://instituteforlegalreform.com/wp-content/uploads/2020/10/Third_Party_Financing.pdf.

[70] Christoph Engel & Keren Weinshall, *Manna from Heaven for Judges: Judges' Reaction to a Quasi-Random Reduction in Caseload*, 17 J. EMPIRICAL LEGAL STUD. 759 (2020), https://doi.org/10.1111/jels.12265 (finding that "Judges working in courts with reduced caseload invested more resources in resolving each case")

[71] Bert I. Huang, *Lightened Scrutiny*, 124 HARV. L. REV. 1109 (2011); *see also* Shay Lavie, *Appellate Courts and Caseload Pressure*, 27 STAN. LAW. POL'Y. REV. 57 (2016).

any less hard—they likely work even harder—but because there are physical constraints on what we can honestly expect of even the most diligent public servant.

What will happen to judicial economy in the age of AI? How can our current system—already burdened by its workload—support a dramatic uptick in the number of cases? This Part lays out the argument that past reactions to litigation surges have been accompanied by adjustments that tended to affect primary and procedural rights.

A useful way to think about judicial economy comes from control theory.[72] The core principle of control theory involves the design and analysis of dynamic systems capable of maintaining desired states despite internal and external disturbances. This is achieved using control components, such as controllers, sensors, and actuators, to regulate system behavior.

A common example is the humble thermostat. The thermostat allows users to set a desired temperature (setpoint). It continuously monitors the actual temperature (process variable) and compares it to the setpoint. If a discrepancy is detected, the thermostat adjusts the heating or cooling mechanisms (actuators) to bring the temperature back to the desired level. This dynamic response to changes, both from user input and environmental factors, distinguishes it from an open-loop system like a fan, which operates without feedback.

The legal system has analogous control mechanisms: procedural and substantive requirements that *effectively* function as tools for managing the volume of litigation.[73] Adjusting these doctrines serves as a feedback mechanism to control litigation flow, akin to how a thermostat regulates temperature. Critically, however, such adjustments invariably impact substantive rights. While there is some scholarly debate regarding the appropriateness of using legal rights as a means to manage judicial resources, our focus here is on the descriptive understanding that such adjustments are common responses to maintaining judicial economy.[74]

A few illustrations would communicate the point.[75] The most salient are court fees. Courts in the United States charge a variety of fees, including filing fees to initiate a case, fees for serving documents, court reporter fees, jury fees, and fees for accessing court records. Filing fees vary based on the type of case and jurisdiction, but can range from under $100 for small claims cases to

[72] For an introductory textbook, *see* KATSUHIKO OGATA, MODERN CONTROL ENGINEERING (5th ed. 2010), https://wp.kntu.ac.ir/dfard/ebook/lc/Katsuhiko%20Ogata-Modern%20Control%20Engineering-Prentice%20Hal%20(2010).pdf; *see also* ROBERT H. BISHOP & RICHARD C. DORF, MODERN CONTROL SYSTEMS (13th ed. 2022).

[73] In a contemporaneous article, Abramowicz considers the use of "automatic stabilizers" to consider doctrinal changes in light of potential productivity changes in lawyering due to AI. Michael Abramowicz, *The Cost of Justice at the Dawn of AI* 73 (2024), 61-62 https://ssrn.com/abstract=4543803. In various ways, his Article completes the analysis proposed here.

[74] Compare Ronen Avraham & William H.J. Hubbard, *Civil Procedure as the Regulation of Externalities: Toward a New Theory of Civil Litigation*, 89 U. CHI. L. REV. 1 (2022), which emphasizes an externality control view of civil procedure with Marin K. Levy, *Judging the Flood of Litigation*, 80 U. CHI. L. REV. 1007, 1010-11 (2013).

[75] While my focus here is on procedural mechanisms, substantive standards also encode judgments on judicial resources, but this argument is beyond the current scope.

over $400 for civil cases in federal court. Court fees work well when they deter cases whose probability of winning is so low that it falls below the fee. The *de minimis* rule has a somewhat similar function, in that it filters out cases whose actual value is low, on the premise that their social value is also low. The problem is that fees and these types of rules also screen out socially important and valuable litigation,[76] and its results tend to be quite regressive.[77] We know that even small access barriers can have large effects. Something like the distance from the courthouse, which might seem like a small concern, has a significant effect on the participation rate of the poor – even for life-changing litigation.[78]

Most procedural thermostats are more indirect. *Lone Pine* orders are an example.[79] These are orders set out in large toxic tort cases that call plaintiffs to present preliminary evidence on questions of injury and causation within a deadline or risk dismissal.[80] These orders are clearly meant as a mechanism "to identify and cull potentially meritless claims."[81] Critiques have decried their inconsistency, [82] expressed concern that they turn into "pseudo-summary judgment motions,"[83] and overall worry that it creates a burden that is "unrealistic" and an "exercise [that] is onerous and unrewarding."[84] Nonetheless, courts find them necessary to manage litigation.[85]

Consider next the doctrine of exhaustion of administrative remedies in the context of prisoner's rights. This broadly applied doctrine requires plaintiffs to navigate agency processes to completion before seeking judicial relief. While this doctrine abides by various logics, litigation control is one of them. As a response to the spike in inmate filings of the early 1990s, [86] Congress enacted The Prison Litigation Reform Act. [87] Senator Orrin Hatch, Chair of the Senate Judiciary

[76] Shmuel I. Becher, Yuval Feldman & Meirav Furth-Matzkin, *Toxic Promises*, 63 B.C. L. REV. 753, 777 (2022)

[77] Joseph Shapiro, *As Court Fees Rise, The Poor Are Paying The Price*, NPR (May 19, 2014, 4:02 PM), https://www.npr.org/2014/05/19/312158516/as-court-fees-rise-the-poor-are-paying-the-price

[78] David A. Hoffman & Anton Strezhnev, *Longer Trips to Court Cause Evictions*, 120 Proc. Nat'l Acad. Sci. e2210467120 (2023), https://doi.org/10.1073/pnas.2210467120.

[79] *See* generally, Nora Freeman Engstrom, *The Lessons of Lone Pine*, 129 YALE L.J. 2 (2019).

[80] *See e.g.*, Claar v. Burlington Northern 4 Railroad Co., 29 F.3d 499, 500 (9th Cir. 1994) ("The district court issued a case management order consolidating the twenty-seven cases for pretrial purposes. The order required plaintiffs to submit affidavits describing their exposure to the chemicals they claim harmed them, and affidavits from physicians listing each plaintiff's specific injuries, the particular chemical(s) that in the physician's opinion caused each injury, and the scientific basis for the physician's conclusions.")

[81] Baker v. Chevron USA, Inc., No. 1:05-CV-227, 2007 WL 315346, at *1 (S.D. Ohio Jan. 30). 2007)

[82] Engstrom *supra* note 79, at 37.

[83] Adinolfe v. United Techs. Corp., 768 F.3d 1161, 1168 (11th Cir. 2014)

[84] Engstrom, *supra* note 79, at 52.

[85] *See e.g.*, In Acuna et al. v. Brown & Root, et al., 200 F.3d 335 (5th Cir. 2000) ("It was within the court's discretion to take steps to manage the complex and potentially very burdensome discovery that the cases would require").

[86] Margo Schlanger, *Inmate Litigation*, 116 HARV. L. REV. 1555, 1578-1587 (2003) (on the reasons). Russell Gold highlights that these filters tend to track claims by marginalized individuals. Russell M. Gold, *Power over Procedure*, 73 ALA. L. REV. 1, 105-6 (2022)

[87] Prison Litigation Reform Act of 1995, Pub. L. No. 104-134, §§ 802–809,

Committee, explained: “This landmark legislation will help bring relief to a civil justice system overburdened by frivolous prisoner lawsuits.”[88] The Supreme Court likewise noted in *McKart v. U.S.*, that exhaustion "serves the twin purposes of protecting administrative agency authority and promoting judicial efficiency."

Empirical evidence suggests that the exhaustion requirement does indeed filter out a significant number of potential claims. A study by the U.S. Government Accountability Office found that, in the context of employment discrimination, fewer than 10 percent of individuals who filed charges with the Equal Employment Opportunity Commission subsequently filed lawsuits in federal court. As AI lowers the barriers to navigating the administrative process, courts may feel the need to recalibrate the exhaustion doctrine to maintain an appropriate balance between access to justice and judicial economy

Standards of proof also second as thermostats. Consider what is necessary to prove to win a retaliation claim under Title VII of the Civil Rights Act.[89] Spurred by concerns about a deluge of lawsuits, the Court decided that the standard of proof would be the but-for test, rather than the more plaintiff-friendly motivating factor test. It argued that “[l]essening the causation standard could also contribute to the filing of frivolous claims, which would siphon resources from efforts by employer[s], administrative agencies, and courts.”[90]

Pleading standards serve as obvious thermostats. Consider *Twombly* and *Iqbal*, two of the most important procedural decisions in modern law. They mark the move from a negative “no set of facts" standard to a positive one requiring a showing of plausibility.[91] The motivation, in large part, was the growing costs of discovery that were enabled by the old standard.[92] The effect has been controlling access to justice, as has been amply argued by the various critiques.[93]

A final illustration of procedural thermostats comes from statutes of limitations. There are, by one count, around seven categories of rationales for these laws.[94] One of them is the protection

110 Stat. 1321

[88] 14 1 CONG. REC. SI4,418 (daily ed. Sept. 27, 1995) (statement of Sen. Hatch).

[89] Title VII of the Civil Rights Act of 1964, 42 U.S.C. §§ 2000e et seq.

[90] Univ. of Tex. Sw. Med. Ctr. v. Nassar, 570 U.S. 338, 358 (2013). For a critique, *see* Daiquiri J. Steele, *Rationing Retaliation Claims*, 13 UC IRVINE L. REV. 993, 1003 (2023) (“While courts should be good stewards of judicial resources, docket reduction should not take precedence over ensuring equal justice under the law.”); *see also* Sandra F. Sperino & Suja A. Thomas, *Fakers and Floodgates*, 10 STAN. J. CIV. RTS. & CIV. LIBERTIES 223, 229 (2014).

[91] Edward D. Cavanagh, *Making Sense of Twombly*, 63 S.C. L. REV. 97 (2011).

[92] Twombly (“it is only by taking care to require allegations that reach the level suggesting conspiracy that we can hope to avoid the potentially enormous expense of discovery”); *see also* Asahi Glass Co. *v.* Pentech Pharmaceuticals*,* Inc*.,* 289 F. Supp. 2d 986, 995 (ND Ill. 2003) (Posner, J., sitting by designation) (“[S]ome threshold of plausibility must be crossed at the outset before a patent antitrust case should be permitted to go into its inevitably costly and protracted discovery phase”).

[93] Matthew A. Shapiro, *Distributing Civil Justice,* 109 GEO. L.J. 1473, 1516 (2021) (“heightened pleading requirements and limits on discovery, have been widely criticized for restricting access to justice”)

[94] *See generally* Tyler T. Ochoa and Andrew Wistrich, *The Puzzling Purposes of Statutes of Limitation*, 28 PAC. L. J. 453, 460-99 (1997)

of the integrity of evidence, aimed to "prevent[] surprises through the revival of claims that have been allowed to slumber until evidence has been lost, memories have faded, and witnesses have disappeared."[95] But Congress sometimes uses statutes of limitations as a means of controlling the volume and quality of litigation,[96] and so do some courts.[97]

The common usage of these procedural thermostats reveals something else. These thermostats work by adding friction to the process, which indeed filters out cases, and the (often unverified) hope is that those are cases of lesser merit.[98] The problem is that some of these frictions are quite vulnerable to the introduction of AI tools. The reasons why people fail to meet statutes of limitations requirements are varied, but some of them depend on access to lawyering and litigation financing.[99] AI can ameliorate such barriers because it can help people process the wrong they suffered through the naming-blaming-claiming process, and then assist them in constructing legal documents. Similarly, AI tools can significantly reduce the costs, hurdles, and frictions associated with exhausting administrative remedies. AI-powered tools could quickly identify relevant agencies, help navigate their process, and draft complaints. Finally, the same tools also apply to pleading standards. Plausibility standards do not only filter cases that are implausible. They also filter cases where people were negligent or unskilled in framing their arguments or lacked the requisite polish. These filtering functions of pleading standards are fragile to AI tools that can mass produce elaborate briefs for even the most tenuous of cases.

3. Legal Strategies that Deal with the AI Litigation Boom

If the diagnosis by access to justice advocates is correct, the prognosis is clear. To the extent AI tools remove frictions and costs in access to justice, we should expect a commensurate increase in civil litigation. And because the size of the access to justice gap is so large, a doubling in the volume of litigation is not implausible.[100] Moreover, litigation would also adjust on other dimensions, with verbosity of filings being an expected effect.

[95] Order of Railroad Telegraphers v. Railway Express Agency 321 U.S. 342, 349 (1944).

[96] *See e.g.,* Sandra F. Sperino & Suja A. Thomas, *Fakers and Floodgates*, 10 STAN. J. CIV. RTS. & CIV. LIBERTIES 223, 230 (2014) (arguing that "Congress inserted numerous procedural and substantive provisions in Title VII that limit the number of claims" which includes the short time to claim.)

[97] Ochoa & Wistrich, *supra* note 94, at 495-99.

[98] Is it the case that a lawsuit filed after 320 days for discrimination less meritorious than one filed within 290 days from the offending act?

Cf, however, the logic expressed in cases such as Chase Sec. Corp. v. Donaldson. 325 U.S. 304, 314 (1945), where the court sees statutes of limitation as tools that "are by definition arbitrary, and their operation does not discriminate between the just and the unjust claim, or the voidable and unavoidable delay."

[99] For a psychological account of delay, *see* Andrew J. Wistrich, *Procrastination, Deadlines, and Statutes of Limitation*, 50 WM. & MARY L. REV. 607 (2008)

[100] Ideally when scholars make prescriptions based on their understanding of the future trajectory of the world – as I do here – they should offer some concrete, refutable predictions on how they perceive future trends to evolve. Here, it's important to acknowledge problems of missing data on present litigation patterns, scope and type of

Historically, courts have reacted to threats to judicial economy by adjusting the thermostats available to them. The goal of this Part is to situate thermostat-adjustment as one of several possible strategic reactions to the expected AI boom. I conclude with a discussion of the policy I consider most prudent: proactive integration. AI has shortcomings and reliability issues, but as explained, some are exaggerated and others manageable, and all should be evaluated vis-à-vis the other realistic alternatives we have on the menu. By using whatever time we have left until the AI boom, we can carefully build, test, and deploy AI tools as part of the judicial process.

3.1. Sit-and-Wait

Judges and judicial administrators are careful by nature, and a rapidly expanding and advertised technology such as AI raises understandable concerns about unjustified hype and empty promises. Technological uncertainty remains a significant concern. While it is evident that AI is transforming the production of legal materials, the full extent of this shift and its implications—particularly the potential for a litigation boom—are not yet fully understood. Historical precedents with earlier waves of legal technologies, such as LexisNexis and LegalZoom, suggest that the legal system can adapt without catastrophic disruptions. Moreover, given the current imperfections in AI technologies, prudence might dictate a period of observation and gradual adaptation. Thus, judges and judicial administrators may wish to wait before they make any adaptations to legal processes, procedures, and doctrines.

Further complicating the decision is the pattern of AI adoption. We do not know yet who the dominant users would be, and that may affect our normative evaluation of the technology. Should AI tools follow the trajectory of previous legal tech innovations, we might witness a surge in litigation activities by firms and commercial entities, rather than under-served individuals.[101] There is also the potential for negative uses, such as harassment or unmeritorious litigation initiated by individual plaintiffs, which could distort the justice system and detract from its core functions.

Despite these considerations, I argue against a passive stance. Current trends, though based on preliminary data, indicate a clear trajectory toward increased AI integration within legal practices.[102] The unreliability of AI, rather than a deterrent, should be a catalyst for judicious development and testing. This proactive approach would not only allow for refinement of the technology but also prepare the judicial system to harness AI's benefits effectively.

barriers, levels of unmet needs, and so on. Still, if it turns out in 5–8 years that there was no discernible and practically meaningful AI effect on litigation patterns, the reader should consider this Essay's central claim disproven. *See also* https://twitter.com/ProfArbel/status/1297327039670898688

[101] *See* Engstrom and Engstrom, *A2J Crisis*, *supra* note 12.

[102] *See supra* Part I.2.

Moreover, even assuming the legal system could absorb the impact of AI without significant structural changes, proactive adaptation could still soften the shock of the transition and enhance its efficiency and effectiveness. Innovations such as video conferencing and digital legal research have already demonstrated the benefits of integrating technology in legal processes even when there was no imminent threat to the volume of litigation.

In conclusion, while the allure of a cautious approach is understandable given the unknowns associated with AI, there are strong reasons to adopt a more proactive engagement. This strategy ensures that the judicial system is not merely reactive but remains at the forefront of technological integration, enhancing its capacity to deliver justice effectively.

3.2. Ban & Mark

There is a growing sentiment, mostly expressed to me in private conversations with judges, that generative AI should be banned in the courtroom. Alternatively, some favor a requirement that lawyers disclose when they are using AI-generated materials.

The sentiment is understandable, but I believe it is wrong to follow it in the long term. A ban would kill in the crib our ability to democratize access to the justice system. It would perpetuate the asymmetries that currently exist, working disproportionality against those who have the most to benefit from the technology.

Disclosure regimes are a hopeless enterprise. As far as we know, and to the displeasure of school administrators everywhere, there is no *reliable* technology that can watermark AI-produced texts. Detection of AI-generated texts is probabilistic and error-prone, and it may—at best—only cover the least sophisticated of its users.[103] Their share is small, and their culpability is no worse than their more sophisticated peers. But most importantly, the expected level of AI integration in law practices suggests that disclosure will be as meaningful as requiring litigants to disclose if they used a search engine or a computer. It will communicate no actionable information to the judge and will become as helpful as "here comes the plaintiff" and other legal boilerplate. Overall, I would caution those judges and judicial administrators who, in good faith, worry about rising rates of litigation against trying a hopeless regime of ban and mark.

[103] *See e.g.*, Manshu Zhang, Liming Wu, Tao Yang, Bing Zhu & Yangai Liu, *RETRACTED: The Three-Dimensional Porous Mesh Structure of Cu-Based Metal-Organic-Framework - Aramid Cellulose Separator Enhances the Electrochemical Performance of Lithium Metal Anode Batteries*, 46 Surfaces & Interfaces 104081 (2024), https://doi.org/10.1016/j.surfin.2024.104081. (a retracted article which opens its introduction with "Certainly, here is a possible introduction for your topic:"). The original version is stored in Reddit, https://i.redd.it/zq0raef1aaoc1.jpeg

3.3. Massive Funding

Justice costs money, and the most direct way of dealing with greater demand for justice is by increasing the resources available for that purpose. The prioritization of justice resources is a question for politicians and exceeds my proffered expertise. What is meaningful for evaluating the prospects of a budget increase, however, is the estimated size of funding. If there is room for a two-fold or a five-fold increase in the volume of litigation, then this gives a general sense of the magnitude of the budget required. Of course, not all—not even the majority—of this potential will translate into lawsuits. But the realism of a budget increase that would even approximate a doubling in the number of judges and judicial administrators appears quite tenuous in our current political reality.

One fact that lends *some* realism to this proposition is that currently civil legal aid benefits from roughly $2.7 billion in overall budgets.[104] So it is conceivable that some of these budgets could be redirected towards the legal system, if legal aid is automated.[105] Yet, the federal court system alone is budgeted at $9.2 billion dollars per annum, so the margins are not broad enough.[106]

3.4. Legal Thermostats: Fees, Pleading Standards, and Substantive Standards

As previously argued, a common historical reaction to a litigation boom is the adjustment of legal thermostats, the various doctrines that double as litigation control actuators.[107] Judges and judicial administrators may feel it is necessary for them to require even higher fees to offset the demand for legal resources, to demand even more elaborate pleading standards, or perhaps go as far as narrowing substantive rights. These thermostats can decrease litigation levels,[108] but they also make it harder to vindicate legitimate claims. As every lawyer knows, being right and being able to prove one's case are not the same.

Two fairly obvious but worrisome implications of such adjustments are the narrowing of civil rights and, functionally, a large subsidy to wrongdoers who could get away with more socially pernicious activity. Less obvious is the problem that these mechanisms are not well AI-proof, and so their effects will be unstable and will require constant adjustments.

USA National Report ILAG Conference 2023, Harvard Law School Center on the Legal Profession (May 2023), https://clp.law.harvard.edu/wp-content/uploads/2023/05/USA-National-Report-ILAG-Conference-2023.pdf. According to the LSC data from 2022, the total funding for LSC funded organizations was 1.72 billion. *By The Numbers 2022: The Data Underlying Legal Aid Programs*, Legal Services Corporation, at 13-14 (2023), https://lsc-live.app.box.com/s/h2bajpr3gps4s4a1iio6fwiddhmu1nwb.

[105] *USA National Report ILAG Conference 2023*, at 4 (Since 2000, LSC has funded more than 859 projects totaling over $81 million in Technology Initiative Grants.")

[106] *The Judiciary: Fiscal Year 2025 Congressional Budget Summary*, Admin. Office of the U.S. Courts (Feb. 2024). https://www.uscourts.gov/sites/default/files/fy_2025_congressional_budget_summary.pdf

[107] *See supra* Part II.

[108] Note, however, that they also invite more accidents, and the net effect on litigation levels depends on a broader set of variables.

The force of some of the common legal thermostats, especially pleading standards, can be thought of as a "proof-of-work" mechanism. Proof of work is familiar from blockchain technology, where they are used to validate claims made by certain network participants. But despite their common association with blockchain, such mechanisms are far more general and common than many realize. In the current context, the litigation process can be thought of as having a front-end (initial claim processing) and a back end (trial). Litigants, presumably, have a sense of the merits of their case. The proof-of-work mechanism leverages it, to set higher front-end requirements. A person who puts in the work, and sinks in the necessary cost to meet plausibility standards in the front end likely has a higher estimate of their case than a person who would be discouraged by such costs. This is the case because the back-end costs are only borne by people who would pursue the case to its completion. Plausibility standards require more work on the front end, but serve the litigants later, and thus act as an effective proof-of-work filter.

Assuming for a moment that this assumption is correct in general, AI tools present a particular problem. Normally, the crafting of effective pleadings requires an effective counsel and an investment of time. A judge can relatively quickly discern plausibility when the case involves low-effort filings. But AI can rapidly and easily convert vague claims to elaborate legal arguments, using perfect grammar and compelling structure. This makes the production cheaper and later validation harder. Ironically, hallucinations can contribute to the facial plausibility of the filings, even when the underlying claim lacks any support. Consider, as an illustration, a request that the AI produce a claim for workplace discrimination. Commentators note that plausibility requirements hamper many such claims.[109] The model, however, could simply generate a set of (fictitious) facts and legal arguments that, while not true, will seem true on their face. If the litigant is not careful and scrupulous enough in reviewing them, it could pass initial muster. As a result, filtering mechanisms that rely on proof of work will become less effective than before. This could result in escalation of front-end investments, until the point where AI cannot provide sufficient utility.

Ultimately, adjusting thermostats will come at a considerable cost to litigants. Higher standards on the back end can undo all the access to justice gains AI will bring to under-served litigants. Worse, some of the thermostats will be ineffective or will need to be adjusted further and further, because AI can circumvent conventional proof-of-work mechanisms.

3.5. Integration

If none of the above strategies can effectively and equitably meet the AI boom, the legal system still has one other important course of action available to it: integration. The objective would be to implement AI in all aspects of the legal process, amplifying the productivity of judges and clerks, and allowing them to work at larger-than-ever scales. If done correctly, this strategy

[109] Joseph A. Seiner, *Plausible Harassment*, 54 UC DAVIS L. REV. 1295, 1310 (2021).

would offer a significant stretching of existing judicial resources, allowing judges to meet increased demand without resorting to adjustment of legal thermostats or sacrificing justice in individual cases.

Rather than a hypothesis, this seems to be organically happening. Judges have started admitting to using generative AI to draft opinions, although the backlash suggests that many others are still in hiding.[110] One British judge made the point succinctly and forcefully: "It is useful, and it will be used."[111] Likewise, Richard Re believes that judges will invariably find AI tools to be "irresistibly attractive."[112]

Most remarkably, in a groundbreaking decision, Judge Newsom of the Eleventh Circuit has written an opinion relying on AI for "generative interpretation." Drawing on our academic work on generative interpretation, he said:

> Those, like me, who believe that "ordinary meaning" is the foundational rule for the evaluation of legal texts should consider—consider—whether and how AI-powered large language models like OpenAI's ChatGPT, Google's Gemini, and Anthropic's Claude might—might—inform the interpretive analysis.[113]

Appeal notwithstanding, there is also significant resistance to integration, at least in its strong forms. While scholars such as Eugene Volokh express cautious optimism about robo-judging,[114] others are less sanguine. Aziz Huq speaks of a right to a "human decision,"[115] and experiments suggest a perceived fairness gap between human and artificial adjudicators.[116] These objections rely in part on empirical objections concerning the capacity of these systems to produce judgments that are as good as a human judge in terms of accuracy, bias, and gamability. They also draw on sensible ethical concerns regarding the ethics of adjudication by those who are neither citizens nor humans. The former set of problems is amenable to practical solutions, while the latter can be mostly remedied by including human judges in the loop.[117]

[110] Hibaq Farah, Court of Appeals Judge Praises 'Jolly Useful' ChatGPT After Asking It For Legal Summary, GUARDIAN (Sept. 15, 2023).

[111] *Id.*

[112] Richard Re, *Artificial Authorship and Judicial Opinions*, GEO. WASH. L. REV. (forthcoming 2024), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4696643.

[113] Snell v. United Specialty Ins. Co., 102 F.4th 1208, 1121 (5th Cir. 2024) (Newsome, J., Concurring) (citing Yonathan A. Arbel & David A. Hoffman, Generative Interpretation, 99 N.Y.U. L. REV. (forthcoming 2024)).

[114] Eugene Volokh, *Chief Justice Robots*, 68 DUKE L.J. 1135 (2019)

[115] Aziz Z. Huq, *A Right to a Human Decision*. 105 VA. L. REV. 611 (2020). *See also* Kiel Brennan-Marquez & Stephen E. Henderson, *Artificial Intelligence and Role-Reversible Judgment*, 109 J. CRIM. L & CRIMINOLOGY 137 (2019).

[116] Benjamin Minhao Chen, Alexander Stremitzer & Kevin Tobia, *Having Your Day in Robot Court*, 36 HARV. J. L. & TECH. 127(2022)

[117] Aziz Z. Huq, *A Right to a Human Decision*. 105 VA. L. REV. 611 (2020). *See* also Kiel Brennan-Marquez & Stephen E. Henderson, *Artificial Intelligence and Role-Reversible Judgment*, 109 J. CRIM. L & CRIMINOLOGY 137 (2019).

When we talk about integration, I'd like to suggest that robo-judging should not be a central frame of thinking about the technology. Provocative and exciting, for sure, but ultimately robo-judging is a distraction from the much more mundane but nonetheless powerful utility of AI in the service of justice. In the remainder of this section, I want to highlight a few of these modalities.

The immense volume of text generated in litigation is staggering, and this will likely increase as parties begin leveraging advanced AI tools to augment their legal processes. To mete out justice, we need some way to compress all this information. In other words, we need a summarization machine, and it turns out that generative AI excels at this task.

Document summarization is among the most explored areas within natural language processing using AI. This technology is divided into two main types: abstractive and extractive summarization. Abstractive summarization creates a new, condensed version of the text that conveys the core meaning of the text, potentially using its own words. Extractive summarization, on the other hand, identifies and compiles key phrases directly from the text. Both approaches can significantly aid judges by highlighting essential information and reducing the amount of material they need to personally review.

An abstractive summary can direct a judge's attention to critical parts of a document, effectively serving as a sophisticated, automated high-level summary of a document. A file management system could mark next to a document the parties filed as "exhibit 182A", the text "Sale agreement of the Tuscaloosa house." Unlike summaries written by any of the litigants, the AI has no incentive to highlight a specific frame—it seeks to offer a robust, useful summary, to the best of its ability.[118]

Extractive summaries, on the other hand, are invaluable for identifying crucial elements within the text. An extractive summary of the sale agreement may include elements such as "seller shall deliver the property on or before January 1st". It could also include such specific pieces of evidence, legal authorities, or specific quotes. These summaries are particularly useful in scenarios where precise language and specific details are pivotal.

Both have their uses. To orient oneself in a stack of documents, abstractive summaries are essential; to locate leading phrases and arguments within a document, extractive summarization is powerful.

The implementation of such summarization technologies in case management systems is straightforward and cost-effective (as simple as any large automation project is, viz, more costly and complicated than anticipated, but ultimately solvable). It would be quite possible to integrate these systems at the case management level, ensuring that every submitted document includes an

[118] The sort of biases that afflict AI systems are often irrelevant to summarization tasks. There are some implicit biases that can creep in nonetheless (such as assumptions that a doctor is male), but clerks may well be subject to similar biases and, in any event, the impact on any actual decision is highly attenuated. What is perhaps most important is that the models have no stake in the case at hand.

automated summary and extraction of key parts. This allows effective attention management on the side of the judge, a way to easily sort and find the appendix dealing with the copy of the sale contract the parties mention, or the document that covers CPI adjustments.

There is a more advanced application, commonly known as "document Q&A". Documents, by their nature, are static entities They contain information, and one has to read through the document to extract it. This becomes unwieldy when dealing with a lengthy document. Search engines offer a greater degree of interactivity. They allow one to filter pieces of a document based on keyword searches. Such keywords can be as simple as searching for "choice of law", or more advanced such as a search for "executive* /w3 decision?". Once located, the system will highlight the relevant text and orient attention to all the relevant "hits." The user is expected to sort through them and find the relevant one.

Using document Q&A is the next step. It allows the judge to ask *specific questions* about the document, and rather than using arcane keywords, the judge can use *ordinary language*. That is, after the AI ingests a filing, the judge can simply ask "does this brief mention a meeting in Switzerland?", "does the plaintiff mention the statute of limitations?", "list the executive decision the document mention and what it means". The AI will then diligently provide an answer, based on the content of the document. The answer itself will be in natural language, e.g., "this document mentions a meeting in Zurich between the CEO of Acme and the CFO of Alpha, although it doesn't discuss its purpose."

Using document Q&A is a radical improvement over our current means of interacting with documents. Search engines direct users to not think about the question they want to answer, but rather, on what queries will most likely produce the context that will answer them. We search for "choice of law" not because we necessarily care about the term, but because we think the term will be in the context of the clause that determines the choice the parties have made. Along the way, we trudge along many irrelevant mentions of the term. Document Q&A allows the user to skip this stage. The judge can simply ask "what is the choice of law in this document?"

Document Q&A methods are not an all-knowing sage, of course. It is perhaps most productive to think of them as an always on-call, diligent, and earnest attorney of middling ability. They will try but often fail to answer complex or subtle legal questions, and their responses may be partial or unintentionally misleading. Critically, they will sometimes hallucinate facts that are not true. The model might say that the parties decreed Tuscaloosa, Alabama as their choice of law, even though the agreement contains no such reference.

Both of these problems are important, but they only repeat the lesson that all tools have limitations rather than any fundamental objection to using tools. There are some helpful correctives to many of their shortcomings. In most general terms, these issues can be dealt with in ways similar to how judges currently utilize legal clerks and assistants. Judges benefit from their assistance yet maintain ultimate responsibility for decision-making. Judges learn which parts of the work they can entrust to their assistants, what type of quality assurance checks they must run, and

which parts they should do only by themselves. If a model says that the meeting took place in Zurich, and this fact is important, then the judge should verify it before proceeding to rely on this stated fact.

While such measures take away some of the efficiencies of both clerks and AI models, they still allow the judge to focus their scarce attention efficiently. As is the case for human clerks, the net time saving from AI would generally be positive—and if not, well, the judge could choose not to use them.

Confidentiality is another concern. Many of the models are currently hosted in the cloud. It will be inappropriate to share confidential information, especially when there is a risk that the owner of the model, often a commercial firm, will use the data for future model training. There are a few evolving solutions here – from on-premise model hosting to secure cloud services with proper data licensing requirements. Several AI labs are developing enterprise solutions that are sensitive to such concerns. Additionally, the formulation of legal standards tailored to the use of AI in the legal sector is critical to addressing these privacy issues and enhancing trust in AI applications.

A stronger form of integration relies on the aforementioned generative interpretation. Large language models are trained to develop complex representations of human language, based on training on datasets that encompass trillions of words. These datasets are far more exhaustive than any amount of text a single human can read in a lifetime of dedicated seclusion. Recent work has shown that judges can use artificial intelligence as a tool of textualist interpretation, drastically improving on tools such as dictionaries or corpus linguistics, not to mention the judge's private language sense. Using generative interpretation, a judge can probe the model's internal language representation, and thus access a cheap, effective, and reproducible mode of ascertaining meaning. Moreover, LLMs are designed to account for meaning *in context*. Unlike any dictionary, LLMs can easily distinguish between various plausible usages of a specific word based on its broader context. The word 'run' has no fewer than 645 meanings, and a dictionary would present them all as equiprobable definitions.[119] An LLM will have no trouble distinguishing between meanings based on context. This is why some believe that generative interpretation is the future of textualist interpretation.

In recent work, Richard Re explored the integration of AI as an opinion-drafting co-pilot. There are clear efficiencies inherent in a drafting tool that can help a judge draft an opinion quickly. Today's technology is akin to adding a cadre of enthusiastic but somewhat dull clerks. Re's account recognizes that judges will inevitably find utility in such tools. But he's also offering a deeper investigation of how such integration can affect the very nature of the adjudicative role. The point

[119] Simon Winchester, *The True Meaning of Mecca*, N.Y. TIMES (May 28, 2011), https://www.nytimes.com/2011/05/29/opinion/29winchester.html.

is that in separating opinion writing from adjudication, something—potentially very important—is lost. In Re's retelling, broad adoption will dull the edge of writing opinions; the rhetoric will turn to sophistry; the judgments generated by models majoritarian and unfirm; judicial ownership will become diffused; and deliberation and reason will decline. Moreover, the consumers of judicial opinions – the public and legal professionals – will come to view such opinions with a certain distaste. A fancy form of lifeless boilerplate.

Re is careful enough not to romanticize extant practices. He readily acknowledges that even today judges do not craft each decision from first principles and that they rely on precedent and clerks. [120] But he does view AI as a threat to the authenticity of the process.

Re's arguments are reasonable enough and become ever more reasonable when integration of AI drafting becomes closer to the robo-judging end of the spectrum. It has no real bite on the other extreme where AI is more akin to an overly engineered spell-check. Integration into authorship that helps the judge spot typos, come up with examples or metaphors, or offer variations on formulaic language, are all activities that are barely exposed to his critique. Perhaps having AI suggest legal arguments on specific issues nears the other extreme, but the point is that there are simply so many steps along this spectrum where AI is either non-problematic or that, all things considered, its integration is still a net benefit. Judges should be acutely aware of the dangers of this road, but given the immense practical pressure that looms ahead, they should not abandon it altogether.

***

I have outlined here a few modalities of reaction to the AI moment and emphasized various modes of integration into the legal process. Taken not as a method of outsourcing adjudication to algorithms, and in clear view of the limitations of AI, the recommendation that emerges from this analysis is one that favors integration. By integrating AI into the judicial process, judges will enjoy levels of support that are necessary to meet the AI moment and the potential sharp increase in litigation.

Some are not comfortable putting algorithms near human-life affecting decisions. The message of this essay is directed especially at them. Short of massive funding runs, the real decision the AI moment presents is not *whether* but *between* algorithms of sorts. As AI increases access, it will strain judicial resources. Judges may find themselves pushed to adjust the only thermostat available to them. Worse, politicians may seize the moment to adjust the thermostat against plaintiffs they disfavor on political grounds. They will say that this group uses AI to leech resources from those who really need them (and happen to belong to their favored groups).

[120] Drawing on Posner, Re reminds us that the integration of previous waves of technology have already led to tensions. RICHARD POSNER, THE FEDERAL COURTS: CRISIS AND REFORM 102 (1985).

Adjusting the legal thermostat by increasing fees, limiting substantive rights, increasing standards of pleadings and such acts as a blind algorithm. They deny access to whoever can't meet them, regardless of their need, their eventual ability to meet them, or the merit of their case. Such thermostat adjustments are often regressive and, ultimately, jeopardize substantive and procedural rights, reinstating the barriers to justice that we can finally topple. A nuanced and thoughtful mode of integration involves algorithms, but ones that are artificially intelligent, and with thoughtful integration could far outdo mechanical and potentially politicized thermostat adjustments.

This opens the stage for a new wave of tool-building scholarship, coming from lawyers and directed at lawyers. Now that scholarship has established many of the shortcomings of algorithms and AI, what positive use-cases are there? How could tools be developed with attention to their inherent limitations? There is a small wave of scholarship that tries to do that, but it is led by technologists and is published outside of law reviews. Legal scholars, cooperating with judges and judicial administrators, should take the lead, and collaborate with technologists.

Conclusion

One way of restating the arc of the argument is this: What if we could solve the access to justice problem? Implicit in much of the scholarship is the notion that reducing barriers would naturally translate to more justice for all. Here, we have adopted a more skeptical approach, based on control theory and historical lessons from past waves of litigation spikes. The Prison Litigation Reform Act serves as a sobering reminder that the barriers to justice are but one, and not necessarily the most important, barrier on the way to substantive justice.

I proposed here that an appropriate response is the proactive integration of AI tools into the legal process. At the moment, there is understandable hesitancy, given stereotypes about the ability of machines to perform legal tasks, integration costs, and the model's bias and potential lack of reliability. Such arguments are both real and exaggerated. Bias and unreliability can be addressed effectively by careful integration into the lower-stakes aspects of the process, where verification is available. More importantly, relative to other alternatives such as substantive hurdles, which bluntly and mechanically suppress litigation, AI tools can offer considerable improvement. Ultimately, judicial economy considerations pose in front of us a hard, but urgent, choice: we have to decide how much justice we want to purchase, and whether we want to stretch these dollars further by providing automation tools to judges.