
THE COLUMBIA
SCIENCE & TECHNOLOGY
LAW REVIEW

VOLUME XXV

STLR.ORG

SPRING 2024

DO CASES GENERATE BAD AI LAW?

Alicia Solow-Niederman*

There's an AI governance problem, but it's not (just) the one you think. The problem is that our judicial system is already regulating the deployment of AI systems—yet we are not coding what is happening in the courts as privately driven AI regulation. That's a mistake. AI lawsuits here and now are determining who gets to seek redress for AI injuries; when and where emerging claims are resolved; what is understood as a cognizable AI harm (and what is not), and why that is so.

This Essay exposes how our judicial system is regulating AI today and critically assesses the governance stakes. When we do not adequately recognize how the generative AI cases being decided by today's judges are already operating as a type of AI regulation, we fail to consider which emerging tendencies of adjudication about AI are likely to make good or bad AI law. For instance, litigation may do good agenda-setting and deliberative work as well as surface important information about the operation of private AI systems. But adjudication of AI issues can be bad too, given the risk of overgeneralization from particularized facts; the potential for too much homogeneity in the location of lawsuits and the kinds of litigants; and the existence of fundamental tensions between social concerns and current legal precedents.

* Associate Professor of Law, The George Washington University. Thank you to Bob Brauneis, Danielle Keats Citron, Julie Cohen, Kristen Eichensehr, Margot Kaminski, Martha Minow, Paul Ohm, Spencer Overton, Richard Re, Dan Solove, and Ari Ezra Waldman for generous feedback and support. I am also grateful for incisive comments and questions from the participants in the 2023 Columbia Science and Technology Law Review Symposium, particularly Shyam Baganesh, and for excellent suggestions from the STLR editors, especially Penina Cohen, Emily Davidson, and Amanda Orbuch. Thank you also to participants in the CS+Law Research Workshop, the University of Virginia Law School Law & Tech Colloquium, and the American University Washington College of Law Faculty Workshop for helpful comments. All remaining errors and omissions are my own. This Essay reflects developments through mid-January 2024, when it was substantively finalized for publication. © Alicia Solow-Niederman.

If we overlook these dynamics, we risk missing a vital lesson: AI governance requires better accounting for the interactive relationship between regulation of AI through the judicial system and more traditional public regulation of AI. Shifting our perspective creates space to consider new AI governance possibilities. For instance, litigation incentives (such as motivations for bringing a lawsuit or motivations to settle) or the types of remedies available may open up or close down further regulatory development. This shift in perspective also allows us to see how considerations that on their face have nothing to do with AI—such as access to justice measures and the role of judicial minimalism—in fact shape the path of AI regulation through the courts. Today’s AI lawsuits provide an early opportunity to expand AI governance toolkits and to understand AI adjudication and public regulation as complementary regulatory approaches. We should not throw away our shot.

I. INTRODUCTION	262
II. AI REGULATION IN AND THROUGH COURTS	268
III. JUDGING GENERATIVE AI TODAY	272
A. <i>Regulation in the Courts: The Good</i>	275
B. <i>Regulation in the Courts: The Bad</i>	278
IV. REPOSITIONING AI ADJUDICATION AS AI GOVERNANCE	283
V. CONCLUSION.....	286

I. INTRODUCTION

There’s an AI governance problem, but it’s not (just) the one you think. The problem is that our judicial system is already regulating the deployment of AI systems—yet we are not coding what is happening in the courts as privately driven AI regulation. That’s a mistake. How courts are processing and deciding claims about AI systems is regulating and ordering social affairs for both affected parties as individuals and for the public as a whole. AI lawsuits here and now are determining who gets to seek redress for AI injuries; when and where emerging claims are resolved; what is understood as a cognizable AI harm (and what is not), and why that is so.¹

This Essay exposes this regulatory role of AI adjudication and contends that AI governance requires better accounting for the interactive relationship between AI adjudication and AI regulation. AI adjudication does more than provide a one-shot opportunity for individuals to seek recourse for a particularized harm through the

¹ See Melissa Heikkilä, *How Judges, Not Politicians, Could Dictate America’s AI Rules*, MIT TECH. REV. (July 17, 2023), <https://www.technologyreview.com/2023/07/17/1076416/judges-lawsuits-dictate-ai-rules/> [<https://perma.cc/Q6NN-XW9N>].

legal system. Judicial resolution of claims against the makers of generative AI systems shapes the collective social impact of this technology.²

Three upfront caveats are in order. First, the “AI adjudication” considered in this Essay does not involve “robot judges” or the use of AI to adjudicate.³ The inquiry is a distinct one: How legal disputes involving AI systems and tools, using existing private law causes of action, are already shaping AI governance.⁴ Second, although most of these emerging cases center on copyright infringement claims, this Essay contends that the issues presented transcend intellectual property. Even if there is something specific about, say, copyright parties’ incentives or the way that these cases tend to be resolved in the courts, these disputes serve as vehicles to understand the relationship between law, technological developments, and society more generally. Third, although this Essay discusses early generative AI lawsuits as a way to gain analytic traction in a rapidly changing space, its insights are not limited to generative AI. Because generative AI is the cutting edge of AI today, its seeming novelty can helpfully illuminate dynamics that may be slower to emerge, or subtler, in other areas. These concrete examples serve as a useful indicator of trends and tendencies that are likely to permeate AI disputes—and thus inform AI

² On the potential societal impact of the common law in “the age of AI,” see Mariano-Florentino Cuéllar, *A Common Law for The Age of AI: Incremental Adjudication, Institutions, and Relational Non-Arbitrariness*, 119 COLUM. L. REV. 1773, 1776 (2019) (“[C]ommon law ideas tend to set the terms for conversations among elites and even the larger public about the way social and economic interactions ordinarily occur, and how public agencies should analyze the problems . . . they are designed to mitigate.”). This Essay builds from this insight and from a longstanding literature on the common law as a form of social regulation to pinpoint the likely regulatory impact of generative AI litigation, assess the merits and demerits of AI adjudication, and consider the broader implications for AI governance.

³ For earlier work on the prospect of “robot judges,” see generally Richard M. Re & Alicia Solow-Niederman, *Developing Artificially Intelligent Justice*, 2 STAN. TECH. L. REV. 242 (2019). For a contemporary example of this phenomenon, see Thomas Germain, *Judges Given the OK to Use ChatGPT in Legal Rulings*, GIZMODO (Dec. 12, 2023), <https://gizmodo.com/uk-judges-now-permitted-use-chatgpt-in-legal-rulings-1851093046> (authorizing UK judges’ use of the “jolly useful” chatbot).

⁴ Other law and technology scholars have suggested that common law liability might play a complementary role alongside public regulation of AI or made the case for new individual rights to contest AI decision-making. This Essay is distinctive in its focus on what already *is* when it comes to AI litigation and adjudication and in its analysis of how this emerging reality might affect AI governance systems in an era of generative AI. For a sampling of relevant scholarship on AI regulation and the potential role of common law liability or other individual rights more generally, see, e.g., Andrew B. Selbst, Suresh Venkatasubramanian, & I. Elizabeth Kumar, *Deconstructing Design Decisions: Why Courts Must Interrogate Machine Learning and Other Technologies*, 85 OHIO STATE L.J. (forthcoming 2024), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4564304 [<https://perma.cc/RPM2-JY3X>]; Margot E. Kaminski, *Regulating the Risks of Artificial Intelligence*, 103 B.U. L. REV. 1347, 1395 n.266 (2023); Margot E. Kaminski & Jennifer M. Urban, *The Right to Contest AI*, 121 COLUM. L. REV. 1957, 1964–65 (2021); Andrew B. Selbst, *Negligence and AI’s Human Users*, 100 B.U. L. REV. 1315 (2020); Frank Pasquale, *Data-Informed Duties in AI Development*, 119 COLUM. L. REV. 1917, 1920 (2019); Gary E. Marchant & Yvonne A. Stevens, *Resilience: A New Tool in the Risk Governance Toolbox for Emerging Technologies*, 51 U.C. DAVIS L. REV. 233, 242–44 (2017); Matthew U. Scherer, *Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies*, 29 HARV. J.L. & TECH. 353, 393–98 (2016).

governance.⁵ In short, early generative AI lawsuits, many of which involve copyright claims, serve as raw material from which to distill nascent trends and better understand the present and potential future of AI regulation writ large.

When we do not adequately recognize how the generative AI cases being decided by today's judges are already operating as a type of AI regulation, we elide two related, central questions about how to craft an AI governance regime. First, we fail to consider which emerging tendencies of AI adjudication are likely to make good or bad AI law, and why.⁶ This oversight leads to a second: We fail to see how the potential and limits of private litigation interact with the potential and limits of more traditional, top-down regulation, as well as how such “bottom-up” and “top-down” lawmaking⁷ can interact over time and, potentially, complement one another.⁸

This Essay foregrounds these questions by focusing on the regulatory function of AI adjudication. Understanding generative AI lawsuits as a form of regulation permits a more honest appraisal of the potential good alongside the bad, as well as

⁵ I have taken a similar stance that social and technological developments can place latent legal challenges and questions in especially stark relief in past work, both alone and with co-authors. *See, e.g.*, Alicia Solow-Niederman, *Information Privacy and The Inference Economy*, 117 N.W. L. REV. 357, 363 n.18 (2022) (“I do not argue that ML is wholly unique or new in revealing . . . challenges [to information privacy law]; rather, my point is that the social and technological dynamics of ML illuminate issues with particular force, to be taken seriously here and now.”); Richard M. Re & Alicia Solow-Niederman, *Developing Artificially Intelligent Justice*, 22 STAN. TECH. L. REV. 242, 247 (2019) (arguing that an analysis of AI in “judicial decision-making . . . sheds light on governance issues that are likely to emerge more subtly or slowly elsewhere”). *See also* LAWRENCE LESSIG, CODE: VERSION 2.0 25 (2006) (contending that technology such as a “worm” can change the underlying costs of conducting a search and “reveal[] . . . ‘a latent ambiguity’ in the original constitutional rule”).

⁶ *See* Frederick Schauer, *Do Cases Make Bad Law?*, 73 U. CHI. L. REV. 883, 884 (2006) (drawing on Justice Oliver Wendall Holmes’s classic statement, “[g]reat cases like hard cases make bad law”); *see also* Mike Ananny, *Seeing Like an Algorithmic Error*, 24 YALE J.L. & TECH. 1, 3 (2022) (citing Justice Holmes and Professor Schauer and “ask[ing] which algorithmic errors are ‘good’ mistakes . . . that point to systematic problems”).

⁷ *See* Jeffrey J. Rachlinski, *Bottom-Up Versus Top-Down Lawmaking*, 73 U. CHI. L. REV. 933 (2006) (identifying and analyzing the “bottom-up” and “top-down” forms of lawmaking). *Cf.* Patrick Luff, *Risk Regulation and Regulatory Litigation*, 64 RUTGERS L. REV. 173, 214–15 (2011) (recounting the difficulty in defining “regulatory litigation” and concluding that “[t]he failure to realize that regulatory litigation can be divided into top-down uses of law and bottom-up uses of remedy to influence behavior and address realized risks has led to piecemeal and jumbled attempts to define the phenomenon”).

⁸ Margot Kaminski has argued that algorithmic accountability requires a “binary governance” approach that pairs individual due process rights and a systemic regulatory approach. Margot Kaminski, *Binary Governance: Lessons from the GDPR’s Approach to Algorithmic Accountability*, 92 S. CAL. L. REV. 1529, 1534–35 (2019). This Essay takes a complementary, yet distinct tack, moving one level up to analyze how private adjudication here and now, based on existing common law and statutory causes of action, is itself a form of AI governance.

what we might do about it.⁹ Put differently, seen as an instrument of regulation, is resolution of AI disputes through the judicial system good or bad, and why?

The good is deliberative and instrumental: A spate of generative AI lawsuits brought by private parties might represent “good” law by providing a forum in which to resolve individual disputes while simultaneously fostering agenda-setting and public deliberation.¹⁰ Adjudication underscores emerging social issues, such as artists’ and authors’ concerns that generative AI threatens their livelihoods,¹¹ as well as citizens’ concerns that scraping of their data to train “large language models” (“LLMs”) violates their privacy rights.¹² That’s good—courts can be a forum for dialogue about how technology is affecting society¹³ and, potentially, help to surface concrete issues and harms for policymakers to address.¹⁴ Litigation might also be the only practical option if political gridlock stalls AI-specific

⁹ I adopt Lawrence Lessig’s broad understanding of regulation as “the constraining effect of some action, or policy, whether intended by anyone or not.” Lawrence Lessig, *The New Chicago School*, 27 J. LEGAL STUD. 661, 662 n.1 (1998). See also Alicia Solow-Niderman, *Administering Artificial Intelligence*, 93 S. CAL. L. REV. 633, 643 n.44 (2020) (adopting same definition). I use the term “public regulation” to refer to top-down regulation by the state.

¹⁰ On the potential for tort law to serve both private law “civil recourse” functions and public law functions such as “problem articulation, norm amplification, and intergovernmental signaling,” see Doug Kysar, *The Public Life of Private Law: Tort Law as a Risk Regulation Mechanism*, 9 J. RISK REG. 48, 49–50 (2018). Cf. Guido Calabresi & Spencer Smith, *On Tort Law’s Dualisms*, 135 HARV. L. REV. F. 184, 184 (2022) (arguing that tort law has both public and private faces: “At one level, tort law *is* about wrongs and redress. That is the private side of torts. . . . At another level, tort law *is* about preventing harms or, if you like, about the regulatory needs of society. That is the public side of torts.”).

¹¹ See, e.g., *Authors Guild v. OpenAI, Inc.*, No. 23-CV-08292 (S.D.N.Y. Sept. 19, 2023) (WestLaw); *Tremblay v. OpenAI, Inc.*, No. 23-CV-03223, 2024 WL 557720 (N.D. Cal. June 28, 2023) (WestLaw); *Silverman v. OpenAI, Inc.*, No. 23-CV-03416 (N.D. Cal. July 7, 2023) (WestLaw); *Chabon v. OpenAI, Inc.*, No. 23-CV-04625 (N.D. Cal. Sept. 8, 2023) (WestLaw); *Chabon v. Meta Platforms, Inc.*, No. 23-CV-04663 (N.D. Cal. Sept. 12, 2023) (WestLaw); *Kadrey v. Meta Platforms, Inc.*, No. 23-CV-03417, 2024 WL 235199 (N.D. Cal. Jan. 22, 2024); *Andersen v. Stability AI Ltd.*, No. 23-CV-00201, 2023 WL 7132064 (N.D. Cal. Oct. 30, 2023); *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 23-CV-00135 (D. Del. Feb. 3, 2023) (WestLaw).

¹² See, e.g., *A.T. v. OpenAI LP*, No. 23-CV-04557 (N.D. Cal. Sept. 5, 2023) (WestLaw); *L. v. Alphabet Inc.*, No. 23-CV-03440 (N.D. Cal. July 11, 2023) (WestLaw). See also *P.M. v. OpenAI LP*, No. 23-CV-03199 (N.D. Cal. June 28, 2023) (WestLaw) (case dismissed without prejudice by plaintiffs).

¹³ See Cuéllar, *A Common Law for The Age of AI*, *supra* note 2, at 1776 (“[C]ommon law ideas tend to set the terms for conversations among elites and even the larger public about the way social and economic interactions ordinarily occur, and how public agencies should analyze the problems . . . they are designed to mitigate.”).

¹⁴ Kysar, *The Public Life of Private Law*, *supra* note 10, at 49–50.

legislation¹⁵ or if regulatory interventions cannot effectively address all relevant aspects of AI systems.¹⁶

The bad is substantive and intrinsic: Generative AI litigation and adjudication may produce low-quality decisions,¹⁷ as well as fail to address and redress collective harms.¹⁸ In addition to generally applicable concerns about unrepresentative facts producing bad legal rules¹⁹ and inequities in litigation,²⁰ the facts presented in emerging generative AI lawsuits may limit available causes of action under entrenched precedents.²¹ For example, courts tend to dismiss privacy-related claims, leaving plaintiffs unable to obtain relief unless they can frame their injury in sufficiently concrete terms.²² By way of further illustration, copyright

¹⁵ For an overview of the “positive political theory” literature on the relationship between political division and litigation and an empirical analysis of the relationship at the state level, as mediated by economic factors, see Tonja Jacobi, *The Role of Politics and Economics in Explaining Variation in Litigation Rates in the U.S. States*, 38 J. LEGAL STUD. 205 (2009).

¹⁶ Cf. David Freeman Engstrom, *The Automated State: A Realist’s View*, 56 GEO. WASH. L. REV. (forthcoming 2024) (manuscript on file with author) (contending that AI accountability in the public sector will be litigated, not legislated).

¹⁷ I follow Professor Schauer’s framework and reserve a separate set of concerns about decisional legitimacy. See Schauer, *Do Cases Make Bad Law?*, *supra* note 6, at 914 (distinguishing decisional quality from decisional legitimacy, and focusing on the former).

¹⁸ Indeed, while I reserve further treatment for future work, there are profound system-level questions about how neoliberal managerial judicial processes engage with technological harms at scale. See JULIE COHEN, BETWEEN TRUTH AND POWER: THE LEGAL CONSTRUCTIONS OF INFORMATIONAL CAPITALISM 143–69 (2019) (“The gradual but accelerating movement to informational capitalism has confronted the judicial system with two large and interrelated problems: a proliferation of asserted harms that are intangible, collective, and highly informationalized; and an unmanageably large and ever-increasing number of claimants and interests.”).

¹⁹ See discussion *infra* Part III.B.

²⁰ For a classic account from the law and society literature, see generally Marc Galanter, *Why the “Haves” Come Out Ahead: Speculations on the Limits of Legal Change*, 9 L. & SOC. REV. 95 (1974).

²¹ This is not to say that the doctrines necessarily ought to be this way, but rather to emphasize what presently is and how this state of affairs is likely to affect generative AI adjudication on the ground. See COHEN, BETWEEN TRUTH AND POWER, *supra* note 18, at 144 (“[J]udges cannot simply proclaim that courts are to be remade as institutions for resolving mass disputes involving intangible, collective harms without colliding headlong with long-standing institutional traditions.”).

²² On courts’ failures to recognize privacy harms and tendency to credit only “visceral and vested harms,” see Daniel J. Solove & Danielle Keats Citron, *Risk and Anxiety: A Theory of Data Breach Harms*, 96 TEX. L. REV. 737, 754–55 (2018); Danielle Keats Citron & Daniel J. Solove, *Privacy Harms*, 102 B.U. L. REV. 793, 796–99 (2022). On the “injury-in-fact” prong of standing doctrine as a “generally reactionary (i.e., noninterventionist) institutional response to the problem of harm,” see COHEN, BETWEEN TRUTH AND POWER, *supra* note 18, at 146–47.

Even at this early stage, and considering just the fourteen causes of action analyzed in this Essay, see *infra* note 39 and accompanying text, at least one generative AI litigation case speaks to the judicial system’s general tendency to fail to recognize privacy harms and to credit only “visceral and vested” harms. Specifically, in *Doe 1 v. GitHub*, the court dismissed the plaintiffs’ privacy-based claims on the grounds that plaintiffs had not amply established the injury-in-fact prong of standing. *Doe 1 v. Github, Inc.*, No. 22-CV-06823, 2023 WL 3449131, at *4 (N.D. Cal. May 11, 2023) (“Plaintiffs do not identify the specific sensitive or private information at issue[, and] . . .

disputes expose questions about the equity of labor displacement in ways that may sit uneasily with received understandings of the law.²³ The consolidation of similar generative AI lawsuits in just a few jurisdictions heightens these concerns.²⁴ On the ground, a trial judge faced with, say, a set of three practically identical class action lawsuits involving copyright infringement, unfair competition, negligence, and unjust enrichment claims against the same big tech company that is headquartered in their district will have to strike a difficult balance between resolving the particularized facts of the dispute before the court and speaking to the social outcomes that the AI system makes more or less likely.²⁵ The decisions ultimately rendered may not reflect broader societal discussions or the perspectives of non-litigants, even though the legal dispute sets forth a rule that may both be relevant in subsequent disputes and implicate collective considerations.

Recognizing how courts may be more or less auspicious sites for the emergence of AI governance, and why, allows us to be strategic in crafting public regulation to complement existing common law and statutory causes of action. In particular, it creates space to consider how litigation dynamics (such as motivations for bringing a lawsuit, or motivations to settle); the breadth of a holding in a particular

therefore do not allege facts sufficient for the Court to infer that Plaintiffs' privacy interests are implicated by the alleged misuse of such information.”).

²³ See, e.g., Matthew Sag, *Fairness and Fair Use in Generative AI*, FORDHAM L. REV., https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4654875 [https://perma.cc/DHN9-X6DE] (forthcoming) (manuscript at 9–11); Mark Lemley & Bryan Casey, *Fair Learning*, 99 TEX. L. REV. 743, 766–69 (2021); Benjamin L.W. Sobel, *Artificial Intelligence's Fair Use Crisis*, 41 COLUM. J. L. & ARTS, 45, 75–79 (2017). Cf. Pamela Samuelson, *Fair Use Defenses in Disruptive Technology Cases*, UCLA L. REV., https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4631726 [https://perma.cc/H8VV-67Z] (forthcoming) (manuscript at 63–64) (focusing on generative AI lawsuits involving the use of copyrighted works as training data and asserting that “generative AI developers would seem to have a reasonable chance to prevail with fair use defenses” based on “computational use precedents,” but also noting that “[w]ith many authors and artists railing against them for “stealing” their works, generative AI developers may encounter more resistance in the courts than they may expect”); Katherine Lee, A. Feder Cooper, & James Grimmelman, *Talkin' 'Bout AI Generation: Copyright and the Generative-AI Supply Chain*, J. COPYRIGHT SOC'Y (forthcoming 2024) (manuscript at 30) (“[C]ontent creators and their original works are part of the generative-AI supply chain, whether they would like to be or not.”).

²⁴ See discussion *infra* Part III.B.

²⁵ This is not a hypothetical proposition. As of January 2024, when this Essay was substantively finalized, there was already considerable consolidation. See *Tremblay v. OpenAI*, *Silverman v. OpenAI*, and *Chabon v. OpenAI* (related before same judge in N.D. Cal.); *Chabon v. Meta Platforms* and *Kadrey v. Meta Platforms* (related before different judge in N.D. Cal.). See also *Authors Guild v. OpenAI*, 23-CV-08292 (S.D.N.Y. Sept. 19, 2023); *Alter v. OpenAI, Inc.*, No. 23-CV-10211 (S.D.N.Y. Nov. 21, 2023) (also captioned as *Sancton v. OpenAI*); *N.Y. Times Co. v. Microsoft Corp.*, No. 23-CV-11195 (S.D.N.Y. Dec. 27, 2023); and *Basbanes v. Microsoft Corp.*, No. 24-CV-00084 (S.D.N.Y. Jan. 5, 2024) (presenting similar generative AI copyright claims and all related to *Authors Guild v. OpenAI* and *Alter v. OpenAI* in S.D.N.Y.). The S.D.N.Y. court presiding over both *Authors Guild v. OpenAI* and *Alter v. OpenAI* granted the parties' motion to consolidate the cases in late January 2024, and the *Basbanes* plaintiffs have moved to consolidate their lawsuit with the *Authors Guild / Sancton (Alter)* litigation. See, respectively, *Alter v. OpenAI*, No. 23-CV-10211 (S.D.N.Y. Nov. 21, 2023), Docket No. 31 (Jan. 22, 2024) (order granting motion to consolidate *Authors Guild* and *Alter*); (motion to consolidate *Basbanes* with *Authors Guild* and *Sancton*); *Basbanes v. Microsoft Corp.*, No. 24-CV-00084 (S.D.N.Y. Jan. 5, 2024), Docket No. 21 (Jan. 22, 2024).

case; or the types of remedies available may open up or close down further regulatory development.²⁶ This insight works in the opposite direction, too: it allows us to consider judicial interventions that might correct for potential limitations of “top-down” regulatory interventions. Indeed, this way of understanding AI governance is powerful because it allows us to push past a static dichotomy of courts and legislatures as mutually exclusive options and to instead consider how issues might percolate back and forth between the judicial and legislative system in salutary ways.

This Essay proceeds in three parts. Part II briefly makes the case that contemporary AI adjudication is a form of AI regulation, invoking a longstanding legal literature on statutory and common law judging as well as judicial rulemaking. Part III then turns to recently filed generative AI complaints. Part III.A pinpoints the good agenda-setting, deliberative work that this litigation is doing and suggests that this role might promote constructive interactions between courts and other public regulators. Part III.B identifies how these cases might be bad from a decisional quality standpoint and from a collective perspective, focusing on the risk that AI claims are predominantly being pursued by economically lucrative sectors, such as established authors and artists; the potential for path dependence and stasis in the resolution of suits; and the tension between social concerns and current legal precedents. Part IV sketches potential implications, including the counterintuitive proposal that we may need more litigation to mitigate the bad outcomes of AI adjudication. Part V concludes by urging public officials focusing on AI governance to situate AI adjudication and public regulation as complementary, and interactive, regulatory approaches.

II. AI REGULATION IN AND THROUGH COURTS

This Part assesses how courts and the adjudicatory process act as regulation and lays a foundation for Part III’s evaluation of pending generative AI adjudication.

Framing AI adjudication as regulation might seem obvious to anyone familiar with the common law or with relatively open-textured “common law statutes.”²⁷ Yet too often, the literature on legal options to regulate risk and guard against

²⁶ A corollary question involves extra-judicial remedies, such as “machine unlearning” that might permit an LLM to remove just one affected person’s data. *See* Fabian Pedregosa & Eleni Triantafillou, *Announcing the First Machine Unlearning Challenge*, Google Research Blog (June 29, 2023), <https://blog.research.google/2023/06/announcing-first-machine-unlearning.html> [<https://perma.cc/U3KL-EER4>]. I reserve further consideration of such technological remedies, and their relationship to judicial remedies, for future work.

²⁷ On federal common law statutes and the statutory interpretation challenges they present, *see generally* Charles Tyler, *Common Law Statutes*, 99 NOTRE DAME L. REV. 669 (2023).

societal harms in general,²⁸ as well as on AI governance in particular,²⁹ seems to position *ex ante* regulation through the legislature and *ex post* liability through the courts as distinct options. That positioning misses a lot. Adjudication can and does play a powerful regulatory role.

Two especially significant factors contribute to AI adjudication as regulation in the sense intended in this Essay: The regulatory impact of a judge as they resolve a particular dispute and the regulatory incentives created by the outcome of a particular case.³⁰

First, and most straightforwardly, common law adjudication entails the making of law.³¹ As Frederick Schauer explains, the “[c]ommon law method is not simply the discovery of immanent law, but rather an approach in which the decision of live disputes in concrete contexts guides the lawmaking function.”³² In a stylized sense, because such lawmaking does not promulgate a general rule that applies prospectively, it might be considered distinct from a legislative enactment.³³ Nonetheless, as Jeffrey Rachlinski discusses in his response to Professor Schauer, “[c]ase-by-case adjudication produces law when courts adopt general principles to decide the outcome of individual disputes.”³⁴ Of course, this lawmaking function can be stronger or weaker in a given instance.³⁵ Its force will depend on factors

²⁸ For one classic analysis positioning liability and regulation of safety as opposing archetypes, see Steven Shavell, *Liability for Harm versus Regulation of Safety*, 13 J. LEGAL STUD. 357 (1984). See also, e.g., Christopher H. Schroeder, *Lost in the Translation: What Environmental Regulation Does that Tort Cannot Duplicate*, 41 WASHBURN L.J. 599 (2002) (identifying and assessing procedural, remedial, and substantive differences between tort liability and environmental regulation). But see Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249, 1278–81 (2008) (exposing how public sector automated systems can blur rulemaking and adjudication).

²⁹ For recent work that “identifies the arguments for using largely *ex ante* risk regulation, as opposed to focusing on *ex post* liability,” for AI governance, see Kaminski, *Regulating the Risks of AI*, *supra* note 4, at 1354 (emphasis added). In past work, I have implicitly reinforced this dichotomy by focusing on *ex ante* interventions for AI governance. See *id.* at 1368 n.94 (citing Alicia Solow-Niederman, *Administering Artificial Intelligence*, 93 S. CAL. L. REV. 633, 688 (2020)).

³⁰ Further substantiating the argument that what happens in the judicial system has a regulatory effect far beyond the four walls of the courthouse, there is also a voluminous literature on “regulation by litigation,” “regulation through litigation,” and myriad other variations on these terms. See, e.g., Diego A. Zambrano, *Discovery as Regulation*, 119 MICH. L. REV. 71, 75–76 n.13 (2020) (compiling “long line of works that describe litigation more generally as a form of regulation”).

³¹ This claim may once have been contentious, but it is now commonly accepted, at least when it comes to common law claims. See Schauer, *supra* note 6, at 883.

³² *Id.*

³³ Rachlinski, *supra* note 7, at 933.

³⁴ *Id.* at 993. See also Emily Sherwin, *Judges as Rulemakers*, 73 U. CHI. L. REV. 919, 920 (responding to Schauer, *supra* note 7) (“[M]odern judges admit frankly that in the course of explaining their decisions, they create law for future cases.”).

³⁵ See Schauer, *supra* note 6, at 889 (recognizing that the argument that courts “make law” “is stronger insofar as the announced norm is precise—a rule and not a standard—and highly stringent—having great precedential force and significant weight against countervailing interests”).

such as how strongly a court sees itself as bound by precedent³⁶ and whether the court's resolution establishes a more flexible or vaguer standard, or instead stipulates a firmer or more explicit rule.³⁷ But the bottom line is that adjudication can regulate.

Second, there is a rich literature that underscores the regulatory function of the kinds of common law claims that appear in pending AI litigation.³⁸ For example, one generative AI lawsuit presses a negligence claim for OpenAI's use of personal data in training Chat-GPT;³⁹ another raises a defamation claim for Chat-GPT's allegedly libelous outputs.⁴⁰ Notably, common law tort claims such as these are

³⁶ See *id.* at 889. In broad strokes, the extent to which the resolution of a specific dispute will “make law” will depend on how binding a later court believes the earlier court's pronouncement to be. Suppose that courts give zero weight to precedent. If so, then resolution of a particular dispute might be said to make no law: The court would resolve that dispute, and it could articulate its reasons for doing so, but the judge would not announce any rule to guide parties in future cases. In a world of limited judicial resources, however, there would be sharp practical costs to this path. And there might be good reasons for a jurist to follow precedent. See, e.g., Richard M. Re, *Precedent as Permission*, 99 TEX. L. REV. 907, 910 (2021) (suggesting that there are advantages to making it “permissible” for judges to follow precedent and identifying “precedent as shortcut,” wherein a judge “might opt in favor of it for the sake of ease or efficiency” and “precedent as shield,” wherein “precedent's permissibility could help a judge resist political and other pressures to rule in more adventurous ways”); Sebastian Lewis, *Precedent and the Rule of Law*, 41 OXFORD J. LEG. STUD. 873, 874 (2021) (arguing that a court “always has a reason” to resolve a “precedent-governed dispute” by following precedent, because doing so promotes the values of stability, reliability, and equality, which are “part and parcel of the rule of law ideal”).

³⁷ For further discussion of rules and standards in adjudication, see Schauer, *supra* note 6, at 889 (“[F]or all but the most vacuous of norms—‘do the right thing,’ for example—a norm set forth by the deciding court will operate as constraining law for future cases.”) (citations omitted).

³⁸ For prescient pre-generative AI work on negligence and AI systems, see generally Selbst, *Negligence and AI's Human Users*, *supra* note 4.

³⁹ *A.T. v. OpenAI LP*, No. 23-CV-04557 (N.D. Cal. Sept. 5, 2023) (WestLaw). See also *P.M. v. OpenAI LP*, No. 23-CV-03199 (N.D. Cal. June 28, 2023) (WestLaw) (voluntarily dismissed without prejudice by plaintiffs). Another case, *Doe 1 v. GitHub*, also included negligence claims for alleged privacy violations, which the judge dismissed with leave to amend in May 2023. *Doe 1 v. GitHub, Inc.*, No. 22-CV-06823, 2023 WL 3449131 (N.D. Cal. May 11, 2023), ECF No. 95. The amended complaint did not include any claims for negligence arising from alleged privacy invasions, though it does allege negligent interference with prospective economic relations as well as other tort and common law claims. See First Amended Complaint, *Doe 1 v. GitHub*, No. 22-CV-06823, 2023 WL 3449131 (N.D. Cal. June 8, 2023), ECF No. 98. In January 2024, the court dismissed the *Doe 1* plaintiffs' state law claims. Order Granting in Part and Denying in Part Motion to Dismiss, *Doe 1 v. GitHub*, No. 22-CV-06823, 2023 WL 3449131, at *16 (N.D. Cal. Jan 22, 2023), ECF No. 195 (“Plaintiffs’ state law claims for intentional and negligent interference with prospective economic relations, unjust enrichment, negligence, and unfair competition are dismissed with prejudice.”) & *16 n.9 (dismissing unfair competition claim “only to the extent that it is predicated on [plaintiffs’] state law claims”).

⁴⁰ *Walters v. OpenAI, L.L.C.*, No. 23-CV-03122 (N.D. Ga. July 14, 2023) (WestLaw); *Walters V. OpenAI, L.L.C.*, No. 23-13843 (11th Cir. Nov. 17, 2023) (WestLaw). For analysis of the potential liability of AI companies for libelous outputs, see *Legal Explainer, Can I Sue a GenAI Company for Defamation if its Tool Generates False Information About Me?*, KNOWING MACHINES (Oct. 24, 2023), <https://knowingmachines.org/knowing-legal-machines/legal-explainer/questions/can-i-sue-a-genai-company-for-defamation-if-its-tool-generates-false-information-about-me> [https://perma.cc/F8LN-RA8T]; Eugene Volokh, *Large Libel Models? Liability for AI Output*, 3 J.

often said to promote social objectives like deterrence, with the threat of liability mediating the behavior of potential defendants.⁴¹ From this law and economics-inflected perspective, tort cases act prospectively, regulating future risk.⁴² The very common law causes of action that feature prominently in AI suits, and the potential outcomes of those legal claims, are part of the regulatory equation, too.

To be sure, most early generative AI lawsuits involve intellectual property claims that derive from the Copyright Act of 1976.⁴³ This might seem to disprove or at least diffuse the claim that AI adjudication is a form of common law-like judicial regulation. Statute-based law is often juxtaposed with the common law, after all.⁴⁴ Such lawmaking through legislation “build[s] law from the top down by creating general principles that cover future disputes,” in contrast to “bottom-up” lawmaking through the courts.⁴⁵ Claims arising from the statutory mandates of the Copyright Act might therefore seem quite different from claims that sound in the standards of the common law.

But this conclusion is misguided because the Copyright Act is a “common law statute.”⁴⁶ Resolving copyright claims is not a simple act of applying a text that affords no judicial discretion. To the contrary, as Shyamkrishna Balganesh has

FREE SPEECH L. 489, 510 n.75 (2023) (concluding that Walters’ case is unlikely to prevail under current doctrine).

⁴¹ See, e.g., Richard A. Posner, *Regulation (Agencies) versus Litigation (Courts): An Analytical Framework*, in REGULATION VS. LITIGATION: PERSPECTIVES FROM ECONOMICS AND LAW 11 (Daniel P. Kessler, ed. 2010) (“Economic analysis of law treats common law fields, especially tort law—which provides legal remedies for . . . wrongful conduct—as forms of regulation.”). On the centrality of deterrence and risk regulation for modern tort law, see, e.g., Catherine M. Sharkey, *Modern Tort Law: Preventing Harms, Not Recognizing Wrongs*, 134 HARV. L. REV. 1423, 1425 (2021) (reviewing JOHN C.P. GOLDBERG & BENJAMIN C. ZIPURSKY, *RECOGNIZING WRONGS* (2020)) (“[L]aw-and-economics, deterrence-based theory holds the most promise for judges facing two primary challenges of modern tort law: (1) containing risks at the cutting edge of the regulatory state and (2) addressing widespread harms.”).

⁴² See Kysar, *supra* note 10, at 49. See also Marchant & Stevens, *supra* note 4, at 242–44; Kaminski, *Regulating the Risks of AI*, *supra* note 4, at 1395 n.266.

⁴³ Several of the pending intellectual property claims also include claims for negligence under state law; however, this Essay assumes *arguendo* that they will be preempted by the federal Copyright Act and therefore does not consider them here. Indeed, early dispositions indicate as much. See, e.g., Order on Motions to Dismiss and Strike, *Andersen v. Stability AI Ltd.*, No. 23-CV-00201 (N.D. Cal. Oct. 30, 2023), ECF No. 117; Order Granting Motion to Dismiss, *Kadrey v. Meta*, No. 23-CV-03417 (N.D. Cal. Nov. 20, 2023), ECF No. 56. In the event they are not preempted, it would only bolster this Essay’s claims about the salience of common law causes of action as a form of generative AI regulation.

⁴⁴ See, e.g., Richard A. Posner, *Chapter 8 – The Common Law versus Statute Law*, in THE PROBLEMS OF JURISPRUDENCE 247, 247–61 (1990).

⁴⁵ Rachlinski, *supra* note 7, at 934.

⁴⁶ Shyamkrishna Balganesh, *Causing Copyright*, 117 COLUM. L. REV. 1, 10 n.51 (2017); Shyamkrishna Balganesh, *Debunking Blackstonian Copyright*, 118 YALE L.J. 1126, 1167–68 (2009). See also Joseph P. Liu, *Regulatory Copyright*, 83 N.C. L. REV. 87, 99–101 (2004); Margaret H. Lemos, *Interpretive Methodology and Delegations to Courts: Are “Common-Law Statutes” Different?*, in INTELLECTUAL PROPERTY AND THE COMMON LAW 89, 90 (Shyamkrishna Balganesh ed., 2013) (citing Balganesh, *Debunking Blackstonian Copyright*, at 1167).

persuasively argued, “[c]opyright adjudication is both rulemaking and rule enforcing, in a contextual setting.”⁴⁷ A court resolving a generative AI intellectual property controversy cannot mechanistically reference the text of the Act to determine the metes and bounds of the “idea-expression” dichotomy and the limits of non-expressive use of copyrighted works in the age of AI. The judge must make the call in context, filling in the gaps of the Copyright Act in a manner that, like common law determinations, makes law.⁴⁸ What’s more, content creators’ allegations of copyright infringement in generative AI suits are subject to a fair use defense. And because fair use began as a common law doctrine and retained the evolutionary character of the common law even when the Copyright Act codified this domain, fair use remains common law-like.⁴⁹ Just as other common law determinations will regulate AI’s impact on society, so too will the resolution of these copyright infringement disputes shape the path of AI law.

* * *

As legislative and regulatory bodies seek to govern AI through new or expanded regulatory interventions, the practical reality is that courts are already acting as a site of AI regulation. The key questions, then, are how this regulation is emerging from the bottom up, what the promise and perils of this regulatory mode might be, and how interaction between top-down and bottom-up approaches might improve AI governance. The next two Parts take up these issues.

III. JUDGING GENERATIVE AI TODAY

This Part draws from the legal scholarship discussed in Part II to consider what pending generative AI litigation might tell us about the present and potential future of AI regulation. Part III.A focuses on the good, and Part III.B turns to the bad. The analysis does not purport to offer comprehensive, up-to-the-minute coverage of every development in every case, nor does it endeavor to canvass the procedural or substantive details associated with every complaint. Rather, this Essay uses examples from early generative AI litigation to illustrate the potential merits and

⁴⁷ Balganes, *Debunking Blackstonian Copyright*, *supra* note 46, at 1171.

⁴⁸ Cf. Shyamkrishna Balganes, *Copyright as Legal Process: The Transformation of American Copyright Law*, 168 U. PA. L. REV. 1101, 1116 (2020) (arguing that copyright is a “legal process” that “rel[ies] on the private enforcement (of individual rights) to realize its public-regarding (collectivist) goals” and drawing parallels between this shift to “public law” and similar shifts in tort law and contract law).

⁴⁹ This understanding is consistent with the House Report on the Copyright Act of 1976, which first codified the fair use doctrine. See H.R. REP. NO. 94-1476, at 66 (1976) (“The bill endorses the purpose and general scope of the judicial doctrine of fair use, but there is no disposition to freeze the doctrine in the statute . . . Beyond a very broad statutory explanation of what fair use is and some of the criteria applicable to it, the courts must be free to adapt the doctrine to particular situations on a case-by-case basis. Section 107 is intended to restate the present judicial doctrine of fair use, not to change, narrow, or enlarge it in any way.”). See also Liu, *supra* note 46, at 100 (noting that the 1976 Act codified the fair use defense “while leaving room for the courts to continue to develop the defense in a case-by-case manner”).

demerits of common law judging about generative AI, which places the regulatory stakes of AI adjudication in stark relief.⁵⁰

To develop this picture, consider the initial complaint in *Silverman v. OpenAI*, a class action suit filed against the maker of Chat-GPT.⁵¹ The heart of the complaint is simple: The stand-up comic and author Sarah Silverman, along with two other authors, alleges that OpenAI ingested their books without permission and used the authors' raw material as part of the massive dataset required to train its LLM,⁵²

⁵⁰ To gain traction and permit concrete analysis of a rapidly developing space, this Essay focuses on the fourteen private law causes of action involving generative AI that had been filed in American courts as of October 1, 2023. To identify these cases, I began with the cases logged in the AI Litigation Database as of that date. See *AI Litigation Database*, ETHICALTECH@GW, <https://blogs.gwu.edu/law-eti/ai-litigation-database-search> [<https://perma.cc/CSR7-ZMS4>] (last visited Oct. 1, 2023). I excluded one case brought in the U.K. and one case sounding in the Administrative Procedure Act and then conducted a supplementary Bloomberg Law search for cases filed against Alphabet, GitHub, Google, Meta, Microsoft, OpenAI, or Stability AI, or which involved the search term “generative AI.” This search revealed one additional case, *Battle v. Microsoft Corp.*, No. 23-CV-01822 (D. Md. July 7, 2023) (Westlaw), that was dismissed without prejudice in August 2023. Order Dismissing the Complaint, *Battle v. Microsoft Corp.*, No. 23-CV-01822 (D. Md. July 7, 2023, ECF No. 9. In addition, one case, *P.M. v. OpenAI*, No. 23-CV-03199 (N.D. Cal. June 28, 2023), was voluntarily dismissed without prejudice by all plaintiffs in September 2023. I excluded these two dismissed cases from the analysis. As reported in the AI Litigation Database, four additional lawsuits were filed in American courts between October 2023 and January 2024: *UMG v. Anthropic*, No. 23-CV-01092 (M.D. Tenn. Oct. 18, 2023); *Huckabee v. Meta Platforms, Inc.*, No. 23-CV-09152 (S.D.N.Y. Oct. 17, 2023); *Sancton v. OpenAI*, now known as *Alter v. OpenAI*, No. 23-CV-10211 (S.D.N.Y. Nov. 21, 2023); and *N.Y. Times v. Microsoft*, No. 23-CV-11195 (S.D.N.Y. Dec. 12, 2023). Some of the claims in the S.D.N.Y. suits in the *Huckabee* case were subsequently transferred to the Northern District of California. *Huckabee v. Meta Platforms, Inc.*, No. 23-CV-06663 (N.D. Cal. Dec. 28, 2023). See discussion *infra* note 88. In addition, a review of Bloomberg Law dockets revealed one additional case, *Basbanes v. Microsoft Corp.*, No. 24-CV-00084 (S.D.N.Y. Jan. 5, 2024), that was filed in January 2024. As of this writing, there had been partial consolidation of some of these suits, with possible further consolidation pending. See discussion *supra* note 25. Because further consolidation would only strengthen the claims in this Essay, I reserve further consideration of the procedural details and emphasize what we can discern from early signals about the likely path of AI adjudication.

⁵¹ For an accessible and informative explainer of the lawsuit, see Victoria Bekiempis, *Can Sarah Silverman's AI Lawsuit Save Us from Robot Overlords?*, VULTURE (Aug. 11, 2023), <https://www.vulture.com/article/sarah-silverman-openai-meta-lawsuit-explainer.html> [<https://perma.cc/V3PP-ZU42>].

⁵² Complaint at ¶¶ 22–27, *Silverman v. OpenAI*, No. 23-CV-03416 (N.D. Cal. July 7, 2023), ECF No. 1. The same parties have also sued Meta for what is functionally the same legal theory. See *Kadrey v. Meta Platforms*, No. 23-CV-03417, 2024 WL 235199 (N.D. Cal. Jan. 22, 2024). Moreover, a different set of named plaintiffs, represented by a different firm, has also sued OpenAI with nearly identical language to the *Silverman* filing in their complaint. See Complaint at ¶¶ 22–27, *Tremblay v. OpenAI, Inc.*, No. 23-CV-03223 (N.D. Cal. June 28, 2023), ECF No. 1. The language in the *Tremblay v. OpenAI* complaint is quite similar to the text of a different suit, *Chabon v. OpenAI, Inc.*, No. 23-CV-04625 (N.D. Cal. Sept. 8, 2023). Separately, the attorneys in *Chabon v. OpenAI* have filed a similar suit against Meta. See *Chabon v. Meta Platforms*, No. 23-CV-04663 (N.D. Cal. Sept. 12, 2023). For purposes of this Essay, suffice it to say that these claims overlap. For discussion of the consolidation of these lawsuits and the potential broader effects of this phenomenon for the regulatory path of AI adjudication, see text accompanying notes 88 to 97, *infra*.

thereby gaining the capacity to spit out “very accurate summaries” of the plaintiffs’ creative works.⁵³

Critically, there are both legal and social questions presented. The legal questions include copyright infringement, negligence, and unjust enrichment.⁵⁴ The social questions involve the fundamental relationship between generative AI systems, the companies developing and releasing them, and content creators. Moreover, the issues presented extend beyond the pages of the complaint: No matter what happens in the litigation itself, and even if particular claims are ultimately dismissed, they represent public considerations that resonate far more widely.⁵⁵ The case is more than a “retail,” individualistic, *ex post* claim for relief. It is a “wholesale” call for a collective reckoning about generative AI systems and their impact on society.⁵⁶ Indeed, other parties have brought a series of similar lawsuits against OpenAI as well as against Meta Platforms (the maker of the LLaMA, or “Large Language Model Meta AI” system); Stability AI (the maker of a text-to-image generative AI system, Stable Diffusion); and Google (the maker of the Bard chatbot as well as several other generative AI tools). Among these other litigants is the Authors Guild, which has separately filed a class action against OpenAI⁵⁷ and also issued an open letter, signed by thousands of authors, decrying the way that “[m]illions of copyrighted books, articles, essays and poetry provide the ‘food’ for AI systems, endless meals for which there has been no bill[.]”⁵⁸

⁵³ Complaint at ¶ 42, *Silverman v. OpenAI*, No. 23-CV-03416 (N.D. Cal. July 7, 2023), ECF No. 1.

⁵⁴ The original *Silverman* complaint includes allegations of copyright infringement, negligence, unjust enrichment, unfair competition, and alleged removal of copyright management information (“CMI”) in violation of the Digital Millennium Copyright Act. *Id.* at ¶¶ 52–87. Other pending suits include a similar mix of claims. Recognizing that the state law claims in the intellectual property cases are likely preempted by the Copyright Act, this Essay lists all the pending claims both to provide a sense of what legal theories are being asserted in these cases and to emphasize the regulatory role that the judge must play in filtering out which claims may (not) properly proceed.

⁵⁵ For an argument that generative AI must be understood as a “public problem,” drawing on John Dewey’s definition of this concept, see Mike Ananny, *To Reckon with Generative AI, Make It a Public Problem*, ISSUES IN SCI. & TECH. 88, 88 (Winter 2024) (“Although generative AI has been framed as a technical problem, recasting it as a public problem offers new avenues for action. Generative AI is quickly becoming a language for telling society’s collective stories and teaching us about each other.”).

⁵⁶ I am hardly the first scholar to suggest that the retail, “liability in tort” model of enforcement tends to be a bad fit for AI systems, though this Essay is distinct in connecting this point to generative AI adjudication by private actors. For arguments in the public law setting, see, e.g., Aziz Z. Huq, *Constitutional Rights in the Machine-Learning State*, 105 CORNELL L. REV. 1875, 1937–38 (2020); Engstrom, *supra* note 16 (manuscript at 4) (citing Citron, *Technological Due Process*, *supra* note 28); DANIEL E. HO ET AL., QUALITY ASSURANCE SYSTEMS IN AGENCY ADJUDICATION: EMERGING PRACTICES AND INSIGHTS (Nov. 30, 2021) (report to the Admin. Conf. of the U.S.); David Ames et al., *Due Process and Mass Adjudication: Crisis and Reform*, 72 STAN. L. REV. 1 (2020); Daniel E. Ho & Sam Sherman, *Managing Street-Level Arbitrariness: The Evidence Base for Public Sector Quality Improvement*, 13 ANN. REV. L. & SOC. SCI. 251 (2017).

⁵⁷ *Authors Guild v. OpenAI*, No. 23-CV-08292 (S.D.N.Y. Sept. 19, 2023).

⁵⁸ Letter from The Authors Guild, to Sam Altman, CEO, OpenAI, Sundar Pichai, CEO, Alphabet, Mark Zuckerberg, CEO, Meta, Emad Mostaque, CEO, Stability AI, Arvind Krishna,

Notably, the social considerations are complex, because there are collective values on the other side, too. The countervailing social challenge is how to preserve vital space for generative creation and remixing of others' content in a world that does not wish to lock down the free flow of information and in which transformative uses of copyrighted works are a social good. These social considerations thus go to the central purpose of the underlying substantive law, in general, and how to construe it in an age of evolving digital technologies, in particular.

Moreover, these entangled social and legal considerations are not limited to intellectual property. Emerging cases that center on privacy and property claims also implicate fundamental questions about the relationship between generative AI systems, their makers, and citizens.⁵⁹ Because these social battle lines are being channeled through the judicial system and transformed into discrete legal disputes,⁶⁰ they are resolved against the backdrop of existing information privacy law and copyright law jurisprudence and common law doctrines. That makes it essential to attend to the regulatory role of the judge, the use of courts to execute this regulatory function, and the ways in which this can be good and bad.

A. Regulation in the Courts: The Good

In the face of such thorny social challenges, adjudication might foster good AI law by improving both individual decisions and collective outcomes.

For one, framing a generative AI question as a concrete legal dispute might produce more specific and well-tailored rules. A private lawsuit requires the parties to identify a clearly bounded, live controversy and a particularized injury. This legal vehicle can focus attention on the details of the dispute, permitting a resolution that avoids speculative conclusions and instead provides a rule that accounts for the impact of a decision on the ground.⁶¹ For example, rather than attend to hypotheticals about economic loss or labor displacement in a future world where generative AI systems replace artists, a lawsuit might focus attention on much more concrete and immediate issues. Such issues could include whether the specific parts of an artist's book that were copied and used to train an AI system are expressive or non-expressive, and if they are expressive, whether fair use would apply.⁶²

CEO, IBM, Satya Nadella, CEO, Microsoft (July 2023), <https://authorsguild.org/app/uploads/2023/07/Authors-Guild-Open-Letter-to-Generative-AI-Leaders.pdf> [https://perma.cc/3SZ5-BWW3].

⁵⁹ See Complaint at ¶¶ 71–78, A.T. v. OpenAI, No. 23-CV-04557 (N.D. Cal. Sept. 5, 2023), ECF No. 1.

⁶⁰ See William L.F. Felstiner, Richard L. Abel & Austin Sarat, *The Emergence and Transformation of Disputes: Naming, Blaming, Claiming . . .*, 15 L. & SOC. REV. 631, 632–33 (1980) [hereinafter Felstiner et al.]. For more detailed analysis of the emergence and transformation of generative AI disputes and its broader implications for AI governance, see discussion *infra* Part III.

⁶¹ See Schauer, *supra* note 6, at 892. See also Rachlinski, *supra* note 7, at 936–37.

⁶² “Might” is an important qualifier. As Professors Lemley and Casey suggest, writing about machine learning before the latest generative AI surge, there is a “very strong, if often unarticulated, anti-free riding instinct in courts. And the appeal of that argument is likely to be strengthened by the

Relatedly, the ability to carefully appraise arguments made by actual litigants in context might improve the judge's understanding of the relevant facts.⁶³ As one example from the putative class action *A.T. v. OpenAI*, rather than speculate about generative AI's threats to individual privacy, concrete details about how Plaintiff A.T. used ChatGPT as well as "a variety of websites and social media applications" might help the court to grasp A.T.'s specific concerns about how he posted his political views and how he shared technical advice based on his experience as a project manager online.⁶⁴ These details might allow the court to better evaluate whether ChatGPT's scraping and use of this data on this scale, without consent, amounts to a cognizable injury to A.T.'s privacy or property interests.⁶⁵

Furthermore, even if the court disposes of the case on narrow grounds or otherwise concludes that this injury is not legally cognizable, the very act of adjudicating can be instrumentally valuable for both individuals and for the public. As Douglas Kysar argues, "judicial engagement with the details of [a plaintiff's] claim helps to frame her suffering as a legible subject of public attention and governance."⁶⁶ AI adjudication might thus not only provide dignitary benefits to the individual, but also translate an individual's asserted harm into a more generalizable social concern.⁶⁷ Especially over time and in the aggregate, the process of resolving these disputes within the judicial system might make underlying social issues appear visible and pressing in ways that warrant further regulatory action outside of the courtroom. This effect seems especially likely where, as here, lawsuits are arriving in tandem with other public relations efforts, such as the Authors Guild's public letter, as well as alongside testimony before elected representatives.⁶⁸ The public attention created by a flurry of related lawsuits might also spur inter-institutional dialogue, as seems to be occurring with

commercial nature of many ML applications." Lemley & Casey, *supra* note 23, at 765 (citations omitted).

⁶³ See Schauer, *supra* note 6, at 892 n.45 (citing Wendy Anton Fitzgerald, *Maturity, Difference, and Mystery: Children's Perspectives and the Law*, 36 ARIZ. L. REV. 11, 107 n.577 (1994); Martha Minow, *The Supreme Court 1986 Term: Foreword: Justice Engendered*, 101 HARV. L. REV. 10, 89 (1987)).

⁶⁴ See *A.T. v. OpenAI Complaint*, *supra* note 59, at ¶¶ 20–26, 172–274.

⁶⁵ *Id.*

⁶⁶ Kysar, *supra* note 10, at 50.

⁶⁷ On law's expressive function in the context of cyber gender harassment, see Danielle Keats Citron, *Law's Expressive Value in Combating Cyber Gender Harassment*, 108 MICH. L. REV. 373, 407 (2009) ("Law educates the public about what is socially harmful. This legitimates harms, allowing the harmed party to see herself as harmed. It signals appropriate behavior. . . . Because law creates and shapes social mores, it has an important cultural impact that differs from its more direct coercive effects." (citations omitted)).

⁶⁸ See *Artificial Intelligence and Intellectual Property: Hearing Before the Subcomm. on Intellectual Property of the S. Comm. on the Judiciary*, 118th Cong. 10 (2023) (statement of Karla Ortiz, named plaintiff in *Andersen v. Stability AI* lawsuit).

generative AI as agencies such as the FTC⁶⁹ enter the conversation alongside Congress, the Executive Branch,⁷⁰ and the courts.

Operating as part of an inter-institutional conversation, the adjudicatory process might play an information-forcing role for both courts and other public regulators. At its best, the information surfaced through the courts has the potential to influence subsequent efforts at public regulation of generative AI in productive ways. Litigation might, for example, require AI companies to publicly disclose more about their infamously secretive and closed internal operations⁷¹ or the nature of their corporate structure.⁷² For example, in *Walters v. OpenAI*, a pending libel suit that OpenAI attempted to remove to federal court, the district court found that the company had not adequately established diversity citizenship and ordered OpenAI to “affirmatively identify by name each member of an LLC and then allege whatever specific facts are necessary to establish the citizenship of that member.”⁷³ Rather than do so, OpenAI withdrew its notice of removal, stating that it “is not in a position to provide further information beyond its prior filing.”⁷⁴ The case has now been remanded to state court. Although information was not forced out in this instance, *Walters* suggests how litigation can put disclosure demands on private actors. A court that engages in careful judicial management of the lawsuit can play a positive regulatory role, shaping the ways that social questions are translated into legal questions within our judicial system.

Zooming out, the potential for adjudication to reveal information might also foster beneficial interactions between courts and policymakers. The information that is surfaced in a particular lawsuit might be directly helpful if it reveals the need

⁶⁹ See, e.g., Press Release, FTC, FTC to Host Virtual Roundtable on AI and Content Creation (Sept. 22, 2023), <https://www.ftc.gov/news-events/news/press-releases/2023/09/ftc-host-virtual-roundtable-ai-content-creation> [<https://perma.cc/U4TB-48G9>].

⁷⁰ See, e.g., Exec. Order No. 14110, 88 Fed. Reg. 75191 (Oct. 30, 2023).

⁷¹ On “open” AI and the often confusing contemporary use of the terms “open” and “open source,” using the company OpenAI as an example, see David Gray Widder, Meredith Whittaker, & Sarah Myers West, *Open (for Business): Big Tech, Concentrated Power, and the Political Economy of Open AI* (Aug. 17, 2023) (unpublished manuscript), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4543807 [<https://perma.cc/4Z34-2BKF>]. On OpenAI’s decision to disclose minimal information about GPT-4, see James Vincent, *OpenAI Co-Founder on Company’s Past Approach to Openly Sharing Research: ‘We Were Wrong’*, THE VERGE (Mar. 15, 2023, 10:59 AM) <https://www.theverge.com/2023/3/15/23640180/openai-gpt-4-launch-closed-research-ilya-sutskever-interview> [<https://perma.cc/B8AE-L5GB>]. On the challenge of gaining access to information about automated systems in general, see, e.g., Hannah Bloch-Wehba, *Access to Algorithms*, 88 FORDHAM L. REV. 1265 (2020).

⁷² See, e.g., *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 23-CV-00135 (D. Del. Feb. 3, 2023) (featuring jurisdictional discovery dispute that implicates questions of corporate structure).

⁷³ Order to Show Jurisdiction, *Walters v. OpenAI*, No. 23-CV-03122 (N.D. Ga. Sept. 22, 2023). See also Peter Henderson (@peterhenderson.bsky.social), BLUESKY (Sept. 27, 2023, 1:34 PM), <https://bsky.app/profile/peterhenderson.bsky.social/post/3kafehwcpw223> [<https://perma.cc/3GK2-7RCF>] (noting that the pending suit may require the company to reveal a great deal about its corporate structure).

⁷⁴ Defendant OpenAI’s Response to Order Dated September 22, 2023, *Walters v. OpenAI*, No. 23-CV-03122, at *1 (N.D. Ga. Oct. 6, 2023).

for a specific kind of targeted regulatory intervention. If a litigant reveals that its training data was not amply tested for a particular form of bias, or that it failed to check whether the training data contained copyrighted works, for example, then that revelation might suggest that a particular kind of public regulation is especially important for the legislature to prioritize. The information that is surfaced in a particular private lawsuit might also be indirectly helpful for other public actors. For instance, if the kind of information about corporate structure that the *Walters* Court asked OpenAI to produce is helpful in order for public regulators to determine the appropriate scope of AI legislation, then policymakers might consider codifying additional disclosure requirements.⁷⁵ The conditions that are more and less likely to promote such movement of issues between courts and other regulators, translating individualized, retail-level disputes into matters for collective, systemic resolution, is an essential question to which Part IV returns.

B. Regulation in the Courts: The Bad

Notwithstanding the ways that adjudication might contribute positively to the path of AI law, the role of the courts in AI disputes might also be for the worse. This Essay focuses on two overlapping concerns: First, concerns with the quality of a judicial decision itself,⁷⁶ and, second, concerns that the process of private adjudication will fall short of its deliberative potential.

Start with decisional quality. There are several reasons to worry that focusing on a single case is a suboptimal lens for resolving AI disputes. For one, this case-specific analysis might lead a judge to focus on nonrepresentative issues, announcing a rule that's ill-suited for the broader context.⁷⁷ As Mike Ananny explains, legal scholars have long worried that “the specific ‘*this-ness*’ of a case limits lawmaking because its facts are so technical, peculiar, or idiosyncratic that its reasoning cannot be generalized.”⁷⁸ By way of illustration, take one pending intellectual property suit, *Getty Images v. Stability AI*, which includes many specific visual examples in the complaint. Among these images is the following

⁷⁵ Indeed, information disclosure, especially in the face of public-private information asymmetries, is one instrumental reason that policymakers and scholars call for algorithmic audits and impact assessments. On algorithmic impact assessments as a regulatory response to AI, see Andrew Selbst, *An Institutional View of Algorithmic Impact Assessments*, 35 HARV. J. L. & TECH. 117, 119–25 and sources cited therein (2021). On the value of transparency in regulation of AI systems, see Margot E. Kaminski, *Understanding Transparency in Algorithmic Accountability*, in THE CAMBRIDGE HANDBOOK OF THE LAW OF ALGORITHMS 121 (Woodrow Barfield ed., 2020).


⁷⁶ See note 17 and accompanying text. See also Schauer, *supra* note 6, at 914 (distinguishing focus on “decisional quality” from other kinds of arguments, such as those focused on “decisional legitimacy”).

⁷⁷ See Schauer, *supra* note 6, at 894–95; Ananny, *supra* note 6, at 1 (citing Schauer, *supra* note 6, at 884).

⁷⁸ Ananny, *supra* note 6, at 1 (citing Schauer, *supra* note 6, at 884). This concern has an historic pedigree: Professor Ananny cites Professor Schauer, who is himself building from Justice Holmes’s classic adage that hard cases make bad law.

juxtaposition of a Getty Images photograph from a football match (on the left) and an output from Stability AI (on the right):⁷⁹

Figure 1: Images from *Getty Images v. Stability AI*.



Here, factors like the unnatural bend in the player's leg and the two athletes' blended arms in the Stability AI image on the right might subconsciously affect the analysis. The risk is that the court would be moved by the aesthetically unappealing transformation, regardless of whether that reflects the applicable legal standard for copyright infringement or a fulsome balancing of the factors in the fair use defense. Alternatively or in addition, the seeming inclusion of a blurry Getty Images mark in the image on the right might enter the copyright analysis when it is most relevant for the trademark claim, even though *Getty Images* is the only generative AI case to date that features a trademark claim.⁸⁰ In other words, these details may be visually alluring, yet legally irrelevant or unrepresentative of the broader universe of text-to-image generation claims, let alone copyright infringement and fair use claims writ large. But the rules that are announced could potentially appear to cover a more general set of cases.

⁷⁹ Complaint at ¶ 52, *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 23-CV-00135 (D. Del. Feb. 3, 2023). See also Matthew Sag, *Copyright Safety for Generative AI*, 61 HOUS. L. REV. 295, 311–12 (2023) (discussing images to illustrate “The Problem of ‘Memorization’” and to consider how copyright precedents might apply to arguments concerning the “potential similarity between outputs from the Stable Diffusion model and Getty’s own copyrighted images,” as well as to “broader” concerns about LLM models’ memorization of training data and potential “communicat[ion of] original expression from the training data via its output”).

⁸⁰ Cf. Schauer, *supra* note 6, at 899–900 (“If judges have a hard time avoiding what they see as the right result for the particular case in all of its contextual richness, and if they are at the same time making law for future cases, then the combination of the salience of the particular case and the pull to decide it correctly may produce a rule that is unrepresentative of the full range of future cases that can be expected to be decided under it.”).

In addition, the potential for path dependence presents a closely related concern.⁸¹ The basic worry is that a decision at “time one” mediates an outcome at “time two,” even though this later outcome may not be the product of an optimal decision-making process.⁸² For instance, a court might invoke precedent to resolve a subsequent case, even though that latter case is properly distinguished on its facts. In addition, regulated parties might alter their behavior based on their read of a time one decision, whether or not that earlier holding actually covers their conduct.

Thomson Reuters v. Ross Intelligence, a Third Circuit case filed nearly two years before the rise of chatbots like Chat-GPT, illustrates these issues.⁸³ This copyright suit alleges that the defendant made unauthorized use of the plaintiff’s legal “headnotes”⁸⁴ to train a machine-learning powered legal search engine.⁸⁵ The goal of such a “natural language” legal search engine differs from the goal of a generative AI system,⁸⁶ and the specific text at issue is a very particular kind (legal headnotes). Despite these particularized considerations, the outcome in *Thomson Reuters* might speak to whether the use of this text is copyright infringement, as well as whether it is fair use, in ways that affect subsequent generative AI cases.⁸⁷ Particularly if the earlier holding is written in an especially broad manner, the case

⁸¹ On path dependence in common law legal systems, see Oona A. Hathaway, *Path Dependence in the Law: The Course and Pattern of Legal Change in a Common Law System*, 86 IOWA L. REV. 101 (2001). See also Alicia Solow-Niederman, *Can AI Standards Have Politics?*, 71 UCLA L. REV. DISCOURSE (forthcoming 2024) (manuscript at 16 and sources cited therein) (on file with author) (assessing potential path dependence in AI standards development).

⁸² This phenomenon might also be understood as a form of “precedential cascades,” referring to a situation “in which a sequence of judges each rule on a sequence of similarly situated cases.” Eric Talley, *Precedential Cascades: An Appraisal*, 73 S. CAL. L. REV. 87, 92 (1999). See generally *id.* (discussing precedential cascades and concluding that the necessary conditions for such cascades are implausible in most cases).

⁸³ *Thomson Reuters Enter. Ctr. GmbH v. Ross Intelligence Inc.*, No. 20-CV-00613, 2023 WL 6210901 (D. Del. May 6, 2020). Thank you to Bob Brauneis for helpful discussions about this case.

⁸⁴ See Mem. Op. at 2, *Thomson Reuters v. Ross Intelligence*, No. 20-CV-00613 (D. Del. Sept. 25, 2023).

⁸⁵ See *id.* at 2–3.

⁸⁶ In the future, it is likely that generative AI will itself shape the search experience. But that is not the issue in *Thomson Reuters*. Here, Ross Intelligence was attempting to build a natural-language, semantic search engine. Complaint at ¶ 26, *Thomson Reuters v. Ross Intelligence*, No. 20-CV-00613 (D. Del. May 6, 2020). This would represent an advance over keyword search, but it would still deliver lists of legal sources for the searcher to parse. A generative AI system might, in contrast, provide direct answers to a query (ideally with citations). For further discussion of AI search capabilities and potential future developments, see Ege Gurdeniz & Kartik Hosanagar, *Generative AI Won’t Revolutionize Search — Yet*, HARV. BUS. REV. (Feb. 23, 2023), <https://hbr.org/2023/02/generative-ai-wont-revolutionize-search-yet> [<https://perma.cc/9VQ5-WJWR>].

⁸⁷ To be sure, there is an essential threshold question: In September 2023, the court denied the defendant’s motion for summary judgment and found that the jury must decide whether the headnotes are copyrightable or not. Order., *Thomson Reuters Enter. Ctr. GmbH v. Ross Intelligence Inc.*, No. 20-CV-00613 (D. Del. Sept. 25, 2023), ECF No. 548. If the jury concludes that the headnotes are not copyrightable, then there will not necessarily be a generalizable outcome that is relevant for later cases involving other AI systems.

might shape the later court's judgment about the propriety of training AI models with vast quantities of scraped data, without fully analyzing how the data is distinct from headnotes, or how each element of the fair use test should apply to the new facts. In addition, a conclusion that there is (or is not) fair use in *Thomson Reuters* might affect the present and future behavior of AI companies more generally.⁸⁸ If no liability is found, then commercial actors might read the case as a permission slip to scrape similar kinds and quantities of data. Conversely, if liability is found, then other parties might decide to settle any claims that arise, or even decide that the legal exposure is too high and that they should not develop an AI system at all.

Outcomes such as these might or might not reflect desirable social policy, but it's hard to argue with a straight face that they reflect robust deliberation about how the law should apply to specific facts. Instead, adjudication may regulate a topic with vast and far-reaching social implications based on timing and coincidences. Though the magnitude of this risk depends to some extent on factors such as the speed of adjudication and the number and type of AI lawsuits, there are reasons to be especially wary of this possibility because, as discussed next, early evidence suggests that there are a limited quantity and variety of cases on point.

Indeed, regulating generative AI through adjudication seems to risk substantial homogeneity in ways that may compromise the deliberative ideal. One concern involves who can bring these cases. Notably, many of the pending suits involve well-resourced authors, artists, or large companies. There are obvious worries about which claims do and do not get their day in court, and whether they amply represent the interests of other affected parties. To be sure, generative AI claims may not be any less representative along these or other dimensions than private lawsuits as a whole. But doing poorly across the board seems like a flimsy reason to write off unrepresentativeness in AI lawsuits, especially because of the interwoven social and legal challenges involved, and especially where the earliest cases do exhibit such homogeneous qualities.

In addition, a second concern about homogeneity involves where these claims are being resolved. The majority of pending generative AI cases are arising in the Northern District of California, where many of the big tech firms are headquartered.⁸⁹ To the extent that jurists are impartial decision-makers,

⁸⁸ See Scott Nover, *A.I. May Not Get a Chance to Kill Us if This Kills It First*, SLATE (Oct. 17, 2023, 10:00 AM) <https://slate.com/technology/2023/10/artificial-intelligence-copyright-thomson-reuters-ross-intelligence-westlaw-lawsuit.html> [<https://perma.cc/C4DN-M59>] (discussing the *Thomson Reuters* case and its potential impact on future AI development: "There is a disaster scenario for OpenAI and other companies funneling billions into A.I. models: If a court found that a company was liable for copyright infringement, it could completely halt the development of the offending model").

⁸⁹ Other than the consolidated set of cases associated with *Authors Guild v. OpenAI, Inc.*, No. 23-CV-08292 (S.D.N.Y. Sept. 19, 2023); *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 23-CV-00135 (D. Del. Feb. 3, 2023); and *Walters v. OpenAI, L.L.C.*, No. 23-CV-03122 (N.D. Ga. July 14, 2023), all of the pending lawsuits filed as of October 1, 2023, sit in the Northern District of California. See *supra* note 25 for discussion. Time will tell if subsequent cases, such as a suit filed in mid-October in the Middle District of Tennessee, counteract this pattern. See *UMG v. Anthropic*, No. 23-CV-01092 (M.D. Tenn. Oct. 18, 2023). Early trends, albeit trends drawn from a small sample

independent of location, this point may be moot. However, if jurors from the community are called on to decide questions of fact, like fair use, and local norms differ from norms in other jurisdictions;⁹⁰ if judges draw on their localized experience in deciding cases; and/or if courts become familiar with the habits of “repeat players,” whether litigants or law firms,⁹¹ in ways that affect their analysis, then where these cases arise matters. What’s more, centralization of cases in a single district might make anxieties about homogeneity more acute because it channels claims through only one part of the judicial system, rather than relying on many courts, nationwide, to resolve these questions by incrementally developing the common law.⁹² Now, the class action posture of many of the pending generative AI lawsuits might alleviate this concern insofar as the facts and claims presented adequately represent the interests of all similarly affected parties, wherever they are located. But it still cannot resolve an underlying issue: A private lawsuit is an individualized vehicle through which to approach a harm with collective dimensions.⁹³ As multiple law and technology scholars have argued, there are reasons to be wary of enforcement litigation in the context of algorithmic injuries.⁹⁴

size, suggest that there may also be consolidation in the Southern District of New York. *See Huckabee v. Meta Platforms*, No. 23-CV-09152 (S.D.N.Y. Oct. 17, 2023) (complaint filed in Southern District of New York in October 2023); *Alter v. OpenAI*, No. 23-CV-10211 (S.D.N.Y. Nov. 21, 2023) (filed in Southern District of New York in November 2023, and consolidated with *Authors Guild v. OpenAI*); *N.Y. Times Co. v. Microsoft Corp.*, No. 23-CV-11195 (S.D.N.Y. Dec. 27, 2023) (filed in Southern District of New York in December 2023, and related to *Authors Guild v. OpenAI*); *Basbanes v. Microsoft Corp.*, No. 24-CV-00084 (S.D.N.Y. Jan. 5, 2024) (filed in Southern District of New York in January 2024, and related to *Authors Guild v. OpenAI*). In *Huckabee*, the claims pertaining to Meta Platforms and Microsoft Corporation were transferred to the Northern District of California and the claim against one defendant was voluntarily dismissed, leaving claims against other defendants pending in the Southern District of New York. *See Huckabee v. Meta Platforms, Inc.*, No. 23-CV-06663 (N.D. Cal. Dec 28, 2023), ECF Nos. 68 & 96.

⁹⁰ For instance, during her book tour, the author Kashmir Hill reported that audiences in San Francisco and New York City had very different responses to “having their face recognized by augmented reality glasses if such a thing existed.” *See* Kashmir Hill (@kashhill.bsky.social), TWITTER (Oct. 23, 2023, 11:20 AM) <https://bsky.app/profile/kashhill.bsky.social/post/3kbfszmunzf2b> [<https://perma.cc/56MW-44BK>].

⁹¹ On litigation and repeat players, *see* Galanter, *supra* note 20. Many of the pending generative AI suits involve the same “repeat player” attorneys and companies.

⁹² It may also pose a host of other challenges. For instance, given the substantial overlap between claims, as well as overlap between parties, as discussed *supra* note 52, many of these cases have been related to one another. But there remain further questions. Consider the fact that *Chabon v. Meta* and *Kadrey v. Meta* have been related to one another before one judge and that *Silverman v. OpenAI*, *Tremblay v. OpenAI*, and *Chabon v. OpenAI* presently sit before a different judge in the same district. Recall, moreover, that many of the allegations in *Chabon v. Meta* are quite similar to those in *Chabon v. OpenAI*. Ensuring that the parties in each of these cases receive a fair and full hearing in court, given the overlapping topics, firms, and litigants, is thorny, to say the least.

⁹³ I follow Professors Felstiner, Abel, and Sarat in thinking of even class actions as “collections of individual disputes, aggregated for reasons of convenience and efficiency, rather than a form of collective action aimed at achieving a group objective[.]” Felstiner et al., *supra* note 60, at 648 n.13.

⁹⁴ *See, e.g.,* Mariano-Florentino Cuéllar & Aziz Z. Huq, *The Democratic Regulation of Artificial Intelligence*, at 2, 18 (Knight Inst., Data & Democracy Ser., 2022), <https://s3.amazonaws.com/kfai-documents/documents/c43514f2e3/1.31.2022-Huq---Cuellar.pdf>

* * *

In the face of these challenges, at least some might be tempted to conclude that AI adjudication is doomed to make bad law and seek alternate avenues. But that's too hasty, and too facile, a resolution. As discussed above, adjudication has notable benefits, and ignoring its role cuts out an important institutional component of our legal system while casting aside the potential of the common law to accommodate social and technological change. Just as importantly, as a practical matter, regulation of generative AI through the courts is already happening. It blinks reality to pretend it is not, even if one concludes that the risk of the bad outweighs the prospect of the good. The vital question, to which the next Part turns, is what to do about it.

IV. REPOSITIONING AI ADJUDICATION AS AI GOVERNANCE

This Part accepts AI adjudication as an indelible component of AI governance and reframes the challenge in institutional terms: First, how can we maximize bottom-up lawmaking's promise and minimize its perils within the judicial system, and second, how can we position judicial lawmaking with respect to individual litigants and individual cases as the start, but not the end, of a broader collective conversation with legislators and with affected publics?

AI adjudication will naturally do some things more or less well. Not every weakness can be corrected; adjudication as a whole has well-documented downsides,⁹⁵ and sometimes, an individual, private law fix is just a bad fit for a collective harm with social ramifications.⁹⁶ Still, it's worth considering internal fixes to the adjudication process.

The most auspicious internal fixes require recognizing that many of the issues with early generative AI cases are not inherently about adjudication. Rather, several pressing concerns are by-products of consolidation of a limited number of cases in particular districts, before just a few judges. This trend might change as more cases are filed; nonetheless, especially because early decisions can dictate future legal understandings of how to engage with an emerging technology,⁹⁷ early

[<https://perma.cc/GZ33-C4CZ>] (arguing that “the governance of AI systems is not well pursued through the management of binary interpersonal relations” and advocating a “policy” framework rather than a “rights framework” for AI systems); Julie E. Cohen, *How (Not) To Write a Privacy Law*, at 17 (Knight Inst., Data & Democracy Ser., 2021), <https://s3.amazonaws.com/kfai-documents/documents/306f33954a/3.23.2021-Cohen.pdf> [<https://perma.cc/U4TA-W5K2>] (“[B]ecause enforcement litigation is predominantly atomistic in its identification and valuation of harms, it cannot effectively discipline networked phenomena that produce widely distributed, collective harms manifesting at scale. The mismatch is most obvious for private remedial litigation, which takes individual injury as the proper frame of reference even when claims are aggregated using class-action devices[.]”).

⁹⁵ See, e.g., discussion *supra* Part III.B and sources cited therein.

⁹⁶ See discussion *supra* Part III.B.

⁹⁷ The history of Section 230 of the federal Communications Decency Act of 1996 is illustrative. The statute, which is perhaps the most important single text in shaping the contemporary internet, was written in response to two early state law cases. See *Section 230: Legislative History*, ELECTRONIC FRONTIER FOUND., <https://www.eff.org/issues/cda230/legislative-history>

consolidation is risky. It comes at the expense of allowing individual causes of action to percolate nationwide. It inhibits experimentation and increases the risk that outcomes will be homogenous.⁹⁸ And it simultaneously heightens the stakes of other potential risks of adjudication; for instance, if there is a smaller set of cases on point, and many of those cases share a common judge, then it matters far more if the rule set out in that case is in some way unrepresentative of other disputes.

A potential intervention is thus a counterintuitive one: Within the judicial system, minimizing the risk of bad regulation through adjudication might require *more* adjudication.⁹⁹ For instance, increasing the number and variety of cases might make a given case less likely to dominate the field. Because another issue with generative AI adjudication is the way that well-resourced litigants feature disproportionately in early lawsuits, an especially auspicious response is to think of this issue as an access to justice problem. As William Felstiner, Richard Abel, and Austin Sarat explain, a legal dispute does not spring into existence; rather, disputes are “social constructs,” and they require an individual to perceive an injurious experience, attribute it to another party or entity, and voice it, requesting a remedy.¹⁰⁰ In intellectual property suits, there might be adequate financial incentives and clear enough existing legal hooks for parties to bring claims. But relying on this mechanism alone risks entrenching power dynamics as they stand today, rather than ensuring that legal processes can represent diverse stakeholder

[<https://perma.cc/R9KR-2UNU>] (last visited Dec. 28, 2023) (discussing *Cubby, Inc. v. CompuServe, Inc.* and *Stratton Oakmont, Inc. v. Prodigy Servs. Co.*). After enactment, in *Zeran v. AOL*, a single appellate court construed its key provisions in ways that remain encoded in courts’ understandings of the statute to this day. See Bobby Allyn, *The Story of Section 230 Goes Back to An AOL Troll. Now The Law May Be Undone*, NPR (Feb. 22, 2023, 11:38 AM), <https://www.npr.org/2021/05/11/994395889/how-one-mans-fight-against-an-aol-troll-sealed-the-tech-industrys-power> [<https://perma.cc/ALP6-LQ2Z>] (discussing Ken Zeran’s experience and disposition of case). For further history on Section 230 and a critique of how courts have interpreted it to immunize the actions of “Bad Samaritans,” see DANIELLE KEATS CITRON, *THE FIGHT FOR PRIVACY: PROTECTING DIGNITY, IDENTITY, AND LOVE IN THE DIGITAL AGE* 84–89 (2022). For an argument that Section 230 “obviates, at least in principle, the need for courts to apply the First Amendment to carriage and moderation decisions,” and that this phenomenon has produced “interpretive debt” and “legislative debt,” see Blake E. Reid, *Section 230’s Debts*, FIRST AMENDMENT L. REV. (forthcoming) (manuscript at 8–9), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4624865 [<https://perma.cc/JZ6E-AXF2>]. Thank you to Mark Lemley for suggesting this Section 230 example.

⁹⁸ For instance, in resolving motions to dismiss in two different AI lawsuits before two judges in the same district, each court offered similar statements concerning copyright law’s “substantial similarity” analysis. See Order on Motions to Dismiss and Strike, *Andersen v. Stability AI*, No. 2023-CV-00201 (N.D. Cal. Oct. 30, 2023), ECF 117; Order Granting Motion to Dismiss, *Kadrey v. Meta*, No. 23-CV-03417 (N.D. Cal. Nov. 20, 2023), ECF 56. That is not to say that the judges are necessarily wrong, but rather to highlight how they are interpreting a judge-made doctrine that is not part of the text of the Copyright Act and applying it to generative AI in similar ways.

⁹⁹ Cf. Jonathan Zittrain, *The Copyright Cage*, LEGAL AFFAIRS (Jul.–Aug. 2003), https://www.legalaffairs.org/issues/July-August-2003/feature_zittrain_julaug03.msp [<https://perma.cc/B5DB-ZCUX>] (recounting Woody Allen’s classic pan of a bad restaurant: “The food at this place is really terrible . . . and such small portions”).

¹⁰⁰ Felstiner et al., *supra* note 60, at 631, 633–36 (describing the “naming,” “blaming,” “claiming” process).

interests. The unresolved challenge, moreover, is what to do about social interests that are not so clearly tied to financial interests. For this set of injuries, there might not be a natural “transformation” of an individual’s “perceived injurious experience”¹⁰¹ into a grievance that is resolved through the courts. Whether it is possible to calibrate such “transformation” to ensure that individually felt AI harms can be translated into collective conversations, without overloading the legal system, is a rich subject for future research.

A complementary set of responses to the risks posed by judicial resolution of AI issues involves courts, rather than parties. For instance, courts might consider crafting narrower holdings when they are the first to consider a novel technological question.¹⁰² Such judicial minimalism might decrease the reach of early precedent, diminishing the risk that initial decisions functionally preclude subsequent development. Whether this approach is satisfying in practice, in the sense that it permits remedies that respond to legal as well as social challenges and fosters outcomes that do justice both as between the parties and with respect to society—will depend on the practical availability of viable causes of action that speak to felt AI harms. So, too, will it depend on whether there is robust adjudication that reflects a range of societal perspectives and whether there are other available channels for redress of social harms.

In addition, when a court announces a rule in litigation, it must carefully contemplate whether that rule should apply retroactively to the litigants before it, or prospectively to bind future actors. This consideration will affect not only the outcome for the parties in the case, but also the collective impact on society over time. If the announced rule applies retroactively, then the court might face pressure to craft a rule that fits contemporary social expectations, or to hold in a way that preserves the status quo—lest the parties be held to an unfair standard after the fact, when it is too late to modify their behavior. But perhaps the best future-oriented public policy would upset contemporary social expectations and set a new standard; if so, then fairness might dictate only prospective application of the rule.¹⁰³

There is no obvious answer here, only a principled balancing of tradeoffs. For generative AI, this challenge will likely be sharpest in suits where the plaintiffs seek injunctive relief and systemic change of business practices, as is the case in many

¹⁰¹ Felstiner et al., *supra* note 60, at 633–36.

¹⁰² See *Carpenter v. United States*, 138 S. Ct. 2206, 2220 (2018) (“As Justice Frankfurter noted when considering new innovations in airplanes and radios, the Court must tread carefully in such cases, to ensure that we do not ‘embarrass the future.’” (quoting *Northwest Airlines, Inc. v. Minnesota*, 322 U. S. 292, 300 (1944))).

¹⁰³ This legal conundrum is certainly not unique to AI and has been especially pressing in constitutional adjudication. While the Supreme Court has barred the retroactive application of constitutional holdings in criminal cases, see *Griffith v. Kentucky*, 479 U.S. 314 (1987), it has split on the result in civil cases. For an accessible summary of Supreme Court precedents on retroactivity of civil decisions, see *ArtIII.S1.7.3.3 Retroactivity of Civil Decisions*, CORNELL LAW SCHOOL LEGAL INFO. INST. <https://www.law.cornell.edu/constitution-conan/article-3/section-1/retroactivity-of-civil-decisions> [<https://perma.cc/ZL9J-VEBE>].

of the pending lawsuits.¹⁰⁴ Applying a purely prospective rule may be technologically impossible, however, because changing behavior for the future requires unwinding past behavior: The contested data in the underlying AI model cannot be removed without re-training the defendant's entire generative AI model.¹⁰⁵ And so, at least unless and until the technological state of the art changes, a court may be stuck between a small-c conservative outcome that maintains the status quo, yet potentially does not address systemic considerations, or a more sweeping decision that produces a new social baseline, yet risks unfairness to the private defendants.

These sorts of challenges underscore why AI adjudication cannot be the end of the governance conversation: It must be part of a broader, bi-directional dialogue with affected stakeholders and with other public regulators. That's even more important considering the reality that legislatures are also not platonic ideals of governance; to the contrary, they have their own political and procedural pathologies. Bottom-up regulation through courts and top-down public regulation must be complementary and interactive, seeking to position each modality to correct for potential deficiencies of the other.

V. CONCLUSION

Given that adjudication that relies on existing causes of action is already regulating generative AI, courts will continue to govern our relationship to the technology in the present and future. Whether that's a good or a bad thing, in normative terms, may turn on one's faith in the common law system to respond to emerging social and technological developments. There is certainly evidence that courts have adapted in the past, with adjudication affording opportunities for incremental experimentation across substantive domains; bringing greater salience to precise kinds of harm; and producing information that is instrumentally necessary to craft better-targeted regulatory interventions. Even for skeptics of the common law, however, this Essay contends that failing to recognize the regulatory function of generative AI adjudication is a mistake. It means missing an early window to think strategically about the best ways to develop public regulation to complement what AI adjudication does more or less well. It also misses an important opportunity to expand AI governance toolkits by thinking strategically about non-AI points of intervention.

An approach that is attuned to the reality of AI adjudication on the ground is not perfect. Still, recognizing that these cases can make both "good" and "bad" law, all at once, opens up fresh avenues for research and reform. As one example, if

¹⁰⁴ See discussion *infra* text accompanying notes 50 to 58.

¹⁰⁵ As of this writing, there is not a viable technological path for widespread and economical "machine unlearning," in the sense of deleting a given data point as well as "[f]ully erasing the influence of the data requested to be deleted," including "erasing the influence of that data on other artifacts such as trained machine learning models." Pedregosa & Triantafillou, *supra* note 26. See also Tiffany C. Li, *Algorithmic Destruction*, 75 SMU L. REV. 479, 482 (2022) (coining the term the "algorithmic shadow" to refer to the lingering influence of data on the trained machine learning model).

transparency is seen as necessary, if not sufficient, to promote AI accountability, then one option is to push for AI-specific regulation that requires disclosure or auditing of AI firms. However, given the technical, institutional, social, and political challenges inherent in passing new laws or enabling new causes of action,¹⁰⁶ a pragmatic complementary path is to leverage the adjudicatory process as it is to both permit private redress and achieve collective objectives that are instrumental to good AI governance. For instance, if there are concerns about what AI information is (not) accessible, then a partial answer is to ensure that incentives to settle or prevent disclosure during litigation do not impede the benefits of access to this information. If there are concerns about who can bring AI grievances to light, then a partial answer is promoting access to justice for other kinds of litigants. If there are concerns about where AI cases come to light, given the trend of consolidation in a single jurisdiction, before a single judge, then a partial answer is to ensure that regulators provide other opportunities to deliberate about the kinds of questions being decided in these cases. This might not be the AI governance we think we want. But it might be just what we need.¹⁰⁷

¹⁰⁶ For analysis of technical and institutional challenges in four common types of AI regulatory proposals, see Daniel E. Ho et al., *AI Regulation Has Its Own Alignment Problem: The Technical and Institutional Feasibility of Disclosure, Registration, Licensing, and Auditing*, GEO. WASH. L. REV. (forthcoming 2024), https://dho.stanford.edu/wp-content/uploads/AI_Regulation.pdf [<https://perma.cc/E45T-HG88>].

¹⁰⁷ See Mick Jagger & Keith Richards, *You Can't Always Get What You Want*, in LET IT BLEED (1969).