



WORKING PAPER SERIES

The Terms and Conditions of Generative AI

John M. Newman
Andres Sawicki

[View Report Online](#)

August 2026

©2026 by John M. Newman and Andres Sawicki. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full © credit, including notice, is given to the source.

www.equitablegrowth.org

601 13th Street NW, 12th Floor
Washington, D.C. 20005

(202) 545-6002



THE TERMS AND CONDITIONS OF GENERATIVE AI

John M. Newman* & Andres Sawicki†

Generative AI raises pressing questions for myriad doctrinal areas. Policymakers and scholars have responded rapidly, devoting substantial attention to these difficult legal and policy questions. But, as a practical matter, firms offering generative AI tools can unilaterally resolve or sidestep many issues using online terms and conditions. These terms, though seldom read or understood by users, can radically alter rights, shift legal risks and costs, and eliminate entire causes of action. Yet despite their fundamental implications for law and policy, little is known about the content of generative AI terms and conditions.

This Article offers the first empirical examination, at scale, of the digital strings attached to widely available generative AI tools. Our study reveals some surprisingly user-friendly findings, as well as novel variations on familiar drafting tools like mandatory-arbitration provisions. On the whole, however, the results are troubling. Broad, royalty-free generative AI licenses threaten to erode authors' ability to exercise copyrights in their original works. Contrary to public messaging from high-profile providers, their terms and conditions do not indemnify ordinary users against infringement risk. More broadly, an array of dispute-resolution terms limit access to adjudication, eliminate long-established causes of action, tilt adjudication in providers' favor, and sharply limit monetary remedies for injured users.

These terms and conditions appear likely to have negative individual-, class-, and system-level impacts. Though the study's primary contribution is empirical, the findings support immediate policy responses. We thus conclude by identifying a range of prescriptive reforms—voiding certain one-sided provisions and, under certain circumstances, imposing penalties to disincentivize their use—to ameliorate the harms currently facilitated by the terms and conditions of generative AI.

* Herff Chair of Excellence, University of Memphis School of Law.

† Professor of Law, University of Miami School of Law; Director, Data Ethics Law + Society, Frost Institute for Data Science & Computing.

The authors would like to thank Lea Bishop, Andrea Boyack, Chris Buccafusco, Andrew Dawson, Sarah Fackrell, David Hoffman, Stacey Lantagne, Jake Linford, Sarah Wasserman Rajec, Guy Rub, Victoria Schwartz, Cathay Smith, Alicia Solow-Neiderman, and Christopher Yoo. The Article benefited immensely from feedback gathered at the M3 Workshop at Chicago-Kent College of Law, the University of Miami IDSC Day, the Works-in-Progress IP Colloquium at Boston University School of Law, the Competition Law & Academics Network Academic Workshop at the University of Melbourne Law School, and the Antitrust Enforcement Symposium at Seoul National University School of Law. Francesca Cappelletti, Jeffrey A. Miller, William Fitzgerald, and Carolina Sanchez provided excellent research assistance. We are grateful to Jayaram Law for sharing essential data with us. We would also like to thank the Frost Institute for Data Science & Computing and the Washington Center for Equitable Growth for grant funding to support this research.

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	SURVEY MOTIVATION AND METHODOLOGY	5
A.	Sample	6
B.	Terms Selected for Analysis	7
III.	EMPIRICAL RESULTS	9
A.	Copyright Terms	10
B.	Dispute-Related Terms	21
IV.	IMPACT ANALYSIS	30
A.	The Effects of Copyright Terms	31
B.	Individual- and Class-Level Harms from Dispute-Related Terms	38
C.	System-Level Harms	40
V.	RESPONSIVE POLICY REFORMS	42
A.	The Limited Potential of Enhanced “Consent” Regimes	42
B.	Voiding Problematic Terms	43
C.	Penalizing the Cumulative Use of Unenforceable Terms	45
VI.	CONCLUSION	47

I. INTRODUCTION

Generative artificial intelligence affects a large, rapidly growing array of social and market interactions.¹ But it comes with strings attached: like many other digital tools and services,² generative AI firms typically require users to “agree” to lengthy terms and conditions before accessing the desired product. Very few users read these terms.³ Fewer still understand their legal implications.⁴

Nevertheless, these digital strings—the terms and conditions of generative AI—carry far-reaching implications for U.S. law and policy.⁵ Scholars and policymakers are grappling with generative AI’s implications for intellectual property, privacy, antitrust, free speech, and many other doctrinal areas.⁶ But generative AI providers can use terms and conditions to opt out of existing legal requirements, acquire valuable property rights, pick and choose among various bodies of law, seize home-field advantage for dispute resolution, and more, all while arrogating to themselves the power to change any terms unilaterally at any time. The upshot: while debates persist over how best to govern this powerful technology, those debates may be resolved unilaterally by self-interested suppliers in fine print.⁷

As an example, consider copyright law.⁸ There are fierce ongoing controversies over who ought to own the outputs of generative AI systems and who ought to be liable for infringement associated with the use of generative AI.⁹ But generative AI providers already have the ability to decide for themselves how these fundamental

¹ See, e.g., Matthew Sag, *Fairness and Fair Use in Generative AI*, 92 FORDHAM L. REV. 1887, 1888–1890 (2024).

² Andrea J. Boyack, *Abuse of Contract: Boilerplate Erasure of Consumer Counterparty Rights*, 110 IOWA L.J. 497, 500–01 (2025).

³ E.g., Yannis Bakos et al., *Does Anyone Read the Fine Print? Consumer Attention to Standard-Form Contracts*, 43 J. LEGAL STUD. 1, 3 (2014) (finding that just 0.2% of internet users in large sample clicked through to hyperlinked T&Cs).

⁴ WENDY WAGNER, INCOMPREHENSIBLE! A STUDY OF HOW OUR LEGAL SYSTEM ENCOURAGES INCOMPREHENSIBILITY, WHY IT MATTERS, AND WHAT WE CAN DO ABOUT IT 7–8 (2019); Melvin Aron Eisenberg, *The Limits of Cognition and the Limits of Contract*, 47 STAN. L. REV. 322, 354 (1995).

⁵ See, e.g., Alicia Solow-Niederman, *AI & Doctrinal Collapse*, 78 STAN. L. REV. __ (forthcoming 2026); Woodrow Hartzog & Jessica Silbey, *How AI Destroys Institutions*, 77 U.C. L.J. __ (forthcoming 2026); Daryl Lim & Peter K. Yu, *The Antitrust–Copyright Interface in the Age of Generative Artificial Intelligence*, 74 EMORY L.J. 847 (2025); Ayelet Gordon-Tapiero & Yotam Kaplan, *Generative AI Training as Unjust Enrichment*, 86 OHIO ST. L.J. 285 (2025).

⁶ See sources cited *supra* note 5; *infra* note 8.

⁷ See *infra* Part III (reporting the results of our empirical study).

⁸ See, e.g., Sag, *supra* note 1; Edward Lee, *Fair Use and the Origin of AI Training*, 63 HOUS. L. REV. 105, 112 (2025) (“[T]he stakes in the copyright lawsuits against AI companies in the United States could not be higher.”).

⁹ For a sampling of recent contributions to a large literature, see Katherine Lee, A. Feder Cooper, & James Grimmelman, *Talkin’ ‘Bout AI Generation: Copyright and the Generative AI Supply Chain*, 72 J. COPYRIGHT SOC’Y 251, 325 (2025); Mark Lemley, *How Generative AI Turns Copyright Upside Down*, 25 COLUM. SCI. & TECH. L. REV. 21 (2024); Oren Bracha, *The Work of Copyright in an Age of Machine Production*, 38 HARV. J.L. & TECH. 171, 181–82 (2024); Edward Lee, *Prompting Progress: Authorship in the Age of AI*, 76 FLA. L. REV. 1445 (2024); Ryan Abbott & Elizabeth Rothman, *Disrupting Creativity: Copyright Law in the Age of Generative AI*, 75 FLA. L. REV. 1141, 1144 (2023); Matthew Sag, *Copyright Safety for Generative AI*, 61 HOUS. L. REV. 295 (2023); Mala Chatterjee & Jeanne C. Fromer, *Minds, Machines, and the Law: The Case of Volition in Copyright Law*, 119 COLUM. L. REV. 1887 (2019).

questions will be answered.¹⁰ Because copyrights are alienable¹¹ and U.S. courts treat terms and conditions as contracts,¹² generative AI providers can use terms and conditions to allocate copyright interests and infringement risks as between themselves and their users.¹³

Nor is copyright the only doctrinal field implicated by generative AI terms and conditions. A great many generally applicable conduct-regulating laws—tort, contract, antitrust, consumer protection, privacy, state statutes banning nonconsensual pornography,¹⁴ and more—depend on private enforcement for operationalization and effect.¹⁵ Even hybrid public-private enforcement regimes often rely far more on private plaintiffs than resource-constrained public enforcers.¹⁶ While scholars and stakeholders are busily debating how each of these fields should approach generative AI, providers might already be using terms and conditions to resolve those debates in their own favor.

Despite all this, relatively little is known about the actual prevalence or content of generative AI terms and conditions.¹⁷ This Article describes the results of the first systematic, empirical analysis of the generative AI terms and conditions that are already being deployed in a variety of market settings. Our study reveals some surprisingly user-friendly practices, as well as novel variations on familiar drafting tools like mandatory-arbitration provisions. On the whole, however, the results are troubling. In the near term, generative AI terms and conditions threaten

¹⁰ Nancy S. Kim, *AI and the Fine Print Disruption of Copyright*, 129 PENN. ST. L. REV. 577, 593-94 (2025).

¹¹ See *infra* Parts II.B.1–.2.

¹² See Mark A. Lemley, *Terms of Use*, 91 MINN. L. REV. 459, 459 (2006) (“Every court to consider the issue has found ‘clickwrap’ licenses, in which an online user clicks ‘I agree’ to standard form terms, enforceable.”).

¹³ See *infra* Part II.B.3.

¹⁴ See e.g., Gregory Smith et al., *Liability for Harms from AI Systems: The Application of U.S. Tort Law and Liability to Harms from Artificial Intelligence Systems*, RAND (2024), https://www.rand.org/pubs/research_reports/RRA3243-4.html (tort); Blake Brittain, *Lawsuits Accuse AI Content Creators of Misusing Copyrighted Work*, REUTERS (Jan. 17, 2023), <https://www.reuters.com/legal/transactional/lawsuits-accuse-ai-content-creators-misusing-copyrighted-work-2023-01-17> (describing lawsuits by web publishers against generative AI providers) (contract); *United States v. Google LLC*, No. 20-cv-3010 (APM, 2025 WL 2523010 (D.D.C. Sept. 2, 2025); (antitrust remedy opinion); *Complaint, Chegg, Inc. v. Google LLC*, No. 1:25-cv-00543 (D.D.C. Feb. 24, 2025) (antitrust substantive opinion); Press Release, FTC, *FTC Launches Inquiry into AI Chatbots Acting as Companions* (Sept. 11, 2025), <https://www.ftc.gov/news-events/news/press-releases/2025/09/ftc-launches-inquiry-ai-chatbots-acting-companions> (consumer protection); *Calhoun v. Google*, 4:20-CV-05146-YGR, Nos. 1058, 1060, 1077 & 1081 (N.D. Cal. June 9, 2025) (denying motion for class certification in private-plaintiff privacy lawsuit) (privacy); *The TAKE IT DOWN Act: A Federal Law Prohibiting the Nonconsensual Publication of Intimate Images*, CONG. RES. SERV. (May 20, 2025), <https://www.congress.gov/crs-product/LSB11314> (nonconsensual pornography).

¹⁵ On the difference between conduct-regulating and loss-allocation rules, see, e.g., WILLIAM LANDES & RICHARD POSNER, *THE ECONOMIC STRUCTURE OF TORT LAW* 64–66 (1987).

¹⁶ For example, Congress authorized both federal and private enforcers to bring antitrust lawsuits. 15 U.S.C. §§ 9, 15. In a typical year, plaintiffs file hundreds of private actions, while federal enforcement agencies bring at most a few dozen cases. See, e.g., Daniel A. Crane, *Toward a Realistic Comparative Assessment of Private Antitrust Enforcement*, in *RECONCILING EFFICIENCY AND EQUITY: A GLOBAL CHALLENGE FOR COMPETITION POLICY* 341, 343 (Damien Gerard & Ioannis Lianos eds., 2019).

¹⁷ For a survey of extant literature, see *infra* nn. 54-56 and accompanying text.

ongoing harm to users on an individual- and class-level basis. Over the longer run, these terms are likely to hinder judicial and legislative doctrinal development.¹⁸

Although not unique to this setting, these threats are especially acute in the context of a nascent, still-developing technology that raises particularly novel, far-reaching legal questions. Both the path of innovation and the path of the law are still being charted for generative AI.¹⁹ Unless policymakers intervene, self-interested technology firms will play an outsized role in determining their directions. And once path-dependence sets in, course correction will be unlikely.

We begin, in Part II, by articulating the motivation for our study and our methodological approach. Part II.A describes our sample of generative AI products; Part II.B then provides the rationale for selecting the terms we analyzed. These can be helpfully divided into three general groups: terms that affect copyrights and infringement liability risks, terms that affect dispute resolution, and terms that permit ex post unilateral modification of contractual provisions.

Part III reports our findings. As an initial matter, the use of generative AI terms and conditions is ubiquitous. Providers attached online terms and conditions to each of the 86 generative AI tools in our selected sample.

As to content, the generative AI terms and conditions in our study include an array of one-sided copyright-related clauses. The overwhelming majority grant providers nonexclusive licenses to any material a user submits to the generative AI system.²⁰ And they often reach broadly: in many instances, providers can redeploy user-submitted material to virtually any purpose, even where the user owned their copyright before using the generative AI system at all. Similar terms apply to system outputs, and they are similarly pervasive.²¹ Generative AI terms and conditions also frequently include representations and warranties by which users purportedly promise that they have the rights required to submit copyrighted material and will not infringe third-party copyrights when using these systems.²² If users break those promises, even unwittingly, indemnification provisions that appear in virtually all of these terms and conditions enable the generative AI firm to demand compensation from the user.²³

The dispute-related generative AI terms and conditions in our sample diverge somewhat from those in other industries. Mandatory arbitration and class waivers, for example, are appreciably less common in these generative AI terms and

¹⁸ These effects occur against a broader U.S. backdrop of limited access to justice. *See, e.g.*, Drew Simshaw, *Interoperable Legal AI for Access to justice*, 134 YALE L.J. F. 795, 795–97 (2025).

¹⁹ *E.g.*, Joshua S. Gans, *Regulating the Direction of Innovation*, Nat'l Bureau of Econ. Rsch. Working Paper 32741 (Jan. 2025).

²⁰ *Id.* at Part III.A.1.

²¹ *Id.* at Part III.A.2.

²² *Id.* at Part III.A.3.

²³ *Id.*

conditions than in other digital contexts.²⁴ And many of the arbitration clauses in generative AI terms and conditions include carveouts for IP-related disputes—a novel approach not often seen in other market settings.²⁵ But on the whole, these generative AI terms and conditions are—like the copyright-related provisions described above—restrictive and one-sided. A substantial majority require dispute resolution on the provider’s home turf, void negligence (and sometimes all) claims against providers, void warranties and extrinsic promises by providers, impose stringent caps on damages awards, and eliminate entire categories of damages.²⁶

Moreover, the overwhelming majority of generative AI terms in our sample permit providers—but not users—to change terms at any time. Many do so while expressly disclaiming any duty to provide meaningful consent. Consequently, even if users were to read and understand their IP and dispute-related rights and duties before “agreeing” to them, those terms are subject to change at any time.

Part IV turns to normative analysis. Generative AI terms and conditions threaten to undermine authors’ statutory rights in their works. Given the stated ambitions of generative AI firms and the rapid diffusion of the technology across the creative economy,²⁷ these licenses pose a significant threat to authors’ ability to meaningfully enjoy their copyright interests.

More troublingly, the one-way indemnification provisions in generative AI terms and conditions turn users into copyright crumple zones, exposing them to potentially ruinous liability from the clash between generative AI and the content industries.²⁸ Fundamental questions remain over ordinary uses of generative AI tools that produce text, images, songs, or videos that are substantially similar to copyrighted works.²⁹ But whatever courts eventually decide, the indemnification provisions in generative AI terms and conditions mean that, as a practical matter,

²⁴ Compare *infra* Part III.B.1 (finding that 51% of surveyed T&Cs include arbitration clauses and 47% include class waivers), with Boyack, *supra* note 2, at 517 (reporting that 66% of all online T&Cs studied included arbitration clauses and 65% include class waivers).

²⁵ See Boyack, *supra* note 2 (surveying T&Cs across a wide variety of industry sectors, not reporting any IP carveouts from arbitration clauses).

²⁶ See *infra* Part III.B (reporting findings on dispute-related terms).

²⁷ See Sam Altman, *The Gentle Singularity*, <https://blog.samaltman.com/the-gentle-singularity> (June 10, 2025) (claiming that “we have recently built systems that are smarter than people in many ways”); Dario Amodei, *The Adolescence of Technology*, <https://www.darioamodei.com/essay/the-adolescence-of-technology> (Jan. 2026) (asserting that “it cannot possibly be more than a few years before AI is better than humans at essentially everything”).

²⁸ We borrow the “crumple zone” framing from Madeline Clare Elish’s influential article, *Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction*, 5 ENGAGING SCI., TECH., & SOC’Y 40 (2019).

²⁹ See Peter Henderson et al., *Foundation Models and Fair Use*, 24 J. MACHINE LEARNING 1, 22 (2023) (reasoning that any “non-transformative material” that a model provides to a user would “undergo the same assessment as in any other case”); Michael P. Goodyear, *Artificial Infringement*, 77 U.C. L.J. __ (forthcoming 2026) (proposing a novel causation-based framework for analyzing infringement by generative AI system outputs); Lee, Cooper, & Grimmelmann, *supra* note __ at 350-55 (describing three paradigm scenarios); Chatterjee & Fromer, *supra* note __ (developing a volition-based framework for answering these questions).

users could bear the financial risks of infringement.³⁰ And copyright law's relatively harsh remedial provisions make that risk considerable.³¹

As to dispute-related terms, our study revealed a few encouraging findings, but—especially when the cumulative effects of various provisions—are taken into account, generative AI terms and conditions appear to seriously constrain users' ability and incentive to receive effective redress for legal wrongs. That effect, in turn, threatens to erode the vitality of a wide range of common-law doctrines and democratically enacted statutory regimes, distort the development of law, and stunt its ability to meet the challenges raised by generative AI.

Part V concludes by identifying a range of potential policy responses. Given the growing consensus that notice-and-consent regimes tend to underperform, we recommend, at minimum, voiding certain one-sided terms. Even then, prior research finds that voiding terms does not necessarily eliminate their real-world, *in terrorem* effects, leaving in place an incentive for firms to continue using them. Given that, we recommend policymakers consider affirmatively penalizing the (over)use of one-sided generative AI terms and conditions. Particularly when firms deploy multiple types of these terms in combination, their negative effects merit a robust intervention. We conclude by identifying multiple avenues—legislative and administrative, federal and state—for potential reform efforts.

II. SURVEY MOTIVATION AND METHODOLOGY

Internet-based firms frequently use lengthy terms and conditions ("T&Cs") to govern their relationships with their users.³² These T&Cs can be quite powerful; critics have worried that, if enforced, they could produce "mass governance by information that we do not see and cannot manage."³³ Those concerns have been buttressed by recent empirical research demonstrating that problematic terms are widespread in online T&Cs.³⁴ Still, the most wide-ranging of these studies have not reached generative AI firms,³⁵ whose products and business models raise a

³⁰ We explain why generative AI providers might enforce these provisions, notwithstanding the business considerations that might counsel otherwise, *infra* text accompanying notes __.

³¹ See *infra* text accompanying notes __ (describing copyright's statutory damages provisions).

³² These are sometimes alternatively known as "terms of service" or "terms of use." We will alternately refer to them as "terms and conditions" or "T&Cs."

³³ Margaret Jane Radin, *Access to Justice and Abuses of Contract*, 33 WINDSOR YEARBOOK ACCESS TO JUST. 177, 177 (2016).

³⁴ Empirical studies tended to focus on a particular type of term, rather than T&Cs more generally. See, e.g., David Horton & Andrew Cann Chandrasekher, *After the Revolution: An Empirical Study of Consumer Arbitration*, 104 GEO. L.J. 57 (2015); see also Shmuel I. Becher & Uri Benoliel, *Sneak in Contracts*, 55 GA. L. REV. 656 (2021) (surveying use of unilateral-modification terms).

³⁵ Boyack, *supra* note 2 (surveying 100 firms across several sectors that do not include generative AI); Shmuel I. Becher & Uri Benoliel, *Sneak In Contracts*, 55 GA. L. REV. 657 (2021); Florencia Marotta-Wurgler & Robert Taylor, *Set in Stone? Change and Innovation in Consumer Standard-Form Contracts*, 88 N.Y.U. L. REV. 240 (2013) (examining a sample of 500 online contracts but focusing only on the presence of unilateral-modification clauses); studying changes in online T&Cs from 2003 to 2010, before the rise of generative AI products) John M. Newman, *Innovation*

plethora of important—and novel—legal questions.³⁶ This gap in the literature provided the motivation for our study: to examine, at scale, the presence and content of generative AI firms’ T&Cs. In this Part, we describe our sample, the terms we selected for analysis, and the process we used for analyzing them.

A. Sample

The generative AI landscape is rapidly changing. Although foundation models offered by OpenAI, Anthropic, and Google tend to dominate public attention, hundreds of products vie for users across a wide range of market niches. We sought to assemble a sample including an array of market participants of varying sizes.

To do so, we began constructing our sample in late 2024 with two lists of generative AI products. The first, compiled and made publicly available by the Silicon Valley venture-capital firm a16z, identified the top fifty generative AI web products by monthly unique visits.³⁷ Because some important GenAI products may have relatively few monthly visits, we supplemented the a16z list with another one compiled by Jayaram Law, a technology-oriented law firm; the Jayaram list identified generative AI products that its clients were most interested in tracking.³⁸ Accounting for overlap among the two lists, this process yielded a total sample of 101 generative AI products.³⁹

Our research assistants (“RAs”) then began collecting T&Cs associated with each of these products in January 2025. By that point, one of the products in our initial list (Inferkit) had shut down, leaving a total of 100 products in our sample. Our RAs downloaded local PDF copies of each set of T&Cs (including applicable privacy policies, tool-specific T&Cs, and any other document that was incorporated by reference into the T&Cs). Where different T&Cs were deployed for zero-price and paid product versions, we asked our RAs to download local copies of both sets of T&Cs. The collection process was completed in May 2025.

Policy for Cloud-Computing Contracts, in RESEARCH HANDBOOK ON DIGITAL TRANSFORMATIONS 380 (F. Xavier Olleros & Majlinda Zhegu eds., 2016) (focusing only on choice-of-law clauses in cloud-services provider T&Cs).

³⁶ A few contributions have explored generative AI T&Cs, but these cover only a few suppliers and focus on a relatively limited set of potentially concerning terms. The most relevant prior contributions include Peter Henderson & Mark Lemley, *The Mirage of Artificial Intelligence Terms of Use Restrictions*, 100 IND. L.J. 1 (2025) and Nancy Kim, *supra* note __. Henderson and Lemley describe how a few of the best-known generative AI firms purport to restrict certain uses of their tools. Henderson & Lemley, 100 IND. L.J. 1. Their analysis explains why those terms will not be effective at their stated purpose of privately regulating use of powerful generative AI models; they do not explore the impact of these terms on users’ own copyright interests and liability risks, nor on the dispute-related terms we evaluate. Kim addresses some of the copyright terms we evaluate; like Henderson and Lemley, however, she focuses on a relatively small number of generative AI firms. Kim, *supra* note __.

³⁷ See The Top 100 Gen AI Consumer Apps, 3d Edition, available at <https://a16z.com/100-gen-ai-apps-3/> (August 21, 2024). This a16z list separately identified the top 50 web products and the top 50 mobile apps. We focused on the web environment because mobile apps may present different issues regarding the enforceability of T&Cs.

³⁸ This list is on file with the authors.

³⁹ Some products, like ChatGPT, appeared on all three lists.

We then reviewed that collection to ensure that our sample complied with two inclusion criteria: (1) the GenAI firm appeared to rely primarily on its publicly available T&C to govern its relationship with its customers; and (2) a broadly similar set of copyright ownership and liability issues were likely implicated by the nature of the GenAI product. We excluded eight GenAI products for failure to meet the first criterion.⁴⁰ We excluded another six for failure to meet the second criterion.⁴¹ This left us with a total of 86 products in our sample.

B. Terms Selected for Analysis

To identify provisions within the T&Cs for our analysis, we consulted prior theoretical and empirical work, reviewed a randomly selected subset of the T&Cs in our sample, and received helpful feedback and suggestions from subject-matter experts.⁴² This yielded the T&C provisions identified and discussed below.⁴³

1. Copyright Terms

We began our study with copyright-related provisions. Although generative AI has profound implications for information law and policy,⁴⁴ copyright law has quickly emerged as an especially hotly contested battleground. High-profile infringement lawsuits against generative AI firms have attracted considerable media attention,⁴⁵ the U.S. Copyright Office has issued controversial rulings and reports on generative AI,⁴⁶ and Congress has held contentious hearings on the topic.⁴⁷

⁴⁰ Deeper Insights and Shortly AI did not have a publicly available T&C governing customer relationships. Four enterprise GenAI offerings—Data Robot, IBM Watson, SAP Leonardo, and You.com—typically rely heavily on private documents alongside any publicly available T&Cs. Google Imagen was not available as a standalone product (Google Gemini, which could use Imagen to generate images in response to user prompts, remained in our sample). Finally, Doubao did not have an English-language T&C available.

⁴¹ Five of these were GenAI cloud infrastructure products: Amazon AWS AI, Google Cloud AI, Microsoft Azure AI, OctoAI, and Oracle AI. The sixth was Kite, which develops blockchain infrastructure for AI agents.

⁴² See Uri Y. Hachohen, Amit Elazari, & Talia Schwartz-Maor, *A Penny for their Creations—Apprising Users' Value of Copyrights in their Social Media Content*, 36 BERKELEY TECH. L.J. 511, 515 (2021); Peter Henderson et al., *supra* note __; Henderson & Lemley, *supra* note __; Kim, *supra* note __; Lee, Cooper, & Grimmelmman, *supra* note __.

⁴³ See *infra* Parts I.B & I.C.

⁴⁴ See, e.g., Woodrow Hartzog & Jessica Silbey, *How AI Destroys Institutions*, 77 U.C. L.J. __ (forthcoming 2026); Alicia Solow-Niederman, *AI & Doctrinal Collapse*, 78 STAN. L. REV. __ (forthcoming 2026); Ayelet Gordon-Tapiero & Yotam Kaplan, *Generative AI Training as Unjust Enrichment*, 86 OHIO ST. L.J. 285 (2025); Daryl Lim & Peter K. Yu, *The Antitrust—Copyright Interface in the Age of Generative Artificial Intelligence*, 74 EMORY L.J. 847 (2025).

⁴⁵ In fact, some of these lawsuits have directly involved media corporations. See, e.g., Michael M. Grynbaum & Ryan Mac, *The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work*, N.Y. TIMES (Dec. 27, 2023), <https://www.nytimes.com/2023/12/27/business/media/new-york-times-open-ai-microsoft-lawsuit.html>.

⁴⁶ See U.S. COPYRIGHT OFF., *Copyright and Artificial Intelligence*, COPYRIGHT.GOV, <https://www.copyright.gov/ai/> (last visited Feb. 23, 2026).

⁴⁷ S. Hrg. 119-202, *Too Big to Prosecute?: Examining the AI Industry's Mass Ingestion of Copyrighted Works for AI Training*, July 16, 2025.

Both copyright ownership and liability can be impacted by GenAI T&Cs. Many GenAI products allow users to submit copyrighted content, including photographs, poems, and songs, to the system. Copyright interests might also vest in the output of a GenAI product. And because copyright interests are alienable, T&Cs might (purport to) allocate these interests in inputs and outputs as between GenAI firms and their users.

We accordingly sought to identify provisions that purported to do the following: (1) limit what users could submit to the GenAI system; (2) convey to GenAI firms all or some of the copyright interests that users hold in any submitted material; (3) dictate what GenAI firms could do with submitted material; (4) allocate copyright interests in GenAI system outputs; and (5) dictate what users or GenAI firms could do with system outputs.

Meanwhile, because GenAI products might copy or distribute copyrighted works as part of their ordinary functioning, GenAI firms might use their T&Cs to shift infringement risks to their users (or to shield users from those risks). They could do so either by including explicit promises regarding relevant information or by directly shifting responsibility for copyright-related costs and damages.

We accordingly sought to identify the following provisions: (1) representations and warranties made by users on copyright-related issues; (2) indemnifications of GenAI firms by users for copyright-related costs and damages; (3) representations and warranties made by GenAI firms on copyright-related issues; and (4) indemnifications of users by GenAI firms for copyright-related costs and damages.

2. Dispute-Related Terms

Because generative AI technology offers many varied end uses, it implicates a wide range of general conduct-regulating laws. Disputes between users and generative AI providers can involve theories of harm based on tort, contract, privacy, consumer protection, antitrust, and more.⁴⁸ As with copyright, however, the actual impact of these various bodies of law will ultimately depend heavily on the content of any applicable T&Cs—and, in particular, any dispute-related terms.

Over time, sophisticated drafters have developed a toolkit of contractual terms that purport to affect the resolution of any disputes that subsequently arise between the parties.⁴⁹ Drawing from existing theoretical and empirical work on dispute-related terms,⁵⁰ we identified sixteen types of potentially problematic

⁴⁸ See *supra* notes 16–21 and accompanying text (collecting real-world examples).

⁴⁹ For a few seminal examples of scholarship describing and analyzing these terms, see MARGARET JANE RADIN, *BOILERPLATE: THE FINE PRINT, VANISHING RIGHTS, AND THE RULE OF LAW* (2013); Boyack, *supra* note 2; Mark Lemley, *Terms of Use*, 91 MINN. L. REV. 459 (2006); Todd D. Rakoff, *Contracts of Adhesion: An Essay in Reconstruction*, 96 HARV. L. REV. 1173 (1983).

⁵⁰ See Boyack, *supra* note 2, at 508–16; Marco Loos & Joasia Luzak, *Wanted: A Bigger Stick. On Unfair Terms in Consumer Contracts with Online Service Providers*, 39 J. CONSUMER POL'Y 63 (2016); (identifying distant-venue

terms for inclusion in our study. These can be helpfully organized into four categories.

First, T&Cs can limit an affected party's ability to pursue formal dispute resolution. Within this category, we sought to identify (1) mandatory-arbitration provisions; (2) class-action waivers; (3) contractual limitations periods; (4) forum-selection clauses; and (5) jurisdiction waivers.

Second, T&Cs can explicitly or effectively eliminate some causes of action. Within this category, we sought to identify terms that purport to (1) void negligence claims; (2) exclude warranties; and (3) void prior promises and statements.

Third, several commonly used terms can shift the balance of power in favor of the drafting party. Within this category, we sought to identify (1) unilateral modification provisions; (2) confidential-proceeding requirements; (3) choice-of-law provisions; (4) choice-of procedure provisions; and (5) cost-shifting provisions.

And *fourth*, T&Cs can reduce the amount of money that a prevailing party can expect to recover. Within this category, we sought to identify (1) liquidated damages provisions; (2) terms that eliminate specified categories of damages (e.g., incidental, consequential, or punitive damages); and (3) jury trial waivers.

* * *

Between June 2025 and February 2026, research assistants conducted an initial analysis of the contents of each T&C in the downloaded repository. Each T&C was reviewed by at least two RAs. During their analysis, RAs hand-coded all T&Cs according to instructions developed by the authors.⁵¹ The authors then reviewed the coding provided by the RAs, resolved any differences in the coding, and independently conducted supplemental reviews of T&Cs.

III. RESULTS

The following discussion describes and visually illustrates our empirical findings. We first discuss the nature and prevalence of T&C terms that purport to allocate copyright interests and liability. We then evaluate the representations, warranties, and indemnification provisions. Next, we turn to our dispute-related terms. Given the potentially cumulative and multiplicative effects of one-sided

provisions, choice-of-foreign-law terms, liability-limiting terms, and unilateral-termination or -modification terms as potentially unfair under applicable EU regulations); *see also* Marco Lippi et al., 27 *ARTIFICIAL INTELLIGENCE & L.* 117, 119 (2019) (additionally identifying mandatory-arbitration provisions as potentially unfair under applicable EU regulations).

⁵¹ These instructions were designed to guide RAs' assessments of the T&Cs so they could more easily find the provisions of interest. The instructions are available upon request.

terms, we also report how many of the generative AI T&Cs in our sample include multiple types of terms within each of the four categories described above.

A. *Copyright Terms*

1. Rights to User Inputs

Users own the copyright in a substantial amount of content that they submit through generative AI prompt interfaces. To be sure, mundane prompts fall short of even the low bar for copyrightability.⁵² But users also submit a vast array of expressive works—lengthy essays, complex images, and sophisticated music—to generative AI systems, and these will ordinarily be protected by copyright law. Generative AI T&Cs might then attempt to convey the user’s preexisting copyright interests to the provider.

We accordingly explored three key questions: (1) what is the character of the conveyance, if any; (2) which of copyright’s exclusive rights fall within the scope of the conveyance; and (3) what are the purposes for which the generative AI provider can exercise the rights it obtains from the conveyance.

Type of conveyance. In standard configurations, a generative AI system deployed via a hosted service will make copies of the material the user submits.⁵³ The provider would accordingly need authorization from the copyright owner in order to avoid infringement liability. Authorization can take the form of an assignment, an exclusive license, or a nonexclusive license. The most natural form in this setting would be a nonexclusive license: the user would retain her interests in the submitted material while granting the generative AI provider permission to use it in some manner. We therefore expected that generative AI T&Cs would include nonexclusive licenses covering user-submitted material.⁵⁴

Consistent with this expectation, 84.9% of the T&Cs we study purport to grant the generative AI provider a nonexclusive license to material that the user submits to the model. To our surprise, 14% do not grant any explicit authorization to the provider, apparently opening those providers up to at least a possible risk of

⁵² See *Feist*, 499 U.S. at 345 (noting that the level of creativity required to obtain a copyright interest is extremely low, but holding that an alphabetical order of telephone directory entries failed to meet it).

⁵³ See Lee, Cooper, & Grimmelmann, *supra* __, at 325 (explaining that “a generative-AI system’s use of a user-inputted prompt is a ‘reproduction’ within the meaning of copyright law”).

⁵⁴ A “transfer of copyright” must be in writing and signed by the owner. 17 U.S.C. § 203. The statute defines a “transfer of copyright” as including assignments and exclusive licenses, while excluding nonexclusive licenses. 17 U.S.C. § 101. Unlike assignments and exclusive licenses, then, nonexclusive licenses can arise implicitly, and they likely would in this scenario; however, there’s no apparent reason for a generative AI provider to rely on the fallback of an implied license when they have the opportunity and ability to obtain an explicit one. See *infra* text accompanying notes __.

infringement suits from their users.⁵⁵ Only one provider—Chatbot App—might be read as attempting to assign user submissions to the provider.⁵⁶

Of the seventy-three T&Cs that include provisions purporting to grant providers nonexclusive licenses to user-submitted material, we assessed whether the license was worldwide, royalty-free, perpetual, irrevocable, sublicensable, or transferable. The overwhelming majority (69.9%) were worldwide.⁵⁷ The same amount (69.9%) asserted that they were royalty-free; the remainder were silent on this score.⁵⁸ All but five of the providers that included one of those provisions also included the other.⁵⁹ Just under half (45.2%) of the licenses are perpetual. Just over half (52.1%) assert that they are irrevocable; one (Getty Images) explicitly states that is revocable; and one (Veed) states that it is revocable only by paid users. In total, 71.2% explicitly gave the generative AI provider some ability to give others the rights it had under the license, via a transfer or sublicensable or both.

Conveyance Type

N = 86 products

	<i>n</i>	%	
Assignment	1	1.2	
Exclusive license	0	0.0	
Non-exclusive license	73	84.9	
Silent	12	14.0	
Total	86	100.0	

Figure 1: Conveyance of User Inputs to Generative AI Provider

⁵⁵ Those suits would likely fail, for reasons we explain below. See text accompanying notes ____.

⁵⁶ The Chatbot App T&C states that user submissions “will become the sole property of Chatbot App.” We suspect that this language was meant to refer to submissions to the firm’s email address regarding the product itself, rather than to the materials a user might submit to a generative AI system. See Chatbot App T&C at 9 (claiming ownership arises when “you submit or post any ideas, creative suggestions, designs, photographs, information, advertisements, data or proposals, including ideas for new or improved products, services, features, technologies or promotions”). But an entirely separate clause in a different section of the T&Cs already purports to grant ownership in such materials to Chatbot App. See *id.* at 4 (“Any feedback, comments, ideas, improvements or suggestions . . . provided by you to Chatbot App with respect to the app shall remain the sole and exclusive property of Chatbot App”). In light of the broad language of the T&C and its failure to otherwise say anything about copyright interests in user inputs, we think the quoted language could be the basis for an assertion of ownership; still, the language may fall short of what’s required to effect an assignment, although we cannot be certain. See *Thomsen v. Famous Dave’s of Am., Inc.*, 606 F.3d 905, 908 (8th Cir. 2010) (holding that author assigned his copyright interests in a contract that stated he “released . . . all copyright, proprietary design, and sign work”). In any event, none of this changes the main story told by the data, which reveals the prevalence of non-exclusive licenses as the dominant type of conveyance.

⁵⁷ We included here licenses described either as worldwide or global.

⁵⁸ We included here licenses described either as royalty-free or fully paid-up.

⁵⁹ Inworlds and Luminar AI were worldwide but not royalty-free; Canva, Kafkai, and Pixelcut were royalty-free but not worldwide.

License Characteristics

N = 73 non-exclusive licenses to user inputs

	<i>n</i>	%	
Worldwide	51	69.9	
Royalty-free	51	69.9	
Perpetual	33	45.2	
Irrevocable	38	52.1	
Sublicensable	46	63.0	
Transferable	31	42.5	

Figure 2: Characteristics of Licenses in User Inputs

Section 106 rights conveyed. 17 U.S.C. § 106 establishes six exclusive rights. Five have relatively general application: the reproduction, derivative works, distribution, public performance, and public display rights.⁶⁰ The sixth, which applies only to sound recordings, grants the owner exclusivity in public performances via digital audio transmissions.⁶¹

At least some of these rights might be implicated in the ordinary operation of a generative AI system; as noted above, user submissions will typically cause the system to create a copy, including any copyrighted material in it. But the conveyances might reach further to include rights that would not be implicated in the ordinary operation of a generative AI system; the public performance and display rights, for example, wouldn't be implicated by interactions between the system and a single user because they both require multiple third-parties to access the work through the performance or display. We therefore assessed which of the Section 106 rights were included within the scope of the conveyances to see how far the T&Cs reach beyond the rights necessary for ordinary operations.

The reproduction right was the most common Section 106 right explicitly identified in the T&Cs; 83.6% of the T&Cs that included a license referred to reproducing, storing, hosting, or copying user-submitted material. There was only one TOS that did not refer to the reproduction right but that did explicitly refer to other Section 106 rights—SeaArt's T&C explicitly grants it only the right to display and perform user-submitted material. The distribution right is explicitly identified in 75.3% of the T&Cs with a license; the display right appears in 74%; and the performance right in 45.2%. Meanwhile, 71.2% of T&Cs with a license referred

⁶⁰ The precise mix depends on the type of work at issue. All types of works receive the same three core rights: reproduction, derivative works, and distribution rights. 17 U.S.C. § 106(1), (2), & (3). Literary, musical, dramatic, and choreographic works receive the public performance and public display rights. 17 U.S.C. § 106(4) & (5). Pantomimes, which are statutorily grouped in with choreographic works, also receive these rights, but it would mess up the structure of the list to call them out there. *Id.* Motion picture and other audiovisual works receive the public performance right; pictorial, graphic, and sculptural works receive the public display right. *Id.*

⁶¹ *Id.* at § 106(6).

explicitly either to the derivative works right or to acts that would fall within its scope.⁶²

One provider—Luma—distinguishes free from paid users on this dimension. With respect to material submitted by users relying on the free service, Luma’s T&C explicitly conveys a nonexclusive license covering all of the Section 106 rights; with respect to material submitted by paid users, however, it asserts that Luma has only the ability to “use, host, and store” it “solely as reasonably necessary” to fulfill certain (broadly described) purposes.

§106 Rights in User Inputs

N = 73 non-exclusive licenses

	<i>n</i>	%	
Reproduction	61	83.6	
Derivative works	52	71.2	
Distribution	55	75.3	
Performance	33	45.2	
Display	54	74.0	

Figure 3: Section 106 Rights in User Inputs Conveyed to Provider

Purposes for which providers can exercise their rights as grantees. Copyright owners have considerable flexibility in specifying the purposes for which a grantee can exercise whatever rights are being conveyed.⁶³ It seems appropriate for the provider to obtain from a user the authorization it needs in order to provide her with the services that she wants to use. It is another thing entirely for the provider to obtain copyright interests it could deploy to completely unrelated purposes, including to create new expressive works that might compete with the user’s works.

We therefore examined whether generative AI T&Cs circumscribed the purposes for which the provider could exercise the interests it obtained in the user’s inputs. We identified four possible approaches: T&Cs might state that the provider could exercise the copyright interests it obtains from the user (1) only to provide the service to the user; (2) only to provide the service to other users; (3) to conduct any of its business activities; or (4) for any purpose whatsoever.

Of the 73 generative AI T&Cs that conveyed nonexclusive licenses in user-submitted material to the provider, a handful (6.8%) limited themselves to using those interests only to provide the service to that user. A plurality (47.9%) asserted the ability to exercise their interests in order to provide the service to other users.

⁶² Luma’s T&C purported to include these rights in its license from free users but not in its license from paid users.

⁶³ See 17 U.S.C. § 201(d) (allowing owners to transfer their copyrights “in whole or in part,” including by “subdivi[ding] any of the rights specified by section 106”).

Nearly one-third (30.1%) stated that they could use the licensed material for any business purpose. Finally, 13.7% of the nonexclusive licenses were either described as unlimited or did not contain any explicit limitations.

We note one interesting multi-tiered structure regarding the purposes for which the provider could exercise its license. Veed can use input from its free users for any business purpose; its paid users, however, will only have their input used so that Veed can provide the service to that and other users (in other words, Veed cannot use input from its paid users for Veed’s general business purposes).⁶⁴

Permitted Uses of User Inputs

N = 73 non-exclusive licenses

	<i>n</i>	%	
Only to serve this user	5	6.8	
To serve this and other users	35	47.9	
Any business activity	22	30.1	
No limitation	10	13.7	
Varies by tier	1	1.4	

Figure 4: Permitted Uses of User Inputs by Generative AI Provider

2. Rights to System Outputs

The second category of terms we study relate to copyright interests in system outputs. Our analysis of these terms largely tracks our analysis of copyright interests in user inputs.

In general, we considered whether generative AI T&Cs allocate copyright interests in system outputs to users, providers, both, or neither. To be sure, it may turn out that there are no such interests to allocate: the Copyright Office has taken the position that most system outputs would fall into the public domain. Even under the Office’s view, however, copyright interests will arise when “the traditional elements of human authorship in the work” are “conceived and executed” by a human operator.⁶⁵ And it is not clear that the Office’s position will endure.⁶⁶ Given the unsettled state of copyright in generative AI outputs, we concluded that it was worth investigating what generative AI T&Cs had to say on the matter.

When generative AI T&Cs address copyright interests in system outputs, we evaluated: (1) whether the user owns, has an exclusive license, or has a non-

⁶⁴ We counted Leonardo as an “unlimited” license because that is the most expansive category into which it falls, notwithstanding the fact that paid users can opt for a more restrictive approach.

⁶⁵ Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence, 88 Fed. Reg. 16190 (proposed Mar. 16, 2023) (to be codified at 37 C.F.R. pt. 2).

⁶⁶ See Edward Lee, *Prompting Progress: Authorship in the Age of AI*, 76 FLA. L. REV. 1445, 1450 (2024) (explaining “why the Copyright Office’s restrictive approach to AI-generated works is wrong”).

exclusive license in system outputs; (2) whether the generative AI provider owns, has an exclusive license, or has a non-exclusive license in system outputs; (3) which of the Section 106 rights fall within the scope of conveyances to generative AI providers; and (4) what are the purposes for which the generative AI provider can exercise the rights it obtains from the conveyance. As we'll explain momentarily, users rarely obtain licenses to system outputs, so parts (3) and (4) of our analysis here were relevant only with respect to generative AI providers.

User interests in system outputs. Most (66.3%) generative AI T&Cs are written such that users will own any copyrights that arise in generative AI outputs.⁶⁷ They typically do so with language that is framed as though the T&Cs are leaving intact copyright's default ownership rules.⁶⁸ A few use express assignment language.⁶⁹ Most generative AI T&Cs that include express assignment language include qualifiers, like "to the extent we acquire any rights," that reflect the prevailing uncertainty regarding the copyrightability of generative AI outputs. Meanwhile, a small fraction (3.5%) of generative AI T&Cs asserted that users had only a license to system outputs,⁷⁰ while nearly a quarter of generative AI T&Cs (24.4%) are silent as to whether users obtain any copyright interests in system outputs.

Four generative AI T&Cs attempt to make the user's interests turn on whether they have a free or paid subscription to the service.⁷¹ As an example, consider Leonardo's T&Cs. With respect to paying users, the T&Cs attempt to leave all copyright interests in system outputs in the user's hands.⁷² With respect to free users, Leonardo's T&Cs attempt to leave all copyright interests in system outputs in the provider's hands.⁷³ Another three asserted that users of the paid service obtained ownership of system outputs while users of the free service had only a license.⁷⁴

⁶⁷ Four of those draw distinctions between free and paid users.

⁶⁸ See, e.g., Jasper AI T&C ("We claim no ownership rights over Customer Property. The Customer Property remains yours.").

⁶⁹ See, e.g., Ideogram T&C ("As between us and you, to the extent we acquire any rights in any User Output, we hereby assign to you all right, title and interest in and to such User Output.").

⁷⁰ InVideo AI's language is representative: "Subject to these Terms, InVideo grants you a royalty-free, perpetual, irrevocable, worldwide, non-exclusive license to use, sub-license, reproduce, modify, adapt, publish, translate, create derivative works from, distribute, derive revenue or other remuneration from, and communicate to the public, perform and display your Output (in whole or in part) worldwide and/or to incorporate it in other works in any form, media or technology now known or later developed."

⁷¹ These are Leonardo, Suno, AI Studios, and Yodayo.

⁷² The relevant provision states "If you are on one of our paid Subscriptions (Paid Subscriber), as between the Parties, subject to your compliance with these Terms, ownership of all Intellectual Property Rights in any Content you, or your Authorised Users, create while using the Platform will vest in you upon creation, and to the extent that ownership of such Intellectual Property Rights does not automatically vest in you, we hereby assign all such Intellectual Property Rights to you and agree to do all other things necessary to assure your title in such rights."

⁷³ The relevant provision states "If you are on one of our free Subscriptions (Free Subscriber), as between the Parties, ownership of all Intellectual Property Rights in any Output you, or your Authorised Users, create while using the Platform will vest in us upon creation, and to the extent that ownership of such Intellectual Property Rights does not automatically vest in us, you hereby assign all such Intellectual Property Rights to us and agree to do all other things necessary to assure our title in such rights."

⁷⁴ These were Suno, AI Studios, and Yodayo.

User's Interest in System Output

N = 86 products

	<i>n</i>	%	
Owns	57	66.3	
Exclusive license	0	0.0	
Non-exclusive license	3	3.5	
No interests	1	1.2	
Silent	21	24.4	
Varies by tier	4	4.7	
Total	86	100.0	

Figure 5: User Interests in System Output

Generative AI provider interests in system outputs. The most common arrangement (62.8%) is for the generative AI provider to obtain a non-exclusive license in system outputs. In one instance, Leonardo, the provider obtains a non-exclusive license with respect to outputs generated by users with paid subscriptions and ownership with respect to outputs generated by users with free subscriptions.⁷⁵ In very few (4.7%) generative AI T&Cs, the provider purports to obtain or retain ownership of system outputs.⁷⁶ A similar fraction (5.8%) explicitly state that the provider obtains no interests at all. The remaining generative AI T&Cs (25.6%) do not say anything about the provider's copyright interests in system outputs.

Provider's Interest in System Output

N = 86 products

	<i>n</i>	%	
Owns	4	4.7	
Exclusive license	0	0.0	
Non-exclusive license	54	62.8	
No interests	5	5.8	
Silent	22	25.6	
Varies by tier	1	1.2	
Total	86	100.0	

Figure 6: Provider Interests in System Output

Section 106 rights conveyed to generative AI providers. Of the 54 generative AI T&Cs that conveyed a non-exclusive license in system output to the provider, most

⁷⁵ The other providers that purported to make user ownership turn on subscription status—SunO, AI Studios, and Yodayo—all state that the provider only obtains a non-exclusive license to system outputs, leaving unresolved the question of who owns system outputs for users of their free services.

⁷⁶ This count does not include Leonardo, which, as noted in the preceding text, obtains ownership for outputs generated by free users but not for outputs generated by paid users.

(79.6%) either explicitly identified the reproduction right or described actions that would fall within its scope, like hosting output.⁷⁷ The distribution right similarly appeared in many (64.8%) T&Cs that license system outputs to providers, as did the derivative works right (70.4%) and the public display right (72.2%). Public performance rights, meanwhile, appeared in 40.7% of generative AI T&Cs that conveyed non-exclusive licenses in system outputs to providers.

§106 Rights in System Outputs

N = 54 non-exclusive licenses

	<i>n</i>	%	
Reproduction	43	79.6	
Derivative works	35	64.8	
Distribution	38	70.4	
Performance	22	40.7	
Display	39	72.2	

Figure 7: Section 106 Rights in System Outputs Conveyed to Provider

Purposes for which providers can exercise their rights as grantees. As was the case with respect to user inputs, we examined whether generative AI T&Cs circumscribed the purposes for which the provider could exercise the interests it obtained in the system outputs. We identified four possible approaches. T&Cs might state that the provider could exercise its interests in system outputs: (1) only to provide the service to the user; (2) only to provide the service to other users; (3) to conduct any of its business activities; or (4) for any purpose whatsoever.

Of the 54 generative AI T&Cs that conveyed nonexclusive licenses in system outputs to the provider, just Tabnine limited its use of those interests only to providing the service to that user. Exactly half asserted the ability to exercise their interests in order to provide the service to other users. A little over one-third (35.2%) stated that they could use the licensed material for any business purpose. Finally, 13% of the nonexclusive licenses were either described as unlimited or did not contain any explicit limitations.

⁷⁷ In addition, Luma's T&C states that it has the ability to store or host system outputs produced by users of its paid service, but that it has all of the Section 106 rights with respect to users of its free service. We include Luma in our counts of licenses that identify each of the Section 106 rights.

Permitted Uses of System Outputs

N = 54 non-exclusive licenses

	<i>n</i>	%	
Only to serve this user	1	1.9	
To serve this and other users	27	50.0	
Any business activity	19	35.2	
No limitation	7	13.0	

Figure 8: Permitted Uses of System Outputs by Generative AI Provider

3. Allocating Copyright Infringement Risks

Our third category of terms concerns efforts to allocate the risk of copyright infringement. We focus on risks associated with the use of generative AI models.⁷⁸ Generative AI T&Cs can allocate these risks in two ways. The first relies on promises about copyright-relevant facts or future copyright-relevant conduct. The second relies on taking—or shifting—responsibility for losses incurred as a result of infringement claims.

Representations & Warranties. Representations and warranties are essential components of copyright licenses.⁷⁹ Because of copyright’s expansive liability provisions and harsh remedies, a party to a copyright license might be concerned about incurring liability as a result of their counterparty’s conduct. For example, licensees need some assurances that their licensors in fact have and can convey the rights they purport to convey; otherwise, the licensees’ use of the material conveyed could incur liability.⁸⁰ Representations and warranties offer a way to elicit explicit assurances on these matters. And in the context of a negotiated agreement, they are unobjectionable—if a potential licensor refuses to represent that they own the copyright they are offering to license, then that would be good reason for the potential licensee to walk away from the deal.

But representations and warranties are also pervasive in mass consumer contracts, including in online T&Cs.⁸¹ Most users do not know about or understand the representations and warranties that they are supposedly making.⁸² The

⁷⁸ The high-profile copyright lawsuits against generative AI defendants primarily concern whether model training infringed the plaintiffs’ interests. But model use can infringe whether or not training does. *See infra*, __.

⁷⁹ *See* 3 MILGRIM ON LICENSING § 23.00 (describing “representations, warranties, and covenants” as “the most critical nonmonetary part of a license agreement,” other than the scope of the license itself). A representation is a statement regarding the existence (or not) of some fact at the time of the agreement; a warranty is a guarantee regarding a future course of conduct or the existence of some fact in the future. *Id.*

⁸⁰ *See id.* (offering the following as an example: “The Licensor is the legal owner of the Licensed Patent”).

⁸¹ Boyack, *supra* note __ at 500–01.

⁸² Yannis Bakos et al., *Does Anyone Read the Fine Print? Consumer Attention to Standard-Form Contracts*, 43 J. LEGAL STUD. 1, 3 (2014) (finding that just 0.2% of internet users in large sample clicked through to hyperlinked T&Cs); WENDY WAGNER, *INCOMPREHENSIBLE! A STUDY OF HOW OUR LEGAL SYSTEM ENCOURAGES INCOMPREHENSIBILITY, WHY IT MATTERS, AND WHAT WE CAN DO ABOUT IT* 7–8 (2019); Melvin Aron Eisenberg, *The Limits of Cognition and the Limits of Contract*, 47 STAN. L. REV. 322, 354 (1995).

assertions in online T&Cs thus do not correspond to any real statements that consumers make. Instead, they are best understood as efforts—justified or not—by firms to shift liability risk to their customers. They do so because indemnification provisions are typically triggered by, among other things, a breach of any clause in the T&Cs; the representations and warranties thus establish a hook for indemnification.⁸³ As a result, these provisions turn ordinary users into copyright crumple zones, absorbing some of the impact from the collision between generative AI and the copyright regime.⁸⁴

We accordingly examined whether generative AI T&Cs contained the following representations and warranties: (1) that the user owns or has a license to the copyright in submitted material; (2) that the user has the right to convey her interests in submitted material; and (3) that no third-party copyright interests will be infringed by the user's use of the generative AI service. Although the first two may be sensible, the third is decidedly not—the stochastic nature of these products makes potential infringement challenging if not impossible to anticipate.

Nearly all (83.7%) of generative AI T&Cs include representations and warranties from their users. Of those, 76.4% include a statement that the user owns or has a license to material that the user submits. Nearly all generative AI T&Cs that include representations and warranties (88.9%) include one referring to the user's authority to grant permission to the provider to use the submitted material. Meanwhile, exactly half (50%) require that the user represent and warrant that they will not use the service in a way that infringes copyright or other IP rights.⁸⁵

Representations Made

N = 72 products imposing representations

	<i>n</i>	%	
Owns/licensed copyright	55	76.4	
Right to convey	64	88.9	
No infringement	36	50.0	

Figure 9: Representations and Warranties Made by Users

Indemnifications. In negotiated agreements, indemnification provisions are a sensible way to deal with the possibility that an innocent party could be held liable because of its counterparty's decisions.⁸⁶ A book retailer could be held liable for

⁸³ We address those provisions *infra* ____.

⁸⁴ Cf. Madeline Clare Elish, *Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction*, 5 ENGAGING SCI., TECH., & SOC'Y 40 (2019).

⁸⁵ These often also include representations and warranties regarding privacy torts, rights of publicity, and moral rights.

⁸⁶ Copyright law has not been found to provide for a right of contribution, so parties instead rely on indemnification provisions in their contracts. See JORGE CONTRERAS, INTELLECTUAL PROPERTY LICENSING AND

selling an infringing novel even though it had nothing to do with its creation.⁸⁷ To protect against this possibility, a retailer would demand that a publisher indemnify it against such claims (and the publisher, in turn, would sensibly demand a similar assurance from the author).

Indemnification provisions in online T&Cs, however, have a different character. Unlike the parties to a negotiated agreement, users will rarely know that they could incur liability through ordinary use of generative AI services; given the stochastic nature of these systems and the fact that copyright is a strict liability regime, users may also be unable to avoid incurring liability. To be sure, it may be unfair for a provider to bear responsibility when a user deliberately evades guardrails that the provider implemented to prevent infringement. But indemnification provisions could reach quite broadly, forcing users to bear liability for a wide range of conduct that is far more within the provider's control than their own—liability for infringement of the reproduction right, for example, could turn on technical details of system design, like how long a user's input is retained by the system.⁸⁸ Users might also be misled about infringement liability because of the generative AI provider's public statements or product design.⁸⁹ And users could be bankrupted by liability that providers could easily bear.⁹⁰

We therefore examined the following aspects of the indemnification provisions in generative AI T&Cs: (1) the triggers that create the indemnification obligation; and (2) the losses covered by the indemnification.

Seventy-five generative AI T&Cs include a clause requiring the user to indemnify the generative AI provider. These clauses typically incorporate an array of triggers. Of the generative T&Cs that include an indemnification provision, virtually all (85.3%) state that breach of the T&Cs can trigger an indemnification obligation, while 80% state that it can be triggered by use of the service generally. Three-quarters (74.7%) specifically identify user content as the source of the indemnified liability. A small number (13.3%) include AI-generated output as a potential basis for indemnification of the provider by the user.

TRANSACTIONS: THEORY AND PRACTICE 299 (2022) (noting that “most sophisticated IP transactions include provisions by which the parties seek to allocate the risk of third-party infringement” among themselves because “IP licensees cannot assume that the rights that they license will permit them to practice a particular technology, or that they will not later become subject to infringement claims by third parties”); 4 NIMMER ON COPYRIGHT § 31E.12 (D)(2)(b) (citing district court cases holding that federal copyright law does not include a right to contribution or indemnification). While the common law in some states may provide for indemnification, there's little reason to rely on that backstop because all states would generally enforce indemnification provisions consistent with general contract law. *See id.* (reviewing federal and state causes of action for contribution or indemnification in the context of copyright liability and concluding that “[w]hen a party voluntarily undertakes an obligation to indemnify, it stands to reason that the state law of contracts renders the obligation enforceable”).
⁸⁷ *See* 17 U.S.C. § 106(3) (conferring upon the author the exclusive right “to distribute copies . . . to the public by sale”).

⁸⁸ *Compare* *MAI Sys. Corp. v. Peak Comput., Inc.*, 991 F.2d 511 (9th Cir. 1993) (holding that a computer program loaded into random-access memory constituted a “copy”) *with* *Cartoon Network, LP v. CSC Holdings, Inc.*, 536 F.3d 121 (2d Cir. 2009) (holding that a video stream present in a buffer did not constitute a “copy”).

⁸⁹ *See infra* text accompanying notes ____.

⁹⁰ *See infra* text accompanying notes ____.

Nearly all (93.3%) indemnification clauses cover claims, demands, or causes of action. In addition, most ultimately (57%) include the broadest possible exposure, encompassing (1) claims, demands, or causes of action; (2) damages, losses, or liability; (3) costs or expenses; and (4) attorneys' fees. Another fourteen percent include three out of those four, with roughly equal numbers omitting each one.

Indemnification Triggers

N = 75 products requiring indemnification

	<i>n</i>	%	
Breach of T&Cs	64	85.3	
Use of the service	60	80.0	
User content	56	74.7	
AI output	10	13.3	

Figure 10: Triggers Giving Rise to Indemnification Obligation on Users

Losses Covered

N = 75 products requiring indemnification

	<i>n</i>	%	
Claims/demands/causes of action	70	93.3	
Damages/losses/liability	61	81.3	
Costs/expenses	59	78.7	
Attorneys' fees	59	78.7	

Figure 11: Indemnification Coverage

Meanwhile, only two T&Cs include a clause by which the generative AI provider promises, in certain circumstances, to indemnify the user for claims that a system's output infringes a third-party's copyright interests.⁹¹ The upshot is that, when users elicit outputs from generative AI systems, the T&Cs will typically require users to bear all resulting copyright liability. And this is so even if black-letter copyright doctrine would impose liability on the generative AI provider.

B. *Dispute-Related Terms*

1. Limiting Access to Adjudication

The first category of boilerplate dispute-related terms covered by our study can limit or even effectively eliminate parties' ability to access formal adjudication. We begin with mandatory-arbitration clauses, perhaps the single most heavily

⁹¹ These are Adobe Firefly and SummarizeBot.

publicized and criticized type of boilerplate dispute-resolution term. We turn then to class waivers, provisions that purport to shorten limitations periods, forum selection clauses, and jurisdiction waivers.

Mandatory arbitration. Mandatory-arbitration clauses receive great force in U.S. jurisdictions due to the Supreme Court’s recent interpretations of the Federal Arbitration Act.⁹² Just over half (54.7%) of the generative AI T&Cs surveyed include an arbitration provision governing claims by users against providers. Nearly all of these are mandatory, though a few contain limited opt-outs.⁹³ Interestingly, the majority of the T&Cs with arbitration provisions include carveouts for certain types of disputes and/or venues. Most of the T&Cs that mandate arbitration nonetheless allow parties to pursue certain claims in court (n = 36). The most common carveout permits lawsuits in small-claims court (n = 31),⁹⁴ followed by suits involving intellectual-property rights (n = 28),⁹⁵ and suits seeking only injunctive relief (n = 8).⁹⁶ A few T&Cs also granted carveouts for pursuing miscellaneous other types of disputes.⁹⁷

Class Waivers. When coupled with a mandatory-arbitration term—as they often are⁹⁸—class waivers are virtually always enforced.⁹⁹ Even standalone class waivers are often binding, though some courts historically refused to enforce them on policy grounds.¹⁰⁰ Just over half (51.2%) of the T&Cs we surveyed include a class-action waiver. These are virtually all mandatory; only one permits an opt-out for users.¹⁰¹ The generative AI T&Cs that include a class-action waiver (n = 44) nearly always couple it with an arbitration provision (n = 41).

Contractual Limitations Periods. Adjudicators often give binding force to terms that purport to shorten the available time period in which injured parties can

⁹² *AT&T Mobility LLC v. Concepcion*, 536 U.S. 333 (2011); Jodi Wilson, *How the Supreme Court Thwarted the Purpose of the Federal Arbitration Act*, 63 CASE W. RES. L. REV. 91 (2012).

⁹³ Open AI’s T&Cs, for example, let users opt out of arbitration within a thirty-day window after the user creates their account or Open AI updates its T&Cs.

⁹⁴ The T&Cs for Microsoft Copilot go on to state that “either you or we may sue the other party in small claims court seeking only individualized relief, so long as the action meets the small claims court’s requirements and remains an individual action seeking individualized relief.”

⁹⁵ For example, the T&Cs for Microsoft Copilot state that “if we do [have a dispute], you and we agree to binding individual arbitration The term ‘dispute’ is as broad as it can be. It includes any claim or controversy . . . except disputes relating to the enforcement or validity of your, your licensors’, our, or our licensors’ intellectual property rights.”

⁹⁶ The Adobe Firefly T&Cs, for example, state that “either party shall be entitled to apply for preliminary injunctive remedies . . . in any jurisdiction.”

⁹⁷ A few T&Cs also identified other types of disputes. Janitor AI, for example, extended its arbitration carveout to claims for “defamation [and] violation of the Computer Fraud and Abuse Act” in addition to IP infringement claims.

⁹⁸ Christopher R. Leslie, *The Arbitration Bootstrap*, 94 TEX. L. REV. 265, 274–81 (2015).

⁹⁹ See *id.* at 281 (“The combination of *Conception* and *Italian Colors* makes it exceedingly difficult for the victims of illegal conduct to challenge the enforceability of class action waivers [coupled with mandatory-arbitration clauses].”).

¹⁰⁰ See, e.g., *Homa v. Am. Express Co.*, 558 F.3d 225, 230 (3d Cir. 2009) (“[T]he [unconscionability] defense is a general contract defense, one that applies to all waivers of class-wide actions, not simply those that also compel arbitration.”).

¹⁰¹ The Tavus T&Cs permitted users to opt out.

initiate legal proceedings.¹⁰² A great many federal and state statutes impose outside limits on when injured parties can bring various types of claims.¹⁰³ Because these caps are statutory, drafters cannot lengthen,¹⁰⁴ but can seek to shorten, the limitations period.¹⁰⁵ Even so, relatively few (11.6%) of the generative AI T&Cs specify a contractual limitations period. Nearly all of these (n = 8) set one year as the outside time limit for commencing claims.¹⁰⁶

Forum Selection. U.S. courts usually give binding force to terms that purport to select a forum for any disputes.¹⁰⁷ Forum-selection clauses are quite common in our sample: 81.4% of these generative AI T&Cs contain a provision that specifies a geographic forum for dispute resolution. Northern California was by far the most commonly selected region (n = 33).¹⁰⁸ Given the prevalence of generative AI providers headquartered in the area, this suggests that drafters are generally seeking home-court advantage.¹⁰⁹ A substantial minority (n = 19) of T&Cs specify a non-U.S. jurisdiction for dispute resolution, despite these likely being unenforceable in the United States.¹¹⁰

Jurisdiction Waivers. Nearly half (46.5%) of the generative AI T&Cs in our study include some form of jurisdiction waiver provision. Jurisdiction-waiver provisions can take multiple forms. Some state that one or both parties waive their right to contest whether a chosen forum actually has jurisdiction to hear a dispute.¹¹¹ Others state that parties submit “exclusively” to a particular jurisdiction.¹¹² An additional set of T&Cs did not separately select a forum but included an

¹⁰² *E.g.*, *Heimeshoff v. Hartford Life & Accident Ins. Co.*, 571 U.S. 99 (2013) (affirming enforceability of a privately imposed limitations clause that specified both a period and a start date). Although these terms often receive binding force, exceptions do exist. *See, e.g.*, *Logan v. MGM Grand Detroit Casino*, 939 F.3d 824 (6th Cir. 2019) (holding that a privately drafted limitations clause could not shorten a statutorily created limitations period).

¹⁰³ *E.g.*, 28 U.S.C. § 1658(a) (imposing default limit of four years after the cause of action accrues for civil actions brought under federal statutes).

¹⁰⁴ *See, e.g.*, *Home Ins. Co. v. Dick*, 281 U.S. 397, 405–06 (1930) (deeming it unconstitutional to enforce a state law that would have invalidated a one-year contractual limitations period as against an out-of-state insurer with no notice that that state’s law would apply).

¹⁰⁵ *See, e.g.*, *Shaw v. Aetna Life Ins. Co.*, 395 A.2d 384, 386–87 (Del. Super. Ct. 1978).

¹⁰⁶ One set of T&Cs, Speechify’s, sets the limitations period at “one (1) year and twenty-eight (28) days after such claim or cause of action arose.”

¹⁰⁷ John F. Coyle, “Contractually Valid” *Forum Selection Clauses*, 108 IOWA L. REV. 127, 161 (2022) (finding that federal courts gave force to forum-selection clauses 88% of the time in a sample of 658 reported decisions).

¹⁰⁸ Janitor AI’s T&Cs, which were perhaps themselves generated by an AI tool, appear to have hallucinated a “Netherlands County” in California that does not exist. The term in question was somewhat ambiguously worded, but appears to contemplate both a forum selection and jurisdiction waiver in favor of this (nonexistent) county court system.

¹⁰⁹ This is occasionally explicit: AI Studios’s T&Cs, for example, specify that disputes are to be brought “in the county in California, USA, where the Company is located.” Another four sets of T&Cs specified California generally, for a total of 37 selecting some part or all of California.

¹¹⁰ For example, Pixlr’s T&Cs state that all disputes must be “settled by arbitration in Singapore.”

¹¹¹ Ken Adams, *Jurisdiction provisions, Part 1*, ADAMS ON CONTRACT DRAFTING (Aug. 24, 2019), <https://www.adamsdrafting.com/stop-using-consent-to-jurisdiction-provisions/>.

¹¹² *Id.* Adams contends that jurisdiction waivers are formally redundant of forum-selection clauses, citing some supportive case law. *Id.* (citing *Jacobsen Constr. Co. v. Teton Builders*, 106 P.3d 719, 728 (Utah 2005)). But uncertainty persists, such that a jurisdiction waiver may increase both the formal force and any in terrorem effects of a forum-selection clause.

“exclusive” jurisdiction provision. Effectively, then, mandatory forum selection is even more common, with 47.7% (n = 41) sets of T&Cs imposing it.

Figure 12 graphically displays these findings. Understanding the practical effect of these provisions requires considering them holistically; consequently, we also report cumulative use. Nearly half (45%) of the generative AI T&Cs in our sample include three or more of these access-limiting terms.

Limiting Access to Adjudication

N = 86 products

	<i>n</i>	%	
Mandatory Arbitration	47	54.7	
Class Waiver	44	51.2	
Contractual Limitations Period	10	11.6	
Forum Selection	70	81.4	
Jurisdiction Waiver	40	46.5	
At Least 3	40	46.5	

Figure 12. Prevalence of Access-Limiting Terms

2. Eliminating Causes of Action

The second group of generative AI T&Cs our study analyzed can effectively eliminate causes of action. The discussion that follows begins with terms, sometimes known as “assumption of risk” provisions, that avoid negligence claims. It turns next to terms that exclude warranties and concludes with terms (“merger” or “integration” clauses) that void prior promises and statements.

Voiding Negligence Claims, Assumption of Risk Provisions. U.S. courts generally—and increasingly—give legal force to terms that void negligence claims.¹¹³ And firms often deploy these terms even when they are formally unenforceable, suggesting they can have an in *terrorem* effect on users.¹¹⁴ These provisions are pervasive in the generative AI T&Cs in our sample: 84.9% include a provision that purports to void negligence claims. The provisions took multiple forms, likely due in part to underlying doctrinal muddiness. Some language invokes the tort-defense concept of users “assuming the risk.”¹¹⁵ Other language, sometimes

¹¹³ These clauses trace their history to a common-law tort defense, assumption of risk, but have increasingly morphed into contractual waivers—instead of creating a tort defense, they contractually eliminate the tort cause of action. Takshid, *supra* note __. That shift was significant. Courts using a tort framework evaluated enforceability using a relatively plaintiff-friendly public-policy test. *Id.* at 2192. Courts using a contract lens instead generally ask whether the clause is unconscionable, a high hurdle to clear. *Id.* at 2196.

¹¹⁴ Edward K. Cheng, *Unenforceable Waivers*, 76 VAND. L. REV. 571, 582–88 (2023).

¹¹⁵ Claude’s T&Cs provide, under the heading “Disclaimer of warranties, limitations of liability, and indemnity,” that “Your use of the services and materials is solely at your own risk.” SeaArt’s T&Cs state that “You assume the risk of using” user-generated content.

appearing in the same set of T&Cs, more forcefully purports to void or “disclaim” any legal liability whatsoever resulting from an individual’s use of the provider’s AI tool.¹¹⁶ These terms purport to go further than any previous studies appear to have uncovered. Their enforceability is, as of yet, unclear.

Excluding Warranties. U.S. courts generally enforce warranty-excluding terms to eliminate claims for breach of contract, negligent misrepresentation¹¹⁷ and, increasingly, even claims for intentional, fraudulent misrepresentation.¹¹⁸ In the generative AI T&Cs in our study, warranty-excluding provisions are nearly ubiquitous: 97.7% of T&Cs in our sample include provisions that exclude or limit warranties. Drafters generally expressly extended these to both “warranties” and “representations,” contemplating the elimination of potential contract and tort claims. The T&Cs also generally extended these to both “express” and “implied” warranties. And they generally meet the legal requirements to be “conspicuous,” mention the word “merchantability,” and include the “as is” language provided as an example disclaimer by the Uniform Commercial Code..

Voiding Prior Promises, Merger Clauses. U.S. courts generally enforce terms that purport to nullify extrinsic promises and statements so as to eliminate any breach-of-contract claims based on those promises and statements.¹¹⁹ Clauses that seek to void prior promises and agreements, often known as “merger clauses” or “integration clauses,” are quite common in our sample: 67.4% of the generative AI T&Cs in our study include a merger clause.

Figure 13 below displays these findings. As to cumulative use, 90.7% of generative AI T&Cs in our sample include at least two types of claim-limiting terms. And half of the generative AI T&Cs in our sample impose all three types of claim-eliminating terms on their users.

¹¹⁶ The Adobe Firefly T&Cs, for example, state that the “Covered Parties specifically disclaim all liability for any actions resulting from your [i.e., the user’s] use of any Services and Software.”

¹¹⁷ See Glenn D. West & W. Benton Lewis, Jr., *Contracting to Avoid Extra-Contractual Liability—Can Your Contractual Deal Ever Really Be the “Entire” Deal?*, 64 BUS. LAW. 999, 1021 (2009).

¹¹⁸ *Id.* at 1027–29.

¹¹⁹ See, e.g., *Ungerleider v. Gordon*, 214 F.3d 1279 (11th Cir. 2000) (applying the parol evidence rule to void an alleged prior promise because, inter alia, the parties’ written agreement stated that it was a complete integration of the contractual terms). The only contrary examples generally involve extremely incomplete writings that purport to be the complete statement of the terms. See, e.g., *L.S. Heath & Son, Inc. v. AT & T Info. Sys., Inc.*, 9 F.3d 561 (7th Cir. 1993) (holding that, although “the presence of a merger clause is strong evidence that the parties intended the writing to be the complete and exclusive statement between them,” a written agreement that failed to identify any prices, products, or services was incomplete).

Eliminating Causes of Action

N = 86 products

	<i>n</i>	%	
Assumption of Risk	73	84.9	
Excluding Warranties	84	97.7	
Merger/Integration Clauses	58	67.4	
At Least 2	78	90.7	
At Least 3	50	58.1	

Figure 13. Prevalence of Claim-Eliminating Terms

3. Shifting the Balance of Power

The third category of dispute-resolution terms we analyze tend to shift the balance of power in favor of the drafting party. This category includes provisions that empower one party to modify the T&Cs, confidential-proceeding requirements, choice-of-law provisions, choice-of-procedure provisions, and cost-shifting provisions.

Unilateral modification, with or without notice requirement. Terms that permit unilateral modification of the contract are in tension with foundational contract notions like mutuality of exchange.¹²⁰ That tension causes some judges to balk at enforcing them,¹²¹ but most U.S. courts enforce these provisions, especially when they identify some obligation (however minimal) for the empowered party to notify the other of any changes.¹²² Nearly all (96.5%) of the T&Cs in our sample allow one party to modify the terms without first obtaining the other party's agreement or consent. Most of these (*n* = 76) are expressly one-sided: in other words, 88.4% of the T&Cs expressly allow providers, not users, to modify the terms.¹²³ Over half (*n* = 47) allow unilateral modification without any meaningful notice requirement.¹²⁴ Many of these include no notice requirement at all.¹²⁵ Others provide for minimal, but not meaningful, notice.¹²⁶

¹²⁰ RESTATEMENT (SECOND) OF CONTRACTS § 71 (bargained-for exchange theory of consideration).

¹²¹ *E.g.*, *Johnson v. Continental Fin. Co.*, No. 23-2049, at *15 (4th Cir. Mar. 11, 2025) (holding that a unilateral-modification clause rendered an arbitration agreement illusory and nonbinding under state law); *Cheek v. United Healthcare of Mid-Atl., Inc.*, 835 A.2d 656, 658 (Md. 2003) (same).

¹²² *E.g.*, *Holloman v. Circuit City Stores, Inc.*, 894 A.2d 547, 554 (Md. 2006).

¹²³ For example, Yodayo's T&Cs state, "We reserve the right, at Our sole discretion, to modify or replace these Terms at any time."

¹²⁴ For example, Vidnoz's T&Cs "reserve[] the right to update these terms and conditions . . . without prior notice."

¹²⁵ *Id.*

¹²⁶ Some of the T&Cs (You.Com, for example) stated that the provider would simply change the "last revised" date at the top of the webpage containing the T&Cs. Particularly in light of research indicating that very few users ever click through to the T&Cs webpage even upon initial sign-up, it is reasonable to conclude that virtually zero users periodically check T&Cs webpages for updates and revisions.

Confidential-proceeding requirements. Terms that purport to require confidential proceedings are generally enforceable as to arbitration,¹²⁷ though they are not binding on U.S. courts.¹²⁸ A minority (n = 13) of the T&Cs we surveyed include a requirement that dispute resolution be kept confidential. Unsurprisingly, given that confidential-proceeding terms are generally not enforced by courts, nearly all (n = 12) of these T&Cs also include an arbitration provision.

Choice-of-law provisions. Terms that purport to select governing law are virtually always binding in arbitration¹²⁹ and generally enforceable in court.¹³⁰ A clear majority (91.9%) of the generative AI T&Cs sampled include a choice-of-law provision. California law was most commonly selected by far (n = 36); as with choice-of-forum provisions, this suggests firms are selecting home-court law. New York is the next most commonly selected U.S. state (n = 6), followed by Delaware (n = 4). Only two T&Cs in our sample choose the law of a user's home jurisdiction.¹³¹

Choice-of-procedure provisions. In practice, drafters rarely attempt to choose procedural rules for litigated disputes,¹³² instead confining themselves to identifying rules for arbitration.¹³³ About half (51.2%) of the T&Cs include a choice-of-procedure provision. All of the T&Cs (n=44) that identify a chosen set of procedural rules also include an arbitration provision. In other words, generative AI providers in our sample do not attempt to select procedural rules for in-court dispute resolution, only for arbitration. The American Arbitration Association's

¹²⁷ See *id.*

¹²⁸ Arbitration is generally not a confidential process by default, but some arbitration institutions impose confidentiality rules, as do some state laws governing arbitration. See, e.g., *Confidentiality in US Arbitration*, PRACTICAL LAW ARBITRATION (last visited Jan. 22, 2026) (available via Westlaw, resource ID W-002-7851).

¹²⁹ An arbitrator may even apply a set of religious beliefs if T&Cs prescribe it See, e.g., Debra Cassens Weiss, *Buying a Bamboo Floor? This Company Requires Religious Arbitration; Courts Uphold Clauses*, ABA J. (Nov. 3, 2015), http://www.abajournal.com/news/article/buying_a_bamboo_floor_this_company_requiresreligious_arbitration_courts_up.

¹³⁰ This is especially true as to "interpretation" issues. See RESTATEMENT (SECOND) OF CONFLICT OF LAWS § 187. It is true even as to fundamental "validity" issues given at least some relationship between the parties and the chosen law. See *id.* (requiring a "substantial relationship to the parties or transaction" or some "other reasonable basis" for the chosen law); UNIFORM COMM. CODE § 1-301 (requiring only a "reasonable relationship"). Choice-of-law terms provisions can therefore displace the generally applicable default choice-of-law rules that exist in every U.S. state. See John F. Coyle, William S. Dodge & Aaron D. Simowitz, *Choice of Law in the American Courts in 2024: Thirty-Eighth Annual Survey*, 73 AM. J. COMPARATIVE L. (forthcoming 2025).

¹³¹ The Microsoft Copilot T&Cs state that, as to U.S.-based users, "the laws of the state where you live (or, if a business, where your principal place of business is located) govern all claims." The Rephrase T&Cs state that disputes are to be resolved "in accordance with the applicable laws of your jurisdiction."

¹³² Interestingly, most commentators who have weighed in on the issue believe that parties do have the power to choose procedural rules. Robert G. Bone, *Party Rulemaking: Making Procedural Rules Through Party Choice*, 90 TEX. L. REV. 1329, 1346 (2012) (noting that this commentary mostly relies on a single case for support).

¹³³ For example, terms can specify that the UN Commission on International Trade Law ("UNCITRAL") Arbitration Rules apply. *UNCITRAL Arbitration Rules*, UN COMM'N ON INT'L TRADE L., <https://uncitral.un.org/en/texts/arbitration/contractualtexts/arbitration> ("The UNCITRAL Arbitration Rules provide a comprehensive set of procedural rules upon which parties may agree for the conduct of arbitral proceedings arising out of their commercial relationship and are widely used in ad hoc arbitrations as well as administered arbitrations.").

rules are most common (n = 21), followed by JAMS (formerly Judicial Arbitration and Mediation Services, Inc.) procedural rules (n = 11), and National Arbitration and Mediation's procedural rules (n = 4).

Cost-Shifting Provisions. Courts often give binding force to these terms,¹³⁴ which vary the default U.S. rule that each side must bear its own costs.¹³⁵ A substantial minority (41.9%) of the generative AI T&Cs in our sample include a provision addressing the allocation of dispute-resolution costs. All of these, except one outlier,¹³⁶ also include an arbitration provision. So the generative AI providers almost never try to allocate dispute-resolution costs for litigation, only arbitration. Most of the cost-shifting provisions (n = 21) analyzed contemplate some form of cost-splitting between the parties. The moving party pays at least some of the costs in six of the T&Cs in our sample. Providers agree to pay at least some of the costs in five of the T&Cs in our sample. The losing party pays in three of the T&Cs in our sample. Only one set of T&Cs expressly put costs on users, making cost-shifting provisions the most neutral—indeed, even user-friendly—of any of the dispute-resolution terms analyzed in our study.

Figure 14 below displays these findings, as well as cumulative-use statistics. We excluded from the latter cost-shifting terms, which were (as noted above) neutral-to-user-friendly in the generative AI T&Cs we studied. Even so, nearly all of the generative AI T&Cs in our study include at least two power-shifting terms. More than half deploy at least three.

Shifting Balance of Power

N = 86 products

	<i>n</i>	%	
Unilateral Modification	83	96.5	
Confidential Proceedings	13	15.1	
Choice of Law	79	91.9	
Choice of Procedure	44	51.2	
Cost-Shifting	36	41.9	
At Least 2 [†]	80	93.0	
At Least 3 [†]	46	53.5	

[†] Excluding cost-shifting.

Figure 14. Prevalence of One-Sided, Power-Shifting Terms

¹³⁴ *E.g.*, *In re Stewart v. Capital City Morg. Corp.*, No. 00-00046, 2004 WL 3130573 (Bankr. D.D.C. Nov. 10, 2004).

¹³⁵ This is generally known as the “American Rule.” Nierengarten, *supra* note __, at 227.

¹³⁶ The Cleanup Pictures T&Cs allow users the option to attempt dispute resolution via mediation. Cleanup Pictures will, according to its T&Cs, pay the costs of mediation.

4. Limiting Remedies

The final category of terms analyzed herein tend to limit remedies for prevailing parties to disputes. This category includes liquidated-damages provisions that set a fixed figure or upper bound for damages awards, terms that avoid entire categories of damages like consequential or punitive damages, and jury-trial waivers that prevent juries (who tend to award higher amounts than judges) from deciding damage figures.

Liquidated-damages clauses. Terms that expressly fix an amount for—or cap on—monetary damages are generally enforceable.¹³⁷ These provisions are quite common among providers included in our study, with 83.7% of T&Cs including some form of liquidated-damages clause. These clauses calculate damages figures in three different ways: the amount(s) paid by users to providers, a fixed sum, or a hybrid of the two.¹³⁸ Most of the T&Cs that impose a liquidated-damages provision in our study use a hybrid structure (n = 41).¹³⁹ Next most common was limiting damages to the amount(s) paid by users (n = 24). Relatively few use only a fixed sum (n = 7).¹⁴⁰ The highest of the fixed-sum figures is \$500,¹⁴¹ and the lowest is \$50.¹⁴²

Avoiding categories of damages. Whether U.S. courts will enforce terms that purport to avoid entire damages categories remains somewhat unsettled, especially as to user-facing online T&Cs.¹⁴³ Nonetheless, nearly all (96.5%) of the generative AI T&Cs in our sample prohibit certain categories of damages. The lists of covered categories are generally quite extensive. Prohibited categories of

¹³⁷ U.S. courts generally require only that damages were ex ante difficult to predict, and (somewhat paradoxically) the specified amount was a reasonable estimate of actual damages that would accrue. *See, e.g.,* Deborah S. Coldwell et al., *Liquidated Damages*, 29 FRANCHISE L.J. 211, 211 (2010). These two requirements operate on a sliding scale: the more difficult damages were to estimate, the more willing judges are to deem the specified amount “reasonable.”

¹³⁸ Given the prevalence of zero-price generative AI product offerings, the popularity of this approach in our sample may reflect drafters’ concerns about potential unenforceability. On zero-price business models generally, see John M. Newman, *Attention Capitalism: The Law and Political Economy of Attention Markets*, 78 STANFORD L. REV. 415 (2026); John M. Newman, *The Myth of Free*, 86 GEO. WASH. L. REV. 513 (2018). For zero-price products, a “what the user paid” term would tend to yield a liquidated amount of \$0, potentially violating the general U.S. rule that liquidated damages provisions must represent a reasonable estimate of actual damages.

¹³⁹ Adobe Firefly’s T&Cs, for example, provide that damages are “limited to the greater of (A) US \$100; or (B) the aggregate amount that you paid for access to the Services and Software during the three-month period preceding the event giving rise to the liability.”

¹⁴⁰ For example, Speechify’s T&Cs provide that “[i]n no event shall the total aggregate liability. . . exceed, in the aggregate, the amount, if any, paid by you to Speechify for your use of the Services.”

¹⁴¹ Luminar AI’s T&Cs state that damages awards shall not be “in excess of \$500.00.”

¹⁴² InVideo’s T&Cs cap damages at a “MAXIMUM OF \$50.”

¹⁴³ Courts generally enforce these terms in business-to-business contexts. *See, e.g.,* Far East Aluminum Works Co. v. Viracon, Inc., No. 21-1875 (8th Cir. Mar. 9, 2022) (enforcing limit on consequential damages in business-to-business contract); Tenn. Gas Pipeline Co. v. Technip USA Corp., No. 01-06-00535-CV, 2008 Tex. App. LEXIS 6419, at *2 (Tex. App. Aug. 21, 2008) (same).

damages include “incidental,” “indirect,” “consequential,” “moral,” “exemplary,” “punitive,” and more.¹⁴⁴

Jury-trial waivers. Jury-trial waivers were less common than the other types of remedy-limiting terms, with 39.5% of T&Cs containing a jury waiver. The context in which these appear is counterintuitive: most of the T&Cs that include a jury waiver also contain an arbitration provision. Yet arbitration itself effectively eliminates the possibility of a jury trial.¹⁴⁵ Follow-on qualitative research might helpfully explore the strategy (if any) behind this practice.

Figure 15 below displays these findings, as well as cumulative-use statistics. Most (82.6%) of the generative AI T&Cs in our sample contain at least two types of remedy-limiting terms. Nearly forty percent contain all three types of remedy-limiting terms.

Limiting Remedies

N = 86 products

	<i>n</i>	%	
Liquidated Damages	72	83.7	
Avoiding Categories	83	96.5	
Jury Waiver	34	39.5	
At Least 2	71	82.6	
All 3	34	39.5	

Figure 15. Prevalence of Remedy-Limiting Terms

IV. IMPACT ANALYSIS

The empirical findings set forth above shed light on a number of foundational questions that, until now, remained largely unanswered. From a social-welfare perspective, at least some data are encouraging. But most of our findings are troubling. The discussion that follows normatively evaluates the most salient issues raised by our empirical results. And, as to particularly problematic practices identified by our study, we offer responsive policy prescriptions.

¹⁴⁴ For example, ChatGPT’s T&Cs state, “NEITHER WE NOR ANY OF OUR AFFILIATES OR LICENSORS WILL BE LIABLE FOR INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL, OR EXEMPLARY DAMAGES, INCLUDING DAMAGES FOR LOSS OF PROFITS, GOODWILL, USE, OR DATA OR OTHER LOSSES . . .”

¹⁴⁵ See Lionel M. Schooler & Luke Gillman, *Private Juries Within the Arbitration Framework: A Third Path in Dispute Resolution*, 21 PEPPERDINE DISPUTE RESOLUTION L.J. 457, 458 (2021) (proposing increased use of private juries in arbitration while conceding that the practice is currently “exceedingly rare”). Many of the arbitration provisions in question do include carveouts permitting some dispute resolution in court. But these carveouts are for small-claims court (where a jury request typically moves the case to a higher court, rendering the contractual carveout inapplicable) and civil actions seeking only injunctive relief (where plaintiffs have no right to a jury). These jury-trial waivers may simply reflect a belt-and-suspenders approach, a hypothesis supported further by the fact that all of the T&Cs with a jury-trial waiver also contain a liquidated-damages term and a prohibition on categories of damages.

A. *The Effects of Copyright Terms*

1. On Control over User Inputs

When generative AI burst into public consciousness with the release of ChatGPT, user inputs were mostly simple instructions and back-and-forth chats.¹⁴⁶ Since then, consumer-facing generative AI systems have evolved so that users can submit much more sophisticated content. These systems can now take as inputs a vast array of expressive works—lengthy essays, complex images, and sophisticated music. All of this material will ordinarily be protected by copyright. In many instances, then, users themselves will own copyrights before they begin using the generative AI system. The question of copyright interests in prompts has accordingly become urgent.¹⁴⁷

Our analysis reveals that, when users submit their copyright-protected materials through generative AI interfaces, they are typically also granting the provider a worldwide, royalty-free non-exclusive license to those materials.¹⁴⁸ That license will often be perpetual, irrevocable, and transferable or sublicensable to others.¹⁴⁹ The license typically encompasses the reproduction, display, distribution, and derivative works rights, and almost half the time it can be deployed to any of the provider’s business purposes.¹⁵⁰

These provisions accordingly have the potential to eviscerate authors’ exclusive rights in any works they submit to generative AI systems. To see why, suppose a photographer, call her Lynn, takes a portrait of a famous musician in her studio. She then uploads that portrait to Midjourney, where she asks for various modifications until she is satisfied with the end result.

At this point, Midjourney can share Lynn’s original portrait—the one she took in her studio with no contribution by Midjourney whatsoever—with anyone for any reason.¹⁵¹ Midjourney could allow any of its users to use Lynn’s portrait to

¹⁴⁶ In this period, short context windows and limited user interfaces prevented users from submitting much aside from direct instructions to generative AI systems. To be sure, some users continue to use generative AI systems this way. When they do, the prompts will typically not be protected by copyright law. Although copyright law’s originality requirement is famously weak, it still exists—prompts like “tell me a story” or “amuse me” would fall short of it. Moreover, even prompts that are original must not run afoul of various exclusions from the scope of copyrightable subject matter, like the idea-expression dichotomy and the “method of operation” exclusion. 17 U.S.C. § 102(b) (providing that “copyright protection for an original work of authorship [does not] extend to any idea [or] method of operation,” among other things); *Baker v. Selden*, 101 U.S. 99 (1879) (establishing the idea-expression dichotomy). In these instances, we need not be concerned about who owns the copyright interests in material submitted by a user, for the simple reason that there are no such copyright interests.

¹⁴⁷ Cf. Mark A. Lemley, *How Generative AI Turns Copyright Upside Down*, 25 COLUM. SCI. & TECH. L. REV. 21 (2024) (exploring the possibility of “center[ing] copyright around whether enough human creativity goes into the prompt”).

¹⁴⁸ See *supra* text accompanying notes ____.

¹⁴⁹ *Id.*

¹⁵⁰ *Id.*

¹⁵¹ Midjourney’s license conforms to the common approach: “You grant to Midjourney, its affiliates, successors, and assigns a perpetual, worldwide, non-exclusive, sublicensable no-charge, royalty-free, irrevocable copyright

produce other generative AI outputs. It could display Lynn’s portrait on its front page to attract new users.¹⁵² It could even open a brick-and-mortar store where it sells Lynn’s portrait at the marginal cost of a copy.¹⁵³ Or it can simply open an online gallery of user-submitted works, including not just Lynn’s but all of the works submitted by all of its users. Lynn’s ability to charge a copyright-enabled premium for her original portrait will have been eroded to nonexistence. And lest these examples seem fanciful, note that they are not far from the actual conduct of social media platforms relying on their T&Cs to exploit user-generated content.¹⁵⁴

To be sure, Midjourney, like any other generative AI provider, has good reason to obtain some sort of non-exclusive license in submitted material. That is because the ordinary provision of a generative AI system will regularly implicate some of copyright’s exclusive rights. Consider the reproduction right. When a generative AI system is deployed in a web-accessible application, the system makes copies (in the copyright sense) of user inputs in order to generate outputs.¹⁵⁵ The generative AI provider would accordingly need authorization from the copyright owner in order to avoid infringement liability when it does the very thing the user is asking it to do—use the submitted content as the basis for a generation. It seems entirely sensible to seek such authorization explicitly from its users.¹⁵⁶

license to reproduce, prepare derivative works of, publicly display, publicly perform, sublicense, and distribute the Content You input into the Services, as well as any Assets produced by You through the Service. This license survives termination of this Agreement by any party, for any reason.” Midjourney’s T&Cs do not contain any provisions that purport to limit its ability to exercise its rights under this license.

¹⁵² Cf. Uri Y. Hacohen, Amit Elazari, & Talia Schwartz-Maor, *A Penny for their Creations—Appraising Users’ Value of Copyrights in their Social Media Content*, 36 BERKELEY TECH. L.J. 511, 515 (2021) (describing litigation between Facebook and one of its users over a Facebook affiliate’s use of the user’s profile picture in an advertisement).

¹⁵³ See *id.* at 516-17 (discussing Richard Prince’s infamous Frieze Art Fair exhibition featuring other Instagram users’ photographs).

¹⁵⁴ See Kim, *supra* note __, at 588-92 (describing controversies involving Facebook and Instagram’s use of photos uploaded by their users, in which the firms relied on their TOS to defeat copyright claims from users).

¹⁵⁵ See Katherine Lee, A. Feder Cooper, & James Grimmelman, *Talkin’ ‘Bout AI Generation: Copyright and the Generative AI Supply Chain*, 72 J. COPYRIGHT SOC’Y 251, 325 (2025) (explaining that “a generative AI system’s use of a user-inputted prompt” will often be a “‘reproduction’ within the meaning of copyright law”). Some models may be deployed such that they likely do not implicate the reproduction right—for Meta’s Llama family, for example, inference takes place on the user’s device, so the reproduction right would not be implicated for copyrights that the user herself owns. It’s also plausible that a generative AI system accessible through a web-based interface will replicate user input for such a short period that it does not meet the duration requirement and would therefore not constitute a copy in the copyright sense. See *Cartoon Network, LP v. CSC Holdings, Inc.*, 536 F.3d 121 (2d Cir. 2008) (holding that content held in a buffer for less than two seconds was not a copy that implicated the reproduction right).

¹⁵⁶ There are alternatives. Courts would probably conclude that users grant generative AI providers an implied license at least to use material they submit in order to generate the desired output. See *Effects Assocs. v. Cohen*, 908 F.2d 555, 558-59 (9th Cir. 1990) (concluding that an implicit license arose because the copyright owner “created a work at defendant’s request and handed it over, intending that defendant copy and distribute it”). Although most commonly invoked when the defendant requested the creation of the work, an implied license can arise whenever “the totality of the circumstances indicates an intent to grant” one. *Estate of Hevia v. Portrio Corp.* 602 F.3d 34, 41 (1st Cir. 2010); see also, Kim, *supra* note __, at 588 (explaining that sites relying on user-generated content need licenses from their users “to copy and distribute that content,” but that such licenses are implied from their users’ conduct). Still, there would be little to gain and much to lose from the provider’s perspective if they relied on implicit licenses when they are in position to obtain an explicit one. Generative AI providers might also try to rely on the DMCA’s safe harbors, especially the one regarding storage of material at the direction of the user. See 17 U.S.C. § 512(c) (immunizing a “service provider” from infringement liability arising “by reason of the storage at the direction of a user of material that resides on a system or network controlled or operated by

Still, that is no justification for the breadth of the licenses we document. Generative AI providers do not need the right to distribute copies of user-submitted content for any business purpose whatsoever simply in order to provide a user with the output she seeks. We also doubt that users are aware that they are effectively granting generative AI providers the ability to redeploy their copyrighted material without restriction.¹⁵⁷ These non-exclusive licenses thus extend far beyond what generative AI providers need and what we might reasonably think users would agree to.

Moreover, the stated ambition of leading generative AI providers is to develop products that are capable of exceeding human performance across a broad range of creative and intellectual domains.¹⁵⁸ If they succeed, it would mean that much or most of creative production would occur through generative AI services. And we doubt that most human users would willingly agree to let generative AI providers use their copyrighted, non-AI-generated works to displace or even replace those same human users. But that is precisely what many generative AI T&Cs permit.

2. On Control over System Outputs

A long-running scholarly debate has questioned whether copyright ought to attach to expressive works produced using artificial intelligence and, if so, in whom (or what) that copyright should vest.¹⁵⁹ Generative AI systems are (as yet) insensitive to the incentives that copyright provides; for many, then, the question is whether users or generative AI providers (or both) would respond to copyright's incentives in a way that justifies their social costs.¹⁶⁰

In that regard, there is some good news in our analysis of generative AI T&Cs. We worried that generative AI providers would moot this question by including in their T&Cs explicit assignments of user copyrights in system outputs.¹⁶¹ If so,

or for the service provider"). But the uncertain application of that provision to the generative AI context similarly makes it an "all-risk, no-reward" proposition for the generative AI provider, at least as compared to obtaining an explicit license from the user. See Lee, Cooper, & Grimmelmann, *supra* note __, at 361-65 (explaining how the Section 512 safe harbors apply at various points in the generative AI supply chain).

¹⁵⁷ Cf. Hachon, Elazari, & Schwartz-Maor, *supra* note __, at __ (reporting survey evidence indicating that users of social media platforms that rely on user-generated content are not aware of the scope of the licenses in those platforms' T&Cs).

¹⁵⁸ See Sam Altman, *The Gentle Singularity*, <https://blog.samaltman.com/the-gentle-singularity> (June 10, 2025) (claiming that "we have recently built systems that are smarter than people in many ways"); Dario Amodei, *The Adolescence of Technology*, <https://www.darioamodei.com/essay/the-adolescence-of-technology> (Jan. 2026) (asserting that "it cannot possibly be more than a few years before AI is better than humans at essentially everything").

¹⁵⁹ See *supra* note __.

¹⁶⁰ See Lemley, *Upside Down*, *supra* note __, at 28 (noting that it's "not obvious that generative AI needs the incentive of copyright law").

¹⁶¹ There is some uncertainty regarding the efficacy of such attempted assignments. Copyright law requires that a transfer of ownership be in a written instrument of conveyance, note, or memorandum that is signed by the owner. 17 U.S.C. § 204(a). Courts have held that online T&Cs can satisfy this requirement, but that position is not

then even if black-letter copyright doctrine vested copyrights in system outputs to users, as a practical matter, all of those copyrights would be owned by generative AI providers. Users would give up their interests unwittingly and without meaningful compensation. As it turns out, however, the vast majority of generative AI T&Cs include some language specifically indicating that the user will own whatever copyright interests result from their use of the system.¹⁶²

But we are not entirely out of the woods. As described *supra*, many generative AI T&Cs include provisions stating that they could be changed by the provider at any time.¹⁶³ If and when generative AI systems become involved in large-scale production of commercially valuable expressive works, we would not be surprised to see some such modifications. And even in their current form, the T&Cs typically undermine user interests in system outputs in much the same way that they undermine user interests in submitted material. Generative AI T&Cs frequently grant providers broad nonexclusive licenses to use system outputs.¹⁶⁴ Any user that expects to exert exclusive control in expressive works they produce using generative AI will be sorely disappointed.

We want to be precise here. We are not asserting that users ought to be the copyright owners of generative AI system outputs—we take no position on that question. The appropriate role of copyright in fostering creativity, supporting a democratic society, and enabling human flourishing has always been hotly contested.¹⁶⁵ And it may turn out that society would be best served if generative AI providers (or the systems themselves) were deemed copyright owners of works produced with the assistance of generative AI system. Or that no copyright should vest in such works at all.

Our point is that this debate ought to be resolved through the open, deliberative processes by which weighty questions are traditionally resolved in our representative democracy. Courts ought to have the opportunity to work through the issue in a common law manner, as they have done for centuries of technological changes that impact copyright law. Congress may ultimately conclude that amendments to the copyright statute are warranted, as it has previously done over centuries of technological changes. Private ordering

without controversy. See *Metro Reg'l Info. Sys. v. Am. Home Realty Net.*, 722 F.3d 591, 601-03 (4th Cir. 2013) (affirming preliminary injunction in favor of grantee notwithstanding alleged infringer's argument that grantor's "electronic agreement to [grantee's] TOU does not operate as an assignment" for failure to "satisfy Section 204's writing and signature requirements"); but see 3 NIMMER § 103(A)(1) (criticizing decisions for concluding that terms and conditions satisfy the signature requirement).

¹⁶² See *supra* text accompanying notes ____.

¹⁶³ See *supra* text accompanying notes ____.

¹⁶⁴ See *supra* text accompanying notes ____.

¹⁶⁵ See, e.g., JESSICA M. SILBEY, *AGAINST PROGRESS: INTELLECTUAL PROPERTY AND FUNDAMENTAL VALUES* (2022); MADHAVI SUNDER, *FROM GOODS TO A GOOD LIFE: INTELLECTUAL PROPERTY AND GLOBAL JUSTICE* (2012); Barton C. Beebe, Bleistein, *the Problem of Aesthetic Progress, and the Making of American Copyright Law*, 117 COLUM. L. REV. 319 (2017); Yochai Benkler, *Free as the Air to Common Use: First Amendment Constraints on Enclosure of the Public Domain*, 74 N.Y.U. L. REV. 354 (1999); Neil Weinstock Netanel, *Copyright and a Democratic Civil Society*, 106 YALE L.J. 283 (1996).

undoubtedly has a crucial role to play as well. But generative AI T&Cs are hardly worthy of that label—they are instead efforts to dictate by fiat in fine print who gets to control the unprecedented quantity of creative output that this technology enables.

3. On Bearing the Risks of Copyright Infringement

We turn now to the problem of infringement. As previously mentioned, high-profile lawsuits against generative AI firms ask whether model training constitutes copyright infringement.¹⁶⁶ We focus here on the next wave of litigation, which will grapple with who, if anyone, ought to be liable when generative AI systems produce artifacts that appear substantially similar to copyrighted material.

This issue presents difficult and unresolved legal questions about the proper allocation of liability when multiple actors with varying degrees of involvement participate in activities that, collectively, seem to tread upon a copyright owner's exclusive territory.¹⁶⁷ Similar questions arose at the dawn of the Internet era. In the 1990s, courts confronted cases involving copies of protected material that were made on devices owned by internet service providers as a result of user instructions.¹⁶⁸ ISPs were concerned that they could be held liable even when they were passive conduits for activity directed by their users.¹⁶⁹

In a landmark decision, the Northern District of California established the volitional conduct test for direct liability.¹⁷⁰ To summarize that test, a store that makes a photocopier available for its customers would not be directly liable if its customers use the photocopiers to make copies of copyrighted books; on the other

¹⁶⁶ See *supra* text accompanying notes ____.

¹⁶⁷ See *CoStar Grp., Inc. v. LoopNet, Inc.*, 373 F.3d 544, 550 (4th Cir. 2004) (summarizing the volitional conduct test as requiring “actual infringing conduct with a nexus sufficiently close and causal to the illegal copying that one could conclude that the machine owner himself trespassed on the exclusive domain of the copyright owner”); *Goodyear*, *supra* note ____; *Lee, Cooper, & Grimmelmann*, *supra* note ____, at 530-36; *Chatterjee & Fromer*, *supra* note ____.

¹⁶⁸ See e.g., *Sega Enter. Ltd. v. MAPHIA*, 857 F. Supp. 679 (N.D. Cal. 1994) (holding bulletin board liable for content uploaded by its users; *Religious Tech. Cent. v. Netcom Online Servs.*, 907 F. Supp. 1361, 1365-66 (N.D. Cal. Nov. 21, 1995) (concluding that ISP could only be held indirectly liable, not directly liable, for copies stored on its servers following instructions from its users).

¹⁶⁹ See *Terrica Carrington*, *Twenty Years of the DMCA: Notice and Takedown in Hindsight (Part I)*, COPYRIGHT ALLIANCE (Oct. 26, 2018), available at <https://copyrightalliance.org/dmca-notice-and-takedown/> (“Internet service providers and platforms (ISPs) were concerned about their own potential liability for copyright infringement committed by their users or subscribers.”)

¹⁷⁰ See *Religious Tech.*, 907 F. Supp. at 1365-66. Defendant Netcom was an internet service provider; defendant Dennis Erlich was a former Scientology minister; in response to Erlich's instructions, Netcom's servers made unauthorized copies of Scientology's texts. *Id.* at 1367-68. The court refused to impose direct liability on Netcom because its automatic response to Erlich's instructions meant that its “system [was] merely used to create a copy by a third party.” *Id.* at 1369-70.

hand, a store that makes the copies for the customers would be directly liable because it is actively choosing to comply with its customers' requests.¹⁷¹

Generative AI will put renewed pressure on the volitional conduct test because neither the user nor the generative AI provider has quite so direct a role in choosing content.¹⁷² A user's prompt elicits specific content only with some probability—this is not the same as a customer choosing to download copies of her favorite artist's new album. In some cases, the probability may be sufficiently high that she can be said to have chosen the output; in other cases, the user's prompt may have only a very attenuated connection to the output.¹⁷³ Similarly, the generative AI provider's choices in developing and deploying the system will sometimes be more and other times less closely connected to the output.¹⁷⁴ As a result, the appropriate allocation of liability between users of generative AI systems and their providers seems ripe for common law adjudication, enabling case-by-case assessment of the circumstances that ought to make one or another liable.

Our research again reveals that generative AI T&Cs might, as a practical matter, effectively moot this important question of copyright law and policy. Nearly half of the T&Cs include representations and warranties by which the user promises that their use of the system will not infringe third-party copyright interests.¹⁷⁵ Meanwhile, the overwhelming majority include indemnification provisions that would be triggered either by breach of those representations and warranties or, more generally, by the user's use of the generative AI system.¹⁷⁶ Whatever courts decide about where liability should fall in the first instance, T&Cs with expansive indemnification provisions triggered by broad representations and warranties mean that users will bear ultimate responsibility.

Making matters worse, users will typically have no idea that they made these promises or that their conduct might violate them. The scope of copyright law often diverges from the lay understanding of it.¹⁷⁷ Users may think they are doing nothing wrong when they elicit sequels to their favorite films from a generative AI service, but copyright law would beg to differ.¹⁷⁸

¹⁷¹ See *Cartoon Network LP v. CSC Holdings, Inc.*, 536 F.3d 121, 131-32 (2d Cir. 2008) (distinguishing copy shop cases in which "a company . . . merely makes photocopiers available to the public on its premises" and therefore do not satisfy the volitional conduct requirement from those in which a customer "mak[es] a request to a human employee, who then volitionally operates the copying system to make the copy").

¹⁷² See *Goodyear*, *supra* note __, at __ (arguing that unlike "with previous technologies that created infring[ing]" copies, "the developer and user of generative AI does not know the identity of the infringing work in advance").

¹⁷³ *Id.*

¹⁷⁴ *Id.*

¹⁷⁵ See *supra* text accompanying notes __.

¹⁷⁶ See *supra* text accompanying notes __.

¹⁷⁷ See Gregory N. Mandel, *The Public Perception of Intellectual Property*, 66 FLA. L. REV. 261 (2014) (reporting survey regarding how laypeople understand copyright law).

¹⁷⁸ See *Anderson v. Stallone*, 11 U.S.P.Q.2d 1161 (C.D. Cal. 1989) (holding that unauthorized script for a sequel to *Rocky III* infringed the derivative works right).

Moreover, generative AI firms themselves have encouraged users to think that they are insulated from copyright liability when they use generative AI services. For example, OpenAI attracted considerable favorable headlines for its promise to indemnify its users for copyright infringement.¹⁷⁹ OpenAI's "Copyright Shield," however, extends only to its enterprise and API users.¹⁸⁰ For most customers, reality is the mirror image of this carefully cultivated perception. Not only do generative AI providers fail to indemnify ordinary users; instead, generative AI T&Cs overwhelmingly demand that such users indemnify the providers.¹⁸¹

It's tempting to dismiss these concerns on the grounds that business considerations make it unlikely that generative AI providers would engage in a widespread litigation campaign that would effectively bankrupt their paying customers. If the natural consequence of demanding indemnification is that users will stop using generative AI systems at all, then we should expect that generative AI providers will not demand indemnification, even if the T&Cs give them the right to do so.

That appears correct, as far as it goes. But it does not go as far as it may seem. Similar considerations did not stop the Recording Industry Association of America from pursuing such a litigation strategy in the analogous context of peer-to-peer file sharing in the first decade of the 2000s.¹⁸² Furthermore, users may end up owing indemnification obligations to entities that would not be so circumspect. It seems plausible that at least some generative AI providers would become insolvent as a result of copyright infringement litigation against them; in a liquidation, the trustee must maximize the value of the bankruptcy estate for the benefit of its creditors.¹⁸³ That trustee may conclude that not only can it sue the provider's (former) users to recover the copyright liability the firm incurred as a result of their (seemingly innocuous) prompts, but that it has a fiduciary obligation to do so.¹⁸⁴

¹⁷⁹ See, e.g., Blake Montgomery, *OpenAI Offers to Pay for ChatGPT Customers' Copyright Lawsuits*, THE GUARDIAN (Nov. 6, 2023), available at: <https://www.theguardian.com/technology/2023/nov/06/openai-chatgpt-customers-copyright-lawsuits>.

¹⁸⁰ *Id.*

¹⁸¹ See *supra* text accompanying notes ____.

¹⁸² See Genan Zilkha, *The RIAA's Troubling Solution to File-Sharing*, 20 FORDHAM INTELL. PROP. MEDIA & ENT. L. J. 667, 668-89 (2009) (describing \$1.92 million verdict against an individual user of a file-sharing service for downloading 24 songs in 2005).

¹⁸³ See Richard A. Booth, *Stockholders, Stakeholders, and Bagholders (Or How Investor Diversification Affects Fiduciary Duty)*, 53 BUS. LAW. 429, 431 n.3 (1998) ("If the corporation becomes bankrupt, creditors can, through the trustee, pursue claims that the corporation would have had so as to maximize the assets which will be available to pay creditor claims.")

¹⁸⁴ See *In re Microwave Prod.*, 102 B.R. 666, 672 (Bankr. W.D. Tenn. 1989) (holding that debtor-in-possession breached fiduciary duty by failing to pursue claims against insider-shareholders); C.R. Bowles, Jr., & Nancy B. Rapoport, *Has the DIP's Attorney Become the Ultimate Creditors' Lawyer in Bankruptcy Reorganization Cases?*, 5 AM. BANKR. INST. L. REV. 47 (1997) (characterizing failure to disclose that debtor-in-possession refused to pursue claims as a breach of its fiduciary duty).

And the damages in a lawsuit for works associated with even a single user could easily be massive—copyright law includes statutory damages provisions that can reach \$30,000 per work even for innocent infringements of the sort we describe, and up to \$150,000 for willful infringement.¹⁸⁵ Recall that this liability would be imposed in the first instance in a lawsuit by the copyright owner against the generative AI provider—it is easy enough to imagine a jury reaching a \$300,000 verdict against the provider for ten infringing outputs. The user could then face a \$300,000 indemnification claim. These previously unnoticed provisions thus expose users to potentially crippling liability.

B. Individual- and Class-Level Harms from Dispute-Related Terms

Many already-existing legal regimes conceivably implicated by generative AI depend heavily on private enforcement for their efficacy. Antitrust, privacy, tort, and a host of other conduct-regulating laws are, in part or in whole, operationalized by private enforcement.¹⁸⁶ As a result, eliminating or substantially suppressing private enforcement can effectively limit the scope and dampen the impact of the underlying laws, thereby denying individuals and classes of similarly situated individuals a remedy for injuries caused or facilitated by generative AI tools.

Counterparties are an especially important source of private enforcement for these and many other U.S. legal regimes. As a direct trading partner, a counterparty typically has direct visibility into their supplier's behavior. And counterparties often bear the direct cost of harmful conduct.¹⁸⁷ Lawmakers and judges frequently craft legal institutional structures and rulings based on the assumption that private counterparties will be especially active enforcers, ready and able to vindicate the rights ostensibly created by generally applicable conduct-regulating laws.¹⁸⁸

When dispute-related generative T&Cs are legally binding—and they often are—they can have dramatic implications for counterparties' ability and incentive to vindicate those rights across a wide swath of conduct-regulating rules. And even where these terms are not formally enforceable, they can have significant in *terrorem* effects on injured parties.¹⁸⁹

A given user of the generative-AI tools covered by our survey is likely bound by multiple one-sided dispute-related terms across multiple categories. Suppose,

¹⁸⁵ 17 U.S.C. § 504(c) (establishing statutory damages for copyright infringement).

¹⁸⁶ See *supra* notes __ and accompanying text.

¹⁸⁷ To be sure, counterparties that act as intermediaries (e.g., distributors) can sometimes pass on the costs of harmful conduct to others.

¹⁸⁸ See, e.g., *Ill. Brick Co. v. Ill.*, 431 U.S. 720 (1977) (denying standing for downstream customers to bring private actions under the Sherman Act, reasoning that direct counterparties are best situated as private enforcers).

¹⁸⁹ Cf. Evan Starr et al., *The Behavioral Effects of (Unenforceable) Contracts*, 36 J. L. ECON. & ORG. 633 (2020) (empirically finding in *terrorem* effects of unenforceable noncompete clauses).

for example, a hypothetical user were bound only by the types of terms that more than half (50+%) of the T&Cs in our study deploy. That user would lose access to the court system, lose access to local adjudication, forfeit any claims for negligence, forfeit any claims for misrepresentation or breach-of-contract based on extracontractual promises, and be subject to the provider's chosen law in arbitration.¹⁹⁰ If the user could—despite all of that—somehow prevail, any damages award would likely be no more than \$100 and could not include punitive, consequential, or incidental damages.¹⁹¹ Throughout all of this, the provider could change the T&Cs at any time without first obtaining the user's consent.

Dispute-related generative AI T&Cs will likely have—indeed, are likely already having—a substantial, cumulative effect. Otherwise-meritorious cases will not be brought at all. The few cases that survive will face an uneven playing field with little prospect of meaningful recovery. In general, these T&Cs will markedly constrict the reach and lessen the effectiveness of a wide swath of substantive laws as to generative AI providers and their user-facing conduct.

To illustrate, consider the U.S. antitrust enterprise. Industry observers have already begun to sound the alarm over potentially anticompetitive conduct in generative AI markets.¹⁹² For example, three startup firms—Character AI, Inflection AI, and Adept AI—were at one point working to develop foundation models that would have competed against the models currently offered by Google, Microsoft, Amazon, and Meta, along with those incumbents' "investment partners" like OpenAI and Anthropic.¹⁹³ But each of the three startups then announced a deal with an incumbent: Google quasi-acquired Character AI, Microsoft quasi-acquired Inflection AI, and Amazon quasi-acquired Adept AI.¹⁹⁴ Soon after these announcements, each of the three target startups discontinued work on foundation models.¹⁹⁵

These quasi-mergers raise both important doctrinal questions and substantial antitrust concerns.¹⁹⁶ Suppose they caused users to miss out on additional choice and quality, thereby violating the federal Clayton Act.¹⁹⁷ But federal enforcement agencies have not, under current leadership, initiated any antitrust proceedings

¹⁹⁰ See *supra* Pt. III.B.1–.4.

¹⁹¹ *Id.* Pt. III.B.4.

¹⁹² See Lemley & Henderson, *supra* note __; FTC OFF. OF TECH., *supra* note __.

¹⁹³ Alexandros Kazimirov, *Mergers and Cooptive Acquisitions*, AM. ANTITRUST INST. Working Paper No. 25-01 (Apr. 3, 2025), <https://www.antitrustinstitute.org/wp-content/uploads/2025/04/AAI-Paper-Mergers-Cooptive-Acquisitions.pdf>.

¹⁹⁴ *Id.* at 19.

¹⁹⁵ *Id.*

¹⁹⁶ See *id.*; see also Colleen Cunningham, Florian Ederer & Song Ma, *Killer Acquisitions*, 129 J. POL. ECON. 649 (2021) (describing and empirically studying the phenomenon of incumbents acquiring and shutting down rivals).

¹⁹⁷ Section 7 of the Clayton Act prohibits mergers whose effect "may be substantially to lessen competition," 15 U.S.C. § 18, or, perhaps relevantly for present purposes, to "tend to create a monopoly." See Robert H. Lande, John M. Newman & Rebecca Kelly Slaughter, *The Forgotten Anti-Monopoly Law: The Second Half of Clayton Act Section 7*, 103 TEX. L. REV. 785, 831–32 (2025).

involving generative AI markets. State attorneys general face severe resource constraints, a perennial problem that worsens during periods of federal underenforcement. In theory, excluded rivals not subject to user-facing online T&Cs can bring antitrust lawsuits. But here, the would-be rivals voluntarily agreed to remunerative quasi-acquisitions, making them unlikely (and awkwardly situated) potential enforcers. That leaves user lawsuits, most likely class actions, as the lone fallback. Character,¹⁹⁸ Inflection,¹⁹⁹ and Adept,²⁰⁰ however, all include mandatory-arbitration and class-waiver provisions, among other one-sided terms, in their user-facing T&Cs. Private antitrust enforcement by users would thus face exceedingly high hurdles.

C. *System-Level Harms*

In addition to their short-run effects on users, generative AI T&Cs also carry significant, negative implications for the judicial and legislative development of doctrine to address this nascent technology. Many general, conduct-regulating laws are supposed to be “supple enough” to account for nascent technologies and practices.²⁰¹ Although these laws arose before the widespread adoption of both online T&Cs and generative AI applications,²⁰² they are generally supposed to be able to expand in scope to address new marketplace realities.

But generative AI T&Cs threaten to stunt and distort the common-law process of application and refinement for a wide range of general conduct-regulating laws. This is most obviously a problem for pure common-law causes of action. But it also implicates laws that are nominally statutory, where, in the U.S. legal tradition, accretive judicial gap-filling and ambiguity-resolving still play a critical role.²⁰³

Consider again the example of antitrust law. Congress enacted the primary federal antitrust statutes in 1890 and 1914. These statutes’ institutional design depends on private plaintiffs who, along with state attorneys general and federal agencies, form an institutional “three-legged stool” of enforcement.²⁰⁴ Despite a variety of judicially created anti-enforcement rules that now stand in their way,²⁰⁵ private enforcers still account for the vast majority of U.S. antitrust lawsuits filed

¹⁹⁸ Character AI was included in our study.

¹⁹⁹ *Terms of Service*, INFLECTION AI (updated Dec. 22, 2025), <https://inflection.ai/privacy-policy#terms-of-service>.

²⁰⁰ *Terms of Service*, ADEPT AI (updated Nov. 9, 2025) <https://adeptai.org/terms-of-service>.

²⁰¹ *Id.* at 111 (quoting Richard A. Posner, *Antitrust in the New Economy*, 68 ANTITRUST L.J. 925, 925 (2001) (internal quotation marks omitted)).

²⁰² Tort law, for example, has centuries-old roots. See Robert J. Kaczorowski, *The Common-Law Background of Nineteenth-Century Tort Law*, 51 OHIO ST. L.J. 1127, 1128 (1990) (noting that judges deciding tort cases in the 1800s built upon a tradition that extended back to the 1600s).

²⁰³ See generally Charles W. Tyler, *Common Law Statutes*, 99 NOTRE DAME L. REV. 669 (2023) (cataloguing statutes that are supposedly resistant to the standard U.S. statutory interpretation toolkit).

²⁰⁴ E.g., Michael Kades & Randy Stutz, *The Three-Legged Stool of U.S. Antitrust Enforcement*, AM. ANTITRUST INST. (Nov. 24, 2025), <https://www.antitrustinstitute.org/work-product/the-three-legged-stool-of-u-s-antitrust-enforcement-a-conversation-with-michael-kades/>.

²⁰⁵ *Brunswick Corp. v. Pueblo Bowl-O-Mat, Inc.*, 429 U.S. 477 (1979) (requiring separate proof of “antitrust injury”); *Ill. Brick Co. v. Ill.*, 431 U.S. 720 (1977) (requiring proof that plaintiffs were “direct purchasers” from the violator(s)).

each year.²⁰⁶ Most serious scholars agree that private enforcement has generally been an important, salutary part of the antitrust enterprise.²⁰⁷ Federal and state enforcers are relatively few in number and chronically underfunded. And public enforcement ebbs and flows with the prevailing political tides.

In situations like these, as we saw above, generative AI T&Cs can prevent potentially harmful conduct from being addressed. The ubiquity and force of these T&Cs thereby raise serious rule-of-law questions. The underlying conduct-regulating laws are, for the most part, understood to apply generally. They are supposed to “treat like things alike,”²⁰⁸ a fundamental, long-recognized rule-of-law principle. To the extent that generative AI tools cause conceptually similar harm as prior offline technologies, then, these laws are supposed to apply with equal force. But generative AI T&Cs can eliminate or substantially blunt their impact.

These T&Cs can also prevent public adjudication of important questions. Even if such claims were to reach court, they would likely be decided in the defendant’s preferred, home-field venue.²⁰⁹ Generative AI T&Cs can thus both stunt and distort the common-law process of developing doctrine for generative AI markets.

Dispute-related generative AI T&Cs can also sever an important causal mechanism between the judicial and legislative branches. Controversial decisions by the judiciary often play a causal role by prompting Congress to act. Continuing the antitrust example, a series of U.S. Supreme Court decisions publicly adopted a defendant-friendly approach under the Sherman Act in the early Twentieth Century. Congress passed the Clayton Act of 1914 as a direct response to those decisions.²¹⁰ Stepping back, judicial dispute resolution generally produces a publicly available decision that can prompt federal or state legislative action, as judges themselves are well aware.²¹¹ By preventing dispute resolution of some claims and shunting others into confidential arbitration, generative AI T&Cs can

²⁰⁶ See Crane, *supra* note 16.

²⁰⁷ See, e.g., Robert H. Lande & Joshua P. Davis, *Benefits from Private Enforcement: An Analysis of Forty Cases*, 42 U.S.F. L. REV. 879 (2008).

²⁰⁸ See John M. Newman, *Antitrust in Zero-Price Markets: Applications*, 94 WASH. U. L. REV. 49, 58 (2016) (quoting RICHARD A. POSNER, *THE PROBLEMS OF JURISPRUDENCE* 42 (1990) (internal quotation marks omitted)); Richard Abel, *General Damages Are Incoherent, Incalculable, Incommensurable, and Inegalitarian (but Otherwise a Great Idea)*, 55 DEPAUL L. REV. 253, 302 (2006) (“Law claims to treat like things alike—one of its most fundamental reasons for existing.”).

²⁰⁹ Using these three firms as examples, Character AI (included in our study) specifies state and federal courts “located within Santa Clara County, California,” where Character AI is headquartered. Inflection AI similarly requires dispute resolution on its home turf. See *Terms of Service*, INFLECTION AI, *supra* note __ (including an exclusive-jurisdiction provision identifying state and federal courts located nearby Inflection AI’s Silicon Valley headquarters). Amusingly, Adept AI appears to have uploaded a form set of terms without filling in its preferences. *Terms of Service*, ADEPT AI, *supra* note __ (“These Terms are governed by the laws of [State/Country] Any disputes arising from these Terms shall be resolved through binding arbitration The arbitration shall take place in [Location].” (all brackets and bracketed language in original)).

²¹⁰ John M. Newman, *Lawless Antitrust*, 94 FORDHAM L. REV. 1597, 1601 (2026).

²¹¹ See, e.g., Lori Hausegger & Lawrence Baum, *Inviting Congressional Action: A Study of Supreme Court Motivations in Statutory Interpretation*, 43 AM. J. POL. SCI. 162 (1999) (surveying examples of U.S. Supreme Court opinions that expressly invited congressional action on the issues before the Court).

sever that historically important link. Bad judicial decisions are not optimal. But in the long run, resolving disputes in secret—or not at all—can be even worse.

V. POLICY REFORMS

Our findings suggest that generative AI T&Cs raise acute policy concerns. The terms analyzed in our study are in tension with not only traditional contract-law ideals,²¹² but also the goals and principles underlying a wide variety of legal institutional designs, at an important moment in the development trajectory of the underlying technology. Generative AI T&Cs can let AI providers erode democratically enacted statutory rights, stunt the common-law development of rules and standards to govern this nascent technology, and prevent problems from surfacing in public judicial opinions, thereby severing an important link between the judicial and legislative branches.

None of this is inevitable. Innovation is rapidly yielding new products and tools, but copyright, contract, tort, privacy, antitrust, and a swath of other laws need not be passively dragged along in their old forms.²¹³ Legal policies permitted—and indeed incentivized—the current state of affairs.²¹⁴ New policies could steer a different course.²¹⁵ In the discussion that follows, we identify a range of potential policy responses to one-sided generative AI T&Cs. We caution against relying on enhanced “consent,” proposing instead that policymakers at minimum void certain terms and—particularly when generative AI providers deploy cumulative, compounding terms—penalize their use to prevent the *in terrorem* effects that even void terms can cause.

A. *The Limited Potential of Enhanced “Consent” Regimes*

Potential policy reforms in this area fall along a spectrum, from less forceful to more robust. On the weaker end lie moves like requiring more meaningful consent than U.S. contract law currently requires. By way of analogous example, the Fair Credit Reporting Act imposes default privacy requirements for some types of information uses, yet allows individuals to consent to those uses while imposing a relatively high threshold for legal “consent.”²¹⁶ Somewhat more forceful are opt-in-type mechanisms that permit voiding or avoiding some terms, albeit while requiring affirmative action by the empowered individuals to do so. Some U.S.

²¹² See, e.g., Brett Frischmann & Moshe Y. Vardi, *Better Digital Contracts with Prosocial Friction-in-Design*, 65 JURIMETRICS 1, 3–4 (2025).

²¹³ *Id.* at 4 (“The current state of affairs in digital contracting is neither natural nor inevitable.”).

²¹⁴ Cf. John M. Newman, *Attention Capitalism: The Law and Political Economy of Attention Markets*, 78 STAN. L. REV. 415 (2026) (identifying ways in which contingent legal choices have channeled activity into digital attention markets).

²¹⁵ See generally Jedediah Britton-Purdy et al., *Building a Law-and-Political-Economy Framework: Beyond the Twentieth-Century Synthesis*, 129 YALE L.J. 1784 (2020) (describing the recently revitalized law-and-political-economy school of legal thought).

²¹⁶ Nikolas Guggenberger, *Consent as Friction*, 66 B.C. L. REV. 353, 413–14 (2025).

state laws let users affirmatively act to limit particular uses of information, for example.²¹⁷ These mechanisms take a variety of forms, from testing for user understanding to ensure (actual) informed consent to offering menus of options.²¹⁸

Past research shows that disclosure-and-consent requirements are generally of dubious effectiveness.²¹⁹ Even when imposed with good intentions, they can fail to help—and can affirmatively harm—their intended beneficiaries.²²⁰ They also risk normalizing the underlying phenomena that raised concerns in the first place.²²¹ In fact, by creating a clear path to enforceability, these requirements may affirmatively incentivize more firms to deploy the harmful terms in question.²²² Consequently, policymakers should be wary of potential responses that fall short of, at minimum, voiding problematic terms.

B. Voiding Problematic Terms

Toward the more robust end of the policy spectrum are outright prohibitions that void the covered terms, rendering them legally unenforceable. Voiding terms can sidestep the use of increasingly questionable notice-and-consent regimes,²²³ creating protection without requiring any action by affected individuals.²²⁴

The indemnification provisions we document reach quite broadly. Some individual state courts might refuse to enforce them as unconscionable or against public policy.²²⁵ But federal legislative action voiding them is likely warranted to guard against the prospect of individual users facing potentially ruinous, one-on-one litigation against powerful firms, as was true in the peer-to-peer file-sharing

²¹⁷ *Id.* at 414; cf. David A. Hoffman, *Defeating the Empire of Forms*, 109 VA. L. REV. 1367 (2023) (proposing state laws deny enforceability of written forms for consumer transactions under \$100).

²¹⁸ Frischmann & Vardi, *supra* note __, at 11 (pointing to the EU's GDPR requirements for cookie consent options); see also Mark A. Lemley, *The Benefit of the Bargain*, 2023 WIS. L. REV. 237, 268 (proposing that a party seeking to enforce a contractual provision that changes the otherwise-applicable default rule must demonstrate that "both parties had a choice and actually agreed to those terms").

²¹⁹ See generally Omri Ben-Shahar & Carl E. Schneider, *The Failure of Mandated Disclosure*, 159 U. PA. L. REV. 647 (2011).

²²⁰ See, e.g., George Loewenstein et al., *Disclosure: Psychology Changes Everything*, 6 ANN. REV. ECON. 391, 415 (2014).

²²¹ Tess Wilkinson-Ryan, *The Perverse Consequences of Disclosing Standard Terms*, 103 CORNELL L. REV. 117, 165 (2017) (finding that disclosure of unenforceable T&Cs increased test subjects' belief that those terms were enforceable and arguing that increased disclosure could therefore increase firms' incentive to deploy such terms).

²²² See Mark A. Lemley, *Protecting Consumers in a Post-Consent World*, 77 STAN. L. REV. ONLINE 247, 248 (2025) ("Characterizing privacy as a right that parties can contract away has simply meant that parties essentially always contract it away.").

²²³ See, e.g., *id.* at 247.

²²⁴ Kim, *supra* note __, at 600 (proposing that "TOS that reallocate rights over mass new technologies such as consumer generative AI, be considered void").

²²⁵ See Larry A. DiMatteo, Bruce Louis Rich, *A Consent Theory of Unconscionability: An Empirical Study of Law in Action*, 33 FSU L. REV. 1067, 1077 (2006) ("The fact that an unconscionable clause was inserted in a merchant form contract in a merchant-consumer transaction has been a common factor in successful unconscionability cases."); Russell Korobkin, *Bounded Rationality, Standard Form Contracts, and Unconscionability*, 70 U. CHI. L. REV. 1203, 1260 (2003) ("Courts frequently allege that unequal bargaining power resulting from a seller's market power allows the seller to force onerous contract terms on buyers and thus supports a finding of procedural unconscionability.").

era. And shifting the liability risk arising from system outputs back to generative AI firms could have salutary political economy effects. It was the prospect of such liability that led ISPs to seek a legislative compromise with copyright content owners; that dynamic produced the Digital Millennium Copyright Act.²²⁶ If generative AI firms can continue shunting some of their copyright liability risk onto their users, it is doubtful that widely dispersed users will be able to have a comparable voice in any broader legislative response to this technology.

Legislative action at the federal or state level is one potential mechanism for prohibiting one-sided provisions in generative AI T&Cs. The U.S. Congress has enacted several statutes that limit firms' ability to impose one-sided terms in other settings.²²⁷ More recently, Senators Blumenthal and Hawley introduced a bipartisan bill that would void some generative AI T&Cs as to a newly created federal tort claim for "misuse of covered data."²²⁸

The inclusion of this voiding provision is encouraging, inasmuch as it indicates some awareness among lawmakers that T&Cs can pose a hurdle to effective, statutorily created private causes of action. At the same time, however, the voiding provision extends only to mandatory arbitration clauses and class waivers.²²⁹

That said, the voiding provision covers only mandatory-arbitration and class-waiver clauses. Our analysis suggests that these are not the only generative AI T&Cs with potentially significant, one-sided effects on users. For example, a clear majority of T&Cs impose strict limits on monetary damages, eliminating several damages categories and liquidating total award amounts at \$100, \$50, or even \$0.²³⁰ If legally valid, these terms could substantially undercut the effectiveness of legislative efforts like the "misuse of covered data" tort-claim bill. Our findings militate in favor of broadening the voiding lens to include not only arbitration and class-waiver terms, but also other one-sided terms identified in our study, particularly when used cumulatively.

To the extent federal statutory action is unlikely, reform could—and may well be more likely to—take place at the state level. U.S. states have already been actively legislating across a wide range of issues raised by generative AI.²³¹ At least

²²⁶ We recognize that the political economy of the DMCA was not an entirely happy story. See Pamela Samuelson, *Intellectual Property and the Digital Economy: Why the Anti-Circumvention Provisions of the DMCA Need to Be Revised*, 14 BERKELEY TECH. L.J. 519 (1999) (explaining why "the battle in Congress over the anti-circumvention provisions of the DMCA was a battle between Hollywood and Silicon Valley" that did not produce a desirable outcome and advocating for reform). Our point is merely that leaving users to fend for themselves is unlikely to produce better legislative results.

²²⁷ See, e.g., 10 U.S.C. § 987(e)(2) (prohibiting waivers of rights in contracts with covered members of the armed services); 15 U.S.C. § 45b (prohibiting firms from limiting consumers' ability to write negative product reviews).

²²⁸ S. 2367, 119th Cong., 1st Session (July 21, 2025).

²²⁹ S. 2367, 119th Cong., 1st Session (July 21, 2025).

²³⁰ See *supra* notes __ and accompanying text.

²³¹ See Scott Babwah Brennen et al., CTR. ON TECH. POL'Y, STATE TECHNOLOGY POLICY REPORT (2025), https://techpolicy.nyu.org/wp-content/uploads/2026/02/CTP_Final_State-Report2025.pdf (surveying state legislative activity for calendar-year 2025).

some of these laws create or affect private rights of action.²³² And multiple U.S. states have recently taken steps toward general legislative reform in other areas, citing shortcomings in corollary federal statutes and a perceived low probability of federal legislative reform as motivating factors.²³³ Including in state statutes like these voiding provisions on one-sided generative AI T&Cs would allow private enforcers to act as a force multiplier for resource-constrained state attorneys general, helping to ensure that these legislative efforts are effective and meaningful.

Federal and some state administrative agencies also have at their disposal various tools that could both prohibit the use of these terms. For example, the Federal Trade Commission has sought to prohibit specific types of terms under both its “unfair or deceptive acts or practices”²³⁴ and (somewhat more controversially) its “unfair methods of competition”²³⁵ rulemaking authority. The Consumer Financial Protection Bureau (“CFPB”) has issued a notice of proposed rulemaking that would prohibit firms from enforcing one-sided unilateral-amendment terms, terms that purport to waive statutory consumer-protection provisions and terms that purport to limit rights to free expression in the form of (for example) providing negative product reviews.²³⁶ While we see value in generally applicable prohibitions of one-sided generative AI T&Cs, doctrinal and/or industry area-by-area reform like the “misuse of covered data” bill or the analogous CFPB rulemaking may realistically be a more likely avenue for change.

C. *Penalizing the Cumulative Use of Unenforceable Terms*

Past research indicates that even legally void T&Cs can proliferate and cause negative real-world effects. Consequently, policymakers who wish to effectively regulate generative AI tools should consider reforms that go beyond passively voiding undesirable terms. Statutes and regulations that impose affirmative penalties on the use of covered T&Cs lie at the forceful end of the reform spectrum. But particularly when firms supplying generative AI tools deploy online T&Cs with cumulative and compounding effects, such a forceful response may be warranted.

²³² See, e.g., *id.* 32 (noting that California enacted an enhanced cause of action for AI-generated CSAM).

²³³ See, e.g., Justin Wise & Zach Williams, *New York Legislators Push for Antitrust Overhaul, Echoing Biden*, BLOOMBERG LAW, Jan. 27, 2025, <https://news.bloomberglaw.com/antitrust/new-york-legislators-push-for-antitrust-overhaul-echoing-biden>; Titus Wu, *California Policymakers Take Steps to Revamp State Antitrust Law*, BLOOMBERG LAW, Jan. 23, 2025, <https://news.bloomberglaw.com/bloomberg-government-news/california-policymakers-take-steps-to-revamp-state-antitrust-law>.

²³⁴ 49 C.F.R. Pt. 7740 (1984) (prohibiting lenders from, among other things, imposing obligations on consumers that are or contain a “cognovit or confession of judgment . . . , warrant of attorney, or other waiver of the right to notice and the opportunity to be heard”).

²³⁵ 16 C.F.R. Pt. 910 (2024) (prohibiting non-compete restrictions in employment contracts).

²³⁶ 12 C.F.R. Pt. 1027 (Jan. 14, 2025).

Empirical research on other potentially problematic terms finds that firms may continue deploying them despite their legally void status. For example, California has statutorily voided noncompete restrictions in employment contracts since 1872.²³⁷ Nonetheless, California employers continued to insert noncompete restrictions into employment contracts, handbooks, and the like.²³⁸ And these restrictions—though nominally void—continue to exert *in terrorem* effects on workers.²³⁹ Everyday people are not legal experts, and, particularly if a firm brings a term to their attention, may believe it is valid, or at least have enough doubt about its status to affect their conduct.

By adding penalties to disincentivize firms from deploying unenforceable terms—monetary fines, a presumption of liability, or both—laws and rules can disincentivize firms from taking advantage of this dynamic. Continuing the above example, some state statutes impose per-worker financial penalties on firms that deploy legally void noncompete restrictions.²⁴⁰

Whatever the precise mechanism or avenue for change, our findings suggest the need to take a holistic approach to generative AI T&Cs. Voiding or even penalizing the use of certain types of provisions will not level the playing field if others remain available. To illustrate, suppose a state legislature were to ban class-action waivers. Generative AI T&Cs could still effectively keep many, perhaps most, injured users from pursuing redress by eliminating a wide variety of tort and contract claims, capping damages at prohibitively low amounts, and requiring adjudication in the provider's home court without a jury. Moreover, such a ban would do virtually nothing to address the one-sided copyright licensing and risk-allocation terms that appear in many of the generative AI T&Cs in our study.

As in other industry contexts, the effects of one-sided generative AI T&Cs tend to be cumulative. For example, a choice-of-venue/exclusive-jurisdiction clause on its own has less effect than when coupled with a class-action waiver. And even the two terms together have less effect than when a damages cap is imposed. Every set of generative AI T&Cs examined in our study deploys multiple one-sided terms. Thus, for policymakers wishing to effectively govern generative AI markets, the use of cumulatively forceful T&Cs may be an appropriate target for the strongest form of policy intervention—both voiding the terms themselves and penalizing drafters that continue deploying them.²⁴¹

²³⁷ Catherine L. Fisk, *Reflections on the New Psychological Contract and the Ownership of Human Capital*, 34 CONN. L. REV. 765, 782 (2002).

²³⁸ *Id.* at 782–83.

²³⁹ See Evan Starr, J.J. Prescott & Norman Bishara, *The Behavioral Effects of (Unenforceable) Contracts*, 36 J. L. ECON. & ORG. 633, 633 (2020).

²⁴⁰ Amir El-Aswad, *Noncompete Agreements in 2026: A Federal and State Overview*, JD SUPRA, July 14, 2026, <https://www.jdsupra.com/legalnews/noncompete-agreements-in-2026-a-federal-9408064/>.

²⁴¹ Cf. CAL. BUS. & PROF'L CODE § 16600 (voiding contractual noncompete clauses); CAL. LABOR CODE § 432.5 (prohibiting employers from knowingly requiring employees to agree to an unenforceable term or condition); *id.* §§ 23, 433 (making violations of the California Labor Code punishable by a fine of up to \$1,000, six months in prison, or both).

Our aim here is not to conclusively resolve the question of how to address one-sided generative AI T&Cs. Again, this Article's primary contribution is empirical, not prescriptive. That said, our findings, viewed in light of past experience, suggest the need for a policy response, and one that falls toward the more robust end of the spectrum. Our research does raise as many question as it answers—perhaps more. But in light of the likelihood for widespread societal harms described above, waiting to act until all available data is gathered, and all policy debates resolved, would be a mistake.

VI. CONCLUSION

Any societal benefits from generative AI appear likely to come with substantial strings attached. Every generative AI tool included in our study deploys some version of online terms and conditions. Though seldom read or understood by users, most of the generative AI terms we analyzed transfer away valuable rights, allocate potentially enormous liability risks onto users, and prevent users from vindicating a wide range of long-established legal rights, all while allowing providers to change the rules of the game at any point with little or no notice. Any disputes that do manage to reach formal adjudication under such terms will face an uneven playing field and the prospect of minimal remedies. Generative AI terms and conditions threaten to deprive individuals of redress for real-world harms, and to distort the development of doctrine to govern this critically important technology. But none of this is inevitable. Laws currently give these generative AI terms and conditions their force; different laws and rules can tame their excesses.