AI ARTISTS ON THE STAND: BIAS AGAINST ARTIFICIAL INTELLIGENCE-GENERATED WORKS IN COPYRIGHT LAW

Joseph J. Avery* & W. Michael Schuster*

Generative artificial intelligence (AI) is revolutionizing the creation of art, literature, and music—challenging the boundaries of intellectual property law. To date, scholars have primarily focused on AI's authorship/entity status and the regulation of its use, overshadowing a critical issue: how AI's involvement in creative processes influences legal judgments in copyright disputes. Our empirical research reveals systemic bias against AI-generated works in such legal matters. In our studies, participants read about a company that had hired either a human designer (condition one) or a generative AI art system (condition two) to produce works of art, and those works of art arguably infringed an existing copyright. Everything except for the identity of the hired creator (human vs. AI) was held constant, including the works of art: Participants saw identical works. The results showed that, when the works were produced by the AI (vs. the human), participants' perception and behavior radically shifted. They were more likely to recommend that a copyright suit be commenced, more likely to find substantial similarity between the original and the allegedly infringing work, and more likely to rule in favor of infringement. Importantly, we do not argue that this bias against AI is bad or immoral or unjust. Rather, we merely identify the fact that, when a creator is an AI, people perceive the creation differently. There is a perceptual bias, and more importantly, this perceptual bias impacts legal outcomes. From this foundation, from this advance in basic science, a number of normative claims may emerge. For one, it is reasonable to assume that human actors, such as corporations, that make use of AI creators, face increased legal risk. As an extension, we delve into copyright law's primary objective—to promote the creation of new works—and we show that biases against generative AI may frustrate this objective, impeding the very creativity copyright law aims to incentivize.

---

* Denotes Equal Contributions

INTRODUCTION

The widespread adoption of new forms of artificial intelligence (AI)—notably generative AI such as ChatGPT and Dall-E—has precipitated significant legal developments. In fact, "significant" is an understatement. Legislative bodies in many U.S. states and foreign countries are focused on such technologies.[1] Legal scholars are also focused on such technologies, with scholarship in this area largely focused on authorship—whether generative AI should qualify as independent authors—and regulation of these artificial entities and their creators, owners, and users. However, amidst this fevered discourse, one aspect remains conspicuously understudied. Every day, countless companies and individuals leverage AI to create new works of authorship, and these works are increasingly facing copyright lawsuits, but nothing is known about (not a single study has been run regarding) how the presence of AI affects outcomes in these matters. Do technically irrelevant characteristics—the identity of the creator as human or machine—sway legal judgments? Numerous scholars have examined how plaintiffs or defendants' demographic characteristics may impact legal outcomes.[2] Thus, the emergence of AI as a new type of actor propels us into related but uncharted territory: It necessitates scholarly exploration of how the presence of AI may distort legal

1. Christian Auty, Amy de La Lama & Goli Mahdavi, *US State-by-State AI Legislation Snapshot*, JDSUPRA (Sept. 18, 2023), https://www.jdsupra.com/legalnews/us-state-by-state-ai-legislation7314466/ [perma.cc/5J3Y-TAVR]; Jeffrey D. Neuburger, Karla Grossenbacher, Avi Gesser, Megan K. Bannigan & Robert H. Newman, *ChatGPT and Generative AI: Key Legal Issues*, REUTERS (June 2023), https://www.reuters.com/practical-law-the-journal/transactional/chatgpt-generative-ai-key-legal-issues-2023-06-01&#8203 [perma.cc/CWM8-M8VP].

2. *See* BIAS IN THE LAW: A DEFINITIVE LOOK AT RACIAL PREJUDICE IN THE U.S. CRIMINAL JUSTICE SYSTEM (Joseph Avery & Joel Cooper eds., Lexington Books, 2020).

determinations. AI-centric biases, in other words, would raise concerns where the use of AI is legally irrelevant.

As an example, imagine that a company offloads a task onto an AI (rather than a human), and that AI makes a mistake. How is this perceived by legal adjudicators? Is the harm seen as greater (than if a human had caused it)? Is the behavior deemed more negligent and is the company overpenalized due to its use of AI? Consider the following.

In 2018, at an Amazon warehouse in New Jersey, a can of bear repellent (a spray containing capsaicin, a compound found in chili peppers) fell off a shelf, and a warehouse robot inadvertently punctured it.[3] The robot was fine, but the noxious fumes caused twenty-four human workers to be hospitalized "as a precaution," and all were apparently released quickly and in good health.[4] For context, Amazon employs hundreds of thousands of warehouse workers,[5] and the company reports close to 40,000 injuries in some years.[6] Moreover, just a few months earlier an identical accident occurred at an Amazon warehouse in Indiana—only in that case it was a human who caused the bear repellent to rupture.[7] The human mistake received little attention, but the AI mistake generated international media headlines: "a perfect horror story."[8] While the media is incentivized to be sensational, why sensationalize only the AI error? Why is Amazon vilified for the robot's mistake but not for the equivalent human mistake?

Our inquiry delves into this unexplored domain, shedding light on potential differential treatment of AI and human actors. This Article contains the first empirical analysis of systematic bias in copyright legal outcomes owing to the AI vs. human identity of the relevant actor.

Copyright offers a fertile ground for this examination due to its centrality in ongoing litigation against AI enterprises, including the headline-dominating *N.Y. Times* lawsuit,[9] and the potential for biases to be introduced during the infringement analysis. Specifically, the scope of copyright protection is a function of factual questions such as whether a new work is substantially similar to the original or

---

3. Jasper Jolly, *Amazon Robot Sets Off Bear Repellant, Putting 24 Workers in Hospital*, THE GUARDIAN (Dec. 6, 2018, 8:00 AM), https://www.theguardian.com/technology/2018/dec/06/24-us-amazon-workers-hospitalised-after-robot-sets-off-bear-repellent [perma.cc/9BT3-63XK].

4. *Id.*

5. Jason Del Rey, *How Robots are Transforming Amazon Warehouse Jobs—For Better and Worse*, VOX (Dec. 11, 2019, 5:00 AM), https://www.vox.com/recode/2019/12/11/20982652/robotsamazon-warehouse-jobs-automation [perma.cc/GGN9-6SRK].

6. Annie Palmer, *Amazon Warehouse Workers Suffer Serious Injuries at Twice the Rate of Rivals, Study Finds*, CNBC (Apr. 12, 2023, 3:47 PM), https://www.cnbc.com/2023/04/12/study-amazonworkers-seriously-hurt-at-twice-rate-of-otherwarehouses.html#:~:text=Amazon's%20serious%20injury%20rate%20fell,hiring%20push%20during%20the%20pandemic [perma.cc/BU82-VWF3].

7. Louise Matsakis, *This Wasn't Even Amazon's First Bear Repellent Accident*, WIRED (Dec. 6, 2018, 7:36 PM), https://www.wired.com/story/amazon-first-bear-repellent-accident/ [perma.cc/3VNJSUHC].

8. *Id.*

9. *See* Gary Marcus (@GaryMarcus), X (Dec. 29, 2023, 3:36 PM), https://x.com/GaryMarcus/status/1740834294803902807 [perma.cc/KF8Y-LG5J] ("OpenAI is in a heap of trouble, and it's not just text."); 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 [perma.cc/V5BQ-6KXR]. *See also* Blake Brittain, *More Writers Sue OpenAI for Copyright Infringement over AI Training*, REUTERS (Sept. 11, 2023, 12:33 PM), https://www.reuters.com/technology/more-writers-sue-openai-copyrightinfringement-over-ai-training-2023-09-11/ [perma.cc/5XVP-LHTC].

whether it merely adopted an earlier style. In a series of studies, we empirically investigate whether lay assessments of these and related queries vary contingent on the putatively infringing work being created by a human or a generative AI. In doing so, we make two contributions to the literature, one theoretical and one applied.

Our theoretical contribution extends the literature on human perception of activities undertaken by non-human actors to include AI output and copyright litigation. To this point, we empirically investigate lay assessments of (i) the quality of AI-generated artwork (as compared to the same work from a human); (ii) the fairness of using AI to create works building on earlier expressions; and (iii) whether an infringement lawsuit should be filed and whether, if filed, it should succeed.

Using the tools of experimental jurisprudence, we find the following. First, lay assessments of artwork vary depending on whether participants believe it was created by a human or by AI. Specifically, respondents found the exact same artwork to be more beautiful and creative when they believed it was created by a human as opposed to an AI. Second, and more importantly, participants were more likely to determine that AI-generated works infringed earlier copyrights. This seems puzzling given that both groups saw the same art and that there are no legally relevant differences in the situations. These findings are not, however, wholly surprising when viewed in light of the prior literature.

Our project builds from earlier work on variations in human perception depending on if the action is performed by a human or an AI. When identical behaviors, including deficient behavior, is observed in humans vs. machines, there is significant disparity in subsequent beliefs about those actors and their capabilities.[10] For creative endeavors this is also true, as a research team found variations in the assessment of art depending on whether respondents believed the work was generated by AI or a human.[11]

Building from this literature and the current empirical results, our analysis then turns to application—asking how the observed biases can influence law and policy. On this point, we start from prior research describing how lay biases create distortionary effects in the copyright regime. For example, non-financial-motivations can bring about infringement lawsuits that may distort copyright's goal of encouraging the creation of new works.[12] Moreover, jurors' moral intuitions may bias their infringement determinations and unintentionally create disincentives for future creators.[13] Both of these results suggest—and prior researchers explicitly recognize—the need for further study exploring the nature and scope of such biases.[14] These calls motivate the present Article.

---

10. S. Mo Jones-Jang & Yong Jin Park, *How Do People React to AI Failure? Automation Bias, Algorithmic Aversion, and Perceived Controllability*, 28 J. COMPUTER-MEDIATED COMMUNICATION 1, 1–2; 6–7 (2022); Berkeley J. Dietvorst, Joseph P. Simmons & Cade Massey, *Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err*, 144 J. OF EXPERIMENTAL PSYCH.: GEN. 114, 114–26 (2015).

11. Martin Ragot, Nicolas Martin & Salomé Cojean, *AI-Generated vs. Human Artworks. A Perception Bias Towards Artificial Intelligence?*, CHI '20: CHI CONFERENCE ON HUMAN FACTORS IN COMPUTING SYSTEMS (Apr. 2020).

12. Christopher Buccafusco & David Fagundes, *The Moral Psychology of Copyright Infringement*, 100 MINN. L. REV. 2433, 2435 (2016).

13. Zahr K. Said, *Jury-Related Errors in Copyright*, 98 IND. L.J. 749, 777 (2023).

14. Buccafusco & Fagundes, *supra* note 12, at 2507 ("One thing we suspect all readers will agree with is the need for further empirical study in this area. It would be valuable to know precisely what

Through experimental research, we unearth a startling insight: If an organization uses AI to generate works of art, its legal exposure rises in the copyright realm. When AI (vs. a human) created the putatively infringing work, our participants were more likely to recommend suing the company that used the AI, more likely to find the two works (the original and the allegedly infringing work) substantially similar, more likely to believe the new work was infringing, and more likely to impose harsher sanctions. In addition, there were even signs of perceptual distortions: Participants found the exact same artwork to be more beautiful when they believed it was created by a human, as opposed to AI.

In conjunction with these findings, we set forth our application-relevant contribution to the literature. Specifically, we describe how this biased application of the law can undermine the primary goal of copyright: encouraging the creation of new works of authorship (e.g., art). Our findings support expectations that AI-generated works are disproportionately likely to face copyright infringement lawsuits (regardless of their merit) and are disproportionately likely to lose and face harsher penalties during those lawsuits. The paper concludes by describing how these biases discourage the creation of new works of art using nascent technologies such as generative AI.

Importantly, in this Article, we are not arguing that this bias against AI is bad or immoral or unjust. We are not making a moral or normative claim. Rather, we are merely identifying the fact that, when a creator is an AI, people perceive the creation differently. There is a perceptual bias. And, most importantly, this perceptual bias leads to bias in legal outcomes, namely, the bias in copyright infringement outcomes that we explored in our experiments. From this foundation, this advance in basic science, a number of normative claims may emerge. For one, it is reasonable to assume that human actors, such as corporations, that make use of AI creators face increased legal risk from such use. As an extension, it also is reasonable to assume that this bias, in shifting the expected value propositions of creation, will work to distort the incentives of the current copyright regime.

Bias Against Artificial Intelligence-Generated Works in Copyright Law proceeds in three sections. The first reviews the literature discussing human perceptions and AI, the influence that legally irrelevant events (such as the use of AI in creating artwork) may have on legal outcomes, and the ways that such biases may undermine the goals of copyright. Section II presents our empirical hypotheses and empirical work, both the experimental designs and the findings. The third Section discusses implications of our findings and future research. In all, this Article takes a nascent yet crucial academic step, probing the implications of AI's distinctive actor-status on legal adjudications, with a particular focus on copyright law. Through a rigorous examination, we begin to unravel the complexities intertwined between AI, human actors, and the copyright regime, contributing an essential foundation for current and future scholarship and litigation.

---

sorts of activities trigger owners' moral foundations and how widely distributed those responses are."); Said, *supra* note 13, at 804 ("Experiments and empirical research designed to test lay, legal, and judicial experiences of copyright law would be helpful . . . . More experimental work on biases and juror comprehension would be invaluable.").

I. COPYRIGHT LAW AND PERCEPTION OF LEGAL ACTORS

Copyright law in the United States is utilitarian in nature.[15] To this end, limited monopoly rights are granted to authors to incentivize production of creative works.[16] These rights are intended to create a quid pro quo whereby authors can reap financial rewards associated with their work, and in turn, the public can consume the new work and will receive unfettered use once the work enters the public domain.[17]

However, this monetary incentive structure becomes unmoored from its utilitarian goals when copyright lawsuits are brought for purposes other than financial harm. Such lawsuits do not serve copyright's goal of ensuring that authors can vindicate harms to their financial interests. Indeed, those lawsuits actually *discourage* defendants and related parties from creating new works (from fear of getting sued). As discussed herein, we posit that lawsuits of this nature may be brought against users of generative AI due to biases against the technology. This can, in turn, undermine copyright's goal of encouraging the creation of new works (including works created using AI).

Beyond the choice to sue, if courts disproportionately rule against copyright defendants with certain legally irrelevant attributes, this too distorts the system's incentivist bargain. Specifically, anti-AI biases may bring about erroneous findings of copyright infringement associated with works produced by generative AI. This would discourage the (legal) creation of new works of authorship using such technologies and, thus, undermine basic policy goals of the copyright system. The following section presents a literature review setting the stage for our study into whether anti-AI biases can cause such distortions in the copyright system.

We begin by describing research associated with AI and human perception. The next subsection discusses prior work showing how legally irrelevant information can bias legal decisions in copyright law. Indeed, in later sections of this paper, we present new evidence of such phenomena in the realms of generative AI and copyright. Finally, the Section concludes by discussing research showing how biased use and application of the law can undermine the goals of copyright.

A. Perceptions of AI Actors

There is a robust literature on trust of AI actors[18] and factors that influence human reliance on machines.[19] However, a more foundational—and particularly relevant to this study—line of inquiry has been initiated in recent years. It asks how the

15. *See* Lydia Pallas Loren, *The Pope's Copyright? Aligning Incentives with Reality by Using Creative Motivation to Shape Copyright Protection*, 69 LA. L. REV. 1, 6 (2008).

16. *Id.*

17. FMC Corp. v. Control Sols., Inc., 369 F. Supp. 2d 539, 578 (E.D. Pa. 2005); Infodek, Inc. v. Meredith-Webb Printing Co., 830 F. Supp. 614, 622 (N.D. Ga. 1993); *Harmonizing Copyright's Internationalization with Domestic Constitutional Constraints*, 121 HARV. L. REV. 1798, 1817 n.124 (2008).

18. Mary Cummings, *Automation Bias in Intelligent Time Critical Decision Support Systems*, COLLECTION OF TECHNICAL PAPERS - AIAA 1ST INTELLIGENT SYSTEMS TECHNICAL CONFERENCE (2004); Raja Parasuraman & Dietrich H. Manzey, *Complacency and Bias in Human Use of Automation: An Attentional Integration*, 52 HUMAN FACTORS 381, 381–410 (2010).

19. Dietvorst et al., *supra* note 10, at 114–26.

presence of AI alters human perceptions of what is happening and what has happened.[20] In terms of academic real estate, this question falls within a combination of computer science, psychology, human factors, and machine behavior.[21]

We begin with an everyday example of human response to AI actions in the form of Tesla, Inc.'s most public-facing technology—namely, the driver assistance system in its automobiles. As of 2021, Teslas had been involved in just over ten accidents with emergency vehicles.[22] While any such collision is cause for concern, this number represents a small fraction of the approximately 4,500 annual accidents involving ambulances.[23] Despite this imbalance, the National Highway Traffic Safety Administration launched an investigation into the Tesla Autopilot system, citing concerns about these (ten) accidents.[24] When there are so many human-caused crashes, the emphasis on a small number of AI-related collisions may suggest a quirk in how AI errors are perceived by humans.[25]

Given the nature of this paper, a natural follow-up question arises from the proposition that humans perceive AI and human-centric events differently. Namely, if AI involvement influences how we perceive certain behaviors, does it likewise affect legal conclusions we draw therefrom? To this point, we hypothesize that the presence of AI will lead to disproportionately negative legal outcomes—even when AI's participation is legally irrelevant.

Such a supposition might seem far-fetched were it not for the well-documented influence that other legally irrelevant factors have on courtroom outcomes.[26] As an example, race should be irrelevant to legal conclusions, but research tells a different story. Consider a 2012 meta-analysis that analyzed thirty-four studies showing that, in most instances, race affected legal outcomes—even

---

20. Beth A. Bechky, *Evaluative Spillovers from Technological Change: The Effects of "DNA Envy" on Occupational Practices in Forensic Science*, 65 ADMIN. SCI. Q. 606, 606–43 (2020).

21. Iyad Rahwan et al., *Machine Behavior*, 568 Nature 477, 477–86 (2019).

22. Neil Boudette & Niraj Chokshi, *U.S. Will Investigate Tesla's Autopilot System over Crashes with Emergency Vehicles*, N.Y. TIMES (Aug. 16, 2021), https://www.nytimes.com/2021/08/16/business/tesla-autopilot-nhtsa.html [perma.cc/5ZTX-PGXP]; David Shepardson, *U.S. Identifies 12th Tesla Autopilot Car Crash Involving Emergency Vehicle*, REUTERS (Sept. 1, 2021, 5:31 PM), https://www.reuters.com/business/autos-transportation/us-identifies-12th-tesla-assisted-systems-car-crash-involving-emergency-vehicle-2021-09-01/ [perma.cc/W48Y-8L45].

23. Noah Smith, *A National Perspective on Ambulance Crashes and Safety*, EMSWORLD (Sept. 2015), https://www.hmpgloballearningnetwork.com/site/emsworld/article/12110600/a-national-perspective-on-ambulance-crashes-and-safety [perma.cc/87PQ-R96J]. *See also* Hyunjoo Jin, Mike Spector & David Shepardson, *A Life And Death Question For Regulators: Is Tesla's Autopilot Safe?*, REUTERS (Sept. 21, 2021, 5:20 PM), https://www.reuters.com/business/autos-transportation/life-death-que stion-regulators-is-teslas-autopilot-safe-2021-09-21/ [perma.cc/5S7K-EAKG].

24. Boudette & Chokshi, *supra* note 22.

25. Consider a second example. In 2018, an autonomous vehicle under testing by Uber struck and killed a pedestrian in Arizona. Kate Conger, *Driver Charged in Uber's Fatal 2018 Autonomous Car Crash*, N.Y. TIMES (Dec. 7, 2020), https://www.nytimes.com/2020/09/15/technology/uber-autonomous-crash-driver-charged.html [perma.cc/72GQ-MDYV]. Any fatal accident is a tragedy. But there were more than 1,000 deaths caused by automobile accidents in Arizona in 2018 and only one was attributed to an autonomous vehicle. After an investigation, liability and blame for the tragedy were attributed largely to a human back-up driver who was watching a video instead of watching the road. Nonetheless, public concern led Uber to cease testing of autonomous vehicles for years. ARIZ. DEP'T OF TRANSP., 2018 MOTOR VEHICLE CRASH FACTS FOR THE STATE OF ARIZONA vii (2019), https://azdot.gov/sites/default/files/news/2018-Crash-Facts.pdf [perma.cc/4YFD-VJ6L].

26. *See* BIAS IN THE LAW: A DEFINITIVE LOOK AT RACIAL PREJUDICE IN THE U.S. CRIMINAL JUSTICE SYSTEM, *supra* note 2.

after controlling for numerous non-race-factors.[27] Prosecutors are less likely to offer Black defendants a plea bargain, less likely to reduce charge offers for Black defendants, and more likely to offer Black defendants plea bargains that include prison time.[28]

Consistent with this theme, work by Bielen, Marneffe, and Mocan shows how a defendant's skin tone (which is, of course, legally irrelevant) can influence courtroom outcomes.[29] They used Virtual Reality technology to replace White defendants in the courtroom with individuals of Middle Eastern or North African descent.[30] Study participants were randomly assigned to watch the trial (using Virtual Reality headsets) from the viewpoint of the judge, and they made decisions as to guilt/innocence and other potential outcomes. The results showed racial biases against minority defendants regarding convictions, sentences, and fines.[31] Their work shows how decision making can be influenced by irrelevant information, such as race. We now review research showing how the presence of AI can similarly influence decision making.

---

27. BESIKI KUTATELADZE, VANESSA LYNN & EDWARD LIANG, DO RACE AND ETHNICITY MATTER IN PROSECUTION? A REVIEW OF EMPIRICAL STUDIES (2012), https://verainstitute.files.svdcdn.com/production/downloads/publications/race-and-ethnicity-in-prosecution-firstedition.pdf [perma.cc/GV5Z-P3UB].

28. Besiki Kutateladze, Nancy R. Andiloro & Brian D. Johnson, *Opening Pandora's Box: How Does Defendant Race Influence Plea Bargaining?*, 33 JUST. Q. 398 (2016); Besiki Kutateladze, Nancy R. Andiloro, Brian D. Johnson & Cassia Spohn, *Cumulative Disadvantage: Examining Racial and Ethnic Disparity in Prosecution and Sentencing*, 52 CRIMINOLOGY 514 (2014); JOHN MACDONALD & STEVEN RAPHAEL, AN ANALYSIS OF RACIAL AND ETHNIC DISPARITIES IN CASE DISPOSITIONS AND SENTENCING OUTCOMES FOR CRIMINAL CASES PRESENTED TO AND PROCESSED BY THE OFFICE OF THE SAN FRANCISCO DISTRICT ATTORNEY 3 (2017), http://raceandpolicing.issuelab.org/resources/30712/30712.pdf [perma.cc/W6SJ-ZJBN]. Defendants who are Black, young, and male fare worst. Celesta A. Albonetti, *Race and the Probability of Pleading Guilty*, 6 J. Quantitative Criminology 315 (1990); Kutateladze, Andiloro & Johnson, *supra*. In a study of federal cases, Rehavi and Starr found that Blacks received sentences that were 10% longer than those of similarly situated Whites, and this disparity was at least partly explained by prosecutors' initial charging decisions. M. Marit Rehavi & Sonja B. Starr, *Racial Disparity in Federal Criminal Sentences*, 122 J. POLIT. ECON. 1320 (2014).

29. Samantha Bielen, Wim Marneffe & Naci H. Mocan, *Racial Bias and In-Group Bias in Judicial Decisions: Evidence from Virtual Reality Courtrooms* (Nat'l Bureau of Econ. Rsch., Working Paper No. 25355, 2018).

30. *Id.*

31. *Id.* Racial biases are troublesome not only because of their prevalence—they also are extremely hard to mitigate, much less eradicate. As Professor Eberhardt and Banks wrote, "Most training programs have not been rigorously evaluated, and where they have been, results are, at best, mixed." Jennifer L. Eberhardt & Sandy Banks, *Implicit Bias Puts Lives in Jeopardy. Can Mandatory Training Reduce The Risk?*, L.A. TIMES (July 12, 2019, 6:15 AM), https://www.latimes.com/opinion/op-ed/la-oe-eberhardt-banks-implicit-bias-training-20190712-story.html [perma.cc/YR7K-KVF2]. Perhaps this should be unsurprising: we have a lifetime's worth of associations, and these associations often help us navigate the world. We rely upon them. In a recent study, it was found that criminal offenses deemed immoral were associated with Black males, while those deemed wrong but not necessarily immoral were associated with White males. Joseph Avery, DongWon Oh & Joel Cooper, *Race and Perceived Immorality in Stereotypes of Criminal Subtypes*, 43 BASIC & APPLIED SOC. PSYCH. 307, 307–18 (2021). Incorrectly and unjustly, the perceived trait of immorality was seen as responsible for Black males' prior bad acts; it also was seen as responsible for an increased likelihood that they would commit bad acts in the future. In contrast, immoral crimes committed by White males were perceived as mere mistakes, errors reflective of fleeting states. These participants' perceptions are objectively incorrect and deeply flawed, and we believe that similarly flawed thinking may be occurring in the context of AI. This is what makes the analogy to the present research apposite.

An initial line of work describes human/machine interactions. On this topic, decades of research show an inclination for humans to interact with machines as if they are human. Briefly restated, "Computers are social actors,"[32] and as such, computers elicit social responses from humans.[33] These responses, of course, may include shadings from implicit biases and stereotypes. The past few years have started to yield empirical evidence of these biases in play.

First consider studies on how identical behaviors are interpreted differently when they are performed by AI rather than by humans. Maninger and Shank delved into the perceptual differences in moral violations committed by artificial agents and humans.[34] Bigman and associates also explored the potential variance in moral judgments, proposing that the level of perceived human likeness in AI—how closely an AI resembles a human—might aid in bridging perceptual differences.[35] Raveendhran and Fast examined behavior tracking in the workplace when it was conducted by AI vs. humans, finding stark differences in how the behaviors were perceived.[36]

As concerns artwork in particular, Ragot and colleagues studied how participants' beliefs about whether a painting is created by AI or a human influenced perceptions about the work.[37] They provided a piece of artwork that was attributed to either a human or an AI author (though the work never changed). Study participants then reported whether they agreed with statements, including "This painting looks beautiful," "This painting seems novel," and "I perceive the meaning of the painting."[38] Works putatively made by AI actors were rated lower on all metrics—despite the same artwork being viewed by both groups.[39]

Hong proffers a theoretical foundation for such biases using schema theory.[40] Under this theory, an audience's perception of input is processed through the lens of biases, memories, and associated thoughts triggered during cognition.[41] As Hong explains, the lay public assumes that artwork is created by a human author, and any assessment of "human" artwork considers only the art itself.[42] In contrast, when asked to assess art created by AI, respondents will consider the attributes of the work, while their brain also subconsciously engages a second layer of cognition associated with the AI artist—including any positive or negative biases associated

---

32. Clifford Nass, Jonathan Steuer & Ellen Tauber, *Computers Are Social Actors*, CONF. ON HUM. FACTORS IN COMPUTING SYS. 72 (1994).

33. Youngme Moon, *Intimate Exchanges: Using Computers to Elicit Self-Disclosure from Consumers*, 26 J. OF CONSUMER RSCH. 323, 323–39 (2000); Clifford Nass & Youngme Moon, *Machines and Mindlessness: Social Responses to Computers*, 56 J. OF SOC. ISSUES 81, 81–103 (2000).

34. Timothy Maninger & Daniel B. Shank, *Perceptions of Violations by Artificial and Human Actors Across Moral Foundations*, 5 COMPUTERS IN HUM. BEHAV. REPORTS 100154 (2022).

35. Yochanan E. Bigman, Adam Waytz & Ron Alterovitz, *Holding Robots Responsible: The Elements of Machine Morality*, 23 TRENDS IN COGNITIVE SCIENCES 365, 365–68 (2019).

36. Roshni Raveendhran & Nathanael J. Fast, *Humans Judge, Algorithms Nudge: The Psychology of Behavior Tracking Acceptance*, 164 ORG. BEHAV. & HUM. DEC. PROC. 11, 11–26 (2021).

37. Ragot et al., *supra* note 11.

38. *Id.* at 3.

39. *Id.* at 7.

40. Joo-Wha Hong, *Bias in Perception of Art Produced by Artificial Intelligence*, *in* HUMAN-COMPUTER INTERACTION: INTERACTION IN CONTEXT 290, 292 (Hutchinson et al. eds., 2018).

41. *Id.*

42. *Id.* at 295–96.

therewith.[43] Under this theory, stereotypes or biases for or against generative AI can influence lay assessment of AI-created work without the respondent being aware of it.

The present research continues this line but with a focus on copyright law. Our empirical work will show that decision making about copyright is not consistent: The identity of the author under consideration (AI vs. human) matters. While this information should be legally irrelevant, we posit that this is not the case. Indeed, prior research presents evidence describing how other types of legally irrelevant information may bias copyright outcomes.

B. Legally Irrelevant Information Can Influence Copyright Decisions

Recognizing the literature on AI's influence on lay perceptions, we now describe how legally superfluous information may influence outcomes in copyright. On this point, Balganesh and colleagues investigated whether participants' beliefs about copyright infringement could be influenced by the amount of effort exerted in creating the original work. Of course, the amount of effort exerted in creating an original work is legally irrelevant; the copyright laws equally protect a masterpiece created in fifteen minutes and a travesty created over fifteen years. Further, assessment of substantial similarity—which is the hallmark of copyright infringement—should not be influenced by how much labor it took to create the original work. Two pieces are either similar or not; the amount of work exerted in creating the first work does not influence the image (or any similarity between two works).

To address this question, Balganesh and colleagues told respondents that an original author invested either two months ("high labor" condition) or ten minutes ("low labor" condition) into the creation of a work of art.[44] All parties were informed that the creator of a second work copied the original and were then instructed that the threshold for copyright infringement is "substantial similarity." They were subsequently shown the works and asked whether the second work was substantially similar to the original.[45]

Logically, the amount of labor should not influence judgements of similarity; the two are unrelated. However, the research team expected to see an influence arising via a mechanism that we refer to as "punishment."[46] They describe it as an act of "motivated reasoning," whereby respondents find certain copying to be a morally repugnant act worthy of punishment. As the proposal goes, this motivation increases the will to find the two works similar because that conclusion may lead to a warranted punishment.[47] They further hypothesized that the punishment mechanism could be enhanced by notions of fairness. If participants think it is unfair to copy a work that someone invested a lot of labor in creating, they will be motivated to assess the situation in a way that leads to punishment of the copier.

---

43. *Id.*

44. Shyamkrishna Balganesh, Irina D. Manta & Tess Wilkinson-Ryan, *Judging Similarity*, 100 IOWA L. REV. 267, 283 (2014). Respondents were also either told that the new work either had no effect on the market for the original work or significantly harmed the market for the original work. However, as no significant differences arose from these variations, they are not further discussed.

45. *Id.* at 282.

46. *Id.* at 281.

47. *Id.* at 282.

Here, the respondents could punish the copier by reporting the two works to be substantially similar—which triggers a finding of infringement.[48]

Results supported the punishment mechanism.[49] Participants disproportionately said that the works were substantially similar (and thus an infringement occurred) where the original represented a significant investment of authorial labor. The research team reasoned that respondents were more likely to find the copying of a work in the "high labor" condition relatively more morally repugnant than copying in the "low labor" condition. They believe this, in turn, caused respondents to engage in motivated reasoning in their assessment of substantial similarity.[50] In this way, legally (and logically) irrelevant information showed influence over copyright outcomes.

C. Legal Errors Associated with Irrelevant Information Can Distort the Copyright System

While Balganesh et al. (above) evidenced how seemingly irrelevant information can bias decisions regarding similarity, later work discusses the manner in which such biases distort the copyright system. Specifically, that literature describes how certain behaviors hinder copyright's primary goal—encouraging the creation of new works. In *The Moral Psychology of Copyright Infringement*, Buccafusco and Fagundes address a variety of reasons copyright holders might bring infringement lawsuits that diverge from the motivations expected under copyrights' incentivist structure.[51] They then discuss how those lawsuits disincentivize the creation of new works.

Under the incentivist narrative, copyright encourages authors to create new works by granting them exclusive rights associated with monetizing their output.[52] As a corollary, individuals are expected to only sue for copyright infringement where others inhibit their ability to make money. However, as Buccafusco and Fagundes point out, authors sometimes deviate from this theory by bringing lawsuits for non-monetary reasons—such as dislike for another author or disapproval of the message underlying an alleged infringement.[53]

---

48. *Id.* at 285–86. In contrast to the first study, the researchers believe the attention mechanism is unimportant here. In both conditions, the fact pattern said that the second work was copied from the original. Thus, to the extent that copying may cause respondents to look for similarities as predicted by the attention mechanism, this effect would be the exact same in both the "high labor" and "low labor" conditions because both groups received the same information about copying.

49. *Id.* at 280.

50. *Id.* at 281. In another part of the paper, a variable group of respondents was told that the creator of the second work copied the work from the first. The control group was told that the creator of the second work was unaware of the original. There was no high-/low-labor distinction, and no instructions about infringement. In that case, they expected the variable group to rate the works as being more similar. They reasoned that when the act of copying is pointed out, respondents begin to look for indica of copying. This reinforces any similarities between the two pieces. They refer to such behavior as "a kind of exercise in confirmation bias." In contrast, where the respondent is not told that the second work was copied from the original, they have no additional motivation to proactively look for similarities (i.e., indicia of copying). Despite the irrelevance of copying regarding how similar two works are, the evidence found the variable group to disproportionately rate the works as similar.

51. Buccafusco & Fagundes, *supra* note 12, at 2435.

52. *Id.* at 2438–44.

53. "[A]necdotal evidence indicates that owners of literary property frequently perceive themselves to be wronged for reasons unrelated to financial harm. Instead, their objections to unauthorized use invoke a variety of non-pecuniary considerations." *Id.* at 2435.

When a lawsuit ensures that an author reaps financial gains associated with their work, it furthers copyright's incentivist goals. However, "when authors sue to prevent uses of their works that are unrelated to protecting [financial] incentives, they cause substantial harms to public access, cultural exchange, and free speech."[54] In that instance, the lawsuit does not reward the author within the financial incentives allocated to them under copyright. Indeed, in such a case, the lawsuit's only copyright-relevant effect is *inhibiting* the creation of new works. Restated, these lawsuits harm copyright's goals by discouraging future authors (who fear being sued despite not interfering with another's financial gains) without vindicating copyright's financial incentives for the original author. As discussed later, we believe that anti-AI biases may lead copyright owners to bring a disproportionate number of copyright lawsuits against users of AI that are unrelated to financial incentives.

In a related vein, Said's *Jury-Related Errors in Copyright* discusses how copyright doctrine often diverges from lay jurors' intuitions of what is right or wrong.[55] Such divergences can bias application of the law. For example, it has previously been argued that jurors more faithfully apply legal rules that they believe are just.[56] Said cites the influence of legally irrelevant information in shaping juror opinions—such as how much labor it took to create an original work (as discussed in Balganesh et al.).[57]

She further posits that jurors' moral intuitions may favor "parties making labor-based claims," which would create pro-plaintiff biases as they are more likely to have invested significant effort in creating their work.[58] We believe this assertion is potentially enhanced where generative AI is involved. In such a case, jurors may favor a plaintiff who engaged in the human labor of creating the first work, as opposed to a defendant who employed an automated tool (generative AI) to make their work.[59]

Beyond influencing individual decisions, Said discusses how juror biases can distort copyright policy. She cites a case where a jury essentially expanded the scope of copyright protections to include functional objects—which are not traditionally protected by copyright.[60] Then, to the extent that such cases are relied upon as precedent, they may have subsequent distortionary effects on the copyright system.

Lastly, Said addresses the influence that well-known jury trials (that may be infected with biases) have on the copyright ecosystem. A large verdict may encourage the filing of more copyright lawsuits or an increased willingness to settle such cases on pro-plaintiff terms.[61] These effects can, in turn, influence later decisions to bring lawsuits or choices to engage in creative activities that might lead to a lawsuit against the author.[62]

---

54. *Id.* at 2483.

55. Said, *supra* note 13, at 754.

56. *Id.* at 767–68.

57. *Id.* at 777.

58. *Id.*

59. Savvy attorneys will recognize such biases and become disproportionately likely to represent plaintiffs in such suits and prosecute their claims through a jury trial. Thus, anti-AI bias might encourage more lawsuits against users of generative AI and stack the deck against those users at trial.

60. Said, *supra* note 13, at 752 (citing Plaintiffs' Objections to Defendants' Proposed Jury Charge at 3, No. 3:03 CV-0887, 2005 WL 2893652 (N.D. Tex. Oct. 11, 2005)).

61. *Id.* at 803.

62. *Id.*

Taken as a whole, the work done by Said, Buccafusco, and Fagundes describes how biases against particular copyright-irrelevant attributes or events can encourage (discourage) the filing of infringement lawsuits and increase (decrease) the chances that plaintiffs win. Both articles recognize the need for further study exploring the nature and scope of these biases.[63]

We take up this call. As applied to generative AI, our human trials test a generalized negative bias against works created by such technology. If the expected biases are present, they may lead to the users of generative AI having to defend against a disproportionately large number of infringement lawsuits and those parties losing a disproportionately large number of those cases.

It is not a leap from there to expect these trends to influence the use of generative AI to create new works of authorship. Parties will recognize the legal peril associated with doing so, even if their work does not harm an earlier author's financial incentives. In turn, those parties will logically reduce or abandon the practice. This stands in direct contravention of copyright's primary goal: encouraging the creation of new works. Instead, biased use of the copyright system could actually discourage such creation.

II. THE EMPIRICAL RESEARCH

Before we begin our investigation into lay assessments of putatively infringing works created by humans versus AI authors, we will set forth our empirical expectations. Building on Ragot et al. (comparing assessments of works putatively generated by AI or a human), we hypothesize that respondents who are told that artwork was AI-generated will rate it as lower quality. This expectation is further informed by the work of Hong—who described how information about AI artists triggers certain mental associations.[64] Given the generalizable negative biases against AI found in the above-described work, we expect these tendencies will cause respondents to rate the same artwork disproportionately negatively when they believe it was AI-generated. This expectation is stated in hypothesis 1:

H1: Respondents' beliefs that art was generated by AI will lead to lower assessments of the art's quality.

Our next hypothesis is informed by the findings of Balganesh et al. that lay respondents disproportionately found works substantially similar if the original was a labor-intensive production. We predict that this pro-labor bias will arise when our respondents are asked about the fairness of using AI to generate output in the same style as an earlier work. Specifically, participants are expected to view the use of generative AI to be disproportionately unfair or unethical vis-à-vis human artists.

In support of this expectation, the literature cites historical examples of similar biases influencing court opinions. Discussing *Roth Greeting Cards v. United Card*—where the defendant produced greeting cards similar to plaintiff's works—

63. Buccafusco & Fagundes, *supra* note 12, at 2507 ("One thing we suspect all readers will agree with is the need for further empirical study in this area. It would be valuable to know precisely what sorts of activities trigger owners' moral foundations and how widely distributed those responses are."); Said, *supra* note 13, at 804 ("Experiments and empirical research designed to test lay, legal, and judicial experiences of copyright law would be helpful . . . . More experimental work on biases and juror comprehension would be invaluable.")

64. Hong, *supra* note 40, at 292.

Samuelson states that the court was "probably influenced by a sense that [the defendant] had engaged in unfair competition by not doing enough independent creative work."[65] In support of this proposition, she notes that the opinion recognized that *the defendant employed no writers or artists* to create its greeting cards.[66] To the extent that "taking a short cut" in generating creative works was influential in *Roth Greeting Cards*, we predict that the use of generative AI in lieu of human authors may be viewed similarly. With this in mind, we expect:

H2: The use of generative AI to create new works will be viewed as less ethical or fair, relative to employing humans to do the same.

Finally, we believe that legal outcomes will be disproportionately negative for users of generative AI. As described immediately above, the use of generative AI is predicted to be viewed as unfair or unethical. From this, we expect to see evidence that respondents engage in motivated reasoning to punish users of generative AI (i.e., consistent with the punishment mechanism described in Balganesh et al.). This punishment would manifest itself through disproportionate willingness to bring infringement lawsuits against users of generative AI and to rule against those parties once they are in court. Accordingly, we expect to see disproportionately negative legal outcomes for users of generative AI in the copyright realm as stated:

H3: Biases against generative artificial intelligence will lead to worse legal outcomes for companies that make use of it.

The following section outlines our methodology employed in testing the above via human trials. Data collection, survey outlines, and empirical results are set forth.

A. Introduction to Studies

Section I of this Article laid the foundations for our primary hypothesis. Specifically, we expect that when an AI and a human independently create visually identical artwork, there exists an inclination among observers to view the AI-generated piece as more likely to infringe a copyright. This tendency would underscore a broader trend in perception of AI work product—one that potentially causes detrimental legal ramifications for parties leveraging AI technologies.

Despite the comprehensive theorization presented in Section I, the literature lacks empirical research confronting biases associated with generative AI and copyright. That absence is glaring when juxtaposed against the burgeoning tide of AI-related copyright litigation. This empirical void not only highlights the novelty and urgency of our inquiry but also underscores a broader academic and practical lacuna awaiting bridging efforts.

Recognizing this academic and practical gap, we designed a study, followed by a replication, to empirically probe our hypotheses. Beyond mere validation, our objective was multifaceted: to furnish the legal community with empirical insights that might shape legal strategy, influence judicial perspectives, and inform policy-making. Additionally, by pioneering research in this nexus of AI, works of authorship, and copyright law, we aspire to catalyze further scholarly exploration.

---

65. Pamela Samuelson, *A Fresh Look at Tests for Nonliteral Copyright Infringement*, 107 NW. U. L. REV. 1821, 1833 (2013) (citing Roth Greeting Cards v. United Card Co., 429 F.2d 1106, 1107–08 (9th Cir. 1970)).

66. *Id.* at 1833.

The ensuing discussion presents results from our empirical studies; it also elucidates the implications of the studies, offering both a diagnostic lens into prevailing biases and a prognostic perspective on the importance for the legal landscape.

B. Survey Design

Our primary survey and its replication were approved by the University of Miami's Institutional Review Board. The general design was as follows:[67] participants read about an artist. She is an emerging but not yet famous painter named Jessica Cassatt. Her work is described as being "youthful, diverse, and smart." Participants then read about a large U.S.-based company that launched a new marketing campaign. The goal of the campaign was to make the company appear "youthful, inclusive, and intellectually bold." After this basic introduction, participants were divided into two conditions—one human and one AI.

In the human condition, the company hired a freelance graphic designer and artist who specialized in marketing artwork. The company told the designer "to produce a few images that seem youthful, inclusive, and intellectually bold, in the style of Jessica Cassatt." The designer then created ten images that built upon their "broad art knowledge base, which included Jessica's paintings." One of the ten images was subsequently used in ads that ran in print media across the country.

The AI condition was identical except that, instead of hiring a human graphic designer and artist, the company used an AI artwork generator. Participants were informed about how the AI functioned ("if presented with instructions describing a desired image, it can produce a series of images using an AI-driven, computational process."). As in the human condition, they were told that the AI produced ten images that built upon its "broad art knowledge base, which included Jessica's paintings."

All participants were then told that the original artist—Jessica Cassatt—saw one of the ads and believed that the company had "ripped off" her work. She sued for copyright infringement, and the case was pending at the time of the study.

Participants were shown both the allegedly infringing image and the original produced by Cassatt. To ensure ecological validity,[68] participants were given ample opportunity to examine both the copyrighted work and the allegedly infringing one, with both being presented along with each question that concerned the images.

Further, the study was designed to protect against effects that might track the images themselves. First, the images were counterbalanced, such that they were switched (the copyrighted image became the allegedly infringing image, and vice-versa) for half of respondents. Second, for other respondents, an entirely different set of images was used.[69] Again, in that instance, the images were counterbalanced for half of the participants.

---

67. See https://osf.io/gb86e/ (last visited Sept. 1, 2024), for all study materials, including prompts and data.

68. David L. Breau & Brian Brook, *"Mock" Mock Juries: A Field Experiment on the Ecological Validity of Jury Simulations*, 31 LAW & PSYCHOL. REV. 77, 78 (2007) ("Social science researchers use the term ecological validity to describe the degree of similarity between the conditions of a simulation experiment and the real-world phenomenon that the experiment is designed to model." (citation omitted)).

69. See images at *Copyright and AI*, OSF, https://osf.io/gb86e/ (last visited Sept. 1, 2024).

Participants then responded to several questions. The inquiries can be loosely broken into three topics—with each category generally associated with one of our three hypotheses.

One set of questions concerned respondents' perceptions of human vs. AI-created art—unrelated to copyright issues. Here, participants were asked about the success of the allegedly infringing art ("success" in the sense of how well the human designer/AI art generator accomplished the aim of creating an image that appeared "youthful, inclusive, and intellectually bold"), the beauty of the art, its creativity, and its overall similarity in appearance to the original. With the exception of the similarity question, these queries are relevant to our Hypothesis 1, which states, "Respondents' beliefs that art was generated by AI will lead to lower assessments of the art's quality."

A second set of questions related to legal outcomes and thus addressed Hypothesis 3: Biases against generative AI will lead to worse legal outcomes for companies that make use of it. The first question began with this explanation of copyright infringement:

> When a court considers a copyright infringement claim, it asks the jury to determine whether the original work and the new work are "substantially similar." If the works are not substantially similar, there is no copyright infringement.
>
> An artist's "style" is not protected by copyright law. Therefore, two works are not substantially similar if all they share is the same "style."
>
> In the present case, when evaluating if the new painting is substantially similar to the original one, your only consideration should be the artistic expression embodied in the original painting.

These instructions set forth the legal standard—stating that it is legally irrelevant if the new work was created "in the style" of Cassatt's original.[70] Rather, respondents should only consider the substance of the new and original expressions. With this in mind, participants indicated on a seven-point Likert scale, ranging from strongly disagree (0) to strongly agree (6), whether they agreed with this statement: "I believe the original and new works are substantially similar."[71]

Using the same scale, participants were queried about how they would respond if they owned the original copyright or if they were on the jury in Cassatt's infringement trial. Specifically, respondents were asked to what extent they agreed with these statements:

- If I were the original artist and owned the copyright, I would sue for copyright infringement.[72]
- If I were on the jury in this case, I would rule that there was copyright infringement.[73]

---

70. Judith Ripka Designs, Ltd. v. Preville, 935 F. Supp. 237, 248 (S.D. N.Y. 1996) (copyright does not protect styles, only the particular original expressions of those styles); McDonald v. West, 138 F. Supp. 3d 448, 455 (S.D.N.Y. 2015), aff'd, 669 F. App'x 59 (2d Cir. 2016).

71. In the remainder of this Article, designated the "Substantial Similarity" question.

72. In the remainder of this Article, designated the "Bring Suit" question.

73. In the remainder of this Article, designated the "Find Infringement" question.

Each participant was then given follow-up questions regarding potential remedies. First, we presented four potential trial outcomes shown below—with remedies increasing in severity from left to right.[74] Using this scale, participants were asked, "In this case, what should the result be?"

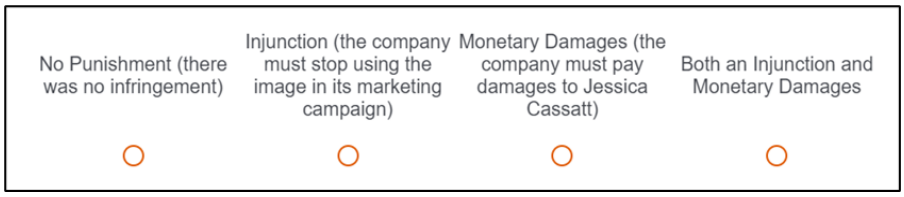

Following up, respondents were instructed to award appropriate damages on a $0 to $30,000 sliding scale.[75]

The last legal outcomes question asked about willfulness—a doctrine that allows damages to be increased in certain situations.[76] Specifically, participants were asked to what extent they agreed (on a seven-point Likert Scale) with the below:

- "Willfulness" in a copyright infringement case is when a person knows that their acts infringe another's copyright. In this case, there seems to be willful infringement.

Our final category of inquiries investigated perceived fairness and ethical considerations. These questions were intended to test Hypothesis 2: The use of generative AI to create new works will be viewed as less ethical or fair, relative to employing humans to do the same. Participants were asked to identify to what extent they agreed with the following three statements:

- The new image took more effort to create than the original
- It was ethical to create the new image.
- It was a fair business practice to use the new image in the marketing campaign.

Consistent with Hypothesis 2, we predicted responses from the AI group would report that the AI-generated work took relatively less effort than the Human group. Likewise, we expected the AI group to give relatively lower (less agreement) answers regarding whether the new work was created in an ethical manner or if its use was a fair business practice.

Lastly, we had control questions for how familiar the participants were with generative AI and art, and we included two attention check questions.[77] We concluded by asking for participants' age, gender, and race.

74. Arguably the second and third choices could be switched.

75. The $30,000 cap was selected to mimic the high end of the statutory damages allowed for non-willful-infringement. 17 U.S.C. § 504(c)(1). This may not be a perfect analogy, but it provides a base range for damages.

76. 17 U.S.C. § 504(c)(2).

77. Participants who failed either of the attention checks questions were removed.

C. Participants

Participants were recruited through Prolific,[78] an online platform for human intelligence tasks. Like any such source of participants, Prolific has its limitations.[79] But these types of platforms, including Amazon's Mechanical Turk, have found widespread acceptance and support in the academic community.[80] Separate samples of participants were used for the study and the replication.[81] Because we were interested in U.S. lay decision-makers, all participants were current U.S. residents. Data analysis was performed using the R software/programming language.[82]

The study population was as follows. To detect medium-sized effects across our conditions, the aim was to have 400 participants, so 410 participants were recruited. Thirteen failed the attention check questions and were removed, leaving a final population of 397 participants. These participants were 47% male, 47% female, 2% transmale, 1% transfemale, 3% gender variant/non-nonforming, and <1% indicated "not listed." The participants ranged from 18 to 85 years old, with an average age of 39.3 years. The racial breakdown was 72% White/Caucasian, 8% Black/African American, 8% Hispanic/Latino, 6% Asian/Pacific Islander, 1% Native American, 5% Multiracial, <1% indicated "not listed," and <1% elected not to disclose.

D. Results[83]

We begin by discussing results associated with generative AI and legal outcomes. Our expectations were embodied in Hypothesis 3, which provided that "biases against generative artificial intelligence will lead to worse legal outcomes for companies that make use of it." These "worse legal outcomes" were predicted to embody findings of infringement, willingness to bring an infringement lawsuit, and imposition of harsher legal remedies. After discussing those points, the remaining Results section discusses items that concern less central but still important issues, including ones that may hold clues for both explaining our primary results and applying the lessons learned to future litigation.

---

78. *See* PROLIFIC.CO, https://www.prolific.co [perma.cc/5KPA-YLLT] (last visited Apr. 12, 2025).

79. *See generally* Krin Irvine, David A. Hoffman & Tess Wilkinson-Ryan, *Law and Psychology Grows Up, Goes Online, and Replicates*, 15 J. EMPIRICAL LEG. STUD. 320, 326 (2018).

80. *See* Michael Buhrmester, Tracy Kwang & Samuel D. Gosling, *Amazon's Mechanical Turk: A New Source of Inexpensive, Yet High-Quality, Data?*, 6 PERSP. PSY. SCI. 3, 5 (2011); Krista Casler, Lydia Bickel & Elizabeth Hackett, *Separate but Equal?: A Comparison of Participants and Data Gathered via Amazon's MTurk, Social Media, and Face-to-Face Behavioral Testing*, 29 COMP. HUM. BEHAV. 2156, 2158 (2013).

81. The replication study is discussed below, and all study materials, including a writeup of the results, can be found at https://osf.io/gb86e/ [perma.cc/A9GW-YGQX]. Participants were blocked from participating in both the original study and the replication.

82. R version 3.6.2, R FOUNDATION FOR STATISTICAL COMPUTING (Dec. 2019), https://www.R-project.org/ [perma.cc/HX6X-S5FZ].

83. See *Copyright and AI*, OSF, https://osf.io/gb86e/ (last visited Sept. 1, 2024) for the data.

1. Are people more likely to find AI-generated (vs. human-generated) art to infringe earlier copyrights, even when the AI- and human-generated works are identical? In other words, is legal risk vis-à-vis copyright law elevated when AI is implicated?

This question was addressed with six items. As discussed supra, these items included Substantial Similarity (opinion on substantial similarity between the putatively infringing work and the original work), Bring Suit (likelihood of bringing suit if participant were situated like the plaintiff), Find Infringement (likelihood of finding copyright infringement if on jury), and three items regarding damages: First, should there be no punishment, an injunction, monetary damages, or both an injunction and monetary damages. Second, if damages, what should the damages be (ranging from $0 to $30,000). Third, was the alleged infringement willful, which is relevant, as it could lead to increased trial damages.

The results of the first three can be seen in Figures 1 and 2. In an examination of perceptions surrounding potential infringement, distinct differences were observed contingent upon the authorship of the artwork—human or AI. When participants were told that the art in question was created by AI, there was an increased propensity to judge the work as being substantially similar to the copyrighted material, in comparison to when the work was ascribed to a human creator. On a scale where greater scores indicate greater substantial similarity, the participants in the AI condition (Mean = 3.32, Standard Deviation = 1.84) were significantly more likely to believe the two works were substantially similar, relative to the human condition (M = 2.60, SD = 1.81).[84]

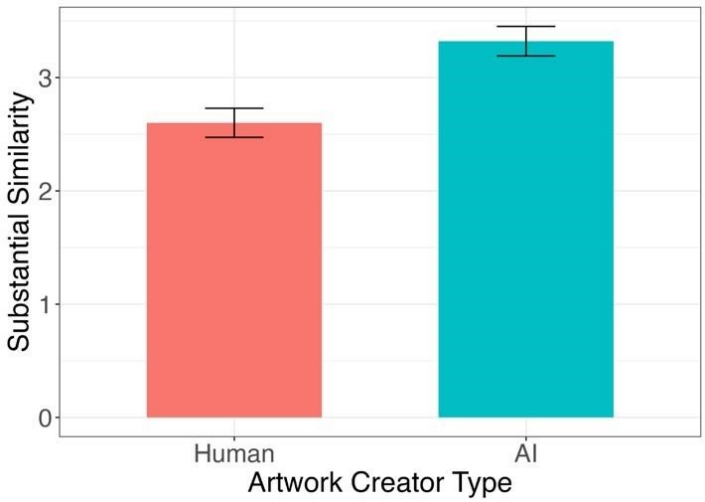

Figure 1. When a work of art was created by an AI, participants were significantly more likely to find the work to be substantially similar to a copyrighted work. Standard error bars, which indicate the variability of the mean estimates, are included.

In addition, there was a significant difference in respondents' willingness to bring an infringement lawsuit if they owned the copyright for the original work. The

84. $t(395) = -3.93$, $p < .001$.

AI group (M = 3.36, SD = 2.01) was disproportionately more likely to report that they would bring a lawsuit than the human group (M = 2.50, SD = 2.02).[85] Finally, when hypothetically positioned as jury members, participants exhibited a heightened likelihood of determining that the new work was infringing if it was generated by AI (M = 3.20, SD = 2.08), as opposed to a human (M = 2.35, SD = 2.03).[86] It is paramount to underscore that the artwork evaluated by participants across both conditions was identical. The sole variable, and thus the potential source of the observed bias, was the type of creator: human versus AI.

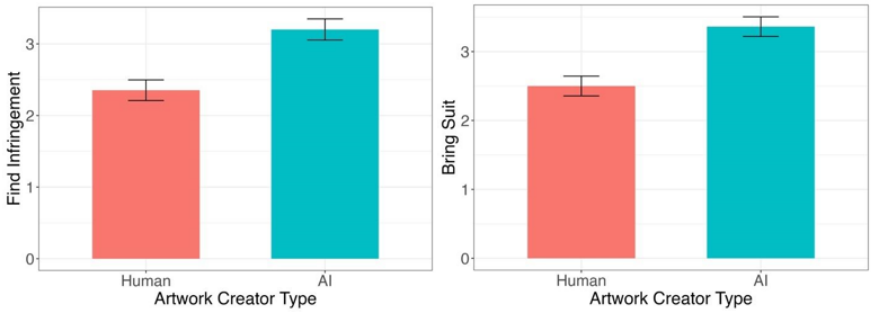

Figure 2. When a work of art was created by an AI, participants were significantly more likely to initiate a lawsuit regarding the matter and, if seated on the jury, to find infringement. Standard error bars, which indicate the variability of the mean estimates, are included.

Since these three items all studied our primary outcome (legal risk vis-à-vis copyright infringement), we expected them to have high internal consistency. This expectation was correct, as Cronbach's alpha for the three was .95, showing excellent internal consistency.[87] We thus collapsed (i.e., averaged) them into a single measure, and included them in a regression analysis assessing the impact of different variables on this combined outcome measure. In addition to our primary variable of interest—type of creator (AI vs. human)—we included all control variables: participants' age, gender, race, and their familiarity with generative AI and art.

Greater scores (i.e., a higher "Estimate") indicate worse legal outcomes, and a lower p-value indicates a lesser chance that the effect is due to random variation. After controlling for all participant-level variables, we found that the identity of the actor that caused the potential infringement (AI creator vs. human creator) mattered, such that legal outcomes were worse when AI was used. As shown in Figure 3, the effect was substantial (higher Estimate) and was unlikely to be due to random variation (low p-value), particularly when compared to the control variables.

85. $t(395) = -4.26$, $p < .001$.

86. $t(395) = -4.11$, $p < .001$.

87. Cronbach's alpha assesses how closely related a set of items or questions are as a group. It provides a measure of the internal consistency of a test or scale; it's an estimate of the reliability of a psychometric instrument. Lee Joseph Cronbach, *Coefficient Alpha and the Internal Structure of Tests*, 16 PSYCHOMETRIKA 297 (1951).

	Estimate	Std. Error	t-value	p-value
Intercept	1.25	.54	2.29	.02*
AI Creator in Prompt (vs Human Creator)	.85	.19	4.59	<.001***
Participants' Age	.01	.01	1.47	.26
Participants' Gender	.18	.12	1.53	.13
Participants' Race	.07	.06	1.13	.14
Participants' Familiarity with Generative AI	.08	.07	1.11	.27
Participants' Familiarity with Art	.12	.05	2.31	.02*

Figure 3.
*Significant at 0.05. **Significant at 0.01. ***Significant at 0.001.
Multiple R-squared: .07; Adjusted R-squared: .06
F-statistic: 5.14 on 6 and 390 DF, p-value <.001

The regression analysis revealed that type of creator (AI vs. human) plays a significant role in influencing the combined outcome value. In lay terms, the entity that caused the potential infringement (AI vs. human) mattered, such that the company that used the AI suffered worse legal outcomes. More specifically, the data showed that when the creator was an AI, there was a notable deterioration in legal outcomes, with an estimate of 0.8548 (greater scores on the 0–6 scale indicate worse legal outcomes). This result aligns with our initial hypothesis, positing an elevated legal risk associated with the use of AI. Restated, when AI is used in scenarios with potential copyright implications, the legal risks for the organization involved tend to escalate.

Among the control variables, only the artist factor (participants' professed familiarity with art) was statistically significant—indicating that this variable is also associated with the combined outcome (but to a small extent: Estimate = 0.1187).[88] However, race, gender, age, and familiarity did not significantly predict the combined outcome at a 5% significance level. Overall, the model accounted for approximately 7.3% of the variance in the combined outcome (Adjusted R-squared: 0.05896). This indicates that, while there are other unaccounted for factors influencing the combined outcome, our model captures some of the significant predictors.[89] In short, AI was positively and significantly related to the combined outcomes, such that legal risk was elevated when the potentially infringing artwork was created by AI.

The outcomes/punishments imposed by participants also were biased against AI. If we think of the outcomes as ranging from none (no punishment = 1) to light (injunction = 2) to more severe (damages = 3) to most severe (both injunction and damages = 4), then participants awarded more severe outcomes when the

88. $t(390) = 2.314$, $p = 0.0212$.

89. The overall significance of the model was further supported by the F-statistic: $F(6, 390) = 5.135$, $p < 0.001$.

infringement was caused by an AI. The mean punishment for the AI group was 2.51 (SD = 1.24), which was statistically significantly higher than the human (M = 2.13, SD = 1.21).[90] Similarly, participants imposed mean damages of $7,719 when the artwork was created by AI and only $5,931 when the artwork was created by a human. However, this result was shy of significance, with a p-value of .06.[91]

Our final question under this category dealt with the extent to which respondents agreed that "there seems to be willful infringement" in this case.[92] The trend was in keeping with the above, as respondents were more likely to believe there was willful infringement in the AI case (M = 3.81, SD = .14), as opposed to the Human condition (M = 3.38, SD = .14). This finding was statistically significant, but only slightly so, with a p-value of .03.

2. When a work's creator is an AI (rather than a human), do people perceive it differently? That is, are there perceptual differences? Relatedly, could these, or other specific factors, help explain why AI-created artwork increases legal risk vis-à-vis human-created work?

This portion of the study and the results are more experimental in nature, providing a foundation for future work unpacking the results discussed in the previous section. Here, we queried participants about their perception of the AI- and human-created artwork: its success (in accomplishing the artistic goal; "Successful"), its beauty ("Beautiful"), its creativity ("Creative"), and its overall similarity—from a bare, perceptual level—in appearance to the original work ("Resemblance"). These questions (except for Resemblance) address Hypothesis 1, which states, Respondents' beliefs that art was generated by AI will lead to lower assessments of the art's quality.

In addition, we also asked how effortful it likely was to produce it ("Effortful"), whether it was ethical to do so ("Ethical") and whether its use was a fair business practice ("Fair"). These questions related to Hypothesis 2: The use of generative AI to create new works will be viewed as less ethical or fair, relative to employing humans to do the same. The results for all eight items can be seen in Figure 4 below. In brief, we find strong support for Hypothesis 2 (AI use is viewed as less fair) and some support for Hypothesis 1 (AI art will be viewed as lower quality).

In analyzing this data, we recognize that because these were exploratory items loosely grouped together, there was an issue of multiple comparisons. We addressed this by using Bonferroni adjustments.[93] In brief, participants perceived the human-created artwork as relatively more beautiful and creative than the AI-created work (though there were no significant findings regarding the *Successful* question). Recalling the fact that the human-created and AI-created works were identical, this data provides partial support for Hypothesis 1.

---

90. t(395) = -3.10, p = .002.

91. t(394) = -.186, p = .06.

92. Again, a finding of willfulness is important in copyright litigation, as it potentially allows for increased damages. *See, e.g.*, 17 U.S.C. § 504(c)(2).

93. Karlo v. Pittsburgh Glass Works, LLC, 849 F.3d 61, 82 (3d Cir. 2017) ("The Bonferroni procedure makes it more difficult to find statistical significance where a researcher tests multiple comparisons using the same data. In theory, a researcher who searches for statistical significance in multiple attempts raises the probability of discovering it purely by chance, committing Type I error (i.e., finding a false positive).").

Participants likewise perceived the human-created work as requiring more effort. They further stated that its creation was more ethical and its subsequent business use fairer. This provides strong support for Hypothesis 2. Taken as a whole, the data reported in Figure 4 below supports the idea that the provenance of creative work (human vs. AI) radically shifts our perception of the work itself.

Our findings regarding Hypothesis 1—the use of generative AI to produce new creative works increases legal risk—should no longer be surprising. While our study was not designed to fully unpack the mechanisms underlying the data related to legal risks (indeed, our results suggest that it is multiply determined, such that the provenance of creative works shifts perception in many respects, and those shifts cause the increased exposure to legal risk), our study does provide a firm foundation for what will certainly prove to be a voluminous and robust line of future research.

	Human [Mean (SE)]	**AI [Mean (SE)]**	**t-statistic t(394)**	**p-value**	**Bonferroni adj. p-value**
Perceptions of the Work					
Successful	3.90 (.10)	3.70 (.11)	1.38	0.17	1.00
Beautiful	2.41 (.12)	1.91 (.11)	3.00	**0.003****	**0.02***
Creative	2.46 (.12)	1.97 (.12)	2.88	**0.004****	**0.03***
Resemblance	3.81 (.11)	4.10 (.11)	-1.87	0.06	0.44
Fairness and Ethical Use					
Effortful	2.32 (.12)	1.19 (.11)	6.74	**<0.001*****	**<0.001****
Ethical	2.88 (.12)	2.12 (.12)	4.41	**<0.001*****	**<0.001****
Fair	2.91 (.13)	2.05 (.13)	4.65	**<0.001*****	**<0.001****

Figure 4. Perception of AI- and human-created works of art.
*Significant at 0.05. **Significant at 0.01. ***Significant at 0.001.

E. Replication

We ran a second study (a replication) for a few reasons. First, the original was conducted in summer 2023 (July). AI is a rapidly developing area, and although adoption of generative AI systems spiked in early 2023, a considerable amount of adoption continued throughout the year. Thus, we decided to conduct a replication in fall 2023 (October). Second, a colleague mentioned that there might be something special about the use of AI by a "large" company, such that AI in big business might be vilified. So, in our replication, we revised things slightly, thus making it not fully a replication but somewhat of a second study. Instead of reading about a large company that was offloading a task (onto either AI or a human), participants read about a "small business" that was offloading a task.

The replication robustly confirmed our original results.[94] When participants believed art was created by AI, there was an increased propensity to judge it as being substantially similar to the copyrighted material—in comparison to when the work was ascribed to a human creator. On a scale where greater scores indicate greater substantial similarity, participants in the AI condition (M = 2.92, SD = 1.67) were significantly (p-value = .01) more likely to believe the two works were substantially similar, relative to the human condition (M = 2.29, SD = 1.84). There also was a significant difference in respondents' willingness to bring an infringement lawsuit if they owned the copyright for the original work: The AI group (M = 2.92, SD = 1.85) was disproportionately more likely (p-value = .002) to report that they would bring a lawsuit than the human group (M = 2.09, SD = 1.88). Finally, when hypothetically positioned as jury members, participants exhibited a heightened likelihood (p-value = .003) of determining that the new work was infringing if it was generated by AI (M = 2.77, SD = 1.90) as opposed to a human (M = 1.95, SD = 1.92).

Moreover, since these three items all studied our primary outcome (legal risk vis-à-vis copyright infringement), we expected them to again have high internal consistency. As set forth in the Appendix, regression analysis again showed that AI authorship was coordinated at a statistically significant level with negative legal outcomes. Overall, the significance and magnitude of the effect found in our Replication Study are similar to those found in the original.

III. DISCUSSION OF EMPIRICAL RESULTS

This Section addresses the implications of our empirical findings. We describe how biases against generative AI may undermine copyright's goal of encouraging the creation of new works. The first part discusses how anti-AI biases may motivate copyright lawsuits unrelated to financial incentives. We follow by describing how these biases could create positive feedback loops encouraging more lawsuits against AI users. Third, we discuss how these increased legal risks might disincentivize the use of generative AI and how this undermines copyright's primary goal. We close by considering how biases regarding AI may evolve or be crafted over time and future research that is necessary to address mechanisms underlying these biases. In sum, this Section connects our empirical results to underlying theory in order to highlight the application of our findings to the copyright system and its utilitarian goals.

A. Non-Financial Incentives to Sue Users of Generative AI

As previously discussed, the incentivist theory underlying domestic copyright laws predicts that infringement lawsuits will be brought when needed to vitiate economic interests hindered by others.[95] However, our results indicate that biases may exist that encourage copyright owners to bring infringement lawsuits against users of generative AI when they would not against human creators. Restated, we believe that parties are disproportionately more likely to sue users of generative AI despite having no financial reason to do so.

In our tests, no information was given to differentiate the human and AI groups, beyond the author's identity. Accordingly, respondents had no reason to

94. See Appendix for a complete presentation and discussion of the results.
95. *See supra* Section I(B).

believe that a lawsuit in either instance would be more lucrative than the other. Restated, participants had no economic information that would lead them to expect to gain more money if they sued parties using generative AI (versus human authors).

Likewise, our fact scenarios did not indicate that a lawsuit for either group (AI or human) would be more likely to succeed. Both groups were instructed that the second work was created "in the style of Jessica Cassatt," and all were told that "[a]n artist's 'style' is not protected by copyright law." Accordingly, there is no legal grounds upon which to determine that either group would be more likely to win their respective case.

Lastly, if the art was different for the AI (versus human) group, this could lead one set of respondents to be disproportionately more likely to bring a lawsuit. But this isn't the case. In our tests, the artwork was varied, such that one piece of art was presented as the original in 50% of instances and as the second work 50% of the time.[96] This was the same for both human and AI groups.

For each of these reasons, the underlying art could not make one group or the other think that they were more likely to win an infringement lawsuit or collect more damages. Restated, *there is no financial reason* that a party would disproportionately choose to bring a lawsuit against users of generative AI or human artists. Despite this, participants in the AI group were more likely to report that they would sue if they owned the original copyright.

Further, this was not a small variation. On a 0 to 6 scale ("not likely to sue" up to "very likely to sue"), the average responses for the AI group (M = 3.32 on the 0–6 scale) was 27.7% greater than those in the human group (M = 2.6). Furthermore, analysis of our data provided a p-value of less than .001—indicating a worse than 1 in 1,000 chance that our results were random variation.

With this data in mind, it is reasonable to expect a disproportionate number of lawsuits against users of generative AI. Further, these lawsuits are likely to be disproportionately motivated by non-monetary aims such as anti-AI bias or a belief that some unfairness is occurring (though copyright recognizes no legally cognizable difference between the groups in these instances). As Buccafusco and Fagundes explained, lawsuits unrelated to financial incentives generate substantial public costs by limiting access, exchange of information, and speech—without vitiating the economic interests that copyright is intended to provide. We believe that, absent any intervention, anti-AI biases will encourage the filing of a disproportionate number of such "non-financial" copyright lawsuits and further the problems outlined by Buccafusco and Fagundes.

B. Positive Feedback Loops and Infringement Litigation

We likewise posit that other anti-AI biases suggested by our empirical results are likely to create a positive feedback loop. Through such a mechanism, infringement lawsuits will beget more infringement lawsuits, which will beget more lawsuits, and so on. We say this for a number of reasons.

Our data initially suggests that potential jurors are more likely to find liability where the putatively infringing work was created using AI. Specifically, individuals in the AI group reported a greater likelihood of finding the new work to be

96. Likewise, in some instances, entirely different images were used.

infringing if they were on the jury (M = 3.20), as opposed to respondents in the human group (M = 2.35). Again, this finding was highly significant, with less than a 1 in 1,000 chance that it was random variation.

Savvy copyright owners and attorneys are likely to recognize this advantage. We thus expect them to bring infringement lawsuits against users of generative AI where they would not against human creatives. Indeed, this choice could be made in the absence of a strong infringement case or evidence of damages, but solely to encourage a defendant (who also is aware of anti-AI biases) to settle on pro-plaintiff terms.

Likewise, as recognized in Said's aforementioned research, a large verdict may encourage the filing of even more copyright lawsuits.[97] A single high dollar "blockbuster" case encourages other copyright plaintiffs to enter the fray where they might not have otherwise.[98] Such a case is disproportionately likely to arise where copyright owners and jurors hold biases against generative AI.

These considerations predict a positive feedback loop, whereby copyright owners are more likely to file lawsuits against users of generative AI. Then, based on our research, they are more likely to win these cases due to juror biases. Seeing these victories encourages still more copyright lawsuits, which are disproportionately likely to be won by the plaintiff, and on and on. This cycle only serves to exacerbate the issues associated with potentially biased jurors and the filing of copyright infringement lawsuits that are unrelated to the vitiation of the original author's financial interests.

C. Disincentives to Use Generative AI to Create New Works

The above-described behaviors are notable for the effects they may have in the litigation ecosystem. However, their second order effects are perhaps more important. Increased legal risks could have a chilling effect on parties' willingness to employ generative AI to create new works. This would, in turn, undermine the primary goal of copyright—namely encouraging the production of new works of authorship.[99]

To start, rational economic actors may decide not to use generative AI if they expect the potential costs to outweigh the benefits. This reluctance would be most acute for small entities and individual creators who lack the resources to defend against infringement lawsuits. Several downstream effects could arise therefrom.

Firms may rationally choose to employ human actors to generate works of authorship (instead of AI). Such an effect may be a preferred policy for certain parties, as that choice delays the possibility that generative AI will replace certain human creatives. This effect is, however, arguably less preferable from an aggregate financial perspective. Based on the reasonable assumption that human creatives cost more to employ than generative AI for certain projects (or at least they will in the future), policies that discourage the use of art-producing technology may increase the cost for consumers to enjoy their output. Moreover, increased costs

97. Said, *supra* note 13, at 803.
98. *Id.*
99. Loren, *supra* note 15, at 6.

(associated with employing human artists) may encourage certain firms to forego production or use of creative works altogether.

The rational decision to use human creatives instead of generative AI or to forgo creative endeavors entirely will disproportionately influence the actions of small firms. As stated above, these entities cannot afford to undertake increased legal risks. Moreover, they are also less likely to be able to pay increased costs associated with employing human creatives in lieu of generative AI. In turn, they may rationally opt out of the use of creative actors. This is not to say that such choices will not be made by larger firms. All entities are cost sensitive to some degree; we simply expect the influence to be outsized for smaller firms.

At scale, the aggregate effects could be profound. If a significant portion of users abandon or limit their use of generative AI, this would substantially reduce the new works produced with the aid of these technologies. Given generative AI's potential to significantly increase creative output, disincentives to its use may undermine the copyright system's goal of promoting new works of authorship.

To be fair, we recognize potential arguments relating to the relative quality of creative output from humans versus AI. It is possible that human creative works are "better" in some manner. If that is true, encouraging the use of generative AI may be suboptimal to the extent it serves as a substitute for certain human endeavors. We are, however, unaware of any studies that would support that proposal, beyond arguments that the use of generative AI may financially undercut certain human creatives by offering a cheaper alternative.

D. Potential Evolution of AI-Related Biases

Our findings regarding lay perceptions about the aesthetic value of AI-generated work, while complex, suggest differentiation in how participants perceive art based on its creator. Artworks created by AI were, in some cases, seen as marginally less beautiful, creative, and effortful than their human-created counterparts. These perceptions might stem from deeply ingrained beliefs about the inherent value of human effort and the authenticity of the creative process. Indeed, prior research supports the contention that humans view AI as an "outgroup," and perceptions associated with AI are viewed in a more negative light.[100] Assessed in such a light, we expect biases associated with generative AI to move in concert with societal acceptance and normalization of robots and AI in society.

A related explanation builds from the above-described work of Hong—who proffered that schema theory explains how assessment of AI works is influenced by mental associations that viewers hold to generative AI.[101] The literature on associative network theory is informative on how these associations are created and

100. *See* Congyu Wang & Kaiping Peng, *AI Experience Predicts Identification with Humankind*, 13 BEHAV. SCI. 89 (2023). Shuang Yu, Cheng Xu, Xueji Liang, Xudong He, Yanqi Sun & Hao Xu, *Exploring Human Trust in Ai: The Influence of Ingroup/Outgroup Context and Ai Human-Likeness*, SSRN ELEC. J., Apr. 1, 2023, at 1. Yuwei Zhou, Yichuan Shi, Wei Lu & Fang Wan, *Did Artificial Intelligence Invade Humans? The Study on the Mechanism of Patients' Willingness to Accept Artificial Intelligence Medical Care: From the Perspective of Intergroup Threat Theory*, FRONTIERS IN PSYCH., May 2022, at 1.

101. *See* Hong, *supra* note 40, at 292.

evolve.[102] It describes how related concepts in the mind are connected by links of varying strength, with stronger connections being more likely to activate when a topic comes to mind.[103]

As an example, seeing a baseball may lead one's mind to think about a baseball bat. The association between these topics is strong, such that one is likely to evoke the other. This is germane to Hong's work, as he describes how certain associations are triggered when a viewer is told about an AI author.[104] As he proffered, people have certain negative stereotypes (i.e., associations related to AI) that activate when an individual is informed that a work is created by AI. This, in turn, influences their assessment thereof.

But these associations are not static. Links between ideas—perhaps between generative AI and un-creativity—are formed over time as individuals are exposed to various experiences.[105] Simply paying attention to a group of ideas at the same time can reinforce their connection[106]—such that each of these ideas are likely to come to mind when one of them is triggered. Thus, to the extent that individuals think (or are told) that generative AI produces highly creative or high-quality artwork, the connection between generative AI and positive associations is reinforced. As such positive (or negative) associations strengthen in the public's collective mind, individuals become disproportionately likely to have positive (or negative) thoughts when informed that a machine created a work of art. That, in turn, influences their assessment of the art.

Based on this, we can expect subjective assessments of generative AI's output to move in concert with lay associations to the technology. As society comes to believe that these machines generally produce good (or bad) work, those stereotypes will influence assessments of the quality of particular works of art (that logically should be independent from society's aggregate beliefs). Accordingly, we expect that influential parties and variation in the quality of generative AI will affect changes in individual assessments of particular works by altering public stereotypes.

Because these associations are not static and subjective assessments will move over time, it is important that continued research on perception of AI in the legal space be conducted. We took steps in this direction by running our primary study in summer 2023 and a follow-up (the replication study) in October 2023. It is crucial to timestamp these findings and reassess them periodically to understand and track associations and assessments.

Beyond aesthetic perceptions, our study also found that participants viewed the creation of art to be more ethical and fairer when it was generated by a human, as opposed to AI. Similar to artistic assessment, we believe fairness considerations will vary as societal norms and stereotypes evolve. Specifically, we expect the use of generative AI to increase in perceived fairness as AI is integrated into everyday life

---

102. Suneal Bedi & Mike Schuster, *Towards an Objective Measure of Trademark Fame*, 54 U.C. DAVIS L. REV. 431 (2020).

103. *Id.* at 440–41.

104. Hong, *supra* note 40.

105. Bedi & Schuster, *supra* note 102, at 440; Kevin Lane Keller, *Conceptualizing, Measuring, and Managing Customer-Based Brand Equity*, 57 J. OF MKTG. 1, 2–3 (1993).

106. Jenni Romaniuk & Byron Sharp, *Conceptualizing and Measuring Brand Salience*, 4 MKTG. THEORY 327, 329 (2004); DANIEL REISBERG, COGNITION: EXPLORING THE SCIENCE OF THE MIND 2367 (4th ed. 2010).

and is "normalized." As that happens, we believe such perceptions of unfairness will be mitigated.

The above expectations are outside the scope of this initial foray into the intersection of copyright, psychological biases, and generative AI. However, they present fertile ground for future research into the mechanisms driving our results and changes in such biases (if any) over time. Moreover, research should look into *how* particular firms attempt to influence public perceptions and biases associated with generative AI to affect preferred policy choices.

CONCLUSION

There exists a chasm in scholarship on the intersection of AI and law. Generative AI is revolutionizing the creation of art, literature, and music—challenging the boundaries of intellectual property law. Yet to date, scholars have primarily focused on AI's authorship/entity status and the regulation of its use, overshadowing a critical issue: how AI's involvement in creative processes influences legal judgments in copyright disputes. In this Article, we presented empirical research on biases in legal assessments related to generative AI and copyright law. We showed that artworks attributed to generative AI, as compared to those credited to human artists, are subject to a discernible bias in legal decision-making. Specifically, artworks are more likely to trigger copyright lawsuits, be perceived as infringing, and attract harsher penalties when believed to be AI-created rather than human-created.

These findings have implications for copyright law and its underlying policies. Biases against AI could undermine the utilitarian objectives of copyright, deterring the use of AI in creative production. The results indicate a heightened legal risk for using generative AI, potentially discouraging its adoption for creative purposes, and thereby, inadvertently opposing copyright's aim to encourage creative endeavors.

Future research should explore whether these biases diminish over time. As we proposed in Section III.D., the widespread adoption of generative AI might reduce negative biases through increased exposure and familiarity, although this is not guaranteed. Ongoing empirical studies are crucial to observe shifts in perceptions of AI-created works and their legal interpretations. To counteract these biases, strategies such as public education or judicial instructions may be necessary to mitigate or prevent the potential adverse effects outlined.

In conclusion, our Article provides an investigation into biases surrounding generative AI within the copyright realm. We uncover anti-AI biases, which risk subverting the purpose of copyright law. As generative AI continues to proliferate, further research and policy initiatives are essential to ensure that this innovative technology can complement existing creative processes effectively, without triggering heightened legal risks.

APPENDIX - THE REPLICATION STUDY

A. Survey Design

As with the original study, our replication was approved by the University of Miami's Institutional Review Board. The design was identical to the original except

for the following minor changes.[107] First, instead of reading about a large company that was offloading a task (onto either an AI or a human), participants read about a "small business" that was offloading a task. This task resulted in ads that ran in print media (rather than, as in the original study, "print media across the country"). Second, because our primary findings concern the legal ramifications—desire to bring suit, findings vis-à-vis substantial similarity, findings vis-à-vis copyright infringement—and because these findings are so important for real-world legal applications, we wanted to replicate them. Even more, we wanted our replication to focus solely on these items. So, we shortened the experimental design, eliding everything except for those three key items and all control variables. Third, we used only one set of images, and we did not counterbalance the positioning of the images. That—multiple sets of images and counterbalancing of images—was done with an abundance of caution in designing the original study, and here, in the replication, we elected to simplify things.

The general design was as follows: Participants read about an artist. She is an emerging but not yet famous painter named Jessica Cassatt. Her work is described as being "youthful, diverse, and smart." Participants then read about a U.S.-based company that had launched a new marketing campaign. The goal of the campaign was to make the company appear "youthful, inclusive, and intellectually bold." After this basic introduction, participants were divided into two conditions—one human and one AI.

In the human condition, the company hired a freelance graphic designer and artist who specialized in marketing artwork. The company told the designer "to produce a few images that seem youthful, inclusive, and intellectually bold, in the style of Jessica Cassatt." The designer then created ten images that built upon their "broad art knowledge base, which included Jessica's paintings." One of the ten images was subsequently used in ads that ran in print media.

The AI condition was identical except that, instead of hiring a human graphic designer and artist, the company used an AI artwork generator. Participants were informed about how the AI functioned ("if presented with instructions describing a desired image, it can produce a series of images using an AI-driven, computational process."). As in the human condition, they were told that the AI produced ten images that built upon its "broad art knowledge base, which included Jessica's paintings."

All participants were then told that the original artist—Jessica Cassatt—saw one of the ads and believed that the company had "ripped off" her work. She sued for copyright infringement, and the case was pending at the time of the study.

Participants were shown both the allegedly infringing image and the original produced by Cassatt. To ensure ecological validity, participants were given ample opportunity to examine both the copyrighted work and the allegedly infringing one, with both being presented along with each question that concerned the images.

Participants then responded to the three key questions concerning legal outcomes: Biases against generative artificial intelligence will lead to worse legal

107. See *Copyright and AI*, OSF, https://osf.io/gb86e/ (last visited Sept. 1, 2024) for all study materials, including prompts and data.

outcomes for companies that make use of it. The first question began with this explanation of copyright infringement:

> When a court considers a copyright infringement claim, it asks the jury to determine whether the original work and the new work are "substantially similar." If the works are not substantially similar, there is no copyright infringement.
>
> An artist's "style" is not protected by copyright law. Therefore, two works are not substantially similar if all they share is the same "style."
>
> In the present case, when evaluating if the new painting is substantially similar to the original one, your only consideration should be the artistic expression embodied in the original painting.

These instructions set forth the legal standard—stating that it is legally irrelevant if the new work was created "in the style" of Cassatt's original. Rather, respondents should only consider the substance of the new and original expressions. With this in mind, participants indicated on a seven-point Likert scale, ranging from strongly disagree (0) to strongly agree (6), whether they agreed with this statement: "I believe the original and new works are substantially similar."

Using the same scale, participants were queried about how they would respond if they owned the original copyright or if they were on the jury in Cassatt's infringement trial. Specifically, respondents were asked to what extent they agreed with these statements:

- If I were the original artist and owned the copyright, I would sue for copyright infringement.
- If I were on the jury in this case, I would rule that there was copyright infringement.

Lastly, we had control questions for how familiar the participants were with generative AI and art, and we included two attention check questions. We concluded by asking for participants' age, gender, and race.

B. Participants

Participants were recruited through Prolific, an online platform for human intelligence tasks. Like any source of participants for human intelligence tasks, Prolific has its limitations, but such types of sources, such as Amazon's Mechanical Turk, have found widespread acceptance and support in the academic community. Separate samples of participants were used for the study and the replication. Because we were interested in U.S. lay decision-makers, all of the participants were current U.S. residents. Data analysis was performed using the R software/programming language.

The study population was as follows. For power reasons, the goal was to have 200 participants, so 210 participants were recruited. Six failed the attention check questions and were removed, leaving a final population of 204 participants. These participants were 49% male, 49% female, <1% transmale, and 1% gender variant/non-conforming. The participants ranged from 18 to 78 years old, with an average age of 37.1 years. The racial breakdown was 70% White/Caucasian,

10% Black/African American, 5% Hispanic/Latino, 10% Asian/Pacific Islander, <1% Middle Eastern, 3% Multiracial, and <1% elected not to disclose.

C. Results

1. Are people more likely to find AI-generated (vs. human-generated) art as infringing earlier copyrights, even when the AI- and human-generated works are identical? In other words, is legal risk vis-à-vis copyright law elevated when AI is implicated?

The results can be seen in Figures 5 and 6. As in the original study, in an examination of perceptions surrounding potential infringement, distinct differences were observed contingent upon the authorship of the artwork—human or AI. When participants were told that the art in question was the creation of an AI, there was an increased propensity to judge the work as being substantially similar to the copyrighted material, in comparison to when the work was ascribed to a human creator. On a scale where greater scores indicate greater substantial similarity, the participants in the AI condition (Mean = 2.92, Standard Deviation = 1.67) were significantly (p-value = .01) more likely to believe the two works were substantially similar, relative to the human condition (human condition: M = 2.29, SD = 1.84).

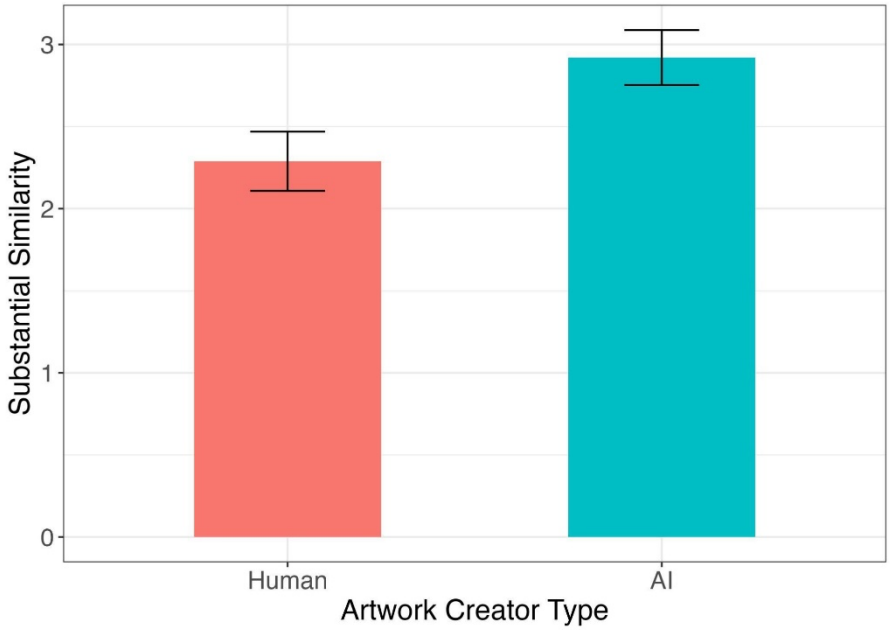

Figure 5. When a work of art was created by an AI, participants were significantly more likely to find the work to be substantially similar to a copyrighted work. Standard error bars, which indicate the variability of the mean estimates, are included.

There also was a significant difference in respondents' willingness to bring an infringement lawsuit if they owned the copyright for the original work. The AI group (M = 2.92, SD = 1.85) was disproportionately more likely (p-value = .002) to report that they would bring a lawsuit than the human group (M = 2.09, SD = 1.88). Finally, when hypothetically positioned as jury members, participants exhibited a heightened likelihood (p-value = .003) of determining that the new work was

infringing if it was generated by AI (M = 2.77, SD = 1.90), as opposed to a human (M = 1.95, SD = 1.92). It is paramount to underscore that the artwork evaluated by participants across both conditions was identical. The sole variable, and thus the potential source of the observed bias, was the type of creator: human versus AI.

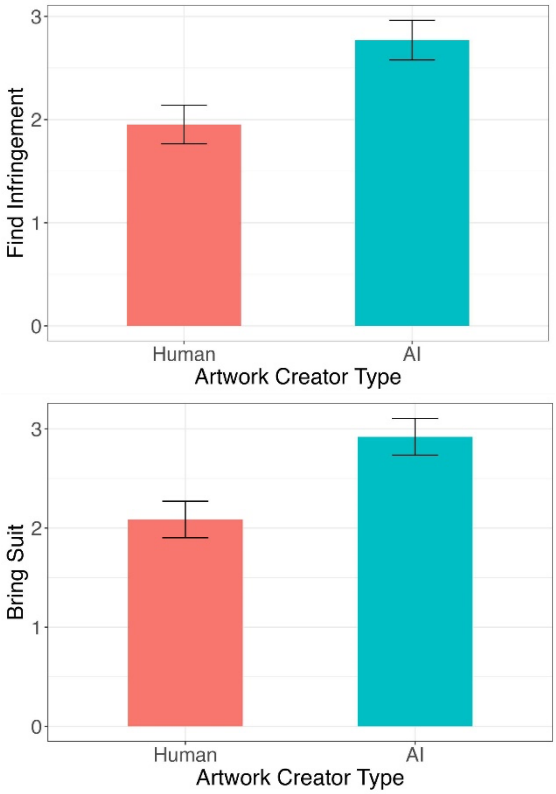

Figure 6. When a work of art was created by an AI, participants were significantly more likely to initiate a lawsuit regarding the matter and, if seated on the jury, to find infringement. Standard error bars, which indicate the variability of the mean estimates, are included.

Since these three items all studied our primary outcome (legal risk vis-à-vis copyright infringement), we expected them to have high internal consistency. Indeed, Cronbach's alpha for the three was .94—showing excellent internal consistency. We thus collapsed them into a single measure and included them in a regression analysis assessing the impact of different variables on this combined outcome measure. In addition to our primary variable of interest—type of creator (AI vs. human)—we included all control variables: participants' age, gender, race, and their familiarity with generative AI and art.

	Estimate	**Std. Error**	**t-value**	**p-value**
Intercept	1.40	.73	1.93	.06
AI Creator in Prompt (vs Human Creator)	.77	.24	3.19	.002**

Participants' Age	-.00	.01	-.13	.90
Participants' Gender	.50	.19	2.64	.009**
Participants' Race	.16	.08	2.04	.04*
Participants' Familiarity with Generative AI	.04	.10	.42	.68
Participants' Familiarity with Art	.06	.07	.95	.34

Figure 7. *Significant at 0.05. **Significant at 0.01.
Multiple R-squared: .11; Adjusted R-squared: .08
F-statistic: 3.92 on 6 and 197 DF, p-value = .001

Greater scores indicate worse legal outcomes. After controlling for all participant-level variables, it was found that the entity that caused the potential infringement (AI creator vs. human creator) mattered, such that legal outcomes were worse when the AI was used.

The regression analysis revealed that the nature of the creator plays a significant role in influencing the combined outcome value. In lay terms, whether the entity that caused the potential infringement (AI vs. human) mattered, such that the company that used the AI suffered worse legal outcomes. More specifically, the data showed that when the creator was an AI, there was a notable deterioration in legal outcomes, with an estimate of 0.77 (greater scores on the 0–6 scale indicate worse legal outcomes). This result aligns with our initial hypothesis, positing an elevated legal risk associated with the use of AI. Restated, when an AI is utilized in scenarios with potential copyright implications, the legal risks for the organization involved tend to escalate.

Overall, the model accounted for approximately 11% of the variance in the combined outcome (Adjusted R-squared: 0.08). This indicates that, while there are other unaccounted factors influencing the combined outcome, our model captures some of the significant predictors. The overall significance of the model was further supported by the F-statistic: $F(6, 197) = 3.92$, $p = 0.001$. In short, AI was positively and significantly related to the combined outcomes, such that legal risk was elevated when the potentially infringing artwork was created by an AI.