

## REGULATING HIDDEN AI AUTHORSHIP

*Jacob Noti-Victor\**

*With the rapid emergence of high-quality generative artificial intelligence (“AI”), some have advocated for mandatory disclosure when the technology is used to generate new text, images, or video. But the precise harms posed by nontransparent uses of generative AI have not been fully explored. While the use of the technology to produce material that masquerades as factual (“deepfakes”) is clearly deceptive, this Article focuses on a more ambiguous area: the consumer’s interest in knowing whether works of art or entertainment were created using generative AI.*

*In the markets for creative content—fine art, books, movies, television, music, and the like—producers have several financial reasons to hide the role of generative AI in a work’s creation. Copyright law is partially responsible. The Copyright Office and courts have concluded that only human-authored works are copyrightable, meaning much AI-generated content falls directly into the public domain. Producers thus have an incentive to conceal the role of generative AI in a work’s creation*

---

\* Associate Professor, Cardozo School of Law. For generative conversations and comments on earlier drafts, thanks to Amy Adler, BJ Ard, Shyam Balganes, Bruce Boyden, Oren Bracha, Robert Brauneis, Christopher Buccafusco, Mala Chatterjee, Bryan Choi, Carys Craig, Deven Desai, Stacey Dogan, Rebecca Eisenberg, Jeanne Fromer, Kristelia García, Katrina Geddes, Andrew Gilden, Myriam Gilles, Elizabeth Goldman, Eric Goldman, Michael Goodyear, James Grimmelmann, Paul Heaton, Michael Herz, Camilla Hrdy, Daphne Ippolito, Patricia Judd, Aniket Kesari, Jocelyn Kestenbaum, Rachel Landy, Jyh-An Lee, Mark Lemley, Melanie Leslie, Tiffany Li, Nathaniel Lubin, Florencia Marotta-Wurgler, Michael Meurer, Christina Mulligan, Yaran Noti-Victor, Michael Pollack, Nicholson Price, Blake Reid, Alex Reinert, Matthew Sag, Pam Samuelson, Sepehr Shahshahani, Jessica Silbey, David Simon, Ben Sobel, Chris Sprigman, Ed Stein, Stewart Sterk, Haochen Sun, Xiyin Tang, Charlotte Tschider, Rebecca Tushnet, Daniel Victor, Saurabh Vishnubhakat, Polk Wagner, Melissa Wasserman, Maggie Wittlin, Felix Wu, Christopher Yoo, and participants in the 2023 Junior IP Scholars Association (“JIPSA”) Workshop, the 2024 AALS IP Emerging Scholars Panel, the 2024 Tri-State Region IP Workshop, the Cardozo Summer Faculty Workshop, the 2024 Works in Progress in IP (“WIPI”) Conference, the 2024 Emory Law School AI Roundtable, the 2024 Copyright Scholars Roundtable, the 2024 Chinese University of Hong Kong Machine Lawyering Conference, the 2024 Intellectual Property Scholars Conference, the 2024 Junior Faculty Forum for Law and STEM, the Michigan Law School AI and the Law Seminar, and the Boston University Law School IP Workshop. Special thanks to Greyson Cohen and Samantha Moody for their excellent research assistance.

*because disclosure could jeopardize their ability to secure copyright protection and monetize the work.*

*Whether and why this obfuscation harms consumers is a different matter. The law has never required disclosure of the precise ways a work is created; indeed, failing to publicly disclose the use of a ghostwriter or other creative assistance is not actionable. But AI authorship is different for several reasons. There is growing evidence that consumers have strong ethical and aesthetic preferences for human-created works and understand the failure to disclose AI authorship as deceptive. Moreover, hidden AI authorship is normatively problematic from the perspective of various theories of artistic value. Works that masquerade as human-made destabilize art’s ability to encourage self-definition, empathy, and democratic engagement, turning all creative works into exclusively entertainment-focused commodities.*

*This Article also investigates ways to facilitate disclosure of the use of generative AI in creative works. Industry actors could be motivated to self-regulate, adopting a provenance-tracking or certification scheme. And Federal Trade Commission (“FTC”) enforcement could provide some additional checks on the misleading use of AI in a work’s creation. Intellectual property law could also help incentivize disclosure. In particular, doctrines designed to prevent the overclaiming of material in the public domain—such as copyright misuse—could be used to raise the financial stakes of failing to disclose the role of AI in a work’s creation.*

INTRODUCTION ..... 141

I. HIDDEN AI AUTHORSHIP IN THE CREATIVE INDUSTRIES ..... 147

    A. *Generative AI and Creative Content: An Overview* ..... 149

    B. *Incentives to Hide* ..... 152

        1. *Marketability* ..... 153

        2. *Copyrightability* ..... 154

    C. *Obfuscation in Practice* ..... 162

II. THE HARMS OF HIDDEN AI AUTHORSHIP ..... 164

    A. *Consumer “Process Preferences” for Human-Authored Content* ..... 166

        1. *Ethics* ..... 166

        2. *Authenticity* ..... 168

        3. *Fandom* ..... 170

2025]	<i>Regulating Hidden AI Authorship</i>	141
	<i>B. Authorship Transparency and the Social Role of Art</i> .....	173
	1. <i>Authorship and Readership as Social Dialogue</i> .....	176
	2. <i>From Obfuscation to (Complete) Commodification</i> .....	179
	<i>C. Implications and Objections</i> .....	182
	1. <i>Distinguishing the Human Ghostwriter</i> .....	183
	2. <i>Disentangling the AI Author from the AI Tool</i> .....	184
	3. <i>Leaving Space for the Responsible Proliferation of AI-Generated Content</i> .....	189
III.	SOLUTIONS.....	190
	<i>A. Affirmative Disclosure</i> .....	191
	1. <i>Private Ordering</i> .....	191
	i. <i>Provenance Tracking</i> .....	192
	ii. <i>Certification</i> .....	193
	2. <i>Legislation</i> .....	195
	<i>B. Penalizing Deceptive Omissions via FTC Enforcement</i> .....	197
	<i>C. Using Intellectual Property Law to Incentivize Disclosure</i> .....	199
	1. <i>Copyright Misuse</i> .....	200
	2. <i>Trademark and the Right of Publicity</i> .....	207
	CONCLUSION.....	209

## INTRODUCTION

When Marvel Studios’ big-budget series, *Secret Invasion*, premiered in June 2023, most viewers did not give a second thought to the show’s opening credits, which featured angular alien faces, a toothless Samuel L. Jackson, and swirling green cityscapes. Shortly after the show premiered, however, director and executive producer Ali Salim made an unusual admission: the credit sequence’s visuals had been generated using artificial intelligence (“AI”).<sup>1</sup> Public outcry was swift. Many criticized the use of generative AI by a wealthy studio as “unethical,” especially in light of Hollywood labor disputes driven by the possible effects of generative AI on acting and writing jobs.<sup>2</sup> Others argued that Marvel’s use of AI was

<sup>1</sup> Zosha Millman, Yes, *Secret Invasion*’s Opening Credits Scene Is AI-Made—Here’s Why, Polygon (June 22, 2023, 7:16 PM), <https://www.polygon.com/23767640/ai-mcu-secret-invasion-opening-credits> [https://perma.cc/5WPN-93TY].

<sup>2</sup> See Angela Watercutter, Marvel’s *Secret Invasion* AI Scandal Is Strangely Hopeful, Wired (June 23, 2023, 9:00 AM), <https://www.wired.com/story/marvel-secret-invasion-artificial-intelligence/>.

lazy, yielding images devoid of artistic merit.<sup>3</sup> The criticism ultimately prompted Marvel to walk back its admission, explaining that “AI is just one tool among the array of tool sets our artists used. No artists’ jobs were replaced by incorporating these new tools; instead, they complemented and assisted our creative teams.”<sup>4</sup>

With the dramatic arrival of high-quality generative AI systems,<sup>5</sup> scholars and policy-makers have begun debating the potential harms posed by the technology’s many possible applications. Much of this debate has centered on generative AI’s ability to create materials that masquerade as factual—in particular, false photorealistic images and audiovisual content, commonly known as “deepfakes”—which can harm individual reputations or further misinformation that undermines public trust.<sup>6</sup> But the *Secret Invasion* controversy illustrates an underexplored dimension of the lack of transparency in many uses of generative AI: Are consumers also deceived by nonfactual, AI-generated creative works that masquerade as human-made? Put another way, does hidden AI “authorship”—that is, the undisclosed use of AI to produce expression

---

<sup>3</sup> See Dani Di Placido, *Marvel’s AI-Generated ‘Secret Invasion’ Sequence Sparks Backlash*, *Forbes* (June 23, 2023, 11:47 AM), <https://www.forbes.com/sites/danidiplacido/2023/06/21/the-big-backlash-against-marvels-secret-invasion-explained/?sh=2ef04c17344e>.

<sup>4</sup> Carolyn Giardina, *‘Secret Invasion’ Opening Using AI Cost “No Artists’ Jobs,” Says Studio That Made It*, *Hollywood Rep.* (June 21, 2023, 8:12 PM), <https://www.hollywoodreporter.com/tv/tv-news/secret-invasion-ai-opening-1235521299/> [<https://perma.cc/3DTD-U46W>].

<sup>5</sup> For a full discussion of what I mean by “generative AI,” see *infra* Section I.A.

<sup>6</sup> See generally Lisa Macpherson, *Lies, Damn Lies, and Generative Artificial Intelligence: How GAI Automates Disinformation and What We Should Do About It*, *Pub. Knowledge* (Aug. 7, 2023), <https://publicknowledge.org/lies-damn-lies-and-generative-artificial-intelligence-how-gai-automates-disinformation-and-what-we-should-do-about-it/> [<https://perma.cc/9BNL-QN9S>] (discussing the potential harm to the integrity of news through the use of generative AI); Adam Satariano & Paul Mozur, *The People Onscreen Are Fake. The Disinformation Is Real*, *N.Y. Times* (Feb. 7, 2023), <https://www.nytimes.com/2023/02/07/technology/artificial-intelligence-training-deepfake.html> (describing how deepfakes make it difficult to separate reality from forgeries and enable the spread of propaganda by foreign governments); Todd C. Helmus, *Artificial Intelligence, Deepfakes, and Disinformation: A Primer*, *RAND Corp.* (July 6, 2022), <https://www.rand.org/pubs/perspectives/PEA1043-1.html> [<https://perma.cc/HD55-62L5>] (observing that, in an increasingly polarized and fact-resistant political climate, deepfakes pose a potent threat). See also Council Regulation 2024/1689, art. 134, 2024 O.J. (L) 1, 34 (EU) (requiring creators of deepfakes to disclose use of AI).

that we generally consider to be within the purview of human creators<sup>7</sup>—pose harm to the public?

This Article provides the first comprehensive treatment of this question, as well as the problem of hidden AI authorship more generally. In so doing, this Article makes three contributions. The first contribution is descriptive: the Article examines how and why producers of commercial creative works—visual art, books, television, music, films, and more—might choose to hide the role of generative AI in the production of new content. Though not immediately obvious, this phenomenon is deeply intertwined with intellectual property law, and copyright law in particular. As copyright decision-makers increasingly find that AI-generated works are largely unprotectable, producers have an incentive to hide their use of the technology. The second contribution is normative: the Article argues that the hiding of AI authorship indeed poses harm to consumers, albeit a less straightforward form of harm than the clear problems posed by deceptive deepfakes. This harm must be understood by examining the strong evidence that many consumers prefer human-created works of art and entertainment, as well as the broader social significance of human authorship in art’s ability to foster self-definition, empathy, and political engagement. The third contribution is prescriptive: the Article identifies various regulatory options, including existing consumer protection and intellectual property law regimes, that could be used to encourage greater transparency among the sellers of AI-generated content, enabling better informed consumer choice.

How might generative AI come to be frequently, but nontransparently, used to create new works of art and entertainment? As Part I explores, this problem is already emerging. Generative AI is quickly being incorporated into content creation, leading large content producers to encounter a dilemma like the one faced by Marvel: whether or not to disclose the role of the technology in a work’s creation. As the creators of *Secret Invasion* discovered, many consumers seem to bristle at the use of the technology. Indeed, recent empirical research suggests that many consumers consider “AI-generated” works inferior, even if they cannot tell from the work

---

<sup>7</sup> As discussed below, AI cannot be an author for legal purposes, so when I say “AI authorship,” I refer to situations in which AI has accomplished the creative work that we generally associate with human authorship. See *infra* Subsections I.B.2, II.C.2.

itself that AI had a role in its creation.<sup>8</sup> This yields a clear financial incentive to hide the AI provenance of a work from the public.

Current trends in copyright law compound producers' incentives to hide their use of generative AI. It is black-letter law that a work created by a nonhuman is ineligible for copyright protection.<sup>9</sup> Courts and the Copyright Office<sup>10</sup> have emphasized the importance of "elements of human creativity" when assessing whether an AI-generated work can be registered, such as human-made decisions about how to organize and structure AI-generated material in a final work.<sup>11</sup> For content producers, this means that highlighting the role of generative AI in a work's production can compromise efforts to achieve copyright registration.<sup>12</sup> Failure to obtain protection essentially means that new content immediately falls into the public domain and cannot be monetized. Thus, if trends in the law continue in their likely direction, content producers will increasingly try to hide the role of AI in new creative works to ensure such works remain protectable.<sup>13</sup>

Should this obfuscation be considered a problem? After all, if a consumer enjoys a work like the *Secret Invasion* credit images, does it matter whether they know that work was produced using generative AI?

---

<sup>8</sup> See *infra* Subsection I.B.1; see, e.g., Lucas Bellaiche et al., *Humans Versus AI: Whether and Why We Prefer Human-Created Compared to AI-Created Artwork*, 8 *Cognitive Rsch.*, 2023, at 1, 3 (observing that, across multiple artistic mediums, study participants preferred art labelled "human-created" over art labelled "AI-created").

<sup>9</sup> See, e.g., *Naruto v. Slater*, 888 F.3d 418, 426 (9th Cir. 2018) (monkey that took "selfie" photos could not claim copyright for lack of standing); see also *infra* Subsection I.B.2 (discussing limits on copyright protections for works not produced by human beings).

<sup>10</sup> The Copyright Office is a regulatory body housed in the Library of Congress that is responsible for registering new works. See *infra* Subsection I.B.2.

<sup>11</sup> See, e.g., Letter from Robert J. Kasunic, U.S. Copyright Off. Rev. Bd., to Van Lindberg, Taylor Eng. Duma LLP, *Zarya of the Dawn* (Registration # VAu001480196), at 5–8 (Feb. 21, 2023), <https://www.copyright.gov/docs/zarya-of-the-dawn.pdf> [<https://perma.cc/X54U-MQ4F>] (finding that AI-generated images in a comic book were unprotectable but the comic book could still be thinly protected as a compilation).

<sup>12</sup> See, e.g., *Thaler v. Perlmutter*, 687 F. Supp. 3d 140, 149–50 (D.D.C. 2023) (upholding denial of copyright registration where "the record designed by plaintiff from the outset of his application for copyright registration . . . [showed] the absence of any human involvement in the creation of the work," precluding copyrightability for AI-generated image); see also *infra* Subsection I.B.2 (describing instances in which acknowledging the role of AI in a work's creation foreclosed copyright protection).

<sup>13</sup> Disclosure to the Copyright Office and public disclosure are interrelated. The Copyright Office has begun listing registrations that explicitly state whether a work is a product of generative AI (noting that the AI-produced elements are unprotectable). These registrations are easily publicly searchable. See *infra* Section I.C.

Part II addresses this question, arguing that consumers seem to have a range of “process preferences”<sup>14</sup>—that is, preferences that relate to how a work was created, rather than just the work itself—that implicate the use of generative AI. One issue is ethical: consumers may prefer human-created works because of ethical concerns over AI supplanting human labor.<sup>15</sup> Another issue is aesthetic: as showcased in old debates regarding the use of “mechanical reproduction” in art, human creation can confer an element of authenticity on a creative work that a machine-generated work lacks.<sup>16</sup> Finally, consumers have deep-seated connections to specific artists, born out of a sense of fandom, which are undermined by the more specific example of AI-generated works that mimic an artist’s voice or style.<sup>17</sup> Considering these preferences, obscuring the role of generative AI in a work’s creation may prevent a consumer from making an informed decision about whether to consume it.

But just because consumers have these preferences does not mean the law must respect them.<sup>18</sup> Part II thus also provides a separate normative case for why consumers who care about human authorship should be taken seriously. As aesthetic and ethical theorists have argued, authorship and readership<sup>19</sup> are fundamentally social activities; through art, the public can engage in ongoing “dialogic” processes of self-definition, ethical development, and political engagement. Novelist and journalist Jay Caspian Kang has recently put it more plainly: “[T]he reason we read books and listen to songs and look at paintings is to see the self in another self, or even to just see what other people are capable of creating.”<sup>20</sup> Thus, even if, as some have argued, the author’s “intent” lacks significance,<sup>21</sup>

---

<sup>14</sup> See Douglas A. Kysar, *Preferences for Processes: The Process/Product Distinction and the Regulation of Consumer Choice*, 118 *Harv. L. Rev.* 525, 529 (2004) (examining how “consumer preferences may be heavily influenced by information regarding the manner in which goods are produced”).

<sup>15</sup> See *infra* Subsection II.A.1.

<sup>16</sup> See *infra* Subsection II.A.2.

<sup>17</sup> See *infra* Subsection II.A.3.

<sup>18</sup> See *infra* Part II (examining difficulties of determining whether consumer preferences should give rise to regulatory action).

<sup>19</sup> By “readership,” I mean the experience of engaging with a work of art or entertainment. See *infra* Section II.B.

<sup>20</sup> Jay Caspian Kang, *What’s the Point of Reading Writing by Humans?*, *New Yorker* (Mar. 31, 2023), <https://www.newyorker.com/news/our-columnists/whats-the-point-of-reading-writing-by-humans>; see also *infra* note 199 (surveying other writers’ similar perspectives).

<sup>21</sup> See *infra* Subsection II.B.1 (noting that postmodern theorists have questioned the importance of the individual author). See generally Roland Barthes, *The Death of the Author*, in *Image, Music, Text* 142 (Stephen Heath trans., 1977) (criticizing literary critics’

the author's and reader's basic shared humanness can be essential to allowing art to play a meaningful social and ethical function. The undisclosed use of generative AI in authoring a work<sup>22</sup> fundamentally destabilizes this dialogue between author and reader, robbing art of its social value and turning it into an exclusively entertainment-focused commodity.<sup>23</sup> The argument here is not that AI authorship is inherently immoral; indeed, AI might yield a range of works that consumers enjoy. Rather, it is that such use must be disclosed in order to allow consumers to choose whether and on what terms they wish to engage with a work.

The obvious solution to the problem of undisclosed AI authorship is to provide consumers with information about a work's provenance, so that they can make an informed choice. Part III explores various regulatory options for fostering transparency, examining their benefits and shortcomings. An affirmative disclosure regime could come about through industry self-regulation; if it is true that some consumers prefer human-made works, the market would logically step in to provide this information.<sup>24</sup> A legislative transparency mandate would also—and more thoroughly—accomplish this task.<sup>25</sup> In lieu of a comprehensive affirmative disclosure regime, the FTC could also target specific instances in which producers deceptively omit information about a work's origins so as to mislead consumers.<sup>26</sup>

An additional, and perhaps more politically feasible,<sup>27</sup> set of options is offered by existing intellectual property law. Such an approach would look to IP's existing doctrines as a way of forcing information, making it costlier for producers to hide the fact that they used generative AI to produce works. In particular, litigants could take advantage of the often-

---

preoccupation with individual authors and instead emphasizing the importance of readers as recipients and interpreters of literary texts).

<sup>22</sup> Importantly, I distinguish between truly "authoring" a work—that is, generating something that is at the creative heart of the work—and using AI in a merely assistive role. Copyright's doctrinal distinction between these two concepts corresponds well to this normative distinction. See *infra* Subsection II.C.2. I also explain why the use of a human ghostwriter does not pose the same problems as undisclosed AI authorship. See *infra* Subsection II.C.1.

<sup>23</sup> See *infra* Subsection II.B.2 (exploring how lack of knowledge regarding a work's provenance forces consumers to *exclusively* engage with works of art on market-based, rather than social, terms).

<sup>24</sup> See *infra* Subsection III.A.1 (examining private ordering solutions, such as provenance tracking and a certification regime).

<sup>25</sup> See *infra* Subsection III.A.2.

<sup>26</sup> See *infra* Section III.B.

<sup>27</sup> See *infra* Section III.C (exploring barriers to legislation and FTC enforcement).



ignored doctrine of copyright misuse to police those who assert that an entire work was human-created, when, in fact, it was a product of AI. Such assertions should fall within one of the categories of copyright misuse: the overclaiming of material that is in the public domain.<sup>28</sup> Using the copyright misuse doctrine in litigation would raise the financial risks of surreptitiously using AI-generated materials, incentivizing rightsholders to disclose (and disclaim) this content.<sup>29</sup> Trademark law and the right of publicity could also play a role in raising the financial stakes of nontransparency. For the specific subset of AI-generated works that mimic a human artist's voice or likeness,<sup>30</sup> trademark and the right of publicity provide causes of action that could subject producers and distributors to damages.<sup>31</sup> In combination, these various tools could ideally achieve a world in which information about most works' provenance is readily accessible to consumers.

#### I. HIDDEN AI AUTHORSHIP IN THE CREATIVE INDUSTRIES

The exact definition of generative AI, especially as distinct from other forms of AI, remains in flux. Following the lead of recent scholarship in this area, I consider generative AI to be a “catch-all name for a massive ecosystem of loosely related technologies, including conversational text chatbots like ChatGPT, image generators like Midjourney and DALL-E, coding assistants like GitHub Copilot, and systems that compose music, create videos, and suggest molecules for new medical drugs.”<sup>32</sup> Though these systems can employ very different architectures, they all rely on machine learning and massive amounts of data to produce sophisticated, high-quality, and (generally) unique text, image, and/or sound-based outputs.<sup>33</sup> Matthew Sag notes that one distinguishing feature of the new

---

<sup>28</sup> See *infra* Subsection III.C.1; see also *infra* Subsection I.B.2 (noting that AI-generated material is inherently in the public domain due to lack of authorship, rendering many AI-derived works unprotectable or only thinly protectable).

<sup>29</sup> See *infra* Subsection III.C.1. In particular, a copyright misuse finding prevents a rightsholder from enforcing even legally protectable aspects of their work, rendering the work completely uncopyrightable and essentially valueless. See *infra* Subsection III.C.1.

<sup>30</sup> See *infra* Subsection II.A.3 (discussing examples such as “fake Drake”).

<sup>31</sup> See *infra* Subsection III.C.2.

<sup>32</sup> Katherine Lee, A. Feder Cooper & James Grimmelmann, Talkin' 'Bout AI Generation: Copyright and the Generative-AI Supply Chain, *J. Copyright Soc'y U.S.A.* (forthcoming) (manuscript at 4), [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4523551](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4523551) [<https://perma.cc/YL3V-LR49>].

<sup>33</sup> See *id.* (manuscript at 7–16); Matthew Sag, Copyright Safety for Generative AI, 61 *Hous. L. Rev.* 295, 313–25 (2023).

generative AI systems is that they are capable of passing the Turing Test; that is, they can generate materials that the human observer could believe was created by another human.<sup>34</sup>

While generative AI technology has a range of different applications, this Article is predominantly concerned with systems that users can prompt in order to produce text, images, or sounds.<sup>35</sup> Within this broad category, the Article focuses on AI-generated content that is used for *creative* works that are marketed to the public, including, but not limited to, novels and other textual works, visual art, television, films, photograph-like imagery, and music. This excludes a significant (and dangerous) use-case for generative AI technology: the creation of works that masquerade as factual, such as false news articles, false voice recordings, or false photorealistic images or video (often referred to as “deepfakes”). Though the use of generative AI to create deepfakes and other factually deceptive material poses significant challenges, they are distinct from the harms considered in this Article. Others have considered this problem separately.<sup>36</sup>

There are many ways generative AI technology can be employed to create works of art and entertainment. The first Section of this Part examines some of these use-cases, focusing in particular on the ways humans can modify or interact with generative AI systems in the service of content production. The second Section considers the market for generative AI-produced creative works, arguing that producers have clear incentives to hide the role of the technology from consumers. Marketplace evidence has shown that consumers are more likely to consume a work if they believe it was human-made. And doctrinal trends in copyright law

---

<sup>34</sup> Sag, *supra* note 33, at 299.

<sup>35</sup> The precise forms of machine learning involved in these technologies can vary significantly. See Lee et al., *supra* note 32 (manuscript at 22–27) (discussing difference between diffusion-based models and transformer-based architectures such as large language modeling (“LLM”)).

<sup>36</sup> See generally Hearing on “Advances in Deepfake Technology” Before the Subcomm. on Cybersecurity, Info. Tech. & Gov’t Innovation of the H. Comm. on Oversight & Accountability, 118th Cong. (2023) (statement of Spencer Overton, Patricia Roberts Harris Rsch. Professor of L., Geo. Wash. Univ.), <https://oversight.house.gov/wp-content/uploads/2023/11/Overton-Testimony-on-Advances-in-Deepfake-Technology-11-8-23-1.pdf> [<https://perma.cc/H37W-JMZG>] (outlining problem and existing scholarship); Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts, at 14–15, COM (2021) 206 final (Apr. 21, 2021) (outlining problem of AI-generated deepfakes).

have made it risky for producers to disclose their use of generative AI, lest they lose copyright protection. The third Section examines the ease with which content producers can obfuscate their use of generative AI.

*A. Generative AI and Creative Content: An Overview*

There are many use cases for generative AI in the creation of works of art and entertainment. Generative AI systems can be used to create unique textual works (such as screenplays or stories), visual images (including highly photorealistic imagery), sound recordings (including original songs), and audiovisual works (including, for example, a perpetually running version of the television show *Seinfeld*).<sup>37</sup>

While it may be relatively easy for a user to prompt a tool like DALL-E to create a new image, humans can also interact with generative AI systems in more complex ways. Katherine Lee, Feder Cooper, and James Grimmelman have helpfully analyzed the “generative AI supply chain,” that is, the “interconnected set of stages that transform training data (millions of pictures of cats) into generations (a new, potentially never-seen-before picture of a cat that has never existed).”<sup>38</sup> The variety of ways human agents can interact with the system at these stages has implications for broader questions of copyrightability discussed below.

Using the concept of the generative AI supply chain, the remainder of this Section examines the different ways that users might employ generative AI systems to produce creative works.

*Model development and fine-tuning.* Most recent discussions of generative AI have focused on corporations—such as OpenAI, Google, and Anthropic—that have trained their own models and released user-facing versions to the public. But this focus on public-facing generative AI systems neglects the ways that content companies can develop their own proprietary generative AI systems, built on their own specific datasets, for their own private use.<sup>39</sup> Consider, for example, the *Marvel Secret Invasion*<sup>40</sup> credit sequence described above. It is highly possible Marvel’s visual effects company, Method Studios, has developed its own

---

<sup>37</sup> Emlyn Travis, *An A.I. Version of Seinfeld Exists and—Yada Yada Yada—It’s Extremely Bizarre*, Ent. Wkly. (Feb. 3, 2023), <https://ew.com/tv/a-i-seinfeld-exists-extremely-surreal/> [<https://perma.cc/T8CE-GALM>].

<sup>38</sup> Lee et al., *supra* note 32 (manuscript at 2).

<sup>39</sup> See *id.* (manuscript at 41–42).

<sup>40</sup> See *supra* notes 1–4 and accompanying discussion.

pecially trained system for creating visual art for television and film projects.<sup>41</sup>

It is also possible that Marvel and Method Studios might have licensed an existing generative AI model and chosen to “fine tun[e]” it to better accomplish the specific tasks required of it. Fine tuning involves a process of “modifying a preexisting, already-trained model,” the goal of which is to take that preexisting model and “mak[e] it better along some dimension of interest.”<sup>42</sup>

In the *Secret Invasion* example, a fine-tuned model might have incorporated various data drawn from Marvel films and comics in order to create a model better equipped to produce imagery specific to Marvel’s needs.

*Prompting and prompt engineering.* “Prompting” is the most straightforward example of how a human user can utilize a fully trained and deployed<sup>43</sup> generative AI system to produce new content. When receiving a prompt from a user, the generative AI system will generate one or multiple textual, visual, audio, or audiovisual “outputs.”<sup>44</sup>

While a prompt can be as simple as “generate the image of a cat,” prompting can also be highly complex. “Prompt engineering” involves the crafting of detailed prompts, specifically designed to elicit material that best meets the user’s needs.<sup>45</sup> Prompting can also be iterative; some generative AI systems allow users to refine, change, or augment an initial output using further prompting.<sup>46</sup>

As explained further below, prompting, as the primary way humans directly interact with a generative AI system to produce a specific work, directly implicates the question of whether generative AI works are protectable by copyright.

*Revising and/or incorporating generated material.* In the broader economic ecosystem in which works are created and marketed to the

---

<sup>41</sup> See Charles Pulliam-Moore, Unfortunately, *Secret Invasion*’s AI Credits Are Exactly What We Should Expect from Marvel, *The Verge* (June 27, 2023), <https://www.theverge.com/2023/6/27/23770133/secret-invasion-ai-credits-marvel> [<https://perma.cc/2DRP-L3QD>].

<sup>42</sup> Lee et al., *supra* note 32 (manuscript at 42–43).

<sup>43</sup> That is, a trained model that has been deployed with an interface that is accessible by users. *Id.* (manuscript at 46–47).

<sup>44</sup> *Id.* (manuscript at 49).

<sup>45</sup> What Is Prompt Engineering?, McKinsey & Co. (Mar. 22, 2024), <https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-prompt-engineering> [<https://perma.cc/MN52-W7B9>] (“Skilled prompt engineers design inputs to interact optimally with other inputs in a generative AI tool.”); Lee et al., *supra* note 32 (manuscript at 66).

<sup>46</sup> See discussion *infra* Subsection I.B.2.

public, the use of generative AI does not occur in a vacuum. Generated material can be further refined and revised by human creators—writers, visual artists, and the like. Indeed, in the *Secret Invasion* example, Marvel claimed that the credit scene was produced using a “highly collaborative and iterative” process, involving a combination of generative AI and “conventional techniques.”<sup>47</sup> Some companies have begun releasing tools explicitly designed to make the process of human editing of generative AI outputs more straightforward; for example, Adobe’s Firefly generative AI system was designed to be integrated into the existing image-editing features of Adobe Photoshop.<sup>48</sup>

While some AI-generated works, such as visual art, can be (for the most part) standalone, some are incorporated into complex combinations of AI and human creativity. This reflects the fact that many works of art or entertainment are the product of numerous creative inputs generally—a film, for example, is considered a singular work but often involves screenwriting, acting, directing, and visual effects produced by numerous different contributors. In such works, generative AI material may be combined with the contributions of human creators to produce a final work.<sup>49</sup> For example, it is easy to imagine a film that incorporates an AI-generated screenplay but still uses human actors.<sup>50</sup> Indeed, a fear of such an arrangement seems to have motivated some of the demands in the 2023 Writers Guild of America (“WGA”) labor dispute.<sup>51</sup> Human and AI creativity may be combined in simpler works as well; for example, the combination of human-written text with a compilation of AI-generated images to create a comic book.<sup>52</sup>

---

<sup>47</sup> Millman, *supra* note 1.

<sup>48</sup> Shelby Putnam Tupper, Hands On: Adobe Firefly in Photoshop, PC Mag., <https://www.pcmag.com/news/hands-on-adobe-firefly-in-photoshop> [<https://perma.cc/3HHW-FJNX>] (last updated Sept. 13, 2023); Pam Clark, Adobe Photoshop Streamlines Power and Precision for the Web, Adobe (Sept. 27, 2023), <https://blog.adobe.com/en/publish/2023/09/27/photoshop-streamlines-power-precision-web> [<https://perma.cc/77U4-SGS8>].

<sup>49</sup> See *infra* Subsection II.C.2 (discussing various examples).

<sup>50</sup> It is possible such a film has already been created. See Rachel Handler, *Irish Wish* Is a Crypto-Fascist, AI-Generated Harbinger of Doom, Vulture (Mar. 15, 2024), <https://www.vulture.com/article/netflix-irish-wish-is-an-ai-generated-harbinger-of-doom.html> (speculating that the movie *Irish Wish* “could have only been written by a malevolently programmed artificial intelligence”).

<sup>51</sup> Jamie Burton, Your Favorite TV Show May Never Be Canceled Again, Thanks to AI, Newsweek (Aug. 22, 2023, 5:00 AM), <https://www.newsweek.com/favorite-tv-show-never-canceled-thanks-artificial-intelligence-wga-sag-strikes-1820173>.

<sup>52</sup> See *infra* Subsection I.B.2 (discussing example of the AI-generated comic book *Zarya of the Dawn*).

Relatedly, generative AI technology can be used to refine and modify existing material, especially images or audiovisual content. The technology can be used, for example, to fill in gaps in images, to remove unwanted elements, or to convert existing material to a higher resolution.<sup>53</sup> This kind of “post-production” editing looks quite similar to the use of more conventional digital tools, but generative AI streamlines and speeds up a process that can otherwise be very labor intensive.<sup>54</sup>

As discussed further below, the combination of human creativity with AI-generated material creates difficult line-drawing questions in determining what should constitute an “AI-generated work,” both for copyrightability purposes and with respect to the broader normative question of when the use of AI should be disclosed.<sup>55</sup>

### *B. Incentives to Hide*

As generative AI is still in its early stages, it remains somewhat unclear how most consumers perceive the technology. Indeed, it is sometimes difficult to disaggregate consumers’ perception of the technology itself from consumers’ frustration with the flaws that many AI-generated works display on their face. Indeed, many AI-generated images still exhibit visual cues that point to the use of the technology—such as uncanny, cartoonish effects—that consumers dislike.<sup>56</sup> But as the technologies improve, it is becoming less and less obvious to the lay observer whether a creative work was produced by conventional means, generative AI, or some combination.<sup>57</sup>

---

<sup>53</sup> Tom Davenport, Cuebric: Generative AI Comes to Hollywood, *Forbes* (Mar. 13, 2023, 3:56 PM), <https://www.forbes.com/sites/tomdavenport/2023/03/13/cuebric-generative-ai-comes-to-hollywood/?sh=55c5e274174b> (discussing these examples).

<sup>54</sup> See *id.*

<sup>55</sup> See *infra* Subsection II.C.2.

<sup>56</sup> Laura Pitcher, Why Is AI Art So Cringe?, *Vice* (Jan. 20, 2023, 4:45 AM), <https://www.vice.com/en/article/m7gynq/why-is-ai-art-so-bad> [<https://perma.cc/V5FS-MSEC>]; Keith Edward, Why Does All AI Art Look Like That?, *Medium* (July 5, 2023), <https://medium.com/@keithkisser/why-does-all-ai-art-look-like-that-f74e2a9e1c87> [<https://perma.cc/PDG5-BURX>].

<sup>57</sup> See Harsha Gangadharbatla, The Role of AI Attribution Knowledge in the Evaluation of Artwork, 40 *Empirical Stud. Arts* 125, 125 (2022) (empirical study showing that most viewers were unable to accurately identify whether an artwork was AI-generated); see also Meaghan Tobin & Cade Metz, China Is Closing the A.I. Gap With the United States, *N.Y. Times* (July 25, 2024), <https://www.nytimes.com/2024/07/25/technology/china-open-source-ai.html> (discussing Kuaishou, a highly realistic video generative AI system).

In a world in which a work itself does not reveal any information about its origins, producers face a choice: disclose their use of generative AI technology or hide it. This Section examines why producers are likely to choose to conceal the role of generative AI in a work's creation, rather than disclose it.

### *1. Marketability*

Especially as the novelty of generative AI technology wanes, choosing not to disclose its role in the creation of a new work may be a sound business decision. There is growing empirical evidence that consumers are simply less likely to choose to purchase a work if they believe that it was produced by generative AI.

Even before the rise of current generative AI systems, one 2009 study used functional MRI technology to measure neurological responses to the same set of abstract artworks that subjects were told were either from a prestigious art gallery or computer-generated. The study clearly found a preference for the gallery works—both in subjects' stated preferences as well as in reactions from the brain's pleasure centers.<sup>58</sup> More recent studies have found similar results in the specific generative AI context. For example, two 2023 studies exposed subjects to the same works of AI-generated visual art, human-made visual art, and AI-generated music, but labeled each as either "human artist" or "AI artist." The results showed a clear bias in favor of works labeled "human artist."<sup>59</sup> Subjects also stated they would be less likely to purchase "AI artist" works.<sup>60</sup> Another study of AI-generated visual artworks across different artistic genres (again labeled "human-created" or "AI-created") also found evidence of preferences for works labeled "human-created."<sup>61</sup>

Notably, almost all the works used in these studies were in fact AI-generated. The label "human" versus "AI" seemed to be doing all of the

---

<sup>58</sup> Ulrich Kirk, Martin Skov, Oliver Hulme, Mark S. Christensen & Semir Zeki, *Modulation of Aesthetic Value by Semantic Context: An fMRI Study*, 44 *NeuroImage* 1125, 1125, 1129–30 (2009).

<sup>59</sup> Kobe Millet, Florian Buehler, Guanzhong Du & Michail D. Kokkoris, *Defending Humankind: Anthropocentric Bias in the Appreciation of AI Art*, 143 *Computs. Hum. Behav.*, 2023, at 1, 3–4, 7; Cinzia Di Dio et al., *Art Made by Artificial Intelligence: The Effect of Authorship on Aesthetic Judgments*, *Psych. Aesthetics Creativity & Arts*, Aug. 2023, at 1, 8.

<sup>60</sup> Millet et al., *supra* note 59, at 6; see also Gangadharbatla, *supra* note 57, at 137 (finding that in the case of non-abstract, representational artwork, respondents were more likely to attribute the work to humans and express more favorable evaluations and purchase intentions).

<sup>61</sup> Bellaiche et al., *supra* note 8, at 5, 7, 14.

work in mediating the subjects' responses.<sup>62</sup> There may be several explanations for this phenomenon. The findings might reflect what psychologists call the “effort heuristic,” the tendency for consumers to view the effort that goes into a work as a proxy for its quality.<sup>63</sup> On this account, it is the ease and speed of creating generative AI works that makes consumers skeptical of their quality. Consumers may also have ingrained anthropocentric biases when it comes to their understanding of what makes something “art.”<sup>64</sup> Consumers may also have ethical or aesthetic preferences that steer them away from AI-generated material. Some of these explanations are discussed in Part II, below.<sup>65</sup>

Whatever the explanation, these studies suggest that when a consumer cannot tell from the face of a work whether or not it is AI-generated, they are more likely to consume it if they believe it to be human-made. This creates a clear financial incentive for producers to hide the role of AI in a work's creation from the public.

## 2. Copyrightability

A second incentive for content producers to hide their use of generative AI relates to ongoing debates about the copyrightability of creative works produced by non-humans.<sup>66</sup> Copyright grants various exclusive rights in creative works, including the right to control copying, distribution, performance, and the creation of “derivative works.”<sup>67</sup> Without copyright protection, the producers of creative works could not convincingly demand any compensation—for purchase or licensing—from the public; all works would be free to use. The Copyright Office, a regulatory body

---

<sup>62</sup> *Id.* at 7; Millet et al., *supra* note 59, at 3, 4.

<sup>63</sup> Justin Kruger, Derrick Wirtz, William Altermatt & Leaf Van Boven, *The Effort Heuristic*, 40 *J. Experimental Soc. Psych.* 91, 91–92 (2004).

<sup>64</sup> See Millet et al., *supra* note 59, at 1–2; Bellaiche et al., *supra* note 8, at 20; see also Gabriel Lima, Assem Zhunis, Lev Manovich & Meeyoung Cha, *On the Social-Relational Moral Standing of AI: An Empirical Study Using AI-Generated Art*, 8 *Frontiers Robotics & AI*, Aug. 2021, at 1, 10 (“Study 1 suggests that nudging participants to think about an AI system's mind could lead to a lower appreciation of AI-generated art. A possible interpretation is that machine creativity is not valued to the same extent as its human counterparts, particularly when AI systems' lack of humanness and mind becomes apparent.”).

<sup>65</sup> See *infra* Subsection II.A.1.

<sup>66</sup> Copyright law is of course implicated at other stages of the generative AI supply chain, including, for example, when determining whether the incorporation of copyrighted works into a data set infringes those works. See generally Sag, *supra* note 33 (considering this and related problems).

<sup>67</sup> 17 U.S.C. § 106.



housed in the Library of Congress, plays an important role in clarifying whether a work is protectable by copyright. One may register any creative work with the Copyright Office—a relatively straightforward process involving an online application—and, if approved, receive a certificate attesting to the work’s protectability.<sup>68</sup> Though registration is not a prerequisite to copyright protection, it is required before the start of infringement litigation.<sup>69</sup> Immediate registration also provides the copyright owner with additional protections.<sup>70</sup> Thus, registration is very common, especially for works that are expected to have any commercial viability.<sup>71</sup>

Courts, and the Copyright Office, have long interpreted the Copyright Act’s protection of “original works of authorship,”<sup>72</sup> as implicitly creating a *human authorship* requirement for copyright protection.<sup>73</sup> Thus, for example, the U.S. Court of Appeals for the Seventh Circuit denied copyright protection for a cultivated garden since, among other reasons, “[a]uthors of copyrightable works must be human; works owing their form to the forces of nature cannot be copyrighted.”<sup>74</sup> The Ninth Circuit found that Naruto, a crested macaque who took several striking “selfie” photos using a photographer’s camera, did not have statutory standing to sue under the Copyright Act.<sup>75</sup> And courts have even noted that claims that a work was divinely written may pose a bar to copyrightability.<sup>76</sup>

The stakes of this issue are high. Without copyright protection, content producers cannot stop infringement, meaning they lack any leverage to demand compensation. Essentially, the work is automatically placed in the public domain. Indeed, Naruto’s famous selfie photos now appear all over the internet, much to the chagrin of the camera owner, who claims

---

<sup>68</sup> Id. § 408. Registration creates a rebuttable presumption of copyright protection. Id. § 410 note (1976) (House Rep. No. 94–1476).

<sup>69</sup> Id. § 412.

<sup>70</sup> See *infra* notes 118–20 (discussing incentives).

<sup>71</sup> Dotan Oliar, Nathaniel Pattison & K. Ross Powell, Copyright Registrations: Who, What, When, Where, and Why, 92 *Tex. L. Rev.* 2211, 2212–16 (2014).

<sup>72</sup> 17 U.S.C. § 102.

<sup>73</sup> *Thaler v. Perlmutter*, 687 F. Supp. 3d 140, 146–47 (D.D.C. 2023) (surveying history), *argued*, No. 23-5233 (D.C. Cir. Sept. 19, 2024).

<sup>74</sup> *Kelley v. Chi. Park Dist.*, 635 F.3d 290, 304 (7th Cir. 2011).

<sup>75</sup> *Naruto v. Slater*, 888 F.3d 418, 420, 426 (9th Cir. 2018) (“Naruto—and, more broadly, animals other than humans—lack statutory standing to sue under the Copyright Act.”).

<sup>76</sup> See *Urantia Found. v. Maaherra*, 114 F.3d 955, 958 (9th Cir. 1997).

that he has lost thousands of dollars in licensing revenue thanks to the court's holding.<sup>77</sup>

The courts' and the Copyright Office's<sup>78</sup> recent forays in the world of generative AI have found that generative AI presents similar barriers to copyrightability as other non-human-creation cases. The primary case addressing this issue, *Thaler v. Perlmutter*, involved Stephen Thaler, the creator of a generative AI system called the "Creativity Machine."<sup>79</sup> Thaler attempted to register a visual artwork produced by the system, asserting that Creativity Machine was the "author" of the work and that he, by virtue of being the creator of Creativity Machine, was the copyright owner.<sup>80</sup> The Copyright Office rejected this argument, finding that work lacked human authorship and thus could not be protected by copyright.<sup>81</sup>

Thaler challenged this outcome at the U.S. District Court for the District of Columbia, but the court agreed with the Copyright Office, finding that "United States copyright law protects only works of human creation" and has "never stretched so far . . . as to protect works generated by new forms of technology operating absent any guiding human hand."<sup>82</sup> Thaler sought to analogize generative AI to technologies like the camera,<sup>83</sup> which the Supreme Court in 1884 found was a tool through which a photographer could produce "an original work of art, the product of [the photographer's] intellectual invention, of which [the photographer] is the author."<sup>84</sup> But the court rejected this analogy, finding that, in the photography cases, "[h]uman involvement in, and ultimate creative control over, the work at issue was key to the conclusion that the new type of work fell within the bounds of copyright."<sup>85</sup>

---

<sup>77</sup> Photographer 'Lost £10,000' in Wikipedia Monkey 'Selfie' Row, BBC (Aug. 7, 2014), <https://www.bbc.com/news/uk-england-gloucestershire-28674167> [<https://perma.cc/6X2D-5TRP>].

<sup>78</sup> The Copyright Office can release letters explaining its reasoning for rejecting an application; these letters can provide some legal guidance to future registrants. U.S. Copyright Off., Circular 20: Requests for Reconsideration 1 (Mar. 2021), <https://www.copyright.gov/circulars/circ20.pdf> [<https://perma.cc/J7N8-YVMQ>].

<sup>79</sup> *Thaler v. Perlmutter*, 687 F. Supp. 3d 140, 142 (D.D.C. 2023).

<sup>80</sup> *Id.* at 143.

<sup>81</sup> *Id.*

<sup>82</sup> *Id.* at 146.

<sup>83</sup> *Id.*

<sup>84</sup> *Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53, 60 (1884).

<sup>85</sup> *Thaler*, 687 F. Supp. 3d at 146. The conclusion that an AI system itself (or the programmer of the AI) cannot be an "author" is consistent with ample scholarship that has made similar arguments. See, e.g., Daniel J. Gervais, *The Machine as Author*, 105 *Iowa L. Rev.* 2053, 2094, 2106 (2020); James Grimmelman, *There's No Such Thing as a Computer-*

The court left open the possibility that with greater “human involvement in the creation of the work” (not present in the facts of this case), copyright could potentially vest in the human *user* of the generative AI system.<sup>86</sup> The Copyright Office has echoed this idea, explaining in recent guidance that

[i]n the case of works containing AI-generated material, the Office will consider whether the AI contributions are the result of “mechanical reproduction” or instead of an author’s “own original mental conception, to which [the author] gave visible form.” The answer will depend on the circumstances, particularly how the AI tool operates and how it was used to create the final work. This is necessarily a case-by-case inquiry.<sup>87</sup>

In several Copyright Office decisions, this general guidance has been somewhat expanded to yield several conclusions.

*An output from a generative AI system is likely not protectable simply because of human creativity at the prompt phase.* In one recent decision, the Copyright Office considered a work of visual art, *Théâtre D’opéra Spatial*.<sup>88</sup> This work had already received some notoriety for winning the 2022 Colorado State Fair’s Fine Arts Competition.<sup>89</sup> The putative author, Jason Allen, attempted to register the work, without disclosing the use of generative AI.<sup>90</sup> But when the Copyright Office caught wind of the art fair controversy, it requested additional information regarding the artwork’s

---

Authored Work—and It’s a Good Thing, Too, 39 Colum. J.L. & Arts 403, 403 (2016); Annemarie Bridy, Coding Creativity: Copyright and the Artificially Intelligent Author, 2012 Stan. Tech. L. Rev. 1, 25–26.

<sup>86</sup> *Thaler*, 687 F. Supp. 3d at 149–50.

<sup>87</sup> Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence, 88 Fed. Reg. 16190, 16192 (Mar. 16, 2023) (to be codified at 37 C.F.R. pt. 202) [hereinafter Copyright Registration Guidance] (citations omitted), [govinfo.gov/content/pkg/FR-2023-03-16/pdf/2023-05321.pdf](https://www.govinfo.gov/content/pkg/FR-2023-03-16/pdf/2023-05321.pdf) [<https://perma.cc/8LMQ-8XZV>]. See generally U.S. Copyright Off., Copyright and Artificial Intelligence—Part 2: Copyrightability (Jan. 2025) [hereinafter Copyright and Artificial Intelligence], <https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf> [<https://perma.cc/FM7N-S6BK>] (expanding on this guidance).

<sup>88</sup> Letter from Suzanne V. Wilson, Maria Strong & Jordana Rubel, U.S. Copyright Off. Rev. Bd., to Tamara Pester, Tamara S. Pester, LLC, Second Request for Reconsideration for Refusal to Register *Théâtre D’opéra Spatial* 1 (Sept. 5, 2023), <https://www.copyright.gov/rulings-filings/review-board/docs/Theatre-Dopera-Spatial.pdf> [<https://perma.cc/HKL5-8HSB>].

<sup>89</sup> *Id.* at 2.

<sup>90</sup> *Id.*

provenance.<sup>91</sup> Allen claimed that he “input[ted] numerous revisions and text prompts at least 624 times to arrive at the initial version of the image” using Midjourney, and then made some changes to the image using Adobe Photoshop.<sup>92</sup> The Office ultimately rejected Allen’s request for copyright registration<sup>93</sup> and went on to explain that “[i]f all of a work’s ‘traditional elements of authorship’ were produced by a machine, the work lacks human authorship, and the Office will not register it.”<sup>94</sup>

Allen attempted to argue that his series of detailed prompts evinced human creativity that was more than *de minimis* and thus rendered him the author of the final image.<sup>95</sup> But the Office rejected this argument, acknowledging “that the process of prompting can involve creativity . . . . But that does not mean that providing text prompts to Midjourney ‘actually form[s]’ the generated images,” such that Allen could be considered the author.<sup>96</sup>

From this and similar decisions,<sup>97</sup> the Copyright Office appears to have arrived at a general principle that creativity at the prompt phase is not itself sufficient to confer protection on an output. The Copyright Office guidance explains: “[W]hen an AI technology receives solely a prompt from a human and produces complex written, visual, or musical works in response, the ‘traditional elements of authorship’ are determined and executed by the technology—not the human user” and the AI-generated work is not protectable.<sup>98</sup>

---

<sup>91</sup> *Id.*

<sup>92</sup> *Id.*

<sup>93</sup> *Id.* at 1.

<sup>94</sup> *Id.* at 4.

<sup>95</sup> *Id.* at 3.

<sup>96</sup> *Id.* at 7.

<sup>97</sup> See, e.g., Letter from Suzanne V. Wilson, Maria Strong & Mark T. Gray, U.S. Copyright Off. Rev. Bd., to Alex P. Garens, Day Pitney, LLP, Second Request for Reconsideration for Refusal to Register SURYAST 8 (Dec. 11, 2023), <https://www.copyright.gov/rulings-filings/review-board/docs/SURYAST.pdf> [<https://perma.cc/32X5-F3YC>] (holding that the use of an AI system that can redesign an image in a new “style” does not yield a copyrightable work because the human contribution is *de minimis*).

<sup>98</sup> Copyright Registration Guidance, *supra* note 87, at 16192; see also Copyright and Artificial Intelligence, *supra* note 87, at 18–22 (examining the issue in more detail and confirming the conclusion that prompting alone cannot yield protectability). This element of the Copyright Office’s approach has been particularly controversial. Some scholars have suggested that detailed prompt engineering should be sufficient to yield copyright protection in the output. See, e.g., Letter from Pamela Samuelson, Matthew Sag & Christopher Sprigman to Suzanne V. Wilson, Gen. Couns. & Assoc. Reg. of Copyrights, U.S. Copyright Off., Comments in Response to the Copyright Office’s Notice of Inquiry on Artificial Intelligence and Copyright 3 (Oct. 30, 2023) (“[T]here is no reason in principle why prompts couldn’t be

*Compilations of outputs may yield (thinly) protectable works.* In a second recent decision, the Office considered a comic book written by Kristina Kashtanova titled *Zarya of the Dawn*.<sup>99</sup> Kashtanova used Midjourney to generate images and then compiled them into a comic book, with the addition of text.<sup>100</sup> Kashtanova did not disclose the use of AI in their application, and the Office registered the work.<sup>101</sup> But after becoming aware of social media statements describing their use of AI, the Office canceled the original registration and replaced it with a registration that significantly narrowed the scope of the copyright.<sup>102</sup>

For reasons similar to the *Théâtre D’opéra Spatial* decision, the Office found that the images themselves were not protectable since they were not products of human authorship.<sup>103</sup> The Office, however, was willing to confer some copyright protection on the work as a whole based on Kashtanova’s additions to these images. First, the Office noted that the text of the comic book was certainly registrable, as this was a “product of human authorship” that was sufficiently creative.<sup>104</sup> The Office also found that “selection and arrangement” of the images and text could confer copyright protection on the work as a whole.<sup>105</sup> Applying familiar principles regarding the use of public-domain material in new works, the Office concluded that *Zarya of the Dawn* was protectable as a kind of compilation.<sup>106</sup> In practice, such compilation copyrights are considered thin, in that they generally protect only verbatim or near-verbatim copies of the compilation itself, rather than works that borrow from the compiled

---

detailed enough to meet the traditional threshold of authorship in some cases. This will depend on the circumstances. Sophisticated prompts that specify details of an image should be sufficient to meet the requirement that the work that results from and reflects a person’s original conception of the expression.”); Mark Lemley, *How Generative AI Turns Copyright Upside Down*, 25 Colum. Sci. & Tech. L. Rev. 190, 201–02 (2024); Lee et al., *supra* note 32 (manuscript at 59–60); see also Dan L. Burk, *Thirty-Six Views of Copyright Authorship*, by Jackson Pollock, 58 Hous. L. Rev. 263, 268, 317–18 (2020) (discussing how “concepts of actual and proximate cause, intent, and volition” can be used to determine whether an individual’s input into a work amounts to copyright authorship).

<sup>99</sup> Kasunic, *supra* note 11, at 1, 12.

<sup>100</sup> *Id.* at 2.

<sup>101</sup> *Id.*

<sup>102</sup> *Id.* at 1–2, 12; *Zarya of the Dawn*, Registration No. TXu002356581 (effective Sept. 15, 2022).

<sup>103</sup> Kasunic, *supra* note 11, at 1, 12.

<sup>104</sup> *Id.* at 4.

<sup>105</sup> *Id.* at 5.

<sup>106</sup> *Id.*

material.<sup>107</sup> For example, the copyright interest identified by the Copyright Office likely would not prevent third parties from making copies of specific images in the comic book or from creating their own sequels that utilize the characters.<sup>108</sup>

From this decision, the Office has arrived at the general principle that “a human may select or arrange AI-generated material in a sufficiently creative way that ‘the resulting work as a whole constitutes an original work of authorship’”; however, “copyright will only protect the human-authored aspects of the work, which are ‘independent of’ and do ‘not affect’ the copyright status of the AI-generated material itself.”<sup>109</sup>

*Direct human changes to outputs may yield protectable works, but the required degree of change is unclear.* In both of the decisions discussed above, the Copyright Office also considered whether direct edits or changes to generated content by a human author may also be sufficient to confer copyrightability. In the *Théâtre D’opéra Spatial* registration dispute, the Office noted the possibility that Allen’s revisions to the Midjourney-produced image could be sufficiently creative to confer a thin copyright that “exclude[d the] non-human authorship elements,” but ultimately could not register such a copyright because Allen refused to disclaim the Midjourney-created material.<sup>110</sup> In *Zarya of the Dawn*, however, Kashtanova also attempted to argue that certain images could be protectable because of post-facto changes using Photoshop.<sup>111</sup> But one such change, involving “modif[ication to] the rendering of Zarya’s lips and mouth” was deemed “too minor and imperceptible to supply the necessary creativity for copyright protection.”<sup>112</sup> Thus, while it remains conceivable that a more comprehensive set of modifications could confer

---

<sup>107</sup> See *Feist Publ’ns, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340, 349 (1991) (“[T]he copyright in a factual compilation is thin. Notwithstanding a valid copyright, a subsequent compiler remains free to use the facts contained in another’s publication . . . so long as the competing work does not feature the same selection and arrangement.”).

<sup>108</sup> Cf. *Key Publ’n, Inc. v. Chinatown Today Publ’g Enters.*, 945 F.2d 509, 514 (2d Cir. 1991) (“[T]he components of a compilation are generally in the public domain, and a finding of substantial similarity or even absolute identity as to matters in the public domain will not suffice to prove infringement.”); *Apple Comput., Inc. v. Microsoft Corp.*, 35 F.3d 1435, 1442 (9th Cir. 1994) (noting that “only ‘thin’ protection, against virtually identical copying, is appropriate”).

<sup>109</sup> Copyright Registration Guidance, *supra* note 87, at 16192–93; see also Copyright and Artificial Intelligence, *supra* note 87, at 24–27 (providing additional detail).

<sup>110</sup> Wilson et al., *supra* note 88, at 7.

<sup>111</sup> Kasunic, *supra* note 11, at 10.

<sup>112</sup> *Id.* at 10–11.

some protection<sup>113</sup>—perhaps protection sufficient to protect direct copying of the images—the Office has yet to supply clear examples of what kind of changes are required.

\* \* \*

The Copyright Office has clarified that “applicants have a duty to disclose the inclusion of AI-generated content in a work submitted for registration and to provide a brief explanation of the human author’s contributions to the work.”<sup>114</sup> But, in light of the courts’ and Copyright Office’s current approach to the authorship requirement, few applicants will have much of a motivation to disclose. As outputs are generally unprotectable, regardless of the level of creativity at the prompt phase, most AI-generated works will not be protectable at all. Compilations may be protectable, but such a thin copyright would not prevent direct copying of specific images or pieces of text in the compilation. And while it is possible that direct changes by a human author to an output can yield copyrightability, the Copyright Office has not provided clear guidance about the degree of change that is necessary.

The stakes of failing to register a copyright are high; in the absence of registration, applicants would be required to turn to the courts (a costly prospect) in order to establish a valid copyright. And, in light of the *Thaler v. Perlmutter* decision, it seems that the courts have, at least for now, adopted the same approach as the Copyright Office. In the absence of copyright protection, or if granted only a thin compilation copyright, producers would face significant hurdles to monetizing their creations. The inability to prevent others from selling the exact same work yields a race to the bottom, where the cost of the work ultimately approaches the cost of producing new copies—which is zero in a digital environment. Indeed, it is for this reason that most public domain works have little market value.<sup>115</sup>

The risks of losing copyright in potentially lucrative books, television shows, films, and works of art would likely motivate most producers to simply fail to disclose the role of AI in the work’s creation. The next

---

<sup>113</sup> Copyright Registration Guidance, *supra* note 87, at 16192–93 (“[A]n artist may modify material originally generated by AI technology to such a degree that the modifications meet the standard for copyright protection.”).

<sup>114</sup> *Id.* at 16193.

<sup>115</sup> See generally Paul J. Heald, *The Public Domain*, in *Handbook on the Economics of Copyright* 93 (Richard Watt ed., 2014) (discussing theoretical and empirical research regarding the public domain).

Section explains how this might happen in practice, and sets out the few legal mechanisms available to ensure compliance with Copyright Office disclosure requirements.

### *C. Obfuscation in Practice*

As the last Section explained, producers of new creative content have strong incentives to hide their use of generative AI. Revealing this information to the public might jeopardize sales, and revealing it to the Copyright Office might jeopardize the ability to secure copyright protection.

To be clear, public disclosure and disclosure to the Copyright Office are interrelated. Copyright Office registrations are publicly available and easily searchable. Thus, if a registrant reveals information about the use of AI to the Copyright Office, this information is available to the public. Indeed, several very recent publicly available registrations specifically note that material in the registered work that was “generated by artificial intelligence” is excluded from protectability.<sup>116</sup> Similarly, revealing information to the public but attempting to hide it from the Copyright Office is also unfeasible. Indeed, in the *Zarya of the Dawn* dispute, the Copyright Office discovered the use of AI from Kashtanova’s social media marketing, even though they had failed to reveal their use of the technology in their application.<sup>117</sup>

It is possible that copyright registrations will begin to decline as content producers realize the high bar to registering AI-generated work. This, however, would be a limited strategy for a content producer; registration is a prerequisite to any infringement action, meaning any purported rightsholder would need to attempt to register their work if they ever sought to stop infringement.<sup>118</sup> Furthermore, copyright owners are strongly incentivized to register well before any possible infringement. Acts of infringement that occurred after registration can result in an award of statutory damages or attorneys’ fees.<sup>119</sup> And early registration provides

---

<sup>116</sup> See, e.g., Kegio: Prayer Warrior, Registration No. TXU002396886 (effective Sept. 30, 2023) (registering only text accompanying an image, and not the “artwork generated by artificial intelligence”).

<sup>117</sup> See *supra* note 11 and accompanying text.

<sup>118</sup> 17 U.S.C. § 411.

<sup>119</sup> *Id.* § 412.



a copyright owner with certain evidentiary advantages when litigating an infringement claim.<sup>120</sup>

It is more likely that rightsholders will continue to register but simply fail to disclose their use of AI to the Copyright Office. Although the Office has stated that “applicants have a duty to disclose the inclusion of AI-generated content in a work submitted for registration and to provide a brief explanation of the human author’s contributions to the work,”<sup>121</sup> it has few tools to enforce disclosure. The Copyright Act criminalizes making “a false representation of a material fact” in a registration application, but the penalty is a mere \$2,500 fine.<sup>122</sup> The provision is almost never enforced.<sup>123</sup> In a situation where a copyright owner brings an infringement claim, a defendant may also assert “fraud on the Copyright Office,” if there is evidence of inaccuracies in the registration application.<sup>124</sup> But fraud on the Copyright Office is a “heavy burden” to meet, requiring a showing that the omission was material and that copyright owner intended to defraud.<sup>125</sup> Moreover, copyright owners are generally permitted to refile their registration with the error corrected, and even recommence litigation.<sup>126</sup> The availability of fraud on the Copyright Office as a defense thus provides little assurance that accurate registration information will be submitted.<sup>127</sup>

Indeed, the ease with which registrants can hide information regarding a work’s AI origins can be seen in both the *Théâtre D’opéra Spatial* and *Zarya of the Dawn* disputes. In both of these instances, the ostensible

<sup>120</sup> Id. § 410(c) (“In any judicial proceedings the certificate of a registration made before or within five years after first publication of the work shall constitute prima facie evidence of the validity of the copyright and of the facts stated in the certificate.”).

<sup>121</sup> Copyright Registration Guidance, supra note 87, at 16193.

<sup>122</sup> 17 U.S.C. § 506(e). A similarly paltry fine is also in place for the fraudulent use of a copyright notice on material that is not protectable. Id. § 506(c).

<sup>123</sup> Jason Mazzone, Copyfraud, 81 N.Y.U. L. Rev. 1026, 1036 (2006).

<sup>124</sup> 5 William F. Patry, Patry on Copyright § 17:126, Westlaw (database updated Sept. 2024).

<sup>125</sup> 2 Melville B. Nimmer & David Nimmer, Nimmer on Copyright § 7.20[B][1] (2024); cf. *Unicolors, Inc. v. H&M Hennes & Mauritz, L. P.*, 142 S. Ct. 941, 945 (2022) (finding that, for purposes of the requirement to register before commencing litigation, either a mistake of law or of fact can excuse an error in a copyright registration application).

<sup>126</sup> 5 Patry, supra note 124, § 17:126; see also, e.g., *Morgan, Inc. v. White Rock Distilleries, Inc.*, 230 F. Supp. 2d 104, 108–09 (D. Me. 2002) (finding problematic mistakes in a copyright registration, but nonetheless dismissing complaint without prejudice “since a copyright owner may register his claim at any time during the life of the copyright”).

<sup>127</sup> Part III, *infra*, examines some ways that other doctrines, especially copyright misuse, could be leveraged to provide greater assurance.

copyright owners failed to disclose and disclaim the AI-generated materials in their works.<sup>128</sup> The Copyright Office only became aware of the issue because the registrants later chose to publicly reveal their use of the technology: *Théâtre D'opéra Spatial* having won an art competition,<sup>129</sup> and *Zarya of the Dawn* having been discussed on social media.<sup>130</sup> Had this information not been disclosed, it seems unlikely that the Office would have ever discovered the truth.

None of this is to say that existing law is completely toothless in dealing with producers who hide their use of AI from the public. Copyright, as well as trademark, right of publicity, and competition law, all contain some legal mechanisms that could potentially be used to force out this information by raising the costs of obfuscation. After the next Part diagnoses the specific harms of failing to disclose AI authorship, Part III examines these areas of existing law and how they might be used to help accomplish this goal.

## II. THE HARMS OF HIDDEN AI AUTHORSHIP

When the law (and copyright law in particular) considers the public's relationship to art and entertainment, the assumption is generally that consumers choose whether to purchase a creative work solely based on the enjoyment they derive from the work itself.<sup>131</sup> But this perspective ignores what Douglas Kysar has called “preferences for processes.”<sup>132</sup> Process preferences are ubiquitous in other areas of economic life; a consumer may not know whether coffee was produced using ethical labor

---

<sup>128</sup> Wilson et al., supra note 88, at 1–2; Kasunic, supra note 11, at 1–2.

<sup>129</sup> Wilson et al., supra note 88, at 2.

<sup>130</sup> Kasunic, supra note 11, at 2.

<sup>131</sup> See Joseph P. Liu, Copyright Law's Theory of the Consumer, 44 B.C. L. Rev. 397, 403 (2003) (discussing the view that “consumers see copyrighted works as commodities”); see also Michael Abramowicz, A New Uneasy Case for Copyright, 79 Geo. Wash. L. Rev. 1644, 1646 (2011) (examining how copyrighted works can be substitutes for one another). See generally William W. Fisher III, Reconstructing the Fair Use Doctrine, 101 Harv. L. Rev. 1659 (1988) (examining assumptions that tend to underlie economic analyses of copyright law).

<sup>132</sup> Kysar, supra note 14, at 529 (“[C]onsumer preferences may be heavily influenced by information regarding the manner in which goods are produced. Such information—which this Article refers to as ‘process information’—can include the labor conditions of workers who produce a consumer good, the environmental effects of a good's production, the use of controversial engineering techniques such as genetic modification to create a good, or any number of other social, economic, or environmental circumstances that are related causally to a consumer product, but that do not necessarily manifest themselves in the product itself.” (footnote omitted)).

practices simply based on the taste of the coffee, but this information might still play an important role in a purchasing decision. Indeed, as Kysar notes, processes can implicate a range of different consumer concerns that are separate from the utility derived from the product itself.<sup>133</sup> Most obviously, this can include a desire to instrumentally influence production processes that consumers find objectionable.<sup>134</sup> But, more subtly, process preferences can also reflect consumers' desire to express a voice in public policy-making<sup>135</sup> or to further ethical commitments that are essential to self-identity.<sup>136</sup>

Building on the concept of process preferences, this Part argues that some consumers appear to have strong preferences for human-made creative works, that these preferences are worth taking seriously, and that, accordingly, the hiding of information regarding a work's AI provenance can constitute a kind of deceptive harm.

I make this case over three separate Sections. First, I examine the nature of consumer process preferences for non-AI-generated works, exploring how the question of AI authorship implicates broader consumer concerns over ethically produced goods, authenticity, and artist-fan relationships. Second, I argue that these preferences are worth taking seriously. Broader theoretical accounts of the value of human authorship generally suggest that a world of ubiquitous undisclosed AI authorship is one in which a social, ethical, and political role for art would be diminished. Third, I examine several implications of the conclusion that hidden AI authorship is harmful by exploring the difference between human ghostwriting and hidden AI authorship, the line between AI-as-author and AI-as-tool, and the possible value that AI-generated works might have for at least some consumers.

---

<sup>133</sup> *Id.* at 581–601.

<sup>134</sup> *Id.*

<sup>135</sup> See, e.g., *id.* at 604 (“[I]f, for instance, consumers believe that shade-grown coffee production helps to protect migratory birds—then process-labeled products provide consumers with a vehicle for expressing their belief in the overall public policy significance of the public good. Utility for the consumer on this account is not necessarily derived from effecting change in the world—from actually saving migratory birds—but from participating in a process whereby one is able to express a ‘vote’ in favor of such change, whether or not it actually occurs.” (footnote omitted)).

<sup>136</sup> See, e.g., *id.* at 616–17 (citing example of kosher products and other consumption choices that are “directed toward personal moral practice” (emphasis omitted)).

*A. Consumer “Process Preferences” for Human-Authored Content*

AI-generated content appears to be an arena where information about a work’s process of creation can play a meaningful role in a consumer’s consumption decision. As noted in Part I, there have now been several empirical studies that suggest that consumers value a creative work more when they believe it was human-made, compared to AI-generated.<sup>137</sup>

This Section highlights three broader areas of consumer concern that appear to be at play in consumers’ preferences for human-created works over AI-generated works: ethical concerns about labor markets, beliefs about “authenticity,” and fan-based commitments to specific artists. These three perspectives can help conceptualize why information about the process that yielded a creative work is meaningful to a consumer’s consumption choice, even if a work cannot be distinguished as human-made or AI-generated on its face.<sup>138</sup>

*1. Ethics*

The 2023 simultaneous SAG-AFTRA and WGA strikes caused significant disruption to the entertainment industry and an estimated \$6 billion in losses.<sup>139</sup> As many sources reported, a major impetus for both strikes was the growing use of generative AI by large Hollywood studios.<sup>140</sup> The WGA sought guarantees that writers’ jobs would not be lost to AI,<sup>141</sup> and SAG-AFTRA sought protections designed to prevent

---

<sup>137</sup> See supra notes 58–64 and accompanying text.

<sup>138</sup> Cf. Kysar, supra note 14, at 532 (“[C]onsumer products—even when physically indistinguishable—are not perfect substitutes to the extent that they are produced using different processes about which consumers have strong feelings.”).

<sup>139</sup> Natasha Chen & Jack Hannah, Entertainment Strikes Pushing Toward \$6 Billion in Losses: ‘It Just Gets Worse Each Day,’ CNN, <https://www.cnn.com/2023/09/20/entertainment/entertainment-strikes-6-billion> (last updated Sept. 21, 2023, 5:32 PM).

<sup>140</sup> See Kalia Richardson, One Year After the Actors’ Strike, AI Remains a Persistent Threat, *Rolling Stone* (July 14, 2024), <https://www.rollingstone.com/tv-movies/tv-movie-features/actors-strike-sag-aftra-ai-one-year-later-1235059882/>; see also Noam Scheiber, Will A.I. Upend White-Collar Work? Consider the Hollywood Editor, *N.Y. Times* (July 30, 2024), <https://www.nytimes.com/2024/07/30/business/economy/artificial-intelligence-hollywood-unions.html> (reporting on additional threats AI poses to traditional human jobs within Hollywood studios).

<sup>141</sup> Alissa Wilkinson & Emily Stewart, The Hollywood Writers’ Strike Is Over—and They Won Big, *Vox* (Sept. 28, 2023, 9:45 AM), <https://www.vox.com/culture/2023/9/24/23888673/wga-strike-end-sag-aftra-contract> [<https://perma.cc/7585-STUZ>].

the unauthorized use of actors' likenesses using generative AI.<sup>142</sup> A recent strike by video game actors is seeking similar guarantees that actors' work will not be supplanted by generative AI.<sup>143</sup>

Over the course of these strikes, the actors' and writers' demands surrounding the use of generative AI were treated sympathetically by many members of the general public, suggesting that many consumers view the disruption of creative labor markets by generative AI as a problem. The *Secret Invasion* controversy described above is illustrative; after the use of generative AI in the credit scene was publicly disclosed, many objected to Marvel's use of the technology in light of the potential supplanting of graphic designer jobs.<sup>144</sup> The musical artist, Washed Out, released the first entirely AI-generated music video and faced similar backlash.<sup>145</sup>

Commentators have also begun noting in general terms that consumption of AI-generated works implicates questions about the equitable compensation of artists and other creative professions.<sup>146</sup> Such ethical concerns are arguably an outgrowth of growing scrutiny of labor practices in the creative industries in general.<sup>147</sup> Thanks to public controversies involving artists and the corporations that commonly own

<sup>142</sup> Wes Davis, *The Screen Actors Guild's Strike-Ending Deal Has Entered Its Final Step*, *The Verge* (Nov. 11, 2023, 1:24 PM), <https://www.theverge.com/2023/11/11/23956660/sag-aftra-actors-strike-deal-approved-streaming-revenue-ai> [<https://perma.cc/5PHG-2UL7>].

<sup>143</sup> Mandalit del Barco, *AI Is Changing Video Games—and Striking Performers Want Their Due*, *NPR* (Aug. 14, 2024, 12:27 PM), <https://www.npr.org/2024/08/14/nx-s1-5072638/video-game-strike-ai-animation-sag-aftra> [<https://perma.cc/B78Q-C9YK>].

<sup>144</sup> Watercutter, *supra* note 2; Pulliam-Moore, *supra* note 41.

<sup>145</sup> Brian Hiatt, *Washed Out Made an AI Music Video. The Backlash Was Swift*, *Rolling Stone* (May 9, 2024), <https://www.rollingstone.com/music/music-features/washed-out-a-i-video-backlash-hardest-part-interview-1235017826/>.

<sup>146</sup> Samuel Shaibu, *AI-Generated Art: The Ethical Implications and Debates, Becoming Hum.: AI Mag.* (Mar. 24, 2023), <https://becominghuman.ai/ai-generated-art-the-ethical-implications-and-debates-6f0132d158c7> [<https://perma.cc/GS6C-ZESC>] (“As AI-generated art becomes more popular and accessible, there’s a concern that it could lead to a decline in demand for human artists and result in the loss of jobs in the art industry.”); see also Electra Nanou, *The Ethical Pros and Cons of AI Art Generation, MakeUseOf* (July 6, 2023), <https://www.makeuseof.com/ai-art-generation-ethical-pros-cons/> [<https://perma.cc/8Y98-C33Y>] (“Besides AI engines and their developers taking advantage of copyrighted art without repercussions, artists also have to compete with free and easy-to-use art-generating software that makes potential customers turn to AI instead of human artists.”).

<sup>147</sup> See generally Xiyin Tang, *Intellectual Property Law as Labor Policy*, *N.Y.U. L. Rev.* (forthcoming), <https://ssrn.com/abstract=4761809> [<https://perma.cc/V2YY-LBR6>] (examining role of labor concerns in IP and arguing that individual laborers must be better prioritized in IP-intensive industries).

copyrights in creative works—record labels, book publishers, movie studios—some consumers have begun factoring underlying labor conditions into their decision whether to consume a specific work.<sup>148</sup> Generative AI shines a spotlight on these perceived inequities by showcasing just how dispensable human labor is in the creative-content industrial model.

## 2. Authenticity

Understanding why consumers might prefer human-made works on ethical grounds is relatively straightforward. But there are more subtle, aesthetics-oriented reasons why consumers might prefer to refrain from consuming an AI-generated work. Understanding these aesthetic preferences requires unpacking a murky but pervasive concept in debates about art and culture: authenticity.

As Amy Adler notes, “[t]he term ‘authenticity’ signals two complex, overlapping concepts in art. First, it signals originality, usually but not exclusively in the sense of uniqueness. Second, authenticity signals authorship—an authentic work is ‘by’ an artist and can be attributed to her.”<sup>149</sup> Both of these conceptions of authenticity are often implicated when society seeks to distinguish between new uses of “mechanical” or machine-driven forms of production and older methods.<sup>150</sup>

---

<sup>148</sup> For example, Taylor Swift has publicly discussed her fraught relationship with her prior record label and the current owners of many of her sound recording copyrights (“master” recordings). Many consumers now focus on the “Taylor’s Version” recordings, which are owned by Swift herself. Hugh McIntyre, Taylor Swift’s Re-Recorded Albums Are Huge Successes—But They Were a Real Risk, *Forbes* (July 27, 2023, 9:00 AM), <https://www.forbes.com/sites/hughmcintyre/2023/07/27/taylor-swifts-re-recorded-albums-are-huge-successes-but-they-were-a-real-risk/?sh=71091ed41074>.

<sup>149</sup> Amy Adler, *Artificial Authenticity*, 98 *N.Y.U. L. Rev.* 706, 715 (2023) (footnote omitted).

<sup>150</sup> See, e.g., Stefan Bechtold & Christopher Jon Sprigman, *Intellectual Property and the Manufacture of Aura*, 36 *Harv. J.L. & Tech.* 291, 291, 336–43 (2023) (examining nonfungible tokens (“NFTs”) through the lens of debates about authenticity); see *infra* note 153 (discussing the authenticity of handmade artisanal products). The link between authenticity and machine-based production is often attributed to Walter Benjamin’s 1936 essay, *The Work of Art in the Age of Mechanical Reproduction*. In that essay, Benjamin examines the effect of easy “mechanical” reproducibility on the aesthetic experience (as well as commercial value and political role) of works of art. Walter Benjamin, *The Work of Art in the Age of Mechanical Reproduction*, in *Illuminations* 217, 221 (Hannah Arendt ed., Harry Zohn trans., Schocken Books 1969) (1955); see also Barton Beebe, *Intellectual Property Law and the Sumptuary Code*, 123 *Harv. L. Rev.* 809, 842, 868 n.337 (2010) (noting links between Benjamin’s work and modern conceptions of authenticity).

Along these lines, Dan Burk has directly confronted how understandings of authenticity are likely to affect public perception of works created by generative AI.<sup>151</sup> Burk notes that the perception that only human-made works can be “authentic” is far older than the generative AI debate, going back to the rise of industrialization generally.<sup>152</sup> For example, the introduction of machine-produced furniture and ornamental goods during the Victorian Era led many to attach newfound importance to hand-crafted goods.<sup>153</sup> Burk predicts that the same will be true for AI-generated content: “[A]s the supply of such [machine-produced] artifacts becomes routine and settled, society begins to place an increasing value on the fruits of human production, imperfect and inefficient as they may be.”<sup>154</sup>

Indeed, recent empirical work suggests that consumers’ preference for human-created works on authenticity grounds is not just speculative. Several of the studies described in the last Part, which found clear preferences for works of visual art and songs labeled “human-made” (even if the works were actually AI-generated) noted that human-labeled works induced greater feelings of “awe”<sup>155</sup> and “[b]eauty, [p]rofundity, and [w]orth”<sup>156</sup> compared to AI-labeled works. Interestingly, one study concluded that when subjects were asked questions prompting them to think about whether a generative AI system could be a moral or artistic agent—i.e., questions regarding the system’s ability to “create” or “experience”—they were then *less likely* to assign a high monetary value to works produced by generative AI.<sup>157</sup> This study implies that directly confronting consumers with the mechanical nature of AI generation—as

---

<sup>151</sup> Dan L. Burk, Cheap Creativity and What It Will Do, 57 Ga. L. Rev. 1669, 1686–87 (2023).

<sup>152</sup> Id. at 1686, 1690 (emphasis omitted).

<sup>153</sup> Id. at 1683–87; see also Krzysztof Pelc, AI Will Make Human Art More Valuable, Wired (Mar. 16, 2023, 9:00 AM), <https://www.wired.com/story/art-artificial-intelligence-history/> (summarizing the success of English artists selling handmade furniture during the Victorian era and comparing that phenomenon with AI).

<sup>154</sup> Burk, supra note 151, at 1682.

<sup>155</sup> Millet et al., supra note 59, at 2, 7 (“Awe concerns feelings of wonder, sublime or impressiveness.”).

<sup>156</sup> Bellaiche et al., supra note 8, at 7.

<sup>157</sup> Lima et al., supra note 64, at 7–8; see also Joo-Wha Hong & Nathaniel Ming Curran, Artificial Intelligence, Artists, and Art: Attitudes Toward Artwork Produced by Humans vs. Artificial Intelligence, 15 ACM Transactions on Multimedia Computing Commc’ns & Applications, July 2019, at 1, 11 (“[P]eople with the stereotype ‘AI cannot produce art’ gave significantly lower ratings [to works labeled as AI-generated] compared to people without the stereotype.”).

contrasted with prevailing beliefs about the human artistic process—leads them to further discount the value of AI-generated works.

### 3. *Fandom*

Just about any AI-generated work of art or entertainment can implicate consumers' beliefs about ethics or authenticity. But there is also a subset of works that implicate a third area of importance to consumers: a preference for works by a *specific author* born out of a sense of fan loyalty.

There are now documented examples of AI being used to mimic the voice, likeness, or style of a specific artist. In one example, a song using AI-generated vocals that compellingly mimicked the artists Drake and The Weeknd went viral before being taken down.<sup>158</sup> While the song's use of AI was not hidden, it was of high enough quality that consumers could easily have been misled. Ongoing litigation brought by several visual artists against various generative AI systems allege, among other things, that generative AI systems can “generate images in the style of particular artists,” such that consumers might choose to consume AI-generated works because they believe these works were produced by their preferred artists.<sup>159</sup>

Of course, unlike more general uses of generative AI, a work that uses the voice or likeness of an artist or mimics a distinctive style plausibly harms a specific person: the imitated artist herself. Indeed, there are several areas of law that conceivably allow artists to police the unauthorized use of their voice, likeness, or style by a generative AI system. Some states' right of publicity statutes may allow an artist to sue for appropriation of her voice or likeness.<sup>160</sup> Trademark and unfair competition law may also protect an artist from the use of her name or

---

<sup>158</sup> Joe Coscarelli, An A.I. Hit of Fake ‘Drake’ and ‘The Weeknd’ Rattles the Music World, *N.Y. Times* (Apr. 14, 2023), <https://www.nytimes.com/2023/04/19/arts/music/ai-drake-the-weeknd-fake.html>.

<sup>159</sup> First Amended Complaint at 72–73, *Andersen v. Stability AI Ltd.*, No. 23-cv-00201 (N.D. Cal. Nov. 29, 2023).

<sup>160</sup> Jennifer E. Rothman, Univ. of Pa. Carey L. Sch., Comment Letter on Artificial Intelligence, Copyright, and Right of Publicity 3 n.7 (Oct. 25, 2023), [https://rightofpublicityroamap.com/wp-content/uploads/2023/10/Prof-Rothman-Comments-to-Copyright-Office-on-Right-of-Publicity-and-AI\\_October-2023.pdf](https://rightofpublicityroamap.com/wp-content/uploads/2023/10/Prof-Rothman-Comments-to-Copyright-Office-on-Right-of-Publicity-and-AI_October-2023.pdf) [<https://perma.cc/9K6L-NDLZ>].



likeness in an AI-generated work if the use suggests the artist endorsed or participated in the work.<sup>161</sup>

The question of style is more complicated; some case law has tacitly endorsed copyright protection for style but such holdings have been criticized for overextending copyright.<sup>162</sup> In *Andersen v. Stability AI Ltd.*, in which a group of artists alleged that their visual style was reproducible by generative AI systems, the plaintiffs appeared to rely on a novel combination of copyright and trade dress protection to make their claim; their amended complaint stated that generative AI was being used to create “substantially similar substitutes . . . that imitate the trade dress of particular artists—including Plaintiffs. This is already damaging the market for Plaintiffs’ artwork and labor . . . .”<sup>163</sup>

These various causes of action are specific to the impersonated artist herself—reflecting the specific harm to her brand, personhood, or established markets—and cannot be asserted by a deceived consumer. But consumers also have a specific interest in knowing whether a work is, indeed, the product of a favored artist. Defining the nature of this harm raises familiar questions: if a consumer enjoys, for example, a song that features the AI reproduction of Drake’s voice, does it truly matter whether or not they know that song was actually sung by him?

Answering this question requires considering the unique relationships that consumers have with specific artists. In its simplest terms, as several commentators have noted, many consumers treat authorship as a kind of trademark-like identifier of source. Accurate authorship information is needed to allow consumers to understand the type or quality of the work they are considering purchasing.<sup>164</sup> On this account, the use of AI to

---

<sup>161</sup> Id. at 3; see also *infra* Subsection III.C.2 (exploring trademark and right of publicity laws in more detail). The recently proposed NO FAKES Act would bolster these protections by creating a federal cause of action to combat unauthorized “digital replicas.” See Justin Hughes, NO FAKES Act: Unpacking the New Bipartisan Bill on Digital Replicas, George Mason Univ.: Ctr. for Intell. Prop. x Innovation Pol’y (Aug. 1, 2024), <https://cip2.gmu.edu/2024/08/01/no-fakes-act-unpacking-the-new-bipartisan-bill-on-digital-replicas/> [https://perma.cc/327U-KJ6Y].

<sup>162</sup> See Sag, *supra* note 33, at 342 n.154 (noting the criticism that protection of style compromises copyright’s idea-expression dichotomy). But see Benjamin L.W. Sobel, Elements of Style: Copyright, Similarity, and Generative AI, 38 Harv. J.L. & Tech. 49, 49 (2024) (arguing that style should be protectable).

<sup>163</sup> First Amended Complaint at 3, *Andersen*, No. 23-cv-00201.

<sup>164</sup> Greg Lastowka, The Trademark Function of Authorship, 85 B.U. L. Rev. 1171, 1179–80 (2005); see also Jane C. Ginsburg, The Author’s Name as a Trademark: A Perverse Perspective on the Moral Right of “Paternity”?, 23 Cardozo Arts & Ent. L.J. 379, 381 (2005)

mimic the markers of an artist's distinctive brand (such as voice or style) could be considered a kind of deceptive "passing off"—an attempt to induce consumers to purchase a work they might otherwise choose not to purchase.

But there is also a more subjective relationship between consumers and artists that suggests that consumers value knowledge about the precise authorship of a work even when authorship is not functioning as a meaningful brand signifier. In an article criticizing posthumous product endorsements by celebrities—licensed by the celebrity's family member or, more problematically, commercial entities that have purchased trademarks or publicity rights—Andrew Gilden has noted that

fans often carry strong emotional and psychological attachments to the celebrity, and the primary way fans express such attachment is through consumption of goods, services, and causes associated with the celebrity. Public identification with prominent celebrities is broadly understood to be both an act of self-definition and a reflection of an authentic, if sometimes stigmatized, bond with the celebrity.<sup>165</sup>

In the more specific context of creative works, such as television shows and fiction, Rebecca Tushnet has also documented the ways that fandom can shape self-definition and community engagement.<sup>166</sup>

This recognition of the powerful social role of fandom demonstrates the problem with undisclosed AI mimicry of an artist's voice, likeness, or style. Many consumers' decisions to consume a specific creative work can be driven by a strong identification with the artist herself. If a consumer believes that only Drake can produce a true "Drake song," then a song that features an AI-generated version of Drake's voice would never be understood as the same, even if the consumer found the song otherwise enjoyable. Indeed, the FTC has recently suggested consumers may

---

(arguing that an author's name can function as a de facto trademark, allowing consumers to distinguish between creative works).

<sup>165</sup> Andrew Gilden, *Endorsing After Death*, 63 *Wm. & Mary L. Rev.* 1531, 1578 (2022) (footnote omitted).

<sup>166</sup> Rebecca Tushnet, *Legal Fictions: Copyright, Fan Fiction, and A New Common Law*, 17 *Loy. L.A. Ent. L.J.* 651, 655–56 (1997) ("Fan culture, like traditional folk culture, constructs a group identity, articulates the community's ideals, and defines its relationship to the outside world." (quoting Henry Jenkins, *Textual Poachers: Television Fans & Participatory Culture* 273 (1992))); see also Betsy Rosenblatt & Rebecca Tushnet, *Transformative Works: Young Women's Voices on Fandom and Fair Use*, in *eGirls, eCitizens* 385, 386–88 (Jane Bailey & Valerie Steeves eds., 2015) (surveying ways in which fan-made works can serve as modalities of personal development and community building for young women).

experience this use of AI as a kind of deception “when authorship does not align with consumer expectations, such as when a consumer thinks a work has been created by a particular musician or other artist, but it has been generated by someone else using an AI tool.”<sup>167</sup>

In this respect, undisclosed AI-mimicry is a problem that also implicates consumers, rather than just the mimicked artist. But the general public has little legal recourse. A fan does not have a cause of action under the right of publicity, trademark, or copyright to go after the distributor of an AI-generated work in situations where the mimicked artist chooses not to.<sup>168</sup>

Furthermore, some rightsholders might begin *licensing* artists’ publicity rights, trademarks, or style-based copyright interests to generative AI companies, especially posthumously.<sup>169</sup> YouTube, for example, has licensed the use of several artists’ voices for its experimental Dream Track AI music generator.<sup>170</sup> In situations where the license is undisclosed, fans may still experience a kind of deception, even though the artist or her assignee has essentially approved of the use.<sup>171</sup>

### *B. Authorship Transparency and the Social Role of Art*

Identifying the three areas that seem to inform consumer preferences for human-created works does not necessarily tell us whether the law should care if the AI provenance of a work is not disclosed. After all, consumer information can be inaccurate, and lawmakers have often

---

<sup>167</sup> Fed. Trade Comm’n, Comment of the United States Federal Trade Commission, Artificial Intelligence and Copyright 5 (Oct. 30, 2023), [https://www.ftc.gov/system/files/ftc\\_gov/pdf/p241200\\_ftc\\_comment\\_to\\_copyright\\_office.pdf](https://www.ftc.gov/system/files/ftc_gov/pdf/p241200_ftc_comment_to_copyright_office.pdf) [<https://perma.cc/29NR-VBRB>].

<sup>168</sup> See *infra* Subsection III.C.2 (discussing this issue further).

<sup>169</sup> See, e.g., Peter C. Baker, The Beatles Are Still Charting the Future of Pop. It Looks Bleak., N.Y. Times (Nov. 21, 2023), <https://www.nytimes.com/2023/11/21/magazine/the-beatles-now-and-then.html?smid=nytcore-ios-share&referringSource=articleShare>.

<sup>170</sup> Lyor Cohen & Toni Reid, An Early Look at the Possibilities as We Experiment with AI and Music, YouTube Blog (Nov. 16, 2023), <https://blog.youtube/inside-youtube/ai-and-music-experiment/> [<https://perma.cc/BLT6-BS5Q>].

<sup>171</sup> This, of course, raises questions—discussed further below—about why such licensed use would not be considered the same as existing scenarios where the law permits a work to be marketed as having been created by someone who did not actually create it, such as a ghostwriter. See *infra* Subsection II.C.1.

declined to privilege consumer preferences that seem to reflect inefficient, protectionist, or bigoted impulses.<sup>172</sup>

Some of the preferences described above are indeed contestable. Ethical concerns about AI supplanting human labor may be overblown; as some have argued, such concerns have emerged after every new technological revolution and are often based more on fear of change than meaningful public policy concerns.<sup>173</sup> And several scholars have noted, that authenticity policing is often used to create artificial scarcity in a world of abundance.<sup>174</sup> Seen through this lens, a consumer preference for human-made content might be no different than any preference for “craft” or “artisan” products.<sup>175</sup> These preferences may be benign, but they may also intertwine with class-based inequality.<sup>176</sup>

While consumers’ preferences for human-created works might appear somewhat contestable, they may also be undergirded by a stronger set of normative claims that are worth taking seriously. Kysar notes that this can often be true of process-based preferences:

Rather than being scientifically unfounded, nakedly protectionist, or ethically inconsistent, consumer process preferences instead offer an important vehicle through which individuals influence the world, express their views on public issues, and fashion their moral identity in an era of extraordinary interconnectedness, complexity, and dynamism

---

<sup>172</sup> See Kysar, *supra* note 14, at 586–89 (examining these arguments and citing examples, such as regulators’ occasional hostility to consumer preferences for locally produced products).

<sup>173</sup> Rebecca Stropoli, *A.I. Is Going to Disrupt the Labor Market. It Doesn’t Have to Destroy It.*, *Chi. Booth Rev.* (Nov. 14, 2023), <https://www.chicagobooth.edu/review/ai-is-going-disrupt-labor-market-it-doesnt-have-destroy-it> [<https://perma.cc/N9CD-9U7D>] (“[A]s with past tech revolutions that have disrupted the workforce, the creation of new jobs as a result of A.I. advances could help offset the losses . . .”). But see generally Brian Merchant, *Blood in the Machine: The Origins of the Rebellion Against Big Tech* (2023) (reconsidering the Luddites as a kind of early labor-protection movement).

<sup>174</sup> Beebe, *supra* note 150, at 870; see also Adler, *supra* note 149, at 715 (noting that the concept of authenticity is inherently unstable and often redefined to fit new consumer preferences or artistic practices); Burk, *supra* note 151, at 1690 (“In general we may define authenticity as a discursive construct connoting legitimacy and social value.”).

<sup>175</sup> Burk, *supra* note 151, at 1693 (quoting Thomas Thurnell-Read, *A Thirst for the Authentic: Craft Drinks Producers and the Narration of Authenticity*, 70 *Brit. J. Socio.* 1448, 1451 (2019)).

<sup>176</sup> See Jeremy N. Sheff, *Veblen Brands*, 96 *Minn. L. Rev.* 769, 775 (2012) (examining and critiquing the role of trademark law in reinforcing markets for luxury goods); see also Zahr K. Said, *Fables of Scarcity in IP*, 7 *Frontiers Rsch. Metrics & Analytics*, Mar. 2023, at 1, 1, 10–13 (examining costs of NFTs as artificially scarce goods).

in the market. Although consumers undoubtedly suffer from some informational deficiencies with regard to the meaning and significance of various manufacturing practices, the case for wholesale irrationality or unreliability of process preferences is unpersuasive.<sup>177</sup>

Indeed, in other markets, various process preferences can reflect a complex combination of political engagement and personal self-definition. For example, as Kysar notes, a consumer may seek to avoid “conflict diamonds” because of a desire to actively influence the diamond industry but also because avoiding complicity is essential to their sense of self.<sup>178</sup>

This Section attempts to identify a broader “moral identity” that may at least partially underlie consumer preferences for human-made creative works. First, I draw on various theories of aesthetic value to argue that many works of art are fundamentally social, reflecting an ongoing dialogue between author and reader. This dialogue is essential to allowing art to play a role in individual self-definition, as well as political and ethical engagement. Second, I argue that hidden AI authorship destabilizes this social role by forcing consumers to engage with art only as a kind of entertainment-focused commodity.

To be sure, the goal of this Section is not to provide any kind of comprehensive theory of art or authorship, or even necessarily an argument that *all* consumers should care about the distinction between human-made and AI-generated creative works. Rather, my aim is to show that those consumers who articulate a preference for human-created works are worth taking seriously. On this account, when consumers espouse ethical-labor, authenticity, or fandom concerns—such as those noted above—they may in fact be articulating a broader conception of the value of art in a market-oriented society, one which rests on a long line of aesthetical, ethical, and political theory.

---

<sup>177</sup> Kysar, *supra* note 14, at 624.

<sup>178</sup> *Id.* at 615 (quoting Tracey Michelle Price, *The Kimberley Process: Conflict Diamonds, WTO Obligations, and the Universality Debate*, 12 *Minn. J. Glob. Trade* 1, 1 (2003)); *id.* at 615–16. The former can be more easily rebutted by empirical evidence—for example, that consumption choices have little impact on diamond sourcing—but the latter simply reflects “the premise that consumption often is an intensely personal activity.” *Id.* at 616–17.

*1. Authorship and Readership as Social Dialogue*

At its most basic, the notion that an author's (human) touch matters to the reader's<sup>179</sup> experience of a work is intuitive. In Arthur Danto's famous example, Andy Warhol's Brillo Boxes may be aesthetically indistinguishable from a commercially manufactured Brillo box, but the viewer only understands the former as "art."<sup>180</sup>

Such an assumption might strike some as reflecting romanticism, which, with its emphasis on "heroic solitary genius," glorifies the individual creator as the source of artistic significance.<sup>181</sup> Postmodern theorists have questioned romanticism's emphasis on the individual author, instead seeking to situate art in the social, cultural, and political context within which it is created. Most famously, Roland Barthes declared the "[d]eath of the [a]uthor," arguing that "[t]o give a text an Author is to impose a limit on that text, to furnish it with a final signified, to close the writing."<sup>182</sup>

But caring about authorial *identity* (or, at the very least, authorial humanness) is not the same as defining a work's meaning with reference to authorial *intent*.<sup>183</sup> Indeed, even those operating in the reader response tradition—which, by and large, shares Barthes's skepticism of the primacy of the author over that of the reader—have noted that the act of

---

<sup>179</sup> I use this term to describe the person who experiences creative works of any kind.

<sup>180</sup> Arthur Danto, *The Artworld*, 61 *J. Phil.* 571, 580–81 (1964) (distinguishing between "a Brillo box and a work of art consisting of a Brillo Box"). Building on this intuition, Mala Chatterjee has argued that creative works are "author-individuated," that is, "there is an important sense in which an author's expression—and, resultantly, her artistic work—*uniquely* is hers, even if others might have produced something that looks just like it themselves, simply in virtue of the fact that it came from her." Mala Chatterjee, *Understanding Intellectual Property: Expression, Function, and Individuation*, 70 *J. Copyright Soc'y U.S.A.* 57, 71 (2023); see also *id.* at 77–78 (discussing Brillo boxes example).

<sup>181</sup> Barton Beebe, *Bleistein*, *The Problem of Aesthetic Progress, and the Making of American Copyright Law*, 117 *Colum. L. Rev.* 319, 365 (2017); see also Martha Woodmansee, *On the Author Effect: Recovering Collectivity*, 10 *Cardozo Arts & Ent. L.J.* 279, 280 (1992) (arguing the modern notion of an author derives in part from Romantic ideals emphasizing novelty in writing by departing from tradition).

<sup>182</sup> Barthes, *supra* note 21, at 147; see also Michel Foucault, *What Is an Author?*, in *The Foucault Reader* 101, 108–13 (Paul Rabinow ed., 1984) (considering the manner in which conceptions of authorship have changed through history and discussing the complexity of the "author function" today).

<sup>183</sup> Chatterjee, *supra* note 180, at 84.

interpretation will invariably carry with it a set of assumptions about the real or imagined identity of the author.<sup>184</sup>

In this respect, it is possible to argue for the meaningfulness of human authorship to the reader's experience of a work without necessarily embracing the assumptions of romanticism. Carys Craig and Ian Kerr, in particular, have provided a defense of human authorship that avoids romanticism<sup>185</sup> and instead embraces a communitarian vision of the social nature of the self.<sup>186</sup> Craig and Kerr argue that "authorship" is a fundamentally social activity, a "discursive interaction" or "dialogic exchange" through which both authors and readers engage in a process of self-definition.<sup>187</sup> A machine by its very nature cannot participate in this process because "[a]uthorship, as discursive interaction, necessarily occurs in the domain of relatedness—a domain alien to the romantic author, of course, and likewise foreign to the machine."<sup>188</sup> Thus, an AI-produced work is fundamentally "different in kind" even if the reader cannot tell whether the work was AI-generated on its face.<sup>189</sup>

A dialogic account of the value of human authorship and readership—that is, a theory that emphasizes the fundamental social nature of both the

---

<sup>184</sup> See, e.g., Stanley Fish, *Biography and Intention*, in *Contesting the Subject* 9, 14 (William H. Epstein ed., 1991) [hereinafter Fish, *Biography and Intention*] ("The only way to read unbiographically would be to refrain from construing meaning—to refrain, that is, from regarding the marks before you as manifestations of intentional behavior; but that would be not to read at all."); see also Stanley Fish, *Is There a Text in this Class? The Authority of Interpretive Communities* 14 (1980) (elaborating on the concept of the "interpretive community" in which norms of interpretation are constructed).

<sup>185</sup> Carys Craig & Ian Kerr, *The Death of the AI Author*, 52 *Ottawa L. Rev.* 31, 51–52 (2020–2021) ("By freeing ourselves of the ideology of the romantic author, we can understand the text as circulating discourse and concern ourselves with the place and function of the speaking subject in discursive relations to and through the text.")

<sup>186</sup> *Id.* at 83 (citing Charles Taylor, *The Dialogical Self*, in *The Interpretive Turn: Philosophy, Science, Culture* 304, 314 (David R. Hiley, James F. Bohman & Richard Shusterman eds., 1991)).

<sup>187</sup> *Id.* at 83–85 ("[T]he human subject as embedded in social networks of interdependence, but also as possessing autonomy—autonomy that is properly conceptualized in relational terms: it is only *through* relationships that genuine autonomy is made possible.")

<sup>188</sup> *Id.* at 85.

<sup>189</sup> *Id.*; cf. Chatterjee, *supra* note 180, at 62 ("[T]wo distinct acts of authorship—even ones resulting [in] works that are 'structurally' identical—cannot result in the very same artistic work, because the identity of the author in part *makes* the work *what it is*."). Chris Buccafusco's theory of authorship, which argues that "authorship" requires both the "inten[t] to produce some mental effect in an audience" and the creation of the "text, object, or medium that is capable of producing that mental effect," provides some tacit support for this conclusion as well, as machines are not capable of intent. See Christopher Buccafusco, *A Theory of Copyright Authorship*, 102 *Va. L. Rev.* 1229, 1232 (2016).

creation of works of art and the experience of these works—gains purchase in theoretical accounts that emphasize the importance of art toward self-definition, empathy, and moral development. Richard Rorty, for example, has argued that literature plays an essential role in “liberat[ing] one from one’s own previous ways of thinking about the lives and fortunes of individual human beings.”<sup>190</sup> More specifically:

Most novels tell us how other erring mortals think of themselves, how they contrive to put the actions that appall us in a good light, how they give meaning to their miserable or tragic or banal lives. Seen in this light, what novels do for us is to let us know how people quite unlike ourselves think of themselves, how they contrive to put actions that appall us in a good light, how they give their lives meaning. The problem of how to live our own lives then becomes a problem of how to balance our needs against theirs, and their self-descriptions against ours. To have a more educated, developed and sophisticated moral outlook is to be able to grasp more of these needs, and to understand more of these self-descriptions.<sup>191</sup>

Other theorists note that it is precisely this ability to foster the recognition of others’ subjectivity that gives art a necessarily political role. Martha Nussbaum, in particular, has argued that the novel furnishes “the ability to imagine what it is like to live the life of another person who might, given changes in circumstance, be oneself or one of one’s loved ones.”<sup>192</sup> Art thus enables the kind of empathy and compassion that furthers better civic engagement,<sup>193</sup> and it makes citizens more committed to social equality and justice.<sup>194</sup> Rorty similarly argues that literature is an essential element in fostering “recognition” of oppressed groups by

---

<sup>190</sup> Richard Rorty, *Redemption from Egotism: James and Proust as Spiritual Exercises*, in *The Rorty Reader* 389, 390 (Christopher J. Voparil & Richard J. Bernstein eds., 2010).

<sup>191</sup> *Id.* at 393; see also Alasdair MacIntyre, *After Virtue: A Study in Moral Theory* 216–18 (3d ed. 2007) (“[M]an is in his actions and practice, as well as in his fictions, essentially a story-telling animal. He is not essentially, but becomes through his history, a teller of stories that aspire to truth.”).

<sup>192</sup> Martha C. Nussbaum, *Poetic Justice: The Literary Imagination and Public Life* 5 (1995).

<sup>193</sup> *Id.* at 10–12.

<sup>194</sup> *Id.* at 79–99; see also Martha C. Nussbaum, *Exactly and Responsibly: A Defense of Ethical Criticism*, 22 *Phil. & Lit.* 343, 345–50 (1998) (describing literature as “a valuable source of deliberative enrichment”).



acknowledging that the members of such groups have “as rich an inner life as one does oneself.”<sup>195</sup>

## 2. *From Obfuscation to (Complete) Commodification*

In all theoretical accounts that focus on artistic engagement as social, ethical, or political activity, the understanding of the reader that they are engaged in dialogue with *another human* is taken as a given. As Craig and Kerr note, “[t]he dialogic act of authorship cannot be separated from a social context because the ‘[u]tterance, as we know, is constructed between two socially organized persons.’”<sup>196</sup> Theorist Noel Carroll notes that an ethical role for art<sup>197</sup> requires that the “author and the audience . . . share a common background of beliefs about the world and about human nature, as well as a relatively common emotional life.”<sup>198</sup> Thus, even if, as many have argued (echoing Barthes), the author’s “intent” lacks significance for the reader, the author’s and reader’s basic shared participation in the same (human) social world remains essential to allowing art to fulfill a social, ethical, or political function.<sup>199</sup>

<sup>195</sup> Richard Rorty, Is “Cultural Recognition” a Useful Concept for Leftist Politics?, in *The Rorty Reader*, supra note 190, at 463, 465; see also Richard Rorty, Human Rights, Rationality, and Sentimentality, in *The Rorty Reader*, supra note 190, at 351, 360–65 (arguing that empathy derived from this recognition undercuts tribalistic tendencies).

<sup>196</sup> Craig & Kerr, supra note 185, at 82 (quoting V.N. Voloshinov, *Marxism and the Philosophy of Language* 85 (Ladislav Matejka & I.R. Titunik trans., Harvard Univ. Press 1986) (1929)) (“The word is oriented toward an addressee, toward *who* that addressee might be. . . . Or, as analytic philosophers argue, there is no such thing as a private language.” (quoting Voloshinov, supra, at 85)); cf. Fish, *Biography and Intention*, supra note 184, at 15 (“[C]riticism can only proceed when . . . notions of agency, personhood, cause, and effect are already assumed and are already governing.”).

<sup>197</sup> Carroll is skeptical that *all* art has a moral valence but agrees that many forms do. Noël Carroll, *Art, Narrative, and Moral Understanding*, in *Aesthetics and Ethics* 126, 137–39, 156 n.5 (Jerrold Levinson ed., 1998).

<sup>198</sup> *Id.* at 139.

<sup>199</sup> Several creative writers have also begun to express similar sentiments. See Kang, supra note 20 (“[T]he reason we read books and listen to songs and look at paintings is to see the self in another self, or even to just see what other people are capable of creating.”); Ted Chiang, *Why AI Isn’t Going to Make Art*, *New Yorker* (Aug. 31, 2024), <https://www.newyorker.com/culture/the-weekend-essay/why-ai-isnt-going-to-make-art> (“Whether you are creating a novel or a painting or a film, you are engaged in an act of communication between you and your audience. What you create doesn’t have to be utterly unlike every prior piece of art in human history to be valuable; the fact that you’re the one who is saying it, the fact that it derives from your unique life experience and arrives at a particular moment in the life of whoever is seeing your work, is what makes it new.”); Vivian Lam, *Human Art Already Has So Much in Common with AI*, *Wired* (Feb. 24, 2023, 8:00 AM), <https://www.wired.com/story/generative-art-algorithms-creativity/> (A work of art is “something that bears the weight

Some consumers may of course have no interest in engaging with a creative work as anything other than an object for entertainment. But for consumers who do care, a world in which AI authorship is routinely used but not disclosed is a world that is not compatible with a social, ethical, or political role for art. In a world in which a consumer lacks information as to whether a novel, song, or painting is a product of a machine or of the human mind, all art must be viewed skeptically, as presumptively machine-made. This robs art of its social dimension; a consumer can *only* assess the work based on its entertainment value.<sup>200</sup>

This, in essence, turns art into what Margaret Radin might call a “complete commodity.” As Radin explains, commodification is not simply a question of whether something has a market-based price, it is also a “worldview” that treats “the person as a commodity-trader” and “[a]ll social and political interactions . . . as exchanges for monetizable gains.”<sup>201</sup> In the case of works of art, a commodified worldview would essentially treat the value of art as coextensive with the utility (i.e., entertainment value) a consumer derives from its consumption, rather than reflecting a desire to engage dialogically with an author as part of a process of self-development or ethical enrichment.

---

and warmth of another’s life. Every trope is made unique when run through the sieve of the individual. . . . [W]hat devalues art is the negation of the artist and community that imbue a work with significance beyond its mere existence.”); Ahmed Elgammal, *The Robot Artists Aren’t Coming*, N.Y. Times (May 27, 2020), <https://www.nytimes.com/2020/05/27/opinion/artificial-intelligence-art.html> (“While the definition of art is ever-evolving, at its core it is a form of communication among humans. Without a human artist behind the machine, A.I. can do little more than play with form, whether that means manipulating pixels on a screen or notes on a musical ledger. These activities can be engaging and perceptually intriguing, but they lack meaning without interaction between artist and audience.”).

<sup>200</sup> This phenomenon partially reflects what Sherry Turkle has noted is the human tendency to “change our metric for assigning value and place a higher and higher value on what is always available. And that will always be a machine.” Sherry Turkle, *Empathy Machines: Forgetting the Body*, in *A Psychoanalytic Exploration of the Body in Today’s World* 17, 22 (Vaia Tsolas & Christine Anzieu-Premmereur eds., 2018). Thus, for example, as people grow accustomed to the idea of caretaking robots for the elderly, “[e]fficiency and cost become the new metrics and the inanimate creature becomes better than anything, more reliable than what life could ever provide.” *Id.* at 21–22. Turkle views this trend as contributing to an “assault” on the value of genuine human empathy. *Id.* at 18; see also *id.* at 27 (“Empathy machines raise the concern that we will model who we are on their limitations and human absence will come to seem sufficient unto the day.”).

<sup>201</sup> Margaret Jane Radin, *Contested Commodities* 5–6 (2001). In so doing, this worldview “assimilates personal attributes, relations, and desired states of affairs to the realm of objects by assuming that all human attributes are possessions bearing a value characterizable in money terms.” *Id.* at 6.

This, again, is not to deny that many works can *also* be assessed based on their entertainment value—and that some works may only be worth assessing through this lens. As Radin notes, “commodified understandings of certain transactions can coexist with noncommodified understandings,” but this requires that the public has both the ability and freedom to engage with an object in different ways.<sup>202</sup> The problem, then, with undisclosed AI authorship is the total privileging of a commodified understanding of art and the destruction of any social dimension. Information regarding a work’s provenance restores art’s ability to function both socially *and* as a commodity by enabling consumers to choose whether and on what terms—as entertainment commodity, as meaningful source of dialogic engagement, or both—to engage with the work.

Thus, to summarize, it is the *undisclosed* use of AI that is troubling here. Lack of disclosure presumptively forces all art into the category of entertainment commodity, robbing it of its ability to not only entertain, but also promote a reader’s self-definition, empathy, and political development through participation in a dialogue with other human agents.

\* \* \*

This argument provides an additional normative framework for understanding why some consumers seem uncomfortable with undisclosed AI authorship, as discussed in the last Section. Indeed, consumers’ ethical, authenticity-based, and fandom-derived preferences for human-made works can be taken more seriously when considered through the lens of a dialogical account of authorship.<sup>203</sup> An ethical preference for art by humans may not simply reflect an outmoded understanding of labor markets,<sup>204</sup> but can be based on the argument that

---

<sup>202</sup> Id. at 102–03. For example, “[t]he same person can understand an interaction in different, and conflicting, ways, as, analogously, she can both feel a painting is priceless and yet have it appraised for insurance purposes.” Id.; cf. Elizabeth Anderson, *Value in Ethics and Economics* 207 (1993) (“[T]he fact that there are markets in aesthetic goods does not show that market prices comprehensively measure their values. Market prices measure only how much people value goods for exclusive appropriation and consumption.”).

<sup>203</sup> Cf. Kysar, *supra* note 14, at 624 (“Rather than being scientifically unfounded, nakedly protectionist, or ethically inconsistent, consumer process preferences instead offer an important vehicle through which individuals influence the world, express their views on public issues, and fashion their moral identity . . . . Although consumers undoubtedly suffer from some informational deficiencies with regard to the meaning and significance of various manufacturing practices, the case for wholesale irrationality or unreliability of process preferences is unpersuasive.”).

<sup>204</sup> See *supra* notes 172–75 and accompanying text.

human-made works of art play a particularly important social and moral function and thus must be elevated. A belief that only human-made art can be “authentic” need not simply reflect classist preferences for scarcity,<sup>205</sup> but can be motivated by fundamental assumptions about the necessity of human subjectivity in the creation of art. And a fan-based preference for works by specific artists<sup>206</sup> may in fact be an indication that meaningful social bonds are being developed through a shared dialogue with a specific author.

That being said, it is important to reiterate that the argument here is not that the use of AI in creative endeavors is inherently problematic; the technology is likely to yield a range of works that many consumers might find enjoyable and choose to consume. Nor is the argument that all consumers *should* care about whether a human or a machine is the source of a work of art or entertainment. Rather, my goal is to demonstrate that those consumers who value human-created works are on sound—and indeed well-developed—normative ground, rather than operating based on a set of irrational preferences. For these consumers, AI-generated works that masquerade as human-made must be considered a kind of deceptive harm. The next Section builds on this implication and explores several objections to it.

### *C. Implications and Objections*

The last two Sections made the case that consumer preferences for non-AI-generated creative works exist and are worth taking seriously. For those consumers who care about human authorship, obfuscation of a work’s AI provenance frustrates these preferences and should be considered harmful. While this harm may not rise to the level of the deceptive harm posed by AI-generated deepfakes, we can still conclude that disclosure of the AI origins<sup>207</sup> of both factual and non-factual AI-generated works is important.

Before addressing various ways disclosure can be encouraged (in Part III), this Section lays out some objections and difficulties posed by disclosure. First, I argue why the law’s tolerance for undisclosed ghostwritten works should not extend to undisclosed AI-generated works.

---

<sup>205</sup> See *supra* notes 172–75 and accompanying text.

<sup>206</sup> See *supra* Subsection II.A.3.

<sup>207</sup> Below, I examine when a work should be considered to have AI origins. See *infra* Subsection II.C.2.

Second, I examine the line between true AI authorship and the minimal use of AI as an assistive tool in the service of human authorship. Third, I explain why a disclosure regime would not harm those consumers who do not care about the AI origins of a work.

### *1. Distinguishing the Human Ghostwriter*

Any normative contention that consumers receive accurate information regarding a work's origins faces a strong counterargument: What about ghostwriting? The law has never treated the use of a "ghostwriter"—that is, a person who agrees to create a work that is later attributed to an entirely different person—as deceptive. Why should the use of generative AI be considered different? One response might be that the two are not different, and that the use of a ghostwriter should in fact be considered a form of deception, as some have argued.<sup>208</sup>

But the undisclosed use of generative AI is also problematic in ways distinct from human ghostwriting. The use of a ghostwriter is a question of disclosed *attribution*, which is normatively different from the question of disclosed *authorship*. Indeed, most of the dialogic accounts of authorship, described above, do not particularly care that a reader knows exactly who the creator was; anonymously or pseudonymously produced art can be just as capable of eliciting self-definition, empathy, or solidarity as any other work. Rather, it is the understanding that the reader is in dialogue with another human agent, who is capable of experiencing the world in a similar way (and drawing on those experiences), that yields a social or ethical dimension for art.<sup>209</sup> The consent of the ghostwriter to forego the normal expectations of attribution does not necessarily alter this outcome.

Ghostwriting also occupies an ambiguous cultural space in which consumers are understood as having consented to the charade. As Christopher Miller has noted, "the roman à clé, the pseudonym or nom de plume, the found journal, the ghostwriter—[all] are at least partially transparent and purposefully so; they exist within a socially sanctioned implied contract or pact with the reader."<sup>210</sup> This runs in contrast to a true

---

<sup>208</sup> See Lastowka, *supra* note 164, at 1227–28 (arguing that the undisclosed use of a ghostwriter is deceptive).

<sup>209</sup> See *supra* Subsection II.B.1.

<sup>210</sup> Christopher L. Miller, *Impostors: Literary Hoaxes and Cultural Authenticity* 8 (2018); see also Laura A. Heymann, *Naming, Identity, and Trademark Law*, 86 *Ind. L.J.* 381, 383 (2011) (noting that the public has a "shared understanding of when the disjunction between

“hoax,” which violates the cultural norms surrounding authorship and attribution.<sup>211</sup>

An element of this cultural understanding is that a ghostwriter is often understood to be *assisting* a person in translating their history or perspective into prose. In these cases, a consumer may actually (and reasonably) still understand the attributed author to be the true “author,” even if another person actually prepared most of the work.<sup>212</sup> The ghostwriter might be considered more akin to a creative assistant, like an editor. A similar cultural understanding exists in the case of fine artists, such as Jeff Koons or Dale Chihuly, who envision a work and then task assistants with physically constructing it.<sup>213</sup>

In contrast, the undisclosed use of generative AI, at least in many of the uses described above—the creation of a work of fine art, the writing of a screenplay, the creation of a comic book—does not rest on such tacit cultural norms. In such uses, the AI surreptitiously accomplishes the tasks that are understood to be fundamental to authorship.<sup>214</sup> That being said, there may indeed be cases where a human author’s use of generative AI could be understood as closer to this kind of creative assistance supplied by a ghostwriter or editor or art-studio assistant. As the next Subsection explains, a dialogic account of authorship can help with this line drawing, paving the way for reasonable policy proposals as to how disclosure of the use of AI can be incentivized.

## 2. *Disentangling the AI Author from the AI Tool*

As Part I explored, generative AI has a range of applications in the production of creative works. In some of these applications, nearly the entire work is generated by AI, with little creative agency by a human,

---

producer and attribution is appropriate (a celebrity autobiography ghostwritten by another writer is sanctioned, but plagiarism by a college student is not . . . )”); Catherine L. Fisk, *Credit Where It’s Due: The Law and Norms of Attribution*, 95 *Geo. L.J.* 49, 65–67 (2006) (“We participate in a willing suspension of disbelief when we attribute speeches to the candidate or politician, and treat the speeches as a window into the heart and mind of the speaker.”).

<sup>211</sup> Miller, *supra* note 210, at 8–10.

<sup>212</sup> See Ginsburg, *supra* note 164, at 381–82 (suggesting that in the case of ghostwriters, “the ‘author’ is the person who presents herself as such, who succeeds in persuading the public that her personality pervades the work, even if someone else wrote it”).

<sup>213</sup> See Rosie Lesso, *How Does Jeff Koons Make His Art?*, *Collector* (Aug. 31, 2022), <https://www.thecollector.com/how-does-jeff-koons-make-his-art/> [<https://perma.cc/95H8-NY8P>] (“I’m basically the idea person. I’m not physically involved in the production.”); see also *infra* note 215 (discussing Dale Chihuly).

<sup>214</sup> See *supra* Section II.B; *infra* notes 219–26 and accompanying text.

but in others, the AI serves more as a tool or assistant, supplementing and bolstering the creative efforts of a human author.

Intuitively, these two distinct uses do not feel the same. The use of creative assistance—editing, research assistance, assistance with the physical creation of fine art<sup>215</sup>—is common in most creative industries. Moreover, many works, especially films, involve the collaboration and combination of many people’s labor to create a “unitary” work.<sup>216</sup>

Many uses of generative AI seem to echo these more assistive uses. Film provides a particularly relevant example. Many current generative AI use-cases involve the enhancement or supplementing of existing imagery during post-production of a film.<sup>217</sup> For example, the AI generation tool Runway was used as part of the visual effects process for the film *Everything Everywhere All at Once*.<sup>218</sup>

Other generative AI uses, however, would seem to supplant labor that we would consider to be properly within the purview of the (human) author. To continue with the example of film, consider the use of a generative AI system to write an entire script, such as with the (hypothetical) AI-generation system SHOW-1, which purportedly will be able to craft “narrative arcs (i.e., an entire episode for a TV series) that are consistent with the characters and canon of an existing, pre-trained intellectual property.”<sup>219</sup> Open-AI’s new Sora tool goes a step further, by

<sup>215</sup> See *Moi v. Chihuly Studio, Inc.*, 846 F. App’x 497, 499 (9th Cir. 2021) (holding that an assistant in the studio of fine artist Dale Chihuly was not a “coauthor” within the meaning of copyright law when the artist “controlled everything from the artistic design to the decision of whether the finished paintings would be signed, marketed, and sold or merely discarded”); see also Taylor Dafoe, *A Judge Has Rejected a \$20 Million Lawsuit by a Former Studio Assistant Who Claims Famed Sculptor Dale Chihuly Didn’t Properly Credit His Work*, Artnet (June 24, 2019), <https://news.artnet.com/art-world/dale-chihuly-lawsuit-dismissed-1582921> [<https://perma.cc/V9LW-W46C>] (“Moi could not properly demonstrate that he was a co-author of the paintings and therefore was owed none of the proceeds from their sale, the judge ruled.”).

<sup>216</sup> See *Garcia v. Google, Inc.*, 786 F.3d 733, 737 (9th Cir. 2015) (rejecting a theory of copyright law that would allow “any contributor from a costume designer down to an extra or best boy to claim copyright in random bits and pieces of a unitary motion picture”).

<sup>217</sup> Davenport, *supra* note 53 (discussing this use-case).

<sup>218</sup> Jazz Tangcay, ‘Hollywood 2.0’: How the Rise of AI Tools Like Runway Are Changing Filmmaking, *Variety* (Feb. 22, 2023, 3:39 PM), <https://variety.com/2023/artisans/news/artificial-intelligence-runway-everything-everywhere-all-at-once-1235532322/> [<https://perma.cc/U7RV-F24U>].

<sup>219</sup> Doug Shapiro, *AI Use Cases in Hollywood*, Medium (Sept. 18, 2023), <https://dougshapiro.medium.com/ai-use-cases-in-hollywood-362707e899f1> [<https://perma.cc/2A4P-RX36>]; see also Philipp Maas et al., *To Infinity and Beyond: SHOW-1 and Showrunner Agents in Multi-Agent Simulations*, <https://fablestudio.github.io/showrunner-agents/> [<https://perma.cc/>]

creating realistic videos, including “complex scenes with multiple characters” based on text prompts.<sup>220</sup> Sora has already been used to generate a commercial, and OpenAI is reportedly marketing it to Hollywood studios.<sup>221</sup>

These distinct use-cases—the use of an AI system for post-production visual enhancement versus its use for writing a screenplay or generating an entire film—are intuitively different and indeed would be considered different under the normative accounts of authorship described above. A dialogic theory of artistic value requires that a reader understand a human agent as the creative source of the work of art—another human who can be questioned, empathized with, argued with, i.e., someone who shares our social world. This does not necessarily mean that this human agent must be the *exclusive* source of the work, but they must be the guiding agent. As novelist and critic Ted Chiang has noted, “art requires making choices at every scale; the countless small-scale choices made during implementation are just as important to the final product as the few large-scale choices made during the conception.”<sup>222</sup> When a human supplies only a high-level creative vision (such as via a prompt), but the machine makes all of the creative choices that substantiate that vision, the machine has ceased being a mere assistive tool.<sup>223</sup>

As an area of law that must similarly grapple with the line between human-derived and machine-derived creativity, copyright law is a natural place to seek guidance on how to draw this line. Numerous cases have held that providing creative assistance to an author<sup>224</sup> or providing a small

---

TB74-3MWD] (last updated July 24, 2023) (outlining SHOW-1’s “approach of using multi-agent simulation and large language models for generating high-quality episodic content”).

<sup>220</sup> Creating Video from Text, OpenAI, <https://openai.com/index/sora/> [<https://perma.cc/QJ74-TRSG>] (last visited Oct. 29, 2024); see also Tobin & Metz, *supra* note 57 (discussing similar generative AI video technology created by Chinese company Kuaishou).

<sup>221</sup> Thomas Buckley, Lucas Shaw & Shirin Ghaffary, OpenAI Courts Hollywood in Meetings with Film Studios, Directors, Bloomberg (Mar. 22, 2024, 11:51 AM), <https://www.bloomberg.com/news/articles/2024-03-22/openai-courts-hollywood-in-meetings-with-film-studios-directors>; Sean Hollister, The Owner of Toys ‘R’ Us Just Used OpenAI’s Sora to Animate the Zombie Brand, The Verge (June 25, 2024, 7:30 PM), <https://www.theverge.com/2024/6/25/24186051/toys-r-us-openai-sora-brand-film> [<https://perma.cc/ZNX4-ZDTP>].

<sup>222</sup> Chiang, *supra* note 199.

<sup>223</sup> Cf. *id.* (discussing the importance of “small-scale” choices toward crafting art across literary and visual media).

<sup>224</sup> *Thomson v. Larson*, 147 F.3d 195, 197, 203–05 (2d Cir. 1998) (finding that “dramaturgical assistance” in refining and improving a dramatical work does not create joint authorship); *Moi v. Chihuly Studio, Inc.*, No. 17-cv-00853, 2019 WL 2548511, at \*3 (W.D. Wash. June 20, 2019), *aff’d*, 846 F. App’x 497 (9th Cir. 2021).



creative contribution that is incorporated into a larger work by another,<sup>225</sup> does not supplant authorship; the author remains the person who “‘superintends’ the work by exercising control. . . . ‘[T]he inventive or master mind’ who ‘creates, or gives effect to the idea.’”<sup>226</sup>

But “superintendence” still requires that the author make the expressive choices that substantiate their idea. When another human provides these elements, they will instead be considered the author (or, if multiple people shared this role, the work will be considered jointly authored).<sup>227</sup> And when a non-human animal or the forces of nature provide the creative substance of a work, copyright will not recognize authorship at all, and the work is considered to be in the public domain.<sup>228</sup> As discussed in Part I, the Copyright Office has been grappling explicitly with these precedents with respect to generative AI, insisting that a work can only be considered “authored” (and thus copyrightable), if a human, and not the generative AI system, exercises more than *de minimis* creative control over the final work.<sup>229</sup> Indeed, the Copyright Office has recently clarified that “[w]hile assistive uses that enhance human expression do not limit copyright protection, uses where an AI system makes expressive choices” will generally not yield a copyrightable work.<sup>230</sup>

While it is beyond the scope of this Article to consider whether this line drawing is normatively sound within the four corners of copyright law,<sup>231</sup>

<sup>225</sup> *Aalmuhammed v. Lee*, 202 F.3d 1227, 1233 (9th Cir. 2000) (“A creative contribution does not suffice to establish authorship of the movie.”); *Garcia v. Google, Inc.*, 786 F.3d 733, 742–43 (9th Cir. 2015).

<sup>226</sup> *Aalmuhammed*, 202 F.3d at 1234 (quoting *Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53, 61 (1884)); see also *Cnty. for Creative Non-Violence v. Reid*, 490 U.S. 730, 737 (1989) (“As a general rule, the author is the party who actually creates the work, that is, the person who translates an idea into a fixed, tangible expression entitled to copyright protection.”).

<sup>227</sup> *Aalmuhammed*, 202 F.3d at 1233–34.

<sup>228</sup> See *supra* Subsection I.B.2 (demonstrating that courts have consistently rejected non-human claims for copyright protection).

<sup>229</sup> See *supra* Subsection I.B.2.

<sup>230</sup> Copyright and Artificial Intelligence, *supra* note 87, at 11–12.

<sup>231</sup> Several commentators have made doctrinal and normative arguments that a predominantly AI-generated work should not be protectable by copyright. See Pamela Samuelson, *Allocating Ownership Rights in Computer-Generated Works*, 47 *U. Pitt. L. Rev.* 1185, 1201 (1986) (“[W]hen the user’s instructions become increasingly brief or general and the role of the computer in the design or arrangement process becomes correspondingly greater, the authorship of the user becomes increasingly difficult to defend.”); see also Daniel J. Gervais, *The Machine as Author*, 105 *Iowa L. Rev.* 2053, 2100 (2020) (arguing that the protectability of a work should focus on “originality causation,” that is, whether the original protectable expression can be causally traced to a human agent or whether it stems from

it does correspond well to the account of artistic value explored above.<sup>232</sup> If consumer preferences are primarily focused on there being a human mind behind the creative decisions that go into the creative substance of a work,<sup>233</sup> then the Copyright Office's distinctions seem reasonable.<sup>234</sup> From the consumer's perspective, there is a fundamental difference between a *de minimis* human addition to an otherwise AI-generated work and a work that is predominantly human-superintended but uses AI for editing or refinement.<sup>235</sup> In this respect, copyright law might prove useful in developing a framework for defining the line between AI-as-author and AI-as-tool, not only for purposes of copyright protection but also for purposes of disclosure obligations.<sup>236</sup>

---

"machine-made choices"); Jane C. Ginsburg & Luke Ali Budiardjo, *Authors and Machines*, 34 *Berkeley Tech. L.J.* 343, 433 (2019) ("If the user of the machine supplies her creative contribution without influencing *how* the machine translates that contribution into a final work, then the user does not *execute* the final work and thus cannot claim authorship.").

<sup>232</sup> See *supra* Section II.B.

<sup>233</sup> See *supra* Section II.B.

<sup>234</sup> That being said, the Copyright Office's current approach is only a starting point; the Office's guidance is woefully incomplete and would benefit from additional clarifications as to when the merging of human and AI-generated materials creates a protectable work. See *supra* notes 110–13 and accompanying discussion.

<sup>235</sup> Katrina Geddes has argued that essentially all AI-generated works maintain some human touch because a "human user must have an idea, express that idea in one or more prompts, and continue prompting the AI system until it produces the desired result. There is still a discursive, relational, dialogic exchange taking place between creator and audience, just with the assistance of a powerful machine." Katrina Geddes, *Generative AI's Public Benefit 17* (Working Paper, Feb. 1, 2025), [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4865510](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4865510) [<https://perma.cc/GT2L-HJGQ>]. But this conclusion overstates the prompter's contributions. While a prompt can define the general parameters of an output, much of the work's substance is better attributed to the AI's algorithmic decision-making. This is a primary reason the Copyright Office has denied protection to AI-authored works. See Kasunic, *supra* note 11, at 9 ("Rather than a tool that [a user] controlled and guided to reach [their] desired image, Midjourney generates images in an unpredictable way. . . . As the Supreme Court has explained, the 'author' of a copyrighted work is the one 'who has actually formed the picture,' the one who acts as 'the inventive or master mind.' A person who provides text prompts to Midjourney does not 'actually form[]' the generated images and is not the 'master mind' behind them." (quoting *Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53, 61 (1884))); see also Chiang, *supra* note 199 (arguing that AI does not produce art because artistry "results from [the artist] making a lot of choices"); Copyright and Artificial Intelligence, *supra* note 87, at 19 ("The gaps between prompts and resulting outputs demonstrate that the user lacks control over the conversion of their ideas into fixed expression, and the system is largely responsible for determining the expressive elements in the output.").

<sup>236</sup> Though beyond the scope of this Article, First Amendment law may also provide some guidance. In examining why a navigation chart would be considered a "product" but a work like Marcel Duchamp's *The Fountain* (a repurposed urinal) would be considered speech, Robert Post explains that First Amendment-protected speech "presuppose[s] and embod[ies]

### *3. Leaving Space for the Responsible Proliferation of AI-Generated Content*

Some have argued that concerns over the rise of easily accessible generative AI tools neglect the technology’s potential for democratizing cultural production.<sup>237</sup> On this account, the ease and widespread availability of generative AI tools can help foster “creative play,” especially among communities that have previously lacked the resources to experiment with artistic creation, or that have aesthetic sensibilities that have been poorly reflected in dominant cultural narratives.<sup>238</sup>

This Article’s primary normative conclusion—that the meaningful use of AI in the creation of a work must be disclosed to consumers—does not necessarily dispute such arguments. In this respect, it is important to clarify what a disclosure obligation would or would not accomplish. A disclosure obligation would inform consumers about the provenance of a work, but it would not stop the work’s creation or dissemination. Those consumers who simply do not care about AI authorship—or even prefer

---

a certain kind of relationship between speaker and audience. We might roughly describe that relationship as dialogic and independent. . . . The Fountain assume[s] that [its] audience[] will autonomously query [its] meaning and authority. Navigation charts do not receive First Amendment protection, however, because we interpret them as speaking monologically to their audience, as inviting their audience to assume a position of dependence and to rely on them.” Robert Post, *Recuperating First Amendment Doctrine*, 47 *Stan. L. Rev.* 1249, 1253–55 (1995).

<sup>237</sup> See, e.g., Geddes, *supra* note 235, at 4 (“Generative AI promises to democratize cultural production by distributing powerful and accessible tools to previously excluded creator communities.”); *id.* at 1 (“Ordinary individuals can now create sophisticated synthetic media by modifying, remixing, and transforming cultural works without any artistic training or skills. The democratization of cultural production will have profound implications for the distribution of semiotic power.”).

<sup>238</sup> *Id.* at 36–38, 43–46 (citing recent examples); see also, e.g., Martin Nebelong, *AI Art Is Only a Threat if We Let “Prompt-Jockeys” Take Control*, *Creative Bloq* (Apr. 15, 2024), <https://www.creativebloq.com/news/ai-art-is-too-generic> [<https://perma.cc/N7U4-LFRF>] (arguing that “generative AI art tools can be an ally for artists, helping [them] improve workflows, find new ideas and speed up creativity”); Natalie Stoclet, *For These 5 Creatives, AI Has Become an Artistic Ally*, *Forbes* (Oct. 18, 2023, 9:31 AM), <https://www.forbes.com/sites/nataliestoclet/2023/10/17/for-these-5-creatives-ai-has-become-an-artistic-ally/> (showcasing various artists, including Marie Laffont, who thinks of AI as “a tool that lets young designers like [Laffont], who don’t have a lot of resources, turn their dreams into a reality”); Elisa Rae Shupe, *Don’t Let the Copyright Fight Take Away My Ability to Create*, *Tech Pol’y Press* (Apr. 30, 2024), <https://www.techpolicy.press/dont-let-the-copyright-fight-take-away-my-ability-to-create/> [<https://perma.cc/KUU5-FC6F>] (discussing her own struggles with cognitive disabilities and how she used generative AI to write and publish her first novel).

it—remain free to enjoy such content or to experiment with generative AI technology themselves.

That being said, disclosure obligations may certainly affect *the ability to profit* from AI-generated works. As noted in Part I, the two major incentives to obfuscate both relate to profitability: the ability to market to consumers who may otherwise choose not to consume a work if they knew it was AI-generated and the ability to gain copyright protection.<sup>239</sup> However, as both of these strategies essentially involve misleading (consumers and/or the Copyright Office) for financial gain, this is hardly justification to not impose a disclosure requirement.

Moreover, if it is indeed true that AI has particular value in enabling small-scale creative play,<sup>240</sup> the inability to monetize a creative work may not prove much of an impediment to the technology's most socially valuable use-cases. As several scholars have noted, small-scale cultural production and experimentation is often not financially motivated, and is driven instead by norms of sharing and community building.<sup>241</sup>

### III. SOLUTIONS

This Part looks to solutions to the problem of hidden AI authorship, examining several ways that content producers could be incentivized or required to disclose the meaningful use of generative AI in a work's creation.

First, I examine the possible ways that the AI provenance of a work could be affirmatively disclosed to the public, looking at proposed industry self-regulatory solutions, as well as the possibility of a legislative transparency mandate.<sup>242</sup> Second, I consider regulatory options short of a comprehensive transparency requirement. The FTC, under its mandate to protect consumers from deception, could police specific instances in which the omission of AI authorship information has been shown to be misleading. As these options all face some barriers to implementation, I

---

<sup>239</sup> See *supra* Section I.B.

<sup>240</sup> See *supra* note 237 and accompanying text.

<sup>241</sup> See, e.g., Rebecca Tushnet, *Payment in Credit: Copyright Law and Subcultural Creativity*, 70 *Law & Contemp. Probs.* 135, 135 (2007); Betsy Rosenblatt, *Belonging as Intellectual Creation*, 82 *Mo. L. Rev.* 91, 128–30 (2017).

<sup>242</sup> Of course, any affirmative disclosure regime must grapple with what kind of uses of AI must be disclosed and what kinds need not be disclosed. As noted above, the analytic framework used to assess human authorship for copyrightability purposes may be able to fill this role. See *supra* Subsection II.C.2.

finally consider a possible role for existing intellectual property law—copyright, trademark, and the right of publicity—in incentivizing disclosure.

### *A. Affirmative Disclosure*

#### *1. Private Ordering*

From “fairtrade” clothing, to “organic” food, to “small-batch” beer, consumers are often willing to pay a premium for a product that (they believe) was made using a process they approve of.<sup>243</sup> In such industries, the market steps in to provide process-based information. This form of private ordering is indeed a possible outcome in the case of hidden AI authorship. If some consumers truly value human authorship enough, some content producers may publicly commit to avoiding AI, and publicize that fact, in order to reach this segment of the buying public.<sup>244</sup>

That being said, such forms of self-regulation are notoriously unreliable. Many producers simply lack the incentives to make voluntary disclosures when doing so makes it more difficult to use cheaper production techniques.<sup>245</sup> And even in situations where information is voluntarily disclosed, consumers may lack the resources to assess the reliability of this information.<sup>246</sup>

The remainder of this Section surveys some proposed mechanisms for voluntary disclosure and notes their limitations.

---

<sup>243</sup> Jordan Bar Am, Vinit Doshi, Anandi Malik & Sherry Frey, *Consumers Care About Sustainability—and Back It Up with Their Wallets*, McKinsey & Co. (Feb. 6, 2023), <https://www.mckinsey.com/industries/consumer-packaged-goods/our-insights/consumers-care-about-sustainability-and-back-it-up-with-their-wallets/> [<https://perma.cc/2XNY-8WYS>].

<sup>244</sup> For example, the graphics design app Procreate recently issued a pledge to never incorporate generative AI into its system. This pledge seems, in part, to be a way to distinguish its products from competitors such as Adobe. See Aisha Malik, *Procreate Takes a Stand Against Generative AI, Vows to Never Incorporate the Tech into Its Products*, TechCrunch (Aug. 19, 2024, 7:20 AM), <https://techcrunch.com/2024/08/19/procreate-takes-a-stand-against-generative-ai-vows-to-never-incorporate-the-tech-into-its-products/> [<https://perma.cc/M9DK-WLX8>]; Katelyn Chedraoui, *Why Procreate’s Anti-AI Pledge Is Resonating with Its Creators*, CNET (Aug. 30, 2024, 5:30 AM), <https://www.cnet.com/tech/services-and-software/why-procreates-anti-ai-pledge-is-resonating-with-its-creators/> [<https://perma.cc/QNU6-N28R>].

<sup>245</sup> Kysar, *supra* note 14, at 625–28.

<sup>246</sup> *Id.*; see also Jeanne C. Fromer, *The Unregulated Certification Mark(et)*, 69 *Stan. L. Rev.* 121, 158–67 (2017) (discussing this problem in relation to certification marks).

*i. Provenance Tracking*

In reaction to the threat of deepfakes, AI technology companies, content producers, and platforms have proposed a variety of technical measures that might be used to identify AI-generated misinformation. Many of these technical solutions could also be used to help identify AI-generated creative works generally.

One approach is the use of AI-detection software, often known as an AI “classifier.” These tools generally employ their own AI systems trained to identify features unique to AI-generated text or images.<sup>247</sup> But the technology has proven unreliable. For example, OpenAI suspended its own classifier because of its “low rate of accuracy.”<sup>248</sup>

A more promising avenue is offered by technical standards that could be used to mark AI-generated materials at their inception. The Coalition for Content Provenance and Authenticity (“C2PA”), a consortium of technology and media companies, has launched an interoperable standard that allows for tracking the provenance of any content, including whether the content is AI-generated.<sup>249</sup> Another approach is watermarking, which physically modifies generated text or images in ways imperceptible to the human eye, to allow cryptographic techniques to determine whether the content is AI-generated.<sup>250</sup>

Both provenance verification and watermarking are imperfect solutions. Many producers may be too tempted by the prospect of

---

<sup>247</sup> Tiffany Hsu & Steven Lee Myers, Another Side of the A.I. Boom: Detecting What A.I. Makes, *N.Y. Times* (May 19, 2023), <https://www.nytimes.com/2023/05/18/technology/ai-chat-gpt-detection-tools.html>.

<sup>248</sup> New AI Classifier for Indicating AI-Written Text, OpenAI (Jan. 31, 2023), <https://openai.com/blog/new-ai-classifier-for-indicating-ai-written-text> [<https://perma.cc/V553-7JAE>]; Emilia David, OpenAI Can’t Tell if Something Was Written by AI After All, *The Verge* (July 25, 2023, 5:30 PM), <https://www.theverge.com/2023/7/25/23807487/openai-ai-generated-low-accuracy> [<https://perma.cc/8LYX-75J9>]; see also Geoffrey A. Fowler, Detecting AI May Be Impossible. That’s a Big Problem for Teachers, *Wash. Post* (June 2, 2023, 12:30 PM), <https://www.washingtonpost.com/technology/2023/06/02/turnitin-ai-cheating-detector-accuracy/> (discussing the reliability problems of Turnitin’s AI detector).

<sup>249</sup> C2PA Explainer, <https://c2pa.org/specifications/specifications/1.2/explainer/Explainer.html> [<https://perma.cc/6ECT-ZRMK>] (last visited Oct. 29, 2024).

<sup>250</sup> See, e.g., Sayash Kapoor & Arvind Narayanan, How to Prepare for the Deluge of Generative AI on Social Media, *Knight First Amend. Inst.* (June 16, 2023), <https://knightcolumbia.org/content/how-to-prepare-for-the-deluge-of-generative-ai-on-social-media> [<https://perma.cc/8HMM-RF9Z>]; see also John Kirchenbauer et al., A Watermark for Large Language Models (May 1, 2024), <https://arxiv.org/pdf/2301.10226> [<https://perma.cc/WDB7-32QJ>] (studying watermarking and proposing a watermarking framework that more accurately identifies AI-generated text).

producing cheap AI-generated material to agree to opt in to a technical measure that would easily reveal that information to the public.<sup>251</sup> Watermarks are also easy to remove or modify after the fact.<sup>252</sup> Some kind of regulatory oversight—similar, perhaps, to the legal protection for technical measures designed to limit access to copyrighted works—would likely be necessary to make watermarks and other technical solutions more reliable.<sup>253</sup> This possibility, and recent state legislation that has mandated the use of watermarks and other provenance-tracing technologies, is discussed further below.

### *ii. Certification*

Content production companies could also commit to certifying that their content was not produced by generative AI. In contrast to watermarks, a certification scheme is explicitly public-facing: it would require the creation of a certification mark that could be affixed to specific media products to signify that generative AI was not used in their creation.<sup>254</sup> The use of certification marks to reflect consumer concerns over the process of a good's creation—what Margaret Chon has termed “marks of rectitude”<sup>255</sup>—is not new. “Fairtrade,” “organic,” and a variety of environment-focused ecolabels are ubiquitous in the modern economy.<sup>256</sup>

---

<sup>251</sup> Cf. Kapoor & Narayanan, *supra* note 250 (noting limitations of any opt-in approach). Indeed, OpenAI is internally divided over whether to release a watermarking feature in ChatGPT out of fear over losing users. See Wes Davis, OpenAI Won't Watermark ChatGPT Text Because Its Users Could Get Caught, *The Verge* (Aug. 4, 2024, 8:18 PM), <https://www.theverge.com/2024/8/4/24213268/openai-chatgpt-text-watermark-cheat-detection-tool> [<https://perma.cc/Y6AM-VWTR>].

<sup>252</sup> Mehrdad Saberi et al., Robustness of AI-Image Detectors: Fundamental Limits and Practical Attacks 3 (2024), <https://arxiv.org/pdf/2310.00076> [<https://perma.cc/WDB7-32QJ>]; Kate Knibbs, Researchers Tested AI Watermarks—and Broke All of Them, *Wired* (Oct. 3, 2023, 6:00 AM), <https://www.wired.com/story/artificial-intelligence-watermarking-issues/>.

<sup>253</sup> See generally 17 U.S.C. § 1201 (penalizing the circumvention of copyright protection systems). But see U.S. Copyright Off., Section 1201 of Title 17: A Report of the Register of Copyrights 51–52 (June 2017), <https://www.copyright.gov/policy/1201/section-1201-full-report.pdf> [<https://perma.cc/ZS6C-T8CP>] (noting criticism that Section 1201 has in fact failed to meaningfully curb circumvention).

<sup>254</sup> Certification marks are explicitly contemplated under U.S. trademark law. See 15 U.S.C. § 1054.

<sup>255</sup> See Margaret Chon, Marks of Rectitude, 77 *Fordham L. Rev.* 2311, 2311–12 (2009).

<sup>256</sup> See Introduction to Ecolabels and Standards for Greener Products, EPA, <https://www.epa.gov/greenerproducts/introduction-ecolabels-and-standards-greener-products> [<https://perma.cc/L6QM-9KFB>] (last visited Oct. 29, 2024).

At least one such certification scheme has already been launched to certify that television and film content has not been produced using AI.<sup>257</sup> While it does not appear this mark has been widely adopted, opting into this kind of certification scheme could in fact ultimately become a sound business decision. Studies have shown that consumers are willing to pay more for goods and services that carry desirable certifications, such as the “organic” label.<sup>258</sup>

But certification marks are only effective at protecting consumer interests if consumers and producers believe they convey accurate information.<sup>259</sup> This generally requires the creation of some kind of organization committed to maintaining the standards and accuracy of the certification. But such oversight can be difficult to maintain effectively. Privately managed certification standards are notoriously opaque and prone to manipulation by stakeholders who may seek to redefine the standard to serve their business interests or exclude competitors.<sup>260</sup> Moreover, in the absence of a centralized regulatory standard, multiple certification schemes can proliferate. This “clutter” can make it prohibitively difficult for consumers to understand precisely what a specific certification mark indicates.<sup>261</sup>

So, as with the use of technical measures, self-certification would likely be effective only if accompanied by significant opt-in and some regulatory standardization and oversight.<sup>262</sup> That being said, online platforms may be well-equipped to create, police, and even mandate their own certification standards via their terms of use. For example, the sales platform, Etsy, recently revised its terms to require that sellers mark AI-generated materials with the label “seller-prompted AI creation” in order

---

<sup>257</sup> About, Credo23, <https://credo23.com/ABOUT> [<https://perma.cc/8KFX-5GWG>] (last visited Oct. 29, 2024).

<sup>258</sup> See, e.g., Ellen Van Loo, Vincenzina Caputo, Rodolfo M. Nayga Jr., Jean-Francois Meullenet & Steven C. Ricke, Consumers’ Willingness to Pay for Organic Chicken Breast: Evidence from Choice Experiment, 22 *Food Quality & Preference* 603, 603 (2011); Dan Yu, The Value of Certification to Consumers: Evidence from the U.S. Organic Food Market 1 (Univ. of Minn., Twin Cities, Working Paper, Dec. 22, 2022), <https://ssrn.com/abstract=4299523> [<https://perma.cc/MU68-XSRE>].

<sup>259</sup> Fromer, *supra* note 246, at 158–67.

<sup>260</sup> *Id.* at 160–63.

<sup>261</sup> See Chon, *supra* note 255, at 2341–46 (identifying this problem and documenting it in the context of fair-trade coffee).

<sup>262</sup> Cf. Fromer, *supra* note 246, at 181–95 (articulating various forms of regulatory oversight to ensure that certification marks are reliable and not abused).



to avoid misleading consumers.<sup>263</sup> Etsy is somewhat unique in its commitment to human-made products,<sup>264</sup> but this example showcases how platforms could play a role in creating more reliable and enforceable certification standards.

## 2. Legislation

If, as the last Parts have argued, the undisclosed use of generative AI in content creation is both harmful and prevalent, one solution could be to mandate disclosure via legislation. The recent European Union (“EU”) AI Act is one such approach, though it falls short of mandating a comprehensive disclosure requirement for all AI-generated material. The EU AI Act includes a transparency requirement but limits it to an “AI system that generates or manipulates image, audio or video content constituting a deep fake.”<sup>265</sup> The EU AI Act, however, also provides a more general requirement that generative AI systems mark all outputs so that they can be detectable as AI-generated material using the kinds of provenance-tracing technical measures described above.<sup>266</sup> While this provenance-tracing requirement may apply to all AI-generated material, including creative content, it is unclear if the EU AI Act would, in practice, create a full disclosure regime for non-deepfake content.<sup>267</sup>

---

<sup>263</sup> Etsy Staff, What’s Etsy’s Stance on AI Creations?, Etsy (July 9, 2024), <https://www.etsy.com/seller-handbook/article/1275449912004> [<https://perma.cc/GZB9-CV8U>].

<sup>264</sup> See Keep Commerce Human, Etsy, <https://www.etsy.com/about?ref=fr> [<https://perma.cc/759T-96V5>] (last visited Oct. 29, 2024) (“In a time of increasing automation, it’s our mission to keep human connection at the heart of commerce.”).

<sup>265</sup> Council Regulation 2024/1689, ch. IV, art. 50, para. 4, 2024 O.J. (L) 1 (EU) (“Deployers of an AI system that generates or manipulates image, audio or video content constituting a deep fake, shall disclose that the content has been artificially generated or manipulated.”). The EU AI Act defines a “deep fake” as an “AI-generated or manipulated image, audio or video content that resembles existing persons, objects, places, entities or events and would falsely appear to a person to be authentic or truthful.” *Id.* at ch. I, art. 3, para. 60.

<sup>266</sup> *Id.* at ch. IV, art. 50, para. 2 (“Providers of AI systems, including general-purpose AI systems, generating synthetic audio, image, video or text content, shall ensure that the outputs of the AI system are marked in a machine-readable format and detectable as artificially generated or manipulated.”); see also *supra* Paragraph III.A.1.i (discussing provenance-tracing as a possible solution to discovering undisclosed AI).

<sup>267</sup> See Gernot Fritz, Theresa Ehlen & Tina Fokter Cuvan, EU AI Act Unpacked #8: New Rules on Deepfakes, Freshfields Bruckhaus Deringer (June 26, 2024, 2:29 PM), <https://technologyquotient.freshfields.com/post/102jb19/eu-ai-act-unpacked-8-new-rules-on-deepfakes> [<https://perma.cc/Y24X-3NV8>] (noting the Act’s primary focus on deepfakes, and in particular that only deepfakes seem to require direct disclosure of the use of AI to the public).

In the United States, it is unclear if Congress will pass legislation imposing a comprehensive disclosure requirement. Introduced in late 2023, the Schatz-Kennedy AI Labeling Act would have mandated that all generative AI content include “a clear and conspicuous notice, as appropriate for the medium of the content, that identifies the content as AI-generated content.”<sup>268</sup> This bill, however, has not gained political traction.

At least one state, however, has begun to embrace mandatory disclosure requirements. In late 2024, California became the first state to pass legislation requiring any large generative AI company to provide the public with a “detection tool” capable of ascertaining whether an image, video, or audio recording was generated by the company’s technology, as well as providing information regarding the material’s provenance or history of modification.<sup>269</sup> The new law also requires large companies to include a “latent” disclosure (i.e., a watermark that is detectable but not visible) of the AI origins of any image, video, or audio content, and the option for users to include a “conspicuous” (i.e., visible) disclosure.<sup>270</sup> This legislation would essentially create a mandatory watermarking and provenance detection regime, of the kind described above,<sup>271</sup> at least for non-text-based AI content.

In contrast to a voluntary opt-in disclosure system, AI transparency legislation, such as California’s law, may face hurdles. In particular, some have argued that a broad disclosure requirement would constitute compelled speech in violation of the First Amendment.<sup>272</sup> While an

---

<sup>268</sup> AI Labeling Act, S. 2691, 118th Cong. § 2(B)(i) (2023).

<sup>269</sup> California AI Transparency Act, Cal. Bus. & Prof. Code §§ 22757–22757.6 (West 2023); Tayla Pinto, California’s New AI Laws: What They Mean for Developers, Michalsons (Oct. 21, 2024), <https://www.michalsons.com/blog/californias-new-ai-laws-what-they-mean-for-developers/76026> [<https://perma.cc/6ATE-YGZ>].

<sup>270</sup> *Id.*

<sup>271</sup> See *supra* Paragraph III.A.1.i.

<sup>272</sup> See, e.g., Artificial Intelligence, Free Speech, and the First Amendment, FIRE, <https://www.thefire.org/research-learn/artificial-intelligence-free-speech-and-first-amendment> [<https://perma.cc/Z9NE-H7BZ>] (last visited Oct. 29, 2024) (“In most cases, the government can no more compel an artist to disclose whether they created a painting from a human model as opposed to a mannequin than it can compel someone to disclose that they used artificial intelligence tools in creating an expressive work. Government may be able to compel disclosure in a few narrow circumstances, such as in election ads, or when seeking a copyright, but these circumstances are exceptions to the general rule.”); cf. R. Polk Wagner, Filters and the First Amendment, 83 *Minn. L. Rev.* 755, 778 (1999) (examining First Amendment implications of internet content labeling); Eric Goldman, The Constitutionality of Mandating Editorial Transparency, 73 *Hastings L.J.* 1203, 1217–18 (2022) (“Mandatory editorial

assessment of this argument's viability is beyond the scope of this Article,<sup>273</sup> it suggests that any mandatory disclosure regime may end up subject to lengthy litigation.<sup>274</sup>

### *B. Penalizing Deceptive Omissions via FTC Enforcement*

An affirmative disclosure regime, of the kinds described above, would resolve the problem of hidden AI authorship by ensuring that all AI-generated creative goods are marked as such. But more targeted regulatory solutions could also address this problem. Rather than impose blanket disclosure requirements, regulators could target specific instances where the omission of information regarding a work's provenance materially misleads consumers. The FTC, which is explicitly empowered to protect consumers from deceptive trade practices, is the most likely regulator to take on this task.

---

transparency laws also should qualify as 'compelled speech.'"). See generally Eugene Volokh, Mark A. Lemley & Peter Henderson, Freedom of Speech and AI Output, 3 J. Free Speech L. 651 (2023) (discussing applicability of First Amendment to AI generally); Valerie C. Brannon, Victoria L. Killion, Whitney K. Novak & L. Paige Whitaker, Cong. Rsch. Serv., IF12388, First Amendment Limitations on Disclosure Requirements (2023), <https://crsreports.congress.gov/product/pdf/IF/IF12388> [<https://perma.cc/WL6E-4QLY>] (articulating various standards applied by courts to determine when mandatory disclosure requirements become unconstitutional compelled speech, for commercial disclosure and product labeling, for example).

<sup>273</sup> In the context of privacy notices, Felix Wu has argued that compelled speech concerns may be unfounded because such notices "involve a type of speech that implicates neither the listener interests at stake in commercial speech restrictions nor the speaker interests at stake in compelling individuals to speak." Felix T. Wu, The Constitutionality of Consumer Privacy Regulation, 2013 U. Chi. Legal F. 69, 88; cf. David S. Ardia & Evan Ringel, First Amendment Limits on State Laws Targeting Election Misinformation, 20 First Amend. L. Rev. 291, 298 (2022) (concluding state statutes regulating election misinformation that target defamatory or fraudulent speech, speech that harms the election process, or speech that intimidates voters is likely constitutional, but statutes that target other types of speech will likely run afoul of the First Amendment).

<sup>274</sup> The legislation also does not provide much guidance on when material should be considered "AI-generated." See *supra* Subsection II.C.2 (discussing the many different ways AI can be implicated in a work's creation). Failure to adequately define exactly when disclosure is required may ultimately make the legislation ineffective. Scholars have argued that mandatory disclosure regimes are far less likely to succeed when definitions are vague or overbroad. See Omri Ben-Shahar & Carl E. Schneider, The Failure of Mandated Disclosure, 159 U. Pa. L. Rev. 647, 684–85 (2011); see also Omri Ben-Shahar, Carcinogenic Laws: Coffee Shop Cancer Warnings Do More Harm Than Good, *Forbes* (Mar. 31, 2018, 10:21 PM), <https://www.forbes.com/sites/omribensshahar/2018/03/31/carcinogenic-laws-coffee-shop-cancer-warnings-do-more-harm-than-good/> (criticizing the overbreadth of California's carcinogen-warning law).

The FTC’s authority to police “unfair or deceptive acts or practices”<sup>275</sup> is expansive.<sup>276</sup> The FTC may target any material representation, omission, or practice that would likely mislead a reasonable consumer.<sup>277</sup> At its disposal is the ability to promulgate rules defining acts that are unfair or deceptive, to bring civil actions for injunctive relief on behalf of consumers, to (in certain instances) impose financial penalties, and more.<sup>278</sup>

The FTC has already set its sights on generative AI but has generally focused on antitrust concerns, the use of the technology to facilitate consumer fraud, exploitation of personal data in machine learning, and the unauthorized use of creative works in machine learning.<sup>279</sup> However, the FTC has also noted that consumers feel “deceived when authorship does not align with consumer expectations, such as when a consumer thinks a work has been created by a particular musician or other artist, but it has been generated by someone else using an AI tool.”<sup>280</sup> If the FTC

<sup>275</sup> 15 U.S.C. § 45(a)(1).

<sup>276</sup> See Woodrow Hartzog & Daniel J. Solove, *The Scope and Potential of FTC Data Protection*, 83 *Geo. Wash. L. Rev.* 2230, 2249–50 (2015) (“The concept of unfairness is thus quite intentionally broad and subject to refinement over time. Instead of specific categories, the FTC’s unfairness authority is limited to instances where there is an actual unavoidable harm or likelihood thereof, which is not outweighed by countervailing benefits to consumers or competition.”).

<sup>277</sup> Letter from James C. Miller III, Chairman, U.S. Fed. Trade Comm’n, to John D. Dingell, Chairman, Comm. on Energy & Com., U.S. House of Reps., *FTC Policy Statement on Deception* (Oct. 14, 1983) [hereinafter *Miller Letter*].

<sup>278</sup> *A Brief Overview of the Federal Trade Commission’s Investigative, Law Enforcement, and Rulemaking Authority*, Fed. Trade Comm’n (revised May 2021), <https://www.ftc.gov/about-ftc/mission/enforcement-authority> [<https://perma.cc/448C-AY4B>]; see also Hartzog & Solove, *supra* note 276, at 2277–83 (discussing privacy and data security cases brought by the FTC based on a more flexible understanding of harm).

<sup>279</sup> Lina M. Khan, *Opinion, Lina Khan: We Must Regulate A.I. Here’s How*, *N.Y. Times* (May 3, 2023), <https://www.nytimes.com/2023/05/03/opinion/ai-lina-khan-ftc-technology.html> (articulating these issues); U.S. Fed. Trade Comm’n, *Generative Artificial Intelligence and the Creative Economy Staff Report: Perspectives and Takeaways 5–7* (2023) [hereinafter *FTC Staff Report*], [https://www.ftc.gov/system/files/ftc\\_gov/pdf/12-15-2023AICEStaffReport.pdf](https://www.ftc.gov/system/files/ftc_gov/pdf/12-15-2023AICEStaffReport.pdf) [<https://perma.cc/U9UH-TRD8>]; Press Release, U.S. Fed. Trade Comm’n, *FTC Proposes New Protections to Combat AI Impersonation of Individuals* (Feb. 15, 2024), <https://www.ftc.gov/news-events/news/press-releases/2024/02/ftc-proposes-new-protections-combat-ai-impersonation-individuals> (proposing rule that would target AI systems that enable the creation of deceptive deepfakes) [<https://perma.cc/P8WU-7MNN>].

<sup>280</sup> Fed. Trade Comm’n, *supra* note 167, at 5; see also *id.* at 7 (“Participants said that when generative AI tools use artists’ faces, voices, and performances without permission to make digital impersonations, it can . . . create consumer confusion . . .”); *FTC Staff Report*, *supra* note 279, at 33–35 (reporting direct quotes from participants of the FTC’s October 2023

indeed begins to consider this use of AI to be a form of consumer deception, it could begin targeting this behavior using its arsenal of enforcement tools.

While the FTC has, thus far, seemed to limit its attention to the consumer deception caused by AI-mimicry of an artist’s voice, likeness, or style, it could also potentially tackle the problem of undisclosed AI use in creative works more generally. As discussed above, many consumers have clear preferences for human-made creative works.<sup>281</sup> By motivating consumers to purchase a product they may not otherwise purchase, the omission of information regarding a work’s provenance could be considered deceptive.<sup>282</sup> While such an omission is likely not presumptively material,<sup>283</sup> the empirical documentation of consumers’ preference for human-made works<sup>284</sup> (and, if available, evidence that a producer deliberately omitted the information to increase sales<sup>285</sup>) could bolster the case for materiality.<sup>286</sup>

### *C. Using Intellectual Property Law to Incentivize Disclosure*

Targeting hidden AI authorship through a legislative disclosure mandate or FTC enforcement faces barriers. Legislation is politically fraught and could face First Amendment scrutiny.<sup>287</sup> FTC enforcement could also be hamstrung by judicial oversight, especially now that the Supreme Court has dismantled the *Chevron* deference framework.<sup>288</sup>

This Section considers a third option that could be used to motivate disclosure of AI authorship: private intellectual property litigation. As

---

Creative Economy and Generative AI roundtable concerning the theme of “Deepfakes, Impersonations, and/or Fan confusion”).

<sup>281</sup> See *supra* Subsection II.A.2.

<sup>282</sup> See Miller Letter, *supra* note 277, at 5 (“A ‘material’ misrepresentation or practice is one which is likely to affect a consumer’s choice of or conduct regarding a product.”).

<sup>283</sup> See *id.* (discussing limited circumstances where an omission is presumptively material, such as omissions regarding the health and safety of a product).

<sup>284</sup> See *supra* notes 58–64 and accompanying discussion.

<sup>285</sup> See *supra* Subsection I.B.1 (discussing financial incentives to obfuscate).

<sup>286</sup> Miller Letter, *supra* note 277, at 5 (noting that the FTC “may require evidence that the claim or omission is likely to be considered important by consumers”).

<sup>287</sup> See *supra* note 272 and accompanying text.

<sup>288</sup> *Loper Bright Enters. v. Raimondo*, 144 S. Ct. 2244, 2247–48 (2024); see also Tamara Brightwell et al., “*Chevron* Is Overruled”: How *Loper Bright* Will Change the Regulatory Law Landscape, Wilson Sonsini (July 8, 2024), <https://www.wsgr.com/en/insights/chevron-is-overruled-how-loper-bright-will-change-the-regulatory-law-landscape.html> [<https://perma.cc/3ZJC-J3PW>] (noting difficulties that *Loper Bright* poses for FTC enforcement).

noted above, hidden AI authorship already implicates copyright law, and, to a lesser extent, trademark and the right of publicity. These areas of law could potentially be used to raise the financial stakes of non-disclosure. While IP may not alone be sufficient to fully address the problem of hidden AI authorship, it presents a readily useable framework, one that could potentially operate in conjunction with some of the solutions described above.

### *1. Copyright Misuse*

As the last Parts examined, copyright is already intertwined with the problem of undisclosed AI authorship. Disclosure to the public and disclosure to the Copyright Office are two sides of the same coin; a content producer cannot feasibly do one without the other.<sup>289</sup> And the risk of receiving no copyright registration, or a thin, barely enforceable registration, compounds content producers' incentives to keep secret their use of generative AI.<sup>290</sup> In light of copyright's existing role in the problem of undisclosed AI authorship, copyright's internal doctrines may also offer some solutions.<sup>291</sup>

What Jason Mazzone has called “copyfraud”—the false assertion of copyright in public domain materials—is not a new phenomenon.<sup>292</sup> In the past, however, copyfraud was limited to a relatively small number of scenarios. Someone might publish a work whose copyright has expired (like a copy of an old painting) or an unprotectable government document (like a judicial opinion), and falsely assert that the work was under copyright.<sup>293</sup> In these scenarios, however, there is relatively little economic upside to the decision to engage in copyfraud, as an unprotectable work can be published by anyone, yielding a race to the bottom.

But we are reaching a point where, for perhaps the first time in history, *new valuable content is unprotectable by copyright*. AI-generated content is in the public domain, not because its copyright has expired, but because

---

<sup>289</sup> See *supra* notes 116–17 and accompanying text.

<sup>290</sup> See *supra* Subsection I.B.2.

<sup>291</sup> As I have argued in previous work, copyright law often plays both an explicit and tacit role in regulating industry behavior. See generally Jacob Victor, *Reconceptualizing Compulsory Copyright Licenses*, 72 *Stan. L. Rev.* 915 (2020) (examining varieties of ways that copyright law regulates licensing to promote socially valuable new technologies).

<sup>292</sup> Mazzone, *supra* note 123, at 1028–29.

<sup>293</sup> *Id.* at 1039–42, 1046.

of its inherent lack of human authorship. Unlike the false claiming of old public domain materials, the false claiming of such AI-generated materials could be highly lucrative.<sup>294</sup>

The mechanisms designed to prohibit such overclaiming could be used to tackle this problem, and, in turn, disincentivize the hiding of AI authorship from consumers.<sup>295</sup> Unfortunately, such mechanisms are relatively few. As noted above, the Copyright Act contains a small number of criminal penalties for, among other things, the placement of a copyright notice on unprotectable materials with fraudulent intent and the false representation of material facts to the Copyright Office in a copyright registration application.<sup>296</sup> But the minimal penalty for these acts—a fine of \$2,500—renders the provisions virtually toothless.<sup>297</sup> These provisions could be reformed to increase the financial penalties for fraudulent overclaiming, but this would require legislative change.<sup>298</sup>

Litigation, however, could provide some checks on obfuscation. In court, Copyright Office registrations “constitute prima facie evidence of the validity of the copyright and of the facts stated in the certificate.”<sup>299</sup> But this evidence is rebuttable. If evidence is presented that a registrant failed to disclose meaningful AI contributions to a registered work, a court could decline to find copyrightability and/or narrow the scope of the copyright on lack-of-authorship grounds.<sup>300</sup> However, the putative copyright owner would face few, if any, financial penalties if this were to occur. And they could still prevail in infringement litigation if it turns out that even the narrowed copyright was infringed.<sup>301</sup>

But there is a more powerful litigation-based tool that could be used to financially disincentivize copyright owners from hiding the role of AI in a work: copyright misuse. Copyright misuse is a common law doctrine that is derived from patent misuse, itself a common law doctrine derived

---

<sup>294</sup> See *supra* Part I.

<sup>295</sup> As noted above, disclosure for purposes of determining copyrightability and disclosure to the public go hand-in-hand. See *supra* notes 116–17 and accompanying text.

<sup>296</sup> 17 U.S.C. § 506(c), (e).

<sup>297</sup> *Id.*; see *supra* Section I.C.

<sup>298</sup> Such legislative change could be worthwhile, but the main goal of this Part is to examine how *existing* law could be used to tackle the problem of undisclosed AI authorship.

<sup>299</sup> 17 U.S.C. § 410.

<sup>300</sup> *Thaler v. Perlmutter*, 687 F. Supp. 3d 140, 145–50 (D.D.C. 2023).

<sup>301</sup> See *infra* notes 322–24 and accompanying text (discussing such a hypothetical scenario in more detail).

from the broad equitable principle of unclean hands.<sup>302</sup> Patent misuse emerged in the mid-twentieth century as a means of preventing patent holders from overreaching: utilizing a patent's grant of exclusive rights to police activities beyond the scope of the patent or to restrain trade more broadly.<sup>303</sup> Today, patent misuse is often (though not always) understood to involve the use of a patent in ways that violate antitrust laws.<sup>304</sup>

Copyright misuse shares patent misuse's concern with overreach, but is less tethered to antitrust law.<sup>305</sup> Rather, as one treatise notes, many copyright "misuse cases are best thought of as abuse-of-process cases."<sup>306</sup> Indeed, several cases have found that an attempt to assert copyright protection over public domain materials can constitute copyright misuse, even in the absence of an antitrust violation.<sup>307</sup> Such public domain materials have included unprotectable architectural design elements,<sup>308</sup> unprotectable sales techniques,<sup>309</sup> unprotectable tax data,<sup>310</sup> unprotectable

<sup>302</sup> Brett Frischmann & Dan Moylan, *The Evolving Common Law Doctrine of Copyright Misuse: A Unified Theory and Its Application to Software*, 15 *Berkeley Tech. L.J.* 865, 867 (2000); Thomas F. Cotter, *Misuse*, 44 *Hous. L. Rev.* 901, 901–02 (2007).

<sup>303</sup> Dan L. Burk, *Anticircumvention Misuse*, 50 *UCLA L. Rev.* 1095, 1115 (2003).

<sup>304</sup> See Herbert Hovenkamp, Mark D. Janis, Mark A. Lemley, Christopher R. Leslie & Michael A. Carrier, *IP and Antitrust: An Analysis of Antitrust Principles Applied to Intellectual Property Law* § 3.02[D] (3d ed. Supp. 2023) (outlining the debate over patent misuse's connection to antitrust law).

<sup>305</sup> See *id.* § 3.04[B][1], at 3-66 to -67; Deepa Varadarajan, *The Uses of IP Misuse*, 68 *Emory L.J.* 739, 759 (2019).

<sup>306</sup> Hovenkamp et al., *supra* note 304, § 3.04[B][2][a], at 3-69.

<sup>307</sup> As discussed further below, if found to have engaged in copyright misuse, the copyright owner cannot recover for even the *otherwise protectable elements* of their work. See *infra* notes 322–24 and accompanying text.

<sup>308</sup> See *Design Basics, LLC v. Petros Homes, Inc.*, 240 F. Supp. 3d 712, 721 (N.D. Ohio 2017) (denying summary judgment because "[d]efendants have produced specific information that, if true, could be construed by a jury to support the misuse of copyright defense"); *Home Design Servs., Inc. v. Park Square Enters.*, No. 02-cv-00637, 2005 WL 1027370, at \*12 (M.D. Fla. May 2, 2005).

<sup>309</sup> *Huthwaite, Inc. v. Randstad Gen. Partner (US), LLC*, No. 06-cv-01548, 2006 WL 3065470, at \*9 (N.D. Ill. Oct. 24, 2006) (refusing to strike copyright misuse affirmative defense when plaintiff alleged that the defendant had copied sales techniques, which were unprotectable ideas, suggesting plaintiff is "attempting to use its copyrighted books to cover the unprotectable ideas within [protectable] books").

<sup>310</sup> See *Assessment Techs. of WI, LLC v. WIREdata, Inc.*, 350 F.3d 640, 647 (7th Cir. 2003) (Posner, J.) (noting in dicta that "[t]he data in the municipalities' tax-assessment databases are beyond the scope of AT's copyright. . . . The argument for applying copyright misuse . . . is that for a copyright owner to use an infringement suit to obtain property protection, here in data, that copyright law clearly does not confer, hoping to force a settlement or even achieve an outright victory over an opponent that may lack the resources or the legal sophistication to resist effectively, is an abuse of process").



microprocessor cards,<sup>311</sup> unprotectable articles of clothing,<sup>312</sup> unprotectable medical records,<sup>313</sup> and more.<sup>314</sup> This use of copyright misuse—as a means of policing the overclaiming of public domain materials—may be something of a recent trend, and some courts may still need convincing that it is within the scope of the doctrine.<sup>315</sup> But several scholars, especially Judge Richard Posner, have made compelling arguments that understanding copyright misuse as a means of policing overclaiming is an essential way of ensuring the doctrine can further copyright’s policy agenda.<sup>316</sup>

<sup>311</sup> See *Alcatel USA, Inc. v. DGI Techs., Inc.*, 166 F.3d 772, 793 (5th Cir. 1999) (“A reasonable juror could conclude, based on the licensing agreement, that ‘DSC has used its copyrights to indirectly gain commercial control over products DSC does not have copyrighted,’ namely, its microprocessor cards.”).

<sup>312</sup> See *Galiano v. Harrah’s Operating Co.*, No. 00-cv-00071, 2004 WL 1057552, at \*14 (E.D. La. May 10, 2004) (“Two dimensional artwork for costumes is copyrightable. As addressed supra, the clothing/uniform designs depicted by that artwork in the Collection is not copyrightable. . . . If the facts show that Gianna . . . knew that the Copyright Office would not have given a registration certificate on the Collection if the nature of the work had been accurately described as ‘designs for uniforms and costumes’ rather than two dimensional ‘artwork for wearing apparel’, then the copyright may be invalid.” (emphases omitted)). Note that this decision seems to treat copyright misuse and fraud on the copyright office as synonymous. *Id.* at \*1 (describing “fraud and misuse of the Copyright Office”).

<sup>313</sup> See *Shloss v. Sweeney*, 515 F. Supp. 2d 1068, 1080–81 (N.D. Cal. 2007) (“Plaintiff undertook to write a scholarly work on Lucia Joyce—the type of creativity that the copyright laws exist to facilitate. . . . Plaintiff was allegedly intimidated from using (1) non-copyrightable fact works such as medical records and (2) works to which Defendants did not own or control copyrights, such as letters written by third parties. The Court finds that Plaintiff has sufficiently alleged a nexus between Defendants’ actions and the Copyright Act’s public policy of promoting creative expression to support a cause of action for copyright misuse.”).

<sup>314</sup> See *F.E.L. Publ’ns, Ltd. v. Cath. Bishop of Chi.*, No. 81-cv-01333, 1982 WL 19198, at \*5 n.9 (7th Cir. Mar. 25, 1982) (noting in dicta that “it is copyright misuse to exact a fee for the use of a musical work which is already in the public domain”). Courts have also held that when the creator of a derivative work attempts to assert copyright in material that is in fact part of the underlying work, this can be copyright misuse. See *qad. Inc. v. ALN Assocs., Inc.*, 770 F. Supp. 1261, 1267 (N.D. Ill. 1991), *aff’d*, 974 F.2d 834 (7th Cir. 1992) (“Here qad’s misuse was even more egregious: It used its copyright to sue ALN and to restrain it from the use of material over which qad itself had no rights. That is a misuse of both the judicial process and the copyright laws.”).

<sup>315</sup> 5 Patry, *supra* note 124, § 17:128; see also, e.g., *Michael Grecco Prods., Inc. v. Valuewalk, LLC*, 345 F. Supp. 3d 482, 510 (S.D.N.Y. 2018) (noting that copyright misuse “has not been firmly established in the Second Circuit and it is unclear whether an antitrust violation is required for its use”).

<sup>316</sup> See William F. Patry & Richard A. Posner, *Fair Use and Statutory Reform in the Wake of Eldred*, 92 Calif. L. Rev. 1639, 1659 (2004) (“But where the warning grossly and intentionally exaggerates the copyright holder’s substantive or remedial rights, to the prejudice of publishers of public-domain works, the case for invoking the doctrine of copyright misuse seems to us compelling.”); *Assessment Techs. of WI, LLC v. WIREdata, Inc.*, 350 F.3d 640,

The interrelated doctrines of fraud on the Copyright Office and unclean hands are sometimes conflated with copyright misuse, even though they are distinct.<sup>317</sup> Citing fraud on the Copyright Office and/or unclean hands, some courts have indeed held that material omissions to the Copyright Office can render a copyright registration invalid.<sup>318</sup> As mentioned above, however, this is a high bar to meet, requiring both materiality and evidence of fraudulent intent.<sup>319</sup> Both doctrines also require that the defendant be *directly* harmed by the plaintiff's false representation.<sup>320</sup> Moreover, a finding of fraud on the Copyright Office does not constitute an absolute bar to recovery, but just a temporary hiccup; the plaintiff may file a new registration and commence the lawsuit again.<sup>321</sup>

In contrast, as one treatise explains: copyright “[m]isuse . . . properly focuses on the registrant’s behavior, not defendant’s, and the adverse impact on the public.”<sup>322</sup> Consistent with this rationale, a finding of copyright misuse *completely* prevents enforcement of a copyright interest,

---

647 (7th Cir. 2003) (Posner, J.) (adopting this reasoning in dicta); see also Varadarajan, *supra* note 305, at 770–71 (discussing Judge Posner’s argument for using copyright misuse to address overclaiming and other scholars’ views on using it to deter copyright owners from threatening litigation); David S. Olson, *First Amendment Based Copyright Misuse*, 52 Wm. & Mary L. Rev. 537, 537–38 (2010) (arguing that copyright misuse has deterrent effects that will better protect free speech rights by discouraging copyright holders from abusing their rights to control others’ expression).

<sup>317</sup> 5 Patry, *supra* note 124, §§ 17:127–17:128.

<sup>318</sup> See, e.g., *Original Appalachian Artworks, Inc. v. Toy Loft, Inc.*, 684 F.2d 821, 828 (11th Cir. 1982) (finding “omissions or misrepresentations in a copyright application can render the registration invalid”).

<sup>319</sup> See *supra* notes 125–26 and accompanying text.

<sup>320</sup> See 5 Patry, *supra* note 124, § 17:128 (“Misuse differs, significantly, from fraud on the Copyright Office and unclean hands by not requiring that the defendant have been prejudiced by the conduct in question . . . .”); see also *Mitchell Bros. Film Grp. v. Cinema Adult Theater*, 604 F.2d 852, 863 (5th Cir. 1979) (“The maxim of unclean hands is not applied where plaintiff’s misconduct is not directly related to the merits of the controversy between the parties . . . . The alleged wrongdoing of the plaintiff does not bar relief unless the defendant can show that he has personally been injured by the plaintiff’s conduct.”); 4 Nimmer & Nimmer, *supra* note 125, § 13.09[B]. In contrast, as noted below, copyright misuse employs a much broader conception of harm. See 5 Patry, *supra* note 124, § 17:128; *infra* note 322 and accompanying text.

<sup>321</sup> 5 Patry, *supra* note 124, § 17:126 (“[T]he applicant may resubmit a new, accurate application and refile the suit, and is disadvantaged only by its inability to receive statutory damages and attorney’s fees.”); see also *Morgan, Inc. v. White Rock Distilleries, Inc.*, 230 F. Supp. 2d 104, 109 (D. Me. 2002) (dismissing a copyright infringement claim because the copyright was not lawfully registered, but explaining that “[t]he Court finds that the dismissal should be without prejudice since a copyright owner may register his claim at any time during the life of the copyright”).

<sup>322</sup> 5 Patry, *supra* note 124, § 17:128.

unless and until the misuse is cured.<sup>323</sup> It is this feature that renders the doctrine a potentially powerful tool in combatting the hiding of AI authorship. The *Zarya of the Dawn* comic book, discussed above, provides a helpful example. Recall that the Copyright Office found this comic book could only be protected by a thin compilation copyright; the AI-generated illustrations were unprotectable. Consider a hypothetical scenario: the comic book’s putative author, Kristina Kashtanova, successfully registers the entire comic book without disclosing their use of AI (as they indeed attempted to do).<sup>324</sup> They later bring an infringement lawsuit against an unauthorized publisher, but during litigation, it emerges that the images in the comic book were AI-generated. A court would be within its rights to narrow the scope of the copyright interest to cover only the compilation of the public-domain images (as the Copyright Office indeed found). But, if that were to happen, the unauthorized publisher could still be held liable for infringing this compilation-based copyright. However, if the defendant successfully asserts copyright misuse, based on the overclaiming of public domain materials, Kashtanova would be barred from enforcing *even their valid copyright interest*, meaning they would have absolutely no recourse to prevent any unauthorized distribution. This would make the comic book essentially worthless.

A finding of copyright misuse not only renders the copyright owner completely unable to enforce their copyright (even a valid interest), but also carries a potential financial penalty. While misuse generally functions as an affirmative defense to copyright infringement claims,

---

<sup>323</sup> Id. (“Misuse . . . [does] not requir[e] that the defendant have been prejudiced by the conduct in question; instead, it is enough that a public policy has been violated, such as . . . extending control over public domain elements through illegitimate derivative registrations, or impermissibly extending the term of copyright through such registrations.”); Hovenkamp et al., *supra* note 304, § 3.04[B][2], at 3-73 to -74 (“A party that overclaims its copyrights . . . may lose not only the too-broad claim but also the right to enforce the copyright against those who clearly do infringe. Conceived in this way, the rationale for this set of misuse cases is not concerned so much with competition as with integrity, both of the judicial process and of copyright policy.”); see *Lasercomb Am., Inc. v. Reynolds*, 911 F.2d 970, 979 (4th Cir. 1990) (finding a misuse defense available as an absolute bar to copyright enforcement even if defendant was not personally harmed by the plaintiff’s anticompetitive conduct). If the plaintiff cures the misuse, they are still barred from bringing any claim that accrued during the period the misuse was occurring, though some courts have suggested otherwise. See also 4 Nimmer & Nimmer, *supra* note 125, § 13.09[A][2][b] nn.40.1–40.2 (showing courts’ differing views regarding whether a copyright claimant can reassert their infringement claim after curing the misuse).

<sup>324</sup> See *supra* Section I.C.

there is an open question as to whether it can also be brought as a *counterclaim* for damages.<sup>325</sup> Such an approach would allow a defendant to treat misuse as a species of common law unfair-competition tort that warrants compensatory or punitive damages. Commentators, however, have criticized the counterclaim approach, arguing that misuse should be considered purely an affirmative defense (or, at most, a counterclaim that only warrants declaratory relief).<sup>326</sup>

But even if a misuse counterclaim for damages is unavailable, the successful assertion of a misuse defense creates prime conditions for a court to order that the plaintiff pay the defendant's attorney's fees and costs. Though the Copyright Act is one of the few areas of American law where fee-shifting is explicitly contemplated,<sup>327</sup> in practice it can be quite difficult to obtain a fee-shifting order.<sup>328</sup> But "some element of moral blame" or "bad faith" on the part of the losing party can often be key to determining fee-shifting.<sup>329</sup> And there is indeed precedent that a finding of misuse can support a large award of attorneys' fees and costs.<sup>330</sup>

The risk of a copyright misuse finding could help incentivize content producers to disclose their use of generative AI to the Copyright Office (and, by default, to the public).<sup>331</sup> The complete unenforceability of any copyright interest (even a narrowed one), as well as additional financial penalties, is a risky prospect, that counterbalances any potential upside of keeping AI authorship hidden.<sup>332</sup>

---

<sup>325</sup> See *Amaretto Ranch Breedables, LLC v. Ozimals, Inc.*, 790 F. Supp. 2d 1024, 1033 (N.D. Cal. 2011) (examining this question); *Harrington v. 360 ABQ, LLC*, No. 22-cv-00063, 2022 WL 1567094, at \*3 (D.N.M. May 18, 2022) (same); *Nielsen Co. (US) v. Truck Ads, LLC*, No. 08-cv-06446, 2011 WL 221838, at \*6–7 (N.D. Ill. Jan. 24, 2011) (same).

<sup>326</sup> 4 Patry, *supra* note 124, § 10A:1; see 4 Nimmer & Nimmer, *supra* note 125, § 13.09 n.53.2 (suggesting the declaratory relief approach); see also Hovenkamp et al., *supra* note 304, § 3.06[B], at 3-92 (same).

<sup>327</sup> 17 U.S.C. § 505.

<sup>328</sup> 5 Nimmer & Nimmer, *supra* note 125, § 14.10[D] n.211 (outlining confusing precedent in this area).

<sup>329</sup> *Id.*

<sup>330</sup> See, e.g., *Omega S.A. v. Costco Wholesale Corp.*, 776 F.3d 692, 694, 696 (9th Cir. 2015) (affirming award of approximately \$400,000 fees to defendant who prevailed because of plaintiff's copyright misuse).

<sup>331</sup> See *supra* notes 116–17 and accompanying discussion.

<sup>332</sup> Moreover, in the context of licensing agreements between producers and distributors, a contract claim could also function as a *de facto* misuse cause of action. Large distributors, such as streaming services, generally require that licensors warrant that they own the licensed materials, free of any issue that could interfere with exploitation. See, e.g., *Netflix, Inc. & Nat'l Lampoon, Inc., License Agreement for Internet Transmission* § 5.2 (Apr. 21, 2009), <https://www.sec.gov/Archives/edgar/data/798078/000115752309003675/a5960803ex10.txt>

Of course, a copyright misuse defense is only available if the putative copyright owner chooses to sue for infringement, which may be relatively infrequent.<sup>333</sup> However, in many areas of both real and intellectual property law, the prospect of litigation (and, in particular, a loss of entitlement by operation of law) can influence behavior more broadly, even if specific litigation is infrequent.<sup>334</sup>

## *2. Trademark and the Right of Publicity*

There is a subset of AI-generated works that raises particular deception concerns: the use of an artist's name, voice, or likeness in an AI-generated work. As discussed above, the recent use of AI-generated technology to mimic the voices of Drake and The Weeknd shows how easily technology can coopt the specific markers of an artist's identity.<sup>335</sup> This particular use of generative AI opens up several additional causes of action that, if utilized, could also financially penalize producers that make use of generative AI.

The unauthorized use of an artist's name is the most straightforward. Marketing any work using an artist's name, without their permission, likely raises causes of action under both trademark law<sup>336</sup> and some states' right of publicity protections.<sup>337</sup> Thus, if an AI-generated song is

---

[<https://perma.cc/RFW2-ZGWZ>] (last visited Oct. 29, 2024). The undisclosed use of AI (and concomitant overclaiming of material unprotectable by copyright) might constitute a breach of such warranties, granting the distributor the right to terminate the agreement and potentially deny payment of royalties.

<sup>333</sup> Some courts, however, have held that a copyright misuse declaratory judgment action may be brought directly against copyright owners who have engaged in misuse. See Hovenkamp et al., *supra* note 304, § 3.06[B] at 3-92 n.362 (examining this issue).

<sup>334</sup> Doctrines such as adverse possession, trademark abandonment, and copyright fair use all play this role. Cf. Thomas W. Merrill, Property Rules, Liability Rules, and Adverse Possession, 79 *Nw. U. L. Rev.* 1122, 1130 (1985) ("A . . . reason commonly advanced in support of a system of adverse possession is that it punishes [owners] who 'sleep on their rights.' Under this view, the shift in entitlement acts as a penalty to deter [owners] from ignoring their property or otherwise engaging in poor custodial practices. Since forfeiture is a stiff penalty (frequently deemed unconscionable in other contexts) presumably the objective will be realized in most cases."). See generally Jacob Noti-Victor, Copyright's Law of Dissemination, 44 *Cardozo L. Rev.* 1769 (2023) (outlining the role of various copyright defenses in regulating rightsholder behavior).

<sup>335</sup> See *supra* Subsection II.A.3.

<sup>336</sup> 15 U.S.C. § 1125 (causes of action for false endorsement and dilution); see also Ginsburg, *supra* note 164, at 386–87 (examining trademark protection for author's names).

<sup>337</sup> See, e.g., Cal. Civ. Code § 3344 (West 1984) ("Any person who knowingly uses another's name, voice, signature, photograph, or likeness, in any manner, on or in products, merchandise, or goods, or for purposes of advertising or selling, or soliciting purchases of,

explicitly marketed using the name “Drake,” Drake would clearly have causes of action against whoever produced and distributed the song.<sup>338</sup> The Visual Artist Rights Act also grants artists an additional cause of action to “prevent the use of his or her name as the author of any work of visual art which he or she did not create.”<sup>339</sup>

A somewhat more complicated scenario would involve the AI-mimicry of an artist’s voice or likeness, without the use of their name. Deliberate mimicry of a singer or actor’s voice can implicate a right of publicity protection.<sup>340</sup> And Jennifer Rothman has recently suggested that, in the case of most artists or celebrities, many states’ right of publicity statutes would indeed protect against unauthorized mimicry of a voice or likeness by an AI system.<sup>341</sup> Trademark law may also provide protection for the cooption of the voice or likeness of an artist<sup>342</sup>—such as a singer or actor—but this protection would require that consumers be plausibly confused by the appropriation, which courts have held hinges on recognizability.<sup>343</sup> Famous celebrities’ likenesses would likely be protectable,<sup>344</sup> but many other artists’ may not be.

All of these laws require that the mimicked artist herself pursue a claim; the public, even if they are indeed deceived by the AI mimicry, has no independent cause of action. Thus, if Drake were to license his voice or name for the use of AI-generated songs,<sup>345</sup> the public would have no way of policing such use under the right of publicity or trademark law. This,

---

products, merchandise, goods or services, without such person’s prior consent, or, in the case of a minor, the prior consent of his parent or legal guardian, shall be liable for any damages sustained by the person or persons injured as a result thereof.”); see also *id.* § 3344.1 (providing similar protection posthumously).

<sup>338</sup> See, e.g., *King v. Innovation Books*, 976 F.2d 824, 828–29 (2d Cir. 1992) (holding that the use of a writer’s name when the writer “had no involvement in, and gave no approval” of the work violates the Lanham Act).

<sup>339</sup> 17 U.S.C. § 106A(a)(1)(B).

<sup>340</sup> See *Midler v. Ford Motor Co.*, 849 F.2d 460, 462–63 (9th Cir. 1988).

<sup>341</sup> See Rothman, *supra* note 160, at 3; see also *supra* note 161 (discussing expansion of this protection in the proposed NO FAKES Act).

<sup>342</sup> See, e.g., *Waits v. Frito-Lay, Inc.*, 978 F.2d 1093, 1107 (9th Cir. 1992); *Wendt v. Host Int’l, Inc.*, 125 F.3d 806, 812 (9th Cir. 1997).

<sup>343</sup> See, e.g., *Electra v. 59 Murray Enters.*, 987 F.3d 233, 258 (2d Cir. 2021) (affirming dismissal of Lanham Act claim related to appropriation of various models’ likenesses in a Strip Club’s advertising on the grounds that most of the models’ likenesses were not recognizable to the public); *White v. Samsung Elecs. Am., Inc.*, 971 F.2d 1395, 1400 (9th Cir. 1992).

<sup>344</sup> See Rothman, *supra* note 160, at 3.

<sup>345</sup> See *supra* notes 169–70 (examining some examples of licensed uses of an artist’s voice).

coupled with the fact that general AI-generated works—i.e., works that do not involve the specific cooption of a celebrity name, voice, or likeness—do not implicate trademark or the right of publicity, means these areas of law will invariably play a limited role as information-forcing tools.<sup>346</sup>

#### CONCLUSION

Generative AI technology promises a world of seamless content creation at the press of a button. While some of this content may prove valuable and enjoyable, lawmakers and scholars have begun to note the many risks the technology poses, and in particular its potential to aid deception and fraud. As this Article has argued, any discussion of generative AI's potential harms must also consider the consumer's interest in knowing whether works of art and entertainment are AI-generated. Though sometimes neglected in debates about entertainment markets, many consumers have clear process-focused preferences for works that are human-made, born out of both ethical and aesthetic concerns. And these preferences are defensible. The hidden use of AI destabilizes art's social role—the dialogue between author and reader that furthers self-definition, empathy, and ethical and political engagement. By removing this social, and fundamentally human, element, all creative works become valuable only for their potential to entertain. Members of the public deserve to know the role of AI in a work's creation so they can choose whether, and on what terms, to engage with it.

These potential harms are not just speculative. The lack of copyright protection for most AI-generated content is leading to a world in which,

---

<sup>346</sup> Indeed, Dan Burk has noted that trademark law, as currently constituted, has limited utility in addressing authorized uses of generative AI, and argues that trademark law needs “restructuring” to enable a kind of certification system for human-made works—a “trademark certification of authentic [non-AI] origin.” Burk, *supra* note 151, at 1708–10 (“Control of the trademark prevents outside imitators from staging an inauthentic Queen performance, but there is nothing to prevent the trademark holders *themselves* from staging an inauthentic performance.”); cf. Paul Ohm, *Branding Privacy*, 97 *Minn. L. Rev.* 907, 944 (2013) (arguing that trademark protection should be tethered to the privacy guarantees made by a company at the time of the mark's inception). This approach, however, would not only require significant legislative change but would also present a dramatic reshaping of trademark law away from its foundations as a means of policing unfair trade practices between competitors, toward a predominantly consumer-protection-focused regime. See Mark P. McKenna, *The Normative Foundations of Trademark Law*, 82 *Notre Dame L. Rev.* 1839, 1858–60, 1866 (2007) (discussing the origins of trademark as unfair competition law and discussing tension between this traditional role and consumer-protection justifications).

for perhaps the first time, many valuable new works will be immediately in the public domain. This creates incentives for producers of such works to hide their use of AI. Private ordering, legislation, and FTC enforcement all may have a role to play in making sure the public has access to accurate information about a work's origins. IP may also provide some partial solutions. By penalizing producers who fail to disclose their use of the technology, the doctrine of copyright misuse can help force out information regarding a work's provenance. Trademark and the right of publicity litigation can also help facilitate disclosure.

The use of IP, and copyright in particular, to safeguard consumers' ethical and aesthetic preferences may strike some as counterintuitive. Copyright is primarily concerned with authors' incentives to create; consumers' interests, if any, are often framed around the problems copyright can create when it comes to access, pricing, or follow-on creativity. But, as this Article has suggested and future work will explore more fully,<sup>347</sup> copyright policy is intertwined with both how and why consumers choose to engage with creative works on aesthetic, ethical, and other terms. While recognizing this link may raise some tensions—in particular, between the market-oriented approach of copyright law and the anti-commodification impulses that drive much aesthetic appreciation—copyright clearly has a role to play in moderating the ways that consumers can satisfy their preferences for creative works in the age of generative AI.

---

<sup>347</sup> See Jacob Noti-Victor, *Copyright and the Consumer* (work in progress) (on file with author).