COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

VIRGINIA LAW REVIEW ONLINE

VOLUME 112 JULY 2026 135–161

ESSAY

THE ONLYFANS ECONOMY: INTELLECTUAL PROPERTY'S PIVOT FROM SCARCITY TO AUTHENTICITY

*Vincent Joralemon**

Generative AI is destabilizing the foundational assumption of intellectual property law: that creation is difficult, is expensive, and requires legal inducement. When machines produce text, images, and code at near-zero marginal cost, the utilitarian justification for copyright and patent protection begins to collapse. This Essay argues that what emerges in its place is a regime organized not around the scarcity of creation, but around the scarcity of verification—a shift from an incentive paradigm to a source identification paradigm dominated by trademark, rights of publicity, and platform-controlled authentication infrastructure.

OnlyFans provides a revealing case study. Despite AI-generated pornography flooding the market, the platform paid over $5.8 billion to human creators in 2024. Consumers pay not for content, which AI can approximate, but for provenance—the verified knowledge that they are interacting with a specific, authenticated person. This Essay identifies a "triple-lock" structure underlying this economy: verified identity draws consumers in, proprietary infrastructure ensures that access requires staying inside the walled garden, and aggressive

* Director, Berkeley Life Sciences Law & Policy Center; Senior Fellow, Berkeley Center for Law & Technology.

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

copyright enforcement destroys unauthorized copies that would otherwise undermine authenticity's value.

Applying a Law and Political Economy lens, the Essay argues that this emerging regime systematically advantages those who already possess recognized brands, legal departments, and capital for proprietary infrastructure—while offering little to individual creators who lack pre-existing fame. It concludes by proposing that the Library of Congress serve as a public "digital notary," preventing the infrastructure of authenticity from becoming a private toll road.

INTRODUCTION 136
I. THE COLLAPSE OF THE INCENTIVE PARADIGM 139
A. The Human Authorship Barrier 139
B. The Wholesale Shift: Data Licensing as Institutional Capture 141
C. Copyright as Border Patrol 143
II. THE EMERGING REGIME OF PROVENANCE AND PLATFORM CONTROL 144
A. The Walled Garden 144
B. The Front End: Trademark and the Market for Authenticity 145
C. The Trademark Paradox: Hallucinations vs. Erasure 147
D. The Human as Brand: Rights of Publicity 148
III. THE POLITICAL ECONOMY OF AUTHENTICITY 150
A. Whose Authenticity? 151
B. The Infrastructure of Trust 153
IV. GOVERNANCE AS THE REMAINING VARIABLE 156
A. The Structural Case for Public Infrastructure 156
B. The Library of Congress as Digital Notary 157
CONCLUSION 159

INTRODUCTION

AI-generated pornography is flooding the internet; *The Economist* reports that the market for synthetic adult content will reach $2.5 billion in 2025 and is projected to grow at twenty-seven percent annually.[1] And yet, OnlyFans, the online adult content platform designed around *human*

[1] AI Is Upending the Porn Industry, The Economist, Nov. 29, 2025, at 55, 55.

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

content creators, generated $7.22 billion in revenue in 2024.[2] The platform has paid over $25 billion to creators since 2016.[3]

These facts sit uneasily together. If generative AI can produce unlimited synthetic explicit content on demand, customized to any specification, and cheaply available—why would anyone pay for content from a specific human being? The answer illuminates something important about the future of intellectual property. What consumers pay for on OnlyFans is not the content itself, which AI can now approximate with increasing fidelity. They pay for the *provenance*—the verified knowledge that they are interacting with a specific, authenticated person. The content is abundant; the source is scarce.

But provenance alone is not enough. OnlyFans does not merely verify identity; it protects that identity within a fortress of interlocking legal and technological defenses. Verified identity draws consumers in. Proprietary infrastructure—DRM encryption, anti-scraping measures, behavioral data collection—functions as enclosure. And aggressive copyright enforcement destroys unauthorized copies that would otherwise flood the market with "authentic" content. This triple-lock structure—identity as the hook, infrastructure as the lock, enforcement as border control—is the architecture of the post-incentive economy.

The intellectual property system has long operated on a different premise: that creation is difficult, expensive, and scarce, and therefore creators require the inducement of a temporary monopoly to produce. This utilitarian bargain underwrites the patent and copyright clause of the U.S. Constitution.[4] Yet generative artificial intelligence has begun to unravel this foundational assumption. When the marginal cost of producing competent text, images, and code approaches zero, the economic logic of incentivizing creation collapses. Something must fill the void.

What emerges is a regime organized not around the scarcity of creation, but around the scarcity of verification. We are witnessing a transition from what might be called the *incentive paradigm* of

[2] Id. at 56; Todd Spangler, OnlyFans Gross Revenue Rises 9% to $7.2 Billion in 2024, Variety (Aug. 22, 2025, at 05:01 PT), https://variety.com/2025/digital/news/onlyfans-fiscal-2024-revenue-earnings-1236495750/.

[3] Rose Henderson, OnlyFans Has Paid Creators $25 Billion Since 2016, CEO Says, Bloomberg (Oct. 21, 2025, at 11:06 ET), https://www.bloomberg.com/news/articles/2025-10-21/onlyfans-has-paid-creators-25-billion-since-2016-ceo-says.

[4] U.S. Const. art. I, § 8, cl. 8.

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

intellectual property to a *source identification paradigm* dominated by trademark, rights of publicity, and platform-controlled verification infrastructure. In this new order, the legal system's primary function shifts from encouraging the production of goods to certifying their origin.

This transformation is often presented as a neutral response to technological change.[5] This Essay argues that such framing obscures the distributive stakes. The emerging regime favors those who already possess recognized brands, access to legal departments, and the capital to build proprietary infrastructure. It promises to entrench the dominance of large platforms and legacy media organizations while offering little to individual creators who lack pre-existing fame.[6]

While legal scholarship has extensively diagnosed the erosion of the incentive paradigm, it has yet to fully articulate replacing it. Mark Lemley has persuasively argued that the digital age renders IP's artificial scarcity inefficient,[7] and Pamela Samuelson has warned against rushing to create *sui generis* rights for AI outputs.[8] Yet, the critical question remains unanswered: if the economic logic of incentivizing creation is collapsing, how is value being enclosed in its absence? This Essay provides the first unified account of that replacement regime, offering both a diagnostic tool for the post-incentive economy and a prescriptive path to prevent the authentication regime from becoming a form of private enclosure.

This Essay proceeds in four Parts. Part I documents the collapse of the incentive paradigm. Part II analyzes the emerging regime of provenance and private control. Part III turns to political economy,

[5] See, e.g., C2PA Founding Press Release, Coal. for Content Provenance & Authenticity (Feb. 22, 2021), https://c2pa.org/c2pa-founding-press-release/ [https://perma.cc/47CF-75QN] ("With the collective expertise of this group, we will accelerate the critical work of rebuilding the public's trust in online content through broad and open adoption of a provenance standard at scale."); Exec. Order No. 14110, 3 C.F.R. 657 (2024) (treating content provenance and watermarking standards as self-evidently beneficial transparency measures, without addressing governance of the underlying infrastructure); Council Regulation 2024/1689, ch. IV, art. 50, ¶ 2, 2024 O.J. (L) 82 (EU) (mandating machine-readable labeling of AI-generated outputs that are "effective, interoperable, robust and reliable" without note of the infrastructure ownership).

[6] For the Law and Political Economy framework applied here, see Jedediah Britton-Purdy, David Singh Grewal, Amy Kapczynski & K. Sabeel Rahman, Building a Law-and-Political-Economy Framework: Beyond the Twentieth-Century Synthesis, 129 Yale L.J. 1784 (2020).

[7] Mark A. Lemley, IP in a World Without Scarcity, 90 N.Y.U. L. Rev. 460, 470–71 (2015).

[8] See Pamela Samuelson, Generative AI Meets Copyright, 381 Sci. 158, 160–61 (2023).

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

arguing that the triple-lock structure systematically advantages incumbents over entrants. Part IV proposes that the Library of Congress could serve as a public root of trust—a "digital notary" that prevents the infrastructure of authenticity from becoming a private toll road.

I. THE COLLAPSE OF THE INCENTIVE PARADIGM

The traditional utilitarian justification[9] for intellectual property tied to creation—copyright and patent—rests on assumptions about human creative labor that generative AI directly contradicts. Courts and administrative agencies have responded by erecting doctrinal walls that exclude AI-generated outputs from protection,[10] but these exclusions do not restore the old equilibrium. They instead accelerate a flight toward alternative legal regimes.

A. The Human Authorship Barrier

The U.S. Copyright Office and federal courts have been unequivocal: copyright protection requires human authorship. In *Thaler v. Perlmutter*, the D.C. Circuit affirmed the denial of registration for a work created by an AI system designated as the sole author, grounding its holding in statutory text demonstrating that Congress intended "author" to mean a biological human.[11] The Copyright Office's subsequent treatment of *Zarya of the Dawn*[12]—a graphic novel whose human-authored text was paired with AI-generated images—reinforced this boundary: the human author's text and selection were protectable, but the AI-generated images were not.[13]

[9] The utilitarian account is one of several justifications for intellectual property and is most pertinent to the economic justifications herein. For the canonical statement of the utilitarian account, see William M. Landes & Richard A. Posner, The Economic Structure of Intellectual Property Law 186–231 (2003). For an overview of others, see William Fisher, Theories of Intellectual Property, *in* New Essays in the Legal and Political Theory of Property 168, 169–72 (Stephen R. Munzer ed., 2001).

[10] See infra Section I.A (discussing *Thaler v. Perlmutter* and the *Zarya of the Dawn* registration decision); Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence, 88 Fed. Reg. 16190 (Mar. 16, 2023) (to be codified at 37 C.F.R. pt. 202).

[11] Thaler v. Perlmutter, 130 F.4th 1039, 1043–45 (D.C. Cir. 2025).

[12] Kris Kashtanova, Zarya of the Dawn, https://aicomicbooks.com/wp-content/uploads/2022/09/kashtanova-full-book-v1-small.pdf [https://perma.cc/QER9-NPZ2] (last visited Apr. 16, 2026).

[13] U.S. Copyright Off., Letter from Robert J. Kasunic, Assoc. Reg. of Copyrights & Dir. of the Off. of Registration Pol. & Prac., to Van Lindberg, Taylor English Duma LLP 1–8 (Feb.

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

The immediate implication is that vast categories of commercially valuable AI output are born into the public domain. But the deeper significance lies in what this exclusion reveals about copyright's underlying theory. If the regime exists to incentivize creation, and machines now create without incentive, then copyright's justification evaporates for an expanding share of cultural production.

The incentive story simply cannot account for a world in which the creative act requires no inducement.[14] As Pamela Samuelson has argued, this "gap" is not a crisis but a recurrence of a historical pattern: from player pianos to VCRs, disruptive technologies have periodically thrown the incentive paradigm into disarray.[15] But, as she warns, rushing to fill this void with *sui generis* protections would be a mistake; the public domain status of these works is a feature of the system, not a bug.[16]

Mark Lemley, in his prescient article *IP in a World Without Scarcity*, anticipated this unraveling.[17] Lemley argued that when production and distribution costs approach zero, IP's artificial scarcity becomes economically inefficient and socially obstructive.[18] We do not need to incentivize a machine; it generates text and images for the cost of electricity.[19] In this context, laws creating artificial scarcity become barriers to access rather than engines of production. Patent law faces analogous strain, but the question across both regimes is the same: What replaces the incentive logic when creation no longer requires inducement?

21, 2023), https://www.copyright.gov/docs/zarya-of-the-dawn.pdf [https://perma.cc/TRW2-L4SN].

[14] This collapse is less surprising when one recognizes that the incentive account was always contestable. Robert Merges has argued that intellectual property can be justified on Lockean labor, Kantian autonomy, and Rawlsian distributive grounds, or what he terms "midlevel principles" including proportionality, non-removal, and dignity, none of which depends on incentive effects. See Robert P. Merges, Justifying Intellectual Property 31–101, 128–33, 139–58 (2011).

[15] Samuelson, supra note 8, at 158.

[16] Id. at 158–59, 161.

[17] Lemley, supra note 7, at 461.

[18] Id. at 462–65.

[19] Lemley's subsequent work on "faith-based intellectual property" documented how the incentive framework has persisted despite lacking empirical support. Mark A. Lemley, Faith-Based Intellectual Property, 62 UCLA L. Rev. 1328, 1338 (2015). This is a triumph of ideology over evidence that AI may finally expose.

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

B. The Wholesale Shift: Data Licensing as Institutional Capture

Defenders of the traditional paradigm might point to the burgeoning market for training data licenses as evidence that the incentive system is adapting, not collapsing. Major AI developers have struck high-profile deals with legacy media organizations: OpenAI has announced partnerships with *The Atlantic*, *Vox Media*, and other major publishers.[20]

But this market reveals something darker than adaptation. It reveals copyright's transformation from a tool that rewards creators into a toll that enriches aggregators. News Corp's agreement with OpenAI (worth up to $250 million over five years) licenses decades of journalism from *The Wall Street Journal* and *New York Post*.[21] Unlike traditional syndication or database licensing, which distributes journalism to readers who encounter it as such, AI training licenses surrender entire archives as raw material for systems designed to generate competing content—and do so under existing publishing contracts that afford individual journalists neither consent nor compensation.[22] The journalists who wrote that journalism do not benefit as they would from traditional syndication; instead, AI companies are training models designed to replace, in whole or in part, the work of the journalist. As *Nieman Lab* reported in September 2025: "In the U.S., when a publisher signs a licensing deal with an AI company, newsroom staffers don't get a cut."[23] The compensation flows to the entity that owns the aggregation, not the humans who created the content being aggregated. Copyright, designed to incentivize the author, now incentivizes the archive.

---

[20] OpenAI Signs Content Deals with The Atlantic and Vox Media, Reuters (May 29, 2024, at 13:25 ET), https://www.reuters.com/business/media-telecom/openai-signs-content-deals-with-atlantic-vox-media-2024-05-29/.

[21] Todd Spangler, News Corp Inks OpenAI Licensing Deal Potentially Worth More Than $250 Million, Variety (May 22, 2024, at 14:14 PT), https://variety.com/2024/digital/news/news-corp-openai-licensing-deal-1236013734/.

[22] Frank Pasquale & Haochen Sun, Consent and Compensation: Resolving Generative AI's Copyright Crisis, 110 Va. L. Rev. Online 207, 207–10 (2024); see also Matthew Sag, The False Hope of Content Licensing at Internet Scale, ProMarket (Nov. 19, 2025), https://www.promarket.org/2025/11/19/the-false-hope-of-content-licensing-at-internet-scale/ [https://perma.cc/375T-Q3BX]; Mark A. Lemley & Jacob Noti-Victor, Anticompetitive Acquiescence in AI Content Licensing, ProMarket (Nov. 18, 2025), https://www.promarket.org/2025/11/18/anticompetitive-acquiescence-in-ai-content-licensing/ [https://perma.cc/WV9T-GWDA].

[23] Andrew Deck, Some French Publishers Are Giving AI Revenue Directly to Journalists. Could That Ever Happen in the U.S.?, Nieman Lab (Sep. 4, 2025, at 16:01 ET), https://www.niemanlab.org/2025/09/in-france-ai-revenue-is-going-directly-to-journalists-could-that-happen-in-the-u-s/ [https://perma.cc/QF62-G9CN].

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

This is the market working as designed. Transaction costs make direct licensing impossible for individual authors, inevitably favoring large aggregators. The Copyright Office's 2025 report acknowledges this with remarkable candor: "Direct licensing is most common and most promising with respect to corporate entities with catalogs of high-quality and easily identifiable content."[24]

The stock photography market offers stark quantification. Shutterstock earned $104 million in AI licensing revenue in 2023 from deals with OpenAI, Meta, and others.[25] One analysis of payments to 58 photographers found the median compensation was $0.007 per image: less than a penny for work that will train systems capable of replacing them.[26] Contributors had no ability to opt out; AI licensing was covered under existing terms of service.[27]

The photographer's copyright became Shutterstock's revenue stream. The legal regime ostensibly designed to protect the "romantic genius"[28] now functions to ensure that romantic genius sees nothing while the firm that employed her collects the rent. Copyright no longer incentivizes creation; it incentivizes accumulation. This "wholesale shift" does not save the incentive paradigm; it exposes its capture by institutional players the system was never meant to privilege.

---

[24] U.S. Copyright Off., Copyright and Artificial Intelligence, Part 3: Generative AI Training 70 (May 2025), https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-3-Generative-AI-Training-Report-Pre-Publication-Version.pdf [https://perma.cc/3UA7-44A5].

[25] Brody Ford, Shutterstock's AI-Licensing Business Generated $104 Million Last Year, Bloomberg (June 4, 2024, at 08:00 ET), https://www.bloomberg.com/news/articles/2024-06-04/shutterstock-s-ai-licensing-business-generated-104-million-last-year.

[26] Matt Growcoot, Shutterstock May Have Paid Out Over $4 Million From Its AI Contributor Fund, PetaPixel (July 12, 2023), https://petapixel.com/2023/07/12/shutterstock-may-have-paid-out-over-4-million-from-its-ai-contributor-fund/ [https://perma.cc/6V6F-YZFL].

[27] Shutterstock, Shutterstock Data Licensing and the Contributor Fund, https://submit.shutterstock.com/help/en/articles/10594694-shutterstock-data-licensing-and-the-contributor-fund [https://perma.cc/9C5D-5YGJ] (last updated Oct. 24, 2024) (explaining that AI data licensing is "covered under Section 1a of our Contributor Terms of Service, which grants Shutterstock the right to develop new features and products," and noting that an opt-out function was added only in January 2023).

[28] See, e.g., Martha Woodmansee, The Genius and the Copyright: Economic and Legal Conditions of the Emergence of the 'Author,' 17 Eighteenth-Century Stud. 425, 426–28 (1984) (tracing the eighteenth-century construction of the author as "original genius" whose work merits property protection); Peter Jaszi, Toward a Theory of Copyright: The Metamorphoses of "Authorship," 1991 Duke L.J. 455, 462–71 (arguing that the "Romantic 'authorship' construct" continues to shape copyright doctrine).

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

C. Copyright as Border Patrol

Yet even as copyright's *incentive* function collapses, its *enforcement* function has become more valuable than ever. For the OnlyFans creator, copyright is no longer about the utilitarian bargain with the public; it is a border patrol tool. Experts estimate fifty to seventy percent of paid OnlyFans content is stolen or spread without permission.[29] The challenge is not creating content worth protecting (the creator handles that) but maintaining the infrastructure to detect infringement and execute takedowns at scale.

This enforcement capacity has become a service that only platforms and specialized intermediaries can effectively provide. OnlyFans maintains a dedicated copyright team and partners with DMCAForce as an official Digital Millennium Copyright Act agent, providing digital fingerprinting and automated monitoring.[30] But creators who want meaningful protection must pay for third-party services charging between $150 and $500 per month for continuous monitoring and automated enforcement.[31] One creator described paying "several hundred dollars a month" for protection; another characterized the work as "an endless game of Whac-a-Mole."[32]

Copyright enforcement has thus become "Copyright-as-a-Service": a professionalized infrastructure that individual creators cannot replicate and must purchase from intermediaries. The creator who once needed a publisher to distribute her work now needs a platform to protect it. The dependency has changed form but not structure. And the distributional

[29] Mayya Chernobylskaya, Kristina Böhmer, Polina Bachlakova & Apolena Rychlíková, Consent Violated: How Creators Pay the Price for Stolen OnlyFans Content, Unbias The News (Sep. 4, 2025), https://unbiasthenews.org/consent-violated-how-creators-pay-the-price-for-stolen-onlyfans-content/ [https://perma.cc/JS2Y-N9YB].

[30] DMCAForce, OnlyFans DMCA Service, https://dmcaforce.com/onlyfans-dmca-content-protection/ [https://perma.cc/G6E6-GQ82] (last visited Apr. 16, 2026).

[31] See DMCAForce, DMCAForce Pricing, https://dmcaforce.com/dmcaforce-pricing/ [https://perma.cc/T86Y-5HME] (last visited Apr. 16, 2026) (listing plans starting at $150 per month, with some clients at $45 per month at time of publication). CopyrightShark and BranditScan offer plans starting at $59 and $69 per month respectively. See Maria Harutyunyan, Top DMCA Protection Tools for OnlyFans Content in 2026, She's SINGLE Mag. (Feb. 20, 2026), https://www.shessinglemag.com/blogs/resources/top-dmca-protection-tools-for-onlyfans-content-in-2026 [https://perma.cc/M2W8-WT65].

[32] Courtney Kocak, My OnlyFans Was a Fun Way for Me to Make Money. Then My Content Got Stolen., Slate (Nov. 15, 2024, at 10:27 ET), https://slate.com/technology/2024/11/onlyfans-content-stolen-privacy-security-internet-dark-side.html [https://perma.cc/6AQV-UNVA].

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

consequences are predictable: creators with resources can afford the subscription; creators without cannot. The right exists; the capacity to vindicate it does not.

II. THE EMERGING REGIME OF PROVENANCE AND PLATFORM CONTROL

As the incentive paradigm falters, legal and commercial actors are migrating toward frameworks that prioritize verification over creation. What emerges is a two-layered regime: on the front end, trademark, unfair competition, and rights of publicity establish *who* produces content; on the back end, proprietary platform infrastructure ensures that verification cannot be replicated outside the walled garden. The logic is simple: if you cannot protect what you create, you protect who you are and ensure no one can access that identity without going through you.

A. The Walled Garden

Provenance has value only if it cannot be replicated outside the platform. OnlyFans protects verified identity with multiple layers of technological protection: encryption shields uploaded content; per-subscriber watermarking embeds identifying metadata into any screen-captured video, tracing every leak back to its source; and digital fingerprinting tags each piece of content with a unique identifier that enables automated detection of unauthorized redistribution across the wider web.[33] The platform is a fortress designed to ensure that "authentic" access requires staying inside.

The walled garden protects more than content. OnlyFans accumulates extensive operational data—creator profiles, message archives, subscriber lists, vault media, transaction data, sales and renewal events, and revenue statistics—inside platform-mediated infrastructure.[34] Even if a creator exports their records and walks away, they leave behind the messaging rails through which they communicated with subscribers, the content protection stack that hosted their library, and the discovery

[33] How OnlyFans Provides Content Protection for Content Creators, DMCAForce, https://dmcaforce.com/how-onlyfans-provides-content-protection/ [https://perma.cc/PWJ3-TYRF] (last visited Apr. 16, 2026).

[34] OnlyFans, Privacy Policy, https://onlyfans.com/privacy [https://perma.cc/TND8-2KH8] (last visited Apr. 16, 2026); Jonas S., OnlyFans API: A Game-Changer for Creators & Agencies in the Digital Realm, List25 (Mar. 23, 2026), https://list25.com/onlyfans-api-creator-economy-game-changer/ [https://perma.cc/WBJ8-DBBR].

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

infrastructure that connected them to their audience.[35] This infrastructural asymmetry is the platform's true moat. The creator provides the identity; the platform controls the intelligence layer through which that identity becomes revenue.

In *Waidhofer v. Cloudflare, Inc.*, OnlyFans and Patreon creators sued Cloudflare, which allegedly provided services to Thothub, a piracy site that aggregated leaked paywalled content, asserting copyright infringement, RICO violations, and unfair competition claims.[36] Thothub shut down within days—before any judgment, before any injunction.[37] The walled garden's defenses are not merely technical but legal; overlapping claims create zones of protection that no single doctrine could provide.

B. The Front End: Trademark and the Market for Authenticity

The Lanham Act's—the federal trademark statute's[38]—legislative history is explicit: trademark law exists "to protect the public so it may be confident that, in purchasing a product bearing a particular trademark which it favorably knows, it will get the product which it asks for and wants to get."[39] Unlike copyright and patent, trademark has never been primarily about incentivizing creation; the source-identification paradigm is thus a logical expansion of trademark's domain to fill the void left by AI's disruptive effects on legacy intellectual property models.[40]

---

[35] See sources cited supra note 34.

[36] No. 20-cv-06979, 2021 WL 8532943 (C.D. Cal. Sep. 29, 2021). The original complaint, filed August 3, 2020, named Thothub.tv and Cloudflare as defendants. See Complaint at 2, *Waidhofer*, No. 20-cv-06979, 2021 WL 8532943. The First Amended Complaint alleged that Thothub "apparently ceased operations days after the original complaint in this case was filed." And the court's later order repeated that allegation. *Waidhofer*, 2021 WL 8532943, at *2 n.5; First Amended Complaint at 3 n.1, *Waidhofer*, No. 20-cv-06979, 2021 WL 8532943; see Samantha Cole, 'Leaked' Porn Site Thothub Goes Dark Following Lawsuit, Vice (Aug. 6, 2020, at 11:38 ET), https://www.vice.com/en/article/leaked-porn-site-thothub-down-lawsuit/ [https://perma.cc/8U4Z-SZQW] (reporting that the site "went down on Wednesday evening" and, at time of publication, was replaced by a placeholder page; quoting plaintiff's counsel that Thothub had "no choice but to shut down its criminal activity").

[37] See Cole, supra note 36.

[38] 15 U.S.C. §§ 1051–1141n.

[39] H.R. Rep. No. 79-219, at 2 (1945).

[40] See, e.g., Robert G. Bone, Trademark Functionality Reexamined, 7 J. Legal Analysis 183, 187 (2015).

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

Trademark and unfair competition law are rising to manage the front end of the authenticity economy. Generative AI creates what economist George Akerlof calls a classic "Market for Lemons" dynamic: when buyers cannot distinguish authentic content from synthetic imitations, trust collapses across entire categories.[41] As AI floods the internet with synthetic pornography, human performers face this problem in its starkest form. One might expect abundance to destroy the market for human-created content. Yet OnlyFans paid $5.8 *billion* to creators in fiscal year 2024.[42]

What explains this resilience? Researchers have found that viewers are aroused by AI-generated images, but less so than by those they believe show real people.[43] The value proposition rests precisely on verified authenticity.

In a lemons market, the solution is signaling. A trademark functions as a credibility proxy: "*The New York Times*" attached to an article signals verification by an organization with reputational stakes.[44] As the cost of content generation falls, the value of the brand rises. Trademark creates artificial scarcity not of the good itself, but of the identity associated with it. In an age of AI abundance, this identity scarcity may be the only scarcity that matters.

[41] George A. Akerlof, The Market for "Lemons": Quality Uncertainty and the Market Mechanism, 84 Q.J. Econ. 448, 488–89 (1970).

[42] Spangler, supra note 2.

[43] Marco Marini et al., Real Is the New Sexy: The Influence of Perceived Realness on Self-Reported Arousal to Sexual Visual Stimuli, 38 Cogn. & Emotion 348, 353, 356–57 (2024).

[44] The Supreme Court has long located trademark law in the source-identification function. In *Dastar Corp. v. Twentieth Century Fox Film Corp.*, the Court explained that federal trademark law "has no necessary relation to invention or discovery," but rather, by preventing competitors from copying a "source-identifying mark," "reduce[s] the customer's costs of shopping and making purchasing decisions" and helps ensure that the producer, not an "imitating competitor," reaps the "reputation-related rewards associated with a desirable product." 539 U.S. 23, 34 (2003) (first quoting Trade-Mark Cases, 100 U.S. 82, 94 (1879); and then quoting Qualitex Co. v. Jacobson Prods. Co., 514 U.S. 159, 163–64 (1995)). More recently, *Jack Daniel's Properties, Inc. v. VIP Products LLC* describes a trademark's "primary" function as identifying the origin or ownership of the article to which it is affixed, and explained that a mark "tells the public who is responsible for a product." 143 S. Ct. 1578, 1583 (2023) (quoting Hanover Star Milling Co. v. Metcalf, 240 U.S. 403, 412 (1916)). McCarthy's leading treatise captures the doctrine's dual orientation: trademark law protects "consumers from deception and confusion over trademarks *as well as* to protect the plaintiff's infringed trademark as property." J. Thomas McCarthy, McCarthy on Trademarks and Unfair Competition, § 2:2 (5th ed. 2024) (emphasis added).

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

C. The Trademark Paradox: Hallucinations vs. Erasure

Plaintiffs are responding by pivoting from copyright to Lanham Act theories, but trademark law is ill-equipped to handle the full scope of the provenance crisis. The divergent outcomes in recent litigation reveal a fault line: trademark can protect against false attribution, but it cannot compel correct attribution.

The *New York Times* lawsuit against Perplexity AI exemplifies the viable frontier.[45] The *Times* alleges that Perplexity generates "hallucinations" and falsely attributes them to the *Times*, engaging in "false designation of origin" under 15 U.S.C. § 1125.[46] This theory gained traction in *Advance Local Media LLC v. Cohere Inc.*, where the Southern District of New York allowed similar Lanham Act claims to proceed on the grounds that misattributed hallucinations can create a likelihood of consumer confusion and reputational harm.[47]

However, the attempt to remedy erasure of identity has hit a structural wall. In *Doe 1 v. GitHub, Inc.*, plaintiffs argued that when Copilot reproduces open-source code without license credits, it strips the code of its provenance.[48] Defendants invoked *Dastar Corp. v. Twentieth Century Fox Film Corp.*, where the Supreme Court held that "origin of goods" in the Lanham Act refers to the producer of the tangible product, not the author of the ideas embodied within it.[49] The district court dismissed the Lanham Act claims.[50]

The law effectively protects against "defamation-by-AI" (false attribution) but lacks a mechanism to prevent "erasure-by-AI" (attribution to no one at all).[51] The individual creator whose work is

[45] Dow Jones & Co. v. Perplexity AI, Inc., 797 F. Supp. 3d 305, 315–16 (S.D.N.Y. 2025).

[46] Id.

[47] No. 25-cv-01305, 2025 WL 3171892 (S.D.N.Y. Nov. 13, 2025).

[48] 672 F. Supp. 3d 837, 846–47 (N.D. Cal. 2023).

[49] 539 U.S. 23, 37 (2003). Justice Scalia's opinion for the Court warned that extending § 43(a) to creative authorship claims "would create a species of mutant copyright law that limits the public's 'federal right to "copy and to use" ' expired copyrights." Id. at 34 (quoting Bonito Boats, Inc. v. Thunder Craft Boats, Inc., 489 U.S. 141, 165 (1989)).

[50] *GitHub*, 672 F. Supp. 3d at 847, 861–62.

[51] This is the post-*Dastar* attribution gap problem: federal law recognizes some attribution-related claims, particularly false-attribution, false-advertising, and copyright-management-information claims, but does not provide a general freestanding Lanham Act remedy for copying or reuse that simply omits attribution. Justin Hughes, American Moral Rights and Fixing the *Dastar* "Gap", 2007 Utah L. Rev. 659, 669–704; *Dastar Corp.*, 539 U.S. at 37–38. The Copyright Office has acknowledged this gap with remarkable candor, recommending that Congress "consider adopting an amendment to section 43(a) of the

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

ingested into a training corpus without attribution finds that *Dastar* blocks any claim for identity erasure.

D. The Human as Brand: Rights of Publicity

Beyond trademark, the right of publicity is expanding to protect the human persona itself. Deepfake technology can now produce a convincing explicit image from a single clear photograph in minutes; so-called "nudify" applications require no technical expertise, only an internet connection and a social media photo.[52] Explicit AI-generated images of Taylor Swift went viral in 2024; increasingly, the victims are ordinary people with no legal team.[53]

The right of publicity's application to AI-generated content builds on a doctrinal foundation established decades before generative AI existed. In *Midler v. Ford Motor Co.*, the Ninth Circuit held that Bette Midler could recover when Ford hired a soundalike to imitate her distinctive voice for a commercial, reasoning that "[a] voice is as distinctive and personal as a face" and that "[t]o impersonate her voice is to pirate her identity."[54] *Waits v. Frito-Lay, Inc.* extended this protection and added Lanham Act false endorsement liability, establishing that deliberately imitating a distinctive, widely known voice for commercial purposes communicates falsely about endorsement.[55] And in *White v. Samsung Electronics America, Inc.*, the court extended publicity rights beyond name and likeness to protect against appropriation of a celebrity's entire "identity"—holding that even a robot evoking Vanna White's persona triggered liability.[56]

Lanham Act that would expand the unfair competition protections to include false representations regarding authorship of expressive works." U.S. Copyright Off., Authors, Attribution, and Integrity: Examining Moral Rights in the United States 4 (2019).

[52] See Jonathan Vanian, 5 Takeaways from CNBC's Investigation into 'Nudify' Apps and Sites, CNBC (Sep. 28, 2025, at 09:16 ET), https://www.cnbc.com/2025/09/28/5-takeaways-from-cnbcs-investigation-into-nudify-apps-and-sites.html [https://perma.cc/M833-JXG6]; Margi Murphy, 'Nudify' Apps That Use AI to 'Undress' Women in Photos Are Soaring in Popularity, Time (Dec. 8, 2023, at 20:10 ET), https://time.com/6344068/nudify-apps-undress-photos-women-artificial-intelligence/ [https://perma.cc/P46N-ZLSP].

[53] Kate Conger & John Yoon, Explicit Deepfake Images of Taylor Swift Elude Safeguards and Swamp Social Media, N.Y. Times (Jan. 26, 2024), https://www.nytimes.com/2024/01/26/arts/music/taylor-swift-ai-fake-images.html.

[54] 849 F.2d 460, 463 (9th Cir. 1988).

[55] 978 F.2d 1093, 1098–1102 (9th Cir. 1992).

[56] 971 F.2d 1395, 1398–99 (9th Cir. 1992).

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

Tennessee's "ELVIS Act," effective July 1, 2024, creates a property right in one's voice, including AI simulations.[57] The ELVIS Act represents the first U.S. legislation specifically targeting AI voice cloning. It extends publicity rights to include any voice "readily identifiable and attributable to a particular individual," regardless of whether living or deceased—with post-mortem rights running for at least ten years and continuing for as long as heirs commercially exploit them—and explicitly covers AI-generated simulations and creates secondary liability for platforms and tools that facilitate unauthorized cloning.[58] The proposed federal "NO FAKES Act" would extend similar protections nationally, creating a right in voice and visual likeness persisting for up to seventy years after death, contingent on active and authorized public use of the protected attribute.[59]

In *Lehrman v. Lovo, Inc.*, voice actors sued an AI company for cloning their voices;[60] the court dismissed copyright claims—holding that copyright in a sound recording does not extend to imitative AI outputs that merely simulate the original—but allowed right of publicity claims to proceed.[61] The *Lehrman* court's analysis reveals how different IP regimes address authenticity through distinct mechanisms. The court dismissed the plaintiff's Lanham Act false endorsement claims because the cloned voices were used as products (the AI voices themselves) rather than as source identifiers suggesting the plaintiffs endorsed Lovo's services.[62] But the New York right of publicity claims survived because state law protects the authentic human source of the voice regardless of how it is used commercially.[63] The distinction matters: federal trademark law protects source identification, while state personality rights protect the source itself. In an AI economy, both functions are necessary—and neither maps onto copyright's incentive logic. The human persona becomes another form of source identifier—protectable not because it was created, but because it can be verified.

Scholars have debated whether publicity rights should be understood through trademark's source-identification lens or as a distinct species of

[57] Tenn. Code Ann. §§ 47-25-1101 to -1108 (2024).
[58] Id. §§ 47-25-1102, -1104, -1105.
[59] NO FAKES Act, S. 1367, 119th Cong. (2025).
[60] 790 F. Supp. 3d 348, 355 (S.D.N.Y. 2025).
[61] Id. at 379–80.
[62] Id. at 372–73.
[63] Id. at 379.

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

property. Stacey Dogan and Mark Lemley argue that trademark's focus on consumer confusion and false endorsement provides a better framework than copyright's incentive-based approach: publicity rights function to ensure consumers are not deceived about who is endorsing or appearing in content.[64] Mark McKenna offers a different grounding, arguing that publicity rights protect "autonomous self-definition" (the individual's interest in controlling how her identity is publicly presented) applicable to celebrities and ordinary individuals alike.[65] On either account, the right protects authentic human identity, not creative output.

Yet civil liberties organizations have raised substantial First Amendment concerns. The Foundation for Individual Rights and Expression warns that "almost all of the covered 'replicas' are fully protected by the First Amendment" and creates a "digital heckler's veto whereby one complaint can erase lawful speech."[66] California courts have attempted to navigate this tension through a "transformative use" test borrowed from copyright fair use. In *Comedy III Productions, Inc. v. Gary Saderup, Inc.*, the California Supreme Court held that publicity rights must yield when a work "adds significant creative elements so as to be transformed into something more than a mere celebrity likeness or imitation."[67] But the transformative use standard, developed for human expression, fits awkwardly with AI-generated content that may be highly "creative" in a technical sense while still exploiting identity without consent. The doctrinal tools developed for human expression may require rethinking in an era of synthetic media. The tension between protecting identity and protecting speech remains unresolved.

III. THE POLITICAL ECONOMY OF AUTHENTICITY

The shift to source identification systematically advantages incumbents over entrants, capital over labor, and organizations over individuals. A caveat: the copyright regime being displaced was never a

[64] Stacey L. Dogan & Mark A. Lemley, What the Right of Publicity Can Learn from Trademark Law, 58 Stan. L. Rev. 1161, 1164, 1187–88 (2006).

[65] Mark P. McKenna, The Right of Publicity and Autonomous Self-Definition, 67 U. Pitt L. Rev. 225, 227–29 (2005).

[66] John Coleman, The NO FAKES Act Is a Real Threat to Free Expression, Found. For Individ. Rts. & Expression (Nov. 13, 2025), https://www.thefire.org/news/no-fakes-act-real-threat-free-expression [https://perma.cc/PT9J-3ACN].

[67] 21 P.3d 797, 799 (Cal. 2001).

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

paradise for individual creators. Publishers, labels, and studios have always captured the lion's share of copyright's value. This is not a new observation. Scholars have long documented how copyright law has consistently been shaped by industry incumbents through legislative processes that excluded individual creators.[68] The romantic vision of copyright protecting the "solitary genius" has always obscured a regime designed by and for intermediaries. The critique here is not nostalgic; it is that the emerging regime offers even less to the unaffiliated creator, not that its predecessor offered enough.

A. Whose Authenticity?

In a world where content is abundant and undifferentiated, the winners are those who already possess recognized brands. *The New York Times* can sue to protect its trademark; a freelance journalist has no comparable asset. The *Times* can invest in technical provenance frameworks (e.g., C2PA credentials); an independent creator cannot.[69]

The apparent counterexample, namely individual OnlyFans creators monetizing their identities, deserves careful treatment. The platform has paid over $25 billion to creators since 2016, and genuinely self-made success stories exist.[70] Sophie Rain, a Miami-area model in her early twenties who built her audience after a stint as a minimum-wage waitress, reportedly earned approximately $43 million in her first year on OnlyFans and has since documented $101 million in lifetime earnings.[71] Her rise was platform-native: a fired waitress when she launched her account, Rain built parallel presences on TikTok, Instagram, X, and Snapchat to funnel viewers to her $5–$10 OnlyFans subscription.[72] For creators like Rain, the OnlyFans economy represents

[68] See, e.g., Jessica Litman, Digital Copyright 35–69 (2001).

[69] See infra note 83 for an explanation of C2PA.

[70] Henderson, supra note 3.

[71] Eileen Cartter, Behind the Scenes with Sophie Rain, the Clean Queen of the Dirty Internet, GQ (Mar. 23, 2026), https://www.gq.com/story/sophie-rain-profile; Joe Price, OnlyFans Model Sophie Rain Earned $43 Million in Her First Year on the Platform, Complex (Nov. 29, 2024), https://www.complex.com/pop-culture/a/backwoodsaltar/onlyfans-sophie-rain-earnings [https://perma.cc/3YA5-FN6Q]; Maggie Ekberg, Sophie Rain Drops Receipts Showing More Than $101M in OnlyFans Earnings, Complex (Jan. 26, 2026), https://www.complex.com/pop-culture/a/maggie-ekberg/sophie-rain-net-worth-100m [https://perma.cc/P4CS-6YQZ].

[72] Cartter, supra note 71.

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

a genuine pathway to wealth accumulation that traditional entertainment industries never offered.

Yet the aggregate figures obscure a troubling distribution. According to industry data, the top one percent of creators capture roughly thirty percent of platform revenue, while the median creator earns approximately $131 per month—fewer than fifteen dollars a day, before taxes and platform fees.[73] The platform's top earner lists remain dominated by those who brought pre-existing celebrity: Cardi B, Bella Thorne, and Iggy Azalea.[74] Sophie Rain herself attributes her success to "an IMMENSE amount of luck" and warns that average creators earn almost nothing.[75] The model *can* reward the unknown creator who grinds her way to recognition, but the odds are closer to a lottery than a labor market.

More fundamentally, even successful creators depend on the platform as a verification intermediary. OnlyFans certifies that the account belongs to the person it claims to be; the platform extracts twenty percent for providing this authentication service.[76] This rate is dramatically more favorable than what YouTube retains (30–55%),[77] what major record labels historically captured (80–90%),[78] or what

[73] OnlyFans Statistics 2026 – Number of Creators & Top Earners, StatisticsOnly.Fans, https://ofstats.net [https://perma.cc/D77T-4U9R] (last visited Apr. 21, 2026). This inequality is not unique to OnlyFans: it characterizes the creator economy as a whole. Brooke Erin Duffy's research on "aspirational labor" documents how content creation holds the promise of social and economic capital while delivering highly uneven rewards, noting approximately eight percent of creators earn enough to support themselves, while the vast majority engage in what amounts to uncompensated reputation-building. Brooke Erin Duffy, (Not) Getting Paid to Do What You Love: Gender, Social Media, and Aspirational Work 4–7, 197–212 (2017).

[74] OnlyFans Statistics 2026, supra note 73 (listing Cardi B, Bella Thorne, and Iggy Azalea among the platform's highest-earning creators and identifying Thorne by her role as "CeCe Jones in Disney's Shake It Up").

[75] Ekberg, supra note 71.

[76] Fenix Int'l Ltd., Annual Report and Consolidated Financial Statements for the Year Ended 30 November 2023, at 2 (2024), https://www.upmarket.co/content/files/2024/09/Fenix-International--OnlyFans--September-5-2024-Filing.pdf [https://perma.cc/KG7A-MU95].

[77] YouTube Partner Earnings Overview, YouTube, https://support.google.com/youtube/answer/72902?hl=en#zippy=%2Cwhats-my-revenue-share [https://perma.cc/GJW5-NNQZ] (last visited May 17, 2026).

[78] Ben Sisario, A New Spotify Initiative Makes the Big Record Labels Nervous, N.Y. Times (Sep. 6, 2018), https://www.nytimes.com/2018/09/06/business/media/spotify-music-industry-record-labels.html.

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

studios kept in traditional adult entertainment (100%, with performers receiving only flat per-scene fees).[79]

The music industry offers a cautionary parallel. Spotify paid more than $11 billion to rights holders in 2025, yet of an estimated thirteen million artists who have uploaded music to the platform, only about 81,000—fewer than one percent—earned $10,000 or more, and tracks with fewer than 1,000 annual streams generated no royalties at all under the platform's monetization threshold.[80] The shift from physical to digital distribution promised democratization; what emerged was platform oligopoly. OnlyFans' 80/20 revenue split is dramatically more favorable than Spotify's economics, but the structural pattern—a small number of superstar creators capturing the vast majority of platform revenue while the median creator earns poverty wages—replicates across content markets.

The improvement is real. But the platform's economics reveal where value concentrates: in fiscal year 2024, OnlyFans' sole owner, Leonid Radvinsky, extracted $497 million in dividends—equivalent to the annual earnings of approximately 317,000 average creators.[81]

The illustrator whose style is replicated by a generative model captures the broader problem. Under the old regime, copyright doctrine was at least open to claims grounded in the artist's distinctive personality and style. Under the new regime, the protective shift toward trademark and source-identification doctrines systematically advantages those who have already accumulated the brand capital that those doctrines presuppose.[82] Her name is not famous; her "trade dress" is not registered. She is, in effect, unprotectable: a casualty of a system organized around institutional identity rather than individual creativity.

B. The Infrastructure of Trust

The triple-lock structure described in this Essay's introduction—identity as the hook, infrastructure as lock, enforcement as border

[79] See Carrie Weisman, The Porn Industry Is a Lot Less Lucrative Than You Might Think, Salon (Jan. 31, 2016, at 01:00 ET), https://www.salon.com/2016/01/31/how_much_money_do_porn_stars_really_make_partner/ [https://perma.cc/MPR6-22QC].

[80] Your Questions, Answered, Spotify: Loud & Clear, https://loudandclear.byspotify.com/#faq (last visited June 14, 2026).

[81] Spangler, supra note 2; OnlyFans Statistics 2026, supra note 73.

[82] See, e.g., Note, Recovering Personality in Copyright's Originality Inquiry, 138 Harv. L. Rev. 1123 (2025).

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

patrol—is already being generalized beyond adult content through the C2PA standard.[83] For example, the EU AI Act's transparency mandates point toward a regulatory future in which AI outputs must be watermarked and labeled. The EU AI Act requires that providers of generative AI systems ensure outputs are "marked in a machine-readable format and detectable as artificially generated or manipulated," and that deployers disclose when content depicts deepfakes or AI-generated text on matters of public interest.[84] The December 2025 Draft Code of Practice signals a shift from "'principle-based' transparency to specific technical requirements involving multi-layered watermarking, metadata standards, and user-facing icons."[85] This regulatory convergence suggests the authenticity infrastructure question is not whether provenance verification will become mandatory but who will control the systems that provide it. This sounds like consumer protection, and in part it is. But mandatory labeling also consolidates the position of platforms as arbiters of authenticity.

Hardware implementing C2PA at the point of capture remains confined to expensive professional equipment: the Leica M11-P was the first camera to natively embed Content Credentials at capture, and the Nikon Z6III added compatibility through Nikon's Authenticity Service in 2025. The Leica retails for over $10,000, and the Nikon for approximately $2,700.[86] While the Google Pixel 10 introduced the first

[83] The Coalition for Content Provenance and Authenticity ("C2PA") is an industry consortium, founded by Adobe, Arm, BBC, Intel, Microsoft, and Truepic, that develops open technical standards for certifying the origin and editing history of digital content through cryptographically signed provenance metadata. See About Coalition for Content Provenance and Authenticity (C2PA), Coal. For Content Provenance & Authenticity, https://c2pa.org/about/ [https://perma.cc/8NZ7-KYJJ] (last visited May 17, 2026); FAQs, Coal. For Content Provenance & Authenticity, https://c2pa.org/faqs/ [https://perma.cc/A8T4-34VV] (last visited May 17, 2026); Content Credentials: C2PA Technical Specification, Coal. For Content Provenance & Authenticity, https://spec.c2pa.org/specifications/specifications/2.3/specs/C2PA_Specification.html [https://perma.cc/Z644-UUFW] (last visited May 17, 2026).

[84] Council Regulation 2024/1689, ch. IV, art. 50, ¶¶ 2, 4, 2024 O.J. (L) 1 (EU).

[85] EU AI Act: First Draft Code of Practice on AI-Generated Content Transparency, Shibolet Law Firm (Dec. 24, 2025), https://www.shibolet.com/en/eu-ai-act-first-draft-code-of-practice-on-ai-generated-content-transparency/ [https://perma.cc/DU7W-5A8V].

[86] M11-P Technical Specification, Leica Camera AG, https://leica-camera.com/en-US/photography/cameras/m/m11-p-black/technical-specification [https://perma.cc/A73X-SJN3] (last visited May 17, 2026); Nikon Authenticity Service, Nikon USA, https://www.nikonusa.com/content/nikon-authenticity-service [https://perma.cc/L6LT-HR7M] (last visited May 17, 2026); Nikon Z6III: Tech Specs, Nikon USA, https://www.nikonusa.com/p/z6iii/1890/overview#tech-specs [https://perma.cc/PR3P-LTNV] (last visited May 17, 2026).

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

mainstream smartphone support for C2PA in late 2025, it remains a solitary island in a mobile market where affordable consumer devices still lack native capture-level credentials.[87] Furthermore, maintaining credentials through the editing process requires ongoing Adobe Creative Cloud subscriptions—often ranging from $60 to over $100 per month.[88] Because the system's trust hierarchy is anchored by third-party Certificate Authorities who charge fees under their own commercial terms, the very architecture of digital authenticity is effectively being privatized.[89]

Human rights organizations have flagged the emerging divide. WITNESS warns of a coming digital world "split between content that does have . . . verifiable provenance and content that does not," cautioning that "evidence of human rights violations could be dismissed in legal instances for not including cryptographic metadata" while "those speaking truth to power could be forced to choose between visibility and safety."[90] *Fortune* puts the equity concern bluntly: "if small media outlets, indie journalists, or independent creators don't make the list, their work could be penalized or dismissed."[91] The infrastructure of authenticity is being built by those who can afford to participate in its governance. Who is trusted, and who decides, are governance questions. And the current answer is Adobe, Microsoft, Leica, and the platforms that read their credentials.

[87] See Eric Lynch & Sherif Hanna, How Pixel and Android Are Bringing a New Level of Trust to Your Images with C2PA Content Credentials, Google Online Sec. Blog (Sep. 10, 2025), https://security.googleblog.com/2025/09/pixel-android-trusted-images-c2pa-content-credentials.html [https://perma.cc/8ZK2-3B7B].

[88] Plans and Pricing for Creative Cloud Apps and More, Adobe, https://www.adobe.com/creativecloud/plans.html (last visited May 17, 2026) (listing the all-apps Creative Cloud Pro plan, which includes Content Credentials functionality, at $69.99 per month for individual subscribers).

[89] The C2PA's own conformance program acknowledges that "the actual fees are set by individual [Certificate Authorities] and subject to their commercial terms." FAQs, supra note 83; see also Kate Kaye & Pam Dixon, Privacy, Identity and Trust in C2PA: A Technical Review and Analysis of the C2PA Digital Media Provenance Framework, World Privacy Forum (Sep. 2025).

[90] WITNESS is an international human rights organization that trains and supports the use of video and technology in human rights documentation and advocacy. Tomorrow's Great Digital Divide: Content With or Without Provenance, WITNESS Blog (Mar. 5, 2025), https://blog.witness.org/2025/03/tomorrows-great-digital-divide/ [https://perma.cc/W7S6-AVFP].

[91] Sharon Goldman, Big Tech's Standard for Fighting AI Fakes Puts Privacy on the Line, Fortune (Sep. 18, 2025, at 10:46 ET), https://fortune.com/2025/09/18/big-techs-c2pa-content-credentials-standard-for-fighting-deepfakes-puts-privacy-on-the-line/.

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

IV. GOVERNANCE AS THE REMAINING VARIABLE

If the shift from incentive to verification is inevitable, the remaining question is governance. Who will control the infrastructure of authenticity? Currently, the answer is a consortium of hardware manufacturers and software giants. A public option for verification is essential to prevent the aristocracy of authenticity from becoming permanent.

A. The Structural Case for Public Infrastructure

The C2PA standard's technical architecture embeds a particular answer to the governance question: trust flows downward from corporate root authorities through intermediate certificates to end users.[92] The entity that controls the root controls who can be trusted. This is a policy choice masquerading as a technical necessity.

Alternative architectures exist. The KERI protocol enables self-certifying identifiers that do not depend on any centralized authority.[93] Decentralized identity standards developed by the World Wide Web Consortium offer portable credentials not locked to any single platform.[94] A public provenance system could adopt such approaches while providing a state-backed anchor for those who want institutional verification and, critically, for those who cannot afford the private alternative.

The case for public infrastructure is structural rather than ideological. If verification is the new locus of value, and if access to verification determines who can participate in content markets, then verification infrastructure is too important to leave entirely to private gatekeepers. The same logic that justifies public roads, public courts, and public vital records systems justifies a public option in digital trust. A public verification layer would not eliminate C2PA or Adobe credentials; it

[92] See Content Credentials: C2PA Technical Specification, supra note 83, §§ 6, 14 (describing the hierarchical trust model in which trust anchors, typically corporate certificate authorities, issue signing credentials to content creators, and verifiers check signatures against a configurable trust list); Kaye & Dixon, supra note 89 (providing an independent technical review and privacy analysis of the C2PA framework trust architecture).

[93] KERI, https://keri.one/ [https://perma.cc/THA6-2JRQ] (last visited May 17, 2026).

[94] Decentralized Identifiers (DIDs) v1.0, World Wide Web Consortium (July 19, 2022), https://www.w3.org/TR/did-1.0/. The Consortium is the principal international standards organization for the World Wide Web.

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

would ensure that the capacity to be believed is not monopolized by their successors.

B. The Library of Congress as Digital Notary

The natural home for this infrastructure is the Library of Congress. For 150 years, the Library of Congress has served as the nation's centralized registry of creative claims through the Copyright Office.[95] In the print era, it managed the "scarcity of creation" through copyright registration. In the AI era, it should manage the "scarcity of verification" through identity credentialing.

The institutional capacity already exists; what is required is an expansion of mission. The Library of Congress is actively modernizing its technical infrastructure, with major contracts underway for network systems, software development, and the Copyright Office's Enterprise Copyright System.[96] The Copyright Public Record System launched in June 2025, and a new registration system is under development.[97] The Copyright Office has explicitly acknowledged the need to adapt its registration function for the digital age. Its 2024 report on digital replicas recommended urgent federal legislation to address "gaps in existing laws" for individuals whose likenesses are replicated without consent, recognizing that the current registration system, designed for fixed works, cannot address harms rooted in identity appropriation.[98] A public verification registry would extend this institutional evolution: just as the Copyright Office adapted from registering physical deposits to accepting digital uploads, it could adapt from certifying authorship claims to certifying identity claims. The pipes are being laid; the question is what flows through them.

This proposal aligns with emerging global models of Digital Public Infrastructure. These models embody a principle increasingly recognized in digital governance: that foundational infrastructure enabling market participation should be provided as a public layer, with

[95] History of U.S. Copyright Office, Libr. of Cong.: U.S. Copyright Off., https://www.copyright.gov/history/copyright-exhibit/history-co/ [https://perma.cc/7RE5-RBBH] (last visited May 17, 2026).

[96] Continuous Development of IT Systems, Libr. of Cong.: U.S. Copyright Off., copyright.gov/copyright-modernization/ [https://perma.cc/7YXL-6MHQ] (last visited May 17, 2026).

[97] Id.

[98] U.S. Copyright Office, Copyright and Artificial Intelligence, Part 1: Digital Replicas 4 (2024).

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

private innovation occurring on top.[99] The analogy to physical infrastructure is precise. The interstate highway system did not eliminate private trucking companies; it ensured that the capacity to move goods was not monopolized by railroad barons who controlled the only existing infrastructure. Estonia's e-Residency program offers state-verified digital identity to non-citizens, enabling participation in European markets without physical presence.[100] These models demonstrate that the state can provide the base layer of verification without controlling the applications built upon it. By providing that base layer, the state prevents platform monopolies from enclosing the market for trust.

A Library of Congress-anchored verification system would offer concrete functions: a voluntary registry for creators to establish timestamped identity claims; a public API for platforms to verify those claims without becoming gatekeepers themselves; and a fee structure ensuring access is not conditioned on ability to pay Adobe's subscription fees or Leica's hardware prices. Critically, a public system need not displace existing private infrastructure: it could interoperate with it. The C2PA specification is designed to accommodate multiple trust anchors—there is no technical barrier to a Library of Congress certificate authority operating alongside Adobe's Content Authenticity Initiative.[101] What a public option would provide is an alternative for those excluded from or unable to afford private credentialing: the independent journalist, the whistleblower, the creator in a developing country, the dissident artist whose government pressures private platforms to exclude her. The goal is not to replace the private infrastructure of authenticity but to ensure it does not become the only infrastructure. The independent journalist could obtain credentials as robust as the *New York Times*'s. The whistleblower could establish provenance without purchasing a $9,000 camera. The dissident artist

[99] See generally Rahul Matthan, The Third Way: India's Revolutionary Approach to Data (2023) (describing that digital system in India's public infrastructure).

[100] Bruce Sterling, Estonian E-Residency, an Estonian Primer, Wired (Oct. 18, 2017, at 06:38 ET), https://www.wired.com/beyond-the-beyond/2017/10/estonian-e-residency-estonian-primer/.

[101] Content Credentials: C2PA Technical Specification, supra note 83, § 14.4 (providing that verifiers may configure their own "trust lists" of recognized certificate authorities, such that a certificate from the Library of Congress could be recognized alongside existing private-sector trust anchors without requiring changes to the core specification).

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

could prove authenticity without routing trust through corporations that may have their own reasons to exclude her.

The details (credentialing procedures, liability frameworks, interoperability with existing C2PA infrastructure) require development that exceeds this Essay's scope. But the essential point is that these are design choices, not technological inevitabilities. The infrastructure of authenticity can be public or private. The certificate chain can root in Adobe or in the Library of Congress. We are choosing, right now, by failing to choose. The private option is being built. The public option is not. That asymmetry will harden into permanence unless the governance question is asked before it answers itself.

CONCLUSION

Intellectual property's organizing logic is transforming. The old regime asked: "Did you create this?" The new regime asks: "Can you prove you are who you say you are?"—and, just as importantly, can you afford the infrastructure to make that proof stick?

To observe these dynamics is not to deny that this economy creates genuine winners. Twenty-five billion dollars has flowed to creators—more annually than most major professional sports leagues pay their athletes.[102] For some performers, particularly those who would have received flat per-scene fees under the old studio system, the 80/20 split and content ownership represent a revolutionary improvement in bargaining power. Sophie Rain built a hundred-million-dollar fortune through grinding cross-platform promotion, not inherited celebrity. The question is not whether *anyone* benefits, but *who* benefits systematically, and at whose expense.

The incentive paradigm that AI has destabilized was always only one of several competing justifications for IP protections. Personality theory, the Kantian and Hegelian view that creative works embody and externalize the creator's identity, offers a framework in which protecting authentic human source is not an adaptation to technological change but

[102] See, e.g., What Is the Highest Paid Sport in the World?, World Sports Network, https://www.wsn.com/blog/highest-paid-sport/ [https://perma.cc/F3FW-J9KK] (last visited May 17, 2026), (reporting total annual player salaries in 2023 as approximately $4.68 billion for the National Basketball Association, $3.4 billion for Major League Baseball, and $2.6 billion for the National Hockey League); Henderson, supra note 3.

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

a return to first principles.[103] On this account, the emerging source-identification regime vindicates what IP should always have protected: not the work as commodity, but the work as expression of personhood. The question is whether that vindication will be available to all, or only to those who can afford the infrastructure to prove who they are.

The emerging architecture channels value toward those who already possess brand capital and infrastructural capacity, but also permits a narrow class of self-made creators to accumulate real wealth through verification of their identity—offering individual performers ownership rights that traditional entertainment industries never provided. And it positions platforms as essential intermediaries in a market for trust, intermediaries whose owners extract hundreds of millions annually while the median creator earns poverty wages and pays monthly subscriptions for the privilege of enforcing her own copyrights. The adult-content industry has historically served as the canary in the coal mine for digital technology shifts. VHS, secure payments, streaming video: where adult content went, the broader economy followed. The triple-lock architecture now emerging on OnlyFans will not remain confined to adult content. When AI-generated text floods journalism, when synthetic voices saturate podcasting, when machine-generated code populates repositories, every content market will face the same provenance crisis—and the same pressure to construct walled gardens of authenticated identity. The question is whether those gardens will be privately owned toll roads or public infrastructure.

The incentive paradigm is collapsing; what replaces it will be decided by choices about governance, not technology. A public root of trust—a Library of Congress reimagined as digital notary—could ensure that the capacity to be believed is not conditioned on the capacity to pay Adobe's subscription fees or Leica's hardware prices. The alternative is an aristocracy of authenticity, hard-coded into certificate chains controlled by the same platforms that already dominate digital distribution.

The OnlyFans economy is not a dystopia for everyone, nor a meritocracy for anyone. It is a system that rewards identity differently than its predecessors—better for some, worse for others, and controlled by neither creators nor public. The infrastructure of authenticity is being

[103] See Justin Hughes, The Philosophy of Intellectual Property, 77 Geo. L.J. 287, 330–50 (1988); Margaret Jane Radin, Property and Personhood, 34 Stan. L. Rev. 957, 959–61 (1982).

COPYRIGHT © 2026 VIRGINIA LAW REVIEW ASSOCIATION

built right now. The question is whether anyone other than its owners will have a voice in its design.