



**Stanford – Vienna  
Transatlantic Technology Law Forum**

A joint initiative of  
Stanford Law School and the University of Vienna School of Law



# **TTLF Working Papers**

**No. 154**

**Algorithmic Couture:  
Trademark Protection and the Rule of Law  
in the Rule of Code  
Across the EU, US, and China**

**Maria Lucia Passador**

**2026**

# **TTLF Working Papers**

**Editors: Siegfried Fina, Mark Lemley, and Roland Vogl**

## **About the TTLF Working Papers**

TTLF's Working Paper Series presents original research on technology, and business-related law and policy issues of the European Union and the US. The objective of TTLF's Working Paper Series is to share "work in progress". The authors of the papers are solely responsible for the content of their contributions and may use the citation standards of their home country. The TTLF Working Papers can be found at <http://ttlfs.stanford.edu>. Please also visit this website to learn more about TTLF's mission and activities.

If you should have any questions regarding the TTLF's Working Paper Series, please contact Vienna Law Professor Siegfried Fina, Stanford Law Professor Mark Lemley or Stanford LST Executive Director Roland Vogl at the

Stanford-Vienna Transatlantic Technology Law Forum  
<http://ttlfs.stanford.edu>

Stanford Law School  
Crown Quadrangle  
559 Nathan Abbott Way  
Stanford, CA 94305-8610

University of Vienna School of Law  
Department of Business Law  
Schottenbastei 10-16  
1010 Vienna, Austria

## **About the Author**

Maria Lucia Passador is an Assistant Professor of Corporate Law and Financial Markets Regulation, Bocconi University, Milan, Italy; Adjunct Professor, SUPSI; Affiliate, Transatlantic Technology Law Forum (TTLF), Stanford Law School.

## **General Note about the Content**

The opinions expressed in this paper are those of the author and not necessarily those of the Transatlantic Technology Law Forum or any of its partner institutions, or the sponsors of this research project.

## **Suggested Citation**

This TTLF Working Paper should be cited as:  
Maria Lucia Passador, Algorithmic Couture: Trademark Protection and the Rule of Law in the Rule of Code Across the EU, US, and China, Stanford-Vienna TTLF Working Paper No. 154, <http://ttlfs.stanford.edu>.

## **Copyright**

© 2026 Maria Lucia Passador

## Abstract

Trademark protection has undergone a paradigmatic shift—from reactive, document-based enforcement to an anticipatory, algorithmic regime. This shift materializes differently across consumer marketplaces, social-commerce environments, and B2B verification infrastructures, where enforcement is exercised by platforms, brand owners, or hybrid arrangements with distinct incentives, data access, and error-cost profiles. Artificial intelligence is no longer an auxiliary instrument; it has become the primary operational layer of brand protection, translating legal standards into operational code. Drawing on doctrinal analysis, technical case studies, and comparative legal insights, the paper maps the entire lifecycle of AI deployment in trademark enforcement—from unsupervised threat detection and predictive vendor scoring to generative adversarial mimicry and algorithmic evidence.

It makes three core contributions. First, it reveals a feedback loop between legal legibility and algorithmic precision: the accuracy of AI enforcement depends on the informational quality of the rights it protects. Second, it conceptualizes the emerging “algorithmic cold war,” in which generative models weaponize imitation to fabricate counterfeit brand narratives and evade detection at scale. Third, it examines how courts in Germany, the United States, the United Kingdom, and China are assimilating AI-generated evidence, and how the EU Artificial Intelligence Act is beginning to constitutionalize the legitimacy of algorithmic enforcement.

Taken together, these contributions illuminate how efficiency, legitimacy, and institutional design converge in the algorithmic governance of trademarks, transforming enforcement from a reactive process into a constitutional architecture of legality by design.

The Article concludes by advancing a theory of algorithmic trademark resilience—an institutional framework anchored in explainability, auditability, and constitutional design. In the age of machine enforcement, it argues, the legitimacy of trademark protection depends less on technical speed than on institutional foresight: on building systems that can detect infringement at scale while remaining open to judicial scrutiny and capable of preserving the rule of law by design through verifiable documentation, contestability, and human oversight—recognizing that the evidentiary admissibility and procedural legitimacy of AI-generated outputs remain contested across jurisdictions.

## CONTENTS

|    |                                                                            |    |
|----|----------------------------------------------------------------------------|----|
| 1. | INTRODUCTION                                                               | 1  |
| 2. | PREDICTIVE ENFORCEMENT AND THE PRIVATIZATION OF DETECTION                  | 10 |
| 3. | OPERATIONAL PREVENTION: PLATFORMS AS ALGORITHMIC REGULATORS                | 15 |
| 4. | COLLABORATIVE INTELLIGENCE: FEDERATED ENFORCEMENT AND INSTITUTIONAL DESIGN | 18 |
| 5. | EMPIRICAL FRONTIERS: AI TRADEMARK ENFORCEMENT IN PRACTICE                  | 23 |
| 6. | ALGORITHMIC EVIDENCE AND THE POLITICAL ECONOMY OF PROOF                    | 29 |
| 7. | EMBEDDING AI ACROSS THE TRADEMARK LIFECYCLE                                | 35 |
| 8. | THE AI COUNTER-OFFENSIVE: WEAPONISED INTELLIGENCE IN BRAND INFRINGEMENT    | 40 |
| 9. | CONCLUSION                                                                 | 45 |

### 1. Introduction

Some secrets do not hide in shadows; they hide in plain sight. In *The Da Vinci Code*, Dan Brown once observed that the most intricate codes are often cloaked in familiarity, not obscurity. Much the same can be said of brand threats today. Counterfeiting no longer sits at the periphery of commerce. It is embedded in the infrastructure of e-commerce and social media, scaled through search, recommendation, and automated listing tools. The modern vector of infringement is therefore not the press or the dye, but the algorithmic pathways that manufacture visibility and trust. This shift does not emerge *ex nihilo*. It builds on earlier waves of automated notice-and-takedown systems—initially manual, later boutique, and eventually industrialized—as extensively documented in the intellectual property literature, particularly in the copyright context. What distinguishes the present phase, however, is not automation *per se*, but the migration of

enforcement into a predictive, infrastructural regime in which algorithmic systems no longer merely execute legal judgments, but actively structure visibility, risk, and compliance *ex ante*.<sup>1</sup>

Trademark enforcement is moving from an interpretive, court-centered system to a computational, platform-centered one. AI is no longer a peripheral aid; it increasingly operates as the primary layer through which detection, prioritization, takedown, and evidentiary packaging occur. That shift does not merely increase speed. It relocates where legal standards are operationalized—from doctrine and procedure to datasets, model thresholds, and governance controls. What Dan Brown called codes hidden in plain sight now describes trademark law’s new terrain: enforcement has moved from courtrooms to codebases. This Article argues that algorithmic enforcement is transforming trademark law from an interpretive system into a computational one, requiring a new framework of legitimacy grounded in explainability and constitutional design. In that sense, trademark protection in the AI era marks the embedding of legality into technical design-law translated into architecture. What was once a matter of judicial reasoning now unfolds as an architecture of compliance, where legality must be engineered rather than merely declared.

A caveat is necessary. Assertions that algorithmic enforcement can be scaled without sacrificing accuracy, fairness, or legal defensibility are not neutral descriptions; they are contested normative claims. At least three fault lines recur in the literature and in institutional practice: (i) accuracy varies materially by domain, data quality, and adversarial pressure, and scale can amplify false positives alongside true positives; (ii) opacity and asymmetric access to models can undermine contestability and shift error costs onto smaller sellers and users; and (iii) evidentiary ‘defensibility’ depends not merely on predictive performance but on explainability, reproducibility, and

---

\* Assistant Professor of Corporate Law at Bocconi University; Associated Researcher, European Banking Institute; Research Fellow, Baffi Centre on Economics, Finance and Regulation; Transatlantic Technology Law Forum Fellow, Stanford Law School.

This paper builds on the author’s remarks originally presented at the 13th Brand Protection Excellence Forum, held in Milan on July 3, 2025, and incorporates additional bibliographic references and analytical developments prepared specifically for this written contribution.

<sup>1</sup> See, e.g., JENNIFER M. URBAN & JOE KARAGANIS, NOTICE AND TAKEDOWN IN EVERYDAY PRACTICE, (2016) <https://dx.doi.org/10.2139/ssrn.2755628>.

procedural pathways that allow the affected party to understand and challenge the basis of enforcement.<sup>23</sup>

Trademark protection has moved from procedures to prediction: AI both shields and threatens brand integrity, narrowing the line between detection and deception.<sup>4</sup> This Article does not assume that such predictive systems operate uniformly across contexts. OCR applied to high-volume B2B documentation flows differs institutionally and legally from behavioral detection of fraud in consumer marketplaces. Accordingly, references to “algorithmic trademark enforcement” in this paper denote a structured but heterogeneous set of deployments whose commonality lies in the automation of detection, prioritization, and proof-generation, rather than in any single technical method.

Algorithms now operationalize the categories of confusion, dilution, and bad faith long before courts or agencies pronounce upon them.<sup>5</sup> Legal standards are no longer operationalized only through ex post adjudication; they are increasingly encoded ex ante in datasets, thresholds, and false-positive tolerances.<sup>6</sup> This inversion marks a profound epistemic realignment: what once depended on interpretive reasoning now emerges from probabilistic computation.<sup>7</sup> More precisely,

---

<sup>2</sup> Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1, 7–8 (2014).

<sup>3</sup> See OECD Principles on Artificial Intelligence, OECD (last visited May 17, 2026) [<https://perma.cc/HF2T-RKQ2>].

<sup>4</sup> See, e.g., George-Mihai Irimescu, *Artificial Intelligence and Trademark Protection*, 16 CHALLENGES OF THE KNOWLEDGE SOC’Y 532 (2023); Christine Haight Farley, *Trademarks in an Algorithmic World*, 98 WASH. L. REV. 1123 (2023); Alpana Roy & Althaf Marsoof, *Removing the Human from Trademark Law*, 55 INT’L REV. INTELL. PROP. & COMPETITION L. 727 (2024); Dev Saif Gangjee, *Panoptic Brand Protection? Algorithmic Ascendancy in Online Marketplaces*, 46 EUR. INTELL. PROP. REV. 448 (2024).

<sup>5</sup> The delegation of interpretive authority to algorithmic systems reshapes not only enforcement but norm production itself. In platform environments, the categories of “confusion” or “bad faith” are converted into machine-readable variables—similarity scores, clickstream overlaps, and complaint ratios—thereby creating quasi-doctrinal metrics embedded directly in code.

<sup>6</sup> Every detection threshold embeds an implicit welfare function. Raising sensitivity lowers expected infringement loss while increasing expected exclusion error. In this sense, algorithmic trademark enforcement operationalizes a regulatory choice about acceptable error rates—one rarely made explicit in doctrinal analysis.

<sup>7</sup> Modern enforcement systems approximate Bayesian updating: prior beliefs about infringement risk are continuously revised based on observed seller behavior. The result is not adjudication, but sequential risk management under uncertainty.

algorithms now operate as a primary operational locus of enforcement, often preceding and conditioning judicial interpretation, rather than supplanting it.

From the standpoint of economic theory, this shift is not surprising. Trademark law has always mediated information asymmetries between producers and consumers. Its central function—signaling origin and assuring quality—rests on the Coasean intuition that credible signaling reduces transaction costs in markets of imperfect information. Yet in a digital marketplace governed by machine curation, those asymmetries have changed form. The relevant information is no longer solely about product quality or source identity, but about the data provenance, authenticity, and traceability of digital representations themselves.<sup>8</sup> A counterfeit today is not simply a misbranded good; it is an engineered distortion in the information environment—a false signal optimized for algorithmic visibility.

This reconceptualization forces a deeper confrontation with the institutional economics of trademark enforcement. The traditional model treated enforcement as a public good administered by courts, customs, and agencies. The emerging model is privatized, automated, and increasingly governed by economic incentives. Platform intermediaries—Amazon, Alibaba, TikTok—now function as quasi-regulatory institutions,<sup>9</sup> wielding algorithmic tools that determine what is visible, legitimate, or delisted.<sup>10</sup> Recent scholarship on Amazon’s Brand Registry argues that platform-operated registries can function as a parallel, quasi-administrative trademark system that reshapes incentives to register marks and reallocates enforcement power away from public institutions.<sup>11</sup> Their role is not merely technical but constitutive: they allocate enforcement rights, set evidentiary

---

<sup>8</sup> Control over provenance data creates a secondary monopoly over enforcement credibility. The political economy concern is not counterfeit suppression, but the consolidation of epistemic authority in a small number of data-rich intermediaries.

<sup>9</sup> Privatized enforcement reallocates not only cost but power. Platforms gain agenda-setting authority over which infringements matter, when, and at what scale—an archetypal feature of regulatory capture, albeit exercised through infrastructure rather than lobbying.

<sup>10</sup> Visibility functions as a rationed economic resource. Algorithmic enforcement reallocates visibility ex ante, determining market access before any transaction occurs. This resembles licensing rather than liability—a shift with significant political-economic implications.

<sup>11</sup> Jeanne C. Fromer & Mark P. McKenna, *Amazon’s Quiet Overhaul of the Trademark System*, 113 CALIF. L. REV. 1169 (2025).

thresholds, and translate legal concepts into computational heuristics. In effect, they perform a form of delegated regulation, one that blurs the line between public authority and private governance.<sup>12</sup>

This privatization embodies a Coasean internalization of enforcement costs<sup>13</sup>—now refracted through platform architectures. Platforms internalize the externalities of counterfeit trade<sup>14</sup>—consumer confusion, reputational loss, regulatory scrutiny—by embedding enforcement directly into their architecture.<sup>15</sup> The result is a quasi-market in compliance, where detection, takedown, and verification operate as tradable services. This logic echoes the broader evolution of platform governance identified by Zuboff and Pasquale, where informational asymmetry is not merely corrected but commodified. Yet this same efficiency generates new frictions: asymmetries of power between rightsholders and small sellers,<sup>16</sup> the opacity of automated decision-making, and the displacement of judicial oversight by platform protocols. What appears efficient in a microeconomic sense may prove destabilizing in an institutional one.<sup>17</sup>

This dynamic reveals a paradox at the heart of algorithmic enforcement. AI systems depend for their accuracy on the legibility of the rights they are meant to protect. Where trademarks are unregistered, poorly categorized, or legally ambiguous, machine learning models

---

<sup>12</sup> The phenomenon resembles functional delegation in financial regulation, where private intermediaries assume quasi-public responsibilities without formal administrative status. The legal tension lies not in delegation per se, but in the absence of corresponding fiduciary and procedural constraints.

<sup>13</sup> Ronald Coase's insight that firms internalize externalities when transaction costs permit takes on a distinctly architectural form in platform governance. Enforcement is no longer merely internalized—it is embedded, with code functioning as the boundary of the firm. See R.H. Coase, *The Problem of Social Cost*, 3 J.L. & ECON. 1 (1960).

<sup>14</sup> While algorithmic enforcement dramatically lowers marginal transaction costs of monitoring, it raises fixed governance costs—model development, auditability, and dispute resolution. The economic question is not whether enforcement is cheaper, but whether governance costs scale as efficiently as detection.

<sup>15</sup> This form of internalization parallels how financial intermediaries embed compliance into transaction infrastructure: enforcement becomes endogenous to the market design rather than an external corrective imposed by the state.

<sup>16</sup> Automated enforcement redistributes risk downward. Small sellers face higher relative error costs because they lack the scale, legal resources, or data access needed to contest algorithmic decisions—an effect analogous to regressive compliance costs in financial regulation.

<sup>17</sup> From an error-cost perspective, algorithmic enforcement reallocates rather than eliminates social cost. Gains in detection accuracy reduce false negatives but often increase false positives, shifting error burdens from counterfeiters to lawful sellers. See generally Louis Kaplow & Steven Shavell, *Accuracy in the Assessment of Damages*, 39 J. LAW & ECON. 191 (1996).

lose precision; they cannot “see” what the law itself has failed to define. The effectiveness of private enforcement thus presupposes the clarity and completeness of the public registry—a reminder that information design, not just legal doctrine, underwrites the economic value of marks. In this new terrain, visibility itself becomes a precondition of validity:<sup>18</sup> what cannot be rendered legible to code risks falling outside the reach of legal protection.<sup>19</sup>

The phenomenon also complicates the canonical narrative of law’s hierarchy. Courts, agencies, and private actors have always shared the work of norm articulation: commercial practice, platform rules, industry standards, and technical constraints co-produce the “law in action” that structures markets. What has changed is sequencing and scale. In the algorithmic marketplace, the traditional order—statute, adjudication, compliance—is increasingly inverted or bypassed. Code enforces first; law interprets later. Doctrinal categories such as “confusion” or “use in commerce” are not merely adjudicated *ex post*, but instantiated *de facto* through platform governance and automated workflows. As economic sociologists might say, the locus of norm creation is migrating from the courtroom to the interface.

To keep the analysis concrete, consider a common enforcement workflow on a major marketplace: (i) automated models flag listings based on multimodal similarity and seller-behaviour signals; (ii) internal tooling converts that output into a takedown packet; (iii) the platform’s escalation rules determine whether human review occurs; and (iv) the resulting action (removal, account restriction, or reinstatement) becomes training feedback that shifts future thresholds. This is the paper’s core institutional point in miniature: legality is not merely applied to facts after the fact—it is increasingly instantiated by design choices in the enforcement pipeline.

---

<sup>18</sup> Algorithmic enforcement renders registrability and data hygiene economically salient. Rights that are doctrinally valid but informationally illegible risk functional extinction—a result foreign to classical trademark theory but familiar to information economics.

<sup>19</sup> Legibility here refers not merely to registration formalities, but to data ontologies—how trademark registries are structured, annotated, and interoperable with platform databases. Machine unreadability can thus function as a new species of legal invisibility.

This transformation challenges foundational assumptions about the division of labor between law and markets. In the classical Chicago framework, trademark law is justified as a market-corrective mechanism that supports efficient exchange by reducing search costs. But when the market itself becomes the site of legal enforcement—when detection, verification, and sanctioning occur within proprietary systems—the corrective becomes endogenous. Trademark law ceases to correct the market; it becomes one of its internal design parameters.<sup>20</sup> The resulting equilibrium is no longer purely legal or purely economic, but *sociotechnical*: an equilibrium produced through the interaction of statutes, algorithms, and incentives.

Comparatively, jurisdictions are beginning to diverge in their responses. In the European Union, the Artificial Intelligence Act may apply to components of AI-enabled trademark enforcement depending on the system’s intended purpose and whether it falls within the Act’s high-risk categories (notably Annex III) or other applicable obligations. Even where a system is not classified as high-risk, the Act’s governance logic reinforces expectations around documentation, transparency, and human oversight in automated enforcement workflows.<sup>21</sup> The United States, by contrast, relies on sectoral delegation, allowing platforms and rights-holders to self-regulate under general doctrines of notice and takedown.<sup>22</sup> China’s Internet Courts have gone further, incorporating blockchain and machine evidence into formal adjudication.<sup>23</sup> Each model reflects a different balance between efficiency and accountability, between market governance and legal legitimacy.

What unites them, however, is a common recognition: enforcement has become architectural. In a world where data can be forged and models inverted, the sustainability of

---

<sup>20</sup> The classic Chicago-school justification of trademark law as a market-corrective device presupposes an external legal intervention. Algorithmic enforcement collapses that distinction, rendering law endogenous to market design.

<sup>21</sup> Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act), OJ L, 2024/1689, 12.7.2024 [hereinafter AI Act].

<sup>22</sup> Urban et al., *supra* note 1.

<sup>23</sup> Mark Jia, *Special Courts, Global China* 62 VA. J. INT’L L. 559 (2022)

trademark protection depends less on procedural formality than on the *design* of information systems—on traceability, explainability, and institutional foresight. The brand, once a static emblem of identity, now functions as a living data structure<sup>24</sup>—a continuously updated ledger of authenticity and trust.

The challenge for trademark law is therefore not to resist this transformation, but to govern it. The task is to ensure that algorithmic enforcement remains compatible with the constitutional virtues of transparency, proportionality, and due process. It is to reassert that the efficiency of private ordering cannot displace the legitimacy of public reasoning. The future of trademark protection lies not in stronger fences but in smarter institutions—institutions capable of translating economic rationality into normative coherence, and technological design into the rule of law.

In the end, the question is not whether AI can protect trademarks more effectively. The real question is whether the systems we build to safeguard markets can themselves be trusted as part of the market’s moral infrastructure. In that sense, the brand has always been more than a sign of commerce; it is a signal of credibility. To preserve that signal in the algorithmic age, law must learn not only to read the code—but to write within it.

Institutionally, algorithmic enforcement is exercised by three overlapping actors. Platforms police infringement to mitigate consumer harm, reputational risk, and regulatory exposure—while simultaneously optimizing for traffic and liquidity. Brand owners and their agents prioritize integrity and scarcity but often lack access to platform-level behavioral metrics. Hybrid arrangements seek to combine these capacities, redistributing evidentiary power and enforcement discretion. These incentive structures shape not only enforcement outcomes, but the operational meaning of trademark concepts such as confusion, bad faith, and use.

The paper offers an integrated framework that moves from marketplace context to technical deployment and finally to legal legitimacy. First, it explores how machine learning (ML)

---

<sup>24</sup> This reconceptualization aligns trademark governance with contemporary corporate data-governance debates, where value increasingly attaches not to static assets but to continuously curated informational systems.

models—particularly unsupervised detectors, transformer-based natural language processors, and Convolutional Neural Networks (CNNs)—now operate as the vanguard of *ex post* surveillance. Second, it turns to *ex ante* applications, including vendor risk scoring, geo-temporal forecasting, and takedown prioritization—tools that shift enforcement from reaction to anticipation. Third, the analysis reveals a structural paradox: that the effectiveness of such systems is inextricably tied to the quality and accessibility of the very rights they are designed to protect. Fourth, it addresses the rise of AI-driven infringement—the “counter-offensive”—where generative models are now used to fabricate not only products but entire brand narratives, evading detection through adversarial mimicry. Fifth and finally, the article interrogates the evidentiary implications of AI-enforced trademark protection: the rising demand for auditability, transparency, and procedural fairness in light of the EU AI Act and cross-border judicial scrutiny.

What emerges is not merely a new set of tools, but instead a new theory of brand resilience—one premised on algorithmic traceability, federated intelligence, and legal readiness. In an ecosystem where every digital trace can be forged, and every model inverted, enforcement must be reimagined as an act of architectural design.<sup>25</sup> And in this new landscape, the brand is no longer a static asset, it stands as a living cipher: a pattern of data that must be made legible, defensible<sup>26</sup> and, above all, anticipatory. As in Brown’s narrative, the question is not simply *who* holds the truth, but *who can see it in time*.

In short, the future of trademark protection will hinge not on stronger rights, but on the architecture of their intelligibility.

---

<sup>25</sup> The metaphor is deliberate. As with complex financial systems, meaning here emerges from interaction rather than authorship—a reminder that in algorithmic governance, no single actor “writes” the law, yet all are bound by it.

<sup>26</sup> In what follows, ‘legal defensibility’ refers to a bundle of constraints—admissibility (or probative acceptance), reliability and reproducibility of the pipeline, transparency sufficient for adversarial testing, and procedural contestability for affected parties—rather than to model accuracy alone. See Pamela Langham, *Applying Daubert and Frye to AI Evidence*, MD. STATE BAR ASS’N. (Oct. 23, 2024), [https://www.msba.org/site/site/content/News-and-Publications/News/General-News/Applying\\_Daubert\\_and\\_Frye\\_to\\_AI\\_Evidence.aspx](https://www.msba.org/site/site/content/News-and-Publications/News/General-News/Applying_Daubert_and_Frye_to_AI_Evidence.aspx) [<https://perma.cc/5YDC-PKLD>].

## 2. Predictive Enforcement and the Privatization of Detection

AI-enabled monitoring has shifted trademark enforcement from episodic, reactive intervention to continuous surveillance and prioritization. Unsupervised anomaly detectors, multimodal similarity models, and behavioral risk scores now shape which listings are escalated, which sellers are restricted, and which disputes never reach a court. In that practical sense, AI systems do not just detect infringement; they structure the enforcement pipeline and thereby influence how infringement is defined in practice. At the core of AI-enabled threat detection lies a subtle but profound shift: from reactive enforcement to predictive surveillance. The difference, as any economist might note, is not semantic but structural. Predictive systems invert enforcement: they forecast risk rather than react to harm.<sup>27</sup>

This predictive logic thrives on asymmetry—the same asymmetry that justified trademark protection in the first place. Classic trademark doctrine existed to solve an informational problem: consumers lacked reliable signals about the origin or quality of goods. The mark served as a shortcut for trust, a legal prosthesis compensating for imperfect information. In the platform economy, however, the information asymmetry has mutated. The consumer may still be confused, but the confusion no longer stems from lack of data; it stems from too much of it. The signal itself risks drowning in noise—and the algorithm steps in to reconstruct trust by inference.

The detection stack described in this Part often originates in general anti-abuse tooling—scam, spam, and fraud classifiers trained on behavioral and transactional anomalies. In this article, however, those signals are analysed only to the extent that platforms and rightsholders use them as trademark-enforcement proxies: to surface counterfeiting, passing off, and other bad-faith marketplace conduct that operationally maps onto “use in commerce,” consumer confusion, and dilution. Put differently, the model may be trained to answer “does this look like illicit automation or coordinated abuse?”, but the legal question this paper tracks is narrower: when and how such

---

<sup>27</sup> In predictive enforcement, the legal act is anticipated by the probabilistic forecast. This creates an inversion of temporal order—where suspicion precedes harm, and mitigation replaces adjudication.

detections are translated into trademark-relevant determinations and remedies (delisting, prioritised review, evidentiary packaging) rather than generic fraud suppression.

Unsupervised models now flag anomalies—metadata, price patterns, or image hashes—as statistical symptoms of illicit commerce<sup>28</sup> that often correlate with trademark harm. The counterfeit, in short, is treated as a trademark-salient anomaly in the dataset—a statistical pathology within the marketplace graph that platforms triage for IP review.<sup>29</sup>

This development is both ingenious and troubling, and it transforms the evidentiary burden: suspicion becomes statistical, and legality is inferred rather than demonstrated.

The logic of *likelihood of confusion*—a human inquiry into perception—is displaced by the logic of *likelihood of deviation*—a machine’s inference from correlation. What was once an interpretive standard becomes a statistical one. Somewhere between *Polaroid v. Polarad* and *PageRank*, the doctrinal figure of the “average consumer” has been quietly recoded as a probabilistic entity<sup>30</sup> — a construct inferred from behavioural data rather than direct human perception.

The technical mechanics may sound arcane— Density-Based Spatial Clustering of Applications with Noise (DBSCAN) clustering, transformer embeddings, anomaly detection—but their legal implications are anything but.<sup>31</sup> The system learns not what infringement *is*, but what it

---

<sup>28</sup> Not all fraud signals are coextensive with trademark infringement. Behavioral anomaly detection may capture payment fraud, account takeover, review manipulation, or spam that is orthogonal to trademark doctrine. The overlap is operational rather than conceptual: platforms frequently route anti-abuse outputs into IP workflows because the same behavioral fingerprints often accompany counterfeit distribution and passing off.

<sup>29</sup> This redefinition translates a normative judgment (“infringement”) into a data deviation. In practical terms, liability becomes a question of outlier detection—a profoundly different epistemology of wrongdoing.

<sup>30</sup> This marks a quiet but profound doctrinal mutation: the “average consumer” of trademark law—long a proxy for human perception—reappears here as a statistical construct inferred from behavioral data. The shift mirrors parallel developments in antitrust and securities regulation, where enforcement increasingly targets patterns rather than intent.

<sup>31</sup> For instance, see Govind A & Rohith Syam, *Using DBSCAN to Identify Customer Segments with High Churn Risk on Amazon Consumer Behavior Data*, TECHRXIV (Jan. 25, 2024), <https://www.techrxiv.org/users/715598/articles/700262/master/file/data/DBSCAN1/DBSCAN1.pdf?inline=true> [https://perma.cc/D4GP-EEL2].

*looks like*.<sup>32</sup> It is the jurisprudence of plausibility: patterns become proxies for wrongdoing.<sup>33</sup> MarqVision’s density-based algorithms, for instance, do not identify a single infringing product; they map constellations of sellers behaving too similarly to be innocent.<sup>34</sup> Their deterrent power lies less in accuracy than in omnipresence. When detection feels ubiquitous, deterrence becomes ambient. Law becomes infrastructure—embedded in the architectures of code that now govern markets.<sup>35</sup>

Linguistic mimicry adds another layer to this computational semiotics. Counterfeiters long ago discovered that perfect imitation is less profitable than strategic resemblance. Hence, the proliferation of “Chanelle,” “Luis Vutton,” and “Güccci”—a taxonomy of strategic approximation. Transformer-based language models, trained on oceans of text, now detect this mimicry not by spelling but by semantics.<sup>36</sup> They recognize when a phrase functions as a brand identifier in an e-commerce title, even when it appears in non-standard, phrase-like, or keyword-

---

<sup>32</sup> Andrew Tausz, *Similarity Clustering to Catch Fraud Rings*, STRIPE (Feb 20, 2020), <https://stripe.com/blog/similarity-clustering> [<https://perma.cc/X6EM-D4EC>].

<sup>33</sup> Algorithmic plausibility functions as a proxy for probabilistic harm. In economic terms, “likelihood of confusion” is no longer a qualitative standard but a risk-allocation device, approximating the Kaplow–Shavell insight that optimal enforcement minimizes the sum of enforcement costs and expected harm rather than pursuing factual certainty.

<sup>34</sup> See MARQVISION, <https://www.marqvision.com/abm-2/growve#:~:text=45.39,a%20high%20traffic%20marketplace> [<https://perma.cc/GGW7-G8MZ>]; *AI Anti-Counterfeiting*, PRISM SUSTAINABILITY DIR. (Apr. 9, 2025), <https://prism.sustainability-directory.com/term/ai-anti-counterfeiting/#:~:text=,norms%2C%20can%20highlight%20potential%20counterfeits> [<https://perma.cc/5L39-45XX>].

<sup>35</sup> The phrase suggests a diffusion of legal influence through technological mediation—rules are not applied as discrete commands but experienced as ambient constraints within digital infrastructures.

<sup>36</sup> NLP techniques, particularly those employing transformer-based architectures such as BERT (Bidirectional Encoder Representations from Transformers), RoBERTa (A Robustly Optimized BERT Pretraining Approach), and DeBERTa (Decoding-enhanced BERT with Disentangled Attention), have become instrumental in brand protection contexts. Trained on vast datasets comprising product descriptions, user reviews, and advertising text, these models excel in detecting subtle forms of linguistic imitation — for instance, phonetically similar brand references, intentionally misspelled trademarks, and euphemistic descriptions that evoke brand identity without infringing on protected terms directly. This semantic nuance is particularly valuable in combating “near-infringement” strategies employed by counterfeiters, who often avoid verbatim use of trademarks while still exploiting their semantic field to mislead consumers. Such models can be fine-tuned to surface deceptive patterns across multilingual platforms and dynamically evolving content ecosystems. See, e.g., Yinhan Liu et al., *RoBERTa: A Robustly Optimized BERT Pretraining Approach*, ARXIV (2019); Pengcheng He et al., *DeBERTa: Decoding-Enhanced BERT with Disentangled Attention*, ARXIV (2021); Zecong Wang et al., *Implementing BERT and Fine-Tuned RoBERTa to Detect AI-Generated News by ChatGPT*, ARXIV (2023).

based form.<sup>37</sup> A review saying “luxury inspired by Parisian heritage” may trigger a flag precisely because it resembles machine-generated brand-adjacent promotional language.<sup>38</sup>

It is important to clarify that the detection of linguistic mimicry by natural language processing (NLP) systems does not, in itself, amount to a determination of trademark infringement. From a doctrinal perspective, the use of suggestive or phonetically proximate terms in product listings—such as “Chanelle” or “luxury inspired by Parisian heritage”—will not automatically qualify as use of a sign as a trademark, particularly where such terms appear in descriptive text rather than as indicators of commercial origin.

Rather, the function of NLP-based detection lies upstream of legal qualification. These systems operate as risk-identification mechanisms, flagging patterns of semantic proximity, stylistic imitation, or narrative alignment that may warrant closer legal scrutiny. In this sense, algorithmic detection does not replace the doctrinal inquiry into “use as a mark in the course of trade,” but restructures the enforcement pipeline by identifying contexts in which such use may plausibly arise—such as where linguistic mimicry is coupled with visual similarity, product category overlap, or platform-level presentation that collectively signals origin association.

Accordingly, AI-driven linguistic analysis should be understood not as asserting infringement, but as operationalizing pre-adjudicatory relevance: it narrows the universe of listings requiring human or legal assessment, without collapsing the distinction between deceptive suggestion and legally actionable trademark use.

This shift also reframes error. Trademark doctrine asks whether consumers are likely to be confused; automated enforcement systems ask whether a listing is sufficiently similar or statistically

---

<sup>37</sup> Zhen Wang et al., *Continuous Prompt Tuning Based Textual Entailment Model for E-commerce Entity Typing*, ARXIV (2022). In a notable deployment, the model was able to *escalate* risk scores based on cross-linguistic cues—a Chinese-language listing using specific idioms often linked with counterfeit promotion was flagged and removed before any human escalation was required. *See also* Fred Rocafort & Dan Harris, *Protecting YOUR Intellectual Property From Alibaba and Other Chinese E-Commerce Sites* (July 3, 2023), <https://harris-sliwoski.com/chinalawblog/protecting-your-intellectual-property-from-alibaba-and-other-chinese-e-commerce-sites/?utm> [https://perma.cc/V9C6-YSEX].

<sup>38</sup> Clinton Amos & Lixuan Zhang, *Consumer Reactions to Perceived Undisclosed ChatGPT Usage in an Online Review Context*, 93 *TELEMATICS & INFORMATICS* 102163 (Sept. 2024).

anomalous to warrant intervention. The risk is that correlation becomes a substitute for justification: patterns become proxies for wrongdoing, and false positives can scale with the same ease as true positives. Accordingly, any account of ‘efficient’ algorithmic enforcement must incorporate contestability, calibration, and safeguards against over-removal as integral design variables—not afterthoughts. Computer-vision pipelines add further capacity. A CNN’s computer-vision pipeline trained on authorized brand imagery can detect deviations in logos, packaging geometry, or typography. Coupled with Optical Character Recognition (OCR) and multimodal embedding models (e.g., Contrastive Language–Image Pre-training (CLIP)-style architectures), these systems can parse large volumes of listings, extract identifiers, cross-check them against internal databases, and generate structured takedown packets. The operational consequence is that private enforcement becomes both faster and more standardized—while the justificatory basis for individual decisions can become harder to reconstruct unless logging and explanation are designed in from the outset. Yet the elegance of this automation conceals a deep institutional asymmetry. Predictive systems rely on visibility, and visibility belongs to platforms. The datasets that make detection possible are privately held; the algorithms that interpret them are proprietary. Enforcement, once an act of public authority, now unfolds as a function of private infrastructure. In economic terms, enforcement has been *verticalized*. Platforms have internalized the cost of surveillance because it maximizes their reputational capital, but in doing so they have also monopolized the epistemic basis of legal truth.<sup>39</sup> The counterfeit risks becoming, in operational terms, whatever the algorithm classifies as such.

---

<sup>39</sup> This dynamic reflects the Coasean logic of internalization, whereby firms absorb external enforcement costs to minimize transaction frictions within their own boundaries. See R. H. Coase, *The Nature of the Firm*, 4 *ECONOMICA* 386 (1937); OLIVER E. WILLIAMSON, *MARKETS AND HIERARCHIES* (1975). In the platform economy, such internalization manifests as vertical integration of trust functions: by embedding enforcement directly into architecture, intermediaries convert legal uncertainty into reputational capital—what Reinier Kraakman terms the “gatekeeper” function of reputation. See R. Kraakman, *Gatekeepers: The Anatomy of a Third-Party Enforcement Strategy*, 2 *J.L. ECON. & ORG.* 53 (1986). Yet this efficiency produces a paradox familiar to information economics: as platforms minimize transaction costs, they centralize the informational rents of enforcement, effectively monopolizing the epistemic basis of legal truth. Cf. George Akerlof, *The Market for Lemons*, 84 *Q.J. ECON.* 488 (1970); SHOSHANA ZUBOFF, *THE AGE OF SURVEILLANCE CAPITALISM* (2019).

This privatization of vision challenges the very logic of trademark as a public right in reputation. The mark's traditional role was to translate private goodwill into public information. But in the algorithmic marketplace, the direction of translation reverses: private data becomes the new source of public meaning. The platform—not the registry, not the court—decides what constitutes authenticity, and it does so in code. One might say, updating Lessig, that *code is not only law; it is evidence*.

Trademark law thus finds itself in a curious position. It must now govern technologies that govern it. The future of threat detection will depend not only on the volume of data collected, but on the institutional capacity to preserve contestability and doubt. For if every anomaly becomes an infraction, then the distinction between vigilance and control dissolves—and with it, the fragile legitimacy of algorithmic enforcement.<sup>40</sup>

AI can predict risk; the question is whether law can still predict justice.<sup>41</sup> Yet the more enforcement migrates from human judgment to machine inference, the more detection merges into governance itself. The next section examines this transformation—how predictive systems evolve into operational prevention, where platforms function as algorithmic regulators rather than passive enforcers.

### 3. Operational Prevention: Platforms as Algorithmic Regulators

In the modern political economy of trademark protection, *operational prevention* marks the point where enforcement ceases to be reactive and becomes infrastructural. Detection locates the breach; prevention redesigns the conditions under which breaches occur. In this respect, artificial

---

<sup>40</sup> In contrast to traditional views that treat transparency as a constraint on efficiency, error-cost theory suggests the opposite. Where enforcement is automated, explainability reduces expected social costs by lowering uncertainty, enabling behavioral adjustment, and constraining false positives. In this sense, transparency operates not as a moral add-on but as an efficiency input.

<sup>41</sup> This line signals a structural transformation rather than a rhetorical one. As predictive architectures replace human deliberation with probabilistic inference, the locus of legal reasoning shifts from normative justification to computational optimization. The question is therefore not merely philosophical but institutional: whether justice, traditionally articulated through procedure and public reasoning, can remain intelligible when expressed in the language of models and metrics.

intelligence does not simply make enforcement faster—it makes it *structural*. It converts what was once an episodic function of litigation into a continuous process of governance. The shift is not technological alone: it reconfigures the institutional foundations of trademark law itself, displacing adjudicatory authority into the logic of platforms and predictive systems.

The underlying logic can be traced to the economic theory of the trademark. Although Coase did not address trademarks directly, his transaction-cost framework provides the conceptual foundation for understanding trademarks as devices that reduce the costs of market exchange. Landes and Posner later developed this insight explicitly, arguing that trademarks reduce information asymmetry between producers and consumers by operating as reliable signals of quality and origin.<sup>42</sup> A mark economizes on search and verification costs; it converts uncertainty into trust. Yet this elegant efficiency model assumes a world of bilateral exchange and observable conduct—conditions that digital markets have upended. In e-commerce ecosystems, informational asymmetry arises less from the product itself than from the *metadata* surrounding it: the provenance of listings, the authenticity of imagery, the behavioral fingerprints of sellers, and the algorithmic pathways that bring goods to consumers.

Against this backdrop, AI-driven prevention systems operationalize what economists would call *risk internalization*. They do not merely detect infringing behavior *ex post*; they price reputational and legal risk into participation *ex ante*. Vendor-risk scoring models—combining metadata histories, IP reputation, upload cadence, linguistic mimicry, and image duplication—produce dynamic *trust coefficients* that determine whether a seller is fast-tracked, shadow-monitored, or excluded altogether.<sup>43</sup> What emerges is a Coasean market in compliance: enforcement becomes

---

<sup>42</sup>See Coase, *supra* note 13 (providing the transaction-cost framework); William M. Landes & Richard A. Posner, *Trademark Law: An Economic Perspective*, 30 J.L. & ECON. 265 (1987) (developing an economic theory of trademark law based on search-cost reduction and source-identification functions); WILLIAM M. LANDES & RICHARD A. POSNER, THE ECONOMIC STRUCTURE OF INTELLECTUAL PROPERTY LAW 166–209 (2003).

<sup>43</sup> Such models effectively quantify trust, turning qualitative legal judgments into scalar probabilities. The “trust coefficient” functions as a private analogue to the state’s licensing power. FINANCIAL SERVICES INFORMATION SHARING AND ANALYSIS CENTER (FS-ISAC), GENERATIVE AI VENDOR RISK ASSESSMENT GUIDE (2024) [https://www.fsisac.com/hubfs/Knowledge/AI/FSISAC\\_GenerativeAI-VendorEvaluation&QualitativeRiskAssessment.pdf](https://www.fsisac.com/hubfs/Knowledge/AI/FSISAC_GenerativeAI-VendorEvaluation&QualitativeRiskAssessment.pdf) [https://perma.cc/8VSJ-XDHL]. See, generally, R. Marshal, *A*

a form of transaction-cost minimization, but one administered privately by platforms rather than publicly by courts.

This privatization of enforcement reflects a larger institutional realignment. As Fromer and McKenna have shown in *Amazon's Quiet Overhaul of the Trademark System*,<sup>44</sup> platforms increasingly function as *delegated regulators* of brand legitimacy, setting the procedural and evidentiary baselines that determine visibility, removability, and liability. Their algorithms instantiate norms of “use in commerce,” “likelihood of confusion,” and “bad faith” long before—and sometimes instead of—judicial interpretation. This creates a de facto regime of algorithmic self-regulation in which platforms become primary sites of operational norm-setting, translating legal standards into enforcement rules, evidentiary thresholds, and removal practices upstream of judicial review.

Yet the efficiency of this model reveals its own limits. Economic theory teaches that while internalizing externalities can reduce social costs, it can also generate new monopolies of information and power. The same platforms that enforce IP norms also define the rules of access to the marketplace. Their dual role as regulator and participant gives rise to institutional conflicts of interest: exclusionary practices may be disguised as risk management, and automated takedowns may occur without procedural safeguards. From the standpoint of comparative law, the challenge resembles that faced by financial regulation in the early age of compliance automation: how to ensure accountability in systems where *code performs governance*. This analogy is more than rhetorical; it underscores that algorithmic enforcement, like early RegTech, demands a fusion of technical validation and normative oversight if legitimacy is to accompany efficiency.

Yet even the most sophisticated platform-based prevention remains structurally bounded by informational silos and fragmented incentives. A single marketplace may internalize risk, but counterfeit networks arbitrage across platforms, jurisdictions, and modalities—listings, social

---

*Robust and Implementable Approach for AI Vulnerability Risk Scoring*, 1 INT’L J. OF CYBERNETICS AND CYBER-PHYSICAL SYSTEMS 349 (2024).

<sup>44</sup> Jeanne C. Fromer & Mark P. McKenna, *Amazon's Quiet Overhaul of the Trademark System*, 113 Ca. L. Rev. (2025) 1169.

commerce, and synthetic media—exploiting the gaps between proprietary detection stacks. At that point, the problem is no longer merely one of delegated regulation within a platform, but of coordination across institutions. The next Part therefore turns to collaborative intelligence: how federated enforcement architectures can pool detection capacity without collapsing confidentiality, competition, or due process—and how institutional design can convert cooperative efficiency into public-facing legitimacy. Seen through this lens, trademark enforcement becomes a laboratory for reimagining legality itself: a site where economic rationality and normative authority must be continuously co-produced. The challenge is not simply to automate enforcement, but to constitutionalize it—to ensure that the speed of code does not outrun the deliberative tempo of law. Historical precedent bears this out. The Internet’s evolution through rough consensus and running code, and the Asilomar moratorium on recombinant DNA, offer two imperfect but useful analogies for productive restraint: one procedural, the other substantive. In both cases, governance emerged not from premature legal closure, but from a period of disciplined deliberation in which expert communities stabilised technical and ethical norms before formal rules took firmer shape.<sup>45</sup> Yet even this hybrid of private ordering and public oversight remains limited by informational silos. The following section explores how collaborative intelligence—through federated enforcement networks—can overcome those silos and distribute both efficiency and legitimacy across institutions.

#### **4. Collaborative Intelligence: Federated Enforcement and Institutional Design**

This Part moves from platform-centered prevention to network-centered enforcement, asking how cooperation can be engineered without reproducing either data centralization or unchecked private power.

---

<sup>45</sup> See Pete Resnick, *On Consensus and Humming in the IETF*, INTERNET ENGINEERING TASK FORCE (2014); Michael J.D. Vermeer, *Historical Analogues That Can Inform AI Governance*, RAND CORP. (2024).

Digital counterfeiters exploit network effects; enforcement must therefore evolve into a coordinated, network response.

From the standpoint of law and economics, algorithmic enforcement functions as a transaction-cost technology—internalizing coordination failures that once made collective vigilance prohibitively expensive. Yet this very efficiency invites a new legitimacy dilemma: the same infrastructures that enable cooperation also concentrate epistemic authority.

Federated learning, originally developed for privacy-preserving analytics in health and finance, now supplies a model for collective enforcement without full data disclosure. Participants retain local control of proprietary datasets but contribute encrypted parameter updates to a shared model that continuously refines its detection capacity. This enables rights-holders to realize economies of scope in enforcement—the capacity to learn from others’ experience without bearing the disclosure cost of cooperation. A fashion conglomerate monitoring image-based counterfeits, a pharmaceutical firm tracking gray-market APIs, and a consumer-electronics company mapping QR-code frauds can each train a joint classifier that identifies behavioral or linguistic anomalies common to all. The output is not a centralized database but a distributed algorithmic equilibrium: a shared intelligence layer that reduces redundancy and amplifies deterrence.<sup>46</sup>

Federated learning converts enforcement from isolated vigilance into cooperative equilibrium—not by pooling secrets, but by engineering trust through decentralized computation, controlled aggregation, and, where appropriate, cryptographic safeguards.<sup>47</sup> In institutional-

---

<sup>46</sup> Such architectures exploit positive network externalities while preserving data sovereignty—a balance that traditional IP enforcement frameworks were not designed to manage.

<sup>47</sup> What is lacking in trademark protection is not the technology — it is the institutional willingness to operationalise collaboration at scale. See Jeyimmy Lorena Viracacha Peña, *Trademarks Have No Passport: Toward a Transnational Protection Based on Reputation and Legal Best Practices* (May 1, 2025), <https://ssrn.com/abstract=5257116>. The future of IP protection lies in interdependence by design: systems built to learn from each other, not in spite of competition, but because of it. The notion of interdependence by design reflects a paradigmatic shift in the architecture of IP enforcement. Traditional models have favoured siloed approaches—each rights-holder maintaining its own proprietary detection, takedown, and litigation strategies. However, the increasingly decentralized and networked nature of IP threats—ranging from generative counterfeiting and adversarial SEO to cross-platform infringement networks—demands a coordinated, intelligence-sharing response. Interdependence does not imply the forfeiture of

economic terms, it lowers the “price” of cooperation while preserving the incentives for self-help, generating what Ostrom would call polycentric governance of enforcement.<sup>48</sup> Law, in this sense, ceases to prescribe cooperation; it becomes the framework within which cooperation is technologically feasible. This transformation suggests that normativity itself is becoming infrastructural — encoded within interoperability protocols as much as in statutory texts.

The institutional implications are profound. Historically, trademark enforcement was administered through hierarchical legal processes: courts, customs authorities, and administrative agencies. Those institutions imposed *ex post* order—reacting to infringements once visible. Today, however, enforcement unfolds within platform governance regimes that operate *ex ante* and at scale. Platforms such as Amazon, Alibaba, and TikTok function as regulatory intermediaries: they design the algorithms that determine visibility, credibility, and takedown thresholds. Their rules—what is flagged, what is verified, what is delisted—constitute, in Lessig’s language, a new modality of norm production where code is law. But unlike the state, these actors internalize both the costs and benefits of enforcement, aligning their incentives with the economics of trust and transaction efficiency rather than with formal legality. The result is an economically endogenous legal order—an enforcement equilibrium governed by risk-weighted optimization rather than adjudicated doctrine. To preserve legitimacy within such endogenous orders, legal scholarship must articulate evaluative criteria that translate efficiency metrics into principles of accountability and fairness.

Seen through the prism of institutional economics, this amounts to a migration from adjudicative legitimacy to institutional legitimacy by design. In the traditional model, the legitimacy of enforcement derived from its procedural form—judicial neutrality, due process, and public

---

competitive advantage, but rather a strategic convergence of technical infrastructures and enforcement insights. This includes the use of federated learning systems, shared threat databases, and collaborative AI pattern recognition to detect emerging threats earlier and more accurately. The premise is simple yet radical: IP systems should not merely survive competition but evolve through it, allowing diverse actors to co-develop anticipatory capabilities that are individually stronger because they are collectively informed. This vision aligns with emergent models of “collaborative resilience” in cybersecurity and brand protection, where ecosystem-wide vigilance replaces reactive, actor-specific containment.

<sup>48</sup> ELINOR OSTROM, *GOVERNING THE COMMONS: THE EVOLUTION OF INSTITUTION FOR COLLECTIVE ACTION* (1990).

accountability. In the emerging model, legitimacy derives from performance metrics: speed, accuracy, false-positive rates, and compliance efficiency. These are functional substitutes for fairness, measurable in real time but conceptually alien to the doctrinal imagination of trademark law. The legal system, in turn, confronts a coordination paradox: to preserve the public character of a norm that is increasingly enforced through private infrastructure.

The empirical trajectory already reflects this hybridization. Amazon’s Counterfeit Crimes Unit and BMW Group’s joint civil action in Spain (2024) demonstrated the productivity of cross-institutional collaboration: combining platform-based AI detection with corporate evidence to obtain injunctive relief and damages. Similarly, Homeland Security’s Operation Stolen Promise, launched during the COVID-19 pandemic, illustrates the operational value of public-private cooperation in combating counterfeit and fraudulent pandemic-related goods. Through partnerships with industry actors such as 3M, Pfizer, Amazon, Alibaba, Citi, and Merck, the initiative combined investigative capacity, sector-specific expertise, and market intelligence to address COVID-19-related illicit trade.<sup>49</sup> These initiatives embody an emergent logic of collective enforcement by design, in which public authority, private expertise, and platform infrastructure converge to overcome enforcement externalities that neither sector could resolve alone.

This evolution also invites caution. Collective intelligence may enhance deterrence, but it risks re-concentrating enforcement power in a few dominant intermediaries. When platforms act simultaneously as marketplace, adjudicator, and enforcer, they occupy the position of the “least-cost avoider”—but without the public constraints that legitimize that role. Economically efficient enforcement can thus drift into oligopolistic governance. The challenge is to design accountability mechanisms that preserve competition, transparency, and procedural fairness within federated

---

<sup>49</sup> See U.S. IMMIGRATION AND CUSTOMS ENFORCEMENT, *HSI Partners with Pfizer, 3M, Citi, Alibaba, Amazon, Merck to Protect Consumers Against COVID-19-Related Fraud* (May 5, 2020), <https://www.ice.gov/news/releases/hsi-partners-pfizer-3m-citi-alibaba-amazon-merck-protect-consumers-against-covid-19> [https://perma.cc/T2E9-PCHM] (noting that industry experts from Pfizer, 3M, Citi, Alibaba, Amazon, and Merck joined HSI-led IPR Center criminal investigators in a public-private partnership to combat fraud and other illegal activity surrounding COVID-19).

systems. In antitrust terms, the task is to balance the allocative efficiency of data sharing with the dynamic efficiency of market entry and innovation.

Comparative experience suggests diverging paths. The European Union, through the AI Act and the Digital Services Act, is attempting to re-embed algorithmic enforcement within a framework of ex ante transparency and human oversight. The United States, by contrast, continues to rely on sectoral delegation and ex post accountability—trusting private enforcement to operate under the general law of contracts, torts, and competition. China’s Internet Courts, meanwhile, have integrated machine evidence and blockchain provenance into judicial procedure, illustrating a model of technocratic formalism. These variations reveal that the same technological substrate can sustain distinct institutional equilibria: privatized in the U.S., co-regulated in Europe, bureaucratized in China.<sup>50</sup> Trademark enforcement, once harmonized by TRIPS, is thus fragmenting along lines of governance rather than doctrine.

The economic rationale for cooperation—reducing information costs, internalizing spillovers, optimizing enforcement—is inseparable from its normative dimension: preserving the conditions of legal legitimacy within algorithmic governance.

The path forward, therefore, demands an integrated institutional response. Effective trademark protection in the algorithmic age will depend on governance frameworks that transform efficiency into legitimacy, interoperability into accountability. The aim is not to retreat from collaboration, but to civilize it—to embed within federated architectures the safeguards that law once supplied through procedure.

In a networked market, resilience is no longer achieved by exclusion, but by the capacity to share intelligence under the discipline of law. The brand’s integrity, like the market’s trust, becomes a collective asset—one that survives only through governance by design.

---

<sup>50</sup> The divergence highlights that algorithmic governance, like monetary or competition policy, embeds local constitutional values into global infrastructures of control.

This dynamic can be visualized as a lifecycle of algorithmic enforcement—an integrated circuit in which data, law, and institutional design co-produce legitimacy. Table 1 (“The Algorithmic Trademark Enforcement Lifecycle”) maps this lifecycle, tracing how enforcement evolves from data capture to normative feedback. Each phase—technical, economic, and legal—functions as a node within a single governance architecture, illustrating how the rule of law becomes embedded in the rule of code.

But to render that governance credible, law must also confront a subtler frontier: the status of *algorithmic evidence* itself. The next part turns from institutional architecture to epistemic authority—how machines now generate, certify, and contest proof.

## 5. Empirical Frontiers: AI Trademark Enforcement in Practice

The normative analysis developed thus far acquires analytical traction only if tested against the empirical configurations through which algorithmic enforcement already operates. The contemporary landscape of trademark protection reveals a clear transition from doctrinal aspiration to institutional reality. Figure 2 maps the institutional ecology of algorithmic enforcement. It translates normative trade-offs into efficiency–legitimacy metrics and shows how different enforcement modalities distribute efficiency and legitimacy across institutional forms.

Three domains—judicial adjudication, private enforcement, and administrative automation—offer distinct yet convergent manifestations of this evolution.

In *Hermès International v. Rothschild*, 590 F. Supp. 3d 647 (S.D.N.Y. 2023), the Southern District of New York confronted the first major conflict between generative art and trademark protection. The defendant, operating under the pseudonym “Mason Rothschild,” marketed a series of “MetaBirkin” non-fungible tokens depicting algorithmically generated fur-covered versions of Hermès’ iconic Birkin bag. Relying on the *Rogers v. Grimaldi* framework, Rothschild claimed First Amendment protection for expressive works.

The jury’s verdict—finding trademark infringement and awarding damages—illustrates the judiciary’s struggle to apply analog doctrines to synthetic media. The harm was not the sale of counterfeit goods, but the algorithmic simulation of a brand’s semiotic identity. From an institutional-economic perspective, *MetaBirkin* exposes the limits of human adjudication when infringement becomes probabilistic rather than perceptual. The case suggests that the “likelihood of confusion” inquiry, long rooted in consumer psychology, is beginning to confront a new problem: the statistical likelihood of generative resemblance. Courts remain anchored to the semantics of deception, while the market operates through the mathematics of simulation.

At the opposite pole of institutional practice stands Amazon’s Counterfeit Crimes Unit (CCU), established in 2020. Staffed by former federal prosecutors, data scientists, and investigators, the CCU epitomizes the privatization of enforcement authority. According to Amazon’s 2023 *Brand Protection Report*, Amazon stopped more than 700,000 attempted new selling-account creations by bad actors in 2023 and identified, seized, and disposed of more than seven million counterfeit products worldwide. Its Counterfeit Crimes Unit has also pursued extensive litigation and referrals, including the joint Amazon/BMW civil action in Spain, in which the European Union Trademark Court in Alicante ruled against four defendants who attempted to sell counterfeit BMW parts and accessories across Europe. The case illustrates the integration of platform-based machine-learning detection with brand-owner evidence and civil enforcement, rather than proving that the counterfeit network itself was AI-assisted.<sup>51</sup>

By embedding detection, adjudication, and sanction within its architecture, the platform minimizes enforcement costs—but efficiency introduces constitutional tension: the same private entity that defines authenticity also exercises quasi-adjudicatory power. The CCU thus operates as a de facto administrative agency without the procedural safeguards of administrative law. The comparative-institutional question is not whether such enforcement is scalable—the reported

---

<sup>51</sup> AMAZON, 2023 *Brand Protection Report* (2024); AMAZON, *Amazon and BMW Group Win First Civil Lawsuit in Spain to Shut Down Counterfeiters* (Feb. 8, 2024), <https://www.aboutamazon.eu/news/policy/amazon-and-bmw-group-win-first-civil-lawsuit-in-spain-to-shut-down-counterfeiters> [<https://perma.cc/UJ9G-HXTY>].

throughput is substantial<sup>52</sup>—but whether it is accurate and procedurally legitimate in the legally relevant sense: whether error rates, escalation pathways, and contest mechanisms are sufficiently transparent to satisfy principles of accountability and due process. but whether it is accurate and procedurally legitimate in the legally relevant sense: i.e., whether error rates, escalation pathways, and contest mechanisms are sufficiently transparent to satisfy principles of accountability and due process.

The European Union Intellectual Property Office’s *AI for TMView* project marks the complementary phenomenon of bureaucratic automation. The initiative integrates natural-language processing and computer-vision tools into trademark classification and similarity assessments across EU registries. Preliminary reports indicate a twenty-percent increase in inter-examiner consistency and a significant reduction in manual review time.

While these results confirm administrative efficiency, they also instantiate a paradigmatic shift from interpretive discretion to algorithmic delegation. The examiner’s judgment is pre-structured by model output; human reasoning becomes a form of post-hoc validation. The EU’s legislative response—particularly the AI Act and the Digital Services Act—reflects an explicit attempt to constitutionalize automation by embedding transparency, documentation, and human-oversight obligations ex ante.

The European approach therefore seeks to preserve legitimacy through design rather than through ex post adjudication.

The luxury industry provides an applied illustration of algorithmic enforcement as an operational infrastructure. LVMH’s Brand Protection Department, in partnership with Entrupy and Microsoft, reported the automated removal of more than sixty-five million infringing listings

---

<sup>52</sup> “Amazon reports large-scale automated detection and removal volumes in its 2023 Brand Protection Report; such metrics evidence enforcement capacity and investment, but they are not equivalent to independently validated accuracy rates or procedurally meaningful contestability.” AMAZON, BRAND PROTECTION REPORT 7 (2024), <https://assets.aboutamazon.com/24/cd/c7e7ccd94abeae59dec904c2b637/2023-brand-protection-report-final-english.pdf> [<https://perma.cc/L3QB-MAKG>].

in 2023. Kering’s 2023 memorandum of understanding with Alibaba and Tencent established a shared AI platform achieving a ninety-five-percent automated detection rate of infringing content. Richemont’s Geneva Digital Hub integrates blockchain authentication and multimodal image forensics across brands such as Cartier and Montblanc.

These practices demonstrate that enforcement has become a *RegTech* function internal to corporate governance. The conglomerates effectively perform quasi-regulatory roles, translating trademark integrity into compliance architecture. In economic terms, algorithmic traceability functions as a competitive asset: the ability to verify authenticity becomes a differentiating factor in markets where reputational capital substitutes for tangible scarcity. Yet, from a normative perspective, the vertical integration of enforcement also concentrates informational power. The capacity to define, detect, and delist infringing content now resides in private infrastructures largely shielded from public oversight.

Across jurisdictions, algorithmic enforcement reflects distinct constitutional equilibria—European design, American contract, Chinese code.

53

| Jurisdiction      | Enforcement<br>Modality | Institutional Logic                | Principal<br>Instruments                       | Efficiency–<br>Legitimacy<br>Tradeoff           |
|-------------------|-------------------------|------------------------------------|------------------------------------------------|-------------------------------------------------|
| European<br>Union | Co-regulation           | Constitutionalization<br>by design | AI Act,<br>Digital Services<br>EU IPO AI Pilot | High legitimacy;<br>slower adaptive<br>learning |

---

<sup>53</sup> Unless independently audited, platform-, vendor-, or consortium-published performance statistics should be read as indicators of operational capacity rather than dispositive measures of legal reliability. In particular, ‘accuracy’ in detection tasks is not isomorphic with legal defensibility, which additionally turns on explainability, reproducibility, chain-of-custody integrity, and the availability of meaningful challenge mechanisms.

| Jurisdiction  | Enforcement<br>Modality    | Institutional Logic       | Principal<br>Instruments                                | Efficiency–<br>Legitimacy<br>Tradeoff    |
|---------------|----------------------------|---------------------------|---------------------------------------------------------|------------------------------------------|
| United States | Privatized self-regulation | Delegation by contract    | DMCA, Lanham Act, Platform Brand Registries             | High efficiency; low procedural symmetry |
| China         | Administrative technocracy | Bureaucratization by code | Internet Courts, Centralized Blockchain Evidence Regime | High coherence; limited transparency     |

What these divergences reveal is not merely regulatory pluralism but a new *constitutional cartography* of enforcement. Each jurisdiction embeds its theory of legitimacy into code, transforming technological design into a reflection of constitutional identity.

In the European Union, algorithmic enforcement undergoes *constitutionalization by design*. The AI Act and the Digital Services Act instantiate legality as architecture: human oversight, transparency, and proportionality are no longer principles invoked after the fact, but engineering constraints built into systems ex ante. This reflects the European constitutional habit of embedding legitimacy in process — a procedural constitutionalism that transforms the rule of law into the rule of design.

The United States, by contrast, practices *privatization by contract*. Safe-harbor regimes under the DMCA and Section 230 outsource enforcement to market intermediaries under the presumption that competition, litigation, and tort correction will approximate accountability. Here, legitimacy is ex post and market-corrective: due process is secured through adversarial procedure rather than encoded into algorithmic architecture.

In China, we see *bureaucratization by code*. Judicial AI integration, blockchain-evidence protocols, and automated administrative adjudication together construct a technocratic legality rooted in state rationality rather than procedural autonomy. The algorithm does not replace bureaucracy; it extends bureaucratic authority through infrastructure.

These three archetypes—European design, American contract, Chinese code—represent distinct constitutional economies of enforcement. Each balances efficiency and legitimacy along different dimensions: *process versus competition versus control*. Yet their interaction now defines the global enforcement ecosystem. As platforms operate transnationally, their models of due process, transparency, and redress migrate across borders, producing hybrid regimes of algorithmic legality that escape any single jurisdictional logic.

This transnational diffusion invites a deeper inquiry into the *constitutional law of enforcement architecture*. When the same algorithm flags counterfeit listings across multiple legal orders, whose procedural standards apply? When a takedown executed in California removes content visible in Berlin or Shanghai, which conception of due process governs? Algorithmic enforcement thus transforms traditional *forum shopping* into *regime shopping*: actors may select not only favourable forums, but also the most permissive enforcement architectures.

At the multilateral level, institutions such as WIPO and the OECD could serve as normative anchors for this emerging constitutional pluralism. A baseline charter of algorithmic transparency and accountability—analogue to TRIPS but focused on enforcement architectures—could harmonize procedural legitimacy without flattening diversity. Such a framework would not harmonize substantive rights, but the *conditions of their algorithmic enforcement*: transparency by default, auditability ex ante, and contestability across borders.

In this light, trademark law becomes a case study in *constitutional interoperability*: the effort to translate legality across architectures. The challenge ahead is not to achieve a single global standard but to design protocols of mutual recognition—ensuring that algorithmic enforcement developed under one constitutional paradigm can be trusted under another. Only through such

interoperability can legality remain legible in a world where enforcement itself is transnational, data-driven, and designed by code.

Empirically, these cases substantiate the claim that algorithmic trademark enforcement forms an autonomous institutional order. Judicial proceedings test the adaptability of doctrinal standards; platform initiatives reveal the efficiency of privatized enforcement; administrative pilots demonstrate the feasibility—and peril—of automated bureaucracy. The normative challenge for comparative law and economics is to design hybrid mechanisms capable of reconciling efficiency with legitimacy. As subsequent sections argue, reconciliation depends on evidentiary governance—on transforming algorithmic output into accountable legal proof. In short, the frontier of trademark protection now lies not in detection, but in justification.

## 6. Algorithmic Evidence and the Political Economy of Proof

As artificial intelligence becomes interwoven into the machinery of trademark protection, its role in shaping legal knowledge itself becomes impossible to treat as merely instrumental. What we are witnessing is not simply the automation of enforcement, but the *automation of evidentiary authority*. AI systems do not merely assist human investigators—they increasingly define what counts as evidence, who can produce it, and at what cost<sup>54</sup>. This subtle but radical shift relocates the center of gravity of trademark law from human interpretation to computational generation, from ex post narrative to ex ante inference.

In the classical structure of adjudication, evidence performs a dual function: epistemic and economic. It allows the law to know, and it disciplines the cost of knowing. Legal economists from

---

<sup>54</sup> See Allen Waxman, Jason Kort & Marcelo Barros, *Proving Admissibility of AI Outputs Centers on Authenticity*, BLOOMBERG LAW (Feb. 25, 2025), <https://news.bloomberglaw.com/business-and-practice/proving-admissibility-of-ai-outputs-centers-on-authenticity> [<https://perma.cc/3TKN-ED4C>] (clarifying that “a new rule was proposed, Rule 707, which would subject AI-generated outputs to the same requirements as the admissibility of expert testimony, governed by Rule 702. To ensure Rule 702’s requirements are met, the committee contemplates that the courts would examine the inputs used by the AI system, guarantee that the objecting party has adequate access to the AI system to assess its functionality, and determine whether the process has been validated under sufficiently similar circumstances”).

Coase to Kaplow have long emphasized that adjudication is not only a truth-seeking process but a mechanism for allocating information costs among parties. In trademark law, these costs are central: the entire doctrine rests on controlling information asymmetries in the marketplace—ensuring that signals of origin, quality, and authenticity remain reliable enough to sustain efficient exchange. Yet when evidence itself becomes algorithmic, the asymmetry no longer lies between producer and consumer, but between those who control the infrastructure of proof and those who do not.

Algorithmic enforcement, from this vantage, introduces a new *market for credibility*. Platforms and large rightsholders acquire the capacity to generate legally persuasive data through proprietary AI pipelines, while smaller actors and public authorities must purchase access to these same systems or risk procedural inferiority. The very efficiencies that make AI attractive—speed, scale, pattern recognition—translate into institutional asymmetries in evidentiary production. What Coase once described as the optimal internalization of externalities becomes, in this context, the vertical integration of truth-making itself. Enforcement ceases to be a public adjudicative function; it becomes a semi-private production process, optimized for proof rather than persuasion.<sup>55</sup>

The resulting transformation is epistemological as much as legal. Machine learning systems are not witnesses in the traditional sense; they are generators of probabilistic correlations whose outputs often lack interpretability. The question is no longer whether a fact occurred, but whether the model's representation of that fact can be trusted as both accurate and explainable. In doctrinal terms, this reframes evidentiary reliability as a design problem rather than a testimonial one — the architecture of proof becomes as consequential as its content<sup>56</sup>. Procedurally, this represents a shift from *testimonial reliability* to *architectural reliability*—from trusting persons to trusting systems. The

---

<sup>55</sup> When evidence generation is internal to the enforcement process, the distinction between investigation and adjudication blurs. This fusion risks displacing procedural safeguards traditionally attached to the act of proof.

<sup>56</sup> Ljupcho Grozdanovski, *The Explanation One Needs for the Explanation One Gives: The Necessity of Explainable AI (XAI) for Causal Explanations of AI-Related Harm—Deconstructing the “Refuge of Ignorance” in the EU’s AI Liability Regulation 10 n.28* (Feb. 27, 2024), <https://ssrn.com/abstract=4740419>.

evidentiary virtues of honesty and memory are replaced by those of transparency and auditability. Yet these new virtues, unlike the old ones, are contingent on design choices and computational governance. Where those architectures are proprietary, accountability becomes as scarce as access.

To mitigate this opacity, a parallel economy of verification has emerged. Blockchain-based timestamping, cryptographic hashing, and digital evidence-preservation systems—tools once peripheral to IP management—are increasingly used to support the authenticity, integrity, and temporal ordering of digital evidence, although their probative force still depends on jurisdiction-specific rules of admissibility, chain of custody, and judicial review<sup>57</sup>.

Cryptographic timestamping, blockchain-based hash anchoring, and watermarking or provenance tools increasingly support evidentiary claims concerning priority, integrity, and traceability. They should not, however, be treated as self-sufficient guarantees of legal authenticity: absent auditability, contestability, and institutional oversight, they risk converting evidentiary legitimacy into a merely technical assertion.<sup>58</sup>

---

<sup>57</sup> Samuel Rutz, Matthias Hafner, Felix Wüthrich & Beatrix Marosvölgyi, *DLT and the Intellectual Property Ecosystem of Switzerland*, SWISS ECONOMICS 4, (2023), [https://www.swiss-economics.ch/publications.html?file=files/content/dokumente/publikationen/2022\\_RutzHafnerW%C3%BCthrichMarosv%C3%B6lgyi\\_DLT\\_and\\_IP.pdf](https://www.swiss-economics.ch/publications.html?file=files/content/dokumente/publikationen/2022_RutzHafnerW%C3%BCthrichMarosv%C3%B6lgyi_DLT_and_IP.pdf) [<https://perma.cc/Z53F-YRFM>]; Bryan Yip & Delia Horst, *Where Blockchain Meets Intellectual Property: Creating Transformative Chain Reactions*, World Intellectual Property Organization 6 (2025), [https://www.wipo.int/edocs/mdocs/wild/en/wild\\_1/wild\\_1\\_t06\\_2.pdf](https://www.wipo.int/edocs/mdocs/wild/en/wild_1/wild_1_t06_2.pdf) [<https://perma.cc/5XKF-ANT3>]. For evidentiary treatment of electronic timestamping and blockchain-verified evidence, see Regulation (EU) No 910/2014 of the European Parliament and of the Council of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market and repealing Directive 1999/93/EC, OJ L 257, 28.8.2014, pp. 73–114 (providing that electronic time stamps shall not be denied legal effect or admissibility as evidence solely because they are electronic, and that qualified electronic time stamps enjoy a presumption as to date, time, and integrity); Zuigao Renmin Fayuan Guanyu Hulianwang Fayuan Shenli Anjian Ruogan Wenti de Guiding (受高人民法院关于互联网法院审理案件若干问题的规定) [Provisions of the Supreme People's Court on Several Issues Concerning the Hearing of Cases by Internet Courts] (promulgated by Sup. People's Ct., Sep. 6, 2018, effective Sep. 7, 2018) (WestlawChina).

<sup>58</sup> See Regulation (EU) No 910/2014, *supra* note 57, at art. 41, on the legal effect of electronic time stamps, providing that electronic time stamps shall not be denied legal effect or admissibility solely because of their electronic form and that qualified electronic time stamps enjoy a presumption as to the accuracy of the date and time indicated and the integrity of the data to which they are bound; EUROPEAN TELECOMMUNICATIONS STANDARDS INSTITUTE, *ELECTRONIC SIGNATURES AND INFRASTRUCTURES (ESI): POLICY AND SECURITY REQUIREMENTS FOR TRUST SERVICE PROVIDERS ISSUING TIME-STAMPS* (2023); STUART HABER & W. SCOTT STORNETTA, *HOW TO TIME-STAMP A DIGITAL DOCUMENT*, 3 J. CRYPTOLOGY 99 (1991); WORLD INTELLECTUAL PROPERTY ORGANIZATION, *Digital Date-and-Time-Stamping: The Evidentiary Value and Practical Significance of WIPO PROOF* (2020); NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY, *Reducing Risks Posed by Synthetic Content: An Overview of Technical Approaches to Digital Content Transparency* (2024). On blockchain-based judicial evidence, see also Guanyu guifan he jiaqiang rengong zhineng

This tension calls for a normative framework capable of reconnecting technical rationality with the internal morality of law. Lon Fuller’s principles of legality — consistency, publicity, and clarity — must be reimagined for code, transforming engineering constraints into vehicles of legitimacy. In parallel, Teubner’s theory of reflexive law and Sabel and Simon’s experimentalist governance offer institutional templates for federated AI enforcement, where self-regulation is disciplined by transparency and adaptive feedback. Finally, in Rawlsian terms, algorithmic enforcement becomes a site of public reason — a forum where justificatory reciprocity between private code and public norms must be continuously renewed. Together, these strands outline a theory of algorithmic constitutionalism: legality not imposed upon technology, but co-produced through it. In that co-production lies the constitutional fate of the rule of law itself.

Different legal cultures internalize this tension. Germany continues to treat authenticity as a formal precondition to probative value, insisting that digital artifacts bear the marks of qualified certification—an echo of the civil law’s deep trust in procedural ritual as a source of legitimacy<sup>59</sup>. The United States, by contrast, follows a functionalist pragmatism under *Daubert*<sup>60</sup>, transforming

---

zai sifa lingyu yingyong de yijian (关于规范和加强人工智能在司法领域应用的意见) [Opinions on Regulating and Strengthening the Applications of Artificial Intelligence in the Judicial Fields] [(promulgated by Sup. People’s Ct., Dec. 8, 2022, effective Dec. 8, 2022), <https://www.chinajusticeobserver.com/law/x/the-supreme-people-s-court-the-opinions-on-regulating-and-strengthening-the-applications-of-artificial-intelligence-in-the-judicial-field-20221208/enchn> [<https://perma.cc/7F2C-NTS5>].

<sup>59</sup> Oberlandesgericht Frankfurt am Main [OLG Frankfurt Mar. 27, 2025, 16 U 9/23, <https://www.rv.hessenrecht.hessen.de/bshe/document/LARE250000472> [<https://perma.cc/U7ER-6JWQ>]. See also HESSEN ORDENTLICHE GERICHTSBARKEIT, *Hohe Anforderungen an die Prüfung der Zuverlässigkeit einer Quelle* [High demands are placed on the reliability testing of a source], Press Release No. 17/2025, <https://ordentliche-gerichtsbarkeit.hessen.de/presse/hohe-anforderungen-an-die-pruefung-der-zuverlaessigkeit-einer-quelle#:~:text=Die%20Beklagten%20h%C3%A4tten%20nicht%20nachweisen,der%20Zuverl%C3%A4ssigkeit%20und%20Richtigkeit%20der> [<https://perma.cc/CLD8-6GH8>].

<sup>60</sup> The *Daubert* standard, established by the United States Supreme Court in *Daubert v. Merrell Dow Pharmaceuticals, Inc.*, 509 U.S. 579 (1993), governs the admissibility of expert testimony under Rule 702 of the Federal Rules of Evidence. The decision marked a departure from the earlier *Frye* standard, introducing a more rigorous, trial judge-centered gatekeeping role. Under *Daubert*, expert testimony must rest on a reliable foundation and be relevant to the task at hand. The Court outlined several non-exhaustive factors for determining reliability: (i) whether the theory or technique can be and has been tested; (ii) whether it has been subjected to peer review and publication; (iii) its known or potential error rate; (iv) the existence and maintenance of standards controlling its operation; and (v) whether the methodology has gained general acceptance in the relevant scientific community. These criteria were later extended in *Kumho Tire Co. v. Carmichael*, 526 U.S. 137 (1999), to encompass not only scientific knowledge, but also technical and other specialized knowledge. As codified in the amended Rule 702 (2023), courts must find it “more likely than not” that the expert’s opinion is the product of reliable methods reliably applied to sufficient facts or data. In practice, *Daubert* demands that expert evidence—particularly when derived from novel sources such as algorithmic

the court into a marketplace of epistemic claims: reliability is purchased through empirical validation and adversarial testing.<sup>61</sup> China's Internet Courts have formalized a technocratic hybrid, combining blockchain verification with judicial oversight—technology as servant of bureaucracy rather than its substitute.<sup>62</sup> The United Kingdom, characteristically pragmatic, mediates between these poles, grounding admissibility in intelligibility: evidence must not only be accurate, but comprehensible to the fact-finder.<sup>63</sup>

Behind these divergences lies a common problem: *who pays the epistemic price of automation?* Each system assigns the burden differently—civil law to institutional procedure, common law to adversarial testing, hybrid regimes to administrative supervision—but all face the same structural challenge. As AI compresses the marginal cost of data generation, it inflates the fixed cost of interpretability. What once required witnesses now requires engineers. The locus of trust moves from the courtroom to the codebase, from procedure to infrastructure.

---

or AI-generated outputs—demonstrate scientific validity, methodological transparency, and relevance to the legal issue in dispute. Failure to meet these standards may render the evidence inadmissible, regardless of the expert's credentials or experience.

<sup>61</sup> Shane G. Ramsey, *Safeguarding the Courtroom from AI-Generated Evidence: Federal Rule of Evidence 707 Approved by Judicial Conference*, NELSON MULLINS (June 11, 2025), <https://www.nelsonmullins.com/insights/blogs/red-zone/news/safeguarding-the-courtroom-from-ai-generated-evidence-federal-rule-of-evidence-707-approved-by-judicial-conference#:~:text=At%20its%20core%2C%20Rule%20707,The%20rule%20states> [<https://perma.cc/33PB-XT8A>].

<sup>62</sup> See Vivien Chan & Anna Mae Koo, *Blockchain Evidence in Internet Courts in China: The Fast Track for Evidence Collection for Online Disputes*, LEXOLOGY (July 15, 2020), <https://www.lexology.com/library/detail.aspx?g=1631e87b-155a-40b4-a6aa-5260a2e4b9bb#:~:text=In%20September%20of%20the%20same,popular%20short%20video%20streaming%20platforms> [<https://perma.cc/GCR9-GLUS>]; see also Beijing Weibo Shijie Keji Youxian Gongsi Su Baidu Zaixian Wangluo Jishu (Beijing) Youxian Gongsi (北京微播视界科技有限公司与百度在线网络技术北京) 有限公司) [Beijing Weibo Shijie Tech. Co. v. Baidu Online Network Tech. (Beijing) Co.], Jing 0491 Min Chu No. 1 (Beijing Internet Ct. 2018) (China).

<sup>63</sup> While the doctrinal architecture is clear, it must be acknowledged that UK case law on AI-generated evidence in trademark enforcement remains embryonic. The present analysis thus relies on analogical inference from adjacent domains (e.g. regulatory proceedings, criminal forensics, administrative law) and on interpretive projections grounded in procedural norms and judicial reasoning patterns. This is a necessary methodological choice given the rapid technological evolution outpacing formal legal developments. As such, conclusions regarding the likely reception of AI-generated evidence in UK trademark litigation are predictive, not descriptive, and should be read accordingly. See e.g., *R v. Luttrell* [2004] EWCA (Crim) 1344; *Singham v. General Medical Council* [2023] EWHC (Admin) 1349; *R v. Dlugosz* [2013] EWCA (Crim) 2; *R (Bridges) v. Chief Constable of South Wales Police* [2020] EWCA (Civ) 1058; *Lifestyle Equities CV v. Amazon UK Services Ltd* [2024] UKSC 8; DEP'T FOR SCI., INNOVATION & TECH., A PRO INNOVATION APPROACH TO AI REGULATION, 2023, Cm. 815 (UK) (<https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper> [<https://perma.cc/Z3PD-T82B>]).

Trademark law, which has always functioned as the law of the marketplace's *language of trust*, finds itself once again at the center of a broader transformation in the political economy of information. A trademark is, at its core, a signal designed to economize on search and verification costs; it is a promise of reliability in an uncertain world. Algorithmic evidence, paradoxically, reintroduces uncertainty at a higher level: it makes the signal itself dependent on a machine whose logic few can audit. The result is a recursive asymmetry—an information economy built upon informational opacity.

This is why the evidentiary dimension of AI in trademark enforcement cannot be reduced to technical compliance. It implicates the very legitimacy of the enforcement order. The EU AI Act, with its obligations of transparency, documentation, and human oversight, can be read as an attempt to restore the rule of law not over technology, but through technology—to convert procedural fairness into a design constraint.<sup>64</sup> What emerges is a new synthesis of legal and technical rationality: enforcement systems must now not only detect infringement but *explain* themselves.<sup>65</sup> They must generate proofs that are both machine-verifiable and human-meaningful.

In this sense, evidentiary governance becomes a species of constitutional design. Just as trademark law historically balanced private incentives with public information goals, the algorithmic era demands a similar equilibrium—between efficiency and intelligibility, between automation and accountability. The legitimacy of trademark protection will no longer depend solely on the accuracy of detection but on the institutional quality of explanation. Yet explanation is not merely a procedural virtue; it is an infrastructural one. The next section examines how AI's *embedding across the trademark lifecycle* reconfigures this infrastructure—from monitoring to proof—as a continuous system of governance.

---

<sup>64</sup> Maria Lucia Passador, *AI in the Vault: AI Act's Impact on Financial Regulation*, 56 LOY. U. CHI. L. REV. (forthcoming 2025), <https://ssrn.com/abstract=4898828> [<https://perma.cc/6S3P-8TAJ>].

<sup>65</sup> See e.g., Lilian Edwards & Michael Veale, *Slave to the Algorithm? Why a "Right to an Explanation" Is Probably Not the Remedy You Are Looking For*, 16 DUKE L. & TECH. REV. 18 (2017).

## 7. Embedding AI Across the Trademark Lifecycle

The transformative promise of artificial intelligence in trademark enforcement does not lie in the spectacle of new tools, but in the quiet architecture of their integration. The decisive question is not whether AI can detect counterfeits more efficiently, but how its embeddedness across the trademark lifecycle reorganizes the relationship among information, institutions, and incentives. Once inserted into the day-to-day operations of brand surveillance, AI ceases to be a mere instrument of enforcement and becomes a system of *governance by design*.

At the technical frontier, machine-learning infrastructures now cover every link of the enforcement chain. Automated monitoring engines crawl the open web, social media, and gray-market ecosystems, flagging suspect listings through multimodal analysis that fuses computer vision, OCR, and natural-language inference. Reinforcement models continuously refine detection accuracy based on feedback loops from past takedowns, while data-fusion pipelines correlate visual, textual, and behavioral anomalies across platforms. Yet this technological sophistication, while remarkable, is merely the surface expression of a deeper institutional reconfiguration: the conversion of enforcement into an *anticipatory system*.

Historically, trademark enforcement has followed a reactive model. Rights holders observed infringement ex post, initiated notice-and-takedown proceedings, and ultimately sought adjudication. That model reflected a world in which transaction costs of information gathering were high, verification required human expertise, and enforcement was bounded by the capacity of courts. Algorithmic enforcement reverses each of those premises. It reduces search costs to near zero, automates verification through pattern recognition, and substitutes continuous compliance for episodic adjudication. Platforms now embody enforcement within their architectures—a form of self-optimization that enhances efficiency while displacing public oversight.

This evolution is not uniquely “AI-era.” A significant empirical literature on automated notice-and-takedown shows that IP enforcement workflows have long been moving along a

trajectory from manual, case-by-case practice, to boutique service providers, to heavily industrialised systems driven by high-volume automation. Urban, Karaganis and Schofield's work on notice-and-takedown documents how, for some major actors, automated ("bot") systems leave little room for human review, and how enforcement scale reshapes institutional incentives and error management. Although their empirical focus is copyright, the institutional logic maps closely onto contemporary trademark enforcement infrastructures (including platform brand registries and automated listing moderation), where complaint-handling and delisting increasingly operate as routinised, data-driven pipelines.

Once enforcement becomes infrastructural, its logic becomes economic rather than doctrinal. Platforms behave as rational minimizers of reputational risk and regulatory scrutiny. They optimize for throughput and deterrence rather than for the subtleties of legal doctrine. The resulting hybrid—part legal enforcement, part risk management—exemplifies what institutional economists call *private ordering under bounded rationality*. Enforcement norms arise not from adjudicated principles but from adaptive heuristics designed to minimize collective loss.<sup>66</sup> What follows is an emergent equilibrium: the brand, the platform, and the consumer co-produce a state of probabilistic compliance, held together less by legal command than by algorithmic trust.

For luxury markets, this algorithmic trust acquires an additional, paradoxical dimension. In Veblenian terms, the value of exclusivity derives from controlled scarcity and social signaling—from being both recognizable and unattainable. Yet algorithmic enforcement, by maximizing visibility and accessibility, risks eroding the very asymmetry on which luxury signaling rests. The automation of authenticity transforms reputation from a scarce relational good into a quantified metric. Paradoxically, the more perfectly authenticity is secured by code, the less costly—and thus less prestigious—it becomes. Understanding this dynamic requires a fusion of trademark theory with the economics of distinction: the algorithm as both the guardian and the equalizer of luxury

---

<sup>66</sup> Algorithmic heuristics perform the function once reserved for case law—they embody past experience as predictive guidance, thereby converting jurisprudence into computation.

value. This paradox invites a deeper semiotic reading of algorithmic enforcement. In Veblen's classic logic of distinction, scarcity functions not only as an economic condition but as a social signal: exclusivity communicates status through inaccessibility. Bourdieu's notion of symbolic capital refines this insight—authenticity is not a property but a relational performance, sustained by collective belief in difference.

AI-based authentication systems, by optimizing visibility and eradicating uncertainty, threaten to collapse that symbolic differential. When authenticity becomes measurable, it ceases to be aspirational. The code that guarantees genuineness also democratizes it, transforming luxury from an experience of distinction into a commodity of transparency.

In this sense, the algorithm over-produces the very value it was designed to protect. Perfect authenticity devalues itself: a condition of *semiotic inflation*, where the abundance of verified truth undermines the allure of rarity. For luxury brands, therefore, the governance of authenticity is not merely a matter of enforcement efficiency but of curating interpretive ambiguity—preserving the interpretive space in which prestige can still reside. In institutional terms, this dynamic also redefines the internal governance of luxury itself. What begins as brand protection through algorithmic certainty ends as a regime of procedural control: *algorithmic over-authentication transforms luxury governance into compliance governance*. The pursuit of flawless verification, once a mark of exclusivity, becomes a managerial function—aligning the aesthetics of authenticity with the bureaucracies of accountability.

This dynamic echoes George Akerlof's theory of the "market for lemons." Counterfeits function as informational pollutants that erode consumer confidence and thus destroy surplus. In such a setting, automated detection can be seen as an institutional response to an information asymmetry problem: by screening low-quality signals, it sustains the signaling value of the mark. Yet it also generates a new asymmetry—between those who design the screening algorithms and those subject to them. The producer of the algorithm, not the public regulator, now defines what

counts as credible information. Enforcement thereby becomes a private mechanism for restoring informational efficiency, but at the cost of procedural transparency.

Law and economics have long taught that the efficient allocation of enforcement authority depends on comparative institutional advantage. Courts offer legitimacy and procedural fairness but operate with high fixed costs. Platforms offer immediacy and data granularity but lack accountability. AI-mediated enforcement reveals a third hybrid form—what might be called *computational institutionalism*: a regime in which algorithms allocate enforcement effort based on continuous optimization. In such a model, enforcement becomes an output of Bayesian updating.

The doctrinal implications of this institutional shift are far-reaching. Trademark law’s traditional architecture relies on human interpretation of “confusion,” “use,” and “distinctiveness”—concepts saturated with context. Algorithmic enforcement translates them into operational metrics: similarity thresholds, sentiment polarity, virality coefficients. These metrics do not merely measure compliance; they *construct* it. Once a platform defines a confusion score of 0.85 as a trigger for takedown, the threshold itself becomes a legal norm, shaping how infringement is recognized and contested. What was once a matter of judicial interpretation becomes an artifact of design choice.

Economically, the move from reactive to predictive enforcement resembles the shift from ex post liability to ex ante regulation. Just as tort law internalizes harm after the fact, algorithmic enforcement pre-emptively modulates behavior through probabilistic threat. The deterrence effect, however, depends on visibility and perceived fairness. If enforcement appears arbitrary or inscrutable, it may provoke strategic evasion or reputational backlash. Hence, transparency—long treated as a matter of procedural virtue—becomes a variable of economic performance. Enforcement accuracy without explainability undermines deterrence because agents cannot infer the rule being enforced.

This insight leads to a broader normative claim: *explainability is the new efficiency*.<sup>67</sup> In markets where legitimacy drives compliance, the marginal value of interpretability rivals that of speed. The rational firm will prefer a slower but intelligible enforcement system to a fast but unpredictable one,<sup>68</sup> because the former minimizes uncertainty across the supply chain. This logic mirrors the “rule of law premium” in economic development literature: predictability reduces transaction costs even when it limits discretion. Thus, the future of AI in trademark protection will hinge not only on better models, but on better *governance of models*.

Platforms now function as delegated regulators, blending risk management with normative discretion.

From a comparative perspective, different jurisdictions are now crystallizing divergent responses to this privatized enforcement ecology. The European Union, through the AI Act, seeks to constitutionalize transparency and human oversight as *ex ante* obligations for high-risk systems, effectively embedding procedural fairness into design. The United States, in contrast, relies on tort-based after-the-fact correction—delegating primary enforcement to platforms under notice-and-takedown rules and addressing abuse through litigation. China’s approach reflects technocratic pragmatism: blockchain timestamping and automated evidence recognition in its Internet Courts have normalized machine-generated proof within judicial proceedings. These variations illustrate competing models of institutional adaptation—constitutionalization, privatization, and bureaucratic automation—all orbiting around the same gravitational problem: how to preserve legitimacy when enforcement becomes code.

Seen through the lens of institutional economics, these developments suggest that trademark enforcement is evolving into a *multi-layered governance market*. Rights holders, platforms,

---

<sup>67</sup> In this emergent landscape, explainability functions as a new form of efficiency: a governance dividend that converts transparency into predictability and thereby sustains market trust.

<sup>68</sup> Empirical compliance literature supports this claim. Procedural-justice studies show that voluntary compliance depends not only on expected sanctions or decisional speed, but also on whether regulated actors perceive decision-making as fair, predictable, and legitimate. See Glenn D. Walters & P. Colin Bolger, *Procedural Justice Perceptions, Legitimacy Beliefs, and Compliance with the Law: A Meta-Analysis*, 15 J. EXPERIMENTAL CRIMINOLOGY 341 (2019) (reviewing 64 studies and finding significant associations among procedural justice, legitimacy, and compliance).

regulators, and consumers participate in a dynamic game of incentives, where information asymmetry, data access, and algorithmic transparency define payoffs.

Future research must therefore move beyond doctrinal analysis to model these interactions empirically. First, by quantifying the welfare effects of privatized enforcement—does algorithmic efficiency offset the costs of procedural opacity? Second, by examining competitive distortion—do dominant platforms’ enforcement systems entrench market concentration by privileging incumbents with data access? Third, by exploring institutional complementarity—under what conditions can public courts and private algorithms coexist as credible substitutes or complements?

The next frontier lies in constructing a *law-and-economics of algorithmic governance*. Such a framework would integrate signalling theory, principal-agent models, and behavioral insights to explain how algorithmic enforcement affects compliance, innovation, and competition. It would also require comparative inquiry: how do differing legal cultures mediate the trade-off between efficiency and accountability? And normatively, what institutional design—public auditability, algorithmic due process, or collective redress mechanisms—can re-embed democratic oversight in this new ecosystem of code-based law?

Ultimately, the embedding of AI across the trademark protection lifecycle signals the emergence of a new jurisprudential object: enforcement as infrastructure. Legal meaning no longer travels exclusively through courts, statutes, or treaties; it circulates through systems of data classification and risk allocation. To borrow from Bruno Latour and Niklas Luhmann, law is now co-produced with technology—it operates not above the network, but within it. The challenge ahead is to ensure that the rule of law does not become an afterthought in a world where the rule of code is already the rule of the market.

## **8. The AI Counter-Offensive: Weaponised Intelligence in Brand Infringement**

Artificial intelligence, at its core, is an amplifier of cognition; yet amplification is not itself a normative good. It magnifies the purposes, incentives, and institutional constraints embedded in its deployment.<sup>69</sup> The same architectures that extend the reach of legitimate enforcement can also extend the reach of deception. Trademark law thus confronts an inversion of purpose: technologies designed to protect authenticity are increasingly capable of fabricating it. This inversion is not merely technological. It is institutional and economic, reallocating informational advantage among brand owners, intermediaries, and counterfeiters.

From an economic standpoint, counterfeiting has always been an archetype of information asymmetry. Consumers cannot directly observe product quality or provenance; trademarks operate as costly signals that convey reliability and reduce search costs. Coasean analysis long treated enforcement as a transaction-cost problem—an effort to internalize externalities of confusion through credible commitment devices. Yet the rise of generative AI destabilizes this equilibrium. When models can replicate not only the product but the signal of quality itself—its imagery, tone, linguistic style, and cultural connotations—the separating equilibrium that once sustained the informational value of trademarks collapses toward a pooling equilibrium in which genuine and counterfeit signals become indistinguishable. The result is not merely more counterfeiting; it is the erosion of the epistemic infrastructure that makes trademarks economically meaningful. This collapse of informational separation erodes not only consumer trust but also the legal semantics through which authenticity is recognised and enforced.

Generative Adversarial Networks and diffusion architectures have transformed imitation into simulation. StyleGAN-based image generators, initially developed for benign applications in design and medical imaging, can now produce photorealistic advertising materials and “brand

---

<sup>69</sup> See Giulio Corsi, Kyle Kilian & Richard Mallah, *Considerations Influencing Offense-Defense Dynamics from Artificial Intelligence*, ARXIV 2 (2024) (observing that “advanced AI systems are inherently dual-use and generally agnostic to offensive or defensive orientations”). The point is not that AI has an intrinsic offensive or defensive character, but that its effects depend on the purposes, incentives, and governance structures that shape its deployment. In the trademark context, the same capabilities that enable automated detection, authentication, and monitoring may also be repurposed to generate synthetic brand imagery, deceptive product narratives, and adversarially adapted counterfeit signals.

campaigns” indistinguishable from corporate originals.<sup>70</sup> In economic terms, the marginal cost of producing persuasive counterfeit signals has approached zero. The brand’s function as a guarantee of authenticity—the classic Spencean signal—loses credibility when synthetic equivalence can be achieved without incurring reputational cost or production investment. This dynamic undermines the incentive structure on which the law of trademarks rests: if authenticity can be faked cheaply, the premium on genuine reputation diminishes, eroding both deterrence and consumer welfare.

What follows is a shift in enforcement from detection of objects to detection of representations. Modern infringers exploit AI not to copy logos but to replicate patterns of trust: promotional aesthetics, influencer narratives, and the “visual grammar” of a brand. In doctrinal terms, these practices occupy the gray zone between confusion and dilution—deceptive similarity without literal misappropriation.<sup>71</sup> They produce the appearance of legitimacy without direct infringement, challenging the conventional predicates of “use in commerce” and “likelihood of confusion.”<sup>72</sup> The consumer’s perception, once the core of trademark analysis, is now displaced by algorithmic mediation: search engines, recommender systems, and social-media feeds act as the new cognitive interface through which confusion occurs. The average consumer is no longer a

---

<sup>70</sup> Konrad Olsson, *The AI Tools Fashion Is Talking About Right Now* (Mar. 27, 2025), <https://scandinavianmind.com/the-ai-tools-fashion-is-talking-about-right-now/#:~:text=In%20a%20viral%C2%A0Medium%20article%2C%20creative,garments%2C%20and%20stylised%20editorial%20backdrops> [https://perma.cc/25V8-CDLV].

<sup>71</sup> Compare Jabbar Hussain, Magnus B  th & Jonas Ivarsson, *Generative Adversarial Networks in Medical Image Reconstruction: A Systematic Literature Review*, 191 COMPUT. BIOLOGY & MED. 110094 (2025), <https://doi.org/10.1016/j.compbio.2025.110094> [https://perma.cc/6HVG-JM7K] and Jelmer M. Wolterink et al., *Generative Adversarial Networks: A Primer for Radiologists*, 41 RADIOGRAPHICS 3 (2021), <https://pubs.rsna.org/doi/abs/10.1148/rg.2021200151> [https://perma.cc/DK7A-MEPB] with Kuldeep Pal, Rapti Chaudhuri, Suman Deb & Ashim Saha, *Artistic Essence of Generative Adversarial Networks: Analyzing Training Data’s Impact on Performance*, 235 PROCEDIA COMPUT. SCI. 2577 (2024) and Lisa-Marie Mueller, Charalampos Andriotis & Michela Turrin, *Using Generative Adversarial Networks to Create 3D Building Geometries*, in 1 *Data-Driven Intelligence: Proceedings of the 42nd Conference on Education and Research in Computer Aided Architectural Design in Europe (eCAADe 2024)* 479, 479–88 (Odysseas Kontovourkis, Marios C. Phocas & Gabriel Wurzer eds., eCAADe 2024), <https://doi.org/10.52842/conf.ecaade.2024.1.479> [https://perma.cc/SH72-DH5Q] and Mousumi Bose, Lilly Ye & Yiming Zhuang, *How to Create a Fake and Catch the Fake: Generative Adversarial Networks in Marketing*, in *The Impact of Digitalization on Current Marketing Strategies (Marketing & Technology: New Horizons and Challenges)* 39 (Luis Matosas-L  pez ed., Emerald Publ’g Ltd. 2024), <https://doi.org/10.1108/978-1-83753-686-320241003>.

<sup>72</sup> The resulting category—semantic simulation—exposes the limitations of traditional infringement tests that presuppose visual or linguistic copying rather than behavioral mimicry.

legal fiction but a computational agent—an algorithm trained to predict relevance. In this environment, confusion is not psychological but statistical.

The implications for legal design are profound. Trademark law has historically assumed that enforcement would be reactive and ex post—courts adjudicating disputes once harm has crystallized. The digital marketplace reverses that temporality. Enforcement must occur ex ante, through automated filters and AI-driven moderation systems that operate at platform scale. Yet these very systems are now targets of adversarial manipulation. Machine-learning classifiers that detect counterfeit listings are themselves being gamed by adversaries who understand their decision boundaries. The result is a recursive “arms race” between generative and discriminative models: each side training against the other, each iteration narrowing the information differential that once sustained deterrence.

Private enforcement maximizes efficiency but erodes accountability, operating without the feedback loops of public adjudication. In Coasean terms, the internalization of enforcement costs by platforms creates new externalities—opaque decision-making, over-removal, and the concentration of normative authority in market actors whose incentives align with volume and velocity rather than due process.

This privatized architecture invites a comparison to earlier regulatory analogues. In securities markets, algorithmic trading prompted the emergence of “market microstructure” regulation to discipline informational asymmetry among high-frequency actors. Trademark law has yet to develop an equivalent regime. The challenge is not only evidentiary but allocative: determining which institution—court, platform, regulator—bears the marginal cost of error in an environment where both false positives (removing lawful content) and false negatives (allowing counterfeit signals) impose social costs. Law and economics would frame this as a problem of optimal enforcement under uncertainty: the minimization of aggregate error costs when detection itself is endogenous to the strategic behavior of regulated parties.

Comparatively, jurisdictions are beginning to distribute this burden in divergent ways. The European Union's Digital Services Act imposes traceability obligations on platforms, internalizing some externalities of counterfeit trade through due-diligence duties. The United States retains a more laissez-faire, contractual model anchored in notice-and-takedown procedures, effectively privatizing both detection and remedy. China's administrative enforcement system represents the opposite pole, fusing algorithmic surveillance with state-driven market governance. Each approach reflects a different equilibrium between public oversight and private capability—and each economically entails a different trade-off between enforcement cost and legal legitimacy.

As the AI counter-offensive accelerates, the key variable is no longer technological sophistication but institutional design. The normative goal should not be perfect enforcement—a theoretical impossibility—but resilient governance: systems capable of absorbing adversarial shocks without collapsing informational trust. This entails layered safeguards: adversarial testing (“red-teaming”) of enforcement algorithms, secured provenance of training data, and independent auditability of model decisions. Such mechanisms parallel the due process of machines—procedural constraints that preserve contestability and reduce error costs within automated environments. They suggest that due process must now be reconceived as a design parameter—encoded in model governance rather than adjudicated post hoc. In this sense, explainability and transparency are not merely ethical preferences; they are efficiency conditions for sustaining credible signaling in the market for authenticity.

Ultimately, the algorithmic cold war in trademark enforcement reflects a deeper tension within law and economics itself: the balance between private ordering and public legitimacy. As enforcement becomes computational, its efficiency may increase but its normative grounding weakens. The economic logic of deterrence must therefore be complemented by a constitutional logic of accountability. The law's task is not simply to close the asymmetry of information between consumers and counterfeiters, but to close the asymmetry of power between code and law. Only

by embedding procedural rationality into algorithmic enforcement can trademark protection remain both economically efficient and normatively legitimate.<sup>73</sup>

The future of brand protection will not hinge on who possesses the most advanced model, but on who builds the most governable one. In the long run, resilience—the capacity to maintain trust under adversarial conditions—will prove a stronger competitive advantage than speed or secrecy. The market for authenticity, like the market for information, depends not merely on innovation, but on institutions that can make truth auditable again.

## 9. Conclusion

We must look beyond the machinery of enforcement to the architecture of meaning that sustains it. The question is no longer whether artificial intelligence can enhance trademark protection, but what kind of *institutional order* this transformation portends. We are not merely confronting new techniques of surveillance or detection; we are confronting a structural inversion of the trademark system itself. The very instruments designed to secure authenticity, reliability, and trust are now capable of fabricating them. The counterfeit is not a forgery of form but of signal. The infringer is no longer merely a bootlegger; infringement increasingly operates as a feedback loop among adversarial learning, platform incentives, and data abundance.

In this new ecology, infringement does not emerge gradually, nor does it announce itself at the border. It replicates through networks, accelerates through automation, and arrives pre-optimized for invisibility. The adversary is not a person but a process—the emergent by-product of adversarial learning, platform incentives, and data abundance. The result is a world in which authenticity itself becomes a probabilistic claim.

Responding to this evolution demands more than regulatory agility; it demands a rethinking of the *economic and legal foundations* of trademark law. Traditionally, trademark protection has been justified

---

<sup>73</sup> Procedural rationality in this context denotes a system's capacity to justify its outputs by reference to auditable reasoning chains, enabling contestability without halting automation.

on informational grounds: as a mechanism for reducing search costs and mitigating information asymmetry between producers and consumers. In Coasean terms, the trademark system lowered the transaction costs of trust by enabling reputational capital to serve as a substitute for direct inspection. But in a data-driven economy, those informational asymmetries have mutated. Consumers no longer rely primarily on brand signals transmitted through packaging or advertising; they rely on platform metadata—ratings, reviews, recommendation algorithms—whose design lies beyond the reach of trademark doctrine. The “search costs” that once justified protection have been re-engineered by platforms themselves.

If traditional trademark law functioned as an *exogenous guarantee* of information quality, algorithmic commerce now internalizes that function. Enforcement, too, has been absorbed into the platform architecture. What began as a reactive process administered by courts and customs authorities has become a continuous, automated regime of classification and exclusion. Amazon’s Brand Registry, Alibaba’s IPP Platform, and Meta’s content-moderation systems collectively perform what public institutions once did: they decide which signals of authenticity circulate and which do not. In doing so, they have transformed enforcement into a form of delegated regulation—a hybrid of private ordering and administrative law carried out by code.

The next generation of scholarship must therefore bridge doctrinal analysis with institutional economics. How should the law allocate responsibility among brands, platforms, and regulators when detection and adjudication are partially automated? How should courts assess algorithmic “reasonableness” in light of traditional standards like good faith or proportionality? Should platforms be analogized to landlords policing trespass, to utilities providing infrastructure, or to quasi-administrative agencies exercising delegated authority? Each analogy implies a distinct legal regime and a different equilibrium of incentives.

Comparative perspectives reveal that no jurisdiction has yet resolved these tensions. The European Union’s *AI Act* introduces procedural duties—human oversight, explainability, record-keeping—that resemble due-process safeguards in administrative law. The United States, more

reliant on tort analogies and safe-harbor doctrines, allows greater latitude for self-regulation under the Digital Millennium Copyright Act and Section 230. China's Internet Courts, by contrast, have begun to formalize algorithmic enforcement through judicial recognition of blockchain-verified evidence. Each system reflects a distinctive synthesis of efficiency, accountability, and sovereignty—but none provides a complete blueprint for reconciling automation with legality. Beneath these institutional divergences lies a deeper normative fault line—one of epistemic justice. As enforcement becomes increasingly data-driven, the question is no longer merely who enforces, but who owns the infrastructure of truth. The datasets, models, and validation pipelines that define authenticity have become new sites of epistemic power, determining whose perceptions of genuineness are encoded and whose are excluded. When authenticity itself is computationally constructed, control over data curation translates into substantial control over the conditions under which authenticity is recognized. The challenge for law is thus to democratize not only enforcement, but epistemology—to ensure that the capacity to define truth remains a public good rather than a proprietary privilege.

Economically, these differences suggest varying tolerances for *error allocation* and *regulatory externalities*. Where the EU prioritizes ex-ante compliance and transparency, the U.S. model privileges ex-post adjudication and market correction. The comparative law of algorithmic trademark enforcement may thus become a laboratory for testing deeper theories of institutional design: how to balance precision against procedural fairness, speed against legitimacy, and global coordination against national jurisdiction.

The sociological dimension of this shift is equally profound. Legal authority no longer radiates solely from legislatures or courts; it diffuses through the interfaces of platforms, the practices of developers, and the data infrastructures that sustain them. As Luhmann would describe it, the legal system is becoming “autopoietic” with its technical environment, continuously

adapting to informational complexity rather than merely responding to it.<sup>74</sup> Jasanoff's notion of "co-production"<sup>75</sup> captures the same dynamic: law and technology now evolve together, each shaping the epistemic boundaries of the other. Trademark protection, in this sense, is no longer a static set of prohibitions but a living grammar of market communication — one written simultaneously in statutes, in algorithms, and algorithmic systems. To sustain coherence within this living grammar, legal institutions must cultivate what Sunstein calls "incompletely theorised agreements": adaptive frameworks that permit coordination without foreclosure of contestation.

The normative task is therefore not to resist privatized enforcement, but to civilize it. AI systems that govern authenticity must themselves be governed — through transparency mandates, auditability standards, and institutionalized avenues for contestation. The procedural safeguards of due process must migrate into the design of digital enforcement, ensuring that algorithmic decisions can be explained, reviewed, and appealed. This is not a matter of nostalgic humanism; it is a matter of preserving the rule of law under conditions of informational opacity.

Future research must take up this agenda across three intersecting dimensions. First, **economic design**: how to model the optimal allocation of enforcement functions among public and private actors when monitoring costs are endogenous to platform structure. Second, **comparative governance**: how different legal systems operationalize accountability in algorithmic enforcement—from the EU's rights-based model to the U.S. market-corrective model to China's hybrid of judicial and technocratic oversight. Third, **institutional jurisprudence**: how doctrines of evidence, due process, and administrative review can evolve to recognize algorithmic decision-making without surrendering their normative core.

These inquiries converge on a single insight: the legitimacy of trademark protection in the algorithmic age will not turn on the sophistication of its models but on the *architecture of its*

---

<sup>74</sup> NIKLAS LUHMANN, LAW AS A SOCIAL SYSTEM (Fatima Kastner et al. eds., Klaus A. Ziegert trans., Oxford Univ. Press 2004).

<sup>75</sup> STATES OF KNOWLEDGE: THE CO-PRODUCTION OF SCIENCE AND SOCIAL ORDER (Sheila Jasanoff ed., Routledge 2004).

*accountability*. The challenge is not to perfect the machine, but to design the institutions that govern it. Trademark law, at its best, has always been a law of trust—of signals made credible through public norms. The rule of law must learn to inhabit the very architectures it seeks to regulate. The future of trademark protection will hinge on this constitutional migration of legality into design — the rule of law rendered, quite literally, as the rule of code. That vocation endures, but its terrain has shifted. To sustain it, we must treat code not as a technical substrate to be regulated after the fact, but as a constitutional domain in its own right: a site where legality, economics, and technology co-produce the conditions of market order.

If there is a line to be drawn, it is not between law and technology, but between automation and abdication. Law must inhabit the architecture of enforcement rather than chase its consequences.<sup>76</sup> Its task is to bind intelligence to integrity, to tether innovation to legality, and to reaffirm that protection, at its constitutional core, is not a proprietary privilege but a collective covenant safeguarding the informational commons on which both markets and democracies rely.

This analysis yields several implications for the regulatory architecture of AI-enabled trademark enforcement, where the true challenge lies not in efficiency but in legitimacy. Translating legality into design requires that the principles of due process, accountability, and constitutional coherence migrate from doctrine into code. Four propositions follow.

First, Algorithmic Due Process. Algorithmic enforcement must be auditable, explainable, and reviewable by design. Procedural fairness can no longer depend on post hoc adjudication; it must be engineered ex ante into the systems that detect, decide, and delist.

Second, Fiduciary Accountability of Platforms. Platforms and AI vendors operating as de facto regulators should bear fiduciary-like duties of fairness, proportionality, and transparency toward both rightsholders and users. The governance of authenticity is, in essence, a governance of trust.

---

<sup>76</sup> To inhabit enforcement architecture is to embed normative constraints at the design phase—what computer scientists call “governance by architecture”—ensuring that legality is not an ex post audit but an ex ante property of the system.

Third, Constitutionalization by Code. Legality must be embedded within the technical infrastructures of enforcement, ensuring that algorithmic systems internalize rule-of-law principles as operational constraints rather than rhetorical ideals.

Fourth, Institutional Interoperability. As enforcement now unfolds across a lattice of transnational architectures, procedural legitimacy must acquire the grace of portability. The law's authority can no longer end at the border of a jurisdiction if its enforcement travels at the speed of code. Algorithmic determinations must therefore remain contestable across borders and intelligible across constitutional idioms—capable of translating fairness as fluently as they translate data.

Taken together, these propositions form a single imperative: to reimagine trademark protection not as an act of private vigilance, but as a civic architecture of accountability—one in which the rule of law survives precisely because it has learned to encode itself within the grammar of machines.

Like a cipher that reveals its meaning only to those who know where to look, the future of trademark law will depend on our capacity to recognise authenticity even when truth itself is written in the language of algorithms.

## Appendix

### Table 1: The Algorithmic Trademark Enforcement Lifecycle.

Enforcement evolves from data acquisition to normative feedback through an integrated socio-technical architecture. Each phase converts information into institutional order: from the microeconomics of detection to the constitutionalization of design. The lifecycle visualizes how legitimacy migrates from procedure to infrastructure—the rule of law embedded in the rule of code.

**Purpose:** to visualize how enforcement evolves from *data capture* to *normative feedback*, illustrating the co-production of efficiency and legitimacy across technical, institutional, and constitutional layers.

#### I. Data Capture → The Informational Substrate

**Actors:** Platforms, rightsholders, web crawlers, vendors

**Mechanisms:**

- Multimodal scraping (images, text, metadata)
- Blockchain provenance and serial IDs
- Sensor and IoT-based authentication streams

**Economic Function:** Reduces search costs and internalizes discovery costs (Coasean efficiency).

**Legal Tension:** Ownership of raw data; privacy and GDPR constraints; admissibility of machine-collected evidence.

→ *Feeds into automated detection engines.*

#### II. Automated Detection → The Machine Jurisprudence

**Actors:** AI detection engines (CV, NLP, multimodal embeddings), platform enforcement units

**Mechanisms:**

- Pattern recognition (CNNs, transformer embeddings)
- Vendor risk scoring, anomaly detection, semantic mimicry
- Adversarial filtering and clustering (e.g., DBSCAN, CLIP)

**Economic Function:** Converts enforcement from ex post liability to ex ante risk forecasting.

**Legal Tension:** False positives, opacity, “algorithmic paranoia,” procedural fairness.

→ *Triggers automated flagging, delisting, or escalation to human review.*

### **III. Adjudication Layer → Hybrid Governance**

**Actors:** Platforms (as delegated regulators), courts, administrative agencies.

**Mechanisms:**

- Platform-based dispute resolution (Amazon CCU, Alibaba IPP)
- Judicial oversight (Hermès v. Rothschild, CNIPA Internet Courts)
- Regulatory co-supervision (EUIPO AI Pilot, DSA transparency duties)

**Economic Function:** Minimizes transaction costs of enforcement through vertical integration.

**Legal Tension:** Privatization of adjudication; erosion of due process; asymmetry of access to algorithms.

→ *Outcomes feed into evidentiary infrastructures.*

### **IV. Evidentiary Generation → Algorithmic Proof**

**Actors:** AI vendors, blockchain certifiers, rightsholders, forensic auditors.

**Mechanisms:**

- Cryptographic timestamping, digital watermarking
- Model logs and audit trails (explainability metrics)
- Human verification layers for contestability

**Economic Function:** Creates a new market for credibility—algorithmic evidence as an asset.

**Legal Tension:** Reliability, interpretability, admissibility under Daubert and EU AI Act; evidentiary asymmetry.

→ *Feeds verified data into model retraining and governance reporting.*

## **V. Feedback and Normative Adaptation → Institutional Reflexivity**

**Actors:** Legislators, regulators, standards bodies (WIPO, OECD, ISO), academia.

**Mechanisms:**

- Feedback loops from enforcement outcomes to model design
- Transparency reports (DSA Art. 42), third-party audits
- Cross-platform federated learning for adaptive detection


**Economic Function:** Reduces systemic enforcement costs via cooperative learning (Ostrom polycentricity).


**Legal Tension:** Accountability diffusion, transnational interoperability, constitutional legitimacy.

→ *Completes the cycle: data capture standards and design obligations are updated.*

## **Suggested Visual Structure**

**Format Suggestion (linear-to-circular hybrid):**





### The Algorithmic Trademark Enforcement Lifecycle

Each node can include **three concentric rings**:

- Inner ring → *Technical layer* (AI functions)
- Middle ring → *Economic rationale*
- Outer ring → *Legal & institutional implications*

**Figure 2: Comparative Empirical Map of Algorithmic Trademark Enforcement**

| <b>Domain</b>               | <b>Entity / Case</b>                                | <b>Operational Modality</b>               | <b>Core Metrics (2023–24)</b>                                                       | <b>Institutional Logic</b>                     | <b>Efficiency–Legitimacy Frontier</b>                                            |
|-----------------------------|-----------------------------------------------------|-------------------------------------------|-------------------------------------------------------------------------------------|------------------------------------------------|----------------------------------------------------------------------------------|
| <b>Platform Enforcement</b> | <b>Amazon Counterfeit Crimes Unit (CCU)</b>         | AI-driven detection + joint civil actions | >700,000 counterfeit listings removed weekly; >50 joint litigations (US, EU, China) | Privatized quasi-administrative enforcement    | <b>High efficiency,</b> low transparency; due-process gap in automated takedowns |
|                             | <b>Alibaba IPP Platform &amp; Kering MOU (2023)</b> | Federated AI training with brand partners | 95% automated detection accuracy; 30% reduction in manual reviews                   | Co-regulation through platform–brand consortia | <b>Balanced,</b> co-produced legitimacy; procedural opacity persists             |
|                             | <b>Meta Rights Manager</b>                          | ML-based visual/textual cross-            | ~5 billion pieces of content                                                        | Self-regulation via algorithmic moderation     | <b>Extreme efficiency,</b> weak contestabilit                                    |

|                                                  |                                                                       |                                                                   |                                                                            |                                                                |                                                                                                   |
|--------------------------------------------------|-----------------------------------------------------------------------|-------------------------------------------------------------------|----------------------------------------------------------------------------|----------------------------------------------------------------|---------------------------------------------------------------------------------------------------|
|                                                  |                                                                       | platform<br>detection                                             | monitored<br>monthly                                                       |                                                                | y; over-<br>enforcemen<br>t<br>externalities                                                      |
| <b>Luxury<br/>Sector<br/>Collaboratio<br/>ns</b> | <b>LVMH–<br/>Entrupy–<br/>Microsoft<br/>“Authentici<br/>ty Cloud”</b> | Vision<br>models +<br>blockchain<br>provenance                    | 65+ million<br>listings<br>removed in<br>2023; <0.5%<br>error<br>tolerance | Internalized<br>enforcement as<br>compliance<br>infrastructure | <b>Efficient<br/>but<br/>verticalized</b><br>;<br>information<br>al monopoly<br>risk              |
|                                                  | <b>Richemont<br/>Geneva<br/>Digital Hub</b>                           | Cross-brand<br>federated<br>blockchain<br>& AI image<br>forensics | Integrated<br>for Cartier,<br>Montblanc,<br>Van Cleef &<br>Arpels          | Corporate<br>RegTech<br>institutionalizati<br>on               | <b>Moderate<br/>efficiency,</b><br>strong<br>procedural<br>control;<br>high fixed<br>cost barrier |
|                                                  | <b>Chanel–<br/>Tencent<br/>Pilot (2024)</b>                           | Adversarial<br>AI for<br>WeChat<br>counterfeit<br>detection       | 92%<br>detection<br>precision;<br>automated<br>takedown<br>latency <2s     | Platform–brand<br>public-private<br>hybridization              | <b>Near-<br/>optimal<br/>deterrence,</b><br>yet opaque<br>thresholds<br>for<br>“confusion”        |

|                                        |                                                   |                                                     |                                                               |                                |                                                                    |
|----------------------------------------|---------------------------------------------------|-----------------------------------------------------|---------------------------------------------------------------|--------------------------------|--------------------------------------------------------------------|
| <b>Judicial / Administrative Cases</b> | <b>Hermès Int’l v. Rothschild (S.D.N.Y. 2023)</b> | Court adjudication of AI-generated expressive use   | Jury verdict for Hermès; damages awarded                      | Ex post doctrinal adjudication | <b>High legitimacy,</b><br>low adaptability;<br>procedural latency |
|                                        | <b>EUIPO TMView AI Pilot</b>                      | AI-assisted trademark similarity and classification | +20% inter-examiner consistency; 30% time reduction           | Bureaucratic co-regulation     | <b>Legitimate by design,</b><br>moderate efficiency                |
|                                        | <b>CNIPA / Chinese Internet Courts (2022–24)</b>  | Blockchain evidence + judicial AI analytics         | Tens of thousands of cases annually; automated authentication | Technocratic bureaucratization | <b>Centralized efficiency,</b><br>limited transparency             |

Analytical Note: Error-Cost Economics and the Efficiency–Legitimacy Frontier

Algorithmic enforcement transforms trademark protection into a continuous optimization problem under conditions of uncertainty. In Calabresi’s framework of *The Costs of Accidents*, deterrence efficiency is achieved when the marginal cost of additional accuracy equals the marginal social cost of error. In AI-driven enforcement, those errors manifest as false positives (lawful

sellers wrongly delisted) and false negatives (counterfeits undetected). Each generates distinct *legal externalities*:

- False positives shift error costs to small sellers and consumers, undermining procedural legitimacy and competition;
- False negatives erode deterrence and consumer trust, imposing informational costs on markets.

Following Kaplow–Shavell’s model of optimal enforcement, the expected social cost  $C$  of AI enforcement can be expressed as:

$$C = p_{fp} \cdot L_{fp} + p_{fn} \cdot L_{fn} + K$$

where  $p_{fp}$  and  $p_{fn}$  are the probabilities of false positives and false negatives,  $L_{fp}$  and  $L_{fn}$  are their respective loss magnitudes (reputational, competitive, or regulatory), and  $K$  represents the fixed cost of system maintenance and auditability.

Empirically,<sup>77</sup> publicly available brand-protection reports and EU pilot data indicate that marginal efficiency gains flatten once detection accuracy surpasses ~95%, while legitimacy costs rise non-linearly due to opacity and over-removal. This produces an efficiency–legitimacy frontier—a curve along which incremental automation yields diminishing deterrence returns but escalating governance deficits. The optimal enforcement equilibrium lies not at maximum accuracy but at maximal explainability subject to acceptable error variance.

From a law-and-economics perspective, this implies that transparency and auditability are efficiency variables, not merely ethical desiderata. Algorithmic explainability lowers the expected social cost of enforcement by converting uncertainty (error variance) into contestability (reviewable error). In short, legitimacy is an input to efficiency, not its constraint.<sup>78</sup>

---

<sup>77</sup> See Amazon, *Brand Protection Report* (2023); Alibaba Group, *Intellectual Property Protection Annual Report* (2023); EUIPO, *AI in IP Enforcement: Pilot Results* (2023); OECD, *Misuse of AI in Counterfeit Trade* (2022).

<sup>78</sup> This restates, in institutional terms, the core insight of Kaplow–Shavell: legal rules are efficient when they minimize total social costs, including uncertainty and error. In algorithmic enforcement, legitimacy performs that minimization function.