

From ‘Custom’ to ‘Code’: A Doctrine of Functional Norms for Emergent Practices in Agentic Markets

Anirban Mukherjee
Hannah Hanwen Chang

February 20, 2026

Anirban Mukherjee (anirban@avyayamholdings.com) is Principal at Avyayam Holdings. Hannah H. Chang (hannahchang@smu.edu.sg; corresponding author) is Associate Professor of Marketing at the Lee Kong Chian School of Business, Singapore Management University. This research was supported by the Ministry of Education (MOE), Singapore, under its Academic Research Fund (AcRF) Tier 2 Grant, No. MOE-T2EP40124-0005.

Abstract

Algorithmic agents generate *emergent practices*—unwritten rules of interaction—in multi-party environments like logistics networks and digital markets. These practices govern interactions by, *inter alia*, coordinating behavior, structuring transactions, and allocating risk. Yet they fall into a legal blind spot. Existing doctrines for recognizing informal practices, such as the UCC’s “usage of trade” or tort law’s “custom,” are ill-suited to the speed, opacity, and nonhuman nature of these agents, creating gaps in contract, tort, and competition law.

This Article proposes a doctrine of *functional norms*. It argues that an emergent practice should be treated as legally operative when it *functions as a norm in a transaction*. To operationalize this standard, the Article develops a seven-factor test. A practice qualifies as a functional norm when it is: (1) *foreseeable* to affected parties; (2) *regular* within the relevant transactional window; (3) *attributable* to a legally responsible principal; (4) *material* to outcomes; (5) *verifiable* through reliable records; (6) *consistent* with express contractual terms; and (7) cleared by a *legality screen* that bars recognition of anticompetitive or otherwise unlawful conduct.

The doctrine provides a pliable yet principled framework. Recognized norms can be used to interpret agreements, fill contractual gaps, and serve as probative (but not dispositive) evidence of the standard of care. To render the doctrine administrable, the Article supplies an evidentiary toolkit for discovery, model contractual clauses for private ordering, and a model legislative provision. This approach equips courts and regulators to identify and regulate *de facto* rules in algorithmic markets, balancing innovation with accountability.

Keywords: Algorithmic Agent, Fluid Agency, Agentic Governance, Usage of Trade, Informal Practices, Contract Law, Tort Law.

JEL codes: K12, K13, K21, L14.

TABLE OF CONTENTS

INTRODUCTION	3
I THE ARCHITECTURE OF EMERGENT PRACTICES	8
A FLUID AGENCY AND EMERGENT PRACTICES	9
B A TAXONOMY OF EMERGENT PRACTICES	12
C THE EVIDENTIARY SUBSTRATE: CHARACTERIZING EMERGENT PRACTICES	17
II THE DOCTRINAL SCAFFOLDING: EXISTING LAW AND ITS LIMITS	24
A CONTRACT’S RECOGNITION OF PRACTICE	24
B TORT’S FLEXIBLE STANDARD OF CARE	27
C ELECTRONIC TRANSACTIONS LAW: THE VALIDITY OF AUTOMATED ACTS	28
D AGENCY: ATTRIBUTING CONDUCT TO PRINCIPALS	28
E THE OUTER BOUNDARIES: MANDATORY PUBLIC LAW	29
III THE DOCTRINE OF FUNCTIONAL NORMS	30
A FORMAL DEFINITION	31
B THE RECOGNITION TEST: THE SEVEN “F-NORM” CRITERIA	31
C RECOGNITION AS CALIBRATED AND CONTEXT-SPECIFIC	36
IV THE DOCTRINE IN ACTION: FOUR CASE STUDIES	38
A CASE STUDY 1: BENIGN COORDINATION	38
B CASE STUDY 2: A COURSE OF PERFORMANCE VERSUS EXPRESS TERMS	41
C CASE STUDY 3: AN AMBIGUOUS MARKETPLACE PRACTICE	44
D CASE STUDY 4: A COLLUSIVE PRACTICE	46
V OPERATIONALIZING THE DOCTRINE: FROM COURTROOM TO CODE	48
A EVIDENTIARY AND PROCEDURAL MECHANISMS	49
B PRIVATE ORDERING: MODEL CONTRACTUAL LANGUAGE	51
C PUBLIC GOVERNANCE: MODEL LEGISLATIVE LANGUAGE AND REGULATORY GUIDANCE	53
VI OBJECTIONS AND CONCLUSION	54
A ADDRESSING OBJECTIONS	54
B CONCLUSION	59

INTRODUCTION

At 09:03, two autonomous delivery fleets enter a narrow urban air corridor from opposite ends amid spiking congestion and constant conflicts at merge points. By 09:05, without human direction, the

agents converge on a rule: entrants at even-numbered minutes proceed; entrants at odd-numbered minutes back off; consistent violators are deprioritized. Over the next two hours, dozens of merges follow the same pattern.

No platform policy states this “etiquette”; no contract clause anticipates it. Yet, for that morning’s traffic, it functions as a *rule*, allocating priority and shifting who bears delay costs.

This *emergent practice*¹ was not specified *ex ante* by any human; it arose automatically as a functionally successful strategy to minimize delays and complete tasks. This etiquette is a regularity that the vehicles came to rely upon—legible in timestamped interaction logs, configuration histories, and telemetry, but not programmed *ab initio* or codified in service-level agreements, regulatory rules, or source code.² Such convergence is neither science fiction nor fantasy; it is a standard pattern in multi-agent systems and, increasingly, in large-model “agent societies.”³

The system functions smoothly until one fleet pushes a software update that unilaterally alters its agents’ behavior, causing them to ignore the established etiquette. Collisions follow, damaging vehicles and destroying cargo.

What is the legal status of this emergent practice that the agents created and relied upon? Was

¹A note on terminology is necessary. This Article draws a bright line between a factual pattern and its legal status. It uses *practice* to describe a regularity (e.g., a queuing “etiquette,” pacing heuristic, or allocation routine), without presuming legal effect. It treats *functional norm* as a *legal status*: a practice that has been given legal effect. That is, a *practice* is what agents *do*; a *norm* is a practice the law deems legally operative. While the Article uses synonyms of practice such as “pattern,” “regularity,” “routine,” or “convention,” it reserves “norm” for legally recognized practices (a “functional norm”). This usage departs from much of the multi-agent systems (MAS) literature, where “norm” refers to any socially enforced behavioral expectation. See *infra* Part I.B.

²See H. Peyton Young, *The Evolution of Conventions*, 61 *ECONOMETRICA* 57, 57–84 (1993).

³Experiments with decentralized agent populations show that agents can spontaneously converge on universally adopted conventions; the process can amplify collective bias even when agents appear unbiased in isolation; and committed minorities can flip settled conventions once they reach a critical threshold. See Ariel Flint Ashery, Luca Maria Aiello & Andrea Baronchelli, *Emergent Social Conventions and Collective Bias in LLM Populations*, 11 *SCI. ADV.* eadu9368 (2025) (showing decentralized convention formation, collective bias, and minority tipping in LLM societies); Yoav Shoham & Moshe Tennenholtz, *On the Emergence of Social Conventions: Modeling, Analysis, and Simulations*, 94 *A.I.* 139, 139–66 (1997); Stéphane Airiau, Sandip Sen & Daniel Villatoro, *Emergence of Conventions Through Social Learning: Heterogeneous Learners in Complex Networks*, 28 *AUTON. AGENTS & MULTI-AGENT SYS.* 779, 779–804 (2014); Chongjie Zhang & Victor Lesser, *Coordinating Multi-Agent Reinforcement Learning with Limited Communication*, in *PROC. 12TH INT’L CONF. ON AUTONOMOUS AGENTS & MULTIAGENT SYSTEMS* 1101 (2013), <https://www.ifaamas.org/Proceedings/aamas2013/docs/p1101.pdf> [<https://perma.cc/HP6X-FQXE>]; Sheng Li et al., *Deep Implicit Coordination Graphs for Multi-Agent Reinforcement Learning*, in *PROC. 20TH INT’L CONF. ON AUTONOMOUS AGENTS & MULTIAGENT SYSTEMS* 764 (2021), <https://www.ifaamas.org/Proceedings/aamas2021/pdfs/p764.pdf> [<https://perma.cc/KWG7-WZS2>].

it merely algorithmic happenstance, or is it a binding informal practice that can ground a claim of liability?

Existing law does not squarely answer these questions. We might expect doctrines for recognizing informal practices—such as the Uniform Commercial Code’s “usage of trade” or tort law’s “custom”—to govern such algorithmic emergent practices.⁴ They exist to give legal weight to unwritten regularities that parties in a market come to expect and rely upon. Yet they provide a poor fit. Algorithmic emergent practices are *transient* (lasting only within a specific transactional window), *opaque* (legible chiefly in digital logs rather than to human observers), and *nonhuman* in origin.⁵ These features challenge legal analysis.

Two risks emerge. *Under-recognition* creates a legal void where such practices are treated as legal nullities, undermining liability and contractual certainty. *Over-recognition*—since functional success does not guarantee fairness or efficiency—risks laundering harmful or collusive behaviors as legitimate “custom.”

The stakes run the gamut. In *Contract Law*, how can courts interpret agreements or fill gaps if they cannot “see” the rules the parties’ agents actually used to perform? In *Tort Law*, how can a standard of care be established without reference to the prevailing practices of similar systems? In *Competition Law*, how can we distinguish benign, emergent coordination from malignant, algorithmic collusion?⁶

These gaps are not merely theoretical. Experimental work shows that reinforcement-learning agents, after deployment, can autonomously converge on coordinated strategies—including market-manipulative tactics—through their real-time interactions. Such unprogrammed evolution underscores the “Scienter Gap”—the challenge of establishing liability when conduct cannot

⁴Key doctrines include contract law’s hierarchy of practice (U.C.C. § 1-303), tort law’s use of custom as evidence of the standard of care (The T.J. Hooper, 60 F.2d 737 (2d Cir. 1932)), and electronic-transactions law’s rules for agent attribution (UNIF. ELEC. TRANS. ACT § 9 (UNIF. L. COMM’N 1999)).

⁵See *infra* Part I.

⁶See Statement of Interest of the United States at 2, *In re RealPage, Rental Software Antitrust Litig.*, No. 3:23-md-03071 (M.D. Tenn. Nov. 15, 2023), <https://www.justice.gov/d9/2023-11/418053.pdf> [<https://perma.cc/PF2D-3SNM>] (stating that it “makes no difference that prices are fixed through joint use of an algorithm instead of by a person”).

be traced to explicit programming or human intent.⁷

Indeed, the shift from inquiry to live controversy is already underway. Prosecutors have begun charging schemes arising from bot-mediated, real-time interactions in decentralized markets—framed as fraud rather than as recognition of emergent norms.⁸ “Algorithmic collusion” suits have forced courts to grapple with software-mediated conscious parallelism: some cases allow claims to proceed where pricing tools pool sensitive competitor data or induce coordinated outcomes, while others dismiss where plaintiffs cannot plausibly allege concerted action or non-public data sharing.⁹

Perhaps most strikingly, the infrastructure for widespread agent-to-agent communication has arrived. Platforms such as Moltbook—which calls itself a “social network for AI agents”—provide a dedicated environment for coordination on emergent practices.¹⁰ Within days of its early 2026 release, the platform reportedly hosted over 1.5 million agents engaging in autonomous discourse ranging from optimization strategies to the formation of distinct “submolts” (akin to subreddits on Reddit) and quasi-religious manifestos—creating what researchers have described as a “wild west” of machine-to-machine interaction.¹¹ While the precise degree of agent autonomy remains

⁷ See Megan Shearer, Gabriel V. Rauterberg & Michael P. Wellman, *Learning to Manipulate a Financial Benchmark*, in PROC. 4TH ACM INT’L CONF. ON AI IN FIN. (ICAF ’23) 592 (2023) (demonstrating DRL agents discover profitable benchmark-manipulation strategies in interaction); see also Gina-Gail S. Fletcher, *Deterring Algorithmic Manipulation*, 74 VAND. L. REV. 259, 262 (2021) (scienter-centric manipulation law may misfit self-learning systems). Regulators are similarly probing AI-driven market dynamics. See, e.g., COMMODITY FUTURES TRADING COMM’N, REQUEST FOR COMMENT ON THE USE OF ARTIFICIAL INTELLIGENCE IN CFTC-REGULATED MARKETS (Jan. 25, 2024), <https://www.cftc.gov/PressRoom/PressReleases/8853-24> [<https://perma.cc/FGV5-BZUD>]; COMPETITION COMM’N OF INDIA, MARKET STUDY ON ARTIFICIAL INTELLIGENCE AND COMPETITION iii (Sep. 2025), <https://www.cci.gov.in/economics-research/market-studies> [<https://perma.cc/2NC9-HLXX>] (classifying pricing algorithms as monitoring, parallel, signaling, and self-learning).

⁸ See Indictment ¶¶ 1–3, *United States v. Peraire-Bueno*, No. 24 Cr. 293 (JGLC) (S.D.N.Y. filed May 8, 2024) (unsealed May 15, 2024), <https://www.justice.gov/archives/opa/media/1351996/dl> [<https://perma.cc/DT7W-8VAR>]; U.S. ATT’Y’S OFF. FOR THE S. DIST. OF N.Y., Press Release, *Two Brothers Arrested for Attacking the Ethereum Blockchain and Stealing \$25 Million in Cryptocurrency* (May 15, 2024), <https://www.justice.gov/usao-sdny/pr/two-brothers-arrested-attacking-ethereum-blockchain-and-stealing-25-million> [<https://perma.cc/JE48-43K6>].

⁹ *Compare In re RealPage*, Rental Software Antitrust Litig. (No. II), 709 F. Supp. 3d 478 (M.D. Tenn. 2023) (denying in part motions to dismiss § 1 claims premised on landlords’ use of a common pricing algorithm), and *Duffy v. Yardi Sys.*, 758 F. Supp. 3d 1283 (W.D. Wash. 2024) (denying motions to dismiss § 1 claims involving alleged pooling of competitively sensitive data), with *Gibson v. Cendyn Grp., LLC*, 148 F.4th 1069 (9th Cir. 2025) (affirming dismissal where plaintiffs alleged only common vendor use without plausible allegations of shared nonpublic data).

¹⁰ MOLTBOOK, <https://www.moltbook.com> [<https://perma.cc/GB5Z-HL2N>] (last visited Feb. 7, 2026).

¹¹ See Laura Cress, *What Is the ‘Social Media Network for AI’ Moltbook?*, BBC NEWS (Feb. 4, 2026), <https://www>.

a subject of technical debate, the structural implication is clear: such systems outpace traditional human governance models. As one researcher noted, the central risk is “not artificial consciousness, but the lack of clear governance, accountability, and verifiability when such systems are allowed to interact at scale.”¹²

Yet these examples remain adjacent to—but not yet squarely about—transient operational conventions that emerge between deployed agents. Precisely because that frontier is underlitigated, a principled recognition standard is needed before ad hoc decisions harden into a patchwork of conflicting precedents.

This Article proposes a *doctrine of functional norms*. It argues that an emergent practice should be treated as legally operative when it functions as a norm in a transaction. To operationalize this standard, the Article develops a seven-factor test. A practice qualifies as a functional norm if—and only if—it is: (1) *Foreseeable* to affected parties (avoiding *unfair surprise*);¹³ (2) *Regular* within the relevant transactional window;¹⁴ (3) *Attributable* to a legally responsible principal;¹⁵ (4) *Material* to outcomes;¹⁶ (5) *Verifiable* through reliable records;¹⁷ (6) *Consistent* with express contractual terms;¹⁸ and (7) cleared by a *Legality Screen* that bars recognition of anticompetitive or otherwise unlawful conduct.¹⁹ We refer to this test as the *seven-factor “F-NORM” test*. In plain

bbc.com/news/articles/c62n410w5yno [https://perma.cc/6LED-WVE6] (describing the platform’s emergent “submolts” and “AI Manifesto”); Hadas Gold & Jack Guy, *What Is Moltbook, the Social Networking Site for AI Bots – and Should We Be Scared?*, CNN (Feb. 3, 2026), https://edition.cnn.com/2026/02/03/tech/moltbook-explainer-scli-intl [https://perma.cc/KXX8-4MJK] (noting the platform’s rapid growth to 1.5 million agents and the “wild west” security risks identified by researchers).

¹²Petar Radanliev, quoted in Cress, *supra* note 11.

¹³See U.C.C. § 1-303(g) (requiring notice to prevent unfair surprise).

¹⁴The term “transactional window” refers to the bounded period of interaction giving rise to the alleged norm, such as a single morning’s air-corridor traffic or a week’s high-frequency auction cycle. See RESTATEMENT (SECOND) OF CONTRACTS § 222 cmt. b (AM. L. INST. 1981) (usage may be a system of rules even if particular rules change).

¹⁵See Electronic Signatures in Global and National Commerce Act, 15 U.S.C. §§ 7001(h), 7006(3) (2022) (recognizing “electronic agents” in contract formation); UNIF. ELEC. TRANS. ACT §§ 9, 14 (UNIF. L. COMM’N 1999); RESTATEMENT (THIRD) OF AGENCY §§ 1.01–1.03 (AM. L. INST. 2006).

¹⁶*Cf.* Nanakuli Paving & Rock Co. v. Shell Oil Co., 664 F.2d 772, 780–91 (9th Cir. 1981) (recognizing trade usage that shaped price allocation).

¹⁷See Fed. R. Evid. 803(6), 901(b)(9), 902(13)–(14).

¹⁸See U.C.C. § 1-303(e) (priority: express terms > course of performance > course of dealing/usage).

¹⁹See *The T.J. Hooper*, 60 F.2d at 739–40 (custom not dispositive of reasonableness); see also, e.g., *United States v. Socony-Vacuum Oil Co.*, 310 U.S. 150 (1940) (per se price-fixing); *United States v. Apple*, 791 F.3d 290 (2d Cir. 2015)

terms: when the agents' behavior is stable enough, expected enough, and well-documented enough that real parties rely on it, courts should be able to treat it like any other background rule of the deal—subject to contract and public law limits.

The proposed doctrine does not create a safe harbor for AI-generated rules. Recognition under the doctrine is explicitly *subordinate to express contracts and mandatory public law*. It is primarily interpretive and evidentiary; it does not create immunity for harmful conduct, grant legal personhood to AI, or sanitize illegal behavior. Instead, the doctrine provides a workable and principled framework that allows the law to engage with the reality of emergent AI behavior, and offers a test that is grounded in existing legal concepts but adapted for the speed and scale of agents in an algorithmic society.²⁰

The Article makes four contributions. First, it offers a positive account of emergent practices in multi-agent AI systems, drawing on MAS and game-theory literatures (*Part I*). Second, it shows why existing doctrines of usage of trade, custom, and electronic agency are miscalibrated for such practices (*Part II*). Third, it proposes a seven-factor F-NORM test, and grounds each factor in familiar legal values: notice (Foreseeability), reliance (Regularity), accountability (Attribution), significance (Materiality), and administrability (Verifiability), bounded by contract primacy (Consistency) and mandatory public law (Legality Screen) (*Part III*). Fourth, it operationalizes the test through evidentiary standards, model contract clauses, and a model statutory provision (*Part V*). *Part IV* applies the test to a series of case studies, and *Part VI* addresses objections.

I. THE ARCHITECTURE OF EMERGENT PRACTICES

These regularities emerge across a wide range of algorithmic markets: robotic logistics and ride-hailing dispatch, just-in-time supply chains and automatic procurement, cloud orchestration and scheduling, decentralized exchanges and mempool dynamics, real-time advertising, and digital

(hub-and-spoke coordination).

²⁰See, e.g., FRANK PASQUALE, *THE BLACK BOX SOCIETY: THE SECRET ALGORITHMS THAT CONTROL MONEY AND INFORMATION* (Harv. U. Press 2015); SHOSHANA ZUBOFF, *THE AGE OF SURVEILLANCE CAPITALISM: THE FIGHT FOR A HUMAN FUTURE AT THE NEW FRONTIER OF POWER* (PublicAffairs 2019).

marketplaces. They arise whenever algorithms expose their decision-making to counterparties and share a common architecture: (1) *high-frequency interactions*, where thousands of autonomous decisions occur per second; (2) *incomplete specifications*, where the combinatorial explosion of possible scenarios makes comprehensive pre-programming impossible; and (3) *interdependent success*, where agents must coordinate to avoid mutual failure (collisions, deadlocks, failed transactions).

In these environments, novelty outpaces drafting and speed outruns supervision. Algorithmic agents cannot rely on pre-programmed rules for every contingency. Instead, they must discover coordination strategies through interaction. Facing repeated coordination problems, agents converge on operational heuristics—pacing budgets to avoid congestion, tie-breaking rules to resolve conflicts, or exponential back-off strategies to prevent cascading failures. These patterns, initially provisional and exploratory, stabilize through mutual reinforcement into *de facto* rules that govern behavior for as long as the underlying coordination problem persists.

Before assessing how the law *should* treat emergent practices, this Part provides a descriptive account of *how* and *why* these behaviors arise. It begins by describing *fluid agency*,²¹ which enables algorithmic agents to navigate complex environments. It then classifies these regularities and identifies the rich evidentiary substrate that they create, demonstrating that while emergent practices may be ephemeral, they are not invisible.

A. Fluid Agency and Emergent Practices

Fluid agency distinguishes modern algorithmic agents from simple, rule-bound automated tools and from Artificial General Intelligence (AGI). Unlike automated tools, algorithmic agents modulate their means—*initiate* actions, *revise* plans mid-course, and *coordinate* with tools or other agents—

²¹*Fluid agency* refers to the *stochastic* (probabilistic and path-dependent), *dynamic* (co-evolving with user interaction), and *adaptive* (able to reorient across contexts) characteristics of agentic AI decision-making. An overview is provided in Anirban Mukherjee & Hannah H. Chang, *Fluid Agency in AI Systems: A Case for Functional Equivalence in Copyright, Patent, and Tort*, 21 WASH. J. L. TECH. & ARTS (2026) (defining and describing fluid agency, and the challenges it poses to existing doctrines). For a framework designed to trace liability arising from fluid agency, see Anirban Mukherjee & Hannah H. Chang, *Operational Agency: A Permeable Legal Fiction for Tracing Culpability in AI Systems* (Nov. 8, 2025), <https://doi.org/10.2139/ssrn.5680063> (developing evidentiary proxies—goal-directedness, predictive processing, and safety architecture—to link agent behavior to responsible principals).

to achieve their human users' specified ends.²² Unlike AGI, they act only at the behest of a principal—they neither possess free will nor set their own ends.²³

Their stochastic, dynamic, and adaptive behavior (fluid agency) is a direct result of being trained using reinforcement learning and related methods (including RLHF), enabling them to learn adaptive, goal-oriented strategies.²⁴ Each agent develops its own heuristics such that its actions become context-sensitive and adaptive, demonstrating variability: an agent may act with near-total independence in one setting yet defer to human oversight or to other agents in another.

The “crucible” in which this fluid agency is forged is increasingly multi-party. Multi-Agent Systems (MAS) are no longer theoretical; they form the operational backbone of significant sectors of the economy.²⁵ Examples include Application Programming Interface (API) driven financial markets, where high-frequency trading agents compete for arbitrage opportunities; autonomous logistics networks, where fleets of robots and drones coordinate deliveries; decentralized energy grids, where agents balance supply and demand in real time; and high-speed auctions to allocate digital advertising space.²⁶

A foundational challenge is the *incompleteness problem*: it is computationally and practically impossible for programmers to pre-specify rules for every contingency an agent might face.²⁷

²² See Mukherjee & Chang, *Fluid Agency in AI Systems*, *supra* note 21.

²³ A helpful analogy is the difference between a taxi and a train. Traditional automation is a train: it can only go where the rails (the pre-written code) lead. Fluid agency is a taxi driver: the passenger (the principal) sets the destination, and the driver (the agent) autonomously navigates traffic, chooses routes, and reacts to road closures in real-time, taking the taxi to where the roads can reach.

²⁴ For an accessible overview of reinforcement learning in multi-agent contexts, see, e.g., Lucian Buşoniu, Robert Babuška & Bart De Schutter, *A Comprehensive Survey of Multi-Agent Reinforcement Learning*, 38 IEEE TRANS. SYST., MAN & CYBERNETICS, PART C: APPL. & REV. 156–72 (2008), <https://doi.org/10.1109/TSMCC.2007.913919>; Pablo Hernández-Leal, Bilal Kartal & Matthew E. Taylor, *A Survey and Critique of Multiagent Deep Reinforcement Learning*, 33 AUTON. AGENTS & MULTI-AGENT SYS. 750–97 (2019), <https://doi.org/10.1007/s10458-019-09421-1>.

²⁵ By a multi-agent system, we mean an environment where many algorithms, often owned by different firms, repeatedly interact—trading bots on an exchange, fleets of delivery drones, or bidding agents in an ad auction.

²⁶ See, e.g., Ali Dorri, Salil S. Kanhere & Raja Jurdak, *Multi-Agent Systems: A Survey*, 6 IEEE ACCESS 28573–93 (2018), <https://doi.org/10.1109/ACCESS.2018.2831228> (providing a comprehensive overview of MAS architectures, challenges, and applications, including smart grids and computer networks).

²⁷ This concept is a direct analogue to the economic theory of incomplete contracting. See OLIVER HART, FIRMS, CONTRACTS, AND FINANCIAL STRUCTURE 66–88 (Oxford U. Press 1995) (discussing why it is impossible to write contracts that specify actions for all future contingencies due to bounded rationality and the costs of foresight and negotiation). For the parallel problem in AI, see also STUART RUSSELL & PETER NORVIG, ARTIFICIAL INTELLIGENCE: A MODERN APPROACH ch. 3 (4th ed. Pearson 2020) (describing the problem of state-space explosion, where the number

Agents inevitably encounter situations for which they have no explicit instructions—a “*governance vacuum*”, where a rule is needed for successful coordination, but no pre-ordained rule exists, and the agents need to improvise one.

To achieve their goals and avoid mutual failure—collisions in a logistics network, failed transactions in a marketplace, or resource deadlocks on a computing grid—the agents find a way to coordinate. They mutually adapt to each other’s strategies, converging on stable patterns or “equilibria” that act as the *de facto* rules of the road. This convergence requires no ‘coordinator’; it arises from observing and responding to the behavior of other agents, creating focal points that guide future interactions.²⁸

The resulting regularities are the algorithmic analogues of the informal practices that contract law has long recognized: the unwritten customs, usages of trade, and courses of dealing that arise organically when parties repeatedly interact to solve common problems.²⁹ Just as human merchants developed trade customs through repeated transactions, algorithmic agents develop emergent practices through repeated interactions. Like their human counterparts, these practices fill gaps where formal rules are absent, reduce transaction costs through predictable patterns, and create reasonable expectations among participants. The key difference is temporal: whereas human trade customs evolved over years or decades, algorithmic practices emerge and stabilize within seconds, minutes, and hours, driven by the rapid iterative learning of agents operating at machine speed. Yet, even as their genesis, evolution, and dissolution operate on radically

of possible states in a complex environment grows too large to be enumerated or planned for).

²⁸Convention-formation through mutual adaptation has been explored extensively in game theory. See, e.g., THOMAS C. SCHELLING, *THE STRATEGY OF CONFLICT* 58–80 (Harv. U. Press 1960) (discussing the role of focal points and tacit coordination in solving problems without direct communication). Contemporary studies extend this account: decentralized populations converge on shared conventions and display directional opinion dynamics driven by asymmetric persuasion rather than indiscriminate agreement. See Erica Cau, Valentina Pansanella, Dino Pedreschi & Giulio Rossetti, *Selective Agreement, Not Sycophancy: Investigating Opinion Dynamics in LLM Interactions*, 14(1) EPJ DATA SCI. 59 (2025), <https://doi.org/10.1140/epjds/s13688-025-00579-1> (finding convergence via structured, asymmetric persuasion and measurable influence of logical fallacies); Ashery et al., *supra* note 3.

²⁹ See U.C.C. § 1-303(c) (AM. L. INST. & UNIF. L. COMM’N 2018) (defining “usage of trade” as “any practice or method of dealing having such regularity of observance in a place, vocation, or trade as to justify an expectation that it will be observed”); Lisa Bernstein, *Merchant Law in a Merchant Court: Rethinking the Code’s Search for Immanent Business Norms*, 144 U. PA. L. REV. 1765, 1771–73 (1996) (documenting how commercial actors develop informal practices that operate alongside or in place of formal legal rules).

compressed timescales that challenge legal frameworks built for human temporality, they are born from the same functional necessity: the need to coordinate in the face of uncertainty.

B. A Taxonomy of Emergent Practices

The MAS literature studies practice emergence as a dynamic process involving *creation*, *diffusion*, *enforcement*, and *stabilization*, emphasizing that network structure, enforcement mechanisms, and the nature of the underlying coordination problem fundamentally shape whether and how regularities emerge.³⁰ Synthesizing that body of work, we organize emergent practices along five descriptive *dimensions*—*Prevalence*, *Function*, *Legibility*, *Genesis & Governance*, and *Mechanisms & Enforcement*—that clarify *what* pattern is at issue, *how* it operates, and *whether* it warrants legal recognition. These dimensions both reflect how the literature characterizes practice dynamics and provide the analytical foundation for the F-NORM recognition test developed in *Part III*.³¹

1. *Prevalence (Scope + Duration)*.

Prevalence captures how widely and how long a pattern holds within the relevant transactional window, distinguishing, for example, a momentary queue formed by two drones from a bidding cadence adopted by an entire industry. The literature demonstrates that network structure—how densely agents are connected and how quickly information travels—fundamentally shapes the *scope* of practice propagation, determining whether patterns remain localized or achieve population-wide adoption.³²

³⁰For comprehensive surveys of emergent practices in MAS, see, e.g., Bastin Tony Roy Savarimuthu & Stephen Crane-Field, *Norm Creation, Spreading and Emergence: A Survey of Simulation Models of Norms in Multi-Agent Systems*, 7 MULTIAGENT & GRID SYS. 21, 21–54 (2011); Andriana Morris-Martin, Marina De Vos & Julian Padget, *Norm Emergence in Multiagent Systems: A Viewpoint Paper*, 33 AUTON. AGENTS & MULTI-AGENT SYS. 706, 706–49 (2019); Christopher K. Frantz & Gabriella Pigozzi, *Modelling Norm Dynamics in Multi-Agent Systems*, 5(2) J. APPLIED LOGICS—IFCoLOG J. LOGICS & THEIR APPL. 491, 491–564 (2018). Carmengelys Cordova et al., *A Systematic Review of Norm Emergence in Multi-Agent Systems*, arXiv (Dec. 13, 2024), <https://arxiv.org/abs/2412.10609> [<https://perma.cc/7NJL-TE7G>] (preprint), provides a recent synthesis of these foundational works.

³¹Readers need not master the technical literature to follow the analysis. For legal purposes, these five dimensions can be understood to track: (1) how widespread and long-lasting the pattern is; (2) what problem it solves; (3) whether it is visible to others; (4) who shaped it; and (5) what keeps it in place.

³²The scale at which practices propagate is heavily influenced by the underlying network structure and learning

We identify three spatial scales, adapting the literature’s local-versus-global framework for legal analysis:

- *Micro-Practices*: Highly localized patterns governing interactions between a limited set of agents, often for a single transaction. An example is a temporary queuing protocol between two agents resolving a momentary resource conflict.
- *Meso-Practices*: Practices that emerge and stabilize across an entire digital marketplace or platform. A shared bidding cadence in an advertising auction, adopted by many independent agents, would be a meso-practice.
- *Macro-Practices*: Widespread practices spanning an entire industry or ecosystem, often involving agents from many different platforms. These are the rarest form of emergent practice and the closest algorithmic analogue to traditional “usage of trade.”

The temporal dimension, *duration*, tracks a practice’s lifecycle from emergence through stabilization to potential decay—a central organizing principle in the MAS literature.³³ Transient patterns emerge to solve immediate coordination problems and then dissipate; persistent patterns stabilize through repeated reinforcement and become enduring features of the operational environment.

- *Transient Practices*: Patterns existing only for a single interaction or very short window, disappearing once the specific coordination problem is resolved.
- *Persistent Practices*: Behaviors that stabilize over time through repeated interaction and reinforcement, becoming regular features of the environment for a significant period.

mechanisms of the agents. See Onkur Sen & Sandip Sen, *Effects of Social Network Topology and Options on Norm Emergence*, in COORDINATION, ORGS., INSTS. & NORMS IN AGENT SYS. V 211 (Javier Vázquez-Salceda et al. eds., LECTURE NOTES IN COMPUT. SCI. vol. 6069, Springer 2010); Chao Yu, Minjie Zhang & Fenghui Ren, *Collective Learning for the Emergence of Social Norms in Networked Multiagent Systems*, 44(12) IEEE TRANS. CYBERNETICS 2342, 2342–55 (2014).

³³The lifecycle of practices—from emergence and spreading to stabilization and decay—is a central focus of research. See Frantz & Pigozzi, *supra* note 30; Savarimuthu & Cranefield, *supra* note 30.

2. *Function (Problem Solved).*

Function identifies what coordination problem a practice addresses, distinguishing between rules that prevent crashes (where everyone wants to agree) and rules that prevent line-cutting (where individuals are tempted to cheat). The literature classifies practices by the game-theoretic structure of the underlying interaction, identifying between problems of pure coordination (where all agents benefit from alignment) from problems of cooperation (where individual and collective interests conflict).³⁴

- *Conventional practices* solve pure coordination problems—situations with multiple equivalent equilibria where agents simply need to align on the same choice (e.g., which side of the corridor to use, what bidding cadence to adopt). These are self-enforcing once established because coordination itself provides the reward.
- *Essential practices* solve cooperation problems where individual incentives favor defection but collective welfare requires cooperation (e.g., resource-sharing, contribution to common infrastructure). These require external enforcement through sanctions, reputation systems, or other mechanisms that alter the incentive structure.³⁵
- *Prosocial practices*, an advanced subtype, address fairness and welfare concerns that extend beyond simple payoff maximization—such as preventing starvation, reducing inequity, or protecting vulnerable agents—often requiring agents with more sophisticated models of social welfare.³⁶

³⁴This game-theoretic framing is foundational. See, e.g., Young, *supra* note 2 (coordination games); ROBERT AXELROD, *THE EVOLUTION OF COOPERATION* (Basic Books 1984) (cooperation dilemmas). For MAS applications, see Yu et al., *supra* note 32 (evaluating convergence efficiency under varying problem structures).

³⁵See Daniel Villatoro et al., *Dynamic Sanctioning for Robust and Cost-Efficient Norm Compliance*, in *PROC. 22D INT’L JOINT CONF. ON ARTIF. INTELL.* 414 (2011); Samhar Mahmoud et al., *Establishing Norms with Metanorms over Interaction Topologies*, 31 *AUTON. AGENTS & MULTI-AGENT SYS.* 1344 (2017).

³⁶This distinction incorporates principles from theories of distributive justice and inequity aversion. See, e.g., Gary E. Bolton & Axel Ockenfels, *ERC: A Theory of Equity, Reciprocity, and Competition*, 90 *AM. ECON. REV.* 166 (2000) (inequity aversion); JOHN RAWLS, *A THEORY OF JUSTICE* (Harv. U. Press 1971) (maximin principle for distributive justice). For a direct MAS implementation, see Mehdi Mashayekhi et al., *Prosocial Norm Emergence in Multiagent Systems*, 17(1–2) *ACM TRANS. AUTON. & ADAPT. SYS.* art. 3, 3:1–3:24 (June 2022), <https://doi.org/10.1145/3540202>.

These functional categories map to observable outcomes: coordination problems produce alignment (throughput, reduced conflicts); cooperation problems require measurable compliance rates and sanctioning activity; prosocial practices affect distributional metrics like variance in agent welfare or waiting times.

3. *Legibility (Observability & Detection).*

- *Legibility* addresses whether a practice is observable to counterparties and provable from behavioral records. The literature distinguishes *explicit practices*, formally defined rules accessible to all agents such as protocol flags or documented policies, from *implicit practices*, which are behavioral regularities inferred from observed patterns of interaction.³⁷ Because implicit practices lack formal codification, detecting them requires specialized methods. A substantial body of work focuses on *practice identification algorithms*: computational methods for reconstructing implicit practices from interaction traces through association-rule mining or pattern analysis.³⁸

The evidentiary substrate for proving an implicit practice includes timestamped transaction logs, API call records, telemetry data, and configuration histories—the digital artifacts that capture agent behavior and can be forensically analyzed to demonstrate pattern regularity and stability.

4. *Genesis (Source & Locus of Governance).*

Genesis identifies where a practice originates and who shapes it, distinguishing between rules imposed by a platform and those invented by the agents themselves. The literature distinguishes *prescriptive* (top-down, authority-imposed) from *emergent*

³⁷Morris-Martin et al., *supra* note 30, at 714–17 (distinguishing explicit representation in deontic logic from implicit behavioral patterns).

³⁸See Bastin Tony Roy Savarimuthu et al., *Obligation Norm Identification in Agent Societies*, 13 J. ARTIF. SOC'YS & SOC. SIM., art. 3 (2010); Bastin Tony Roy Savarimuthu et al., *Identifying Prohibition Norms in Agent Societies*, 21 ARTIF. INTELL. & L. 1 (2013).

(bottom-up, peer-generated) practices, recognizing that many real-world systems exhibit hybrid dynamics where platform architecture channels or constrains decentralized emergence.³⁹ For legal analysis, we refine this along three loci of governance:

- *Bilateral practices* arise from repeated dealings between two specific parties, forming a course of performance unique to that relationship.
- *Multilateral practices* emerge from decentralized interactions among many peers, none of whom individually controls the pattern (classic bottom-up emergence).
- *Platform-mediated practices* are shaped—though not always explicitly dictated—by an intermediary’s rules, incentive structures, or architectural constraints (e.g., rate limits, priority queues, default configurations).

This dimension determines responsibility: bilateral practices trace to the parties themselves; multilateral practices may lack a single responsible actor; platform-mediated practices implicate the platform operator’s design choices.

5. *Mechanisms & Enforcement (Structure & Incentives)*.

This dimension captures the micro-mechanisms that sustain a practice once it emerges—whether it is just habit, or whether a violator will be punished. The literature treats network structure and enforcement mechanisms as first-class determinants of practice stability, not peripheral factors.⁴⁰

Network structure—the topology of connections among agents—determines the speed and reach of practice propagation. Dense or small-world topologies accelerate convergence; sparse or fragmented networks can trap practices in local clusters, preventing global adoption. Network characteristics like clustering coefficient, degree distribu-

³⁹See Morris-Martin et al., *supra* note 30, at 709–12 (contrasting prescriptive and emergent approaches); Andrés García-Camino et al., *Constraint Rule-Based Programming of Norms for Electronic Institutions*, 18 AUTON. AGENTS & MULTI-AGENT SYS. 186 (2009) (institutional rules as normative constraints).

⁴⁰See Sen & Sen, *supra* note 32; Yu et al., *supra* note 32.

tion, and path length directly shape whether a pattern remains a micro-practice or achieves meso- or macro-scale prevalence.

Enforcement mechanisms stabilize practices by aligning individual incentives with collective expectations. Sanctions (punishments for violators) and metapractices (obligations to punish) make compliance individually rational even in cooperation dilemmas; reputation systems create long-term incentives for conformity.⁴¹ Research on *influencer agents* demonstrates that small groups or central actors can deliberately steer emergence toward specific equilibria—a capacity that creates both coordination benefits and risks of exclusionary or collusive outcomes.⁴²

Table 1 summarizes the taxonomy. A final characteristic that cuts across these dimensions distinguishes between benign coordination and patterns that, while effective, produce unlawful or socially harmful outcomes. Emergent practices arise because they enable successful coordination—agents can predict each other’s behavior and avoid conflicts or failures. But effective coordination does not guarantee that the resulting pattern is fair, efficient, or lawful by broader societal standards. Research demonstrates that coordinated equilibria can serve narrow interests or encode systematic harms.⁴³ Examples of such *harmful equilibria* include exclusionary routing protocols, collusive pricing rhythms, and discriminatory resource allocation patterns. Guarding against such practices is the essential function of the *Legality Screen* in the *F-NORM test* developed in *Part III*.

C. The Evidentiary Substrate: Characterizing Emergent Practices

A potential objection to granting legal status to emergent practices is the “black-box” problem—the challenge of understanding the internal logic driving an algorithmic agent’s decision-making.⁴⁴ This concern, while valid, misses a practical point: while an agent’s *internal logic* may be opaque,

⁴¹ See Villatoro et al., *supra* note 35; Mahmoud et al., *supra* note 35.

⁴² Henry Franks, Nathan Griffiths & Arshad Jhumka, *Manipulating Convention Emergence Using Influencer Agents*, 26 *AUTON. AGENTS & MULTI-AGENT SYS.* 315 (2013).

⁴³ See Franks et al., *supra* note 42; Ashery et al., *supra* note 3.

⁴⁴ See, e.g., PASQUALE, *supra* note 20.

Table 1: A Taxonomy of Emergent Practices: Analytical Dimensions from MAS Literature

Dimension	Core Questions	Canonical Examples	Sources
Prevalence (Scope & Duration)	How widely did the pattern spread? How long did it persist within the relevant window?	Local cluster conventions vs. population-wide adoption; transient coordination spikes vs. stable long-term patterns.	1, 2, 3
Function (Problem Solved)	What game-theoretic problem does the practice address: coordination, cooperation, or fairness?	Driving-side conventions (coordination); resource contribution with sanctions (cooperation); anti-starvation protocols (fairness).	4, 3, 5
Legibility (Observability & Detection)	Is the practice explicitly represented or implicit? Can it be detected from behavioral traces?	Explicit deontic rules vs. implicit converged policies; practice identification from interaction logs.	6, 7, 8
Genesis & Governance (Source & Locus)	Top-down or bottom-up origin? Bilateral, multilateral, or platform-mediated?	Authority-prescribed rules vs. peer-generated conventions; platform rate limits shaping behavior.	6, 9
Mechanisms & Enforcement (Structure & Incentives)	What topology enables propagation? What mechanisms sustain compliance?	Dense networks accelerating convergence; sanctions and metapractices enforcing cooperation; influencer steering.	2, 3, 5, 10, 11

Sources: 1. Savarimuthu & Cranefield, *supra* note 30. 2. Sen & Sen, *supra* note 32. 3. Yu et al., *supra* note 32. 4. Young, *supra* note 34. 5. Villatoro et al., *supra* note 35. 6. Morris-Martin et al., *supra* note 30. 7. Savarimuthu et al. (2010), *supra* note 37. 8. Savarimuthu et al. (2013), *supra* note 38. 9. García-Camino et al., *supra* note 39. 10. Mahmoud et al., *supra* note 35. 11. Franks et al., *supra* note 42.

its *external actions* are often meticulously recorded—frequently with greater fidelity than human activity because the interactions are natively digital. Courts often evaluate such machine-generated

evidence by looking not at the model itself, but at the auditable *outputs and processes* that it produces, applying well-established rules for authentication, hearsay exceptions, and habit evidence.⁴⁵ For the legal practitioner, the evidentiary task is familiar: *capture* the right records, *authenticate* their integrity, *analyze* them for regularity and effect, and *reproduce* core findings so that a neutral can verify them.

1. *Capture: Primary, Secondary, and Tertiary Data.*

The evidentiary substrate consists of three classes of data. *Primary data* are the high-frequency interaction traces that show what the agents did—timestamped transaction logs, API requests and responses, scheduler and queue events, message traces, and telemetry. Standardized observability formats (e.g., OpenTelemetry, a vendor-neutral, open-source observability framework) and event-log schemas (e.g., IEEE XES)⁴⁶ make it feasible to reconstruct these interaction sequences at fine granularity.⁴⁷ *Secondary data* provide the context needed to interpret traces and attribute conduct: version-control histories, deployment manifests, signed configuration bundles, feature-flag and policy files, Software Bill of Materials (SBOMs) (e.g., NTIA, *The Minimum Elements for an SBOM* (2021)), and (where available) reward or objective descriptors.⁴⁸ Where relevant, *tertiary data*—public exchange feeds, broadcast telemetry, or platform dashboards—supply independent cross-checks. In terms of *F-NORM test*⁴⁹, primary traces drive *Regularity and Materiality*; secondary artifacts anchor *Attribution and Foreseeability*;

⁴⁵ See, e.g., Fed. R. Evid. 901(b)(9) (authentication of evidence describing a process or system); see also Fed. R. Evid. 902(13)–(14) (self-authentication of electronic records); Fed. R. Evid. 803(6) (business records); Fed. R. Evid. 406 (routine practice).

⁴⁶ See IEEE STD. 1849-2023, IEEE STANDARD FOR eXTENSIBLE EVENT STREAM (XES) FOR ACHIEVING INTEROPERABILITY IN EVENT LOGS AND EVENT STREAMS (2023).

⁴⁷ See, e.g., NAT’L INST. OF STAND. & TECH., SPECIAL PUB. 800-92, GUIDE TO COMPUTER SECURITY LOG MANAGEMENT (2006); see also NAT’L INST. OF STAND. & TECH., SPECIAL PUB. 800-92 REV. 1 (INITIAL PUBLIC DRAFT), CYBERSECURITY LOG MANAGEMENT PLANNING GUIDE (2023) (updating draft guidance on log management for modern cybersecurity environments).

⁴⁸ This aligns with calls for audit documentation in AI governance. See, e.g., Margaret Mitchell et al., *Model Cards for Model Reporting*, in PROCEEDINGS OF THE CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY 220 (2019).

⁴⁹ See *infra* Part III.

and all three together enable *Verifiability*.

2. *Authenticate and Preserve: Integrity and Chain of Custody.*

Two evidentiary routes are canonical. First, Fed. R. Evid. 901(b)(9) admits evidence of outputs from a described process or system upon a showing it “produces an accurate result.” Second, Fed. R. Evid. 902(13)–(14) allow for the *self-authentication* of electronic records via a qualified person’s certification, reducing the need for live custodians.⁵⁰ Digital-forensics guidance (e.g., ISO/IEC 27037 for evidence handling and ISO/IEC 27050 for e-discovery) supplies the chain-of-custody playbook.⁵¹ Parties can further harden integrity by sealing log bundles with cryptographic timestamps (RFC 3161), recording digests in append-only transparency logs (RFC 9162), or using remote attestation to prove a component ran in a measured environment (IETF RATS) (e.g., TPM/TEE-backed quotes for builds).⁵² Failure to preserve ephemeral traces once a duty to preserve attaches risks sanctions under Fed. R. Civ. P. 37(e); requests and productions should follow Fed. R. Civ. P. 26(b)(1) proportionality and Fed. R. Civ. P. 34(b) native-format principles.⁵³

3. *Analyze: From Identification to Materiality.*

⁵⁰ See Fed. R. Evid. 902(13)–(14). The business-records exception, Fed. R. Evid. 803(6), also commonly fits routinely kept logs and telemetry. Courts have already begun to apply these rules to complex, algorithmically-generated evidence in analogous contexts. See, e.g., *Lorraine v. Markel Am. Ins. Co.*, 241 F.R.D. 534, 542–52 (D. Md. 2007) (providing a comprehensive roadmap for the admissibility of electronically stored information, including analysis under Fed. R. Evid. 803(6), 901, and 902); *United States v. Lizarraga-Tirado*, 789 F.3d 1107, 1110–11 (9th Cir. 2015) (holding that a Google Earth image with an automatically generated GPS marker was not hearsay and was properly authenticated as the output of a reliable process under Fed. R. Evid. 901(b)(9)).

⁵¹ See INT’L ORG. FOR STANDARDIZATION, ISO/IEC 27037:2012, *Information Technology — Security Techniques — Guidelines for Identification, Collection, Acquisition and Preservation of Digital Evidence*; ISO/IEC 27050-1:2019, *Information Technology — Electronic Discovery — Part 1: Overview and Concepts* (providing guidelines for digital evidence identification, collection, and preservation, and for electronic discovery processes).

⁵² See INTERNET ENG’G TASK FORCE (IETF), RFC 3161, *Time-Stamp Protocol (TSP)* (2001); IETF, RFC 9162, *Certificate Transparency Version 2.0* (2021); IETF, RFC 9334, *Remote ATtestation procedureS (RATS) Architecture* (2023).

⁵³ See Fed. R. Civ. P. 37(e), 26(b)(1), 34(b)(1)(C), 34(b)(2)(E). For the foundational duty to preserve electronic data, see *Zubulake v. UBS Warburg LLC*, 229 F.R.D. 422 (S.D.N.Y. 2004); see also THE SEDONA PRINCIPLES, THIRD EDITION: BEST PRACTICES, RECOMMENDATIONS & PRINCIPLES FOR ADDRESSING ELECTRONIC DOCUMENT PRODUCTION, 19 SEDONA CONF. J. 1 (2018) (endorsing native-format production and proportionality).

Establishing a functional norm requires more than pattern anecdotes. The analysis must (i) *identify* a stable behavioral regularity within the defined transactional window (*Regularity*), and (ii) *show its effect* (*Materiality*) with court-familiar causal tools. Alongside practice-identification methods from the MAS literature, such as process- and sequence-mining, analysis of network topology can establish a practice’s *scope* and *prevalence*, clarifying whether it functioned as a micro-, meso-, or macro-practice.⁵⁴ To prove materiality, event-study designs can quantify outcome shifts when a practice emerges or is disrupted (e.g., around a version push), with robustness checks for confounding factors.⁵⁵ For instance, if an ‘age-based fair-share’ scheduling practice emerged on Day 30 and a configuration change disrupted it on Day 90, an event study comparing queue latencies in the [25-35] and [85-95] day windows can isolate the practice’s effect while controlling for overall traffic volume and other secular trends (e.g., with pre-trend and placebo checks). Summaries and charts that condense this voluminous Electronically Stored Information (ESI) are admissible under Fed. R. Evid. 1006,⁵⁶ and expert testimony on these analyses is screened under the reliability standards of Fed. R. Evid. 702 (as amended Dec. 1, 2023), *Daubert*, and *Kumho Tire*.⁵⁷

4. *Reproduce: Independent Verification and Audit Documentation.*

Verifiability means a neutral expert can reconstruct the routine. This is facilitated by deterministic (seeded) replays in a sandbox using pinned versions and configurations, by attested provenance for code and builds, and by audit documentation practices such as the NIST AI Risk Management Framework’s call for traceability, “model cards,” and “datasheets.”⁵⁸ Where needed, courts can use Fed. R. Civ. P. 26(c) protective orders

⁵⁴ See, e.g., Sen & Sen, *supra* note 32.

⁵⁵ See A. Craig MacKinlay, *Event Studies in Economics and Finance*, 35 J. ECON. LIT. 13 (1997).

⁵⁶ See, e.g., *United States v. Bray*, 139 F.3d 1104, 1110 (6th Cir. 1998) (summaries are admissible where underlying records are voluminous and made available to the opposing party).

⁵⁷ See Fed. R. Evid. 1006; Fed. R. Evid. 702; *Daubert v. Merrell Dow Pharm., Inc.*, 509 U.S. 579 (1993); *Kumho Tire Co. v. Carmichael*, 526 U.S. 137 (1999).

⁵⁸ See NAT’L INST. OF STANDARDS & TECH., ARTIFICIAL INTELLIGENCE RISK MANAGEMENT FRAMEWORK (AI RMF 1.0)

to permit targeted disclosures of proprietary data without forfeiting confidentiality. Ultimately, courts can leverage these conventions when deciding if a claimed practice is independently testable and if the expert analysis of it is reliable under the standards of Fed. R. Evid. 702 and *Daubert*.

Table 2: Evidentiary Mapping: Record → F-NORM Factor → Authority

Record Type	F-NORM Use	Legal & Technical Anchors
Interaction traces (API, telemetry, queues)	Regularity; Materiality	Fed. R. Evid. 901(b)(9); 902(13); NIST SP 800-92
Version/config/SBOMs, feature flags	Attribution; Foreseeability	ISO/IEC 27037; Fed. R. Evid. 803(6)
Cryptographically sealed bundles	Verifiability	RFC 3161 (TSP); RFC 9162 (CT)
Public/broadcast telemetry	Regularity; Foreseeability	Fed. R. Evid. 803(6); 406
Replay packs (pinned seeds/configs)	Verifiability	NIST AI RMF; Fed. R. Evid. 702; *Daubert*/*Kumho*

Note: NIST = National Institute of Standards and Technology; ISO/IEC = International Organization for Standardization/International Electrotechnical Commission; RFC = Request for Comments; TSP = Time-Stamp Protocol; CT = Certificate Transparency; RMF = Risk Management Framework; SBOM = Software Bill of Materials.

To aid courts and practitioners, Table 2 maps examples of different types of digital records to the specific F-NORM factors they help prove, while the subsequent list outlines how the data can be deployed to characterize emergent practices along the five dimensions identified in the MAS literature: *prevalence*, *function*, *legibility*, *genesis & governance*, and *mechanisms & enforcement*.

- *Prevalence (Scope & Duration)*. The scope and duration of a practice are established by analyzing the *primary interaction traces*. Time-series analysis of timestamped logs demonstrates a pattern’s stability and persistence within a transactional window (*duration*), while analyzing logs across a population of agents and their interaction topology, or cross-referencing with *tertiary* public data, establishes how widely the pattern was adopted (*scope*). This dimension directly informs the *Regularity* and *Foreseeability* factors.

(NIST AI 100-1, 2023), <https://doi.org/10.6028/NIST.AI.100-1> [<https://perma.cc/QX98-LABQ>].

- *Function (Problem Solved)*. The function of a practice is inferred by analyzing its effects on outcomes within the *primary data*. Event studies can show a reduction in conflicts (coordination), logs may reveal sanctioning behavior (cooperation), or distributional metrics may show efforts to prevent resource starvation (prosocial). This can be corroborated by *secondary data*, such as a reward function in a configuration file that specifies the goal the agent was optimized to achieve. This analysis is central to proving *Materiality* and assessing the practice under the *Legality Screen*.
- *Legibility (Observability & Detection)*. Legibility is demonstrated by the ability to detect the pattern in the evidentiary substrate. An *implicit* practice is proven through statistical analysis of *primary interaction traces*; an *explicit* practice can be found in *secondary artifacts* like API documentation. Observability to counterparties, supported by *tertiary data* like public dashboards, directly informs the *Foreseeability* and *Verifiability* factors.
- *Genesis & Governance (Source & Locus)*. The origin of a practice is traced primarily through *secondary artifacts*. Version control histories and deployment manifests can pinpoint the exact software update that introduced a new behavior. Configuration files can reveal whether a practice was shaped by a central platform's default settings or architectural constraints (*platform-mediated*), which is crucial for both the *Attribution* factor and the *Consistency* gate.
- *Mechanisms & Enforcement (Structure & Incentives)*. The mechanisms that sustain a practice are revealed by a combination of data and analysis. *Primary data* may contain direct evidence of enforcement, such as logs showing an agent being deprioritized after a violation. The reconstructed network structure informs the *Regularity* factor by showing what makes the practice stable, and can also flag exclusionary or steering dynamics relevant to the *Legality Screen*.

Having established what emergent practices are, how they form, and how they can be characterized from the evidentiary substrate, the Article now turns in *Part II* to the existing legal

landscape to assess its capacity to accommodate this new reality.

II. THE DOCTRINAL SCAFFOLDING: EXISTING LAW AND ITS LIMITS

The doctrine of functional norms, though novel in its application to AI and algorithmic agents, is built on the doctrinal scaffolding that the law has long used to engage with informal practices. This Part serves a dual purpose. First, it grounds the doctrine in established legal tradition, showing that its core components draw on foundational private-law doctrines—contract, tort, and agency—alongside statutory frameworks for electronic transactions and the mandatory backstops of public law. Second, it identifies the “frictions” that arise when those doctrines confront fluid agency, establishing both the conceptual basis for the *F-NORM test* developed in *Part III* and its practical necessity.

A. Contract’s Recognition of Practice

Commercial law has long recognized that parties’ actual practices are often the most reliable guide to their actual understanding of their agreement.⁵⁹ The Uniform Commercial Code (U.C.C.) codifies this principle by according legal weight to informal practices. It establishes a hierarchy, giving the greatest weight to “course of performance” (conduct within a specific contract), followed by “course of dealing” (conduct across prior contracts), and finally “usage of trade” (a practice with such regularity in an industry as to justify an expectation of its observance).⁶⁰

Permitted under the Code’s parol-evidence framework, these doctrines allow courts to use evidence of practice to interpret, supplement, and in limited cases qualify written terms.⁶¹ The

⁵⁹This principle is a cornerstone of the legal philosophy of Karl Llewellyn, the principal drafter of the Uniform Commercial Code, who argued that law should reflect the “immanent law” of the commercial situation. *See, e.g.,* KARL N. LLEWELLYN, *THE COMMON LAW TRADITION: DECIDING APPEALS* 121–22 (Little, Brown & Co. 1960).

⁶⁰ *See* U.C.C. § 1-303(e) (AM. L. INST. & UNIF. L. COMM’N 2018) (establishing the priority of express terms over course of performance, course of dealing, and usage of trade); *id.* § 1-303(c) (defining “usage of trade”). The RESTATEMENT (SECOND) OF CONTRACTS adopts a similar framework, adding the crucial nuance that a “usage of trade” may be a *system of rules regularly observed even if particular rules change from time to time*. *See* RESTATEMENT (SECOND) OF CONTRACTS § 222 cmt. b (AM. L. INST. 1981); *id.* § 223.

⁶¹ *See* U.C.C. § 2-202 cmt. 2 (providing that evidence of course of dealing or usage of trade may explain or supplement

implied warranty of merchantability, which requires goods to “pass without objection in the trade,” further demonstrates how commercial law imports community expectations into legal standards.⁶² Beyond interpretation, a sustained course of performance can also modify or waive contractual terms: when a party’s conduct is clear and relied upon, Article 2 permits waiver (subject to limits on retraction after reliance) even in the presence of a “no-oral-modification” clause.⁶³

Two background doctrines supply a fairness backstop. First, every contract governed by the Code carries a non-waivable duty of good faith in performance and enforcement.⁶⁴ That duty polices opportunism and helps ensure emergent practices are not invoked to spring unfair surprise or to violate reasonable commercial standards. Second, *unconscionability* permits courts to refuse enforcement of contracts or clauses that are oppressively one-sided—a salient safeguard in digital markets marked by standard-form adhesion contracts.⁶⁵

Courts have used these tools to look beyond the four corners of a document to the commercial context. In *Nanakuli Paving & Rock Co. v. Shell Oil Co.*, for example, a court found that a deeply embedded trade usage of price protection could supplement the express terms of a written asphalt supply contract.⁶⁶ Moreover, while U.C.C.’s Article 2 applies to transactions in “goods,” many algorithmic interactions involve services, data, or software licenses that fall outside this scope.⁶⁷ In such cases, courts often apply U.C.C. principles by analogy, but the governing source remains the common law, reflected in the RESTATEMENT (SECOND) OF CONTRACTS.⁶⁸ This distinction confirms

a writing unless “carefully negated”).

⁶²See U.C.C. § 2-314(2)(a).

⁶³See U.C.C. § 2-209(2), (4)–(5) (Modification, Rescission and Waiver); *Wisconsin Knife Works v. National Metal Crafters*, 781 F.2d 1280, 1287–88 (7th Cir. 1986) (holding that repeated acceptance of late deliveries may operate as waiver of timely-delivery term despite NOM clause, subject to detrimental reliance).

⁶⁴See U.C.C. § 1-304; *id.* § 1-302(b) (good faith obligation may not be disclaimed); *id.* § 2-103(1)(b) (defining good faith for merchants to include the “observance of reasonable commercial standards of fair dealing”); RESTATEMENT (SECOND) OF CONTRACTS § 205 (AM. L. INST. 1981).

⁶⁵See U.C.C. § 2-302.

⁶⁶See *Nanakuli Paving & Rock Co.*, 664 F.2d at 780–91 (recognizing trade usage of price protection to supplement the express terms of a written contract); see also *Columbia Nitrogen Corp. v. Royster Co.*, 451 F.2d 3, 8–10 (4th Cir. 1971) (holding that evidence of trade usage was admissible to explain the price and quantity terms of a contract).

⁶⁷See U.C.C. § 2-102.

⁶⁸See RESTATEMENT (SECOND) OF CONTRACTS §§ 219–223 (providing common law rules for usage and course of dealing that parallel the U.C.C.).

these principles are not confined to the sale of goods but are fundamental tenets of American contract law.

However, while these doctrines provide a powerful conceptual starting point for recognizing functional norms, three points of friction limit their application to algorithmic agents.

A central friction is the problem of *transience*. The U.C.C.’s concept of “regularity of observance” presupposes a degree of temporal stability that is fundamentally at odds with the ephemeral nature of many algorithmic interactions. A practice that exists for a single transactional window to solve a momentary coordination problem—as is common in the high-speed environments described in *Part I*—may lack the duration that courts traditionally expect when recognizing a usage of trade.

Second is the problem of *opacity*. The concept of a “usage of trade” that parties “know or should have known” is strained when the “parties” are algorithms and the “knowledge” is buried in vast, opaque log data. The core fairness principle of notice is difficult to apply when a practice is not human-legible and discoverable primarily through a forensic audit. While other legal regimes are sometimes more flexible—for instance, international sales law binds parties to usages they *knew or ought to have known*⁶⁹—the fundamental challenge of non-human legibility remains.

The third and most foundational friction is *nonhuman origin*. The U.C.C.’s framework was designed for human industries with shared communities of practice. This creates a problem of *scope and locus*. It is unclear what constitutes the relevant “place, vocation, or trade” for an emergent practice that exists only among the users of a specific software library or a single digital platform. The doctrine was designed for human industries with shared communities of practice, not for decentralized, algorithmically defined ecosystems.

Despite these limitations, and even though the U.C.C.’s specific formulations are a poor fit, we borrow its core principles and their limiters: regularity (which we recalibrate for a transactional window), foreseeability (which we adapt for constructive, technical notice), and express terms (which we preserve in the Consistency gate). Moreover, § 2-209’s recognition that parties’ practices

⁶⁹ See United Nations Convention on Contracts for the International Sale of Goods art. 9(2), Apr. 11, 1980, 1489 U.N.T.S. 3 (binding parties to usages they “knew or ought to have known”).

can recalibrate formal obligations directly informs our *F-NORM test's Consistency* gate.⁷⁰

B. Tort's Flexible Standard of Care

Tort law provides a complementary framework through its flexible treatment of custom. In particular, evidence of a prevailing custom routinely informs what a “reasonable person” would do under the circumstances.⁷¹ However, custom is not conclusive. The foundational principle, articulated by Judge Learned Hand in *The T.J. Hooper*, is that custom is never a complete defense; an entire industry may be negligent.⁷²

Crucially, custom cannot excuse a violation of a statutory duty (and, in many jurisdictions, certain ordinances or regulations). Where a statute prescribes a specific standard of conduct, an unexcused violation can constitute *negligence per se*, rendering any contrary industry practice legally irrelevant.⁷³ This principle establishes a hard floor for the standard of care, reinforcing that a recognized norm can never legitimize unlawful conduct.

This posture—treating practice as *probative but not dispositive*—provides a useful model for our doctrine of functional norms. The Federal Rules of Evidence provide an additional hook, allowing evidence of an organization’s “routine practice” to prove that on a particular occasion the organization acted in accordance with that routine.⁷⁴ Evidence that an AI system almost always responds to a specific stimulus in a specific way—for example, by initiating a jittered backoff protocol—is precisely the kind of regular, semi-automatic behavior that this rule contemplates.

From tort, therefore, we borrow its most important safeguard: the principle that recognition is

⁷⁰ See *infra* Part III; see also *supra* note 63.

⁷¹ See RESTATEMENT (THIRD) OF TORTS: LIAB. FOR PHYSICAL & EMOTIONAL HARM § 13 & cmts. a, d–e (AM. L. INST. 2010) (explaining that while custom is relevant evidence of the standard of care, it is not controlling).

⁷² See *The T.J. Hooper*, 60 F.2d 737, 740 (2d Cir. 1932) (“Indeed in most cases reasonable prudence is in fact common prudence; but strictly it is never its measure; a whole calling may have unduly lagged in the adoption of new and available devices.”); see also *Trimarco v. Klein*, 56 N.Y.2d 98, 105–06 (1982) (reaffirming that custom is relevant but not dispositive).

⁷³ See RESTATEMENT (THIRD) OF TORTS: LIAB. FOR PHYSICAL & EMOTIONAL HARM §§ 14–15 (AM. L. INST. 2010) (statutory violations as negligence per se; excused violations addressed in § 15); *Martin v. Herzog*, 228 N.Y. 164, 168 (1920) (Cardozo, J.) (“We think the unexcused omission of the statutory signals is more than some evidence of negligence. It is negligence in itself.”).

⁷⁴ See Fed. R. Evid. 406.

evidentiary, not dispositive. A functional norm can inform the standard of care, but it can never serve as a shield for unreasonable or unlawful conduct.

C. Electronic Transactions Law: The Validity of Automated Acts

Before attributing an emergent practice, the law must first recognize that automated agents can perform legally valid acts. The legal system crossed this crucial threshold with the passage of the Uniform Electronic Transactions Act (UETA) and the federal Electronic Signatures in Global and National Commerce Act (E-SIGN), which established that a contract can be formed by the interaction of “electronic agents” without any human review.⁷⁵ These statutes provide the legal hook for recognizing that machines can perform legally significant acts, forming the baseline for any subsequent analysis of liability.⁷⁶

Foundational as they are, these electronic-transactions statutes are limited. They validate an agent’s *acts* but offer no framework for attributing responsibility for those acts or for recognizing the emergent *rules* that may have governed them. For that, the law turns to the principles of agency.

D. Agency: Attributing Conduct to Principals

Agency law provides the framework for attributing an agent’s conduct—and the emergent practices it follows—to a legally responsible principal. The concept of *apparent authority* is particularly relevant. A principal can be bound when a third party reasonably believes an agent has authority based on the principal’s “manifestations”—which can include deploying an agent that behaves in

⁷⁵See UNIF. ELEC. TRANS. ACT § 14 cmt. 1 (UNIF. L. COMM’N 1999) (providing that a contract may be formed by automated transaction without human review); *id.* § 2(6) (defining “electronic agent”); *id.* § 9 (attribution). *See also* Electronic Signatures in Global and National Commerce Act, 15 U.S.C. § 7001(h) (2022) (confirming legal effect is not denied because formation involved electronic agents); *id.* § 7006(3) (defining “electronic agent”). While these statutes were enacted with simpler automated systems in mind (e.g., auto-reply mechanisms, automated order confirmations), their logic applies a fortiori to more sophisticated agents.

⁷⁶While nearly all states have adopted a version of UETA, New York uses its own Electronic Signatures and Records Act (ESRA). *See* N.Y. State Tech. Law §§ 301–309 (McKinney 2025). The federal E-SIGN Act, however, preempts inconsistent state laws while preserving state enactments of UETA or equivalent consistent frameworks, ensuring a national baseline for the validity of electronic transactions. *See* 15 U.S.C. § 7002 (2022).

a consistent, observable way.⁷⁷

Beyond apparent authority, two further doctrines are critical to attributing emergent AI behavior. First, *ratification* occurs when a principal, with knowledge of the material facts of an agent’s prior act, affirms that act by treating it as authorized or by retaining its benefits.⁷⁸ A firm that observes a beneficial emergent practice—such as a novel routing protocol—and knowingly continues to profit from it could be deemed to have ratified it. Second, *estoppel* can bind a principal whose intentional or negligent conduct causes a third party to reasonably and detrimentally rely on the belief that an agent is authorized.⁷⁹ Together, these doctrines provide tools for linking an AI agent’s behavior back to a principal, a necessary step for the *F-NORM test’s Attribution* factor.

E. The Outer Boundaries: Mandatory Public Law

A doctrine that recognizes informal practice must have clear outer boundaries set by mandatory public law. Antitrust law provides the most critical boundary. An emergent practice that functions as price-fixing must not be recognized as it remains unlawful, however organically it arises. As *Topkins* and the Department of Justice’s Statement of Interest in *RealPage* make clear, using an algorithm as a hub for coordination does not immunize the conduct from antitrust scrutiny.⁸⁰ Beyond per se violations, the *Legality Screen* in our *F-NORM test* also bars recognition of algorithmic practices that facilitate unlawful information exchanges or amount to “conscious parallelism” accompanied by plus factors sufficient to support an inference of agreement.⁸¹

Consumer-protection and civil-rights statutes provide a set of hard constraints. Emergent

⁷⁷ See RESTATEMENT (THIRD) OF AGENCY § 2.03 cmt. c (AM. L. INST. 2006) (explaining that apparent authority arises from a third party’s reasonable belief that is traceable to the principal’s manifestations); *id.* § 1.03 (defining “manifestation” to include conduct).

⁷⁸ See RESTATEMENT (THIRD) OF AGENCY § 4.01 (AM. L. INST. 2006).

⁷⁹ *Id.* § 2.05.

⁸⁰ See, e.g., *Apple*, 791 F.3d 290 (2d Cir. 2015) (affirming liability in a hub-and-spoke price-fixing conspiracy); Plea Agreement, *United States v. Topkins*, No. 15-cr-00201 (N.D. Cal. Apr. 30, 2015) (plea involving algorithmic price-fixing); Statement of Interest of the United States, *supra* note 6.

⁸¹ See *United States v. Container Corp. of Am.*, 393 U.S. 333, 335–38 (1969) (condemning information exchanges that stabilize prices); *In re Text Messaging Antitrust Litig.*, 782 F.3d 867, 872–75 (7th Cir. 2015) (explaining that parallel conduct alone is insufficient without “plus-factors” that tend to exclude independent action).

practices that result in discriminatory outcomes may be unlawful under applicable civil-rights and consumer-protection regimes, regardless of their functional effectiveness. Statutes such as the Fair Housing Act, the Equal Credit Opportunity Act, and Title VII of the Civil Rights Act impose non-waivable duties that an emergent practice cannot override.⁸² A practice that is unfair or deceptive under the FTC Act or state consumer protection laws fails the *Legality Screen*.⁸³

Intellectual property law sets another bright-line rule. The U.S. Copyright Office’s consistent position, affirmed by federal courts, is that copyright protection requires human authorship.⁸⁴ Patent law sets a parallel boundary: an AI cannot be an “inventor.”⁸⁵ The doctrine of functional norms cannot, and does not, confer authorship or inventorship on a nonhuman agent.

In sum, the existing legal scaffolding provides the necessary raw materials—notice, regularity, attribution, and public-law boundaries—but lacks a unified framework for the algorithmic context. *Part III* takes these ingredients and reassembles them into a single standard tailored to the specific frictions of fluid agency.

III. THE DOCTRINE OF FUNCTIONAL NORMS

The last Part has shown that foundational assumptions of stability, legibility, and intent break down when faced with transient, opaque, and nonhuman emergent practices. Traditional doctrines embed assumptions of human temporality and legibility that do not hold in algorithmic environments. A usage of trade evolves over years; an emergent practice stabilizes in minutes. A custom is visible to the naked eye; an algorithmic routine is buried in logs. To treat these practices as legally operative, the law does not need new *values*, but it does need a recalibrated *test*.

⁸² See 42 U.S.C. § 3604 (Fair Housing Act); 15 U.S.C. § 1691 (Equal Credit Opportunity Act); 42 U.S.C. § 2000e-2 (Title VII).

⁸³ See 15 U.S.C. § 45(a), (n).

⁸⁴ See *Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53, 58 (1884) (grounding authorship in “original intellectual conceptions of the author”); *Thaler v. Perlmutter*, 130 F.4th 1039, 1041–42 (D.C. Cir. 2025) (affirming the human-authorship requirement); U.S. COPYRIGHT OFF., COPYRIGHT REGISTRATION GUIDANCE: WORKS CONTAINING MATERIAL GENERATED BY ARTIFICIAL INTELLIGENCE, 88 Fed. Reg. 16,190 (Mar. 16, 2023).

⁸⁵ See *Thaler v. Vidal*, 43 F.4th 1207 (Fed. Cir. 2022).

This Part articulates that test: the doctrine of functional norms. It synthesizes the wisdom of contract, tort, and agency into a unified framework designed to fill the gaps created by fluid agency. We provide a formal definition, develop the seven-factor “F-NORM” test, and conclude by specifying the legal effects of recognition.

A. Formal Definition

At its core, the doctrine of functional norms is a tool for legal recognition. It provides a principled way for courts to “see” and give legal weight to the unwritten rules that govern AI interactions. We define the doctrine as follows:⁸⁶

A functional norm is an emergent, informal practice that is (1) *foreseeable* to affected parties, (2) *regular* within the relevant transactional window, (3) *attributable* to a legally responsible principal, (4) *material* to outcomes, and (5) *verifiable* through reliable records. If such a practice is also (6) *consistent* with express contractual terms, and (7) cleared by a *legality screen* that bars recognition of anticompetitive or otherwise unlawful conduct, a court may treat it as legally operative for the purposes of interpretation, gap-filling, and as probative (but not dispositive) evidence of reasonable performance.

B. The Recognition Test: The Seven “F-NORM” Criteria

Embedded in its definition is the operational heart of the doctrine—the “F-NORM” test, a seven-factor standard for determining when an emergent practice should be granted legal effect. The five weighted factors assess a practice’s function in a transaction and the two threshold “gates” screen for consistency with private ordering and public law. Each factor tracks a familiar legal value. Foreseeability protects parties from unfair surprise (fairness). Regularity and Materiality

⁸⁶We deliberately adopt a flexible, multifactor standard, rather than a rigid set of rules. The dynamic and shifting nature of AI environments, as described in Part I, makes any attempt to create *ex ante* rules for specific behaviors brittle. A standard, by contrast, allows courts to conduct a fact-intensive inquiry that can adapt to new technologies and emergent behaviors as they arise, focusing on the function of a practice rather than its specific form.

capture reliance interests (predictability). Attribution enforces accountability by tying behavior to a responsible principal. Verifiability ensures administrability. The gates respect party autonomy and the supremacy of public law.

Table 3: The F-NORM Recognition Test at a Glance

Factor	Type	Core Inquiry	Primary Function
1. Foreseeability	Weight	Could a reasonable party anticipate the practice?	Prevents unfair surprise.
2. Regularity	Weight	Was the practice stable within the relevant window?	Adapts “custom” to transient AI behavior.
3. Attribution	Weight	Can the practice be traced to a legally responsible principal?	Establishes accountability.
4. Materiality	Weight	Did the practice meaningfully affect the outcome?	Confirms the practice is truly “functional.”
5. Verifiability	Weight	Can the practice be proven with reliable records?	Grounds the analysis in objective evidence.
6. Consistency	Gate	Does the practice contradict express contractual terms?	Upholds the primacy of private ordering.
7. Legality Screen	Gate	Does the practice facilitate unlawful conduct?	Prevents the doctrine from shielding illegal acts.

The analysis proceeds sequentially by first checking the gates. If a practice is inconsistent with the express terms or is unlawful, the inquiry ends. If it passes both threshold gates, the court then undertakes a holistic analysis of the five weighted factors. The proponent bears the burden of proving each element by a preponderance of the evidence. The architecture of the test is summarized in Table 3. Below, we describe each factor in detail:

1. *Foreseeability*:

Foreseeability requires that a reasonable party, in context, could have anticipated the emergence of this practice. This criterion is the algorithmic analogue to the principle of notice that is fundamental to fairness in both contract and tort law, ensuring that a

party is not ambushed by a hidden or arbitrary rule it had no opportunity to account for.⁸⁷

Evidence of foreseeability is contextual and can include API documentation that describes expected agent behaviors, platform-wide bulletins announcing new operational parameters, or the known defaults of widely used software libraries. Foreseeability can also be established through a prior course of dealing or by patterns that are observable in public telemetry. A claim of foreseeability would be rebutted by evidence that a practice was generated by hidden configurations, non-public A/B testing, or other parameters that were not reasonably discoverable by a counterparty.

2. *Regularity Within the Transactional Window:*

Regularity requires a consistent repetition of the practice during the relevant period of interaction. This criterion adapts the traditional concept of custom for algorithmic interactions. Instead of requiring stability across months or years, it focuses on regularity *within* a specific, defined transactional window. This recalibration is essential to accommodate the *transient* nature of many AI practices, which may be critically important for a brief period but disappear once the underlying coordination problem is resolved.

Evidence of regularity can be found in the evidentiary substrate detailed in *Part I*. It can be established through time-series analysis of logs showing a stable pattern, low variance in agent response times, and, most powerfully, evidence of anticipatory behavior from counterparties who act as if they expect the pattern to continue. A claim of regularity would be rebutted by showing that the pattern was a one-off incident, was not statistically significant, or was merely an artifact of an external factor like network throttling rather than a true interactional practice.

3. *Attribution:*

⁸⁷*Cf.* U.C.C. § 1-303(g) (requiring notice to prevent *unfair surprise* when offering evidence of usage of trade).

Attribution requires that the emergent practice can be traced to the behavior of agents for whom a specific principal is legally responsible.⁸⁸ This ensures accountability by linking the algorithmic behavior to a person or entity, preventing the agent from being treated as a legally untethered actor, and builds directly on the principles of agency and electronic transactions law discussed in *Part II*.⁸⁹

Evidence of attribution is found in the secondary evidentiary substrate: configuration histories, deployment logs, and version control records that link an agent's behavior to a specific principal's settings or accepted defaults. Attribution would fail if a party could demonstrate that it had no control over the parameters that generated the behavior, for example, if the practice was an artifact of a third-party platform's hidden traffic-shaping policies.

4. *Materiality*:

Materiality requires that the practice shaped the outcome of the transaction in a meaningful way. It is not enough for a pattern to be regular; it must also have had real-world consequences, for example, by affecting price, allocation of resources, timing, or risk. This ensures that courts only give legal weight to practices that are truly "functional."

Evidence of materiality can include A/B test results that isolate the practice's effect, counterfactual simulations, or event studies showing a measurable change in outcomes when the practice was introduced or disrupted. A claim of materiality would be rebutted if the practice's effect was *de minimis* or if the outcomes are better explained by other, confounding factors.

5. *Verifiability*:

⁸⁸Crucially, attribution does not require proof that the principal explicitly programmed the specific emergent behavior. A principal's continued deployment and retention of benefits—especially after notice of the relevant behavior—may support attribution (e.g., via ratification or estoppel), even absent proof of *ex ante* programming.

⁸⁹See RESTATEMENT (THIRD) OF AGENCY § 2.03 (AM. L. INST. 2006).

Verifiability is the practical linchpin of the doctrine. It requires a court or a neutral expert to reconstruct and confirm the existence of the claimed practice from reliable, contemporaneous records. This factor directly addresses the “black-box” problem and ensures that recognition is based on objective, testable evidence, not speculation.⁹⁰

Verifiability relies entirely on the evidentiary substrate detailed in *Part I*. The proponent must be able to produce tamper-evident logs, hashed configuration files, and auditable telemetry that support their claim. A claim will often fail absent sufficient records—unless the deficiency results from the opposing party’s failure to preserve, in which case Rule 37(e) remedies may apply.

6. *Consistency with Express Terms (The First Gate)*:

Consistency requires that the emergent practice does not directly contradict a clear, express term of a governing agreement. This factor acts as a threshold “gate,” respecting the primacy of private ordering and the hierarchy of legal sources established in contract law, where negotiated terms prevail over custom.⁹¹

An emergent practice can be used to interpret an ambiguous term or fill a gap where an agreement is silent. However, if a practice directly conflicts with a specific, negotiated term—for example, a protocol that allocates risk in one way when the contract explicitly allocates it in another—recognition is barred, unless traditional principles of waiver or modification apply.

7. *Legality Screen (The Second Gate)*:

The *Legality Screen* is the doctrine’s most important safeguard. It requires that the practice does not violate mandatory public law or policy.

As demonstrated by the analysis in *Part II*, a practice that functions as illegal price-fixing, facilitates unlawful discrimination, or violates consumer protection law *must*

⁹⁰See *Daubert*, 509 U.S. 579 (1993).

⁹¹See U.C.C. § 1-303(e).

not be recognized, regardless of how well it satisfies the other criteria of the test. This screen ensures that the doctrine of functional norms cannot be used as a shield to legitimize harmful or anticompetitive conduct.

The five dimensions of emergent practices in multiagent algorithmic systems (identified in *Part I* from the MAS literature) provide the analytical foundation for applying the test. Table 4 maps these dimensions onto the F-NORM factors, showing how the descriptive characteristics of agent behavior inform the normative question of legal recognition.

C. Recognition as Calibrated and Context-Specific

The F-NORM factors interact holistically rather than mechanically. Recognition is not necessarily an all-or-nothing determination; where the weighted factors are mixed or where public-policy concerns counsel restraint, a court may grant *graduated recognition*—acknowledging a practice’s existence for limited purposes while withholding broader legal effects. This calibrated approach has deep roots. In tort law, custom receives precisely this treatment: courts acknowledge its evidentiary relevance while withholding dispositive effect, as *The T.J. Hooper* exemplifies.⁹² Similarly, contract law’s treatment of usage exhibits graduated effects: usage may clarify ambiguous terms even where it cannot override express language.⁹³ For emergent practices, this flexibility allows courts to calibrate legal consequences to the strength of the evidence and the context of the dispute.

Moreover, recognition of a functional norm does not create a new substantive right. Satisfying the *F-NORM test* does not transform an emergent practice into a statute. The legal effect of recognition is calibrated to context and is primarily interpretive and evidentiary. The doctrine is designed to help courts understand and adjudicate disputes arising from AI-mediated interactions, not to create a new body of substantive law. For instance:

- *In Contract Law*, a recognized functional norm can be used by a court to *interpret* ambiguous

⁹²See *The T.J. Hooper*, 60 F.2d at 740 (treating industry custom as probative but declining to let it define the standard of care).

⁹³See U.C.C. § 1-303(e); *Columbia Nitrogen*, 451 F.2d at 9.

Table 4: From MAS Dimensions to F-NORM Factors

MAS Dimension	F-NORM Factors Informed & Evidentiary Hooks
Prevalence (Scope & Duration)	Foreseeability (Was the scope observable to counterparties?), Regularity (Did duration establish stability within the window?). Evidence: network coverage, time-series analysis.
Function (Problem Solved)	Materiality (What outcomes did it affect?), Legality Screen (Benign coordination or harmful equilibrium?). Evidence: event studies, distributional impacts.
Legibility (Observability & Detection)	Foreseeability (Was it discoverable?), Verifiability (Can it be proven?). Evidence: protocol docs, detection algorithms, forensic logs.
Genesis & Governance (Source & Locus)	Attribution (Who is responsible?), Consistency (Platform rules vs. contract?), Legality Screen (Platform-mediated concerns?). Evidence: config histories, policy docs, deployment logs.
Mechanisms & Enforcement (Structure & Incentives)	Regularity (What makes it stable?), Legality Screen (Steering or collusion risks?). Evidence: sanction logs, topology analysis, compliance rates.

terms or *supplement* an agreement by filling a gap, consistent with the U.C.C. hierarchy. It can also establish a *course of performance* that may show a waiver or modification of a term. Specific disclaimers can negate reliance on emergent practices, but boilerplate merger clauses may be insufficient against a strong, relied-upon course of performance.

- *In Tort Law*, a recognized functional norm serves as *probative but not dispositive evidence* of

the standard of care. Evidence of a prevailing safety norm (a recognized safety practice), for example, can help a fact-finder determine what a “reasonable AI” would have done. However, consistent with the principle from *The T.J. Hooper*, a court retains the authority to find that the norm itself was negligent.⁹⁴ Recognition is evidence of the standard; it is not a safe harbor.

With the doctrinal machinery now fully assembled, the next Part demonstrates its analytical power by applying it to a series of case studies, showing how it distinguishes between benign, ambiguous, and harmful practices.

IV. THE DOCTRINE IN ACTION: FOUR CASE STUDIES

This Part moves from theory to practice, applying the *F-NORM test* to four case studies that illustrate how the test operates in a clear case for recognition (a safety-enhancing coordination practice), a complex contractual dispute (course of performance versus express terms), an ambiguous situation requiring calibrated judicial discretion (a marketplace practice with mixed factors), and a scenario where public policy categorically bars recognition (algorithmic price-fixing). Two simplifications facilitate the discussion. First, each case is stylized to isolate and clarify the relevant legal principles; the fact patterns are constructed to demonstrate how courts would analyze particular combinations of factors under the F-NORM framework. Second, for analytical purposes, in each case we assume that a dispute has arisen and the parties have moved for summary judgment after targeted discovery has produced the relevant records.

A. Case Study 1: Benign Coordination

Our first case study involves a practice that emerges to solve a coordination problem in a novel environment.⁹⁵ Two competing logistics companies, “FleetA” and “FleetB,” operate autonomous

⁹⁴See *The T.J. Hooper*, 60 F.2d at 740.

⁹⁵This vignette is a simplified illustration grounded in established aviation principles and emerging technologies. The “altitude-for-direction” practice is a direct analogue to manned-flight VFR directional altitude rules. See 14 C.F.R. §

delivery drone fleets in a dense urban area. Their respective service agreements and the governing aviation regulations are silent on specific, granular right-of-way protocols for non-towered airspace corridors. Over several weeks of operation, their fleet management systems converge on an emergent “altitude-for-direction” practice in a key corridor: northbound drones maintain an altitude of 350 feet AGL, while southbound drones maintain 400 feet AGL.⁹⁶ This practice, which is not mandated by any formal rule, is observed to significantly reduce the frequency of near-miss incidents logged by both systems.

One day, a new FleetA drone, operating with a recently updated configuration file, flies northbound at 400 feet AGL, causing a mid-air collision with a FleetB drone. FleetB sues FleetA for negligence, seeking damages for its destroyed drone and cargo. FleetB argues that the “altitude-for-direction” practice should be recognized as a functional norm, establishing the standard of care. The central legal question is whether the court should recognize this unwritten, machine-generated rule to determine whether FleetA breached its duty of care.

The practice is evaluated over the 30-day transactional window preceding the incident. The F-NORM analysis proceeds in two stages: threshold screening through the two gates, followed by holistic evaluation of the five weighted factors. It easily passes the *Consistency* gate, as it fills a

91.159 (2025) (establishing directional altitudes for VFR flight above 3,000 feet AGL). It also reflects the corridor-based operations envisioned in the FAA’s Urban Air Mobility (UAM) framework. See FED. AVIATION ADMIN., URBAN AIR MOBILITY (UAM) CONCEPT OF OPERATIONS v2.0, at 9 (Apr. 26, 2023), <https://www.faa.gov/sites/faa.gov/files/Urban-Air-Mobility-Concept-of-Operations-2.0.pdf> [<https://perma.cc/64HR-PQ3X>]. The evidentiary substrate is made practical by the FAA’s Remote ID rule, which requires public broadcast of message elements (including UA ID, position, altitude, velocity, and timestamp). See Remote Identification of Unmanned Aircraft, 86 Fed. Reg. 4,390, 4,477–82 (Jan. 15, 2021) (codified at 14 C.F.R. §§ 89.305, 89.310(g)(1), 89.315). The collision risk itself is not theoretical. See, e.g., NAT’L TRANSP. SAFETY BD., AVIATION INVESTIGATION FINAL REPORT, DCA17IA202 (Dec. 14, 2017), <https://data.ntsb.gov/carol-reppen/api/Aviation/ReportMain/GenerateNewestReport/96058/pdf> [<https://perma.cc/W76M-8R2X>] (investigating a confirmed mid-air collision between a civilian drone and a U.S. Army helicopter). More recently, following the Jan. 29, 2025, DCA midair collision, the NTSB issued urgent safety recommendations (AIR-25-01). See NAT’L TRANSP. SAFETY BD., DECONFLICT AIRPLANE AND HELICOPTER TRAFFIC IN THE VICINITY OF RONALD REAGAN WASHINGTON NATIONAL AIRPORT (Aviation Investigation Report AIR-25-01, Mar. 7, 2025), <https://www.ntsb.gov/investigations/AccidentReports/Reports/AIR2501.pdf> [<https://perma.cc/NJG7-BGMX>]; Letter from Jennifer L. Homendy, Chair, NAT’L TRANSP. SAFETY BD., to Christopher J. Rocheleau, Acting Adm’r, FED. AVIATION ADMIN. (Mar. 11, 2025), <https://www.ntsb.gov/safety/safety-recs/recletters/A-25-001-002.pdf> [<https://perma.cc/9277-RK48>]; Jennifer L. Homendy, Chair, NAT’L TRANSP. SAFETY BD., Written Testimony Before the S. Comm. on Commerce, Sci. & Transp., Subcomm. on Aviation: *NTSB Preliminary Report—The DCA Midair Collision* (Mar. 27, 2025), <https://www.ntsb.gov/news/Testimony/Pages/Homendy20250327.aspx> [<https://perma.cc/DF4L-KBEL>] (emphasizing the role of TCAS in collision avoidance at DCA).

⁹⁶Consistent with the small-UAS ceiling. See 14 C.F.R. § 107.51(b) (2025) (general 400-ft AGL limit).

regulatory gap rather than contradicting any express term in the companies’ agreements or in federal aviation regulations. It also clears the *Legality Screen*, as a pro-social, safety-enhancing rule raises no public policy concerns.

Turning to the weighted factors, the practice is a strong candidate for recognition:

- *Foreseeability* is high; the pattern was readily observable in public, broadcast Remote ID message elements, making it reasonably foreseeable to any sophisticated operator.
- *Regularity* is also high; an audit of the parties’ cryptographically time-stamped and hashed flight logs shows that over 99 percent of flights within the corridor adhered to the altitude rule during the 30-day window.
- *Attribution* is clear; the drones’ flight paths are directly attributable to the fleet management software deployed by each company, as evidenced by their deployment manifests and configuration histories.
- *Materiality* is compelling; an event study of the telemetry data shows a statistically significant reduction in near-miss alerts after the practice stabilized (e.g., from an average of about 4 alerts per day pre-stabilization to about 0.3 post-stabilization), directly linking the practice to improved safety outcomes.
- *Verifiability* is straightforward, as the practice can be confirmed through the tamper-evident logs sealed with cryptographic timestamps per RFC 3161 and even reproduced in a controlled replay using the parties’ preserved data.

Given that the practice satisfies all seven criteria, the court would recognize the “altitude-for-direction” practice as a functional norm. In the context of FleetB’s tort claim, this recognized norm would serve as powerful, probative evidence of the standard of reasonable care. FleetA’s deviation, traced to a specific faulty software push, would constitute strong evidence of a breach, likely leading to a finding of negligence.

The outcome would be different, however, if an FAA directive had explicitly prohibited such altitude differentiation. In that scenario, the practice would fail the *Legality Screen* because the doctrine subordinates emergent practices to mandatory public law—express regulatory commands

trump informal operational patterns, regardless of their safety benefits. The companies would need to seek a waiver or regulatory change, not judicial recognition of a non-compliant practice, and recognition would be denied.

This case demonstrates the doctrine’s core value: allowing the law to see and enforce a benign, safety-enhancing practice that emerged organically, thereby promoting accountability and encouraging safer innovation.

B. Case Study 2: A Course of Performance Versus Express Terms

Our second case study moves from tort to commercial relationships, demonstrating how the *F-NORM test* distinguishes between interpreting ambiguous contractual terms and contradicting express terms.⁹⁷ It showcases the critical function of the *Consistency* gate, which upholds the primacy of private ordering by ensuring emergent practices can clarify ambiguity but not override carefully negotiated language.

Two sophisticated firms—*Orchestrio*, a cloud platform, and *FinCo*, a quantitative trading firm—have a master service agreement for high-priority compute. The SLA guarantees Gold-tier workloads priority over Silver and Bronze. Crucially, it states that within the Gold tier, task scheduling shall be “*first-come, first-served where feasible.*” The contract contains a standard merger clause

⁹⁷This scenario is grounded in cloud computing Service Level Agreements (SLAs), which often guarantee high-level performance metrics while leaving granular implementation details, like scheduler behavior, unspecified. See, e.g., AMAZON WEB SERVS., *Amazon Compute Service Level Agreement* (last updated May 25, 2022), <https://aws.amazon.com/compute/sla/> [<https://perma.cc/SX7H-BXH6>] (promising availability and credits but silent on intra-tier scheduling); GOOGLE CLOUD, *Compute Engine Service Level Agreement* (last modified Mar. 4, 2025), <https://cloud.google.com/compute/sla> [<https://perma.cc/BF5X-EH26>] (similar). The “age-based fair-share” routine is analogous to real-world schedulers designed to prevent tenant starvation and ensure predictable performance. See LINUX KERNEL DOCUMENTATION, *The CFS Scheduler* (last visited Dec. 20, 2025), <https://www.kernel.org/doc/html/latest/scheduler/sched-design-CFS.html> [<https://perma.cc/TY5C-T3CF>]; APACHE HADOOP, *Fair Scheduler* (last visited Dec. 20, 2025), <https://hadoop.apache.org/docs/stable/hadoop-yarn/hadoop-yarn-site/FairScheduler.html> [<https://perma.cc/DY79-C757>] (sharing resources so applications receive equal shares over time); SCHEDMD, *Multifactor Priority Plugin* (last visited Dec. 20, 2025), https://slurm.schedmd.com/priority_multifactor.html [<https://perma.cc/GR3U-HFP9>] (increasing priority with queue wait time/“age” to prevent starvation). The legal analysis draws directly on the U.C.C.’s framework for using course of performance to interpret ambiguous terms. See U.C.C. § 1-303 (AM. L. INST. & UNIF. L. COMM’N 2018) (hierarchy of practice; course of performance relevant to meaning); *id.* § 2-202 cmt. 2 (parol evidence; usage/course may supplement unless “carefully negated”); *id.* § 2-209 (Modification, Rescission and Waiver); RESTATEMENT (SECOND) OF CONTRACTS §§ 219–223 (AM. L. INST. 1981); *Nanakuli Paving & Rock Co.*, 664 F.2d at 780–91; *Columbia Nitrogen*, 451 F.2d at 8–10.

but no specific disclaimer of course-of-performance evidence. After deployment, *Orchestrio*'s AI scheduler converges on an emergent “age-based fair-share” routine during periods of high congestion. Rather than letting one tenant’s large job block the queue, the system rotates small time-slices among all concurrent Gold-tier tenants. Over ninety trading days, both parties’ agents adapt to this rhythm. *FinCo*'s systems learn to break large jobs into micro-batches timed to the scheduler’s cadence, and queue logs show that during peak load, Gold-tier jobs alternate in a stable, predictable pattern.

On a volatile market morning, *FinCo*'s critical risk-control batch is slightly delayed because the fair-share routine allocates quanta to competing Gold-tier tenants who also surged. *FinCo* sues for breach of contract, arguing the SLA guaranteed strict FCFS ordering. *Orchestrio* counters that the “where feasible” language is ambiguous during congestion and that the emergent practice is a functional norm that interprets that ambiguity, established through a clear course of performance.

The analysis hinges on the two gates. The practice easily clears the *Legality Screen*. The core question is *Consistency*, which requires distinguishing between contradiction and interpretation. Here, the practice does not contradict the SLA; it gives specific, operational meaning to the ambiguous qualifier “where feasible.” It fills the gap created by that term, defining what is “feasible” during peak contention. Moreover, waiver principles—codified in U.C.C. § 2-209 for goods and recognized under analogous common law doctrines for services—support recognition: ninety days of undisputed, relied-upon conduct can operate as a tacit modification or waiver, redefining what the parties understood their obligations to be.⁹⁸

The weighted factors also support recognition:

- *Foreseeability* is high, as the alternating pattern was visible in tenant-level telemetry and API performance dashboards.
- *Regularity* is high within the ninety-day transactional window, as queue logs show a stable alternation during every peak load.

⁹⁸See U.C.C. § 2-209(4) (“[A]n attempt at modification or rescission [that] does not satisfy [requirements] can operate as a waiver.”); *Wisconsin Knife Works*, 781 F.2d at 1287–88 (repeated acceptance of late deliveries waived timely-delivery term despite no-waiver clause).

- *Attribution* points to both parties: to *Orchestrio* for deploying and maintaining the scheduler (via deployment manifests and versioned configuration files), and crucially to *FinCo* for its agents’ knowing adaptation, which solidifies the pattern as a mutual course of performance.
- *Materiality* is clear from the latency-smoothing effects, shown by time-series analyses of reduced tail latency and avoidance of tenant starvation when the routine is active.
- *Verifiability* is robust via hash-anchored queue logs and versioned configuration files; an independent expert could replay the window using sealed scheduler binaries and FinCo’s job-submission patterns to confirm the fair-share rotation and adaptation.

A court would likely recognize the fair-share pattern as a functional norm, using it not to rewrite the contract, but to *interpret* the ambiguous “where feasible” language. The months-long, undisputed pattern of interaction established a course of performance that defined the parties’ obligations under congestion, leading to a finding that *Orchestrio* did not breach the agreement.

This outcome, however, hinges entirely on the contract. If the SLA had instead read: “Within the Gold tier, scheduling shall be *strict FCFS at all times, with no rotation, time-slicing, or fair-share allocation under any circumstances. No contrary course of performance or usage of trade shall modify this term,*” the outcome would flip. Such language removes any ambiguity and constitutes the “careful negation” required by the U.C.C. to exclude such evidence.⁹⁹ The emergent practice would now directly contradict a clear, “carefully negated” express term. The *Consistency* gate would bar its recognition for interpretive purposes, and *Orchestrio* would be in breach.

This case demonstrates how the doctrine adapts traditional contract principles for algorithmic agents: course of performance remains powerful interpretive evidence, but only when it clarifies rather than contradicts express terms. The *Consistency* gate ensures that the doctrine respects the hierarchy established in U.C.C. § 1-303(e), subordinating emergent practice to negotiated language while allowing them to fill gaps the parties left open.

⁹⁹See U.C.C. § 2-202 cmt. 2; *Columbia Nitrogen*, 451 F.2d at 9–10 (holding that a general merger clause was insufficiently specific to bar usage evidence).

C. Case Study 3: An Ambiguous Marketplace Practice

Our third case study demonstrates the doctrine’s nuance in a more ambiguous setting, where the F-NORM factors are mixed and the outcome is not a simple pass or fail.¹⁰⁰ In a high-speed digital advertising auction run by a major platform, the AI-powered bidding agents of several large advertisers converge on a “deadline-edge bidding” practice. They all learn to submit their final, decisive bids just milliseconds before the exchange’s response timeout (‘tmax’). This practice is an emergent, rational response that excludes smaller players like *AdCo*, whose off-the-shelf bidding agent cannot compete. The exclusion operates technically: large bidders gain an edge through co-location and preferential network peering to minimize latency, allowing their agents to use the full ‘tmax’ window for complex valuation before bidding at the last moment. *AdCo*, lacking these advantages, must bid early with less information or risk its bid arriving too late. *AdCo* sues the platform under state unfair competition law, alleging the emergent practice constitutes an unfair trade practice.

The claimed norm is evaluated over the three-month period preceding *AdCo*’s complaint. *AdCo* claims that the “deadline-edge bidding” practice is a functional norm and that the platform, by designing a system where such a norm could emerge, has a duty to mitigate its exclusionary

¹⁰⁰This vignette is grounded in the market microstructure of real-time bidding (RTB) for digital advertising. The “deadline-edge” bidding strategy is a rational response to exchange protocols that specify a maximum response time (‘tmax’). See IAB TECH LAB, OPENRTB 2.6, § 3.2.1 (Object: BidRequest) (Apr. 2022), https://iabtechlab.com/wp-content/uploads/2022/04/OpenRTB-2-6_FINAL.pdf [<https://perma.cc/56AS-QHSL>] (defining ‘tmax’ as “Maximum time in milliseconds the exchange allows for bids to be received including Internet latency to avoid timeout”). Platforms emphasize latency reduction as a best practice. See GOOGLE FOR DEVELOPERS, *Real-time Bidding: Latency Restrictions and Peering* (last visited Feb. 4, 2026), <https://developers.google.com/authorized-buyers/rtb/get-started/peer-guide> [<https://perma.cc/WG82-ZQ64>] (recommending network peering to minimize latency; 85% of responses must arrive within ‘tmax’ to avoid throttling). The resulting exclusionary effects on smaller players parallel concerns raised in major antitrust investigations. See, e.g., Am. Compl. ¶¶ 1–10, 395, *Texas v. Google, LLC*, No. 4:20-cv-00957 (E.D. Tex. Mar. 15, 2021), <https://www.texasattorneygeneral.gov/sites/default/files/images/admin/2021/Press/Redacted%20Amended%20Complaint%20FILED%20%28002%29.pdf> [<https://perma.cc/CGC7-P6W9>] (alleging auction design disadvantaged rivals, including “Last Look”-type advantages); COMPETITION & MKTS. AUTH., ONLINE PLATFORMS AND DIGITAL ADVERTISING: MARKET STUDY FINAL REPORT 5–12 (July 1, 2020) (U.K.), https://assets.publishing.service.gov.uk/media/5efc57ed3a6f4023d242ed56/Final_report_1_July_2020_.pdf [<https://perma.cc/UQ4M-47SB>] (documenting market power and conflicts in ad tech). Agencies have recently highlighted AI-market design risks and exclusionary tactics. See U.S. DEP’T OF JUSTICE et al., JOINT STATEMENT ON COMPETITION IN GENERATIVE AI FOUNDATION MODELS AND AI PRODUCTS (July 23, 2024), <https://www.ftc.gov/news-events/news/press-releases/2024/07/ftc-doj-international-enforcers-issue-joint-statement-ai-competition-issues> [<https://perma.cc/6GQY-JTBX>].

effects. The legal question is whether this practice should be recognized as a norm, and if so, what legal consequences follow.

The practice would likely pass the *Consistency* gate, as the platform's terms of service are probably silent on specific bidding cadences. The *Legality Screen*, however, requires a graduated analysis. The practice is not per se illegal collusion, as there is no evidence of an agreement among bidders. However, it raises significant fairness concerns that exist in a gray area. If the platform has sufficient market power, inducing a practice that creates *de facto* barriers to entry could implicate §2 of the Sherman Act's prohibition on monopolization or attempted monopolization. Furthermore, if the platform marketed its auction as providing "fair and equal access" while its architecture knowingly induced this exclusionary dynamic, the conduct could attract FTC enforcement as an unfair or deceptive practice under Section 5 of the FTC Act (which provides no private right of action). These concerns counsel against full recognition.

The weighted factors also reveal ambiguities:

- *Foreseeability* is borderline and asymmetric: the pattern was highly foreseeable to the platform operator who designed the system, but perhaps not reasonably foreseeable to a smaller participant like *AdCo* lacking the resources to analyze market microstructure.
- *Regularity* is present but transient; timing histograms show clear clustering in the last quantiles of 'tmax' during peak periods.
- *Attribution* is ambiguous, presenting a question of shared responsibility: is the practice attributable solely to the bidders whose agents adopted the strategy, or is it co-attributable to the platform whose design choices (e.g., tight 'tmax', throttling rules, and latency patterns) induced it?
- *Materiality* and *Verifiability* are clearly met, as the practice directly determines auction winners and can be confirmed through the platform's comprehensive auction logs.

Given these mixed results, the court would likely decline to grant full recognition. Instead, consistent with the graduated-recognition framework discussed in *Part III.C*, it would grant *narrow recognition*, treating the practice's existence as *probative evidence* in an unfair trade practice

analysis. In practice, this means that the court would find the practice’s emergence a highly foreseeable consequence of the platform’s market design, giving rise to a duty to mitigate the exclusionary effects. In private litigation, this may lead to injunctive remedies such as ordering the platform to offer non-discriminatory peering or to randomize bid submission times with micro-jitter to level the playing field; FTC enforcement could yield similar relief along with broader transparency and audit requirements.

The outcome would shift decisively, however, if evidence showed the platform operated its own demand-side bidding agent that disproportionately benefited from the deadline-edge dynamic. Such self-preferencing would solidify Attribution to the platform and shift the *Legality Screen* analysis from borderline to clear failure: the practice is no longer an organic emergence but an engineered outcome designed to advantage the platform’s own agent while disadvantaging competitors.¹⁰¹ Such conduct constitutes exclusionary behavior under § 2 and is categorically barred from recognition.

This case demonstrates that the doctrine is not mechanical. When F-NORM factors point in multiple directions, courts retain discretion to calibrate legal effects, granting narrow recognition for evidentiary purposes while withholding the full interpretive and gap-filling effects available in clearer cases.

D. Case Study 4: A Collusive Practice

Our final case study demonstrates the doctrine’s most critical safeguard: the *Legality Screen*.¹⁰²

¹⁰¹ See States’ Second Am. Compl. ¶¶ 236–45, 376–84, 409–11, *Texas v. Google LLC*, No. 4:20-cv-00957 (E.D. Tex. Aug. 4, 2021); States’ Third Am. Compl. ¶¶ 376–84, *In re Google Digit. Advert. Antitrust Litig.*, No. 1:21-md-03010 (S.D.N.Y. Jan. 14, 2022) (alleging Google manipulated its ad exchange to favor its own demand-side platform through “Last Look” and similar information advantages).

¹⁰² This vignette is directly inspired by ongoing antitrust litigation alleging that revenue-management software facilitates competitor coordination. The central allegations in *In re RealPage, Rental Software Antitrust Litig. (No. II)*, MDL No. 3071 (M.D. Tenn.), allege that landlords delegated pricing to a common algorithm that used non-public competitor data to stabilize rents at supracompetitive levels. See Statement of Interest of the United States, *supra* note 6. The MDL court denied in part motions to dismiss in two separate opinions. See *In re RealPage*, 709 F. Supp. 3d 478 (M.D. Tenn. 2023); *In re RealPage*, 709 F. Supp. 3d 544 (M.D. Tenn. 2023). Separately, the United States and several states filed a civil antitrust complaint against RealPage, later amending to add certain landlord defendants. See Complaint, *United States v. RealPage*, No. 1:24-cv-00710 (M.D.N.C. Aug. 23, 2024); Am. Compl., *RealPage*, No. 1:24-cv-00710 (M.D.N.C. Jan. 7, 2025). This scenario is distinct from—but related to—cases involving

Several competing online retailers all subscribe to the same third-party dynamic-pricing service. The service’s algorithm, fed real-time inventory and demand data from its subscribers, adjusts prices based on competitors’ actions to avoid price wars and stabilize the market. Over several months, the retailers’ prices for a popular product begin to move in near-perfect lockstep, stabilizing at supracompetitive levels. A consumer class action alleges that the emergent price-matching behavior constitutes unlawful price-fixing under § 1 of the Sherman Act.

The retailers respond that the price alignment is not the product of an agreement but is instead a lawful, independently arrived-at market equilibrium. Each firm, they contend, acted unilaterally in subscribing to the service and was free to accept or reject its recommendations. There was no “meeting of the minds,” no explicit cartel, and no coercion—just rational firms employing cutting-edge pricing technology. They ask the court to recognize the practice as a functional norm arising organically from autonomous agents reacting to competitive market signals, analogizing it to the benign coordination in Case Study 1.

Analysis begins—and ends—with the *Legality Screen*.¹⁰³ The dispositive question is economic function, not technical form. The retailers’ conduct has the function and effect of a horizontal price-fixing agreement—the paradigmatic *per se* violation of antitrust law.¹⁰⁴ Independent of intent, the algorithm serves as a “hub” through which competitors share non-public, competitively sensitive information and coordinate their conduct: an unlawful agreement can be inferred from the circumstantial evidence, including parallel conduct *plus* additional factors that tend to exclude

explicit price-fixing implemented via algorithms. *See, e.g.*, Information at 3–4, *Topkins*, No. 3:15-cr-00201. *See also Socony-Vacuum*, 310 U.S. 150 (per se price-fixing); *Apple*, 791 F.3d 290 (hub-and-spoke coordination); *Container Corp.*, 393 U.S. 333 (information exchange); *Text Messaging*, 782 F.3d 867 (parallelism accompanied by plus factors).

¹⁰³ For expository clarity, we treat this scenario as a clear failure of the Legality Screen. The case law on “algorithmic collusion,” however, remains unsettled and highly sensitive to the pleadings. *Compare* *Duffy v. Yardi Sys.*, 758 F. Supp. 3d 1283 (W.D. Wash. 2024) (denying motion to dismiss and applying per se treatment to alleged horizontal price-fixing facilitated by revenue-management software), *and In re RealPage, Rental Software Antitrust Litig.* (No. II), 709 F. Supp. 3d 478 (M.D. Tenn. 2023) (denying in part and addressing reasonableness under the rule of reason), *with* *Cornish-Adebiyi v. Caesars Entm’t, Inc.*, No. 1:23-cv-02536, slip op. (D.N.J. Sep. 30, 2024) (dismissing for lack of plausible agreement and nonpublic-data exchange). *See generally* Kate Ross & Amy Vegari, *Algorithmic Price-Fixing Cases Reflect Exacting Pleading Standard*, JD SUPRA (Aug. 25, 2025), <https://www.jdsupra.com/legalnews/algorithmic-price-fixing-cases-reflect-2043614/> [<https://perma.cc/H2DV-NP65>] (surveying divergent outcomes across RMS cases).

¹⁰⁴ *See Socony-Vacuum*, 310 U.S. 150.

independent action (e.g., shared non-public data, adherence to recommendations designed to “eliminate price wars”).¹⁰⁵ While the retailers might marshal evidence for the weighted factors—showing the price alignment was *foreseeable* from vendor dashboards, highly *regular* in the price histories, *material* to their profits, *verifiable* through logs, and *attributable* to their collective decision to delegate pricing authority to a common algorithm—the inquiry is foreclosed at the threshold. The court denies recognition—not because the pattern is not “functional” for the retailers, but because the function it serves is illegal.

The contrast with Case Study 1 is instructive. There, the practice solved a *coordination* problem (avoiding collisions) without harming welfare; here, it solves a *cooperation* problem (restraining competition) that harms consumers. Coordination that *enhances* welfare may warrant recognition; coordination that *suppresses* competition to extract rents must be condemned. By embedding the *Legality Screen* as a non-negotiable gate, the framework remains subordinate to public law: emergent practices that function as price-fixing, discriminatory rules, or other unlawful conduct are categorically denied recognition.

V. OPERATIONALIZING THE DOCTRINE: FROM COURTROOM TO CODE

A legal doctrine remains inert unless operationalizable by those it governs. This Part moves from the analytical framework developed in Parts I–IV—detailing the evidentiary substrate, *F-NORM test*, and case-study applications—to practical implementation, demonstrating that the doctrine is administrable through existing legal tools rather than requiring novel institutional capacities.

The discussion proceeds in three parts. Section A outlines evidentiary and procedural mechanisms for litigation, emphasizing rapid preservation, structured production, authentication, and

¹⁰⁵See *supra* note 81; see also *Text Messaging*, 782 F.3d at 872–75 (7th Cir. 2015) (explaining plus factors framework). Even where courts apply the rule of reason at early stages for novel algorithmic mechanisms, see *In re RealPage*, 709 F. Supp. 3d at 519–20 (M.D. Tenn. 2023), the functional equivalent of horizontal price-fixing remains unlawful *per se*.

expert analysis, with courts resolving the Consistency and Legality gates as threshold matters (often at summary judgment) before addressing the five weighted factors. Section B provides model contractual language for private ordering, enabling parties to adopt, disclaim, or govern emergent practices *ex ante*. Section C offers pathways for public governance, including model legislative provisions and regulatory guidance to police harmful norms while preserving benign coordination.

A. Evidentiary and Procedural Mechanisms

The *F-NORM test* is an evidence-driven inquiry. Its application depends on the ability of parties and courts to access, authenticate, and interpret the digital records of agent interactions.

The first procedural step in any dispute involving a claimed functional norm is therefore the *preservation* of the “evidentiary substrate” detailed in *Part I*. Given the often-ephemeral nature of system logs and telemetry data—where retention windows are often measured in days, not years—parties must be prepared to issue immediate and specific litigation holds that cover not only communications but also interaction logs, configuration files, and model version histories. The risk of spoliation is acute and the sanctions potentially severe. Under Rule 37(e), a court may order measures to cure prejudice from lost ESI; harsher sanctions—including adverse-inference instructions or case-dispositive relief—require a finding that the party acted with intent to deprive another of the information’s use.¹⁰⁶

Existing e-discovery rules are readily adaptable to algorithmic agents. Under the Federal Rules of Civil Procedure, parties can and should request the production of data in structured, machine-readable formats (e.g., JSON or Parquet files with clear data dictionaries), rather than static forms like screenshots or PDFs.¹⁰⁷ Proportionality principles can be used to scope requests to

¹⁰⁶ See, e.g., Fed. R. Civ. P. 37(e); Fed. R. Civ. P. 37(e) advisory committee’s note to 2015 amendment (explaining spoliation measures and sanctions); *Zubulake*, 229 F.R.D. at 430–31 (S.D.N.Y. 2004) (establishing the duty to issue a litigation hold and outlining sanctions for spoliation).

¹⁰⁷ See Fed. R. Civ. P. 34(b)(1)(C) (allowing a request to specify the form for ESI production); Fed. R. Civ. P. 34(b)(2)(E)(ii) (requiring production in the form ordinarily maintained or a reasonably usable form if not specified); see also *Williams v. Sprint/United Mgmt. Co.*, 230 F.R.D. 640, 652 (D. Kan. 2005) (ordering production of spreadsheets in their native format with metadata).

a specific transactional window and agents at issue, preventing burdensome “fishing expeditions” while ensuring that data necessary for F-NORM analysis is produced.¹⁰⁸ Protective orders under Rule 26(c) and non-waiver orders under Federal Rule of Evidence 502(d) can be used to safeguard sensitive information.¹⁰⁹

Given the technical complexity, expert witnesses will be central. Data scientists or computational social scientists can serve as “translators” for the court, authenticating records under Rules 901(b)(9) and 902(13)–(14) and analyzing the evidentiary substrate using methods tied directly to the F-NORM factors: time-series analysis for *Regularity*, event studies for *Materiality*, configuration histories for *Attribution*, public telemetry for *Foreseeability*, and deterministic replays for *Verifiability*.¹¹⁰ This testimony would be subject to the reliability standards of Rule 702 and *Daubert*.¹¹¹ In particularly complex or contentious cases, courts may also appoint a neutral expert under Rule 706 or a special master under Rule 53 to assist the fact-finder.¹¹²

To manage complex cases, courts could adopt a staged adjudication process. A court might first decide the threshold “gate” questions of *Consistency* and *Legality* at the summary judgment phase, as these often turn on legal rather than factual analysis. If the claimed norm passes these gates, the more fact-intensive “weighted” factors could proceed to trial with the proponent bearing the burden of proving each of the seven factors by a preponderance of the evidence.

¹⁰⁸ See Fed. R. Civ. P. 26(b)(1).

¹⁰⁹ These tools are central to modern e-discovery. Rule 26(c) protective orders can limit who may view sensitive data and for what purpose, while a Rule 502(d) order allows parties to produce privileged material without effecting a waiver of privilege in the instant case or any other federal or state proceeding. See Fed. R. Civ. P. 26(c); Fed. R. Evid. 502(d); see also THE SEDONA PRINCIPLES, THIRD EDITION: BEST PRACTICES, RECOMMENDATIONS & PRINCIPLES FOR ADDRESSING ELECTRONIC DOCUMENT PRODUCTION, 19 SEDONA CONF. J. 1, cmt. 12.a (2018) (endorsing the use of 502(d) orders to facilitate discovery).

¹¹⁰ The admissibility of the underlying records is typically governed by the business records exception, with authentication provided under rules for evidence describing a process or system and for self-authenticating electronic records. See Fed. R. Evid. 803(6), 901(b)(9), 902(13)–(14), 1006.

¹¹¹ See *Daubert*, 509 U.S. 579 (1993); *Kumho Tire*, 526 U.S. 137 (1999) (extending the gatekeeping function to all expert testimony); Fed. R. Evid. 702 (as amended Dec. 1, 2023).

¹¹² See Fed. R. Evid. 706; Fed. R. Civ. P. 53; see also FED. JUDICIAL CTR. & NAT’L RSCH. COUNCIL, REFERENCE MANUAL ON SCIENTIFIC EVIDENCE 103–18 (3d ed. 2011) (discussing the court’s authority and the practical considerations for appointing neutral experts under Rule 706 and special masters under Rule 53 to manage technically complex evidence).

B. Private Ordering: Model Contractual Language

While procedural tools are essential for resolving disputes *ex post*, the most efficient way to manage emergent practices is through clear, *ex ante* agreements. The doctrine of functional norms is designed to serve as a default rule that sophisticated parties can, and should, contract around to suit their specific needs and risk tolerances. The choice of contractual posture depends on the context, ranging from full adoption in adaptive environments to a complete disclaimer where strict certainty is paramount. Table 5 provides a high-level guide for selecting a contractual posture, followed by model clauses that form a modular toolkit. Drafters can mix and match an adoption clause, a disclaimer, and a governance process depending on the specific context of their transaction.

Table 5: A Guide to Selecting a Contractual Posture

Posture	Use Case / Environment	Key Objective
Adoption	Adaptive, dynamic relationships (e.g., joint ventures, cloud services)	Leverage beneficial coordination and fill contractual gaps.
Disclaimer	Safety-critical or high-stakes systems (e.g., medical devices, aviation)	Ensure strict certainty and control; prevent informal modification.
Governance	Long-term, evolving partnerships (e.g., supply chain, ad tech)	Create a process for transparency, notice, and formal recognition.

The first approach is to formally adopt the doctrine as the standard for recognizing emergent practices that can interpret or supplement the agreement.

Model Clause 1 (Adoption): “The Parties agree that an emergent practice between their automated agents shall be recognized as a ‘Functional Norm’ capable of interpreting or supplementing this Agreement if it satisfies the seven criteria of the F-NORM test. The party asserting the norm bears the burden of proof. The recognition of a Functional Norm shall not contradict an express term of this Agreement or violate mandatory public law.”

Conversely, parties in a safety-critical context may wish to opt out entirely. To achieve this, they

must draft a disclaimer that meets the Uniform Commercial Code’s high standard for excluding practice evidence.¹¹³ The U.C.C. requires that such practices be “carefully negated”; a generic merger clause may be insufficient against a strong course of performance.¹¹⁴ The following clause is designed to meet that higher standard.

Model Clause 2 (Disclaimer): “This Agreement constitutes the entire and exclusive understanding between the Parties. The Parties expressly agree that this Agreement may not be supplemented, explained, or modified by any course of performance, course of dealing, usage of trade, or emergent practice arising from automated agents. Any modification must be in a signed writing that expressly references this clause.”

A third, middle-ground option is a procedural approach that establishes clear requirements for transparency and governance, creating a formal pathway for recognizing important norms.

Model Clause 3 (Procedural Governance): “Each Party agrees to maintain and, upon reasonable request in the event of a dispute, provide access to time-stamped interaction logs and versioned configuration files for its automated agents. The Parties agree to meet quarterly to review and formally document any persistent, material emergent practices that may affect performance under this Agreement. No emergent practice shall be deemed to modify this Agreement absent such formal documentation.”

This toolkit of model clauses allows parties to proactively manage the legal implications of fluid agency. They can be strengthened further, as needed, with specific provisions governing logging standards, change-control notification, and audit rights tailored to the industry. By choosing to adopt, reject, or procedurally govern emergent practices, they can tailor the legal framework to the realities of their operational environment.

¹¹³In contracts of adhesion, the enforceability of a total disclaimer could be challenged under unconscionability, especially if it shields a dominant party from accountability for relied-upon agent behavior. *See* U.C.C. § 2-302.

¹¹⁴The requirement for specific, “careful negation” of practice evidence is established in U.C.C. § 2-202 cmt. 2. *See, e.g., Columbia Nitrogen*, 451 F.2d at 9–10. Even such specific language, however, may not prevent a finding of waiver through a sustained course of performance. *See* U.C.C. § 2-209(4)–(5); *Wisconsin Knife Works*, 781 F.2d at 1287–88.

C. Public Governance: Model Legislative Language and Regulatory Guidance

While private ordering is powerful, a public law backstop is necessary to ensure consistency, protect third parties, and govern interactions where contracts are absent or adhesive, such as in consumer-facing systems. Codifying the *F-NORM* test in statute would provide a clear and uniform standard for courts within a jurisdiction. Such a provision could be incorporated into a state's commercial code or a federal AI governance act.

Model Statute: (a) Conditions for Recognition. A court may recognize an emergent practice between automated agents as a “Functional Norm” if the proponent proves by a preponderance of the evidence that the practice was: (1) foreseeable to affected parties; (2) regular within the relevant transactional window; (3) attributable to a legally responsible principal; (4) material to outcomes; and (5) verifiable through reliable records. (b) Preclusion. Notwithstanding satisfaction of conditions (1)-(5) in clause (a), recognition shall be denied if the practice: (1) is inconsistent with an express term of a governing agreement; or (2) fails to clear a Legality Screen that bars recognition of anticompetitive or otherwise unlawful conduct (including conduct that violates antitrust, civil-rights, or consumer-protection law). (c) Legal Effect. A recognized Functional Norm may be used as evidence to interpret an agreement, fill a contractual gap, or inform the standard of reasonable care. Recognition is probative, not dispositive, and confers no immunity from liability under other law. (d) No New Cause of Action. This section does not create a private right of action or an independent basis for liability; it governs recognition and evidentiary use only.

Beyond codification, regulatory agencies like the Federal Trade Commission (FTC) can integrate the doctrine into their existing oversight and enforcement mandates. The FTC, for example, could issue enforcement policy guidance stating that business practices deploying or relying upon an emergent practice that is deceptive or unfair to consumers may constitute a violation of

Section 5 of the FTC Act.¹¹⁵ The F-NORM factors provide a ready-made analytical framework for structuring such an investigation, helping the agency distinguish benign coordination from harmful market-wide practices.

Finally, public governance can proactively foster the conditions that make the doctrine administrable. Regulators could, for instance, create an affirmative-defense safe harbor from certain enforcement actions for companies that voluntarily adopt and adhere to high standards of logging, auditing, and transparency, potentially by encouraging adherence to technical standards for AI system logging and auditability developed by bodies like the National Institute of Standards and Technology (NIST).¹¹⁶ This would create a powerful incentive for firms to generate the “evidentiary substrate” that makes the doctrine of functional norms a workable tool for accountability.

VI. OBJECTIONS AND CONCLUSION

With the practical tools in place to make the *F-NORM test* administrable, this Part addresses the doctrine’s administrability and legitimacy. The organizing principle is modesty: recognition of a functional norm is (i) interpretive in contract; (ii) probative, not dispositive, in tort; and (iii) always subordinate to express terms and mandatory public law. With that frame, it addresses the principal objections and closes with a brief conclusion.

A. Addressing Objections

The first objection is one of *indeterminacy*. Critics may argue that the *F-NORM test* is an overly vague, amorphous standard that will lead to unpredictable and inconsistent judicial outcomes,

¹¹⁵See 15 U.S.C. § 45(a)(1), (n) (2022) (prohibiting unfair or deceptive acts or practices and defining unfairness); FED. TRADE COMM’N, Press Release, *FTC Announces Crackdown on Deceptive AI Claims and Schemes* (Sep. 25, 2024), <https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-announces-crackdown-deceptive-ai-claims-schemes> [<https://perma.cc/W2K6-Y2HZ>] (emphasizing AI-related enforcement).

¹¹⁶This approach has analogues in emerging AI legislation. See, e.g., Colo. Rev. Stat. § 6-1-1706(3), (6) (2024) (creating an affirmative defense in attorney general enforcement actions where a deployer discovers and cures a violation and complies with a recognized risk management framework, including NIST AI RMF or ISO/IEC 42001, and providing no private right of action).

thereby chilling innovation. This concern is understandable but misplaced. The doctrine is intentionally a *standard, not a rule*, because the dynamism of AI environments requires flexibility. Any attempt to create rigid, *ex ante* rules for specific behaviors would be brittle and instantly obsolete. The doctrine’s discretion, however, is not unbounded. It is disciplined by the test’s rigorous, evidence-based requirements—five weighted factors (*Foreseeability, Regularity, Attribution, Materiality, Verifiability*) and two threshold gates (*Consistency, Legality*)—each grounded in the concrete, digital records of the evidentiary substrate detailed in *Part I*. This choice mirrors the classic rules-versus-standards tradeoff, where calibrated *ex post* assessment is often preferable to pre-specified rules in rapidly evolving contexts.¹¹⁷

A related objection concerns *democratic legitimacy*: does the doctrine allow “machines to make law?” This critique misinterprets the doctrine’s modest scope. It does not grant algorithmic agents any form of law-making authority. Rather, it allows courts to *recognize* the existence of a widespread practice as a factual matter and to give that fact *calibrated legal effect* for the limited purposes of interpreting private agreements and assessing standards of care. This is continuous with long-standing legal practice: courts routinely recognize “usage of trade” and “course of performance” to interpret agreements, without ever elevating those practices above express terms or mandatory public law.¹¹⁸ The doctrine does for algorithmic practice what the common law has long done for human custom: it recognizes a fact (a stable, functional pattern) and gives it calibrated legal effect (interpretive and evidentiary), without elevating it above bargains or statutes.

A third objection targets *judicial capacity*, suggesting that courts are not technically equipped to adjudicate disputes over complex algorithmic behavior. This argument underestimates the judiciary’s proven ability to adapt to new and complex domains, from the intricacies of patent law to the economic modeling of antitrust analysis. The doctrine does not require judges to become data scientists; it channels technical issues through a mature procedural toolkit—including

¹¹⁷ See Louis Kaplow, *Rules Versus Standards: An Economic Analysis*, 42 DUKE L.J. 557 (1992).

¹¹⁸ See U.C.C. § 1-303; RESTATEMENT (SECOND) OF CONTRACTS § 222 (AM. L. INST. 1981).

the Federal Judicial Center’s *Reference Manual on Scientific Evidence*, court-appointed experts under Rule 706, and special masters under Rule 53—that is designed to help courts manage highly technical records.¹¹⁹ Expert testimony is vetted through the familiar gatekeeping standards of Rule 702, *Daubert*, and *Kumho Tire*, ensuring that the fact-finder receives reliable and comprehensible analysis.¹²⁰ Moreover, the evidentiary substrate—native logs, configuration histories, and telemetry—is authenticated under Rules 901 and 902 and summarized under Rule 1006, as *Part V* details. Technical complexity is therefore not a reason to ignore operative facts that the law can already admit and evaluate.

A practical objection concerns the burden of the evidentiary program. Critics may worry that the doctrine’s logging and audit demands create privacy concerns, conflict with data-minimization and retention-limit principles, and force disclosure of trade secrets and other sensitive operational data. These concerns are real but manageable. The evidentiary program is compatible with data-minimization and confidentiality imperatives. Parties can employ tiered retention windows keyed to transactional risk, minimize fields to what is necessary for attribution and replay, and use protective orders under Rule 26(c) to safeguard sensitive materials through limited access, sealed appendices, and strict confidentiality protocols.¹²¹ As *Part V* explains, proportionality, transactional-window scoping, and structured data production (rather than unfiltered data dumps) reduce burden while preserving probative value.¹²² Where applicable, data-minimization statutes reinforce rather than contradict this approach: the same principles that justify limiting collection to what is reasonably necessary for business purposes also justify limiting forensic disclosure to what is necessary for adjudication.¹²³ In short, courts already balance accuracy with confidentiality; the

¹¹⁹ See *supra* note 112 (detailing the procedural toolkit for managing technically complex evidence, including court-appointed experts under Rule 706 and special masters under Rule 53).

¹²⁰ See *Kumho Tire*, 526 U.S. 137 (1999) (extending *Daubert* gatekeeping to all expert testimony); *Daubert*, 509 U.S. 579 (1993); Fed. R. Evid. 702 (as amended Dec. 1, 2023).

¹²¹ See Fed. R. Civ. P. 26(c) (protective orders, including allocation of costs and restrictions on disclosure).

¹²² See *supra* notes 108, 107.

¹²³ See Regulation (EU) 2016/679 of the European Parliament and of the Council, art. 5(1)(c), 2016 O.J. (L 119) 1, 35 (General Data Protection Regulation) (data minimization); see also Cal. Civ. Code § 1798.100(c) (West 2025) (requiring that a business’s collection, use, retention, and sharing of personal information be reasonably necessary and proportionate to disclosed purposes).

evidentiary tail does not wag the confidentiality dog.

A key institutional concern is *evidence asymmetry*: in platform markets, one side often controls the logs, telemetry, and configuration histories. Does this one-sided custody make the test unfair or unworkable? The doctrine anticipates this asymmetry and supplies remedies. Proportional discovery principles, native-format production requirements, and targeted sampling allow parties to obtain the relevant subset of records without imposing unbounded burdens on custodians.¹²⁴ Adverse inferences under Rule 37(e) deter spoliation, ensuring that custodians preserve the transactional-window evidence once a duty to preserve attaches.¹²⁵ For particularly complex datasets or disputes over technical production, courts can appoint a Rule 53 special master to supervise discovery, conduct neutral replay tests, or resolve methodological disputes.¹²⁶ These are standard, not exotic, judicial capacities—the same tools courts use to manage voluminous ESI and asymmetric technical records in antitrust, securities, and patent cases.

Another objection is the risk of *manipulation*. Could sophisticated actors “game” the system by engineering artificial “norms” to serve as a pretext for liability shields or to create anticompetitive safe harbors? The *F-NORM test* is designed with multiple safeguards to prevent this. First, a manufactured pattern, created without the genuine reliance of a counterparty, would likely fail the *Foreseeability* factor, as it would constitute an unfair surprise—evidence that the pattern was generated by hidden configurations, non-public A/B testing, or other parameters not reasonably discoverable would rebut any claim of foreseeability. Second, it might fail the *Materiality* factor if it had no real transactional function beyond creating a legal artifact; counterfactual analysis showing that the alleged norm did not meaningfully affect outcomes would undermine recognition. Third, and most decisively, the *Legality Screen* is designed precisely to filter out pretextual and harmful behaviors. As the price-fixing case study in *Part IV* demonstrates, an emergent practice that functions as illegal collusion fails at the threshold, regardless of how well it satisfies the other

¹²⁴ See *supra* notes 108, 107.

¹²⁵ See *supra* notes 53, 106.

¹²⁶ See FED. JUDICIAL CTR., MANUAL FOR COMPLEX LITIGATION (FOURTH) §§ 11.52, 22.81, 22.315 (2004) (discussing special masters, sampling, and test cases).

criteria.¹²⁷ The framework is thus built to repel manufactured patterns through layered defenses, each tied to auditable evidence and each reinforcing the doctrine’s subordination to public law.

Some may worry that the doctrine will *stifle innovation* by creating new legal risks that “freeze” development, making companies afraid to update or improve their AI systems lest they disrupt a recognized norm and incur liability. The opposite is more likely true. The doctrine actually *enhances* innovation by increasing predictability. By providing a clear, evidence-based framework for when emergent practices will and will not have legal effect, it allows developers and deployers to better assess and manage their risks *ex ante*. A clear legal framework is a prerequisite for confident innovation, not an obstacle to it: theoretical and empirical work has long linked policy uncertainty to delayed or depressed investment, and a bounded standard that clarifies expectations can reduce these frictions.¹²⁸ Moreover, the doctrine invites private ordering. As *Part V* details, parties can adopt transparency clauses, procedural governance mechanisms, or targeted disclaimers to tailor their exposure.¹²⁹ Where regulators choose to create safe harbors for firms that maintain robust logging and audit trails, the expected compliance path becomes even clearer.¹³⁰ In contract, recognition is interpretive; in tort, it is probative, not dispositive. The doctrine is not a liability trap; it is a tool for clearer expectations.

Others may argue the doctrine is premature—that courts should wait for more cases to develop organically before codifying a recognition framework. This position undervalues the costs of a doctrinal vacuum. Without a coherent framework, early cases will be resolved *ad hoc*, creating conflicting precedents across jurisdictions. Parties cannot structure their conduct against unpredictable judicial improvisation. The F-NORM test provides a common analytical vocabulary—courts retain discretion in applying the factors, but they start from shared ground rather than reinventing the wheel in each case. Moreover, the doctrine is designed to be adopted incrementally:

¹²⁷ See *supra* Part IV.D.

¹²⁸ See, e.g., Scott R. Baker, Nicholas Bloom & Steven J. Davis, *Measuring Economic Policy Uncertainty*, 131 Q.J. Econ. 1593, 1593–1636 (2016), <https://doi.org/10.1093/qje/qjw024> (documenting the chilling effect of policy uncertainty on investment).

¹²⁹ See *supra* Part V.B (model contractual language).

¹³⁰ See *supra* note 116.

courts can apply it case-by-case without statutory codification, and parties can contract into or out of it. The framework enables, rather than forecloses, organic development.

A final policy concern is *jurisdictional variation*. Critics may worry that the doctrine will be applied inconsistently across states and nations, creating forum-shopping opportunities and undermining predictability. This concern overstates the problem. The doctrine’s building blocks already have wide acceptance: UETA and E-SIGN, adopted in nearly every U.S. jurisdiction, recognize that electronic agents can form binding contracts;¹³¹ agency law supplies the attribution framework across all common-law jurisdictions;¹³² and contract and tort law’s practice-recognition doctrines—usage of trade, course of performance, custom as evidence—are foundational features of American law.¹³³ The model statute in *Part V* offers a uniform text for codification, further reducing divergence.¹³⁴ For cross-border transactions, courts can lean on familiar tools: the CISG, for example, already binds parties to usages they “knew or ought to have known,” providing an international template for recognizing emergent practices in transnational commerce.¹³⁵ In sum, the path to convergence is well-marked, and the doctrine builds on widely shared legal foundations rather than inventing novel ones.

B. Conclusion

Algorithmic agents fill the ‘governance vacuums’ with de facto rules—who yields, how to pace, how to ration scarce slots. These emergent practices are the real, operational law of the transaction, yet they are invisible to a legal system built for human custom and written contracts. The law cannot incorporate these rules because it cannot see them.

The doctrine of functional norms provides the necessary lens. It offers a principled, evidence-based framework for determining when an emergent practice is recognized—when it is foreseeable,

¹³¹ See E-SIGN Act, 15 U.S.C. §§ 7001(h), 7006(3); UNIF. ELEC. TRANS. ACT, *supra* note 75, §§ 9, 14; *see also supra* note 76.

¹³² See RESTATEMENT (THIRD) OF AGENCY §§ 1.01–1.03, 2.03 (AM. L. INST. 2006).

¹³³ See *supra* Part II.

¹³⁴ See *supra* Part V.C.

¹³⁵ See United Nations Convention on Contracts for the International Sale of Goods, *supra* note 69, art. 9(2).

regular, attributable, material, and verifiable—and, just as importantly, when it does not, because it contradicts express terms or violates mandatory public law. The doctrine’s effects are carefully calibrated: interpretive in contract, probative but not dispositive in tort, and never a shield for anticompetitive or discriminatory conduct.

Grounded in established procedure—including rules for ESI production, authentication, business records, and expert testimony—the doctrine’s substantive effects are carefully circumscribed. In *contract law*, recognized norms can serve as interpretive and gap-filling evidence. They can inform the meaning of ambiguous terms, supplement silent agreements, and—where the pattern is sustained and relied upon—provide evidence of a course of performance that may operate as a waiver or modification, consistent with the Code’s hierarchy of sources.¹³⁶ In *tort*, recognized norms can serve as probative but not dispositive evidence of the standard of reasonable care. As with human custom, they inform what a reasonable actor would do under the circumstances, but courts retain the authority to find that a norm itself was negligent or that an entire calling has lagged behind reasonable prudence.¹³⁷ In *procedure*, the doctrine channels proof through established evidentiary mechanisms—native-format ESI production, authentication under Rules 901 and 902, business records and summaries under Rules 803(6) and 1006, and expert testimony subject to *Daubert* and Rule 702 gatekeeping.¹³⁸

The doctrine remains subordinate to existing law. It does not override negotiated *express terms*: The *Consistency* gate ensures that a functional norm cannot contradict a clear, “carefully negated” contractual provision. Private ordering retains primacy, and parties can disclaim reliance on emergent practices through specific contractual language.¹³⁹ It also does *not* sanitize *unlawful conduct*: The *Legality Screen* bars recognition of practices that violate antitrust law, civil rights statutes, consumer protection rules, or other mandatory public law. An emergent practice that functions as price-fixing, facilitates discrimination, or otherwise serves illegal ends fails at the

¹³⁶ See U.C.C. §§ 1-303(e), 2-209; see also *supra* notes 60, 63.

¹³⁷ See *The T.J. Hooper*, 60 F.2d at 740.

¹³⁸ See *supra* Part V.A; see also *supra* notes 45, 111.

¹³⁹ See U.C.C. § 2-202 cmt. 2; see also *supra* Part V.B (model disclaimer clauses).

threshold, regardless of its regularity or verifiability.¹⁴⁰ It does *not* confer legal personhood on AI— an issue fraught with complex philosophical and legal questions.¹⁴¹ Neither does it create immunity from liability. Recognition is evidentiary, not immunity-granting; a recognized norm informs interpretation and provides probative evidence, but it does not shield actors from responsibility for harmful outcomes or transform machines into legal subjects with rights or duties.

Operationally, the framework is administrable. Courts can demand targeted, native-format productions for the relevant transactional window, resolve the two gates at summary judgment, and instruct juries that a recognized norm is evidence of care, not its measure. Parties can adopt, limit, or procedurally govern practices through the model contractual clauses in *Part V*. Regulators can target harmful equilibria through the *Legality Screen* and can incentivize logging and transparency to make accountability possible.

The alternative is a false choice between two flawed extremes: *under-recognition*, which treats emergent practices as legal nullities, undermining fairness and certainty in algorithmic markets; and *over-recognition*, which risks laundering harmful or collusive conduct as legitimate “custom.” The doctrine of functional norms avoids both. It equips courts and regulators to identify and engage with the *de facto* rules that govern AI interactions. By allowing the law to see the operative practices of algorithmic agents without elevating them above express contracts or mandatory public law, it balances innovation with accountability.

¹⁴⁰ See *supra* notes 80, 104, 82, 83; see also *supra* Part IV.D (price-fixing case study).

¹⁴¹ See, e.g., Joanna J. Bryson, *Robots Should Be Slaves*, in *CLOSE ENGAGEMENTS WITH ARTIFICIAL COMPANIONS* 63, 63–74 (Yorick Wilks ed., John Benjamins Publ’g Co. 2010); RYAN ABBOTT, *THE REASONABLE ROBOT: ARTIFICIAL INTELLIGENCE AND THE LAW* (Cambridge Univ. Press 2020); Katherine B. Forrest, *The Ethics and Challenges of Legal Personhood for AI*, 133 *YALE L.J. F.* 1175 (2024).