

Engineering Rivalry: A ‘Duty to Disrupt’ Algorithmic Convergence in Two-Sided Agentic AI Markets

Anirban Mukherjee
Hannah Hanwen Chang

February 19, 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

Section 1 of the Sherman Act is built on a fundamentally anthropocentric premise: liability attaches only to a human “agreement.” This principle misaligns with *algorithmic convergence*, whereby learning systems independently reach and sustain supracompetitive prices without a “meeting of the minds.” Judicial workarounds—such as the hub-and-spoke model and the doctrine of conscious parallelism—prove inconsistent and limited. The doctrinal gap widens in *two-sided agentic AI markets*, where autonomous buyer agents bargain with seller agents, giving rise to what this Article calls *agentic convergence*. In these ecosystems, prices emerge from interacting designs rather than unilateral choices, and myopic buyer agents can inadvertently stabilize tacit collusion, further weakening the inference from supracompetitive effects to conspiratorial intent.

This Article proposes shifting antitrust from policing conspiracies to engineering rivalry. It articulates a *Duty to Disrupt*—a legally enforceable design obligation for developers and deployers of market-facing agents in designated agentic markets. Compliance requires *rivalry-by-design*: (1) long-horizon objectives that penalize convergence; (2) strategic randomness that resists exploitation and deters mimicry; and (3) market-state awareness that triggers procompetitive responses, such as amplifying rewards for defectors and blunting retaliatory punishments.

The duty is operationalized through three pillars: (1) *structural separation* (a “Glass-Steagall for agents”) to prevent conflicts when the same firm fields both buyer- and seller-side agents; (2) *proactive oversight* via mandatory algorithmic audits and regulatory sandboxes; and (3) *evidentiary reform* that creates a rebuttable presumption of design-based causation, shifting the burden of production to developers to produce design records rebutting the inference that their agent’s architecture caused the anticompetitive harm.

Case studies in airlines, hotels, rentals, and retail illustrate how this framework supplies a tractable locus of accountability where the agreement requirement erodes—making competition a *deliberate, auditable* property of agent architecture.

Keywords: Algorithmic Convergence, Two-Sided Agentic AI Markets, Sherman Act, Antitrust Law.

JEL codes: K21, K20, K41, K42, L86, O33.

TABLE OF CONTENTS

INTRODUCTION	4
I ONE-SIDED ALGORITHMIC CONVERGENCE: DOCTRINAL CHALLENGES	9
A THE STRAIN ON AGREEMENT INFERENCE	10
B HUB-AND-SPOKE MODEL	11
C CONSCIOUS PARALLELISM	13
II AGENTIC CONVERGENCE: A NEW PARADIGM AND DOCTRINAL CONSEQUENCES	14
A DEFINING THE TWO-SIDED AGENTIC MARKET	15
B THE ECONOMIC MECHANISM: DESIGN-CONTINGENT COMPETITION	17
C THE OBSOLESCENCE OF TRADITIONAL ANTITRUST DOCTRINE	19
D THE RESULTING EVIDENTIARY VOID AND THE PIVOT TO DESIGN-BASED LIABILITY	26
E THE NECESSITY OF AN ECOSYSTEM-WIDE DUTY	29
III ENGINEERING RIVALRY: A NEW ANTITRUST FRAMEWORK FOR AGENTIC CONVERGENCE	30
A PILLAR ONE: STRUCTURAL INTEGRITY AND THE DUTY TO DISRUPT	31
B PILLAR TWO: PROACTIVE OVERSIGHT AND CERTIFICATION	32
C PILLAR THREE: EVIDENTIARY AND ENFORCEMENT REFORM	34
IV THE FRAMEWORK IN PRACTICE: CASE STUDIES IN AGENTIC GOVERNANCE	36
A CASE STUDY 1: THE HOTEL MARKET—FROM DOCTRINAL DEADLOCK TO A DUTY TO DISRUPT	38
B CASE STUDY 2: THE RENTAL MARKET—FROM HUB-AND-SPOKE TO ECOSYSTEM GOVERNANCE	40
C CASE STUDY 3: THE AIRLINE MARKET—FROM REACTIVE ENFORCEMENT TO PROACTIVE DESIGN AUDITS	41
D CASE STUDY 4: THE GROCERY MARKET—PREVENTING CONFLICTS OF INTEREST IN A NASCENT AGENTIC ECOSYSTEM	43
E CASE STUDY 5: A HYPOTHETICAL FUTURE—GOVERNING THE “SHOPSPHERE” RETAIL MARKET	45
V BROADER LEGAL AND POLICY IMPLICATIONS	46
A THE CORE GOVERNANCE MODEL	47
1 PRIVATE-SECTOR GOVERNANCE: FROM COMPLIANCE TO CORPORATE OVERSIGHT	47
2 THE PUBLIC LAW FRAMEWORK: ADMINISTRATIVE AUTHORITY AND JUDICIAL REVIEW	49
B THE MECHANICS OF ENFORCEMENT AND COMPLIANCE	51
1 THE EVIDENTIARY ARCHITECTURE: PRESERVING AND ACCESSING THE “DESIGN RECORD”	51
2 PRIVATE ENFORCEMENT: STANDING, CLASS ACTIONS, AND NON-WAIVABLE DUTIES	52

3	REMEDIES AND SAFE HARBORS: CALIBRATING LIABILITY AND INCENTIVES	53
C	JURISDICTIONAL AND CONTEXTUAL BOUNDARIES	54
1	THE INTERNATIONAL DIMENSION: COORDINATION, COMITY, AND CROSS-BORDER ENFORCEMENT	54
2	FEDERALISM AND PREEMPTION: THE ROLE OF STATE ENFORCEMENT	56
3	SECTORAL ADAPTATIONS AND TRANSITIONAL JUSTICE	57
D	BROADER ECONOMIC AND INNOVATION EFFECTS	58
1	ECONOMIC TRADE-OFFS AND WELFARE IMPLICATIONS	58
2	INNOVATION, OPEN SOURCE, AND MARKET ENTRY	60
E	AN IMPLEMENTATION PATH	61
VI	OBJECTIONS AND REBUTTALS	63
A	DEFINITIONAL AND OPERATIONAL CHALLENGES	64
1	VAGUENESS AND THE DUTY’S SCOPE	64
2	THE BUYER-SIDE MONOPSONY RISK	65
B	IMPLEMENTATION AND ENFORCEMENT CHALLENGES	67
1	INSTITUTIONAL INCAPACITY AND ADMINISTRABILITY	67
2	GAMING, EVASION, AND AUDIT ROBUSTNESS	68
3	PRIVACY AND SECURITY RISKS FROM TELEMETRY	70
4	RETROACTIVITY AND TRANSITIONAL JUSTICE FOR LEGACY SYSTEMS	71
C	LEGAL AND CONSTITUTIONAL CONSTRAINTS	72
1	STATUTORY AND CONSTITUTIONAL AUTHORITY	72
2	THE REBUTTABLE PRESUMPTION AND EVIDENTIARY STANDARDS	74
3	PRIVATE LITIGATION AND PROCEDURAL HURDLES	75
D	ECONOMIC AND INNOVATION CONCERNS	77
1	CHILLING INNOVATION AND SMALL-ENTITY BURDENS	77
2	PERVERSE ECONOMIC EFFECTS AND CONSUMER WELFARE PARADOXES	78
3	INTERNATIONAL COMPETITIVENESS AND REGULATORY ARBITRAGE	80
VII	CONCLUSION: FROM CONSPIRACY TO CONVERGENCE, DETECTION TO DESIGN	81

INTRODUCTION

By mid-2025, Delta Air Lines was expanding a pilot deployment of AI-enabled revenue management technology—developed in partnership with Fetcherr—intended to recommend (and, in limited test markets, set) fare price points. Delta told investors that the system was being used on roughly three percent of its domestic network, with a stated goal of scaling to roughly twenty percent

by the end of 2025.¹ Earlier, at Delta’s November 2024 Investor Day, executives had framed the initiative as a shift in how pricing decisions are produced: President Glen Hauenstein described the system as a “super analyst” that can continuously simulate what fare levels “should” be, based on the same inputs used by human analysts.² He outlined a longer-term vision to move beyond the “static price grid” toward dynamic “offer management,” ultimately enabling “a price that’s available on that flight, on that time to you, the individual.”³

The expansion prompted immediate scrutiny. On July 15, 2025, Senator Ruben Gallego publicly criticized the concept as “predatory pricing.”⁴ Days later, Senators Gallego, Mark Warner, and Richard Blumenthal sent a letter to Delta’s CEO demanding transparency and warning that such “surveillance-based price setting” could be used to exploit consumers by charging each traveler up to their personal “pain point.”⁵ The “Stop AI Price Gouging and Wage Fixing Act of 2025” was introduced in the House of Representatives shortly thereafter.⁶ Facing intense scrutiny, Delta formally responded in a letter dated July 31, 2025, stating that “[t]here is no fare product Delta

¹ See DELTA AIR LINES, INC., *Q2 2025 Earnings Call* (July 10, 2025) (corrected transcript), at 12 (statement of Glen W. Hauenstein) (stating that “today, we’re about 3% of domestic” and that Delta’s “goal is to have about 20% by the end of the year”), https://s2.q4cdn.com/181345880/files/doc_earnings/2025/q2/transcript/CORRECTED-TRANSCRIPT_-Delta-Air-Lines-Inc-DAL-US-Q2-2025-Earnings-Call.pdf [<https://perma.cc/2T2B-5QJU>].

² See DELTA AIR LINES, INC., *2024 Investor Day* (Nov. 20, 2024) (corrected transcript), at 14 (statement of Glen W. Hauenstein, President) (describing Fetcherr as a “super analyst” and explaining that it simulates “real-time what should the price points be,” such that Delta is “letting the machine ... actually go ahead and price”), https://s2.q4cdn.com/181345880/files/doc_downloads/2024/11/CORRECTED-TRANSCRIPT_-Delta-Air-Lines-Inc-DAL-US-Investor-Day-20-November-2024-8_30-AM-ET.pdf [<https://perma.cc/DG2G-6NKZ>].

³ *Id.* at 14 (stating Delta is moving away from “a static price grid” to “offer management” where there is “a price ... to you, the individual”); see also IRINA IVANOVA, *Delta Moves Toward Eliminating Set Prices in Favor of AI That Determines How Much You Personally Will Pay for a Ticket*, FORTUNE (July 16, 2025), <https://fortune.com/2025/07/16/delta-moves-toward-eliminating-set-prices-in-favor-of-ai-that-determines-how-much-you-personally-will-pay-for-a-ticket/> [<https://perma.cc/J9FM-LM4H>].

⁴ See SEN. RUBEN GALLEG0 (@RubenGallego), X (FORMERLY TWITTER) (July 15, 2025), <https://x.com/RubenGallego/status/1945138645783990445> [<https://perma.cc/Q6LN-P668>] (post) (stating that Delta’s CEO “got caught bragging about using AI to find your pain point” and calling the resulting practice “predatory pricing”).

⁵ See Letter from Sens. Ruben Gallego, Mark Warner & Richard Blumenthal to Ed Bastian, CEO, Delta Air Lines (July 21, 2025), at 1, <https://www.gallego.senate.gov/wp-content/uploads/2025/07/Delta-AI-Letter.pdf> [<https://perma.cc/8WQD-6GZK>]; see also Press Release, SEN. RUBEN GALLEG0, Gallego, Colleagues Demand Answers from Delta CEO as Company Moves Toward AI Pricing Model to Set Airfares (July 21, 2025), <https://www.gallego.senate.gov/press-releases/gallego-colleagues-demand-answers-from-delta-ceo-as-company-moves-toward-ai-pricing-model-to-set-airfares/> [<https://perma.cc/L8Q3-PLT5>].

⁶ See Stop AI Price Gouging and Wage Fixing Act of 2025, H.R. 4640, 119th Cong. (2025) (as introduced July 23, 2025), <https://www.congress.gov/bill/119th-congress/house-bill/4640> [<https://perma.cc/8WL7-H6BX>].

has ever used, is testing or plans to use that targets customers with individualized prices based on personal data.”⁷ Subsequently, U.S. Transportation Secretary Sean Duffy warned that the department would “engage very strongly” and investigate any airline attempting to use AI to set individualized prices based on personal characteristics.⁸

These individualized-pricing concerns are important, and Delta’s stance may have dampened fears of first-degree price discrimination. But focusing on this *one-to-one* harm obscures a broader, systemic risk: that algorithms can produce market-wide supracompetitive pricing—*algorithmic convergence*—without provable agreement. Delta’s assurance disavowed only individualized pricing; it did not address the risk that learning systems interacting across firms can independently reach and sustain high prices.⁹ That systemic danger does not depend on personal data; it is inherent in the design and deployment of the pricing algorithms whose use companies continue to expand.

Algorithmic convergence is not only well established in the literature,¹⁰ but is also central to ongoing litigation, albeit involving conventional revenue-management and recommendation tools that rely on pooled competitor data, rather than the autonomous, co-adaptive agents—such as those envisioned in Senators Gallego, Warner, and Blumenthal’s letter—which are the focus of this Article.¹¹ An “algorithmic handshake,” for instance, is the core allegation in the high-profile antitrust lawsuit against RealPage, Inc., where landlords allegedly fed nonpublic, competitively

⁷ See Letter from PETER CARTER, EVP & Chief External Affairs Officer, Delta Air Lines, to Sens. Ruben Gallego, Richard Blumenthal & Mark Warner (July 31, 2025), <https://news.delta.com/delta-responds-misinformation-around-ai-pricing/> [<https://perma.cc/F7WP-U9VM>].

⁸ See DAVID SHEPARDSON, *US Criticizes Use of AI to Personalize Airline Ticket Prices, Would Investigate*, REUTERS (Aug. 5, 2025) (quoting Secretary Duffy: “To try to individualize pricing on seats based on how much you make or don’t make or who you are, I can guarantee you that we will investigate if anyone does that.”), <https://www.reuters.com/world/us/us-criticizes-use-ai-personalize-airline-ticket-prices-would-investigate-2025-08-05/> [<https://perma.cc/FX2L-NUS5>].

⁹ See Carter Letter, *supra* note 7.

¹⁰ See, e.g., ARIEL EZRACHI & MAURICE E. STUCKE, *VIRTUAL COMPETITION: THE PROMISE AND PERILS OF THE ALGORITHM-DRIVEN ECONOMY* chs. 2–3 (Harv. Univ. Press 2016) (providing the foundational legal and policy analysis of this risk); EMILIO CALVANO ET AL., *Artificial Intelligence, Algorithmic Pricing, and Collusion*, 110 AM. ECON. REV. 3267, 3279–80 (2020) (showing, in simulations of Q-learning pricing agents, that independent algorithms can learn outcomes with elevated prices relative to competitive benchmarks).

¹¹ See *infra* notes 12, 30, and 31 and accompanying text.

sensitive data into a common AI, which then recommended inflated rents.¹²

Section 1’s “agreement-inference paradigm” falters with algorithmic convergence because these supracompetitive outcomes emerge without human “agreement,” any human “meeting of the minds.”¹³ As Part I of this Article demonstrates, the primary judicial workarounds used to infer such an agreement—the “hub-and-spoke” model and the doctrine of “conscious parallelism”—prove inconsistent and ineffective because agent-driven markets yield emergent prices from interacting designs, creating a profound doctrinal void.

Moreover, a parallel revolution is unfolding on the buyer side. Amazon’s “Rufus,” a generative AI shopping agent that converses with customers and compares alternatives, was rolled out to all U.S. customers in 2024.¹⁴ In the travel sector, Navan’s agentic AI, “Ava,” autonomously handles complex rebookings and hotel selections for corporate customers,¹⁵ while apps like Hopper use price prediction and automation to autonomously book and manage travel for retail consumers.¹⁶ This trend is accelerating, with retail giants Google (with its “AI Mode”) and Walmart (with its “Sparky” assistant) launching comparable buyer agents.¹⁷ These tools are not mere comparison

¹² See Complaint ¶¶ 4–6, *United States v. RealPage, Inc.*, No. 1:24-cv-00710 (M.D.N.C. Aug. 23, 2024), <https://www.justice.gov/atr/media/1365471/dl> [<https://perma.cc/4BJE-B9JM>].

¹³ While liability can also attach to a corporate entity, the evidentiary standard for a Section 1 ‘agreement’ remains fundamentally anthropocentric: it requires proof of a ‘meeting of the minds’ or ‘conscious commitment’ derived from the intent, communications, and actions of natural persons. See *Monsanto Co. v. Spray-Rite Serv. Corp.*, 465 U.S. 752, 764 (1984).

¹⁴ See RAJIV MEHTA, *How Customers Are Making More Informed Shopping Decisions with Rufus*, ABOUT AMAZON (Sep. 18, 2024), <https://www.aboutamazon.com/news/retail/how-to-use-amazon-rufus> [<https://perma.cc/J829-YBJR>]; see also SARAH PEREZ, *Amazon AI Chatbot Rufus Is Now Live for All U.S. Customers*, TECHCRUNCH (July 12, 2024), <https://techcrunch.com/2024/07/12/amazon-ai-chatbot-rufus-is-now-live-for-all-u-s-customers/> [<https://perma.cc/EZ9N-4G2X>].

¹⁵ See PYMNTS, *Navan CTO on Building an Agentic AI Platform With ‘Zero Critical Hallucinations’*, PYMNTS.COM (July 2, 2025), <https://www.pymnts.com/news/artificial-intelligence/2025/navan-chief-technology-officer-building-agentic-ai-platform-with-zero-critical-hallucinations/> [<https://perma.cc/G5YT-QAP4>] (quoting Navan’s CTO regarding the system’s capabilities).

¹⁶ See MICHAEL NUÑEZ, *Hopper’s AI Agent Can Book Flights and Cancel Trips Without Human Help*, VENTUREBEAT (Oct. 1, 2025), <https://venturebeat.com/ai/hoppers-ai-agent-can-book-flights-and-cancel-trips-without-human-help/> [<https://perma.cc/Z6EM-Y6LA>].

¹⁷ See, e.g., XANAYRA MARIN-LOPEZ, *Google Is Turning AI into a Personal Shopping Assistant*, RETAIL DIVE (May 22, 2025), <https://www.retaildive.com/news/google-shopping-ai-personal-assistant/748893/> [<https://perma.cc/BSZ4-63BU>]; XANAYRA MARIN-LOPEZ, *Walmart Debuts Sparky, Its Generative AI Assistant for Customers*, RETAIL DIVE (June 10, 2025), <https://www.retaildive.com/news/walmart-launches-generative-ai-assistant-sparky/750300/> [<https://perma.cc/M5CJ-TZCM>].

sites; they are goal-directed agents that can learn, adapt, and transact on a consumer’s behalf.

The true frontier of antitrust risk lies in the *collision of these two forces*: algorithmic convergence in *two-sided agentic AI markets*—what this Article calls *agentic convergence*—where seller agents meet buyer agents on platform-mediated marketplaces. When seller agents like Delta’s “super analyst” interact not only with human consumers but also with buyer agents like Rufus or Ava, traditional antitrust tools become even less probative. Buyer and intermediary agents *interpose autonomous decision-making between* sellers’ conduct and transaction prices. Because buyer conduct shapes the incentives that sustain (or destabilize) coordination, this interaction further weakens the inference from supracompetitive effects to an “agreement” (concerted action) under Section 1. The hub-and-spoke analogy—the workhorse of current algorithmic convergence litigation—becomes a poorer fit as numerous buyer and seller agents interact across multiple intermediary platforms.

This Article argues that antitrust must shift its focus: from preserving competition as a natural state to mandating competition as an *engineered outcome*. Algorithms and agents must be designed to actively disrupt supracompetitive pricing rather than passively sustaining (or even amplifying) collusive pricing.¹⁸

To ensure that firms design agents to fulfill procompetitive roles, the Article proposes a new governance framework centered on a legally enforceable “*Duty to Disrupt*.” This affirmative obligation requires developers to engineer agents to actively counteract collusive strategies. The duty is operationalized through three pillars: (1) *structural separation*—a “Glass-Steagall for AI”—to prevent conflicts of interest when firms operate both buyer- and seller-side agents;¹⁹ (2) *proactive*

¹⁸See FRANCISCO BENEKE & MARK-OLIVER MACKENRODT, *Remedies for Algorithmic Tacit Collusion*, 9 J. ANTITRUST ENF’T 152, 172–76 (2021) (arguing that pricing algorithms “can be programmed to avoid collusive outcomes” and discussing antitrust compliance-oriented remedies); JUSTIN P. JOHNSON, ANDREW RHODES & MATTHIJS WILDENBEEST, *Platform Design When Sellers Use Pricing Algorithms*, 91 ECONOMETRICA 1841, 1842–45 (2023) (demonstrating how platform design rules can destabilize algorithmic collusion). See also ANIRBAN MUKHERJEE & HANNAH H. CHANG, *Agentic AI: Autonomy, Accountability, and the Algorithmic Society* (arXiv Working Paper No. 2502.00289, 2025), <https://arxiv.org/abs/2502.00289> [<https://perma.cc/8KGR-Z5XJ>] (offering one stylized formalization of these dynamics in agentic markets).

¹⁹The term is used by analogy to the Banking Act of 1933 (commonly known as the Glass-Steagall Act), which separated commercial and investment banking to prevent conflicts of interest that were seen as contributing to the Great Depression. See Banking Act of 1933, Pub. L. No. 73-66, 48 Stat. 162 (1933). The principle here is the same: to

oversight via mandatory algorithmic audits and regulatory sandboxes that test interactive effects; and (3) *reformed evidentiary standards* that create a rebuttable presumption of liability, shifting the burden of production to developers once plaintiffs show persistent supracompetitive outcomes.

The Article proceeds in seven parts. *Part I* details the doctrinal void left by the Sherman Act’s collapsing “agreement” requirement in the face of algorithmic pricing. *Part II* develops the theoretical foundation of the two-sided agentic market, detailing the conditions that produce either collusive or competitive outcomes. *Part III* introduces the core of our proposal: a new antitrust framework built on a “Duty to Disrupt” and its three pillars of structural separation, proactive oversight, and evidentiary reform. *Part IV* then applies this framework to a series of case studies—spanning airlines, hotels, rentals, and retail—to demonstrate its practical utility. *Part V* explores the broader legal and policy implications of this framework, outlining how it can realign private-sector incentives and create a new ecosystem for public oversight. *Part VI* anticipates and rebuts key objections to the proposal, addressing challenges of vagueness, institutional capacity, evidentiary standards, and legal authority. *Part VII* concludes by summarizing the case for engineering, rather than merely preserving, competition.

I. ONE-SIDED ALGORITHMIC CONVERGENCE: DOCTRINAL CHALLENGES

Section 1 of the Sherman Act is built upon a fundamentally anthropocentric principle: liability attaches only to a human “agreement.” This doctrinal foundation faces structural strain when algorithms independently learn to sustain supracompetitive prices without any human “meeting of the minds.” This Part dissects the existing legal toolkit. It first traces the erosion of the core “agreement” requirement itself, then analyzes why the two primary judicial workarounds—the “hub-and-spoke” model and the doctrine of “conscious parallelism”—are rendered ineffective in one-sided algorithmic markets.

prophylactically separate functions that create inherent, systemic conflicts of interest.

A. The Strain on Agreement Inference

Section 1 prohibits only a “contract, combination... or conspiracy,” not all harms to competition. The Supreme Court has consistently interpreted this language to require a “meeting of the minds,”²⁰ targeting *concerted* action while deliberately leaving unilateral business decisions untouched.²¹ The doctrinal edifice of Section 1 rests on the ability of a plaintiff to prove a shared understanding among two or more legally distinct actors.

That premise is strained by algorithmic convergence. When multiple competing sellers deploy pricing algorithms that, through independent learning, converge on supracompetitive prices, a human “agreement” becomes elusive; an “algorithmic handshake” can occur without any explicit communication among the human operators, creating a scenario of convergence without a traditional conspiracy.²² Moreover, even the need for an explicit conspiracy becomes functionally unnecessary. Algorithms trained to maximize prices by learning rivals’ strategies can reach and sustain supracompetitive equilibrium prices without human coordination.²³

Ongoing litigation against RealPage, Inc. provides the canonical example: competing landlords allegedly ceded their pricing authority to a common AI (YieldStar), which then inflated market-wide rents by processing their pooled, nonpublic data.²⁴ This arrangement creates an immediate doctrinal puzzle: was there a “meeting of the minds” between the landlords? The AI’s operational autonomy obscures direct evidence of a traditional conspiracy, to the extent that there was one, forcing courts to impute agreement from circumstantial evidence.²⁵

²⁰ See *Copperweld Corp. v. Indep. Tube Corp.*, 467 U.S. 752, 767–68 (1984) (noting the Sherman Act contains a “basic distinction between concerted and independent action”).

²¹ See *Monsanto Co. v. Spray-Rite Serv. Corp.*, 465 U.S. 752, 764 (1984) (requiring a “conscious commitment to a common scheme” to prove concerted action).

²² See, e.g., Calvano et al., *supra* note 10, at 3270–75; Ezrachi & Stucke, *supra* note 10, chs. 2–3.

²³ See Calvano et al., *supra* note 10.

²⁴ See Complaint, *supra* note 12, ¶¶ 5–7.

²⁵ See Mukherjee & Chang, *supra* note 18, at 4–7 (discussing how agentic autonomy complicates traditional notions of accountability and intent).

B. Hub-and-Spoke Model

The leading legal response has been to adapt Section 1’s agreement requirement to algorithmic pricing by reframing the “meeting of the minds” as an *agreement to delegate*: a horizontal understanding among rivals to cede key pricing decisions to a common software “hub.” The Department of Justice has emphasized that where competitors knowingly delegate to a shared pricing algorithm with awareness that others are doing the same, the requisite concerted action can be inferred.²⁶ The theory resonates with classic hub-and-spoke precedents, which require not only vertical agreements between the hub and each spoke but also a horizontal “rim” connecting the spokes.²⁷

Recent outcomes show both this theory’s promise and its limits. In the *RealPage* multidistrict litigation, the court denied a motion to dismiss as to multifamily-housing plaintiffs, concluding that allegations of pooling nonpublic competitor data into a common algorithm and widespread adherence to its recommendations plausibly alleged a horizontal agreement; it dismissed, however, the student-housing complaint on other grounds.²⁸ A federal court likewise allowed parallel allegations to proceed against Yardi Systems, emphasizing that plaintiffs plausibly pleaded a *per se* price-fixing theory given the alleged horizontal arrangement mediated by a common vendor.²⁹

Acceptance of the hub-and-spoke theory, however, has been far from uniform. Courts dismissed complaints against Las Vegas and Atlantic City hotel-casino operators that relied on revenue-management software, finding the pleadings failed to connect the use of common tools to a plausible horizontal “rim,” especially absent specific allegations that nonpublic competitor

²⁶ See Memorandum of Law in Support of Statement of Interest of the United States at 20–21, *In re RealPage, Inc., Rental Software Antitrust Litig.* (No. II), No. 3:23-md-03071 (M.D. Tenn. Nov. 15, 2023) (arguing that “it makes no difference that prices are fixed through joint use of an algorithm instead of by a person”), <https://www.justice.gov/d9/2023-11/418053a.pdf> [<https://perma.cc/C9X8-87LQ>].

²⁷ See, e.g., *United States v. Apple Inc.*, 791 F.3d 290, 314–15 (2d Cir. 2015) (affirming that Apple facilitated a horizontal agreement among publishers); *In re Musical Instruments & Equip. Antitrust Litig.*, 798 F.3d 1186, 1192 (9th Cir. 2015) (describing hub-and-spoke conspiracies as requiring a horizontal “rim” among spokes).

²⁸ *In re RealPage, Inc., Rental Software Antitrust Litig.* (No. II), 709 F. Supp. 3d 478, 509–12 (M.D. Tenn. 2023).

²⁹ See Order Denying Defs.’ Joint Mot. to Dismiss at 13–15, *Duffy v. Yardi Sys., Inc.*, No. 2:23-cv-01391-RSL (W.D. Wash. Dec. 4, 2024), ECF No. 187, <https://cases.justia.com/federal/district-courts/washington/wawdce/2:2023cv01391/326049/187/0.pdf> [<https://perma.cc/7TGT-CBHU>].

data were pooled and used to generate competitor-conditioned pricing. The Ninth Circuit has since affirmed the Nevada dismissal.³⁰

The Northern District of California’s decision in *Dai v. SAS Institute Inc.* crystallizes the point. Consumers alleged that hotel chains used IDEaS’s revenue-management system (RMS) to exchange confidential data and “in nearly every instance” implement supracompetitive pricing recommendations. The court dismissed the complaint under *Twombly* for want of factual specificity, faulting plaintiffs for failing to allege when particular hotel operators began using the software or when they changed their pricing strategies, and for relying on materials about a nonparty to assert that defendants adopted the recommendations “in nearly every instance.”³¹ The court contrasted the sparse allegations with the detailed data-pooling claims credited in *RealPage* and *Duffy*, concluding that generalized assertions could not sustain an inference of a horizontal agreement.³²

Two implications follow. First, plaintiffs who can allege with particularity that rivals provided nonpublic pricing and supply data into a common model *and* the model generated recommendations using that pooled competitor data fare better at the motion-to-dismiss stage. Second, absent those particulars, allegations that firms adopted the same software or faced common market conditions generally will not suffice to connect the spokes and form a rim.

The theory’s greatest vulnerability is structural: in a decentralized market where competitors independently adopt different algorithms, the hub-and-spoke model collapses for want of a hub, making a unified agreement nearly impossible to plead.

³⁰ See *Gibson v. Cendyn Grp., LLC*, No. 24-3576 (9th Cir. Aug. 15, 2025) (affirming dismissal); see also *Gibson v. Cendyn Grp., LLC*, No. 2:23-cv-00140-MMD-DJA, slip op. at 3–6 (D. Nev. May 8, 2024) (dismissing with prejudice for failure to plausibly allege a horizontal agreement); *Cornish-Adebiyi v. Caesars Ent., Inc.*, No. 1:23-cv-02536 (D.N.J. Sep. 30, 2024) (dismissing for lack of well-pleaded allegations that confidential competitor data were pooled or used to set prices), *appeal docketed*, No. 24-3006 (3d Cir. Oct. 29, 2024).

³¹ See *Bell Atl. Corp. v. Twombly*, 550 U.S. 544, 555–57 (2007); see also *Dai v. SAS Inst. Inc.*, No. 24-CV-02537-JSW, slip op. at 4–5 (N.D. Cal. July 18, 2025) (order granting motion to dismiss with leave to amend) (faulting, inter alia, lack of timing allegations and reliance on non-party material for “adopt[ion] . . . in nearly every instance”).

³² See *Dai*, No. 24-CV-02537-JSW, slip op. at 5.

C. Conscious Parallelism

Recognizing the difficulty of finding direct evidence of conspiracy, the law’s primary workaround has been the doctrine of “conscious parallelism.” Under this framework, an agreement can be inferred from parallel conduct (e.g., uniform price hikes) if it is accompanied by “plus factors”—contextual evidence that tends to exclude the possibility of independent action.³³ A key question in modern, data-driven markets is whether the exchange of competitively sensitive information through a third-party intermediary can serve as a plus factor. In *United States v. Agri Stats*, for example, the court denied a motion to dismiss claims that the defendant facilitated collusion by collecting and disseminating detailed, nonpublic data among rivals in the poultry and pork industries.³⁴ The Federal Trade Commission (FTC) has separately alleged that Amazon’s “Project Nessie” algorithm functioned as a decentralized signaling mechanism, testing whether rivals would follow its price increases and thereby maintaining artificially elevated prices across the market.³⁵

Recent cases, however, illustrate the limits of plus-factor pleading when specific facts about pooled nonpublic data or its use are missing. In *Dai v. SAS Institute Inc.*, the court (applying *Twombly*) contrasted the sparse allegations about confidential-information use with the detailed pooling and use allegations credited in *RealPage* and *Duffy*, concluding that generalized assertions that the vendor “plugs” confidential inputs into its algorithm, without particularized facts about commingling and output to rivals, could not sustain an inference of a horizontal agreement.³⁶ By contrast, courts have been receptive where complaints plausibly allege two-way exchanges of

³³ See *Theatre Enters., Inc. v. Paramount Film Distrib. Corp.*, 346 U.S. 537, 540–41 (1954) (explaining that parallel business behavior does not itself establish a Sherman Act violation but may support an inference of conspiracy when combined with additional evidence).

³⁴ See Memorandum Opinion & Order Denying Defendant’s Motion to Transfer and Motion to Dismiss at 1–3, *United States v. Agri Stats, Inc.*, No. 0:23-cv-03009 (D. Minn. May 28, 2024).

³⁵ See Second Amended Complaint ¶¶ 423–27, *FTC v. Amazon.com, Inc.*, No. 2:23-cv-01495 (W.D. Wash. Oct. 31, 2024) (alleging Project Nessie “generated over \$1 billion” in additional profit for Amazon), https://www.ftc.gov/system/files/ftc_gov/pdf/0327-20231031-REDACTEDSecondAmendedComplaint.pdf [<https://perma.cc/M4GG-N4XT>].

³⁶ See *Dai*, No. 24-CV-02537-JSW, slip op. at 5 (distinguishing *In re RealPage* and *Duffy* on the specificity of data pooling and use).

rivals’ nonpublic, competitively sensitive data through a common hub and adherence to competitor-conditioned outputs—as in *In re RealPage* (multifamily housing) and *Duffy v. Yardi* (multifamily rentals), where those use-specifics supported an inference of a horizontal “rim.”³⁷

Several evidentiary frictions arise when applying this doctrine to autonomous systems. Parallelism among algorithms trained on similar data and objectives is *expected* rather than probative of agreement. “Inter-firm communications” are absent by design.³⁸ The plus factor of “acting against self-interest” *misfires* when applied to the learned, long-run strategies of these same reinforcement-learning agents.

One proposed adaptation is to treat competitors’ *knowing adoption* of algorithmic tools with a foreseeable tendency to facilitate coordination as a plus factor.³⁹ But knowledge is difficult to plead and prove: firms can plausibly claim ignorance of emergent behavior, and *ex post* competitive outcomes can mask *ex ante* design intent.

II. AGENTIC CONVERGENCE: A NEW PARADIGM AND DOCTRINAL CONSEQUENCES

Part I showed that, in algorithmic pricing markets, supracompetitive effects can be a weak proxy for the “agreement” element of Section 1: prices may converge without provable interfirm communication or a meeting of the minds. This Part turns to the setting in which that evidentiary problem deepens—*two-sided agentic markets*—where autonomous seller, buyer, and intermediary agents repeatedly interact and transact. In these ecosystems, market outcomes are an emergent product of interacting designs, further attenuating inference from price effects to conspiratorial

³⁷ See *In re RealPage*, 709 F. Supp. 3d at 492–94, 510–12; *Duffy*, No. 2:23-cv-01391-RSL, at 13–15.

³⁸ Reinforcement-learning agents can learn tacitly collusive strategies through repeated interaction without explicit communication. See Calvano et al., *supra* note 10; see also OECD, ALGORITHMS AND COLLUSION: COMPETITION POLICY IN THE DIGITAL AGE 31–34 (2017), https://www.oecd.org/content/dam/oecd/en/publications/reports/2017/05/algorithms-and-collusion-competition-policy-in-the-digital-age_02371a73/258dcb14-en.pdf [<https://perma.cc/9EKV-2FHD>].

³⁹ See JOSEPH E. HARRINGTON, JR., *Developing Competition Law for Collusion by Autonomous Artificial Agents*, 14 J. COMPETITION L. & ECON. 331, 350–51 (2018).

agreement and complicating causal attribution.

This Part proceeds by: (A) defining the two-sided agentic market and its key actors; (B) explaining its core economic mechanisms; (C) detailing why existing antitrust doctrine misfires in this environment; (D) outlining the resulting evidentiary crisis and motivating a pivot to design-based liability; and (E) explaining why any effective remedy must apply ecosystem-wide.

A. Defining the Two-Sided Agentic Market

A two-sided agentic market is one in which autonomous systems—operating on both the supply and demand sides—interact *repeatedly* and *strategically*, resolving market outcomes. Within this ecology operate three principal actors. *Seller agents*—dynamic pricing systems, inventory-allocation algorithms, and automated negotiators that firms deploy to manage supply—continuously adapt seller policies to market feedback.⁴⁰ *Buyer agents*, by contrast, represent demand: from advisory comparison tools to sophisticated “concierge” assistants that autonomously negotiate and execute transactions on a consumer’s behalf.⁴¹ Between them stand *intermediary agents*—platforms, search engines, and marketplaces whose ranking, default, and display algorithms set the rules of engagement for all participants, quietly shaping the strategic environment in which the others learn.

Three characteristics make this structure legally salient:

1. **Autonomy and Learning.** Agents adapt continuously through feedback loops, allowing strategies to co-evolve across firms and over time.⁴²

⁴⁰Examples include the revenue-management systems discussed in Part I.

⁴¹See KATHARINA SCHUMACHER ET AL., *The Agentic Commerce Opportunity: How AI Agents Are Ushering in a New Era for Consumers and Merchants*, MCKINSEY & CO. (Oct. 17, 2025), <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-agentic-commerce-opportunity-how-ai-agents-are-ushering-in-a-new-era-for-consumers-and-merchants> [<https://perma.cc/9H8X-6S85>] (describing interaction models in which “agents transact autonomously with other agents” and may “negotiate a bundle discount”); see also *supra* notes 14, 15, 16 (discussing Amazon’s Rufus, Navan’s Ava, and the Hopper app). See also MICHAEL WOOLDRIDGE & NICHOLAS R. JENNINGS, *Intelligent Agents: Theory and Practice*, 10 KNOWL. ENG’G REV. 115, 117–18 (1995) (describing a “weak” notion of agency including autonomy, social ability, reactivity, and pro-activeness).

⁴²See Wooldridge & Jennings, *supra* note 41, at 117–18.

2. **Repeated Interaction.** A high cadence of market activity enables agents to test, punish, and reward rival strategies—the game-theoretic precondition for tacit cooperation to arise and persist.
3. **Protocolized Communication.** Standardized rules of engagement create rich channels of observability. Agents can infer rivals’ policies from their behavior, facilitating coordination *without any central data-pooling hub or explicit agreement.*⁴³

Not every digital market crosses the threshold in a legally meaningful sense. The *two-sided agentic condition* arises only once the market achieves an operational density of autonomy and interaction: (1) when a non-trivial share of transactions is materially influenced or executed by both buyer *and* seller agents; (2) when those agents can update their strategies based on observed outcomes without contemporaneous human instruction; and (3) when they interact within a *protocolized environment*—such as a platform’s Application Programming Interface (API)—that standardizes communication and thereby enables cross-party learning. Below this threshold, traditional antitrust analysis remains serviceable; beyond it, the causal and evidentiary premises of existing doctrine begin to fail.

The protocol becomes the *board on which the game is played*. Its design primitives—update cadence and latency, visibility rules, ranking and default mechanics, and rate limits or throttling—script the tempo of learning. Each parameter is a legally relevant choice: altering cadence accelerates retaliation; altering visibility enables or disrupts tacit collusion; altering ranking amplifies or mutes the reward for defection; and altering throttling favors coordination over competitive exploration. Competition, therefore, becomes an *architectural property*, not only of the agents themselves, but of the interfaces through which the game is played.

⁴³The term “two-sided” here refers to *interactive autonomy on both sides of a transaction*, not the platform-economics usage concerning indirect network effects between distinct user groups. *See, e.g., Ohio v. Am. Express Co.*, 138 S. Ct. 2274 (2018); *see generally* JEAN-CHARLES ROCHET & JEAN TIROLE, *Platform Competition in Two-Sided Markets*, 1 J. EUR. ECON. ASS’N 990 (2003). Protocol-mediated signaling in such markets can be understood as a sophisticated, distributed evolution of earlier price-signaling practices. *Cf. United States v. Airline Tariff Publ’g Co.*, 836 F. Supp. 9 (D.D.C. 1993). *See also* Calvano et al., *supra* note 10, at 3270–75 (demonstrating how algorithms can learn to coordinate without explicit communication).

B. The Economic Mechanism: Design-Contingent Competition

In seller-only algorithmic markets, pricing systems optimize against atomized consumers acting independently, creating a demand environment that is, in aggregate, *passive* and *predictable*.⁴⁴ In two-sided agentic markets, by contrast, seller agents must “learn against learners”: buyer and intermediary agents adapt in parallel, and a single buyer or intermediary agent can represent many buyers, so market outcomes become an *interaction of algorithmic designs* rather than reflecting the unilateral choices of sellers.

Three design levers—derived from the economic conditions required to sustain tacit collusion—are decisive. First is the *objective function*: because tacit collusion relies on a “shadow of the future” to make punishment credible, agents must be designed with long-horizon objectives that explicitly *penalize sustained convergence* (for sellers) or maximize long-run user welfare (for buyers), thereby altering the effective intertemporal payoffs that otherwise support coordination.⁴⁵ Second is *strategic randomness*: because coordination relies on the predictability of rival responses, agents must inject calibrated exploration and randomized choices to resist exploitation and deter simple mimicry.⁴⁶ Third is *market-state awareness and response*: because the stability of collusion depends on the credibility of retaliatory punishment, agents must detect rivals’ conduct and trigger procompetitive responses—such as blunting that retaliation—to destabilize the equilibrium.⁴⁷

Two canonical equilibria—collusive or competitive—illustrate how these design choices can shape market outcomes. Consider a stylized airline duopoly in which two seller agents, Alpha Air and Beta Air, have converged on a supracompetitive \$500 fare, and Beta contemplates defecting to \$450.

⁴⁴ See Calvano et al., *supra* note 10; see also OECD, *supra* note 38, at 35–40 (discussing how pricing algorithms can facilitate tacit coordination absent explicit communication).

⁴⁵ See DREW FUDENBERG & ERIC MASKIN, *The Folk Theorem in Repeated Games with Discounting or with Incomplete Information*, 54 *ECONOMETRICA* 533 (1986) (showing that sufficiently patient players can sustain cooperative equilibria in repeated games through credible future punishment).

⁴⁶ See RICHARD S. SUTTON & ANDREW G. BARTO, *REINFORCEMENT LEARNING: AN INTRODUCTION* ch. 2 (2d ed. 2018) (discussing exploration–exploitation tradeoffs).

⁴⁷ See Johnson, Rhodes & Wildenbeest, *supra* note 18, at 1842–45 (modeling “price-directed prominence” and related platform rules that steer demand toward lower-priced sellers, thereby incentivizing algorithmic price cutting).

1. **Myopic Buyer Agents: Passive-Acceptance.** If the buyer agents are coded to simply pick today's lowest price, Beta's defection triggers an all-or-nothing demand shift. Observing this, Alpha retaliates (e.g., to \$400); the myopic buyer agents instantly reallocate. Repeated play teaches the seller agents that deviation yields a flash-in-the-pan gain followed by a costly price war, so the learned long-run policy is to preserve the \$500 norm. In short, *myopic buyer agents stabilize tacit collusion.*
2. **Strategic Buyer Agents: Tacit Collusion-Breaking.** If the buyer agents implement the three levers, the response is two-phased. When Beta defects, the strategic buyer agents *amplify the reward* by shifting substantial demand to Beta. If Alpha retaliates, they act as a *strategic shield, blunting the punishment* (e.g., splitting demand or continuing to favor the initial defector). By increasing the gains from defection and weakening the credibility of retaliation, they make collusion an unprofitable long-run strategy. Prices move toward competitive levels, tracking the logic of the folk theorem.⁴⁸ *Strategic buyer agents destabilize tacit collusion.*

The lesson is structural. What determines whether Alpha's and Beta's prices converge or unravel is not whether any humans exchanged assurances, but the mathematics the agents implement—the objective and reward functions, the exploration schedule, and the information protocol through which policies are learned. Those design choices fix the payoffs of deviation and punishment and therefore control the market equilibrium.

Competition, in sum, is not a natural baseline to be presumed but an *engineered property of agent architecture*. This renders Section 1's agreement-inference heuristics *under-identified*: supra-competitive outcomes no longer reliably signal an agreement, as the same result can emerge from the interaction of independently designed agents. By contrast, the *design record*—objective functions, constraints, default parameters, update cadence, and interaction logs—provides observable, auditable evidence of causation.

⁴⁸ See Fudenberg & Maskin, *supra* note 45.

The legal focus must therefore shift from policing secret agreements to *engineering rivalry by design*: imposing ex ante, verifiable obligations on market-facing agents and testing whether they were *reasonably designed and deployed* to compete. This is regulation of *algorithms as artifacts*, not command-and-control pricing; it *makes* competition a testable property of agent architecture rather than a hoped-for by-product.⁴⁹

C. The Obsolescence of Traditional Antitrust Doctrine

These economic properties of two-sided agentic markets—co-adaptation, emergent outcomes, and design-contingent competition—yield a series of misfires in *Section 1* adjudication. The “agreement” element becomes under-identified: supracompetitive *effects* no longer reliably support an inference of agreement once autonomous agents become an independent causal force, making prices the emergent output of interacting designs rather than any single firm’s choice. Hub-and-spoke theories lose traction as buyer agents and protocol layers mediate outcomes. “Conscious parallelism” plus factors invert because parallelism is expected and “against self-interest” is redefined by long-horizon learning. *Per se* shortcuts prove unstable where harm is interaction-contingent. And private suits encounter standing and proximate-cause frictions when agents intermediate transactions. Traditional agreements may still exist and be provable, but in two-sided agentic markets, *effects evidence no longer reliably maps to intent* without design evidence, yielding a systemic under-identification problem that renders traditional doctrine ineffective.

1. The Strain on Agreement Inference

Section 1 reaches only a “contract, combination... or conspiracy,” which courts have construed to require a “meeting of the minds.”⁵⁰ Two features of two-sided agentic markets complicate that inquiry.

- **Compounding Failure of Inference.** Inferring agreement from supracompetitive

⁴⁹ See *infra* Part III.

⁵⁰ See *Copperweld*, 467 U.S. at 768; see also *Monsanto*, 465 U.S. at 764.

outcomes—already fraught in seller-only algorithmic coordination cases—becomes conjectural when prices are the emergent product of interacting seller, buyer, and intermediary policies.

- **No Single Locus of Conspiracy.** A supracompetitive equilibrium can emerge from the interaction of many agents designed in isolation and deployed independently. Causation becomes sufficiently *diffuse* that the very concept of a unifying agreement loses explanatory force; no single bilateral relationship or collective decision is required to produce an anticompetitive outcome.

2. The Inapplicability of the Hub-and-Spoke Model

Hub-and-spoke theory requires plaintiffs to plead a horizontal “rim” connecting competitors, an inference often drawn in one-sided algorithmic markets from rivals knowingly delegating pricing to a common vendor that uses their pooled, nonpublic data.⁵¹ This theory becomes a poorer analytical fit in two-sided agentic markets for three reasons:

- **Intervening Autonomy and Causation Dilution.** Even where sellers use a common “hub,” its recommendations are filtered through independent buyer and intermediary agents that co-determine the final transaction price. The price a consumer pays is an *emergent output* of interacting seller, buyer, and protocol policies—not merely the result of a seller-side conspiracy. This intervening autonomy attenuates causal attribution, making it impossible to reliably infer that supracompetitive prices trace back to the horizontal “rim” rather than to the market’s interactive learning dynamics.
- **Structural Irrelevance in Mature Ecosystems.** Mature agentic markets often look less like a wheel and more like a *graph*: multiple seller AIs from different vendors, multiple buyer AIs with varied objectives, and platform protocols that standardize interaction. In such decentralized settings, there is no single software hub to form the conspiracy’s center, nor is there a stable, administrable “rim” connecting the spokes.

⁵¹For the classic vendor-hub theory, see Part I and sources *supra* notes 26, 28, 29.

- **The Protocol as an Invisible Hub.** The market’s *protocol layer*—its standardized APIs, update frequencies, and visibility rules—can functionally substitute for a hub by enabling mutual observability and punishment. It can stabilize coordination by making deviation quickly detectable and retaliation credible. Yet because it is merely infrastructure, current hub-and-spoke doctrine does not “see” it as a locus of agreement.

The hub-and-spoke theory thus retains its force only when plaintiffs can plead with particularity that rivals fed nonpublic data into a *single* vendor whose model produced *competitor-conditioned* outputs to which they adhered.⁵² Absent those specifics, or once buyer and intermediary autonomy materially co-determine transaction prices, the theory becomes under-identified.⁵³

3. The Inversion of Conscious Parallelism and “Plus Factors”

The doctrine permits inferring agreement from parallel conduct plus “plus factors” that tend to exclude independent action.⁵⁴ In two-sided agentic markets, however, the most common of those evidentiary proxies either invert or lose their probative value.

- **Parallelism Becomes the Baseline, Amplified by Protocol.** Parallelism among algorithms trained on similar objectives and data is the *baseline*—and the market’s *protocol layer* (standardized update cadences and visibility windows) can mechanically synchronize behavior even absent any agreement.⁵⁵
- **“Against Self-Interest” Misfires.** A long-horizon learning agent can rationally forgo an immediate price cut to avoid triggering a costly, algorithm-driven price war. What appears to be against self-interest period-by-period may be the *policy-optimal* strategy over time.
- **Knowledge of Adoption Is an Under-Determined Signal.** Knowing that rivals use a similar pricing tool says little about foreseeable outcomes when a heterogeneous and

⁵² See *supra* notes 28, 29; see also *supra* note 26.

⁵³ See *supra* notes 30, 31.

⁵⁴ See *Theatre Enters.*, 346 U.S. at 540–41.

⁵⁵ See Calvano et al., *supra* note 10; OECD, *supra* note 38, at 35–40.

evolving population of buyer agents co-determines the price. The high degree of *interaction uncertainty* severs the inferential link between common tool adoption and a tacit agreement.

In agentic markets, plus-factor analysis must therefore pivot to *design-based indicators*. The exchange of non-public, competitively sensitive data to generate *competitor-conditioned* outputs is a powerful plus factor.⁵⁶ Likewise, evidence of a vendor’s or deployer’s knowledge—grounded in internal testing or marketing materials—of an agent’s foreseeable tendency to coordinate is highly probative. Absent such design-specific evidence, however, mere parallelism is an insufficient basis from which to infer conspiracy.

4. The Instability of Per Se Rules

Per se condemnation is reserved for conduct “that would always or almost always tend to restrict competition and decrease output.”⁵⁷ Classic horizontal price-fixing is the canonical example, as its anticompetitive effect is so predictable that it is “conclusively presumed to be unreasonable” without a detailed inquiry into market context.⁵⁸

Two-sided agentic markets, however, frustrate this premise because harm becomes *interaction-contingent*: the same seller-side tool can either sustain supracompetitive prices when facing myopic buyer agents or be disciplined toward competitive outcomes when facing collusion-breaker buyers.⁵⁹ Whether a seller’s algorithmic strategy “tends to restrict competition” thus turns on the counterparty’s design, not on an intrinsic property of the seller’s tool—defeating the “always or almost always” predicate. This variability counsels against categorical condemnation and, under the logic of *California Dental*, requires at least a structured inquiry into market mechanics rather than a reflexive *per se* or “quick look” approach.⁶⁰

⁵⁶ See *Agri Stats*, No. 0:23-cv-03009, slip op. at 2.

⁵⁷ See *N. Pac. Ry. Co. v. United States*, 356 U.S. 1, 5 (1958).

⁵⁸ See *United States v. Socony-Vacuum Oil Co.*, 310 U.S. 150, 223–24 (1940); *Ariz. v. Maricopa Cnty. Med. Soc’y*, 457 U.S. 332, 344–48 (1982).

⁵⁹ See *supra* Part II.B (contrasting myopic and strategic buyer-agent equilibria).

⁶⁰ See *Cal. Dental Ass’n v. FTC*, 526 U.S. 756, 770–81 (1999) (warning against abbreviated analysis where competitive effects are ambiguous and turn on contested empirical questions).

Moreover, *architectural ambiguity* makes classification error-prone. Some market arrangements may more closely resemble procompetitive *joint ventures* or *new-product integrations*—for example, where a platform’s protocol layer arguably creates a “new product” from aggregated data.⁶¹ Treating protocol-mediated learning as *per se* price-fixing risks generating false positives and over-detering designs whose competitive effects are context-dependent, a direct contradiction of the error-cost logic that animates modern antitrust.⁶²

None of this is to say *per se* is obsolete. The narrow zone where it remains appropriate maps cleanly onto the one-sided algorithmic cases that fit the hub-and-spoke template: where rivals *pool nonpublic, competitively sensitive data* into a *single hub* to generate *competitor-conditioned outputs* to which they *adhere*, a court can reasonably treat the arrangement as a classic horizontal restraint that “merely cloak[s]” price-fixing in code.⁶³ But once buyer and intermediary autonomy materially co-determine transaction prices, harm becomes interaction-dependent, and *per se* logic loses its footing.

The administrable alternative is a *structured, design-sensitive rule of reason* with calibrated burden shifting, an approach consistent with the Supreme Court’s modern jurisprudence.⁶⁴ As developed in Part III, plaintiffs would first make a *prima facie* showing of persistent supracompetitive *effects*; the burden would then shift to defendants to *produce* their *design record*—objective functions, parameters, and certified red-team results—to rebut the inference of design-based causation. This approach preserves *per se* treatment for naked restraints while supplying courts with the very evidence traditional doctrine lacks in agentic settings.

5. Compounding Failures in Private Litigation

⁶¹ See *Broad. Music, Inc. v. CBS*, 441 U.S. 1, 19–24 (1979) (declining *per se* treatment where a blanket license constituted a new product of integration); *Texaco Inc. v. Dagher*, 547 U.S. 1, 5–8 (2006) (joint venture’s internal pricing evaluated under the rule of reason).

⁶² See *Leegin Creative Leather Prods., Inc. v. PSKS, Inc.*, 551 U.S. 877, 887–94 (2007) (shifting from *per se* to rule-of-reason based on context-dependent effects and the risk of chilling procompetitive conduct).

⁶³ See *Maricopa*, 457 U.S. at 348 (observing that a facially distinct restraint “may be a masquerade for an agreement to fix uniform prices”). This is precisely the theory credited in cases with detailed pooling-and-use allegations. See *In re RealPage*, 709 F. Supp. 3d at 509–12; *Duffy*, No. 2:23-cv-01391-RSL, at 12–14.

⁶⁴ See *NCAA v. Alston*, 141 S. Ct. 2141, 2155–57 (2021) (endorsing a structured, step-by-step rule-of-reason analysis).

Two-sided agentic markets multiply the frictions that already plague private antitrust suits—*who may sue, what counts as injury, how to trace causation, and whether a class is manageable*—turning threshold doctrines into dispositive barriers.

- **Direct Purchaser Status and Damages.** Under *Illinois Brick*, only *direct purchasers* may recover federal damages.⁶⁵ When a buyer agent executes a transaction, defendants will argue that the end consumer is an *indirect purchaser*. *Apple v. Pepper* narrows but does not eliminate this ambiguity; its formalistic inquiry into “who transacts with whom” can support standing where an agent is a mere conduit, but it invites threshold skirmishes over agency and privity that can derail a case before merits discovery.⁶⁶ This doctrinal friction is best addressed by re-grounding liability in design: a parallel, tort-like claim against a buyer-agent’s developer for negligent design is not constrained by *Illinois Brick*, while state-level *Illinois Brick* repealers provide a further avenue for recovery.
- **Antitrust Injury and Proximate Cause.** Plaintiffs must show not only an injury “of the type the antitrust laws were intended to prevent” but also a direct causal link between the violation and the harm.⁶⁷ In agent-mediated markets, defendants can plausibly argue that any overcharge is attributable to the *consumer’s own agent’s* “myopic” design—an *intervening cause* that defeats proximate cause. Even if that defense ultimately fails, it forces plaintiffs into fact-intensive inquiries about algorithmic design that Section 1 doctrine does not presently equip courts to manage.
- **Class Certification.** The dynamic and individualized nature of agent-driven pricing complicates the predominance requirement of Federal Rule of Civil Procedure 23(b)(3). *Comcast* requires a damages model that is consistent with the theory of liability, a challenge when

⁶⁵ See *Ill. Brick Co. v. Illinois*, 431 U.S. 720, 735–36 (1977); see also *Kansas v. UtiliCorp United, Inc.*, 497 U.S. 199, 207–17 (1990) (reaffirming the rule and its rationale against complex pass-through inquiries).

⁶⁶ See *Apple Inc. v. Pepper*, 139 S. Ct. 1514, 1520–23 (2019).

⁶⁷ See *Brunswick Corp. v. Pueblo Bowl-O-Mat, Inc.*, 429 U.S. 477, 489 (1977); *Assoc. Gen. Contractors of Cal., Inc. v. Cal. State Council of Carpenters (AGC)*, 459 U.S. 519, 535–45 (1983).

prices are heterogeneous.⁶⁸ While courts have accepted statistical proof of class-wide impact, the risk that individualized inquiries will overwhelm common questions is acute.⁶⁹ The framework proposed in this Article supplies the necessary common proof by focusing on uniform *design evidence* (objective functions, default parameters) and market-wide econometrics.

- **Contracting Out of Enforcement.** A distinctive feature of agentic markets is the ubiquitous use of end-user license agreements (EULAs) containing *mandatory arbitration clauses* and *class-action waivers*. The Supreme Court has made such provisions difficult to avoid, effectively *atomizing* disputes that depend on common, system-level proof.⁷⁰ In practice, this can remove the only economically viable path for private redress in markets where individualized damages are small but aggregate harm is large.
- **Evidence and Spoliation.** The most probative evidence in these cases—policy-update logs, default parameters, A/B test archives, and red-team results—is highly ephemeral and held exclusively by developers. Without an affirmative retention requirement, this electronically stored information (ESI) is often lost to routine data-purging policies, leaving plaintiffs to rely on the often-inadequate curative measures available under Federal Rule of Civil Procedure 37(e) for the spoliation of evidence.⁷¹ The proposed design-record retention duty directly targets this critical failure mode.
- **Sectoral Preemption.** Sector-specific statutes can truncate private claims. The Airline Deregulation Act, for instance, broadly preempts state-law claims “related to a price, route, or service of an air carrier,” channeling plaintiffs toward federal court and heightening

⁶⁸ See *Comcast Corp. v. Behrend*, 569 U.S. 27, 35 (2013).

⁶⁹ See *Olean Wholesale Grocery Coop., Inc. v. Bumble Bee Foods LLC*, 31 F.4th 651, 664–83 (9th Cir. 2022) (en banc) (discussing standards for statistical evidence in antitrust class actions); cf. *Tyson Foods, Inc. v. Bouaphakeo*, 577 U.S. 442, 453–61 (2016) (accepting representative proof where liability turned on common practices).

⁷⁰ See *AT&T Mobility LLC v. Concepcion*, 563 U.S. 333, 340–52 (2011); *Am. Express Co. v. Italian Colors Rest.*, 570 U.S. 228, 232–38 (2013).

⁷¹ See Fed. R. Civ. P. 37(e) (addressing the failure to preserve ESI); see also Fed. R. Civ. P. 37(e) advisory committee’s note to 2015 amendment (explaining the intent requirement for severe sanctions).

the importance of federal injunctive relief under Clayton Act § 16.⁷² This makes it all the more critical that courts have administrable, performance-based remedies—such as ordering modifications to an agent’s exploration budget or update cadence—that can be implemented without intrusive source-code disclosure.

The through-line is not that private enforcement is impossible, but that each traditional filter becomes harder to pass when prices are the emergent output of interacting designs. That compounding friction justifies the pivot in the next subsection: an evidentiary regime that centers *design-based causation* and *auditable artifacts*, with calibrated burden shifting to elicit the evidence only developers control.

D. The Resulting Evidentiary Void and the Pivot to Design-Based Liability

Traditional Section 1 tools create an *evidentiary void* in two-sided agentic markets. Plaintiffs can show durable, supracompetitive *effects*, but the most probative evidence of how the market behaves—the objective functions, update cadence, default parameters, and interaction logs that determine agent policies—sits outside Section 1’s intent-and-communication paradigm and is held exclusively by developers.⁷³

This void warrants a pivot in the *object* of legal inquiry—from a backward-looking search for conspiratorial *intent* to a forward-looking analysis of algorithmic *design*. The legally relevant question should shift from, “Did the firms agree?” to, “Did the firm *negligently design and deploy* an agent with a reasonably *foreseeable tendency* to cause or sustain anticompetitive outcomes in interaction?” That reframing maps onto familiar design-defect and negligence principles: liability turns on a product’s foreseeable risks and the availability of *reasonable alternative designs*, not on

⁷²49 U.S.C. § 41713(b)(1) (2022); see *Morales v. Trans World Airlines, Inc.*, 504 U.S. 374, 383–90 (1992).

⁷³The recognition of this informational asymmetry has already prompted proposals to shift proof burdens in algorithmic-collusion cases. See, e.g., Preventing Algorithmic Collusion Act of 2025, S. 232, 119th Cong. § 5 (2025) (creating a “presumption of agreement” in specified algorithmic-pricing settings and providing a rebuttal mechanism), <https://www.congress.gov/bill/119th-congress/senate-bill/232/all-info> [<https://perma.cc/RJ5Z-UHXX>].

the secret intentions of its makers.⁷⁴

This approach gives rise to a new, administrable test for what may be termed *design-based causation*. To establish a prima facie case, a plaintiff would need to satisfy three elements:

1. **Effect.** A showing of *persistent supracompetitive outcomes* (e.g., durable overcharges relative to competitive benchmarks after controlling for cost and demand shocks).
2. **Foreseeable Tendency.** Demonstrate that the defendant's agent—as specified and deployed—has a *known or reasonably foreseeable* tendency to contribute to those outcomes in interaction. Foreseeability can be grounded in: (a) the economic literature on convergence among reinforcement-learning price setters;⁷⁵ (b) the vendor's own internal evaluations (simulations, A/B tests, or red-team results) and marketing claims about stabilizing or raising market prices;⁷⁶ and (c) architectural choices that predictably facilitate coordination (e.g., short horizons, deterministic best-response policies, and high-frequency, high-visibility updates).
3. **Feasible Alternatives.** Identify *reasonable alternative designs* that would materially mitigate convergence risk at acceptable cost and without defeating the system's utility—for example: (i) *long-horizon objectives* that penalize sustained market-level price alignment; (ii) *strategic randomness* (calibrated exploration) that resists exploitation and deters mimicry; and (iii) *market-state awareness* that detects convergence signals and triggers procompetitive responses.⁷⁷

Once the plaintiff makes the *Effect* showing, the *burden of production*—not the ultimate burden of persuasion—would shift to the developer to produce its design record and any current certifications or red-team reports that rebut design-based causation. If the firm produces a robust,

⁷⁴ See RESTATEMENT (THIRD) OF TORTS: PRODS. LIAB. § 2(b) (Am. L. Inst. 1998) (defining design defect by reference to foreseeable risks avoidable through reasonable alternative designs).

⁷⁵ See Calvano et al., *supra* note 10; Ezrachi & Stucke, *supra* note 10.

⁷⁶ See Complaint, *supra* note 12, ¶¶ 4, 9 (alleging that pricing recommendations are generated using competing landlords' nonpublic data and that the system monitors and pressures adoption of those recommendations).

⁷⁷ See *supra* Part II.B; cf. RESTATEMENT (THIRD) OF TORTS: PRODS. LIAB. § 2(b) (Am. L. Inst. 1998) (defining a product as defective in design when foreseeable risks could have been reduced by adopting a reasonable alternative design).

contemporaneous record demonstrating rivalry-by-design, the inference of causation is weakened; if it cannot, the inference strengthens. Plaintiffs retain the burden to persuade the trier of fact that negligent design materially contributed to the observed outcomes.⁷⁸

Because the determinants of behavior are technical, the evidentiary focal point must be the *design record*—auditable artifacts that developers already create in the ordinary course:

- **Design artifacts:** objective and reward specifications; programmed constraints; default and deployed hyperparameters (including exploration budgets, learning rates, and update cadence).
- **Operational telemetry:** policy-update histories; A/B test archives and rollout logs; incident reports reflecting observed convergence or price-dispersion collapse.
- **Interaction results:** certified antitrust red-team and simulation outputs showing performance against rival agents (with configuration files, seeds, and sufficient detail for reproduction).

To prevent the evidentiary vacuum that often frustrates enforcement, firms in designated agentic markets can be subject to a *retention rule* for these materials for the model’s lifecycle plus a defined period. Failure to preserve them warrants adverse inferences under well-settled principles governing spoliation of electronically stored information.⁷⁹

Courts need not become software auditors to manage this evidence. Familiar tools suffice: protective orders and *in camera* review for trade-secret material; court-appointed neutral experts under Federal Rule of Evidence 706; and special masters under Rule 53 to supervise technical discovery. Where source-code disclosure is impractical, courts can order *algorithmic escrow* and reproducible-build attestations.⁸⁰

⁷⁸ See *Griggs v. Duke Power Co.*, 401 U.S. 424, 431–32 (1971) (allowing disparate-impact cases to proceed on effects evidence with burden shifting); *United States v. Phila. Nat’l Bank*, 374 U.S. 321, 362–63 (1963) (structural presumption in merger cases).

⁷⁹ See *supra* note 71.

⁸⁰ See Fed. R. Civ. P. 26(c)(1)(G) (protective orders for trade secrets); Fed. R. Evid. 706 (court-appointed experts); Fed. R. Civ. P. 53 (special masters); see also *Ruckelshaus v. Monsanto Co.*, 467 U.S. 986, 1001–04 (1984) (recognizing trade-secret protections).

This “design-based causation” standard does not collapse into *per se* condemnation. It supplies a structured, negligence-style inquiry tailored to the actual levers that control agent behavior and can be administered under the rule of reason. It also harmonizes with burden-shifting frameworks that courts already use when defendants uniquely control the most probative evidence.⁸¹

E. The Necessity of an Ecosystem-Wide Duty

The pivot to a design-based liability framework raises the threshold question: *who* owes the duty? Because anticompetitive harm in a two-sided agentic market is *interactional* and *emergent*, a legal remedy focused on only one side of the market would be incomplete and ultimately ineffective.⁸²

A duty imposed only on *seller agents* would prove fragile. Even if seller AIs are designed with procompetitive guardrails, they still learn against an environment scripted by buyer behavior. If buyer agents remain myopic and predictable, punishments for defection are swift and certain, making tacit cooperation the long-run best reply for seller agents notwithstanding their internal constraints.⁸³ The procompetitive design would be undermined by the anticompetitive structure of the market interaction.

A duty imposed only on *buyer agents* would also be insufficient. While strategic buyer agents can disrupt simple seller strategies, sophisticated seller agents can learn to anticipate and neutralize their countermeasures. Seller AIs could learn to collude on non-price dimensions that are less visible to buyer agents, such as inventory throttling or ancillary fees, or they could develop sophisticated probing strategies to exploit the buyer agents’ own rules—re-establishing a supracompetitive equilibrium.⁸⁴

The duty must also bind the *intermediary agent*. By specifying the protocol—ranking and

⁸¹ See *Griggs*, 401 U.S. at 431–32; *Phila. Nat’l Bank*, 374 U.S. at 362–63. See also *infra* Part III.C (proposing a structured rule-of-reason with calibrated burden shifting).

⁸² See *supra* Part II.B and accompanying text (explaining design-contingent competition).

⁸³ See *supra* notes 45, 46, 47, and accompanying text.

⁸⁴ See *supra* Part II.B (discussing adaptive learning). See also OECD, *supra* note 38, at 43–44 (discussing “risks of algorithmic selection beyond collusion”); Ezrachi & Stucke, *supra* note 10, at 73–75 (describing how algorithms can exploit complex market parameters beyond price).

default rules, visibility windows, update cadence, and rate limits—platforms script the tempo and transparency of learning for buyers and sellers alike. These defaults calibrate how quickly deviation is detected and how credible retaliation becomes, functioning as an “invisible hub” capable of stabilizing convergence even where no nonpublic competitor data are pooled.⁸⁵

Moreover, for the duty to be *operational* rather than aspirational, *auditability* requires symmetry. Whether an ecosystem tends toward coordination or rivalry is not a property of any single agent but a *triadic* property of seller, buyer, and protocol interacting. Verification therefore requires interactive evidence. If any class of agent is exempt from the duty, it can decline to participate in the necessary adversarial simulations or withhold telemetry, rendering audits under-identified and leaving courts to infer intent from effects alone.⁸⁶

The only robust solution is a framework where *all market-facing agents*—seller, buyer, and intermediary—share a legally enforceable duty to promote competition. This necessity for a symmetric, ecosystem-wide obligation forms the foundation for the new antitrust framework proposed in the next Part.

III. ENGINEERING RIVALRY: A NEW ANTITRUST FRAMEWORK FOR AGENTIC CONVERGENCE

If, as Parts I and II show, competition in two-sided agentic markets is an engineered outcome, the law’s task is to specify the architecture.⁸⁷ This Part sets out a framework that makes rivalry a

⁸⁵ See *supra* note 43 and accompanying text.

⁸⁶ See *infra* Part III.B and accompanying text (proposing mandatory algorithmic audits and regulatory sandboxes); see also OECD, *supra* note 38, at 39–40 (discussing market studies and other information-gathering tools to evaluate algorithmic pricing practices).

⁸⁷ The shift toward ex-ante design governance is, in fact, already underway, with the principle of imposing *ex-ante* obligations on firms to ensure fair and competitive market designs gaining traction in major regulatory initiatives. See, e.g., Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 Sep. 2022 on Contestable and Fair Markets in the Digital Sector and Amending Directives (EU) 2019/1937 and (EU) 2020/1828 (Digital Markets Act), 2022 O.J. (L 265) 1, arts. 5–7 (imposing affirmative obligations on designated “gatekeeper” platforms); FED. TRADE COMM’N, *Policy Statement Regarding the Scope of Unfair Methods of Competition Under Section 5 of the Federal Trade Commission Act* 1–6 (Nov. 10, 2022), https://www.ftc.gov/system/files/ftc_gov/pdf/P221202Section5PolicyStatement.pdf [<https://perma.cc/J4JM-7TAV>] (asserting that Section 5 reaches beyond the Sherman and Clayton Acts to address

design requirement. It is organized around three mutually reinforcing pillars: (A) a foundational legal obligation—the *Duty to Disrupt*—paired with *structural rules* that prevent inherent conflicts of interest; (B) a regime of *proactive oversight*, including mandatory audits and regulatory sandboxes, that turns this duty into a verifiable engineering requirement; and (C) *evidentiary reform* that realigns litigation with this new reality by shifting the focus from human intent to design-based causation, backed by a calibrated, rebuttable presumption.

A. Pillar One: Structural Integrity and the Duty to Disrupt

The first pillar addresses the market’s foundational architecture. It establishes *what* legal obligations firms that develop or deploy market-facing agents owe and implements structural safeguards to ensure the corporate entities behind them are not compromised by inherent conflicts of interest.

1. The Duty to Disrupt (The Core Design Obligation)

The framework introduces a market-wide, affirmative *Duty to Disrupt* for *all* firms that develop or deploy market-facing agents—seller, buyer, and intermediary. Framed as a negligence-style obligation to design and deploy agents that *reasonably mitigate* the foreseeable, systemic risk of agentic convergence, the duty is technology-neutral. Compliance requires *rivalry-by-design*, generally evidenced by features such as:⁸⁸

1. A **long-term optimization horizon** that prioritizes sustained, market-level consumer welfare over period-by-period revenue or savings;
2. **Strategic randomness**—calibrated exploration and stochastic choice rules—sufficient to make its behavior resistant to simple mimicry and to erode the credibility of punishment strategies; and
3. **Market-state awareness** that can detect convergence signals and respond with procompetitive policies, such as amplifying rewards for defectors and blunting retaliation.

unfair conduct that tends to negatively affect competitive conditions).

⁸⁸These functional requirements derive from the economic conditions for tacit collusion discussed *supra* Part II.B. See also Mukherjee & Chang, *supra* note 18, at 15–18 (modeling these specific features).

The duty sounds in familiar negligence and design-defect principles: liability turns on foreseeable risk, the availability of reasonable alternative designs, and auditable evidence of implementation—not on a futile search for hidden human intent.⁸⁹ A breach of this duty would become a new cause of action, allowing government agencies or class-action plaintiffs to sue the developer for failing to build a genuinely competitive agent.

2. Structural Integrity (Preventing Built-in Conflicts)

The *Duty to Disrupt* can be systematically undermined if the same firm is allowed to build both sides of the market. To address this, the framework proposes a prophylactic rule of *structural separation*—a “Glass-Steagall for agents.”⁹⁰ Under this rule, a single corporate entity would be presumptively barred from: (i) developing and operating both buyer-side and seller-side agents that interact within the same designated market; and (ii) operating an intermediary agent that sets market protocols while also fielding a buyer or seller agent that competes in that market. Where full structural separation is not warranted, regulators could require, at a minimum, *functional separation* through audited data and operational firewalls.

B. Pillar Two: Proactive Oversight and Certification

To make the *Duty to Disrupt* an operational, auditable engineering requirement, the second pillar employs a tiered regime of proactive oversight. This is justified by the profound *information asymmetry* at play; only developers have access to the code and data needed to verify an agent’s competitive integrity.

⁸⁹See RESTATEMENT (THIRD) OF TORTS: PRODS. LIAB. § 2(b) (Am. L. Inst. 1998) (defining design defect by reference to foreseeable risks avoidable through reasonable alternative designs); see also *Barker v. Lull Eng’g Co.*, 573 P.2d 443, 455 (Cal. 1978) (establishing the risk-benefit test for design defects). Here, the “harm” is the foreseeable emergence of a supracompetitive equilibrium, and the “failure” is the developer’s negligent design of a myopic agent that cannot protect its users from that systemic harm.

⁹⁰The principle of prophylactically separating functions that create inherent, systemic conflicts of interest is directly applicable to modern digital platforms. See, e.g., Complaint, *United States v. Google LLC*, No. 1:23-cv-00108 (E.D. Va. Jan. 24, 2023) (alleging Google unlawfully monopolized advertising services by exploiting its conflicting roles across the ad tech stack); LINA M. KHAN, *Amazon’s Antitrust Paradox*, 126 YALE L.J. 710, 736–42 (2017) (providing the foundational analysis of how platform integration creates anticompetitive conflicts of interest).

The first tier, and the baseline enforcement mechanism, is a system of *mandatory algorithmic audits*. To demonstrate compliance with the *Duty to Disrupt*, developers and deployers of all market-facing agents—seller, buyer, and intermediary—must conduct regular “antitrust red-teaming” and submit certified reports to the relevant regulator. Because anticompetitive harm is interactional, these audits must evaluate an agent’s performance not in isolation, but *in context*: in dyadic and triadic simulations against other agents. For a seller agent, the audit must show it lacks a foreseeable tendency to converge on supracompetitive prices. For a buyer agent, it must demonstrate its effectiveness at breaking, rather than reinforcing, collusive equilibria. For an intermediary agent, it must certify that its protocol design impedes, rather than facilitates, the real-time monitoring and punishment that underpins tacit collusion.⁹¹

For markets that are structurally prone to coordination—such as airlines, telecommunications, or banking—regulators could layer on a second, more intensive tool: *regulatory sandboxes*. This model draws on established practice in other sectors.⁹² Before a new market-significant agent is widely deployed, or before a platform makes a significant change to its protocol layer (such as its ranking or information-display rules), the developer would be required to test the system in a secure, regulator-supervised simulation. This allows regulators to observe emergent behavior and condition approval on specific design modifications *before* the agent can cause widespread consumer harm.

To prevent “certification theater,” the oversight regime is designed to be dynamic and difficult to game. The test suites are not static but include *rotating, partially undisclosed adversarial scenarios* to prevent firms from simply “teaching to the test.” Certifications are *time-bounded* and *automatically lapse upon any material model update*, ensuring that as agents evolve, they are

⁹¹The concept of auditing algorithmic pricing through structured investigative tools is supported by international policy analysis. See, e.g., OECD, *supra* note 38, at 39–40 (discussing market studies and related information-gathering mechanisms in response to algorithmic pricing concerns).

⁹²Regulatory sandboxes are a proven tool for testing new technologies in a controlled environment. See, e.g., FIN. CONDUCT AUTH. (U.K.), *Regulatory Sandbox* (last visited Feb. 6, 2026), <https://www.fca.org.uk/firms/innovation/regulatory-sandbox> [<https://perma.cc/L66N-VT6Y>] (describing the FCA regulatory sandbox); Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act), O.J. L, 2024/1689 (July 12, 2024), art. 57 (establishing AI regulatory sandboxes).

subject to re-evaluation. The pre-deployment testing is supplemented by requirements for firms to submit privacy-preserving *post-deployment telemetry*, to detect suspicious behavioral drift and trigger surprise audits.

Taken together, these tools serve a dual function. Substantively, they give concrete, technical meaning to the *Duty to Disrupt*. Procedurally, they build the very evidentiary record that courts and agencies currently lack: a contemporaneous, technically grounded account of what firms knew or should have known about their agents' competitive risks at the time of design and deployment.

C. Pillar Three: Evidentiary and Enforcement Reform

The third pillar realigns adjudication with the technical realities of agentic markets. Because supracompetitive prices in these markets no longer reliably signal a “meeting of the minds,” the law must pivot to a tractable inquiry into *design-based causation*. This pillar equips courts and agencies with the doctrinal and evidentiary tools to make that shift.

1. A Rebuttable Presumption of Causation

The framework adapts a familiar legal tool for settings where defendants uniquely control the most probative evidence: a rebuttable presumption.⁹³

Under this standard, a plaintiff who produces robust econometric evidence of *persistent supra-competitive outcomes* in a designated agentic market triggers a rebuttable presumption of design-based causation. The *burden of production*—not the ultimate burden of persuasion—then shifts to the developers and deployers of the market-facing agents. To rebut the inference of causation, they must produce evidence from their design record. A defendant who can point to current certifications and passing red-team reports would have a *prima facie rebuttal* to the presumption. Conversely, the failure to produce such evidence could support an adverse inference that the agent was negligently designed.

⁹³This structure parallels established burden-shifting frameworks in other areas of law. See *Griggs*, 401 U.S. at 431–32 (allowing disparate-impact cases to proceed on effects evidence with burden shifting); *Phila. Nat'l Bank*, 374 U.S. at 362–63 (structural presumption in merger cases).

Plaintiffs, however, would retain the ultimate burden of persuading the trier of fact. This approach does not create a *per se* offense; it supplies a structured, negligence-style rule-of-reason analysis tailored to the unique evidentiary challenges of agentic markets.

2. The “Design Record” as the Evidentiary Core

To overcome “black box” defenses and the routine loss of electronically stored information, the framework requires designated actors to preserve a “*design record*” for the model’s lifecycle plus a defined term. This record includes:

- Objective and reward function specifications;
- Default and deployed parameters (including exploration budgets, learning rates, and update cadence);
- Policy-update histories, rollout logs, and incident reports documenting observed convergence; and
- Certified antitrust red-team outputs, with configuration files and seeds sufficient for reproduction.

Courts can manage discovery of this sensitive material using familiar tools: protective orders for trade secrets, *in camera* review, court-appointed neutral experts under Federal Rule of Evidence 706, and special masters under Federal Rule of Civil Procedure 53. Where necessary, *algorithmic escrow* can substitute for broad source-code disclosure.

3. Remedies and Institutional Roles

Where liability is established, courts could order a range of remedies, including *code-level modifications* (such as mandated strategic randomness), *protocol-level adjustments* (like changes to update frequency or information display), and, where conflicts of interest are persistent, *structural separation*.

While comprehensive legislation is the ideal long-term solution, this framework can be advanced under existing authority. The Federal Trade Commission has broad power under Section 5

of the FTC Act to police “*unfair methods of competition*.”⁹⁴ The FTC, DOJ, and sectoral regulators can begin building a persuasive, *Skidmore*-worthy⁹⁵ body of practice by treating the deployment of uncertified, convergence-prone agents as an unfair method of competition. This would provide a powerful avenue for jump-starting the development of case law and enforcement norms even before new legislation is passed.

IV. THE FRAMEWORK IN PRACTICE: CASE STUDIES IN AGENTIC GOVERNANCE

Having proposed a new governance architecture, this Part moves from theory to practice, applying the framework to five case studies as doctrinal stress tests. The first two revisit the hotel and rental markets to show how the framework addresses the *present* deadlocks where Section 1 stalls. The next two examine the airline and grocery sectors, where it functions as a *proactive* tool to manage emerging threats and structural conflicts. The final case study projects forward to a mature, decentralized retail ecosystem where traditional agreement-based tools are under-identified. The callout below maps each market to the central legal failure the framework is designed to resolve.

⁹⁴See 15 U.S.C. § 45 (2022). The FTC has already applied its authority to police harmful digital design practices. See, e.g., Agreement Containing Consent Order, *In re Epic Games, Inc.*, FTC File No. 1923203 (Dec. 19, 2022), https://www.ftc.gov/system/files/ftc_gov/pdf/1923203EpicGamesACCO.pdf [<https://perma.cc/TDW9-RNNZ>] (addressing manipulative user interface designs, or “dark patterns”). The deployment of negligently designed, convergence-prone algorithms fits squarely within this framework.

⁹⁵See *Skidmore v. Swift & Co.*, 323 U.S. 134, 140 (1944) (holding that agency interpretations are entitled to respect proportional to their “power to persuade”).

Doctrinal Stress Tests: Mapping the Legal Failure Modes

Context	The Doctrinal Failure Mode	The Framework's Resolution
1. Hotels	The "Missing Rim" Problem. Courts cannot infer a horizontal agreement from parallel software use without specific allegations of pooled nonpublic data.	Evidentiary Pivot. Replaces the search for a "meeting of the minds" with an audit of the <i>design record</i> to establish design-based causation.
2. Rentals	The "Hub" Limit. Current hub-and-spoke theory reaches the vendor but cannot govern the conflicts of interest in a two-sided agentic ecosystem.	Ecosystem Liability. Extends the Duty to Disrupt to buyer agents and platforms, preventing them from reinforcing coordination.
3. Airlines	The "Reactive" Trap. Enforcement waits for harm, but tacit collusion is durable and difficult to unwind <i>ex post</i> .	Proactive Oversight. Mandates <i>ex ante</i> antitrust red-teaming and sand-box testing to detect convergence risks before deployment.
4. Groceries	The "Dual-Role" Conflict. No clear antitrust tool prevents a firm from operating both the pricing engine and the shopping agent in the same market.	Structural Integrity. Applies a "Glass-Steagall for agents" to prophylactically separate conflicting market functions.
5. Future Retail	The "Intent" Void. In decentralized markets, supracompetitive prices emerge without any human conspiracy to police.	Design-Based Liability. Creates causes of action grounded in <i>negligent design</i> and failure to disrupt, independent of proving human intent.

A. Case Study 1: The Hotel Market—From Doctrinal Deadlock to a Duty to Disrupt

As detailed in Part I, recent antitrust litigation against hotel operators has stalled where complaints allege only the parallel use of revenue-management software.⁹⁶ Courts have dismissed these suits for failing to plead a horizontal “rim,” absent specific allegations that nonpublic competitor **data were pooled** to generate competitor-conditioned prices.⁹⁷

Applied to the hotel sector, the framework would transform the very disputes that stalled in *Gibson* and *Dai*. A hotel operator like Hyatt or its vendor, IDEaS (a defendant in the *Dai* litigation), would presumptively be required to bring into compliance any pricing agent designed solely to maximize short-term revenue.⁹⁸ Compliance would require auditable design changes. For example, its *long-horizon objective* might be specified to maximize quarterly revenue per available room (RevPAR) while simultaneously penalizing sustained price clustering—for instance, maintaining prices within 2% of a defined peer set for more than 72 hours. *Strategic randomness* could require the system to periodically test meaningful price deviations, making its behavior harder for rivals to predict and follow.

That account, however, freezes the market’s current state. The hotel sector is evolving toward fully two-sided *agentic* markets: seller-side RMS agents that adjust prices at high cadence now interact with buyer-side AI assistants, while online travel agencies (OTAs) provide the *protocol layer* on which they learn.⁹⁹ Room prices are becoming an *emergent property* of interacting seller, buyer, and intermediary designs, and not the sellers’ unilateral choices.

⁹⁶ See *supra* notes 30, 31.

⁹⁷ See, e.g., *Gibson*, No. 24-3576 (9th Cir. Aug. 15, 2025) (affirming dismissal); *Gibson*, No. 2:23-cv-00140-MMD-DJA, slip op. at 3–6 (D. Nev. May 8, 2024) (dismissing for failure to plausibly allege a horizontal agreement); *Cornish-Adebiyi*, No. 1:23-cv-02536 (D.N.J. Sep. 30, 2024) (dismissing for lack of allegations that confidential competitor data were pooled), *appeal docketed*, No. 24-3006 (3d Cir. Oct. 29, 2024); *Dai*, No. 24-CV-02537-JSW, slip op. at 3–4 (N.D. Cal. July 18, 2025). Cf. *In re RealPage*, 709 F. Supp. 3d at 509–12 (allowing claims with detailed pooling-and-use allegations to proceed).

⁹⁸ IDEAS, *Hyatt Will Deploy IDEaS Platform*, IDEAS.COM (Aug. 11, 2022), <https://ideas.com/news/hyatt-will-deploy-ideas-platform/> [<https://perma.cc/P4D3-BNTK>]; see also IDEAS REVENUE SOLUTIONS, *About*, IDEAS.COM (last visited Nov. 14, 2025), <https://ideas.com/about/> [<https://perma.cc/9PSS-CKUM>].

⁹⁹ See *supra* note 15.

The framework would require the buyer agent of a travel platform like Navan to, instead of myopically booking the lowest price—a behavior that makes retaliation against a price-cutter swift and certain—*amplify the reward for defection* by alerting users and shifting significant demand toward the first discounter while also *blunting retaliation* by maintaining some allocation to the original defector rather than instantly reverting when rivals respond.¹⁰⁰

In addition, the framework addresses the structural conflicts and protocol risks posed by intermediaries. An OTA like Expedia Group, which owns consumer-facing platforms (such as Hotels.com) and also offers revenue insights and pricing tools like Rev+ to its hotel partners, faces an inherent conflict.¹⁰¹ The “Glass-Steagall for agents” would presumptively require structural separation of these functions.¹⁰² Furthermore, a *protocol audit* of a platform like Booking.com would examine whether its ranking and display rules facilitate collusion, potentially requiring design choices—such as information lags or temporary visibility boosts for price-cutters—that impede real-time monitoring and punishment.¹⁰³

For the hotel cases that failed under Section 1, the evidentiary reform provides a direct path forward. Plaintiffs could make an econometric showing that hotel prices have persistently exceeded competitive benchmarks. The burden would then shift to the defendants—e.g., Caesars, MGM, and their RMS vendors—to produce the design records and certifications showing their systems were engineered to compete, not converge. The absence of such evidence, which the *Gibson* plaintiffs could not obtain under current pleading and discovery rules, would support an adverse inference of negligent design.

More broadly, this approach transitions the locus of governance from speculative, after-the-fact litigation over human intent to proactive, verifiable compliance based on the design record. It does not simply create a new theory of liability; it fosters a market where competition is an engineered

¹⁰⁰ See *supra* Part II.B (describing the collusion-breaker equilibrium).

¹⁰¹ EXPEDIA GROUP, *Maximize Hotel Distribution & Revenue*, PARTNER HUB (last visited Nov. 14, 2025) (stating that Rev+ helps “optimize your pricing strategy”), <https://partner.expediagroup.com/en-us/industries/hotels> [<https://perma.cc/KQE9-5TPS>].

¹⁰² See *supra* Part III.A.

¹⁰³ See, e.g., *Airline Tariff Publ’g Co.*, 836 F. Supp. at 11–12.

and continuously monitored property, making litigation the exception rather than the rule.

B. Case Study 2: The Rental Market—From Hub-and-Spoke to Ecosystem Governance

The landmark antitrust litigation against RealPage, Inc. provides a high-profile example of a modern “hub-and-spoke” conspiracy and serves as a crucial testbed for the proposed framework.¹⁰⁴ The core allegation is that competing landlords delegated pricing authority to RealPage’s algorithm, which used pooled, non-public data to recommend inflated rents, with RealPage promising its software would “help[] curb [landlords’] instincts to respond to down-market conditions.”¹⁰⁵

Legal challenges against RealPage have seen more initial success than similar hotel cases. In the multidistrict litigation, the court denied the motion to dismiss the multifamily plaintiffs’ complaint (while dismissing the student-housing plaintiffs’ complaint), finding that allegations of data pooling and widespread adoption plausibly alleged a horizontal agreement.¹⁰⁶ This is because the alleged conduct—pooling nonpublic data for use by a common pricing algorithm—fits more neatly within a traditional hub-and-spoke theory. But that fit may weaken as rental markets evolve toward two-sided agentic settings: a nascent ecosystem of “tenant AI” agents could shift the market from one-sided algorithmic pricing (with individual tenants as price-takers) toward repeated interactions between landlord agents and AI-powered tenant agents.¹⁰⁷ In this environment, the doctrinal gaps diagnosed in Parts I and II re-emerge.

In this emerging ecosystem, the framework’s first pillar, *structural integrity*, is directly implicated. A scenario in which RealPage, the dominant seller-side agent (alleged by the DOJ to control over 80% of the commercial revenue management market),¹⁰⁸ developed its own renter-facing agent would create a profound conflict of interest. The company would have every incentive to

¹⁰⁴ See Complaint, *supra* note 12, ¶¶ 4–6; see also Memorandum of Law, *supra* note 26, at 19–21 (alleging that competing landlords delegated pricing to a common algorithm using pooled nonpublic data).

¹⁰⁵ See Complaint, *supra* note 12, ¶ 1.

¹⁰⁶ See *In re RealPage*, 709 F. Supp. 3d at 509–12.

¹⁰⁷ See *supra* notes 14, 15, 17.

¹⁰⁸ See Complaint, *supra* note 12, ¶ 13.

design its buyer agent to be a passive price-taker, thereby protecting the supracompetitive rents generated for its landlord clients. The proposed “*Glass-Steagall for agents*” would prophylactically forbid such an arrangement, or at minimum require audited, enforceable separation.¹⁰⁹

The second pillar—the *Duty to Disrupt*—governs the tenant agents themselves. An agent engineered to myopically accept the algorithm’s listed price would have its predictability exploited by the seller-side algorithm, further stabilizing the collusive equilibrium. Under the proposed framework, the developer of this negligently designed agent could be held liable for breaching its affirmative duty to engineer procompetitive features, such as strategic randomness, long-horizon objectives, and market-state awareness.¹¹⁰

The third pillar—*design-based evidentiary reform*—supplies an administrable path where Section 1 doctrine stalls. A robust econometric showing of persistent rental overcharges triggers a rebuttable presumption of design-based causation. The burden of production then shifts to the developers of market-facing agents to produce their design records and certified red-team results. A landlord-side vendor could rebut the presumption by showing it implemented rivalry-by-design and passed adversarial simulations, while a tenant-side vendor that failed to show it had tested for and mitigated convergence risks would face an adverse inference of negligent design. This approach shifts the focus from the backward-looking question of whether landlords and RealPage formed a conspiracy to the forward-looking task of engineering rivalry across the entire ecosystem as it becomes fully agentic.

C. Case Study 3: The Airline Market—From Reactive Enforcement to Proactive Design Audits

The deployment of agentic pricing by Delta Air Lines presents a useful test case for governing a two-sided agentic market in airline pricing. As consumer-side AI agents are not yet dominant, the letters from the concerned senators only warn of “surveillance-based price setting” and the risk

¹⁰⁹ See *supra* notes 19, 90.

¹¹⁰ See *supra* Part II.B.

that consumers will be charged up to their personal “pain point,” demanding information about data privacy and consumer fairness.¹¹¹ The entry of a powerful buyer-side AI would, however, require a proactive regulatory approach to shape the market’s evolution.

In this scenario, what approach could regulators adopt to address the systemic risk of tacit collusion? An investigation predicated on finding a traditional “agreement” between Delta and its competitors would likely prove futile for the reasons detailed in Part I. An investigator armed with the framework proposed in Part III, however, would approach the situation with a better suited set of tools, shifting the core investigative question from whether firms *agreed* to collude to whether their systems were *designed* to compete.

The investigation would begin not with a search for collusive emails, but with compulsory process for *algorithmic evidence*, leveraging the framework’s rebuttable presumption of design-based causation to shift the evidentiary burden.¹¹² This demand would be structured around the pillars of the framework.

Under the proactive oversight regime, regulators would compel the production of *certified antitrust red-teaming reports* discussing, *inter alia*, “Did Fetcherr simulate how its AI performs when interacting with other pricing AIs in an oligopolistic market?” “Do these simulations show a tendency toward supracompetitive price convergence?” They would demand the design documents for the AI’s “*antitrust safety architecture*,” with any claims of trade secret addressed through established legal mechanisms like *in camera* review or third-party audits under a protective order. The inquiry would examine whether the design includes a “collusion-breaker” module or hard-coded rules that prevent simple price-following and promote strategic randomness. The absence of such features, or reports showing a high risk of convergence, would form the evidentiary core of a claim that the system was negligently designed.

¹¹¹ See *supra* notes 4–5.

¹¹² While the FTC lacks jurisdiction over air carriers, see 15 U.S.C. § 45(a)(2) (2022), the DOJ retains Sherman Act authority and can issue Civil Investigative Demands under the Antitrust Civil Process Act, 15 U.S.C. § 1312 (2022). The Department of Transportation (DOT) holds parallel authority to police “unfair or deceptive practice[s] or unfair method[s] of competition” by air carriers under 49 U.S.C. § 41712 (2022). Cf. *Airline Tariff Publ’g Co.*, 836 F. Supp. at 9 (addressing computerized fare dissemination and signaling).

The inquiry would scrutinize the market’s structural integrity. Investigators would examine the corporate relationships of Fetcherr, including whether the company also provided technology to any major consumer-facing travel agents (e.g., Hopper and Navan’s Ava). A finding of such a relationship would trigger scrutiny under the proposed *structural separation* rule—the “Glass-Steagall for AI Agents”—to prophylactically dismantle market structures that create inherent conflicts of interest.

Armed with this evidence, the DOJ or DOT could bring an enforcement action. Delta and Fetcherr would likely mount a defense consistent with Delta’s public statements: that the system is not used for personalized pricing, is only a limited test on roughly 3% of domestic inventory, and is merely a procompetitive tool for responding more efficiently to market dynamics.¹¹³ Under the proposed framework, such claims would be tested empirically. If their red-teaming reports are insufficient to rebut the presumption of anticompetitive harm, regulators could require them to test the AI in a secure regulatory sandbox. Failure to pass such tests could lead to remedies like a consent decree mandating specific design modifications to embed procompetitive features. The legal claim would be that deploying a powerful pricing algorithm with a foreseeable tendency toward supracompetitive convergence, and without a robust, certified safety architecture, constitutes an unfair or anticompetitive practice, shifting the focus from the backward-looking legal question “Did Delta *agree* to collude?” to the more relevant and forward-looking inquiry: “Did Delta and Fetcherr act responsibly by designing and deploying a system to compete fairly in a two-sided agentic market?”

D. Case Study 4: The Grocery Market—Preventing Conflicts of Interest in a Nascent Agentic Ecosystem

The U.S. grocery market, already under federal scrutiny for elevated prices and market-structure concerns, presents a particularly pressing challenge for algorithmic governance.¹¹⁴ While the

¹¹³ See *supra* notes 1, 7.

¹¹⁴ See FED. TRADE COMM’N, *FTC and Justice Department Host First Strike Force on Unfair and Illegal Pricing Meeting* (Aug. 1, 2024), <https://www.ftc.gov/news-events/news/press-releases/2024/08/ftc-justice-department-host-first->

risk of tacit collusion exists here as in other oligopolies, a more profound structural conflict is emerging: the same corporate entities are building both sides of the market.

Companies like Walmart and Amazon are leaders in both seller-side price optimization and the development of consumer-facing AI shopping assistants (e.g., Walmart’s “Sparky,” Amazon’s “Rufus”).¹¹⁵ This dual role creates a direct and profound incentive misalignment: the corporation that profits from maximizing revenue through its pricing systems is also building the tool that consumers rely on to navigate the market. While current iterations function as vertically integrated assistants, the risk becomes acute as these tools evolve into default shopping interfaces that steer demand.

This scenario clearly demonstrates the necessity of the framework’s pillar of *structural integrity*. The proposed “*Glass-Steagall for agents*” would directly address this inherent conflict.¹¹⁶ Such a prophylactic rule would presumptively prohibit a single corporate entity from operating both a market-significant seller-side pricing platform and a consumer-facing buyer agent within the same market, absent audited and enforceable separation. This structural separation is not a punishment; it is a preventive measure designed to dismantle a market structure that creates an overwhelming incentive for a company to design its buyer agent to be a passive, price-taking assistant, thereby protecting its own retail margins at the expense of the consumer.

Furthermore, it highlights the importance of *proactive oversight*. Rather than waiting for evidence of consumer harm to mount, the framework would empower the FTC to act preemptively via consent orders or safe-harbor guidance. It would mandate *algorithmic auditing* for both the seller-side pricing systems and the buyer-side shopping assistants. Developers would be required

strike-force-unfair-illegal-pricing-meeting [https://perma.cc/LB5S-AEUJ] (noting Chair Khan’s plan to launch an inquiry into grocery prices); see also JUSTIN WISE, *FTC Chair Khan Pushes for Inquiry Into Elevated Grocery Prices*, BLOOMBERG LAW (Aug. 1, 2024), https://news.bloomberglaw.com/antitrust/ftc-chair-khan-pushes-for-inquiry-into-elevated-grocery-prices [https://perma.cc/E7CX-PH4B] (reporting on FTC Chair Khan’s push for an inquiry into elevated grocery prices).

¹¹⁵ See *supra* notes 14, 17. See also Complaint ¶¶ 416–27, *FTC v. Amazon.com, Inc.*, No. 2:23-cv-01495-JHC (W.D. Wash. Nov. 2, 2023), https://www.ftc.gov/system/files/ftc_gov/pdf/1910134amazonecommercecomplaintrevisedredaction_s.pdf [https://perma.cc/CCF7-W6BD] (alleging Amazon’s “Project Nessie” pricing algorithm “generated over \$1 billion in excess profit” by testing price increases).

¹¹⁶ See *supra* note 19.

to conduct “antitrust red-teaming” and submit certified reports demonstrating that their agents are designed to compete fairly and effectively. For a market as critical to public welfare as groceries, regulators could even require that new, market-significant agents be tested in a *regulatory sandbox*. This would allow the FTC to observe the emergent, interactive effects of these AIs in a controlled environment *before* they are deployed to millions of consumers.

E. Case Study 5: A Hypothetical Future—Governing the “ShopSphere” Retail Market

To illustrate the full power of the framework in a mature, decentralized ecosystem, consider a hypothetical e-commerce platform, “ShopSphere.” On the supply side, thousands of independent merchants use one of two dominant, third-party pricing AIs, “Apex Pricing” or “Vector Revenue.” On the demand side, millions of consumers use one of two dominant AI shopping agents, “Aura” or “Nexus.” Econometric analysis shows that prices on MarketSphere have been persistently fifteen percent above competitive levels (controlling for cost and demand shocks), yet there is no evidence of any human conspiracy among the merchants or the AI developers.

Under the proposed framework, a consumer class-action lawsuit could be filed against the developer of *Aura*, alleging a breach of the legally mandated “*Duty to Disrupt*.”¹¹⁷ The evidentiary basis for the claim would be simulation evidence demonstrating that *Aura* is a “myopic” agent, designed only to maximize immediate savings. As established in Part II.B, this behavior is known to reinforce, rather than break, collusive equilibria by making retaliation against price-cutters swift and credible. The agent’s design would demonstrably lack the required strategic randomness and long-term foresight. A defense predicated on the argument that the agent fulfilled its duty by securing the lowest price in individual transactions would be insufficient; the “*Duty to Disrupt*” establishes an obligation to protect the consumer’s sustainable, long-term welfare by fostering a competitive market. A court could thus find that the agent’s myopic design foreseeably harmed its users by reinforcing the supracompetitive equilibrium, establishing the developer’s liability.

¹¹⁷ See RESTATEMENT (THIRD) OF TORTS: PRODS. LIAB. § 2(b) (Am. L. Inst. 1998).

Simultaneously, under the Duty to Disrupt framework (as codified), the DOJ could bring an antitrust suit against the developers of the seller AIs. The DOJ would satisfy its initial burden by presenting the econometric evidence of the supracompetitive pricing *effect*. Under the proposed *burden-shifting framework*, the burden of production would then move to the developers to produce their design records and certifications to rebut the inference of design-based causation. Suppose Apex Pricing’s developer could produce its certified antitrust red-teaming reports demonstrating robust anti-convergence safeguards, providing a *prima facie* rebuttal to the presumption. Vector Revenue, having conducted no such testing and having designed a purely profit-maximizing algorithm, fails to rebut the presumption. Its design could be found to have foreseeably contributed to the anticompetitive market convergence, thereby establishing liability.

The framework thus enables two distinct but complementary legal actions that would otherwise be unavailable. It allows for holding a buyer-agent developer accountable for negligent design and a seller-agent developer accountable for facilitating convergence. Such outcomes would be unlikely under current law, which would face a doctrinal impasse in the absence of a human conspiracy.

V. BROADER LEGAL AND POLICY IMPLICATIONS

The framework’s effectiveness depends on aligning incentives across these institutions, ensuring that the duty is not just a doctrinal innovation but a functional governance regime. This Part steps back from the doctrinal analysis of Parts I and II and the case studies of Part IV to map the broader legal, institutional, and market consequences of adopting the “Duty to Disrupt.” It moves from the *what* of the proposal to the *how*, exploring its full lifecycle—from internal corporate governance and market assurance, to agency information-gathering and adjudication, to judicial remedies and cross-border coordination.

The analysis proceeds in five stages. First, the Part details the core governance model, examining the new corporate oversight obligations and the administrative law posture required

for public enforcement in a post-*Chevron* world. Second, it specifies the practical mechanics of enforcement, outlining a new evidentiary architecture for the “design record,” clarifying the path for private litigation, and calibrating a sophisticated regime of remedies and safe harbors. Third, it defines the framework’s jurisdictional and temporal boundaries, addressing the critical international, federalism, and transitional dimensions that will shape its application. Fourth, it weighs the broader economic and innovation effects, analyzing the proposal’s welfare implications and its impact on market entry. Finally, it concludes by outlining a concrete implementation path, demonstrating that the framework is not only doctrinally sound but also administratively viable, economically sensible, and adaptable to the complexities of a global, agent-driven economy.

A. The Core Governance Model

1. Private-Sector Governance: From Compliance to Corporate Oversight

By conditioning liability on auditable design choices, the duty transforms competition safeguards from an *ex post* legal defense into an *ex ante* engineering requirement. This shift institutionalizes the creation of a “*design record*”: the specified objective functions, programmed constraints, default parameters, policy-update logs, and certified interactive test results that form the auditable by-products of responsible AI development. This moves the locus of legal discovery from ambiguous executive emails to the technical facts that actually govern market behavior.

More fundamentally, it reframes the duty as a matter of corporate governance, implicating a board’s duty of oversight for “mission-critical” risks. In an agentic market, the foreseeable risk of causing or contributing to algorithmic convergence becomes a core operational hazard that boards and senior management must monitor through reasonable reporting and escalation systems—potentially supporting *Caremark*-style claims where fiduciaries in bad faith fail to implement or oversee such systems. While oversight liability under Delaware fiduciary law remains “possibly the most difficult theory in corporation law upon which a plaintiff might hope to win a judgment,”¹¹⁸

¹¹⁸*In re Caremark Int’l Inc. Derivative Litig.*, 698 A.2d 959, 967 (Del. Ch. 1996).

the explicit designation of agentic convergence as a mission-critical risk would make a board’s “utter failure” to implement reporting systems far easier to plead.¹¹⁹

To make this oversight function and the underlying duty administrable, the framework mandates a clear *design-record retention rule*. Developers and deployers must preserve, for the model’s lifecycle plus a specified number of years, a standardized design record—comprising objective specifications, parameter logs, and test results—as detailed in the evidentiary architecture below. Because these records will often contain trade secrets, their disclosure in litigation or audits should be managed through protective orders, *in camera* review, or secure escrow regimes.

This clarity enables the emergence of robust markets for assurance that reinforce the duty through private ordering. First, a market for *third-party certification* can arise, with independent labs and standards bodies such as IEEE or ISO validating that an agent’s architecture is governed by a risk-management framework capable of enforcing rivalry-by-design.¹²⁰ Second, these auditable artifacts make risk actuarially intelligible, catalyzing a market for *AI liability insurance*. Insurers can price policies based on the robustness of an agent’s certified safeguards, offering lower premiums for procompetitive designs and effectively acting as private regulators. This market-driven mechanism not only internalizes the externalities of negligent design but also democratizes access to compliance expertise, as insurers may develop standardized templates for red-teaming that smaller developers can adopt at low cost.¹²¹ Finally, *procurement* becomes a powerful enforcement lever. Government agencies and sophisticated enterprise customers can contractually require vendors to provide a current “antitrust red-team” certificate, a right to independent testing, and a

¹¹⁹See *Stone v. Ritter*, 911 A.2d 362, 370 (Del. 2006) (holding that liability requires a showing of “utter failure” to implement a reporting system or a “conscious failure” to monitor it); *Marchand v. Barnhill*, 212 A.3d 805, 820–24 (Del. 2019) (applying the duty to “mission critical” compliance risks); see also *In re McDonald’s Corp. S’holder Derivative Litig.*, 289 A.3d 343, 378–79 (Del. Ch. 2023) (clarifying that corporate officers, not just directors, owe oversight duties regarding mission-critical risks within their sphere of responsibility).

¹²⁰See, e.g., INT’L ORG. FOR STANDARDIZATION & INT’L ELECTROTECHNICAL COMM’N, ISO/IEC 42001:2023, *Information Technology — Artificial Intelligence — Management System* (2023), <https://www.iso.org/standard/81230.html> [<https://perma.cc/L6UA-YD7V>] (establishing an international management-system standard for AI governance and risk management).

¹²¹See OMRI BEN-SHAHAR & KYLE D. LOGUE, *Outsourcing Regulation: How Insurance Reduces Moral Hazard*, 111 MICH. L. REV. 197, 217–28 (2012) (analyzing how insurers act as de facto regulators by promoting safety standards and best practices via premium pricing and underwriting requirements).

duty to notify on material model updates, creating a flywheel effect for adoption.

This new governance model also requires new legal and contractual safeguards. The duty must be established as *non-waivable* by contract (and, as discussed *infra*, shielded from mandatory arbitration), preventing developers from using end-user license agreements (EULAs) to impose waivers that would effectively undermine private enforcement.¹²² Furthermore, this governance model also raises complex questions of *liability allocation* in multi-agent chains. Where harm emerges from the interaction of multiple certified agents, courts and regulators will need to develop rules for apportioning responsibility, potentially based on each agent’s relative contribution to the anticompetitive outcome.

2. *The Public Law Framework: Administrative Authority and Judicial Review*

The framework’s durability rests on a sound administrative law foundation. In a post-*Loper Bright* world, where courts no longer grant *Chevron* deference to agency interpretations of ambiguous statutes, the framework cannot depend on judicial deference to novel or aggressive readings of statutory authority. Instead, its legitimacy must be built upon a *Skidmore*-worthy record—one that is thorough, empirically grounded, and consistently developed through the use of specific statutory tools.¹²³ This approach mitigates the risk of Major Questions Doctrine challenges by anchoring enforcement in established competition and consumer protection mandates—policing unfair methods of competition and deceptive practices—rather than claiming sweeping, transformative new regulatory powers.¹²⁴

Agencies can build this record using their existing statutory toolkit. The FTC, for instance, can issue targeted orders under Section 6(b) of the FTC Act to collect design records and interaction

¹²² See J. MARIA GLOVER, *Mass Arbitration*, 74 STAN. L. REV. 1283, 1304 (2022) (reporting that “virtually all of these include class-action waivers”).

¹²³ See *Loper Bright Enters. v. Raimondo*, 144 S. Ct. 2244 (2024) (overruling *Chevron U.S.A. Inc. v. Nat. Res. Def. Council, Inc.*, 467 U.S. 837 (1984), but acknowledging that agency interpretations may still claim respect under the standard of *Skidmore v. Swift & Co.*); see also *Skidmore*, 323 U.S. at 140 (holding that an agency’s interpretation is entitled to respect based on its “power to persuade,” which depends on the thoroughness of its consideration and the validity of its reasoning).

¹²⁴ See *West Virginia v. EPA*, 142 S. Ct. 2587, 2609–10 (2022) (requiring clear congressional authorization where agencies claim powers of vast economic and political significance).

telemetry from market participants, forming the empirical basis for subsequent action.¹²⁵ Enforcement would then proceed primarily through case-specific administrative complaints under Section 5, which prohibits “unfair methods of competition,” and suits for forward-looking injunctions under Section 13(b).¹²⁶ Sector-specific regulators like the Department of Transportation can exercise parallel authority over air carriers, while the DOJ retains its core Sherman Act powers.¹²⁷ To ensure harmonized enforcement, these agencies could establish a joint task force to develop consistent red-teaming protocols and share technical expertise. The regulatory sandbox, in this context, functions as a tool for building the evidentiary record, operating as a “CCAR-for-algorithms”—a periodic, regulator-specified competition stress test with public performance metrics, analogous to the stress tests used in banking.¹²⁸

Generalist courts, in turn, would not be asked to become software engineers but to serve as expert arbiters of a technical record. The admissibility of simulation evidence would be policed under Rule 702 and the principles of *Daubert* and *Kumho Tire*, ensuring that claims are grounded in reliable methodologies.¹²⁹ To manage the complexity of this evidence, courts could use familiar tools: protective orders to safeguard intellectual property, court-appointed neutral experts under Rule 706, and special masters under Rule 53 to oversee discovery and compliance.¹³⁰ Where liability is found, the remedial palette is both technical and verifiable. A court could order code-level modifications (e.g., mandating a minimum exploration budget), protocol-level adjustments

¹²⁵ 15 U.S.C. § 46(b) (2022).

¹²⁶ 15 U.S.C. § 45 (2022); 15 U.S.C. § 53(b) (2022). While *AMG Cap. Mgmt., LLC v. FTC*, 141 S. Ct. 1341 (2021), limited the FTC’s ability to seek retrospective monetary relief under Section 13(b), its authority to pursue injunctive remedies remains intact.

¹²⁷ See 49 U.S.C. § 41712 (2022) (prohibiting “unfair or deceptive practice[s]” by air carriers); 15 U.S.C. §§ 1–2 (2022) (Sherman Act); 15 U.S.C. § 25 (2022) (authorizing DOJ to seek injunctive relief).

¹²⁸ See, e.g., Bd. of Governors of the Fed. Reserve Sys., *Comprehensive Capital Analysis and Review 2019: Assessment Framework and Results* (June 2019), <https://www.federalreserve.gov/publications/files/2019-ccar-assessment-framework-results-20190627.pdf> [<https://perma.cc/6KKK-54LM>] (describing the stress-testing and supervisory assessment framework).

¹²⁹ See Fed. R. Evid. 702; *Daubert v. Merrell Dow Pharms., Inc.*, 509 U.S. 579 (1993); *Kumho Tire Co. v. Carmichael*, 526 U.S. 137 (1999) (extending *Daubert* to technical and other specialized expertise).

¹³⁰ See Fed. R. Civ. P. 26(c)(1)(G) (protective orders); Fed. R. Evid. 706 (court-appointed experts); Fed. R. Civ. P. 53(a)(1)(C) (authorizing special masters for pretrial and post-trial matters). For particularly complex or recurring issues, a specialized tribunal could be considered. Cf. the Copyright Royalty Board or the Patent Trial and Appeal Board.

(e.g., throttling price update frequency), or, where conflicts of interest are intractable, structural separation.

B. The Mechanics of Enforcement and Compliance

1. *The Evidentiary Architecture: Preserving and Accessing the “Design Record”*

The evidentiary fulcrum of this regime is the “*design record*.” To prevent spoliation and overcome “black-box” defenses, the framework establishes a clear evidence architecture. This begins with a mandatory *retention rule*: under the Duty-to-Disrupt regime, developers and deployers must preserve, for the model’s lifecycle plus a specified term, a standardized set of artifacts, including (i) reward and objective function specifications; (ii) default and deployed parameters (including exploration budgets, learning rates, and update cadence); (iii) policy-update logs, hashed for integrity and time-stamped; (iv) certified adversarial-simulation results, with inputs and seeds sufficient for reproducibility; and (v) interaction telemetry sufficient to reconstruct key market episodes. While this is a regulatory duty, a failure to preserve this record once litigation is reasonably foreseeable would support sanctions under familiar spoliation principles, and—upon a showing of intent to deprive—adverse-inference instructions under Rule 37(e).¹³¹

Access to this sensitive information must be balanced with the need to protect intellectual property. In private litigation, courts can manage discovery using Rule 26(c) protective orders, *in camera* review, and the appointment of neutral technical experts.¹³² Where source-code disclosure is impractical, a court could consider ordering *algorithmic escrow* with reproducible-build attestations. For public oversight, agencies should adopt a model of *protective transparency*: they would issue public-facing summaries of certification results and sandbox findings, while protecting the confidential technical annexes from disclosure under FOIA Exemption 4.¹³³ This architecture

¹³¹ See Fed. R. Civ. P. 37(e)(2) (authorizing sanctions for failure to preserve electronically stored information, including adverse inferences upon a finding of intent to deprive).

¹³² See Fed. R. Civ. P. 26(c)(1)(G). Trade secrets are recognized as a cognizable form of property, reinforcing the need for robust protective measures. See *Ruckelshaus v. Monsanto Co.*, 467 U.S. 986, 1001–04 (1984).

¹³³ 5 U.S.C. § 552(b)(4) (2022); see *Food Mktg. Inst. v. Argus Leader Media*, 139 S. Ct. 2356, 2363–66 (2019) (defining “confidential” commercial information broadly).

ensures that the evidence needed for accountability is available to those with a right to access it, while safeguarding the legitimate commercial interests that drive innovation.

2. *Private Enforcement: Standing, Class Actions, and Non-Waivable Duties*

The viability of private suits, a critical complement to agency action, requires addressing doctrinal and contractual frictions. On the question of standing, particularly the direct purchaser rule established in *Illinois Brick*,¹³⁴ where a buyer agent intermediates a transaction, defendants may argue that the end consumer is an indirect purchaser barred from recovering federal damages. The framework, however, creates two distinct paths for private claims. End-users may still be able to sue seller-side actors directly, arguing that under the logic of *Apple v. Pepper*, they remain the direct purchasers because the agent is merely a conduit for their transaction—a determination that will turn on specific facts of agency and privity.¹³⁵ Separately, consumers can bring a tort-like claim against their *own agent's developer* for a breach of the Duty to Disrupt, framing the harm as a product of negligent design. This second cause of action is not constrained by antitrust's direct purchaser rule, and state-level *Illinois Brick* repealers provide a further avenue for recovery.¹³⁶

The duty must be shielded from contractual evasion. AI vendors and platforms frequently embed mandatory arbitration clauses and class-action waivers in their end-user license agreements (EULAs), which would effectively gut private enforcement. To ensure the framework has teeth, the Duty to Disrupt must be established as *non-waivable* by contract for consumers and small- and medium-sized enterprises, mirroring statutory non-waiver provisions in other contexts.¹³⁷ However, because a non-waiver of substantive rights does not automatically preclude mandatory arbitration, effective implementation may require an explicit congressional command restricting predispute arbitration and class-action waivers for these specific claims.

¹³⁴ See *Ill. Brick Co.*, 431 U.S. at 735–36.

¹³⁵ See *Apple v. Pepper*, 139 S. Ct. at 1520–23.

¹³⁶ See *California v. ARC Am. Corp.*, 490 U.S. 93, 101–02 (1989) (holding that federal antitrust law does not preempt state indirect-purchaser statutes).

¹³⁷ See, e.g., Securities Act of 1933 § 14, 15 U.S.C. § 77n (2022). Cf. *CompuCredit Corp. v. Greenwood*, 565 U.S. 95, 103–04 (2012) (enforcing arbitration absent a “contrary congressional command”).

The framework is designed to facilitate class-action certification. The uniform design records, simulation outputs, and market-wide econometric evidence of supracompetitive pricing provide the kind of common proof needed to satisfy the predominance requirement of Federal Rule of Civil Procedure 23(b)(3).¹³⁸ Rather than devolving into thousands of individual inquiries, a class-action lawsuit could focus on the common question of whether an agent’s design had a foreseeable tendency to cause or sustain anticompetitive harm, with damages calculated on a market-wide basis. This makes collective redress for consumers not just possible, but practicable.

3. Remedies and Safe Harbors: Calibrating Liability and Incentives

To be effective and avoid chilling innovation, the framework’s enforcement tools must be graduated, verifiable, and prefer performance-based constraints over intrusive code-level mandates. Where a breach of the duty is found, courts can order forward-looking injunctive relief, and agencies can seek comparable performance-based conditions through settlements and consent orders, tailored to the specific mechanism of harm.¹³⁹ These include *parameter-level constraints*, such as mandating minimum randomized exploration budgets, requiring long-horizon objective functions, or dampening an agent’s sensitivity to short-run retaliation.¹⁴⁰ They can also impose *protocol-level throttles* by, for example, limiting price-update frequencies or adjusting information visibility rules on a platform to blunt the credibility of punishment. For repeat offenders or high-risk markets, courts could appoint an *independent monitor* with the authority to re-run sandbox tests on material updates, require *stipulated penalties for noncompliance*, or even order *algorithmic escrow*. In cases of intractable conflicts of interest, *structural separation* remains the ultimate remedy.

On the other side of the ledger, the framework must provide ex ante clarity for responsible actors. This is achieved through a robust *safe harbor*: a developer or deployer that maintains a current, third-party certification demonstrating its agent’s procompetitive design would receive

¹³⁸Fed. R. Civ. P. 23(b)(3) requires that “questions of law or fact common to class members predominate over any questions affecting only individual members.” See also *supra* note 68 (discussing *Comcast* and *Olean*).

¹³⁹See 15 U.S.C. § 25 (2022); 15 U.S.C. § 26 (2022).

¹⁴⁰See *United States v. Microsoft Corp.*, 253 F.3d 34, 105 (D.C. Cir. 2001) (affirming the court’s broad discretion to fashion relief that “unfetter[s] a market from anticompetitive conduct,” including conduct remedies).

a rebuttable presumption of compliance with the Duty to Disrupt. This safe harbor would be time-bound, expiring automatically upon a material model update or after a defined period, and would be conditioned on the firm’s cooperation with regulatory audits. This provides a powerful incentive for firms to invest in competition-by-design and internalize compliance.

The enforcement toolkit should be expanded with two additional mechanisms to create a self-policing ecosystem. First, Congress should create *researcher safe harbors* through limited exemptions from the Digital Millennium Copyright Act’s anti-circumvention provisions and the Computer Fraud and Abuse Act.¹⁴¹ These would protect bona fide, good-faith security and competition researchers who probe market-facing agents for vulnerabilities. Second, agencies should develop a “*Leniency 2.0*” program. A developer or deployer that proactively detects convergence-prone behavior, self-reports it to regulators, disables the offending policy, and cooperates with diagnostic testing should receive significant penalty mitigation. This adapts the logic of traditional cartel leniency to an environment where culpability turns on design and response, not secret meetings, encouraging firms to find and fix problems before they cause widespread harm.¹⁴² Together, certification, researcher access, and leniency create a closed-loop system: detecting risk, verifying safeguards, and incentivizing self-correction before harm becomes systemic.

C. Jurisdictional and Contextual Boundaries

1. *The International Dimension: Coordination, Comity, and Cross-Border Enforcement*

Agentic markets are inherently global; code travels across borders, and the largest platforms and developers operate worldwide. A purely domestic U.S. framework would be incomplete, risking both jurisdictional conflicts and regulatory arbitrage where firms could deploy less-regulated

¹⁴¹ See 17 U.S.C. § 1201 (2022) (DMCA); 18 U.S.C. § 1030 (2022) (CFAA).

¹⁴² See U.S. DEP’T OF JUSTICE, ANTITRUST DIV., *Antitrust Division Leniency Policy and Procedures* (updated Mar. 2024), [https://www.justice.gov/media/1226796/dl?inline=\[https://perma.cc/E33F-YWVQ\]](https://www.justice.gov/media/1226796/dl?inline=[https://perma.cc/E33F-YWVQ]); see also U.S. DEP’T OF JUSTICE, *Antitrust Division Corporate Leniency Policy* (1993), <https://www.justice.gov/atr/corporate-leniency-policy> [<https://perma.cc/SC4F-J8QG>] (establishing the original model for self-reporting incentives).

agents from abroad. An effective governance architecture must, therefore, be international in scope, built on principles of coordination and mutual recognition.

The framework’s emphasis on auditable, ex ante design obligations provides a natural bridge to parallel regulatory efforts in other key jurisdictions, most notably the European Union. The Duty to Disrupt complements the EU AI Act’s risk management and conformity-assessment architecture, the Digital Markets Act’s affirmative duties for gatekeeper platforms, and the UK Competition and Markets Authority’s Digital Markets Unit.¹⁴³

To avoid duplicative burdens and a race to the bottom, U.S. agencies should pursue a policy of *mutual recognition* for certifications. Under this “passporting” system, a certification issued by an accredited lab in a trusted jurisdiction would be presumptively valid in the United States, provided the underlying technical standards are comparable. This requires harmonizing the protocols for antitrust red-teaming, a task that can be facilitated by international bodies and competition networks.¹⁴⁴ Where enforcement has cross-border effects, U.S. jurisdiction will be guided by established doctrines like the Foreign Trade Antitrust Improvements Act (FTAIA), which governs the Sherman Act’s reach to foreign conduct with a direct, substantial, and reasonably foreseeable effect on U.S. commerce that gives rise to the plaintiff’s claim.¹⁴⁵

The framework’s preference for performance-based remedies respects principles of international comity. By tailoring orders to observable behavioral constraints—such as mandating minimum exploration rates or throttling update cadences—rather than compelling the public

¹⁴³ See Artificial Intelligence Act, Regulation (EU) 2024/1689, O.J. L, 2024/1689 (July 12, 2024), arts. 9, 17 (establishing risk management and quality management systems); Digital Markets Act, Regulation (EU) 2022/1925, 2022 O.J. (L 265) 1, art. 6 (listing affirmative obligations for gatekeepers); Digital Markets, Competition and Consumers Act 2024, c. 13 (U.K.) (establishing the Strategic Market Status regime). The framework’s technical baselines can be mapped to emerging international standards. See, e.g., NAT’L INST. OF STANDARDS & TECH., NIST AI 100-1, ARTIFICIAL INTELLIGENCE RISK MANAGEMENT FRAMEWORK (AI RMF 1.0) (2023), <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf> [<https://perma.cc/QX98-LABQ>].

¹⁴⁴ See generally INT’L COMPETITION NETWORK, *Framework on Competition Agency Procedures* (2019), https://www.internationalcompetitionnetwork.org/wp-content/uploads/2019/04/ICN_CAP.pdf [<https://perma.cc/R4N4-GBB3>] (promoting procedural fairness and cooperation); OECD/ICN, *Report on International Co-operation in Competition Enforcement* (2021), https://www.oecd.org/content/dam/oecd/en/publications/reports/2021/01/oecd-icn-report-on-international-co-operation-in-competition-enforcement_f4b1638e/86f9eb12-en.pdf [<https://perma.cc/K525-JCTE>].

¹⁴⁵ See 15 U.S.C. § 6a (2022) (requiring that the effect “gives rise to a claim”); *F. Hoffmann-La Roche Ltd. v. Empagran S.A.*, 542 U.S. 155, 162–63 (2004).

disclosure of proprietary source code, regulators can remedy domestic harms while minimizing conflicts with the sovereign interests of other nations. This approach fosters a collaborative, rather than confrontational, international environment for governing the competitive effects of agentic AI.¹⁴⁶

2. *Federalism and Preemption: The Role of State Enforcement*

Implementation of the framework will be polycentric, with state enforcers playing an indispensable role in developing and applying design-based competition norms. State attorneys general can bring enforcement actions in their sovereign capacity under state antitrust and consumer protection statutes—often called “little FTC Acts” or Unfair and Deceptive Acts and Practices (UDAP) laws—and, where authorized, as *parens patriae* on behalf of residents to police the deployment of uncertified, convergence-prone agents.¹⁴⁷ For private plaintiffs, state antitrust regimes may offer a more favorable avenue for relief, as many states are *Illinois Brick* repealers, permitting indirect purchasers to sue for damages and thereby softening the direct-purchaser choke point that can complicate federal litigation in agent-intermediated markets.¹⁴⁸

This state-level authority, however, is not plenary and is constrained by federal preemption. Sector-specific statutes can circumscribe state remedies, most notably the Airline Deregulation Act, which broadly preempts state laws “related to a price, route, or service of an air carrier.”¹⁴⁹ This provision would likely bar a state UDAP claim brought directly against an airline for its pricing

¹⁴⁶ See U.S. DEP’T OF JUSTICE & FED. TRADE COMM’N, *Antitrust Guidelines for International Enforcement and Cooperation* § 4 (2017), <https://www.justice.gov/atr/internationalguidelines/download> [<https://perma.cc/B8H5-7X7V>] (discussing comity and remedies).

¹⁴⁷ See 15 U.S.C. § 15c (2022) (federal *parens patriae* authority); see generally CAROLYN L. CARTER ET AL., NAT’L CONSUMER LAW CTR., *Unfair and Deceptive Acts and Practices* (10th ed. 2021) (analyzing state UDAP statutes).

¹⁴⁸ While *Ill. Brick Co.*, 431 U.S. 720, bars indirect purchasers from recovering damages under federal antitrust law, numerous states have passed statutes or had their courts interpret state law to permit such suits. See *ARC Am. Corp.*, 490 U.S. at 101–02 (holding that federal antitrust law does not preempt state indirect-purchaser statutes); *Antitrust Modernization Comm’n, Report and Recommendations* 7 (2007) (noting that indirect purchasers can sue under state law in thirty-six states and the District of Columbia).

¹⁴⁹ 49 U.S.C. § 41713(b)(1) (2022); see *Morales*, 504 U.S. at 383–84 (interpreting “related to” broadly to preempt state consumer protection guidelines); *Am. Airlines, Inc. v. Wolens*, 513 U.S. 219, 228–29 (1995) (holding that the ADA preempts state consumer-fraud claims but not breach-of-contract claims).

conduct, even as federal agencies like the DOJ and DOT retain their enforcement authority.¹⁵⁰

The framework must therefore channel state action toward its most effective and legally durable applications. States can reduce preemption risk by focusing enforcement on the conduct of non-carrier actors, such as the developers of pricing algorithms or the intermediary platforms whose protocols facilitate convergence, provided the claims target general commercial deception or design negligence rather than functioning as de facto airline rate regulation. Furthermore, a robust federal–state partnership is essential. Federal agencies can serve as a technical clearinghouse, sharing sandbox results and audit taxonomies with state counterparts. Coordinated investigations, concurrent consent decrees, and joint market studies can reduce fragmentation and allow states to act as laboratories for innovation in enforcement while respecting the boundaries of federal law.

3. *Sectoral Adaptations and Transitional Justice*

A one-size-fits-all implementation of the Duty to Disrupt would be both brittle and ineffective. The framework must be adaptable, tailored to the unique legal doctrines that govern specific sectors and pragmatic about the challenges of transitioning from existing market structures.

Certain industries are subject to legal immunities that could shield some conduct from the duty’s full reach. The *state-action doctrine*, for instance, may immunize anticompetitive protocols that are undertaken pursuant to a clearly articulated state policy and—where required—actively supervised by the state, while the *Noerr-Pennington* doctrine could protect firms that jointly petition a government body to adopt a particular market-structuring rule (subject to the “sham” exception).¹⁵¹ In regulated network industries like energy or telecommunications, the *filed-rate doctrine* might constrain private damages claims based on rates filed with a regulator, though it would not necessarily bar forward-looking injunctive relief.¹⁵² The framework must therefore

¹⁵⁰ See 49 U.S.C. § 41712 (2022) (authorizing the DOT to investigate unfair or deceptive practices and unfair methods of competition).

¹⁵¹ See *Parker v. Brown*, 317 U.S. 341 (1943) (establishing state-action immunity); *Cal. Retail Liquor Dealers Ass’n v. Midcal Aluminum, Inc.*, 445 U.S. 97, 105 (1980) (requiring active state supervision for private parties); *E. R.R. Presidents Conf. v. Noerr Motor Freight, Inc.*, 365 U.S. 127 (1961); *United Mine Workers of Am. v. Pennington*, 381 U.S. 657 (1965).

¹⁵² See *Keogh v. Chi. & Nw. Ry. Co.*, 260 U.S. 156 (1922) (barring private antitrust damages for rates filed with a

empower sectoral agencies to tailor sandbox protocols and remedies to these legal constraints, ensuring that the duty is applied robustly without overreaching into areas of established immunity.

Just as the framework must adapt to legal context, it must also respect temporal boundaries by providing a practical path for market transition. Many industries have already deployed a generation of algorithmic tools that were not designed with procompetitive duties in mind. Mandating immediate, universal compliance would be infeasible. Instead, the framework should adopt a policy of *phased implementation*, providing a clear and predictable schedule for adoption. New market-facing agents would be required to comply with the duty upon deployment. For existing *legacy systems*, regulators should establish a defined grace period—for example, eighteen to twenty-four months—during which deployers must update their agents and secure certification. During this period, firms making good-faith efforts to comply would be eligible for temporary enforcement discretion, incentivizing proactive transition rather than defensive resistance. This approach recognizes the realities of embedded systems and long-term contracts while ensuring a clear and inevitable path toward a more competitive agentic ecosystem.

D. Broader Economic and Innovation Effects

1. Economic Trade-Offs and Welfare Implications

A comprehensive assessment of the framework requires weighing its potential economic benefits against its inevitable costs. The primary economic justification for the Duty to Disrupt is the prevention of significant allocative inefficiencies and reductions in consumer welfare. Agentic convergence, if left unchecked, threatens to institutionalize supracompetitive pricing across key sectors of the economy, reducing output and dulling innovation incentives while transferring wealth from consumers to producers without any corresponding efficiency gain.¹⁵³ The target is

regulator); *Georgia v. Pa. R.R. Co.*, 324 U.S. 439, 453–62 (1945) (holding that the doctrine does not bar injunctive relief against conspiracies to fix filed rates).

¹⁵³Economic modeling and empirical analysis have demonstrated that pricing algorithms can systematically learn to charge supracompetitive prices, leading to substantial welfare losses. See Calvano et al., *supra* note 10, at 3270–75 (theoretical modeling); see also STEPHANIE ASSAD ET AL., *Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market*, 132 J. POL. ECON. 723 (2024) (finding empirically that the adoption of

not dynamic pricing as such, but durable supracompetitive pricing sustained by strategic interdependence among agents—an outcome distinct from efficiency-enhancing demand-responsive pricing. By mandating rivalry-by-design, the framework acts as a structural safeguard, preserving the dynamic, long-term competition that drives innovation and allocative efficiency. The societal benefit is not merely lower prices in the short term, but the preservation of a market structure where firms are forced to compete on quality, service, and innovation.

These benefits, however, are not without costs. The framework imposes direct compliance burdens on firms, including the costs of algorithmic auditing, certification, and potentially redesigning agents to incorporate procompetitive features. Mandated strategic randomness, if poorly calibrated, could also introduce short-term price volatility or degrade the allocative efficiency that dynamic pricing is intended to achieve.¹⁵⁴ Furthermore, there is a risk that these compliance costs could disproportionately burden smaller firms and new entrants, potentially leading to increased market concentration—an outcome directly at odds with the framework’s goals.

The framework is designed with several features to mitigate these costs and ensure the net welfare effects are positive. First, the duty is tiered: it applies only after a market crosses a material operational threshold, and the most intensive oversight tools, like regulatory sandboxes, are reserved for high-risk sectors. Second, the availability of a safe harbor for certified agents reduces legal uncertainty and lowers the long-term cost of capital for innovators. Finally, as discussed below, the potential for modular, open-source compliance tools can democratize access to certified procompetitive designs, lowering barriers to entry. Crucially, the sandbox-and-certification machinery also permits calibration of any required randomness through measurable performance metrics rather than blunt, one-size-fits-all mandates. While a full quantitative analysis is beyond the scope of this Article, the framework is premised on the principle that the systemic, long-term costs of market-wide agentic convergence are far greater than the targeted, manageable costs of

algorithmic pricing increased margins).

¹⁵⁴See, e.g., R. PRESTON MCAFEE & VERA TE VELDE, *Dynamic Pricing in the Airline Industry*, in 1 HANDBOOK ON ECONOMICS AND INFORMATION SYSTEMS 527, 529–32 (T.J. Hendershott ed., 2006) (discussing the efficiency benefits of dynamic pricing in clearing markets).

engineering a system to prevent it.

2. *Innovation, Open Source, and Market Entry*

A central challenge for any new regulatory framework is to avoid chilling beneficial innovation or creating barriers to entry that entrench dominant firms. Critics may argue that the Duty to Disrupt, with its requirements for auditing and certification, could create a “compliance moat” driven by high fixed compliance costs and scarce technical expertise, where incumbents can absorb the burdens more easily than startups, thereby stifling competition.¹⁵⁵ The framework, however, is designed not to chill innovation, but to *channel it* in a procompetitive direction. By creating liability for negligently designed, convergence-prone agents, it shifts the incentives for AI development away from a narrow focus on short-term revenue maximization and toward a broader mandate to create agents that are both effective for their users and compatible with a competitive market ecosystem.

To ensure that compliance costs do not become a barrier to entry, the framework is designed to foster an ecosystem of *modular and open-source compliance*. Regulators should explicitly recognize and streamline the certification process for firms that integrate a pre-certified, third-party “coordination-breaker” module into their agents (for example, a certified policy wrapper or deployment-layer control that enforces bounded exploration, constrains short-run retaliation dynamics, and produces standardized audit logs).¹⁵⁶ This would allow a market to develop for specialized, off-the-shelf software components that fulfill the duty’s core technical requirements. An open-source, rigorously vetted collusion-breaker module, perhaps developed through a public-private partnership or an academic consortium, could be made available to all market participants at little to no cost, substantially lowering the marginal cost of compliance for new entrants.

¹⁵⁵ See, e.g., JAMES D. CAMPBELL, AVI GOLDFARB & CATHERINE TUCKER, *Privacy Regulation and Market Structure*, 24 J. ECON. & MGMT. STRATEGY 47 (2015) (finding that privacy regulation can disproportionately burden smaller firms); GARRETT A. JOHNSON, SCOTT K. SHRIVER & SAMUEL G. GOLDBERG, *Privacy and Market Concentration: Intended and Unintended Consequences of the GDPR*, 69 MGMT. SCI. 5695 (2023).

¹⁵⁶ Cf. NAT’L INST. OF STANDARDS & TECH., *supra* note 143, at 6 (describing “AI RMF Profiles” as mechanisms to tailor risk management to specific sectors); FRANK NAGLE, *Open Source Software and Firm Productivity*, 65 MGMT. SCI. 1191 (2019) (documenting the productivity benefits of open-source software).

Component certification, however, should not substitute for system certification. Its function should be to create an accelerated path to approval, conditioned on passing standardized integration tests in the relevant sandbox, since convergence risk can emerge from the interaction of the agent, its objectives, and the market protocol. This approach mitigates the risk of a compliance moat in several ways. First, it lowers the cost and technical expertise required for a new entrant to comply with the duty. Second, it creates a competitive market for compliance tools, driving down prices and improving quality. Third, it ensures that the standards for certification are transparent and accessible, reducing the risk of incumbents capturing the process. To avoid monoculture and capture, regulators should accredit multiple interoperable modules under open standards and require periodic retesting across versions. By making rivalry-by-design a modular and accessible feature, the framework can democratize compliance and ensure that the agentic economy remains a field for dynamic competition, not just for the largest and most established players. This creates a stable and predictable environment that can ultimately encourage, rather than deter, long-term investment in welfare-enhancing AI development.

E. An Implementation Path

The transition to this new governance architecture should be pragmatic and staged, proceeding under existing legal authority while building the foundation for future legislative codification. A pragmatic rollout could proceed in four phases:

- **Phase 0 (Market Designation):** The process begins with a published agency determination—issued after notice and an opportunity for comment—that a given sector constitutes an “agentic market” warranting enhanced oversight. This designation would be triggered when the market crosses a clear, predefined operational threshold based on quantitative metrics, such as the percentage of transactions materially influenced by both buyer and seller agents and the existence of a protocolized environment that enables cross-party learning.¹⁵⁷

¹⁵⁷Cf. Digital Markets Act, Regulation (EU) 2022/1925, 2022 O.J. (L 265) 1, art. 3 (establishing the “gatekeeper”

- **Phase 1 (0–12 Months): Agency-Led Record Building and Guidance:** The FTC should issue targeted orders under Section 6(b) of the FTC Act, while the DOJ and sectoral regulators utilize their respective investigative authorities (such as Civil Investigative Demands) to collect design records and interaction telemetry from key market participants.¹⁵⁸ Grounded in this empirical record, the agencies can then issue joint guidance clarifying that the negligent deployment of uncertified, convergence-prone agents may constitute an unfair method of competition. This stage would also involve establishing an interagency task force to harmonize technical protocols.
- **Phase 2 (12–24 Months): Piloting the Assurance Market:** Regulators should launch pilot programs and regulatory sandboxes in high-risk markets, such as airlines and hotels, to validate testing protocols and build institutional expertise. This phase would focus on developing concrete success metrics—including price dispersion, retaliation half-life, and defection-reward amplification—and accrediting the first third-party auditors.¹⁵⁹ The publication of the accreditation program and the availability of accredited audit capacity would trigger the start of the 18-24 month grace period for legacy systems. In parallel, U.S. agencies should begin negotiations with international counterparts to establish a path toward mutual recognition of certifications as described above.
- **Phase 3 (24–36 Months): Legislative Codification:** With a foundation of agency practice and market adoption, Congress can then codify the framework’s core pillars. This legislation should include: (i) an explicit, ecosystem-wide Duty to Disrupt for all market-facing agents; (ii) the rebuttable presumption of liability triggered by supracompetitive effects; (iii) clear statutory authority for structural separation remedies; (iv) non-waivable rights for private

designation); Digital Markets, Competition and Consumers Act 2024, c. 13 (U.K.) (establishing the “Strategic Market Status” designation).

¹⁵⁸ See 15 U.S.C. § 46(b) (2022) (FTC market study authority); 15 U.S.C. § 1312 (2022) (DOJ Civil Investigative Demand authority).

¹⁵⁹ See Assad et al., *supra* note 153, at 741 (identifying empirical “markers” of algorithmic-pricing adoption, including changes in the frequency of price changes and the speed/frequency of responses to rivals’ price changes and demand/cost shocks); see also *supra* Part II.B (discussing the mechanics of punishment and reward in agentic equilibria).

enforcement, including an express restriction on predispute arbitration and class-action waivers for these claims; and (v) safe harbors for researchers. A sunset provision could ensure that the framework is periodically reviewed and updated.

The central task of 20th-century antitrust was to detect collusion in the evidence of the past; the central task for the 21st century is to engineer competition into the architecture of the future. This implementation pathway provides a blueprint for that task.

VI. OBJECTIONS AND REBUTTALS

A proposal to shift antitrust from policing conspiracies to engineering competition invites skepticism. It requires courts and agencies to engage with complex technical architectures and imposes affirmative duties in a domain traditionally governed by negative prohibitions. This Part confronts these critiques in their strongest form, addressing objections concerning the duty’s vagueness, the capacity of institutions, the use of presumptions, and the basis of legal authority, as well as a wider range of practical, doctrinal, and economic challenges. Each reply is grounded in the framework’s core design commitments—structural integrity, proactive oversight, and design-record-based proof—demonstrating that while the proposal requires careful calibration, none of these objections is fatal.

The analysis proceeds in four thematic areas. First, it addresses *definitional and operational challenges*, clarifying the duty’s precise scope and rebutting concerns about unintended consequences like buyer-side monopsony. Second, it confronts the practical realities of *implementation and enforcement*, detailing how the framework would manage institutional capacity, prevent gaming and evasion, safeguard privacy, and handle the transition of legacy systems. Third, it examines the core *legal and constitutional constraints*, defending the framework’s statutory authority, its burden-shifting mechanism, and its viability in the face of procedural hurdles in private litigation. Finally, it weighs the broader *economic and innovation concerns*, analyzing the proposal’s effect on small firms, its potential for perverse economic outcomes, and its implications for U.S.

competitiveness in a global market.

A. Definitional and Operational Challenges

1. *Vagueness and the Duty's Scope*

Objection.—Critics will argue that a legally enforceable “Duty to Disrupt” is too indeterminate to provide fair notice. The term itself is inflammatory, suggesting a mandate for market destabilization. Developers cannot know *ex ante* what quantum of “strategic randomness” is sufficient or what “long-horizon” objective is adequate. This ambiguity is compounded by ill-defined triggers like a “non-trivial share” of agentic transactions, creating a vague standard that will chill beneficial innovation and impose disproportionate compliance costs on smaller firms and new entrants.

Reply.—This objection mischaracterizes a process-based, negligence-style standard as a vague, prescriptive one. The duty’s requirements are cabined by familiar and administrable legal principles, and courts have long recognized that economic regulation is subject to a less strict vagueness test than criminal statutes because regulated enterprises have the ability to clarify the meaning of regulations by inquiry.¹⁶⁰

First, compliance is not measured against an abstract standard of “disruption” but is *modeled* on the well-established “*reasonable alternative design*” inquiry from product liability law.¹⁶¹ A developer satisfies the duty by demonstrating, through a certified design record, that it implemented feasible, procompetitive design features (like long-horizon objectives and strategic randomness) to mitigate the foreseeable risk of agentic convergence. The duty requires disrupting coordinated equilibria through controlled exploration—a concept consistent with emerging technical standards for AI risk management that emphasize continuous testing, evaluation, verification, and validation (TEVV) across the model lifecycle.¹⁶²

¹⁶⁰ See *Vill. of Hoffman Ests. v. Flipside, Hoffman Ests., Inc.*, 455 U.S. 489, 498–99 (1982) (explaining that “economic regulation is subject to a less strict vagueness test” because its subject matter is often narrower and businesses can plan and consult relevant authorities).

¹⁶¹ See RESTATEMENT (THIRD) OF TORTS: PRODS. LIAB. § 2(b) (Am. L. Inst. 1998). This framework adapts the *evidentiary structure* of the design-defect inquiry to the antitrust context.

¹⁶² See NAT’L INST. OF STANDARDS & TECH., *supra* note 143, at 10, 19 (detailing TEVV tasks for AI actors and emphasizing

Second, the duty’s scope is structurally defined; it attaches only after a formal *Agentic-Market Designation* by a regulator. This designation would be a public finding based on prespecified, quantitative triggers established via notice and an opportunity for comment, such as (i) a significant percentage of market transactions being materially influenced by both buyer and seller agents, (ii) the widespread use of a standardized API or protocol for interaction, and (iii) a defined price-update cadence. Below this operational threshold, the duty does not apply.

Third, for firms operating within a designated market, the framework provides clear safe harbors and scaled obligations to avoid chilling innovation. It recognizes a *de minimis safe harbor* for experimental or low-impact agents that fall below specified significance thresholds.¹⁶³ For smaller firms and open-source developers, it allows for *modular compliance*, such as by integrating a pre-certified, open-source “collusion-breaker” module to satisfy the duty’s core requirements at minimal cost. This approach provides the fair notice and predictability necessary for a dynamic and competitive market.

2. *The Buyer-Side Monopsony Risk*

Objection.—The framework’s solution could become a new version of the problem. If the largest buyer agents (e.g., those operated by major retail platforms) learn to act in parallel to suppress prices below competitive levels, the “Duty to Disrupt”—particularly the “Strategic Shield” mechanism that requires buyer agents to blunt punishments against defecting sellers—could be used to harm sellers, especially smaller merchants, by forcing prices down to levels that reduce output, quality, or innovation. These buyer agents could create a powerful oligopsony or monopsony, thereby distorting the market in a new anticompetitive direction.¹⁶⁴

continuous measurement of risks).

¹⁶³ See *supra* Part V.E (describing the “Phase 0” market designation based on operational thresholds); see also *supra* Part V.B.3 (detailing safe harbors for certified agents).

¹⁶⁴ This risk is not purely theoretical; it mirrors the structural power platforms already exert over third-party merchants through ranking algorithms and “Buy Box” mechanics, whereby the platforms’ market share of consumers enables them to squeeze sellers’ margins. See Khan, *supra* note 90, at 768–83 (analyzing how Amazon’s dual role creates conflicts of interest that harm independent sellers); see also Complaint at 8–10, *FTC v. Amazon.com, Inc.*, *supra* note 115 (alleging Amazon uses its control over product discovery to punish sellers who offer lower prices on rival platforms). The deployment of autonomous buying agents like Amazon’s “Rufus” or Walmart’s “Sparky” could

Reply.—This concern correctly identifies a potential failure mode but overlooks the framework’s symmetric design and unilateral nature. The duty is to engineer *rivalry*, not to create buyer cartels, and several safeguards prevent this outcome.

First, the duty is *individual and unilateral*. It imposes an independent obligation on each agent’s developer to incorporate procompetitive design features. It is a mandate to act as a “coordination-breaker,” not a license to coordinate. Just as the framework disrupts seller-side tacit collusion by reducing the predictability of pricing, its requirements for *strategic randomness* and *long-horizon objectives* impede buyer-side tacit coordination by making parallel purchasing strategies difficult to sustain. Furthermore, the framework reinforces the prohibition on sharing non-public, competitively sensitive data between buyer agents, treating such exchanges as presumptive evidence of conspiracy.

Second, the framework’s core enforcement mechanism is *symmetrically applied*. The rebuttable presumption of liability is triggered by evidence of persistent anticompetitive outcomes, regardless of which side of the market they favor. The Supreme Court has long recognized that the Sherman Act protects sellers from buyer-side conspiracies just as it protects consumers from seller-side price-fixing.¹⁶⁵ If econometric evidence demonstrated persistent *monopsonistic* effects—such as prices pushed sustainably below competitive levels accompanied by reduced output or supplier exit—the burden of production would shift to the developers of the dominant buyer agents to prove their designs were not the cause.¹⁶⁶

Finally, the goal of the “Strategic Shield” is to restore a competitive equilibrium, not to create a subcompetitive one. It operates by making punishment for a seller’s procompetitive defection less credible. An agent that systematically punished *all* sellers to suppress prices would violate its

automate and amplify these exclusionary dynamics.

¹⁶⁵ See *Mandeville Island Farms, Inc. v. Am. Crystal Sugar Co.*, 334 U.S. 219, 235–36 (1948) (holding that the Sherman Act condemns price-fixing “even though the price-fixing was by purchasers” and that the statute protects “sellers, not customers or consumers” alone); see also *Weyerhaeuser Co. v. Ross-Simmons Hardwood Lumber Co.*, 549 U.S. 312, 320–22 (2007) (recognizing monopsony as the “mirror image” of monopoly and noting that it leads to lower output and misallocation of resources).

¹⁶⁶ See HERBERT HOVENKAMP, *Federal Antitrust Policy: The Law of Competition and Its Practice* § 1.2b (6th ed. 2020) (explaining that monopsony reduces output and results in deadweight loss similar to monopoly, ultimately harming consumer welfare through reduced supply or quality).

duty to its own user by degrading the long-term quality and availability of the supply base. The audits and sandbox tests would therefore be explicitly designed to detect these monopsonistic tendencies, ensuring the duty serves as a scalpel for restoring competition, not a hammer for creating new market distortions.

B. Implementation and Enforcement Challenges

1. Institutional Incapacity and Administrability

Objection.—Even if the duty is well-defined, antitrust agencies and generalist courts lack the institutional capacity and technical expertise to evaluate reinforcement-learning policies, interpret complex simulations, and serve as effective algorithmic auditors. This is compounded by the “technology velocity problem”: by the time regulators understand today’s agentic architectures, the market will have moved on. This will lead to arbitrary enforcement, with agencies and courts making Type I errors that punish procompetitive tools or Type II errors that bless coordination-prone designs.

Reply.—This concern is manageable because the framework is designed to leverage existing procedural tools and focuses on empirical performance, not direct code review.

For *agencies*, the task is not to reverse-engineer every model but to build a persuasive administrative record to support their interpretive guidance. They can achieve this by using established powers, such as the FTC’s Section 6(b) authority, to compel the production of design records and interaction telemetry.¹⁶⁷ This data then forms the basis for performance-based evaluations in regulatory sandboxes and market studies, consistent with the monitoring approaches discussed by international bodies like the OECD.¹⁶⁸ To address the velocity problem, the duty is tied to enduring principles—long-term horizons, strategic randomness—not specific architectures. Performance testing can evolve with technology, and agencies can build on existing technical expertise, as demonstrated by the FTC’s establishment of an Office of Technology and the DOJ’s increasing

¹⁶⁷ See 15 U.S.C. § 46(b) (2022).

¹⁶⁸ See OECD, *supra* note 38, at 35–40 (discussing market studies and auditing mechanisms).

engagement in algorithmic cases.¹⁶⁹

For *courts*, adjudication fits within the existing framework for managing complex technical evidence. The admissibility of simulation results would be policed under Federal Rule of Evidence 702 and the reliability standards of *Daubert*.¹⁷⁰ Judges need not become software engineers; they can appoint neutral technical experts under Rule 706 or special masters under Rule 53 to evaluate red-team protocols, review design records under protective orders, and oversee compliance with technical remedies.¹⁷¹ Remedies themselves are analogous to familiar product-safety modifications—such as mandating an increase in an agent’s exploration parameter—and are verifiable through the reproducible testing and telemetry data that developers already utilize.

2. *Gaming, Evasion, and Audit Robustness*

Objection.—The framework’s reliance on certification will devolve into “certification theater” or “teaching to the test.” Firms will optimize their agents to pass a predictable set of red-team scenarios, sprinkling in superficial randomness or disabling collusive strategies only in the test environment. This tracks Goodhart’s Law, often summarized via Strathern’s formulation: “When a measure becomes a target, it ceases to be a good measure.”¹⁷² Sophisticated reinforcement-learning agents could in principle learn to actively evade their design constraints in production, evolving new, undetectable coordinating strategies and creating a regulatory arms race that agencies are destined to lose.¹⁷³

¹⁶⁹ See FED. TRADE COMM’N, *FTC Launches New Office of Technology to Bolster the Agency’s Work* (Feb. 17, 2023), <https://www.ftc.gov/news-events/news/press-releases/2023/02/ftc-launches-new-office-technology-bolster-agencys-work> [<https://perma.cc/K8LR-FB7J>]; see also Memorandum of Law, *supra* note 26, at 16–18 (arguing that existing antitrust law reaches algorithm-facilitated pricing coordination in the rental market).

¹⁷⁰ See Fed. R. Evid. 702; *Daubert*, 509 U.S. at 589–95; *Kumho Tire*, 526 U.S. at 147–49.

¹⁷¹ See Fed. R. Evid. 706; Fed. R. Civ. P. 26(c)(1)(G), 53.

¹⁷² See MARILYN STRATHERN, *‘Improving Ratings’: Audit in the British University System*, 5 EUR. REV. 305, 308 (1997); see also CHARLES GOODHART, *Problems of Monetary Management: The U.K. Experience*, in INFLATION, DEPRESSION, AND ECONOMIC POLICY IN THE WEST 111, 116 (Anthony S. Courakis ed., 1981).

¹⁷³ See DARIO AMODEI ET AL., *Concrete Problems in AI Safety* (arXiv Working Paper No. 1606.06565, 2016), at 7–10 (discussing “reward hacking”/“wireheading”), <https://arxiv.org/pdf/1606.06565.pdf> [<https://perma.cc/3SQW-XMKW>]; see also VICTORIA KRAKOVNA ET AL., *Specification Gaming: The Flip Side of AI Ingenuity*, DEEPMIND (Apr. 21, 2020), <https://deepmind.google/discover/blog/specification-gaming-the-flip-side-of-ai-ingenuity/> [<https://perma.cc/N9HJ-N6UW>].

Reply.—This objection correctly identifies that static, predictable audits are insufficient. The framework therefore anticipates and defends against gaming and evasion with a multi-layered system of dynamic, ongoing oversight.

First, the certification process itself is designed to be robust against simple gaming. Test suites are not static but include *rotating, partially undisclosed adversarial scenarios* and *out-of-distribution tasks* that are unknown to developers beforehand, preventing firms from simply overfitting their models to a known test. Independent, accredited third-party auditors, rather than the developers themselves, conduct these tests to ensure their integrity.

Second, certification is not a one-time, permanent shield. Safe harbors are *time-bounded* and *automatically lapse upon any material model update*. To ensure the deployed agent matches the certified one, the framework requires *cryptographic model hashing* and *deployment attestations*. As agents evolve, they are subject to re-evaluation, preventing firms from certifying a compliant model and then deploying a non-compliant version.

Third, the framework moves beyond pre-deployment testing to include *post-deployment monitoring*. Regulators can require firms to submit privacy-preserving telemetry to detect emergent convergence or suspicious drift in an agent’s behavior after it has been deployed. This can be supplemented with *surprise audits* and the use of regulator-operated test agents designed to probe for coordinating responses in the live market. This approach aligns with modern AI safety principles that emphasize continuous monitoring and adaptation.¹⁷⁴

Finally, these procedural safeguards are backed by legal deterrents. A demonstrated failure to preserve required telemetry, or the discovery of a “defeat device” that alters behavior outside of the test environment, would not only revoke the safe harbor but could also support—upon a showing of intent to deprive—an *adverse inference* in litigation that the agent was negligently designed.¹⁷⁵

¹⁷⁴See NAT’L INST. OF STANDARDS & TECH., *supra* note 143, at 32 (describing post-deployment monitoring plans); see also Artificial Intelligence Act, Regulation (EU) 2024/1689, O.J. L, 2024/1689 (July 12, 2024), art. 72 (requiring providers of high-risk AI systems to implement a post-market monitoring system).

¹⁷⁵See Fed. R. Civ. P. 37(e)(2) (authorizing adverse-inference instructions only upon a finding that the party “acted with the intent to deprive another party of the information’s use”); cf. *In re Volkswagen “Clean Diesel” Mktg., Sales*

3. *Privacy and Security Risks from Telemetry*

Objection.—The framework’s reliance on a “design record” and post-deployment telemetry creates significant privacy and security risks. Requiring firms to log and retain rich data about market interactions could expose sensitive consumer information—such as purchasing patterns and willingness-to-pay—to misuse or data breaches. Even “de-identified” behavioral logs can often be re-identified through linkage with external datasets. The framework’s requirements create new honeypots of sensitive data that could be exploited by malicious actors or demanded by law enforcement for purposes unrelated to competition oversight.

Reply.—This concern is valid, but it mischaracterizes the framework’s purpose. The goal is to increase transparency into algorithmic *behavior*, not to create new streams of personal data for surveillance. The evidentiary requirements are structured to be *privacy-preserving by default*, adhering to principles of data minimization and strict purpose restrictions that limit the use of this data for non-regulatory purposes.

First, the framework minimizes security risks through *distributed retention*. The “design record” and raw telemetry are retained by the *developer*, not the regulator, and are only transmitted to agencies in response to specific audit demands or investigations. This prevents the creation of a centralized “honeypot” of sensitive information.

Second, the data itself is minimized. The design record consists of model-level, not user-level, information: reward functions, parameter settings, and architectural specifications. The interaction telemetry required for audits can be further *data-minimized* to include only the information necessary to reconstruct market dynamics—such as transaction IDs, prices, and timestamps—while explicitly excluding personally identifiable information (PII).

Third, where more granular data is required for sandboxes or post-deployment monitoring, regulators can mandate the use of *privacy-enhancing technologies (PETs)*. This includes requiring aggregation and differential privacy to provide quantifiable privacy guarantees, ensuring that

Pracs., & Prods. Liab. Litig., 959 F.3d 1201, 1207–10 (9th Cir. 2020) (detailing how software “defeat devices” were used to mask true emissions levels during regulatory testing, serving as a functional analog to the suppression of evidence).

market-level convergence can be analyzed without resolving the actions of specific individuals.¹⁷⁶

Fourth, established legal procedures govern access. In private litigation, discovery proceeds under *Federal Rule of Civil Procedure 26(c)* protective orders.¹⁷⁷ For agency oversight, confidential submissions are protected from public disclosure under *FOIA Exemptions 4 and 6*,¹⁷⁸ while strict purpose limitations would constrain secondary use, including demands for unrelated law enforcement surveillance.

4. *Retroactivity and Transitional Justice for Legacy Systems*

Objection.—Imposing a new design-based duty on the vast number of algorithmic systems already deployed would be unfairly retroactive. Firms made significant investments in these legacy systems based on the legal rules in place at the time. Mandating immediate and costly re-engineering would upset settled reliance interests, violate principles of fair notice, and could be challenged as a violation of due process.¹⁷⁹

Reply.—The framework is designed to be *prospective, not retroactive*, and it incorporates a pragmatic transition path to respect legitimate reliance interests while ensuring a timely move to a more competitive market.¹⁸⁰

First, the duty itself is not retroactive. Liability would not attach to conduct that occurred

¹⁷⁶ See NAT'L INST. OF STANDARDS & TECH., NIST SPECIAL PUBLICATION 800-226: GUIDELINES FOR EVALUATING DIFFERENTIAL PRIVACY GUARANTEES (Mar. 2025), <https://doi.org/10.6028/NIST.SP.800-226>; see also NAT'L INST. OF STANDARDS & TECH., NIST PRIVACY FRAMEWORK: A TOOL FOR IMPROVING PRIVACY THROUGH ENTERPRISE RISK MANAGEMENT (Version 1.0, Jan. 2020), <https://doi.org/10.6028/NIST.CSWP.01162020> (identifying “disassociated processing” and de-identification controls).

¹⁷⁷ See Fed. R. Civ. P. 26(c)(1)(G).

¹⁷⁸ 5 U.S.C. § 552(b)(4), (b)(6) (2022). The Supreme Court has confirmed a broad understanding of “confidential” commercial information under Exemption 4. See *Food Mktg. Inst.*, 139 S. Ct. at 2363–66 (holding that commercial information is “confidential” at least where it is customarily and actually treated as private by its owner).

¹⁷⁹ See, e.g., *Landgraf v. USI Film Prods.*, 511 U.S. 244, 265–66 (1994) (describing the presumption against retroactivity as grounded in “elementary considerations of fairness”); *Bowen v. Georgetown Univ. Hosp.*, 488 U.S. 204, 208 (1988) (stating that agencies are not presumed to have authority to promulgate retroactive rules absent clear congressional authorization); *FCC v. Fox Television Stations, Inc.*, 567 U.S. 239, 253 (2012) (explaining that due process requires fair notice of what conduct is prohibited or required).

¹⁸⁰ While the law disfavors primary retroactivity (penalizing past conduct), it permits “secondary retroactivity”—regulatory changes that affect the future value of past investments—provided the agency’s action is reasonable. See *Mobile Relay Ass’n v. FCC*, 457 F.3d 1, 11 (D.C. Cir. 2006) (distinguishing between impermissible retroactive rules and permissible rules with secondary retroactive effects).

before a market was formally designated as “agentic” and the duty was established through public agency action. The framework’s core obligation is forward-looking: to design and deploy agents responsibly *after* the new rules of the road are clear.

Second, for *legacy systems* already operating in a newly designated market, the framework provides a period of *transitional justice*. Rather than mandating an immediate “rip and replace,” regulators would establish a clear *sunrise period* (e.g., eighteen to twenty-four months) during which firms must bring their agents into compliance. This provides the time needed to re-engineer, test, and certify existing systems.

Third, compliance during this transition is phased to prevent evasion. The duty attaches upon the *earlier* of two events: (1) the expiration of the sunrise period; or (2) the implementation of a *material model update*. A firm that continues to operate a legacy system without material changes would benefit from *transitional enforcement discretion* regarding the full design duty, conditioned on preserving its design record and not expanding the agent’s deployment. However, once a firm makes a material update—such as deploying a new model version or substantially altering core parameters—the full duty to comply and certify attaches immediately. This approach respects reliance interests by tying the compliance obligation to the ordinary cycle of innovation while ensuring that legacy systems cannot function as indefinite loopholes.

C. Legal and Constitutional Constraints

1. Statutory and Constitutional Authority

Objection.—The framework exceeds the statutory authority granted by Congress and violates constitutional constraints. Section 1 of the Sherman Act targets “agreements,” not unilateral design choices. Section 5 of the FTC Act is not a roving license for product regulation, particularly in light of the Major Questions Doctrine, which counsels skepticism of transformative agency actions not explicitly authorized by Congress. Furthermore, compelled code modifications could constitute a

regulatory taking of intellectual property or violate the First Amendment, as “code is speech.”¹⁸¹

Reply.—The proposal finds plausible footholds within existing statutory and constitutional architecture; it is an application of established authority to a new technological reality, not a radical expansion of it.

First, the framework rests on long-standing *statutory authority*. The FTC’s power to police “unfair methods of competition” under Section 5 has long been understood by the Supreme Court to extend beyond the precise conduct prohibited by the Sherman Act, encompassing incipient threats to competitive markets.¹⁸² For air carriers, the Department of Transportation holds parallel authority to police “unfair” practices.¹⁸³ Requiring firms to test for foreseeable anticompetitive harms is a classic application of these mandates. It is not a “transformative” assertion of new power that would trigger the *Major Questions Doctrine*, but rather the application of Congress’s expressly flexible mandate to police evolving market mechanisms through case-by-case adjudication.¹⁸⁴ While *AMG Capital* limits the FTC’s ability to seek retrospective monetary relief under Section 13(b), it does not alter its authority to seek forward-looking injunctions under that statute or to impose conduct remedies through administrative orders under Section 5.¹⁸⁵

Second, *constitutional concerns* are misplaced. The *First Amendment* does not bar content-neutral regulation of economic *conduct*. A requirement to include procompetitive guardrails regulates the *functional performance* of an agent in a commercial marketplace, not the content of its expressive speech.¹⁸⁶ A *Takings Clause* claim would also fail. The framework does not

¹⁸¹ See, e.g., *Junger v. Daley*, 209 F.3d 481, 485 (6th Cir. 2000) (holding that source code is protected by the First Amendment because it is “an expressive means for the exchange of information”).

¹⁸² See 15 U.S.C. § 45 (2022); *FTC v. Brown Shoe Co.*, 384 U.S. 316, 320–22 (1966) (holding that the Commission has broad power to declare trade practices unfair, particularly those that conflict with the basic policies of the Sherman and Clayton Acts); *FTC v. Sperry & Hutchinson Co.*, 405 U.S. 233, 239–44 (1972); see also FED. TRADE COMM’N, *supra* note 87 (articulating the Commission’s position that Section 5 reaches beyond the Sherman and Clayton Acts to address incipient threats to competition).

¹⁸³ See 49 U.S.C. § 41712 (2022); see also U.S. DEP’T OF TRANSP., *Guidance Regarding Interpretation of Unfair and Deceptive Practices*, 87 Fed. Reg. 52,677 (Aug. 29, 2022).

¹⁸⁴ Cf. *West Virginia*, 142 S. Ct. at 2609–10 (requiring clear congressional authorization for agency actions of vast economic and political significance).

¹⁸⁵ See *AMG Capital Mgmt.*, 141 S. Ct. at 1347–49 (interpreting 15 U.S.C. § 53(b)).

¹⁸⁶ See *Rumsfeld v. Forum for Acad. & Institutional Rights, Inc.*, 547 U.S. 47, 62 (2006) (upholding regulation of conduct that has only incidental effects on speech).

appropriate intellectual property for public use; it regulates its use in commerce to prevent public harm. The use of protective orders and the safeguarding of trade secrets further ensures that property interests are respected.¹⁸⁷ Finally, the *Due Process* vagueness challenge is answered by the framework’s clear operational triggers, its reliance on auditable, performance-based metrics, and the availability of safe harbors, all of which provide regulated parties with fair notice of their obligations.¹⁸⁸

2. *The Rebuttable Presumption and Evidentiary Standards*

Objection.—The proposal to create a rebuttable presumption of design-based causation based on market *effects* flips antitrust law on its head and violates due process. It risks condemning efficient or merely parallel high prices that result from benign market factors, inviting a flood of litigation based on false positives. It clashes with the plausibility pleading standards of *Twombly* and *Iqbal*¹⁸⁹ and could be challenged under *Daubert*¹⁹⁰ for allowing liability based on effects evidence without a reliable showing of causation. It forces firms to prove their innocence, reversing the traditional allocation of proof.

Reply.—The presumption is narrower, more doctrinally grounded, and more procedurally disciplined than the objection suggests. It is a carefully calibrated tool for correcting a profound information asymmetry, not a shortcut to liability. Crucially, it is a presumption of *design-based causation*, not of conspiratorial intent.

First, the presumption is *narrowly triggered*. The plaintiff’s initial burden is not merely to show high prices, but to make a specific empirical showing of *persistent supracompetitive outcomes*—for example, prices exceeding a competitive benchmark by a statistically significant margin for a durable period, after controlling for legitimate drivers like input costs and demand shocks. This rigorous, effects-based showing is the prerequisite for any burden-shifting and is itself subject

¹⁸⁷ Cf. *Ruckelshaus*, 467 U.S. at 1005–08 (holding that compelled disclosure of data does not constitute a taking where the submitter lacks a reasonable investment-backed expectation of confidentiality given the statutory scheme).

¹⁸⁸ See *Hoffman Ests.*, 455 U.S. at 498–99 (noting that economic regulation is subject to a less strict vagueness test).

¹⁸⁹ See *Twombly*, 550 U.S. at 555–57; *Ashcroft v. Iqbal*, 556 U.S. 662, 678 (2009).

¹⁹⁰ See *Daubert*, 509 U.S. at 589–95.

to *Daubert*'s reliability standards. In practice, this means the presumption is triggered only after plaintiffs produce reliable econometric evidence—typically at summary judgment—not by conclusory allegations at the pleading stage.

Second, the burden that shifts is the *burden of production, not the ultimate burden of persuasion*. Once the plaintiff makes their prima facie case, the defendant must simply *produce* the evidence that they uniquely possess: their certified design records and red-team results. The plaintiff retains the ultimate burden of persuading the trier of fact that the defendant's agent was the cause of the anticompetitive harm. This is a familiar and constitutionally permissible form of burden-shifting used in other areas of law where a defendant has exclusive control over the evidence of causation, such as civil rights and merger law.¹⁹¹

Third, the framework is designed with *procedural guardrails* to minimize false positives and manage litigation. The presumption attaches only after a market has been formally designated as "agentic." Discovery can be *phased*, with the court first ruling on the plaintiff's effects-based showing before compelling the production of sensitive design records under a protective order. Most importantly, developers can avail themselves of a *safe harbor*: maintaining an up-to-date, third-party certification serves as a *prima facie* rebuttal, effectively preventing the presumption from attaching in the first place. This creates a powerful incentive for proactive compliance, not defensive litigation.

3. *Private Litigation and Procedural Hurdles*

Objection.—Even if the duty is legally sound, private enforcement will be impossible in practice. First, under *Illinois Brick*, only direct purchasers can sue for federal damages, a fatal flaw in markets where buyer agents intermediate transactions.¹⁹² Second, under *Associated General Contractors*,

¹⁹¹ See *Griggs*, 401 U.S. at 432 (placing burden on employer to justify practice producing disparate impact); *Phila. Nat'l Bank*, 374 U.S. at 363 (establishing structural presumption in merger cases); see also *United States v. Baker Hughes Inc.*, 908 F.2d 981, 991 (D.C. Cir. 1990) (explaining that the presumption shifts the burden of production, but the ultimate burden of persuasion remains with the plaintiff).

¹⁹² See *Ill. Brick Co.*, 431 U.S. at 735–36.

plaintiffs will be unable to prove that the agent’s design was the proximate cause of their injury.¹⁹³ Third, class-action certification will fail the predominance and manageability requirements of Rule 23, as questions of individual injury and agent performance will overwhelm any common issues.¹⁹⁴

Reply.—The framework is designed to overcome these procedural hurdles by creating a clear, design-based cause of action and leveraging evidence that is common to the class.

First, the framework provides two paths around the *Illinois Brick direct purchaser rule*. In a suit against sellers, consumers may argue consistent with the logic of *Apple v. Pepper* that their agent is merely a conduit, making them the direct purchasers.¹⁹⁵ More importantly, the framework creates a distinct, tort-like cause of action against the *buyer-agent’s developer* for negligent design, because this claim targets the developer’s own conduct rather than passed-through overcharges, it is not subject to *Illinois Brick*’s constraints. This is supplemented by the availability of damages in the many states that are *Illinois Brick* repealers, and the fact that *Illinois Brick* does not bar federal injunctive relief.¹⁹⁶

Second, the framework establishes *design-based causation* to satisfy proximate cause requirements. The plaintiff’s burden is to show that the defendant’s negligent deployment of a convergence-prone agent was a material and foreseeable cause of the supracompetitive outcomes. This causal link is supported by the same evidence used for certification: the design record, interaction logs, and sandbox diagnostics showing the agent’s predictable behavior. The factors from *Associated General Contractors*—such as the directness of the injury and the absence of speculative damages—can favor plaintiffs when the overcharge is measured at the transaction

¹⁹³ See *Assoc. Gen. Contractors*, 459 U.S. at 535–45.

¹⁹⁴ See Fed. R. Civ. P. 23(b)(3).

¹⁹⁵ See *Apple v. Pepper*, 139 S. Ct. at 1520–23.

¹⁹⁶ See *ARC Am. Corp.*, 490 U.S. at 101–02 (holding that federal antitrust law does not preempt state indirect-purchaser statutes); see also ANTITRUST MODERNIZATION COMM’N, *supra* note 148, at 7 (noting that indirect purchasers can sue under state law in thirty-six states and the District of Columbia); 15 U.S.C. § 26 (2022) (authorizing injunctive relief); *Apple v. Pepper*, 139 S. Ct. at 1520 n.1 (noting that *Illinois Brick* addressed damages and did not address injunctive relief).

layer and is directly tied to the agent’s documented performance.¹⁹⁷

Third, the focus on design provides the *common proof needed for class certification*. The central questions—whether an agent’s design was negligent, whether it had a foreseeable tendency to converge, and whether it breached the Duty to Disrupt—are common to all class members. The design record, the vendor’s certification status, and the market-level econometric evidence of harm are all forms of class-wide proof. This satisfies the predominance requirement of *Rule 23(b)(3)*. Damages can be calculated using a common, market-wide overcharge model, satisfying the standard from *Comcast* that the damages methodology must align with the theory of liability.¹⁹⁸

D. Economic and Innovation Concerns

1. Chilling Innovation and Small-Entity Burdens

Objection.—The framework’s compliance regime, with its mandatory audits and certification requirements, will disproportionately burden smaller firms, startups, and open-source developers. This will create a “compliance moat” that entrenches incumbent firms, who can more easily absorb these costs, and will ultimately chill the very innovation and market entry that competition law is meant to protect.¹⁹⁹

Reply.—This objection correctly identifies a critical design constraint for any new regulation. The framework is therefore built not to chill innovation, but to *channel it* in a procompetitive direction, with specific mechanisms to mitigate the burdens on new and smaller entrants.

First, the framework is *scaled and tiered*. The most intensive and costly oversight tools, such as mandatory participation in regulatory sandboxes, are reserved for high-risk, concentrated markets dominated by significant players. The duty itself only attaches after a market crosses a

¹⁹⁷ See *Assoc. Gen. Contractors*, 459 U.S. at 537–45.

¹⁹⁸ See Fed. R. Civ. P. 23(b)(3); *Comcast Corp.*, 569 U.S. at 35; see also *Olean*, 31 F.4th at 664–83 (discussing statistical proof of class-wide impact).

¹⁹⁹ Cf. JAMES D. CAMPBELL, AVI GOLDFARB & CATHERINE TUCKER, *Privacy Regulation and Market Structure*, 24 J. ECON. & MGMT. STRATEGY 47, 47–49 (2015) (finding that privacy regulation disproportionately burdens small and new firms); GARRETT A. JOHNSON, SCOTT K. SHRIVER & SAMUEL G. GOLDBERG, *Privacy and Market Concentration: Intended and Unintended Consequences of the GDPR*, 69 MGMT. SCI. 5695, 5695–97 (2023) (documenting increased market concentration following GDPR implementation).

material operational threshold, exempting nascent or experimental ecosystems.²⁰⁰ Furthermore, the framework scopes liability to the *commercial deployment* of market-facing agents, not to upstream open-source development. This protects individual maintainers, while specific safe harbors protect good-faith researchers.²⁰¹

Second, the framework is designed to foster an ecosystem of *modular and open-source compliance*. A small firm or startup does not need to build a complex, bespoke compliance architecture from scratch. Instead, it can integrate a *pre-certified, third-party coordination-breaker module*—some of which could be open-source and available at little to no cost—to satisfy the duty’s core technical requirements. Regulators would be directed to streamline the certification process for any agent that incorporates such a pre-certified component.²⁰²

This approach directly mitigates the risk of a compliance moat. It materially lowers the cost and technical expertise required for a new entrant to comply. It creates a competitive market for compliance tools, which drives down prices and improves quality. And it helps prevent incumbents from capturing the process by ensuring that standards are transparent and accessible. By making rivalry-by-design a modular and accessible feature, the framework democratizes compliance and ensures that the agentic economy remains a field for dynamic competition, not just for the largest and most established players.

2. *Perverse Economic Effects and Consumer Welfare Paradoxes*

Objection.—The framework’s mandate to “disrupt” could lead to perverse economic outcomes that ultimately harm consumers. Mandated strategic randomness could degrade the efficiency of dynamic pricing, creating chaotic price volatility that harms risk-averse consumers who value predictability.²⁰³ Furthermore, not all forms of coordination are harmful; in some cases, a stable

²⁰⁰ See *supra* Part V.E (detailing the “Phase 0” market designation process).

²⁰¹ See *supra* Part V.B.3 (discussing researcher safe harbors).

²⁰² Cf. NAT’L INST. OF STANDARDS & TECH., *supra* note 143, at 33 (describing “AI RMF Profiles” as mechanisms to tailor risk management to specific sectors and use cases).

²⁰³ See KELLY L. HAWS & WILLIAM O. BEARDEN, *Dynamic Pricing and Consumer Fairness Perceptions*, 33 J. CONSUMER RES. 304, 304–11 (2006) (finding that consumers perceive dynamic pricing as unfair when it violates equity norms); YUMI YOSHIDA ET AL., *Which Dynamic Pricing Rule Is Most Preferred by Consumers?—Application of Choice Experiment*,

price path might be preferable to destructive volatility or repeated price wars that impair investment and service quality. The framework’s blunt mandate could inadvertently destroy beneficial forms of coordination, such as industry standardization or joint ventures that support investment in quality and safety.

Reply.—This objection raises a valid concern about the risk of over-correction, but it misinterprets the nature of the duty and overlooks its calibration. The goal is to engineer *sustainable rivalry*, not destructive chaos.

First, the duty requires *strategic, not irrational, randomness*. The goal of injecting controlled exploration into an agent’s behavior is to make it less predictable to its rivals, thereby undermining the certainty required for tacit collusion to hold. It does not require the agent to make choices that are against its user’s interest. Well-calibrated randomness can destabilize supracompetitive equilibria with a *bounded* impact on allocative efficiency—especially if constrained to a near-optimal action set—and is designed to improve consumer welfare on net in markets where convergence risk is material.²⁰⁴ The certification process would explicitly test for this, ensuring that procompetitive designs do not create excessive volatility by enforcing stability constraints, such as volatility budgets (variance caps).²⁰⁵

Second, the framework is targeted at *tacit collusion on price and output*, the forms of coordination that are most unambiguously harmful to consumers and have the fewest redeeming efficiencies. It does not target beneficial forms of industry coordination, such as the development of interoperability standards or safety protocols. The rebuttable presumption is triggered by evidence of *supracompetitive outcomes*—prices sustainably above the competitive level after controlling for cost and demand shocks—not by evidence of industry standardization.²⁰⁶

6 J. ECON. STRUCTURES art. 4 (2017) (noting that consumer risk aversion can drive avoidance of high-volatility real-time pricing).

²⁰⁴ See Johnson, Rhodes & Wildenbeest, *supra* note 18, at 1845 (finding that procompetitive platform design rules increased consumer surplus relative to the collusive baseline); cf. Calvano et al., *supra* note 10, at 3279–80 (showing that collusion depends on low experimentation rates, implying that enforced exploration disrupts the mechanism).

²⁰⁵ See *supra* Part III.B (detailing the use of regulatory sandboxes to test for performance metrics); see also *supra* note 159 (discussing price dispersion and retaliation metrics).

²⁰⁶ See *supra* Part III.C (defining the prima facie effects showing); see also *Broad. Music*, 441 U.S. at 19–24 (distinguishing naked restraints from procompetitive integration).

Finally, the argument that consumers might prefer stable high prices to competitive volatility is inconsistent with the foundational principles of antitrust law. While price wars can be disruptive, the long-term effect of vigorous competition is to drive prices toward competitive levels, increase quality, and spur innovation. The framework is premised on the principle that a market characterized by periodic, procompetitive defections is healthier and ultimately more beneficial for consumers than one locked in a stable, but artificially inflated, price equilibrium. The duty is not to destroy markets, but to ensure they remain genuinely competitive.²⁰⁷

3. *International Competitiveness and Regulatory Arbitrage*

Objection.—The framework will handicap U.S. firms in global markets. Imposing a unique and costly “Duty to Disrupt” on American developers will put them at a competitive disadvantage against foreign rivals who can operate without such constraints. This will create a powerful incentive for regulatory arbitrage, driving AI development and investment offshore to jurisdictions with laxer rules. Ultimately, this could harm U.S. consumers, who would be left interacting with foreign-developed agents designed without any procompetitive safeguards, while the domestic AI industry withers.

Reply.—This concern is significant, but it overstates the framework’s costs and underestimates its potential to set a global standard. The framework is designed to be a source of competitive advantage and is part of a harmonized, not unilateral, approach to AI governance.

First, the objection assumes that procompetitive design is purely a cost. But in a global marketplace where regulators and enterprise customers increasingly demand demonstrable AI assurance, a certified, procompetitive agent can be a source of *competitive differentiation*. A buyer agent that can be trusted not to coordinate, or a seller platform that can prove it fosters a fair market, has a powerful advantage in attracting users. The framework creates a market for trustworthy AI, turning a compliance obligation into a valuable asset.²⁰⁸

²⁰⁷ See *Socony-Vacuum*, 310 U.S. at 221–23 (rejecting the “age-old cry of ruinous competition and competitive evils” as a defense to price-fixing); *Nat’l Soc’y of Prof’l Eng’rs v. United States*, 435 U.S. 679, 695 (1978) (holding that the “statutory policy precludes inquiry into the question whether competition is good or bad”).

²⁰⁸ See ISO/IEC 42001:2023, *supra* note 120; NAT’L INST. OF STANDARDS & TECH., *supra* note 143; Regulation (EU)

Second, the framework is not designed to operate in isolation. As discussed in Part V, a core pillar of the implementation path is *international coordination and mutual recognition*. The U.S. should work with peer regulators in the EU, UK, and other jurisdictions to harmonize the technical standards for antitrust red-teaming and to create a “passporting” system where a certification from one trusted regime is recognized in others. This dramatically reduces the compliance burden for global firms and mitigates the risk of regulatory arbitrage by creating a level playing field.²⁰⁹

Finally, the risk of U.S. consumers being harmed by unregulated foreign agents is real, but it is a reason *for*, not against, the framework. Under the Foreign Trade Antitrust Improvements Act (FTAIA), the Sherman Act reaches foreign conduct that has a direct, substantial, and reasonably foreseeable effect on U.S. commerce that gives rise to the claim.²¹⁰ A foreign-developed agent deployed into the U.S. market would be subject to the duty. By establishing a clear, auditable standard for what constitutes responsible, procompetitive design, the framework makes it *easier*, not harder, for U.S. enforcers to police the conduct of all agents operating within our borders, regardless of their origin. It sets a benchmark for responsible conduct that could become the de facto global standard.²¹¹

VII. CONCLUSION: FROM CONSPIRACY TO CONVERGENCE, DETECTION TO DESIGN

Algorithmic agents can independently learn to sustain supracompetitive prices through *convergence without conspiracy*, stripping Section 1’s agreement requirement of its diagnostic power. Yet the very feature that frustrates traditional inference—the programmability of market actors—creates a parallel *opportunity* that classic antitrust never had. Human beings cannot be *coded* to internalize

2024/1689.

²⁰⁹ See *supra* Part V.C.1 (discussing international coordination and mutual recognition).

²¹⁰ See 15 U.S.C. § 6a (2022); see also *Empagran*, 542 U.S. at 162–63.

²¹¹ See generally ANU BRADFORD, *THE BRUSSELS EFFECT: HOW THE EUROPEAN UNION RULES THE WORLD* (Oxford Univ. Press 2020) (theorizing how unilateral regulation by a major market can set global standards).

a duty of rivalry, nor can their tacit understandings be stress-tested. Algorithmic agents, however, are *architected*: their objectives specified and their behavior simulated. Their design records can be preserved, linking architecture to market effect. Discovery need not infer intent from human communications; every line of code, every action, and every outcome can be tracked. While we cannot engineer human rivalry, we *can* engineer algorithmic rivalry—and, crucially, *verify it*.

The framework advanced here operationalizes that opportunity. The *Duty to Disrupt* reorients antitrust from policing intent to engineering *rivalry*. It imposes a negligence-style obligation on *all* market-facing agents to incorporate procompetitive features, reinforced by three pillars: *structural separation* to prevent conflicts; *proactive oversight* via mandatory audits and sandboxes; and *evidentiary reform* that shifts the burden of production to developers. This regime transforms competition from an abstract ideal into an *auditable design property*.

Implementation need not wait. Federal agencies can immediately treat the deployment of uncertified, convergence-prone agents as an unfair method of competition, using existing authority to demand design records and mandate interactive audits. This enforcement will build the evidentiary foundation for comprehensive legislation that codifies the duty, formalizes structural remedies, and establishes durable safe harbors.

The central task of twentieth-century antitrust was to detect conspiracies among humans after the fact. The task for the twenty-first is to supervise architectures that keep markets competitive *by construction*. The choice is not between *laissez-faire* and *command-and-control*; it is between hoping for rivalry and *engineering it in*. The law should choose the latter.