



May 12, 2026

Categories:

[Agentic Commerce](#)

[Series](#)

[Publications](#)

Authors:

[Nilesh \(Neal\) Patel](#)

[Brad S. Keeton](#)

Agentic Commerce, Price Control, and the New Antitrust Questions

When the Gatekeeper Picks the Agent, Competition Breaks Before It Begins

Agentic commerce arrives in a field that has spent the better part of 40 years catching up to digital markets, and is now catching up to AI. Antitrust law is built almost entirely on judicial interpretation rather than statutory revision, which has historically allowed it to absorb new commercial forms — two-sided platforms, network effects, algorithmic conduct — without waiting on Congress. Agentic commerce is the next adaptation problem, and the doctrine is not yet there. The old framework asks who set the price, who agreed with whom, and who was excluded from the market. Each of those questions assumed a human buyer and a human seller making decisions in real time. The agent collapses that assumption. The buyer is now an optimization layer. The seller is now an API endpoint. The market is now a routing decision made inside a platform's runtime, often by software the platform itself controls.

Prior [articles in this series](#) traced what happens inside an agentic transaction: who is identified, what authority is delegated, how assent is formed, where loss is allocated, what the payment rails recognize. The next layer is the one above the transaction. The same private actors who own the rails also decide which agents are allowed to use them, and which prices, products, and merchants those agents are allowed to see. That is where competition law begins.

The question is no longer whether one merchant can block one agent. The Ninth Circuit will eventually resolve that narrower dispute in [Amazon v. Perplexity](#). The harder question is whether *coordinated* blocking by the major platforms, or *systematic* routing that advantages a platform's own agent over independent ones, is itself an antitrust violation. At an institutional level, the answer is taking shape in four places at once: the DOJ's RealPage consent decree, the FTC's surveillance pricing study and the parallel House

Oversight investigation, the Italian Competition Authority's interim measures against Meta, and the European Commission's first review of the Digital Markets Act. Each addresses a different fault line. Together they describe the architecture of competition in agentic commerce.

Algorithmic Coordination No Longer Requires a Meeting of the Minds

The doctrinal threshold for unlawful coordination under Section 1 of the Sherman Act has always been agreement. The DOJ's November 2025 proposed final judgment with RealPage marks the clearest federal articulation to date of how that threshold applies when the "agreement" runs through software rather than a smoke-filled room.

The settlement, filed November 24, 2025 in the Middle District of North Carolina, does not require RealPage to admit liability. It does require RealPage, for a seven-year term, to stop running its revenue management software on competitors' nonpublic data in runtime operations, to limit model training to backward-looking nonpublic data at least twelve months old, to refrain from generating geographic effects narrower than statewide, to redesign features that limited price decreases or aligned pricing among competing users, to cease conducting market surveys for nonpublic competitive intelligence, and to operate under a court-appointed monitor. The settlement also prohibits RealPage from making identical pricing recommendations to different owners in the same market and requires user-set parameters that allow recommendations to fall below historical floors as readily as they exceed historical ceilings.

The structure of the consent decree is the substance. The DOJ did not need to prove an explicit agreement among landlords. It alleged, and the proposed judgment treats as actionable, that the software itself functioned as the coordinating mechanism. Competitors did not need to talk to each other when an algorithm could read their data and output a synchronized recommendation.

This matters for agentic commerce because the same architecture is emerging in adjacent markets, including things like procurement

agents trained on shared vendor data, pricing agents that consume real-time competitive feeds through common APIs, and travel and hospitality agents whose recommendations are shaped by data shared with the agent's operator. Each of these systems can produce parallel pricing without parallel intent. The RealPage settlement tells the market what counts as the line: nonpublic competitive data plus a runtime feedback loop plus features that suppress downward price movement. According to the DOJ, the presence of those factors is sufficient for it to seek injunctive relief, without proving anyone in the room shook hands. Firms deploying agentic procurement or pricing systems should treat the settlement as a compliance template rather than a sector-specific resolution.




Figure: The Algorithmic Coordination Spectrum. RealPage and Cartwright are reshaping the contested middle.

California's recent Cartwright Act amendments push the same coordination concern further than federal doctrine has, treating coercive use of common pricing algorithms as *per se* unlawful — the first state statute to characterize algorithmic coordination at that end of the antitrust spectrum.

Consumer Reports and Groundwork Collaborative published their joint investigation of Instacart's algorithmic pricing experiments in December 2025, finding that identical items in identical stores at the same moment carried as many as five different price points across 437 test shoppers, with an average high-low differential of thirteen percent and a maximum of twenty-three percent. Instacart wound down the relevant tests within weeks of the report. The

episode is not a Sherman Act case (although it potentially raises Robinson–Patman Act price discrimination concerns). It is instead a preview of where consumer–protection enforcement will land when an agent, rather than a human shopper, sits between the consumer and the price the consumer ultimately pays.

Surveillance Pricing Is the Consumer–Protection Mirror of the Coordination Problem

Algorithmic coordination is the antitrust face of automated pricing. Surveillance pricing is the consumer–protection face. Regulators are pursuing both, and they are doing so on parallel tracks that converge precisely where the agent sits.

The operative distinction, articulated by the FTC and now codified in legislative drafting across the states, is between *dynamic pricing*, which responds to market conditions like inventory, demand, and seasonality, and *surveillance pricing*, which responds to characteristics of the individual consumer rather than the market. Dynamic pricing remains lawful (subject to the requirements of the Robinson–Patman Act). Surveillance pricing is the practice of using detailed consumer data, including location, browsing history, demographics, behavioral inferences, mouse movements, and cart abandonment patterns, to set individualized prices for the same product. The FTC’s January 2025 staff perspective on its Section 6(b) study, based on documents from Mastercard, Accenture, PROS, Bloomreach, Revionics, and McKinsey & Co., found that intermediaries supporting at least 250 retail clients use these signals to determine pricing and promotions at the individual consumer level.

The federal posture has since shifted. The Ferguson FTC has not advanced the Khan-era surveillance pricing inquiry on the same trajectory, and reporting from the National Association of Attorneys General 2026 annual conference indicated that state AGs view themselves as filling the resulting enforcement vacuum. That reading is consistent with the public record. New York’s algorithmic pricing disclosure law took effect in November 2025 and requires merchants to mark prices set by an algorithm using personal data with the disclosure “THIS PRICE WAS SET BY AN ALGORITHM USING

YOUR PERSONAL DATA.” Maryland has enacted the Protection from Predatory Pricing Act, which restricts certain personalized pricing practices and treats violations as deceptive trade practices. California is pursuing surveillance pricing through privacy enforcement under the CCPA framework. Dozens of additional states have introduced parallel bills, with definitions, sectoral scope, and enforcement mechanisms that vary significantly across jurisdictions.

Congressional oversight has filled the federal gap on a different vector. The House Oversight Committee opened an investigation on March 5, 2026 into AI-driven pricing across travel and platform industries, sending document requests to major operators concerning revenue management algorithms, consumer data inputs, A/B price testing protocols, and internal communications about pricing tools. The committee has framed surveillance pricing as a “black box” in which algorithms infer willingness to pay and adjust prices accordingly without consumer awareness. The investigation creates immediate practical exposure independent of any eventual statute, including compelled document production, public hearings, and referrals to the FTC, DOJ, or state AGs.

The doctrinal architecture is now clearer than it was eighteen months ago, and it is not a single body of law. Surveillance pricing risk in the United States today crosses three overlapping state-law regimes (consumer protection through UDAP, data privacy through CCPA-style statutes, and state antitrust through algorithmic-coordination provisions like the amended California Cartwright Act), federal Section 5 enforcement at the FTC, and Congressional oversight by the House. None of these regimes was designed for agent-mediated transactions, and each will be applied to them.

The agentic-commerce implication is the one that should change how counsel scopes the risk. An AI agent that sits between the consumer and the merchant is, definitionally, a data-collection layer. The agent observes, records, and transmits the consumer’s behavior at a granularity that an unmediated browser session cannot match. Every signal the FTC’s 6(b) study identified as surveillance-pricing input, including things like precise location, browsing history, demographic inference, behavioral patterns,

mouse movements, and abandoned carts, is more readily available to a merchant when an agent has been operating in the consumer's session than when the consumer has been operating alone. Deploying or accepting an agent therefore creates a new compliance surface: the agent itself, the data the agent both relies upon and generates, and the use to which that data is put in pricing decisions made downstream. Disclosures designed for human–shopper interactions will not survive scrutiny when the underlying data flow is agent–mediated.

The strategic question for any company deploying an agent in a consumer–facing transaction is whether the agent's data outputs are being used by the merchant, or by the merchant's pricing intermediary, in ways that would trigger surveillance–pricing scrutiny if the consumer learned about them. The question for any merchant accepting agent–mediated traffic is the same in mirror image: what is the merchant's pricing system doing with the agent's data and would the merchant's existing privacy and pricing disclosures survive a House Oversight document request.

Coordinated Blocking Is the Antitrust Frontier, and the Platforms Know It

The central question is whether coordinated blocking by major platforms, or algorithmic routing that advantages affiliated agents, constitutes an antitrust violation. Perhaps the most important development of the past six months is that European competition authorities have begun answering it in the affirmative.

Italy's Competition Authority ("AGCM") opened Case A576 against Meta in July 2025, focusing on the integration of Meta AI into WhatsApp. On November 25, 2025, the AGCM expanded the investigation to include Meta's October 15, 2025 update to the WhatsApp Business Solution Terms, which prohibited general–purpose AI assistants from using the Business API when AI functionality was the primary offering. The terms were scheduled to take full effect on January 15, 2026 and would have removed Microsoft Copilot, OpenAI's ChatGPT, and other third–party assistants from the platform. On December 24, 2025, the AGCM issued interim measures suspending the contested terms in Italy,

finding that the conduct may constitute an abuse of dominance under Article 102 TFEU and that allowing the policy to take effect would cause “serious and irreparable harm to competition” by foreclosing rivals during a formative stage of the conversational-AI market. The European Commission opened a parallel investigation covering the rest of the European Economic Area in the same month.

The agencies’ legal theory is foreclosure of competition through contractual exclusion. The factual predicate is platform dominance plus a written term that removes competing agents from a key distribution surface, and the remedy is interim suspension. Translated into the agentic-commerce context, the AGCM’s order stands for a proposition that has not yet been tested in U.S. federal court: a platform’s contractual decision to bar third-party agents from a service the platform itself uses to distribute its own agent is the kind of conduct that competition authorities will treat as presumptively exclusionary.

The contrast with U.S. doctrine is sharp and worth pulling forward. *Amazon v. Perplexity* is being litigated under the Computer Fraud and Abuse Act and California’s computer-fraud statute, not the Sherman Act. Judge Chesney’s March 9, 2026 preliminary injunction held that Amazon was likely to succeed on a CFAA theory because Comet accessed password-protected portions of the Amazon site “with the Amazon user’s permission, but without authorization by Amazon,” and that user permission was not a substitute for platform authorization. The Ninth Circuit issued an administrative stay on March 16, 2026, and the merits appeal is now docketed as Case No. 26-1444. The dispute is being treated as an access issue, not a market power case, and antitrust claims have played only a minor role. Although Perplexity argues that Amazon is protecting its advertising revenue from bypassing agents, that theory has not yet been successfully pleaded as a Section 2 monopolization claim.

European authorities are using existing competition law to police agent exclusion, while American courts are using a 1984 anti-hacking statute to ratify it. Until and unless a U.S. enforcer or private plaintiff successfully recasts coordinated agent blocking as a Sherman Section 1 or Section 2 case, the asymmetry will continue,

and platforms will route their most aggressive agent-exclusion policies through technical access controls rather than open contractual prohibitions. The strategic implication for any company building or deploying an agent that reaches consumer-facing platforms is that the favorable forum, for now, is in Brussels and Rome.

Pricing Parity Becomes Self-Enforcing When the Agent Is the Comparison Shopper

Most-favored-nation clauses, also called platform parity clauses, have long been the pressure point where contract law meets competition law in digital marketplaces. These types of clauses generally prohibit a merchant that sells through a platform from offering a lower price elsewhere, sometimes only on the merchant's own direct site (a "narrow" MFN), sometimes across all channels (a "wide" MFN). European authorities have aggressively curtailed wide MFNs in the hotel-booking sector since 2015. The Digital Markets Act now bans MFNs entirely for designated gatekeepers. U.S. courts have been more permissive, though Amazon's pricing provisions have survived motions to dismiss in the Western District of Washington and remain under active rule-of-reason scrutiny in *Frame-Wilson v. Amazon.com*.

Agentic commerce changes the operative economics of these clauses in a way the existing case law has not yet fully metabolized. The doctrinal worry about MFNs has always been that they suppress price competition by removing a merchant's incentive to discount on alternative channels: if the merchant must give the platform every price drop, the merchant stops dropping prices anywhere. In practice, that suppression rarely materialized. Human shoppers had the option to comparison-shop across channels but rarely bothered.

The AI agent does not have that limitation. As an optimization layer that scans every available channel and routes the consumer to the lowest available price, this is exactly the consumer-side discipline MFNs were designed to neutralize. When a price-comparison agent becomes the default purchasing surface, the value of an MFN to the platform increases proportionally, because the clause now

suppresses competition the agent would otherwise have surfaced. A wide MFN against an agent-mediated market is functionally a tax on price competition, paid by every merchant on the platform, and collected through the elimination of the agent's most useful function.

The doctrinal consequence is that platform MFNs that were borderline-defensible in a human-shopper market become considerably harder to justify in an agent-mediated one. The pro-competitive efficiency argument, being that the platform invested in matchmaking and needs MFNs to prevent customer free-riding, weakens when the platform's matchmaking is partially automated and the customer never sees most of what the platform shows. The exclusionary argument, being that MFNs deny rival platforms the ability to compete on price, strengthens when the rival platform's best argument for switching is precisely the lower price the MFN forbids.

The closest U.S. authority on this is probably *Ohio v. American Express*, where the Supreme Court upheld anti-steering provisions on a two-sided platform under a rule-of-reason analysis that demanded plaintiffs show net harm across both sides of the platform. Defenders of platform MFNs in agent-mediated markets will lean on *Amex* to argue that the platform's investment in matchmaking, fraud control, and merchant onboarding justifies restrictions on price competition. The counter is that the *Amex* analysis assumed a human cardholder making a real-time payment-method choice. When the choice is being made by an optimization layer that the platform itself cannot see, the two-sided efficiency story becomes harder to plead and easier to attack on the record.

Epic v. Apple is probably the second U.S. anchor and may be the more important one for agent-mediated transactions. The Ninth Circuit's December 2025 contempt affirmation held that Apple's twenty-seven percent commission on off-app purchases and its "scare screen" warnings functioned as evasion of an anti-steering injunction by making the permitted alternative economically and behaviorally unavailable. The court treated *prohibitive commission* and *engineered friction* as cognizable forms of anti-steering, not just outright prohibition. Translated into agentic commerce, the

doctrinal lesson is that a platform cannot defend an MFN, an agent-discrimination rule, or an interface-friction requirement by pointing to a nominal alternative path that no rational counterparty would actually use. Counsel advising clients on platform contracts in 2026 should treat agent-mediated demand as a material change in the antitrust risk profile of any pricing parity clause, regardless of where the clause has previously been challenged.

The Identity and Authority Stack Has Become a Competitive Lever

The second and third articles in our agentic commerce series treated [identity](#) and [delegated authority](#) as threshold legal problems. They are also competitive instruments. A platform that requires agents to identify themselves through user-agent strings, to pre-register through a Trusted Agent Protocol, or to satisfy network-level credential checks has the technical and contractual ability to decide which agents pass and which do not. *Amazon v. Perplexity* makes that explicit: Comet's failure to identify itself as an agent was the predicate for CFAA liability. A platform's agent-identification rule is therefore both a legitimate fraud-control mechanism and a potential exclusion device, and the line between them runs through how the rule is administered.

The diagnostic question is whether the rule is applied symmetrically. A merchant that admits its own agent without registration, or its preferred partner's agent on favorable terms, but requires hostile or unaffiliated agents to satisfy harder authentication conditions, is not running a fraud-control program. It is running a discrimination program with a fraud-control label. The architecture of the access stack, which are the layers identified in Article 6, is the architecture through which discrimination becomes invisible. Each layer is a defensible technical requirement when viewed alone. The combined effect of asymmetric administration is foreclosure that no single layer would establish on its own.

The Sherman Act gives this conduct fewer doctrinal handles than the European framework. *Verizon Communications v. Trinko* foreclosed most refusal-to-deal claims outside the narrow

circumstances of *Aspen Skiing*, where the Court found a unilateral termination of a profitable existing course of dealing actionable, and the Supreme Court has expressly declined to recognize an essential-facilities doctrine as a freestanding theory. The realistic U.S. theories are therefore narrower: anti-steering and friction-engineering claims drawing on *Epic v. Apple*, exclusive-dealing or de facto exclusivity claims where the platform's preferred-agent terms function as practical exclusivity, and Section 5 unfair-method-of-competition theories at the FTC where the conduct evades neat fit under the Sherman Act. Each of these requires factual development that platforms have so far avoided by routing decisions through technical infrastructure rather than written terms. The AGCM's WhatsApp action shows what the case looks like when the discrimination is written down. The harder cases will be the ones where the discrimination is engineered into the access layer instead.

The European Standard Will Set the Floor for Multinational Deployments

The European Commission published its first statutory review of the Digital Markets Act on April 27, 2026. The Commission concluded that the DMA "remains fit for purpose and does not need to be revised," that AI services will not be designated as core platform services in this review cycle, and that enforcement priorities will focus on existing obligations rather than expanded scope. Stakeholder submissions had pressed for designation of dominant AI assistants as "virtual assistants" under the DMA and for new CPS categories covering foundation models. The Commission declined both. The Open Markets Institute Europe and others called the conclusion a missed opportunity. Whatever the merits of that critique, the regulatory implication for agentic commerce points in one direction: the rules already on the books, including the gatekeeper obligations, the MFN ban, and the data-use restrictions of Articles 5 and 6, are the rules that will govern AI-mediated transactions through at least the next review cycle in 2029. New AI-specific obligations will arrive, if they arrive, through Member State enforcement actions like the AGCM's Meta case and through

Commission-level investigations under existing law.

The U.S. picture is much different in shape and significantly more fragmented in execution. The DOJ's RealPage settlement and its statements of interest in adjacent algorithmic-pricing cases set the federal antitrust line on coordination. The FTC's posture on surveillance pricing has shifted under new leadership, and state attorneys general have moved to fill the resulting gap. The House Oversight Committee's March 2026 AI pricing investigation operates as a third federal track, parallel to enforcement and capable of generating practical exposure even where no statute or rule applies. At the state level, California's Cartwright Act amendments work the coordination problem from one direction while the consumer-protection statutes — New York's algorithmic pricing disclosure law, Maryland's Protection from Predatory Pricing Act, similar bills in Colorado, Pennsylvania, and dozens of other states — work the surveillance pricing problem from the other. Both layers sit beneath this federal patchwork and are moving faster than any of its components.

For companies operating across both jurisdictions, the practical effect is that the most demanding regime, the European one, will set the floor, because the cost of building separate agentic systems for separate markets exceeds the cost of compliance with the stricter standard.

Action Items

- **Treat any algorithm that consumes nonpublic competitive data in runtime as a presumptive antitrust risk and move it to historical inputs.** The RealPage consent decree's twelve-month aging requirement is the operative federal benchmark. Counsel should not advise clients that other thresholds are safe absent specific facts that distinguish the system from the one the DOJ challenged. If the system cannot perform without real-time competitive data, the system needs to be redesigned, not legally justified.
- **Conduct an MFN audit on every platform contract, with agent-mediated demand treated as a material change in**

risk. Wide MFNs that survived earlier rule-of-reason analysis on a human-shopper record should be re-evaluated against the empirical reality of price-comparison agents. Where wide MFNs cannot be eliminated, document the procompetitive justifications in a record contemporaneous with renewal, not after a complaint is filed.

- **Map the data flow from any deployed agent into pricing systems and audit it against surveillance-pricing risk on a state-by-state basis.** The analysis runs in two directions. For companies deploying agents, the question is what consumer data the agent is generating and where that data lands. For merchants accepting agent traffic, the question is whether their existing privacy and pricing disclosures cover agent-mediated sessions and whether their pricing systems would survive a House Oversight document request. The architecture of compliance is no longer a single national standard. It is a set of state-specific obligations layered on top of federal Section 5 and Congressional oversight risk, and it requires jurisdictional mapping rather than one-size-fits-all policy.
- **Build agent-identification programs that are administered symmetrically and that document the symmetry.** The legitimacy of any agent-authentication regime as an antitrust matter depends on whether the same rules apply to the platform's own agent and its preferred partners as to unaffiliated competitors. The contemporaneous record of administration is the difference between a fraud-control defense and a foreclosure case.
- **Map every commercial dependency on a dominant platform and assume the contractual terms will change unilaterally.** Meta's October 2025 WhatsApp Business Solution Terms update is the canonical fact pattern: a platform decides that a category of users is now a competitor, rewrites the terms, and gives ninety days' notice. Any business model whose viability depends on continued access to a single dominant platform's API is a business model with a known failure mode. The work of contingency planning

is unglamorous and is also the difference between a regulatory inconvenience and a corporate crisis.

- **For multinational deployments, build to the European standard and document compliance accordingly.** The DMA's gatekeeper rules, the AGCM's interim-measures playbook, and the Commission's prosecutorial posture under Article 102 will continue to set the practical ceiling on platform conduct. Building a U.S.-only compliance posture for an agentic system that touches European users is a false economy.

Please contact the authors if you have questions or comments on this article. You can also reach out to any member of the firm's [Data, Digital Assets & Technology](#) and [Antitrust & Trade Regulation](#) teams for help navigating AI deployments, governance, and antitrust laws.

This article was prepared with the assistance of generative AI tools. The analysis, conclusions, and legal positions are the authors' own.

Agentic Commerce Series

- [Part 1 — AI Agents Are Starting to Act Inside the Transaction, and Commerce Law Is Not Ready](#)
- [Part 2 — The Identity Problem: Authentication, Fraud, and Who's Actually Buying](#)
- [Part 3 — The Authority Problem: When Does an Authorized Agent Become an Unauthorized Buyer?](#)
- [Part 4 — Contracting by Agent: When the Agent Clicks, Who Assents?](#)
- [Part 5 — When the AI Purchase Goes Wrong, Who Pays, and Who Can Prove It?](#)
- [Part 6 — The Gatekeepers of Agentic Commerce](#)
- [Part 6 \(A\) — The Payment Infrastructure Layer: How Network Rules Already Govern Agentic Commerce](#)
- [Part 7 — Agentic Commerce, Price Control, and the New Antitrust Questions](#)
- [Part 8 — What Epic Reveals About Agentic Commerce](#)

© 2026 FBT Gibbons LLP. May be considered attorney advertising in some jurisdictions.