Silent Handshakes: Tacit Collusion by Autonomous Pricing Algorithms and the Limits of Antitrust Law

An Economic and Legal Analysis

Abstract

Antitrust law has historically policed collusion by looking for evidence of communication — a phone call, an email, a handshake at a trade conference. But a growing body of experimental economics and a wave of recent litigation suggest that independent, self-learning pricing algorithms can arrive at supracompetitive prices without ever "talking" to one another, simply by learning from market feedback. This paper examines the economic mechanics of algorithmic collusion, drawing on the landmark Q-learning experiments of Calvano, Calzolari, Denicolò, and Pastorello (2020), and evaluates how this challenges the legal doctrine that has governed price-fixing since the Sherman Act of 1890. Using the U.S. Department of Justice's 2024–2025 case against RealPage Inc. in the multifamily rental housing market as a central case study, alongside emerging European enforcement actions in hotels, airlines, and banking, this paper argues that existing antitrust frameworks — built around proof of *agreement* — are increasingly mismatched to a world where coordination can emerge from parallel, independent optimization rather than intent. The paper concludes by evaluating three policy responses: information-sharing restrictions, algorithmic auditing requirements, and outright regulation of pricing-algorithm design, and argues that a hybrid, data-focused approach is currently the most legally and economically defensible path forward.

1. Introduction

Textbook oligopoly theory has always had an uncomfortable subplot. Even without explicit collusion, a handful of firms competing repeatedly over time can learn to avoid price wars and settle into higher, more profitable prices — a phenomenon economists call **tacit collusion**. For most of the twentieth century, this was treated as a slow-moving, largely theoretical risk, because sustaining tacit collusion requires firms to closely monitor rivals' behavior, punish defectors, and coordinate expectations — all of which are cognitively and organizationally expensive for human managers to do reliably.

Algorithmic pricing removes each of those frictions. Modern pricing software can observe competitor prices in real time, update its own prices dozens of times per day, and — critically — in the case of **reinforcement-learning algorithms**, discover collusive strategies on its own, through trial and error, without a human ever instructing it to do so. This paper's central question is deceptively simple: *if no human ever agreed to fix prices, and no algorithm was ever told to collude, can the resulting high prices still be illegal?*

This is no longer a hypothetical. In August 2024, the U.S. Department of Justice, joined by eight state attorneys general, sued RealPage Inc. — a company whose software recommends rents to landlords managing millions of U.S. apartment units — alleging that its algorithm functioned as a price-fixing hub for competing landlords. The case, along with dozens of related state and private lawsuits, settled in November 2025, with RealPage agreeing to significant restrictions on how it uses competitor data, though admitting no wrongdoing. Parallel investigations are now underway across Europe, in sectors ranging from hotels and airlines to consumer banking and pharmaceutical wholesaling. Algorithmic collusion has moved from an academic curiosity to one of the defining antitrust issues of the mid-2020s.

---

2. The Economics of Algorithmic Collusion

2.1 Two very different kinds of "pricing algorithm"

It is important to distinguish between two categories of software, since they raise very different legal and economic issues.

Rule-based pricing software follows human-specified logic ("match the lowest competitor price minus 2%," "raise price by 5% if occupancy exceeds 90%"). If such rules are explicitly designed to coordinate with competitors, this is functionally identical to traditional price-fixing — the algorithm is just the messenger.

Reinforcement-learning (RL) algorithms, by contrast, are given only an objective (maximize long-run profit) and a set of actions (possible prices), and they learn, through repeated trial and error against the market environment — including other algorithms — which strategies work best. Crucially, they are not told anything about competitors, collusion, or punishment strategies. Whatever behavior emerges, emerges endogenously.

2.2 The experimental evidence: Calvano et al. (2020)

The most influential empirical demonstration of this second category comes from Calvano, Calzolari, Denicolò, and Pastorello, published in the *American Economic Review* in 2020. The authors built a simulated oligopoly market and populated it with independent Q-learning algorithms — a standard reinforcement-learning technique — with no communication channel between them and no knowledge of the underlying economic environment. Their central finding was striking: the algorithms **consistently and reliably learned to charge prices well above the competitive (Nash equilibrium) level, without any communication whatsoever.**

Even more notable, the collusive behavior was not primitive. The algorithms learned genuine punishment strategies: if one algorithm's price dropped in a way that looked like an attempt to undercut rivals, the others would respond with a temporary price war, followed by a gradual, coordinated return to high prices — a dynamic structurally similar to the trigger strategies economists have long modeled as the theoretical basis of human cartel behavior. The authors found this result held up under a wide range of conditions: asymmetric costs, asymmetric market shares, more than two competitors, and various forms of market uncertainty. A related follow-up study by the same authors extended the result to environments with *imperfect monitoring* — where firms cannot directly observe rivals' prices — modeled on the classic Green and Porter (1984) framework, and found that algorithms could still learn to sustain collusive-like pricing using price wars as an implicit punishment signal.

The policy implication is unsettling: this is not a bug that better-designed algorithms will engineer away. It appears to be an emergent property of profit-maximizing learning systems operating in an oligopoly. As more "rungs" of AI sophistication get added, there is no guarantee the problem shrinks — it may in fact grow, because more capable algorithms are better, not worse, at discovering profitable long-run strategies.

2.3 Why algorithms make collusion structurally easier

Beyond the specific finding of Calvano et al., several structural features of algorithmic markets lower the traditional barriers to sustaining collusion:

- **Speed and granularity of monitoring.** Cartels historically collapsed because monitoring cheating was slow and imperfect. Algorithms can observe and react to rivals' prices within seconds.
- **Elimination of the "cheating temptation" problem.** A human sales manager facing a quarterly bonus has an incentive to shade prices down for a quick win. An algorithm optimizing for long-run expected profit has no such short-term temptation, making cooperative equilibria easier to sustain.
- **Plausible deniability.** Because no communication occurs, executives can genuinely — and truthfully — say they never told the algorithm to coordinate with anyone, which complicates the "meeting of the minds" that antitrust law traditionally requires.
- **Third-party aggregation hubs.** When multiple competitors use the *same* third-party software vendor (rather than independent algorithms), the vendor's platform can function as a data-sharing hub even without literal agreement among the competing firms themselves — a dynamic at the center of the RealPage litigation.

---

3. Case Study: United States v. RealPage, Inc.

RealPage provides revenue-management software used by landlords to set rents across a large share of U.S. multifamily housing units. The Department of Justice's amended complaint, filed jointly with several state attorneys general, alleged that RealPage's software pooled **nonpublic, competitively sensitive data** — including real-time rents, lease terms, and occupancy levels — from competing landlords, and used that pooled information to generate pricing recommendations that aligned rents upward across an entire local market, effectively replicating the outcome of a cartel without any direct communication between the landlords themselves.

The theory of harm here is legally significant because it does not rely on the "pure" algorithmic collusion scenario from Calvano et al., where no data is shared at all. Instead, it targets a more common real-world hybrid: independent competitors voluntarily feeding sensitive data into a shared third-party algorithm, which then issues coordinated-looking price recommendations back out. Courts and regulators have increasingly treated this as a "hub-and-spoke" conspiracy — a legal theory with older roots (reinforced by the European Court of Justice's 2016 *Eturas* ruling, which held that Lithuanian travel agencies could be liable for tacit participation in coordination facilitated through a shared booking algorithm) but newly urgent applicability.

In November 2025, the DOJ reached a proposed settlement with RealPage. The settlement's structure is itself an important data point about where regulators believe the legal line sits: RealPage was not fined and admitted no wrongdoing, but agreed to sweeping restrictions, including a prohibition on incorporating nonpublic, current, or granular competitor pricing data into its algorithm's recommendations, a requirement that any data used be at least twelve months old and aggregated no more finely than the state level, and submission to an independent compliance monitor for up to seven years. In parallel, private multidistrict litigation produced more than $140 million in preliminary settlements from landlord defendants who used RealPage's software, and several states — New York prominent among them — have since passed or proposed legislation explicitly extending antitrust liability to landlords who set rents "based on recommendations from a software, data analytics service, or algorithmic device."

The RealPage case illustrates the doctrinal workaround regulators have converged on: rather than trying to prove the harder claim that an *algorithm itself* independently learned to collude (the Calvano et al. scenario), enforcers have found more legal traction in proving that firms shared competitively sensitive, forward-looking data through a common intermediary — a fact pattern closer to traditional information-exchange antitrust violations, just executed through software.

4. The Doctrinal Problem: Section 1 of the Sherman Act Was Not Built for This

Section 1 of the Sherman Act prohibits "every contract, combination...or conspiracy in restraint of trade." Nearly all subsequent case law has interpreted this to require proof of an **agreement** — a meeting of the minds between at least two parties. Purely independent, parallel pricing decisions, even when they produce identical high prices, have traditionally been legal; this is sometimes called the "conscious parallelism" doctrine, most clearly articulated in the Supreme Court's 1993 decision in *Brooke Group* and reaffirmed in *Twombly* (2007), which require "plus factors" beyond mere parallel pricing to infer an actual agreement.

Autonomous algorithmic collusion of the Calvano et al. variety sits in an uncomfortable gap in this framework:

1. There is no explicit agreement between the firms.
2. There is arguably no agreement between the *algorithms* either — they never communicate; they simply each independently learn that cooperative pricing is the profit-maximizing long-run strategy given the observed behavior of the other.
3. Yet the outcome — sustained supracompetitive pricing, enforced through de facto "punishment" price wars — is functionally indistinguishable from a textbook cartel.

Legal scholars, notably Joseph Harrington in his influential analysis of competition law for autonomous artificial agents, have argued that this scenario may fall entirely outside the reach of Section 1 as currently interpreted, because there is no "agreement" in any traditional sense to point to — the collusive outcome emerges from independent optimization, not coordination. This has produced two broad camps in the policy debate:

- **The "current law is sufficient" camp**, which includes the DOJ's own stated position in several *amicus* and statement-of-interest filings: prosecutors argue that using an algorithm as a *vehicle* for exchanging or acting on competitively sensitive data is already covered by existing information-exchange case law, regardless of whether the pricing engine "understands" what it is doing. The RealPage settlement is consistent with this approach — it avoided the harder question of pure algorithmic tacit collusion by focusing on the data-sharing conduct instead.
- **The "new legislation needed" camp**, exemplified by Senator Amy Klobuchar's repeatedly reintroduced Preventing Algorithmic Collusion Act, which would amend the Sherman Act to explicitly prohibit the use of pricing algorithms that *can* facilitate collusion through nonpublic competitor data, and would create a formal audit tool for regulators — shifting the legal test away from proving intent or agreement and toward the algorithm's design and data inputs themselves.

Europe has approached the same problem through its existing "concerted practices" doctrine, which has always been somewhat more permissive than the American "agreement" requirement — the 2016 *Eturas* precedent already established that mere awareness of a shared coordinating mechanism could establish liability. This is one reason European regulators (the UK's Competition and Markets Authority, Italy's AGCM, Poland's UOKiK, and the European Commission itself) have moved comparatively quickly to open sector-specific investigations into hotels, airlines, banking, and pharmaceuticals, even without a RealPage-style smoking gun.

5. Policy Options and Evaluation

Three broad regulatory responses are currently being debated or piloted:

(a) Data-input restrictions (the RealPage-settlement model). Regulators restrict the *type* of data an algorithm can use — barring real-time, granular, nonpublic competitor data, while permitting historical, aggregated, or public data. *Strength:* fits within existing antitrust doctrine on information exchange; does not require new legislation. *Weakness:* does nothing to address the "pure" Calvano-style scenario where no data is shared between firms at all, only independently learned market feedback.

(b) Algorithmic auditing and design mandates. Firms would be required to document and, on request, disclose the objective functions and design parameters of their pricing algorithms to regulators, allowing audits for collusive-prone design choices (e.g., unusually long time horizons in the profit function, which experimental work suggests increases collusive tendencies). *Strength:* targets the mechanism directly, regardless of communication. *Weakness:* technically difficult to audit black-box reinforcement-learning systems; risks exposing trade secrets; may be difficult to distinguish "collusion-prone" design choices from ordinary sound business optimization.

(c) Strict liability for supracompetitive outcomes ("effects-based" regulation). Rather than examining intent or design, this approach would treat sustained, statistically anomalous price elevation in a concentrated market using common software as presumptive evidence of an antitrust violation, shifting the burden to the firm to justify the outcome. *Strength:* closes the doctrinal gap entirely; addresses harm regardless of mechanism. *Weakness:* a significant departure from a century of antitrust jurisprudence built around intent and agreement; risks chilling legitimate, procompetitive dynamic pricing (e.g., algorithms that lower prices during low demand, which benefit consumers).

Given the doctrinal risk of departing too far from established antitrust principles, and the practical difficulty of auditing proprietary AI systems, approach (a) — data-input restriction — currently has the strongest combination of legal defensibility and enforceability, which is consistent with its being the approach regulators have actually converged on in the RealPage settlement. However, it is best understood as a **necessary but not sufficient** response: it addresses the hub-and-spoke, data-sharing version of algorithmic collusion well, but leaves the harder, purely emergent Q-learning scenario from Calvano et al. largely unaddressed. This suggests that legislative proposals like the Preventing Algorithmic Collusion Act, which explicitly contemplate liability based on an algorithm's *capacity* to facilitate collusion rather than proof of an agreement, may ultimately be necessary to close the remaining gap — even though such a shift would represent a genuine and consequential evolution in antitrust law.

---

6. Conclusion

The RealPage litigation and the parallel wave of European investigations confirm that algorithmic pricing collusion is no longer a speculative concern confined to economics working papers — it is an active, evolving front of antitrust enforcement with billions of dollars and millions of consumers at stake. But the underlying economic mechanism identified by Calvano, Calzolari, Denicolò, and Pastorello — that independent, non-communicating reinforcement-learning algorithms can learn to sustain collusive-level prices purely through trial-and-error profit maximization — remains only partially addressed by current enforcement, which has so far succeeded mainly against the more legally tractable case of shared data hubs rather than pure emergent tacit collusion.

This gap matters because it is a preview, not an endpoint. As pricing algorithms become more sophisticated and more widely deployed across more sectors, the conditions under which pure algorithmic tacit collusion could emerge — without any data-sharing vehicle to prosecute — will become more common, not less. Antitrust law was built on the assumption that collusion requires human minds to meet. Whether it can adapt to a world where profit-maximizing outcomes can emerge from machines that never speak to one another at all is likely to be one of the defining competition-policy questions of the next decade.

References

Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial Intelligence, Algorithmic Pricing, and Collusion. *American Economic Review*, 110(10), 3267–3297.

Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2021). Algorithmic collusion with imperfect monitoring. *International Journal of Industrial Organization.*

Harrington, J. E. (2018). Developing Competition Law for Collusion by Autonomous Artificial Agents. *Journal of Competition Law and Economics*, 14(3), 331–363.

Klein, T. (2021). Autonomous algorithmic collusion: Q-learning under sequential pricing. *RAND Journal of Economics*, 52(3), 538–558.

U.S. Department of Justice, Office of Public Affairs. (2025, November 24). *Justice Department Requires RealPage to End the Sharing of Competitively Sensitive Information and Alignment of Pricing Among Competitors.*

U.S. Department of Justice v. RealPage, Inc., No. 24-cv-710 (M.D.N.C., filed Aug. 23, 2024; amended Jan. 7, 2025).

Wilson Sonsini. (2025, December). *DOJ Settles Its Algorithmic Price-Fixing Case Against RealPage.*

Paul, Weiss. (2025, December). *Practical Takeaways From the DOJ's Algorithmic Pricing Settlement.*

Paul, Weiss. (2026, March). *Enforcement on Algorithmic Pricing Starts to Bite in Europe.*

Akin. (2026, March). *When Bots Set Prices: CMA Highlights Real World Risks of Algorithmic Pricing.*

Court of Justice of the European Union, *Eturas UAB and Others v. Lietuvos Respublikos konkurencijos taryba*, Case C-74/14 (2016).

Note: This paper synthesizes publicly reported litigation and regulatory developments through early 2026. Readers using this paper for submission should independently verify docket numbers, settlement terms, and case citations against primary legal sources before final submission, as antitrust litigation of this kind remains actively developing.