Yale University
Thurman Arnold Project

Department of Justice Conference
Paper 2

Addressing The Potential and Pitfalls of Dynamic Pricing Algorithms

May 2023

Kart Kandula[1]

[1] The author acknowledges and thanks Arjun Talpallikar for his contributions to an early draft of this paper.

Dynamic pricing algorithms, which utilize real-time data to rapidly update prices, have fundamentally altered the landscape of business operations.[3] Enabled by the emergence of digital market platforms and the ubiquity of data-driven analytics, these sophisticated algorithms equip firms with a competitive edge in the marketplace.[4] Despite their inherent benefits, pricing algorithms harbor the potential to foster anticompetitive conduct, posing detrimental consequences for market competition as a whole.

For example, RealPage's YieldStar software, a "software that uses a mysterious algorithm to help landlords push the highest possible rents on tenants," has recently come under scrutiny.[5] An investigation by ProPublica found that YieldStar was being used to manage rent for nearly 19.7 million rental units in 2020, with the software relying on private information about nearby competitors' rental rates to make anonymized recommendations to its users.[6] This has engendered apprehension regarding the potentially anticompetitive ramifications stemming from software that facilitates indirect price coordination.[7]

The implications of dynamic pricing algorithms for market competition are not yet fully understood or adequately addressed by existing regulatory frameworks.[8] Traditional antitrust laws and regulations may not be sufficient to address the unique challenges presented by these pricing algorithms.[9] I propose the development of a simulation environment capable of assessing the behavior of pricing algorithms and their response to market conditions.

The proposed simulation environment would use conventional metrics from the antitrust literature to evaluate the potential for anticompetitive behavior. By analyzing the behavior of pricing algorithms within a controlled and simulated environment, regulators would be better equipped to understand the implications of these algorithms for market competition. Furthermore, such an approach could facilitate the development of more tailored and effective regulatory interventions, aimed at preventing anticompetitive behavior without unnecessarily stifling innovation.

As a proof-of-concept for this approach, I present a prototype sandbox environment. The sandbox environment enables researchers to simulate a variety of pricing algorithm scenarios, including scenarios that might lead to anticompetitive behavior. By running these simulations, researchers can identify potential antitrust concerns and assess the efficacy of potential regulatory interventions.

---

[3] *See* Ariel Ezrachi & Maurice E. Stucke, *Artificial Intelligence & Collusion: When Computers Inhibit Competition*, 2017 U. ILL. L. REV. 1775, 1780 (2017) ("Pricing algorithms dominate online sales of goods . . . and are widely used in hotel booking, and the travel, retail, sport, and entertainment industries.").
[4] *See* Michal S. Gal, *Algorithms as Illegal Agreements*, 34 BERKELEY TECH L.J. 67, 79 (2019).
[5] Heather Vogell, *Rent Going Up? One Company's Algorithm Could Be Why*, PROPUBLICA (Oct. 15, 2022, 5:00 AM EDT), https://www.propublica.org/article/yieldstar-rent-increase-realpage-rent.
[6] *Id.*
[7] *See id.*
[8] *See* Ariel Ezrachi & Maurice E. Stucke, *Sustainable and Unchallenged Algorithmic Tacit Collusion*, 17 NW. J. TECH. & INTELL. PROP. 217, 221 (2020).
[9] *Id.*

I. Models of Algorithmic Collusion

There are several models by which pricing algorithms can facilitate collusion. One way is through a hub-and-spoke model in which a single vendor provides pricing algorithms to multiple competitors who do not interact directly with each other. The vendor may collect and analyze data from all the competitors to determine optimal pricing strategies, leading to anticompetitive outcomes. Under this model, the vendor acts as the hub, providing pricing algorithms to every competitor, which act as the spokes. The vendor collects data from each spoke, including sales data, pricing data, and other relevant information. Using this data, the vendor can identify market trends and determine optimal pricing strategies for each competitor. This approach parallels the tactics allegedly employed within the meat processing industry wherein independent companies utilize data from the analytics firm Agri Stats to orchestrate pricing across the market.[10] The result is a coordinated pricing strategy that reduces competition and harms consumers.

Another scenario where pricing algorithms can lead to anticompetitive outcomes is when they are designed to use data derived from other competitors. In this scenario, pricing algorithms collect and analyze data from competitors in the market, including prices, sales volumes, and other relevant information. The algorithms then use this data to determine optimal pricing strategies that can increase profits and reduce competition. One well-known example of this type of algorithmic collusion occurred on Amazon's marketplace.[11] Two booksellers used pricing algorithms that continually increased the price of an out-of-print book until it was priced at $23.6 million.[12] The algorithms were designed to use data from each other to determine optimal prices, leading to a pricing spiral that resulted in the book's astronomical price.[13] Although the manifestly preposterous price that resulted for the book served as an exemplary demonstration of this occurrence, intricate algorithms would yield more nuanced collusive consequences that remain less discernible.[14]

Finally, pricing algorithms can also lead to anticompetitive outcomes even in the absence of explicit communication. This can occur when sophisticated algorithms are trained on significant amounts of data and optimize prices in an opaque way, leading to collusive outcomes without being explicitly instructed or designed to operate collusively.[15] This behavior can be particularly challenging to detect and prevent. Moreover, the complexity of these algorithms can

---

[10] Eli Hoff, 'Is This Legal?': Why an Obscure Data Service Has Been Sued Nearly 100 Times for Facilitating Anti-competitive Behavior, INVESTIGATE MIDWEST (July 29, 2021), https://investigatemidwest.org/2021/07/29/is-this-legal-why-an-obscure-data-service-has-been-sued-nearly-100-times-for-facilitating-anti-competitive-behavior/.
[11] Olivia Solon, *How a Book About Flies Came to Be Priced $24 Million on Amazon*, WIRED (Apr. 27, 2011, 3:35 PM), https://www.wired.com/2011/04/amazon-flies-24-million.
[12] *Id.*
[13] *Id.*
[14] *See* Ezrachi & Stucke, *supra* note 6, at 218.
[15] *See* Stephanie Assad, Robert Clark, Daniel Ershov & Lei Xu, *Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market*, QUEEN'S UNIV.: DEPT. OF ECON.S (Aug. 2020), https://www.econ.queensu.ca/sites/econ.queensu.ca/files/wpaper/qed_wp_1438.pdf (Queen's Econ. Dept., Working Paper No. 1438, 2020).

make it difficult for regulators to identify collusive behavior, as they may not be able to fully understand how the algorithms arrive at their pricing decisions.

II. Algorithmic Collusion as Tacit Collusion

In the absence of formal agreements, interdependent actions that soften competition are categorized as tacit collusion. Given that algorithmic collusion lacks the explicit agreements of prototypical cartel behavior, it is likely to be classified as tacit collusion under conventional doctrine.

Tacit collusion, although recognized for its detrimental effects on competition, is generally deemed to be permissible under antitrust laws because it lacks the "agreement" that is so critical in Section 1 of the Sherman Act. Furthermore, tacit collusion is difficult to identify and measure. We expect skilled managers to react and respond to rivals' actions as they try to position themselves in the marketplace.[16] Empirical evidence such as price parallelism can often be attributed to factors other than tacit collusion, such as changes in costs or demand.[17]

For example, a decrease in consumer demand for a product can lead to firms in the market reducing their output to avoid a surplus of inventory. This reduction in output may be observed across all firms in the market without being caused by coordinated behavior. Similarly, a rise in input costs may precipitate a rise in price for all the firms in a given market as they accommodate the altered cost structure. For example, if lumber prices were to rise, home builders are very likely to independently raise prices irrespective of their awareness of competitors' actions.

On the other hand, parallel behavior among firm in a given market may be evidence of tacit collusion and yet still fall outside the scope of antitrust liability. This is because even when conduct among firms can be attributed to consciously parallel decisions, condemning those decisions is not always appropriate. Asking market actors to disregard known information regarding pricing strategies of their competitors would be to ask them to behave irrationally, as such knowledge can be a critical factor in making strategic business decisions. This notion was elegantly articulated by Donald Turner:

> In a significant sense, the behavior of the rational oligopolist in setting his price is precisely the same as that of the rational seller in an industry consisting of a very large number of competitors. . . . The rational oligopolist simply takes one more factor into account—the reactions of his competitors to any price change that he makes. . . . [H]e is simply taking another factor into account, which he has to take into account because the situation in which he finds himself puts it there.[18]

---

[16] *In re* Text Messaging Antitrust Litig., 782 F.3d 867, 872 (7th Cir. 2015) ("Express collusion violates antitrust law; tacit collusion does not.").
[17] William E. Kovacic, Robert C. Marshall, Leslie M. Marx & Halbert L. White, *Plus Factors and Agreement in Antitrust Law*, 110 MICH. L. REV. 393, 395 (2011).
[18] Donald F. Turner, *The Definition of Agreement Under the Sherman Act: Conscious Parallelism and Refusals to Deal*, 75 HARV. L. REV. 655, 665-66 (1962).

III. The Dangers of Algorithmic Collusion

While algorithmic collusion and tacit collusion are similar in that both are forms of coordination without explicit agreement, there are several important distinctions between the two that render algorithmic collusion especially dangerous.

One of the key distinctions is that pricing algorithms can analyze vast amounts of data quickly and accurately, allowing businesses to make pricing decisions on a much more frequent basis. Oftentimes, this means prices can be adjusted based on real-time market conditions, such as changes in supply and demand. These algorithms can also consider data on individual customers, such as their purchase history, preferences, and behavior, and use this information to set prices tailored to each customer. This power, speed, and tailoring potentially makes the algorithm much more effective than human price-setting.

In tacit collusion, patience and trust are important considerations in maintaining a collusive equilibrium. For example, in a duopolistic market, two firms may tacitly behave to maintain supercompetitive prices. Both firms must be patient and willing to forego short-term gains to maintain the collusive equilibrium over the long term. If at any time firm A decides to undercut firm B, firm A will be able to make more sales and be rewarded with outsized profits for the short amount of time before firm B is able to retaliate by lowering its own price. Forbearing from engaging in such undercutting requires a high degree of trust among the firms, as each firm must believe that the other firms will also be patient and maintain the agreed-upon behaviors.[19] When humans are making these decisions, the level of trust they have in each other may limit their ability to cooperate. Human emotions may overwhelm the ability to choose the right strategy. Managers may also lack "patience," in the sense that they do not value the future, perhaps because they will be retired by then or their stock options will have been exercised.

Trust also limits the effects of conventional conscious parallelism in markets with many competitors. While a firm may be willing to trust one or two others to establish and maintain a supracompetitive equilibrium prices, this becomes much more difficult to do when there are many more actors involved.

Pricing algorithms, however, can largely eliminate the issues of patience and trust. These algorithms can be programmed to optimize profits over the long term and to not engage in short-term thinking or impatience. They can act as *homo economicus*—individuals of unbounded rationality - that are not subject to the cognitive shortcomings that afflict human decision-makers. As such, the proliferation of pricing algorithms will create situations where there can be pronounced anticompetitive effects absent explicit agreement, even in markets with many

[19] Marc Ivaldi, Bruno Jullien, Patrick Rey, Paul Seabright & Jean Tirole, *The Economics of Tacit Collusion: Final Report for DG Competition, European Commission* (IDEI, Working Paper No. 186, 2003), https://ec.europa.eu/competition/mergers/studies_reports/the_economics_of_tacit_collusion _en.pdf ("What is robust is that 'no collusion' is sustainable if firms are highly impatient (very small discount factor, δ close to zero) and that 'full collusion' (i.e., monopoly outcome) is sustainable when firms are very patient (large discount factor, δ close to 1).").

competitors. Even relatively unsophisticated pricing algorithms could be better equipped to tacitly collude than the human price-setters they replace.

The strategies utilized by pricing algorithms may involve subtle changes in pricing patterns that are not easily detectable, but which can communicate information to other algorithms about the pricing intentions of the firm. For example, an algorithm may adjust its prices in a specific pattern or timing that signals to other algorithms to follow suit or to hold off on making price adjustments. As such, these algorithms can reduce uncertainty and enable horizontal coordination between competitors. In fact, one study found that collusive outcomes from algorithmic pricing may be an inevitable outcome from their use absent regulation:

> To analyze the possible consequences, we study experimentally the behavior of algorithms powered by Artificial Intelligence (Q-learning) in a workhorse oligopoly model of repeated price competition. We find that the algorithms consistently learn to charge supracompetitive prices, without communicating with one another. The high prices are sustained by collusive strategies with a finite phase of punishment followed by a gradual return to cooperation. This finding is robust to asymmetries in cost or demand, changes in the number of players, and various forms of uncertainty.[20]

IV. The Shortcomings of Proposed Solutions

One solution that has been proposed to ameliorate the problem of algorithmic collusion is to prevent significant shares of a market from using the same algorithm to make pricing decisions. The idea is that if firms are using a different algorithm, they will not be able to effectively coordinate their pricing strategies, ameliorating the anticompetitive issues this paper has described. Such a solution is unlikely to work, however. First, liability could attach to a firm when its rival switched to the same algorithm; and sharing information as to which competitor is using which algorithm would likely create competition problems, not solve them. Additionally, even if every firm in a market is using a different algorithm, these algorithms will learn to raise prices to collusive levels if their objective is to maximize firm profits.[21]

As an alternate solution, Alexander Mackay and Samuel Weinstein have put forth a proposal advocating for the implementation of robust firewalls. This would ensure that algorithms can't depend on competitors' pricing information:

> In many cases, [a centralized algorithm] regulator would be tasked with rooting out pernicious racial and gender bias in algorithms. But such a regulator also could oversee pricing algorithms. In all these contexts, firms (and governmental agencies, in some cases) would submit their algorithms to the regulator for review. . . . In this context, reviewing pricing algorithms to determine if they are relying on competitors' prices would seem a relatively simple task, compared, for

[20] Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò & Sergio Pastorello, *Artificial Intelligence, Algorithmic Pricing, and Collusion*, 110 AM. ECON. REV. 3267, 3267 (2020).
[21] *See* Ezrachi & Stucke, *supra* note 6, at 250-55.

instance, to evaluating whether an algorithm produces biased results, especially if that bias is unintentional.[22]

What this approach doesn't account for, however, is that a pricing algorithm can very likely determine rivals pricing information even if it is not explicitly fed such data. There is a rich body of literature showing that "sensitive information"—in this case rival pricing information—can often be inferred by complex algorithms through other proxies.[23] In this case, a pricing algorithm could determine competitor prices through detailed data on consumer behavior and market conditions.

This limitation underscores the necessity of evaluating pricing algorithms' conduct in relation to fluctuating market circumstances, rather than solely scrutinizing the types of data incorporated into a pricing algorithm.

Clearly, finding an antitrust enforcement solution to tacit collusion by algorithm is going to be difficult. One intermediate step that would be helpful is for a regulator to know just how dangerous to competition any given algorithm is. Does the algorithm move industry prices to monopoly levels? Or is it just being efficient in responding to cost shocks for example? Having a measure of "collusiveness" would be really helpful as society and regulators develop optimal policies.

V. A Path Forward: Simulation Environments

We know that pricing algorithms are likely to create anticompetitive market conditions in novel ways that conventional competition law is not currently situated to tackle. Herein lies the crux of our quandary: how do we accurately determine the potential harm any given algorithm might pose? Is any given algorithm driving an industry towards monopoly pricing, or it improving efficiency, expertly responding to shifts in costs and market conditions?

In the absence of transparent evaluation mechanisms, we risk either misunderstanding the nature of algorithmic performance or failing to identify when it crosses the line from healthy competition to harmful collusion. Thus, the challenge lies not only in determining the potential harms but also in distinguishing between the outcomes of these computational processes.

While we have highlighted the anticompetitive consequences that can result from algorithmic pricing, these algorithms also open up a possibility for regulation that was not previously possible: the formalization of regulatory standards for pricing strategies. It is widely acknowledged that the opaque nature of machine learning tools can increase issues with

---

[22] Alexander MacKay & Samuel N. Weinstein, *Dynamic Pricing Algorithms, Consumer Harm, and Regulatory Response*, 100 WASH. U. L. REV. 111, 117 (2022).

[23] *See, e.g.*, Arvind Narayanan & Vitaly Shmatikov, *Robust De-Anonymization of Large Sparse Datasets*, *in* PROC. OF THE 2008 IEEE SYMP. ON SECURITY AND PRIVACY 111; Yves-Alexandre de Montjoye, Cdsar A. Hidalgo, Michel Verleysen & Vincent D. Blondel, *Unique in a Crowd: The Privacy Bounds of Human Mobility*, 3 SCI. REP., Mar. 25, 2013; Will Knight, THE APPLE CARD DIDN'T 'SEE' GENDER—AND THAT'S THE PROBLEM, Wired (Nov. 19, 2019, 9:15 AM), https://www.wired.com/story/the-apple-card-didnt-see-genderand-thats-the-problem.

accountability.[24] Unlike the dispersed human judgments that have typically comprised pricing strategies, however, algorithmic pricing tools are particularly well-suited for testing and experimentation, thereby providing a "focal point" for regulation.[25] This means that pricing strategies do not need to be evaluated post hoc but can be examined before they are implemented in real market settings.

In this Section, I aim to provide a regulatory framework for pricing algorithms that leverages this property. As a solution, we propose a simulation framework that can be used as a sandbox environment to test out pricing algorithms. Through a simulation environment, a pricing algorithm could be assessed in varying market conditions. By examining the pricing algorithms in the context of these varying conditions, a regulator could determine the extent to which an algorithm raises prices above supracompetitive levels in a variety of contexts.

In the context of regulating pricing algorithms, the use of simulation environments can provide significant benefits over traditional methods of monitoring and oversight. In a simulation environment, regulators can create a controlled environment where they can observe how an algorithm behaves under various market conditions. This eliminates the need for regulators to have detailed knowledge of the precise market conditions that the algorithm will be operating in.

In real-world markets, it can be challenging to determine characteristics of supply and demand such as the elasticity of demand marginal cost of production. While surveys and other tools can provide some insight into consumer behavior, these methods have serious external validity limitations. Additionally, determining supply characteristics can be even more difficult, as regulators may not have access to detailed information about production costs from firms. By contrast, in a simulation environment, regulators can directly choose and specify the characteristics of supply and demand that the algorithm will be operating under. This allows them to observe how the algorithm responds to different scenarios and make informed decisions about its performance. Instead of extrapolating determinations of collusion based on inferences of real market conditions, an oftentimes unwieldy task, regulators can directly stipulate the market conditions in a simulation environment, thereby removing factors that impute significant uncertainty into their tasks.

Additionally, because the environment is controlled, regulators can repeat experiments and adjust parameters to refine their understanding of how the algorithm operates. For instance, a regulator could specify both demand and supply characteristics, and vary these properties to learn how the algorithm responds. Parameters that could be varied include distributions of firm marginal costs, own- and cross-elasticities of demand, and shocks to marginal costs. This control does not preclude the regulator from using evidence to determine characteristics of the market and choosing default parameter values that match those of the market that the algorithm would

[24] Jenna Burrell, *How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms*, BIG DATA & SOC'Y, Jan.–June 2016.
[25] David Freeman Engstrom, Daniel E. Ho, Catherine M. Sharkey & Mariano-Florentino Cuéllar, *Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies* 28 (2020).

be deployed in. This can allow for more representative tailoring of parameters within the simulation environment to reflect realistic conditions.

In a simulated environment, complete information about market conditions also allows for welfare and competition analysis with a high degree of precision. A regulator can determine consumer surplus, producer surplus, and deadweight loss resulting from the use of a pricing algorithm results in, calculations that are often contestable and speculative in real market environments. These calculations could be used in litigation to demonstrate harms from a pricing algorithm. An even more proactive approach would be to set standards for a maximum level of deviation from competitive price. The FTC could promulgate such standards through notice-and-comment rulemaking pursuant to its Section 5 authority to prevent unfair or deceptive acts or practices in commerce.

In assessing pricing algorithms in these simulation environments, we draw on a long literature in empirical IO that nests models of monopoly, static Nash, and all the outcomes in between. The key outcome is a parameter (θ) that indicates how close the market outcome is to the monopoly price[26]. In this paper I follow the particular model of Nathan Miller and Matthew Weinberg in their analysis of retail beer prices in response to the MillerCoors joint venture.[27] The econometrician calculates the conduct parameter θ that measures deviation from one-shot Nash-Bertrand pricing equilibrium. In the classic method, this conduct parameter would equal zero under Nash-Bertrand competition and would equal one under fully collusive pricing. Thus, the magnitude of θ quantifies how far the market has moved towards the monopoly outcome.

VI. Shifting from Process to Outcomes

s

By turning to the regulatory system proposed in this article, regulators could bring algorithmic collusion under the umbrella of the consumer welfare standard. Rather than assessing the process by which collusion occurs, algorithmic collusion can be analyzed directly for its impact on consumer welfare.

In analyzing cases of coordination, antitrust law typically focuses on the process by which the anticompetitive conduct came about: the agreement itself. Often, the actual effects on consumers are not a determinative factor, at least as to whether the conduct will be condemned as being illegal. For instance, the Supreme Court has described tacit collusion in the following way:

> Tacit collusion, sometimes called oligopolistic price coordination or conscious parallelism, describes the *process*, not in itself unlawful, by which firms in a concentrated market might in effect share monopoly power, setting their prices at

[26] Timothy F. Bresnahan, *Competition and Collusion in the American Automobile Industry: The 1955 Price War*, 35 J. INDUS. ECON. 457 (1987).

[27] Nathan H. Miller & Matthew C. Weinberg, *Understanding the Price Effects of the MillerCoors Joint Venture*, 85 ECONOMETRICA 1763 (2017); *see also* Bresnahan, *Competition and Collusion*, *supra* note 25.

> a profit-maximizing, supracompetitive level by recognizing their shared economic interests and their interdependence with respect to price and output decisions.[28]

Similarly, Jonathan Baker has noted that "[n]ot every parallel pricing outcome constitutes an agreement because not every such outcome was reached through the process to which the law objects: a negotiation that concludes when the firms convey mutual assurances that the understanding they reached will be carried out."[29] In the context of algorithmic pricing, if we cannot condemn supracompetitive outcomes based on behavior, we need a method such as the one I proposed here to measure the outcomes themselves.

This emphasis on the process of collusion has been understandable within these contexts given the problems that arise from attempting to condemn conscious parallelism, described above. However, this focus on the process of collusion is out of step with the consumer welfare standard that is underpins much of antitrust law. The consumer welfare standard is based on the idea that antitrust law should protect consumer welfare by promoting competition and preventing anticompetitive behavior. This involves looking at the outcomes of competitive processes in order to determine whether they are likely to benefit consumers. As such authorities could effectively situate algorithmic collusion under the purview of the consumer welfare standard through by embracing the proposed regulatory framework. In lieu of scrutinizing the mechanisms through which collusion transpires, algorithmic collusion can be appraised explicitly for its influence on consumer welfare.

VII. Conclusion

While dynamic pricing algorithms offer many benefits for firms, they can also facilitate anticompetitive behavior. It is essential for regulators to have a better understanding of how pricing algorithms work and the impact they have on market competition to ensure fair and competitive markets. The proposed simulation environment could be a useful tool for regulators to evaluate the potential for anticompetitive behavior and develop effective policies to prevent it. Pricing algorithms would need to go through an approval process where they are executed in a regulatory testbed. This testbed environment would allow regulators to manipulate and analyze factors that are hidden in real market environments, such as demand, marginal revenue, and marginal cost.

---

[28] Brooke Grp. Ltd. v. Brown & Williamson Tobacco Corp., 509 U.S. 209, 227 (1993) (emphasis added).

[29] Jonathan B. Baker, *Two Sherman Act Section 1 Dilemmas: Parallel Pricing, the Oligopoly Problem, and Contemporary Economic Theory*, 38 ANTITRUST BULL. 143, 179 (1993).