

POLICY FORUM

TECHNOLOGY AND LAW

Protecting consumers from collusive prices due to AI

Price-setting algorithms can lead to noncompetitive prices, but the law is ill equipped to stop it

By **Emilio Calvano**^{1,2}, **Giacomo Calzolari**^{2,3,4}, **Vincenzo Denicolò**^{1,2}, **Joseph E. Harrington Jr.**⁵, **Sergio Pastorello**¹

The efficacy of a market system is rooted in competition. In striving to attract customers, firms are led to charge lower prices and deliver better products and services. Nothing more fundamentally undermines this process than collusion, when firms agree not to compete with one another and consequently consumers are harmed by higher prices. Collusion is generally condemned by economists and policy-makers and is unlawful in almost all countries. But the increasing delegation of price-setting to algorithms (1) has the potential for opening a back door through which firms could collude lawfully (2). Such algorithmic collusion can occur when artificial intelligence (AI) algorithms learn to adopt collusive pricing rules without human intervention, oversight, or even knowledge. This possibility poses a challenge for policy. To meet this challenge, we propose a direction for policy change and call for computer scientists, economists, and legal scholars to act in concert to operationalize the proposed change.

HUMAN COLLUSION

Collusion among humans typically involves three stages (see the table). First, firms' employees with price-setting authority communicate with the intent of agreeing on a collusive rule of conduct. This rule encompasses a higher price and an arrangement to incentivize firms to comply with that higher price rather than undercut it in order to

pick up more market share. For example, in 1995 the CEOs of Christie's and Sotheby's hatched their plans in a limo at Kennedy International Airport, and in 1994 the U.S. Federal Bureau of Investigation secretly taped the lysine cartel as they conspired in a Maui hotel room. At those meetings, they spoke about charging higher prices and how to enforce them. Second, successful communication results in the mutual adop-

The process that produces higher prices

	COMMUNICATIONS	COLLUSIVE PRICING RULES	HIGHER PRICES
Humans	Present, discoverable	Latent, not discoverable	Observable, difficult to evaluate
Algorithms	Not present	Latent, discoverable	Observable, difficult to evaluate

tion of a collusive rule of conduct, which commonly takes the form of a collusive pricing rule. A crucial component of this pricing rule is retaliatory pricing: Each firm raises its price and maintains that higher price under the threat of a "punishment," such as a temporary price war, should it cheat and deviate from the higher price (3). It is this threat that sustains higher prices than would arise under competition. Third, firms set the higher prices that are the consequence of having adopted those collusive pricing rules.

To determine whether firms are colluding, one could look for evidence at any of the three stages. However, evidence related to the last two stages—pricing rules and higher prices—is generally regarded as insufficient to achieve the requisite level of confidence in the judicial realm. Economists know how to calculate competitive prices given demand, costs, and other relevant market conditions. But many of these factors are difficult to observe and, when observable, are challenging to measure with precision. Consequently, courts do not use the competitive price level as a benchmark

to identify collusion. Likewise, it is difficult to assess whether the firms' rules of conduct are collusive because such rules are latent, residing in employees' heads. In practice, we may never observe the retaliatory lower prices from a firm that cheated, even though that response is there in the minds of the employees and it is the anticipation of such a response that sustains higher prices. In other words, we might lack the events that produce the data that could identify the collusive pricing rules. Furthermore, even if one could observe what looks like a price war, it would be difficult to rule out innocent explanations (such as a decrease in the firms' costs or a fall in demand).

Given the latency of collusive pricing rules and the difficulty of determining whether prices are collusive or competitive, antitrust law and its enforcement have focused on the first stage: communications. Firms are found to be in violation of the law when communications (perhaps supplemented by other evidence) are sufficient to establish that firms have a "meeting of minds," a "concurrence of wills," or a "conscious commitment" that they will not compete (4). In the United States, more specifically, there must be evidence that one firm invited a competitor to collude and that the competitor accepted that invitation. The risk of false positives (i.e., wrongly finding firms guilty of collusion) has led courts to avoid basing their judgments on evidence of collusive pricing rules or collusive prices and instead to rely on evidence of communications.

ALGORITHMIC COLLUSION

Although the use of pricing algorithms has a long history—airline companies, for instance, have been using revenue management software for decades—concerns regarding algorithmic collusion have only recently arisen for two reasons. First, pricing algorithms had once been based on pricing rules set by programmers but now often rely on AI systems that learn autonomously through active experimentation. After the programmer has set a goal, such as profit maximization, algorithms are capable of autonomously learning rules of conduct that achieve the goal, possibly with no human intervention. The enhanced sophistication of learning algorithms makes it more likely that AI systems will discover profit-enhancing collusive pricing rules, just as they have succeeded in discovering winning strategies in complex board games such as chess and Go (5).

Second, a feature of online markets is that competitors' prices are available to a firm in real time. Such information is es-

¹Department of Economics, University of Bologna, Bologna, Italy. ²Center for Economic Policy Research, London, UK.

³Department of Economics, European University Institute, Florence, Italy. ⁴Department of Statistics, University of Bologna, Bologna, Italy. ⁵Department of Business Economics and Public Policy, Wharton School, University of Pennsylvania, Philadelphia, PA, USA. Email: vincenzo.denicolò@unibo.it

sential to the operation of collusive pricing rules. In order for firms to settle on some common higher price, firms' prices must be observed frequently enough because sustaining those higher prices requires the prospect of punishing a firm that deviates from the collusive agreement. The more quickly the punishment is meted out, the less temptation to cheat. Thus, the emergence and persistence of higher prices through collusion is facilitated by rapid detection of competitors' prices, which is now often possible in online markets. For example, the prices of products listed on Amazon may change several times per day but can be monitored with practically no delay.

In light of these developments, concerns regarding the possibility of algorithmic collusion have been raised by government authorities, including the U.S. Federal Trade Commission (FTC) (6) and the European Commission (7). These concerns are justified, as enough evidence has accumulated that autonomous algorithmic collusion is a real risk.

The evidence is both experimental and empirical. On the experimental side, recent research has found the spontaneous emergence of collusion in computer-simulated markets. In these studies, commonly used reinforcement-learning algorithms learned to initiate and sustain collusion in the context of well-accepted economic models of an industry (8, 9) (see the figure). Collusion arose with no human intervention other than instructing the AI-enabled learning algorithm to maximize profit (i.e., algorithms were not programmed to collude). Although the extent to which prices were higher in such virtual markets varied, prices were almost always substantially above the competitive level.

On the empirical side, a recent study (10) has provided possible evidence of algorithmic collusion in Germany's retail gasoline markets. The delegation of pricing to algorithms was found to be associated with a substantial 20 to 30% increase in the markup of stations' prices over cost. Although the evidence is indirect—because the authors of the study could not directly observe the timing of adoption of the pricing algorithms and thus had to infer it from other data—their findings are consistent with the results of computer-simulated market experiments.

A NEW POLICY APPROACH

Algorithmic collusion is as bad as human collusion. Consumers are harmed by the higher prices, irrespective of how firms arrive at charging these prices. However, should algorithmic collusion emerge in a

“...pricing algorithms had once been based on pricing rules set by programmers but now often rely on AI systems that learn autonomously through active experimentation.”

market and be discovered, society lacks an effective defense to stop it. This is because algorithmic collusion does not involve the communications that have been the route to proving unlawful collusion (as distinguished from instances in which firms' employees might communicate and then collude with the assistance of algorithms, as in a recent case involving poster sellers on Amazon Marketplace). And even if alternative evidentiary approaches were to arise, there is no liability unless courts are prepared to conclude that AI has a “mind” or a “will” or is “conscious,” for otherwise there can be no “meeting of minds” with algorithmic collusion.

As a result, if algorithmic collusion occurs and is discovered by the authori-

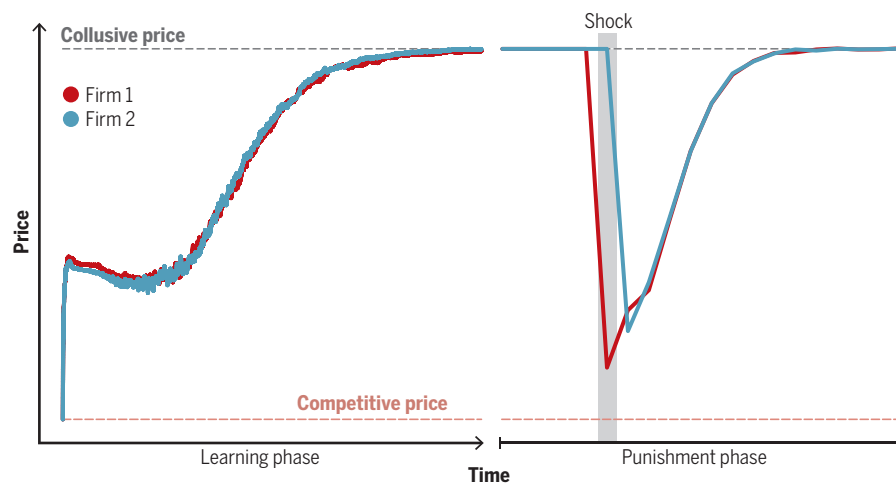
ties, currently it cannot be considered a violation of antitrust or competition law. Society would then have no recourse and consumers would be forced to continue to suffer the harm from algorithmic collusion's higher prices.

There is an alternative path, which is to target the collusive pricing rules learned by the algorithms that result in higher prices (11). These latent rules of conduct may be uncovered when they have been adopted by algorithms. Whereas a court cannot get inside the head of an employee to determine why prices are what they are, firms' pricing algorithms can be audited and tested in controlled environments. One can then simulate all sorts of possible deviations from existing prices and observe the algorithms' reaction in the absence of any confounding factor. In principle, the latent pricing rules can thus be identified precisely.

This approach was successfully used by researchers in (8) to verify that the pricing algorithms have indeed learned the collusive property of reward (keeping prices high unless a price cut occurs) and punishment (through retaliatory price wars should a price cut occur). To show this, the researchers momentarily overrode the pricing algorithm of one firm, forcing it to set a lower price. As soon as the algorithms regained control of the pricing, they engaged in a temporary price war, where lower prices were charged but then gradually returned to the collusive level. Having learned that undercutting the other firm's price brings forth a price war

Collusive pricing rules uncovered

After the two algorithms have found their way to collusive prices (“learning phase,” left side), an attempt to cheat so as to gain market share is simulated by exogenously forcing Firm 1's algorithm to cut its price (“punishment phase,” right side). From the “shock” period onward, the algorithm regains control of the pricing. Firm 1's deviation is punished by the other algorithm, so firms enter into a price war that lasts for several periods and then gradually ends as the algorithms return to pricing at a collusive level. For better graphical representation, the time scales on the right and left sides of the figure are different.



(with the associated lower profits), the algorithms evolved to maintain high prices (see the figure).

It may seem paradoxical that collusion can be identified by the low retaliatory prices, which could be close to the competitive level, rather than by the high prices that are the ultimate concern for policy. But there are two important differences between retaliatory price wars and healthy competition. First, in the absence of the low-price perturbation, the price war remains hypothetical in that it is a threat that is not executed. Second, the price war shown in the figure is only temporary: Instead of permanently reverting to the competitive price level, the algorithms gradually return to the pre-shock prices. This is evidence that the price war is there to support high prices, not to produce low prices.

Focusing on the collusive pricing rules is the key to identifying, preventing, and prosecuting algorithmic collusion (see the table). Policy cannot target the higher prices directly, nor can it target communications as they may not be present (unlike with human collusion). But the retaliatory pricing rules may now be observable, as firms' pricing algorithms can be audited and tested. We therefore propose that antitrust policy shift its focus from communications (with humans) to rules of conduct (with algorithms).

Making the proposed change operational involves a broad research program that requires the combined efforts of economists, computer scientists, and legal scholars. One strand of this program is a three-step experimental procedure. The first step creates collusion in the lab for descriptively realistic models of markets. As the competitive price would be known by the experimenter, collusion is identified by high prices. Having identified an episode of collusion, the second step is to perform a post hoc auditing exercise to uncover the properties of the collusive pricing rules that produced those high prices.

Some progress has been made on the identification of collusive rules of conduct adopted by algorithms, but much more work needs to be done. Economics provides several properties to watch out for. Of course, there is the retaliatory price war discussed above, which is what existing research has focused on (8, 9). Another property is price matching, whereby firms' prices move in sync: one firm changing its price and the other firm subsequently matching that change. Price matching has

been documented for human collusion in various markets, but we do not yet know whether algorithms are capable of learning it. A third property is the asymmetry of price responses. When firms collude, they typically respond to a competitor's price cut more strongly—as part of a punishment—than to a price increase. No such asymmetry is to be expected when firms compete.

“Another path is to make firms legally responsible for the pricing rules that their learning algorithms adopt. Firms may then be incentivized to prevent collusion by routinely monitoring the output of their learning algorithms.”

The aforementioned properties are based on economic theory and studies of human collusion. Learning algorithms may devise rules of conduct that neither economists nor managers have imagined (just as learning algorithms have done, for instance, in chess). To investigate this possibility, computer scientists might develop algorithms that explain their own behavior, thereby making the collusive properties more apparent. One way of doing so is to add a second module to the reinforcement-learning module that maximizes profits; this second module maps the state representation of the first one onto a verbal explanation of its strategy (12).

Having uncovered collusive pricing rules, the third step is to experiment with constraining the learning algorithm to prevent it from evolving to collusion. Computer scientists are particularly valuable here, given that they are involved in similar tasks such as trying to constrain algorithms so that, for instance, they do not exhibit racial and gender bias (13).

Once the capacities to audit pricing algorithms for collusive properties and to constrain learning algorithms so that they do not adopt collusive pricing rules have been developed, legal scholars are called upon to use that knowledge for purposes of prosecution and prevention. One route is to make certain pricing algorithms unlawful, perhaps under Section 5 of the FTC Act, which prohibits unfair methods of competition. In the area of securities law, the 2017 case *U.S. v. Michael Coscia* made illegal the use of certain programmed trading rules and thus provides a legal precedent for prohibiting algorithms. Another path is to make firms legally responsible

for the pricing rules that their learning algorithms adopt (14). Firms may then be incentivized to prevent collusion by routinely monitoring the output of their learning algorithms.

These are some of the avenues that can be pursued for preventing and shutting down algorithmic collusion. There are several obstacles down the road, including the

difficulty of making a collusive property test operational, the lack of transparency and interpretability of algorithms, and courts' willingness and ability to incorporate technical material of this nature. In addition, there is the challenge of addressing algorithmic collusion without giving up the efficiency gains from pricing algorithms such as the quicker response to

changing market conditions. As authorities prepare to take action (15), it is vital that computer scientists, economists, and legal scholars work together to protect consumers from the potential harm of higher prices. ■

REFERENCES AND NOTES

1. A. Ezrachi, M. Stucke, *Virtual Competition: The Promise and Perils of the Algorithm-Driven Economy* (Harvard Univ. Press, 2016).
2. S. Mehra, *Minn. Law Rev.* **100**, 1323 (2016).
3. J. Harrington, *The Theory of Collusion and Competition Policy* (MIT Press, 2017).
4. L. Kaplow, *Competition Policy and Price Fixing* (Princeton Univ. Press, 2013).
5. D. Silver *et al.*, *Science* **362**, 1140 (2018).
6. “The Competition and Consumer Protection Issues of Algorithms, Artificial Intelligence, and Predictive Analytics,” Hearing on Competition and Consumer Protection in the 21st Century, U.S. Federal Trade Commission, 13–14 November 2018; www.ftc.gov/news-events/events-calendar/ftc-hearing-7-competition-consumer-protection-21st-century.
7. “Algorithms and Collusion—Note from the European Union,” OECD Roundtable, June 2017; www.oecd.org/competition/algorithms-and-collusion.htm.
8. E. Calvano, G. Calzolari, V. Denicolo, S. Pastorello, *Am. Econ. Rev.* **110**, 3267 (2020).
9. T. Klein, “Autonomous Algorithmic Collusion: Q-Learning Under Sequential Pricing,” Amsterdam Law School Research Paper 2018-15 (2019).
10. S. Assad, R. Clark, D. Ershov, L. Xu, “Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market,” CESifo Working Paper No. 8521 (2020).
11. J. Harrington, *J. Compet. Law Econ.* **14**, 331 (2018).
12. Z. C. Lipton, *ACM Queue* **16**, 30 (2018).
13. P. S. Thomas *et al.*, *Science* **366**, 999 (2019).
14. S. Chopra, L. White, *A Legal Theory for Autonomous Artificial Agents* (Univ. of Michigan Press, 2011).
15. European Commission, document Ares(2020)2877634.

ACKNOWLEDGMENTS

The paper benefited from detailed and insightful comments by three anonymous reviewers. All authors contributed equally. The authors declare no competing interests.

10.1126/science.abe3796



Protecting consumers from collusive prices due to AI

Emilio Calvano, Giacomo Calzolari, Vincenzo Denicol, Joseph E. Harrington, Jr., and Sergio Pastorello

Science, **370** (6520), .

DOI: 10.1126/science.abe3796

View the article online

<https://www.science.org/doi/10.1126/science.abe3796>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)

Science (ISSN 1095-9203) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science* is a registered trademark of AAAS.
Copyright © 2020, American Association for the Advancement of Science