ALGORITHMIC COLLUSION: A REAL PROBLEM FOR COMPETITION POLICY?

BY EMILIO CALVANO, GIACOMO CALZOLARI, VINCENZO DENICOLÒ & SERGIO PASTORELLO[1]

   

1 Emilio Calvano (Department of Economics, University of Bologna, CEPR), Giacomo Calzolari (Department of Economics, European University Institute, Florence, CEPR), Vincenzo Denicolò (Department of Economics, University of Bologna, CEPR), Sergio Pastorello (Department of Economics, University of Bologna).

Visit www.competitionpolicyinternational.com for access to these articles and more!

www.competitionpolicyinternational.com
Competition Policy International, Inc. 2020© Copying, reprinting, or distributing this article is forbidden by anyone other than the publisher or author.

I. INTRODUCTION

In the last few years, antitrust authorities have started to worry about the possible consequences of algorithmic pricing. While the agencies generally recognize that automated pricing can enhance economic efficiency by allowing quicker responses to changing market conditions, they are also concerned that algorithms might harm consumers.

There seem to be two main concerns. First, antitrust authorities have noted that algorithms can collect and process vast amounts of personal data and thus can engage in price discrimination on an unprecedented scale. Second, agencies worry about the possibility of algorithmic collusion.[2] Even though the two issues may be related (the scope for price discrimination is small in any case, if vigorous competition drives prices towards costs), in this article we focus on the latter.

Relatively little is known about algorithmic collusion, and thus reasonable people may disagree about whether algorithms can really collude, and whether such collusion, if it materializes, poses new challenges to competition policy. Some claim that collusion among pricing algorithms is inevitable and is destined to become prevalent as the diffusion of the software increases.[3] Others point to the paucity of cases brought so far and claim that these cases can be dealt with by means of the traditional tools of competition policy.[4]

It is the contention of this article that algorithmic collusion is a real risk, the seriousness of which is still difficult to assess, but that should not be dismissed lightly by antitrust agencies. The article also discusses the specific new challenges that algorithmic collusion poses.

The argument proceeds in several steps. First, we document that pricing algorithms are already widely used and argue that they are likely to become even more prevalent in the future. Second, we discuss various ways in which algorithms may facilitate collusion without creating any genuinely new antitrust issue. However, we also argue that pricing algorithms may learn to collude "autonomously" and without explicitly communicating with one another. Such form of collusion would pose a challenge to current policy, which typically targets only explicit agreements. By reviewing the existing experimental literature, we show that the risk of autonomous

2 For example, the possibility of algorithmic collusion has been extensively discussed at the 7th session of the FTC Hearings on competition and consumer protection (November 2018) and has been the subject of white papers issued by the Canadian Competition Bureau and the British Competition and Market Authority. See also the 2019 report for the European Commission by Crémer, Jacques, Yves-Alexandre de Montjoye, & Heike Schweitzer. "Competition policy for the digital era," European Commission. *Luxembourg: Publications Office of the European Union* (2019).

3 See for instance Ezrachi, Ariel, & Maurice E. Stucke. "Artificial intelligence & collusion: When computers inhibit competition." *U. Ill. L. Rev.* (2017): 1775.

4 See for instance Schwalbe, Ulrich. "Algorithms, machine learning, and collusion," *Journal of Competition Law & Economics* 14.4 (2018): 568-607 and Schrepel, Thibault. "The Fundamental Unimportance of Algorithmic Collusion for Antitrust Law," *Harvard Journal of Law and Technology* (2020).

algorithmic collusion is real and not just a remote theoretical possibility. By reviewing the existing empirical literature, we show evidence that autonomous algorithmic collusion may have already materialized in certain industries. Finally, we discuss the differences between algorithmic collusion and collusion among humans from the viewpoint of competition policy.

II. THE PREVALENCE OF ALGORITHMIC PRICING

Algorithmic pricing is not new. Airline companies and hotels, for instance, have used revenue management software for decades.[5] However, the use of software algorithms as pricing tools has boomed with the advent of online marketplaces. Studying Amazon, for instance, Chen et al. find that in a sample of over 1,600 best-selling items, more than a third of the sellers had already automated their pricing in 2015.[6] Since then, algorithmic pricing has become so widespread that Amazon has started to advertise its WebService API stressing that the API facilitates pricing automation.[7]

Firms increasingly resort to algorithmic pricing in more traditional sectors also, such as for instance gas stations (an industry that we shall discuss more extensively below).[8] A repricing software industry has indeed grown that supplies turnkey pricing systems to small firms and customizes software for the big ones.[9] As this industry keeps developing, and the software keeps improving, algorithmic pricing is likely to become even more prevalent in the future.

III. HOW PRICING ALGORITHMS CAN FACILITATE COLLUSION

There are several ways in which pricing algorithms can facilitate collusion. The early literature has stressed the fact that algorithms can change prices more frequently than humans.[10] As a result, deviations from collusive agreements can be punished more promptly, reducing the temptation to cheat and thus stabilizing the cartel. As Hal Varian put it long ago, pricing algorithms might come to react to one another so quickly that buyers can make no purchases before a cheater's price cut is matched by its rivals.[11] In this limiting case, the short-run gain from deviations would vanish, and cartels would be perfectly deviation-proof.

Often, however, what limits collusion seems to be not the promptness of punishments but rather the difficulties of coordinating the behavior of heterogeneous players in complex economic environments, especially in the absence of explicit communication.[12] If this, and not the speed of punishments, is indeed the critical factor, then in most cases a higher frequency of oligopolistic interactions would not make collusion any more likely.

Another possibility is that algorithms may be used as a tool for implementing collusive agreements that are reached in traditional ways. The majority of the (few) antitrust cases that so far have involved pricing algorithms seems to fall in this category.[13] Even if the management of cartels is a delicate endeavor (price wars sometimes start by mistake, just as military wars do), the role of algorithms here seems mainly practical. In this respect, algorithms raise no new conceptual issue.

---

5 Revenue management software typically comprises an estimation module and an optimization module. The estimation module estimates a structural model of the industry (which is however often taken to be a monopoly, abstracting from strategic interactions). The optimization module then calculates the optimal prices. See Talluri, Kalyan T., & Garrett J. Van Ryzin. *The theory and practice of revenue management*. Vol. 68. Springer Science & Business Media, 2006 for a comprehensive analysis of revenue management systems.

6 See Chen, Le, Alan Mislove, & Christo Wilson. "An empirical analysis of algorithmic pricing on amazon marketplace." *Proceedings of the 25th International Conference on World Wide Web*. 2016.

7 See https://developer.amazonservices.com/, last accessed May 24, 2020.

8 See Schechner, Sam "Why do gas station prices constantly change? Blame the algorithms," The Wall Street Journal, May 8, 2017.

9 See for instance https://emerj.com/ai-sector-overviews/ai-for-pricing-comparing-5-current-applications/, last accessed May 24, 2020.

10 Ezrachi, Ariel, and Maurice E. Stucke. "Virtual competition" Journal of European Competition Law & Practice 7.9 (2016): 585–586.

11 See Hal Varian. "When commerce moves online, competition can work in strange ways." The New York Times, August 24, 2000.

12 Kuhn, K.-U. & S.Tadelis,. "The Economics of Algorithmic Pricing: Is collusion really inevitable?" (2018) mimeo.

13 For a discussion of these cases, see Gata, João E. "Controlling Algorithmic Collusion: short review of the literature, undecidability, and alternative approaches" (2019) mimeo.

www.competitionpolicyinternational.com
Competition Policy International, Inc. 2020© Copying, reprinting, or distributing
this article is forbidden by anyone other than the publisher or author.

Yet another possibility is that various competitors delegate their pricing decisions to the same software company, which would then come to jointly manage the pricing strategies of a group of firms. Acting as a common agent for its clients, that company may then coordinate their behavior instead of making them compete. For example, Decarolis & Rovigatti find evidence of this phenomenon in Google's advertising auctions.[14] But again, the role of algorithms in implementing this form of collusion seems limited. What should be blamed here is the hub-and-spoke cartel, rather than the algorithms in themselves.

Finally, algorithms may be directly instructed by the programmers to collude. For example, they may be so designed as to play a "tit-for-tat" strategy, or some other strategy conducive to cooperation rather than competition. Likewise, algorithms may be endowed with an objective function that places a positive weight on rivals' profits, which automatically lessens competition.

However, insofar as the software follows mechanically the instructions of the programmer, we are still on traditional ground. The firms remain effectively in charge of the strategic choices, and collusion may be achieved only by means of some sort of agreement (be it explicit or tacit) among the firms themselves. The only difference with standard cartels is that the agreement is made at the programming stage rather than at the pricing stage. For one thing, this could perhaps help to mask the cartel. For another, however, it could facilitate detection: since the collusive instructions must be incorporated into the software, they should be easier to verify.

IV. AUTONOMOUS ALGORITHMIC COLLUSION

Things have changed recently, however, with the advent of a new vintage of pricing algorithms, based on artificial intelligence and reinforcement learning. These new algorithms are radically different from their rule-based predecessors. Instead of following fixed rules, they learn from experience, engaging in active experimentation and adopting more frequently (and thus reinforcing) those strategies that performed best in the past. In this learning process, the algorithms are completely autonomous.[15] Of course, programmers must endow them with a criterion to evaluate past performance (i.e. an objective to maximize, such as the firm's profits). They must also enjoin the algorithms to explore new strategies. But apart from that, the algorithms learn to solve the pricing problem on their own, by trial and error.

Since collusion increases profits, it might happen that such "smart" reinforcement-learning algorithms learn to collude just because in their experience collusion happens to be better than competition. They need not be instructed, either directly or indirectly, to cooperate with rivals. The only instruction they may have received is to maximize profits, which is what they also ought to do to guarantee healthy competition.

From a purely economic viewpoint, this form of algorithmic collusion closely resembles tacit collusion among humans. However, it may be completely unpremeditated. Humans that collude, albeit tacitly, must be perfectly aware of what they are doing. Algorithmic collusion, by contrast, may be a totally unintended consequence of the algorithms' learning. This marks an important difference from an antitrust viewpoint. For example, it seems that inadvertent algorithmic collusion could hardly be regarded as a crime.

The problem, then, is whether such autonomous algorithmic collusion is a real possibility or not. This is not an easy question to answer; as said, little is known on this crucial issue. However, some progress has recently been made both in the experimental and in the empirical literature. This allows tentative hypotheses to be put forward. We shall now discuss each of these literatures in turn.

14 See Decarolis, F. & Rovigatti, G. "From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising" (2019), mimeo.

15 Such machine learning algorithms are increasingly used not only as pricing tools, but also as recommender systems. In this capacity, they increasingly determine, for instance, the music we listen to, the movies we watch, and the news and opinion pieces we read.

www.competitionpolicyinternational.com
Competition Policy International, Inc. 2020© Copying, reprinting, or distributing
this article is forbidden by anyone other than the publisher or author.

V. THE EXPERIMENTAL LITERATURE

Given the difficulties of predicting theoretically what competing pricing algorithms might do, both computer scientists and economists have resorted to an "experimental" approach. That is, scholars code reinforcement algorithms and then deploy them in computer-simulated marketplaces, where their behavior can be observed. Since reinforcement-learning algorithms by design must explore randomly, the experiment should ideally be repeated a large number of times in order to smooth out uncertainty.

In this framework, the emergence of some degree of uncompetitively high prices was documented by an early computer science literature already in the late 1990s.[16] For the most part, this and the subsequent computer science literature has adopted a framework derived from Maskin & Tirole's model of staggered pricing.[17] In this model, there are two firms that alternate in moving. When its turn comes, a firm sets a price to which it commits for two periods. Each firm can condition its price only on rival's current price, and not on past prices.[18]

The best executed paper in this line of research is probably Klein (2019), who focuses on one of the most popular classes of reinforcement learning algorithms, namely, Q-learning.[19] In accordance with Maskin & Tirole's theory, Klein finds that two Q-learning algorithms converge either to a constant, supra-competitive price or to a pattern that resembles an Edgeworth cycle. In his baseline case with 6 possible price levels, profits are roughly mid-way between the competitive and the monopoly level. Enlarging the action set to 12 or 100 possible prices increases the average price and makes it more likely that the algorithms converge to a price cycle rather than to a constant price.

While this framework of staggered pricing is interesting, it is not the canonical model of collusion. Most often, economists model collusion by means of a repeated game where firms price simultaneously in each stage and may condition their current prices on the entirety of past events.

This environment is more challenging for Q-learning algorithms, as the collection of variables which the current price may depend on (technically, the state space) is much bigger. In a recent article, however, we have analyzed the behavior of Q-learning pricing algorithms precisely in this more standard (but, for the algorithms, more exacting) economic framework.[20] We have shown that even in this case, relatively simple reinforcement learning algorithms systematically learn to play sophisticated collusive strategies. These strategies involve punishments of deviations, i.e. price wars that start whenever a firm tries to cheat. These price wars are crucial to stabilize the implicit collusive agreement. They are finite in duration, however, with a gradual return to the original, high prices.

Our algorithms learn these complex punishment strategies purely by trial and error. They are not designed or instructed to collude, they do not communicate with one another, and they have no prior knowledge of the environment in which they operate. What is worse, their propensity to collude appears to be stubborn. In keeping with theory, the degree of collusion decreases as the number of competitors increases, but substantial collusion continues to prevail even when the sellers are three or four in number, when they are asymmetric, and when they operate in stochastic environments.

This experimental evidence is certainly not yet conclusive, and in the paper cited above we point out various limits of the results currently available. However, these results suggest that autonomous algorithmic collusion is a real possibility.

---

16 Greenwald, Amy R., Jeffrey O. Kephart, & Gerald J. Tesauro. "Strategic pricebot dynamics." *Proceedings of the 1st ACM conference on Electronic commerce*. 1999.

17 Maskin, Eric, & Jean Tirole. "A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs." *Econometrica* (1988): 549-569.

18 Technically speaking, the model is "Markovian."

19 Klein, Timo. "Autonomous algorithmic collusion: Q-learning under sequential pricing." *Amsterdam Law School Research Paper* 2018-15 (2019): 2018-05. This paper also provides a critical review of the previous literature in this line of research.

20 Calvano, E., Calzolari, G., Denicolo, V., & Pastorello, S. (2019). "Artificial intelligence, algorithmic pricing and collusion" (2019), mimeo.

www.competitionpolicyinternational.com
Competition Policy International, Inc. 2020© Copying, reprinting, or distributing
this article is forbidden by anyone other than the publisher or author.

VI. THE EMPIRICAL LITERATURE

To the best of our knowledge, so far only one paper has looked at the impact of pricing algorithms on competition empirically: Assad et al. (2020).[21] This paper exploits, as a sort of natural experiment, the extended adoption of algorithmic pricing that occurred in 2017 in Germany's retail gasoline market. In assessing the impact of this event, the authors take advantage of a comprehensive high-frequency database on retail prices available from governmental agency GMTU (Germany Market Transparency Unit).

One preliminary problem the authors face is that firms generally do not disclose whether they use software algorithms in their pricing process or not. To circumvent this problem, the authors identify the adoption of algorithmic pricing by looking for structural breaks in the stations' pricing behavior. They submit that a switch to algorithmic pricing increases the number of price changes and reduces the reaction time to changes in the prices of rivals (i.e. nearby stations). Based on this reasonable assumption, they infer that a switch to algorithmic pricing has occurred when such structural changes are observed.

Having thus empirically identified the adoption of pricing algorithms, the paper then looks at its impact on market outcomes. Here, the authors face another problem, namely, the possible endogeneity of the switch to algorithmic pricing. That is, since adoption decisions are made by station managers, they may be triggered by changes in local market conditions that confound the effects of the adoption itself. To deal with this issue, the authors use brand-level adoption of pricing algorithms at the national level as an instrument for station-level adoption. They then use a "diff-in-diff" approach to compare the price levels in local markets where pricing algorithms have or have not been deployed.

Proceeding in this way, the authors find that in local markets where only one station has switched to algorithmic pricing, prices do not exhibit any significant change. If anything, prices in fact go down, with the sole adopter charging the lowest price. However, things change when at least one nearby competing station also switches to algorithmic pricing. In this case, prices show a significant increase. The estimated effect is sizeable, amounting to about 20-30 percent of the stations' average margin.

Since Assad et al. (2020) can only observe market outcomes and not the underlying strategies, they cannot identify the exact mechanism behind the price increase. Nonetheless, their work suggests that autonomous algorithmic collusion may have already materialized in some markets.

VII. CONCLUSION

The experimental and empirical evidence discussed above suggests that autonomous algorithmic collusion is a real possibility, if not already a fact. This may raise doubts that our current policy towards collusion is no longer adequate in the age of Artificial Intelligence.

As noted, autonomous algorithmic collusion is reminiscent of tacit collusion among humans. Today, the prevalent approach towards such tacit collusion is rather tolerant. Even though some jurisdictions take a more aggressive stance, in most countries (including the U.S. and Europe) tacit collusion in practice is not regarded as illegal.

The rationale for this is threefold. First, tacit collusion is viewed as a chimera: illusory and very hard to achieve. Second, it is believed that if tacit collusion nevertheless occurred, it would be hard to detect. These presumptions imply that an aggressive antitrust policy would entail many false positives, while a tolerant one just a few false negatives. Third, if tacit collusion existed and were detected, it is not clear what the possible remedies might be.

With the advent of autonomous pricing algorithms, however, each of these arguments might have to be reconsidered. First, there are reasons to believe that autonomous algorithmic collusion might be more likely than tacit collusion among humans. Admittedly, there is no clear evidence to this effect yet, but an analogy with board games may be suggestive. In board games such as chess and go, today reinforcement learning algorithms vastly outperform humans.[22] In games of pricing, performance would be assessed (from the firms' viewpoint) by means of the profit level and hence by the extent to which players manage to collude instead of competing. If the analogy with board games is not misleading, we should then expect that algorithms may be better than humans also at colluding. And we should also expect that things may get worse as

21 Assad, S., Clark, R., Ershow, D. & L. Xu. "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market" (2020), mimeo.

22 See e.g. Silver, D., et al. "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play," *Science* 362.6419 (2018): 1140-1144.

www.competitionpolicyinternational.com
Competition Policy International, Inc. 2020© Copying, reprinting, or distributing
this article is forbidden by anyone other than the publisher or author.

smarter and smarter algorithms are developed in the future.

Second, when it comes to detection, there is an important difference between collusion among algorithms and among humans. If antitrust agencies and the courts suspect collusion when algorithms oversee the pricing, they can subpoena and test the algorithms in synthetic environments that closely replicate the specific industry under investigation. With humans, this is not possible. This means that the risk of aggressive antitrust enforcement producing too many false positives may be reduced.

Third, there might be more effective remedies. For example, agencies may request firms not to use any more algorithms that have displayed a strong propensity to collude.[23] Humans, in contrast, cannot forget how to collude tacitly once they have learned to do so.

Whether, and to what extent, all this calls for a change in policy is still difficult to tell. But algorithmic collusion seems to be a real problem, which deserves close attention. Antitrust agencies are wise to be concerned.

23 Harrington Jr, Joseph E. "Developing competition law for collusion by autonomous price-setting agents." *Journal of Competition Law & Economics*, 14(3), 331-363.

www.competitionpolicyinternational.com
Competition Policy International, Inc. 2020© Copying, reprinting, or distributing
this article is forbidden by anyone other than the publisher or author.

CPI Subscriptions

CPI reaches more than 35,000 readers in over 150 countries every day. Our online library houses over 23,000 papers, articles and interviews.

Visit competitionpolicyinternational.com today to see our available plans and join CPI's global community of antitrust experts.