(MIS)UNDERSTANDING ALGORITHMIC COLLUSION

BY TIMO KLEIN[1]

1 Timo Klein, Ph.D. is a Consultant at Oxera Consulting LLP and Lecturer in Competition Economics at Utrecht University. E-mail: timo.klein@oxera.com. Any opinions expressed in this article are my own and do not necessarily represent those of my affiliations. I am not aware of any potential conflicts of interest in writing this article. I am grateful for valuable comments from Thibault Schrepel.

"I think we need to make it very clear that companies can't escape responsibility for collusion by hiding behind a computer program."

– Margrethe Vestager, European Commissioner for Competition.[2]

I. INTRODUCTION

In their well-known 2016 book *Virtual Competition*, Ariel Ezrachi & Maurice Stucke prophesized that algorithms may lead to the end of competition as we know it.[3] For instance, the rise of pricing algorithms may enable competing firms to much more effectively implement price agreements. Moreover, advanced self-learning algorithms may at some point even learn by themselves that it is optimal to collude, without any prior instructions. Authorities risk becoming powerless in the face of advanced AI cartels.

This provocative thesis has proven immensely effective in bringing the exciting world of artificial intelligence into competition policy debates. There have been several specialized reports by various authorities,[4] as well as contribution from various economic consultancies and law firms.[5] For some time now, panels on algorithms and competition have been main acts during many of the major competition policy conferences, including panels specifically on algorithmic collusion.

Nevertheless, in spite of the hype (or perhaps because of it), a sizable group of sceptics and contrarians argue that this is still much ado about nothing. For instance, algorithmic collusion has already been described as "the closest ever our field came to science-fiction"[6] and as "fundamentally unimportant."[7]

More specifically, the criticism voiced against all this attention to algorithmic collusion usually takes the following form.

2 Vestager, M. (2017), "Algorithms and Competition," Speech at the Bundeskartellamt 18th Conference on
Competition, Berlin.

3 Ezrachi, A. & Stucke, M.E. (2016), *Virtual Competition: The Promise and Perils of the Algorithm-Driven Economy*, Harvard University Press.

4 To present, these include in particular OECD (2017), "Algorithms and Collusion: Competition Policy in the Digital Age"; UK Competition and Markets Authority (2018), "Pricing Algorithms: Economic Working Paper on the Use of Algorithms to Facilitate Collusion and Personalised Pricing"; Autoridade da Concorrência (2019), "Digital Ecosystems, Big Data and Algorithms"; and Autorité de la Concurrence and Bundeskartellamt (2019), "Algorithms and Competition."

5 These include, among others, Oxera (2017), "When Algorithms Set Prices: Winners and Losers," *Agenda*, June; RBB Economics (2018) "Automatic Harm to Competition? Pricing Algorithms and Coordination," RBB Brief 55, February 2018.

6 Petit, N. (2017) "Antitrust and Artificial Intelligence: A Research Agenda." *Journal of European Competition Law & Practice*, 8(6), pp. 361–2.

7 Schrepel, T. (2020) "The Fundamental Unimportance of Algorithmic Collusion for Antitrust Law," *Harvard Journal of Law and Technology.*

1. "There are basically no cases involving algorithmic collusion, so why worry?"

2. "And even if there were more cases, our current competition laws can deal with them."

3. "And even if there are hypothetical types of algorithmic collusion that our current competition laws cannot deal with, these are still science fiction."

In this article, I argue that this criticism — while sharpening the debate — may be based on several misunderstandings.

I first provide a simple typology on pricing algorithms, in order to avoid much of the confusion that is still surrounding the debate. I then discuss each misunderstanding in turn and close off with a discussion on the need for more empirical evidence.

II. A TYPOLOGY OF PRICING ALGORITHMS

There are different types of pricing algorithms, each with their own capabilities and possible concerns. Before talking about concerns around pricing algorithms and collusion, it is good to specify exactly which type of pricing algorithm we are talking about — thereby ensuring that everyone is in fact talking about the same thing.

Generally, algorithmic pricing applications can be divided into three broad types: simple pricing rules, static optimization algorithms, and dynamic optimization algorithms.

A. Simple Pricing Rules

With simple pricing rules, firms can readily implement automatic price changes contingent on observed variables, such as competitor price, cost, or inventory. Many e-commerce platforms (including Amazon) offer retailers the option to use simple pricing rules to set their prices. There is even a growing market for specialized pricing software that enables firms to more effectively implement increasingly complex or tailor-made pre-programmed contingent pricing.

As with any pricing algorithm, there may be obvious pro-competitive benefits to simple pricing rules. For instance, they may decrease operational costs in managing prices or allow for quicker adaptation to changing market conditions.

In the case of simple pricing rules, nothing changes fundamentally relatively to an analogue setting, because the seller is still in full control of pricing. A collusive outcome would still need to be explicitly instructed. However, simple pricing rules can nevertheless be used to implement explicit or tacit price agreements much more effectively — but more on that later.

B. Static Optimization Algorithms

Going beyond simple pricing rules, there also exist algorithms that aim to estimate and maximize some objective function by themselves. For instance, they may estimate profit as a function of costs, prices or other variables and suggest a price that maximizes this estimated profit function. This is also referred to as adaptive learning, and it can make use of techniques from supervised machine learning and reinforcement learning.

To collect sufficient data to estimate an objective function, the algorithm may experiment by itself with (slightly) higher or lower prices. In that sense, pricing itself is dynamic: prices may change automatically and continuously, depending on things like the rate of experimentation or demand fluctuations.[8]

8 A survey on static optimization algorithms in the operations research and management science literature is provided by Den Boer, A.V. (2015), "Dynamic Pricing and Learning: Historical Origins, Current Research, and New Directions," *Surveys in Operations Research and Management Science*, 20(1), pp. 1–18.

While prices are dynamic, these algorithms are still static in the sense that they only estimate a one-period objective function. They generally ignore next-period profit and strategic responses by competitors. This means that, according to microeconomic theory, static optimization can never lead to collusive outcomes by itself; it is incapable of learning the threat of retaliation necessary to discipline collusive behavior.[9]

However, this theoretical result breaks down once the algorithms are not specified correctly. For instance, research shows that incorrectly optimizing a monopoly model (in which relevant competitors are ignored) may inadvertently lead to supra-competitive prices.[10] This may be relevant in practice, because pricing software often only take into account competing products that are exactly identical (e.g. products with an equivalent SKU, or Stock Keeping Unit), while ignoring products that are slightly differentiated.

Additionally, static optimization algorithms can of course facilitate collusion if they are instructed to maximize joint profit instead of own profit. It is unclear how firms would coordinate on the use of such collusive optimization algorithms in the first place, but this could in theory be facilitated by, for instance, a common third-party pricing software supplier.

C. Dynamic Optimization Algorithms

Finally, there also exist dynamic optimization algorithms that go beyond static optimization by explicitly taking into account strategic responses and the long-run consequences of current behavior.

The theory behind dynamic optimization is slightly more complicated, but the basic idea is that the algorithm learns iteratively, through autonomous trial-and-error exploration, which kind of behavior maximizes its long-run payoffs. In doing so, it makes use of reinforcement learning techniques.

The main advantage of dynamic optimization algorithms in pricing applications would be that they learn to optimize all future profits, not just one-period profit, and to explicitly take into account strategic responses by competitors. In principle, this may enable them to learn retaliation strategies that can sustain a collusive outcome.

Two recent papers show how dynamic optimization (more specifically an algorithm called "Q-learning") has the capacity to learn collusive strategies in stylized environments of simulated competition.[11] This is achieved without any communication or even explicit instructions to collude.

The main limitation of dynamic optimization algorithms, however, is their application in practice — but more on that later.

9 The seminal contribution that proves this is Milgrom, P. & Roberts, J. (1990), "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica*, 58(6), pp. 1255–77.

10 This is shown for instance by Cooper, W.L., Homem-de-Mello, T., & Kleywegt, A.J. (2015), "Learning and Pricing with Models that do not Explicitly Incorporate Competition," *Operations Research*, 63(1), pp. 86–103 and Hansen, K., Misra, K. & Pai, M. (2020), "Algorithmic Collusion: Supra-Competitive Prices via Independent Algorithms," CEPR Discussion Paper No. DP14372.

11 Calvano, E., Calzolari, G., Denicolò, V. & Pastorello, S. (2019b), "Artificial Intelligence, Algorithmic Pricing and Collusion," SSRN No. 3304991; Klein, T. (2019), "Autonomous Algorithmic Collusion: Q-Learning Under Sequential Pricing," Amsterdam Center for Law & Economics Working Paper No. 2018-05.

III. MISUNDERSTANDING 1: NO CASES MEANS THAT THERE IS NO PROBLEM

In the 2015 and 2016 *Topkins* (U.S.) and *GB Eye Trod* (UK) case, online poster retailers were found to have used simple pricing algorithms to coordinate prices.[12] This was a case of explicit human collusion, effectively implemented using algorithms but prosecuted based on "old-fashioned" incriminating e-mail correspondence.

But apart from this case, there is not much out there. There only seem to be cases that either only vaguely involve some sort of pricing algorithm (such as the 2016 *Eturas* case, in which an online travel platform facilitated collusion by restricting discounts;[13] or reports that several Spanish real estate firms used a common brokerage software to potentially coordinate prices),[14] or cases that have never gone beyond unproven allegations (such as the 2018 allegation that Accenture used a pricing software to help competing car manufacturers coordinate the prices for spare parts).[15]

As Thibault Schrepel rightly observes:

> [E]mpirical studies documenting the frequency of [algorithmic collusion] in the real-world remain to be produced. One cannot find any quantification of algorithmic collusion in official publications coming from antitrust and competition agencies, in any of the reports given to these agencies, or in the OECD publications. When having a look at the litigation brought in the U.S. and in Europe, algorithmic collusion is virtually non-existent.[16]

Schrepel suggests though, in part based on this, that all this attention to algorithmic collusion concerns is unwarranted — but an absence of proof is not the same as a proof of absence.

There are two key points here. First, it is perhaps even suspicious that we are seeing so few cases when the potential profits from collusion are huge and collusion seems as easy as the click of a button. Surely the *Topkins* poster infringement has not been unique? How do we know that the lack of cases is not driven by an increased ability to evade detection and prosecution?

Second, pricing algorithms may be used to sustain a tacit understanding to keep prices high — which is difficult to document or prosecute in the first place. For instance, the Portuguese report discusses how commercial pricing software companies actively advertise their capacity to avoid price wars.[17] As RepricerExpress communicates: "Instead of worrying so much about having the lowest costs among your competitors, RepricerExpress recommends avoiding a price war as a technique for coming out on top. [. . .] Within RepricerExpress, there are features to help sellers detect and avoid a price war."[18]

Of course, the use of pricing software may have pro-competitive justifications. However, at face value, any attempt to avoid price wars raises potential competition concerns, even if this is perfectly legal. It is therefore not surprising that competition authorities are starting to pay close attention.

12 U.S. Department of Justice (2015), "Former E-Commerce Executive Charged with Price Fixing in the Antitrust Division's First Online Marketplace Prosecution," press release, April 6, 2015; UK Consumers and Markets Association (2016), "Online seller admits breaking competition law," press release, July 21, 2016.

13 Case C-74/14 *Eturas*, ECLI:EU:C:2016:42.

14 Comisión Nacional de los Mercados y la Competencia (2020), "The CNMC opens antitrust proceedings against seven firms for suspected price coordination in the real estate intermediation market," press release, February 19, 2020.

15 Reuters (2018), "Software and Stealth: How Carmakers Hike Spare Parts Prices," June 3, 2018.

16 Schrepel, T. (2020), "The Fundamental Unimportance of Algorithmic Collusion for Antitrust Law," *Harvard Journal of Law and Technology*.

17 Autoridade da Concurrência (2019), "Digital Ecosystems, Big Data and Algorithms."

18 RepricerExpress, "How to Avoid a Price War on Amazon," https://www.repricerexpress.com/avoid-price-war-amazon, accessed April 23, 2020.

IV. MISUNDERSTANDING 2: LEGALLY, NOTHING CHANGES

Former FTC Commissioner Maureen Ohlhausen suggests the following to deal with algorithmic collusion: "[e]verywhere the word 'algorithm' appears, please just insert the words 'a guy named Bob.'"[19] In other words, if something is illegal under current competition law for an employee named "Bob," it is illegal for an algorithm.

EU Commissioner for Competition Margrethe Vestager takes a more activist approach, but basically says the same when she argues that firms "can't escape responsibility for collusion by hiding behind a computer program."[20] They need to send their algorithm to law school first.

Both may miss the point, however, for several reasons.

Yes, the *Topkins* case was a classic cartel involving managers explicitly colluding. Once discovered, the law can readily deal with these. However, it seems that the pricing algorithm did facilitate the kind of collusion that would otherwise have been very difficult to sustain. Thus, even in those cases where the law can be readily applied by replacing the word "algorithm" by "a guy named Bob," the fact remains that the act of the infringement may become a lot easier — which is on itself already worrying.

Additionally, firms may require less and less communication to reach and sustain a collusive agreement in the first place. The *Topkins* case involved email correspondence as smoking-gun evidence. One can easily imagine, however, that when pricing algorithms are commonplace, managers need much less of such communication to implement their scheme — making detecting and prosecuting cartel cases more difficult (perhaps to the point that prosecution becomes impossible because of a tacit nature of the collusive understanding, as mentioned before).

More importantly, however, the "antitrust compliance by design" argument seems to ignore the fact that competition law has been developed for human suspects, not algorithms. To be precise, competition law focuses on some form of communication as the necessary evidentiary burden of collusion — i.e. the collusion has to occur explicitly rather than tacitly, even though the economic impact of explicit and tacit collusion is theoretically equivalent. Justice Stephen Breyer argues:

> [T]hat is not because [tacit collusion] is desirable (it is not), but because it is close to impossible to devise a judicially enforceable remedy for "interdependent" pricing. How does one order a firm to set its prices without regard to the likely reactions of its competitors.[21]

However, with algorithms, one may actually be able to devise a judicially enforceable remedy for collusive interdependent pricing — irrespective of whether it is tacit or explicit. This is because unlike with humans, one can (in principle) look inside the head of an algorithm to see its contingent pricing strategy and possibly determine whether it is anti-competitive or not.[22]

---

19 Ohlhausen, M.K. (2017), "Should We Fear The Things That Go Beep In the Night? Some Initial Thoughts on the Intersection of Antitrust Law and Algorithmic Pricing," Remarks from the Concurrences Antitrust in the Financial Sector Conference, New York, NY.

20 Vestager, M. (2017), "Algorithms and Competition," Speech at the Bundeskartellamt 18th Conference on Competition, Berlin.

21 *Clamp-all Corporation v. Cast Iron Soil Pipe Institute*, et al., 851 F.2d 478, 484 (1st Cir. 1988).

22 This argument is developed further in Harrington, J.E. (2019), "Developing Competition Law for Collusion by Autonomous Price-Setting Agents," *Journal of Competition Law and Economics*, 14(3), pp. 331–63.

V. MISUNDERSTANDING 3: TRULY AUTONOMOUS COLLUSION IS SCIENCE FICTION

The capacity of artificial intelligence may seem limitless. Intuitively, it may therefore also seem obvious that artificial intelligence will also, at some point, be able to collude by itself. If collusive strategies are so profitable, surely smart algorithms will find a way to achieve them.

Ezrachi and Stucke (2017) summarize this position by suggesting the following situation: "Two Artificial Neural Networks and one Nash equilibrium meet in an online (pub) hub. After a few milliseconds, a unique silent friendship is formed…"[23]

Many are skeptical. For instance, Kai-Uwe Kühn and Steve Tadelis argue that the correct version of the above situation would be: "Two Artificial Neural Networks meet a multi-dimensional continuum of subgame-perfect Nash equilibria in an online hub. With unbelievable speed the two Artificial Neural Networks react and say 'huh?'"[24]

The point Kühn and Tadelis are making is that it is unclear exactly how algorithms (even really smart ones) would be able to coordinate on any one particular collusive outcome. Algorithms may be able to enforce collusive strategies much more effectively — through better monitoring and quicker retaliation — but what proof is there that algorithms also have the capacity to coordinate on any one particular collusive strategy?

As already discussed, static optimization algorithms cannot by themselves lead to collusive outcomes (unless they are wrongly specified), because they are inherently one-shot. However, dynamic optimization algorithms may learn to collude by themselves simply through trial and error.

Research into this question has only just started, but recent research by Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò and Sergio Pastorello in one paper and by myself in another shows that autonomous collusion by dynamic pricing algorithms is not science fiction.[25] Both papers use Q-learning — a standard and well-developed type of reinforcement learning, which in turn is the type of machine learning in which the algorithm does not learn from some existing dataset, but through autonomous trial-and-error exploration.

These papers show in different simulated environments how Q-learning is indeed able to learn collusive reward — punishment strategy: keep the price high as long as my competitor does, but punish with lower prices if my competitor deviates from this. Punishment in these cases is generally temporary.

Many practical limitations exist. In particular, the algorithm needs many periods of costly experimentation. Additionally, it requires the environment to be stable. If there are structural breaks (such as entry, exit or demand shifts), it may have to relearn its behavior.

The existence of these practical limitations means that any immediate concerns around truly autonomous algorithmic collusion would indeed be overblown. However, these papers do show that autonomous algorithmic collusion is not, in principle, science fiction. There are practical limitations, but these do not seem unsolvable and both papers discuss possible routes towards practical implementation, for instance by specifying a model of the competitive environment and letting the learning take place in an offline, simulated environment instead of in the real world, or by transferring the learning from one product to another product.

23 Ezrachi, A. and Stucke, M. (2017), "Two Artificial Neural Networks Meet in an Online Hub and Change the Future (Of Competition, Market Dynamics and Society)," Oxford Legal Studies Research Paper No. 24/2017.

24 Kühn, K.U & Tadelis, S. (2017), "Algorithmic Collusion," presentation prepared for CRESSE 2017.

25 Calvano, E., Calzolari, G., Denicolò, V. & Pastorello, S. (2019b), "Artificial Intelligence, Algorithmic Pricing and Collusion," SSRN No. 3304991; Klein, T. (2019), "Autonomous Algorithmic Collusion: Q-Learning Under Sequential Pricing," Amsterdam Center for Law & Economics Working Paper No. 2018-05.

VI. THE NEED FOR MORE EMPIRICAL EVIDENCE

The specialized reports by various competition authorities have been very valuable in identifying exactly where the main theoretical concerns around pricing algorithms may be. The problem at this stage, however, is the lack of clear empirical evidence that the use of certain pricing algorithms indeed also leads to anti-competitive outcomes in practice. The theoretical concern is surely there, but the empirical evidence is still limited.

A better empirical understanding requires three things. First, it requires more cases by competition authorities around allegations of explicit algorithmic collusion by humans (as in the *Topkins* case), to see exactly how they may operate in different cases.

Second, we need market studies to know more about the actual use and motivation of different pricing algorithms in practice. These studies could then aim to see whether market outcomes really are collusive, even in the absence of any explicit agreements (i.e. tacit, and hence generally not illegal). Alternatively, these studies could also reassure authorities that certain algorithmic pricing practices are actually pro-competitive.

Third, we need more empirical academic research. There seems to be a limited amount of empirical academic work undertaken on this topic — which is surprising, given the potentially large amount of data that is available.[26]

To close off, it needs to be stressed again that pricing algorithms can bring a huge amount of benefits to competition, such as (much) more effective market clearing, lower costs of price adjustments, and more efficient inventory management. Pricing algorithms can be very much pro-competitive. Arguing for a by-object infringement or imposing cross-sector regulations in the case of algorithmic pricing would therefore surely be problematic.[27]

However, the theoretical concerns around algorithmic collusion are there. Over just a few years, we have come a long way in our understanding of these concerns; however, there is still much to do, both in terms of removing existing misunderstandings and expanding our empirical understanding.

26 Exceptions include, for example, Chen, L., Mislove, A. & Wilson, C. (2016), "An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace," Proceedings of the 25th International Conference on World Wide Web, 1339–1349; Brown, Z. & MacKay, A. (2020), "Competition in Pricing Algorithms," SSRN No. 3485024.

27 A comprehensive review of the legal discussion around algorithmic collusion and a proposal on potential regulatory responses is provided by Gal, M.S. (2019), "Algorithms as Illegal Agreements," *Berkeley Technology Law Journal*, pp. 67–118. An interesting discussion on algorithms and competition law is provided by Schrepel, T. & Gal, M.S. (2020) "Algorithms & Competition Law: Interview of Michael Gal," Concurrences, e-Competitions Special issue Algorithms, May 14, 2020.