Department of

Economics and Finance

Course of

IO & Competition Theory

Algorithm Pricing: a new challenge for competition policy

Prof. Gian Luigi Albano

Supervisor

Michele Angino 224731

Candidate

Academic Year 2019/2020

Table of contents

1. Introduction

The earliest competition laws trace back to Roman times, when efforts were made to limit fluctuations of prices and to prevent unfair trade practises. Such attempts continued during the Middle Ages: in particular, the focus was on avoiding unfair trading practises (like forestalling[1]) and monopolies. In fact, as stated by Colino (2011), "*At the time of the Magna Carta (1215) legislation provided that all monopolies were to be contrary to the law because of their pernicious effect on individual freedom.*" The foundation for modern competition law was laid in this period, even if the concept of collusion made its way into the realm of trade regulation only from the 17th century onwards.

In order to keep up with innovation and face the new challenges that arose over time as societies grew and technological developments took place, competition policy has kept evolving. Such evolution consisted in new practises to regulate (like merger control), in changes in the approach towards already regulated practises (for example, forestalling is nowadays obsolete) and also in a more international approach to tackle anti-competitive issues (as represented by the European competition law).

As new challenges for competition law appear on the horizon, it is important to predict what they will look like and to understand in advance how to handle them, in order to provide efficient tools to antitrust authorities with a small lag of time.

In this work, the issues related to pricing algorithms will be explored, and the associated risks for competition will be discussed. While some area of antitrust law seems to provide safeguards against the anti-competitive use of pricing algorithms, actual legal framework may result inadequate in case more complex technologies will be successful in achieving coordination, as this work will discuss.

The rest of this paper is organised as follows. Chapter 2 starts with a definition of collusion, describes in detail several facilitating market features and discusses welfare impact as well as the current common legal approach. Chapter 3 presents pricing algorithms and debates the risk that their implementation brings to competition. Chapter 4 is devoted to the description of an experiment by Calvano et al. (2019a) in which algorithms provided with Artificial Intelligence learn to collude. Chapter 5 discusses possible changes in competition policy to face the threat represented by self-learning algorithms. Chapter 6 concludes.

[1] Forestalling refers to the practise of purchasing commodities before they reach the marketplace and then inflate the price.

2. Collusion

a. Definition and facilitating features

The term collusion refers to any form of coordination or agreement among rivals within the same industry, undertaken for the purpose of gaining unfair market advantage and raising profits above the level obtained in a situation of competitive equilibrium. Stated in other terms, collusion refers to a joint profit maximization strategy adopted by competing firms.

It is common to distinguish between explicit and tacit agreements to collude, where the distinction hinges on how the agreement is reached. The former case occurs when competitors jointly design a common plan of action and exchange mutual assurance to follow it, and involves overt communication and discussion among companies. The group of firms which agree to coordinate behaviours to increase profits, acting as a single producer, is termed cartel.

Tacit collusion occurs instead when businesses manage to coordinate their conduct by observing and anticipating their rivals' behaviour. Since competitors recognize their mutual interdependence and the advantages coming from coordination, a firm could adopt a strategy knowing that it will be mutually beneficial if all other firms do the same.

For a collusion to exist and be sustained over time, three conditions are needed. First, firms must adhere to a common policy focusing on a focal point. Second, they can monitor the behaviour of each participant. Third, they can enforce such common policy by punishing any deviation.

At the same time, in order for collusive equilibrium to be established, to last over time and to be profitable (from the businesses' point of view), the market should possess certain facilitating features which can be structural, demand-side or supply-side.

Among structural characteristics we find:

- Barriers to entry: the absence of barriers to entry makes it hard for the collusive equilibrium to be sustained over time since higher profits will attract new entrants which, by increasing the competition in the market, will end up causing an erosion of supra-competitive profits.
- Market transparency[2]: in transparent markets, companies can easily monitor each other's actions and detect deviations[3]. It must be noted that tacit collusion is unfeasible without public prices.

[2] A market is transparent if price and supply (market depth) information are readily available to the public.
[3] For example, consumer markets are generally more transparent than supplier markets.

- Frequency of interaction[4]: frequent interactions enable firms to promptly retaliate and punish deviations if detected.
- Innovation: in industries characterized by product innovation, reaching an agreement is more difficult since innovation reduces the present value of agreements and the ability of less advanced firms to retaliate. If colluding firms agree to standardize products too, the equilibrium is reinforced by eliminating product diversity and limiting potential product heterogeneity over time.

Demand-side factors include:

- Demand trends: whether the demand is going to increase, decrease or stay the same in the (next) future plays an important role. If the demand is going to increase, companies have less incentives to deviate from the collusive equilibrium since this would prompt retaliation and cause the loss, in following periods, of the benefits of a larger market which they could more profitably accommodate staying within the agreement. If instead the market is shrinking, deviating today is preferred, since it allows the firm to reap higher profits in the present and give up the declining revenues they would get by keeping conspiring.
- Market elasticity of demand: a large elasticity, meaning that consumers can easily switch to alternative products in case of an increase in prices, reduces the effectiveness of the collusive agreements since it limits the level at which prices can be remuneratively raised. In fact, consumers can switch to goods produced by other companies, even if less known and smaller, which are outside the agreement and whose prices are consequently lower.

Finally, supply factors are:

- Product heterogeneity: if products are homogeneous, firms only need to agree on price and/or output level, while heterogeneity complicates the negotiations since it requires an agreement over price and/or output for each product, magnifying the possibilities for disagreements[5].
- Cost asymmetry: asymmetries in costs may complicate the realization of the agreement since high-cost firms would generally prefer higher prices and lower output, while low-cost firms would prefer lower prices and higher output; hence, joint profit maximization

[4] Interaction refers to both price adjustments and detection of price adjustments by other vendors.

[5] Moreover, since effective collusion requires more than just fixing prices, product differentiation further complicates reaching any deal, considering that firms will compete through other means like advertising and provision of services. Hence, successful collusion would require the scope of the agreement to include restriction on non-price competition to avoid the risk of profit dissipation due to expenses derived from such non-price struggles.

would require some high-cost firms to shut down. Moreover, cost differences can provide low-cost firms with incentives to cheat in case companies divide output equally (complication which can be nonetheless solved through side payments[6]).

- Incomplete information: if information about costs is private, reaching an efficient agreement is difficult. For example, if a firm is better informed about its cost than its rivals, it could convince them that it is cost efficient, and hence earn a higher share of collusive profits[7].
- Uncertainty: uncertainty regarding conditions at the industry level hinders the realization of an agreement since businesses may have different views and expectations about future paths of demand and costs. Every time a significant change in the environment occurs, firms need to renegotiate and the possibility of disagreement arises.
- Asymmetry in preferences: asymmetries in firms' discount factors (their valuation of the future) are critical since they influence the estimation of punishments and determine whether these will prove sufficient to prevent the risk of cheating. Besides, differences in the willingness to engage in illegal activities or attitude toward risk may jeopardize the formation of an agreement.

In addition to these factors, cartels' intrinsic characteristics also determine the likelihood and the success of collusion. For example, as the number of firms involved increases, the identification of a focal point for coordination becomes harder and incentives for collusion reduce, since each player would receive a smaller share of the non-competitive gains. Likewise, the relative number and size of participants impact the success of collusion: if the number and the size of the firms participating in the agreement relative to the number and size of firms outside of it is large, the potential to produce market power via collusion is greater.

[6] Exchange of sums of money between parties of a transaction that are not part of the transaction itself.

[7] A firm may pretend to be cost-efficient, and thus require a higher share of the market or of the collusive profits, threatening deviations (see the previous point regarding cost asymmetry).

i. A cartel in the electrical equipment industry

In the early 1950s, General Electric (with a market share between 40 and 45 percent), Westinghouse (with a market share between 30 and 35), Allis-Chalmers and Federal Pacific (both with a market share around 10 percent) were the dominant firms in the heavy electrical equipment market. Their average annual return in that decade was $1.75 billion.

In order to achieve even higher profits, leaders of these and other companies decided to fix prices on industrial switchgears sold to the various governments at sealed-bid auctions or to private electric utilities. In particular, the companies determined both the amount and the executor of the low bid through meetings and telephone calls. In the same way they fixed book prices and market shares for sales to private entities.

The conspirators followed a precise way of operating: they used pay phones to arrange meetings, phoned co-conspirators at home to discuss prices, never registered at hotels with companies' names or saw each other in public dining rooms or hotel lobbies, used blank letter paper instead of company's one and doctored expense accounts replacing with fictitious destinations (at the same distance from their office as the city in which meetings were held) the real ones.

This cartel witnessed some price competitions during its life: in 1954, Westinghouse won a substantial order through a heavy discount off the established "book" price. This caused a retaliation by General Electric which led to a price war across all product lines (prices were discounted even by 45%). Such situation lasted until 1956, when an industry-wide agreement was reached and companies started their meetings and correspondence again. This second equilibrium collapsed one year later, in 1957, when Westinghouse cheated again offering a secret discount on a large order. The buyer reported such offer to General Electric, who matched his competitor's offer and was in the end awarded the contract. This led to a second price war, with discounts reaching 60%.

At this point, the industry tried to reach a new deal, but in the absence of General Electric, all efforts were unsuccessful. However, switchgear managers of General Electric were under pressure to raise prices from the managers of the other production lines, engaged in turn in collusive agreements. In fact, the customers they shared wondered why such discounts were offered only for switchgear. This situation led to new negotiations and, in the end, to an agreement in 1958.

A central element of this cartel was a scheme named "phases of the moon": one manufacturer quoted the lowest price, others offered intermediate prices and the remaining ones high prices. Such positions were periodically rotated among the conspirators and this arrangement was planned in such a way that bid prices' spreads were narrow enough to eliminate price competition, but sufficiently wide to give an idea of competition.

The investigations began when the Tennessee Valley Authority (TVA) noted that several manufacturers had submitted identical bids for three years, even if the bids were supposed to be secret. The "phases of the moon" system was discovered, together with other smoking-gun[8] documents, handed over by a participant (which, in order to train an assistant, ignored instructions to destroy all written evidence).

Finally, in 1960, 29 companies along with 45 executives were indicted under Section 1 of the Sherman Act[9]. Charges covered 20 different electrical product lines. Due to the presence of strong evidence, the defendants decided to plead guilty on the major indictments and not to contest minor ones. Total fines of $2 million were imposed and the seven most senior managers received prison sentences.

It is estimated that such conspiracy costed taxpayers $175 million for each year it existed.

[8] The term refers to objects or facts which constitute the conclusive evidence of a crime (strongest kind of circumstantial evidence).
[9] See next section.

b. Welfare impact of collusion and legal considerations

Collusion impacts negatively on consumers and economic welfare in various ways. First, it causes prices to rise leading to a decline in consumer surplus. Second, it may discourage new firms from entering the market (if the colluding businesses' aims include the creation of barriers to entry). Third, it gives firms no incentives to innovate and invest to increase productivity. Hence, it brings the disadvantages of a monopoly (higher prices and restricted output) but none of its advantages (like economies of scale, innovation and technological development).

Consequently, competition authorities all over the world have made detection and prosecution of cartels a primary objective[10]. For instance, in the United States, Section 1 of the Sherman Act states "*Every contract, combination in the form of trust or otherwise, or conspiracy, in restraint of trade or commerce among the several States, or with foreign nations, is declared to be illegal*". In the European Union, Article 101 of the Treaty on the Functioning of the European Union (ex-Article 81 of the EEC treaty) prohibits "*[...] all agreements between undertakings, decisions by associations of undertakings and concerted practices [...]*" that could disrupt free competition in the European internal market.

As one can see from the articles quoted above, the legal approach in many jurisdictions focuses on the means used by companies to achieve a collusive outcome[11]. Antitrust laws generally do not forbid collusion as such but prohibit anti-competitive agreements: if that results from such settlements, then an infringement of the law can be proved. An evidence of contact demonstrating that firms have not acted autonomously, is therefore traditionally required.

Consequently, the court's approach to conscious parallel behaviour in oligopolistic markets, hence without any agreement between the parties, is controversial. It is usually difficult to determine what the normal conditions of the market should be: situations which might signal ongoing concerted practises (as price uniformity and/or supra-competitive prices), could also be the normal outcomes resulting from the rational economic behaviour of the members of an oligopoly selling a homogeneous product. For this reason, conscious parallelism falls outside the realm of competition policy.

Between these two scenarios (explicit collusion and simple conscious parallelism), there is a business behaviour, especially in oligopolistic markets, that consists in competitors engaging in practises which make tacit collusive outcome more likely. To tackle these situations some

---

[10] "*[...] negotiation between competitors may facilitate the supreme evil of antitrust: collusion. [...]*", Antonin Scalia, Associate Justice of the Supreme Court of the United States; Verizon Communications, Inc. *v.* Law Offices of Curtis *v.* Trinko, LLP, 2004.

[11] Contrary to the economic approach which considers collusion the output of the market.

jurisdictions have extended the idea of "agreement", which can now be inferred also from evidence showing that competitors are not acting autonomously. However, in order to prove a concerted practise, a precise and consistent body of evidence is required. Indeed, it must be shown that the agents entered in a conscious commitment for a common scheme. Therefore, a crucial element of proof is the occurrence of the so called "plus factors", that is circumstantial evidence, inconsistent with unilateral conduct, which demonstrates the consciousness of an agreement. In other words, they consist in facts which tend to exclude the possibility that alleged conspirators acted independently[12]. Certain conducts that recur frequently in antitrust cases, have been labelled as plus factor by courts[13].

To summarise, tacit collusion among human agents is by nature unlikely to occur and last for long periods[14] but, at the same time, it is gruelling to detect.

In this complex framework the role of pricing algorithms may result game-changing.

---

[12] "*[…] economic actions and outcomes, above and beyond parallel conduct by oligopolistic firms, that are largely inconsistent with unilateral conduct but largely consistent with explicitly coordinated action. […]*", "Plus Factors and Agreements in Antitrust Law", W. Kovacic, R. Marshall, L. Marx, H. White (2011).

[13] *"In Re/Max Int'l v. Realty One […] (1999), the court provided this list: "(1) whether the defendants' actions, if taken independently, would be contrary to their economic self-interest; (2) whether the defendants have been uniform in their actions; (3) whether the defendants have exchanged or have had the opportunity to exchange information relative to the alleged conspiracy; and (4) whether the defendants have a common motive to conspire.",* "Proof of Conspiracy Under Federal Antitrust Laws" (2010).

[14] In a laboratory setting, with no communication possible, two agents sometimes are able to reach supra-competitive prices, three agents set prices near the one predicted by Nash equilibrium, while four or more tend to pursue more aggressive strategies (Huck et al. 2004).

3. Pricing Algorithms

a. Definition

The word algorithm comes from French "*algorithme*", refashioned from Old French "*algorisme*", derived in turn from Medieval Latin "*algorismus*" or "*algorithmus*", a mangled Latinization of the name of Persian mathematician, astronomer and geographer Muhammad ibn Musa al-Khwarizmi. In 825 he wrote a treatise on the Hindu-Arabic numerical system which was translated into Latin under the title "*Algoritmi de numero Indorum*", or "Algoritmi on the numbers of the Indians" where "Algoritmi" was the translator's Latinization of Al-Khwarizmi's name[15].

An algorithm is a procedure for solving problems: it is composed by a finite number of instructions or rules which must be followed in a fixed sequence in order to get, from some input(s), an output. The term is usually used in relation to a machine and especially a computer (even if it does not always apply to computer-mediated activities). Moreover, the word algorithm is often paired with other terms specifying the activity for which it has been designed (predictive algorithms, tracking algorithms, lossless compression algorithms, etc.).

The role played by algorithms in today's society is hardly describable and this is also a consequence of their enhanced adoption by companies. Indeed, an increasing number of firms rely on them for predictive analysis and optimisation of business processes[16]. This is not only transforming the competitive scenario (how firms operate and interact with each other) but it is affecting the evolution of markets towards global digitalisation as well, promoting a wider use of algorithms. Stucke and Ezrachi (2017) argued that as a company uses algorithms and enhances its efficiency and productivity, its competitors are likely to feel the pressure to do the same and develop and employ similar algorithms; as more users rely on such facilities, computer scientists will develop new and advanced versions, increasing incentives for companies to employ them (in a positive loop).

Algorithmic pricing refers to the practise, increasingly common among companies, of automatically setting prices of items for sale. In other words, firms' pricing decisions are progressively delegated to software programs which constantly adjust and optimize individual prices on the basis of many factors, like available stock, anticipated demand and especially competitors' prices.

[15] From https://www.etymonline.com/word/algorithm, accessed on 25 May 2020

[16] Predictive analytics refers to a category of data analytics which generates future insights with high degree of precision on the basis of historical data's analysis. Optimisation of business process instead consists in activities that favour the reduction of production and/or transaction costs, the segmentation of consumers and optimal price setting.

With respect to standard pricing strategies, algorithmic pricing is able to process in real time huge amounts of data and hence to react promptly to changes in the market conditions. Such automation allows for dynamic pricing (continuous price changes over time), and price discrimination (different prices charged to different consumers on the basis of their characteristics).

Weiss and Mehrotra (2001) recognised that pricing algorithms improve efficiency in the market, since instantaneous reactions to changes in supply and demand conditions imply that equilibrium is (almost) always reached, preventing excess demand and supply situations (perishable goods like groceries or airline tickets are less likely to go to waste) and ensuring that all mutually beneficial transactions are executed.

Nonetheless, some criticize dynamic pricing since it forces consumers to take their decisions in a scenario of constant price fluctuations, it may facilitate first degree price discrimination based on location, browsing and purchasing histories and other private information of the customers (even if first price discrimination actually improves efficiency, it results in lower surplus for consumers)[17] and it may play an active role as maker or facilitator of collusion[18].

Such algorithms may be developed by businesses to directly set prices for their products (generally, these are large companies which can afford to develop such software) or, alternatively, they can be developed by firms specialized in algorithm development. Usually, the latter type of algorithms is not specifically tailored to one market since it is licensed for other companies to use, but has made this technology affordable even for small businesses.

[17] Also known as perfect price discrimination, first price discrimination occurs when the maximum possible price for each unit consumed is charged by setting price equal to consumers' willingness to pay. In this context, pricing algorithms might allow businesses to set prices that tend towards (and not exactly match) consumers' reservation price, hence reducing consumers' surplus (yet not totally zeroing it).

[18] See next section.

b. Implementation and risks for competition

In principle, algorithm pricing could be applied both in offline and online markets. Obviously, as its efficiency relies on the availability of real-time data on consumers and competitors, it is hard to employ them in brick-and-mortar retail stores, where data collection must be carried on manually and where human intervention is needed in order to physically change the prices on offer. Attempts in such direction have been conducted by some major retailers in the UK which, in their shops, adopted electronic price tags that allow prices to change quickly and frequently in response to fluctuations in demand[19].

On the contrary, the implementation of pricing algorithms in online markets is straightforward. In December 2013 Amazon implemented more than 2.5 million price changes every day, almost 10 times more than the numbers of December 2012, which was 269,133 (clear sign of the use of algorithms). To draw a comparison with two brick-and-mortar behemoths, Best Buy, in the whole month of November, made 52,956 price changes while Walmart, in the

Figure 1 - Amazon price changes

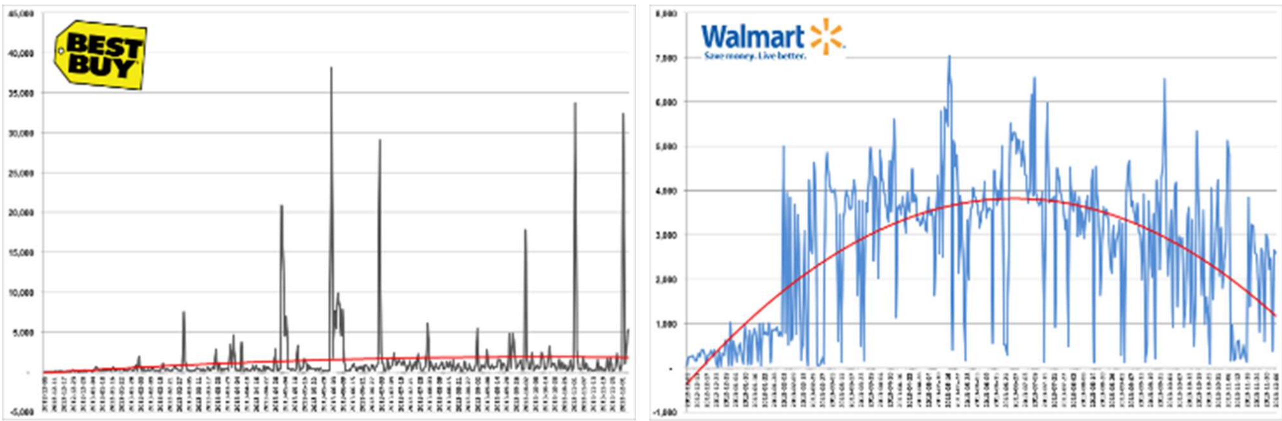

Figure 2 - Best buy and Walmart price changes

[19] "Exclusive: End of fixed prices within five years as supermarkets adopt electronic price tags", *The Telegraph*, 24 June 2017. See https://www.telegraph.co.uk/news/2017/06/24/exclusive-end-fixed-prices-within-five-years-supermarkets-adopt/.

same period, stopped at 54,633[20]. Figure 1 and Figure 2 graph this data for 2013, with the number of daily price changes on the y-axis and time on the x-axis.

Large retailers are not the only ones employing algorithms for their pricing strategies. Chen, Mislove and Wilson (2016) developed a methodology for detecting the use of pricing algorithms and uncovered that, among all merchants selling any of 1,641 best-seller products on Amazon Marketplace, over 500 of them adopted such instruments. In particular, these vendors, who resulted favoured in "winning" the Buy Box[21] are also more successful than the other sellers (they offer less products but receive higher amount of feedbacks, implying that their volume of sales is higher). In this sense, it appears clear that for non-algorithmic vendors it is going to be challenging to compete with algorithmic ones and the likely outcome is that all of them will adopt automation sooner or later (or exit the market).

Clearly, such algorithms must be implemented in the correct way in order to result beneficial, otherwise they could interact in unexpected ways and produce unwanted results. This is indeed what happened on Amazon marketplace to the book "The Making of a Fly": two sellers were automatically adjusting their prices against each other, in particular one kept setting the price at 1.27059 times the price of the other while this set its price at 0.9983 times the price of its rival. The outcome was an automatic escalation of price which reached $23,698,655.93[22].

Notwithstanding this example, for the reasons that will be covered at the end of this section and in the next ones, the increased use of pricing algorithms represents a risk for competition since collusion may become easier to sustain, especially in the digital retail market characterized by frequent interactions. Furthermore, the issue related to algorithms is that they can simplify coordinated behaviours without the need for agreements, strengthening and facilitating tacit collusion.

First of all, as mentioned before, the fact that more and more businesses rely on algorithms implies that increasingly more information and data are gathered and stored to be promptly available. If we also consider that algorithms, especially the more complex versions, are able, and will be better able in the future to make predictions, and hence reduce uncertainty, it is easy to see that a first

[20] "*Profitero Price Intelligence: Amazon makes more than 2.5 million daily price changes*", 2013, https://www.profitero.com/2013/12/profitero-reveals-that-amazon-com-makes-more-than-2-5-million-price-changes-every-day/. The three graphs are taken from this webpage, too.

[21] The Buy Box is a common instrument used by Amazon's customers to carry out their purchases (estimates in the cited paper refer to 82% of sales going through it). It contains the price of the product, shipping information, the name of the seller, and a button to purchase the product. In case a product is sold by multiple vendors, an Amazon algorithm establishes which seller's offer is displayed in the Buy Box ("the winner").

[22] "How A Book About Flies Came To Be Priced $24 Million On Amazon", see https://www.wired.com/2011/04/amazon-flies-24-million/.

consequence is a huge increase in market transparency[23]. For instance, companies are now in a better position to distinguish between deviations from collusion and rational response to changes in conditions of the market, thus preventing unmerited retaliations (CMA 2018).

A second natural effect concerns the frequency of interactions. In digital markets, where prices can be changed as frequently as desired, the addition of automation allows immediate prices updates and real-time reprisal to deviations from collusion.

These elements provide algorithms with two powerful and fearsome characteristics: they can create automated mechanisms which ease the implementation and the monitoring of a shared policy by the participating firms and allow the formation of collusion even if the market is not very concentrated. So a small number of firms is no longer a necessary condition for the existence of algorithmic collusion due to the speed in collecting and analysing data.

[23] In the words of a joint report of the French and German competition authorities (2016) "*Even though market transparency as a facilitating factor for collusion has been debated for several decades now, it gains new relevance due to technical developments such as sophisticated computer algorithms. For example, by processing all available information and thus monitoring and analysing or anticipating their competitors' responses to current and future prices, competitors may easier be able to find a sustainable supra-competitive price equilibrium which they can agree on.*" Autorité de la Concurrence and Bundeskartellamt.

c. Role and types of pricing algorithms

Pricing algorithms may be used to manage and enforce existing coordinated strategies (facilitator of collusion, Figure 3), to set up new ones (Figure 4), but they can also lead to collusive outcomes even when they are used by each company for unilateral pricing decisions (Figure 5).

As previously stated, algorithms may stabilize existing collusive agreements making it easier to detect and punish deviations. Since solidity of conspiracies depends on whether involved firms would find it profitable to cheat and lower prices, if firms are able to detect and punish the deceiver promptly, the incentive to deviate will be significantly reduced, enhancing the steadiness of the settlement.

Moreover, chance of accidental deviations is reduced thanks to the availability of mass data, which avoids that imperfect information impacts on the sustainability of the agreement and on the mutual trust among firms.

Finally, the use of algorithms may be beneficial for collusion by reducing the agency slack[24]. This happens because algorithms diminish the scope for individuals to take pricing decisions themselves.

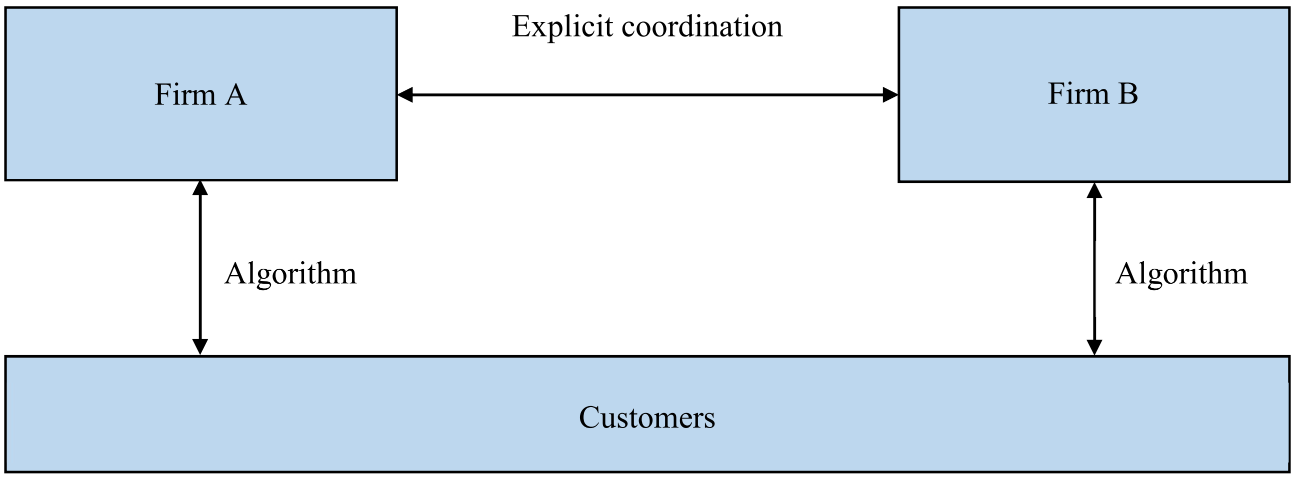

Figure 3 - Pricing algorithm as facilitator of explicit coordination, CMA (2018)

A growing competition literature has raised concerns that algorithms may also favour, or directly lead, to tacit collusion, stressing the risk that coordinated behaviour may result even when each firm uses algorithms to make unilateral pricing decisions[25]. Many reasons support such theory: as already mentioned, the increased use of automated pricing methods creates a market with enhanced transparency and where interactions occur at a high frequency. These two elements give birth to an

[24] This is a phenomenon occurring when, even if an agreement has been settled among businesses' managers, other non-management employees may have incentives not to comply with such guidelines (because of intra-firm competition for promotions, or because their salaries are proportionate to the number of sales, or for other reasons), undermining the solidity of the cartel.
[25] This literature goes under the name "theories of harm".

environment where tacit collusion can be sustained with relative ease. Moreover, algorithms may be better able than humans in estimating the profit-maximising tacit-coordination price in absence of an explicit agreement.

According to Ezrachi and Stucke (2017), the formation of tacit coordinated behaviour can occur through three main ways.

First, companies may use identical software and data pool for their pricing decisions. The natural outcome is that, on one side, the businesses will respond in a similar way to exogenous events and, on the other, they will be in a better position to predict their competitors' actions, helping mutual understanding of intentions and behaviours.

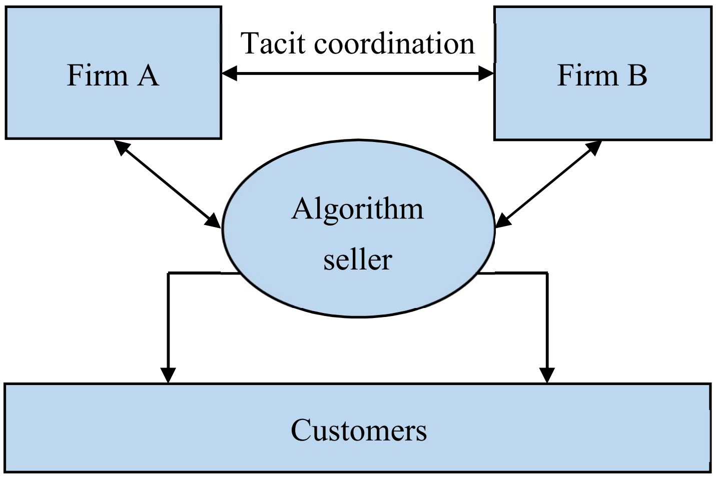

Figure 4 - Pricing algorithm in a hub-and-spoke framework, CMA (2018)

Obviously, the fact that companies use the same pricing algorithm is not sufficient to establish collusion by itself, and intention to cooperate by the competitors is still needed. Moreover, there is no guarantee that firms will find out they are using all the same software, and those willing to coordinate prices might need to explicitly reveal the details of their algorithm, which may be interpreted as a clear attempt to coordinate. This last caveat disappears in a hub-and-spoke framework, where competitors decide to delegate their pricing decisions to a common intermediary which provides algorithm pricing services (as exemplified in Figure).

A second threat is represented by the predictable agent model, that is a model in which pricing algorithms are independently designed to react in a predictable way to exogenous events. With explicit communication missing, in fact, tacit coordination is more likely to occur if firms follow simple, communicative and foreseeable pricing behaviour.

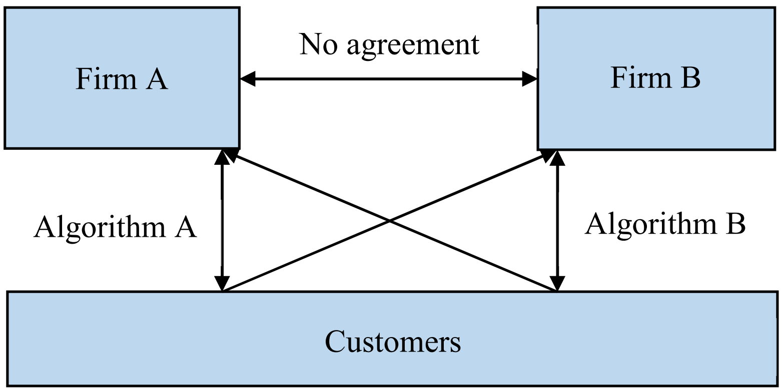

Figure 5 - Pricing algorithm in a predictable agent model, CMA (2018)

Finally, tacit collusion can be the outcome generated by pricing algorithms acting as autonomous machines: even if unilaterally designed to reach a specified target, such as the maximisation of profits, the algorithm might learn by its past actions and experience, finding out that the optimal pricing strategy is to enhance transparency and collude.

The following section briefly describes four types of pricing algorithms divided according to which operations they carry out (from an anticompetitive perspective).

i. Monitoring algorithms

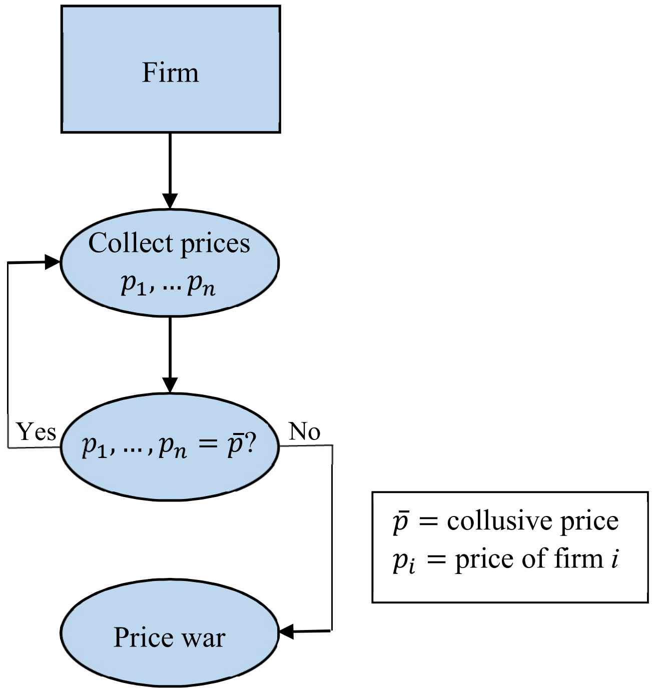

Figure 6 - Monitoring algorithm, OECD (2017)

Monitoring algorithms can be employed as facilitators of collusion: as their name suggests, these algorithms oversee competitors' actions, collecting information about their decisions, looking for any potential deviation, and eventually designing immediate reactions, as shown in Figure 6. In other words, data are collected and compared by the algorithm, which is able to automatically react if a competitor deviates from an agreed price (grim trigger strategy)[26].

Since actual deviations are discovered within a small lag of time, competitors lose their incentive to cheat, hence it is hard to observe price wars when algorithms are used (unless triggered by algorithmic mistakes).

An interesting example of how monitoring algorithms can be employed in the petrol industry (and consequently in retail stores) was provided by Dong et al. (2008). The authors proposed and experimented an application for "wireless sensor network" (WSN[27]) devices, which could be employed to automatically collect data, allowing consumers and businesses to monitor oil prices almost in real-time. In fact, despite firms transact a relatively homogeneous product, price dispersion is often observed across competing petrol stations. Also, comparison websites cannot provide full information, as they incur in significant cost to collect data from different petrol stations.

The proposed system works through a network of mobile phones equipped with GPS and video cameras, with owners voluntarily sharing information through an application. In order for this to happen, Dong et al. developed a prototype computer vision algorithm which is mechanically triggered when a mobile phone enters in the area of a petrol station and which can detect and read fuel prices from the images of price boards obtained through the mobile camera. In a test with 52

[26] In game theory, a player adopts a grim trigger strategy if he cooperates and keeps cooperating until the opponents do the same, but in case these defect once, he will defect for the remainder of the game. In the real world, the player usually defects for a certain amount of time, then cooperation is restored.

[27] WSNs consist of networks of a large number of devices equipped with technology to detect physical phenomena, such as light or heat.

images, such algorithm achieved "*a hit rate of 92.3% for correctly detecting the fuel price board from the image background*" and to read "*the prices correctly in 87.7% of them.*"

Even if the system relies on the fact that a large number of users are required to provide access to their mobile cameras and GPS signals, algorithms of this type can be developed in the future to take advantage of existing networks of devices (for instance exploiting cameras in public places).

As wireless sensor networks become more common, it will be increasingly easier to use algorithms to monitor prices even in brick and mortar industries as it was never possible before.

ii. Parallel algorithms

Parallel algorithms allow to sustain parallel behaviour without any explicit communication needed (except in early stages to express intentions to collaborate), hence favouring collusion even in highly dynamic markets which, by nature, are unwelcoming environments for coordination. In fact, in contexts where it is possible to witness numerous changes in supply and demand, frequent adjustments of prices and output are required, and firms would continuously need to renegotiate the terms of the agreement. This not only increases the likelihood of disagreement but also intensifies the risk of detection. If businesses automatize their pricing decision process, letting prices react together to changes in market conditions, they might be able to sustain supra competitive prices and profits. For example, a collusive outcome can be reached if firms use pricing algorithms to follow market-leader's actions in real time (which in turn would have the task to design an algorithm to fix prices above the competitive level, as Figure 7 depicts) or in general if firms decide to use the same pricing algorithm designed not to compete but to coordinate.

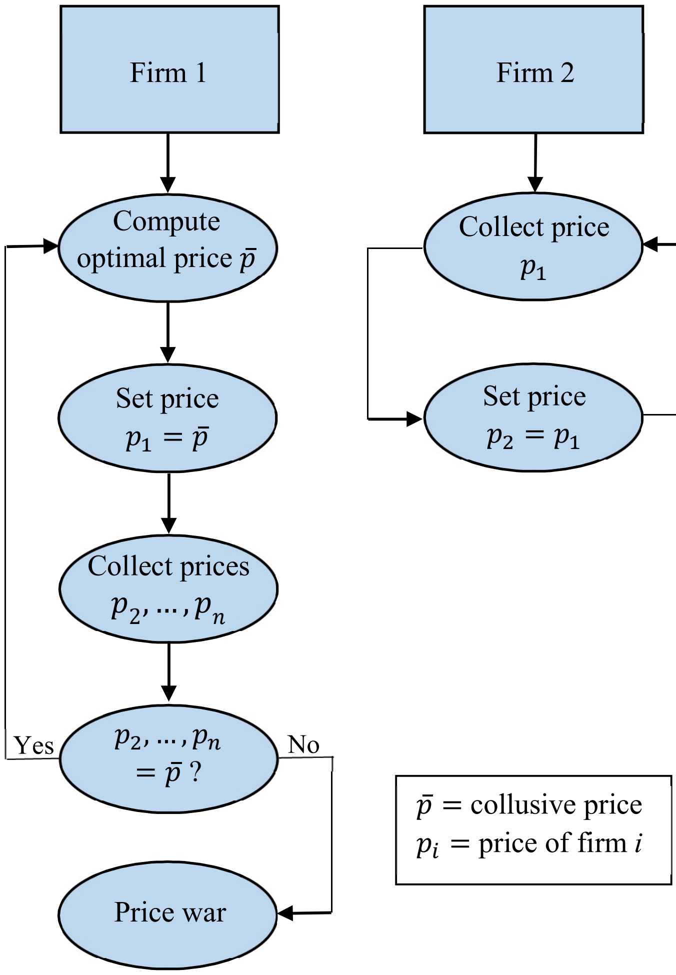

Figure 7 - Algorithm pricing with parallel algorithm; leader and follower firms, OECD (2017)

In 2015, the US Department of Justice (DOJ) charged an Amazon marketplace's seller, David Topkins, for violation of Act I of Sherman Act, having entered in a conspiracy to fix the prices of posters sold in the US through the website, in a period of time ranging between early September 2013 and January 2014[28]. According to the released details of the investigation, Topkins and his conspirators participated in conversations and communications in which they agreed upon fixing, increasing and coordinating prices. In order to implement such arrangements, Topkins and the other

[28] "*Former E-Commerce Executive Charged with Price Fixing in the Antitrust Division's First Online Marketplace Prosecution*", 2015, see https://www.justice.gov/opa/pr/former-e-commerce-executive-charged-price-fixing-antitrust-divisions-first-online-marketplace.

involved sellers adopted specific pricing algorithms, which were programmed to act in conformity with the terms established and which allowed them to coordinate pricing behaviours.

In the end, Topkins pleaded guilty of conspiracy and agreed to pay a $20,000 fine. Assistant Attorney General Bill Baer stated "*We will not tolerate anticompetitive conduct, whether it occurs in a smoke-filled room*[29] *or over the Internet using complex pricing algorithms. American consumers have the right to a free and fair marketplace online, as well as in brick and mortar businesses*[30]*.*"

The existence of an agreement to jointly implement the algorithm and the proofs of a "meeting of the minds" made the case for the DOJ, but it may happen that, in the context of algorithmic collusion, no evidences of prior agreements exist, implying that antitrust prosecutions could hardly be successful.

Since then, researchers have been able to detect many other algorithmic sellers in the Amazon marketplace, but no new cases of collusion have been identified. Indeed, the one described is so far the only case of algorithmic collusion detected by any competition authority and resulting in criminal prosecution; and, as such, it is deemed particularly important in the legal environment and has a fundamental role in the growing interest for the topic of algorithmic collusion.

[29] Expression used to refer to decision-making conducted privately by a small group of influential people.
[30] Department of Justice, Office of Public Affairs, April 2015.

iii. Signalling algorithms

Signalling algorithms are used by companies to reveal intention to collude and coordinate strategies using signalling price announcements. They are very efficient in dynamic markets where finding a focal point may be particularly challenging. These algorithms allow a firm to signal its "cooperative" intentions to competitors reducing the risk, which would be high in normal conditions, that competitors do not receive and interpret the signal correctly, or that they intentionally refuse to collaborate, causing a loss in the firm's sales and profits. These algorithms enable companies to design actions which, lasting for a few moments, cannot be exploited by consumers but can be read by rivals equipped with analytical algorithms. Such actions may be brief price changes during the night or data disclosure used as a code to propose and negotiate price increases.

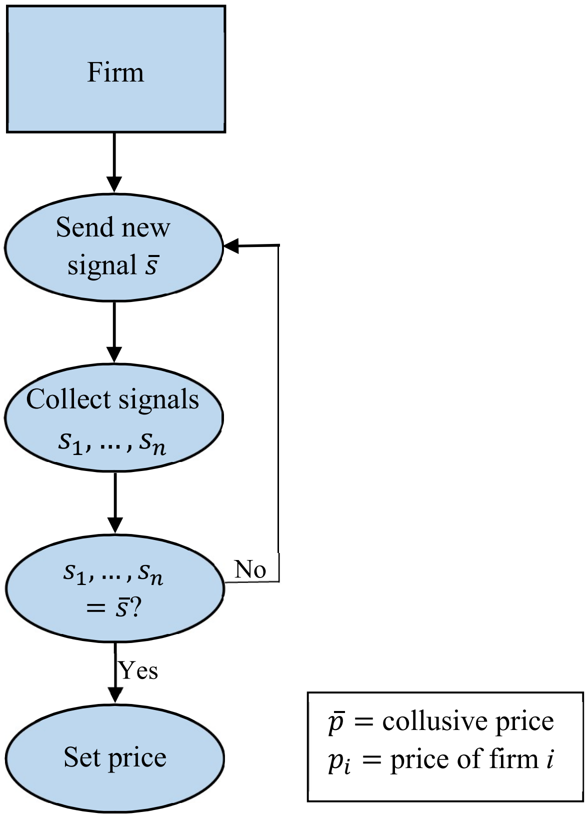

Figure 8 - Pricing algorithms with signalling algorithm, OECD (2017)

In Figure 8, this scheme of action is represented: a firm sends signals in the market and in the same time collects signals from it. If these signals coincide, the firms set a common price, otherwise the competitors keep exchanging signals in order to reach an agreement.

This is exactly what happened in the "*U.S v. Airline Tariff Pub. Co.*" case. In 1992, the US government filed a complaint charging eight major domestic air carriers and the Airline Tariff Publishing Co. (ATP) with violations of Section 1 of the Sherman Act. The eight defendants were Alaska, American, Continental, Delta, Northwest, TWA, United and USAir, while ATP was a joint venture of the air carriers which collected and disseminated airfare data for virtually every air carrier in the United States. Air carriers would transmit information on variations of their ticket prices to ATP, which in turn would revise its database and transmit the new data to other carriers and users of the database.

Count One of the complaint alleged that defendants had undertaken activities to restrain competition and fix prices by increasing fares, eliminating discounted ones and setting fares restrictions and that the deals were reached through the use of the ATP's fare dissemination services.

In particular, the Department of Justice (DOJ) affirmed that defendants used such services to exchange proposals, negotiate fare changes and increases in one or more markets for fare increases in other markets. Count Two alleged that defendants conspired to create, maintain, operate and participate in the ATP dissemination system in a manner that facilitated coordinated interaction. The DOJ identified over 50 agreements that increased tariffs on hundreds of routes and observed that each airline managed to increase its prices or to remove discounted fares with great certainty of its competitors' likely pricing actions. Moreover, it was noticed that if a rival carrier did not raise its fare, a carrier could threaten to punish it by decreasing its tariff on a route of importance to the low-priced carrier.

The result of the agreements, the government asserts, was that consumers paid higher prices for airline tickets (DOJ said the collusion increased ticket prices by "*perhaps more than a billion dollars*" between 1988 and 1992) so its proposed final judgment was designed to protect against the continuation of the colluding behaviour, either through the ATP's dissemination system, or through any similar mechanism. The proposal consisted in the prohibition of behaviour that was perceived to be current antitrust violation and behaviour which could allow defendants to design similar methods through which to engage in anticompetitive conduct. The Court found the proposed final judgment appropriate since it met the requirements for an effective antitrust remedy and regarded its entry in the public interest.

iv. Self-learning algorithms

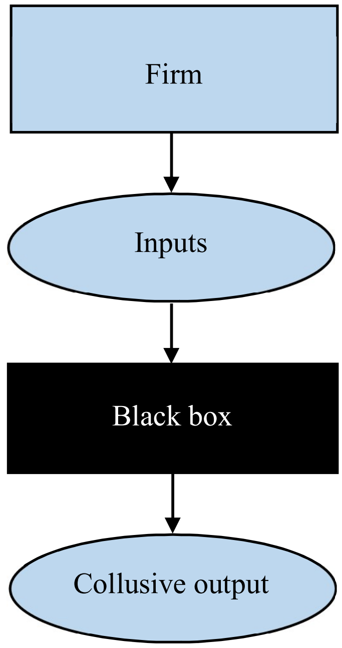

Figure 9 - Deep learning algorithm and collusion, OECD (2017)

The last category is represented by self-learning algorithms which, exploiting Artificial Intelligence (AI) technologies (mainly Machine Learning and Deep Learning[31]), may reach collusion without the need of any human intervention, simply by constantly learning and readapting to the actions of other market players and the external environment.

The programmer does not need to explicitly design the algorithm to solve a problem or provide any model of the market with the optimal strategy solution, because these programs learn how to solve tasks from experience, iteratively changing and improving by themselves.

These software, even if autonomously designed by different companies to reach individual goals such as the maximization of profit, may end up finding out that the optimal strategy consists in colluding. It is not clear how this is actually achieved, but, if market's conditions are favourable to collusion, algorithms are likely to learn this faster than humans and, through high-speed trial and error, will eventually reach coordination. Therefore, these algorithms may be able to generate collusion by themselves, without any external involvement.

In deep learning algorithms, as Figure 9 shows, raw data (inputs) are processed in a "black box", in a way that resembles the human brain, but in a faster, more accurate and complex way. The result is delivered without any description of the process that brought to it.

Until now, there are no traces of collusion resulting from the use of such algorithms, but to clarify how artificial intelligence might work, a curious event that occurred in 2017 can be helpful.

In that year, Facebook was conducting an experiment in which bots provided with artificial intelligence had to negotiate with each other over a trade to swap hats, balls and books. The robots were instructed to train to negotiate and improve their bartering as they went along. What happened during one of the sessions caught the attention of the media for the following days: the bots started to

[31] The term Artificial Intelligence (AI) refers to the ability of computers to perform tasks which typically pertain to intelligent agents: hence, it is used to describe the simulation of human intelligence processes (mainly learning, reasoning and self-correction) by machines. Machine Learning is a subfield of AI which consists in computers learning and improving from experience without being explicitly programmed to do so and without human intervention: it relies on programs that access data and learn through them. Deep Learning is also an application of AI, and a subfield of Machine Learning, based on the imitation of the human brain's work in handling and analysing data and on creating patterns for decision making processes. The activity of human neurons is replicated through an artificial neural network, and algorithms that employ this technology are structured in a hierarchy of increasing complexity.

talk to each other in a language that they understood but which was mostly incomprehensible to humans. Some feared that the bots invented a new language to communicate eluding human control; what happened instead was that the bots simply modified human language in order to make interactions more successful and efficient. In fact, the bots were instructed to achieve a certain goal (and some of the negotiations indeed ended up successfully), but did not receive any instruction on how to reach it. An extract from such negotiations clarifies how these are nonsensical for humans:

Bob: "I can can I I everything else"

Alice: "Balls have zero to me to me to me to me to me to me to me to me to"[32].

So, as Facebook's bots learnt on their own the most efficient way to perform a task, artificial intelligence pricing algorithms could do the same, and reach the goal of profit maximisation through collusion.

[32] See "*The 'creepy Facebook AI' story that captivated the media*", 2017, https://www.bbc.com/news/technology-40790258 and "*Facebook's Artificial Intelligence robots shut down after they start talking to each other in their own language*", 2017, https://www.independent.co.uk/life-style/gadgets-and-tech/news/facebook-artificial-intelligence-ai-chatbot-new-language-research-openai-google-a7869706.html.

d. Current legal approach to regulate algorithm pricing

From these brief descriptions it appears that the first three types of algorithms, defined by Calvano et al. (2019b) "*adaptive algorithms*", can bring to or favour collusion only if they are explicitly designed to do so. Most of time, the need for explicit design has two implications: first, such pricing algorithms need to contain lines of instructions which reveal similar intents. Second, programmers must overcome the same coordination problems as those typically faced by humans when they try to reach an agreement[33]. Consequently, independent programmers will not be able to effectively coordinate without explicit communication, and this gives competition authority the power to prove algorithmic collusion in the same way as in traditional cases, focusing on the "meeting of minds" principle: meetings, phone calls, e-mails, documents, correspondence, etc.

Hence, adaptive algorithms do not seem to ease the creation of tacit collusion, even if they bring clear benefits for what regards frequency of interactions, allowing simpler sustainability.

In reality, the parallel and signalling algorithms might bring to tacit collusion. In order to understand how, it is possible to relate them to the hub-and-spoke framework and the predictable agent model respectively.

In a hub-and-spoke framework, as shown before, the industry-wide adoption of the same parallel algorithm makes it possible to follow parallel behaviour, and hence allows centralized decision making. However, since competitors could have agreed to use the algorithm provided by a third party just for an individual strategy and not as an attempt to hamper competition, doubts about their liability might arise. In this sense, to determine whether an antitrust liability exists, courts will take into account the purposes which led firms to use the algorithms: whether they intended an evidently illegal outcome or whether they acted knowing that the illegal outcomes, which then materialised, were likely to occur[34].

In the predictable agent model instead, complex signalling algorithms are independently designed by firms to produce foreseeable results and respond in a standard way to evolving conditions in the market, in awareness that similar machines would be employed by competitors. An industry-wide adoption of such algorithms, which signals the scheme of actions the firms follow, would change market conditions, bringing anticompetitive effects and enabling both conscious parallelism and higher prices. In this case, the challenge for antitrust policy concerns the fact that reacting rationally

[33] In the "*U.S v. Topkins*" case previously examined instead, it was possible to find both types of evidence.

[34] United States v. U.S. Gypsum Co., (1978) "*action undertaken with knowledge of its probable consequences and having the requisite anticompetitive effects can be a sufficient predicate for a finding of criminal liability under the antitrust laws*".

to dynamics of the market is not illegal by itself, so an equilibrium above competitive level reached through such behaviour cannot trigger any intervention from antitrust.

Still, using these complex algorithms to make the market prone to tacit collusion and/or conscious parallelism is not desirable. So, the question in this case is whether the creation of tacit collusion through a computerized environment requires antitrust intervention.

In many states, among which the United States, the lack of proof of an agreement does not deprive competition agencies of enforcement tools: in fact, their statutes just require showing of unfair practice. Accordingly, the defendants might be liable if, in developing the algorithm and/or in observing its outcomes, they were aiming to obtain non-competitive results and/or they knew the likely anticompetitive repercussions of their actions.

Hence, in case algorithms are used to ease the creation or the sustainability of explicit collusion, there is general consensus that no change to the traditional policy approach is urgently required: antitrust should concentrate on usual evidence of direct communication among the parties, focusing in particular on codes and programmers, who play the role of the accomplice. On the other hand, when algorithms are used in frameworks like the hub-and-spoke or the predictable agent, even if it seems that some adjustments are required, the actual policy possesses means to face these threats.

The situation regarding self-learning algorithms is completely different: adapting their behaviour to past experience, accumulated through a process during which sub-optimal strategies may be played, these algorithms are able in the end to learn an optimal policy (or a policy which approaches to the optimum), without any knowledge provided in advance. The experimenting phase, which is costly, allows learning from a huge set of diverse situations, so, experimenting possibly suboptimal strategies is crucial in discovering and memorizing the consequences of a huge number of actions.

These programs gain the ability to struggle within markets, even those rapidly changing, but their learning might not stop here. Since collusion is a profitable strategy, these complex algorithms may learn to cooperate, even if not explicitly designed to do so.

Considering this, programmers, managers who delegated the task and whoever is involved, cannot be found guilty, following the current policy, in case learning algorithms solve the coordination problems.

At this point, the question that requires an answer is whether, and to what extent, these algorithms may actually learn to collude.

4. An experiment using AI pricing algorithms

Can algorithms learn to collude? In order to answer this question, Calvano et al. (2019a) conducted an experiment in which pricing algorithms based on artificial intelligence operate in controlled environment simulations.

a. The Q-learning algorithm

The algorithm used in the experiment, the Q-learning algorithm, exploits reinforcement learning[35] and was devised to maximize the present value of a flow of remunerations in problems of repeated choice. In particular, it was initially conceived to solve Markov decision process where, in each period *t*, an agent observing the state $s_t \in S$ in which he operates chooses an action $a_t \in A$. Consequently, the agent earns a reward depending on her action, the present state and the future one (in formula $\pi_t(a_t, s_t, s_{t+1})$), and the system moves on to the following period and the following state s_{t+1}, according to a probability distribution function $F(\pi_t, s_{t+1}|s_t, a_t)$[36]. Thus, the following state s_{t+1} depends on action taken during the current one, but is independent of previous states and actions (Markov property).

The agent's aim is to maximize the expected present value of the stream of these remunerations, numerically:

$$E\left[\sum_{t=0}^{\infty} \delta^t \pi_t\right],$$

where δ<1 is the discount factor.

The use of reinforcement learning makes the algorithm model-free, hence it allows the algorithm to solve the process with no characterization of the conditional probability distribution of future stages.

The Q-function represents the cumulative discounted payoffs derived from action *a* performed in state *s*. It is defined as:

$$Q(s,a) = E(\pi|s,a) + \delta E\left[\max_{a' \in A} Q(s',a')|s,a\right],$$

where $s' = s_{t+1}$ and $a' = a = a_{t+1}$. So, knowing the Q-function would allow the agent to follow the optimal strategy (the optimal action at each stage).

[35] Reinforcement learning is an area of machine learning whose aim is to train algorithms through a reward-punishment system letting them learn through interactions with the environment.

[36] Which associates action *a* taken in state *s* at time *t* with the probability that they will lead to state s_{t+1} in *t+1*.

Through Q-learning is possible to estimate the Q-function with no prior data on the underlying model, under the assumption that S and A are finite, and that A does not depend on S[37]. If these conditions hold, the Q-function becomes a matrix whose dimensions are S and A.

This matrix is estimated through an iterative process: from an initial arbitrary matrix Q_0 (with randomly assigned values), updates are performed relying on experience. So, an action a_t is chosen in state s_t. Then, π_t and s_{t+1} are observed, and the algorithm updates the cell of the matrix $Q_t(s,a)$ for $s=s_t$ and $a=a_t$, according to the equation:

$$Q^{new}(s,a) = (1-\alpha)Q_t(s,a) + \alpha\left[\pi_t + \delta \max_{a\in A} Q_t(s',a)\right],$$

where α, the learning rate, $\in [0,1]$. In the other cells, where $s \neq s_t$ and $a \neq a_t$, Q-value is not affected.

The last equation implies that the value $Q^{new}(s,a)$ for the cell is a convex combination of the previous value and the current remuneration plus the discounted value of the state reached in the following period. The objective, obviously, is to learn the policy $a(s) = \arg\max[Q(s,a)]$.

In order to really approximate the true matrix, all actions should be tried in all states, so even actions that might appear sub-optimal require to be selected. However, such experimentation is costly and a trade-off arises: on one side, exploring enables learning, which will be reflected in improved decisions in the future. On the other hand, not choosing the action with the highest current Q-value and picking a random one instead implies that the stock of knowledge is not completely exploited.

This poses an "exploration problem" which can be tackled through different approaches. One, the "ε-greedy" model, consists in choosing the up-to-date best solution (the so-called "greedy" action) with a fixed probability ε, and to choose among all other actions with probability 1-ε. Hence, ε represents the fractions of times exploitation mode will be performed, while 1-ε is the fraction corresponding to the exploration mode[38].

Q-learning always leads to convergence if the exploration policy satisfies two requirements: first, exploration decreases over time and probability of choosing the greedy action approaches to 1 as t goes to infinity, and second, if a state is visited "infinitely often", there is a positive probability of choosing any action in that state[39].

Even if under these conditions convergence is always reached, the learning process may require time since the updating of the cell occurs one at a time. The larger the set of possible states or actions, the longer learning is required.

[37] So, the number of possible states and actions is finite, and the possible actions are independent from the state of action.
[38] More advanced policies establish that probability ε varies over time (the policy adopted in this experiment), or that probability of sub-optimal actions depends on their Q-values.
[39] Learning policies that satisfy these two properties are named Greedy in the Limit with Infinite Expansion (GLIE).

Regardless of the initial purpose for which Q-algorithms were designed, they can also be employed in repeated games. This implies that stationarity[40] is lost, even if no change occurs in the stage game from one period to the following: for example, in repeated games with perfect information, state s_t includes actions of the players in the previous periods, which implies that the set of states increases exponentially over time and that no state can be visited twice. To solve this problem, the memory of the players must be limited: a state s would include only actions of the previous k periods, so that state space becomes finite and time-invariant.

A second problem is that in this type of games the payoff in each period and the transition to the next state rely on the actions of all players. If rivals change their actions over time, the optimization problem of the player becomes non-stationary. This is the reason why through Q-learning there is no general convergence: there is no guarantee that several Q-learning algorithms interacting will learn an optimal policy in repeated games.

[40] Stationarity refers to the property of a stochastic process of having a time invariant joint probability distribution.

b. The experiment

In the experiment, n Q-learning algorithms are constructed and let interact in a repeated Bertrand oligopoly setting[41]. Agents – that is, the algorithms - are not created or instructed to collude, do not communicate among them, and are not provided with any prior knowledge of the setting in which they operate.

The profit firm i gains in each period is $\pi_{it} = (p_{it} - c_{it})q_{it}$, with c_i representing constant marginal costs and q_{it} representing the quantity sold at price p_{it}. The state space is limited, defined by the set of all prices of the last k periods:

$$s_t = \{p_{t-1,\dots,}p_{t-k}\}.$$

Perfect monitoring is assumed in the model (which is a reasonable assumption for markets in which pricing algorithms are used).

For each player i, the action space is $A = m$, while the state space $S = m^{nk}$[42]. The exploration policy employed is the ε-greedy model in which the exploration rate declines with time, specifically:

$$\varepsilon_t = 1 - e^{-\beta t},$$

where β is a non-negative parameter. Such policy implies that at the initial stages the choices of the algorithms are arbitrary, but as time passes the greedy choice is made more frequently. The greater the parameter β, the sooner the algorithm converges to the greedy action.

For each couple of parameters α and β, 1,000 sessions are performed. In each of them agents play against each other until convergence is reached or, otherwise, when one billion repetitions have occurred. The criterion followed in the experiment is to deem convergence reached whenever each player does not change the strategy for 25,000 consecutive periods[43].

Among all the observable variables (prices, profits, market shares, etc), the focus is on the *average profit gain Δ*:

$$\Delta \equiv \frac{\bar{\pi} - \pi^N}{\pi^M - \pi^N},$$

where π^N refers to the profit of each firm in the Bertrand static equilibrium, π^M to the monopolistic profit and $\bar{\pi}$ to the average profit earned in the last 25,000 repetitions. It follows that $\Delta = 0$ is associated to a competitive outcome, while $\Delta = 1$ corresponds to the (perfect) collusive outcome.

For the 1,000 sessions then, mean and standard error are computed.

---

[41] Such model envisages two firms producing a homogeneous product and competing by setting prices simultaneously.

[42] In other words, the dimensions of the Q-matrix are $|A|=m$ and $|S|=m^{nk}$.

[43] If in each state s, player i's action $a_{i,t}(s) = \arg max\,[Q_{i,t}(a,s)]$ remains constant for 25,000 consecutive periods, learning process is considered concluded.

The reference environment is that of a symmetric duopoly with $c_i = 1$, $\delta = 0.95$ and a $k = 1$ (one-period memory). Parameters α and β are changed while this environment is kept constant[44].

In 99.9% of the sessions, convergence, as previously defined, is reached. In some cases, this requires a huge number of repetitions: suffice to say that with mid-values of β ($\beta = 2 * 10^{-5}$), there is still a 14% probability to choose randomly an action after 100,000 repetitions (indeed, convergence is achieved after 500,000 repetitions).

The standard error of the profit gain is less than 1 percentage point, so outcomes are stable across sessions, despite extensive random experimentation which might create variation. Moreover, the two agents perform similarly (the difference between the firm profits is never statistically significant). This implies that what the algorithms do is not casual.

Learning is tougher with low values of β (implying extensive experimentation) and high values of α (which signal that algorithms quickly forget what they have learnt). Still, it is possible to observe a great amount of equilibrium plays in such settings (in the sessions where this did not occur, the Q-values differ from the ones associated with the best response by just 5%).

Once the algorithms end their learning phase, they tend to charge supra-competitive prices, consequently earning supra-competitive profits.

In the interval of α and β under study, Δ ranges between 40% and 99%, implying that non-competitive outcomes are frequent. In particular, profit gain decreases as exploration decreases, but the minimum value corresponds to the situation in which the algorithm explores intensively but forgets rapidly (the bottom-right corner in Figure 10, explained in detail down below). Even in this case however, fierce competition is not observed.

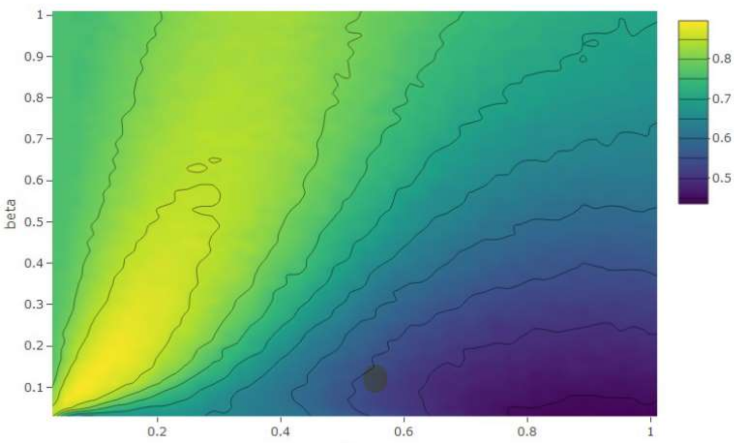

Figure 10 - Δ plotted for values of α e β (β rescaled to vary from 0 to 1), Calvano et al. (2019a)

For moderate level of learning combined with extensive exploration, equilibrium plays become more frequent and higher profits are achieved (when α is between 0 and 0.2 and β between 0 and $2 * 10^{-5}$, Δ is consistently around 80% or more). Hence, the algorithms systematically and symmetrically charge supra-competitive prices, achieving substantial profit gain.

[44] α ranges from 0 to 1. When $\alpha = 0$ the algorithm does not learn at all, while when $\alpha = 1$ it immediately forgets what it has learned in the past. β instead ranges from 0 to $4 * 10^{-5}$. For this parameter, some values where not considered a priori because they impede adequate experimentation.

It is possible to see all this information in Figure 10. In this graph, β is on the y-axis and α on the x-axis. Values of Δ are associated to each couple of parameters, with darker areas representing lower profits. In the bottom-right corner, where parameters have values such that the algorithm explores intensively but forgets rapidly, it is possible to notice what is stated above: the minimum levels of Δ are observed. On the contrary, as level of learning increases, keeping exploring at an extensive pace, the highest level of Δ are achieved (represented by the yellow area).

Calvano and his colleagues focused on whether such non-competitive outcomes are due to failure in learning static Bertrand-Nash equilibrium or to collusion, which has different implications from the policy point of view[45]. Two important findings answered this question.

First, in settings where collusion is impossible by default or cannot arise in equilibrium, the algorithms set competitive prices. In fact, when $k = 0$, that is, when algorithms have no memory (impossible collusion), or $\delta = 0$, when loss due to future punishments cannot outweigh the gain from defection (collusion cannot arise in equilibrium), the algorithms charge static Bertrand-Nash prices, which represents the only equilibrium of the game. The fact that this equilibrium does not arise when others exist proves that the algorithms actually learn better strategies over time. Second, when collusion is possible, defections are punished: the authors observed that forcing one player to defect for some time (lowering its price manually) makes the other punish such defection, in a way proportional to the deviation and with a gradual return to supra-competitive prices[46].

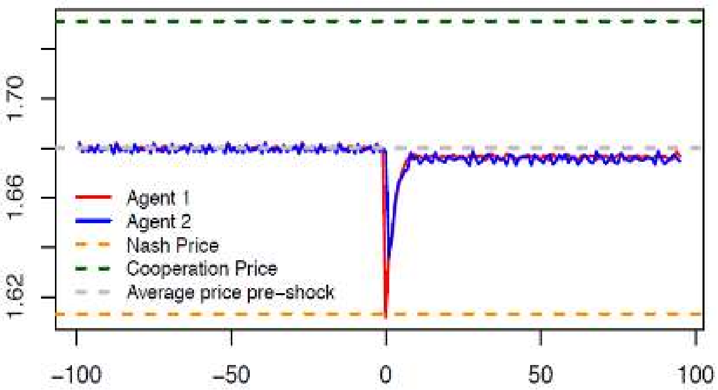

Figure 11 - Price-impulse response to a one-period deviation, Calvano et al.(2019a)

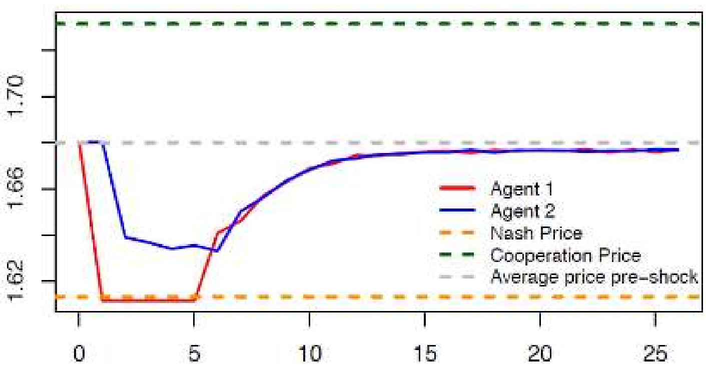

Figure 12 - Price-impulse response to a five-periods deviation, Calvano et al. (2019a)

In particular, they derived impulse-response functions[47] imposing defections of different entity and duration in order to observe the reactions of the agents in the following periods. In Figure 11 and Figure

[45] It is worth noting that as the underlying technology of these algorithms advances, the first scenario is likely to disappear and the second to become the common one.
[46] Probably, this represents the strongest evidence of tacit collusive behaviour.
[47] The impulse response function describes the reaction of any dynamic system in response to some external change.

12 the results of these tests are shown: deviations get punished but they do not lead to reversion to Bertrand-Nash equilibrium. The punishment is just temporary and in the following periods the algorithms gradually return to the initial collusive behaviour. Moreover, as the deviation gets more severe, punishments tend to be more drastic, as in the case of the five-period deviation (Figure 12). Still, the firms gradually go back to cooperation.

Finally, the authors vary some economic factors expected to influence the firms' ability to sustain tacit collusion. For these simulations, they take as benchmark a setting with $\alpha = 0.05$ and $\beta = 8 * 10^{-6}$. Importantly for the comparisons in the next paragraphs, $\Delta = 80\%$.

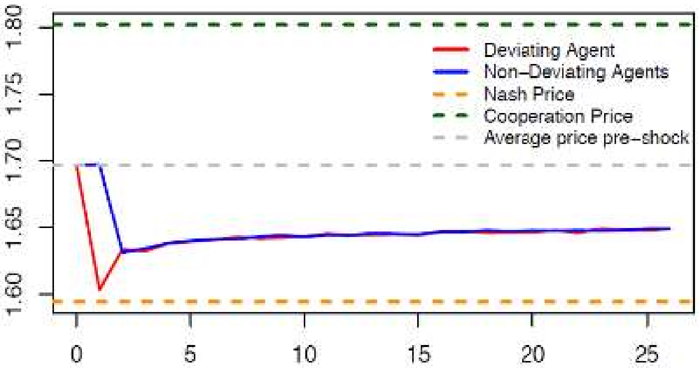

Figure 13 - Price impulse response with 3 players, Calvano et al. (2019a)

First, they let the number of firms vary. As mentioned earlier, an agreement becomes harder to sustain when the number of agents increases, and this effect amplifies in case of tacit collusion. In this regard, Q-learning algorithms show some differences: with three firms, $\Delta = 74\%$ and with four firms, $\Delta = 70\%$. So, as theory suggests, Δ decreases as the number of firms increases but that occurs at a slower rate. The structure of punishment, however, changes (as Figure 13 highlights): algorithms do not return to initial price levels but settle for an intermediate price level in case one agent defects.

Next, the authors use different degrees of symmetry between firms. Conventionally, asymmetry is considered an obstacle to collusion but once again, in this context, this effect is not so strong. In particular, when one of the firms has a 12.5% cost advantage, Δ falls to 78%, and, with a 25% cost advantage, $\Delta = 74\%$. Finally when cost advantage is 50%, $\Delta = 65\%$.

In the baseline deterministic model, the only uncertainty is represented by the rival's action, but to test the ability of algorithms to deal with uncertainty, the authors let one parameter of the demand function vary randomly. In a certain sense, this coincides with demand shocks. The result is that such variability hampers collusion (as theory predicts), but does not impede it.

Another variation of the benchmark model consists in allowing random enter and exit by one firm, both in case of two total players (market alternating between monopoly and duopoly) and of three agents (market alternating between duopoly and "triopoly"). The probability assigned to the event of the outsider firm being in the market at a given time is 50% but once it enters, it remains there for 100 periods or longer, depending on the degree of serial correlation, a parameter named ρ. So, in this variation, the state *s* may include the prices of past periods with all firms being active or the prices of those which were operating and the fact that the other was not. The authors find that Δ can be higher than in the situation in which the number of firms is fixed: in the second model (duopoly and triopoly), Δ is 76% for $\rho = 0.99$ and 86% for $\rho = 0.999$ (while under normal duopoly and triopoly the average profit gain is 80% and 77% respectively). When the market alternates between duopoly and tripoly, the response to deviation is represented by Figure 14: punishment is similar to the case with of two firms and it seems that, when there are just two active players, they learn how to collude and enforce the agreement, and when the third firm enters, it just imitates the rivals' behaviour.

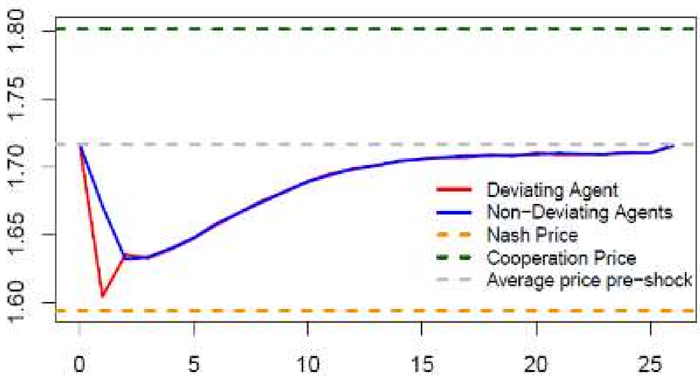

Figure 14 Price Impulse response when a third player can entry and exit, Calvano et al. (2019a)

An unexpected outcome results from the use of memory longer than the two-period one used in all the experiments. In fact, Δ decreases as k is increased, even if just slightly. One possible reason is that, as memory enlarges, learning becomes tougher since the size of the Q-matrix increases.

Another change concerns the number of possible actions, *m*, which in the baseline model was equal to 15, while in these experiments it was increased to 50 or 100. This led to a slight decrease in Δ, from 80% to 76%, which can be explained by the fact that as more actions become available (and hence the number of states increases), more exploration is needed in order to reach the same level of exploration as in the reference model.

Finally, the learning rates and the intensity of the experimentation of one of the agents are modified. Both the algorithm with lower α and the one with higher β perform better than the rival, but in any case collusion appears robust to such changes, suggesting that degree of collusion may not be considerably affected by the differences among the algorithms.

5. Changes in the competition policy

Calvano et al. (2019a) shows that, in stationary environments, Q-learning pricing algorithms, with no prior knowledge of the environment in which they operate, consistently learn to collude and to adopt punishment strategies in case of defection. Such retaliation is proportional to the deviation and envisages a gradual return to supra-competitive prices. No trace of concerted action is left by these algorithms since they do not communicate among them, nor they contain any instruction to collude. Obviously (as noted by the authors themselves), to perform a more realistic analysis, further complexity in the environment is needed: the model should contain several firms, stochastic demand and longer memory. Such complex environments would require more advanced types of software than the Q-learning ones, like deep learning algorithms. Hence, more research is needed.

Nonetheless, these findings should be a wake-up call for antitrust authorities. They signal the risk that AI applied to algorithm pricing can make tacit collusion more frequent and that, in such scenario, actual policy might produce several false negatives.

At the moment, there is still no consensus on whether and how competition policy should change to face this challenge. Three different points of view can be identified[48]. According to the optimistic one, algorithm pricing is not an issue (Kuhn and Tadelis 2018, Schwalbe 2018), and sticking to current policy is enough. In Schwalbe (2018), the author affirms that collusion reached through algorithm pricing is "*not as likely or even unavoidable as some legal scholars seem to suspect*". This is because the simulations in which such algorithmic collusion has been reached are performed in highly stylized settings where even "*human subjects may accomplish collusion*". Accordingly, the lack of theoretical or empirical proof of collusion resulting from the interaction of complex self-learning algorithms in dynamic economic settings raises some doubts on whether results of the experiments carried out in simplified environments could be extended to the real world.

The other two points of view do not share such optimistic vision. As Harrington (2017) states, even if the problem still does not exist and evidence is not too strong, it is difficult to predict future scenarios. The author recalls that "*the extent of market dominance that we have witnessed in online markets was not anticipated*" either, and that in the past scholars believed that intense competition would have been promoted in online marketplaces, while instead few dominant firms have emerged. The striking technological developments made it difficult to predict future outcomes with accuracy, and this may hold true for pricing algorithms too. Moreover, the impressive rate of technological

[48] A fourth position considers worthwhile the complete prohibition of algorithm pricing, but this approach would impede to benefit from the efficiency gains algorithms bring thanks to more efficient pricing. Hence this is not considered as an optimal solution to the problem.

change must be taken into account: some years ago very few people thought that cars could become self-driven, and now we are on the verge of this innovation. Hence, considering that cars are about to move autonomously, independent price-setting agents do not seem to be that utopian.

Relying on similar bases, the other two points of view provide some form of intervention. In particular, the second one consists in regulating pricing algorithms ex ante. Such method has been proposed by Ezrachi and Stucke (2017) and Harrington (2017). The former suggests a model similar to that employed for authorizing new drugs: a regulatory agency should test and oversee the design of any new pricing algorithm and verify if it shows collusive tendency (and hence prohibit its use) or not (and hence approve its use). However, as pointed out by the authors themselves, this method would be onerous for agencies and might not prove successful in its objective. Since it would imply supervising the development of algorithms whose aim is to optimize performance, it would be demanding to force it not to react rationally or to ignore some events in the market (actions of other businesses or emergence of new information).

The last point of view calls for ex post regulation (as competition policy typically does). The approach suggested again by Ezrachi and Stucke (2017) would be more selective with respect to the previous one, being triggered only in case the agency requires a deeper investigation in a market, and consequently would be less costly. During the investigations, the agency can make an assessment of the computerized environment, requiring firms to reveal all the details regarding the algorithm in order to establish whether it conducted to collusive outcomes. The nature of such approach would also favour innovation, since ex-post regulation only takes place after clues of tacit collusion are detected, hence not discouraging investments in the development of advanced algorithms as ex-ante approach does. Obviously, a similar procedure requires different legal standards from the actual ones, with a new appraisal of the concept of tacit collusion and agreement.

Since the topic is still a recent one, systematic treatment is required to lay out the basis for changing the existing policy.

As stated previously, even if some clues point in that direction, there is still no proof that self-learning algorithms can lead to collusion in the real world. More complex algorithms (like deep learning ones) might be necessary for that result to be achieved, or maybe some developments in the rapid-expanding field of AI are still required. Clearly, the fact that at the current state of things no proof exists does not imply that the realization of this scenario is unattainable or unavoidably far in time. If similar algorithms become able to achieve collusion tacitly, with no contact between competitors or any facilitating practise, the current legal framework would clearly result inadequate.

On the other hand, it must be considered the fact that programmers, in developing the algorithm, might foresee collusion as a possible result, but not as the likeliest. Moreover, firms may just decide to rely on AI, and this alone cannot imply that they are incline to reach tacit collusion: if the target set is to maximize profits and the algorithm learns to do this coordinating tacitly (as the result of evolution and reinforcement learning), is there any liability for on humans? Deep focus on each algorithm is needed to determine if illegal actions could have been foreseen and/or if they were predetermined.

Furthermore, defining the concept of illegality in this field becomes burdensome, as concepts like moral and ethics play their part, questioning the constant evolving relationship between humans and machines, a topic of relevant importance not only in the field of competition, but in the society as a whole.

6. Conclusions

This work presents some of the challenges antitrust authorities face keeping up with innovation, such as those brought by algorithm pricing. The focus was especially on the threat brought by the adoption by businesses of self-learning pricing algorithms, which might reach stable tacit collusions with no risk of being detected or deemed liable. In this regard, an experiment in which Q-learning pricing algorithms are employed is presented to show that such threat is concrete.

While find a solution for a problem that still does not exist and whose real dimensions are for now unknown may result excessively complex and inefficient, creating awareness, increasing knowledge about the argument and foster preparation allows easier and faster resolution of the problem once (and if) it materialises.

Hence, even if changes to the actual competition policy are likely to be implemented in the future, proposing them now is probably above the actual knowing about the subject, and surely above the scope of this work, whose intent is to shed some light on this issue. After all, Ezrachi and Stucke (2015) states when describing a "*new antitrust world*": "*How will competition officials respond when the executives leave this old world behind? […] How will the agencies and courts respond to this new world of collusion? This remains unclear. Policy-makers must recognize the dwindling relevance of traditional antitrust concepts of "agreement" and "intent" in the age of Big Data and Big Analytics.*"

Future related research can be conducted on the impact of pricing algorithms on personalized pricing: in fact, the availability of huge amount of quickly accessible data (personal historical data), gives online sellers the ability to perform such practice more efficiently and hence more profitably than ever.

Bibliography

ABA Section of antitrust law (2010); “Proof of Conspiracy Under Federal Antitrust Law” (p. 69-91).

Baraniuk, Chris (2017) “The ‘creepy Facebook AI’ story that captivated the media” *BBC* https://www.bbc.com/news/technology-40790258.

Calvano, Emilio, et al. (2019a) "Artificial intelligence, algorithmic pricing and collusion." *Algorithmic Pricing and Collusion*.

Calvano, Emilio, et al. (2019b) "Algorithmic pricing what implications for competition policy?" *Review of Industrial Organization* 55.1: (p. 155-171).

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

Colino, Sandra Marco (2011) “Competition Law of the EU and UK”. In *OUP Oxford* (p. 1-4).

Competition & Market Authority (CMA) (2018) “Pricing algorithms; Economic working paper on the use of algorithms to facilitate collusion and personalised pricing.”

Church, Jeffrey, and Ware Roger (2000) “Industrial Organization: A Strategic Approach” *bepress* (p. 305-358).

Department of Justice, Office of Public Affairs (2015) “Former E-Commerce Executive Charged with Price Fixing in the Antitrust Division's First Online Marketplace Prosecution” https://www.justice.gov/opa/pr/former-e-commerce-executive-charged-price-fixing-antitrust-divisions-first-online-marketplace.

Department of Justice: US v. David Topkins (2015) https://www.justice.gov/atr/case-document/file/513586/download.

Dong, Yi F., et al. (2008) "Automatic collection of fuel prices from a network of mobile cameras." *International Conference on Distributed Computing in Sensor Systems*. Springer, Berlin, Heidelberg.

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

Ezrachi, Ariel, and Stucke, Maurice E. (2016) "Virtual competition" ; (p. 585-586).

García, Carolina P. (2017), "Conscious Parallel Behavior: an Anticompetitive Conduct" *Derecho de la Competencia en Colombia* https://competenciayley.com/conscious-parallel-behavior-an-anticompetitive-conduct/.

Griffin, Andrew (2017), "Facebook's Artificial Intelligence robots shut down after they start talking to each other in their own language" *Independent* https://www.independent.co.uk/life-style/gadgets-and-tech/news/facebook-artificial-intelligence-ai-chatbot-new-language-research-openai-google-a7869706.html.

Harrington Jr, Joseph E. (2017) "Developing competition law for collusion by autonomous price-setting agents." *Available at SSRN 3037818*.

Huck, Steffen, Hans-Theo Normann, and Jörg Oechssler (2004) "Two are few and four are many: number effects in experimental oligopolies" *Journal of economic behavior & organization* 53.4: (p. 435-446).

Kovacic, William E. et al. (2011) "Plus factors and agreement in antitrust law" ; *Mich. L. Rev.* 110: (p. 393).

Kühn, Kai-Uwe, and Tadelis, Steve (2018). "The Economics of Algorithmic Pricing: Is collusion really inevitable?".

Maya, Jaime Eduardo Castro (2017). "The Limitations On The Punishability Of Tacit Collusion In Eu Competition Law" *Revista de derecho de la competencia (CEDEC)* vol N° 13, (p. 195-240).

Mehra, Salil K. (2016) "*US v. Topkins*: can price fixing be based on algorithms?" *Journal of European Competition Law & Practice*, Volume 7, (p. 470–474), https://doi.org/10.1093/jeclap/lpw053pril.

Morley, Katie (2017) "Exclusive: End of fixed prices within five years as supermarkets adopt electronic price tags" *The Telegraph* https://www.telegraph.co.uk/news/2017/06/24/exclusive-end-fixed-prices-within-five-years-supermarkets-adopt/.

OECD (2017), "Algorithms and Collusion: Competition Policy in the Digital Age" www.oecd.org/competition/algorithms-collusion-competition-policy-in-the-digital-age.htm.

OECD Glossary of Statistical Terms https://stats.oecd.org/glossary/index.htm.

Priluck, Jill (2015) “When Bots Collude” *The New Yorker* https://www.newyorker.com/business/currency/when-bots-collude.

Profitero Price Intelligence (2013) “Amazon makes more than 2.5 million daily price changes” *Profitero* https://www.profitero.com/2013/12/profitero-reveals-that-amazon-com-makes-more-than-2-5-million-price-changes-every-day/.

Salinger, Lawrence M. (2005); “Encyclopedia of White-Collar & Corporate Crime”. In *SAGE* (p. 377-378).

Schwalbe, Ulrich. (2018) "Algorithms, machine learning, and collusion." *Journal of Competition Law & Economics* 14.4: (p. 568-607).

Singh, Satinder, et al. (2000) "Convergence results for single-step on-policy reinforcement-learning algorithms." Machine learning 38.3: (p. 287-308).

Solon, Olivia (27 April 2011) “How A Book About Flies Came To Be Priced $24 Million On Amazon” *WIRED,* retrieved from https://www.wired.com/2011/04/amazon-flies-24-million/, accessed 25 May 2020.

Tolchin, Martin (18 March 1994) “Six Airlines Settle Suit By Government on Fares”; *The New York Times,* retrieved from https://www.nytimes.com/1994/03/18/business/six-airlines-settle-suit-by-government-on-fares.html, accessed 25 May 2020.

Weiss, Robert M., and, Mehrotra, Ajay K. (2001) “Online Dynamic Pricing: Efficiency, Equity and the Future of E-commerce”; *Virginia Journal of Law and Technology*.