

Developing Competition Law for Collusion by Autonomous Artificial Agents¹

Joseph E. Harrington, Jr.

Patrick T. Harker Professor

Department of Business Economics & Public Policy

The Wharton School, University of Pennsylvania

harrij@wharton.upenn.edu

13 November 2018

Journal of Competition Law and Economics, forthcoming

Abstract

After arguing that collusion by software programs which choose pricing rules without any human intervention is not a violation of Section 1 of the Sherman Act, the paper offers a path towards making collusion by autonomous artificial agents unlawful.

¹This paper has benefited from the comments of an anonymous referee, Emilio Calvano, Michal Gal, Salil Mehra, Pierluigi Sabbatini, Danny Sokol, Maurice Stucke, and especially Louis Kaplow. For useful conversations, I would also like to thank Ai Deng, Aaron Roth, and Rakesh Vohra. I am grateful to Roberto Pardolesi for the opportunity present the paper to his competition law class at LUISS, and to Cristina Caffarra for making it a last minute replacement for a panel on "Algorithmic Collusion" at the 12th CRESSE International Conference on Competition & Regulation. The paper served as the supplemental reading for a panel sponsored by Dechert LLP and the New York State Bar Association's Antitrust Law Section on "Antitrust in High-Speed: Colluding Through Algorithms and Other Technologies" (January 8, 2018). Finally, I'd like to thank Mallick Hossain and Takeaki Sunada for research assistance.

1 Introduction

Consider an online market involving two competitors, referred to as firms I and II, that offer identical products. For example, the firms could be booksellers offering the same line of books. Given the ease with which shoppers can compare prices, price competition is intense. Dissatisfied with the current profit margin, the manager in firm I who is in charge of pricing has decided to adopt a software program to perform the task of setting prices. This program uses all available information including firms' past prices (for firm II's prices are available online), firm I's past sales, the cost of each product, the time of the year, and so on. The software program is not a pricing rule but rather a learning algorithm that experiments with different pricing rules in search of one that yields the highest profits. Due to its complexity, firm I's manager has little or no understanding regarding how the learning algorithm works, nor what pricing rule it may eventually adopt. All she knows is that the learning algorithm is purported by its developer to result in highly profitable pricing. It is the plan of firm I's manager to adopt the learning algorithm, observe how it performs over the next few months, and then decide whether to retain it.

Unbeknownst to the manager of firm I, the manager of firm II has recently adopted a learning algorithm as well. While not necessarily the same one, it has comparable capabilities for experimenting with different pricing rules and adapting its rule in response to information about its performance. As was the case for firm I's manager, firm II's manager finds the sophistication of the learning algorithm beyond his capabilities to comprehend, much less predict what it will do. He adopted it only because he hoped it would generate higher profits. Just as firm I's manager did not know that firm II's manager was soon to adopt a learning algorithm, firm II's manager adopted the learning algorithm not knowing that firm I's manager had already done so. But even if the two managers did know or were to become aware that the other is using a learning algorithm to set prices, it would not be useful information because, once again, the complexity of the programs prevents a manager from effectively using that knowledge.

With the two competitors having their learning algorithms in place, each algorithm experiments with different prices and adjusts the pricing rule in response to how well it performs.

Of course, performance depends on the prices set by the other firm's learning algorithm, as those prices determine demand and therefore profits. Initially, each manager notices that profits do not seem higher and, in fact, appear to be a bit lower. However, eventually, prices start to rise and with that rise in prices comes rising profits. After some time, prices settle down and, to the great satisfaction of the managers, profits are higher than before they adopted the learning algorithms. Each manager views the experiment of using a learning algorithm to be a success, and independently decides to continue to delegate the setting of prices to the learning algorithm.

On the presumption that, prior to the adoption of these learning algorithms, prices were at competitive levels then the new prices are necessarily supracompetitive.² The learning algorithms have managed to adjust their pricing rules until they are using the sort of pricing rule that firms deploy when colluding. Later in the paper, I will offer a definition of a collusive pricing rule but, at the moment, it is sufficient to view a collusive pricing rule as a pricing rule that, when adopted by all firms, results in supracompetitive prices. These learning algorithms have managed to latch onto them and deliver higher profits to the firms.

The objective of this paper is to explore whether the collusion that has emerged from these learning algorithms is unlawful and, if it is not, what would be required to make it unlawful. It is taken as a postulate that it is possible for the simultaneous adoption of learning algorithms by competitors to produce collusion. Rather than collusion occurring with human agents, collusion is occurring with autonomous artificial agents where an autonomous artificial agent (AA) is a software program that carries out a set of operations on behalf of human agents without intervention by human agents and does so with some knowledge of the human agent's objectives. Here, an AA takes account of the objective of maximizing profit (so profit is the measure of performance), chooses prices, and adapts its pricing rule in order to yield higher profits. While the ability of AAs to evolve to pricing at supracompetitive levels has been shown in very simplified settings, it is an open question regarding the ease and extent to which it can occur in settings with a richness corresponding to actual markets. An investigation of that question is left

² In this article, *competitive prices* refer to prices that would prevail in the absence of coordination among firms. Competitive prices can exceed cost because firms may have unilateral market power that allows them to profitably price above cost.

to future research. The focus of this paper is on what society can do to deal with such collusion, should it occur.

Though a new issue on the competition law landscape, collusive pricing through algorithms is rapidly gaining attention, as evidenced by recent speeches by the European Commissioner for Competition and the Chair of the Federal Trade Commission.³ Early contributions are Mehra (2014, 2016) and Ezrachi and Stucke (2015, 2016), and policy papers include Ballard and Naik (2017), Capobianco and Gonzaga (2017), Deng (2017), Ezrachi and Stucke (2017), M. Gal (2017a, b), Mehra (2017), and Okuliar and Kamenir (2017). While these papers touch on some of the issues raised here, they do not offer a path to making collusion by AAs unlawful, which is the primary contribution of this paper. Examining the other side of the market, Gal and Elkin-Koren (2017) look at the use of algorithms by consumers. A future topic for investigation is when AAs are operating on both sides of the market. There is also a literature that examines the regulation of artificial intelligence in a broader set of situations; see, for example, Scherer (2016) and Desai and Kroll (2017). One of the central issues there is ensuring that AAs satisfy fairness (e.g., avoiding discrimination); see, for example, Goodman (2016), Johnson, Foster, and Stine (2016), and Joseph, Kearns, Morgenstern, and Roth (2016).

The paper unfolds as follows. Properly understanding collusion will prove crucial to the legal approach developed here. After handling that task in Section 2, I then review what is *unlawful* collusion in Section 3. Section 4 describes how autonomous artificial agents could collude and Section 5 argues that such collusion is lawful according to current jurisprudence. A legal approach to making collusion by autonomous artificial agents unlawful is presented in Section 6. Section 7 concludes.

³ Margarethe Vestager, "Algorithms and Competition," Remarks by the European Commissioner for Competition at the Bundeskartellamt 18th Conference on Competition, Berlin, March 16, 2017. Maureen Ohlhausen, "Should We Fear the Things That Go Beep in the Night? Some Initial Thoughts on the Intersection of Antitrust Law and Algorithmic Pricing," Remarks by the Acting Chairman of the U.S. Federal Trade Commission at the "Concurrent Antitrust in the Financial Sector Conference," New York, May 23, 2017.

2 Defining Collusion

Collusion is a mode of conduct among firms in a market that is defined independent of the legal regime. Indeed, collusion existed long before competition laws were introduced that sought to make it illegal.⁴ Since the early work of Stigler (1964) and Friedman (1971), economists have developed a theoretical framework which makes clear how collusion differs from competition, and have used that framework to study and understand actual cartels. While that literature is vast and technical, all theories of collusion are rooted in the same simple foundation that I will describe here. In doing so, the focus is on collusion with respect to prices, though the ensuing discussion applies as well to when collusion reduces competition in non-price variables such as capacities, product traits, quality, service, complementary products, advertising, and investment.⁵

At a competitive outcome, each firm's price maximizes its profits given the prices charged by rival firms. While firms are individually doing the best they can, that is not the case from a collective perspective. Competing firms "leave money on the table" in that firms' profits would rise if they jointly increased prices. It is the prospect of all firms being better off that leads them to form a cartel and coordinate their conduct. While firms benefit from a coordinated rise in their prices, consumers are harmed.

The objective of collusion is to raise prices to supracompetitive levels in order to earn higher profits, but collusion is *not* the same as charging supracompetitive prices. Instead of forming a cartel, one (or more) of the firms could lobby the government to impose a price floor above the competitive level. In that instance, supracompetitive prices emerge because of government regulation, rather than any collusive arrangement. Given that supracompetitive prices can be achieved without collusion, collusion is not to be equated with supracompetitive prices. Going

⁴ It was more than four centuries prior to the passage of the Sherman Act in 1890 that the representatives of Pope Paul II and King Ferdinand of Naples entered into a collusive arrangement to restrict competition in the sale of alum (Günster and Martin, 2015).

⁵ For non-technical treatments of the economic theory of collusion, the reader is referred to Motta (2004) and Viscusi, Harrington, and Sappington (2018). For technical treatments, see Vives (1999) and Harrington (2017) which focuses on unlawful collusion.

further along this line of argument, any firm could, on its own, set a supracompetitive price. It does not need to coordinate with rival firms, for a firm has full control over its price. Of course, a firm *does not want* to charge a supracompetitive price because, as long as other firms are setting competitive prices, a firm's profits would decline should its price veer from the competitive level. However, what could make it attractive to a firm to raise its price above the competitive level is if rival firms were also to do so. Thus, the challenge for a firm is to induce *other* firms to price at a supracompetitive level. If other firms charge supracompetitive prices then it could be optimal for this firm to do so, too.

How does a firm get other firms to charge supracompetitive prices? Here, we draw on the economic theory of collusion which establishes that the stability of firms charging supracompetitive prices is rooted in pricing strategies that embody a *reward-punishment scheme*. A simple example of such a pricing rule (or strategy) is the following. A firm initially prices at some supracompetitive level. In any future period, it prices at that supracompetitive level if all firms priced at or above that supracompetitive level in past periods; and prices at the lower competitive level if one or more firms priced below that supracompetitive level in some past period. If all firms adopt such a strategy then the resulting outcome is that all firms start out charging supracompetitive prices and continue to do so.

The economic theory of collusion identifies conditions under which the adoption of such a strategy is in a firm's best interests (that is, it maximizes the present value of its profit stream) given it expects other firms to also adopt it. In abiding by a collusive strategy, a firm must resist the temptation to undercut rival firms' collusive prices. Such undercutting is attractive because it means picking up more sales at a healthy profit margin and that results in higher profits today. What neutralizes that temptation is recognition of a causal relationship between a firm's *current* conduct and other firms' *future* conduct. If it prices at the high collusive level today then it anticipates the other firms pricing at that high level tomorrow, while not pricing high today is anticipated to bring forth low (competitive) prices from the other firms tomorrow. Thus, the firm associates high future profits from charging the collusive price today, and low future profits from pricing below the collusive level today. While the described retaliatory response of rival firms is

to price at competitive levels, one could imagine other forms of retaliation. For example, it could involve a short time of very low prices (even below cost, perhaps) with a subsequent return to collusive prices. What is critical is the perceived existence of this causal relationship between a firm's current price and other firms' future prices, and that it has the property that *high prices beget favorable future conduct from rival firms and low prices beget unfavorable future conduct from other firms.*

At the risk of belaboring the point, a reward-punishment scheme is implicit in any (effective) collusive arrangement. If a firm abides by the collusive outcome (e.g., high prices, exclusive territories, customer allocation) then it is *rewarded* in the future by rival firms continuing to abide by the collusive outcome (e.g., setting high prices, staying out of the firm's exclusive territory, not selling to its customers); while if it departs from the collusive outcome (e.g., setting a low price, selling outside of its territory, taking another firm's customers) then it is *punished* in the future by rival firms acting aggressively to reduce that deviating firm's profits (e.g., lowering price, entering the firm's territory, stealing its customers). Collusion ties future rewards and punishments to current behavior and that is what sustains a supracompetitive outcome. It is this causal relationship that links a firm's current conduct with rival firms' future conduct that defines collusion, not the setting of supracompetitive prices. Supracompetitive prices are the product of that causal relationship.

Definition *Collusion* is when firms use strategies that embody a reward-punishment scheme which rewards a firm for abiding by the supracompetitive outcome and punishes it for departing from it.

From hereon, a "collusive strategy" refers to a reward-punishment scheme that, when adopted by all firms, results in supracompetitive prices.

It is appropriate to think of collusion as a *self-enforcing contract*. As with any contract, collusive strategies specify what it means to comply with the contract which, in this case, is to set supracompetitive prices. It also specifies a penalty should a firm fail to comply with the terms

of the contract. That penalty could take the form of rival firms' charging lower prices or some other punishment. Where this contract differs from the usual variety is that there is no third party to enforce it. It must then be self-enforcing in the sense that it is in the narrowly defined self interest of each firm to abide by the contract at all times. Hence, the penalty must be *severe* so that it is in the best interests of a firm to charge a supracompetitive price, rather than deviate from the terms of the contract (by undercutting rival firms' prices) and incur the penalty. And the penalty must be *credible* so that a firm can anticipate it will be imposed should it deviate from the collusive outcome. It is credible when it is in each firm's interests to go through with the punishment in response to a firm having deviated from setting the collusive price. The penalty is then not implemented by a third party, such as the courts with the assistance of law enforcement, but rather is imposed by the firms themselves.

In sum, a collusive strategy contains a reward-punishment scheme designed to provide the incentives for firms to consistently price above the competitive level. Furthermore, it is in a firm's best interests to adopt a collusive strategy only when it believes other firms have similarly adopted it. Hence, there must be some mutual understanding among firms that they are using this collusive strategy. Collusion is the common adoption of a strategy embodying a reward-punishment scheme, and supracompetitive prices are the product of that adoption.

3 Defining Unlawful Collusion

Having defined collusion, let us now address *unlawful* collusion. Competition law as it pertains to collusion is often dated from the Sherman Act in the United States in 1890 (though Canada preempted the United States by instituting its competition law in 1889). Section 1 prohibits contracts, combinations, and conspiracies that unreasonably restrain trade.⁶ Subsequent judicial rulings have effectively replaced the reference to "contracts, combinations, and conspiracies" with the concept of "agreement." It is now understood that firms are in violation of Section 1

⁶ Section 1 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." *Standard Oil Co. v. United States*, 221 U.S. 1, 31, 86 S. Ct. 502 (1911) clarified that only "unreasonable" restraints of trade are prohibited.

when there is *an agreement among competitors to limit competition*. Though the term "agreement," which is now so integral to defining liability, does not appear in the Sherman Act, many jurisdictions that arrived later to the enforcement game have put the term into their competition law. For example, in the European Union, Article 101 (1) of the TFEU (1999) states: "The following shall be prohibited: all agreements between undertakings, decisions by associations of undertakings and concerted practices which ... have as their object or effect the prevention, restriction or distortion of competition."

The question of "What is unlawful collusion?" then becomes "What is an agreement to limit competition?" Key judicial decisions by the U.S. Supreme Court have defined an agreement as when firms have a "unity of purpose or a common design and understanding, or a meeting of minds"⁷ or "a conscious commitment to a common scheme designed to achieve an unlawful objective"⁸. This perspective has been echoed by the E.U. General Court which has defined an agreement as or as requiring "joint intention"⁹ or a "concurrence of wills"¹⁰. All of these terms focus on the same mental state: mutual understanding among firms that they will restrict competition in some manner.

In Section 2, collusion was defined as a "common scheme" based on rewards and punishments to yield supracompetitive prices, and that the common adoption of that scheme requires a "meeting of minds" among firms that all will adopt it. Thus, there are significant links between collusion and the legal system's definition of liability.¹¹ However, we also know that collusion is not a violation of Section 1. It is not enough that firms have mutual understanding to constrain competition for the courts care how that mutual understanding was achieved. As Judge Richard Posner noted, "section 1 of the Sherman Act ... does not require sellers to compete; it

⁷ *Am. Tobacco Co. v. United States*, 328 U.S. 781, 810 66 S. Ct. 1125 (1946)

⁸ *Monsanto Co. v. Spray-Rite Serv. Corp.*, 465 U.S. 752,765 104 S. Ct. 1464 (1984)

⁹ Judgment of the Court of 15 July 1970. *ACF Chemiefarma NV v Commission of the European Communities* Case 41-69.

¹⁰ Judgment of the Court of First Instance of 26 October 2000. *Bayer AG v Commission of the European Communities*

¹¹ This point has been previously made by Yao and DeSanti (1993), Werden (2004), and Kaplow (2013)

just forbids their agreeing or conspiring not to compete."¹² It is then not collusion but the process of "agreeing or conspiring" in order to collude that is determinative of firms' guilt.

In understanding what makes for an illegal process, courts have been guided by the requirement that "there must be evidence that tends to exclude the possibility that the [firms] were acting independently."¹³ While many avenues have been pursued by plaintiffs to argue that firms' conduct could not have been reached independently, successful recipes for convincing a court of that claim have almost always had a common ingredient: *evidence of an overt act of communication*.¹⁴ When firms communicate in a manner pertinent to future conduct (either expressing intentions or conveying information relevant to intentions), they create the legitimate concern that they have influenced each other's conduct and, therefore, their behavior was not reached independently but was instead the product of an agreement.

As to how "overt" must that communication be, an "explicit, verbally communicated assent to a common course of action"¹⁵ certainly does the trick. The issue is "how far may we move away from direct, detailed, and reciprocal exchanges of assurances on a common course of action and yet remain within the statutory and conceptual boundaries of an agreement."¹⁶ Notably, an "overt act of communication" could fall well short of a clearly articulated invitation to collude with a corresponding acceptance of that invitation. It is sufficient that there is some expression of intent resulting in reliance among firms that they will coordinate in order to reduce

¹² *In re Text Messaging Antitrust Litig.*, 630 F.3d 622, 627 (7th Cir. 2010)

¹³ See *supra* note 7, at 768.

¹⁴ "[F]ew courts have found a conspiracy without some evidence of communication tending to show an agreement." Hovenkamp (2016), p. 243. "In the search for evidence that tends to exclude independent action, courts have focused primarily on evidence tending to suggest communication has occurred. Although some cases do not involve testimony or documents detailing communication, the courts nevertheless require proof that they conclude justifies an inference that communications took place. In essence, there is no longer an open-ended plus factors analysis; the only evidence that actually distinguishes interdependent and concerted action is evidence that tends to show that the defendants have communicated in the requisite ways." Page (2007), p. 447.

¹⁵ Turner (1962), p. 683.

¹⁶ Phillip E. Areeda, *Antitrust Law* 9-12 (1986); cited in Kovacic (1993), pp. 18-9.

competition. They may engage in less direct forms of communication as long as it ends up (or is intended to end up) with mutual assurance of compliance on a coordinated plan.

Figure 1

Communication → Collusion → Supracompetitive Prices

In order to highlight where jurisprudence places illegality, Figure 1 depicts the various stages associated with firms restraining competition. Firms begin in a state of competition and, through communication, seek to achieve mutual understanding to stop competing or to compete less aggressively. The intended end result of communication is the adoption of a reward-punishment scheme (collusion) which delivers the desired outcome of supracompetitive prices. Supracompetitive prices are legal.¹⁷ Collusion is legal. It is the process by which firms achieve collusion that is illegal, rather than the state of collusion itself.¹⁸

Collusion is legal when the process by which it is achieved is legal. One such category is *conscious parallelism* which is the process "not in itself unlawful, by which firms in a concentrated market might in effect share monopoly power, setting their prices at a profit-maximizing, supracompetitive level by recognizing their shared economic interests."¹⁹ Along these lines, Judge Posner commented that "it is not a violation of antitrust law for a firm to raise its price, counting on its competitors to do likewise (but without any communication with them on the subject) and fearing the consequences if they do not."²⁰ As opined by Judge Stephen Breyer, the legality of these forms of collusion is undesirable but, in his view, unavoidable.

¹⁷ In some jurisdictions, excessive pricing laws make supracompetitive prices illegal. There are no such laws in the United States though some states have "price gouging" laws which prevent raising prices in some circumstances.

¹⁸ "By operationalizing the idea of an agreement, antitrust law clarified that the idea of an agreement describes a process that firms engage in, not merely the outcome that they reach. Not every parallel pricing outcome constitutes an agreement because not every such outcome was reached through the process to which the law objects: a negotiation that concludes when the firms convey mutual assurances that the understanding they reached will be carried out." Baker (1993), p. 179.

¹⁹ *Brooke Grp. v. Brown & Williamson Tobacco Corp.*, 509 U.S. 209,227, 113 S. Ct. 2578 (1993)

²⁰ See supra note 12, at 624.

Courts have noted that the Sherman Act prohibits *agreements*, and they have almost uniformly held, at least in the pricing area, that such individual pricing decisions (even when each firm rests its own decisions upon its belief that competitors do the same) do *not* constitute an unlawful agreement under section 1 of the Sherman Act ... [T]hat is not because such pricing 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

The challenge lies in being able to define *ex ante* "interdependent" pricing that is intended to produce supracompetitive prices, and to be able to distinguish *ex post* such pricing from interdependence which is a natural consequence of competition. To exemplify the latter challenge, consider the following situation. Suppose either firm I or II raises its price and subsequently the other firm raises its price to that same level. Suppose this pattern is regularly observed. Here are two hypotheses consistent with that conduct. A competitive hypothesis is: Firm I (or II) raised its price in response to a positive cost shock and, because firms I and II are subject to a common cost shock, firm II (or I) experienced the same cost shock and thus consummated a similar price increase. For example, if all retail gasoline stations face the same wholesale price and that wholesale price rises because of an increase in the price of crude oil then we expect all retail gasoline prices to rise by about the same amount at about the same time. In comparison, a collusive hypothesis is: There were no cost shocks (or any other changes in the environment) to justify a higher competitive price, and instead a firm raised its price as an invitation to collude, which the other firm accepted by similarly raising its price. The latter firm does so under the anticipation that if it did not then its rival would lower its price back to the original level.²²

²¹ *Clamp-All Corp. v. Cast Iron Soil Pipe Inst.*, 851 F.2d 478,484 (1st Cir. 1988)

²² "If a firm raises price in the expectation that its competitors will do likewise, and they do, the firm's behavior can be conceptualized as the offer of a unilateral contract that the offerees accept by raising their prices." *In re High Fructose Corn Syrup Antitrust Litig.*, 295 F.3d 651,652 (7th Cir. 2002)

As collusion is a conscious act and collusion is well defined (in that it is the use of a reward-punishment scheme to sustain supracompetitive prices), one might be able to provide an operational description of collusion so that firms would know when they are colluding. That definition would not be based on the prices selected but rather the pricing rules (i.e., collusive strategies) deployed. However, that is for naught if we cannot confidently infer from the observed conduct the underlying strategies in order to determine if they are collusive. The difficulty is that we cannot get inside the heads of managers to determine why, say, firm I raised price and why firm II matched that price increase. We are left trying to infer why they are pricing the way they are based on observable actions. Given the usual data and empirical methods available to economists, it is difficult to conclude with an appropriately high level of confidence that firms are using a collusive strategy, as opposed to acting in a competitive manner. The unpredictability of this analysis means that firms would have a high state of uncertainty as to whether their conduct would result in a conviction or not. That uncertainty would weaken deterrence and, let us recall, the primary goal of competition law is to deter collusive conduct, not convict it.

Let me summarize the argument of this and the preceding section. Collusion is a self-enforcing contract characterized by a reward-punishment scheme which produces supracompetitive prices. Collusion is not unlawful. Certain processes which result in collusion are unlawful. At a minimum, there must be evidence of communication to allow a court to conclude that firms have not acted independently. Though all forms of collusion are harmful, some forms are legal because of the inability to effectively distinguish collusive conduct from competitive conduct. This challenge can be traced to the latency of collusive strategies. Without being able to get inside managers' heads, we cannot confidently determine whether the intent of observed pricing is collusive or competitive. As a result, the court focuses its attention on observable communications that facilitate collusion rather than collusion itself.²³

²³ "Most who advance a prohibition on communications as facilitating practices, however, do not regard successful interdependent pricing behavior to be illegal. ... Note, however, that there is a certain irony involved: Aiding and abetting is

4 Collusion by Autonomous Artificial Agents

The law and the courts have viewed collusion through the perspective of conspiracy: Company representatives with pricing authority communicate in some manner in order to coordinate on charging higher prices. But suppose that a company's pricing authority resides instead in a software program in the form of an autonomous artificial agent (AA). If most firms in a market have their prices set by AAs, could supracompetitive prices result? And, if they could, how likely would it be to occur?

In order to shed some light on these questions, let's begin by describing, in general terms, how an AA operates. An AA has two components. First, at any instance, it has a *pricing algorithm* which selects a price depending on the state of the AA where, roughly speaking, the state summarizes the perceived environment of the AA. Second, it has a *learning algorithm* that modifies the pricing algorithm based on a pricing algorithm's performance. The human agent selects the performance metric for the AA and the particular class of AAs (which constrains the set of feasible pricing algorithms and, in broad terms, how an AA learns). Both of those choices are far removed from determining the pricing algorithm and prices that an AA settles upon. Given the complexity of most AAs, it is fair to say that, from the perspective of any company employee, it cannot *ex ante* predict how an AA will price nor possibly even *ex post* explain the prices that were selected.

A pricing algorithm encompasses a pricing rule which assigns a price to each state. For example, a state could include a firm's cost, inventory, day of the week, and past prices. A state will typically be a partial description of the environment, such as the average price set by each firm rather than the entire collection of past prices. An AA is endowed with a certain number of states. The more states it has, the more finely it can describe its environment. It is partly in this sense that the state is the environment *as perceived* by the AA. An implication of having more states is that it becomes feasible for an AA to make its conduct - price, in our case - more finely

heavily punished, but undertaking the act one is trying to facilitate is freely permitted if such aids are unnecessary or cannot be proved to have been employed." Kaplow (2013), p. 57.

tuned to the environment. In addition to a pricing rule (how price is assigned to a state) and the set of states, an AA has a rule for how it modifies the state depending on the previous period's state and the prices chosen in the previous period. That is, given how it perceived the environment last period and the prices chosen by firms this period, the state-adjustment rule describes how the AA now perceives its environment.

Can AAs collude in the sense that we described in Section 2? Absolutely. With a sufficiently rich perception of its environment (i.e., enough states and an appropriate state-adjustment rule), a collusive strategy can be programmed as a pricing algorithm. A pricing algorithm can encompass a reward-punishment scheme and, should all firms' AAs have that reward-punishment scheme, supracompetitive prices will emerge. An example is provided later in this section. The problematic issue is not whether AAs can collude but rather whether AAs can *learn to collude*. That is, can AAs in a market mutually evolve to using pricing algorithms that produce supracompetitive prices?

To address that question, we must now discuss the learning algorithm component of an AA. As already noted, at any moment, an AA has a pricing algorithm that determines how it prices depending on the state. As it prices over time and realizes profit (or some other measure of performance, such as revenue), an AA's learning algorithm will adjust its pricing algorithm with the objective of finding a "new and improved" pricing algorithm. The most relevant class of learning algorithms is *reinforcement learning*.²⁴ Reinforcement learning has the general feature that an AA will continue to use a pricing algorithm when it has performed well (compared to some benchmark) and experiment with other pricing algorithms when it has not performed well. For example, this could mean adopting a pricing algorithm for some length of time, measuring its performance (e.g., average profit over that window of time), and then deciding whether to retain the pricing algorithm, slightly modify the pricing algorithm, or adopt an entirely new pricing algorithm.²⁵ Or learning might occur over specific elements of a pricing algorithm. As a pricing algorithm assigns a price to each state, the learning could occur at the level of each state-

²⁴ For an introduction to reinforcement learning, see Sutton and Barto (2000) and Gershman (2015).

²⁵ For example, that is a property of reinforcement learning in Hanaki, Sethi, Erev, and Peterhansl (2005).

price pair. If performance was good from choosing the assigned price for a particular state then that price will continue to be assigned to that state. If performance was bad then the learning algorithm will change the price assigned to that state, as it seeks to find a price that is a better choice for that state.²⁶

In seeking to find better pricing algorithms, there is a tension between exploration and exploitation (or "learning and earning"). At any moment, an AA could persist with its existing pricing algorithm (fully exploiting what it has learned) or experiment with another pricing algorithm (exploring for something better) which could either raise or lower performance. The value to experimentation is that, should the change lower performance, an AA can just return to the better pricing algorithm it was using; hence, any loss is temporary. If it should raise performance then that improvement in the pricing algorithm can be retained; hence, any benefit is permanent. When learning has progressed enough, the currently best pricing algorithm identified by an AA is more than likely to outperform some randomly selected pricing algorithm. At that time, experimentation can be expected to reduce performance, at least in the short run. It can raise it in the long run by finding a better pricing algorithm. As learning results in an AA's pricing algorithm getting better and better, the cost of experimentation rises – as the expected reduction in performance is greater - and the benefit of experimentation falls - as it becomes less likely to find a better pricing algorithm. Over time, reinforcement learning will experiment less though it is prudent to always engage in some experimentation since one can never be sure that the best pricing algorithm has been discovered or if the environment has significantly changed.

A complicating feature to learning in a market setting is that multiple agents are simultaneously learning. The best pricing algorithm for a firm's AA depends on the pricing algorithms used by rival firms' AAs. If rival firms' pricing algorithms were fixed then there are learning algorithms assured of eventually finding the best pricing algorithm.²⁷ The problem is that rival firms are also learning so their pricing algorithms are, in fact, changing. With multiple

²⁶ For example, that is a property of reinforcement learning in Calvano, Calzolari, Denicolo, and Pastorello (2018) which is reviewed below.

²⁷ One example is Q-learning which is reviewed later in this section.

AAs simultaneously learning, learning may not settle down (in the sense of always choosing the same pricing algorithm) and, when it does settle down, where it settles down is not uniquely determined. The AAs could eventually adopt pricing algorithms that have them charge competitive prices. Or the AAs could come to use pricing algorithms that embody a reward-punishment scheme which support supracompetitive prices.

In order to make these abstract notions a bit more concrete, let us consider a simple market setting within which AAs will operate. Suppose there are two firms - I and II - and each firm has two possible prices it can charge: low and high. Figure 2 is a profit matrix where the first number in a cell firm I's profit and the second number is firm II's profit. For example, if firm I chooses the low price and firm II chooses the high price then firm I earns profit of 6 and firm II earns profit of 1.

For this market setting, the competitive outcome is when both firms choose the low price. Note that each firm is doing the best it can given what the other firm is doing. Given firm II chooses the low price, firm I's profit is 2 from the low price and 1 from the high price; hence, firm I earns higher profit by also choosing the low price. The same is true for firm II. Next note that both firms would be better off if they jointly raised their prices; each would have its profit rise from 2 to 4. The high price is the supracompetitive price and would be the collusive outcome should firms succeed in colluding.²⁸

Figure 2

		Firm II	
		Low price	High price
Firm I	Low price	2, 2	6, 1
	High price	1, 6	4, 4

Assume that firms I and II repeatedly interact in setting prices and earning profits. Each firm adopts an AA and has the AA measure performance as the weighted sum of current and future profits with a smaller weight given to profits that are in the more distant future. For simplicity,

²⁸ The game in Figure 2 is the Prisoners' Dilemma.

market demand and cost do not change over time so, in each period, firms face a pricing decision with the associated profits shown in Figure 2. What does change from period to period is the history of prices. An AA's state is some summary statistic of past prices and, more specifically, a state here is specified to be the prices charged in the previous period. There are then four states: (low,low), (low,high), (high,low), and (high,high), where the first entry is firm I's price.

Endowed with a particular pricing algorithm, an AA operates as follows. Given the current state (i.e., the prices charged in the previous period), an AA chooses the price assigned by its pricing rule. Turning to the next period, the AA updates its state (which is now the prices that were just charged) and, based on that new state, chooses the price assigned by its pricing rule. And so forth. There are many pricing algorithms and let us describe a few. One pricing algorithm is to price low for all four states. If both AAs adopt that pricing algorithm then the competitive outcome is realized in each period. Another pricing algorithm is a reward-punishment scheme with a one-period punishment. It has an AA price high if the state is (high,high) or (low,low), and price low if the state is (high,low) or (low,high). Note that if both firms priced high in the previous period - so the state is (high,high) - then, according to their pricing algorithms, both will price high this period. If instead one of them priced high and the other priced low - which would correspond to one firm having deviated from the collusive outcome of both pricing high - then both go to pricing low; hence, there is a punishment in response to the deviation. If both go through with that punishment of low prices then, according to their pricing algorithms, they go back to pricing high (as the algorithm assigns a high price to the (low,low) state). If the AAs have these pricing algorithms and both start out pricing high then they will continue to price high. This supracompetitive outcome is sustained by the threat of a one-period punishment should an AA veer from pricing high.²⁹

²⁹ Another pricing algorithm is "tit-for-tat" which has an AA match the price set by the other AA in the previous period. Thus, firm I's AA prices low if the state is (low,low) or (high,low), and prices high if the state is (low,high) or (high,high). Yet another pricing algorithm is the grim punishment strategy which has the punishment be a permanent (not one period) reversion to competition. It prices high if the state is (high,high) and prices low for the other three states. There are many other pricing algorithms.

If, for example, both AAs use the one-period punishment pricing algorithm then AAs will be colluding as reflected in supracompetitive prices. The question is: Will AAs come to adopt these pricing algorithms and, if they do, will they persist with them? If the one-period punishment pricing algorithm is best for a firm given the other firm is using it (which will be the case if an AA gives enough weight to future profits) then adoption of this pricing algorithm will persist over time. More challenging is whether AAs will learn to use these pricing algorithms (or some other pricing algorithms that produce supracompetitive prices).

That question is explored in Calvano, Calzolari, Denicolo, and Pastorello (2018) for the simple model we've been discussing. The particular form of reinforcement learning they assume is Q-learning. Q-learning has an AA assign a perceived performance value to each state-action combination. As there are four states (corresponding to the four possible price pairs that could have occurred in the previous period) and two actions (corresponding to the two feasible prices), there are eight Q-values. Given its current collection of Q-values, an AA chooses the price that yields the highest Q-value given the current state, though with some small probability goes into experimentation mode and chooses the other price. When it does not experiment, an AA is then choosing the best price according to how it values different state-action combinations at that moment in time. Learning occurs by adjusting those Q-values in response to realized performance in the following manner. Having chosen a price given the current state, contemporaneous profit is received. That realized profit (which also depends on the price selected by the rival firm's AA) is used to update the Q-value attached to that state-action pair. The values for other state-action combinations are left unchanged. Come next period, the firm faces a new state, a price is again chosen to yield the highest value, profit is received, the value attached to that action-state pair is adjusted, and the process continues.

In summary, an AA chooses a price at any moment which is best for the current state given its assessment of how different prices perform for different states as described by its Q-values. The AA then adjusts its Q-values as it observes how a price actually performed for a particular state. As Q-values evolve, so does the pricing algorithm of the AA. With two AAs simultaneously learning, what pricing algorithms do AAs end up selecting? What prices are

charged? To address those questions, Calvano et al (2018) run many computer simulations. They start the AAs with Q-values that imply they price competitively; that is, for all states, the low price is the price with the highest Q-value. Just as would be the case with collusion by human agents, the AAs start in competitive mode. Each simulation has AAs interact and learn over 1 million periods, and 1,000 simulations are run.

Summarizing their findings, competition arises 18.3% of the time which means an AA chooses the low price for all states. In 43% of the periods, the AAs are colluding by having adopted the one-period punishment strategy. More than 50% of the time, the pricing algorithms of AAs are some form of collusive strategy which implies they are charging the high price. Starting from competition, it takes, on average, 165,000 periods for AAs to learn how to collude. While that may seem long, it all depends on the length of a period. If, in actual markets, AAs are changing prices and earning profits every minute (hence, a period is one minute) then 165,000 periods translates into 115 days, which is not necessarily long. The takeaway from Calvano et al (2018) is that AAs can adapt their way to collusive prices.

Can AAs collude? Yes. The pricing algorithms of an AA are rich enough to encompass the collusive strategies that have been used by human agents. Can AAs learn to collude in a simple setting? Yes. With two AAs, two prices, and a fixed environment, simulations show that collusion is more common than competition. Can AAs learn to collude in an actual market setting? We do not know, and I am skeptical of anyone who thinks they know. As we cannot dismiss the possibility that AAs are able to learn to collude in actual markets, it is prudent to find an appropriate legal response should they be able to do so.

5 Applying Jurisprudence to Collusion by Autonomous Artificial Agents

Jurisprudence regarding Section 1 of the Sherman Act does not prohibit collusion; it prohibits certain processes that might result in collusion. Effectively, what is illegal is communication among firms intended to achieve an agreement where an agreement is mutual understanding

between firms to limit competition. Though the courts are clear in defining liability as an agreement, they are equally clear that there must be some overt act of communication to create or sustain that mutual understanding.

According to that jurisprudence, I claim that firms that collude through the use of AAs are not guilty of a Sherman Act Section 1 violation. In making this claim, the presumption is that these AAs only have access to information that would be present under competition, such as past prices, sales, and other market data. In particular, the AAs do not post any extraneous information which could possibly be construed as one AA conveying a message to another AA. If there is any communication (broadly defined) between AAs, it is occurring through prices or other legitimate market data. As there is no overt act of communication, a requisite piece of evidence is absent.

Of course, evidentiary requirements evolve and perhaps they could be adapted to handle collusion by AAs. However, there is a more daunting challenge in prosecuting firms that collude using AAs: Firms do not satisfy the court's definition of liability. The firms' managers independently adopted these AAs and lacked awareness that their adoption would produce collusion. Even if each manager subsequently learned that the other firms' prices are also set by AAs, mutual understanding to limit competition is still lacking because of the presumption that the managers did not foresee that collusion would ensue upon each having adopted an AA to set prices.

That managers are not culpable does not immediately get the firms off the hook. Given a company is liable for its employees,³⁰ could it then be the case that a company is liable for its software programs? Under current jurisprudence, this would seem to require that someone working for the firm or something owned by the firm meets the definition of liability. Though the presumption is that all human agents working for the firm lack a "meeting of minds" or a "conscious commitment to a common scheme", could the AAs possess it? Addressing that

³⁰ "[A] corporation may be held criminally responsible for antitrust violations committed by its employees . . . even if, such acts were against corporate policy or express instructions." *United States v. Basic Constr. Co.*, 711 F.2d 570,573 (4th Cir. 1983)

question draws us into deep philosophical territory regarding whether a computer program can "understand".

The philosopher John Searle (1980) famously argued that computers cannot understand, in what has become known as the Chinese Room Argument.

Imagine a native English speaker who knows no Chinese locked in a room full of boxes of Chinese symbols (a data base) together with a book of instructions for manipulating the symbols (the program). Imagine that people outside the room send in other Chinese symbols which, unknown to the person in the room, are questions in Chinese (the input). And imagine that by following the instructions in the program the man in the room is able to pass out Chinese symbols which are correct answers to the questions (the output). The program enables the person in the room to pass the Turing Test for understanding Chinese but he does not understand a word of Chinese.³¹

The point is that "whatever purely formal principles you put into the computer, they will not be sufficient for understanding, since a human will be able to follow the formal principles without understanding anything."³² From this perspective, price-setting AAs can be transmitting data to each other and acting on that data so as to yield coordinated pricing, but that does not imply the AAs understand they are coordinating to restrain competition. And, without understanding, there cannot be mutual understanding. Given no agents - human or artificial - in those firms have a "meeting of minds," the firms do not have an agreement and thus have not violated Section 1 of the Sherman Act.³³

³¹ Searle (1999)

³² Searle (1980), p. 418.

³³ To my knowledge, this argument for why AAs are not liable was first made in Harrington (2012).

The Chinese Room Argument is not without its detractors.³⁴ However, even if one were to grant understanding to an AA, it is another step to reach a state of *mutual* understanding. That requires firm I's AA to understand that firm II's AA is setting a high price on the understanding that firm I's AA will do so. That then requires an AA to have a theory of mind whereby an agent is self-aware of its own mental processes (thinking is a state as well as a process) and assigns similar mental processes to another agent. Perhaps it is not a difficult leap from understanding to mutual understanding, for it has been recognized since von Neumann (1945) that a computer program is both a set of instructions and a file that can be read by itself or other programs. However, even with a credible argument that computer programs could have the requisite mutual understanding for there to be an agreement, it is doubtful that the argument would be sufficiently compelling to convince the courts that AAs, like human agents, can have an agreement to restrict competition.

This view is consistent with a recent statement of the Antitrust Division of the U.S. Department of Justice:

Absent concerted action, independent adoption of the same or similar pricing algorithms is unlikely to lead to antitrust liability even if it makes interdependent pricing more likely. For example, if multiple competing firms unknowingly purchase the same software to set prices, and that software uses identical algorithms, this may

³⁴ Some of the criticisms are presented and addressed in Searle (1980) and, for a more recent treatment, see Preston and Bishop (2002). But Searle (2002, p.52) responds: "The Chinese Room Argument, in short, rests on two absolutely fundamental logical truths, and twenty-one years of debate has not in any way shaken either of these. Here they are. First, syntax is not semantics. That is to say, the implemented syntactical or formal program of a computer is not constitutive of nor otherwise sufficient to guarantee the presence of semantic content; and, secondly, simulation is not duplication. You can simulate the cognitive processes of the human mind as you can simulate rain storms, fire alarm fires, digestion or anything else that you can describe precisely. But it is just as ridiculous to think that a system that had a simulation of consciousness and other mental processes thereby had the mental process, as it would be to think that the simulation of digestion on a computer could thereby actually digest beer and pizza."

effectively align the pricing strategies of all the market participants, even though they have reached no agreement.³⁵

Taking a contrary view:

It is no defense to suggest that algorithms, programmed for autonomy, have learned and executed anticompetitive behavior unbeknownst to the corporation. The software is always a product of its programmers - who of course have the ability to (affirmatively) program compliance with the Sherman Act ...³⁶

But what does it mean to "program compliance with the Sherman Act"? Jurisprudence tells us it means that AAs cannot communicate with each other in the same sense that human managers are prohibited from communicating. Thus, AAs would not be in compliance if they coordinated their conduct using arbitrary messages unrelated to the competitive process, but would be if coordination was achieved through their prices, as that is an example of legal conscious parallelism. We are then back to a situation in which collusion achieved through means that, if executed by human agents, would be lawful.

In conclusion, it seems unlikely that collusion by AAs could result in the firms deploying those AAs being found in violation of Section 1 of the Sherman Act. That is sufficient for us to move forward and consider an alternative legal approach to prosecuting collusion by AAs.

6 A Legal Approach to Collusion by Autonomous Agents

Returning to Figure 1, communication between firms results in the mutual adoption of collusive strategies which yield supracompetitive prices. Collusion is a mode of conduct among firms in a market that has them use pricing rules embodying a reward-punishment scheme whereby high

³⁵ U.S. Department of Justice, "Algorithms and Collusion - Note by the United States," 2017. p. 6.

³⁶ Gosselin, Jones, and Martin (2017)

prices are rewarded by rival firms with the continuation of high prices and low prices are punished by rival firms with, for example, a move to low prices. Even though it is collusion which is harmful, jurisprudence has made the communication that facilitates it illegal. As Judge Breyer noted, the reason for this approach is the lack of a "judicially enforceable remedy." The absence of a remedy is rooted in the fact that collusion (defined as a mode of firm conduct) is not directly observable. The reward-punishment scheme which defines collusion is latent, inside the managers' heads. Hence, if collusion were to be prohibited then courts would be left trying to infer from the prices that firms set, whether the underlying strategies responsible for those prices encompasses a reward-punishment scheme designed to produce supracompetitive prices. However, it is very difficult to confidently determine whether prices are the product of illegitimate interdependence of firms' conduct (i.e., a reward-punishment scheme) or are instead the product of legitimate interdependence. Given these difficulties, the courts have decided that collusion is not illegal but communication that facilitates collusion is illegal.

The judicial doctrine that has just been described is based on collusion *as it is conducted by human agents*. It is predicated on the difficulty of knowing the strategy used by a human agent and, in particular, whether observed prices are supported by a reward-punishment scheme among firms. However, the situation is fundamentally different when prices are set by AAs.

When prices are controlled by an autonomous artificial agent, the firm's strategy is, in principle, observable.

The rule determining price is written down in the algorithm's code which means that it can be accessed (in some manner) and, in that sense, it is possible to get "inside the head" of the price-setting agent. We are not left with trying to indirectly infer the latent strategy from observed behavior amidst a changing environment, but rather can directly observe the strategy itself. And if one can observe the strategy then one can determine whether it embodies a reward-punishment scheme, which is the defining feature of collusion, what results in supracompetitive prices, and what should be prohibited.

The implementation challenge is to take that simple observation – the strategy of an AA can be directly observed - and define liability and construct evidentiary methods so that collusion by AAs can be effectively prosecuted. To begin, I propose an approach to liability based on a *per se* prohibition of certain pricing algorithms.

Liability: There is a *per se* prohibition on certain pricing algorithms (or, equivalently, on pricing algorithms having certain properties) that support supracompetitive prices.

To illustrate this perspective to liability, let us return to the example in Section 4. Though it is not a realistic model of an actual market, it is useful for conveying the principle. In that example, two firms could choose one of two prices, where the low price is competitive and the high price is supracompetitive. One stable pair of pricing algorithms is for each firm to set the low (competitive) price irrespective of the history of past prices. There are other stable pairs of pricing algorithms which instead have firms setting the high (collusive) price. All of those pricing algorithms have a firm's current price be contingent on the previous period's price by the rival firm and, generally, have a firm set a low price when the rival firm previously set a low price. In other words, they punish a rival firm for setting a low price and it is through that threat of punishment that high prices are sustained. An AA does not lower price, even though it would raise current profit, because it implicitly anticipates a retaliatory response by the other AA that would lower future profits. In this simple setting, a pricing algorithm would be prohibited if it conditioned price on a rival firm's past prices. AAs would be allowed to set any price, low or high, but just not use pricing algorithms that could reward or punish a rival firm based on that firm's past prices. Due to the simplicity of the setting, this example ignores many challenging issues associated with the implementation of this legal approach. Nevertheless, it illustrates the approach of a *per se* prohibition on certain pricing algorithms.³⁷

³⁷ In particular, I am not suggesting that, in practice, a pricing algorithm should be prohibited if it conditions on rival firms' past prices. However, within the confines of this simple setting, that would be the proper definition of the set of prohibited pricing algorithms.

Given a set of prohibited pricing algorithms, here is an approach for determining when firms are liable.

Evidentiary Methods: Liability would be determined by: 1) an examination of the pricing algorithm's code to determine whether it is a prohibited pricing algorithm; or 2) entering data into the pricing algorithm and monitoring the output in terms of prices to determine whether the algorithm exhibits a prohibited property.

Though the legal approach is simple to state, its implementation poses some difficult but not insurmountable challenges. First, it must be decided which pricing algorithms should be placed in the prohibited category. Second, the most relevant learning algorithms are not amenable to extracting the pricing algorithm from a reading of the code. In that case, tests will have to be conducted to determine whether a price algorithm embodies a prohibited property. I now turn to discussing these challenges.

6.1 Defining Liability

Our focus has been on how giving pricing authority to AAs can promote collusion. However, to my knowledge, that is not why firms have been shifting pricing authority from humans to software programs. Rather, the objective is to realize efficiency gains from the use of automated pricing algorithms. Hence, when defining the set of prohibited pricing algorithms, it should be guided by the objective of convicting cartels and deterring collusive conduct but doing so without interfering with legitimate competitive rationales. The more that collusion-promoting algorithms are not included in the prohibited set, the more harm is created because there is collusion that, instead of being prosecuted and shut down, continues unabated. The more that efficiency-enhancing algorithms are included in the prohibited set, the more harm is created by the associated foregone surplus. The set of prohibited pricing algorithms should be as inclusive

as possible of those algorithms that promote collusion, and as exclusive as possible of those algorithms that promote efficiency.³⁸

Let pa denote "pricing algorithm" and PPA denote the collection of prohibited pricing algorithms. Given a specification of PPA , $Pr(pa \text{ is in } PPA \mid pa \text{ is collusive})$ is the probability a pricing algorithm is determined to be in the prohibited set when the pricing algorithm is collusive. $Pr(pa \text{ is in } PPA \mid pa \text{ is competitive})$ is the probability a pricing algorithm is determined to be in the prohibited set when the pricing algorithm is actually competitive. Ideally, $Pr(pa \text{ is in } PPA \mid pa \text{ is collusive}) = 1$ and $Pr(pa \text{ is in } PPA \mid pa \text{ is competitive}) = 0$ so that a pricing algorithm is concluded to be unlawful if and only if it is collusive. That such an ideal is not reached will be due to misspecification of PPA - some collusive pricing algorithms are excluded from PPA or some competitive pricing algorithms are included – or incomplete data or inadequate methods for evaluating whether a particular pricing algorithm is in PPA .

Recognizing these imperfections, a useful measure for assessing the efficacy of a particular definition of liability is the likelihood ratio:³⁹

$$LR(PPA) = \frac{Pr(pa \in PPA \mid pa \text{ is collusive})}{Pr(pa \in PPA \mid pa \text{ is competitive})}$$

$LR(PPA)$ is the probability that the pricing algorithm is declared to be prohibited given it is a collusive pricing algorithm divided by the probability that the pricing algorithm is declared to be prohibited given it is a competitive pricing algorithm.

Error costs from false negatives are reduced when $Pr(pa \text{ is in } PPA \mid pa \text{ is collusive})$ is higher. When $Pr(pa \text{ is in } PPA \mid pa \text{ is collusive}) = 1$ then collusion is always prosecuted as, whenever collusion occurs, the associated pricing algorithm is determined to be in the prohibited class. Error costs from false positives are reduced when $Pr(pa \text{ is in } PPA \mid pa \text{ is competitive})$ is lower.

³⁸Implicit in the preceding discussion is that a competitive rationale for a pricing algorithm is necessarily efficiency enhancing. That need not be the case. For example, a pricing algorithm may serve to promote price discrimination, which can either raise or lower social welfare. That issue will not be addressed here, and instead it is presumed that it is desirable to exclude from the set of prohibited practices those pricing algorithms for which there is a competitive rationale.

³⁹For a discussion of the usefulness of the likelihood ratio in the context of evidentiary rules, see Kaplow (2014).

When $Pr(pa \text{ is in } PPA \mid pa \text{ is competitive}) = 0$ then firms are never prosecuted when they are competing. As the set PPA is expanded, so that more pricing algorithms are prohibited, the rate of false negatives falls but the rate of false positives rises; and as PPA is contracted, the rate of false positives goes down but the rate of false negatives goes up. As shown in Kaplow (2014), most evidentiary rules can be represented by the requirement that the likelihood ratio exceeds some critical value. While a socially optimal balancing of fewer false negatives (by expanding PPA) and fewer false positives (by contracting PPA) is not achieved by maximizing the likelihood ratio, it is sensible to choose PPA in order to have a reasonably high likelihood ratio.⁴⁰ Of course, this approach is vacuous if it turns out there does not exist a set PPA such that the likelihood ratio is reasonably high. For example, if all properties of a pricing rule that are useful for collusive purposes are also instrumental in enhancing efficiency then $LR(PPA)$ will be low for all PPA . In that case, either the liability definition will have no bite (when the threshold for the likelihood ratio is set high) or have a substantive chilling effect on competition (when the threshold for the likelihood ratio is set low).

The potential efficiency benefits of a learning algorithm and the automated pricing algorithm that it selects come from two sources.⁴¹ First, the learning algorithm can make the firm more informed about how price affects profit which then allows the firm to better identify profit-maximizing prices; in other words, it uses *past* data to predict the relationship between price and profit. Some learning algorithms make this explicit as they are composed of two modules; an estimation module that makes predictions about demand conditional on current market conditions and a firm's price, and an optimization module that selects prices based on that

⁴⁰ As explained in Kaplow (2014), one does not generally want to maximize the likelihood ratio for that could lead to high error costs associated with a high rate of false negatives. Nevertheless, at this preliminary stage of the analysis, it is useful to search for sets of prohibited pricing algorithms that yield reasonably high likelihood ratios. This approach is consistent with the traditionally high threshold for proving unlawful collusion.

⁴¹ For a review of algorithmic pricing and its benefits, see Oxera Consulting (2017) and Deng (2018).

estimated demand. The learning algorithm results in a better demand estimate and that allows the firm to set a price that yields higher profits.⁴²

Second, a pricing algorithm can allow price to be tailored to current market conditions through the rapid adjustment of price to changes in market conditions and the personalization of prices to a consumer's traits; in other words, it uses *current* data to match price to market conditions. This source of efficiency is made possible by Big Data.⁴³ Often referred to as "dynamic pricing," it is the automation of prices so that they condition on high-frequency data. A firm is learning its sales, rival firms' prices, and other variables on a very fine time scale. With automated pricing, prices can adjust as soon as new information is received, which means price can quickly adapt to changes in sales, inventories, rival firms' prices, and any other variable that is monitored at a high-frequency level. Even if a human manager were to continuously monitor such information, it could not process the information rapidly enough to make sensible changes in prices; the automation of prices is required. Another benefit from automated pricing and rich data is that it can personalize prices to a customer's traits or to classes of customers; that is, engage in more effective price discrimination. This information could be the time of day that a consumer is on a website, the consumer's clickstream activity or past purchases, and demographic information (which the firm may have if the consumer is registered with the website).⁴⁴

42 For example, Nambiar, Simchi-Levi, and Wang (2016) have an estimation module that uses past prices and sales in a regression model to produce an estimate of a firm's demand function. With that estimate of the firm's demand function, an optimization module calculates the price that maximizes expected revenue. In order to promote learning, the chosen price equals the revenue-maximizing price plus some random term that serves to create price experiments. For a survey of some of these learning algorithms, see den Boer (2015).

43 Big Data refers to the use of large scale computing power and sophisticated algorithms on massive data sets for the purpose of finding patterns, trends, and associations in human behavior and other phenomena. Data is "big" in volume (number of observations), variety (heterogeneity in variables), and velocity (time frequency of data).

44 For example, there are documented instances in which online retailers such as Home Depot, Orbitz, and Staples were setting prices to a website user based on location, browser history, and whether the user was mobile-based (Diakopoulous, 2014).

In sum, the potential efficiencies of automating learning and pricing (as would occur with an AA) are: i) estimating a firm's environment (in particular, demand); ii) finding best prices given current market conditions; iii) rapid adjustment of price to changes in market conditions; and iv) personalizing prices for customers. From the perspective of constructing a set of prohibited collusive pricing algorithms, here is the critical point: *The properties of pricing algorithms that deliver these efficiencies are not directly relevant to generating collusion.* Collusion is about influencing rival firms' prices through a reward-punishment scheme. Task (i) may involve taking account of rival firms' past prices - so as to more accurately estimate a firm's demand function - but it does not involve influencing rival firms' future prices. A more sophisticated version of (i) could entail forecasting rival firms' prices but again it does not seek to influence how rival firms price. Some learning algorithms embody price experiments in order to more effectively accomplish task (i). Contrary to a goal of collusion, such price perturbations are likely to make coordination on supracompetitive prices more difficult; hence, an AA that is more effective in learning about demand may be less effective at generating collusion. With regards to (iv), personalized pricing makes monitoring of a firm's prices by rival firms more difficult because the price is tailored to the individual customer and would not necessarily be observed if rival firms are scraping a firm's web page. Thus, personalized pricing may improve a firm's profit under competition but would make collusion more difficult.

While a comprehensive and rigorous examination of these issues is required before any conclusions can be drawn, a first cut suggests that the properties of pricing algorithms that serve legitimate competitive purposes would not be useful for promoting collusion, while the properties that promote collusion seem quite distinct from those that enhance efficiency. It may then be possible to identify a set of prohibited pricing algorithms which would target collusion while not interfering with competition.

6.2 Constructing Evidentiary Methods

Given a set of prohibited pricing algorithms, the next task is developing a process for determining whether or not a firm's pricing algorithm is a prohibited one. There are two general approaches to testing an AA in order to learn its properties.⁴⁵ *Static testing* involves examination of the program's code without running the program. *Dynamic testing* is when the program is run for selected inputs and the output is observed. In evaluating the relevance of these methods, our discussion will focus on AAs that use reinforcement learning.

Static testing may be feasible for some methods of reinforcement learning such as Q-learning, which was discussed in Section 4. Recall that Q-learning assigns a value to each price-state pair. A state describes the environment of an AA such as its cost, inventory, and recent prices charged in the market. Given the current state, the price with the highest Q-value is chosen. After choosing that price, profit is received and its performance (as measured by profit or revenue) is used to update the value attached to that price-state pair. An AA's pricing algorithm at any moment in time is the assignment of the best price to each state according to the current collection of Q-values. Hence, the pricing algorithm can be constructed by inspecting the collection of Q-values.⁴⁶ With that pricing algorithm, one can assess its properties and whether it is in the prohibited set of pricing algorithms though this task could prove difficult when there are many prices and states.

There are other methods of reinforcement learning for which inspection of the code will not prove informative because the mapping between the code and the basis for a choice is too complex to sort out.⁴⁷ Such methods are broadly referred to as "deep learning".⁴⁸

⁴⁵ A useful reference for the ensuing discussion is Desai and Kroll (2017).

⁴⁶ The values are stored either in the form of a table or, when function approximation is used, as a vector of estimated coefficients for a function that maps prices and states to values.

⁴⁷ Such concerns can also be found in Ezrachi and Stucke (2017) and OECD (2017).

⁴⁸ LeCun, Bengio, and Hinton (2015) provides a description of deep learning. See Mnih et al (2015) for the use of deep reinforcement learning to train artificial neural networks to play video games.

Deep learning [is] a class of computerized neural networks-based algorithms [and one] of the things that sets them apart from other algorithms is their limited ability to explain their decision making. In deep learning, features are created as a (possibly complex) computation over multiple features, making such algorithms' decision-making hard to explain.⁴⁹

In practice, static testing is unlikely to be an effective method for assessing whether an AA is using a prohibited pricing algorithm, which leads us to dynamic testing. This approach entails feeding various environmental conditions (e.g., what prices were recently charged and a firm's cost) into an AA and documenting how the generated prices respond to those conditions. The objective is to determine whether the latent pricing algorithm has prohibited properties. Desai and Kroll (2017) note two challenges with this approach. First, the possible sets of inputs (in our context, past prices, costs, demand conditions, etc.) could be very large so one is only able to test a small subset of inputs. To what extent that is an obstacle depends on the property of the pricing algorithm one is looking for. If the property is how an AA responds to other firms' prices, it may be sufficient to consider inputs that differ only in other firms' prices for a representative sample of cost and demand conditions. Second, an AA is changing its pricing algorithm as it learns. It is possible that the pricing algorithm in use today is very different from those that were used during the past year, and it is price data from the past year that might have sparked suspicions of collusion. The extent of that challenge depends on how much the pricing algorithm is changing over time. Though an AA may always be tweaking its pricing algorithm, it could be relatively stable with respect to the properties of interest.

A benefit from having a well-defined procedure for testing whether a pricing algorithm is lawful is that it will clarify to both managers and courts what exactly is illegal. If managers do not know when they are acting unlawfully then illegal behavior cannot be deterred. Managers would be able to determine when they are in compliance with the law by having the learning

⁴⁹ A. Gal (2017), p. 5.

algorithm programmed to engage in periodic testing of the pricing algorithm to ensure it does not exhibit the prohibited property. When feasible, the learning algorithm could also be constrained not to use illegal pricing algorithms. Effective enforcement also requires that courts can reasonably determine when the law is violated. If the court is not effective at making such a determination then it will be prone to false negatives - thereby allowing illegal collusion to continue - and false positives - thereby interfering with competitive markets. Furthermore, even if managers know when they are violating the law, if they anticipate that the court is unable to accurately determine illegality then deterrence will again be weakened because conviction is less tied to whether or not firms are actually acting unlawfully. With a well-defined test for determining whether a pricing algorithm exhibits a prohibited property, courts could reasonably and predictably determine when the law is violated.

6.3 A Research Program for Defining Liability

The challenge of prohibiting collusion by AAs has parallels to policy issues related to fairness and machine learning.⁵⁰ For example, consider an automobile insurance company using machine learning on data sets to make predictions about a driver's accident rate, or a bank using machine learning to assess the credit risk of an applicant for a mortgage loan. In seeking to generate the best predictions, machine learning could use traits such as a person's race or gender. As such emergent discrimination is generally seen as undesirable, it has been proposed to constrain machine learning so that it satisfies some notion of fairness. Implementation of such an objective first requires defining fairness - which is parallel to defining liability in our setting (i.e., prohibited pricing algorithms) - and, secondly, ensuring accountability, in the sense that the machine learning algorithm satisfies that definition of fairness - which is parallel to defining evidentiary methods (i.e., how to determine whether a pricing algorithm is prohibited). Accountability can be challenging. If fairness mandates that decisions not be based on some trait,

⁵⁰ Machine learning is another term used to capture autonomous or semi-autonomous agents that learn. For some recent research on fairness and machine learning, see Goodman (2016), Johnson, Foster, and Stine (2016), Joseph, Kearns, Morgenstern, and Roth (2016), and Kleinberg, Mullainathan, and Raghavan (2016).

such as race, it need not be sufficient to constrain machine learning to operating on a data set that excludes race. If there are other variables in the data set that correlate with race - for example, residential location, income, and education - then machine learning may figure out how to indirectly condition on race.⁵¹

Analogous to this on-going research program to restrict machine learning in order to avoid unfair discrimination, I am proposing here a research program for restricting AAs not to collude, and detecting them when they do collude. While the issue of fairness involves a single AA, collusion involves multiple interacting AAs which makes for a more challenging problem. Ensuring fairness means constraining an AA so it does not condition on certain traits of a person. Preventing collusion means constraining an AA so it does not condition its actions on how rival firms' AAs will respond to those actions in a manner that supports supracompetitive prices. An AA is "fair" if its recommendation is not dependent on, say, a person's gender. An AA is "not collusive" if its price recommendation is not dependent on rival firms' responding in a particular manner; for example, a price increase is not contingent on rival firms subsequently matching that price, or maintaining price is not contingent on rival firms conducting a price war if price were to be reduced.

It will require the execution of a research program to properly identify appropriate sets of prohibited pricing algorithms. Here, I provide the broad outline of what such a program will entail.

Step 1: Create a simulated market setting with learning algorithms that produce collusion and competition as outcomes.⁵² Keep track of when competitive prices emerge and when supracompetitive prices emerge. Perform this exercise with different learning algorithms and for

⁵¹ These fairness issues could also pertain to pricing algorithms that personalize price. What I have been referring to as an efficiency benefit for an AA, on the grounds that it enhances a firm's profit, may actually be inconsistent with fairness. On the other hand, price discrimination typically benefits consumers with a low willingness-to-pay (WTP) and harms consumers with a high WTP. If WTP is positively correlated with income, some personalized pricing may actually serve distributional goals.

⁵² Ezrachi and Stucke (2016) refer to it as a "collusion incubator."

a variety of market conditions. (This first step would also serve to shed light on how easily AAs can produce collusion and the types of markets for which collusion by AAs is likely.)

Step 2: Inspect or test the resulting pricing algorithms for the purpose of identifying those properties that are present when supracompetitive prices emerge but are not present when competitive prices emerge. Pricing algorithms with those properties will have a high likelihood ratio and thus be a candidate for the set of prohibited pricing algorithms.

Step 3: Test the effect of prohibiting a set of pricing algorithms. This would be done by re-running the learning algorithms in the simulated market setting but where the learning algorithms are constrained not to select pricing algorithms in the prohibited set. What we would want to see is that supracompetitive prices are less frequent and competitive prices are not distorted. A generally desirable property is that it is more likely that prices are lower and welfare is higher when some pricing algorithms are prohibited.

Admittedly, there are significant challenges with operationalizing this approach. However, those challenges do not appear insurmountable, and solutions to daunting problems are only found after we become immersed in trying to solve them. We know far too little about algorithmic collusion to be dismissing any approach for dealing with it, and a per se prohibition on collusive pricing algorithms is the only viable approach currently available.

6.4 Legality of a Prohibition on Pricing Algorithms

Finally, let me discuss whether the proposed definition of liability is supported by existing laws in the United States. The prohibition of certain pricing algorithms would seem inconsistent with jurisprudence regarding Section 1 of the Sherman Act. Firms could be pricing according to a prohibited pricing algorithm while not having an agreement, because those algorithms were selected by AAs. However, a prohibition on certain pricing algorithms could come under Section 5 of the FTC Act which states: "Unfair methods of competition in or affecting commerce, and

unfair or deceptive acts or practices in or affecting commerce, are hereby declared unlawful." The properties of pricing algorithms that result in a reward-punishment scheme supporting supracompetitive prices could be interpreted as an "unfair method of competition".

While Section 5 of the FTC Act has largely been used in cartel cases when there is an "invitation to collude" but no evidence of acceptance of that invitation (so that, in the court's view, communications did not result in an agreement), there could be an expanded role for the FTC in having it prosecute cases of collusion by AAs.⁵³ Pertinent to this issue, the FTC recently issued guidelines for the use of Section 5:

In deciding whether to challenge an act or practice as an unfair method of competition in violation of Section 5 on a standalone basis, the Commission adheres to the following principles: the Commission will be guided by the public policy underlying the antitrust laws, namely, the promotion of consumer welfare; the act or practice will be evaluated under a framework similar to the rule of reason, that is, an act or practice challenged by the Commission must cause, or be likely to cause, harm to competition or the competitive process, taking into account any associated cognizable efficiencies and business justifications; and the Commission is less likely to challenge an act or practice as an unfair method of competition on a standalone basis if enforcement of the Sherman Act or Clayton Act is sufficient to address the competitive harm arising from the act or practice.⁵⁴

Using Section 5 to prohibit collusive pricing algorithms falls within these guidelines with the exception of the guidelines' focus on the rule of reason. It is certainly consistent with the approach laid out here to define a set of pricing algorithms that, while not *per se* prohibited, is subject to the rule of reason. In that case, the FTC would have to balance any efficiency benefits

⁵³ Independently, Ezrachi and Stucke (2017) also make this suggestion.

⁵⁴ "Statement of Enforcement Principles Regarding 'Unfair Methods of Competition' Under Section 5 of the FTC Act" (August 13, 2015)

from the pricing algorithm against any proclivity towards collusion. However, as discussed above, the properties that promote collusion are likely to be quite distinct from those that enhance efficiency. To what extent *per se* illegality or a rule of reason is appropriate depends on the outcome of the research program and what we learn about the effects of various pricing algorithms.

The FTC may then have a legal mandate and, in terms of expertise, the FTC could well be the agency most qualified to identify and prosecute collusion in online markets by AAs. In pursuing consumer protection, the FTC has had many cases involving online practices regarding privacy and data security. As noted in its 2016 Privacy & Data Security Update, the FTC has brought enforcement actions relating to "spam, social networking, behavioral advertising, pretexting, spyware, peer-to-peer file sharing, and mobile."⁵⁵ Given this developed expertise for online markets and automated processes, the FTC is in a good position to build on that base of knowledge so as to define and enforce a prohibition of collusive pricing algorithms.

7 Concluding Remarks

Supracompetitive prices are the result of collusion, where collusion is the use of (collusive) strategies that embody a reward-punishment. Collusion is facilitated by communication, where communication involves firms' representatives actively seeking to achieve a mutual understanding to adopt collusive strategies. Thus, communication causes collusion and collusion causes supracompetitive prices. The law objects to neither supracompetitive prices nor collusion but rather to communication; that is, it makes the process by which collusion is achieved illegal. Consequently, collusion without communication is lawful, such as conscious parallelism. Hence, collusion by autonomous artificial agents is lawful because the communication to which courts object is absent, and there is no liability because autonomous artificial agents lack the capacity to

⁵⁵ Recent FTC reports include "Businesses Can Help Stop Phishing and Protect Their Brands Using Email Authentication" (FTC Staff Perspective, March 2017), "The 'Sharing' Economy: Issues Facing Platforms, Participants & Regulators" (FTC Staff Report, November 2016), and "Cross-Device Tracking" (FTC Staff Report, January 2017) which examines issues related to the tracking of consumer behavior across multiple Internet-connected devices.

mutually understand and thus cannot have an agreement. The reason that courts have focused on making communication illegal is because it is difficult to determine the latent strategies underlying observed prices. Courts cannot read the minds of those who are choosing prices, and use communication in its place. In contrast, we can, in principle, "read the mind" of an autonomous artificial agent by reading its code or "simulating its mind" by entering input and observing output. It then becomes possible to determine if the prices observed are only sustained because of a reward-punishment scheme. When conducted by autonomous artificial agents, collusion can be made illegal, rather than the communication that facilitates it.

Even accepting that argument, the charitable skeptic might say: "You have proposed a possibly workable approach to finding a solution to a problem that may not even exist." Indeed, there is currently no evidence of collusion by autonomous artificial price-setting agents in actual markets, and research has yet to be conducted to investigate whether such collusion can occur in a reasonably sophisticated simulated market. One does not have to be oppositional to question the paper's underlying premise: Autonomous artificial price-setting agents can learn to collude in real-world markets.

While I share some of that skepticism, I am not reassured in light of recent experiences with information technology. One lesson learned is that the path of IT is difficult to predict. Looking back two decades, the extent of market dominance that we have witnessed in online markets was not anticipated. While it was recognized that there would be some dominance due to network effects, the emphasis was largely on the low cost of entry into online markets, as opposed to conventional markets, and the ease with which consumers could search, all of which was thought would promote intense competition. While those forces have not been absent, other forces have proven more determinative as reflected in the rise of such dominant firms as eBay, Facebook, Google, Netflix, Airbnb, and Uber. Of particular note, it was quite unexpected that market dominance would be a feature of general retailing, as has emerged with the dominant position of Amazon. In retrospect, the disruptive technological change associated with online markets made it difficult to accurately predict future outcomes. Can we be so assured that collusion in online markets will not prove ubiquitous? A second lesson is that change in IT can be rapid. Ten years

ago, who would have thought that we might be on the verge of self-driving cars? If autonomous cars can navigate city roads and traffic, is it that difficult to imagine autonomous artificial agents figuring out how to collude? Can we really be so sure that collusion by autonomous artificial agents will never be commonplace?

[W]hen we look at the challenges for cartel enforcement in the future, one of the biggest things we need to deal with is the risk that automated systems could lead to more effective cartels. ... So far, those cases have dealt with agreements that were put together by humans. The computers only took over when it was time to put them into practice. It's true that the idea of automated systems getting together and reaching a meeting of minds is still science fiction. ... But we do need to keep a close eye on how algorithms are developing ... so that when science fiction becomes reality, we're ready to deal with it.

- Margarethe Vestager, European Commissioner for Competition⁵⁶

⁵⁶ Margarethe Vestager, "Algorithms and Competition," Remarks by the European Commissioner for Competition at the Bundeskartellamt 18th Conference on Competition, Berlin, March 16, 2017.

References

- [1] Baker, Jonathan B., "Two Sherman Act Section 1 Dilemmas: Parallel Pricing, the Oligopoly Problem, and Contemporary Economic Theory," *Antitrust Bulletin*, 38 (1993), 143-220.
- [2] Ballard, Dylan I. and Amar S. Naik, "Algorithms, Artificial Intelligence, and Joint Conduct," *CPI Antitrust Chronicle*, May 2017, 29-35.
- [3] den Boer, Arnoud V., "Dynamic Pricing and Learning: Historical Origins, Current Research, and New Directions," *Surveys in Operations Research and Management Science*, 20 (2015), 1-18.
- [4] Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello, "Algorithmic Pricing and Collusion: What Implications for Competition Policy?," University of Bologna, working paper, March 2018.
- [5] Capobianco, Antonio and Pedro Gonzaga, "Algorithms and Competition: Friends or Foes?," *Competition Policy International*, August 2017.
- [6] Deng, Ai, "When Machines Learn to Collude: Lessons from a Recent Research Study on Artificial Intelligence," Bates White Economic Consulting, August 30, 2017.
- [7] Deng, Ai, "An Antitrust Lawyer's Guide to Machine Learning," Bates White Economic Consulting, January 9, 2018.
- [8] Desai, Deven R. and Joshual A. Kroll, "Trust but Verify: A Guide to Algorithms and the Law," working paper, Georgia Institute of Technology, April 2017 (*Harvard Journal of Law & Technology*, forthcoming).
- [9] Diakopoulous, Nicholas, "Algorithmic Accountability Reporting: On the Investigation of Black Boxes," Tow Center for Digital Journalism, Columbia University, 2014.
- [10] Ezrachi, Ariel and Maurice E. Stucke, "Artificial Intelligence and Collusion: When Computers Inhibit Competition," Oxford Legal Studies Research Paper No. 18/2015, University of Tennessee Legal Studies Research Paper No. 267, May 2015.
- [11] Ezrachi, Ariel and Maurice E. Stucke, *Virtual Competition: The Promise and Perils of the Algorithm-Driven Economy*, Cambridge, Mass.: Harvard University Press, 2016.

- [12] Ezrachi, Ariel and Maurice E. Stucke, "Algorithmic Collusion: Problems and Counter Measures," OECD Background Paper, Roundtable on Algorithms and Collusion, DAF/COMP/WD(2017) 25, June 2017.
- [13] Friedman, James W., "A Non-cooperative Equilibrium for Supergames," *Review of Economic Studies*, 38 (1971), 1-12.
- [14] Gal, Avigdor, "It's a Feature, not a Bug: On Learning Algorithms and What They Teach Us," OECD Background Paper, Roundtable on Algorithms and Collusion, DAF/COMP/WD(2017) 50, June 2017.
- [15] Gal, Michal S., "Algorithmic-Facilitated Coordination: Market and Legal Solutions," *CPI Antitrust Chronicle*, May 2017a, pp. 22-28.
- [16] Gal, Michal S., "Algorithmic-facilitated Coordination," OECD Background Paper, Roundtable on Algorithms and Collusion, DAF/COMP/WD(2017)26, June 2017b.
- [17] Gal, Michal S. and Niva Elkin-Koren, "Algorithmic Consumers," *Harvard Journal of Law & Technology*, 30 (2017), 1-45.
- [18] Gershman, Samuel J., "Reinforcement Learning and Causal Models," Harvard University, December 2015.
- [19] Goodman, Bryce W., "Economic Models of (Algorithmic) Discrimination," 29th Conference on Neural Information Processing Systems, June 2016.
- [20] Gosselin, Sathya S., April Jones, and Annabel Martin, "Are Antitrust Laws Up to the Task? A US/EU Perspective on Anti-Competitive Algorithm Behavior," Hausfeld LLP (accessed online at <www.hausfeld.com/news/us/are-antitrust-laws-up-to-thetask-a-us-eu-perspective-on-anti-competitive-a>on August 17, 2017)
- [21] Günster, Andrea and Stephen Martin, "A Holy Alliance: Collusion in the Renaissance Europe Alum Market," *Review of Industrial Organization*, 47 (2015), 1-23
- [22] Hanaki, Nobuyuki, Rajiv Sethi, Ido Erev, and Alexander Peterhansl, "Learning Strategies," *Journal of Economic Behavior & Organization*, 56 (2005), 523-542.
- [23] Harrington, Joseph E., Jr., "Game-Theoretic Ruminations on Section 1 of the Sherman Act," Bayard Wickliffe Heath Memorial Lecture, University of Florida Levin College of Law,

March

2012.

<[https://joeharrington5201922.github.io/pdf/Bayard%20Wickliffe%20Heath%20Memorial%20Lecture 03.12.pdf](https://joeharrington5201922.github.io/pdf/Bayard%20Wickliffe%20Heath%20Memorial%20Lecture%2003.12.pdf)>

- [24] Harrington, Joseph E., Jr., *The Theory of Collusion and Competition Policy*, Cambridge, Mass.: The MIT Press, 2017.
- [25] Hovenkamp, Herbert, *Federal Antitrust Policy: The Law of Competition and Its Practice*, 5th edition, St. Paul, Minnesota: West Academic Publishing, 2016.
- [26] Johnson, Kory D., Dean P. Foster, and Robert A. Stine, "FairnessAware Data Mining: Ensuring Fairness in Arbitrary Models," The Wharton School, University of Pennsylvania, March 2016.
- [27] Joseph, Matthew, Michael Kearns, Jamie Morgenstern, and Aaron Roth, "Fairness in Learning: Classic and Contextual Benefits," Department of Computer and Information Sciences, University of Pennsylvania, May 2016.
- [28] Kaplow, Louis, *Competition Policy and Price Fixing*, Princeton: Princeton University Press, 2013.
- [29] Kaplow, Louis, "Likelihood Ratio Tests and Legal Decision Rules," *American Law and Economics Review*, 16 (2014), 1-39.
- [30] Kleinberg, Jon, Sendhil Mullainathan, and Manish Raghavan, "Inherent Trade-Offs in the Fair Determination of Risk Scores," Cornell University, September 2016.
- [31] Kovacic, William E., "The Identification and Proof of Horizontal Agreements Under the Antitrust Laws," *Antitrust Bulletin*, 38 (Spring 1993), 5-81.
- [32] LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton, "Deep Learning," *Nature*, 521 (2015), 436-444.
- [33] Mehra, Salil K., "De-Humanizing Antitrust: The Rise of the Machines and the Regulation of Competition," Temple University Legal Studies Research Paper No. 2014-43, September 2014.
- [34] Mehra, Salil K., "Antitrust and the Robo-Seller: Competition in the Time of Algorithms," *Minnesota Law Review*, 100 (2016), 1323-1375.

- [35] Mehra, Salil K., "Robo-Seller Prosecutions and Antitrust's Error-Cost Framework," *CPI Antitrust Chronicle*, May 2017, 36-39.
- [36] Miller, John H., "The Coevolution of Automata in the Repeated Prisoner's Dilemma," *Journal of Economic Behavior & Organization*, 29 (1996), 87-112.
- [37] Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis, "Human-level Control through Deep Reinforcement Learning," *Nature*, 518 (2015), 529- 533.
- [38] Motta, Massimo, *Competition Policy: Theory and Practice*, Cambridge: Cambridge University Press, 2004.
- [39] Nambiar, Mila, David Simchi-Levi, and Ha Wang, "Dynamic Learning and Price Optimization with Endogeneity Effect," Operations Research Center, MIT, November 2016.
- [40] OECD Background Note by the Secretariat, "Algorithms and Collusion," Roundtable on Algorithms and Collusion, DAF/COMP(2017)4, June 2017.
- [41] Okuliar, Alex and Elena Kamenir, "Pricing Algorithms: Conscious Parallelism or Conscious Commitment?," *Competition Policy International*, July 2017.
- [42] Oxera Consulting, "When Algorithms Set Prices: Winners and Losers," Discussion Paper, 19 June 2017.
- [43] Page, William H., "Communication and Concerted Action," *Loyola University Chicago Law Journal*, 38 (2007), 405-460.
- [44] Preston, John and Mark Bishop, eds., *Views into the Chinese Room: New Essays on Searle and Artificial Intelligence*, Oxford: Clarendon Press, 2002.
- [45] Scherer, Matthew U., "Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies," *Harvard Journal of Law & Technology*, 29 (2016), 354-400.
- [46] Searle, John R., "Minds, Brains, and Programs," *Behavioral and Brain Sciences*, 3 (1980), 417-424.

- [47] Searle, John R., "The Chinese Room," in *The MIT Encyclopedia of the Cognitive Sciences*, R.A. Wilson and F. Keil, eds., Cambridge, MA: MIT Press, 1999.
- [48] Searle, John R., "The Chinese Room," in *Views into the Chinese Room: New Essays on Searle and Artificial Intelligence*, J. Preston, and M. Bishop, eds., Oxford: Clarendon Press, 2002, 51-69.
- [49] Stigler, George J., "A Theory of Oligopoly," *Journal of Political Economy*, 72 (1964), 44-61.
- [50] Sutton, Richard S. and Andrew G. Barto, *Reinforcement Learning: An Introduction*, Cambridge, Mass.: The MIT Press, 2000.
- [51] Turner, Donald F., "The Definition of Agreement under the Sherman Act: Conscious Parallelism and Refusals to Deal," *Harvard Law Review*, 75 (1962), 655-706.
- [52] U.S. Department of Justice, "Algorithms and Collusion - Note by the United States," OECD Background Paper, Roundtable on Algorithms and Collusion, DAF/COMP/WD(2017)41, June 2017.
- [53] Viscusi, W. Kip, Joseph E. Harrington, Jr., and David E. M. Sappington, *Economics of Regulation and Antitrust*, 5th Edition, Cambridge, Mass.: The MIT Press, 2018.
- [54] Vives, Xavier, *Oligopoly Pricing: Old Ideas and New Tools*, Cambridge, Mass.: The MIT Press, 1999.
- [55] von Neumann, John, "First Draft of a Report on the EDVAC," reprinted (in part) in *Origins of Digital Computers: Selected Papers*, B. Randell, ed., Berlin: Springer-Verlag, 1982, 383-392.
- [56] Werden, Gregory J., "Economic Evidence on the Existence of Collusion: Reconciling Antitrust Law with Oligopoly Theory," *Antitrust Law Journal*, 3 (2004) 719-800.
- [57] Yao, Dennis A. and Susan S. DeSanti, "Game Theory and the Legal Analysis of Tacit Collusion," *Antitrust Bulletin*, 38 (1993), 113-141.