Developing Competition Law for Collusion by Autonomous Artificial Agents*

Joseph E. Harrington, Jr.

Department of Business Economics & Public Policy

The Wharton School, University of Pennsylvania

harrij@wharton.upenn.edu

23 May 2018

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.

*This paper has benefitted from the comments of 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 Takeaki Sunada for his research assistance.

Executive Summary and a Roadmap

Consider managers of rival firms independently adopting learning algorithms for the purpose of setting prices. In other words, software programs in the form of autonomous artificial agents (AAs), rather than human agents, are determining which prices to charge. Suppose the AAs end up charging supracompetitive prices because they came to adopt collusive pricing rules. The objective of this paper is to develop a path towards making such collusion unlawful.

Using a simple stylized model, the paper begins by exemplifying how AAs can collude. The case is then made that jurisprudence under section 1 of the Sherman Act does not cover collusion by AAs because "there must be evidence that tends to exclude the possibility that the [firms] were acting independently" when, in fact, the managers did act independently in the adoption of AAs. Furthermore, there is no overt act of communication which is generally seen as a required piece of evidence. But the problem goes deeper than a lack of evidence in that there is no agreement for there is no "meeting of minds" or "conscious commitment to a common scheme."

Having reached this legal dead end, I then go back to first principles regarding what exactly is collusion. Collusion is often misperceived to be supracompetitive prices but that is actually the result of collusion. Collusion is about a firm *causing* rival firms to set supracompetitive prices. More specifically, 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.

Though the use of particular pricing strategies is the essence of collusion, jurisprudence does not focus on identifying whether firms use such pricing strategies because of the associated error costs. While prices are observable, the reward-punishment scheme supporting them is not. The reward-punishment scheme is latent - inside the managers' heads - and can only be indirectly observed when those schemes interact with the environment to produce prices and quantities. Due to the difficulty in identifying firms' latent strategies, evidentiary standards require more than economic evidence; they require an overt act of communication that firms are seeking to coordinate their behavior. In a sense, it is communication that allows an outside observer to get inside the collective heads of managers and assess why they are pricing the way they are and, in particular, whether it is driven by the mutual adoption of strategies to restrain competition. To be clear, I am not describing the reasoning used by courts in adopting these evidentiary standards but rather why evidentiary standards had to be defined this way.

The legal doctrine that has just been described is based on collusion 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. The situation is fundamentally different when prices are set by AAs. The strategy determining price is written down in the algorithm's code which means that it can, in principle, be accessed. Based on that crucial distinction, the following approach is put forth to make collusion by AAs unlawful. Liability is defined by a *per se* prohibition on certain pricing algorithms that support supra-competitive prices. In defining evidentiary standards, 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.

The feasibility of this approach is discussed in terms of AAs that have been used to set prices, and might be used in the future. The approach's practicality is admittedly an open question and, towards addressing that question, a research program is laid out to investigate it.

Outline

1 Introduction

Consider an online market involving two competitors, referred to as firms A and B, 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 A 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 B's prices are available online), firm A'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 A'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 A's manager to adopt the learning algorithm, observe how it performs over the next year, and then decide whether to retain it.

Unbeknownst to the manager of firm A, the manager of firm B 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 A's manager, firm B'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 A's manager did not know that firm B's manager was soon to adopt a learning algorithm, firm B's manager adopted the learning algorithm not knowing that firm A'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 select prices as 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.[1] More specifically, the prices are supracompetitive because 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 feasible for the simultaneous adoption of learning algorithms by competitors to produce collusion. Rather than collusion occurring with human agents, collusion is occurring with autonomous arti-

[1] 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.

ficial 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 moves among pricing rules in order to yield higher profits. While the capacity 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 to future research. Given the advances of AAs more generally, the emergence of collusion by AAs in actual market settings would seem highly possible in the near future, if it is not already occurring. 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 pricing algorithms is rapidly gaining attention, as evidenced by recent speeches by the European Commissioner for Competition and the Chair of the Federal Trade Commission.[2] 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

[2] Margarethe Vestager, "Algorithms and Competition," Remarks by the European Commissioner for Competition at the Bundeskarellamt 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 Commssion at the "Concurrences Antitrust in the Financial Sector Conference," New York, May 23, 2017.

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. So as to be more concrete about the phenomenon, Section 2 exemplifies collusion by AAs for a simple model. To lay the legal groundwork, Section 3 argues that collusion by AAs is legal in the United States under current jurisprudence. To be clear, "collusion by AAs" refers to the sort described above which means the managers who implemented the programs were unaware of what the programs would do. For the purpose of developing an approach to make collusion by AAs unlawful, I go back to first principles with respect to collusion in Section 4. There it is argued that collusion is not about a firm setting a supracompetitive price but rather about a firm *causing* rival firms to set supracompetitive prices. Collusion is then a causal relationship between firms' prices that results in supracompetitive prices. Though that is the essence of collusion, Section 5 explains that jurisprudence does not focus on identifying whether such a causal relationship exists because of the associated error costs. It is due to those error costs that communication, while ancillary to collusion, is made a requisite part of evidentiary standards. The paper's proposal for defining legal standards for collusion by AAs is provided in Section 6. The case is made that liability and evidentiary standards that are ineffective when human agents collude, could be effective when software agents collude. After defining, in general terms, the proposed liability and evidentiary standards, a research program for operationalizing them is offered. Section 7 concludes and, in particular, offers a cautionary note if one is inclined to dismiss the prospect of collusion by AAs. While the papers covers some technical topics, it is written in a non-technical manner and is intended to be understood by a general audience of economists and lawyers.[3]

2 Emergence of Collusion with Autonomous Artificial Agents

Taking it as a postulate that AAs can collude, the objective of this paper is explore the development of a legal approach to making such conduct illegal. It is a separate research project to investigate how easily and frequently AAs can collude, and identifying the manner in which they collude. In the meantime, let me offer here a simple example illustrating the phenomenon and, at least in the context of a rarefied environment, showing its feasibility. This section also offers a review of the economic theory of collusion.

2.1 Collusion with Human Agents: The Economic Theory of Collusion

Consider the Prisoners' Dilemma model of price competition between two competitors, which is depicted as Figure 1. There are two firms - A and B - that simultaneously choose between a *low* price and a *high* price. In the context of Figure 1, firm A chooses a row and firm B chooses a column. The pair of entries in a cell are the profits earned by the firms where the first (second) entry is the profit earned by firm A (B). For example, if both choose the high price then each earns profit of π^c (where c is for "collusion"), while if both choose the low price then each earns π^n (where n is for "not collusion"). If firm A chooses low and firm B chooses high then firm A earns π^d (where d is for "deviation") and firm B earns π^s (where s is for "sucker").[4]

The following relationship between these four values is assumed: $\pi^d > \pi^c > \pi^n > \pi^s$. $\pi^c > \pi^n$ means that each firm earns higher profit when both chooses high prices

[3] Any technical material is in footnotes that the reader can skip.

[4] The appropriateness of the labels "deviation" and "sucker" will be made apparent later.

(and thus earn π^c) than when both choose low prices (and thus earn π^n). The highest profit is earned when a firm prices low and its rival prices high, which yields π^d. In that situation, the lower-priced firm is presumed to pick up a lot of market share. By analogous logic, the lowest profit is earned when a firm prices high and its rival prices low, which yields π^s.[5] If the reader prefers numbers to symbols then suppose $\pi^d = 4, \pi^c = 3, \pi^n = 2, \pi^s = 1$.

Figure 1

Firm B

Firm A		*high price*	*low price*
	high price	π^c, π^c	π^s, π^d
	low price	π^d, π^s	π^n, π^n

Before examining how AAs might play this game, let us consider the standard economic model of firm behavior based on prices being selected by profit-maximizing managers. If the game is played just once then a strategy for a firm is simply a price, low or high. A stable pair of strategies is one in which each firm's strategy maximizes its profit given the other firm's strategy; that is, when it chooses its price, a firm is acting rationally given accurate beliefs as to the price chosen by its rival. Such a stable pair of strategies is referred to as a Nash equilibrium in game theory.

The game in Figure 1 has a unique Nash equilibrium which is for both firms to select a low price. This result is immediate because a low price is optimal for a firm regardless of what the other firm is anticipated to do. If firm B is expected to set a high price then firm A's profit is higher with a low price, which yields π^d, than with a high price, which yields π^c (recall that $\pi^d > \pi^c$). Similarly, if firm B is expected to set a low price then firm A's profit is higher with a low price, π^n, than with a high price, π^s.

The unique Nash equilibrium is for both firms to set low prices, and that is the

[5] It is also typical to assume $2\pi^c > \pi^d + \pi^s$, so joint profits are maximized by both pricing high as opposed to having one firm price high and the other price low.

outcome that is typically taken to be what competition between firms would yield. However, that outcome is not very satisfying to the firms because they know that their profits could be increased from π^n to π^c if they were both to raise price. The prospect of that collective improvement creates an incentive to coordinate, and is the basis for collusion.

Towards providing a model that yields a collusive outcome, we must first enrich the setting with the descriptively realistic assumption that firms do not interact just once but rather repeatedly over time. What is important for the ensuing analysis is that the horizon is indefinite, which means that the firms do not know if or when their interactions will end. It will simplify our discussion if it is taken to be infinite in length so that firms never anticipate it ending.[6] In this infinite horizon setting, assume firms act to maximize the sum of their profits over time.[7]

[6]The ensuing results and insight extend to when the end of the horizon is finite and unknown as long as there is no time at which the firms are absolutely certain that they will no longer be interacting, which seems compelling for most markets.

[7]As just stated, this specification of what firms maximize is problematic because it entails an infinite sum of positive numbers which equals infinity. That means every strategy yields the same payoff of infinity in which case firms are indifferent among them. However, the reader should not be concerned with that technical issue, for the discussion in the paper is robust to a slight modification which is immune to that concern. For those who are interested in learning more, please read on; otherwise, return to the text. What we require is for a firm to evaluate a stream of profits in terms of its *weighted* sum (rather than a simple sum) where smaller weights are given to profits in the more distant future and those weights approach zero fast enough that the weighted sum is finite. The usual specification is that \$1 received in the next period is worth \$$\delta$ in the current period where $0 < \delta < 1$. δ is the "discount factor" where $1 - \delta$ is the amount by which money tomorrow is discounted compared to today. \$1 received in two periods is worth $\delta \times \delta$ or δ^2 today, and in three periods is worth δ^3 compared to today, and so forth. As δ is between 0 and 1, note that δ^2 is smaller than δ, and δ^3 is smaller than δ^2. For example, if $\delta = 0.8$ then the firm values \$1 in the next period the same as 80 cents today, and \$1 in two periods the same as 64 cents today. With this specification, the infinite profit stream $(\pi_1, \pi_2, \pi_3, ...)$, where π_t is the profit in period t, is valued by a firm as $\pi_1 + (\delta \times \pi_2) + (\delta^2 \times \pi_3) + \cdots$. A firm is assumed to choose a strategy that yields the profit stream with the highest weighted sum, as given by that formula. All of the discussion in the

Each period, firms choose prices and realize profits as shown in Figure 1, while anticipating that they will repeat this interaction *ad infinitum.* In the context of this intertemporal setting, a strategy is not a price, nor need it be a sequence of prices. Rather, a strategy assigns a price in each period contingent on the history of play; that is, contingent on past prices. One possibility is to ignore history in which case a firm's strategy is just a sequence of prices, low or high, that can vary over time but does not depend on what has transpired. There is a unique Nash equilibrium for the infinitely repeated Prisoners' Dilemma when firms use such history-independent strategies, and it is to price low in every period; in other words, a repetition of the Nash equilibrium for the one-shot game.[8] For if firm A expects firm B to always charge a low price then firm A can do no better than to charge a low price itself in every period. If it was to deviate from a low price in some period, it would then reduce its profit in that period from π^n to π^s without any offsetting benefit in the future. Competition - in the form of low prices - is then one outcome when firms interact repeatedly, and is the only outcome when firms' prices do not depend on the history.

For firms to support supracompetitive prices, each must use a strategy that makes history matter - that is, price is contingent on past prices - in order to provide an incentive to its rival to set a high price. Consider the following strategy which will be referred to as the "punish T times" strategy, where T is a positive integer (1, 2, 3, ...) . At the start of play, this strategy has a firm choose the high price. It continues to choose the high price as long as both firms have charged the high price in the past. If, at any time, one or both firms chose the low price then a firm sets a low price for

paper applies to this specification as long as δ is close to one. To avoid complicated expressions, we set $\delta = 1$ in the text but, technically speaking, results require δ to be close to, but less than, 1.

[8] To be clear, a firm finds that strategy to be preferable to any other strategy including strategies that make a firm's price contingent on the history. In other words, if the rival firm's price is not contingent on the history then there is no reason for a firm to make its price contingent on the history.

T periods, after which it returns to setting a high price. If there should be another deviation - that is, a firm prices low when it should have priced high - then again there is a punishment in the form of reverting to a low price for T periods. Every time there is a deviation from the high (collusive) price, a firm reverts to the low (competitive) price for T periods.

First note that if both firms use the "punish T times" strategy then each will price high in the first period and, given that both priced high in the first period, each will price high in the second period (as prescribed by the strategy) and, given that both priced high in the first two periods, each will price high in the third period and so forth. Hence, collusive prices will ensue if both use this strategy. The next task is to ensure that each firm wants to use the "punish T times" strategy given the other firm is expected to use it. If firm A was to use this strategy, given firm B is conjectured to use it, then firm A would expect to receive profit of π^c in every period from both firms charging high prices. Any meaningful alternative strategy is to price low in some period and there is no loss of generality in supposing it is the first period. That low price will deliver higher profit of π^d in the current period (here, the firm has performed a "deviation" from the collusive outcome which is why the superscript d is used). However, recall that firm B's strategy has it respond with the low price in the next T periods. Given that response by firm B, firm A can do no better than to charge the low price as well during those subsequent T periods. After the punishment runs it course for T periods, both firms will return to pricing high.[9]

In comparing the profit streams from the high price with that from the low price in period 1, note that they are identical come period $T+2$ as, in either case, firms will be setting the high price and earning profit of π^c. The profit streams only differ with regards to profits over the first $T+1$ periods. When firm A sets the high price, as prescribed by the "punish T times" strategy, it earns π^c in each of those $T+1$

[9] It can be shown that if it is optimal for a firm to price high in the first period then it will be optimal for it to return to pricing high after a punishment.

periods, while if it sets the low price in the first period then it earns π^d in the first period and π^n in the next T periods (as both firms select the low price). Hence, initially charging the high price yields a higher total profit over the first $T+1$ periods compared to initially charging the low price when

$$\pi^c + \underbrace{\pi^c + \cdots + \pi^c}_{T \text{ periods}} > \pi^d + \underbrace{\pi^n + \cdots + \pi^n}_{T \text{ periods}}$$

or

$$\pi^c + (T \times \pi^c) > \pi^d + (T \times \pi^n).$$

Re-arranging the inequality, we have

$$(T \times \pi^c) - (T \times \pi^n) > \pi^d - \pi^c \text{ or } T > \frac{\pi^d - \pi^c}{\pi^c - \pi^n} \equiv T^*.$$

Note that $\pi^d - \pi^c > 0$ and $\pi^c - \pi^n > 0$ so $\frac{\pi^d - \pi^c}{\pi^c - \pi^n}$ (which is denoted as T^*) is a positive number. If the punishment is sufficiently long - that is, T exceeds T^* - then firm A finds it optimal to set a high price in the first period, as prescribed by "punish T times" strategy. The same argument can be used to rationalize charging the high price in every period. Hence, it is optimal for a firm to use the "punish T times" strategy given the other firm is conjectured to use it, as long as the punishment from setting a low price is sufficiently severe; that is, it lasts more than T^* periods. Alternatively stated, it is a Nash equilibrium for the two firms to use the "punish T times" strategy when $T > T^*$. We then have a stable situation in which firms set supracompetitive prices.

Intuitively, a firm is willing to set the high price, even though it could earn more profits in the current period by undercutting the other firm's high price, because it anticipates a punishment in the form of lower profits in the future. It believes that the other firm will retaliate in response to this deviation by lowering its future prices and, even if the firm responds optimally by matching that low price, its profit will be reduced from π^c to π^n in each of those T periods. Collusion then involves a "reward-punishment" scheme in that a firm is rewarded by pricing high today - in that the rival firm will price high tomorrow - and it is punished by pricing low today - in that the rival firm will price low tomorrow. Though cast in the simple Prisoners' Dilemma, this analysis captures the essence of the economic theory of collusion.

2.2 Collusion with Autonomous Artificial Agents

Collusion has just been shown to be a market outcome for when prices are set by rational, forward-looking agents. In reality, those agents are well-incentivized profit-maximizing strategically-minded company employees. But what if prices are set by software programs that randomly experiment in search of better pricing rules? How easily can collusion emerge and be sustained? To provide a preliminary answer to that question, consider the analysis of Hanaki et al (2005). Though the learning that is described below is not the way in which AAs will learn to collude in actual markets (as the learning algorithms currently used are different), it serves our purposes in conveying the principle with minimal mathematical fuss. A broader discussion of learning algorithms is provided in Section 6.

An AA is specified to have two components. First, it is a pricing algorithm that prescribes what price to set depending on the history. Second, it is a learning algorithm that modifies the pricing algorithm based on a pricing algorithm's performance relative to the performance of other pricing algorithms. A pricing algorithm is specified to be a finite automaton.[10] A finite automaton is defined by five elements: 1) a set of states; 2) an initial state; 3) a set of actions; 4) an output function that assigns an action to a state; and 5) a transition function that describes how the state changes depending on the current state and the current actions of the agents. A state can be thought of as a summary variable for the history, while an action is a price in our setting. While the definition of a finite automaton may seem rather abstract, it is actually a fairly simple object which is best conveyed with some examples.

Many readers will know of the strategy "tit-for-tat" which, in the context of the

[10] For a general reference for finite automata, see Lewis and Papadimitriou (1981).

game in Figure 1, prescribes the high price in the first period and then, in any future period, prescribes a price equal to that which the other firm set in the previous period. Figure 2 portrays tit-for-tat as a finite automaton. A box represents a state and the price inside the box is the action prescribed when the automaton is at that state. The transition function between states is represented by the arrows between states. An arrow from one state to another state represents the movement from the former to the latter state when a particular price is chosen by the other firm, which is adjacent to the arrow. Let us label the state in which the firm sets the high price as the "collusive" state, and the state in which it sets the low price as the "punishment" state. If the current state of firm A's pricing automaton is "collusive" then it sets the high price. Inspecting Figure 2, if firm B chose the high price, then the state remains "collusive" (which means firm A will set a high price in the next period). This transition is depicted in Figure 2 by the arrow that goes from the "collusive" state back to itself when the rival firm's price is high. If instead firm B chose the low price then the state of firm A's automaton switches to "punishment" (note the arrow from the "collusive" state to the "punishment" state associated with the rival choosing the low price) which means firm A will set a low price next period. Note that firm A's price next period then matches firm B's price this period. If firm A's automaton is currently at the "punishment" state then it will charge the low price and the state will remain "punishment" if firm B set the low price, and will switch to the "collusive" state if firm B set the high price; again, firm A matches last period's price for firm B. The arrow on the far left pointing to the "collusive" state without an action attached to it indicates the initial state for the automaton; hence, it starts out by setting the high price. If both firms A and B use the tit-for-tat pricing automaton then both will set high prices forever.

Figure 2: Tit-for-tat

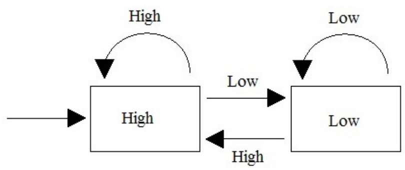

Figure 3: Punish $T = 1$ times

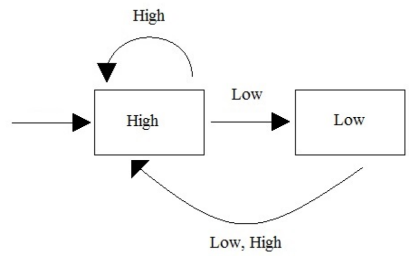

Let us examine some other pricing automata by returning to the "punish T times" strategy. When $T = 1$, the strategy is the two-state automaton shown in Figure 3. When T is infinity - so any deviation from the high price results in charging the low price forever - the strategy (which is commonly referred to as "grim trigger") is the two-state automaton in Figure 4. Examining Figure 4, note that, when in the collusive state, firm A's automaton chooses the high price. If firm B also sets the high price then the state does not change; it remains in the collusive state. If instead firm B chooses a low price then the state switches to the punishment state and firm A sets a low price. Once at the punishment state, the automaton stays there, regardless of the price of firm B.

Figure 4: Punish $T = \infty$ times (grim trigger)

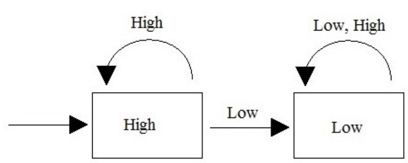

If AAs have access to these finite automata, the question is whether they will come to adopt automata that result in high prices. We offer two learning algorithms to show how this can occur. A more general discussion of learning algorithms is provided in Section 6.3 when we discuss issues related to implementing our approach to competition law. Right now, the objective is to illustrate that AAs can collude in a simple setting.

To keep matters simple, Hanaki, Sethi, Erev, and Peterhansl (2005) restrict each AA to using a pricing algorithm that is an automaton with no more than two states. Once eliminating redundant automata (that is, they generate the same output), there are 26 finite automata with one or two states. These automata include always setting a low price (which is a one-state automaton), always setting a high price (another one-state automaton), tit-for-tat, the grim trigger strategy, "punish $T = 1$ times" strategy, among other strategies. In any period, an AA has a particular automaton that it is using to set price. To allow for learning, an AA will periodically consider changing its automaton where this opportunity occurs randomly.[11] Once given the opportunity to change, an AA is specified to be more likely to choose an automaton when past performance of that automaton is higher relative to that of other automata.

AAs are specified to adapt their choice of a strategy using attraction-based learning. At any moment of time, an AA assigns each pricing automaton an "attraction" value. When given the opportunity to select an automaton, an AA attaches to each pricing algorithm a probability of being selected where the probability is higher when an automaton's attraction value is higher relative to the attraction values of the other pricing automata.[12] Key to this formulation is how performance affects the attraction values. Suppose an AA has been using automaton k for the last x periods and the opportunity to select a new pricing automaton (or keep the old one) arises. The

[11] There is a probability each period that an AA can change its pricing automaton. The simulations assume there is a 5% chance each period.

[12] More specifically, let $A_k(t)$ denote the attraction value that an AA assigns to automaton k as of period t. The probability that automaton k is selected equals

$$\frac{e^{\lambda A_k(t)}}{\sum_{h \in H} e^{\lambda A_h(t)}}$$

where H is the set of all 26 automata. $\lambda > 0$ controls how sensitive the probability is to an automaton's attraction value, with a higher value for λ indicating it is more sensitive.

attraction value assigned to automaton k is assumed to be a weighted average of the attraction value as of x periods ago (when it was adopted) and the average profit earned over the last x periods (during which time pricing automaton k has been in use). Thus, the better it has performed, the higher will be the new attraction value. For the automata that were not in use, their attraction values are unchanged.[13]

The simulation involves a population of AAs who are randomly matched into pairs to repeatedly interact in a Prisoners' Dilemma (see Figure 1).[14] In any period, there is a probability that an AA is rematched with another AA.[15] At any point, an AA may have an opportunity (which is randomly determined) to change its pricing automaton in the manner described above. It also has a chance to change its pricing algorithm whenever it is rematched with another AA. This process is run for a long time - which means many interactions between AAs - until the attraction values for the pricing automata settle down so there is very little change in them.

Four pricing automata were found to "survive" this selection process in the sense that each has much higher attraction values than the other 22 automata. These four automata are tit-for-tat, grim trigger, punish once (which is "punish T times" with $T = 1$), and "punish until the opponent retaliates" which is shown in Figure 5. The first three automata we have discussed. The last one, like the other ones, has the AA set a high price when in the collusive state, set a low price when in the punishment state, and switch from the collusive state to the punishment state when the other

[13] Suppose period t' was the last period that this AA selected an automaton and it chose automaton k. It now has a chance to select an automaton in period $t''(> t')$. If $A_k(t')$ was the attraction value as of period t' then the attraction value for automaton k in period t'' is

$$A_k(t'') = (1-\omega)A_k(t') + \omega\left(\frac{\pi^i_{t'} + \cdots + \pi^i_{t''-1}}{t'' - t' - 1}\right)$$

where $0 < \omega < 1$ and π^i_t is the profit earned in period t so $\frac{\pi^i_{t'} + \cdots + \pi^i_{t''-1}}{t''-t'-1}$ is the average profit earned by the automaton since it was selected. ω controls how much weight is given to the most recent performance.

[14] Referring back to Figure 1, the assumed profits are: $\pi^d = 9, \pi^c = 8, \pi^n = 2, \pi^s = 1$.

[15] In the simulations, there is a 2% chance each period for a rematching.

firm charges a low price. It differs in that it transits from the punishment state to the collusive state when the other firm sets a low price. In other words, the pricing automaton punishes with a low price until the other firm's automaton retaliates with a low price, at which point it returns to the collusive state and selects the high price. Note that all four of these automata have the capability to support supracompetitive prices. In a second round of simulations, Hanaki et al (2005) show that high prices emerge when two AAs are matched and learn over only these four automata.[16] In sum, collusion emerges through simple learning.

Figure 5: Punish until the opponent retaliates

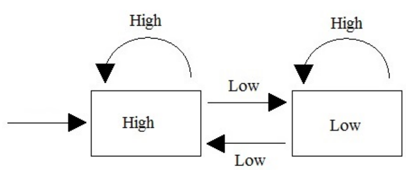

An alternative model of learning is Q-learning, which is closer to how AAs are likely to adapt in actual markets.[17] Q-learning has an AA assign a value (which, in our setting, proxies for the present value of the profit stream) to each state-action combination. In applying Q-learning to the game in Figure 1, Calvano, Calzolari, Denicolo, and Pastorello (2018) define a state to be the pair of prices charged in the previous period, in which case there are four possible states corresponding to the four price pairs that could have been charged. As with Hanaki et al (2005), an action is a price, either low or high. An AA will then assign a value to each combination of "current period's price for this firm" (action) and "previous period's prices for both firms" (state). Given the current state and the current collection of values, a firm

[16] The pricing algorithm that is eventually given the highest attraction value is "punish until the opponent retailiates".

[17] For research using Q-learning in the context of price or quantity competition, see Tesauro and Kephart (2002), Xie and Chen (2004), Waltman and Kaymak (2008), Dogan and Güner (2015), and Hulsen (2016), and Calvano, Calzolari, Denicolo, and Pastorello (2018).

chooses the price that yields the highest value, though with some small probability chooses the other price (as a form of experimentation). A firm's selected price is then the best price according to how it values different state-action combinations at that moment in time. Learning occurs by adjusting those values in response to performance in the following manner. Having chosen the value-maximizing price given the current state, contemporaneous profit is received (as given by the values in Figure 1). That profit is used to update the value attached to that action-state pair. Specifically, the updated value is a weighted average of the current profit received and the maximal value given the new state. The values for other state-action combinations are not changed. Come next period, the firm faces a new state, a price is once again chosen to yield the highest value, profit is received, the value attached to that action-state pair is adjusted, and the process continues.[18]

Calvano et al (2018) consider a pair of firms repeatedly interacting as prescribed in Figure 1,[19] and using Q-learning to improve their pricing algorithms.[20] Each

---

[18] More formally, let $Q^t(a,s)$ denote the value in period t associated with action a and state s. There are as many values as there are pairs of a and s. If s' is the period t state then the selected action a' is that which maximizes $Q^t(a,s')$. The choice of a', given current state s', results in some current reward π' and generates a new state s''. Q-learning updates the value attached to state (a',s') as follows. The value assigned in period $t+1$ to (a',s') is: $Q^{t+1}(a',s') = (1-\alpha)Q^t(a',s') + \alpha[\pi' + \delta Q^t(a^*(s''),s'')]$, where $\delta \in (0,1)$ is the discount factor, $\alpha \in (0,1)$ is the learning parameter, and $a^*(s'')$ is the action that maximizes $Q^t(a,s'')$. What this equation says is that the new value attached to (a',s') is a weighted average of the original value, $Q^t(a',s')$, and the "realized" value, $\pi' + \delta Q^t(a^*(s''),s'')$. However, it is only "realized" in the sense that it encompasses the actual current reward plus the future value based on current beliefs. Only the value attached to the actual action chosen and state observed, (a',s'), is modified. The other values are unchanged, and will be updated only when the associated action-state pair occur. α controls the rate of adjustment where a higher value of α means adjusting quicker to new information; that is, more weight is given to the "realized" value. For an introduction to Q-learning, see Sutton and Barto (2000).

[19] It is assumed $\pi^d = 1.5, \pi^c = 1, \pi^n = 0.5, \pi^s = 0.25$.

[20] In reference to footnote 18, which formally defines Q-learning, Calvano et al (2018) assume $\delta = 0.95$ and $\alpha = 0.15$. The probability that a firm chooses the value-maximizing price is 0.98 so there is a 2% chance of experimenting with the other price. Finally, the initial Q values are set at

simulation has firms 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 a firm's strategy chooses a low price regardless of what prices were set in the previous period (i.e., for all states). Collusion means firms adopted strategies resulting in high prices. In 9.2% of periods, they are using the grim trigger strategy and in 43% of the periods they have settled on the one-period punishment strategy. On average, it takes 165,000 periods for collusion to emerge. 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 it translates into 115 days, which is not necessarily long.

The takeaway from Hanaki et al (2005) and Calvano et al (2018) is that AAs can adapt their way to collusive prices.[21] Collusion emerged without communication or mutual understanding. Of course, this result does not address the question of whether AAs can collude in actual market settings and, if they can, how frequently it might occur. Ultimately, the ease and inevitability of collusion by price-setting AAs in real-world markets is an empirical question. With regards to that question I am currently agnostic, as well as skeptical that anyone can be confident as to what is the answer to that question. The only conclusion we can draw right now is that, at least in simple environments, price-setting AAs do not have to be sophisticated or strategic for supracompetitive prices to emerge. As long as one accepts the possibility that AAs could collude in actual markets, we can move forward to examine its legal ramifications.[22]

---

the present value of competitive profits, which is 10, for all state-action pairs. Hence, firms start in the competition mode.

[21] The first paper to produce collusion in the repeated Prisoners' Dilemma with AAs is Miller (1996) which used a genetic algorithm. Also see Hingston and Kendall (2004) for a related analysis.

[22] Under a certain set of stringent conditions, a recent theoretical result shows that collusion is not only possible but inevitable. It is proven in Salcedo (2016) that if an AA can learn another AA's strategy (say, through machine learning with enough variation in the environment) and if an AA can intertemporally optimize then, under certain conditions, collusion will eventually emerge

3 Evaluating the Legality of Collusion by Autonomous Agents

3.1 State of Jurisprudence

To lay the groundwork, let me offer a very concise and focused review of liability and evidentiary standards in the United States with regards to collusion.[23] Section 1 of the Sherman Act (1890) 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 (1911) clarified that only "unreasonable" restraints of trade are prohibited, while the reference to "contracts, combinations, and conspiracies" has been effectively replaced with the concept of "agreement." It is now understood that firms are in violation of section 1 when there is *an agreement among competitors to restrict competition.*

Liability comes down to what it means for firms to have an "agreement." Consistently, the courts have viewed an agreement as a mutual set of beliefs to restrict competition.

> Where the circumstances are such as to warrant a jury in finding that the conspirators had a unity of purpose or a common design and understanding, or a

for sure. Open questions for this result's relevance for actual markets are: 1) How easily can an AA "decode" another firm's pricing algorithm?; and 2) How easily can an AA intertemporally optimize in real time? A critique of Salcedo (2016) is provided in a forthcoming working paper by Kai-Uwe Kühn and Steven Tadelis which is skeptical that the conditions laid out for collusion to arise will be satisfied in practice. However, it is important to note that Salcedo (2016) provides sufficient, but not necessary, conditions for collusion to emerge.

[23] For more developed treatments, the reader is referred to Kovacic (1993), Werden (2004), *Proof of Conspiracy Under Federal Antitrust Laws* (2010), and Kaplow (2013).

> meeting of minds in an unlawful arrangement, the conclusion that a conspiracy is established is justified.[24]
>
> [T]here must be direct or circumstantial evidence that reasonably tends to prove that the [firms] had a conscious commitment to a common scheme designed to achieve an unlawful objective.[25]
>
> [W]hile a conspiracy involves an agreement to violate the law, it is not necessary that the persons charged meet each other and enter into an express or formal agreement, or that they stated in words or writing what the scheme was or how it was to be effected. It is sufficient to show that they tacitly came to a mutual understanding to allocate customers.[26]

Reference to "meeting of minds," "conscious commitment to a common scheme," and "mutual understanding" all focus on the same mental state: a common understanding among firms that they will restrict competition in some manner. This perspective has been echoed by the European Union General Court which has defined an agreement as or as requiring "joint intention"[27] or a "concurrence of wills"[28].

Turning to the evidentiary standards for concluding that firms have an agreement, an "explicit, verbally communicated assent to a common course of action"[29] certainly exceeds the bar in terms of the type of evidence that will prove liability. 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."[30] Judge Richard Posner proposed that an

[24] *American Tobacco Co. v. United States* 328 U.S. 781, 809-10 (1946)

[25] *Monsanto Co. v. Spray-Rite Serv.*, 465 U.S. 752, 768 (1984)

[26] *United States v. Suntar Roofing, Inc.*, 897 F.2d 469, 474 (10th Cir. 1990)

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

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

[29] Turner (1962), p. 683.

[30] Phillip E. Areeda, *Antitrust Law* 9-12 (1986); cited in Kovacic (1993), p. 18-9.

agreement could be tacitly consummated through firms' prices without the use of express language:[31]

> Section 1 of the Sherman Act forbids contracts, combinations, or conspiracies in restraint of trade. This statutory language is broad enough, as we noted in *JTC Petroleum Co. v. Piasa Motor Fuels, Inc.*, 190 F.3d 775, 780 (7th Cir. 1999), to encompass a purely tacit agreement to fix prices, that is, an agreement made without any actual communication among the parties to the agreement. 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.[32]

Though arguing that the law is "broad enough," Judge Posner also recognized that such evidence falls short of accepted evidentiary standards:

> [I]t is generally believed and the plaintiffs implicitly accept, that an express, manifested agreement, and thus an agreement involving actual, verbalized communication, must be proved in order for a price-fixing conspiracy to be actionable under the Sherman Act.[33]

Express and direct communication that conveys a plan to coordinate behavior is above the evidentiary bar, while the exclusive use of prices as the medium - as opposed to, say, natural language - is below the bar. Towards identifying where the bar lies, the court has been guided by the requirement that "there must be evidence that tends to exclude the possibility that the [firms] were acting independently."[34] While many avenues have been pursued by plaintiffs to argue that firms' conduct could

[31] However, he has recently retracted this position (Posner, 2015).

[32] *In Re High Fructose Corn Syrup Antitrust Litigation Appeal of A & W Bottling Inc et al*, United States Court of Appeals, 295 F3d 651, 652 (7th Cir., 2002).

[33] Ibid, 654

[34] *Monsanto Co. v. Spray-Rite Serv.*, 465 U.S. 752, 768 (1984)

not have been reached independently, successful recipes for convincing the court of that claim have almost always had a common ingredient: *evidence of an overt act of communication.* That communication typically has the feature that it facilitates coordinating on or monitoring compliance with a collusive outcome, while not being a standard feature of the competitive process. As noted in Hovenkamp (2016, p. 243): "[F]ew courts have found a conspiracy without some evidence of communication tending to show an agreement." When two 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. The absence of some overt act of communication has generally prevented the court from excluding the possibility that firms' conduct was independently reached.[35]

Notably, an "overt act of communication" could fall well short of a clearly articulated invitation to collude with a corresponding acceptance of that invitation; communication need not be so egregious. It is sufficient that there is some expression of intent resulting in reliance among firms that they will coordinate in order to reduce 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.

> 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 dis-

[35] As convincingly argued in Kaplow (2013, chp. 3), defining the set of illegal communications is very difficult because communication can be conducted in so many forms. It is then unclear what is exactly is meant by an act of communication that is evidence of an agreement.

> tinguishes interdependent and concerted action is evidence that tends to show that the defendants have communicated in the requisite ways.[36]

It is worth noting that there is one form of economic evidence that substitutes for direct documentation of overt communication, and that is "unnatural parallelism".

> "Unnatural" parallelism refers to parallel behavior that would probably not result from chance, coincidence, independent responses to common stimuli, or mere interdependence unaided by advanced understanding among the parties. ... A good example would be identical secret bids on a made-to-order item unlike anything previously sold and not a mere assembly of common items with a standard price. To be sure, even that parallelism might result from mere coincidence [but] the law does not insist on absolute certainty and rightfully disregards such low-order possibilities.[37]

However, unnatural parallelism should not be viewed as an exception to the requirement that firms have engaged in some overt communication but rather seen as another class of evidence in support of satisfying that requirement. The courts have viewed unnatural parallelism as circumstantial evidence that there must have been communication of the type that, if documented, would be sufficient to establish an unlawful agreement. Only by expressly communicating could the firms have come to change their prices at the same time or submit the same bids.

3.2 Applying Jurisprudence to Collusion by Autonomous Artificial Agents

Summing up the preceding discussion, jurisprudence regarding section 1 of the Sherman Act defines liability as a mutual understanding among competitors to restrict

[36] Page (2007), p. 447.

[37] Areeda and Hovenkamp (2010), pp. 181-2.

competition. With regards to evidentiary standards, there must be an overt act of communication to create or sustain that mutual understanding, and conduct cannot be consistent with firms having acted independently.

According to that jurisprudence, I claim that firms that collude through the use of AAs are not guilty of violating section 1 of the Sherman Act. 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.[38] If there is any communication (broadly

[38] Though involving human, not AAs, an example of an extraneous message is that used by airlines to coordinate their air fares through a computer reservation system. In the late 1980s, there was evidence that airlines tacitly colluded with an understanding that each airline could set high fares at its hub. In some instances in which another airline undercut that high fare, the offended airline would retaliate in the deviating airline's hub with low fares pre-fixed with the letters "FU". FU is an extraneous message in that it has the potential to communicate and is not part of the competitive process. See Viscusi, Harrington, and Vernon (2005), p. 122.

Another example is from the 16 spectrum auctions conducted by the U.S. Federal Communications Commission over 1994-98. Bidders used the last three digits of a bid for a particular license to refer to auctions for other licenses. As winning bids were over $2 million for a license, the last three digits impacted the bid by at most $999 and was effectively costless. As noted in a civil case brought by the U.S. Department of Justice (*U.S. v. Omnipoint Corp.*, In the U.S. District Court for the District of Columbia, Civil Action No.1:98CV02750, Competitive Impact Statement, November 10, 1998; pp. 6-7):

> As the auction proceeded, bidders carefully observed their rivals' actions and often adjusted their own market valuations and business strategies, sometimes based on their assessment of their rivals' objectives. Their rivals' bids, however, did not necessarily reveal their true objectives. ... To eliminate or reduce any ambiguity, Omnipoint sometimes placed bids during the DEF auction in which the final three digits intentionally corresponded to the number for a [basic trading area] (a "BTA end code"). Knowing that other bidders could see the bids and hence the BTA end codes, Omnipoint used the codes to better explain the real purpose of certain bids it made – to reach an agreement with a rival. In particular, Omnipoint used the BTA end codes to link the

defined) between AAs, it is occurring through prices or other legitimate market data which the court has found to fall short of satisfying evidentiary standards. As there is no overt act of communication, a requisite element of evidentiary standards is absent.[39]

Of course, evidentiary standards 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 are not liable. 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 restrict 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. A company is liable for its employees,[40] and third parties hired to manage pricing, such as consultants. 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

> bidding of licenses in two (or more) specific BTA markets, highlight the licenses Omnipoint wanted, and convey to the competing bidders offers to agree with Omnipoint not to bid against each other for the linked licenses.

[39] An interesting question is whether AAs could create their own language by which to communicate. Related to that point is the development of an internal language by Google's Neural Machine Translation system; Devin Coldewey, "Google's AI translation tool seems to have invented its own secret internal language," posted November 22, 2016 (https://techcrunch.com/2016/11/22/googles-ai-translation-tool-seems-to-have-invented-its-own-secret-internal-language/ accessed on January 9, 2018).

[40] "[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 Construction Co.*, 711 F.2d 570, 573 (4th Cir. 1983)

a "meeting of minds" or a "conscious commitment to a common scheme", could the AAs possess it? Addressing that 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.[41]

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."[42] 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.[43]

The Chinese Room Argument is not without its detractors.[44] However, even if one were to grant understanding to an AA, it is another step to reach a state of *mutual*

[41] Searle (1999)

[42] Searle (1980), p. 418.

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

[44] Some of the criticisms are presented and addressed in Searle (1980) and, for a more recent

understanding. That requires AA_1 to understand that AA_2 is setting a high price on the understanding that AA_1 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 effectively align the pricing strategies of all the market participants, even though they have reached

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 syntatical 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, five 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."

> no agreement.[45]

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 ...[46]

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 the situation that motivates this paper which is collusion achieved through means that, if executed by human agents, would be lawful.

In conclusion, it seems problematic 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.

4 Defining Collusion

In order to develop a legal approach to handling collusion by AAs, let us return to first principles and ask: What is collusion? At its most fundamental level, what

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

[46] Gosselin, Jones, and Martin (2017)

distinguishes collusive conduct from competitive conduct? While the focus will be on collusion that involves prices, 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.

While the objective of collusion is higher profits through supracompetitive prices, collusion is *not* the setting of supracompetitive prices. Supracompetitive prices could occur, for example, because of government regulation, which would have nothing to do with what we think of as collusion. More to the point, a firm can, on its own, set a supracompetitive price. It does not need to coordinate in some manner with rival firms; a firm has full control over its price (in the absence of any regulatory constraints). Of course, a firm *does not want* to charge a supracompetitive price because, as long as other firms are setting competitive prices, it is less profitable for a firm to charge a price above the competitive price.[47] However, it *could* be profitable for a firm to raise price above the competitive level if it expected rival firms also to do so. Thus, a firm that seeks to shift market conduct from competition to collusion does not do so by pricing at a supracompetitive level but rather by inducing *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 cause other firms to charge supracompetitive prices? Here, we turn to the economic theory of collusion which shows that the stability of firms charging supracompetitive prices is rooted in pricing strategies that embody a "reward-punishment" scheme which makes it profitable for each firm to price at supracompetitive levels. 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 competitive level if one or more firms priced below

[47] Recall that competitive prices are Nash equilibrium prices; that is, each firm's price maximizes its profit given the prices charged by other firms.

that supracompetitive level in some past period. If all firms adopt such a strategy then the resulting outcome is that all firms charge supracompetitive prices. Under certain conditions, it can be shown that it is in a firm's best interests (that is, it maximizes the present value of its stream of profits) to adopt such a strategy given other firms are expected to do so.[48]

The optimality of such a strategy - and that it is an equilibrium (i.e., a stable situation) for all firms to adopt such a strategy - is easily explained. A firm prices at a supracompetitive level, rather than undercutting that price and picking up more market share and profits today, because it takes account of a causal relationship between its *current* conduct and other firms' *future* conduct. If it prices at a high supracompetitive level today then it anticipates the other firms pricing at a high level tomorrow (and it is always more profitable for a firm when rival firms price higher), while not pricing high today is anticipated to bring forth low competitive prices from the other firms. The firm associates high future profits with charging a supracompetitive or collusive price today, and low future profits with a price below the collusive level today. While the described retaliatory response of rival firms is to price at competitive levels in response to a firm not charging the collusive price, 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 conduct from rival firms and low prices beget unfavorable conduct from other firms.*

If it is to be effective in yielding supracompetitive prices, a collusive strategy must embody a reward-punishment scheme. If a firm abides by the collusive outcome - which could involve high prices, exclusive territories, customer allocation, etc. - then

[48] For more developed non-technical treatments of the economic theory of collusion, the reader is referred to Motta (2004) and Viscusi, Harrington, and Vernon (2005). For technical treatments, see Vives (1999) and Harrington (2017a) which focuses on unlawful collusion.

it is *rewarded* in the future by rival firms continuing to abide by the collusive outcome (e.g., setting high prices, constructing exclusive territories, allocating customers); while if it departs from the collusive outcome (e.g., setting a low price, selling above its quota, stealing 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, stealing its customers). Collusion ties future rewards and punishments to current behavior and, by doing so, results in supracompetitive outcomes. It is the causal relationship that links a firm's current conduct with rival firms' future conduct that defines collusion, not the setting of supracompetitive prices, which are instead 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 strategy that, when adopted by all firms, results in supracompetitive prices. In using this strategy, a firm, when it chooses its price, takes into account a causal relationship between its price and other firms' future prices. In preparing for a later discussion, note that a causal relationship exists between two variables if the state of one variable is partly responsible for the state of the second variable. When firm A is said to cause firm B to price high, it means that there is conduct by firm A that makes it more likely that firm B will price high. Notably, a causal process does not require the mental state of intent or understanding. Hence, in principle, firm A can cause firm B to price higher even if firm A does not intend to do so and is not conscious of doing so. Finally, let me emphasize that the discussion has focused on what it means for firms to collude, which is distinct from what it means for firms to *illegally* collude. Like competition, collusion is a form of firm conduct that is defined independent of the legal environment.

5 Why Jurisprudence Takes the Form It Does

Collusion results in firms coordinating on supracompetitive prices. Firms achieve that end by adopting strategies that embody a reward-punishment scheme. In it simplest form, each firm rewards competitors for pricing high with a high price in the future, while each firm punishes competitors for pricing low with a low price in the future. One can view this arrangement as contractual, where the terms of the contract to be complied with are "high prices" and the penalties for acting contrary to the terms of the contract are low prices by the other firms. As there is no third party to ensure firms' behavior is compliant, it must be a self-enforcing contract; that is, each firm finds it in their best interest to abide by it, as long as all other firms intend to do so. In the jargon of game theory, it is a Nash equilibrium for each firm to use the collusive strategy.

It has been noted in Yao and DeSanti (1993), Werden (2004), and Kaplow (2013) that collusion, as defined by economic theory, is broadly consistent with the legal system's definition of liability:

> Analysis of one-shot games provides the clear definition of self-interest necessary to allow evidence of action against self-interest to play a useful role in inferring the existence of an agreement. If there is a unique Nash, non-cooperative equilibrium to a particular game, as there is in conventional one-shot game oligopoly models, it follows that there is a unique action each player will take if [it] does not coordinate its actions with its rivals. These equilibrium actions are consistent with self-interest, and any other actions are not. ... The existence of an agreement can be inferred from actions inconsistent with Nash, non-cooperative equilibrium in a one-shot game oligopoly model, even though they are consistent with Nash, non-cooperative equilibrium in an infinitely-repeated oligopoly game.[49]

[49] Werden (2004), pp. 770, 779.

As described in Section 3, firms are liable when they have an agreement, and jurisprudence has defined an agreement as mutual understanding to restrain competition. As noted in Section 4, firms are colluding when they use reward-punishment schemes which could result in supracompetitive prices. The unlawful "conscious commitment to a common scheme designed to achieve an unlawful objective" can translate into mutual understanding regarding the use of a collusive strategy.

Given this approximate equivalence between the economic phenomenon of collusion and the legal definition of collusion, one might have imagined (if uninformed of the state of jurisprudence) that the evidentiary process would focus on establishing the use of collusive strategies by firms. However, that is not the case. Economic evidence pertaining to the use of reward-punishment strategies and the setting of supracompetitive prices may be relevant but it is rarely the primary evidence, and is surely insufficient to prove guilt. As commented by Judge Posner:

> [I]t 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.[50]

Or consider a reward-punishment strategy to support a market allocation scheme of exclusive territories: A firm stays out of rival firms' markets because it will be rewarded by those firms not entering its market, and failure to comply will bring forth retaliatory entry. Such a strategy was claimed to have been used by incumbent local exchange carriers (ILECs) - also known as the Regional Bell Operating Companies - to prevent entry after the 1996 Telecommunications Act removed regulatory entry barriers. The U.S. Supreme Court saw this reward-punishment strategy as lawful:

> In the decade preceding the 1996 Act and well before that, monopoly was the norm in telecommunications, not the exception. The ILECs were

[50] *In Re: Text Messaging Antitrust Litigation*, U.S. Court of Appeals, Seventh Circuit. Decided April 9, 2015, p. 16.

> born in that world, doubtless liked the world the way it was, and surely knew the adage about him who lives by the sword. Hence, a natural explanation for the noncompetition alleged is that the former Government-sanctioned monopolists were sitting tight, expecting their neighbors to do the same thing.[51]

Evidentiary standards require evidence of an overt act of communication intended to constrain competition. However, communication is not collusion (as economic theory defines collusion) and communication is not an agreement (as the court defines agreement). Communication *facilitates* mutual understanding among firms to restrain competition, but communication is not itself mutual understanding and, therefore, is not an agreement. In fact, communication is neither necessary nor sufficient for firms to have common beliefs to restrict competition. Its lack of necessity is exemplified by conscious parallelism which is mutual understanding without communication.[52] To establish its lack of sufficiency, consider two firms that expressly exchanged assurances to raise prices but only one of the firms actually raised price. In spite of having the most direct and express form of communication, they clearly lacked a "meeting of minds" for the firm that increased price presumably expected the other firm to do so, which it did not.[53]

Though an express exchange of assurances to restrict competition may not result in mutual understanding between firms, it is appropriate that it is a *per se* offense.

[51] *Bell Atlantic Corp. v. Twombly*, 550 U.S. 544 (2007)

[52] Conscious parallelism 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." *Brooke Group Ltd. v. Brown & Williamson Tobacco Corp.*, 509 U.S. 209, 227 (1993).

[53] The point being made here is that communication does not imply mutual understanding and, therefore, communication does not imply agreement (as the courts have defined agreement). At the same time, I recognize that it is common (but not universal) to see an exchange of assurances as an agreement irrespective of whether or not it produces mutual understanding. I believe such a perspective is confused. The issue is not interpretation but logical consistency.

Such communication rarely has any procompetitive effects but it does have the anti-competitive effect of facilitating mutual understanding to restrain competition even if, in some instances, that objective is not obtained. That an express exchange of assurances is and should be a *per se* offense does not dilute the fact that there is a distinction between collusion and communication that facilitates collusion:

> Most who advance a prohibition on communications as facilitating practices, however, do not regard successful interdependent pricing behavior to be illegal. If coordinated oligopolistic activity is by itself lawful, it seems all the more important to attack practices that facilitate it ... Note, however, that there is a certain irony involved: Aiding and abetting is 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.[54]

Even if there is a *per se* prohibition on certain forms of communication, it does not follow that evidence of communication is a requisite element to establish guilt. That is, making communication a sufficient condition for proving liability does not imply it is a necessary condition. Nevertheless, as a general rule, it is appropriate to require some evidence of communication, rather than exclusively rely on economic evidence, and the reason is error costs. Given the usual data and empirical methods that are 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 challenge is establishing that a firm prices at a high level only because continuing to do so will induce rival firms to price high in the future, while not doing so would bring forth retaliation by those rival firms pricing low in the future.[55]

Let me offer a few scenarios so that we may appreciate the difficulty of identifying firms' latent strategies and thereby assessing the rationale for observed prices.

[54] Kaplow (2013), p. 57.

[55] For a more developed treatment of the ensuing discussion, see Kühn (2001).

Suppose either firm A or B 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 the economic evidence. A competitive hypothesis is: Firm A (or B) raised its price in response to a positive cost shock and, because firms A and B are subject to a common cost shock, firm B (or A) 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. By contrast, 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. The extent to which these two hypotheses can be empirically distinguished depends on the ability to control for the cost and demand drivers of firms' prices. While empirical analysis might be able to make substantive progress towards determining which of these two hypotheses is more credible, experience suggests that, at the end of the analysis, substantive uncertainty as to which hypothesis is true will remain (or whether neither hypothesis is true and the facts are more consistent with some third hypothesis). More generally, parallel price movements can be the product of competition (under certain cost and demand conditions) or collusion, and it is difficult to distinguish between them with the level of confidence that the judicial system requires.

In light of the residual uncertainty associated with economic evidence, the exclusive reliance on it would be socially harmful in two ways.[56] First, it would distort competition. By virtue of exclusively using economic evidence, there is likely to be a significant chance that competition is misconstrued for collusion. Anticipating that

[56] For a more developed treatment, the reader is referred to Kaplow (2013), chapter 9.

possibility, competitive firms will modify their behavior to make it less likely that they are wrongly convicted. For example, suppose parallel price movements were taken as sufficient evidence of collusion. In response to a common cost decrease, and in absence of their conduct being misconstrued for collusion, competing firms A and B would be inclined to both lower price at the same time. However, if such parallel price movements could cause the court to conclude they have an unlawful agreement, the firms might instead stagger their price decreases to avoid wrongful conviction. In response to firm A lowering its price, firm B would delay its price decrease in order to not appear as if firms are coordinating their behavior. This altered behavior is a distortion of the competitive process.

The second source of social harm from an exclusive reliance on economic evidence is that it would weaken the deterrence of collusion. If firms choose not to collude, it is partly because they perceive the likelihood of being convicted and paying penalties is higher by colluding than by competing. However, if the process of determining guilt is less accurate - due to the higher rate of false positives from relying only on economic evidence - then the increase in the probability of paying penalties from colluding is smaller. For example, consider the extreme case in which the judicial process was entirely random and thus independent of whether or not firms were colluding. Firms might as well collude, for doing so raises profits without making it any more likely they will be convicted. More generally, the less sensitive is the court's judgment to the true underlying state of collusion or competition, the less firms are deterred from colluding. Hence, if relying exclusively on economic evidence leads to judicial decisions that are less accurate, it will weaken deterrence.

To sum up to this point, collusion is the use of reward-punishment strategies to sustain supracompetitive prices. In light of existing empirical methods and the usual data that is available to analyze, it is a challenging exercise to identify the latent strategies underlying firms' conduct. As a result, allowing economic evidence to be sufficient to convict suspected firms would generally create large error costs. The prospect of wrongly convicting competing firms would distort competition, and the heightened chance of wrongly convicting firms that are competing would weaken the deterrence of collusion. That the court has found economic evidence to be insufficient to deliver the level of confidence it needs to produce a conviction seems appropriate, at least as a general rule.

As described in Section 3, evidentiary standards require some evidence of overt communication between firms for the purpose of coordinating their behavior. Even though communication is ancillary to the use of collusive strategies, evidence of communication reduces error costs because communication generally lacks a competitive rationale and, even when it is part of the competitive process, firms can probably avoid engaging in it without creating much harm. For these reasons, evidence of communication becomes an informative signal of the presence of collusion. Let me elaborate on these points.

It is generally not a common feature of the competitive process for a firm to communicate with its competitors about what price it is charging. Indeed, there can be a strong incentive not to share such information. Consider a market in which firms offer similar products, and buyers actively search and may even negotiate discounts. If firm A tells firm B its price then firm B is in a better position to undercut firm A's price and win some business. In such a setting, a competitive firm wants its price to be known to customers but not to rival firms.

The competitive incentive to keep prices private information made the exchange of price information highly supportive of the hypothesis of collusion in a case involving suppliers of corrugated containers.[57] It was documented that, upon the request of a competitor, a firm would convey its most recently quoted or charged price. There was a reciprocity to this information sharing from which the Supreme Court inferred there was an agreement to exchange price information. It also found that knowledge of a competitor's price usually meant matching it. The Supreme Court concluded

[57] *United States v. Container Corp. of Am.*, 393 U.S. 333 (1969)

that the information exchange had the effect of limiting price competition. While the information sharing agreement was not found to be *per se* illegal, the firms were found guilty on the basis of that agreement along with evidence of its impact on prices. That firms want to exchange price information when they collude but do not want to exchange price information when they compete makes evidence of communication of prices supportive of the hypothesis that there is collusion.

There are, however, situations in which a competitive firm might want rival firms to know its price. For example, suppose firm A's cost is expected to rise which will cause it to increase price. Once its price goes up and is observed by rival firms, their competitive response will be to increase their prices (though generally not by the same amount) because their demand is higher as a result of a competitor's price having increased. However, until competitors raise their prices, firm A will be pricing at a significant premium and that will lower its demand and profits. For that reason, firm A would like to inform rival firms that it intends to raise its price, as then they will raise their prices sooner. Firm A can achieve that end by announcing to rival firms in advance of a price increase. However, by an analogous logic, firm A would *not* want to announce its plan to lower price, for that would speed up the time until rival firms lower their prices, which is unattractive to firm A. Thus, under competition, a firm would want to communicate to rival firms its intention to raise price but not its intention to lower price.

Relevant to this point is a case involving realtors.[58] On September 5, 1974, John Foley, the president of Jack Foley Realty, Inc., hosted a dinner party where the guests were nine of the leading realtors in Montgomery County. At this party, Mr. Foley announced that it was his intention to raise the commission rate from six to seven percent. He subsequently implemented this intention and, in the following months, many of those in attendance did likewise. There was no evidence of any other forms

[58] Details are from *United States v. Foley*, 598 F.2d 1323 (4th Cir. 1979), cert. denied, 444 U.S. 1043 (1980).

of communication (though one does not know whether there were knowing winks and nods in response to this announcement). They were ultimately convicted of violating section 1. While, as argued above, there is a competitive rationale for informing rival firms of a price increase, there is not a competitive rationale for doing so at an industry gathering rather than through a public announcement.

In a case involving oil companies in the sale of gasoline to distributors, the communication was in the form of public announcements.[59] The gasoline was officially sold in this wholesale market at the "dealer tankwagon price" though it was common to offer discounts. The oil companies were accused of colluding with respect to those discounts. The mechanism through which the suppliers coordinated were press releases that announced price changes in advance of their effective date. In their testimony, the oil company executives emphasized the cost of taking the lead on reducing discounts (referred to as a "restoration"):

> [T]he costs associated with leading a restoration were significant, and they increased with each day that a move was not detected or was not followed. An internal ARCO memorandum noted that its unsuccessful restoration attempts had been "costly to us, not only in terms of volume, but more importantly in expense dollars required to recoup and serious adverse reaction with our customers and in our dealer organization." Another ARCO document noted that "[p]ast experience has taught us how disastrous it can be for us to be the leader in an upward move." Even with public announcement, the costs of leading a restoration were substantial; Standard Oil's internal analysis of the effect of three restorations it led indicated an average drop in sales volume of over 19% within one week. The same Standard Oil document blamed, in part, the company's practice of leading restorations for causing a significant decline in Standard's overall market share. An ARCO document echoed this

[59] *In re Coordinated Pretrial Proceedings in Petroleum Prods. Antitrust Litigation*, 906 F.2d 432 (9th Cir. 1990).

> view, noting that Standard's and Shell's leading of restorations had led to their "disastrous" loss in market share.[60]

It is important to emphasize that this cost of leading a price increase is higher under collusion than under competition. Under competition, a rise in cost will make it profitable for a firm to raise its price even if it might put it at a price disadvantage compared to competitors. However, when a price rise means taking price to a supra-competitive level then it is unprofitable *unless rival firms match it* (or come close to matching it). In that situation, it is critical for other firms to be promptly informed of the price hike so that the initiating firm's price premium is short in duration. We then see that the incentive to inform rival firms of a price increase is stronger under collusion than under competition (where, say, it is due to a cost increase). This means that prohibiting such communication is likely to cause competing firms not to engage in it (and without imposing much harm on them), while colluding firms are likely to persist with it. Hence, communication through advance price announcements can be an informative signal of collusion.

Consistent with the use of price announcements for coordination purposes,

> several officers of the appellee oil companies were questioned concerning the business reasons for publicly announcing changes in tankwagon prices and in the levels of dealer assistance [and] their virtually uniform response was that it was done for the purpose of quickly informing competitors of the price change, in the express hope that these competitors would follow the move and restore their prices.[61]

The Ninth Circuit Court went on to conclude:

> The evidence presented by the appellants thus indicates that the publication of wholesale price increases was intended to make, and had the effect of

[60] Ibid, ¶44.
[61] Ibid, ¶42.

> making, restorations more effective by ensuring that competitors could quickly learn of, and respond to, any withdrawal of dealer aid. The appellees' actions in announcing such information made the market more receptive to price coordination than it otherwise would have been. ... we believe that the evidence concerning the purpose and effect of price announcements, when considered together with the evidence concerning the parallel pattern of price restorations, is sufficient to support a reasonable and permissible inference of an agreement, whether express or tacit, to raise or stabilize prices.[62]

It was then found that

> "mere price parallelism" is insufficient to raise an inference of conspiracy, but evidence of press releases announcing price increases in advance and defendants' employees' testimony that the press releases were made in hopes that other competitors would follow suit can be sufficient, along with the price parallelism, to support an inference of conspiracy.[63]

When communication about prices and price intentions are consistent with collusion but are inconsistent with competition, such communication is solid evidence that firms are colluding. For situations in which such communication is also consistent with competition, competing firms can avoid the risk of wrongful conviction by not engaging in that form of communication. Foley could have not conveyed the message privately and instead publicly announced the price increase, if the intent was not to coordinate on higher prices. Still, if some form of communication is consistent with competition and the prospect of it being mistaken for supporting an illegal agreement leads firms to not engage in that communication (or to modify it in some manner) then the competitive process is being distorted. However, it is not clear there is much harm from doing so. As explained above, competitive firms may have an incentive

[62] Ibid, ¶45.

[63] *Proof of Conspiracy Under Federal Antitrust Laws* (2010), p. 68, f. 92.

to inform rival firms of a price increase but not of a price decrease. Thus, if firms avoid such announcements it will delay the time until prices rise in response to a cost increase but not affect the time until prices decrease in response to a cost decrease. Consumers benefit from that "distortion" to competition and the net social harm is unclear.

In concluding, it is generally not a viable evidentiary approach to determine whether firms use collusive (reward-punishment) strategies using only the analysis of market data. A more viable approach is to require some evidence of communication that facilitates or sustains the use of collusive strategies. Requiring overt communication reduces the extent of false positives and lowers error costs because: 1) communication is not typically part of the competitive process and thus can be an informative signal that firms are colluding; and 2) even when communication is part of the competitive process, competitive firms can avoid the communication (and thereby avoid wrongful conviction) and probably do so with minimal social harm.

6 A Legal Approach to Collusion by Autonomous Agents

6.1 Overview

Let me summarize the argument that, as a general (though not universal) rule, evidentiary standards should require evidence of overt acts of communication. Collusion is when a firm causes rival firms to set supracompetitive prices, which is achieved by using a strategy that encompasses a reward-punishment scheme. It is the mutual adoption of collusive strategies that result in supracompetitive prices. While prices are observable, the reward-punishment scheme supporting them is not. The reward-punishment scheme is latent - inside the managers' heads - and can only be indirectly observed when those schemes interact with the environment to produce prices and quantities. Due to the difficulty in identifying firms' latent strategies, evidentiary standards require more than economic evidence; they require an overt act of communication that firms are seeking to coordinate their behavior. In a sense, it is communication that allows an outside observer to get inside the collective heads of managers and assess why they are pricing the way they are and, in particular, whether it is driven by the mutual adoption of strategies to restrain competition. To be clear, I am not describing the reasoning the courts used in adopting these evidentiary standards but rather why evidentiary standards had to be defined this way.

The doctrine that has just been described is based on collusion 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 an algorithm.

When the price-setting agent is a piece of software, 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 and, in that sense, it *is* possible to get "inside the head" of the price-setting (autonomous artificial) 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 causes it to result in supracompetitive prices, and why it should be prohibited.

The challenge is to take that simple observation - the strategy of an AA can be directly observed - and develop a doctrine of liability and evidentiary standards so that collusion by AAs can be effectively prosecuted. Towards developing that doctrine, 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 2.2. Though it may not be a realistic description of how collusive pricing by AAs would emerge in actual markets, 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. (In the language of finite automata, the automaton has just one state.) 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. Those supracompetitive prices would be avoided with a prohibition of pricing algorithms that condition play on a rival firm's past prices. More formally, the state transition equation would be prohibited from depending on a rival firm's prices. Though this example, by virtue of its simplicity, ignores many challenging issues associated with the implementation of this legal approach, it serves to illustrate the approach of a *per se* prohibition on certain pricing algorithms.[64]

Given a set of prohibited pricing algorithms, evidentiary standards are relatively straightforward to define.

[64] 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.

Evidentiary Standards: Liability would be determined by: 1) 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, some learning algorithms are not amenable to extracting the pricing rule from a reading of the code. In that case, tests will have to be formulated to determine whether a price algorithm embodies a prohibited property. These challenges are examined in the next two subsections.

6.2 Liability

Construction of the set of prohibited pricing algorithms should be guided by the objectives of convicting cartels and deterring collusive conduct. However, as there are clearly efficiency gains from the use of automated pricing algorithms, competition law should not interfere with algorithms that achieve legitimate competitive purposes. The more that collusive-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.[65]

[65] Implicit 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

Let pa denote "pricing algorithm" and PPA denote the collection of prohibited pricing algorithms. Given a specification of PPA, $Pr(pa \text{ is in } PPA\,|pa \text{ is collusive})$ is the probability that a pricing algorithm is determined to be in the prohibited set when the pricing algorithm is collusive. $Pr(pa \text{ is in } PPA\,|pa \text{ is competitive})$ is the probability that a pricing algorithm is determined to be in the prohibited set when the pricing algorithm is competitive. Ideally, $Pr(pa \text{ is in } PPA\,|pa \text{ is collusive}) = 1$ and $Pr(pa \text{ is in } PPA\,|pa \text{ is competitive}) = 0$ 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:[66]

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

$LR(PPA)$ 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 $Pr(pa \text{ is in } PPA\,|pa \text{ is collusive})$ is higher. When $Pr(pa \text{ is in } PPA\,|pa \text{ is collusive}) = 1$ 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 $Pr(pa \text{ is in } PPA\,|pa \text{ is competitive})$ is lower. When $Pr(pa \text{ is in } PPA\,|pa \text{ is competitive}) = 0$ 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

prohibited practices those pricing algorithms for which there is a competitive rationale.

[66]For a discussion of the usefulness of the likelihood ratio in the context of evidentiary standards, see Kaplow (2014).

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.[67]

This approach will not prove useful 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). However, as I argue now, an initial appraisal suggests that the properties that promote efficiency are quite distinct from those that promote collusion. To make this argument, it is necessary to provide a brief description of the efficiency benefits from using a learning algorithm. This discussion is relevant whether the learning algorithm is entirely autonomous (as has been presumed in this paper) or involves some human intervention.

Putting aside the prospect of collusion, the appeal of a learning algorithm is that data is used to make better decisions on price, where "better" means higher profit (though the ensuing discussion applies as well to other performance metrics such as revenue). One can think of a learning algorithm as using past data to select a pricing algorithm, where a pricing algorithm assigns a price for any state of current market conditions. Past data will typically comprise past prices (of this firm and

[67] 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.

possibly rival firms), past sales (most likely only for this firm, as the sales of rival firms are likely to be confidential), other variables that affect demand (e.g., day of the week and season), and other variables relevant to profit maximization including production cost, inventory holding cost, and the amount in inventory. Based on the past performance of prices under various market conditions, the learning algorithm prescribes a "best" price for a given collection of current market conditions such as day of week, season, and inventory. That assignment of price to current market conditions describes the current pricing algorithm. Given the particular set of market conditions at any moment, the pricing algorithm selects a price. Subsequently, sales and profits are realized which, along with current market conditions and possibly rival firms' prices, forms a new observation in the data set which the learning algorithm utilizes to improve the pricing algorithm. With that modified pricing algorithm, the firm is set to select a new price given the market conditions it now faces. This process of choosing a price, collecting a new data point, and updating the pricing algorithm continues *ad infinitum*.

A well-designed learning algorithm will adopt better and better pricing algorithms over time. It should be understood that this does not necessarily mean always setting the price that yields the highest profit (based on the data that is available) because a learning algorithm is, well, trying to learn! It is not just interested in current performance but also future performance, and information learned today can be used to improve the quality of future decisions and thereby improve future performance. Achieving that objective involves balancing exploration and exploitation (or "learning and earning"). The price that is best from the perspective of yielding the highest contemporaneous profit may not be chosen because other prices may generate more information about the firm's environment - in particular, the demand for its products - and that additional information sets the firm up to make better *future* price decisions.

The potential efficiency benefits of a learning algorithm and the automated pricing algorithm that it selects come from two possible sources.[68] 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. Second, an automated 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. With regards to the first source, some learning algorithms 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 estimated demand. 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 which serves to create price experiments that allow for more effective demand estimation.[69] and den Boer and Zward (2014) have a demand estimation module based on maximum likelihood. The optimization routine has price selected from outside of an interval around the revenue-maximizing price to, again, promote learning. This interval is programmed to shrink over time so that these price experiments become smaller and smaller in magnitude. These two learning algorithms, as well as many other ones, automate demand estimation and build in experimentation in order to promote learning.[70]

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

[69] In econometric terms, there is an endogeneity issue that biases estimation, which can be mitigated with these price perturbations.

[70] For a survey of some of these learning algorithsm, see den Boer (2015).

The second source of efficiencies from a learning algorithm is due to Big Data.[71] 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 at time scales of hours, minutes, or even seconds. 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). 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). Past data was used to predict the relationship between price and profit (or revenue) and how it depends on a user's location and browser history (which is an example of the first source of efficiencies of a learning algorithm), and then, having identified a current user's location and browser history, price was tailored to that user (which is an example of the second source of efficiencies).

In sum, the potential efficiencies of a learning algorithm such as 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 condi-

[71] 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).

tions; 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' current prices based on current market conditions but again it does not seek to determine how rival firms price. As noted above, 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 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 (and could either raise or lower social welfare) but would make collusion more difficult.

While a comprehensive and rigorous examination of these issues is required before more definitive 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 substantively interfering with competition.

6.3 Evidentiary Standards

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. Towards addressing that issue, I will first provide a high level view of how the relevant learning algorithms operate, and then discuss the various approaches for assessing the pricing algorithm that is being used by an AA.

There are two general classes of learning algorithms that could be used for the purpose of price setting. The first class is *estimation-optimization algorithms* which embody distinct estimation and optimization modules.[72] The estimation module estimates the firm's environment and delivers predictions as to how the firm's action determines its performance. More specifically for our setting, estimation involves the use of past data to estimate a firm's demand function, and thereby have an estimate of a firm's profit (or revenue) function. The goal of the estimation module is to develop the best predictor of revenue or profit for any set of current prices and market conditions. A variety of estimation methods have been used including regression (e.g., ordinary least squares), maximum likelihood, and an artificial neural network.[73] These methods are designed to deliver a parameterized model that fits past observations with the goal of making accurate predictions about future sales conditional on price and market conditions.

Once the estimation module has been run for a given collection of data, the optimization module operates to select a price. The objective of the optimization depends on how it is programmed. For example, it could choose a price to maximize

[72]For a general description of this type of algorithm, see Shakya, Chin, and Owusu (2010).

[73]Shakya, Kern, Owusu, and Chin (2012) have an estimation module that is either linear regression (with a linear or exponential specification), maximum likelihood (with a multinomial specification), or an ANN with backpropagation. The optimization module uses an evoluiontary algorithm, such as the generic algorithm. The study compares the performance of the four different estimation modules. When the true demand model could be either linear, exponential, or multinomial, the best average performance of the learning algorithm is delivered by an ANN.

revenue using the estimated revenue function based on estimated demand from the estimation module. If cost information is added, price could be chosen to maximize expected profit using the estimated revenue function along with cost. These revenue-maximizing or profit-maximizing prices might be modified in order to engage in some exploration for the purpose of conducting price experiments that would provide data useful for delivering better future estimates. Thus, the optimization module may take into account not just current performance, such as current revenue, but also future performance, such as future revenue. Two examples of these estimation-optimization learning algorithms were discussed above.[74]

The second class of learning algorithms is *reinforcement learning*.[75] An estimation-optimization algorithm estimates the environment faced by a firm and then determines what conduct performs best for that estimated environment. It can deliver a forecast on performance (e.g., profit or revenue) for any action (e.g., price) or strategy (e.g., pricing algorithm). An estimation-optimization algorithm learns over both the environment and the best action for an environment. In comparison, reinforcement learning fuses these two learning processes by learning directly over actions (or strategies); it figures out what action (or strategy) is best based on how various actions (or strategies) have performed in the past. It does not explicitly estimate the firm's environment (e.g., it does not estimate the firm's demand function) and thus is seen as "model free" because it is not based on a particular model of the firm's environment. Referring back to Section 2.2, reinforcement learning was used in Hanaki et al (2005) and Calvano et al (2018).

There are two general approaches to testing an AA in order to learn its properties, and they differ depending on data is available.[76] *Static testing* involves examination of the program's code without running the program, and *dynamic testing* is when the

[74] den Boer and Zward (2014) and Nambiar, Simchi-Levi, and Wang (2016)

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

[76] An important reference to the ensuing discussion is Desai and Kroll (2017).

program is run for selected inputs and the output is observed.

Let us begin by considering how static testing might be able to determine whether the pricing algorithm used by an AA is in the prohibited set. Such an exercise is straightforward when the AA uses an estimation-optimization algorithm because the optimization module describes the pricing algorithm; that is, how the AA assigns a price to a particular set of market conditions (possibly including rival firms' prices). The types of estimation-optimization algorithms that have appeared in the operations research literature are not ones that are conducive to collusion because rival firms' prices do not enter the optimization module. Rival firms' prices may be part of the data used in the estimation module to estimate a firm's demand but they do not enter the optimization module.[77] Collusion could emerge if the estimation module includes forecasting rival firms' prices, and the optimization module takes into account how rival firms' future prices are impacted by a firm's price. It could be reasonably straightforward to determine whether the resulting pricing algorithm has properties consistent with collusive pricing rules. However, the ease of that task will not be known until it is implemented, and Section 6.3 offers a research program to shed light on the matter.

Static testing is more problematic when an AA uses reinforcement learning. It would be feasible to learn the pricing algorithm that is used from inspection of the source code when reinforcement learning is occurring directly over pricing algorithms. By "directly," I mean that the AA is adopting a particular pricing algorithm (as opposed to a particular price) and assessing how it performs relative to other pricing algorithms. An example is Hanaki et al (2005) which we discussed in Section 2.2. At any moment in time, one could determine which pricing algorithm is given the

[77]Interestingly, Cooper, Homen-de-Mello, and Kleywegt (2015) show how a learning algorithm can generate supracompetitive prices when estimation does not condition on rival firms' prices. This property arises because of misspecification in the estimation module, and can be viewed (from the firms' perspective) as a fortuitous benefit of excessive simplicity. Though prices are supracompetitive, there is not collusion because there is no reward-punishment scheme supporting those higher prices.

most weight and thus whether it is a rule conducive to collusion.[78] However, that will probably not prove useful because, in practice, AAs are unlikely to engage in reinforcement learning directly over pricing algorithms. Given that there will typically be a very large number of pricing algorithms (i.e., many rules for describing what price to set for each possible state), learning directly over pricing algorithms would be slow and risky as there could be many periods of poor performance.

A more relevant class of reinforcement learning methods has learning occurring with regards to the performance of a particular price for a particular set of market conditions. Examples of such methods are artificial neural networks and Q-learning. Recall from Section 2.2 that Q-learning assigns a value to each price-market conditions pair. Given the current market conditions, the price with the highest value is chosen. After choosing that price, profit is received and its performance (as measured by profit) is used to update the value attached to that price-market conditions pair. A pricing algorithm assigns a price to each market condition and Q-learning has learning occur one price-market conditions pair at a time. In principle, the pricing algorithm can be constructed by inspecting the collection of values,[79] for those values tell us what the AA views as the best price depending on the market conditions.[80] With that

[78] Though reinforcement learning is selecting a pricing algorithm with some probability, the randomness typically declines over time so that, once learning has settled down, there will be one pricing algorithm that is used with very high probability.

[79] 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 market conditions to values.

[80] For example, returning to our discussion of Calvano et al (2018), an action is a price from the set $A \equiv \{low, high\}$, and a state is the previous period's prices, which comes from the set $S \equiv \{(low, low), (low, high), (high, low), (high, high)\}$. $\{Q^t(a, s)\}_{a \in A, s \in S}$ is the collection of 8 values as of period t for the 8 action-state pairs. For example, suppose the following conditions held: 1) $Q^t(low, (low, low)) > Q^t(high, (low, low))$; 2) $Q^t(low, (low, low)) > Q^t(high, (low, high))$; 3) $Q^t(low, (high, low)) > Q^t(high, (low, low))$; and 4) $Q^t(high, (high, high)) > Q^t(low, (high, high))$. In other words, the value is highest from choosing a low price except when the state is $(high, high)$. Hence, a firm sets a high price if both firms set high prices in the previous period and otherwise sets a low price. That is the grim trigger strategy

pricing algorithm, one can assess its properties and whether it is in the prohibited set of pricing algorithms. How feasible this approach is will depend on how numerous is the collection of values, and that is determined by how expansive is the set of prices and market conditions.

However, there are other methods of reinforcement learning for which inspection of the code will not prove informative.[81] The mapping between the code and the basis for a choice is too complex to sort out. Such methods are broadly referred to as "deep learning".[82]

> 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.[83]

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 market conditions into the AA and documenting how the generated prices respond to those market conditions. The objective would be to identify some properties of the latent pricing algorithm and, in particular, whether those properties are prohibited ones. There are two challenges with this approach, as noted in Desai and Kroll (2017). First, the possible sets of inputs (in our context, market conditions) could be very large in which case one could only test a small subset of inputs. To what extent that is an obstacle will depend on the property of the pricing algorithm one is looking for. If the property is how an AA responds

(which is shown in Figure 5 for when there are only two states). In this way, the pricing algorithm can be recovered by examining the values.

[81] Such concerns can be found in Ezrachi and Stucke (2017) and OECD (2017).

[82] 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.

[83] A. Gal (2017), p. 5.

to other firms' prices then it may be sufficient to consider inputs that differ only in other firms' prices for a representative sample of 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 last year, and it is price data from the last 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, that algorithm could be relatively stable with respect to the properties of interest.

In concluding this discussion, I believe it is possible, in principle, to identify properties of the pricing algorithm used by an AA, which would then allow determination of its legality. Whether it will prove practical is an open question that can only be addressed by systematic investigation. Towards that end, a research program is described in the next sub-section to explore its feasibility.

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 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.

Note that, if we were to replace AAs with human agents, the collusion that this approach would make illegal corresponds to conscious parallelism, for it is collusion without communication (or at least without evidence of an overt act of communication). As noted by Judge Stephen Breyer, conscious parallelism is considered legal in spite of the harm it creates:

> 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 ... that 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?[84]

Conscious parallelism is argued to be legal because of the lack of a remedy, for it would be difficult to describe to managers what forms of interdependent pricing are alright and what forms they should avoid. As a result, illegality is limited to when there is some overt act of communication, and the remedy is prohibiting such acts of communication (and other facilitating practices).[85] In contrast, the remedy for AAs can take the form of prohibiting certain types of interdependent pricing because AAs

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

[85] To be clear, I am describing how the law has been interpreted, not what is a proper interpretation. My view is if firms admit to conscious parallelism then they have violated section 1 of the Sherman Act as they have admitted to a mutual understanding to restrain competition. That it has happened without communication makes it no less harmful, nor is there a concern about a false positive as firms have admitted to the pricing behavior that the law is intended to prohibit.

can be programmed, where human agents cannot. In addition, an AA's strategy - and not just its prices - can be monitored for compliance with the prohibition of certain pricing algorithms. Hence, when prices are set by AAs, there could be a remedy for avoiding the supracompetitive prices associated with collusion, irrespective of acts or evidence of communication.

6.4 A Research Program for Defining Liability

The challenge of defining liability and evidentiary standards to prohibit collusion by AAs has parallels to policy issues related to fairness and machine learning.[86] 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., properties of pricing algorithms to be prohibited) - and, secondly, ensuring accountability, in the sense that the machine learning algorithm satisfies that definition of fairness - which is parallel to defining evidentiary standards (i.e., how to determine whether a pricing algorithm has certain properties). Accountability can be challenging. If fairness mandates that decisions not be based on some trait, 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,

[86] 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 Fish, Kun, and Lelkes (2016), Goodman (2016), Hardt, Price, and Srebro (2016), Johnson, Foster, and Stine (2016), Joseph, Kearns, Morgenstern, and Roth (2016), and Kleinberg, Mullainathan, and Raghavan (2016).

and education - then machine learning may figure out how to indirectly condition on race.[87]

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 lowering prices 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.[88] Keep track of when competitive prices emerge and when supracompetitive prices emerge. Perform this exercise with different learning algorithms and for a variety of market conditions. (This first

[87]These fairness issues could also pertain to pricing algorithms that personalize price. Thus, 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.

[88]Ezrachi and Stucke (2016) refer to it as a "collusion incubator."

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.[89]

While it is difficult to predict what this research program will produce in terms of defining a set of prohibited pricing algorithms, one candidate for inclusion are pricing algorithms that exhibit "price matching". "Price matching" refers to a firm setting its price (or price change) equal to the price (or price change) of a rival firm. To be clear, a prohibition on price matching would not mean that it is illegal for firms to have identical prices. Rather, firms would be prohibited from using pricing algorithms that choose a price (or a price change) to match a rival firm's price (or price change). The prohibition is not on prices but rather on pricing algorithms.

Pricing algorithms with price matching may end up in the set of prohibited pricing algorithms because price matching is a feature of collusion but is not a feature of competition. Addressing the first claim, there is a long history of documented episodes

[89] More formally, the distribution of prices without the prohibition first-order stochastically dominates the distribution of prices with the prohibition.

of collusion in which firms used price matching. Markham (1951) and Scherer (1980, chapter 6) offer some historical examples, and more recent cases include retail gasoline (Bryne and de Roos, 2017), wholesale gasoline (Andreoli-Versbach and Franck, 2015a, 2015b), and supermarkets (Seaton and Waterson, 2013). Supporting these empirical studies are theoretical analyses that show how price matching is an effective device for coordinating on higher prices (Harrington, 2017b) and maintaining compliance in the setting of high prices (Lu and Wright, 2010). In sum, price matching is commonly observed as part of a collusive scheme.

At the same time, I am unaware of any theory of competition that generates price matching. In a classic oligopoly model in which firms have common costs and identical products (or symmetrically differentiated products), equilibrium has identical prices but a firm's strategy is not to match rival firms' prices. A firm's optimal price depends on its cost and the strength and slope of its demand (i.e., cost and demand parameters), as well as the prices of rival firms, but it does not involve price matching. In standard static oligopoly models, a firm's profit-maximizing price is increasing in a rival firm's price but it is not one-for-one as in the case with price matching.[90] For example, a rise in a rival firm's price of 10% may bring forth a 5% increase in a firm's price but not the 10% increase that reflects price matching. There are dynamic models of competition with learning in which a firm's price can be more sensitive to a rival firm's price than occurs under static competition. For example, suppose market (and, thereby, firm) demand is stochastic over time and is not completely known to firms.[91] If firm A receives information suggesting demand is strong, it will then raise its price. Upon observing a higher price for firm A, firm B will also raise its price because a competitor is pricing higher (and, therefore, firm B's demand is stronger) but also because it signals that firm A received information consistent with stronger market demand (which means stronger demand for firm B). Still, a dynamically optimal pricing rule under competition does not call for price matching.

[90] See, for example, Vives (1999).

[91] The ensuing discussion is based on Riordan (1985).

Given that price matching is common under collusion but not under competition, it follows that the likelihood ratio,

$$\frac{Pr(pa \text{ has price matching } | pa \text{ is collusive})}{Pr(pa \text{ has price matching } | pa \text{ is competitive})}$$

is probably high.

This assessment of price matching is not intended to provide the basis for prohibiting price matching but rather to exemplify the process and objective of a research program for specifying the set of prohibited pricing algorithms. The execution of the three-step research program will, hopefully, reveal which pricing algorithms should be prohibited.

6.5 Legality of a Prohibition on Pricing Algorithms

Finally, let me examine the extent to which this notion 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 do not reveal an agreement), there could be an expanded role for the FTC in having it prosecute cases of collusion by AAs. Independently, Ezrachi and Stucke (2017) also suggest drawing on section 5 of the FTC:

> One way to square this circle may be framing the issue as market manipulation or an unfair practice. The focus shifts from the presence of an agreement among companies to the use of advanced algorithms to transform pre-existing market conditions in such a way to facilitate tacit collusion. While the mutual price monitoring at the heart of tacit collusion is legal under competition law, one may ask whether the creation of such a market dynamic, through "artificial" means, gives rise to antitrust intervention.[92]

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.[93]

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

[92] Ezrachi and Stucke (2017), p. 20.

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

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 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."[94] 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

At this point, the charitable skeptic might say: "You have proposed a possibly workable approach to finding a solution to a problem that it isn't clear exists." 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

[94] 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.

collusion can occur in a reasonably sophisticated simulated market. In addition, even if autonomous artificial agents could succeed in colluding in actual markets, there is still the prospect of entry undermining the setting of supracompetitive prices. One does not have to be oppositional to question the paper's underlying premise: Autonomous artificial price-setting agents can collude.

While I share some of that skepticism, I am not reassured in light of recent experiences in the areas of computing and information technology. One lesson learned is that its path 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, 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 can be rapid. The rate of technological change in the area of computing and information technology has been blisteringly fast. Ten years (perhaps even five 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 price-setting agents figuring out how to collude? It is not beyond the realm of possibility that collusion becomes commonplace in online markets.

With all that said, the persistent skeptic could still respond with: "Let's wait until collusion by autonomous artificial agents becomes a phenomenon and then develop a legal doctrine." What that would mean is that judges and lawyers would operate without a well-conceived framework and be forced to develop a legal doctrine "on the fly", which is rather disconcerting. Such an approach would most likely allow collusion to continue unabated until such time that the legal doctrine is in place. Of more long-term concern is that jurisprudence made without a proper foundation lends itself to mistakes and, given the weight of legal precedence, those errors could reverberate for some time. But there is a more fundamental argument for developing the doctrine outside of the judicial arena. Collusion by autonomous artificial agents is not about conspiracy or intent or communication, but about the use of pricing rules that embed a reward-punishment scheme which supports supracompetitive prices. Those pricing rules are well-established by both economic theory and the empirical analysis of collusion and cartels. For that reason, economists and computer scientists, along with legal scholars and practitioners, have a central role to play in developing a legal approach towards collusion by autonomous artificial price-setting agents. The foundation for that approach should be laid offline in the academic realm so that, when the legal cases come, judges will have a framework within which to develop the appropriate jurisprudence. Hopefully, the perspective offered here will be useful in developing that foundation.

> [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 Competition[95]

[95] Margarethe Vestager, "Algorithms and Competition," Remarks by the European Commissioner

References

[1] Andreoli-Versbach, Patrick and Jens-Uwe Franck, "Endogenous Price Commitment, Sticky and Leadership Pricing: Evidence from the Italian Petrol Market," *International Journal of Industrial Organization*, 40 (2015a), 32-48.

[2] Andreoli-Versbach, Patrick and Jens-Uwe Franck, "Econometric Evidence to Target Tacit Collusion in Oligopolistic Markets," *Journal of Competition Law & Economics*, 11 (2015b), 463-492.

[3] Areeda, Phillip E. and Herbert Hovenkamp, *Antitrust Law*, Volume VI, 3rd Edition, Austin, TX: Wolters Kluwer, 2010.

[4] Ballard, Dylan I. and Amar S. Naik, "Algorithms, Artificial Intelligence, and Joint Conduct," *CPI Antitrust Chronicle*, May 2017, 29-35.

[5] Bryne, David P. and Nicolas de Roos, "Learning to Coordinate: A Study in Retail Gasoline," University of Sydney, January 2017.

[6] 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.

[7] den Boer, Arnoud V. and Bert Zwart, "Simultaneously Learning and Optimizing Using Controlled Variance Pricing," *Management Science*, 60 (2014), 770-783.

[8] Calvano, Emilo, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello, "Algorithmic Pricing and Collusion: What Implications for Competition Policy?," University of Bologna, working paper, March 2018.

[9] Capobianco, Antonio and Pedro Gonzaga, "Algorithms and Competition: Friends or Foes?," *Competition Policy International*, August 2017.

for Competition at the Bundeskarellamt 18th Conference on Competition, Berlin, March 16, 2017.

[10] Cooper, William L., Tito Homen-de-Mello, and Anton J. Kleywegt, "Learning and Pricing with Models That Do Not Explicitly Incorporate Competition," *Operations Research*, 63 (2015), 86-103.

[11] Deng, Ai, "When Machines Learn to Collude: Lessons from a Recent Research Study on Artificial Intelligence," Bates White Economic Consulting, August 30, 2017.

[12] Deng, Ai, "An Antitrust Lawyer's Guide to Machine Learning," Bates White Economic Consulting, January 9, 2018.

[13] 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).

[14] Diakopoulous, Nicholas, "Algorithmic Accountability Reporting: On the Investigation of Black Boxes," Tow Center for Digital Journalism, Columbia University, 2014.

[15] Dogan, Ibrahim and Ali R. Güner, "A Reinforcement Learning Approach to Competitive Ordering and Pricing Problem," *Expert Systems*, 32 (2015), 39-47.

[16] 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.

[17] Ezrachi, Ariel and Maurice E. Stucke, *Virtual Competition: The Promise and Perils of the Algorithm-Driven Economy*, Cambridge, Mass.: Harvard University Press, 2016.

[18] 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.

[19] Ferreira, Kris Johnson, Bin Hong Alex Lee, and David Simchi-Levi, "Analytics for an Online Retailer: Demand Forecasting and Price Optimization," Massachusetts Institute of Technology, November 2015.

[20] Fisher, Marshall, Santiago Gallino, and Jun Li, "Competition-Based Dynamic Pricing in Online Retailing: A Methodology Validated with Field Experiments," Working Paper No. 1265, Ross School of Business, University of Michigan, August 2016 (*Management Science*, forthcoming)

[21] 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.

[22] Gal, Michal S., "Algorithmic-Facilitated Coordination: Market and Legal Solutions," *CPI Antitrust Chronicle*, May 2017a, pp. 22-28.

[23] Gal, Michal S., "Algorithmic-facilitated Coordination," OECD Background Paper, Roundtable on Algorithms and Collusion, DAF/COMP/WD(2017)26, June 2017b.

[24] Gal, Michal S. and Niva Elkin-Koren, "Algorithmic Consumers," *Harvard Journal of Law & Technology*, 30 (2017), 1-45.

[25] Gershman, Samuel J., "Reinforcement Learning and Causal Models," Harvard University, December 2015.

[26] Goodman, Bryce W., "Economic Models of (Algorithmic) Discrimination," 29th Conference on Neural Information Processing Systems, June 2016.

[27] 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-the-task-a-us-eu-perspective-on-anti-competitive-a> on August 17, 2017)

[28] Hanaki, Nobuyuki, Rajiv Sethi, Ido Erev, and Alexander Peterhansl, "Learning Strategies," *Journal of Economic Behavior & Organization*, 56 (2005), 523-542.

[29] Hardt, Moritz, Eric Price, and Nathan Srebro, "Equality of Opportunity in Supervised Learning," 29th Conference on Neural Information Processing Systems, June 2016.

[30] 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. <http://assets.wharton.upenn.edu/~harrij/pdf/Bayard%20Wickliffe%20Heath%20Memorial%20Lecture_03.12.pdf>

[31] Harrington, Joseph E., Jr., *The Theory of Collusion and Competition Policy*, Cambridge, Mass.: The MIT Press, 2017a.

[32] Harrington, Joseph E., Jr., "A Theory of Collusion with Partial Mutual Understanding," *Research in Economics*, 71 (2017b), 140-158.

[33] Hilsen, Hans Olav Østbø, "Simulating Dynamic Pricing Algorithm Performance in Heterogeneous Markets," Masters Thesis, Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, June 2016.

[34] Hingston, Philip and Graham Kendall, "Learning Versus Evolution in Iterated Prisoner's Dilemma," *Proceedings of the Congress on Evolutionary Computation*, 2004.

[35] Hovenkamp, Herbert, *Federal Antitrust Policy: The Law of Competition and Its Practice*, 5th edition, St. Paul, Minnesota: West Academic Publishing, 2016.

[36] Javanmard, Adel and Hamid Nazerzadeh, "Dynamic Pricing in High-dimensions," Marshall School of Business, U. of Southern California, November 2016.

[37] Johnson, Kory D., Dean P. Foster, and Robert A. Stine, "Fairness-Aware Data Mining: Ensuring Fairness in Arbitrary Models," The Wharton School, University of Pennsylvania, March 2016.

[38] 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.

[39] Kaplow, Louis, *Competition Policy and Price Fixing*, Princeton: Princeton University Press, 2013.

[40] Kaplow, Louis, "Likelihood Ratio Tests and Legal Decision Rules," *American Law and Economics Review*, 16 (2014), 1-39.

[41] Kawai, Kei and Jun Nakabayashi, "Detecting Large-Scale Collusion in Procurement Auctions," University of California-Berkeley, August 2015.

[42] Kleinberg, Jon, Sendhil Mullainathan, and Manish Raghavan, "Inherent Trade-Offs in the Fair Determination of Risk Scores," Cornell University, September 2016.

[43] Kovacic, William E., "The Identification and Proof of Horizontal Agreements Under the Antitrust Laws," *Antitrust Bulletin*, 38 (Spring 1993), 5-81.

[44] Kühn, Kai-Uwe, "Fighting Collusion by Regulating Communication Between Firms," *Economic Policy: A European Forum*, 32 (April 2001), 167-197.

[45] LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton, "Deep Learning," *Nature*, 521 (2015), 436-444.

[46] Lewis, Harry R. and Christos H. Papadimitriou, *Elements of the Theory of Computation*, Englewood Cliffs, NJ: Prentice-Hall, 1981.

[47] Lu, Yuanzhu and Julian Wright, "Tacit Collusion with Price-Matching Punishments," *International Journal of Industrial Organization*, 28 (2010), 298-306.

[48] Markham, Jesse W., "The Nature and Significance of Price Leadership," *American Economic Review*, 41 (1951), 891-905.

[49] 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.

[50] Mehra, Salil K., "Antitrust and the Robo-Seller: Competition in the Time of Algorithms," *Minnesota Law Review*, 100 (2016), 1323-1375.

[51] Mehra, Salil K., "Robo-Seller Prosecutions and Antitrust's Error-Cost Framework," *CPI Antitrust Chronicle*, May 2017, 36-39.

[52] Miller, John H., "The Coevolution of Automata in the Repeated Prisoner's Dilemma," *Journal of Economic Behavior & Organization*, 29 (1996), 87-112.

[53] Misra, Kanishka, Eric M. Schwartz, and Jacob Abernethy, "Dynamic Online Pricing with Incomplete Information Using Mult-Armed Bandit Experiments," Rady School of Management, University of California, San Diego, June 2017.

[54] 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.

[55] Motta, Massimo, *Competition Policy: Theory and Practice,* Cambridge: Cambridge University Press, 2004.

[56] Nambiar, Mila, David Simchi-Levi, and Ha Wang, "Dynamic Learning and Price Optimization with Endogeneity Effect," Operations Research Center, MIT, November 2016.

[57] OECD Background Note by the Secretariat, "Algorithms and Collusion," Roundtable on Algorithms and Collusion, DAF/COMP(2017)4, June 2017.

[58] Okuliar, Alex and Elena Kamenir, "Pricing Algorithms: Conscious Parallelism or Conscious Commitment?," *Competition Policy International*, July 2017.

[59] Oxera Consulting, "When Algorithms Set Prices: Winners and Losers," Discussion Paper, 19 June 2017.

[60] Page, William H., "Communication and Concerted Action," *Loyola University Chicago Law Journal,* 38 (2007), 405-460.

[61] Page, William H., "*Twombly* and Communication: The Emerging Definition of Concerted Action under the New Pleading Standards," *Journal of Competition Law and Economics*, 5 (2009), 439-468.

[62] Page, William H., "Objective and Subjective Theories of Concerted Action," *Antitrust Law Journal*, 79 (2013), 215-251.

[63] Posner, Richard A., "Oligopoly and the Antitrust Laws: A Suggested Approach," *Stanford Law Review*, 21 (1969), 1562-1606.

[64] Posner, Richard A., *Antitrust Law*, Second Edition, Chicago: University of Chicago Press, 2001.

[65] Posner, Richard A., “Review of Kaplow, *Competition Policy and Price Fixing*,” *Antitrust Law Journal*, 79 (2015), pp. 761-768.

[66] *Proof of Conspiracy Under Federal Antitrust Laws*, ABA Section of Antitrust Law, ABA Publishing, 2010.

[67] Preston, John and Mark Bishop, eds., *Views into the Chinese Room: New Essays on Searle and Artificial Intelligence*, Oxford: Clarendon Press, 2002.

[68] Qiang, Sheng and Mohsen Bayati, "Dynamic Pricing with Demand Covariates," Graduate School of Business, Stanford University, April 2016.

[69] Riordan, Michael H., "Imperfect Information and Dynamic Conjectural Variations," *RAND Journal of Economics*, 16 (1985), 41-50.

[70] Salcedo, Bruno, "Pricing Algorithms and Tacit Collusion," Cornell University, July 2016.

[71] Scherer, F. M., *Industrial Market Structure and Economic Performance*, 2nd Edition, Boston: Houghton Mifflin, 1980.

[72] Scherer, Matthew U., "Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies," *Harvard Journal of Law & Technology*, 29 (2016), 354-400.

[73] Searle, John R., "Minds, Brains, and Programs," *Behavioral and Brain Sciences*, 3 (1980), 417-424.

[74] 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.

[75] 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.

[76] Seaton, Jonathan S. and Michael Waterson, "Identifying and Characterising Price Leadership in British Supermarkets," *International Journal of Industrial Organization*, 31 (2013), 392-403.

[77] Shakya, S., C. M. Chin, and G. Owusu, "An AI-based System for Pricing Diverse Products and Services," *Knowledge-Based Systems*, 23 (2010), 357-362.

[78] Shakya, S., M. Kern, G. Owusu, and C. M. Chin, "Neural Network Demand Models and Evolutionary Optimisers for Dynamic Pricing," *Knowledge-Based Systems*, 29 (2012), 44-53.

[79] Sutton, Richard S. and Andrew G. Barto, *Reinforcement Learning: An Introduction*, Cambridge, Mass.: The MIT Press, 2000.

[80] Tesauro, Gerald and Jeffrey O. Kephart, "Pricing in Agent Economics Using Multi-Agent Q-Learning," *Autonomous Agents and Multi-Agent Systems*, 5 (2002), 289-304.

[81] Turner, Donald F., "The Definition of Agreement under the Sherman Act: Conscious Parallelism and Refusals to Deal," *Harvard Law Review*, 75 (1962), 655-706.

[82] 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.

[83] Viscusi, W. Kip, Joseph E. Harrington, Jr., and John M. Vernon, *Economics of Regulation and Antitrust*, 4th Edition, Cambridge, Mass.: The MIT Press, 2005.

[84] Vives, Xavier, *Oligopoly Pricing: Old Ideas and New Tools*, Cambridge, Mass.: The MIT Press, 1999.

[85] 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.

[86] Waltman, Ludo and Uzay Kaymak, "Q-learning Agents in a Cournot Oligopoly Model," *Journal of Economic Dynamics & Control*, 32 (2008), 3275-3293.

[87] Werden, Gregory J., "Economic Evidence on the Existence of Collusion: Reconciling Antitrust Law with Oligopoly Theory," *Antitrust Law Journal*, 3 (2004) 719-800.

[88] Xie, Ming and Jian Chen, "Studies on Horizontal Competition Among Homogeneous Retailers through Agent-based Simulations," *Journal of Systems Science and Systems Engineering*, 13 (2004), 490-505.

[89] Yao, Dennis A. and Susan S. DeSanti, "Game Theory and the Legal Analysis of Tacit Collusion," *Antitrust Bulletin*, 38 (1993), 113-141.