CHAPTER 4

Price Collusion Using Artificial Intelligence

Joshua Gans
University of Toronto

SYNOPSIS

§ 4.01 Introduction

For many years, computer scientists have been creating environments where computerised agents can play games against each other. One of the most well-known examples is the tournament organised by Robert Axelrod in 1981, where participants programmed agents to play Prisoners' Dilemma games against other agents. The winner of this contest was Anatol Rapoport, who used a 'tit-for-tat' strategy. The tit-for-tat strategy calls for the agent to "cooperate" with any agent that has cooperated in the previous round, but to "defect" if instead the other agent has defected in the previous round. The fact that the tit-for-tat strategy won the contest was interesting because many game theory results tended to focus on more severe 'grim-trigger' strategies, which were not as forgiving as Rapoport's entry. It is worth noting that while grim-trigger strategies can support cooperation (in subgame perfect equilibrium), the same is not true for tit-for-tat, as it may not be credible to forgive the other agent for defecting in a previous round.

It is not surprising that economists and game theorists are interested in exploring how new AI techniques could be used in similar contests or experiments. Pricing algorithms are of particular interest. Antitrust lawyers have raised concerns that if AI algorithms were able to control pricing decisions, were programmed to maximize profits, and had access to public information about competitors' prices, then competing firms might arrive at a collusive outcome without any communication. Indeed, since AI agents aim to maximize the profits of individual firms, there may be no violation of price-fixing or antitrust collusion laws.[1]

Prior to this, the success of using algorithms to price products had been mixed. In some cases, simple algorithms proved disastrous.

> *In the Spring of 2011, two online retailers offered copies of Peter Lawrence's textbook The Making of a Fly on Amazon for $18,651,718.08 and $23,698,655.93 respectively. This was the result of both sellers using automated pricing algorithms. Everyday, the algorithm used by seller 1 set the price of the book to be 0.9983 times the price charged by seller 2. Later in the day, seller 2's algorithm would adjust its price to be 1.27059 times that of seller 1. Prices increased exponentially and remained over one million dollars for at least ten days (!), until one of the sellers took notice and adjusted its price to $106.23.*[2]

Given this, many economists were initially sceptical that AI algorithms, interacting at arm's length and with different interests, could learn and implement collusive outcomes. In many instances, either communication was explicit, or the colluding parties adopted simple rules such as dividing up and allocating different markets.[3]

Despite scepticism, researchers have been inspired to explore certain questions. One area of investigation examines whether AI predictions of outside factors, like market demand, can help firms with pricing decisions made by rational economic agents, ultimately leading to collusion. In this chapter, we will analyze this research direction and discover that while AI adoption might hurt consumers in situations of tacit collusion, it could also weaken the circumstances that allow collusion to occur.

Another research focus has been on whether AI algorithms can learn to collude in typical oligopoly scenarios, such as Bertrand competition. Studies have shown that machine learning algorithms used by competing firms often lead to higher-than-

---

[1] Ezrachi, A. and M. Stucke, VIRTUAL COMPETITION, Harvard University Press (2017); Calvano, E., G. Calzolari, V. Denicolo, and S. Pastorello, *Artificial Intelligence, Algorithmic Pricing, and Collusion*, 110 (10) AMERICAN ECONOMIC REVIEW 3267–3297 (2020).

[2] Salcedo, B., Pricing Algorithms and Tacit Collusion p.2 (2015) (manuscript, on file with Pennsylvania State University).

[3] For example, Byford, M.C. and J.S. Gans, *Collusion at the Extensive Margin*, 37 INTERNATIONAL JOURNAL OF INDUSTRIAL ORGANIZATION 75–83 (2014).

competitive prices, although not always reaching the monopoly outcome that is the usual goal of collusive actions. The latter half of this chapter will provide an overview of the current research in this area.

§ 4.02 What Does AI Do?

Before considering the potential for AI collusion, it is useful to make clear precisely what AI we are talking about. This chapter looks at the types of AI that currently exist under the broad area of machine learning in neural networks. All the recent advances in AI, including for example AlphaGo and ChatGPT, are based on machine learning. In that regard, AI only captures a part of what we would normally consider intelligence.

That part is prediction. All recent advances in AI are advances in the statistics of prediction.[4] Current AI is not about building machines that can think like humans but rather about building machines that can predict better than humans. Thus, the value of AI comes from its ability to predict outcomes more accurately and quickly than humans can.

Current methods of AI are driven by three key components: data, algorithms, and computational power. These three components have improved dramatically in recent years, which has led to significant advances in AI. AI is not a single technology but rather a collection of technologies that work together to enable prediction. In that respect, AI can reduce the uncertainty faced by firms when choosing their prices. In the first part of the chapter, we consider how reducing uncertainty on one key dimension—demand prediction—can impact the conditions for price collusion. The second part looks at whether AI agents can learn to collude without human intervention by predicting the responses of rivals.

§ 4.03 Does AI Prediction of Demand Facilitate Collusion?

[1] In General

First, let us look at how AI predicts demand. In simple terms, market demand in a certain period depends on the price and an uncertain factor. AI prediction provides a signal of the uncertain factor.

To illustrate, consider the following stylized example. Suppose that there are two firms that set prices in each period of an ongoing competition. They sell identical products, so if they have different prices, the firm with the lower price gets all the customers. If they have the same price, they share the customers evenly. The realized market price in each period is always the lowest price.

Without AI prediction, firms must set their prices before knowing the uncertain demand factor. Since they compete on price, the quantity produced adjusts to meet the

---

[4] Ajay Agrawal, Joshua Gans, and Avi Goldfarb, PREDICTION MACHINES: THE SIMPLE ECONOMICS OF ARTIFICIAL INTELLIGENCE (Harvard Business Review Press: Boston, 2018).

market demand.[5] In this situation, the collusive price (when firms cooperate in setting prices) is determined by a profit-maximizing price based on expected demand.[6]

Researchers have analyzed collusion in this environment by assuming that firms follow a "grim-trigger" strategy.[7] That is, firms will continue to charge the collusive price and share the market demand as long as the other firm has done so in the past. If one firm ever deviates and lowers its price (thereby, stealing the other firm's business in that period), both firms will thereafter become competitive and set their prices equal to their costs (in all future periods).

Firms will choose to collude rather than deviate if the short-term gain from deviating (*i.e.*, undercutting the collusive price and serve the entire market) is smaller than the long-term loss from foregoing collusion (in all future periods). Thus, firms will choose to collude rather than deviate if their discount rate is smaller than a certain threshold. The main question is whether AI prediction changes this threshold or other variables, such as the collusive price.

When the firms have access to AI prediction, the collusive price can adjust based on the level of demand predicted by AI. For simplicity, suppose that AI prediction offers a perfect signal of the level of demand. In this case, if both firms adopt AI, they will receive the same signal. Firms might consider deviating from the collusive price after receiving AI predictions. However, they will not deviate if their discount rate is lower than a certain threshold, and that threshold is lower when AI predicts a higher level of demand. Thus, collusion is less likely to be sustainable when AI predicts higher levels of demand. This is consistent with known results that "price wars" are more likely to occur during periods of high demand.[8]

[2] Impact of AI on Collusion

Without AI, collusive firm profits are lower than those with AI, meaning AI increases producer surplus (the difference between the revenues earned by firms and their production costs) when collusion can be sustained (with and without AI). However, when collusion can be sustained, consumer surplus (the difference between what consumers are willing to pay and what they actually pay) falls due to AI adoption. This happens because AI raises consumer surplus when demand is low but reduces it when demand is high. Overall, when collusion can be sustained, total surplus falls when AI is adopted, leading to decreased total welfare.

However, the scope for collusion changes with AI adoption. The discount rate threshold for sustaining collusion is higher when demand is low and lower when demand is high. Therefore, it is possible that collusion may become unsustainable and

---

[5] This is known as the "make-to-order" business model.

[6] Gans, J.S., *Artificial Intelligence Adoption in a Monopoly Market*, MANAGERIAL AND DECISION ECONOMICS, (2023) *forthcoming*.

[7] Miklós-Thal, J. and C. Tucker, *Collusion by Algorithm: Does Better Demand Prediction Facilitate Coordination Between Sellers?* 65 (4) MANAGEMENT SCIENCE 1552–1561 (2019).

[8] *See, e.g.*, Rotemberg, J. and G. Saloner, *A Supergame-Theoretic Model of Price Wars During Booms*, 76 AMERICAN ECONOMIC REVIEW 390–407 (1986).

break down when there is a period of high demand, especially when the probability of low demand in the future increases. Recall that periods of high demand are periods where AI adoption exacerbates the consumer harm from collusion. Thus, while adopting AI under collusion can create some welfare reductions, AI adoption itself may undermine the ability of a cartel to sustain itself.

[3] Outsourced Pricing Algorithms

The previous analysis doesn't specify the origin of the algorithms that allow firms to receive a perfect signal of demand. In many cases, individual firms will not develop their own AI predictions but will instead buy those predictions or buy the algorithms that generate them from third-party providers. For example, one study looked at German gasoline retailers adopting pricing algorithms focused on demand forecasting.[9] They found that in non-monopoly markets, algorithm adoption increased profit margins by 9 percent. In duopoly markets, algorithms had no margin effect if adopted by one firm but a 28 percent increase when both firms adopted them. This is consistent with theoretical predictions about the impact of such algorithms on price competition.[10]

Another researcher explored whether having third parties provide demand-prediction algorithms to competing firms might lead to collusive outcomes.[11] He found that there is a fundamental difference in the design of algorithms that firms would create for themselves versus those intended to be sold to multiple firms. The third-party provider is concerned that its algorithm might compete against itself.

While this might initially seem to cause third parties to provide algorithms that dampen competition, it does not actually lead to higher prices. The reason is that if an algorithm provider developed a product that resulted in a higher average price, this would create value for those who don't adopt the algorithm. Since the provider wants to sell their product to more firms at a higher price, having the algorithm reduce incentives for mutual adoption goes against their goal. Instead, the provider produces predictions that are more sensitive to demand conditions. As previously argued, the resulting price variation makes consumers worse off because they face greater exercise of monopoly power when demand is high. Therefore, consumers are negatively impacted by third-party provision when the firms have some market power, but that provision is not a tool for collusive pricing.

---

[9] Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu *Identifying Algorithmic Pricing Technology Adoption in Retail Gasoline Markets*, AEA PAPERS AND PROCEEDINGS, vol. 112, pp. 457–460 (2022).

[10] Brown, Z. Y. and A. MacKay, *Competition in Pricing Algorithms*, AMERICAN ECONOMIC JOURNAL: MICROECONOMICS (2023).

[11] Harrington, J. E., *The Effect of Outsourcing Pricing Algorithms on Market Competition*, 68 (9) MANAGEMENT SCIENCE 6889–6906 (2022).

§ 4.04 Can AIs Learn to Collude?

[1] Overview

A significant achievement in AI has been its ability to learn to play games against humans. In games like Go or Chess, AI trained with reinforcement learning can now outperform any person. AI can also excel in computer games such as Pong, Atari games, DOTA, and Quake III.

This led economists to question if AI could be trained to obtain better outcomes in games that represent competition between firms.

Reinforcement learning is a process where an AI algorithm is continuously updated based on its experience playing a game. Strategies that worked well against opponents' choices in the past are given more importance, while others are given less. Often, AI gains experience by playing against other AI and learning from shared experiences. For example, AlphaZero, the best AI at Go, was trained in just a few hours by playing millions of games against itself. Could a similar outcome be achieved if AI were trained to play pricing games in oligopolistic markets, particularly those involving repeated games and history-dependent strategies? In other words, could AI learn to collude?

Initially, economists believed this would not be possible. They argued that reinforcement learning, like other machine learning methods, would fall into a category of learning behaviors called adaptive learning.[12] Adaptive learning involves choosing strategies expected to perform best against all possible combinations of strategies based on the competitor's recent play history. As a result, agents would learn which strategies were dominated and avoid playing them, iterating this process over time. In a common class of games,[13] where each agent's best replies were non-decreasing in the strategies of other agents, every adaptive learning process would converge to a unique Nash equilibrium (if it existed), even if players could condition their strategies on their rivals' past play. This would happen even in repeated games, suggesting that AI trained to play pricing games using reinforcement learning wouldn't reach an equilibrium or an outcome with supra-competitive pricing. This is because reinforcement learning doesn't allow agents to coordinate future play based on their play histories. However, despite this initial skepticism, economists started applying reinforcement learning methods to familiar oligopolistic pricing games to observe the resulting price outcomes.

One notable aspect of the economic research into AI collusion is that it focuses on learning price equilibria by using existing reinforcement learning methods rather than examining abstract concepts like perfect prediction or imperfect prediction with a single parameter capturing prediction quality. This approach results in simulated outcomes, which can help improve our understanding of how AI algorithms will interact with one another. In what follows, we will explore Q-learning, a popular research tool, and what economists have learned about AI collusion using this method.

---

[12] Milgrom, P. and J. Roberts, *Adaptive and Sophisticated Learning in Normal Form Games*, 3 (1) GAMES AND ECONOMIC BEHAVIOR 82–100 (1991).

[13] That is, supermodular games.

[2] Q-Learning

Q-learning is a widely used tool for implementing reinforcement learning. The "Q" represents a matrix that records the algorithm's current assessment of the likely payoff going forward for the available choices in every possible situation. This tool has proven useful in dynamic game environments with multiple agents, where they aim to maximize a payoff in a repeated game without any information about the model structure. As a result, economists have become interested in Q-learning to study the potential for AI-driven collusion.

In a typical scenario, two or more firms are involved in a Bertrand game, where their action space consists of pricing options. Each AI agent's goal is to maximize expected discounted profits over an infinite time horizon. The profits for a firm in each period depend on the current price set by the firm and the prices of its rivals. The AI agents play this pricing game without knowledge of the demand function for any player, making it a "model-free" learning situation. The research questions are: (i) can AI agents programmed to use Q-learning achieve a stable pricing strategy between them, and (ii) are those prices similar to competitive prices (*i.e.*, one-shot Nash equilibrium prices) or to optimal collusive prices? A secondary question involves understanding whether the learning strategy involves a fixed price chosen by each agent or a counterfactual strategy, like a grim-trigger strategy or tit-for-tat.

Q-learning in this environment revolves around the Q-matrix, which estimates the expected present value of profits for each agent when choosing a specific action in a particular state of the game. The Q-matrix is initialized with arbitrary starting values but is then updated based on the agent's experience. The updating process takes into account the agent's observed profit from charging a certain price, the current state (previous prices set by all agents), and the future state resulting from the current prices of all agents. Strategies that perform well see their corresponding expected values in the Q-matrix increase, while those that perform poorly experience the opposite. This reinforcement mechanism pushes the agent towards an optimal outcome.

The choice of action is not purely deterministic, meaning it does not always maximize the expected value at the time. Instead, the algorithm includes some "experimentation" and hence randomness that may lead the firm to choose a different price with a certain probability. This randomness can involve more complex choices in richer action spaces.[14] Such experimentation decreases over time, meaning that the prices believed to maximize the expected present value of profits become increasingly more likely to be chosen.

In single-agent decision environments, there are conditions under which Q-learning algorithms will converge to the optimal solution. However, no such conditions have been identified for multi-agent repeated game environments. As a result, research in this area involves running simulations with different environments and parameters to see if a stable, equilibrium-like outcome emerges. The following section will discuss stable and unstable outcomes found by researchers.

---

[14] Such as the Boltzmann algorithm, where the probability of selecting a price is based on the Q-matrix values.

[3] Outcomes of Dynamic Games Between AI Agents

Researchers have explored the outcomes of repeated pricing games played by AI agents.[15] In these simulations, the agents interact in a market where products are differentiated, and each firm aims to maximize profits. The researchers found that the simulations usually converged to a stable outcome, with the average price consistently above the competitive (Nash equilibrium) level, resulting in higher profits for the firms.

The research also investigated whether the outcomes were driven by patterns that normally support collusion, like punishment for deviations in the dynamic game. They found that when one firm artificially lowered its price, the other firm would respond by lowering its price, but not necessarily matching the competitor's price. Both firms would slowly increase their prices until the previous stable outcome was restored. Interestingly, this slow response is inconsistent with the theoretical literature on dynamic competition in repeated games.

[4] Myopic vs. Sophisticated Q-Learning

Further research has explored the limitations of AI algorithms learning to play competitive stage games without payoffs that depend on future outcomes.[16] This research has found that using Q-learning with asynchronous updating, where the Q-matrix updates in a certain way, can result in prices converging to a range of different supra-competitive levels.

Researchers have also explored modifications to Q-learning to allow for more sophisticated strategies, called synchronous updating. In perfect synchronous updating, the algorithm can see the prices of competitors and knows the demand and cost functions. This allows it to calculate what the actual profits would have been had a different price been chosen based on the state. With perfect synchronous updating, even when agents are forward-looking, the simulation outcomes accord with theoretical predictions.

Another more limited form of synchronous updating involves the algorithm updating in a way that contains the knowledge that demand curves slope downwards. This leads to a set of convergent solutions: one consistent with competitive pricing and another with supra-competitive pricing.

These results highlight the impact of the type of updating that algorithms perform on eventual pricing outcomes. A key question, then, is what type of updating rules firms might choose. Research shows that using an asynchronous approach is a dominant strategy for both firms, but understanding what algorithm they might be facing in a competitor is likely to be a significant issue.[17]

---

[15] Calvano et al., *op.cit.*

[16] Asker, J., C. Fershtman, A. Pakes, et al., *Artificial Intelligence, Algorithm Design and Pricing*, AEA Papers and Proceedings, Volume 112, pp. 452–456 (2022).

[17] Asker, J., C. Fershtman, and A. Pakes, *The Impact of AI Design on Pricing* (Working Paper, 2022).

§ 4.05 Responses to AI Collusion

[1] In General

The discussions we have had so far about collusion involve markets without any regulations. Even in these markets, collusion can trigger reactions. For example, some researchers suggest that consumers could use algorithms to create more competitive outcomes.[18] Others think AI could be used by authorities to identify and prevent collusion.[19]

Now, let us look at situations where an interested party designs the market, and that party sets the rules for competing sellers. We will discuss two cases: first, a platform that hosts competing sellers and can decide which sellers to promote; second, an auction environment where one seller is trying to earn more revenue from buyers who compete to buy a product but may also collude to keep prices low. In both cases, we will see that some lessons from using AI prediction to help agents learn can also be used to develop strategies that counteract collusion.

[2] Platforms

In our previous analysis, we looked at traditional collusion environments with several competing firms. Recently, there has been an interest in situations where firms compete on a platform, like a marketplace or exchange that connects buyers and sellers. A platform with a monopoly might want to ensure that competition between sellers works in its favour, so it is important to ask whether platforms can manage competition when sellers use algorithms that could lead to collusion.

One study examined this question by examining platform rules that direct customers to sellers offering the lowest prices.[20] The idea is that the platform promotes a certain number of sellers with the lowest prices, making it harder for others to be noticed. This can create a trade-off for the platform, as promoting low prices may lead to fewer sellers and less variety for customers. The study showed that platforms typically use steering techniques to lower prices enough to compensate customers for the reduced variety. However, if sellers can collude, these steering methods become less effective, and the platform may need to use more complicated mechanisms to encourage competition.

One way to deal with collusion between sellers is for the platform owner to use a dynamic price-directed mechanism to decide which sellers get promoted. In this approach, one seller is chosen for prominence based on their low price, and they can keep that position for a certain amount of time as long as they don't raise their price,

---

[18] Gal, M. S. and N. Elkin-Koren, *Algorithmic Consumers*, 30 HARVARD JOURNAL OF LAW AND TECHNOLOGY 309 (2016).

[19] Calvano, E., G. Calzolari, V. Denicolò, J. E. Harrington Jr. and S. Pastorello, *Protecting Consumers From Collusive Prices Due to AI*, 370 (6520) SCIENCE 1040–1042 (2020).

[20] Justin Johnson, Andrew Rhodes, and Matthijs R. Wildenbeest, "*Platform Design When Sellers Use Pricing Algorithms*" (mimeo. Cornell, 2020), *available at* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3691621.

and no one else undercuts it by more than a set amount. This mechanism makes it harder for sellers to collude because it rewards a seller for deviating from the collusive outcome. The study showed that there are situations where the platform wants to use these mechanisms to break collusion in a way that increases both its own profits and consumer benefits.

What about the role of AI in this context? The same study conducted experiments with sellers using AI algorithms capable of Q-learning to see if the platform's competition-promoting strategies (which can also be implemented algorithmically) could improve platform profits and consumer benefits. The results showed that simple price-directed steering could reduce prices and increase benefits when there's not too much product differentiation. The study also found that when sellers use Q-learning, the dynamic price-directed prominence strategy can significantly lower prices if one seller is promoted to nearly all customers.

[3] Auctions

Auctions are used by people on either the buying or selling side of a market to encourage competition among people on the other side. For example, buyers might use auctions to create competition and hence purchase goods at lower prices, while sellers might use them to obtain higher prices from buyers. The auction rules can impact the benefits for the person who designed the auction. Until now, we have focused on competition among sellers, but in this section, we will discuss the use of AI by buyers in auction settings. This is important because the most intense auctions nowadays are for advertising space, where the bidders are advertisers seeking to buy space. These auctions take place in real time when users search for information or visit websites.

We will look at how AI agents, acting on behalf of buyers, bid in two different types of auctions: first and second-price sealed-bid auctions. In these auctions, bidders submit a price (their bid) to the auctioneer, and the auctioneer sells the product to the highest bidder. In the first-price auction, the winner pays their bid amount, while in a second-price auction, the winner pays the amount of the second-highest bid. When buyers have only private values for the product (meaning their values are not related to the values of other buyers), both auction designs bring in the same expected revenue for the seller.[21] Notably, in a second-price auction, it is a dominant strategy for bidders to submit bids equal to their willingness to pay.

One study used Q-learning algorithms with two competing bidders to see if the bidders could submit bids that resulted in low revenue for the seller.[22] They found that the second-price auction led to more competitive outcomes than the first-price auction. This is not surprising, considering our previous discussions. First, the first-price auction is similar to a type of competition we have seen before which results in higher-than-

---

[21] This *revenue equivalence* result assumes that bidders are risk-neutral and *ex ante* symmetric. Milgrom, Paul R., and Robert J. Weber, "*A Theory of Auctions and Competitive Bidding*," ECONOMETRICA, 1982: 1089–1122.

[22] Banchio, M. and A. Skrzypacz, *Artificial Intelligence and Auction Design*, Proceedings of the 23rd ACM Conference on Economics and Computation, pp. 30–31 (2022).

competitive profits, meaning lower revenue for the seller in the present auction context. Second, the second-price auction has a dominant strategy, and we have seen that Q-learning works well in finding dominant strategies. In the second-price auction, bidding according to one's willingness to pay is the dominant strategy, which leads to high revenue for the seller.

Interestingly, it was found that more competitive outcomes could be created in the first-price auction by giving AI agents information on the winning bid. With this information, the AI agents can update their learning using "what-if" scenarios regarding how different bid choices might have resulted in different outcomes. This approach is similar to one we have seen before in a different context. As in the case of regular market competition, providing more sophistication to AI algorithms leads to more competitive outcomes. However, this does not mean that buyers with AI at their disposal will necessarily choose to use these "smarter" AIs for bidding.

§ 4.06 Conclusion

This chapter has explored the implications and applications of AI and machine learning in various market settings, focusing on collusion and responses to collusion, including platform design and bidding in auction markets. We have seen that AI can both facilitate collusion and be used to counteract it. Platforms can implement specific mechanisms to promote competition and prevent collusion among sellers, while buyers in auctions can utilise AI algorithms to achieve more competitive outcomes.

The insights demonstrate the diverse and powerful ways AI can be employed in modern markets. As AI continues to advance and become more sophisticated, it is crucial for researchers, policymakers, and market participants to understand and harness its potential to promote competitive markets and prevent anti-competitive behaviour. Thus, while it is premature to imagine that there will be significant use of AI to generate collusive outcomes yet, the future remains open, and there may be requirements for antitrust law to evolve and consider such possibilities.