COLLUSIVE OUTCOMES VIA PRICING ALGORITHMS

KARSTEN T. HANSEN [*], KANISHKA MISRA[**], AND MALLESH M PAI[***]

Key Points:

- We provide a survey of mechanisms that can produce non-competitive outcomes in markets where all firms rely on algorithms for pricing.
- Previous literature studies channels where pricing algorithms can respond to competitors' prices to essentially "enforce" collusive outcomes.
- We report new results which show collusive price outcomes may occur even when firms rely on algorithms that have no competitive price information.

[*] RADY SCHOOL OF MANAGEMENT, UNIVERSITY OF CALIFORNIA, SAN DIEGO. K4HANSEN@UCSD.EDU
[**] RADY SCHOOL OF MANAGEMENT, UNIVERSITY OF CALIFORNIA, SAN DIEGO. KAMISRA@UCSD.EDU
[***] DEPARTMENT OF ECONOMICS, RICE UNIVERSITY. MALLESH.PAI@RICE.EDU
Date: January 15, 2021.
Pai gratefully acknowledges financial support from NSF grant CCF-1763349.

1. Introduction

As Machine Learning and Artificial Intelligence based methods proliferate, firms are increasingly delegating management decisions which used to be the domain of human decision makers to algorithms. Management decisions such as how to price products/ services, how much stock to carry (and of which products), and what products/ services to continue/ discontinue have all increasingly been delegated to algorithms. To the upside, these algorithms have allowed companies to achieve and then successfully manage unprecedented scale: for example, Netflix uses machine learning to manage a catalog of tens of thousands of videos, Amazon maintains an inventory of and prices millions of products.

However, recent research has also uncovered an unforeseen downside to algorithmic pricing, in particular a downside that raises several antitrust concerns: the possibility that companies may end up colluding *through algorithms*. Of particular concern is the use of algorithms for pricing — an area where online retailers increasingly use machine learning algorithms.[1] A steadily increasing body of work both in economics, and in the public policy/legal spheres suggests that long-run prices when firms use (certain) algorithms for pricing may result in prices that are above the competitive level, i.e. inconsistent with price competition and more consistent with price-levels that result from collusion.

The purpose of this paper is to survey various known mechanisms by which competing pricing algorithms may nevertheless settle upon prices that are higher than the competitive level. A majority of the literature considers the use of algorithms that observe and respond to competitors' prices. These algorithms, by various means, "learn" to coordinate on high prices, by essentially punishing low prices/undercutting by competitors. We then present a newer strand of the literature that shows that supra-competitive prices are possible even if firms do not observe competitors' prices. The distinction is interesting and important from a policy standpoint for a couple of reasons. Firstly, it suggests that simple remedies, like ensuring that each seller's algorithm is oblivious to competitors' prices, may not have the desired effect in general. Secondly, it suggests that "stable" markets, i.e. ones where day-to-day (or more generally, period-to-period) demand is a relatively predictable function of prices are ones that are susceptible to the latter mechanism, and that such markets may be the ones to focus on in terms of regulatory oversight.

[1] Council of Economic Advisers 'White House'. "Big Data and Differential Pricing". In: February (2015). URL: https://obamawhitehouse.archives.gov/blog/2015/02/06/economics-big-data-and-differential-pricing; Le Chen, Alan Mislove, and Christo Wilson. "An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace". In: *Proceedings of the 25th International World Wide Web Conference (WWW 2016)*. Montreal, Canada, 2016; Kanishka Misra, Eric M Schwartz, and Jacob Abernethy. "Dynamic online pricing with incomplete information using multiarmed bandit experiments". In: *Marketing Science* 38.2 (2019), pp. 226–252.

Throughout the paper, we focus (for simplicity) on the case of a duopoly. The two sellers' goods are imperfect substitutes. To fix ideas consider the product as a service or perishable good— there is no intertemporal substitution or stockpiling by consumers. In each period, sellers announce a price for their product. Demand in that period is a stochastic function that depends on both sellers' prices in the usual way: decreasing in own price, increasing in competitors' price. Several of the results/mechanisms extend more generally, but the basic points are best explained within the context of this barebones setting.

2. Algorithms that Observe Competitors' prices

A majority of the literature that shows the possibility of collusive seeming prices under algorithmic pricing considers settings where competitors are able to observe each other's prices. Each seller's algorithm thus takes past prices by competitors into account when setting price. The literature has identified several mechanisms through which the algorithms may jointly settle on high prices.

In this series, we begin by discussing the thought provoking work of Calvano et. al.[2]— they study a setting where firms each use a particular kind of machine learning algorithm (formally these are called reinforcement learning algorithms) known as Q-learning. These are "model free" learning algorithms, i.e., they do not come in with any preconceived notion of either the demand or how competitors' algorithms are programmed. Instead the algorithms proceed by trial and error, increasing the frequency of strategies (reinforcing) that performed well in the past. The authors show that these algorithms consistently learn to charge supracompetitive prices, without communicating with one another. The high prices are sustained by classic collusive strategies: undercutting is responded to with a punishment (i.e., a price-war), followed by a gradual return to cooperation. Note that none of these collusive strategies are programmed in to the algorithm—the algorithm itself is neutral, and these strategies are instead "discovered" by the algorithm in an unsupervised fashion.

Some papers have tried to study the role of timing of pricing updates: Brown and Mackay[3] consider the implications of timing of pricing updating on equilibrium prices. In their model, firms know the true demand model, but unlike the canonical pricing model, they allow firms to update prices asynchronously. They find that these asynchronously responses can lead to non-Nash prices in equilibrium. A related work is the paper of by

[2]Emilio Calvano et al. "Artificial intelligence, algorithmic pricing, and collusion". In: *American Economic Review* 110.10 (2020), pp. 3267–97.

[3]Zach Brown and Alexander MacKay. "Competition in Pricing Algorithms". In: *Available at SSRN 3485024* (2020).

Salcedo.[4] He considers a model where firms "commit" to an algorithm in the short-run—this partial commitment can again allow for supra-competitive prices to emerge.

Other papers try to understand the effect of firms' understanding of the underlying demand environment on the possibility of collusion. A recent paper is that of Miklós-Thal and Tucker.[5] They study a setting where competitors are unsure about the market fundamentals (in particular, underlying demand). They look to understand how information about the market fundamentals affects the possibility of collusion. Their main tension is whether this assists collusion (competitors know better the collusive prices to aim for), or hinders it (superior information means firms can better evaluate deviations from collusion, and hence may be more tempted to do so). They show conditions under which the latter effect dominates. Subsequent work by O'Connor and Wilson[6] studies a similar tension and points out that superior information also increases the value of collusion, so the net effects may be ambiguous.

3. Algorithms that Do not Observe Competitors' prices

Consider the scenario where firms are not able to observe their competitors' prices and therefore cannot include this information in their pricing algorithms. Instead firms pricing decisions are made using their past prices and their observed demand (i.e., sales and profits). In this scenario it is important to recognize that while firms' algorithm are not incorporating information about competing firms' prices, their observed demand will still be influenced by competitors' pricing strategies.

This setting greatly reduces complexity: it simplifies the problem to that of a single choice variable (own price), and estimating a profit curve which is only a function of the firm's own price. For example, consider a firm running a large field experiment where they vary their price and see the implications on profits.[7] This is commonly referred to as a *misspecified* model, i.e., firms' algorithms are relying on a model of the market that is not fully aligned with the "correct" market model (since it incorrectly omits competitors' prices). The central question of interest is whether it is possible for firms to coordinate on supra-competitive prices via algorithms even when they do not observe each other's prices?

[4]Bruno Salcedo. "Pricing algorithms and tacit collusion". In: *Working paper, Pennsylvania State University* (2015).

[5]Jeanine Miklós-Thal and Catherine Tucker. "Collusion by algorithm: Does better demand prediction facilitate coordination between sellers?" In: *Management Science* 65.4 (2019), pp. 1552–1561.

[6]Jason O'Connor and Nathan Wilson. "Reduced Demand Uncertainty and the Sustainability of Collusion: How AI Could Affect Competition". In: *FTC Bureau of Economics, Working Paper* 341 (2019).

[7]Wang Chi Cheung, David Simchi-Levi, and He Wang. "Dynamic pricing and demand learning with limited price experimentation". In: *Operations Research* 65.6 (2017), pp. 1722–1731; Jean-Pierre Dubé and Sanjog Misra. "Personalized pricing and customer welfare". In: *Available at SSRN 2992257* (2019).

Our focus in this section is to summarize the findings in our recent paper,[8] where firms' experiment[9] using a multi-armed bandit (MAB) algorithm (from reinforcement learning)[10] Consider a simple example where the firm is deciding between two prices: A and B. In a traditional field experiment, for fixed time (say one quarter), every hour the firm would randomly select a price A or B with equal probability. In any given hour the observed profit is an imprecise or noisy estimate of the expected profit for that price, we will refer this the "informational value of the experiment". As opposed to the fixed equal probability between prices, with a MAB every hour the algorithm will decide the price to experiment based on the observed profits so far. Conceptually this is an automated real-time field experiment that learns about potential profit for each price (learning) and set the product's profit-maximizing price (earning). The precise trade off between learning and earning depends on the specific algorithm chosen.

The algorithm we study is called UCB (Upper Confidence Bound) and is guaranteed to achieve the "best" learning rate for any unknown profit function.[11] This is in the class of index algorithms, where in round r, each possible price is assigned an index or score that is comprised of (a) the number of times the price has been charged so far, and (b) profits observed from prior experiments of that price (rounds 1 to $r-1$). The index optimally balancing the need to 'explore' or 'learn' (try out different prices to find the best price for the future) and 'exploit' or 'earn' (charge the best price for current profits). The price charged in round r is simply the price with the highest index. The algorithms observes realized profits, updates the available information, and then proceeds to round $r+1$. Other popular algorithms in this class include the famous Gittins index.[12]

We analyze this setting both theoretically and using a series of market simulations. We find the long run outcome depends on the informational value of the underlying pricing experiments. Price experiments are highly informative if the observed profit measures in any short experiment have very low noise and are close to the true profit. On the other

[8] Karsten T Hansen, Kanishka Misra, and Mallesh M Pai. "Frontiers: Algorithmic Collusion: Supra-competitive Prices via Independent Algorithms". In: *Marketing Science* (2021).

[9] In the Operation Research literature (William L Cooper, Tito Homem-de Mello, and Anton J Kleywegt. "Learning and pricing with models that do not explicitly incorporate competition". In: *Operations research* 63.1 [2015], pp. 86–103)) show a large set of outcomes (include competitive and collusive) are possible if firms repeatedly optimize prices without experimentation.

[10] An "arm" of the MAB in our application refers to a potential price. See the original paper cited above for the technical details and precise assumptions.

[11] Peter Auer. "Using Confidence Bounds for Exploitation-Exploration Trade-offs". In: *Journal of Machine Learning Research* 3 (2002), pp. 397–422.

[12] John C. Gittins. "Bandit Processes and Dynamic Allocation Indices". In: *Journal of the Royal Statistical Society, Series B* 41.2 (1979), pp. 148–177.

hand pricing experiments have low information value, if the observed profit measures have large magnitudes of noise around the true profit. [13]

In our simulations we estimate the long run-market prices when independent firms are running independent pricing algorithms. Figure 1 displays the distribution of these prices, by the informational value of experiments. We see in markets where price experiments have low information value, the resulting long-run prices are statistically indistinguishable from Nash Equilibrium prices, i.e. competitive prices. In the figure we can see the cloud of resultant prices is centered around the competitive price (right chart). Further, in the original paper, we show that the misspecified models achieve nearly first-best (perfectly respond to competitive price) profits for each firm.

However, in markets where price experiments are highly information value, market prices are supra-competitive or above competitive levels. In the figure we can see the cloud of resultant prices is centered around the monopoly price (right chart). To understand the mechanism, notice that in figure 1 price experiments across firms become inadvertently correlated (along the 45-degree line) as the experiments become more information. Since the models underlying the firms' algorithms are misspecified, this results in an omitted variable bias (the omitted variable being the competitor's price), and an upward bias of own price sensitivity. This means that firms believe that markets are less sensitive to prices than they actually are and all firms will therefore increase prices. The absence of sufficiently large demand shocks to force independent experimentation then results in this being self-reinforcing. Both sellers's algorithms then settle on high prices, even though this is neither the equilibrium of the underlying game, or indeed, even the best response given the competitor's strategy. When demand shocks are small and the sellers' algorithms are close to deterministic, such "mis-learning" may occur and the competing firms may settle down on charging collusive prices.

[Figure 1 about here.]

A different, and intriguing mechanism is proposed by Harrington.[14] This considers a slightly different setting than ones considered before: in his model, firms are unable to design their own algorithms, and instead there is a third-party provider of such algorithms. He envisages a setting where demand depends on a "high-frequency" variable that is not observable to firms doing their own pricing. However, the third-party's algorithm is able to condition price on this variable. Firms therefore must choose between pricing on their own (and being unable to tailor their price to the realization of this demand shock) and using the third-party's algorithm. He shows that a third party designer, taking into account

[13] In terms of learning rates, with large noise firms need many pricing experiments to learn true profit, however with small noise firms need few pricing experiments to learn true profits.

[14] Joseph E Harrington. "Third Party Pricing Algorithms and the Intensity of Competition". In: *Available at SSRN 3723997* (2020).

the possibility that multiple competitors are using its pricing algorithm, will effectively design a "collusive" algorithm to maximize their joint profit; i.e. the algorithm will price less competitively than if it was designed for a single firm. In his basic model, whether or not a firm uses the algorithm is exogenous—the third party chooses to maximize joint profits as a way to maximize its own (unmodeled) ability to extract e.g. licensing fees from the firms. In an extension, he considers a setting where the third party chooses both an algorithm and a licensing fee, and then firms endogenously choose whether to adopt the third party's algorithm or price independently. He shows that a similar effect carries through in this richer model. More broadly, regulators need to be aware of the possibility of only a few market data/ analytics/ pricing intelligence providers supplying to firms in an ostensibly competitive scenario.

4. Conclusions and Policy Implications

In this review article we summarize the recent research that show mechanisms by which firms using Artificial Intelligence algorithms for pricing can lead to non-competitive outcomes. The current literature relies on results from theory or stylized simulated markets. Nevertheless these proposed mechanisms appear prima facie plausible, and ones that policy makers and regulators should be aware of as they investigate markets where prices are set algorithmically.

At a practical level, there remain questions about what kinds of markets may be susceptible to such collusion. The existing literature points to two kinds of markets. The first are stable markets, i.e. ones where demand is a relatively predictable function of the prices offered. The stability of the underlying environment facilitates (algorithmic) learning: of the underlying demand environment, and/ or learning about competing algorithms. An early work in the area is by Assad et. al.[15] who study prices in German gasoline stations before and after they adopted pricing algorithms. They show that prices (and margins) increased after adoption in competitive markets, with no change in monopoly markets. A second kind, as identified by Harrington, are markets where the demand is a predictable function of additional high frequency information firms may not have access to. An example of such a market may be hotels, where local franchisees lack the expertise/ market data to accurately predict demand (e.g. they are unable to account for increased/ decreased demand from events, conventions, etc), but a sufficiently data-rich third party can. Evidence of asymmetry in information about market demand conditions in hotels is provided by Leisten.[16]

[15] Stephanie Assad et al. "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market". In: *CESifo Working Paper* (2020).

[16] Matthew Leisten. "Informational Differences Among Rival Firms: Evidence from Hotel Pricing". 2020.

We hope this article highlights the need for competition law to adapt with the rapid adoption of algorithm pricing. Several other recent articles have also raised similar or related antitrust/ policy concerns: see for example, the review article of Harrington[17] on how competition law should adapt, or the discussion of issues in Calvano.[18]

[17]Joseph E Harrington. "Developing Competition Law for Collusion by Autonomous Artificial Agents". In: *Journal of Competition Law & Economics* 14.3 (2018), pp. 331–363.

[18]Emilio Calvano et al. "Protecting consumers from collusive prices due to AI". in: 370.6520 (2020), pp. 1040–1042.

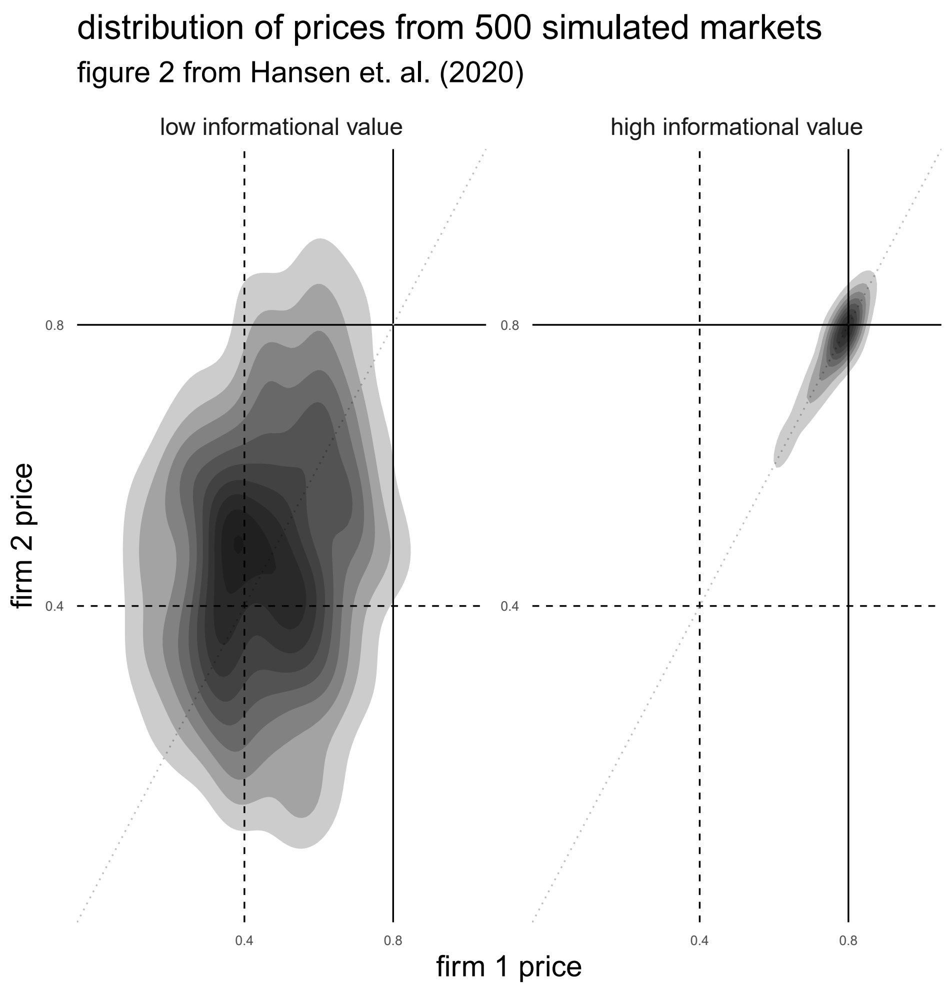

FIGURE 1. 2D Density plot of the distribution of prices for the two firms (x-axis is firm 1 price, y-axis is firm 1 price) after 2 million pricing rounds in 500 simulated markets. Darker colors represent the more observed prices and lighter colors represent fewer observed prices. Each chart represents a market setting described by the informational value of price experiments. The dashed lines reflect the competitive equilibrium prices; the solid lines reflect monopoly prices or collusive prices. The light gray dotted line presents the 45-degree line. [See original paper for more details, SNR = 1/10 and SNR = 10 shown here]