AI, ALGORITHMIC PRICING, AND COLLUSION

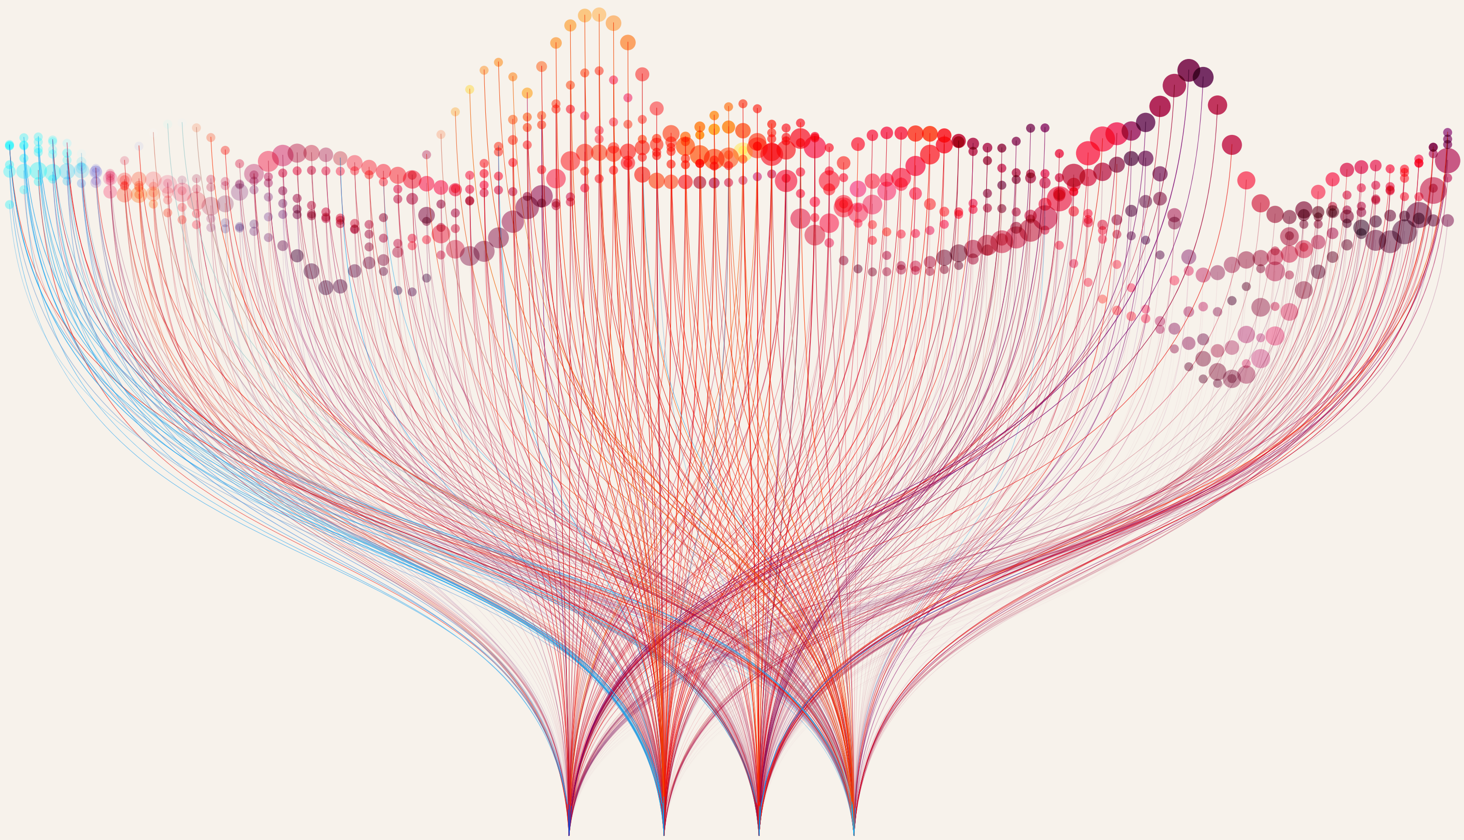

BY JEANINE MIKLÓS-THAL & CATHERINE TUCKER[1]

1 Jeanine Miklós-Thal is the Fred H. Gowen Professor of Economics and Management at the Simon Business School, University of Rochester, and a CEPR research fellow. Catherine Tucker is the Sloan Distinguished Professor of Management at MIT Sloan and an NBER research associate. She has consulted for many technology companies that are detailed at https://mitmgmtfaculty.mit.edu/cetucker/disclosure/.

CPI ANTITRUST CHRONICLE February 2024

AI, ALGORITHMIC PRICING, AND COLLUSION

By Jeanine Miklós-Thal & Catherine Tucker

Advances in artificial intelligence and machine learning and their application to pricing decisions have led to concerns that such innovation could lead to higher prices for consumers. In particular, there are worries that such tools will make it easier for firms within an industry to sustain a collusive outcome. However, we argue that when a close examination of the actual pricing algorithms commonly used by firms, suggests these fears may be overblown. Indeed, in recent research we find that better forecasting of demand may undermine firms' ability to sustain collusive prices. However, we also present evidence suggesting that there are some types of pricing algorithms that may be more problematic - such as ones that allow competitors to simply choose to peg their price to a competitor or a market price. In general, this highlights that the type of pricing algorithm matters a great deal for whether or not there are competition concerns, and that we need to be careful to understand exactly how a specific pricing algorithm works.

I. INTRODUCTION

The rising use of artificial intelligence (AI) and algorithms to enhance pricing by firms has led to concerns among scholars and policymakers around the globe that algorithmic pricing could create new opportunities for collusion.[2] While the adoption of AI and algorithms promises greater efficiencies in price management, a potential concern is that algorithms could also make it easier for sellers to achieve supra-competitive price levels, be it through an explicit illegal agreement or through tacit collusion.

In this article, we discuss this concern based on our academic work in this area. We argue that in fact, building on key insights from the economics of artificial intelligence and the traditional Industrial Organization literature on collusion, it is not clear that such fears reflect the underlying economics. Instead, our game-theoretic analysis shows that pricing algorithms, which allow firms to incorporate more data and base their pricing on better predictions of market conditions, can make collusive outcomes harder to sustain. We also emphasize that it is important for economists and policymakers who are studying these questions to take into account the actual features of commercial pricing algorithms. Our research suggests that rule-based pricing algorithms that allow firms to peg their prices to those of competitors, through a simple price matching algorithm for instance, are more likely to raise concerns than algorithms that improve sellers' pricing performance thanks to better demand forecasts. By modeling how the pricing algorithms that are actually used by firms work, we can start to understand when algorithms may lead to higher prices or lower prices.

II. ARTIFICIAL INTELLIGENCE AND PRICING ALGORITHMS

The key message of economists who study the economics of artificial intelligence and the economics of algorithms is that these technological advances can be best thought of as a drop in the costs of prediction.[3] Better prediction, in turn, may facilitate automation, for example for driverless cars, but also for pricing.

Algorithmic pricing is a broad term that is used to refer to various forms of automated decision-making in price management. There are two main ways in which the pricing algorithms that are commonly deployed in practice (at least at the current time) work. One version is that the pricing algorithm is essentially a demand forecasting tool enabling sellers to better tailor their pricing decisions to market conditions. This fits squarely into the literature on the economics of artificial intelligence in that the costs of predicting future demand have fallen. The question for policymakers then is how improvements in firms' ability to forecast demand will affect pricing outcomes and firms' ability to coordinate on collusive strategies.

There are, however, pricing algorithms, especially those used by smaller sellers on online marketplaces, that do not actually rely on prediction. Instead, they largely automate a pricing rule or heuristic. Such algorithms are distinct from how digital economists usually think about the underlying economics of artificial intelligence as they do not reflect the opportunity of better prediction. Instead, they automate an existing pricing policy such as matching the price of a competitor, or undercutting that price by a fixed amount, or charging more than that price by a fixed amount. These are the type of pricing algorithms which are actually often mentioned as compelling examples of algorithms that drive up prices - for example, a popular example is the instance of an Amazon seller of a book about evolutionary biology and flies ended up charging $23 million dollars for it as a result of setting the price by an algorithm that attempted to charge just a little bit more than other sellers.[4]

As the discussion of our academic research in the next two sections will show, it is important for policymakers to understand the distinctions between these two different versions of algorithmic pricing and the very different effects they may have for either enhancing or reducing the potential for collusive outcomes.

III. DEMAND PREDICTION AND TACIT COORDINATION

A key feature of pricing algorithms compared to traditional pricing methods is that algorithms can incorporate more information and make frequent price changes in response to this information. Algorithms can ingest a broad variety of information about market conditions, including real time data about external factors like web traffic or weather conditions, and data on supply conditions like raw material prices.

2 See, for instance, https://www.gov.uk/government/publications/algorithms-how-they-can-reduce-competition-and-harm-consumers/algorithms-how-they-can-reduce-competition-and-harm-consumers#theories-of-harm, or https://www.oecd.org/competition/algorithmic-competition.htm#june.

3 Agrawal A., Gans J. & A. Goldfarb (2018), *Prediction Machines: The Simple Economics of Artificial Intelligence*, Harvard Business Press, Boston.

4 https://www.wired.com/2011/04/amazon-flies-24-million/.

This means that sellers can use pricing algorithms to try and optimize prices based on forecasts of demand, rather than simply retrospectively analyzing historical data, and calculating price sensitivities. This view of pricing algorithms as improvements in firms' ability to forecast demand and adjust prices accordingly echoes the key message of economists studying artificial intelligence and algorithms that these technological advances can be analyzed as reductions in the costs of prediction, as well as developments in industry emphasizing data, machine learning, and algorithms as tools to improve pricing performance thanks to better demand forecasts.

Motivated by these developments, our published academic work in Miklós-Thal and Tucker (2019) studies how improvements in sellers' ability to predict demand affects the sustainability of collusive strategies.[5] To address this question, we analyze a repeated game of price competition between competing sellers that builds on classic game-theoretic models in the Industrial Organization literature. Demand varies over time due to stochastic fluctuations in market conditions, and in each period, the sellers obtain data that is useful to predict demand. Importantly, the sellers' pricing strategies can depend on their demand forecasts, in addition to the history of past prices. Price changes can thus be triggered by changes in the forecasted demand or by changes in a competitor's price.

Our main question is how improvements in the quality of the sellers' demand forecasts, as they switch from traditional pricing methods to pricing algorithms and these algorithms improve thanks to technological advances, affect the effectiveness of coordination on supra-competitive prices. Do more accurate demand forecasts make it easier or harder for sellers to coordinate on strategies that lead to high prices?

Improvements in the quality of demand predictions have two fundamental effects in our model. First, better demand predictions allow sellers to better tailor prices to demand conditions. The sellers can implement dynamic pricing strategies more effectively when their demand predictions are more accurate. This first effect raises the profits of sellers.

The second fundamental effect of improvement in the quality of demand predictions in our model is on the sellers' incentives to deviate from collusive prices. More accurate demand predictions raise each seller's temptation to undercut price and steal demand from its competitors at times when demand is predicted to be high, and this second effect can make coordination less effective. Consider coordination on the monopoly prices that maximize the sellers' expected joint profit given the demand forecasts in every period of the repeated game. For sufficiently high levels of seller patience, the threat of a breakdown of coordination following a deviation ensures that sellers have no incentive to undercut price. For intermediate levels of seller patience, however, coordination on the monopoly prices breaks down if sellers' demand forecasts become too accurate, because each seller's temptation to undercut price now becomes too strong in periods of better-predicted high demand. To deter price undercutting and still sustain some level of coordination, the collusive price in periods with a high predicted demand would therefore need to be set below the monopoly level in these cases, which makes coordination less profitable.

What, then, is the impact of better demand prediction on consumers? The impact on consumer surplus of the first effect is ambiguous. Although prices rise in periods with favorable demand forecasts as prediction quality improves, prices fall in periods with unfavorable demand forecasts. Overall, consumers can therefore be made better or worse off as a result of this first effect. In markets where sellers would default towards high prices in the absence of accurate demand forecasts, consumers would benefit from more accurate forecasts, but in markets where sellers would default towards low prices in the absence of accurate demand forecasts, more accurate forecasts would harm consumers. The second effect, however, implies that consumers may benefit from better prediction even in cases where they would be harmed by it due to the first effect. Notably, consumers can benefit from better prediction because it undermines the effectiveness of collusion by raising sellers' temptation to undercut price in better-predicted periods of high demand.

In summary, our game-theoretic analysis suggests that pricing algorithms that rely on more data and better demand predictions than traditional pricing methods can, counterintuitively, make collusion less effective. Because of the increased incentives to undercut price at times of more accurate high demand forecasts, better forecasting algorithms can in fact lead to lower prices and higher consumer welfare.

IV. PRICING RULES AND TACIT COORDINATION

A very different type of pricing algorithm abstracts away from advances in AI, machine learning, and forecasting. Instead, these algorithms simply automate the implementation of human heuristics and rules for how to set prices. A commonplace use of such tools are on online marketplaces, as smaller sellers otherwise lack the capacity to be constantly monitoring prices and updating their price frequently. For instance, the pricing rule might be to match the lowest competitor price or to undercut it by a pre-set amount. By contrast to machine-learning powered algorithmic

5 Miklós-Thal, J. & C. Tucker (2019), "Collusion by Algorithm: Does Better Demand Prediction Facilitate Coordination between Sellers?" *Management Science*, Vol. 65(4), pp. 1552-1561. https://pubsonline.informs.org/doi/10.1287/mnsc.2019.3287.

software, rather than ceding pricing decisions to a black-box algorithm, pricing managers retain substantial control. Human managers in firms that adopt rule-based pricing algorithms often choose which pricing rules to use, set the parameters of the rules in the software, and can change the rules and parameters over time if they wish.

To understand the implications of rule-based pricing automation with human control, we (together with Egor Kudriavtcev, University of Rochester) recently ran a small-scale experiment with MBA students specialized in pricing at the Simon Business School, University of Rochester. Our goal was to understand whether price levels differ between markets in which sellers have access to simple rule-based pricing algorithms and markets in which sellers are confined to setting prices in the traditional manual way.

The students in the experiment played a real-time market simulation over a period of two weeks on their own devices. In the simulation, each market featured three sellers, all of them human subjects. Sellers' demands were determined by a logit model that was unknown to the subjects. The subjects also didn't know the identities of the other sellers in the markets in which they were selling, which was important to rule out explicit collusion involving communication between competing sellers.

Each seller was active on two markets simultaneously. First, each seller participated in a so-called manual pricing market. These were markets in which the sellers could set prices in real time, but all prices had to be entered manually. Second, each seller also participated in an automated pricing market, in which all sellers had the option of using simple pricing rules in addition to manual pricing. Importantly, the sellers could change their decisions at any time, and as frequently, or as infrequently, as they wanted.

The pricing rules available to sellers in the automated pricing markets were modeled after pricing tools offered to sellers by online marketplaces as well as third-party repricing software tools available to online marketplace sellers. The seller chooses (i) a reference price, which could be either the lowest competitor's price or the highest competitor's price, and (ii) a rule, which could be match the reference price, stay below it by some fixed amount X, or stay above by it some fixed amount X, and (iii) a minimum price and a maximum prices that limits the range of prices that the rule can lead to.

Our first finding was that the sellers were successful at achieving some level of price coordination both in markets where all sellers used traditional manual pricing and in markets where sellers could use pricing rules. The average price charged by sellers was significantly above the theoretical benchmark of competition (static Nash equilibrium) in both types of markets.

Importantly, however, sellers were significantly more successful at achieving coordination on high prices in markets where they could use pricing rules than in the manual pricing markets. Average prices were closer to the price levels that maximize the sellers' joint profits in the market where sellers were allowed to use rule-based pricing automation than in the markets where they couldn't. Moreover, the difference between the average prices in the two types of markets was larger in the second week of the experiment than in the first week of the experiment, suggesting that sellers learned how to use pricing rules to their advantage over time.

Comparing the average prices across markets that differed in the modal number of sellers who chose to use pricing rules, we find that the automated markets in which the sellers were most successful at achieving high prices typically involved two or three of the three sellers actually using pricing rules. These differences were again more pronounced in the second week of the experiment than in the first week.

Our experiment also reveals some insights into the types of pricing algorithm that sellers who were given the choice between three different classes of pricing rules (matching, undercutting, and overcutting) and manual pricing chose to adopt. The most commonly chosen pricing rule was "match the lowest price", followed by "undercut the lowest price", but other pricing rules such as "match the highest price" were used by a non-negligible share of sellers as well. Moreover, a fair share of sellers, around 35%, chose not to use a pricing rule but to set prices manually even in the market where they had access to rule-based pricing automation.

Our main takeaway from this experiment is that sellers who were given rule-based pricing algorithms did successfully use these algorithms to achieve prices and profits closer to collusive levels. It is important to note that many combinations of pricing rules adopted by competing sellers could have led to highly competitive prices, for instance, all three sellers in a market adopting "undercut the lowest price" would mechanically lead to a race to the bottom. However, that is not what we witnessed in most of the automated pricing markets in our experiments. Sellers active in the same market used combinations of pricing rules that led the entire market to higher prices, and they also made clever use of the minimum price restrictions in the rules to prevent prices from racing to the bottom. In many markets, sellers used pricing rules in ways that reduced incentives to undercut high prices, because sellers understood that their competitors would follow price cuts, but that at the same time strengthened incentive to raise price starting from low levels, because again sellers understood that competitors would follow their price changes.

The results of our experiment, although small in scale, suggest that certain types of pricing algorithms are more likely to raise competition concerns than others. Policymakers should be wary of pricing algorithms whose primary input are competitors' prices and in which a seller's own prices follow competitor's prices.

V. CONCLUSION

In summary, our work highlights that when policy makers and people interested in antitrust and competition policy express general worries about algorithmic collusion, they should pause and ask themselves "What algorithm?". Our work highlights the extent to which the nature of the algorithm is fundamental for assessing both risks and benefits of automated pricing for competition policy. On the one hand, our published research in Miklós-Thal and Tucker (2019) shows that demand-forecasting tools which reflect advances in AI and machine learning can actually make collusive prices harder to sustain. This is because more accurate demand forecasts can raise each firm's temptation to cut price in order to steal business from its competitors. By contrast, if the algorithm is simply an automation of human decision-making rules referencing rivals' prices, then in separate work we show that can be problematic. Even if empirically many sellers choose to use such an algorithm to undercut the lowest competitor price, the adoption of rule-based pricing automation by multiple competing sellers can lead to prices that are supracompetitive.

Therefore, next time you hear a universal statement being made about the potential for algorithmic pricing collusion just ask yourself: What is the algorithm actually doing?