Collusion by Algorithm: Does Better Demand Prediction Facilitate Coordination Between Sellers?

As Published: 10.1287/MNSC.2019.3287

Publisher: Institute for Operations Research and the Management Sciences (INFORMS)

Persistent URL: https://hdl.handle.net/1721.1/134402

Version: Original manuscript: author's manuscript prior to formal peer review

Collusion by Algorithm: Does Better Demand Prediction Facilitate Coordination Between Sellers?*

Jeanine Miklós-Thal[†] Catherine Tucker[‡]

October 5, 2018

Abstract

We build a game-theoretic model to examine how better demand forecasting due to algorithms, machine learning and artificial intelligence affects the sustainability of collusion in an industry. We find that while better forecasting allows colluding firms to better tailor prices to demand conditions, it also increases each firm's temptation to deviate to a lower price in time periods of high predicted demand. Overall, our research suggests that, despite concerns expressed by policymakers, better forecasting and algorithms can lead to lower prices and higher consumer surplus.

*We thank Tony Ke for insightful comments.

[†]Simon Business School, University of Rochester; e-mail: jeanine.miklos-thal@simon.rochester.edu.

[‡]MIT Sloan School of Management; e-mail: cetucker@mit.edu.

1 Introduction

The digital transformation of the economy has allowed the development of algorithms which parse large amounts of consumer data. This means that for key firm decisions such as pricing and operations, human inputs may no longer be needed but instead can be automated by machine learning. One concern, however, is that the use of algorithms for pricing by businesses could have implications for consumer welfare especially if it facilitates collusive outcomes. And indeed, this has recently become a new frontier issue in competition policy and anti-trust regulation. For example, in US vs David Topkins, a poster seller on Amazon was accused of conspiring with other poster sellers to use algorithms, for which he wrote computer code, to coordinate their pricing decisions. However, as of yet there has been little research into whether or how algorithms affect collusive outcomes, a gap that this paper attempts to redress.

We set up a theoretical model that captures one key feature of the use of algorithms for pricing: improvements in demand forecasting. This echoes recent work in economics which emphasizes that machine learning, artificial intelligence (AI) and the spread of algorithms can be analyzed as a fall in costs of prediction technology (Agrawal et al., 2018). And indeed, while much early research on algorithms in operations has focused on the challenge of optimally pricing a fixed inventory (Gallego and Van Ryzin, 1994; Weatherford and Bodily, 1992), recent research has emphasized the role of machine learning techniques in predicting demand in order to optimize pricing algorithms (Ferreira et al., 2015). This academic work is paralleled by developments in industry (Chase Jr, 2013; Feng and Shanthikumar, 2017) that emphasize 'big data' and machine learning as tools that allow firms to improve their pricing performance due to better demand forecasts.[1]

This means that firms now can use pricing algorithms which rather than simply retrospectively looking at historical data and calculating price sensitivities, instead try to optimize prices based on forecasts of demand. For example, regarding its 'Cortana Intelligence Suite', Microsoft states that its artificial intelligence solution "will enable companies to ingest historical transaction data, predict future demand, and obtain optimal pricing recommendations on a regular

[1]Other work has developed algorithms that better adjust to shifts in competitor pricing (Fisher et al., 2017), emphasizing the ease in the new data-driven economy of monitoring and tracking competitors' pricing via algorithmic means.

basis.[2]

Motivated by these developments, we analyze how improvements in sellers' ability to predict demand and price accordingly affect cartel pricing strategies in a repeated game of price competition with time-varying demand. Though our model is highly stylized, it has the advantage of being built upon classic models of drivers of collusive behavior in economics set out by Rotemberg and Saloner (1986); Haltiwanger and Harrington Jr (1991). This allows a direct comparison about how a digital setting characterized by algorithms and better demand forecasting might affect some of the key drivers of collusion.

We find that better prediction affects firms' ability to collude through two channels. First, it improves the cartel's ability to predict the joint profit-maximizing price in each period, which raises the maximum level of profit that the cartel can attain if collusion is successful towards the monopoly level. The impact on consumers of this first effect is *a priori* ambiguous. While the price that maximizes the firms' expected joint profits rises in periods with favorable demand signals as prediction quality improves, it falls in periods with unfavorable signals. Overall, consumers can thus be made better or worse off as a result of this first effect. Second, better prediction also affects the firms' incentives to deviate from a collusive strategy because it increases each firm's temptation to undercut price in periods when consumers are predicted to be willing to pay high prices. To deter such deviations, the cartel may need to set below-monopoly prices in periods with a high predicted demand – similar to Rotemberg and Saloner (1986)'s prediction of "price wars" during booms. The increased incentives to cut price in better predicted periods of high demand imply that consumers can benefit from better prediction even in cases where they would be harmed by it were the firms able to sustain collusion on the prices that maximize their expected joint profits in every period. In summary, our research shows that the effects of improved prediction on firms' ability to collude are nuanced and that better demand forecasting can in fact lead to lower prices and higher consumer welfare.

This research is related to three literatures. First, the broader point that better information can sometimes hurt cooperation has also been made in a number of recent papers in the economics literature, albeit in different contexts. In a model of collusion where the cartel members

[2] See `https://gallery.azure.ai/Solution/Demand-Forecasting-and-Price-Optimization/` as an example

divide the market between themselves, Sugaya and Wolitzky (2018) argue that better information about *past* behavior in another firm's market may allow a firm to better predict future demand conditions in that market, which in turns helps the firm to tailor deviations – in which it deviates by entering its rival's market – to demand conditions. And Hochman and Segev (2010) show that better information about current demand conditions can reduce welfare in a model of self-enforcing international trade agreements. Our contribution to this literature is to explicitly analyze the effects of improved demand prediction on profits and consumer welfare in a simple model of price collusion with perfect monitoring and a single market.

The second literature we contribute to is the nascent literature on digital technologies, algorithms and pricing. Borenstein (2004) provides a discussion of one of the first cases where the use of digital technologies were part of the theory of collusion. Much of this discussion has been led by legal scholars such as Ezrachi and Stucke (2017) who propose that algorithms pose a deep threat to competition. Our work, however, suggests that in fact in some cases the use of algorithms may actually reduce the harm that arises from collusion. Salcedo (2015) provides the only other paper we are aware of which takes an economics approach to this question. Unlike our model, their model of dynamic competition predicts that collusion is unavoidable when firms use a pricing algorithm. However, our model provides a useful counterpoint to this conclusion in emphasizing the extent to which better prediction may influence the effectiveness of a collusive agreement, by incorporating a key new feature of the digital environment and algorithms into a more traditional model of tacit collusion from IO.

The third literature we contribute to is the broader literature about the effects of algorithms for consumer welfare. Mullainathan and Spiess (2017) provides a good overview of this nascent literature, and highlights work such as Kleinberg et al. (2017) which shows that using an algorithm to help guide decisions regarding bail, can help relative to a counterfactual where the judge's judgement could be clouded by the time of day or other external factors. Similarly, Cowgill (2017) shows that algorithmically-based hiring decisions may be less 'biased' than human decision making.[3] Though the economics literature has generally been optimistic about

[3]There is also a complementary literature that looks at how machine learning can be used to tackle policy problems. For example, Ullrich and Ribers (2018) show how consumer welfare can be improved by the use of machine learning to reduce the over-prescription of antibiotics

better algorithms, we study a context where there has been less optimism about algorithms - that is, their use by firms for pricing purposes. Here, there is a perceived risk that algorithms could improve the quality of firm pricing in a way that can be detrimental to consumer welfare. These concerns were expressed in recent policy initiatives such as OECD's 'Big data: Bringing competition policy to the digital era[4]' and the FTC commentary on 'Algorithms and Collusion.[5]' Our results present somewhat reassuring results for antitrust authorities who are worried about the implications for anti-competitive and collusive behavior of the digital environment.

2 The Model

We consider an infinitely repeated game. Two firms ($i = 1, 2$) set prices simultaneously in every period $t = 1, 2, 3, ...$ The goods are perfect substitutes and are produced at constant marginal cost $c \geq 0$. The firms discount future profits at a common discount factor $\delta < 1$.

In each period t, there is mass one of potential buyers with homogeneous willingness to pay v_t. We denote the prices in period t by (p_{t1}, p_{t2}). If $p_{t1} = p_{t2} \leq v_t$, the firms split the total demand evenly, so each has a demand of one half. If firm i undercuts its rival (i.e., $p_{ti} < p_{t-i}$) and $p_{ti} \leq v_t$, all buyers purchase from firm i. At the end of each period, the firms observe each other's prices, making the game one of perfect monitoring.[6]

Consumer willingness to pay is stochastic. There are two possible states of nature in each period. With probability one half, willingness to pay is high: $v_t = \overline{v}$. With probability one half, willingness to pay is low: $v_t = \underline{v}$, where $c < \underline{v} < \overline{v}$. The realizations of willingness to pay are independently and identically distributed (i.i.d.) over time.[7]

Prior to setting prices in each period t, the firms obtain a (common) signal $s_t \in \{\underline{v}, \overline{v}\}$ about

[4]See `http://www.oecd.org/competition/big-data-bringing-competition-policy-to-the-digital-era.htm`

[5]`https://www.ftc.gov/system/files/attachments/us-submissions-oecd-other-international-competition-fora/algorithms.pdf`

[6]Perfect monitoring is a realistic assumption in the context of online pricing, where the prices of competing products can easily be scraped, parsed and used as an input into an algorithm. In a model with imperfect monitoring of pricing decisions, better demand prediction would help firms to distinguish between secret price cuts and low demand states. This collusion-enhancing effect of better prediction would operate in addition to the effects that we analyze in this paper.

[7]The qualitative insights of our analysis would remain unchanged if the demand states corresponded to different downward-sloping demand functions rather than different willingness to pay levels.

the state of demand in period t. The signal has precision ρ:

$$\Pr\{s_t = \overline{v} \,|v_t = \overline{v}\} = \Pr\{s_t = \underline{v} \,|v_t = \underline{v}\} = \rho \in \left[\frac{1}{2}, 1\right].$$

The firms update their belief about the state of demand given the observed signal by Bayes' rule. Given that the two demand states are equally likely ex ante, $\Pr\{v_t = \overline{v} \,|s_t = \overline{v}\} = \Pr\{v_t = \underline{v} \,|s_t = \underline{v}\} = \rho$ and $\Pr\{v_t = \overline{v} \,|s_t = \underline{v}\} = \Pr\{v_t = \underline{v} \,|s_t = \overline{v}\} = 1 - \rho$. The signal precision ρ will be interpreted as the *prediction ability* of the pricing algorithm. As ρ rises, the firms' prediction about buyers' willingness to pay becomes more accurate. For $\rho = \frac{1}{2}$, the signal is uninformative: $\Pr\{v_t = \overline{v} \,|s_t = \overline{v}\} = \Pr\{v_t = \overline{v}\} = \frac{1}{2}$. For $\rho = 1$, the signal is fully diagnostic of the state of the world: $\Pr\{v_t = \overline{v} \,|s_t = \overline{v}\} = \Pr\{v_t = \underline{v} \,|s_t = \underline{v}\} = 1$

We will focus on cases in which the *uninformed monopoly price*, that is, the price that maximizes the firms' expected joint profit given the prior belief, is equal to $\underline{v}$. Formally, we assume that

$$\frac{1}{2}(\overline{v} - c) < \underline{v} - c. \tag{1}$$

This assumption biases our analysis *against* finding a positive effect of improved prediction ability on consumer surplus when the firms collude. If the uninformed monopoly price were $\overline{v}$ (i.e., the inequality in (1) were reversed), then intuitively greater precision could benefit consumers by inducing the cartel to price at $\underline{v}$ (rather than at $\overline{v}$) after an unfavorable signal. In contrast, when the uninformed monopoly price is $\underline{v}$, greater precision intuitively harms consumers because it may induce the firms to feel confident enough about the demand state to set a price above $\underline{v}$ after a favorable signal.

Indeed, there exists a signal precision threshold such that the monopoly price is signal-dependent if and only if signal precision lies above the threshold. After an unfavorable signal, $\Pr\{v_t = \overline{v} \,|s_t = \underline{v}\} = 1 - \rho \leq \frac{1}{2}$, hence (1) implies that a price of $\underline{v}$ maximizes the expected joint profit of the firms in period t. After a favorable signal, $\Pr\{v_t = \overline{v} \,|s_t = \overline{v}\} = \rho$, hence setting a price of $\overline{v}$ rather than $\underline{v}$ yields higher joint profits if

$$\rho(\overline{v} - c) > \underline{v} - c,$$

which is equivalent to

$$\rho > \rho^* \equiv \frac{\underline{v} - c}{\overline{v} - c}.$$

Thus, if $\rho > \rho^*$, the outcome that maximizes expected industry profits in each period is for both firms to charge $\overline{v}$ in periods with a favorable signal and $\underline{v}$ in periods with an unfavorable signal. If $\rho < \rho^*$, expected industry profits are maximized if both firms charge $\underline{v}$ in all periods regardless of the signal.[8] We will refer to the case $\rho > \rho^*$ as "high prediction ability" and to the case $\rho \leq \rho^*$ as "low prediction ability."

In what follows, we will analyze the sustainability of collusion at different levels of signal precision. We are interested in the subgame-perfect equilibria that yield the highest repeated-game payoff per firm among all pure-strategy symmetric equilibria, and will refer to such an equilibrium as a *most cooperative equilibrium.*

3 Analysis

Regardless of the firms' prediction ability ρ, the scope for collusion in our game is maximized when the firms use grim trigger strategies à la Friedman (1971), where any deviation triggers play of the static Nash equilibrium in which both firms charge c and earn zero profits in all future periods. We therefore assume throughout that collusion is sustained by grim trigger strategies.

Low prediction ability ($\rho \leq \rho^*$)

In cases of low prediction ability, the firms' joint profits are maximized when both charge $\underline{v}$ in every period. This outcome is sustainable if, after any signal s_t, the short-term profit gain that a firm could obtain from undercutting $\underline{v}$ in order to obtain the entire demand in that period is outweighed by the discounted future loss in profits from being punished, i.e., if

$$\frac{1}{2}\left(\underline{v} - c\right) \leq \frac{\delta}{1-\delta}\frac{\underline{v} - c}{2}. \tag{2}$$

[8] For $\rho = \rho^*$, expected industry profits are the same regardless of whether the firms charge $\overline{v}$ or $\underline{v}$ in periods with a high signal.

The gain from undercutting the collusive price is independent of the signal realization, because consumers are always willing to buy at price $\underline{v}$. The non-deviation condition (2) is equivalent to

$$\delta \geq \frac{1}{2}, \tag{3}$$

the standard discount factor threshold for collusion in a symmetric Bertrand game with two firms. If $\delta < \frac{1}{2}$, both firms make zero profits in any subgame-perfect equilibrium. If $\delta \geq \frac{1}{2}$, collusion on the price $\underline{v}$ in all periods is sustainable for any precision level ρ, and is optimal for the firms if $\rho < \rho^*$. The expected discounted profit per firm on the equilibrium path in this case is

$$\underline{V} \equiv \frac{1}{2}\frac{\underline{v} - c}{1 - \delta}. \tag{4}$$

Consumer obtain a surplus of $\overline{v} - \underline{v}$ in periods with a high willingness to pay, yielding them an expected discounted consumer surplus on the equilibrium path of

$$\frac{1}{2}\frac{\overline{v} - \underline{v}}{1 - \delta}. \tag{5}$$

Both firm profits and consumer surplus are constant in ρ for $\rho \in \left[\frac{1}{2}, \rho^*\right)$.

High prediction ability ($\rho > \rho^*$)

In cases of high prediction ability, the firms' expected joint profits are maximized when the price in each period follows the signal in that period. When the signal-dependent monopoly prices are sustainable, the expected discounted profit per firm on the equilibrium path is

$$V_{full} \equiv \frac{1}{4}\frac{\rho\left(\overline{v} - c\right) + \underline{v} - c}{1 - \delta}$$

Profits at the monopoly outcome are now increasing in the signal precision ρ, because a more precise signal improves the firms' ability to tailor their pricing to the state of the world.

The expected short-term profit gain from undercutting the monopoly price is $\frac{1}{2}\rho\left(\overline{v} - c\right)$ after a favorable signal and $\frac{1}{2}\left(\underline{v} - c\right)$ after an unfavorable signal. High prediction ability ($\rho > \rho^*$) means that $\rho\left(\overline{v} - c\right) > \left(\underline{v} - c\right)$, hence the temptation to undercut is stronger after a favorable signal than after an unfavorable signal. Thus, the monopoly outcome is sustainable if and only if

$$\frac{1}{2}\rho\left(\overline{v}-c\right) \leq \delta\frac{1}{4}\frac{\rho\left(\overline{v}-c\right)+\left(\underline{v}-c\right)}{1-\delta},$$

which yields a critical discount factor of

$$\delta^*\left(\rho\right) \equiv \frac{2\rho\left(\overline{v}-c\right)}{3\rho\left(\overline{v}-c\right)+\underline{v}-c}.$$

The critical discount factor $\delta^*\left(\rho\right)$ exceeds $\frac{1}{2}$ because $\rho\left(\overline{v}-c\right) > \left(\underline{v}-c\right)$ for all $\rho > \rho^*$. Moreover, it is increasing in the firms' prediction ability. Although better prediction ability raises the expected monopoly profit, it also increases the expected profit gain from undercutting price after a favorable signal, making collusion on the monopoly prices harder to sustain.

For $\delta \in \left[\frac{1}{2}, \delta^*\left(\rho\right)\right)$, the signal-dependent monopoly prices are unsustainable. In this case, the optimal collusive scheme involves a price below $\overline{v}$ after a favorable signal, in the spirit of Rotemberg and Saloner (1986)'s theory of "price wars" during booms. The most cooperative equilibrium outcome is obtained by solving[9]

$$\max_{\overline{p}} \frac{1}{4}\frac{\rho\left(\overline{p}-c\right)+\underline{v}-c}{1-\delta} \tag{6}$$

subject to

$$\frac{1}{2}\rho\left(\overline{p}-c\right) \quad \leq \quad \delta\frac{1}{4}\frac{\rho\left(\overline{p}-c\right)+\left(\underline{v}-c\right)}{1-\delta}, \tag{7}$$

where $\overline{p} \in \left(\underline{v}, \overline{v}\right)$ is the collusive price in periods with a favorable signal.[10] At the solution of the cartel problem, $\overline{p}$ is chosen at the highest level compatible with the non-deviation constraint (7), which yields

$$\overline{p}-c = \frac{\delta}{2-3\delta}\frac{\underline{v}-c}{\rho}. \tag{8}$$

Since $\frac{\delta}{2-3\delta} \geq 1$ for $\delta \in \left[\frac{1}{2}, \delta^*\left(\rho\right)\right)$, we have $\rho\left(\overline{p}-c\right) > \underline{v}-c$ at the optimal $\overline{p}$, hence the collusive

[9]Setting a price different from $\underline{v}$ in periods with an unfavorable signal would reduce profits without helping to sustain collusion. To streamline the exposition, we therefore assume from the outset that the firms charge $\underline{v}$ in periods with an unfavorable signal and focus on the no-deviation constraint in periods with a favorable signal.

[10]We will confirm shortly that the solution to this maximization problem always dominates charging $\underline{v}$ in every period for the firms.

scheme in which the firms set $\overline{p}$, as given in (8), after a favorable signal (and $\underline{v}$ after an unfavorable signal) yields higher profits than a scheme in which the firms charge $\underline{v}$ in every period. The most cooperative equilibrium thus involves signal-dependent prices on the equilibrium path for any $\delta \geq \frac{1}{2}$ when the prediction ability is high, although the monopoly prices can only be sustained for $\delta \geq \delta^* (\rho) > \frac{1}{2}$.

When the monopoly prices cannot be sustained, that is, for $\delta \in \left[\frac{1}{2}, \delta^* (\rho)\right)$, the expected discounted profit per firm on the equilibrium path is equal to

$$V_{partial} \equiv \frac{1}{4} \frac{\frac{\delta}{2-3\delta} (\underline{v} - c) + \underline{v} - c}{1 - \delta},$$

which is invariant with respect to ρ. An improvement in signal precision has two effects on the collusive profits in this case. On the hand, it increases the likelihood that consumers indeed buy when the firms set a high price after a favorable signal, because a favorable signal is a better indication of a high willingness to pay. All else equal, this effect predicts that more precise signals lead to greater profits. On the other hand, greater precision also increases the temptation to undercut in periods with a favorable signal, which means that the firms need to set a lower price to deter deviations, i.e., $\overline{p}$ as given in (8) is decreasing in ρ. The two effects exactly offset each other, so that profits are constant in signal precision. This happens because both the objective function and the no-deviation condition in (7) depend only on $\rho (\overline{p} - c)$, not on ρ and $\overline{p}$ individually; thus, the cartel optimally offsets any change in ρ by a change in $\overline{p}$ that keeps $\rho (\overline{p} - c)$ unchanged and (7) binding.

Denote the cartel's optimal price in periods with a favorable signal by $\overline{p}^* \equiv \min \{\overline{p}, \overline{v}\}$, with $\overline{p}$ as given in (8). Consumers obtain a surplus of $(\overline{v} - \overline{p}^*)$ in periods with a favorable signal and a high willingness to pay and a surplus of $(\overline{v} - \underline{v})$ in periods with an unfavorable signal and a high willingness to pay. Given the frequencies of these two events, the expected discounted consumers surplus is

$$\frac{1}{2} \frac{\rho (\overline{v} - \overline{p}^*) + (1 - \rho) (\overline{v} - \underline{v})}{1 - \delta}. \tag{9}$$

When $\delta \geq \delta^* (\rho)$, the first term is zero because $\overline{p}^* = \overline{v}$, hence consumer surplus is decreasing in ρ. When $\delta \in \left[\frac{1}{2}, \delta^* (\rho)\right)$, however, $\overline{p}^* = \overline{p}$ and the marginal impact of ρ on consumer surplus has the sign of

$$-\left(\overline{p}-\underline{v}\right)-\rho\frac{\partial\overline{p}}{\partial\rho}=\underline{v}-c>0 \tag{10}$$

Consumers thus benefit from greater precision when the firms cannot sustain the monopoly prices. Although greater precision leads to fewer periods with a price $\underline{v}$ in spite of a high willingness to pay, it also enables the consumers to obtain higher surplus in periods with a favorable signal, leading to a higher surplus overall.

Effect of signal precision on firms and consumers

To provide a full characterization of how signal precision affects firm profits and consumer surplus in the most cooperative equilibrium as ρ rises from $\frac{1}{2}$ to 1, we need to distinguish between two different ranges of the discount factor.

Case A: $\delta \geq \delta^*(1) = \frac{2(\overline{v}-c)}{3(\overline{v}-c)+\underline{v}-c}$

In this case, the discount factor is high enough that the firms can sustain the outcome that maximizes expected industry profits at all levels of precision. For $\rho \leq \rho^*$, the firms charge $\underline{v}$ in every period, and for $\rho > \rho^*$, the firms charge $\overline{v}$ after a favorable signal and $\underline{v}$ after an unfavorable signal in the most cooperative equilibrium. Firm profits are constant in ρ for $\rho \leq \rho^*$ and increasing in ρ for $\rho > \rho^*$, with a kink at $\rho = \rho^*$. Consumer surplus is equal to $\frac{1}{2}\frac{\overline{v}-\underline{v}}{1-\delta}$ for $\rho \leq \rho^*$, and equal to

$$\frac{1}{2}\frac{(1-\rho)\left(\overline{v}-\underline{v}\right)}{1-\delta}$$

for $\rho > \rho^*$. Overall, consumer surplus is constant in ρ for $\rho \leq \rho^*$, exhibits a discrete downward jump at $\rho = \rho^*$, and is linearly decreasing in ρ for $\rho > \rho^*$. Figure 1 provides an illustration of profit, consumer surplus, and total welfare in case A.

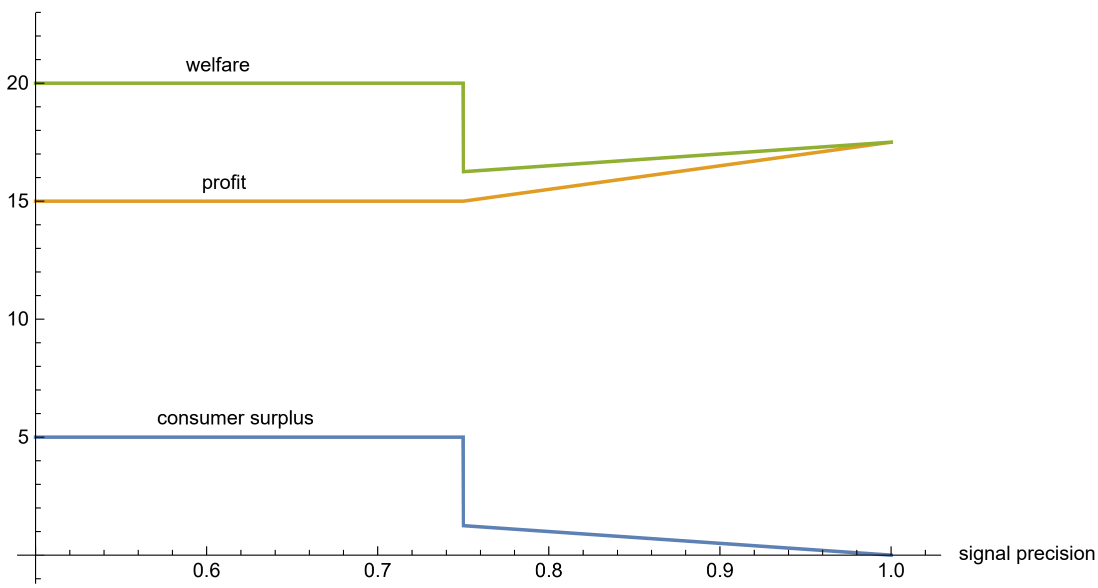

Figure 1: An illustration of profit,consumer surplus, and total welfare in Case A

Case B: $\delta \in \left[\frac{1}{2}, \delta^*(1)\right)$

If $\frac{1}{2} \le \delta < \delta^*(1)$, there exists a unique $\widehat{\rho} \in [\rho^*, 1)$ such that

$$\delta = \delta^*(\widehat{\rho}), \tag{11}$$

which is equivalent to $\widehat{\rho} = \frac{\delta}{2-3\delta}\rho^*$, with $\delta \ge \delta^*(\widehat{\rho})$ if and only if $\rho \le \widehat{\rho}$. The collusive scheme then changes as follows with the belief precision ρ. For $\rho \le \rho^*$, the firms change $\underline{v}$ in every period. For $\rho^* < \rho \le \widehat{\rho}$, the firms charge $\overline{v}$ after a favorable signal and $\underline{v}$ after an unfavorable signal. For $\rho > \widehat{\rho}$, the firms charge $\overline{p}$ as given in (8) after a favorable signal and $\underline{v}$ after an unfavorable signal. Overall, firm profits in the most cooperative equilibrium are first constant, then increasing, and then constant again in ρ. The profit line in Figure 2 provides an illustration.

Consumer surplus behaves non-monotonically in belief precision in Case B. Consumer surplus is constant in ρ for $\rho \le \rho^*$, exhibits a discrete downward jump at $\rho = \rho^*$, and is linearly decreasing in ρ for $\rho \in (\rho^*, \widehat{\rho})$. Importantly, however, for $\rho > \widehat{\rho}$ consumer surplus is increasing in ρ (as implied by equation (10) above). The consumer surplus line in Figure 2 provides an illustration.[11]

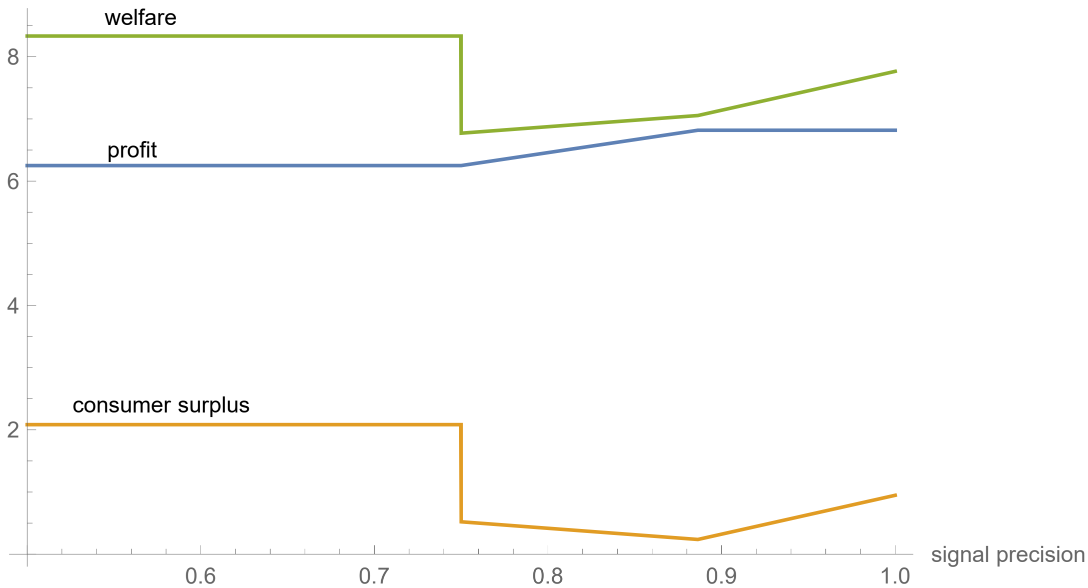

Figure 2: Consumer surplus behaves non-monotonically in signal precision in Case B

[11]Welfare, defined as the sum of profits and consumer surplus, always increases in ρ for $\rho > \rho^*$, because greater precision raises the number of periods with trade when the firms charge signal-dependent prices.

The following proposition summarizes our findings about the effect of the firms' prediction ability on profits and consumer surplus:

Proposition 1 *In a most cooperative equilibrium,*

- *Expected discounted consumer surplus is weakly decreasing in the firms' ability to predict willingness to pay for large discount factors ($\delta \geq \frac{2(\overline{v}-c)}{3(\overline{v}-c)+\underline{v}-c}$)*
- *Expected discounted consumer surplus is non-monotonic in the firms' ability to predict willingness to pay for intermediate discount factors ($\delta \in \left[\frac{1}{2}, \frac{2(\overline{v}-c)}{3(\overline{v}-c)+\underline{v}-c}\right)$).*
- *Expected discounted firm profit is weakly increasing in the firms' ability to predict willingness to pay for any discount factor $\delta \geq \frac{1}{2}$.*

In summary, using a simple setting, our analysis shows that improved prediction ability can benefit consumers, because it raises firms' temptation to undercut price in periods with a high predicted demand, thereby limiting the cartel's ability to set high prices in those periods.

Firms benefit (weakly) from better predictions in the model we have presented, but this need not be true more broadly. If the uninformed monopoly price were $\overline{v}$ (that is, the inequality in (1) were reversed), for instance, then firm profits in the most cooperative equilibrium would be (weakly) decreasing in the prediction ability ρ for intermediate discount factors. As in our baseline model, the temptation to undercut is stronger after a favorable signal than after an unfavorable signal in this case, and that temptation increases as signal precision rises. For intermediate discount factors and high ρ within the range of low prediction ability, the cartel therefore needs to engage in "countercyclical pricing" to sustain collusion, charging *less* than $\overline{v}$ in periods with a *favorable* signal while charging $\overline{v}$ in periods with an *unfavorable* signal. As illustrated in Figure 3, profits thus decrease as the signal becomes more precise in this case, whereas consumer surplus rises. Furthermore, as also shown in Figure 3, consumer surplus exhibits a discrete upward jump at the threshold between low and high signal precision, because the cartel optimally starts to price at $\underline{v}$ after an unfavorable signal once signal precision is sufficiently good. As signal precision improves further, consumer surplus keeps rising for intermediate discount factors, because the optimal price after a favorable signal falls to accommodate the increased

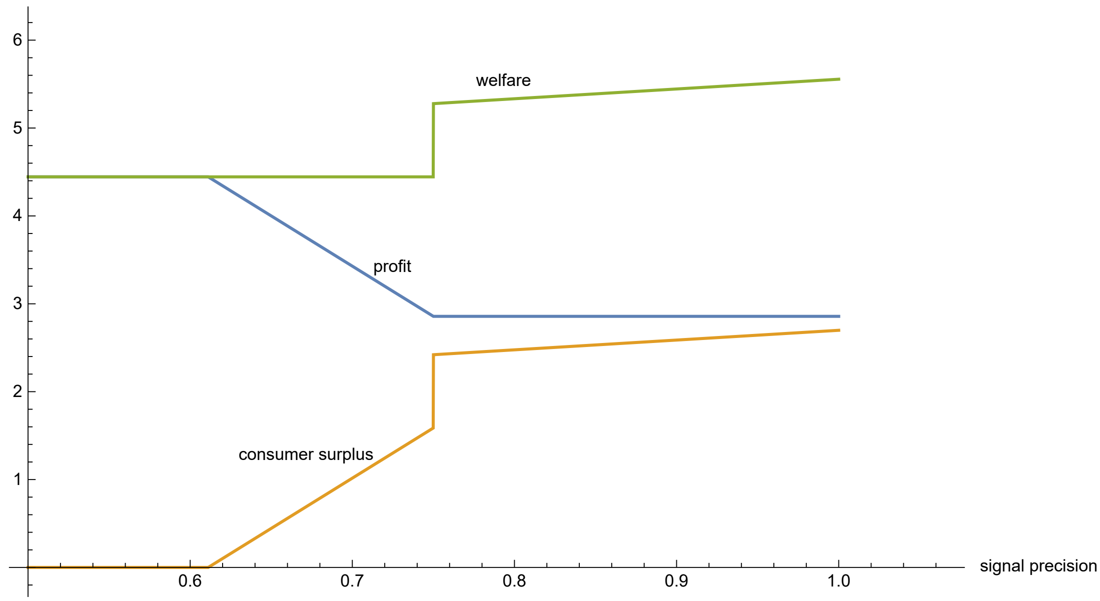

Figure 3: Profit can be decreasing in signal precision when the uninformed monopoly price is high

temptation to deviate in those periods. Overall, consumer surplus is highest when the firms observe a fully informative signal in this scenario.

Our model assumed a low uninformed monopoly price to show that consumers can benefit from higher precision even when firm profits are weakly increasing in precision.

4 Implications

This paper presents a theoretical model that investigates the role of algorithms in potentially facilitating collusive behavior. We build upon classic models of collusion developed in IO, and investigate how they would alter given a key characteristic of a digital economy: the use of algorithms to better predict demand.

Our research suggests that somewhat counterintuitively better forecasting algorithms can lead to lower prices and higher consumer welfare due to increased incentives to cut prices at times of more accurately forecasted high demand. This is important, because in general antitrust authorities have been pessimistic regarding the introduction of algorithms and policy emphasis has been on their role in potentially increasing collusive behavior and raising prices.

There are of course several limitations to this research. First, we do not model frequency of interactions. However, algorithms may facilitate collusion by enabling firm to detect and punish deviations more quickly, although fast arrival of information has also been shown to potentially hinder collusion in games with imperfect monitoring (Sannikov and Skrzypacz, 2007). Second, we do not model hub-and-spoke style collusion with a common algorithm – something which was the focus of recent FTC investigations. Third, this is a theory paper and we do not study any specific empirical context. Notwithstanding these limitations, we believe our focus on algorithms as a forecasting tool is a useful first step for understanding how machine learning and algorithms may effect collusion incentives.

References

Agrawal, A., J. Gans, and A. Goldfarb (2018). *Prediction Machines: The simple economics of artificial intelligence*. Harvard Business Press.

Borenstein, S. (2004). Rapid price communication and coordination: The airline tariff publishing case (1994). *The Antitrust Revolution: Economics, Competition, and Policy 4*.

Chase Jr, C. W. (2013). Using big data to enhance demand-driven forecasting and planning. *The Journal of Business Forecasting 32*(2), 27.

Cowgill, B. (2017). Automating judgement and decision-making: Theory and evidence from résumé screening.

Ezrachi, A. and M. E. Stucke (2017). Artificial intelligence & collusion: When computers inhibit competition. *U. Ill. L. Rev.*, 1775.

Feng, Q. and J. G. Shanthikumar (2017). How research in production and operations management may evolve in the era of big data. *Production and Operations Management*.

Ferreira, K. J., B. H. A. Lee, and D. Simchi-Levi (2015). Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management 18*(1), 69–88.

Fisher, M., S. Gallino, and J. Li (2017). Competition-based dynamic pricing in online retailing: A methodology validated with field experiments. *Management Science 64*(6), 2496–2514.

Friedman, J. W. (1971). A non-cooperative equilibrium for supergames. *The Review of Economic Studies 38*(1), 1–12.

Gallego, G. and G. Van Ryzin (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management science 40*(8), 999–1020.

Haltiwanger, J. and J. E. Harrington Jr (1991). The impact of cyclical demand movements on collusive behavior. *The RAND Journal of Economics*, 89–106.

Hochman, G. and E. Segev (2010). Managed trade with imperfect information. *International Economic Review 51*(1), 187–211.

Kleinberg, J., H. Lakkaraju, J. Leskovec, J. Ludwig, and S. Mullainathan (2017, February). Human decisions and machine predictions. Working Paper 23180, National Bureau of Economic Research.

Mullainathan, S. and J. Spiess (2017, May). Machine learning: An applied econometric approach. *Journal of Economic Perspectives 31*(2), 87–106.

Rotemberg, J. J. and G. Saloner (1986). A supergame-theoretic model of price wars during booms. *The American Economic Review 76*(3), 390–407.

Salcedo, B. (2015). Pricing algorithms and tacit collusion. *Manuscript, Pennsylvania State University*.

Sannikov, Y. and A. Skrzypacz (2007). Impossibility of collusion under imperfect monitoring with flexible production. *American Economic Review 97*(5), 1794–1823.

Sugaya, T. and A. Wolitzky (2018). Maintaining privacy in cartels. *Journal of Political Economy, forthcoming*.

Ullrich, H. and M. Ribers (2018). Battling resistance: using machine prediction to improve antibiotic prescribing.

Weatherford, L. R. and S. E. Bodily (1992). A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking, and pricing. *Operations research 40*(5), 831–844.