

WORKING PAPERS



Reduced Demand Uncertainty and the Sustainability of Collusion: How AI Could Affect Competition

**Jason O'Connor
Nathan E. Wilson**

WORKING PAPER NO. 341

June 2019

FTC Bureau of Economics working papers are preliminary materials circulated to stimulate discussion and critical comment. The analyses and conclusions set forth are those of the authors and do not necessarily reflect the views of other members of the Bureau of Economics, other Commission staff, or the Commission itself. Upon request, single copies of the paper will be provided. References in publications to FTC Bureau of Economics working papers by FTC economists (other than acknowledgment by a writer that he has access to such unpublished materials) should be cleared with the author to protect the tentative character of these papers.

**BUREAU OF ECONOMICS
FEDERAL TRADE COMMISSION
WASHINGTON, DC 20580**

Reduced Demand Uncertainty and the Sustainability of Collusion: How AI Could Affect Competition*

Jason O'Connor

University of Pittsburgh
jason.m.oconnor@gmail.com

Nathan E. Wilson

Federal Trade Commission
nwilson@ftc.gov

June 10, 2019

Abstract

We consider how technologies that eliminate sources of demand uncertainty change the character and prevalence of coordinated conduct. Our results show that mechanisms that reduce firms' uncertainty about the true level of demand have ambiguous welfare implications for consumers and firms alike. An exogenous increase in firms' ability to predict demand may make collusion possible where it was previously unsustainable. However, it also may make collusion impracticable where it had heretofore been possible. The underlying intuition for this ambiguity is that greater clarity about the true state of demand raises the payoffs both to colluding and to cheating. The net effect will depend on a given market's location in a multidimensional parameter space. Our findings on the ambiguous welfare implications of AI in market intelligence applications contribute to the emerging literature on how algorithms and other forms of artificial intelligence may affect competition.

JEL Codes: K12, L13, L40

Keywords: Artificial Intelligence, Uncertainty, Collusion, Price Discrimination, Antitrust

*The views expressed in this article are those of the authors. They do not necessarily represent those of the Federal Trade Commission or any of its Commissioners. We are grateful for comments from Rob Porter, Scott Stern, Yu-Wei Hsieh, Debbie Minehart, Alex MacKay, and attendees of the 2019 International Industrial Organization, DC IO Day, and Georgetown Center for Economic Research Biennial conferences. In addition, we gratefully acknowledge comments and assistance from Dave Balan, Patrick Degraha, Dan Hosken, David Schmidt, and other colleagues at the FTC. The usual caveat applies.

I Introduction

In the real world, competition has always occurred under conditions of at least somewhat imperfect information. Firms generally do not directly observe the true level of demand and all competitively relevant actions by their rivals. However, the recent proliferation of large data sets and the availability of algorithmic tools to analyze them have caused some commentators to conclude that the prevalence of uncertainty may be changing, and that this has potentially baleful implications for competitive intensity. For example, in October 2018, the Assistant Attorney General for antitrust intimated that a price-fixing case involving algorithmic price-setting tools might be brought soon.¹ At roughly the same time, the UK’s Competition and Markets Authority released a white paper describing risk factors associated with the use of data-driven algorithms.²

In this paper, we bring the tools of formal theory to bear on the question of how the use of predictive analytics could alter the incidence and character of collusion. In particular, we use a framework derived from the seminal [Green and Porter \(1984\)](#) model of dynamic competition to consider how changes to demand uncertainty affect the incidence and character of collusion. We specifically consider discrete changes to firms’ uncertainty about the ex ante profitability of choosing different prices. We generally refer to the uncertainty-reducing technology as “AI” in keeping with how the recent literature ([Agrawal et al., 2018](#)) describes the use of algorithms and data to improve “nowcasting” capabilities.

Although clearly a simplification, we believe our modeling framework captures key dynamics of how AI could affect various industries. In particular, markets for intermediate goods are often characterized by bilateral negotiations between buyers and sellers ([Shapiro, 2010](#)). For example, in food service markets, buyers and sellers engage in bilateral negotiations where the prices offered by firms are not observable to their rivals.³ The outcomes of

¹See https://www.broadcastingcable.com/news/delrahim-criminal-case-against-anti-competitive-search-algorithms-coming#disqus_thread.

²See <https://www.gov.uk/government/publications/pricing-algorithms-research-collusion-and-personalised-pricing>).

³See, e.g., <https://www.ftc.gov/system/files/documents/cases/150219syscopt3cmpt.pdf>.

these negotiations will be shifted by expectations about demand conditions, which may be difficult for the seller to precisely observe. An individual seller may have priors about the prices offered by rivals but will not be able to observe the outcome of individual negotiations. Thus, failure to make a sale may reflect an unobserved decline in demand or undercutting by rivals. Sellers will not be able to distinguish these possibilities insofar as they only directly observe their own sales and profits. Algorithms that can reliably collect and analyze data correlated to the true state of demand could reduce some of this uncertainty.

We find that the exogenous adoption of AI by sellers has ambiguous implications for both firms and consumers. On the one hand, reduced demand uncertainty may benefit firms in many cases. By clarifying what the true demand conditions are, AI enables firms to more precisely differentiate rivals' cheating from unobserved negative demand shocks. Furthermore, increased clarity about the true state of demand allows colluding firms to better tailor their prices, increasing the average per period profit earned during periods of collusion. All of these factors may make collusion sustainable when it previously was not, and also may cause the duration of punishment to shrink relative to the pre-AI period (assuming collusion was possible then). On the other hand, once firms have a better knowledge of the true state of demand, they may better time their decision to cheat. All else equal, this pushes equilibrium strategies towards the possibility that collusion might not be sustainable even if it previously had been. We find that the net effect of the pro- and anti-competitive effects depends on where in the parameter space the market is.

The implications of reduced demand uncertainty are also ambiguous for consumers. They clearly benefit if collusion no longer can be sustained, and they unambiguously suffer if collusion becomes sustainable when it was previously not. In the case where collusion is sustainable both before and after the technology is adopted and firms unambiguously benefit from its adoption, consumers *may* not suffer. This is because despite higher average prices during "high" demand periods, sales *may* increase in periods of low demand.

To provide greater insight into how different model parameters interact, we implement

our model numerically in a linear demand setting. In this environment, we find that consumers gain most when the aspect of uncertainty that is reduced is large relative to that which remains. In addition, we find that consumers are more likely to gain, all else equal, when the discount rate is lower. Whenever collusion is sustainable both before and after AI, producer surplus increases while consumer (and total) surplus declines.⁴

Overall, we contribute to the large literature on conditions conducive to collusion and coordinated conduct (Tirole, 1988, Kovacic et al., 2011), particularly with respect to the role played by uncertainty (Robson, 1981, Green and Porter, 1984, Kandori, 1992, Raith, 1996, Athey and Bagwell, 2001). Our work shows that reducing uncertainty has ambiguous effects, making collusion more attractive in some cases, but also sometimes more difficult to sustain.

In addition, our paper contributes to the ongoing debate about how antitrust policy should address the competitive effects of recent developments in data analytics. This literature already contains contributions reflecting a wide variety of opinions (see, e.g., the partial bibliography of Ritter (2017)), but many of the most vocal commentators have suggested that these phenomena threaten consumer welfare. In our view, these critics have focused on two related, but distinct, theories of harm: (1) increased ability to personalize prices will allow firms to better extract consumer surplus (Ezrachi and Stucke, 2016, 2017), and (2) use of algorithms will facilitate collusion (Ezrachi and Stucke, 2016, Mehra, 2015). In the main, we worry that the critics have proceeded to judgement without necessarily establishing a rigorous basis for either theory of harm in the economic or computer science literatures.⁵

By rooting consideration of AI in the game theoretic treatment of uncertainty and coordination, our paper shows that the effects of AI and related technologies are more nuanced

⁴In Section 5.2, we show that if firms are allowed to collude at prices below the monopoly level, consumer surplus may increase even when collusion is sustainable in both the pre-AI and post-AI periods.

⁵Salcedo (2015) is a notable exception in the literature on algorithmic collusion. However, the necessary and sufficient conditions for collusion in the model are restrictive and not obviously consistent with the approach taken by firms actually engaged in algorithmic pricing. Most of the assessments of price discrimination ignore the fact that economic theory indicates that the practice has ambiguous effects, particularly in the presence of competition (Cooper et al., 2004, Carlton and Perloff, 2015). Consistent with this, recent empirical work on the use of algorithmic pricing has found that most consumers benefit (Dubé and Misra, 2017) or that overall welfare substantially increases (Reimers and Shiller, 2018).

than most commentators have heretofore considered. Our results show that even without endogenizing entry, exit, and repositioning on the supply-side, let alone demand-side technological adaptation (Gal and Elkin-Koren, 2016), the implications are varied, and depend heavily on market primitives. This nuanced conclusion stands in contrast to the findings in the policy-oriented literature.

Within the literature on algorithms and competition, our paper fits with other emerging contributions taking a more technical approach (Ittoo and Petit, 2017, Kuhn and Tadelis, 2017, Calzolari et al., 2018, Miklós-Thal and Tucker, Forthcoming). In particular, our work is very similar in motivation to that of Miklós-Thal and Tucker (Forthcoming), who also assume algorithms allow firms to better predict demand rather than directly set prices. However, their model extends the Rotemberg and Saloner (1986) model of collusion where firms rely on a signal of future demand in setting prices, while we modify Green and Porter (1984). Despite the fact that the two foundational models reach different conclusions about when collusion will break down (booms vs. busts), both we and Miklós-Thal and Tucker (Forthcoming) find that AI has ambiguous effects on consumer welfare and profits. That our respective papers reach broadly similar conclusions supports a cautious policy approach on the potential threat of AI to consumer welfare.

The paper is organized as follows. Section II outlines the modeling framework, and then Section III compares the equilibria that will result in the pre- and post-AI states. Section IV presents the results for different parameterizations of the baseline model to further clarify how different parameters affect AI's impact on welfare. Section V briefly discusses the reasonability of thinking of AI adoption as exogenous. Finally, Section VI concludes.

II The Theoretical Model

II.1 Setting

We consider an infinite horizon, discrete-time model of duopolistic competition. The base level of demand is the same in every period. However, it is subject to two random and independent shocks, M and Σ , each period. If either of these shocks occur, they reduce demand. Both are assumed to occur with strictly positive probabilities μ and σ , respectively. The role of AI is to reduce demand uncertainty by assuming that one of the shocks becomes perfectly predictable to firms. For the sake of simplicity, the magnitude of individual shocks is equal, but if both occur then firms are unable to sell their product at any price, and, therefore, make no profits.⁶

The two firms produce homogeneous products, and compete in prices.⁷ Demand slopes down. In the event that one firm has a lower price, it supplies the entire market that is willing to purchase the product at that price. The other firm makes no sales and earns no profits. When the two firms choose the same price, they divide the market equally. Firms observe only their own prices, sales, and profits in each period. They do not directly observe the conduct of their rivals. Thus, if firm i observes that it makes zero profits in one period, it could indicate that both negative demand shocks occurred *or* that its price had been higher than that of firm j . We assume there is no way for a firm to undercut its rival and limit its sales to less than the entirety of the market.

⁶In the context of linear demand, this could either be because each shock is at least half of the intercept for demand or that there is some compounding of demand reduction shocks when both occur.

⁷We believe the assumption of homogeneous products is without loss of generality so long as the extent of differentiation across firms is public and the shocks reduce demand equally for both products. Allowing for differentiation simply changes the profit level of the firms in the different states but does not qualitatively change the welfare effects of introducing AI.

II.2 Pre-AI

Given the nature of competition and the form of demand uncertainty, we model coordination in a manner akin to [Green and Porter \(1984\)](#) in the context of quantity-setting firms and re-formulated for Bertrand competition in [Tirole \(1988, Section 6.7.1\)](#).⁸ Firms' collusive equilibrium strategy takes the following form: The competitors agree ex ante to price at a defined collusive level, p^m (i.e., the monopolist's price given the existence of uncertain demand), until observing a period in which they earn zero profits. When a firm realizes zero profits in some period t , it will price at the competitive level, p^c , for the next T periods before returning to the collusive price in period $t + T + 1$. Because both M and Σ occur with positive probability, price wars (i.e., the punishment phase where both firms price at the competitive level) will occur even if neither firm ever deviates from the collusive strategy.

Assuming a common discount rate of δ and letting π_k^j indicate industry profits if firms charge p^j in demand state k , the expected discounted stream of profits for each firm from participating in the collusive arrangement is:⁹

$$\begin{aligned}
 V &= (1 - \mu)(1 - \sigma) \left(\frac{\pi_h^m}{2} + \delta V \right) + \mu(1 - \sigma) \left(\frac{\pi_l^m}{2} + \delta V \right) \\
 &\quad + (1 - \mu)\sigma \left(\frac{\pi_l^m}{2} + \delta V \right) + \mu\sigma(\delta^{T+1}V) \\
 V &= (1 - \mu)(1 - \sigma) \left(\frac{\pi_h^m}{2} \right) + (\mu + \sigma - 2\mu\sigma) \left(\frac{\pi_l^m}{2} \right) + (1 - \mu\sigma)\delta V + \mu\sigma\delta^{T+1}V, \quad (1)
 \end{aligned}$$

where π_k^m indicates the profits earned if both firms charge the collusive price, p^m , given demand condition k . k can take two levels: high h and low l , reflecting whether no shocks occurred or only one shock occurred, respectively.¹⁰ The identity of the shock is irrelevant, as they are equivalent in magnitude. Rearranging terms in [equation \(1\)](#) shows that the value

⁸As noted above, [Miklós-Thal and Tucker \(Forthcoming\)](#) consider an alternative approach to coordination between firms competing in prices.

⁹Since this is a homogenous Bertrand model, the firms make zero profits in the price war phase.

¹⁰As indicated above, if both shocks occur, then all firms earn zero profits.

of coordinating is:

$$V = \frac{(1 - \mu)(1 - \sigma)\pi_h^m + (\mu + \sigma - 2\mu\sigma)\pi_l^m}{2(1 - (1 - \mu\sigma)\delta - \mu\sigma\delta^{T+1})}. \quad (2)$$

Thus, the payoffs to colluding are declining in the length of the punishment period.

Collusion will be sustainable if neither firm would prefer to undercut the collusive price to earn monopoly-level profits today and trigger a punishment phase with certainty over continuing to split the collusive profits until a punishment phase is triggered by the arrival of both shocks. This incentive compatibility condition can be expressed as:

$$V \geq (1 - \mu)(1 - \sigma)\pi_h^m + \mu(1 - \sigma)\pi_l^m + (1 - \mu)\sigma\pi_l^m + \delta^{T+1}V. \quad (3)$$

The inequality shows that the sustainability of collusion will be a function of the discount rate as in the classic modeling of collusion under conditions of no uncertainty. However, it will also be affected by the probability of shocks and the length of the punishment period.

All else equal, [equation \(3\)](#) shows that it becomes easier to sustain collusion as the duration of the price war goes to infinity. However, since V is decreasing in T (from [equation \(2\)](#)), the optimal punishment period for the firms is the smallest value of T that still satisfies [equation \(3\)](#). This value can be calculated by setting [equation \(2\)](#) equal to [equation \(3\)](#) and solving for T :

$$T = \frac{\text{Log}\left(\frac{1-2\delta+2\delta\mu\sigma}{2\mu\sigma\delta-\delta}\right)}{\text{Log}(\delta)}. \quad (4)$$

Overall, the model setup and results resemble those described in [Tirole \(1988\)](#) with minor modifications to account for the existence of separate random shocks. However, one of the key changes is that the optimal collusive price will no longer equal the monopoly price during periods of “ordinary” (i.e, unshocked) demand. Instead, it will reflect a balancing of the profits to be earned in the different demand states where demand is still greater than

zero for some positive price:

$$p^m \equiv \operatorname{argmax}_p (1 - \mu)(1 - \sigma) \left(\frac{\pi_h^m}{2} \right) + (\mu + \sigma - 2\mu\sigma) \left(\frac{\pi_l^m}{2} \right). \quad (5)$$

The determination of the optimal collusive price will be akin to that made by monopolists using a single price to sell to a pool of different consumer types who cannot be separated. In other words, it resembles the price-setting problem when there are two types of consumers but third degree price discrimination is not possible.

If the likelihood of being in the “low” demand state is high, then the optimal price should be substantially discounted relative to the monopoly price during the high period, even if that means failing to extract significant surplus from consumers willing to purchase during the high demand state. Similarly, if the likelihood of a low demand state is small, then the optimal price may be close to the monopoly level even though that means sacrificing some potentially profitable sales during low demand periods.

II.3 Post-AI

We now suppose that both firms adopt a technological innovation (AI) that allows them to perfectly predict the outcome of the M shock and react accordingly.¹¹ However, the incidence of the Σ shock remains unobserved. The introduction of the new technology can be seen as having reduced – but not completely removed – uncertainty in the market, allowing all firms to have greater insight into the demand conditions facing them in any given period. Thus, in each period, pricing will be optimized based on the public revelation of one of the demand shocks.

Given that firms’ information sets and optimal competitive choices change as a result of the introduction of AI, so, too, will their coordinated strategy. Whereas they had previously coordinated on one collusive price, p^m , now it is optimal to agree on two separate prices that

¹¹We discuss the possibility of endogenous adoption of AI in [Section V](#).

depend on what is commonly observed about the underlying demand conditions. Specifically, the optimal plan would be to charge p^l if M occurs and p^h otherwise, where $p^h \geq p^m \geq p^l$ with at least one of the inequalities being strict.

The elimination of uncertainty around one of the shocks alters more about the structure of coordination than just the number of prices to be selected ex ante. The firms can now perfectly infer if cheating has taken place when M has not occurred. As a result, we assume that there will be two different punishment rules.

Following the general literature, we assume that the punishment rule when a firm knows with certainty that its rival has cheated is an infinite reversion to competitive pricing. In contrast, we assume coordination will be maintained in the “low” demand states in much the same manner as before the arrival of the new technology. After all, the evidence suggests that cheating has possibly occurred, but the firm does not know this with certainty.

Let S be the duration of the new punishment period when a firm realizes zero profits after the M shock has occurred. The duration of the punishment phase S in this state may be different than the pre-AI punishment period T . This is because we assume that the punishment periods are chosen to maximize firm profits while keeping collusion incentive compatible.

We now derive the discounted ex-ante value of coordination, again letting π_k^j indicate industry profits if both firms charge price p^j , $j = l, h$ in demand state k . The demand state now differentiates between the two shocks (the observable M state is listed first): $k = hh$ with probability $(1 - \mu)(1 - \sigma)$, $k = hl$ with probability $(1 - \mu)\sigma$, and $k = lh$ with probability $\mu(1 - \sigma)$. The new discounted value of coordination is:

$$\begin{aligned}
 U = & (1 - \mu)(1 - \sigma) \left(\frac{\pi_{hh}^h}{2} + \delta U \right) + (1 - \mu)\sigma \left(\frac{\pi_{hl}^h}{2} + \delta U \right) \\
 & + \mu(1 - \sigma) \left(\frac{\pi_{lh}^l}{2} + \delta U \right) + \mu\sigma\delta^{S+1}U.
 \end{aligned} \tag{6}$$

Equation (6) can be rearranged to express the value of coordinating in the underlying

market parameters:

$$U = \frac{(1 - \mu)(1 - \sigma)\pi_{hh}^h + \mu(1 - \sigma)\pi_{lh}^l + (1 - \mu)\sigma\pi_{hl}^h}{2(1 - \delta(1 - \mu\sigma) - \mu\sigma\delta^{S+1})}. \quad (7)$$

Like the discounted value of collusion before AI (equation (2)), equation (7) declines in the length of the punishment period, signalling that the optimal duration of punishment periods is the shortest length that still satisfies the firms' incentive compatibility constraints.

As in the pre-AI world, a collusive equilibrium requires that neither firm would prefer to seize all of the profit in a given period even if it triggers a certain price war. However, while there was previously only one incentive compatibility constraint per firm, now there are two. This is because the relative payoffs to deviating from the collusive strategy differ depending on what is observed about the incidence of M .

First, it must be the case that firms would not cheat if they learn that the M shock will not happen. Cheating in this case would lead to a permanent breakdown in coordination - the continuation value is zero. Therefore, the relevant inequality for when M will not occur is:

$$(1 - \sigma)\left(\frac{\pi_{hh}^h}{2} + \delta U\right) + \sigma\left(\frac{\pi_{hl}^h}{2} + \delta U\right) \geq (1 - \sigma)\pi_{hh}^h + \sigma\pi_{hl}^h \quad (8)$$

$$U \geq \frac{1}{2\delta}((1 - \sigma)\pi_{hh}^h + \sigma\pi_{hl}^h).$$

The equation shows that while the long-run stream of payoffs goes to 0 if a firm cheats, deviating from the collusive arrangement leads to higher expected one-period payoffs, relative to the pre-AI deviation condition.

Second, it must also be the case that the analogue to equation (3) holds when all firms observe that the M shock will occur. Therefore, the relevant inequality for coordination when

M will occur is:

$$(1 - \sigma)\left(\frac{\pi_{hl}^l}{2} + \delta U\right) + \sigma\delta^{S+1}U \geq (1 - \sigma)(\pi_{hl}^l + \delta^{S+1}U) + \sigma\delta^{S+1}U \quad (9)$$

$$U \geq \frac{1}{\delta - \delta^{S+1}}\left(\frac{\pi_{lh}^l}{2}\right).$$

In this case, the expected short-term payoffs of deviation are lower than in the pre-AI world. All else equal, this should imply that fewer punishment periods are required to make deviating from the collusive strategy unappealing (i.e. $S < T$).

Comparing the two constraints shows that they are in tension. While a higher S makes it more likely that [equation \(9\)](#) holds, it makes it less likely that [equation \(8\)](#) also holds. This is because [equation \(7\)](#) showed that U was declining in S . Setting [equation \(7\)](#) equal to the two constraints now provides a lower and an upper bound on the values S may take and sustain collusion. These bounds are:

$$\underline{S} = \frac{\text{Log}\left(\frac{(\delta(-1+\mu)(\pi_{hh}^h(-1+\sigma)-\pi_{hl}^h\sigma)+\pi_{lh}^l(-1+\delta(1+\mu-2\mu\sigma)))}{(\delta(\pi_{hh}^h(-1+\mu)(-1+\sigma)-\pi_{hl}^h(-1+\mu)\sigma+\pi_{lh}^l(\mu-2\mu\sigma)))}\right)}{\text{Log}(\delta)} \leq S \quad (10)$$

and:

$$S \leq \frac{\text{Log}\left(\frac{(-\pi_{lh}^l\delta\mu(-1+\sigma)+\pi_{hh}^h(-1+\sigma)(1+\delta(-2+\mu+\mu\sigma))-\pi_{hl}^h\sigma(1+\delta(-2+\mu+\mu\sigma)))}{(\delta\mu\sigma(\pi_{hh}^h(-1+\sigma)-\pi_{hl}^h\sigma))}\right)}{\text{Log}(\delta)} = \bar{S}. \quad (11)$$

III Comparing Equilibria Before and After AI

III.1 Summary

In this section, we focus on what we take to be a central concern of policy-makers: that AI may both widen the scope of collusion and reduce consumer surplus even further where collusion is already possible. If this were the only impact of AI's reduction in demand uncertainty, consumers would be unambiguously worse off with the introduction of AI - a rare case

of technology benefits accruing only to the supply side of the market. However, as implied in the presentation of the model above, the reduction in demand uncertainty has certain commonalities with price discrimination, which the research literature has not shown to be unambiguously bad for consumers. Moreover, the introduction of AI leads to an additional incentive compatibility constraint that further bounds the parameter space in which collusion can take place. In theory, this may mean that AI actually decreases the incidence of collusion. We now outline our model’s implications for consumer and producer welfare with the introduction of AI technologies. All proofs are provided in [Appendix A](#).

III.2 Results

Proposition 1. *Collusion can take place in the pre-AI world.*

Our pre-AI setting largely replicates the original [Green and Porter \(1984\)](#) model but with two independent sources of uncertainty. While the additional shock requires minor modifications to when and for how long price wars take place, the underlying intuition for existence of a collusive equilibrium remains the same. Firms must place at least moderately high weight on future profits, and the likelihood of negative shocks cannot be too high. The payoffs in the different states is unimportant insofar as they do not appear in [equation \(4\)](#).

Proposition 2. *Collusion can, but does not always, take place in the post-AI world.*

We find that it is possible for collusion to be sustainable in the post-AI world. This makes intuitive sense; the incentive to collude is present when per-period profits – conditional on collusion being sustained – at least weakly increase relative to the pre-AI world. However, the various constraints that must be met for a finite, non-negative S (the optimal punishment period length) imply that this will depend not just on the discount rate and probabilities of the shocks but also the profits earned in the different states of the world. Unfortunately, the complexity of the constraints prevent us from delineating the specific part(s) of the parameter space where this is true. We can, and do, show that there are areas where these conditions are met, but also that there are circumstances where they are not satisfied.

Proposition 3. *Collusion can take place where previously not possible.*

Consistent with intuition, and perhaps policy-makers' fears, we find that there are parts of the parameter space where AI makes collusion feasible where it was previously not. This may stem from a decline in the attractiveness of cheating since it will be more detectable. In addition, the payoffs to coordinating are higher since more surplus can be extracted through better tailored pricing.

Proposition 4. *Collusion can cease to be possible where previously possible.*

Interestingly, and in contrast to some stated concerns from policy makers about how AI may affect competition, we also find that there are parts of the parameter space that can no longer sustain collusion after AI's adoption. In other words, the implications of reduced demand uncertainty are not unambiguously better for firms. The intuition is essentially the converse of **Proposition 3**. Because firms can better identify when it may be attractive to cheat, the threat of punishment may have to be more severe, putting upward pressure on the minimum level of S that can sustain collusion.

Proposition 5. *Consumer welfare may not fall even if AI increases the value of coordination.*

To this point we have focused on the implications of AI for firms. Intuitively, producer surplus increases where AI makes collusion possible where previously not and vice versa. Similarly, consumers benefit when collusion ceases to be feasible. In such parts of the parameter space, the effect of AI is to dramatically increase consumer surplus, and maximize total welfare.

However, the consequences for consumer and total welfare may be more nuanced when collusion takes place both prior to and following the arrival of the AI technology. Whether or not consumer welfare increases crucially depends on the incremental sales during periods when one demand shock is observed to occur, and the relative incidence of these events. This is because $p^l < p^m$, which will lead to more sales and greater consumer surplus during these periods. Pushing the opposite direction will be the reduced duration of price wars,

during which consumer surplus is maximized, and the higher prices charged when one shock is observed not to have occurred.

Overall, our results show that the welfare implications of increased transparency are ambiguous and not easily parsed *ex ante*. Depending on how the model is parameterized, either consumers or producers may suffer. Similarly, either group may benefit from the change, and at least one group will be strictly better off. However, the specific relationship between the different model parameters and outcomes is indeterminate given the complexity of the parameter space and the multiple, non-linear constraints required.

IV Linear Demand

The relationships between the different parameters and the prevalence and character of collusion are complex. To develop more precise intuition about how the equilibrium outcomes shift in response to differences in the parameter space, we analyze numerical outcomes for the case of linear demand.¹² Specifically, we assume that the inverse market demand function takes the familiar form of $P = a - bQ$. In the event that either the Σ or M shocks occur, then demand shifts down by c , i.e., $P = (a - c) - bQ$. As indicated above, both duopolists can produce a potentially infinite quantity at zero marginal cost.

To disentangle the relative importance of different parameters, as well as their interactions, we fix $a = 10$ and explore what happens as the other parameter values change. Specifically, we allow b to vary between 1 and 10, c to vary between 2.5 and 8, μ and σ (the independent probabilities of the M and Σ shocks respectively) to vary between 0.1 and 0.9, and δ to vary between 0.5 and 0.9.¹³

Table 1 shows descriptive statistics for key outcome variables of the numerical model. CS_{pre} and CS_{post} represent the net present values of consumer surplus in the pre- and post-AI

¹²Appendix B shows the formulations of elements of interest in the linear demand setting.

¹³Specifically, we vary μ and σ uniformly by 0.002; δ uniformly by 0.1; and b uniformly by 1. We define c as $\frac{10}{x}$ with x varying from 1.25 to 4 by 0.25.

worlds. U and V are, as above, the per-firm amount of discounted profit streams, Δ indicates change, and PS represents total producer surplus. The Table indicates that on average firm values are higher after the adoption of AI leads to reduced uncertainty (i.e. $U > V$). In contrast, the average amount of consumer surplus declines. However, relative to the pre-AI mean levels of surplus, neither average change is particularly large.

Table 1: Descriptive Statistics of Linear Demand Parameterizations

Statistic	N	Mean	St. Dev.	Min	Max
CS_{pre}	3,936,600	16.043	19.140	0.172	179.044
CS_{post}	3,936,600	15.356	17.842	0.172	171.550
V	3,936,600	2.048	4.152	0.000	55.430
U	3,936,600	2.162	4.433	0.000	56.132
ΔCS	3,936,600	-0.688	4.218	-49.140	57.549
ΔPS	3,936,600	0.229	1.374	-19.183	15.447

The relationships between the various elements of the parameter space and economic outcomes are non-linear. However, intuition about modest changes in the different elements that influence consumer and producer surplus can be gleaned by regressing them on a linearly separable function of the different model parameters. The results of these regressions on changes in producer surplus, consumer surplus, and an indicator for an improvement in consumer surplus ($1(\Delta CS > 0)$) are shown in [Table 2](#).

[Table 2](#) helps to clarify the relative salience of different elements. The first two columns suggest that, on average, the interests of consumers and firms are never aligned about the desirability of AI when demand is linear. Any parameter that is associated with an average increase in producer surplus following the arrival of AI is also associated with an average decline in consumer surplus, and vice versa.

In general, the regressions indicate that when demand is more steeply sloped (as measured by b), consumers benefit from AI. Conversely, when the negative shocks are larger, on average, the introduction of AI is associated with greater consumer harm. These results connect to the size of profits in the different states of the world. As b increases, the gaps

between the different profit states decrease. Interestingly, the final two columns indicate that c 's impact is non-linear. In expected value terms, a higher probability of the negative shock revealed by AI leads to lower consumer surplus. However, it increases the odds that consumer surplus rises relative to the pre-AI world.

On average, δ has a positive impact on producer surplus and a negative one on consumers. This suggests that, all else equal, collusion is more easily sustained, more lucrative, or both when firms highly value future profits and there is decreased demand uncertainty.

The two demand shock probabilities have opposite effects. The estimates for σ suggest that larger absolute changes in the level of demand uncertainty are associated with superior coordination. However, greater levels of μ are associated with better outcomes for consumers, and worse ones for producers. We interpret this result as indicating that, on average, sustaining collusion is harder when a high likelihood negative shock becomes observable.

To better disentangle the non-linearities, we now turn to a series of Figures that show how different economic outcomes change in relation to combinations of the parameter values. In these Figures, we fix the level of δ at 0.7, which is in the middle of the examined range. The results are qualitatively similar for other values of the discount rate; however, consistent with [Table 2](#), higher levels of δ are associated with better outcomes for firms and worse outcomes for consumers.

[Figure 1](#) displays how the sustainability of collusion changes from the pre-AI world to the post-AI world for the full range of possible values of the two uncertainty parameters μ and σ . If one or the other measures of uncertainty is low, collusion will be sustainable regardless of whether AI has been introduced. When μ is significantly larger than σ , collusion will only be sustainable in the pre-AI world. The dramatic reduction in uncertainty actually reduces the scope for coordination. In contrast, when σ is substantially larger than μ , there is a chance that the introduction of AI enables collusion when it was previously impossible.

[Figure 2](#) shows how consumer surplus changes due to the implementation of AI by the firms as a function of μ and σ . Qualitatively, the figure closely mirrors [Figure 1](#). This implies

Table 2: Decomposition of Economic Outcomes on Parameter Values

	<i>Dependent variable:</i>		
	ΔPS (1)	ΔCS (2)	$1(\Delta CS > 0)$ (3)
b	-0.058*** (0.0002)	0.174*** (0.001)	0.00000 (0.00004)
c	0.044*** (0.0004)	-0.123*** (0.001)	0.012*** (0.0001)
δ	3.072*** (0.005)	-9.275*** (0.014)	-0.106*** (0.001)
μ	-0.644*** (0.003)	1.900*** (0.009)	0.154*** (0.0004)
σ	0.248*** (0.003)	-0.796*** (0.009)	-0.225*** (0.0004)
Constant	-1.598*** (0.004)	4.831*** (0.013)	0.108*** (0.001)
Observations	3,936,600	3,936,600	3,936,600
R ²	0.131	0.126	0.100
Adjusted R ²	0.131	0.126	0.100
Residual Std. Error (df = 3936594)	1.281	3.944	0.205
F Statistic (df = 5; 3936594)	118,835.300***	113,551.600***	87,357.520***

Note:

*p<0.1; **p<0.05; ***p<0.01

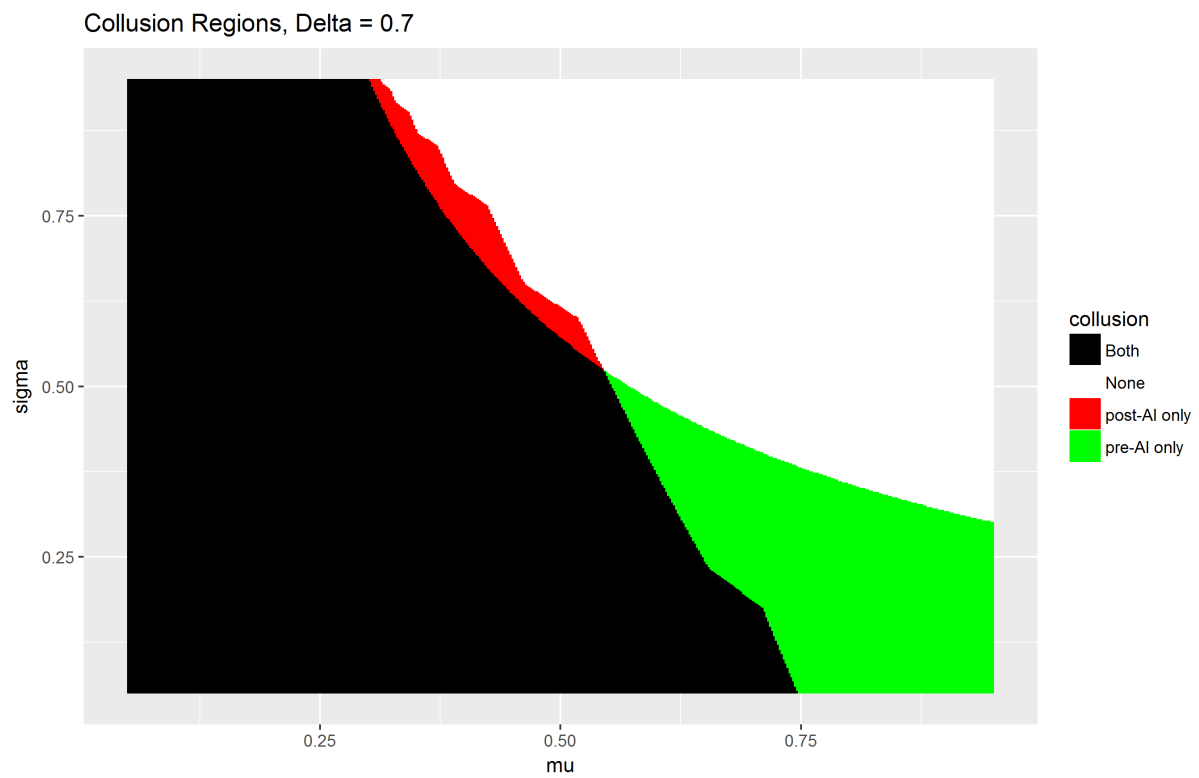


Figure 1: Sustainability of Collusion as a Function of Uncertainty Levels and AI

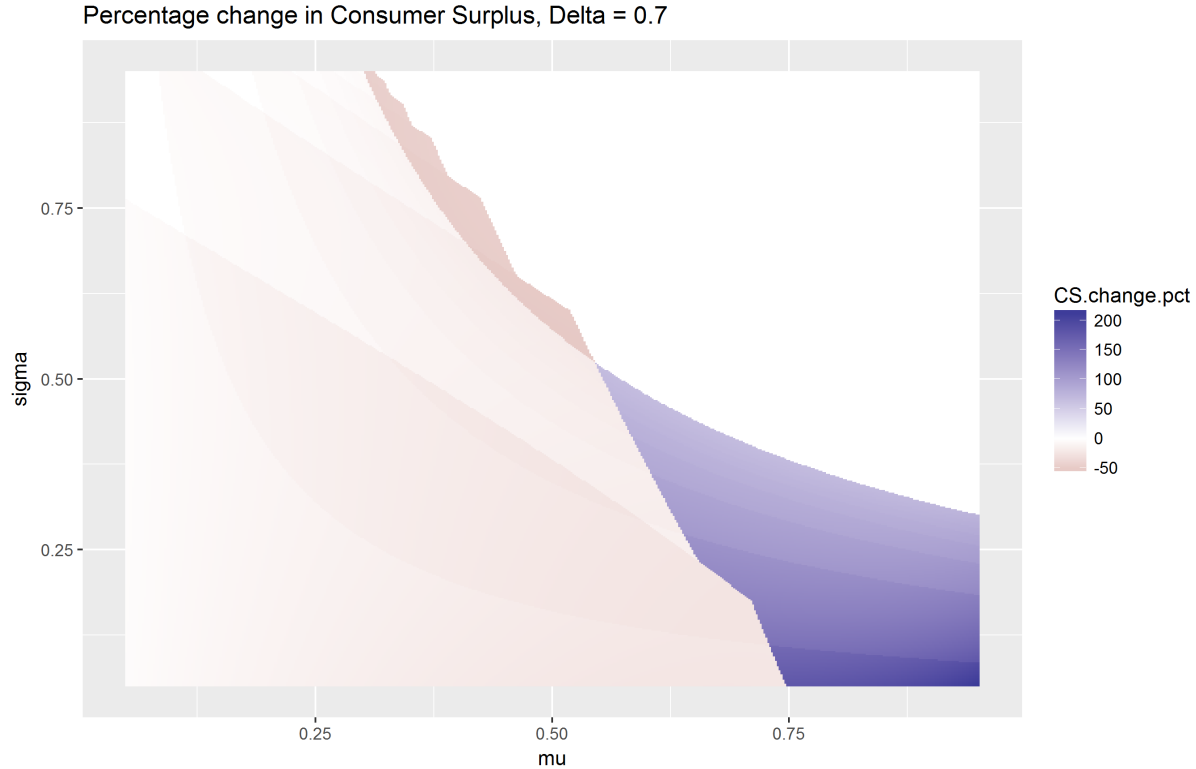


Figure 2: Consumer Surplus as a Function of Uncertainty Levels and AI

that consumer surplus does not improve with the implementation of AI unless AI makes collusion unsustainable. Otherwise, consumers are worse off. In line with expectations, we find that the highest percentage of lost consumer surplus occurs where AI makes collusion possible.

Figure 3 shows the percentage change in total welfare due to the implementation of AI. We find that total welfare changes closely track changes in consumer surplus. That is, the magnitude of changes in consumer surplus swamp the magnitude of the changes in producer surplus. Where consumers are worse off, total welfare falls and vice-versa. Thus, at least for the case of linear demand in the parameter range considered, the potential efficiency gains of being better able to price discriminate intertemporally are dwarfed by the increased incidence of coordination.

Finally, Figure 4 shows the change in optimal price war length that makes collusion




Figure 3: Total Surplus as a Function of Uncertainty Levels and AI

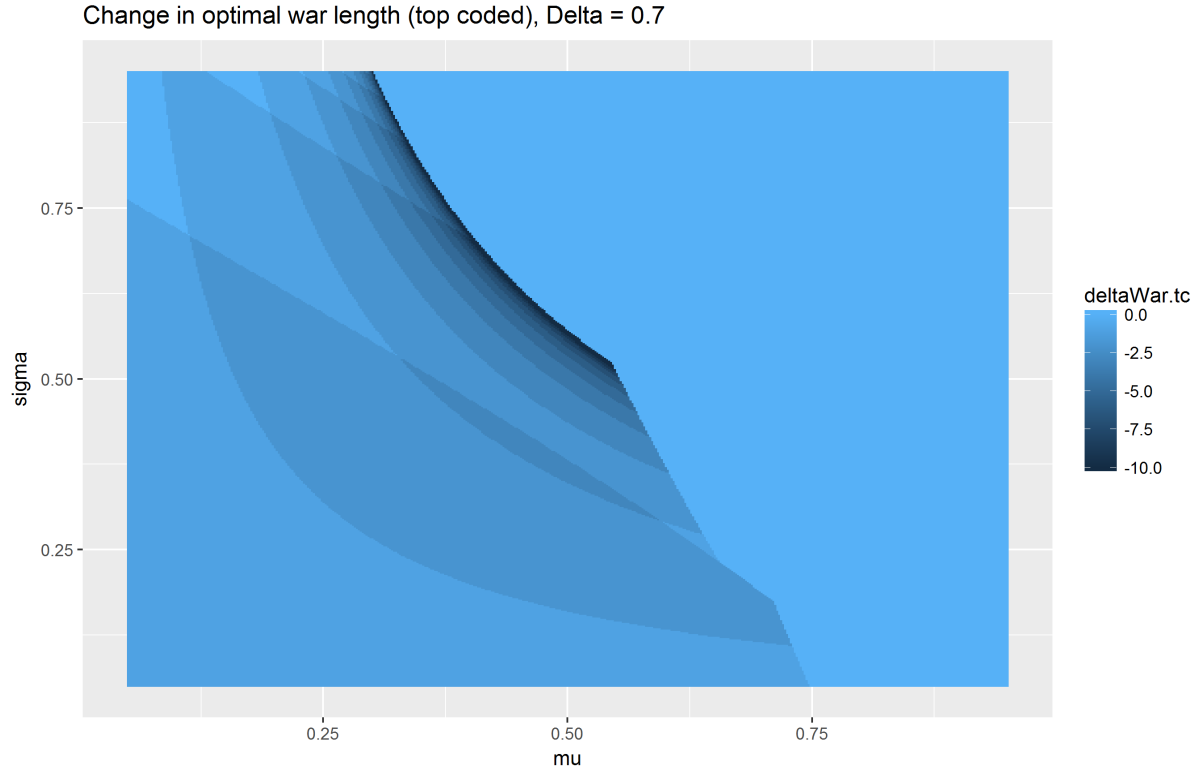


Figure 4: Change in Equilibrium Punishment Duration as a Function of Uncertainty Levels and AI (top coded at 10 periods)

sustainable for parameter values where collusion is sustainable regardless of the presence of AI. For every set of parameter values where collusion is sustainable in both states of the world, the optimal war length decreases with the introduction of AI. Interestingly, changes in consumer surplus do not necessarily track changes in war length. In particular, in the parts of the parameter space where war length drops significantly (e.g., by more than 10 periods), the percentage change in consumer welfare is not noticeably outsized. Relative to other factors, a significantly longer price war length does not significantly increase consumer surplus.

V Discussion

V.1 Why adopt AI if it renders collusion unsustainable?

Given that AI adoption renders collusion unsustainable for some parameter values, it is reasonable to ask why firms would ever adopt AI in these scenarios. We now discuss several plausible reasons why AI might be adopted even if it led to reductions in producer surplus.

First, competitors may find themselves with something of a coordination problem that bears a close resemblance to a standard “entry” game (Farrell, 1987). Consider a market where the two firms have successfully colluded without AI. Now, it becomes possible to adopt AI. It is observable that if both firms adopt the technology, they will no longer be able to successfully collude. However, each firm may privately have an incentive to adopt AI. This is because for some regions of the parameter space, a firm with AI may find it profit-maximizing to depart from the non-AI collusive strategy if the other firm does not have AI. The other firm will not necessarily alter its behavior, even though it now earns lower profits, because it still earns positive profits during “low” demand phases. With no means of coordinating over which firm will adopt and increase its profits at the rival’s expense, however, both might well choose to do so.

Second, firms may choose to invest in developing or adopting AI if it is uncertain ex ante which of the different shocks will become publicly observable. Thus, firms may invest in AI expecting to be able to coordinate significantly better ex post, but subsequently learn that the technology actually makes it no longer sustainable. We believe such misplaced expectations are not unreasonable in connection to AI and other digital means of reducing demand uncertainty. The technology space is evolving extremely rapidly, plausibly making predictions about the capabilities of new innovations difficult.

Third, we believe there may be circumstances where AI adoption may be an ancillary effect of decisions driven by other factors influencing firms’ objective functions. For example, multi-divisional firms may have an incentive to adopt more sophisticated analytical tools to

improve operations that have the additional effect of reducing the scope for coordination in some markets.

In future work, we hope to probe the robustness of our results to more fully endogenizing the development and adoption of the demand reducing technologies.

V.2 What if firms simultaneously set S and the collusive prices?

Pre-AI, the sustainability of collusion does not depend on the collusive payoffs. This is because no profit term appears in [equation \(4\)](#), which provides the criteria for the existence of a finite punishment period. Therefore, firms would never have an incentive to choose prices other than the risk-adjusted optimal price given in [equation \(5\)](#). Post-AI, however, this may not hold. Because the different profit terms appear in [equation \(10\)](#) and [equation \(11\)](#), choosing different collusive prices may have the effect of altering the space in which a finite price war period exists that may sustain collusion. If this were to be true, it might also be the case that by selecting different collusive prices firms could reduce the amount of the parameter space where AI actually works to consumers' benefit, providing support for some of the more forceful policy concerns associated with AI.

To address this question, we numerically examine alternative equilibrium outcomes in our linear demand setting. Specifically, we consider how price war lengths and firms' profits change when the different combinations of collusive prices are chosen. Alternative collusive equilibrium prices are selected on the basis of maximizing firm values.¹⁴

[Figure 5](#) shows how the scope for post-AI collusion may be expanded when non-monopoly prices are chosen. It indicates that in a large portion of the parameter space AI enables some degree of collusion when it was previously impossible. Moreover, allowing firms to simultaneously choose prices and price war durations means that there is no longer any

¹⁴We assess this using a brute force, grid search approach. Because of computational burden of this process, we restrict the considered parameter space to a subset of that used in our full linear simulations. Specifically, we increase the tick size between values of σ and μ from 0.002 to 0.01.

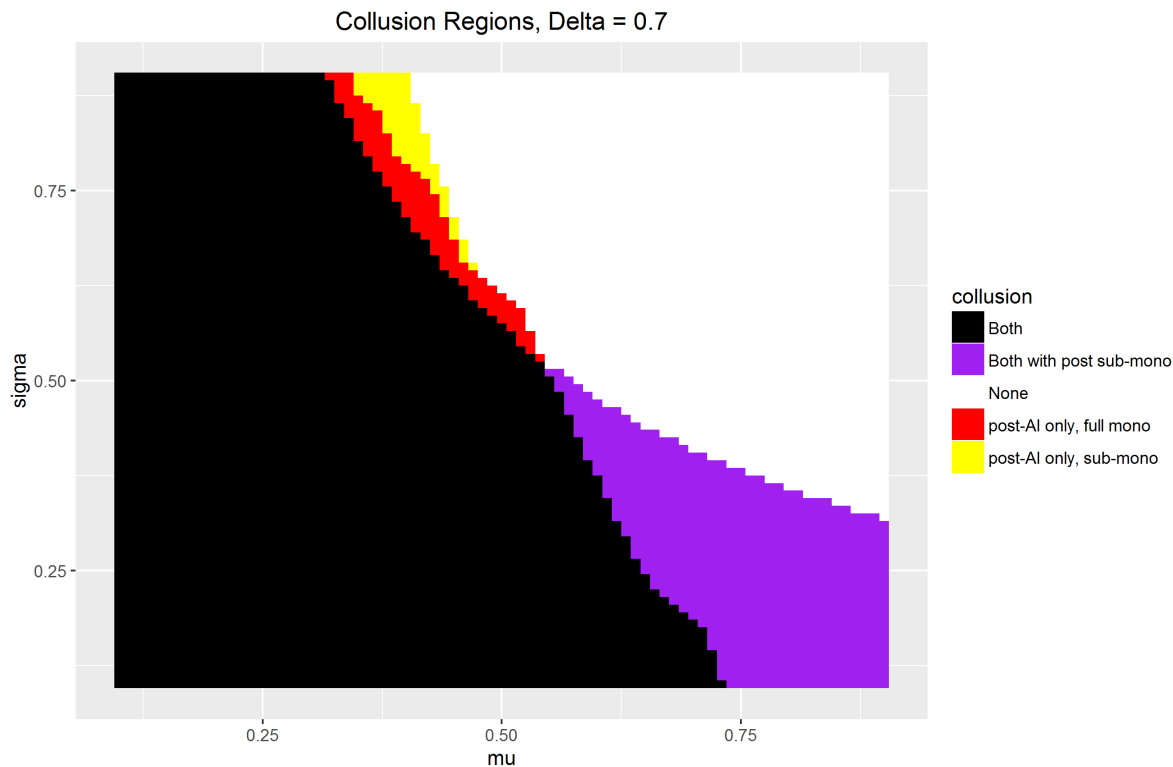


Figure 5: Sustainability of Collusion as a Function of Uncertainty Levels and AI

part of the parameter space where collusion was sustainable pre-AI, but not post-AI. All else equal, this appears to support the more pronounced concerns observed in the legal commentary about the impact of AI on competition.

While Figure 5 shows that varying prices leads to the unambiguous prediction that AI expands the scope for *some* degree of collusion, it does not necessarily undermine the conclusion that AI's welfare implications are ambiguous. To consider this, we plot the sign of changes in consumer welfare from pre-AI to post-AI when prices are allowed to vary in Figure 6. The Figure indicates that relative to the pre-AI world, consumer welfare still increases in parts of the parameter space even when *some* form of collusion is sustainable via non-monopoly pricing in the post-AI world. Therefore, our underlying conclusion that exogenous reductions in demand uncertainty have ambiguous implications for firms and consumers alike remains robust. This conclusion is not an artifact of using "monopoly" prices, which might be easily identified ex ante but sub-optimal from the firms' perspective.




Figure 6: Change in Consumer Welfare as a Function of Uncertainty Levels and AI

VI Conclusion

A number of works from antitrust scholars have claimed that AI and algorithms will lead to a dramatic transfer of surplus from consumers to producers. Often, this is driven by the belief that the algorithms will figure out a way to collude. Other parts of the literature have pushed back, outlining the many practical impediments that the original criticism missed. While we are sympathetic to the counterargument, we think it is important not to lose sight of the fact that the use of AI and/or data processing algorithms could nevertheless affect the incidence and costs of collusion. In this paper, we show how improved predictions on the state of demand, a plausible outcome of greater analytical sophistication, can influence firm conduct. Our results imply that the effect is ambiguous, but that there are parts of the parameter space where the adoption of improved analytical tools could harm consumers when the market structure is held constant and demand-side changes are not allowed.

References

- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb**, *Prediction Machines: The simple economics of artificial intelligence*, Harvard Business Press, 2018.
- Athey, Susan and Kyle Bagwell**, “Optimal collusion with private information,” *RAND Journal of Economics*, 2001, pp. 428–465.
- Calzolari, Giacomo, Emilio Calvano, Vincenzo Denicolò, and Sergio Pastorello**, “Algorithmic Pricing: What Implications for Competition Policy?,” 2018.
- Carlton, Dennis W and Jeffrey M Perloff**, *Modern industrial organization*, Pearson Higher Ed, 2015.
- Cooper, James C, Luke Froeb, Daniel P O’Brien, and Steven Tschantz**, “Does Price Discrimination Intensify Competition-Implications for Antitrust,” *Antitrust LJ*, 2004, 72, 327.
- Dubé, Jean-Pierre and Sanjog Misra**, “Scalable price targeting,” Technical Report, National Bureau of Economic Research 2017.
- Ezrachi, Ariel and Maurice E Stucke**, *Virtual competition*, Oxford University Press, 2016.
- and —, “The rise of behavioural discrimination,” *European Competition Law Review*, 2017, 37, 484–491.
- Farrell, Joseph**, “Cheap talk, coordination, and entry,” *The RAND Journal of Economics*, 1987, pp. 34–39.
- Gal, Michal S and Niva Elkin-Koren**, “Algorithmic consumers,” 2016.
- Green, Edward J and Robert H Porter**, “Noncooperative Collusion under Imperfect Price Information,” *Econometrica*, 1984, 52 (1), 84–100.
- Ittoo, Ashwin and Nicolas Petit**, “Algorithmic pricing agents and tacit collusion: A technological perspective,” *working paper*, 2017.
- Kandori, Michihiro**, “The use of information in repeated games with imperfect monitoring,” *The Review of Economic Studies*, 1992, 59 (3), 581–593.
- Kovacic, William E, Robert C Marshall, Leslie M Marx, and Halbert L White**, “Plus factors and agreement in antitrust law,” *Mich. L. Rev.*, 2011, 110, 393.
- Kuhn, Kai-Uwe and Steve Tadelis**, “Algorithmic Collusion,” *working paper*, 2017.
- Mehra, Salil K**, “Antitrust and the robo-seller: Competition in the time of algorithms,” *Minn. L. rev.*, 2015, 100, 1323.

- Miklós-Thal, Jeanine and Catherine Tucker**, “Collusion by Algorithm: Does Better Demand Prediction Facilitate Coordination Between Sellers?,” *Management Science*, Forthcoming.
- Raith, Michael A**, “Product differentiation, uncertainty and the stability of collusion,” *working paper*, 1996.
- Reimers, Imke and Benjamin R Shiller**, “Welfare Implications of Proprietary Data Collection: An Application to Telematics in Auto Insurance,” *working paper*, 2018.
- Ritter, Cyril**, “Bibliography on Antitrust and Algorithms,” 2017.
- Robson, Arthur J**, “Implicit oligopolistic collusion is destroyed by uncertainty,” *Economics Letters*, 1981, 7 (1), 75–80.
- Rotemberg, Julio J and Garth Saloner**, “A supergame-theoretic model of price wars during booms,” *The American Economic Review*, 1986, 76 (3), 390–407.
- Salcedo, Bruno**, “Pricing Algorithms and Tacit Collusion,” *Manuscript, Pennsylvania State University*, 2015.
- Shapiro, Carl**, “The 2010 horizontal merger guidelines: From hedgehog to fox in forty years,” *Antitrust LJ*, 2010, 77, 49.
- Tirole, Jean**, *The theory of industrial organization* 1988.

A Proofs

Proposition 1: Collusion can take place in the pre-AI world.

For collusion to be sustainable in the pre-AI world, T must be finite and non-negative. These constraints will be met when the numerator of [equation \(4\)](#) is positive but less than 1. This will hold when $2\mu\sigma > 1$ or $(-1 + 2\mu\sigma)(1 + 2\delta(-1 + \mu\sigma)) < 0$ (and all of μ, σ , and δ are all positive but less than 1).

Algebraic manipulation of these constraints shows that they will be met when $\frac{1}{2} < \delta < 1 \wedge \left((0 < \mu \leq \frac{2\delta-1}{2\delta} \wedge 0 < \sigma < 1) \vee \left(\frac{2\delta-1}{2\delta} < \mu < 1 \wedge 0 < \sigma < \frac{2\delta-1}{2\delta\mu} \right) \right)$.

Proposition 2: Collusion can, but does not always, take place in the post-AI world.

For collusion to be sustainable in the post-AI world, \underline{S} and \bar{S} must both be finite and non-negative. Moreover, \underline{S} must be weakly greater than \bar{S} . These constraints may be re-written as five distinct – often non-linear – inequality conditions. In addition, the solution must take into account the fact that $\pi_h^h > \pi_h^l > \pi_l^h$ as well as that all of μ, σ , and δ are all positive but less than 1. The combination of the non-linear quality constraints renders analytically solving for the parameter space in which collusion may take place impracticable. Therefore, we used numerical methods to establish that there exist places in the parameter space that satisfy all relevant constraints. For example, we found that if $\pi_h^h = \frac{270275}{262144}$, $\pi_l^h = \frac{1}{8192}$, $\pi_h^l = 1$, $\delta = \frac{16573}{32768}$, $\mu = \frac{1}{2}$, and $\sigma = \frac{1}{64}$, then all of the constraints hold. As seen below in the proof for [Proposition 4](#), there are also places in the parameter space where collusion cannot be sustained.

Proposition 3: Collusion can take place where previously not possible.

For collusion to be sustainable in the post-AI world, yet not in the pre-AI world, it must be the case that all of the constraints identified above in the discussion of [Proposition 2](#) hold while at least one of the constraints in [Proposition 1](#) do not. Unfortunately, the analytic burden of overlaying further restrictions onto those already necessary for a collusive equilibrium to exist in the post-AI world makes delineating the boundaries of the subset impossible. However, once more, we succeeded in employing numerical methods to show that all of the various constraints could be met in the parameter space. For example, we found that if $\pi_h^h = \frac{65}{64}$, $\pi_l^h = 1$, $\pi_h^l = \frac{99007}{98304}$, $\delta = \frac{536869079}{536870912}$, $\mu = \frac{123419}{131072}$, and $\sigma = \frac{278401}{524288}$, then all of the constraints hold.

Proposition 4: Collusion can stop being possible where previously possible.

As with the previous two propositions, we turn to numerical methods to establish existence. We find at least one example of a point in the parameter space that lies in the parameter space identified in **Proposition 1** but fails to meet one or more of the conditions for **Proposition 2**. Specifically, this will be true when $\pi_h^h = \frac{17}{16}$, $\pi_l^h = \frac{1}{8}$, $\pi_h^l = 1$, $\delta = \frac{3}{4}$, $\mu = \frac{1}{4}$, and $\sigma = \frac{7}{8}$.

Proposition 5: Consumer welfare may not fall even if collusion improves.

The intuition is akin to that for 3rd degree price discrimination. If the pre-AI price is high enough that no sales are made except in the high state of demand, then greater clarity about the true demand may enable selling some units during periods when one demand shock occurs. So long as the ratio of coordinated pricing periods to price war periods does not radically shift, it is possible that consumers could end up better off in aggregate. Thus, the overall welfare implications of increased transparency are ambiguous. Depending on how the model is parameterized, either consumers or producers, but not both, may suffer. At least one group will benefit.

B Linear Demand

Below, we provide the analytical formulas for key elements of interest.

Monopoly Price no shocks

$$\begin{aligned}
 & \max_p p \left(\frac{a}{b} - \frac{1}{b} P \right) \\
 FOC : & 0 = \frac{a}{b} - \frac{2}{b} P \\
 & p^m = \frac{a}{2} \\
 & \pi^m = \frac{a^2}{4b}
 \end{aligned} \tag{12}$$

Optimal Collusive Price with Shocks and No AI

$$\begin{aligned}
 & \max_p (1 - \mu)(1 - \sigma) \left(\frac{a}{b} p - \frac{p^2}{b} \right) + (\mu + \sigma - 2\mu\sigma) \left(\frac{a - c}{b} p - \frac{p^2}{b} \right) \\
 FOC : & 0 = (1 - \sigma - \mu + \mu\sigma) \left(\frac{a}{b} - \frac{2p}{b} \right) + (\mu + \sigma - 2\mu\sigma) \left(\frac{a - c}{b} - \frac{2p}{b} \right) \\
 & 0 = \frac{a}{b} (1 - \mu\sigma) - \frac{2p}{b} (1 - \mu\sigma) - \frac{c}{b} (\mu + \sigma - 2\mu\sigma) \\
 & p^m = \frac{a(1 - \mu\sigma) - c(\mu + \sigma - 2\mu\sigma)}{2(1 - \mu\sigma)}
 \end{aligned} \tag{13}$$

Collusive Quantity - 2 Possibilities

Case 1: No Shocks

$$\begin{aligned}
 & Q_h = \frac{a}{b} - \frac{1}{b} p \\
 Q_h^m &= \frac{a}{b} - \frac{1}{b} \left(\frac{a(1 - \mu\sigma) - c(\mu + \sigma - 2\mu\sigma)}{2(1 - \mu\sigma)} \right) \\
 &= \frac{a + c\mu + c\sigma - a\mu\sigma - 2c\mu\sigma}{2b - 2b\mu\sigma}
 \end{aligned} \tag{14}$$

Case 2: 1 Shock

$$Q_l = \frac{a-c}{b} - \frac{1}{b}p$$

$$Q_l^m = \frac{a-c}{b} - \frac{1}{b} \left(\frac{a(1-\mu\sigma) - c(\mu + \sigma - 2\mu\sigma)}{2(1-\mu\sigma)} \right)$$

$$= \frac{a-2c+c\mu+c\sigma-a\mu\sigma}{2b-2b\mu\sigma}$$
(15)

Collusive Profits - 2 possibilities

Case 1: No Shocks

$$\pi_h^m = \frac{Q_h^m}{2} * p^m$$

$$= \frac{(-a + (a + 2c)\mu\sigma - c(\mu + \sigma))(c(\mu + \sigma - 2\mu\sigma) + a(-1 + \mu\sigma))}{8b(-1 + \mu\sigma)^2}$$
(16)

Case 2: 1 Shock

$$\pi_l^m = \frac{Q_l^m}{2} * p^m$$

$$= \frac{(-c(-2 + \mu + \sigma) + a(-1 + \mu\sigma))(c(\mu + \sigma - 2\mu\sigma) + a(-1 + \mu\sigma))}{8b(-1 + \mu\sigma)^2}$$
(17)

Consumer Surplus

$$CS_{no} = (0.5) * ((1 - \mu) * (1 - \sigma) * (a - p^m) * \frac{Q_h^m}{2} + (\mu + \sigma - 2 * \mu * \sigma) * (a - c - p^m) * \frac{Q_l^m}{2})$$
(18)

Optimal Collusive Price with Shocks and AI

Case 1: Observe negative demand shock μ . In that case, pick monopoly price conditional on one demand shock. Thus, per work above:

$$p^l = \frac{a-c}{2}$$
(19)

Case 2: Observe that at least one demand shock doesn't happen. Thus, optimal price

maximizes expected profits across scenarios when no shock hits and unobservable shock hits.

$$\begin{aligned}
& \max_p (1 - \sigma) \left(\frac{a}{b}p - \frac{p^2}{b} \right) + \sigma \left(\frac{a - c}{b}p - \frac{p^2}{b} \right) \\
FOC : 0 &= (1 - \sigma) \left(\frac{a}{b} - \frac{2p}{b} \right) + \sigma \left(\frac{a - c}{b} - \frac{2p}{b} \right) \\
& 0 = \frac{a}{b} - \frac{2p}{b} - \sigma \frac{c}{b} \\
& p^h = \frac{a - \sigma c}{2}
\end{aligned} \tag{20}$$

Collusive Quantities

Case 1: Observe negative demand shock μ .

$$Q_l^l = \frac{a - c}{2b} \tag{21}$$

Case 2: Observe that at least one demand shock doesn't happen. Two possibilities: one where no shock occurs; one where one does.

$$\begin{aligned}
Q_h^h &= \frac{a + c\sigma}{2b} \\
Q_l^h &= \frac{a + c(-2 + \sigma)}{2b}
\end{aligned} \tag{22}$$

Collusive Profits

Case 1: Observe negative demand shock μ .

$$\begin{aligned}
\pi_l^l &= p^l * \frac{Q_l^l}{2} \\
&= \frac{(a - c)^2}{8b}
\end{aligned} \tag{23}$$

Case 2: Observe that at least one demand shock doesn't happen. Two possibilities: one where no shock occurs; one where one does.

$$\begin{aligned}
\pi_h^h &= p_h^h * \frac{Q_h^h}{2} = \frac{(a - c\sigma)(a + c\sigma)}{8b} \\
\pi_l^h &= p_l^h * \frac{Q_l^h}{2} = \frac{(a + c(-2 + \sigma))(a - c\sigma)}{8b}
\end{aligned} \tag{24}$$

Consumer Surplus

$$\begin{aligned}CS_{AI}^l &= (0.5) * (1 - \sigma) * (a - c - p^l) * \frac{Q_l^l}{2} \\CS_{AI}^h &= (0.5) * (1 - \sigma) * (a - p^h) * \frac{Q_h^h}{2} + \sigma * (a - c - p^h) * \frac{Q_l^h}{2} \\CS_{AI} &= (1 - \mu) * CS_{AI}^h + \mu * CS_{AI}^l\end{aligned}\tag{25}$$