Preliminary: Do Not Cite Without Authors' Permission

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

Jason O'Connor
Federal Trade Commission
joconnor@ftc.gov

Nathan E. Wilson
Federal Trade Commission
nwilson@ftc.gov

January 14, 2019

Abstract

We consider how technologies that reduce demand uncertainty may influence the character and prevalence of coordinated conduct. Our results show that reducing uncertainty for firms has ambiguous implications for both consumers and firms. If coordination can be sustained after the reduction in uncertainty, we find that producer surplus unambiguously increases. For numerical analyses of markets with linear demand, we find that consumers in these markets always do worse. However, this result need not generalize to all demand systems. We also find that reducing demand uncertainty means that some parts of the parameter space can no longer sustain collusion. This occurs because cheating becomes more attractive when firms can tell if demand is high. When collusion ceases to be feasible, firms always suffer while consumers always benefit.

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 and Scott Stern as well as Dan Hosken, Patrick Degraba, David Schmidt, and other colleagues at the FTC. The usual caveat applies.

I Introduction

The role of uncertainty in the sustainability of collusion has long interested economists. This reflects the fact that no market has ever operated under conditions of perfect information. However, the recent proliferation of large data sets, plus the availability of algorithmic tools (”artificial intelligence”, broadly) to analyze them, have caused some commentators to wonder whether that remains the case. These developments naturally lead to the question of what the effect of a dramatic increase in market transparency might be on competitive conduct. While the full effects of these tools are still unclear, recent statements by policy-makers have suggested growing concern about the impact sophisticated analytical tools may have on competition.[1]

In this paper, we consider the implications of reducing uncertainty on the incidence and character of collusion. Using a framework derived from the seminal Green and Porter (1984) model of dynamic competition in an uncertain environment, we compare producer and consumer surplus before and after firms adopt a technology that allows firms to process noisy information about demand more efficiently; the process reduces their uncertainty about the profitability of different prices. A focus on being able to better respond to variation in demand conditions is broadly in keeping with how artificial intelligence is often discussed in the broader economic literature (Agrawal et al., 2018). Therefore, for narrative convenience, we will often refer to the uncertainty-reducing technology as “AI.”

We find that the exogenous introduction of AI has ambiguous implications for both firms and consumers. Intuitively, reduced demand uncertainty benefits firms in many cases. For example, by clarifying what the true demand conditions are, AI enables firms to more precisely differentiate rivals' cheating from unobserved negative demand shocks. Firms also use

[1] For example, in October 2018, the Assistant Attorney General for antitrust intimated that a price-fixing case involving algorithms might be brought soon (https://www.broadcastingcable.com/news/delrahim-criminal-case-against-anti-competitive-search-algorithms-coming#disqus_thread). At roughly the same time, the UK's Competition and Markets Authority released a white paper describing risk factors associated with the use of algorithms (https://www.gov.uk/government/publications/pricing-algorithms-research-collusion-and-personalised-pricing).

the technology to increase collusive per-period profits by setting lower prices when demand is likely to be low. Overall, conditional on collusion being sustainable both before and after the introduction of the uncertainty-reducing technology, we find that the reduction of uncertainty leads unambiguously the intuitive result: higher profits for firms. Their expected per-period profits increase, and the duration of the optimal (i.e., shortest) punishment period falls. Moreover, we find that the increased payoffs to coordination enables collusion in some parts of the parameter space where previously it was not sustainable.

However, contrary to the intuitive result discussed above, the introduction of AI does not invariably benefit firms. Our analysis reveals that reducing uncertainty may sometimes inhibit coordination when it was previously possible. The crucial intuition is that greater clarity about the true demand state raises the incentive for firms to cheat when they observe high demand. In such circumstances, consumer welfare unambiguously increases and firms are always worse off. Interestingly, the implications of reduced demand uncertainty are ambiguous for consumers even when collusion is sustainable both before and after the technology is adopted. While consumers suffer from the reduction in price war duration as well as higher average prices during "high" demand periods, sales increase in periods of low demand. Thus, the overall implications of the introduction of AI will depend on the specific parameterization of demand.

To gain greater insight into how different model parameters interact, we implement our model numerically in a standard linear demand setting. 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.

Although clearly a simplification, we believe our modeling framework captures some of the 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 foodservice markets, buyers and sellers engage in bilateral negotiations where the prices offered by firms is not observable to their rivals.[2] The outcomes of 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.

Overall, we contribute to the large and ever expanding 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, Raith, 1996, Athey and Bagwell, 2001). Our work shows that reducing uncertainty has ambiguous effects, making collusion both more profitable and sometimes more difficult to sustain.

In addition, the paper contributes to the ongoing debate about how antitrust policy should address algorithms and data analytics. This literature already contains contributions reflecting a wide variety of opinions (see, e.g., the partial bibliography of Ritter (2017)), but many commentators have suggested that these phenomena threaten consumer welfare. In our view, the critics have focused on two related, but distinct, theories of harm: the increased ability to personalize prices to better extract consumer surplus (Ezrachi and Stucke, 2016, 2017) and that 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 providing a rigorous basis for either in the economic or computer science literatures.[3]

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

[2] See, e.g., `https://www.ftc.gov/system/files/documents/cases/150219syscopt3cmpt.pdf`.

[3] Salcedo (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. Reducing uncertainty has clear implications for price discrimination, which, in turn, affects the sustainability of collusion. This is because the reduction in uncertainty affects the relative payoffs to coordinating as well as incentive compatibility constraints facing firms. 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. These nuanced results stand in contrast to the qualitative literature, but fit with other emerging contributions taking a more technically grounded approach (Ittoo and Petit, 2017, Kuhn and Tadelis, 2017, Calzolari et al., 2018). These preliminary studies have shown that the use of algorithms has ambiguous effects on the likelihood that pricing algorithms successfully automate tacit collusion.

The paper is organized as follows. Section II outlines the modeling framework, and then Section III compares the equlibria 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. These occur with strictly positive probabilities μ and σ, respectively. For the sake of simplicity, the magnitude of individual shocks are equal, but if both occur then firms are unable to sell their product at any price, and, therefore, make no profits.

The two firms produce homogeneous products, and compete in prices.[4] Demand slopes down. In the event that one firm has a lower price, it supplies the entire market what 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. However, it could also indicate that its price had been higher than that of firm j.

II.2 Pre-AI

Given the nature of competition and the form of demand uncertainty, one form that coordination might take is akin to that first proposed by Green and Porter (1984) in the context of quantity-setting firms and re-formulated for Bertrand competition in Tirole (1988, Section 6.7.1). Firms' collusive equilibrium strategy would take 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 profit for each firm from

[4] We believe the assumption of homogeneous products is without loss of generality so long as the extent of differentiation across firms is public. 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. Allowing for unobserved

participating in the collusive arrangement is:[5]

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

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 occured or only one shock occurred, respectively.[6] Rearranging terms in equation (1) shows that the value of coordinating is:

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

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:

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

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

[5] Since this is a homogenous Bertrand model, the firms make zero profits in the price war phase.
[6] As indicated above, if both shocks occur, then all firms earn zero profits.

duration of the price war goes to infinity. However, given that equation (2) showed that V is decreasing in T, this means that 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:

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

Overall, the model setup resembles that described in Tirole (1988) with modest changes made 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 postiive price:

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

Thus, 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 arriving during the high demand state. Similarly, if the fraction of periods expected to be spent in low demand states 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 that allows them to observe the outcome of M and react accrordingly.[7] However, the incidence of the other shock (Σ) 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 can and will be tailored 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 technological advancement, 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 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 about one of the shocks alters more about the structure of coordination than just the number of prices to be selected ex ante. This is because now the firms can perfectly infer if cheating has taken place when demand has not been reduced by the incidence of M. As a result, it would be rational for there to be two different punishment rules, with the severity being much more acute when it is clear that cheating has taken place.

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, coordination will be maintained in the "low" demand states in much the same manner as before the arrival of the new technology. Let S be the duration of the new punishment period when a firm realizes zero profits but the M shock has occurred. That is, cheating has possibly occurred, but the firm does not know this with certainty. The duration of the punishment phase in this state, S, may be different than the pre-AI punishment period T. Once more letting π_k^j indicate industry profits if both firms charge price p^j in demand

[7] We address the possibilty of endogenous adoption in Section V.

state k, where $b = hh$ with probability $(1-\mu)(1-\sigma)$, $k = hl$ with probability $(1-\mu)\sigma$, and $b = lh$ with probability $\mu(1-\sigma)$, the new discounted value of engaging in coordination is:

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

As before, equation (6) can be rearranged to express the value of coordinating in the underlying market parameters:

$$U = \frac{(1-\mu)(1-\sigma)\pi^h_{hh} + \mu(1-\sigma)\pi^l_{lh} + (1-\mu)\sigma\pi^h_{hl}}{2(1-\delta(1-\mu\sigma) - \mu\sigma\delta^{S+1})}. \tag{7}$$

Like 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, whereas 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. Therefore, the relevant inequality is:

$$\begin{aligned} U|\text{no } M = (1-\sigma)(\frac{\pi^h_{hh}}{2} + \delta U) + \sigma(\frac{\pi^h_{hl}}{2} + \delta U) \geq (1-\sigma)\pi^h_{hh} + \sigma\pi^h_{hl} \\ U \geq \frac{1}{2\delta}((1-\sigma)\pi^h_{hh} + \sigma\pi^h_{hl}). \end{aligned} \tag{8}$$

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

Second, it must also be the case that the analogue to equation (3) holds when all firms observe that M does happen. Thus, coordination requires:

$$\begin{aligned} U|M = (1-\sigma)(\frac{\pi^l_{hl}}{2} + \delta U) + \sigma\delta^{S+1}U &\geq (1-\sigma)(\pi^l_{hl} + \delta^{S+1}U) + \sigma\delta^{S+1}U \\ U &\geq \frac{1}{\delta - \delta^{S+1}}(\frac{\pi^l_{lh}}{2}). \end{aligned} \tag{9}$$

In this case, the expected short-term payoffs 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.

Comparing and contrasting 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:

$$\frac{Log\left(\frac{(\delta(-1+\mu)(\pi^h_{hh}(-1+\sigma)-\pi^h_{hl}\sigma)+\pi^l_{lh}(-1+\delta(1+\mu-2\mu\sigma)))}{(\delta(\pi^h_{hh}(-1+\mu)(-1+\sigma)-\pi^h_{hl}(-1+\mu)\sigma+\pi^l_{lh}(\mu-2\mu\sigma)))}\right)}{Log(\delta)} \leq S \tag{10}$$

and:

$$S \leq \frac{Log\left(\frac{(-\pi^l_{lh}\delta\mu(-1+\sigma)+\pi^h_{hh}(-1+\sigma)(1+\delta(-2+\mu+\mu\sigma))-\pi^h_{hl}\sigma(1+\delta(-2+\mu+\mu\sigma)))}{(\delta\mu\sigma(\pi^h_{hh}(-1+\sigma)-\pi^h_{hl}\sigma))}\right)}{Log(\delta)}. \tag{11}$$

III Comparing Equilibria Before and After AI

Under Revision; New Results Coming Soon

IV Applications

IV.1 Linear Demand

We consider how welfare changes depending on the parameter space for the case of linear demand.[8] Specifically, we assume that the inverse 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., $P = (a - c) - bQ$. Both duopolists can produce a potentially infinite quantity at zero marginal cost.

To consider the relative importance of different parameters, 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 to vary between 0.1 and 0.9, and δ to vary between 0.5 and 0.9.

Table 1 shows descriptive statistics for key outcome variables of the simulations. M and N represent the discounted amount of consumer surplus in the pre- and post-AI worlds. U and V are, as above, the per-firm amount of discounted profit, Δ indicates change, and PS and CS represent total consumer and producer surplus, respectively. The Table indicates that on average firm profits are higher after the adoption of AI leads to reduced uncertainty. In contrast, the average amount of consumer surplus declines.

The relationships between the various elements of the parameter space and economic outcomes are ultimately non-linear. However, intuition about how the different elements influence consumer and producer surplus can be gleaned through linear projections of them on a linearly separable function of the different model parameters. The results of projections of changes in the different amounts of surplus and the indicator for an improvement in consumer surplus are shown in Table 2.

Table 2 helps to clarify the relative salience of different elements. The first two columns

[8] The interested reader is referred to Appendix B, which shows the formulations of elements of interest in the linear demand setting.

Table 1: Descriptive Statistics of Linear Demand Parameterizations

Statistic	N	Mean	St. Dev.	Min	Max
M	48,600	15.806	19.096	0.172	165.625
N	48,600	14.721	17.522	0.172	165.625
U	48,600	2.280	4.599	0.000	56.132
V	48,600	2.099	4.320	0.000	55.430
ΔCS	48,600	-1.085	5.641	-103.031	55.852
ΔPS	48,600	0.362	1.880	-18.617	34.344
$1(\Delta(CS) > 0)$	48,600	0.053	0.223	0	1

suggests that, on average, the interests of consumers and firms are never aligned about the desirability of AI. 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. Interestingly, the final two columns indicate that μ's impact is non-linear. On average, a higher the probability of the negative shock revealed by AI leads to reductions in consumer surplus. However, it increases the odds that consumer surplus rises relative to the pre-AI world.

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 M 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.

Table 2: Decomposition of Economic Outcomes on Parameter Values

	Dependent variable:		
	ΔPS	ΔCS	$1(\Delta(CS) > 0)$
	(1)	(2)	(3)
b	-0.091^{***}	0.274^{***}	0.000
	(0.003)	(0.008)	(0.0003)
c	0.001	-0.002	0.012^{***}
	(0.005)	(0.014)	(0.001)
δ	4.296^{***}	-12.888^{***}	-0.107^{***}
	(0.056)	(0.168)	(0.007)
μ	0.070^{**}	-0.211^{**}	0.165^{***}
	(0.031)	(0.092)	(0.004)
σ	0.885^{***}	-2.654^{***}	-0.214^{***}
	(0.031)	(0.092)	(0.004)
Constant	-2.623^{***}	7.868^{***}	0.101^{***}
	(0.052)	(0.157)	(0.006)
Observations	48,600	48,600	48,600
R^2	0.139	0.139	0.110
Adjusted R^2	0.139	0.139	0.109
Residual Std. Error (df = 48594)	1.745	5.235	0.211
F Statistic (df = 5; 48594)	$1,566.024^{***}$	$1,566.024^{***}$	$1,195.641^{***}$
Note:			*p<0.1; **p<0.05; ***p<0.01

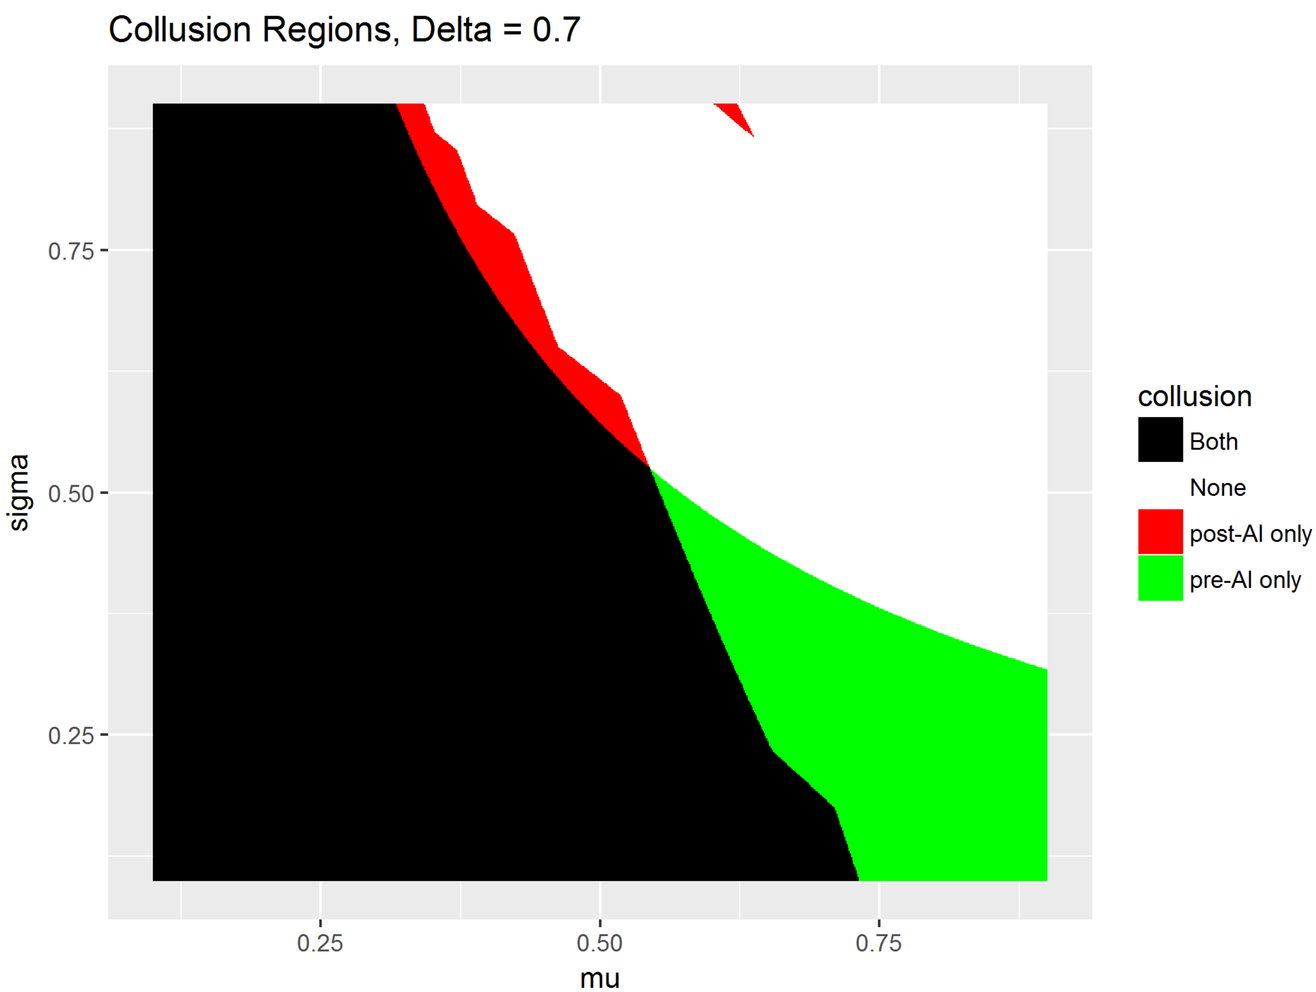

Figure 1: Sustainability of Collusion as a Function of Uncertainty Levels and AI

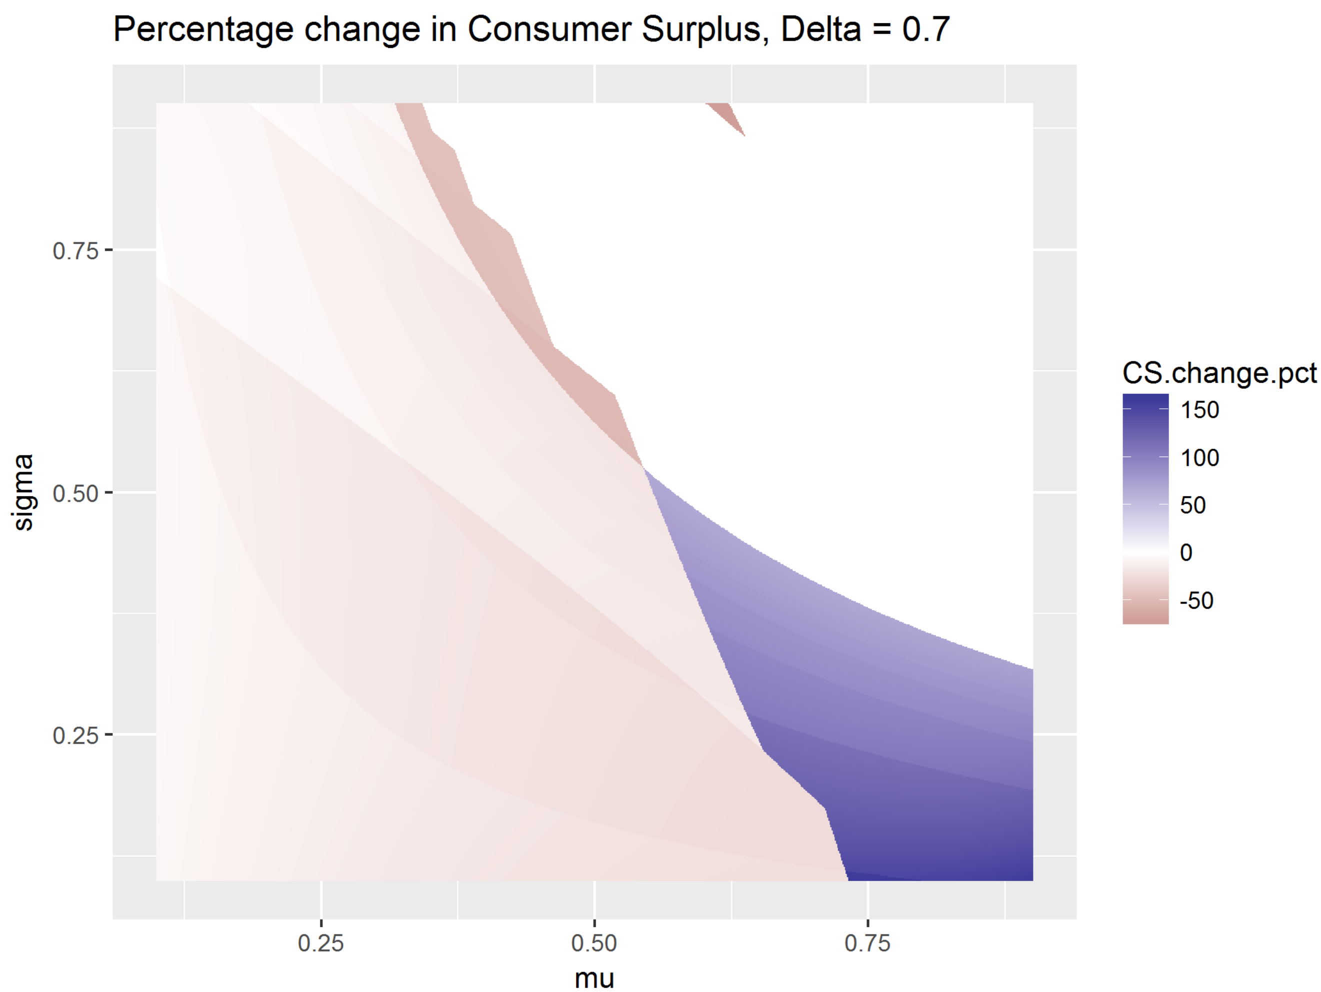

Figure 2: Consumer Surplus as a Function of Uncertainty Levels and AI

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 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

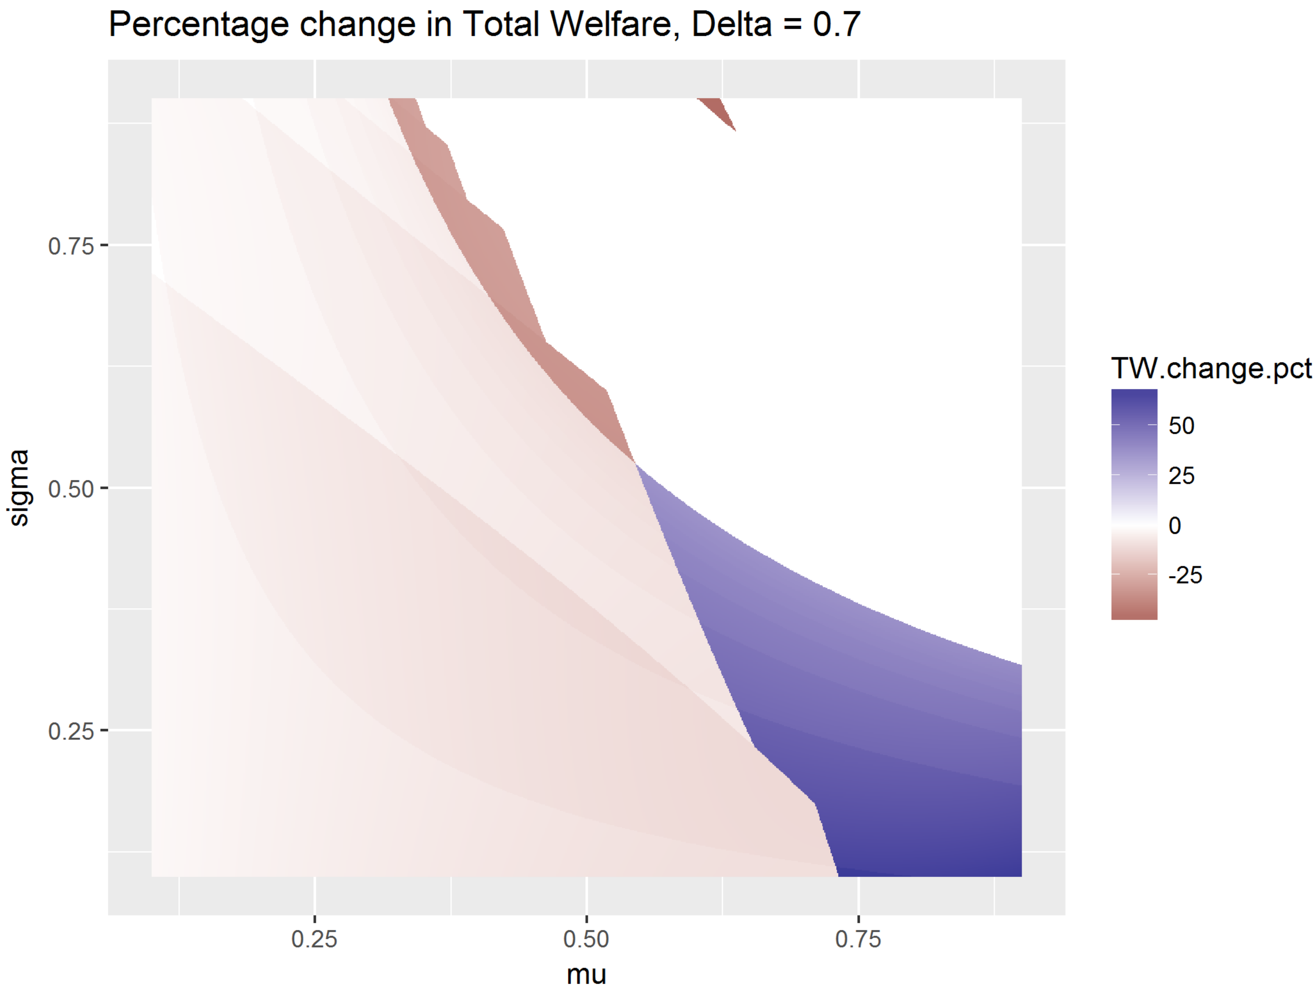

Figure 3: Total Surplus as a Function of Uncertainty Levels and AI

surplus. Where consumers are worse off, total welfare falls and vice-versa. Thus, at least for the case of linear demand, 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 sustainable for parameter values where collusion is sustainable regardless of the presence of AI. Consistent with Proposition 1, 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

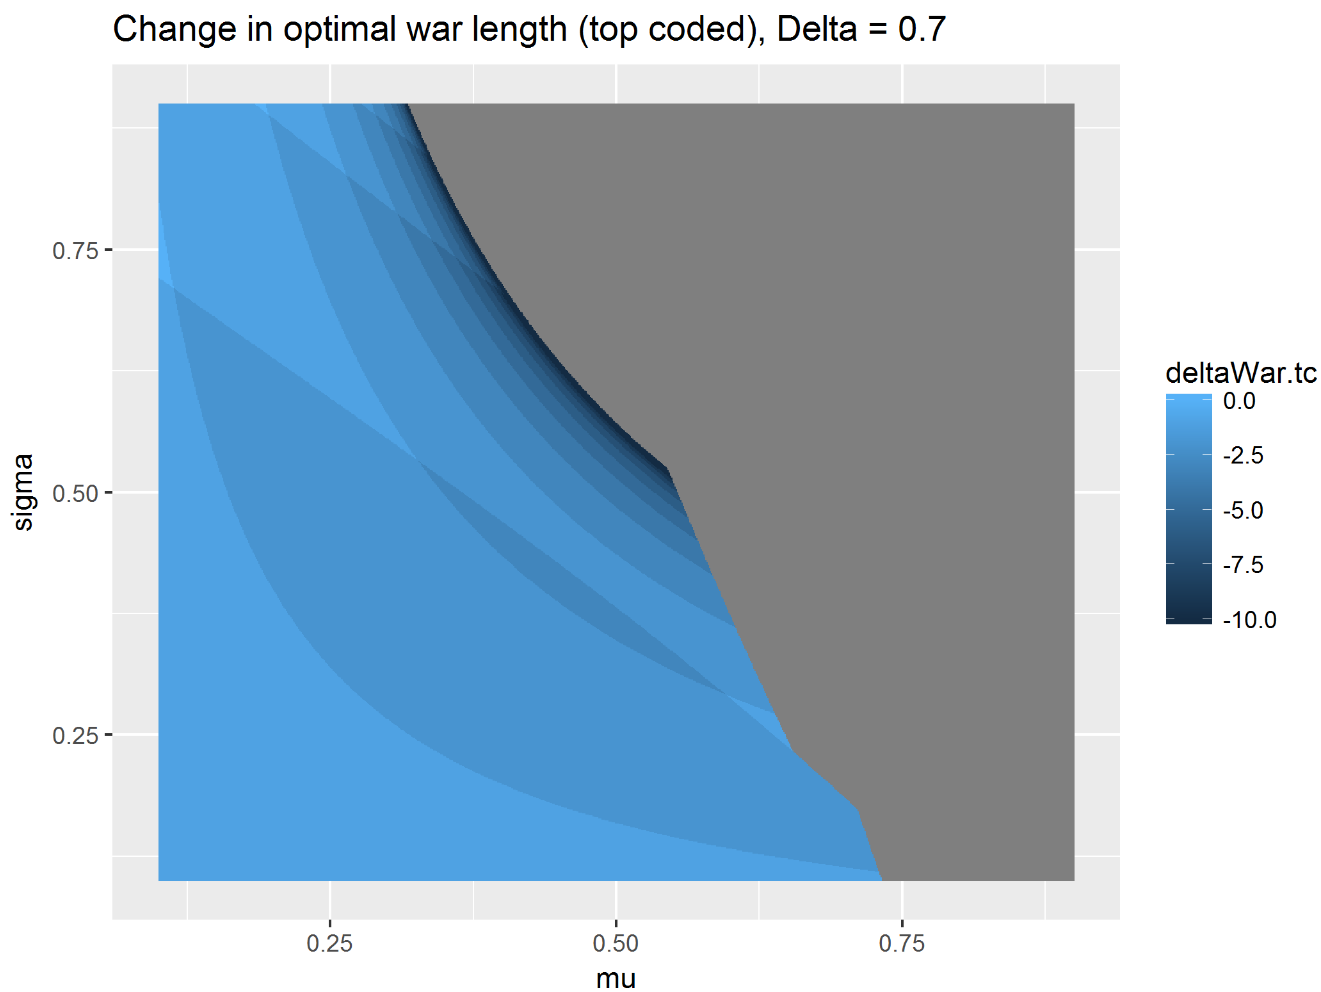

Figure 4: Change in Equilibrium Punishment Duration as a Function of Uncertainty Levels and AI (top coded at 10 periods)

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?

It is reasonable to wonder why firms would ever adopt AI if they were in parts of the parameter space where it renders collusion unsustainable. We agree that this is a valid question. However, we can think of several plausible reasons why AI might be adopted even if it led to reductions in producer surplus.

First, competitors may find themselves in something of a "prisoners' dilemma." 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. It should be noted that this is not a true "prisoner's dilemma," because AI adoption is not a dominant strategy. However, with no means of coordinating over which firm firm will adopt and increase its profits at the rival's expense, 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.

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.

VI Conclusion

Antitrust has seen an avalanche of articles claiming 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 changes in information 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.

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.

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.

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.

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

Under Revision; New Results Coming Soon

B Linear Demand

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

Monopoly Price no shocks

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

Optimal Collusive Price with Shocks and No AI

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

Collusive Quantity - 2 Possibilities

$$\begin{aligned} &\text{Case 1: No Shocks} \\ Q_h &= \frac{a}{b} - \frac{1}{b}p \\ Q_h^m = \frac{a}{b} - \frac{1}{b} & \left(\frac{a(1-\mu\sigma) - c(\mu+\sigma-2\mu\sigma)}{2(1-\mu\sigma)}\right) \\ &= \frac{a + c\mu + c\sigma - a\mu\sigma - 2c\mu\sigma}{2b - 2b\mu\sigma} \end{aligned} \tag{14}$$

Case 2: 1 Shock

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

$$Q_l^m = \frac{a-c}{b} - \frac{1}{b}\left(\frac{a(1-\mu\sigma) - c(\mu+\sigma-2\mu\sigma)}{2(1-\mu\sigma)}\right) \tag{15}$$

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

Collusive Profits - 2 possibilities

Case 1: No Shocks

$$\pi_h^m = \frac{Q_h^m}{2} * p^m \tag{16}$$

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

Case 2: 1 Shock

$$\pi_l^m = \frac{Q_l^m}{2} * p^m \tag{17}$$

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

Consumer Surplus

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

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:

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

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

$$
\begin{aligned}
CS^l_{AI} &= (0.5) * (1-\sigma) * (a-c-p^l) * \frac{Q^l_l}{2} \\
CS^h_{AI} &= (0.5) * (1-\sigma) * (a-p^h) * \frac{Q^h_h}{2} + \sigma * (a-c-p^h) * \frac{Q^h_l}{2} \\
CS_{AI} &= (1-\mu) * CS^h_{AI} + \mu * CS^l_{AI}
\end{aligned} \tag{25}
$$