

Collusion with Private Monitoring and Private Information

Martin Cripps* and Suehyun Kwon†

January 5, 2015

Abstract

We consider a model of price competition and collusion with private monitoring. The sellers have homogeneous good, and the buyer has a unit demand every period over an infinite horizon. The buyer's valuation is his private information, and the sellers start with a common prior. The seller only knows when the buyer comes to him and what price he charges. The buyer knows the price of the seller he visits. We characterize sequential equilibria and focus on equilibria that maximize the sellers' payoffs. When the sellers know the buyer's valuation, the buyer's payoff can be any number between his outside option and the upper bound. With incomplete information, the sellers repeat the static optimum in the best equilibria for the sellers. If the virtual surplus is positive, the buyer buys every period with probability one, and he never buys otherwise. Market sharing is limited by the seller's incentives to charge the equilibrium price.

Keywords: collusion, price competition, private monitoring

1 Introduction

With membership plans and loyalty cards, sellers can keep track of a buyer's purchase history. For example, with a Nectar card, a buyer can collect points when he shops at Sainsbury's or Expedia, and the sellers can in return know what and when the buyer buys. The seller-buyer interaction is not anonymous, and one would think that the buyer as well as the seller behaves strategically in the repeated interaction. However, while the interaction between a particular seller and the buyer keeps track of the history, the seller doesn't necessarily know the interaction between the buyer and the other sellers. The only information the seller has comes when the buyer swipes the loyalty card, and the seller doesn't know if the buyer buys, from whom and at what price, in any period when he doesn't come to the seller.

*Department of Economics, UCL, 30 Gordon Street, London, UK m.cripps@ucl.ac.uk

†Department of Economics, UCL, 30 Gordon Street, London, UK suehyun.kwon@ucl.ac.uk

We consider a model of price competition and collusion with private monitoring. The sellers have homogeneous good, and the buyer has a unit demand every period. The buyer's valuation is his private information. Each period over an infinite horizon, the buyer decides whether to buy and from whom to buy. The seller decides what price to charge if the buyer comes to him. The seller only knows when the buyer comes to him and what price he charges. The buyer knows the price of the seller he visits. We characterize sequential equilibria of the repeated game, and we focus on equilibria that maximize the sellers' payoffs.

There are two strands of literature this paper is most closely related to. The first is models of price competition for a single buyer. The complete information benchmark of our model is closely related to Bergemann and Välimäki (2002). They consider a finite horizon model with differentiated goods, and they consider pure strategy sequential equilibria. In their model, the sellers know the buyer's valuation whereas in our model, the buyer's valuation is his private information. In Bergemann and Välimäki (2006), they consider Markovian equilibria of the finite and infinite horizon models where the stage game payoffs can depend on the past history. The main difference between their models and ours is that the buyer's valuation is his private information in our model.

Another strand of literature is that on collusion. The main difference from the existing literature is that we consider a long-run buyer who is strategic. Harrington and Skrzypacz (2011) consider collusion with private monitoring and communication. They consider semipublic perfect equilibrium whereas we consider sequential equilibria. The demand side in their model is given by a stochastic function and is not strategic. Other papers on collusion include Stigler (1964), Green and Porter (1984), Harrington and Skrzypacz (2007) and Athey and Bagwell (2001, 2008).

If the sellers know the buyer's valuation, for every payoff that is individually rational for the buyer and is below the cutoff, there exists an equilibrium in which the buyer's expected payoff equals that payoff. The upper bound on the buyer's payoff comes from the fact that the sellers can charge the maximum price in a given period and risk losing the buyer forever. The continuation values of the sellers have to be above the cutoff to follow the equilibrium strategy. Since the monitoring technology is private monitoring, the strategies are triggered by the interaction between the seller and the buyer, and in particular, the seller doesn't always learn that a deviation occurred between the other seller and the buyer.

When the buyer's valuation is his private information, the sellers can potentially learn the buyer's type over time. Given the prior on the types and the strategies of each type, the seller can update his posterior after the buyer comes or doesn't come. Furthermore, the sellers' beliefs can diverge off the equilibrium path; for example, if the buyer's strategy was to go to seller 1 but he didn't buy, then seller 1 believes they're off the equilibrium path while seller 2 believes they're on the equilibrium path.

Instead of characterizing the set of all equilibrium payoffs, we focus on equilibria that maximize the sellers' payoffs. First, we consider the relaxed problem for the sellers. The buyer can always mimic another type forever, and the sellers won't be able to detect the deviation. The relaxed problem with the IC for mimicking another type forever pins down the discounted sum of the probabilities the buyer buys from either of the sellers and the expected payment to maximize the sum of the sellers' payoffs. We can transform the problem to one that is analogous to the static adverse selection problem, and the sellers' payoffs are maximized when the sellers sell to the types above the cutoff with probability one every period and never sell to any other type.

The remaining step is to construct a sequential equilibrium with the above properties. The key step is to note that there always exists a sequential equilibrium in which the seller believes with probability one that the buyer has the highest valuation if the seller learns that the buyer has deviated. Given this belief, the seller charges the highest price forever after he learns a deviation, and the buyer prefers to mimic another type rather than to reveal that he has deviated. This reduces the set of ICs to those for mimicking another type forever, and the solution to the relaxed problem maximizes the sellers' payoffs in the full problem.

Another thing to note about the best equilibria for the sellers is that the prices and the probabilities that the buyer buys in the best equilibria coincide with the static optimum. This is because we only consider mimicking another type forever in the relaxed problem. The buyer decides whose strategy to follow in the beginning of the first period, and it is equivalent to the static problem where the buyer chooses a type to report. The best equilibria for the sellers are to repeat the static optimum every period.

The optimal selling strategy and the cutoff doesn't depend on the number of sellers in the market. Competition comes into play when we consider the incentives of the sellers. The continuation value of each seller should be sufficiently high to sustain the equilibrium pricing strategy; the worst punishment for the seller is to lose the buyer forever, and the continuation value on the equilibrium path should be higher than the deviation payoff in the current period. However, we can distribute the surplus between the sellers in anyway by allowing the buyer to choose one seller in the first period and buy from him forever.

The main contributions of this paper are in three folds. First, we consider sequential equilibria with both private monitoring and private information. The existing literature either hasn't allowed for all sequential equilibria or is without private information on the buyer side. Second, we show that with a continuum of types and a continuum of action space for the sellers, the best equilibria for the sellers is to repeat the static optimum. The equilibrium construction also shows the effect of competition on the payoffs. Lastly, the off-the-equilibrium path strategies and beliefs show what types of strategies one needs to

support a sequential equilibrium with private monitoring.

The rest of the paper is organized as follows. Section 2 describes the model, and the main results are presented in Section 3. Section 4 discusses extensions, and Section 5 concludes.

2 Model

There are two sellers and a buyer trading over an infinite horizon, $t = 1, 2, \dots$. The common discount factor is $\delta < 1$. The sellers have homogeneous indivisible good, and the buyer has a unit demand every period. At the beginning of each period, the buyer decides whether to buy and if so, the buyer decides whom to buy from. After the buyer decides, the sellers choose the prices from $[0, \bar{v}]$. If the buyer doesn't buy, everyone gets 0 in this period.¹ If the buyer buys from seller 1 (seller 2), then seller 1 (seller 2) gets the price he charges, seller 2 (seller 1) gets 0, and the buyer gets his valuation subtracted by the price.² The cost for the sellers is zero. There is a continuum of buyers, $v \in [\underline{v}, \bar{v}]$. The buyer's type is his private information, and at the beginning of period 1, the sellers start with a common prior F . F has positive density everywhere and satisfies the monotone hazard rate. Type v values the good at v , and every type values the good above the marginal cost, $\underline{v} \geq 0$.

The monitoring technology is private monitoring. Each seller only knows his own prices and whether the buyer buys from him in which period. The buyer knows his own action and the prices of the seller he buys from. The equilibrium concept is sequential equilibria.

The history of seller 1 is of the form $h^{t1} = (p_1^1, p_2^1, \dots)$ where p_t^1 is the price seller 1 charges in period t . If the buyer doesn't come in period t , then $p_t^1 = \emptyset$. The history of seller 2 is defined respectively. The history of the buyer is $h^{tb} = (d_1^b, p_1^b, d_2^b, \dots)$ where $d_t^b = 0$ if the buyer doesn't buy in period t , 1 if he buys from seller 1 and 2 if he buys from seller 2. p_t^b is the price the buyer bought at time t , and $p_t^b = \emptyset$ if he doesn't buy in period t . The strategy of seller 1 is a mapping $\sigma^1 : \mathcal{H}^1 \rightarrow [0, \bar{v}]$. The strategy of seller 2 is defined similarly. The strategy of the buyer is a mapping $\sigma^b : \mathcal{H}^b \rightarrow \{(\sigma_t^{b1}, \sigma_t^{b2}) | 0 \leq \sigma_t^{b1} + \sigma_t^{b2} \leq 1\}$. The buyer goes to seller 1 with probability σ_t^{b1} and to seller 2 with probability σ_t^{b2} .

Let v_0 be the solution to $v = \frac{1-F(v)}{f(v)}$ and $\hat{v} = \max\{\underline{v}, v_0\}$. Assume $(1 - \delta)(\bar{v} - \hat{v}) \leq \delta\hat{v}$: the seller prefers to charge \hat{v} every period than to charge \bar{v} and lose the buyer.

¹This is a normalization without loss of generality. If the buyer's outside option is \bar{u} , we can reformulate everything with respect to $v - \bar{u}$ where v is the valuation of the buyer. For the sellers, we can consider $p - \hat{u}$ where \hat{u} is the sellers' outside option.

²We assume that the buyer has to pay the price when he goes to seller i . The main results remain the same if we allow the buyer to decide whether to buy after observing the price of seller i .

3 Results

This section presents the main results. Throughout the paper, we focus on equilibria that maximize the sum of the sellers' payoffs. For the complete information benchmark, we focus on equilibria in which the buyer buys every period. Because there is gain from trade for every type of the buyer, it is efficient to trade every period.

In the complete information benchmark, we show that any payoff for the buyer between his outside option and the cutoff can be supported in an equilibrium. When there is more than one seller, competition affects the sellers' incentives to charge the equilibrium price, and the market shares of the sellers are restricted. The continuation value of each seller should be sufficiently high for him to charge the equilibrium price, and there is a lower bound on the expected payoff for each seller. When the buyer's type is unknown to the sellers, the sellers' payoffs are maximized when they repeat the static optimum every period. Similar to the complete information benchmark, competition limits the degree of collusion among the sellers and the market shares they can sustain.

3.1 Complete Information Benchmark

This section characterizes the set of efficient equilibrium payoffs when the sellers know the buyer's type. Since there is positive gain from trade, the buyer buys every period in an efficient equilibrium.

3.1.1 One Seller

When there is only one seller, the set of efficient equilibrium payoffs consists of payoffs such that the buyer's payoff is between his outside option and the upper bound. The upper bound comes from the fact that the worst punishment for the seller is to lose the buyer forever.

Proposition 1. *Suppose there is one seller and the seller knows the buyer's valuation. The set of payoffs where the buyer buys every period is $\{(W^B, W^S) | W^B + W^S = v, W^B \geq 0, W^S \geq (1 - \delta)\bar{v}\}$.*

Proof. Given $(1 - \delta)\bar{v} \leq p \leq v$, $(v - p, p)$ is an equilibrium payoff, and the following strategies form the equilibrium. The seller charges p every period on the equilibrium path, and the buyer buys every period. If the buyer deviates, the seller charges \bar{v} in all future periods, and the buyer never buys again. If the seller deviates, he charges \bar{v} in all future periods, and the buyer never buys in the future. If there has been a deviation and the seller hasn't charged \bar{v} , his strategy is to charge \bar{v} from this period on. If there has been a deviation and the buyer bought after the deviation, his strategy is never to buy in the future.

If the seller has deviated before, regardless of the price he charges this period, the buyer's strategy is never to buy again. The seller doesn't have incentives to lower the price. If the seller has deviated before and the buyer buys, he gets $v - \bar{v}$ in the current period, and the buyer weakly prefers not to buy. If the seller has never deviated before but the buyer has, then the seller charges \bar{v} in all future periods. The buyer weakly prefers not to buy, and the seller has no incentives to lower the price. If no one has deviated before, the seller doesn't deviate from price p because $(1 - \delta)(\bar{v} - p) \leq \delta p$. The buyer gets non-negative payoff on the equilibrium path and zero off the equilibrium path, and he doesn't deviate. The minimum payoff the seller can guarantee himself when the buyer buys every period is $(1 - \delta)\bar{v}$ when he charges \bar{v} this period. \square

The equilibrium is constructed by trigger strategies. There is no trade after a party deviates, and given that the buyer never buys again, the seller charges the maximum price; given that the seller charges the maximum price, the buyer never buys. The worst punishment for the seller is that the buyer never buys again, and the seller's payoff from charging the maximum price this period and losing the buyer sets the lower bound on his payoff.

3.1.2 Two Sellers

When there are two sellers, the set of the payoffs of the buyer coincides with that in the one-seller case, but the sellers' payoffs can be distributed in different ways. The intuition that the buyer's payoff is the same is that we can always construct an equilibrium in which the buyer buys from one seller every period. By the same reasoning as in the one-seller case, if the buyer comes to a seller with a positive probability this period, the expected payoff of the seller has to be above $(1 - \delta)\bar{v}$ for it to be the equilibrium strategy.

Proposition 2. *Suppose there are two sellers, and the sellers know that the buyer's valuation is v . The set of payoffs where the buyer buys every period is*

$$S = \{(W^B, W_1^S, W_2^S) | W^B + W_1^S + W_2^S = v, W^B, W_1^S, W_2^S \geq 0, v - W^B \geq (1 - \delta)\bar{v}\}.$$

Proof. Given (W^B, W_1^S, W_2^S) in S , we can find $\sigma_1, \sigma_2 \in [0, 1]$ such that $W_1^S = \sigma_1(v - W^B)$, $W_2^S = \sigma_2(v - W^B)$. Consider the following strategies:

1. With probability σ_1 , the buyer goes to seller 1 every period. With probability σ_2 , the buyer goes to seller 2 every period. Once the buyer goes to either of the sellers in the first period, he goes to the same seller every period.
2. If there has been no deviation in the seller's private history, he charges $v - W^B$ every period.

3. If the buyer went to seller i , and seller i deviates, the buyer substitutes not buying with buying from seller i in all future periods. If the buyer comes to seller i after a deviation, he charges \bar{v} in all future periods.
4. If the buyer was supposed to come to seller i but doesn't, seller i charges \bar{v} in all future periods.
5. If the buyer wasn't supposed to come to seller i but does, seller i charges \bar{v} in all future periods.
6. If the buyer has deviated and the seller hasn't charged \bar{v} , his strategy is to charge \bar{v} from this period.
7. If the seller has deviated and the buyer has bought after a deviation, the buyer's strategy is never to buy again.

Since the payoffs are bounded, we have continuity at infinity. Suppose there have been a finite number of deviations by all three players. If both sellers' private histories contain deviations, by either the buyer or the seller, and they don't deviate any more, then the sellers charge \bar{v} in all future periods. If the buyer comes to seller i , the buyer never buys from the next period; the sellers don't have incentives to lower the price. The buyer faces the same price from both sellers, and he weakly prefers not buying.

Suppose only seller i 's private history contains a deviation. If the buyer goes to seller i , seller i charges \bar{v} . If the buyer goes to seller j , seller j 's history now has a deviation, and he charges \bar{v} . The buyer faces a weakly lower payoff than not buying in every period, and the buyer prefers not buying. Since the buyer's strategy is never to buy again, the sellers don't have incentives to lower the price.

Suppose no player has deviated before. If the buyer goes to seller i when he's supposed to go to seller j , then both sellers start charging \bar{v} in all future periods, and the buyer is weakly worse off. If the buyer was supposed to buy from seller i but didn't buy, his payoff in all future periods is weakly lower. If the buyer wasn't supposed to buy but goes to seller i , then again, his payoff in all future periods is weakly lower. If a seller deviates, the buyer never buys again, and from $(1 - \delta)\bar{v} \leq v - W^B$, the sellers don't have incentives to deviate.

In the first period, the buyer faces the same price sequence from both sellers and is indifferent between the two.

Therefore, the given strategies form an equilibrium.

On the contrary, for any (W^B, W_1^S, W_2^S) in S , we have $W_1^S, W_2^S \leq v - W^B$. When the buyer comes to seller 1 (seller 2) with a positive probability this period, the seller can always charge \bar{v} , and the maximum punishment is to lose the buyer ever. A necessary

condition for W_1^S to be an equilibrium payoff is $(1 - \delta)\bar{v} \leq W_1^S$. Therefore, it is necessary that $(1 - \delta)\bar{v} \leq v - W^B$.

The set of all efficient equilibrium payoffs is S . □

The equilibrium strategies have the following two properties. Once the buyer deviates, the seller charges the maximum price forever. If a seller deviates, the buyer never buys again. These are the worst punishments one can impose on another, and they are triggered by the private history of each player. Once a deviation occurs, the history after the deviation doesn't matter in a sense that even if there have been multiple deviations, the players' strategies dictate to go back to the strategy after the first deviation in the next instance. These properties allow the strategies to generalize to more-than-two-sellers case.

If the buyer visits a seller, the sellers' payoff is bounded from below by $(1 - \delta)\bar{v}$. But we can generate the set of all efficient equilibrium payoffs by equilibria in which the buyer mixes between the two sellers only in the first period. In those equilibria, once the buyer visits a seller, the seller captures all the surplus from trade, and the seller's IC is most relaxed.

3.2 Incomplete Information

This section characterizes the set of equilibria that maximize the sum of the sellers' payoffs. Both in the one-seller case and the two-sellers case, the best equilibria for the sellers are to repeat the static optimum every period. The multiplicity of equilibria comes from the distribution of surplus between the two sellers. In the static optimum, the sellers sell to types above the cutoff with probability one and don't sell to types below the cutoff.

3.2.1 One Seller

This section characterizes the best equilibrium for the seller when there is only one seller.

Proposition 3. *Suppose the seller's payoff is maximized in an equilibrium. There exists \hat{v} such that the buyer's payoff is $v - \hat{v}$ if $v \geq \hat{v}$, 0 otherwise, and the seller's payoff is $(1 - F(\hat{v}))\hat{v}$. If the solution to*

$$v = \frac{1 - F(v)}{f(v)}$$

is in the support, \hat{v} is the solution to the above equation. If not, $\hat{v} = \underline{v}$. The buyer with $v \geq \hat{v}$ buys every period, and other types never buy. The expected payment is \hat{v} for all $v \geq \hat{v}$.

Proof. Let λ_v and T_v be the discounted sum of the probabilities type v buys and the discounted sum of the payments he makes. One of the IC constraints is that type v and v'

don't mimic each other forever:

$$\begin{aligned}\lambda_v v - T_v &\geq \lambda_{v'} v - T_{v'}, \\ \lambda_{v'} v' - T_{v'} &\geq \lambda_v v' - T_v.\end{aligned}$$

We can combine the two IC constraints as $(\lambda_v - \lambda_{v'})v \geq T_v - T_{v'} \geq (\lambda_v - \lambda_{v'})v'$, and the local ICs are sufficient:

$$\lim_{v' \rightarrow v} (\lambda_v - \lambda_{v'})v = T_v - T_{v'}.$$

The relaxed problem with IC constraints that type v doesn't mimic v' forever is as follows:

$$\begin{aligned}\max & \int T_v dF(v) \\ \text{s.t.} & \lambda_v \in [0, 1] \forall v, \\ \text{IC} & : \lim_{v' \rightarrow v} (\lambda_v - \lambda_{v'})v = T_v - T_{v'}, \\ \text{IR} & : \lambda_{\underline{v}} \underline{v} - T_{\underline{v}} = 0.\end{aligned}$$

From the IC constraint, we get

$$\begin{aligned}\int T_v dF(v) &= \int (v\lambda_v - \int_{\underline{v}}^v \lambda_z dz) dF(v) \\ &= \int (v\lambda_v - \frac{1 - F(v)}{f(v)} \lambda_v) dF(v).\end{aligned}$$

Since λ_v is bounded, we get 0 or 1 as the solution, and from the monotone hazard rate condition, there exists \hat{v} such that $\lambda_v = 1$ if and only if $v \geq \hat{v}$. $\lambda_v = 0$ otherwise.

The next step is to show that there exists an equilibrium of the full problem with the same property. Suppose the seller always charges \hat{v} on the equilibrium path and the buyer buys every period if $v \geq \hat{v}$. The buyer never buys if $v < \hat{v}$. If the seller learns that the buyer has deviated before, the seller believes that the buyer's valuation is \bar{v} and charges \bar{v} every period. The buyer never buys again. If the seller deviates, the buyer believes that the seller will charge \bar{v} every period in the future; the buyer doesn't buy. If there has been a deviation but the seller hasn't charged \bar{v} , his strategy is to charge \bar{v} from this period and on. If the seller deviated but the buyer has bought after the deviation, his strategy is never to buy again.

If the seller has deviated, then regardless of his price, the buyer's strategy is never to buy again. The seller has no incentives to lower the price. When the seller charges \bar{v} , the buyer weakly prefers not to buy. If the seller has never deviated but the buyer has, there are two cases. If the seller learns the deviation, he charges \bar{v} every period, and the buyer weakly

prefers not to buy. The buyer never buys again, and the seller doesn't have incentives to lower the price. If the seller doesn't learn the deviation, he keeps charging \hat{v} ; the buyer weakly prefers not to reveal that he has deviated before. If no one has deviated before, the buyer gets non-negative payoff on the equilibrium path and doesn't deviate. In particular, given the strategies, when the buyer deviates, the seller has to learn the deviation, and the buyer prefers to stay on the equilibrium path. The seller prefers to charge \hat{v} instead of charging \bar{v} by the assumption $(1 - \delta)(\bar{v} - \hat{v}) \leq \delta\hat{v}$.

Given the strategies, the buyer never wants to reveal that he has deviated in the past, and the IC constraints in the relaxed problem are sufficient. \square

Constructing the best equilibrium consists of the following steps. First, we find the solution to the mechanism design problem corresponding to the relaxed problem of the seller. Instead of considering all IC constraints, we consider the problem when the buyer decides which type to mimic and follow the strategy over the infinite horizon. We can transform the problem into an analogue of the static adverse selection problem with bounded support. The supports are bounded because the discounted sum of the probabilities has to be between 0 and 1. Once we find the solution to the relaxed problem, we construct the equilibrium strategies so that all other IC constraints are satisfied. There exists an equilibrium in which the seller believes that the buyer's type is \bar{v} if he catches a deviation. The main step is to construct the trigger strategies so that the buyer never wants to reveal that he has deviated in the past.

3.2.2 Two Sellers

We characterize the best equilibria for the sellers when there are two sellers. The first step is similar to the one-seller case. We consider the relaxed problem with the IC constraints for mimicking another type forever. The solution coincides with the static optimum. Then we construct the equilibrium strategies so that the solution to the relaxed problem is an equilibrium.

Proposition 4. *Suppose the sum of the sellers' payoffs is maximized in an equilibrium. There exists \hat{v} such that the buyer's payoff is $v - \hat{v}$ if $v \geq \hat{v}$, 0 otherwise. The set of the sellers' payoffs is given by $\{(W_1^S, W_2^S) | W_1^S, W_2^S \geq 0, W_1^S + W_2^S = (1 - F(\hat{v}))\hat{v}\}$. If the solution to*

$$v = \frac{1 - F(v)}{f(v)}$$

is in the support, \hat{v} is the solution to the above equation. If not, $\hat{v} = \underline{v}$. The buyer with $v \geq \hat{v}$ buys every period, and other types never buy. The expected payment is \hat{v} for all $v \geq \hat{v}$.

Proof. Let λ_v and T_v be the discounted sum of the probabilities type v buys from either of the sellers and the discounted sum of the payments he makes. One of the IC constraints is that type v and v' don't mimic each other forever:

$$\begin{aligned}\lambda_v v - T_v &\geq \lambda_{v'} v - T_{v'}, \\ \lambda_{v'} v' - T_{v'} &\geq \lambda_v v' - T_v.\end{aligned}$$

We can combine the two IC constraints as $(\lambda_v - \lambda_{v'})v \geq T_v - T_{v'} \geq (\lambda_v - \lambda_{v'})v'$, and the local ICs are sufficient:

$$\lim_{v' \rightarrow v} (\lambda_v - \lambda_{v'})v = T_v - T_{v'}.$$

The relaxed problem with IC constraints that type v doesn't mimic v' forever is as follows:

$$\begin{aligned}\max & \int T_v dF(v) \\ \text{s.t.} & \lambda_v \in [0, 1] \quad \forall v, \\ \text{IC} & : \lim_{v' \rightarrow v} (\lambda_v - \lambda_{v'})v = T_v - T_{v'}, \\ \text{IR} & : \lambda_v \underline{v} - T_v = 0.\end{aligned}$$

From the IC constraint, we get

$$\begin{aligned}\int T_v dF(v) &= \int (v\lambda_v - \int_{\underline{v}}^v \lambda_z dz) dF(v) \\ &= \int (v\lambda_v - \frac{1 - F(v)}{f(v)} \lambda_v) dF(v).\end{aligned}$$

Since λ_v is bounded, we get 0 or 1 as the solution, and from the monotone hazard rate condition, there exists \hat{v} such that $\lambda_v = 1$ if and only if $v \geq \hat{v}$. $\lambda_v = 0$ otherwise.

The next step is to show that there exists an equilibrium of the full problem with the same property. Suppose the sellers always charge \hat{v} on the equilibrium path and the buyer buys every period if $v \geq \hat{v}$. The buyer never buys if $v < \hat{v}$. If a seller learns that the buyer has deviated before, the seller believes that the buyer's valuation is \bar{v} and charges \bar{v} every period. The buyer never buys again from that seller, and the buyer substitutes not buying with buying from that seller. If a seller deviates, the buyer believes that the seller will charge \bar{v} every period in the future; the buyer doesn't buy from that seller again. If there has been a deviation but the seller hasn't charged \bar{v} , his strategy is to charge \bar{v} from this period and on. If the seller deviated but the buyer has bought after the deviation, his strategy is never to buy from that seller again.

If the seller has deviated, then regardless of his price, the buyer's strategy is never to

buy again from that seller. The seller has no incentives to lower the price. When the seller charges \bar{v} , the buyer weakly prefers not to buy than to buy from that seller. If the seller has never deviated but the buyer has, there are two cases. If the seller learns the deviation, he charges \bar{v} every period, and the buyer weakly prefers not to buy. The buyer never buys again, and the seller doesn't have incentives to lower the price. If the seller doesn't learn the deviation, he keeps charging \hat{v} ; the buyer weakly prefers not to reveal he has deviated before. If no one has deviated before, the buyer gets non-negative payoff on the equilibrium path and doesn't deviate. In particular, given the strategies, when the buyer deviates, the seller has to learn the deviation, and the buyer prefers to stay on the equilibrium path. The seller prefers to charge \hat{v} instead of charging \bar{v} by the assumption $(1 - \delta)(\bar{v} - \hat{v}) \leq \delta\hat{v}$.

Given the strategies, the buyer never wants to reveal that he has deviated in the past, and the IC constraints in the relaxed problem are sufficient.

In the first period, the buyer faces the same price sequence from both sellers and is indifferent. The buyer can randomize between the sellers with any probability $\sigma \in [0, 1]$, and the sellers' payoffs can be distributed in anyway between the two. \square

In the mechanism design problem, we consider the sum of the sellers' payoffs, and the relaxed problem doesn't depend on the number of sellers. The maximum payoff for the sellers is attained when the buyer buys if and only if his type is above the cutoff. The number of sellers, and competition, matters when we construct the equilibrium for the solution of the relaxed problem. In order for the sellers to charge the equilibrium price, their expected payoff should be weakly higher than what they can get by charging the maximum price this period. This puts the lower bound on the payoff of the seller the buyer visits with a positive probability this period. The IC constraint of the seller is most relaxed when the seller captures all future surplus from the buyer, i.e., if the buyer keeps buying from the same seller every period. We can distribute the surplus between the sellers in anyway by letting the buyer pick a seller in the first period and buy from him every period.

We should point out that the best equilibria are sustained by the off-the-equilibrium-path belief that if the seller learns that the buyer has deviated before, he believes that the buyer's type is \bar{v} with probability one. Given the beliefs, the sellers are willing to charge \bar{v} off the equilibrium path, and the buyer never buys from them again.

The technical contribution of the paper is to construct equilibrium strategies such that the solution to the relaxed problem can be sustained as an equilibrium under private monitoring. But because it is private monitoring, the strategies generalize to any number of sellers, as long as the seller has incentives to charge the equilibrium price.

4 Discussion

4.1 More than two sellers

The equilibrium strategies in Section 3 are triggered by a deviation in the seller-buyer relationship. When the buyer deviates, the seller charges \bar{v} in the future, and if the seller deviates, the buyer substitutes not buying with going to that seller. The strategies don't depend on the total number of sellers, and the results generalize to more than two-sellers case.

For the sellers to charge \hat{v} , we need the continuation values of the sellers to be above $(1 - \delta)\bar{v}$. We can distribute the surplus in anyway by letting the buyer pick a seller in the first period, but otherwise, the market sharing is restricted by the seller's incentives to charge \bar{v} .

Proposition 5. *Suppose there are n sellers and the sum of the sellers' payoffs is maximized in an equilibrium. There exists \hat{v} such that the buyer's payoff is $v - \hat{v}$ if $v \geq \hat{v}$, 0 otherwise. The set of the sellers' payoffs is given by $\{(W_1^S, \dots, W_n^S) | W_1^S, \dots, W_n^S \geq 0, \sum W_i^S = (1 - F(\hat{v}))\hat{v}\}$. If the solution to*

$$v = \frac{1 - F(v)}{f(v)}$$

is in the support, \hat{v} is the solution to the above equation. If not, $\hat{v} = \underline{v}$. The buyer with $v \geq \hat{v}$ buys every period, and other types never buy. The expected payment is \hat{v} for all $v \geq \hat{v}$.

The proof is same as the proof of Proposition 4 and is in the appendix.

4.2 If the buyer decides after seeing the price

In Section 2, the buyer can choose from not buying and buying from either of the sellers. An alternative model is to allow the buyer to go to a seller and decide whether to buy after observing the price. The main results generalize to this environment. The first part of the proofs for the incomplete information case doesn't depend on when the buyer decides to buy. The optimal mechanism for the sellers is to sell at \hat{v} every period. When the buyer has to go to a seller every period, charging 0 and buying every period is an equilibrium. From the previous equilibrium, if we assume instead that the buyer and the seller go to the marginal cost equilibrium after the seller deviates, the new strategies form an equilibrium of the alternative model. The only difference is the condition to ensure that the seller will charge \hat{v} . The condition is slightly weaker than $(1 - \delta)(\bar{v} - \hat{v}) \leq \delta\hat{v}$, and the distribution of the surplus among the sellers is easier in the alternative model.

5 Conclusion

We developed a model of collusion and price competition when the monitoring technology is private monitoring. When the sellers know the buyer's valuation, the set of efficient equilibrium payoffs consists of all payoffs that add up to the buyer's valuation. We construct equilibria in which the buyer mixes among the sellers in the first period and goes to the same seller every period after the first period. This relaxes the seller's IC constraint to the same constraint as in the one-seller case, and the set of payoffs is convex because there is no restriction on the mixing probabilities. The equilibrium is sustained by trigger strategies, and the players go to the worst equilibrium after one player deviates. However, other sellers don't necessarily learn that a deviation has occurred between a particular seller and the buyer.

When the sellers don't know the buyer's valuation, we focus on the best equilibria for the sellers. The sum of the sellers' payoffs is maximized when the sellers repeat the static optimum. To characterize the best equilibria, we first consider the relaxed problem in which the buyer decides on a type and mimics that type forever. Since the buyer chooses a type in the first period, this is equivalent to the static adverse selection with reporting types. The solution to the relaxed problem is pinned down by the hazard rate, and the sellers sell with probability one to the types with positive virtual surplus and never sell to the types with negative virtual surplus.

The next step is to construct a sequential equilibrium that matches the solution to the relaxed problem. We can construct a sequential equilibrium in which the seller believes that the buyer's type is the highest when he learns that the buyer has deviated before. Then, the seller extracts all the surplus from the buyer off the equilibrium path, and given the deviation strategy, the buyer prefers not to reveal that he has deviated before. If the buyer never reveals that he has deviated before, the only relevant constraints are to mimic another type forever, and the solution to the relaxed problem can be sustained as a sequential equilibrium.

When the buyer comes to a seller, the seller can deviate and charge the maximum price. The worst punishment for the seller is to lose the buyer forever, and this puts a lower bound on the seller's payoff. The IC constraint of the seller is most relaxed when the buyer picks a seller in the first period and goes to the same seller every period. This is how we constructed the set of payoffs. If the sellers are to share the market in the future, their continuation values are lower than the total surplus from trade, and the number of sellers in the market is limited.

We have assumed so far that the buyer has to visit the seller to see the price. One extension is to allow the buyer to observe both prices and decide whether to buy and whom

to buy from. Another extension is to allow for multi-unit demand. We have also focused on the best equilibria for the sellers, and the characterization of all sequential equilibria remains to be done.

A Proofs

Proof of Proposition 5. Let λ_v and T_v be the discounted sum of the probabilities type v buys from either of the sellers and the discounted sum of the payments he makes. One of the IC constraints is that type v and v' don't mimic each other forever:

$$\begin{aligned}\lambda_v v - T_v &\geq \lambda_{v'} v - T_{v'}, \\ \lambda_{v'} v' - T_{v'} &\geq \lambda_v v' - T_v.\end{aligned}$$

We can combine the two IC constraints as $(\lambda_v - \lambda_{v'})v \geq T_v - T_{v'} \geq (\lambda_v - \lambda_{v'})v'$, and the local ICs are sufficient:

$$\lim_{v' \rightarrow v} (\lambda_v - \lambda_{v'})v = T_v - T_{v'}.$$

The relaxed problem with IC constraints that type v doesn't mimic v' forever is as follows:

$$\begin{aligned}\max & \int T_v dF(v) \\ \text{s.t.} & \lambda_v \in [0, 1] \forall v, \\ \text{IC} & : \lim_{v' \rightarrow v} (\lambda_v - \lambda_{v'})v = T_v - T_{v'}, \\ \text{IR} & : \lambda_v \underline{v} - T_v = 0.\end{aligned}$$

From the IC constraint, we get

$$\begin{aligned}\int T_v dF(v) &= \int (v\lambda_v - \int_{\underline{v}}^v \lambda_z dz) dF(v) \\ &= \int (v\lambda_v - \frac{1 - F(v)}{f(v)} \lambda_v) dF(v).\end{aligned}$$

Since λ_v is bounded, we get 0 or 1 as the solution, and from the monotone hazard rate condition, there exists \hat{v} such that $\lambda_v = 1$ if and only if $v \geq \hat{v}$. $\lambda_v = 0$ otherwise.

The next step is to show that there exists an equilibrium of the full problem with the same property. Suppose the sellers always charge \hat{v} on the equilibrium path and the buyer buys every period if $v \geq \hat{v}$. The buyer never buys if $v < \hat{v}$. If a seller learns that the buyer has deviated before, the seller believes that the buyer's valuation is \bar{v} and charges \bar{v} every period. The buyer never buys again from that seller, and the buyer substitutes not

buying with buying from that seller. If a seller deviates, the buyer believes that the seller will charge \bar{v} every period in the future; the buyer doesn't buy from that seller again. If there has been a deviation but the seller hasn't charged \bar{v} , his strategy is to charge \bar{v} from this period and on. If the seller deviated but the buyer has bought after the deviation, his strategy is never to buy from that seller again.

If the seller has deviated, then regardless of his price, the buyer's strategy is never to buy again from that seller. The seller has no incentives to lower the price. When the seller charges \bar{v} , the buyer weakly prefers not to buy than to buy from that seller. If the seller has never deviated but the buyer has, there are two cases. If the seller learns the deviation, he charges \bar{v} every period, and the buyer weakly prefers not to buy. The buyer never buys again, and the seller doesn't have incentives to lower the price. If the seller doesn't learn the deviation, he keeps charging \hat{v} ; the buyer weakly prefers not to reveal he has deviated before. If no one has deviated before, the buyer gets non-negative payoff on the equilibrium path and doesn't deviate. In particular, given the strategies, when the buyer deviates, the seller has to learn the deviation, and the buyer prefers to stay on the equilibrium path. The seller prefers to charge \hat{v} instead of charging \bar{v} by the assumption $(1 - \delta)(\bar{v} - \hat{v}) \leq \delta\hat{v}$.

Given the strategies, the buyer never wants to reveal that he has deviated in the past, and the IC constraints in the relaxed problem are sufficient.

In the first period, the buyer faces the same price sequence from both sellers and is indifferent. The buyer can randomize between the sellers with any probability $\sigma \in [0, 1]$, and the sellers' payoffs can be distributed in anyway between the two. \square

References

- [1] Athey, Susan, and Kyle Bagwell. "Optimal collusion with private information." *RAND Journal of Economics* 32 (2001): 428-465.
- [2] Athey, Susan, and Kyle Bagwell. "Collusion with Persistent Cost Shocks." *Econometrica* 76 (2008): 493-540.
- [3] Bergemann, Dirk, and Juuso Välimäki. "Information acquisition and efficient mechanism design." *Econometrica* 70 (2002): 1007-1033.
- [4] Bergemann, Dirk, and Juuso Välimäki. "Dynamic price competition." *Journal of Economic Theory* 127 (2006): 232-263.
- [5] Green, Edward J., and Robert H. Porter. "Noncooperative collusion under imperfect price information." *Econometrica* 52 (1984): 87-100.

- [6] Harrington, Joseph E., and Andrzej Skrzypacz. "Collusion under monitoring of sales." *RAND Journal of Economics* 38 (2007): 314-331.
- [7] Harrington, Joseph E., and Andrzej Skrzypacz. "Private monitoring and communication in cartels: Explaining recent collusive practices." *American Economic Review* 101 (2011): 2425-2449.
- [8] Stigler, George J. "A theory of oligopoly." *Journal of Political Economy* 72 (1964): 44-61.