

On Communication and Collusion[†]

By YU AWAYA AND VIJAY KRISHNA*

We study the role of communication within a cartel. Our analysis is carried out in Stigler's (1964) model of repeated oligopoly with secret price cuts. Firms observe neither the prices nor the sales of their rivals. For a fixed discount factor, we identify conditions under which there are equilibria with "cheap talk" that result in near-perfect collusion, whereas all equilibria without such communication are bounded away from this outcome. In our model, communication improves monitoring and leads to higher prices and profits. (JEL C73, D43, D83, L12, L13, L25)

Antitrust authorities view any inter-firm communication with concern, but despite this, even well-established cartels meet regularly. In their landmark study of a sugar-refining cartel, Genesove and Mullin (2001) write that cartel members met weekly for almost a decade! Harrington (2006) too reports on the frequent meetings of cartels in many industries: Vitamin A and E (weekly or quarterly), citric acid (monthly), and lysine (monthly). Why do cartels find it necessary to meet so often? Court documents and insider accounts reveal that many issues are discussed at such gatherings: demand conditions, costs, and yes, sometimes even prices. But an important reason firms meet is to *monitor* each other's compliance with the collusive agreement. This is usually done by exchanging sales data, allowing firms to check that their market shares are in line with the agreement. Members of the isostatic-graphite cartel entered their sales figures on a pocket calculator that was passed around the table. Only the total was displayed, allowing each firm to check its own market share while keeping individual sales confidential. Often data are exchanged not in face-to-face meetings but indirectly through third parties, like trade associations, consulting firms, or statistical bureaus. The copper plumbing tubes cartel reported individual sales figures to the World Bureau of Metal Statistics which then disseminated these in aggregate form. Many other cartels—for instance, the amino acids and zinc phosphate cartels—used trade associations for this purpose.¹

*Awaya: Department of Economics, University of Rochester, 280 Hutchison Road, Rochester, NY 14627 (e-mail: yuawaya@gmail.com); Krishna: Department of Economics, Penn State University, 516 Kern Building, University Park, PA 16802 (e-mail: vkrishna@psu.edu). We thank Olivier Compte, Joyee Deb, Ed Green, and Satoru Takahashi for helpful conversations and various seminar audiences for useful comments. Joseph Harrington and Robert Marshall provided detailed comments on an earlier draft and enhanced our understanding of antitrust issues. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Go to <http://dx.doi.org/10.1257/aer.20141469> to visit the article page for additional materials and author disclosure statement(s).

¹Details of these and other cartel cases can be found in Harrington (2006) and Marshall and Marx (2012).

It is clear that in these cases cartel members cannot ascertain each other's sales directly; otherwise, there would be no need for firms to report these figures. Self-reported sales numbers, however, are difficult to verify and firms have the incentive to underreport these. Recognizing this, the lysine cartel even resorted to auditing its members. This was only moderately successful; a leading firm, Ajinomoto, was still able to hide its actual sales. Given that the data provided by firms cannot be verified, what is the point of having them report these? Can such "cheap-talk" communication really help firms monitor each other?

In this paper, we study how unverifiable communication about past sales can indeed facilitate collusion. We adopt, and adapt, Stigler's (1964) classic oligopoly model of secret price cuts in a repeated setting; firms cannot observe each other's prices nor can they observe each other's sales. Each firm observes only its own sales and because of demand shocks, both common and idiosyncratic, these are *noisy* signals of other firms' actions. Absent any communication, such imperfect monitoring limits the potential for collusive behavior because—and this was Stigler's point—deviations from a collusive agreement cannot be detected with confidence. We show, however, that if firms are able to exchange sales information, then collusive outcomes can be sustained. We study situations where the correlation in firms' sales is *sensitive* to prices. Precisely, it is high when the difference in firms' prices is small—say, when both firms charge close to monopoly prices—and decreases when the difference is large: say, because of a unilateral price cut.² For analytic convenience, we suppose that the relationship between sales and prices is governed by log-normal distributions, independently over time.

We show that in the statistical environment outlined above, reports of past sales allow firms to better monitor each other. The information communicated is not only *unverifiable* but also *payoff irrelevant*—it has no direct effects on current or future profits. Our main result is³

THEOREM 0: *For any high but fixed discount factor, when the monitoring is noisy but sensitive enough, there is an equilibrium with communication whose profits are strictly greater than those from any equilibrium without communication.*

The argument underlying our main result is divided into two steps. The first task is to find an effective bound for the maximum equilibrium profits that can be achieved in the absence of any communication. In Proposition 1 we develop a bound on equilibrium profits by using a very simple *necessary* condition—a deviating strategy in which a firm *permanently* cuts its price to an unchanging level should not be profitable. This deviation is, of course, rather naïve; the deviating firm does not take into account what the other firm knows or does. We show, however, that even this minimal requirement can provide an effective bound when the relationship between prices and sales is rather noisy relative to the discount factor. For a fixed discount factor, as sales become increasingly noisy, the bound becomes tighter.

²This can arise quite naturally, for example, in a Hotelling-type model with random transport costs (see Section I).

³A formal statement of the result is in Section IV.

The second task is to show that the bound developed earlier can be exceeded with communication. We directly construct an equilibrium in which firms exchange (coarse) sales reports in every period (see Proposition 2). Firms charge monopoly prices and report their sales truthfully. As long as the reported sales of the firms—whether high or low—are similar, monopoly prices are maintained. If the reported sales are significantly different—say, one firm reports high sales while the other reports low, a price war is triggered. These strategies form an equilibrium because when firms charge monopoly prices, their sales are highly correlated and so the likelihood that their truthful reports will agree is also high. If a firm cuts its price, sales become less correlated and so it cannot accurately predict its rival's sales. Even if the deviating firm strategically tailors its report, the likelihood of an agreement is low. Thus a strategy in which differing sales reports lead to noncooperation is an effective deterrent. In this way, cheap-talk communication allows firms to monitor each other more effectively.

We emphasize that the analysis in this paper is of a different nature than that underlying the so-called “folk theorems” (see Mailath and Samuelson 2006). These show that for a fixed monitoring structure, as players become increasingly patient, near-perfect collusion can be achieved in equilibrium. In this paper, we keep the discount factor fixed and change the monitoring structure in a way that communication improves profits.

Real-World Cartels.—Some key features of our model echo aspects of real-world cartel behavior. First, monitoring is imperfect and difficult. Second, the purpose of communication is to overcome monitoring difficulties and it is unverifiable. Third, communication allows firms to base their future behavior on relative sales—market shares—rather than the absolute level of sales.

The difficulties firms face in monitoring each other have been well-documented. Secret price cuts are hard to track. In the famous Trenton Potteries case, first-quality products were invoiced as being of second quality and so sold at a discount. In the sugar refining case, a firm included some shipping expenses in the quoted price. In each case, other members of the cartel discovered this “price cut” only with difficulty.

Clark and Houde (2014) study the internal functioning of a gasoline distribution cartel via recordings of telephone conversations. They use a “natural experiment”—a public announcement of an investigation that led to a cessation of all communication—to deduce that higher prices and profits result when firms communicate. Clark and Houde (2014) find that the communication helped the members to both coordinate pricing and to monitor each other. Harrington (2006) and Marshall and Marx (2012) study numerous cartels and how these functioned. Again, regular communication for monitoring purposes seems to be a common feature of cartels in many industries. Often the verity of what is communicated must be inferred. Indeed, Genesove and Mullin (2001, p. 389, italics added) write that the meetings of the sugar-refining cartel often served “as a court in which an accused firm might prove its innocence, in some cases on *factual*, in others on *logical*, grounds.” Harrington (2006) writes that the monitoring of sales is even more important than the monitoring of prices. As mentioned above, firms are often informed only about total industry sales—via a calculator or a trade association—because these are sufficient for

firms to calculate their own market shares. In our model too, a price war is triggered by relative sales (market shares) and not by the absolute level of sales. If all firms experience low sales—and so shares are not too far off—then this is attributed to adverse market conditions and the firms continue to collude. Implicitly, a firm that experiences low sales cares about the *reason* that its sales are low and not just the *fact* that they are low. Genesove and Mullin (2001) recount an episode in which the president of a Western firm threatened a price war unless it could be convinced that its low sales were not caused by the actions of Eastern refiners. Notice that this sort of behavior would not be possible without communication. A firm that only knew its own sales would be unable to infer the reason why they were low.

Other Models of Collusion.—The application of the theory of repeated games to understand collusion goes back at least to Friedman (1971). His model assumes that all past actions are commonly observed without any noise. Under such *perfect* monitoring, given any fixed discount factor, the set of subgame perfect equilibrium payoffs with and without communication is the same—there is nothing useful to communicate. Collusion with imperfect monitoring was studied by Green and Porter (1984) in a model of repeated quantity competition with noisy demand. Monitoring is imperfect because firms do not observe each others' output but rather only the market price. Since all firms observe the same noisy signal (the common price), this situation is referred to as one of *public* monitoring. Again given any fixed discount factor, the set of profits from (pure strategy) public perfect equilibria—in which firms' strategies depend only on the history of past prices—with and without communication is also the same. Green and Porter (1984) show that with sufficiently low discounting, public perfect equilibria are enough to sustain collusion. Our model is one of *private* monitoring; different firms observe different noisy signals (their own sales). As we show, in this environment, communication enlarges the set of equilibria. In both models, the firms are necessarily subject to a "type II" error when colluding; price wars are triggered even if no one has deviated. In the Green and Porter (1984) model, firms cannot distinguish between a deviation and an aggregate market shock; both lead to lower profits and result in price wars. In our model, communication allows firms to distinguish between the two; one leads to lower sales for all while the other leads to a discrepancy in sales and market shares. As mentioned above, only the latter triggers a punishment.

In our model, communication improves monitoring and, by enhancing the possibility of collusion, reduces welfare. Of course, not all communication need be welfare reducing. It can play a benign, even beneficial, role in some circumstances by allowing firms to share economically valuable information about demand or cost conditions. Indeed, the courts have long recognized this possibility and opined that such information sharing "... can hardly be deemed a restraint of commerce ... or in any respect unlawful."⁴ Economists and legal scholars have elaborated this argument (Shapiro 1986 and Carlton, Gertner, and Rosenfield 1996). The sharing of privately known costs via cheap talk is studied in a repeated-game model by Athey and Bagwell (2001). They show that, for moderate discounting, cheap-talk

⁴Areeda and Kaplow (1988, pp. 337–38).

communication increases profits. The channel through which this occurs is quite different from ours, however. The Athey and Bagwell (2001) model is one of perfect monitoring but incomplete payoff-relevant information (costs). Communication allows firms to allocate greater market shares in favor of low cost firms and these cost savings are the source of almost all of the profit gains. Thus, communication actually has social benefits; the resulting production efficiencies increase welfare. By contrast, our model is one of imperfect monitoring but complete information. Communication improves monitoring and this is the *only* source of profit gains. Now, communication has no social benefits; the resulting high prices reduce welfare. Another way to see the difference is by considering how a “planner” who can dictate the prices that firms charge would implement collusive outcomes in the two models. In the Athey and Bagwell (2001) model, such a planner must elicit the privately known cost information in order to implement maximally collusive outcomes. The resulting mechanism design problem necessitates communication. It then seems clear that in the absence of a planner, colluding firms would need to communicate with each other as well. In our model, however, there is no need to communicate any information to such a planner; he/she can simply dictate that firms charge the monopoly price. In this case, the fact that in the absence of a planner, maximally collusive outcomes can only be achieved with communication is our main finding. We believe that we are the first to establish that cheap-talk communication about payoff irrelevant information can aid collusion.

The remainder of the paper is organized as follows. The next section outlines the nature of the market. Section II analyzes the repeated game without communication whereas Section III does the same with. The findings of the earlier sections are combined in Section IV to derive the main result. We also calculate explicitly the gains from communication in an example with linear demands. Omitted proofs are collected in the Appendix.

I. The Market

There are two symmetric firms in the market, labeled 1 and 2. The firms produce differentiated products at a constant cost, which we normalize to zero. Each firm sets a price $p_i \in \mathcal{P}_i = [0, p_{\max}]$, for its product and given the prices set by the firms, their sales are *stochastic*. Prices affect the joint distribution of sales as follows. First, they affect expected sales in the usual way; an increase in p_i decreases firm i 's expected sales and increases firm j 's expected sales. Second, they affect how correlated are the sales of the two firms; the more similar are the two firms' prices, the higher is the correlation in sales.

To facilitate the analysis, we will suppose that given the two firms' prices $\mathbf{p} = (p_1, p_2)$, their sales (Y_1, Y_2) are jointly distributed according to a bivariate *log normal* density $f(y_1, y_2 | \mathbf{p})$. Precisely, at prices \mathbf{p} the log sales $(\ln Y_1, \ln Y_2)$ have a bivariate normal density with means $\mu_1(\mathbf{p})$, $\mu_2(\mathbf{p})$ and variance-covariance matrix of the form

$$\Sigma(\mathbf{p}) = \sigma^2 \begin{bmatrix} 1 & \rho(\mathbf{p}) \\ \rho(\mathbf{p}) & 1 \end{bmatrix}.$$

Note that we are assuming that the variance of log sales is unaffected by prices.⁵ We now specify the exact manner in which prices affect the distribution of sales.

A. Expected Sales

The function $\mu_i(p_i, p_j)$ determining the *expected* log sales of firm i is assumed to be a continuous function that is decreasing in p_i and increasing in p_j . The firms are symmetric so that $\mu_i = \mu_j$. Note that the first argument of μ_i is always firm i 's own price and the second is its competitor's price. Because sales are log normally distributed, the expected sales of firm i at prices \mathbf{p} are

$$E[Y_i | \mathbf{p}] = \exp\left(\mu_i(\mathbf{p}) + \frac{1}{2} \sigma^2\right).$$

In what follows, it will be convenient to denote the expected sales $E[Y_i | \mathbf{p}]$ by $Q_i(p_i, p_j)$. The function Q_i is then firm i 's expected demand function.

B. Correlation of Sales

We will suppose the correlation between firms' (log) sales is high when they charge similar prices and low when they charge dissimilar prices. This is formalized as:

ASSUMPTION: *There exists $\rho_0 \in (0, 1)$ and a symmetric function $\gamma(p_1, p_2) \in [0, 1]$ such that $\rho = \rho_0 \gamma(p_1, p_2)$ and γ satisfies the following conditions: (i) for all p , $\gamma(p, p) = 1$; and (ii) for all $p_1 \leq p_2$, $\partial \gamma / \partial p_1 > 0$ and $\partial^2 \gamma / \partial p_1^2 \geq 0$ and so, γ is an increasing and convex function of p_1 .*

Note that $\partial \rho / \partial p_1 = \rho_0 \partial \gamma / \partial p_1$, and so for fixed γ , an increase in ρ_0 represents an increase in the *sensitivity* of the correlation to prices. Some simple examples of the functions γ that satisfy the assumption are: $\gamma(p_1, p_2) = \min(p_1, p_2) / \max(p_1, p_2)$ and $\gamma(p_1, p_2) = 1 / (1 + |p_1 - p_2|)$.

Figure 1 is a schematic illustration of the stochastic relationship between prices and sales.

This kind of correlation structure is quite natural in many settings. Consider, for example, a symmetric Hotelling-type market in which the firms are located at different points on the line and consumers have identical but random "transport" costs. First, consider the case where the two firms charge very similar prices. In this case, their sales are similar—roughly, they split the market—no matter what the realized transport costs are. In other words, when firms charge similar prices, their sales are highly correlated. Next, consider the case where the two firms' prices are dissimilar, say, firm 1's price is much lower than firm 2's price. When transport costs are high, consumers are not that price sensitive and so the firms'

⁵A heteroskedastic specification in which the variance increased with the mean log sales can be easily accommodated.

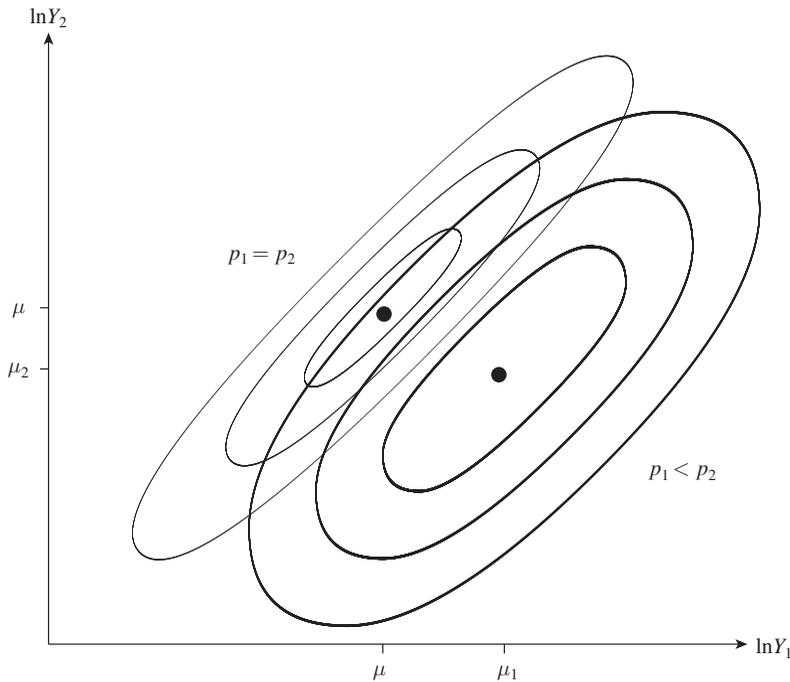


FIGURE 1. THE ASSUMED STOCHASTIC RELATIONSHIP BETWEEN PRICES AND LOG SALES VIA THE CONTOURS OF THE RESULTING NORMAL DENSITIES

Notes: When firms charge the same price, their log sales have the same mean and high correlation. If firm 1 cuts its price, its mean log sales rise while those of its rival fall. Correlation also falls.

realized sales are rather similar. But when transport costs are low, consumers are rather price sensitive and so firm 1's sales are much higher than firm 2's sales. In other words, when firms charge dissimilar prices, the correlation between their sales is low. Of course, the same kind of reasoning applies if we substitute search costs for transport costs.

Another setting in which the kind of correlation structure postulated above arises is the "random utility" model for discrete choice, widely used as the basis of many empirical industrial organization studies. In such a specification, consumers' utility is assumed to have both common and idiosyncratic components and the common component plays a role not unlike that of transport costs in the Hotelling model. Again, consider first the case where firms charge similar prices. Now, regardless of the realization of the common component, their sales are similar and so highly correlated. But when the firms charge different prices, how close their sales are depends on the realization of the common component of utility. When the common component is high, prices are not that important for consumers and so the sales of the two firms will be relatively similar. When the common component is low, however, price differences become important, the lower-priced firm will experience much higher sales than its rival, and so the two firms' sales will be dissimilar. Overall, this means that the correlation between their sales is lower.

C. One-Shot Game

Firms maximize their *expected* profits:

$$\pi_i(p_i, p_j) = p_i Q_i(p_i, p_j),$$

and we suppose that π_i is *strictly concave* in p_i .⁶

Let G denote the one-shot game where the firms choose prices p_i and p_j and the expected profits are given by $\pi_i(p_i, p_j)$. Under the assumptions made above, there exists a symmetric Nash equilibrium (p_N, p_N) of G and let π_N be the resulting profits of a firm.⁷

Suppose that (p_M, p_M) is the *unique* solution to the monopolist's problem:

$$\max_{p_i, p_j} \sum_i \pi_i(p_i, p_j),$$

and let π_M be the resulting profits per firm. We assume that monopoly pricing (p_M, p_M) is not a Nash equilibrium. For technical reasons we will also assume that a firm's expected sales are bounded away from zero.

II. Collusion without Communication

Let $G_\delta(f)$ denote the infinitely repeated game in which firms use the discount factor $\delta < 1$ to evaluate profit streams. Time is discrete. In each period, firms choose prices p_i and p_j and given these prices, their sales are realized according to f as described above. As in Stigler (1964), each firm i observes *only* its own realized sales y_i ; it observes neither j 's price p_j nor j 's sales y_j . We will refer to f as the *monitoring structure*.

Let $\mathbf{h}_i^{t-1} = (p_i^1, y_i^1, p_i^2, y_i^2, \dots, p_i^{t-1}, y_i^{t-1})$ denote the *private* history observed by firm i after $t - 1$ periods of play and let \mathcal{H}_i^{t-1} denote the set of all private histories of firm i . In period t , firm i chooses its prices p_i^t knowing \mathbf{h}_i^{t-1} and nothing else.

A *strategy* s_i for firm i is a collection of functions (s_i^1, s_i^2, \dots) such that $s_i^t : \mathcal{H}_i^{t-1} \rightarrow \Delta(\mathcal{P}_i)$, where $\Delta(\mathcal{P}_i)$ is the set of distributions over \mathcal{P}_i . Thus, we are allowing for the possibility that firms may randomize. Of course, since \mathcal{H}_i^0 is null, $s_i^1 \in \Delta(\mathcal{P}_i)$. A *strategy profile* \mathbf{s} is simply a pair of strategies (s_1, s_2) . A *Nash equilibrium* of $G_\delta(f)$ is strategy profile \mathbf{s} such that for each i , the strategy s_i is a best response to s_j .

The main result of this section provides an upper bound to the joint profits of the firms in *any* Nash equilibrium of the repeated game without communication.⁸ The task is complicated by the fact that there is no known characterization of the set of equilibrium payoffs of a repeated game with private monitoring. Because the players in such a game observe different histories—each firm knows only its own past prices

⁶In terms of the primitives, this is guaranteed if μ_i is sufficiently concave.

⁷If the one-shot game has multiple symmetric Nash equilibria, let (p_N, p_N) denote the one with the lowest equilibrium profits.

⁸Of course, the bound so derived applies to any refinement of Nash equilibrium as well.

and sales—such games lack a straightforward recursive structure and the kinds of techniques available to analyze equilibria of repeated games with *public* monitoring (as in the work of Abreu, Pearce, and Stacchetti 1990) cannot be used here.

Instead, we proceed as follows. Suppose we want to determine whether there is a Nash equilibrium of $G_\delta(f)$ such that the sum of firms' discounted average profits are within ε of those of a monopolist, that is, $2\pi_M$. If there were such an equilibrium, then both firms must set prices close to the monopoly price p_M often (or equivalently, with high probability). Now consider a secret price cut by firm 1 to \bar{p} , the static best response to p_M . Such a deviation is profitable today because firm 2's price is close to p_M with high probability. How this affects firm 2's future actions depends on the *quality* of monitoring, that is, how much firm 1's price cut affects the distribution of 2's sales. If the quality of monitoring is poor, firm 1 can keep on deviating to \bar{p} without too much fear of being punished. In other words, a firm has a profitable deviation, contradicting that there were such an equilibrium.

This reasoning shows that the resulting bound on Nash equilibrium profits depends on three factors: (i) the trade-off between the incentives to deviate and efficiency in the *one-shot* game⁹; (ii) the quality of the monitoring, which determines whether the short-term incentives to deviate can be overcome by future actions; and, of course, (iii) the discount rate.

We consider each of these factors in turn.

A. Incentives versus Efficiency in the One-Shot Game

Define, as above, $\bar{p} = \arg \max_{p_i} \pi_i(p_i, p_M)$, the static best-response to p_M . Let $\alpha \in \Delta(\mathcal{P}_1 \times \mathcal{P}_2)$ be a joint distribution over firms' prices. We want to find an α such that (i) the sum of the expected profits from α is within ε of $2\pi_M$; and (ii) it minimizes the (sum of) the incentives to deviate to \bar{p} . To that end, for $\varepsilon \geq 0$, define

$$(1) \quad \Psi(\varepsilon) \equiv \min_{\alpha} \sum_i [\pi_i(\bar{p}, \alpha_j) - \pi_i(\alpha)]$$

subject to

$$\sum_i \pi_i(\alpha) \geq 2\pi_M - \varepsilon,$$

where α_j denotes the marginal distribution of α over \mathcal{P}_j .

The function Ψ measures the trade-off between the incentives to deviate (to the price \bar{p}) and firms' profits. Precisely, if the firms' profits are within ε of those of a monopolist, then the total incentive to deviate is $\Psi(\varepsilon)$. It is easy to see that Ψ is (weakly) decreasing. Two other properties of Ψ also play an important role. First, Ψ is convex because both the objective function and the constraint are linear in the choice variable α . Second, $\lim_{\varepsilon \rightarrow 0} \Psi(\varepsilon) > 0$. To see this, note that $\Psi(0) > 0$ because at $\varepsilon = 0$, the only feasible solution to the problem above is (p_M, p_M) and

⁹By "efficiency" we mean how efficient the cartel is in achieving high profits and not "social efficiency."

by assumption, this is not a Nash equilibrium. Finally, the fact that Ψ is continuous at $\varepsilon = 0$ follows from the Berge maximum theorem.

Since (p_N, p_N) is feasible for the program defining Ψ when $\varepsilon = 2\pi_M - 2\pi_N$, it follows that $\Psi(2\pi_M - 2\pi_N) \leq 0$. We emphasize that Ψ is completely determined by the one-shot game G . Define the inverse of Ψ by

$$\Psi^{-1}(x) = \sup \{ \varepsilon : \Psi(\varepsilon) = x \}.$$

The bound we develop below will depend on Ψ^{-1} .

B. Quality of Monitoring

Consider two price pairs $\mathbf{p} = (p_1, p_2)$ and $\mathbf{p}' = (p'_1, p'_2)$ and the resulting marginal distributions of firm i 's sales: $f_i(\cdot | \mathbf{p})$ and $f_i(\cdot | \mathbf{p}')$. If these two distributions are close together, then it will be difficult for firm i to detect the change from \mathbf{p} to \mathbf{p}' . Thus, the quality of monitoring can be measured by the maximum “distance” between any two such distributions. In what follows, we use the so-called total variation metric to measure this distance. Since f is symmetric, the quality of monitoring is the same for both firms.

DEFINITION 1: *The quality of a monitoring structure f is defined as*

$$\eta = \max_{p, p'} \|f_i(\cdot | \mathbf{p}) - f_i(\cdot | \mathbf{p}')\|_{TV},$$

where f_i is the marginal of f on \mathcal{Y}_i , the set of i 's sales, and $\|g - h\|_{TV}$ denotes the total variation distance between g and h .¹⁰

It is important to note that the quality of monitoring depends only on the *marginal* distributions $f_i(\cdot | \mathbf{p})$ over i 's sales and not on the joint distributions of sales $f(\cdot | \mathbf{p})$. In particular, the fact that the marginal distributions $f_i(\cdot | \mathbf{p})$ and $f_i(\cdot | \mathbf{p}')$ are close— η is small—does not imply that the underlying joint distributions $f(\cdot | \mathbf{p})$ and $f(\cdot | \mathbf{p}')$ are close. Because of symmetry, η is the same for both firms. When $f(\cdot | \mathbf{p})$ is a bivariate log normal, η can be explicitly determined as

$$(2) \quad \eta = 2\Phi\left(\frac{\Delta\mu_{\max}}{2\sigma}\right) - 1,$$

where Φ is the cumulative distribution function of a univariate standard normal and $\Delta\mu_{\max} = \max_{p, p'} |\ln Q_i(p_i, p_j) - \ln Q_i(p'_i, p'_j)|$ is the maximum possible difference in log expected sales. As σ increases, η decreases and goes to zero as σ becomes arbitrarily large.

¹⁰The total variation distance between two densities g and h on X is defined as $\|g - h\|_{TV} = \frac{1}{2} \int_X |g(x) - h(x)| dx$.

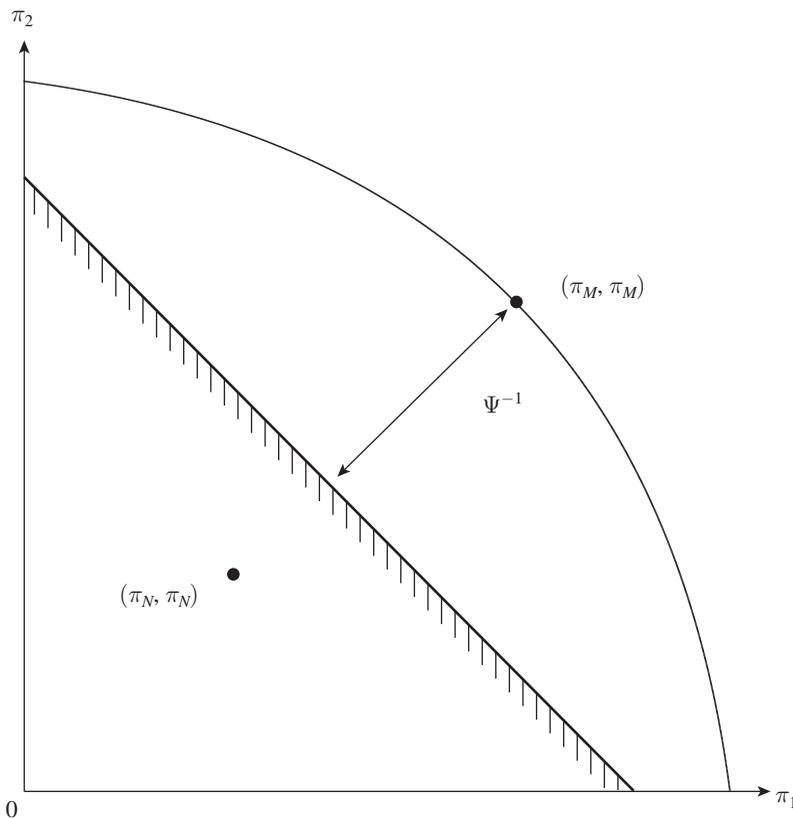


FIGURE 2. THE SET OF FEASIBLE PROFITS OF THE TWO FIRMS AND THE POSITION OF THE NO-COMMUNICATION BOUND

Note: The bound lies between monopoly and one-shot Nash profits and its size depends on the demand structure, the discount factor, and the monitoring quality.

C. A Bound on Profits

The main result of this section, stated below, develops a bound on Nash equilibrium profits when there is no communication. An important feature of the bound is that it is independent of any correlation between firms’ sales and depends only on the marginal distribution of sales.

PROPOSITION 1: *In any Nash equilibrium of the repeated game without communication, the average profits*

$$(3) \quad \pi_1 + \pi_2 \leq 2\pi_M - \Psi^{-1}\left(4\bar{\pi}\eta \frac{\delta^2}{1-\delta}\right),$$

where $\bar{\pi} = \max_{p_j} \pi_i(\bar{p}, p_j)$ and η is the quality of monitoring.

Before embarking on a formal proof of Proposition 1 it is useful to outline the main ideas (see Figure 2 for an illustration). A necessary condition for a strategy profile s to be an equilibrium is a deviation by firm 1 to a strategy \bar{s}_1 in which it

always charges \bar{p} not be profitable. This is done in two steps. First, we consider a fictitious situation in which firm 1 assumes that firm 2 will not respond to its deviation. The higher the equilibrium profits, the more profitable would be the proposed deviation in the fictitious situation; this is exactly the effect the function Ψ captures in the one-shot game and Lemma 1 shows that Ψ captures the same effect in the repeated game as well. Second, when the monitoring is poor— η is small—firm 2's actions cannot be very responsive to the deviation and so the fictitious situation is a good approximation for the true situation. Lemma 5 measures precisely how good this approximation is and quite naturally this depends on the quality of monitoring and the discount factor. Notice that profits both from the candidate equilibrium and from the play after the deviation are evaluated in ex ante terms.

Observe that if we fix the quality of monitoring η and let the discount factor δ approach one, then the bound becomes trivial (since $\lim_{\delta \rightarrow 1} \Psi^{-1}(4\bar{\pi}\eta\delta^2/(1-\delta)) = 0$) and so is consistent with Sugaya's (2013) folk theorem. On the other hand, if we fix the discount factor δ and decrease the quality of monitoring η , the bound converges to $2\pi_M - \Psi^{-1}(0) < 2\pi_M$ and is effective. One may reasonably conjecture that if there were "zero monitoring" in the limit, that is, if $\eta \rightarrow 0$, then no collusion would be possible. But in fact $2\pi_N < 2\pi_M - \Psi^{-1}(0)$ so that even with zero monitoring, Proposition 1 does not rule out the presence of collusive equilibria. This is consistent with the finding of Awaya (2014b).¹¹

PROOF OF PROPOSITION 1:

We argue by contradiction. Suppose that $G_\delta(f)$ has an equilibrium, say \mathbf{s} , whose average total profits¹² $\pi_1(\mathbf{s}) + \pi_2(\mathbf{s})$ exceed the bound on the right-hand side of (3). If we write $\varepsilon(\mathbf{s}) \equiv 2\pi_M - \pi_1(\mathbf{s}) - \pi_2(\mathbf{s})$, then this is equivalent to

$$\Psi(\varepsilon(\mathbf{s})) > 4\bar{\pi}\eta \frac{\delta^2}{1-\delta}.$$

Given the strategy profile s , define

$$\alpha_j^t = E_s \left[s_j^t(\mathbf{h}_j^{t-1}) \right] \in \Delta(\mathcal{P}_j),$$

where the expectation is defined by the probability distribution over $t-1$ joint histories $(\mathbf{h}_i^{t-1}, \mathbf{h}_j^{t-1})$ determined by \mathbf{s} . Note that α_j depends on the strategy profile \mathbf{s} and not just s_j . Let $\alpha_j = (\alpha_j^1, \alpha_j^2, \dots)$ denote the strategy of firm j in which it plays α_j^t in period t following any $t-1$ period history. The strategy α_j replicates the ex ante distribution of prices p_j resulting from s but is *nonresponsive* to histories.

Let \bar{s}_i denote the strategy of firm i in which it plays \bar{p} with probability one following any history. In the Appendix we show that the function Ψ , which measures

¹¹ Awaya (2014b) constructs an example in which there is zero monitoring—the marginal distribution of every player's signals is the same for all action profiles—but, nevertheless, there are nontrivial equilibria.

¹² We use $\pi_i(s)$ to denote the discounted average payoffs from the strategy profile s as well as the payoffs in the one-shot game.

the incentive-efficiency trade-off in the one-shot game can also be used to measure this trade-off in the repeated game as well. Essentially, the convexity of Ψ ensures that the best dynamic incentives are in fact stationary; the argument resembles a consumption smoothing result. Formally, from Lemma 1,

$$(4) \quad \sum_i \left[\pi_i(\bar{s}_i, \alpha_j) - \pi_i(\mathbf{s}) \right] \geq \Psi(\varepsilon(\mathbf{s})).$$

Second, from Lemma 5, we have for $i = 1, 2$

$$(5) \quad \left| \pi_i(\bar{s}_i, s_j) - \pi_i(\bar{s}_i, \alpha_j) \right| \leq 2\bar{\pi}\eta \frac{\delta^2}{1-\delta},$$

which follows from the fact that the total variation distance between the distribution of j 's sales induced by (\bar{s}_i, s_j) and (\bar{s}_i, α_j) in any one period does not exceed η (Lemma 3) and, as a result, the distance between distribution of j 's t -period sales histories does not exceed $t\eta$ (Lemma 4). A simple calculation then shows that the difference in payoffs does not exceed the right-hand side of (5).

Combining (4) and (5), we have

$$\begin{aligned} \sum_i \left(\pi_i(\bar{s}_i, s_j) - \pi_i(\mathbf{s}) \right) &= \sum_i \left(\pi_i(\bar{s}_i, \alpha_j) - \pi_i(\mathbf{s}) \right) + \sum_i \left(\pi_i(\bar{s}_i, s_j) - \pi_i(\bar{s}_i, \alpha_j) \right) \\ &\geq \sum_i \left(\pi_i(\bar{s}_i, \alpha_j) - \pi_i(\mathbf{s}) \right) - \sum_i \left| \pi_i(\bar{s}_i, s_j) - \pi_i(\bar{s}_i, \alpha_j) \right| \\ &\geq \Psi(\varepsilon(\mathbf{s})) - 4\bar{\pi}\eta \frac{\delta^2}{1-\delta}, \end{aligned}$$

which is strictly positive. But this means that at least one firm has a profitable deviation, contradicting the assumption that \mathbf{s} is an equilibrium. ■

In a recent paper, Pai, Roth, and Ullman (2014) also provide a bound on equilibrium payoffs that is effective when monitoring is poor. They derive necessary conditions for an equilibrium by considering “one-shot” deviations in which a player cheats in one period and then resumes equilibrium play. The Pai, Roth, and Ullman (2014) bound is based on how the *joint* distribution of the private signals is affected by players' actions.

In games with private monitoring, “one-shot” deviations affect the deviating players' beliefs about the other player's signals thereafter and so affect his subsequent (optimal) play. The optimal play, quite naturally, exploits any correlation in signals. The deviation we consider—a permanent price cut—is rather naïve but has the feature that future play, while suboptimal, is straightforward. This renders unnecessary any conjectures about the future behavior of the other firm and so our bound does not depend on any correlation between signals (sales); it is based solely on the *marginal* distributions.

When we consider communication in the next section, we will exploit the correlation of signals to construct an equilibrium whose profits exceed our bound (Proposition 2 below). The fact that correlation can vary while keeping the marginal

distributions fixed is key to isolating the effects of communication. The bound obtained by Pai, Roth, and Ullman (2014) applies to both forms of collusion—with and without communication—and cannot distinguish between the two. Sugaya and Wolitzky (2015) also provide a bound but one that is based on entirely different ideas. For a fixed discount factor, they ask which monitoring structure yields the highest equilibrium profits. This, of course, then identifies a bound that applies across all monitoring structures. But this method, while quite general, again cannot distinguish between the two settings.

III. Collusion with Communication

We now turn to a situation in which firms can, in addition to setting prices, communicate with each other in every period, sending one of a finite set of messages to each other. The sequence of actions in any period is as follows: firms set prices, receive their private sales information, and then simultaneously send messages to each other. Messages are costless—the communication is “cheap talk”—and are transmitted without any noise. The communication is unmediated.

Formally, there is a finite set of messages M_i for each firm and that each M_i contains at least two elements. A $t - 1$ period private history of firm i now consists of the complete list of its own prices and sales as well as the list of all messages sent and received. Thus a private history is now of the form

$$\mathbf{h}_i^{t-1} = (p_i^\tau, y_i^\tau, m_i^\tau, m_j^\tau)_{\tau=1}^{t-1},$$

and the set of all such histories is denoted by \mathcal{H}_i^{t-1} . A strategy for firm i is now a pair (s_i, r_i) where $s_i = (s_i^1, s_i^2, \dots)$, the pricing strategy, and $r_i = (r_i^1, r_i^2, \dots)$, the reporting strategy, are collections of functions: $s_i^t : \mathcal{H}_i^{t-1} \rightarrow \Delta(\mathcal{P}_i)$ and $r_i^t : \mathcal{H}_i^{t-1} \times \mathcal{P}_i \times \mathcal{Y}_i \rightarrow \Delta(M_i)$ where \mathcal{Y}_i denotes the set of possible sales realizations.

Call the resulting infinitely repeated game with communication $G_\delta^{com}(f)$. A sequential equilibrium of $G_\delta^{com}(f)$ is a strategy profile (\mathbf{s}, \mathbf{r}) such that for each i and every private history \mathbf{h}_i^{t-1} , the continuation strategy of i following \mathbf{h}_i^{t-1} , denoted by $(s_i, r_i)|_{\mathbf{h}_i^{t-1}}$, is a best response to $E[(s_j, r_j)|_{\mathbf{h}_j^{t-1}} | \mathbf{h}_i^{t-1}]$.¹³

A. Equilibrium Strategies

Monopoly pricing will be sustained using a *grim trigger* pricing strategy together with a *threshold* sales-reporting strategy in a manner first identified by Aoyagi (2002). Since the price set by a competitor is not observable, the trigger will be based on the communication between firms, which is observable. The communication itself consists only of reporting whether one’s sales were “high”—above a commonly known threshold—or “low.” Firms start by setting monopoly prices and continue to do so as long as the two sales reports agree—both firms report “high”

¹³ Sequential equilibrium is usually defined only for games with a finite set of actions and signals. We are applying it to a game with a continuum of actions and signals. But this causes no problems in our model because signals (sales) have full support and beliefs are well defined.

or both report “low.” Differing sales reports trigger permanent noncooperation as a punishment.

Specifically, consider the following strategy (s_i^*, r_i^*) in the repeated game with communication where there are only two possible messages H (“high”) and L (“low”). The pricing strategy s_i^* is:

- In period 1, set the monopoly price p_M .
- In any period $t > 1$, if in all previous periods, the reports of both firms were identical (both reported H or both reported L), set the monopoly price p_M ; otherwise, set the Nash price p_N .

The reporting strategy r_i^* is:

- In any period $t \geq 1$, if the price set was $p_i = p_M$, then report H if $\ln y_i^t \geq \mu_M$; otherwise, report L .
- In any period $t \geq 1$, if the price set was $p_i \neq p_M$, then report H if $\ln y_i^t \geq \mu_i + \frac{1}{\bar{\rho}}(\mu_M - \mu_j)$; otherwise, report L .

$$\left(\mu_M = \ln Q_i(p_M, p_M) - \frac{1}{2} \sigma^2, \mu_i = \ln Q_i(p_i, p_M) - \frac{1}{2} \sigma^2, \text{ and } \mu_j = \ln Q_j(p_M, p_i) - \frac{1}{2} \sigma^2. \right)$$

Denote by (s^*, r^*) the resulting strategy profile. We will establish that if firms are patient enough and the monitoring structure is noisy (σ is high) but correlated (ρ_0 is high), then the strategies specified above constitute an equilibrium. We begin by showing that the reporting strategy r_i^* is indeed optimal.¹⁴

B. Optimality of Reporting Strategy

Suppose firm 2 follows the strategy (s_2^*, r_2^*) and until this period, both have made identical sales reports. Recall that a punishment will be triggered only if the reports disagree. Thus, firm 1 will want to maximize the probability that its report agrees with that of firm 2. Since firm 2 is following a threshold strategy, it is optimal for firm 1 to do so as well. If firm 1 adopts a threshold of λ such that it reports H when its log sales exceed λ , and L when they are less than λ , the probability that the reports will agree is

$$\Pr[\ln Y_1 < \lambda, \ln Y_2 < \mu_M] + \Pr[\ln Y_1 > \lambda, \ln Y_2 > \mu_M],$$

which, for normally distributed variables, is

$$\int_{-\infty}^{\frac{\lambda - \mu_1}{\sigma}} \int_{-\infty}^{\frac{\mu_M - \mu_2}{\sigma}} \phi(z_1, z_2; \rho) dz_2 dz_1 + \int_{\frac{\lambda - \mu_1}{\sigma}}^{\infty} \int_{\frac{\mu_M - \mu_2}{\sigma}}^{\infty} \phi(z_1, z_2; \rho) dz_2 dz_1,$$

¹⁴The strategies described above have been dubbed *semi-public* by Compte (1998) and others. All actions depend only on past public signals—that is, the communication. The communication itself also depends on current private signals.

where $\phi(z_1, z_2; \rho)$ is a standard bivariate normal density with correlation coefficient $\rho \in (0, 1)$.¹⁵ Maximizing this with respect to λ results in the *optimal* reporting threshold:

$$(6) \quad \lambda(p_1) = \mu_1 + \frac{1}{\rho} (\mu_M - \mu_2),$$

where μ_1, μ_2 , and ρ are evaluated at the price pair (p_1, p_M) .

If firm 1 deviates and cuts its price to $p_1 < p_M$, then clearly the expected (log) sales of the two firms will be such that $\mu_1 > \mu_M > \mu_2$. Thus, $\lambda(p_1) > \mu_1 > \mu_M$, which says, as expected, that once firm 1 cuts its price—and so experiences stochastically higher sales—it should optimally *underreport* relative to the equilibrium reporting strategy r_1^* . On the other hand, if firm 1 does not deviate and sets a price p_M , then (6) implies that it is optimal for it to use a threshold of $\mu_M = \lambda(p_M)$ as well.

We have thus established that if firm 2 plays according to (s_2^*, r_2^*) , then following any price p_1 that firm 1 sets, the communication strategy r_1^* is optimal. The optimality of the proposed pricing strategy s_1^* depends crucially on the probability of triggering the punishment and we now establish how this is affected by the extent of a price cut.

If firm 1 sets a price of p_1 , the probability that its report will be the same as that of firm 2 (and so the punishment will not be triggered) is given by

$$(7) \quad \beta(p_1) \equiv \int_{-\infty}^{\frac{\lambda(p_1) - \mu_1}{\sigma}} \int_{-\infty}^{\frac{\mu_M - \mu_2}{\sigma}} \phi(z_1, z_2; \rho) dz_2 dz_1 + \int_{\frac{\lambda(p_1) - \mu_1}{\sigma}}^{\infty} \int_{\frac{\mu_M - \mu_2}{\sigma}}^{\infty} \phi(z_1, z_2; \rho) dz_2 dz_1.$$

Note that while $\lambda(p_1), \mu_1, \mu_2$, and the correlation coefficient ρ depend on p_1, σ does not. Observe also that for any $p_1 \leq p_M, \beta(p_1) \geq \Phi\left(\frac{\mu_M - \mu_2}{\sigma}\right) \geq \frac{1}{2}$ where Φ denotes the cumulative distribution function of the standard univariate normal. In other words, the probability of detecting a deviation is less than one-half. This is because firm 1 could always adopt a reporting strategy in which after a deviation to $p_1 < p_M$, it always says L independently of its own sales, effectively setting $\lambda(p_1) = \infty$. This guarantees that firm 1's report will be the same as firm 2's report with a probability equal to $\Phi\left(\frac{\mu_M - \mu_2}{\sigma}\right)$. Figure 3 depicts the function β when $\rho = \rho_0 / (1 + |p_1 - p_2|)$.

If both follow the proposed pricing and reporting strategies, the probability that their reports will agree is just $\beta(p_M)$, obtained by setting $\lambda(p_1) = \mu_1 = \mu_2 = \mu_M$ and $\rho = \rho_0$ in (7). Sheppard's formula for the cumulative of a bivariate normal (see Tihansky 1972) implies that

$$(8) \quad \beta(p_M) = \frac{1}{\pi} \arccos(-\rho_0),$$

¹⁵The standard (with both means equal to zero and both variances equal to one) bivariate normal density is $\phi(z_1, z_2; \rho) = \exp(-(z_1^2 + z_2^2 - 2\rho z_1 z_2)/2(1 - \rho^2)) \div 2\pi\sqrt{1 - \rho^2}$.

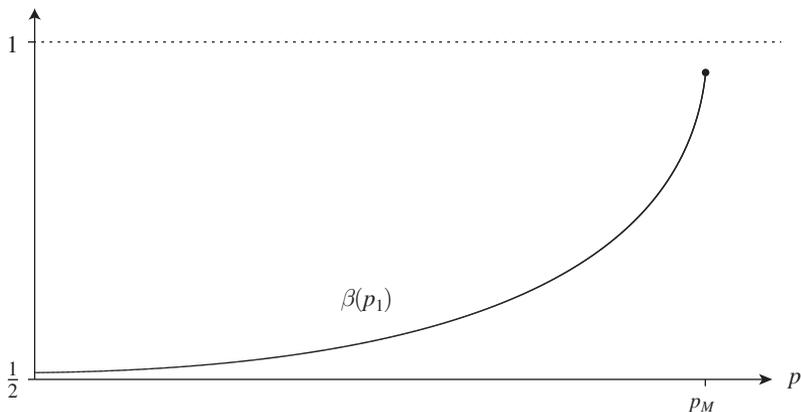


FIGURE 3. THE PROBABILITY THAT SALES REPORTS AGREE AFTER FIRM 1 CUTS ITS PRICE AND REPORTS OPTIMALLY

Notes: This probability lies strictly between $\frac{1}{2}$ and 1. It is highest at the monopoly price and decreases with the extent of the price cut.

which is increasing in ρ_0 and converges to 1 as ρ_0 goes to 1. Importantly, this probability does not depend on σ .

C. Optimality of Pricing Strategy

We have argued above that given that the other firm follows the prescribed strategy, the reporting strategy r_1^* is an optimal response for firm 1. To show that the strategies (s^*, r^*) constitute an equilibrium, it only remains to show that the pricing strategy s_1^* is optimal as well. This is verified next.

PROPOSITION 2: *There exists a $\underline{\delta}$ such that for all $\delta > \underline{\delta}$, once σ and ρ_0 are large enough, then (s^*, r^*) constitutes an equilibrium of $G_\delta^{com}(f)$, the repeated game with communication.*

PROOF:

First, note that the lifetime average profit π^* resulting from the proposed strategies is given by

$$(9) \quad (1 - \delta)\pi_M + \delta[\beta(p_M)\pi^* + (1 - \beta(p_M))\pi_N] = \pi^*.$$

Next, suppose that in all previous periods, both firms have followed the proposed strategies and their reports have agreed. If firm 1 deviates to $p_1 < p_M$ in the current period, it gains

$$(10) \quad \Delta_1(p_1) = (1 - \delta)\pi_1(p_1, p_M) + \delta[\beta(p_1)\pi^* + (1 - \beta(p_1))\pi_N] - \pi^*,$$

where π^* is defined in (9). Thus,

$$\Delta'_1(p_1) = (1 - \delta) \frac{\partial \pi_1}{\partial p_1}(p_1, p_M) + \delta \beta'(p_1) \times [\pi^* - \pi_N].$$

We will show that when σ is large enough, for all p_1 , $\Delta'_1(p_1) > 0$. Since $\Delta_1(p_M) = 0$, this will establish that a deviation to a price $p_1 < p_M$ is not profitable. Now observe that from Lemma 6,

$$\begin{aligned} \lim_{\sigma \rightarrow \infty} \Delta'_1(p_1) &= (1 - \delta) \frac{\partial \pi_1}{\partial p_1}(p_1, p_M) \\ &\quad + \delta \frac{1}{\pi \sqrt{1 - \rho_0^2 \gamma(p_1, p_M)^2}} \rho_0 \frac{\partial \gamma}{\partial p_1}(p_1, p_M) \times [\pi^* - \pi_N] \\ &\geq (1 - \delta) \frac{\partial \pi_1}{\partial p_1}(p_M, p_M) \\ &\quad + \delta \frac{1}{\pi \sqrt{1 - \rho_0^2 \gamma(0, p_M)^2}} \rho_0 \frac{\partial \gamma}{\partial p_1}(0, p_M) \times [\pi^* - \pi_N], \end{aligned}$$

where the last inequality uses the fact that since π_1 is concave in p_1 , $\frac{\partial \pi_1}{\partial p_1}(p_1, p_M) > \frac{\partial \pi_1}{\partial p_1}(p_M, p_M)$, and the fact that $\gamma(p_1, p_M)$ is increasing and convex in p_1 . Let $\sigma(\delta)$ be such that for all $\sigma > \sigma(\delta)$, the inequality above holds.

Let $\underline{\delta}$ be the solution to

$$(11) \quad (1 - \delta) \frac{\partial \pi_1}{\partial p_1}(p_M, p_M) + \delta \frac{1}{\pi \sqrt{1 - \gamma(0, p_M)^2}} \frac{\partial \gamma}{\partial p_1}(0, p_M) \times [\pi^* - \pi_N] = 0,$$

which is just the right-hand side of the inequality above when $\rho_0 = 1$. Such a $\underline{\delta}$ exists since $\frac{\partial \pi_1}{\partial p_1}(p_M, p_M)$ is finite and, by assumption, $\frac{\partial \rho}{\partial p_1}(0, p_M)$ is strictly positive. Notice that for any $\delta > \underline{\delta}$, the expression on the left-hand side is *strictly* positive.

Now observe that

$$\frac{1}{\pi \sqrt{1 - \rho_0^2 \gamma(1, p_M)^2}} \rho_0 \frac{\partial \gamma}{\partial p_1}(0, p_M) \times [\pi^* - \pi_N]$$

is increasing and continuous in ρ_0 (recall that π^* is increasing in ρ_0). Thus, given any $\delta > \underline{\delta}$, there exists a $\rho_0(\delta)$ such that for $\rho_0 = \rho_0(\delta)$:

$$(1 - \delta) \frac{\partial \pi_1}{\partial p_1}(p_M, p_M) + \delta \frac{1}{\pi \sqrt{1 - \rho_0^2 \gamma(0, p_M)^2}} \rho_0 \frac{\partial \gamma}{\partial p_1}(0, p_M) \times [\pi^* - \pi_N] = 0.$$

Note that $\rho_0(\delta)$ is a decreasing function of δ and for any $\rho_0 > \rho_0(\delta)$, the left-hand side is strictly positive.

A deviation by firm 1 to a price $p_1 > p_M$ is clearly unprofitable. ■

Aoyagi (2002) was the first to use threshold reporting strategies in a correlated environment. He shows that for a given monitoring structure (ρ_0 and σ fixed) as the discount factor δ goes to one, such strategies constitute an equilibrium. The idea—as in all “folk theorems”—is that even when the probability of detection is low, if players are patient enough, future punishments are a sufficient deterrent. In contrast, Proposition 2 shows that for a given discount factor (δ high but fixed), as ρ_0 goes to one and σ goes to infinity, there is an equilibrium with high profits. Its logic, however, is different from that underlying the “folk theorems” with communication, as for instance, in the work of Compte (1998) and Kandori and Matsushima (1998). In these papers, signals are conditionally independent and, in equilibrium, players are indifferent among the messages they send. Kandori and Matsushima (1998) also show that with correlation, strict incentives for “truth-telling” can be provided as is also true in our construction. But the key difference between our work and these papers is that in Proposition 2 the punishment power derives not from the patience of the players; rather it comes from the noisiness of the monitoring. A deviating firm finds it very difficult to predict its rival’s sales and hence, even it “lies” optimally, a deviation is very likely to trigger a punishment.

The equilibrium constructed in Proposition 2 relies on using the correlation of signals to “check” firms’ reports. This is reminiscent of the logic underlying the full-surplus extraction results of Crémer and McLean (1988) but, of course, in our setting there is no mechanism designer who can commit to arbitrarily large punishments to induce truth-telling.

In our model, firms communicate *simultaneously*—neither firm knows the other’s report prior to its own—and this feature is crucial to the equilibrium construction. It is the uncertainty about what the other firm will say that disciplines firms’ behavior. Often trade associations help cartels exchange sales data in a manner that is mimicked by the simultaneous reporting in our equilibrium. Cartel members make confidential sales reports to the association which, as mentioned above, disseminates aggregate sales data to the cartel. Individual sales figures are not shared.¹⁶ Their confidentiality is usually preserved by third parties. The pocket calculator ploy mentioned in the introduction also serves the same purpose.

Finally, the “grim trigger” strategies specified above are *unforgiving* in that a single disagreement at the communication stage triggers a permanent reversion to the one-shot Nash equilibrium. Since disagreement occurs with positive probability even if neither firm cheats, this means that in the long-run, collusion inevitably breaks down. Real-world cartels do punish transgressors but these punishments are not permanent. The equilibrium strategies we have used can easily be amended so that the punishment phase is not permanent and that, after a prespecified number of periods, say T , firms return to monopoly pricing. Specifically, this means that the equilibrium profits π^{**} are now defined by

$$(1 - \delta) \pi_M + \delta \left[\beta(p_M) \pi^{**} + (1 - \beta(p_M)) \left((1 - \delta^{T+1}) \pi_N + \delta^{T+1} \pi^{**} \right) \right] = \pi^{**}$$

¹⁶The sharing of aggregate data is not illegal (see the discussion of the *Maple Flooring* case in Areeda and Kaplow 1988).

instead of (9). For any $\delta > \underline{\delta}$ (as defined by (11)), there exists a T large enough so that the forgiving strategy constitutes an equilibrium as well. Note that $\pi^{**} > \pi^*$.

IV. Gains from Communication

Proposition 1 shows that the profits from any equilibrium without communication cannot exceed

$$2\pi_M - \Psi^{-1}\left(4\bar{\pi}\eta \frac{\delta^2}{1-\delta}\right),$$

whereas Proposition 2 provides conditions under which there is an equilibrium with communication with profits $2\pi^*$ (as defined in (9)). From (8) it follows that as $\rho_0 \rightarrow 1$, $\beta(p_M) \rightarrow 1$. Now from (9) it follows that as $\rho_0 \rightarrow 1$, $\pi^* \rightarrow \pi_M$. Combining these facts leads to the formal version of the result stated in the introduction. Let $\underline{\delta}$ be determined as in (11).

THEOREM 1: *For any $\delta > \underline{\delta}$, there exist $(\sigma(\delta), \rho_0(\delta))$ such that for all $(\sigma, \rho_0) \gg (\sigma(\delta), \rho_0(\delta))$ there is an equilibrium with communication with total profits $2\pi^*$ such that*

$$2\pi^* > 2\pi_M - \Psi^{-1}\left(4\bar{\pi}\eta \frac{\delta^2}{1-\delta}\right).$$

In the limit, as $\rho_0 \rightarrow 1$ and $\sigma \rightarrow \infty$, the left-hand side of the inequality goes to $2\pi_M$ and the right-hand side goes to $2\pi_M - \Psi^{-1}(0)$ since $\eta \rightarrow 0$. Thus, the difference in profits between the two regimes is at least $\Psi^{-1}(0) > 0$.

The workings of the main result can be seen in Figure 4, which is drawn for the case of linear demand (see the next section for details). First, notice that the profits with communication depend on ρ_0 and not on σ (but σ has to be sufficiently high to guarantee that the suggested strategies form an equilibrium). The bound on profits without communication, on the other hand, depends on σ and not on ρ_0 . For small values of σ , the bound is ineffective and says only that these do not exceed joint monopoly profits. As σ increases, the bound becomes tighter but at intermediate levels the profits from the equilibrium with communication do not exceed the bound. Once $\sigma > \sigma(\delta)$, communication strictly increases profits.

A. Example with Linear Demand

We now illustrate the workings of our results when (expected) demand is linear. Suppose that¹⁷

$$Q_i(p_i, p_j) = \max(A - bp_i + p_j, 1),$$

¹⁷This specification of “linear” demand is used because $\ln 0$ is not defined.

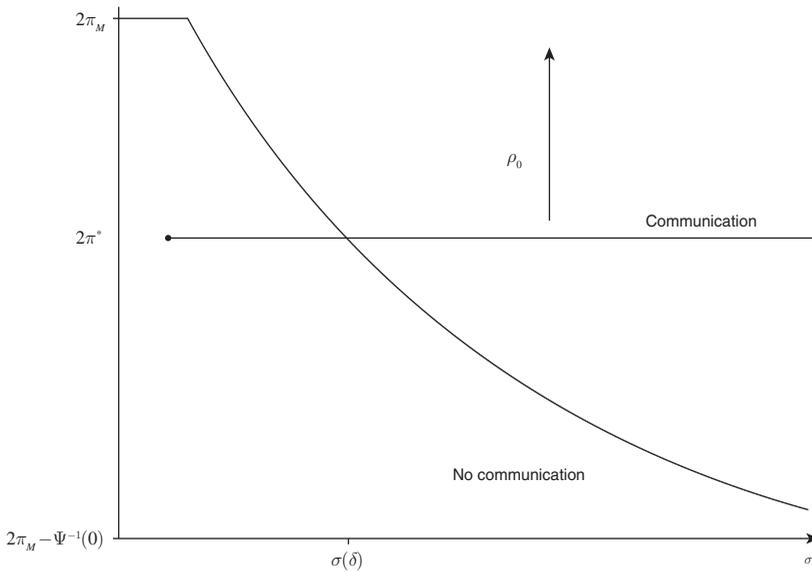


FIGURE 4. THE MAIN RESULT

Notes: The no-communication bound decreases with σ , the standard deviation of log sales, but is independent of ρ_0 , their correlation. The profits from the constructed equilibrium with communication are, on the other hand, independent of σ , but increase with ρ_0 . Thus, for large σ and ρ_0 , the profits with communication exceed those from any equilibrium without.

where $A > 0$ and $b > 1$. For this specification, the monopoly price $p_M = A/2(b - 1)$ and monopoly profits $\pi_M = A^2/4(b - 1)$. There is a unique Nash equilibrium of the one-shot game with prices $p_N = A/(2b - 1)$ and profits $\pi_N = A^2b/(2b - 1)^2$. A firm's best response if the other firm charges the monopoly price p_M is to charge $\bar{p} = A(2b - 1)/4b(b - 1)$. The highest possible profit that firm 1 can achieve when charging a price of \bar{p} is $\bar{\pi} = \pi_1(\bar{p}, p_M) = A^2(2b - 1)^2/16b(b - 1)^2$.

It remains to specify how the correlation between the firms' log sales is affected by prices. In this example we adopt the following specification:

$$\rho = \frac{\rho_0}{1 + |p_1 - p_2|},$$

which, of course, satisfies the assumption made in Section I.

Then, recalling (1), it may be verified that for $\varepsilon \in [0, \pi_M/2b^2]$,

$$\Psi(\varepsilon) = \varepsilon + \frac{A}{8b(b - 1)^2} (A - 2\sqrt{2(b - 1)}(2b - 1)\sqrt{\varepsilon}),$$

which is achieved at equal prices. Note that $\Psi(0) = \pi_M/2b(b - 1)$ and $\Psi^{-1}(0) = \pi_M/2b^2$.

Finally, from (2)

$$\eta = 2\Phi\left(\frac{\Delta\mu_{\max}}{2\sigma}\right) - 1,$$

where $\Delta \mu_{\max} = \ln Q_2(0, p_M) - \ln Q_2(p_M, 0)$.

As a numerical example, suppose $A = 120$ and $b = 2$. Let $\delta = 0.7$ and $\rho_0 = 0.95$. For these parameters, $\pi_M = 3,600$, $\bar{\pi} = 4,050$, and $\Delta \mu_{\max} = 5.19$. Also, the profits from the equilibrium with communication, $\pi^* = 3,524$ (approximately).

Figure 4 depicts the bound on profits without communication as a function of σ using Proposition 1. For low values of σ (approximately $\sigma = 60$ or lower), the bound is ineffective—it equals $2\pi_M$ —and as $\sigma \rightarrow \infty$, converges to $2\pi_M - \Psi^{-1}(0)$. As shown, the profits with communication exceed the bound when $\sigma > \sigma(\delta) = 200$ (approximately).

Figure 5 verifies that the strategies (s^*, r^*) constitute an equilibrium—a deviation to any $p_1 < p_M$ is unprofitable as $\Delta(p_1) < 0$ (as defined in (10)). This is verified for $\sigma = 60$ and, of course, the same strategies remain an equilibrium when σ is higher.

B. On the Necessity of Communication

Theorem 1 identifies circumstances in which communication is *necessary* for collusion in a private monitoring setting. That it is *sufficient* has been pointed out in many studies, beginning with Compte (1998) and Kandori and Matsushima (1998) and then continuing with Fudenberg and Levine (2007), Zheng (2008), and Obara (2009). Of particular interest is the work of Aoyagi (2002) and Harrington and Skrzypacz (2011) in oligopoly settings. Under various assumptions, all of these conclude that the folk theorem holds—any individually rational and feasible outcome can be approximated as the discount factor tends to one. But as Kandori and Matsushima (1998, p. 648, their italics) recognize, “One thing which we did *not* show is the *necessity* of communication for a folk theorem.”

Indeed when players are arbitrarily patient, communication is not necessary for collusion, as a remarkable paper by Sugaya (2013) shows. He establishes the surprising result that in very general environments, the folk theorem holds without any communication. An important component of Sugaya’s proof is that players implicitly communicate via their actions. Thus, he shows that with enough time, there is no need for direct communication.

In a recent paper, Rahman (2014) shows how communication can overcome the “no collusion” result of Sannikov and Skrzypacz (2007) in a repeated Cournot oligopoly with public monitoring. But since the latter result concerns only equilibria in public strategies, this again does not establish the necessity of communication for collusion, only its sufficiency. Moreover, the communication considered by Rahman (2014) is either mediated or verifiable. In our work, as well as in the papers mentioned above, communication is unmediated and unverifiable.

The restriction to public strategies is also a feature of the analysis of Athey and Bagwell (2001) in their model of collusion with incomplete information about costs. Communication is used to sustain collusion, but as shown by Hörner and Jamison (2007), once the restriction to public strategies is removed, communication is no longer necessary for collusion.

Escobar and Toikka (2013) generalize the Athey-Bagwell model to study general repeated games with incomplete, possibly persistent, information with

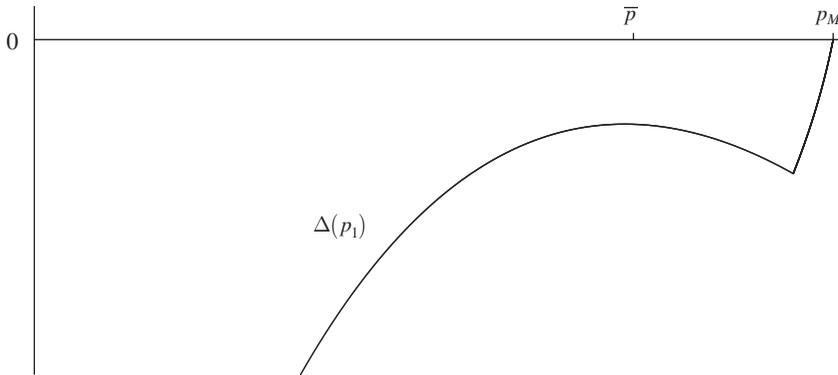


FIGURE 5. THE NET GAINS FROM DEVIATING FROM THE PRESCRIBED STRATEGIES WITH COMMUNICATION

Notes: Small price cuts are not that profitable but are hard to detect. Bigger price cuts are more profitable but easier to detect. This trade-off accounts for the nonmonotonicity of net gains.

communication. They show that with low discounting, efficient outcomes can be approximated. On the other hand, Escobar and Llanes (2015) derive conditions under which efficiency cannot be attained without communication. The two results then isolate circumstances in which, for low discounting, communication is necessary to achieve efficiency.

In a different vein, Awaya (2014a) studies the prisoners' dilemma with private monitoring and shows that for a fixed discount rate, there exist environments in which without communication, the only equilibrium is the one-shot equilibrium whereas with communication, almost perfect cooperation can be sustained. This paper is a precursor to the current one.

Finally, the necessity of communication has been studied in laboratory *experiments* as well (by Fonseca and Normann 2012 and Cooper and Kühn 2014 among others). These experiments, however, do not involve imperfect monitoring and so do not directly address the issues dealt with in this paper.

V. Conclusion

We have provided theoretical support for the idea that even unverifiable communication within a cartel facilitates greater collusion and is detrimental for society. How does “cheap talk” aid collusion? When the monitoring is poor, and there is no communication, a price cut cannot be detected with any confidence and this is the basis of the bound developed in Proposition 1. With communication, however, the probability that a price cut will trigger a punishment is significant relative to the short-term gains. Thus, communication reduces the type II error associated with imperfect monitoring and this is the driving force behind our main result.

It has been argued (Carlton, Gertner, and Rosenfeld 1996) that the exchange of information among firms is beneficial to society because it improves the allocation of resources. Consider a firm that experiences high sales. If it learns, in addition, that other firms have also experienced high sales, then it would rightly infer that its own high sales are the result of a market-wide positive demand shock. If this is likely to persist in the future, a firm can then increase production and investment with

greater confidence. In this paper, however, demand shocks are independent across periods. Last period's sales thus carry information about the rival's past behavior but reveal nothing about future demand. Thus, any positive effect resulting from information exchange is absent. A model in which demand shocks are persistent would incorporate both effects and, perhaps, be able to measure the trade-off. This remains a project for future research.

We emphasize that, in this paper, we have not explored the role of communication in coordinating cartel behavior: agreeing on prices, market shares, and transfers, among other things. Since cartels cannot sign binding contracts enforced by third parties, presumably they can only coordinate on strategies that form an equilibrium. Pre-play communication would then serve to select one among many equilibria (see Green, Marshall, and Marx 2014). But a theory of how communication helps firms select among equilibria, while important for antitrust practice, remains a challenge.

APPENDIX

A. Collusion without Communication

Nonresponsive Strategies.—The ex ante distribution over \mathcal{P}_j in period t induced by a strategy profile \mathbf{s} is

$$(A1) \quad \alpha_j^t(\mathbf{s}) = E_{\mathbf{s}} \left[s_j^t(\mathbf{h}_j^{t-1}) \right] \in \Delta(\mathcal{P}_j).$$

Given a strategy profile \mathbf{s} , recall that α_j denotes the strategy of firm j in which it plays $\alpha_j^t(\mathbf{s})$ in period t following any $t - 1$ period history. The strategy α_j replicates the ex ante distribution of prices resulting from \mathbf{s} but is *nonresponsive* to histories.

The following lemma shows that the function Ψ , defined in (1), which determines the incentives versus efficiency trade-off in the one-shot game, embodies the same trade-off in a *repeated* setting if the nondeviating player follows a nonresponsive strategy. It shows that to minimize the average incentive to deviate while achieving average profits within ε of $2\pi_M$ one should split the incentive evenly across periods. The lemma resembles an intertemporal “consumption smoothing” argument (recall that Ψ is convex).

LEMMA 1 (Smoothing): *For any strategy profile \mathbf{s} whose profits exceed $2\pi_M - \varepsilon$,*

$$\sum_i \left[\pi_i(\bar{s}_i, \alpha_j(\mathbf{s})) - \pi_i(\mathbf{s}) \right] \geq \Psi(\varepsilon),$$

where \bar{s}_i denotes firm i 's strategy in which it sets \bar{p} following any history.

PROOF:

Define

$$\varepsilon(t) = 2\pi_M - E_{\mathbf{s}} \left[\sum_i \pi_i(\mathbf{s}^t(\mathbf{h}^{t-1})) \right]$$

as the difference between the sum of monopoly profits $2\pi_M$ and the sum of expected profits in period t . Now clearly $(1 - \delta) \sum_{t=1}^{\infty} \delta^t \varepsilon(t) \leq \varepsilon$.

Then,

$$\begin{aligned} \sum_i [\pi_i(\bar{s}_i, \alpha_j(\mathbf{s})) - \pi_i(\mathbf{s})] &= E_{\mathbf{s}} \left[(1 - \delta) \sum_{t=1}^{\infty} \delta^t \sum_i [\pi_i(\bar{p}, p_j^t) - \pi_i(p_i^t, p_j^t)] \right] \\ &\geq (1 - \delta) \sum_{t=1}^{\infty} \delta^t \Psi(\varepsilon(t)) \\ &\geq \Psi \left((1 - \delta) \sum_{t=1}^{\infty} \delta^t \varepsilon(t) \right) \\ &\geq \Psi(\varepsilon). \end{aligned}$$

The first equality follows from the fact that the induced distribution over prices p_j^t is the same under (\bar{s}_i, α_j) as it is under \mathbf{s} . The second inequality follows from the definition of Ψ . The third and fourth result from the fact that Ψ is convex and non-increasing, respectively. ■

Weak Monitoring.—For a fixed strategy pair (s_1, s_2) , let λ_j^t be the induced probability distribution over firm j 's private histories $\mathbf{h}_j^t \in \mathcal{H}_j^t = (\mathcal{P}_j \times \mathcal{Y}_j)^t \subset \mathbb{R}^{2t}$. Similarly, let $\bar{\lambda}_j^t$ be the probability distribution over j 's private histories induced by the strategy pair (\bar{s}_i, s_j) .¹⁸ We wish to determine the total variation distance between λ_j^t and $\bar{\lambda}_j^t$.

The *total variation* distance between two distributions G and \bar{G} over \mathbb{R}^n is equal to

$$(A2) \quad \|G - \bar{G}\|_{TV} = \frac{1}{2} \sup_{\|\varphi\|_{\infty} \leq 1} |E[\varphi] - \bar{E}[\varphi]|,$$

where E and \bar{E} denote the expectations with respect to the distributions G and \bar{G} , respectively, and the supremum is taken over all measurable functions φ with sup norm $\|\varphi\|_{\infty} \leq 1$. Note that the definition in (A2) is equivalent to the one in Definition 1. See, for instance, Levin, Peres, and Wilmer (2009).

As a first step, we decompose the total variation distance between two probability distributions into the distance between their marginals and that between their conditionals.

LEMMA 2: *Given two distributions G and \bar{G} over $\mathbb{R}^m \times \mathbb{R}^n$,*

$$\|G - \bar{G}\|_{TV} \leq \|G_X - \bar{G}_X\|_{TV} + \sup_x \|G_{Y|X} - \bar{G}_{Y|X}\|_{TV},$$

where G_X is the marginal distribution of G on \mathbb{R}^m and $G_{Y|X}(\cdot | x)$ is the conditional distribution of G on \mathbb{R}^n given $X = x$ (and similarly for \bar{G}).

¹⁸Recall that \bar{s}_i denotes the strategy of firm i in which it sets \bar{p} with probability one following any history.

PROOF:

Let E and \bar{E} denote the expectations with respect to G and \bar{G} , respectively. Given any function $\varphi : \mathbb{R}^m \times \mathbb{R}^n \rightarrow [-1, 1]$, we have

$$\begin{aligned}
 \frac{1}{2} |E[\varphi] - \bar{E}[\varphi]| &= \frac{1}{2} |E_X[E_{Y|X}[\varphi]] - \bar{E}_X[\bar{E}_{Y|X}[\varphi]]| \\
 &\leq \frac{1}{2} |E_X[E_{Y|X}[\varphi]] - \bar{E}_X[E_{Y|X}[\varphi]]| \\
 &\quad + \frac{1}{2} |\bar{E}_X[E_{Y|X}[\varphi]] - \bar{E}_X[\bar{E}_{Y|X}[\varphi]]| \\
 &\leq \frac{1}{2} \sup_{\|\varphi\|_\infty \leq 1} |E_X[\varphi] - \bar{E}_X[\varphi]| + \frac{1}{2} \bar{E}_X[|E_{Y|X}[\varphi] - \bar{E}_{Y|X}[\varphi]|] \\
 &\leq \|G_X - \bar{G}_X\|_{TV} + \bar{E}_X[\|G_{Y|X} - \bar{G}_{Y|X}\|_{TV}] \\
 &\leq \|G_X - \bar{G}_X\|_{TV} + \sup_x \|G_{Y|X} - \bar{G}_{Y|X}\|_{TV},
 \end{aligned}$$

and so

$$\begin{aligned}
 \|G - \bar{G}\|_{TV} &= \frac{1}{2} \sup_{\|\varphi\|_\infty \leq 1} |E[\varphi] - \bar{E}[\varphi]| \\
 &\leq \|G_X - \bar{G}_X\|_{TV} + \sup_x \|G_{Y|X} - \bar{G}_{Y|X}\|_{TV},
 \end{aligned}$$

which is the required inequality. ■

Next we show that given a history \mathbf{h}_j^{t-1} , the total variation distance between the two conditional distributions cannot exceed the monitoring quality.

LEMMA 3: For any \mathbf{h}_j^{t-1} ,

$$\left\| \lambda_j^t(\cdot | \mathbf{h}_j^{t-1}) - \bar{\lambda}_j^t(\cdot | \mathbf{h}_j^{t-1}) \right\|_{TV} \leq \eta.$$

PROOF:

Let $S_j(\cdot | \mathbf{h}_j^{t-1})$ denote the distribution of prices that j 's strategy s_j induces after j 's private history \mathbf{h}_j^{t-1} . Similarly, let $S_i(\cdot | \mathbf{h}_i^{t-1})$ denote the distribution of prices that i 's strategy induces after i 's private history \mathbf{h}_i^{t-1} . Let $\hat{S}_i(\cdot | \mathbf{h}_j^{t-1}) = E[S_i(\cdot | \mathbf{h}_i^{t-1}) | \mathbf{h}_j^{t-1}]$ denote j 's expectation about i 's distribution of prices, given j 's own history \mathbf{h}_j^{t-1} . Finally, let $\hat{S}(\mathbf{p} | \mathbf{h}_j^{t-1}) = \hat{S}_i(p_i | \mathbf{h}_j^{t-1}) S_j(p_j | \mathbf{h}_j^{t-1})$ denote the joint distribution of prices that j expects given j 's private history \mathbf{h}_j^{t-1} (recall that firms' choices are independent).

Now, for any $\varphi : \mathcal{P}_j \times \mathcal{Y}_j \rightarrow [-1, 1]$

$$\begin{aligned} I(\varphi) &\equiv \left| \int_{\mathcal{P}_j \times \mathcal{Y}_j} \varphi d\lambda_j^t(\cdot | \mathbf{h}_j^{t-1}) - \int_{\mathcal{P}_j \times \mathcal{Y}_j} \varphi d\bar{\lambda}_j^t(\cdot | \mathbf{h}_j^{t-1}) \right| \\ &= \left| \int_{\mathcal{P}_j} \int_{\mathcal{Y}_j} \varphi dF_j(y_j | \mathbf{p}) dS_j(p_j | \mathbf{h}_j^{t-1}) d\hat{S}_i(p_i | \mathbf{h}_j^{t-1}) \right. \\ &\quad \left. - \int_{\mathcal{P}_j} \int_{\mathcal{Y}_j} \varphi dF_j(y_j | \bar{p}, p_j) dS_j(p_j | \mathbf{h}_j^{t-1}) \right| \\ &= \left| \int_{\mathcal{P}} \int_{\mathcal{Y}_j} \varphi dF_j(y_j | \mathbf{p}) d\hat{S}(\mathbf{p} | \mathbf{h}_j^{t-1}) - \int_{\mathcal{P}} \int_{\mathcal{Y}_j} \varphi dF_j(y_j | \bar{p}, p_j) d\hat{S}(\mathbf{p} | \mathbf{h}_j^{t-1}) \right|, \end{aligned}$$

where $\mathcal{P} \equiv \mathcal{P}_1 \times \mathcal{P}_2$ and the second equality follows by integrating

$$\int_{\mathcal{P}_j} \int_{\mathcal{Y}_j} \phi dF_j(y_j | \bar{p}, p_j) dS_j(p_j | \mathbf{h}_j^{t-1})$$

over \mathcal{P}_i using the distribution $\hat{S}_i(\cdot | \mathbf{h}_j^{t-1})$. Thus,

$$\begin{aligned} I(\varphi) &= \left| \int_{\mathcal{P}} \int_{\mathcal{Y}_j} \varphi [f_j(y_j | \mathbf{p}) - f_j(y_j | \bar{p}, p_j)] dy_j d\hat{S}(\mathbf{p} | \mathbf{h}_j^{t-1}) \right| \\ &\leq \int_{\mathcal{P}} \left[\int_{\mathcal{Y}_j} |\varphi f_j(y_j | \mathbf{p}) - \varphi f_j(y_j | \bar{p}, p_j)| dy_j \right] d\hat{S}(\mathbf{p} | \mathbf{h}_j^{t-1}) \\ &\leq 2 \int_{\mathcal{P}} \eta d\hat{S}(\mathbf{p} | \mathbf{h}_j^{t-1}) \\ &= 2\eta, \end{aligned}$$

where the second inequality follows from the definition of total variation.

Thus,

$$\begin{aligned} \left\| \lambda_j^t(\cdot | \mathbf{h}_j^{t-1}) - \bar{\lambda}_j^t(\cdot | \mathbf{h}_j^{t-1}) \right\|_{TV} &= \frac{1}{2} \sup_{\|\varphi\|_{\infty} \leq 1} I(\varphi) \\ &\leq \eta, \end{aligned}$$

as was to be shown. ■

Combining the preceding two lemmas we obtain the key result that the total variation distance between λ_j^t and $\bar{\lambda}_j^t$ is bounded above by a linear function of time.

LEMMA 4: For all t ,

$$\|\lambda_j^t - \bar{\lambda}_j^t\|_{TV} \leq t\eta.$$

PROOF:

The proof is by induction. For $t = 1$, there is no history and Lemma 3 implies the result directly.

Now suppose that the result holds for $t - 1$. Using Lemma 2, we have

$$\begin{aligned} \|\lambda_j^t - \bar{\lambda}_j^t\|_{TV} &\leq \|\lambda_j^{t-1} - \bar{\lambda}_j^{t-1}\|_{TV} + \sup_{\mathbf{h}_j^{t-1}} \|\lambda_j^t(\cdot | \mathbf{h}_j^{t-1}) - \bar{\lambda}_j^t(\cdot | \mathbf{h}_j^{t-1})\|_{TV} \\ &\leq (t-1)\eta + \eta \end{aligned}$$

by the induction hypothesis and Lemma 3. ■

The next result verifies the intuition that when the monitoring quality is low, the profits of a deviator who undertakes a permanent price cut are not too different from those when its rival follows a distributionally equivalent nonresponsive strategy. The importance of the lemma is in quantifying this difference.

LEMMA 5: Let α be the nonresponsive strategy as defined in (A1). Then,

$$|\pi_i(\bar{s}_i, s_j) - \pi_i(\bar{s}_i, \alpha_j)| \leq 2\bar{\pi}\eta \frac{\delta^2}{1-\delta},$$

where $\bar{\pi} = \max_{p_j} \pi_i(\bar{p}, p_j)$.

PROOF:

As above, let λ_j^t be the distribution over firm j 's private histories \mathbf{h}_j^t induced by (s_i, s_j) and let $\bar{\lambda}_j^t$ be the distribution over j 's private histories induced by (\bar{s}_i, s_j) . Then,

$$\pi_i(\bar{s}_i, s_j) = (1-\delta) \sum_{t=1}^{\infty} \delta^t \int_{\mathcal{H}_j^{t-1}} E[\pi_i(\bar{p}, s_j) | \mathbf{h}_j^{t-1}] d\bar{\lambda}_j(\mathbf{h}_j^{t-1}).$$

Also, if $S_j(\cdot | \mathbf{h}_j^{t-1})$ denotes the distribution over j 's prices induced by the strategy s_j following the history \mathbf{h}_j^{t-1} , then

$$\begin{aligned} \pi_i(\bar{s}_i, \alpha_j) &= (1-\delta) \sum_{t=1}^{\infty} \delta^t \int_{\mathcal{P}_j} \pi_i(\bar{p}, p_j) d\alpha_j^t(p_j) \\ &= (1-\delta) \sum_{t=1}^{\infty} \delta^t \int_{\mathcal{P}_j} \pi_i(\bar{p}, p_j) \left(\int_{\mathcal{H}_j^{t-1}} dS_j(p_j | \mathbf{h}_j^{t-1}) d\lambda_j(\mathbf{h}_j^{t-1}) \right) \\ &= (1-\delta) \sum_{t=1}^{\infty} \delta^t \int_{\mathcal{H}_j^{t-1}} E[\pi_i(\bar{p}, s_j) | \mathbf{h}_j^{t-1}] d\lambda_j(\mathbf{h}_j^{t-1}) \end{aligned}$$

since by definition

$$\alpha_j^t(p_j) = \int_{\mathcal{H}_j^{t-1}} \mathcal{S}_j(p_j | \mathbf{h}_j^{t-1}) d\lambda_j(\mathbf{h}_j^{t-1}).$$

Thus,

$$\begin{aligned} & \left| \pi_i(\bar{s}_i, s_j) - \pi_i(\bar{s}_i, \alpha_j) \right| \\ & \leq (1 - \delta) \sum_{t=1}^{\infty} \delta^t \left| \int_{\mathcal{H}_j^{t-1}} E[\pi_i(\bar{p}, s_j) | \mathbf{h}_j^{t-1}] (d\bar{\lambda}_j(\mathbf{h}_j^{t-1}) - d\lambda_j(\mathbf{h}_j^{t-1})) \right| \\ & \leq 2\bar{\pi}\eta(1 - \delta) \sum_{t=1}^{\infty} \delta^t (t - 1) \\ & = 2\bar{\pi}\eta \frac{\delta^2}{1 - \delta}, \end{aligned}$$

where the second inequality is a consequence of Lemma 4 and the fact that, as in (A2), given any two distributions λ and $\bar{\lambda}$, $|E[\varphi] - \bar{E}[\varphi]| \leq 2\|\varphi\|_{\infty} \times \|\lambda - \bar{\lambda}\|_{TV}$ for any bounded measurable function φ . ■

B. Collusion with Communication

LEMMA 6: For any $p_1 \geq 1$,

$$\lim_{\sigma \rightarrow \infty} \beta'(p_1) = \frac{1}{\pi \sqrt{1 - \rho_0^2 \gamma(p_1, p_M)^2}} \times \rho_0 \frac{\partial \gamma}{\partial p_1}(p_1, p_M).$$

PROOF:

First, we derive $\beta'(p_1)$. Since λ is optimally chosen, the envelope theorem guarantees that

$$\frac{\partial \beta(p_1)}{\partial \lambda} = 0,$$

and so

$$\frac{\partial \beta(p_1)}{\partial \mu_1} = -\frac{\partial \beta(p_1)}{\partial \lambda} = 0$$

as well.

Thus, we have

$$\beta'(p_1) = \frac{\partial \beta(p_1)}{\partial \rho} \frac{\partial \rho}{\partial p_1} + \frac{\partial \beta(p_1)}{\partial \mu_2} \frac{\partial \mu_2}{\partial p_1}.$$

Now, since we can, write

$$\beta(p_1) = 2\bar{\Phi}\left(\frac{\lambda(p_1) - \mu_1}{\sigma}, \frac{\mu_M - \mu_2}{\sigma}; \rho\right) + \Phi\left(\frac{\lambda(p_1) - \mu_1}{\sigma}\right) + \Phi\left(\frac{\mu_M - \mu_2}{\sigma}\right) - 1,$$

where

$$\bar{\Phi}(z_1, z_2; \rho) = \Pr[Z_1 \geq z_1, Z_2 \geq z_2] = \int_{-1}^{\rho} \phi(z_1, z_2; \theta) d\theta$$

using Sheppard's formula¹⁹ (see Tihansky 1972), and Φ is the cumulative distribution function of a standard univariate normal. Thus,

$$\frac{\partial\beta(p_1)}{\partial\rho} = 2\phi\left(\frac{\lambda(p_1) - \mu_1}{\sigma}, \frac{\mu_M - \mu_2}{\sigma}; \rho\right),$$

which converges to $2\phi(0, 0; \rho) = 1/\pi\sqrt{1 - \rho^2}$ as $\sigma \rightarrow \infty$.

Finally,

$$\begin{aligned} \frac{\partial\beta(p_1)}{\partial\mu_2} &= \frac{1}{\sigma} \left[\int_{-\infty}^{\frac{\lambda - \mu_1}{\sigma}} \phi\left(z_1, \frac{\mu_M - \mu_2}{\sigma}; \rho\right) dz_1 - \int_{\frac{\lambda - \mu_1}{\sigma}}^{\infty} \phi\left(z_1, \frac{\mu_M - \mu_2}{\sigma}; \rho\right) dz_1 \right] \\ &= \frac{1}{\sigma} \left[2\Phi\left(\frac{1 - \rho^2}{\rho} \frac{\mu_M - \mu_2}{\sigma}\right) - 1 \right] \phi\left(\frac{\mu_M - \mu_2}{\sigma}\right). \end{aligned}$$

This converges to 0 as $\sigma \rightarrow \infty$. ■

REFERENCES

- Abreu, Dilip, David Pearce, and Ennio Stacchetti.** 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58 (5): 1041–63.
- Aoyagi, Masaki.** 2002. "Collusion in Dynamic Bertrand Oligopoly with Correlated Private Signals and Communication." *Journal of Economic Theory* 102 (1): 229–48.
- Areeda, Phillip, and Louis Kaplow.** 1988. *Antitrust Analysis: Problems, Text, Cases*. 4th edition. Boston: Little, Brown and Company.
- Athey, Susan, and Kyle Bagwell.** 2001. "Optimal Collusion with Private Information." *RAND Journal of Economics* 32 (3): 428–65.
- Awaya, Yu.** 2014a. "Private Monitoring and Communication in Repeated Prisoners' Dilemma." <http://www.personal.psu.edu/yxa120/research.html> (accessed November 17, 2015).
- Awaya, Yu.** 2014b. "Cooperation without Monitoring." <http://www.personal.psu.edu/yxa120/research.html> (accessed November 17, 2015).
- Carlton, Dennis W., Robert H. Gertner, and Andrew M. Rosenfield.** 1996. "Communication among Competitors: Game Theory and Antitrust." *George Mason Law Review* 5 (3): 423–40.
- Clark, Robert, and Jean-Francois Houde.** 2014. "The Effect of Explicit Communication on Pricing: Evidence from the Collapse of a Gasoline Cartel." *Journal of Industrial Economics* 62 (2): 191–228.
- Compte, Olivier.** 1998. "Communication in Repeated Games with Imperfect Private Monitoring." *Econometrica* 66 (3): 597–626.
- Cooper, David J., and Kai-Uwe Kühn.** 2014. "Communication, Renegotiation, and the Scope for Collusion." *American Economic Journal: Microeconomics* 6 (2): 247–78.

¹⁹This employs a change of variables to the original formula.

- Crémer, Jacques, and Richard P. McLean.** 1988. "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions." *Econometrica* 56 (6): 1247–57.
- Escobar, Juan F., and Gastón Llanes.** 2015. "Cooperation Dynamics in Repeated Games of Adverse Selection." Centro de Economía Aplicada, Universidad de Chile, Documentos de Trabajo 311.
- Escobar, Juan F., and Juuso Toikka.** 2013. "Efficiency in Games with Markovian Private Information." *Econometrica* 81 (5): 1887–1934.
- Fonseca, Miguel A., and Hans-Theo Normann.** 2012. "Explicit vs. Tacit Collusion—The Impact of Communication in Oligopoly Experiments." *European Economic Review* 56 (8): 1759–72.
- Friedman, James W.** 1971. "A Non-cooperative Equilibrium for Supergames." *Review of Economic Studies* 38 (1): 1–12.
- Fudenberg, Drew, and David K. Levine.** 2007. "The Nash-Threats Folk Theorem with Communication and Approximate Common Knowledge in Two Player Games." *Journal of Economic Theory* 132 (1): 461–73.
- Genesove, David, and Wallace P. Mullin.** 2001. "Rules, Communication, and Collusion: Narrative Evidence from the Sugar Institute Case." *American Economic Review* 91 (3): 379–98.
- Green, Edward J., Robert C. Marshall, and Leslie M. Marx.** 2014. "Tacit Collusion in Oligopoly." In *The Oxford Handbook of International Antitrust Economics*, Vol. 2, edited by Roger D. Blair and D. Daniel Sokol, 464–97. Oxford, UK: Oxford University Press.
- Green, Edward J., and Robert H. Porter.** 1984. "Noncooperative Collusion under Imperfect Price Information." *Econometrica* 52 (1): 87–100.
- Harrington, Joseph E., Jr.** 2006. "How Do Cartels Operate?" *Foundations and Trends in Microeconomics* 2 (1): 1–108.
- Harrington, Joseph E., and Andrzej Skrzypacz.** 2011. "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices." *American Economic Review* 101 (6): 2425–49.
- Hörner, Johannes, and Julian Jamison.** 2007. "Collusion with (Almost) No Information." *RAND Journal of Economics* 38 (3): 804–22.
- Kandori, Michihiro, and Hitoshi Matsushima.** 1998. "Private Observation, Communication and Collusion." *Econometrica* 66 (3): 627–52.
- Levin, David A., Yuval Peres, and Elizabeth L. Wilmer.** 2009. *Markov Chains and Mixing Times*. Providence, RI: American Mathematical Society.
- Mailath, George J., and Larry Samuelson.** 2006. *Repeated Games and Reputations: Long-Run Relationships*. Oxford, UK: Oxford University Press.
- Marshall, Robert C., and Leslie M. Marx.** 2012. *The Economics of Collusion: Cartels and Bidding Rings*. Cambridge, MA: MIT Press.
- Obara, Ichiro.** 2009. "Folk Theorem with Communication." *Journal of Economic Theory* 144 (1): 120–34.
- Pai, Mallesh, Aaron Roth, and Jonathan Ullman.** 2014. "An Anti-Folk Theorem for Large Repeated Games with Imperfect Monitoring." <http://arxiv.org/abs/1402.2801v2> (accessed November 17, 2015).
- Rahman, David.** 2014. "The Power of Communication." *American Economic Review* 104 (11): 3737–51.
- Sannikov, Yuliy, and Andrzej Skrzypacz.** 2007. "Impossibility of Collusion under Imperfect Monitoring with Flexible Production." *American Economic Review* 97 (5): 1794–1823.
- Shapiro, Carl.** 1986. "Exchange of Cost Information in Oligopoly." *Review of Economic Studies* 53 (3): 433–46.
- Stigler, George J.** 1964. "A Theory of Oligopoly." *Journal of Political Economy* 72 (1): 44–61.
- Sugaya, Takuo.** 2013. "Folk Theorem in Repeated Games with Private Monitoring." <https://sites.google.com/site/takuosugaya/home/research> (accessed November 17, 2015).
- Sugaya, Takuo, and Alexander Wolitzky.** 2015. "On the Equilibrium Payoff Set in Repeated Games with Imperfect Private Monitoring." <http://economics.mit.edu/faculty/wolitzky/research> (accessed November 17, 2015).
- Tihansky, Dennis P.** 1972. "Properties of the Bivariate Normal Cumulative Distribution." *Journal of the American Statistical Association* 67 (340): 903–05.
- Zheng, Bingyong.** 2008. "Approximate Efficiency in Repeated Games with Correlated Private Signals." *Games and Economic Behavior* 63 (1): 406–16.