On Tacit versus Explicit Collusion*

Yu Awaya† and Vijay Krishna‡
Penn State University

October 30, 2014

Abstract

Antitrust law makes a sharp distinction between tacit and explicit collusion whereas the theory of repeated games—the standard framework for studying collusion—does not. In this paper, we study this difference in Stigler's (1964) model of secret price cutting. This is a repeated game with private monitoring since in the model, firms observe neither the prices nor the sales of their rivals. For a fixed discount factor, we identify conditions under which there are equilibria under explicit collusion that result in near-perfect collusion—profits are close to those of a monopolist—whereas all equilibria under tacit collusion are bounded away from this outcome. Thus, in our model, explicit collusion leads to higher prices and profits than tacit collusion.

JEL classification: C73, D43

Keywords: Collusion, communication, repeated games, private monitoring

1 Introduction

In ruling on an antitrust case in 1993, the US Supreme Court clearly stated that tacit collusion—the setting of supracompetitive prices without evidence of conspiracy—was not in itself unlawful.[1] When there is evidence of explicit collusion, however, the law provides for severe fines, even prison terms. Antitrust law thus makes a sharp distinction between tacit and explicit collusion. In the former, there is no communication between firms, whereas in the latter there is. The theory of repeated games—the standard framework for studying collusion—does not, however, provide a justification for this distinction since in most models communication does not increase

---

*We thank Olivier Compte, Joyee Deb and Ed Green for helpful conversations and seminar participants at NYU, Warwick and Indiana for comments.

†E-mail: YuAwaya@gmail.com

‡E-mail: vkrishna@psu.edu

[1] See *Brooke Group v. Brown & Williamson*, 509 US 209, US Supreme Court, June 21, 1993.

cartel profits. But as Marshall and Marx (2012) write: "... repeated [tacit] interaction is not enough in practice, at least not for many firms in many industries. Even for duopolies ... explicit collusion was required to substantially elevate prices and profits (p. 3)." Harrington (2005) points to the shortcomings of economic theory in this regard: "There is a gap between antitrust practice—which distinguishes explicit and tacit collusion—and economic theory—which (generally) does not (p. 6)."

In this paper, we study this issue in Stigler's (1964) model of secret price cutting, which is a repeated game with private monitoring. Firms cannot observe each other's prices nor can they observe each other's sales. Each firm only observes its own sales and, because of demand shocks, these are imperfect signals of the other firms' actions. These signals are *noisy* in the sense that given the prices, the marginal distribution of a firm's sales is dispersed. At the same time firms' sales are correlated. We study situations where this correlation is rather *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, when there is a unilateral price cut. This kind of *monitoring structure* can arise quite naturally, for example, in a Hotelling-type model with random transport costs (see Section 2). For analytic convenience, we suppose that the relationship between sales and prices is governed by log normal distributions.

Under tacit collusion, firms base their pricing decisions only on their own history of prices and sales. Because sales are subject to unobserved shocks, it is difficult for firms to detect a rival's price cuts. Under explicit collusion, firms can communicate with each other in every period and pricing decisions are now based on these communications as well as their private histories. The communication is "cheap talk"—the firms exchange non-verifiable sales reports in every period.

Our main result is[2]

Theorem *For any high but fixed discount factor, when the monitoring is noisy but sensitive enough, there is an equilibrium under explicit collusion whose profits are strictly greater than those from any equilibrium under tacit collusion.*

Explicit collusion leads to higher sustainable prices and profits because even unverifiable communication improves monitoring. In their study of the sugar refining cartel, based on internal documents, Genesove and Mullin (2001), point to the monitoring role of the weekly (!) meetings of the firms. Clark and Houde (2014) find similar evidence in the retail gasoline market in Canada. The exchange of sales figures for monitoring purposes seems to have been key to the functioning of cartels in numerous industries, including citric acid, lysine and graphite electrodes (see Harrington, 2005).

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

[2] A formal statement of the result is in Section 5.

under tacit collusion. But the model we study is that of a repeated game with private monitoring and there is no known characterization of the set of equilibrium payoffs. This is because with private monitoring each firm knows only its own history (of prices and sales) and has to infer its rival's history. Since firms' histories are not commonly known, these cannot be used as state variables in a recursive formulation of the equilibrium payoff set. Thus we are forced to proceed somewhat differently. 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 naive—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.

The second task is to show that the bound developed earlier can be exceeded under explicit collusion. This is done by directly constructing an equilibrium in which firms exchange sales reports in every period (see Proposition 2). Firms charge monopoly prices and report truthfully. In this case, firms' sales are highly correlated and so the likelihood that their reports will agree is also high. If a firm were to cut its price, sales become less correlated and 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 non-cooperation is an effective deterrent. When the correlation between firms' sales is high, the chances of triggering a punishment without a deviation are small and so this equilibrium can achieve high profits even for relatively low discount factors. It turns out that noisier sales only make the inference problem for the deviating firm harder and thus decrease the incentive to cut price.

The key to our results is that the bound developed in Proposition 1 depends only on the *marginal* distribution of sales—precisely, on how noisy these are—and not on the *correlation* between sales. The equilibrium constructed in Proposition 2, however, depends on the correlation structure and, as mentioned above, is actually reinforced by noise. Thus we are able to identify conditions under which the bound on tacit collusion is tight while the equilibrium under explicit collusion approximates the monopoly outcome.

We emphasize that the analysis in this paper is of a different nature than that underlying the so-called "folk theorems" (see Sugaya, 2013). 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 to drive a wedge between tacit and explicit collusion. Our goal is only to identify *some* natural circumstances in which this happens—we do not attempt a full identification of monitoring structures which distinguish between the two.

In our model, firms share sales information and this allows them to better monitor each other. There are, of course, other pieces of information that may be communicated to facilitate the cartel. Firms may coordinate on market shares after exchanging privately known cost information (as in Athey and Bagwell, 2001 and Escobar and Toikka, 2013). Another view is that communication allows the cartel to coordinate on one among many repeated game equilibria (see Green, Marshall and Marx, 2013). But there is no formal "meta-theory" of how players coordinate on a single equilibrium. Our explanation of the gains from communication does not rely on equilibrium selection. We exhibit an equilibrium under explicit collusion that dominates *all* equilibria under tacit collusion.

Related literature

There is a vast literature on repeated games under different monitoring assumptions. Under *perfect* monitoring, given any fixed discount factor, the set of perfect equilibrium payoffs with and without communication is the same. Under *public* monitoring, again given any fixed discount factor, the set of (public) perfect equilibrium payoffs with and without communication is also the same. Thus, in these settings there is no difference between tacit and explicit collusion. The reason, of course, is that all relevant information is commonly known and so there is nothing useful to communicate.

Compte (1998) and Kandori and Matsushima (1998) study repeated games with *private* monitoring when there is *communication* among the players. In this setting, they show that the folk theorem holds—any individually rational and feasible outcome can be approximated as the discount factor tends to one. This line of research has been pursued by others as well, in varying environments (see Fudenberg and Levine (2007) and Obara (2009) among others). Aoyagi's (2002) work is, in particular, closely related because he also considers a secret price cutting model with a similar monitoring structure and communication.[3] He shows that monopoly outcomes can be approximated as the discount factor tends to one. Harrington and Skrzypacz (2011) also study explicit collusion but allow for transfers. All of these papers thus show that communication is *sufficient* for cooperation. But as Kandori and Matsushima (1998) recognize, "One thing which we did *not* show is the *necessity* of communication for a folk theorem (p. 648, their italics)."

In a remarkable paper, Sugaya (2013) shows the surprising result that in very general environments, the folk theorem holds without any communication. Thus, in fact, communication is *not* necessary for a folk theorem. The analysis of repeated games with private monitoring is known to be difficult—and more so if communication is absent. Although Sugaya's result was preceded by folk theorems for some limiting cases where the monitoring was almost perfect (or almost public), the fact that such a result holds even when monitoring is of very low quality is quite unexpected. An important component of Sugaya's proof is that players implicitly communicate via

[3] Zheng (2008) explores a similar monitoring structure in the context of general symmetric games.

their actions.[4] Thus, he shows that with enough time, there is no need for explicit communication.

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.

Key to our result is a method of bounding the set of payoffs under tacit collusion. In a recent paper, Pai, Roth and Ullman (2014) also provide a bound on the equilibrium payoffs that is effective when monitoring quality is low. The measure of monitoring quality used by Pai et al. is based on how the *joint* distribution of the private signals is affected by players' actions. But the bound obtained by them applies to the payoffs from explicit as well as from tacit collusion, and so does not help in distinguishing between the two. In contrast, our measure, and hence the bound in Proposition 1, is based solely on the *marginal* distributions and not on any correlation between players' signals (sales). On the other hand, the equilibrium construction in Proposition 2 relies primarily on the properties of the *joint* distribution. The fact that correlation can vary while keeping the marginal distributions fixed is key to our main result. A method developed by Cherry and Smith (2011) is also unable to distinguish between tacit and explicit collusion.

The monitoring structure we study was introduced by Aoyagi (2002) and then also explored in Zheng (2008) and Awaya (2014a). These papers all assume that the correlation between signals depends on actions in a particular way—it is high when players take similar actions and low when they do not. We also follow Aoyagi (2002) in postulating the way that firms communicate. But our Proposition 2, which constructs an equilibrium with communication, is very different in nature from Aoyagi's result. In his paper, the monitoring structure is fixed and an equilibrium is constructed for discount factors tending to one. In our result, the discounting is held fixed and an equilibrium is constructed for noisy but correlated monitoring structures. As noted above, because of Sugaya's (2013) result, the first exercise is unable to show that communication is necessary for collusion—which is, of course, the goal of this paper.

Our paper provides a *theoretical* basis for distinguishing between tacit and explicit collusion. There is strong *empirical* evidence in support of this distinction that comes from the study of cartels in different industries. Genesove and Mullin (2001) examine this question by looking at a cartel in the sugar refining industry and find strong support that higher prices and profits emerge when firms communicate. Clark and Houde (2014) find the same to be true in the retail gasoline market in Canada. The same conclusion has been reached in laboratory *experiments* as well, by Fonseca and Normann (2012) and Cooper and Kühn (2014) among others.

[4] This idea was used by Hörner and Olzewski (2006) to prove a folk theorem with almost perfect monitoring.

The remainder of the paper is organized as follows. The next section outlines the nature of the market. Section 3 analyzes the repeated game under tacit collusion whereas Section 4 does the same when collusion is explicit. The findings of the earlier sections are combined in Section 5 to derive the main result. In Section 6 we calculate explicitly the gains from communication in an example with linear demands. Omitted proofs are collected in an Appendix.

2 The market

There are two firms in the market, labelled 1 and 2. The firms produce differentiated products at a constant cost, which we normalize to zero. Each firm sets a price $p_i \in P_i = [0, p_{\max}]$, for its product and given the pair of prices p_1 and p_2, the sales Y_1 and Y_2 are stochastic. Prices affect sales via two channels. First, they affect expected sales in the usual way—an increase in p_1 decreases firm 1's expected sales and increases firm 2's expected sales. Second, they affect how correlated are the sales of the two firms in a manner specified below. As in Aoyagi (2002), sales are more correlated when the difference in firms' prices is small.

Expected demand. The *expected* demand of firm i is determined as follows:

$$E[Y_i \mid p_1, p_2] = Q_i(p_i, p_j) \tag{1}$$

where Q_i is a continuous function that is decreasing in p_i and increasing in p_j. We will suppose that the firms are symmetric so that $Q_i = Q_j$. Note that the first argument of Q_i is always the firm's own price and the second is its competitor's price. The *expected* profit of firm i is then

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

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

Let G denote the one-shot game where the firms choose prices p_i and p_j and the 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 the resulting one-shot game and let π_N be the resulting profits of a firm.[5]

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.

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

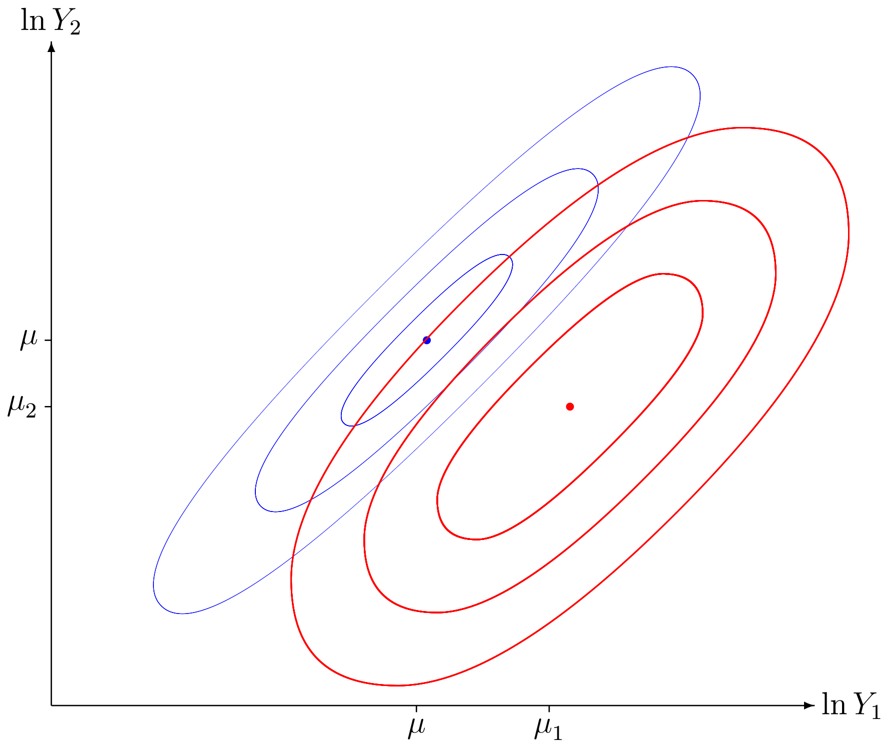

Figure 1: Monitoring Structure

Distribution of sales. We will suppose that given prices (p_1, p_2), the two firms' sales (Y_1, Y_2) are jointly distributed according to a bivariate *log normal* density $f(y_1, y_2 \mid p_1, p_2)$; equivalently, the log sales $(\ln Y_1, \ln Y_2)$ are jointly normally distributed. Prices affect the means μ_i of the normal distribution as well as the correlation coefficient ρ, but, for simplicity, do not affect the (identical) variance σ^2.[6] Specifically, $\mu_i = \ln Q_i(p_i, p_j) - \frac{1}{2}\sigma^2$ so that (1) holds.[7] The correlation between firms' (log) sales is high when they charge similar prices and low when their prices are dissimilar.

For our purposes, a benefit of the (log) normal specification is that the variance and correlation parameters can vary independently.

Figure 1 is a schematic illustration of such a monitoring structure. When both firms charge the same price, their (log) sales have the same mean μ and a high correlation, depicted by the narrower contours of the resulting normal density. When prices are different, say firm 1 charges a lower price, then the (log) sales have different means $\mu_1 > \mu_2$ and low correlation, now depicted by the wider contours. This is formalized as:

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

[7] Recall that if $\ln Y$ is normally distributed with mean μ and variance σ^2, then $E[Y] = \exp\left(\mu + \frac{1}{2}\sigma^2\right)$.

Assumption 1 *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: (1) for all p, $\gamma(p,p) = 1$; and (2) for all $p_1 \leq p_2$, $\partial\gamma/\partial p_1 > 0$ and $\partial^2\gamma/\partial p_1^2 > 0$ and so, γ is an increasing and convex function of p_1.*[8]

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.

This kind of correlation structure can result, for example, in a symmetric Hotelling-type market in which consumers have identical but random "transport costs". When firms charge similar prices, their sales are similar no matter what the realized transport costs are. In other words, when firms charge similar prices, their sales are highly correlated. When firms charge dissimilar prices, their sales are again similar if the realized transport costs are high because consumers are not so price sensitive. But their sales are quite dissimilar if the realized costs are low because now consumers are very price sensitive. 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.

3 Tacit collusion

Let $G_\delta(f)$ denote the infinitely repeated game with *private monitoring* 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 $h_i^{t-1} = \left(p_i^1, y_i^1, p_i^2, y_i^2, ..., p_i^{t-1}, y_i^{t-1}\right)$ denote the *private* history observed by firm i after $t-1$ periods of play and let 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 h_i^{t-1} and nothing else.

A *strategy* s_i for firm i is a collection of functions $(s_i^1, s_i^2, ...)$ such that $s_i^t : H_i^{t-1} \rightarrow \Delta(P_i)$. Of course, since H_i^0 is null, $s_i^1 \in \Delta(P_i)$. We will denote by $s_i^t\left(p_i \mid h_i^{t-1}\right)$ the probability that firm i sets a price p_i following the private history h_i^{t-1}. Thus, we are allowing for the possibility that firms may randomize. A *strategy profile* s is simply a pair of strategies (s_1, s_2).

A *sequential equilibrium* of $G_\delta(f)$ is strategy profile s such that for each i and every private history h_i^{t-1} such that the continuation strategy of i following h_i^{t-1}, denoted by $s_i \mid_{h_i^{t-1}}$, is a best response to $E[s_j \mid_{h_j^{t-1}} \mid h_i^{t-1}]$. Since f has full support, the set of sequential equilibrium outcomes (price paths) is the same as the set of Nash equilibrium outcomes (see Mailath and Samuelson, 2006, p. 396).

We remind the reader that as yet there is no communication between the firms.

[8]Some examples satisfying the assumption are $\gamma(p_1, p_2) = \min(p_1, p_2)/\max(p_1, p_2)$ and $\gamma(p_1, p_2) = 1/1 + |p_1 - p_2|$.

3.1 Equilibrium under tacit collusion

The purpose of this subsection is to provide an upper bound to the joint profits of the firms in *any* equilibrium under tacit collusion. 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 and sales—such games lack a straightforward recursive structure and the kinds of techniques available to analyze (public perfect) equilibria of repeated games with *public* monitoring (see Abreu, Pearce and Stacchetti, 1990) cannot be used here.

Instead, we proceed as follows. Suppose we want to determine whether there is an 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 $\overline{p}$, 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 2 can keep on deviating to $\overline{p}$ 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 bound on equilibrium profits depends on three factors of the market: (1) the trade-off between the incentives to deviate and efficiency in the *one-shot* game[9]; (2) the quality of the monitoring, which determines whether the short-term incentives to deviate can be overcome by future actions; and, of course (3) the discount factor.

We consider each of these factors in turn.

Incentives versus efficiency in the one-shot game. Define, as above, $\overline{p} = \arg\max_{p_i} \pi_i(p_i, p_M)$, the static best-response to p_M. Let $\alpha \in \Delta(P_1 \times P_2)$ 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 $\overline{p}$. To that end, for $\varepsilon \geq 0$, define

$$\Psi(\varepsilon) \equiv \min_{\alpha} \sum_i \left[\pi_i(\overline{p}, \alpha_j) - \pi_i(\alpha)\right] \tag{2}$$

subject to

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

where α_j denotes the marginal distribution of α over P_j.

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

The function Ψ measures the trade-off between the incentives to deviate (to the price $\overline{p}$) 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. Lemma A.1 establishes that it is convex and satisfies $\lim_{\varepsilon \to 0} \Psi(\varepsilon) > 0$.

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 Ψ^{-1} by

$$\Psi^{-1}(x) = \sup \{\varepsilon : \Psi(\varepsilon) = x\} \tag{3}$$

Quality of monitoring. Consider two price pairs $p = (p_1, p_2)$ and $p' = (p'_1, p'_2)$ and the resulting distributions of firm i's sales: $f_i(\cdot \mid p)$ and $f_i(\cdot \mid p')$. If these two distributions are close together, then it will be difficult for firm i to detect the change from p to p'. Thus, the quality of monitoring can be measured by the "distance" between the two distributions. In what follows, we use the so-called total variation metric to measure this distance.

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

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

where f_i is the marginal of f on Y_i 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 $f_i(\cdot \mid p)$ over i's sales and not on the joint distributions of sales $f(\cdot \mid p)$. In particular, the fact that the marginal distributions $f_i(\cdot \mid p)$ and $f_i(\cdot \mid p')$ are close—η is small—does not imply that the underlying joint distributions $f(\cdot \mid p)$ and $f(\cdot \mid p')$ are close. When $f(\cdot \mid p)$ is a bivariate log normal, η can be explicitly determined as

$$\eta = 2\Phi\left(\tfrac{\Delta\mu_{\max}}{2\sigma}\right) - 1 \tag{4}$$

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

3.1.1 A bound on tacit collusion

The main result of this section develops a bound on equilibrium profits under tacit collusion. 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.

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

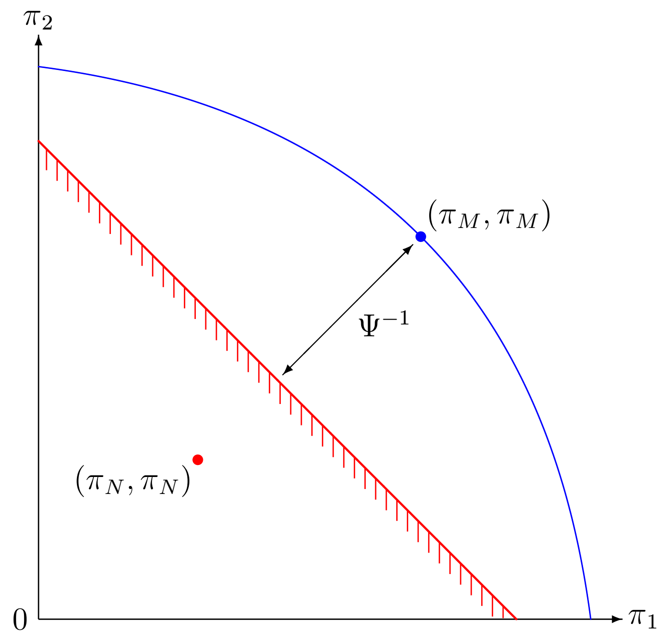

Figure 2: Bound on Profits from Tacit Collusion

Proposition 1 *In any equilibrium of $G_\delta(f)$, the repeated game without communication, the profits*

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

where $\overline{\pi} = \max_{p_j} \pi_i(\overline{p}, p_j)$ *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 that a deviation by firm 1 to a strategy $\overline{s}_1$ in which it *always* charges $\overline{p}$ 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 A.2 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 A.6 measures precisely how good this approximation is and quite naturally this depends on the quality of monitoring and the discount factor.

It is usual to derive necessary conditions for an equilibrium by considering "one-shot" deviations in which a player cheats in one period and then resumes equilibrium play.[11] But in games with private monitoring, such deviations affect the deviating

[11] Pai et al. (2014) consider such a deviation.

players' beliefs about the other player's signals thereafter and so affect his subsequent (optimal) play. The deviation we consider—a permanent price cut—is rather naive but has the feature that future play, while suboptimal, is straightforward. 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 $\lim_{\delta\to 1}\Psi^{-1}\left(4\overline{\pi}\frac{\delta^2}{1-\delta}\eta\right)=0$) and so is consistent with the 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 $\eta \to 0$, 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).[12]

Proof of Proposition 1. We argue by contradiction. Suppose that $G_\delta(f)$ has an equilibrium, say s, whose average profits[13] $\pi_1(s)+\pi_2(s)$ exceed $2\pi_M-\Psi^{-1}\left(4\overline{\pi}\frac{\delta^2}{1-\delta}\eta\right)$. If we write $\varepsilon = 2\pi_M - \pi_1(s) - \pi_2(s)$, then this is equivalent to $\eta < \frac{1-\delta}{\delta^2}\frac{1}{4\overline{\pi}}\Psi(\varepsilon)$.

Given the strategy profile s, define

$$\alpha_j^t = E_s\left[s_j^t\left(h_j^{t-1}\right)\right] \in \Delta(P_j)$$

where the expectation is defined by the probability distribution over $t-1$ joint histories $\left(h_i^{t-1}, h_j^{t-1}\right)$ determined by s. Note that α_j depends on the strategy *profile* s and not just s_j. Let $\alpha_j = \left(\alpha_j^1, \alpha_j^2, ...\right)$ 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 *non-responsive* to histories.

Let $\overline{s}_i$ denote the strategy of firm i in which it plays $\overline{p}$ with probability one following any history. From Lemma A.2

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

From Lemma A.6 we have for $i = 1, 2$

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

[12] Awaya (2014b) constructs an example in which there is zero monitoring—the distribution of a player's signals is the same for all action profiles—but, nevertheless, there are non-trivial equilibria.

[13] 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.

and

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

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

4 Explicit collusion

We now turn to a situation in which firms can explicitly collude. By this we mean that prior to choosing prices in any period, firms can communicate with each other, 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. 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

$$h_i^{t-1} = \left(p_i^\tau, y_i^\tau, m_i^\tau, m_j^\tau\right)_{\tau=1}^{t-1}$$

and the set of all such histories is denoted by H_i^{t-1}. A strategy for firm i is now a pair (s_i, r_i) where $s_i = (s_i^1, s_i^2, ...)$, the pricing strategy, and $r_i = (r_i^1, r_i^2, ...)$, the reporting strategy, are collections of functions: $s_i^t : H_i^{t-1} \to \Delta\left(P_i\right)$ and $r_i^t : H_i^{t-1} \times P_i \times Y_i \to \Delta\left(M_i\right)$.

Call the resulting infinitely repeated game with communication $G_\delta^{com}\left(f\right)$. Sequential equilibrium is defined as before.

4.1 Equilibrium strategies

We will now identify some properties of the monitoring structure f that will allow the firms to achieve near-perfect collusion, that is, the sum of their profits will be close to those of a monopolist. The log-normal monitoring structure has two parameters—the variance of log-sales σ^2 and the correlation between log-sales, ρ_0, when the firms charge identical prices. Of course, the discount factor δ is key parameter as well.

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" or both report "low". Differing sales reports trigger permanent non-cooperation 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 communication strategy r_i^* is:

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

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

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. But before doing this, it is useful to calculate the lifetime average profits if firms follow the proposed strategies.

4.1.1 Optimality of communication 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

$$\Pr\left[\ln Y_1 < \lambda, \ln Y_2 < \mu_M\right] + \Pr\left[\ln Y_1 > \lambda, \ln Y_2 > \mu_M\right] \tag{5}$$

The *optimal* reporting threshold is (see Lemma A.7)

$$\lambda\left(p_1\right) = \mu_1 + \frac{1}{\rho}\left(\mu_M - \mu_2\right) \tag{6}$$

If firm 1 deviated and cut its price to $p_1 < p_M$, then clearly the expected (log) sales of the two firms are 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 *under-report* relatively to the equilibrium reporting strategy r_1^*. On the other hand, if firm 1 did not deviate and set a price p_M, then the (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.

Thus, 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

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

where $\phi(z_1, z_2; \rho_0)$ is a standard bivariate normal density with correlation coefficient $\rho_0 \in (0,1)$.[14] Note that while $\lambda(p_1)$, μ_1, μ_2 and the correlation coefficient ρ depend on p_1, σ is independent of p_1. Observe also that for any $p_1 \le p_M$, $\beta(p_1) \ge \Phi\left(\frac{\mu_M-\mu_2}{\sigma}\right) \ge \frac{1}{2}$ where Φ denotes the cumulative distribution function of the standard univariate normal. This is because firm 1 could always adopt a communication strategy in which after a deviation to $p_1 < p_M$, it always reports 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 $\Phi\left(\frac{\mu_M-\mu_2}{\sigma}\right)$. Since $\mu_M \ge \mu_2$, $\Phi\left(\frac{\mu_M-\mu_2}{\sigma}\right) \ge \frac{1}{2}$. This means that the probability of detecting a deviation is less than one-half.

4.1.2 Equilibrium profits

If both set prices p_M and follow the proposed reporting strategy, 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

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

which is increasing in ρ_0 and converges to 1 as ρ_0 goes to 1.

The lifetime average profit π^* resulting from the proposed strategies is given by

$$(1-\delta)\pi_M + \delta\left[\beta(p_M)\pi^* + (1-\beta(p_M))\pi_N\right] = \pi^* \tag{8}$$

[14] The standard (with both means 0 and both variances 1) bivariate normal density is $\phi(z_1, z_2; \rho) = \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left(-\frac{1}{2(1-\rho^2)}\left(z_1^2 + z_2^2 - 2\rho z_1 z_2\right)\right)$

and it is easy to see that for any fixed δ,

$$\lim_{\rho_0 \to 1} \pi^* = \pi_M$$

4.1.3 Equilibrium with communication

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. 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

$$\Delta_1(p_1) = (1-\delta)\pi_1(p_1, p_M) + \delta[\beta(p_1)\pi^* + (1-\beta(p_1))\pi_N] - \pi^* \tag{9}$$

where π^* is defined in (8). Thus,

$$\Delta_1'(p_1) = (1-\delta)\frac{\partial \pi_1}{\partial p_1}(p_1, p_M) + \delta\beta'(p_1)[\pi^* - \pi_N]$$

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 A.8,

$$\begin{aligned}
\lim_{\sigma\to\infty} \Delta_1'(p_1) &= (1-\delta)\tfrac{\partial \pi_1}{\partial p_1}(p_1, p_M) + \delta\tfrac{1}{\pi\sqrt{1-\rho_0^2\gamma(p_1,p_M)^2}} \times \rho_0 \tfrac{\partial \gamma}{\partial p_1}(p_1, p_M) \times [\pi^* - \pi_N] \\
&\geq (1-\delta)\tfrac{\partial \pi_1}{\partial p_1}(p_M, p_M) + \delta\tfrac{1}{\pi\sqrt{1-\rho_0^2\gamma(0,p_M)^2}} \times \rho_0 \tfrac{\partial \gamma}{\partial p_1}(0, p_M) \times [\pi^* - \pi_N]
\end{aligned}$$

where the last inequality follows from 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 $\underline{\delta}$ be the solution to

$$(1-\delta)\tfrac{\partial \pi_1}{\partial p_1}(p_M, p_M) + \delta\tfrac{1}{\pi\sqrt{1-\gamma(0,p_M)^2}} \times \tfrac{\partial \gamma}{\partial p_1}(0, p_M) \times [\pi^* - \pi_N] = 0 \tag{10}$$

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

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

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 all $\rho_0 = \rho_0(\delta)$

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

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.

This completes the proof. ■

Aoyagi (2002) was the first to introduce threshold reporting strategies. He shows that for a given monitoring structure (ρ_0 and σ fixed) as the discount factor δ goes to one, these strategies constitute an equilibrium. The idea—as in all the "folk theorems"—is that even when the probability of a deviation being detected is low, if players are patient enough, future punishments are a sufficient deterrent even if they are distant.

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". Here the punishment power derives not from the patience of the players; rather it comes from the noisiness of the monitoring. A deviating firm will then find it very difficult to predict its rival's sales and hence, even it "lies" optimally, a deviation is very likely to trigger a punishment.

5 Gains from communication

Proposition 1 shows that the profits from any equilibrium under tacit collusion cannot exceed

$$2\pi_M - \Psi^{-1}\left(4\overline{\pi}\tfrac{\delta^2}{1-\delta}\eta\right)$$

whereas Proposition 2 provides conditions under which there is an equilibrium under explicit collusion that with profits $2\pi^*$ (as defined in (8)). The two results together lead to the formal version of the result stated in the introduction. Let $\underline{\delta}$ be determined as in (10).

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 under explicit collusion with total profits $2\pi^*$ such that*

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

As $\rho_0 \to 1$, $\pi^* \to \pi_M$ and as $\sigma \to \infty$, $\eta \to 0$. Thus, in the limit the difference in profits between explicit and tacit collusion is at least $\Psi^{-1}(0) > 0$.

The workings of the main result can be seen in Figure 3, which is drawn for the case of linear demand (see the next section for details). First, notice that the profits from explicit collusion 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 from tacit collusion, on the other hand, depends on σ and not on ρ_0. For small values

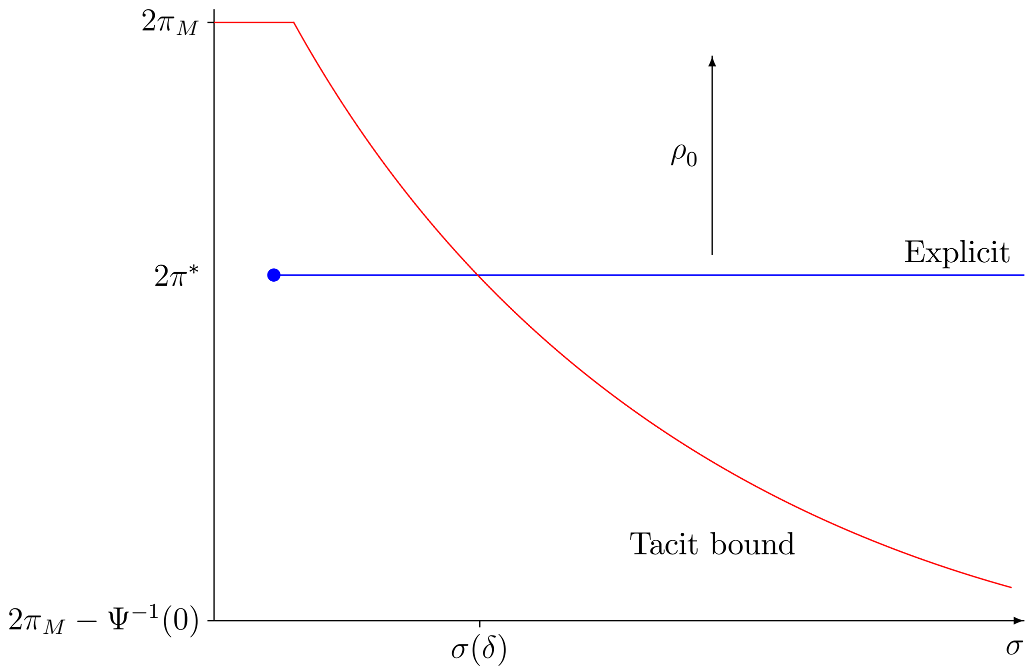

Figure 3: Gains from Communication

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 under explicit collusion do not exceed the bound. Once $\sigma > \sigma(\delta)$, explicit collusion results in strictly higher profits than tacit collusion. In the limit, the bound is $2\pi_M - \Psi^{-1}(0)$.

6 Linear demand

In this section, we illustrate the workings of our results when (expected) demand is linear.

Suppose that[15]

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

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 $\pi_N = A^2 b/(2b-1)^2$. A firm's best response if the other firm charges the monopoly price p_M is $\overline{p} = A(2b-1)/4b(b-1)$. The highest possible profit that firm 1 can achieve when charging a price of $\overline{p}$ is $\overline{\pi} = \pi_1(\overline{p}, p_M) = A^2(2b-1)^2/16b(b-1)^2$.

It remains to specify how the correlation between the firms' log sales is affected

[15] This specification of "linear" demand is used because $\ln 0$ is not defined.

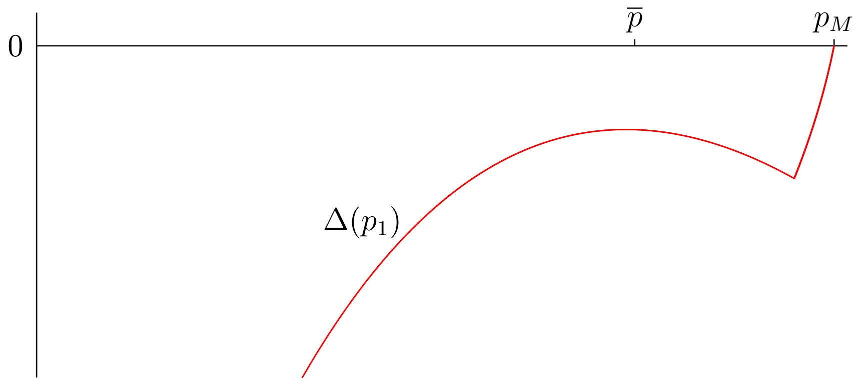

Figure 4: Unprofitable Deviations

by prices. In this example we adopt the following specification:

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

which, of course, satisfies Assumption 1.

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

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

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 (4)

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

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

A numerical example Suppose $A = 120$ and $b = 2$. Let $\delta = 0.7$ and $\rho_0 = 0.95$. For these parameters, $\pi_M = 3600$, $\overline{\pi} = 4050$ and $\Delta\mu_{\max} = 5.19$. Also, the profits from the equilibrium under explicit collusion, $\pi^* = 3524$ (approximately).

Figure 3 depicts the bound on profits from tacit collusion 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 $\sigma \to \infty$, converges to $2\pi_M - \Psi^{-1}(0)$. As shown, the profits under explicit collusion exceed the bound when $\sigma > \sigma(\delta) = 200$ (approximately).

Figure 4 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 (9)). This is verified for $\sigma = 60$ and, of course, the same strategies remain an equilibrium for higher values of σ.

7 Conclusion

We have provided *theoretical* support for the idea that communication facilitates greater collusion. Where do the gains from communication come from? When the monitoring is poor, under tacit collusion, a price cut cannot be detected with any confidence and this is the basis of the bound developed in Proposition 1. Under explicit collusion, 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 the main result.

We conjecture that the fact that communication facilitates greater collusion holds quite generally beyond the circumstances we have identified in this paper—that the monitoring quality be low and this in turn requires that sales be rather volatile. We view our result as only a first step towards distinguishing between the two forms of collusion and recognize its limitations.

First, we did not identify the best equilibrium under explicit collusion; we only constructed *an* equilibrium. This equilibrium was based on very simple grim trigger strategies and these, because of their unforgiving nature, are known to perform badly. Moreover, the communication strategies are not very good at detecting deviations—the probability that a price cut will be trigger a punishment is less than one-half.

Second, the upper bound on profits under tacit collusion provided here bites only when the monitoring quality is rather poor. The development of better payoff bounds for repeated games with private monitoring remains a challenge.

A Appendix

A.1 Tacit collusion

The first lemma derives some simple properties of the function Ψ. This function delineates the trade-off between efficiency and incentives in the one-shot game and is central to the bound for equilibrium payoffs of the repeated game developed below.

Lemma A.1 *Ψ is non-increasing, convex and satisfies* $\lim_{\varepsilon\to 0}\Psi(\varepsilon) = \Psi(0) > 0$.

Proof. The fact that Ψ is non-increasing follows trivially from its definition. To see that Ψ is convex, note that

$$\pi_i(\alpha) = \int \pi_i(p_i, p_j)\, d\alpha(p_i, p_j)$$

Suppose α' is the solution to the program above for $\varepsilon = \varepsilon'$ and similarly, suppose α'' is the solution for $\varepsilon = \varepsilon''$. Then since the constraint is a linear function of α, for any $\theta \in [0,1]$, $\theta\alpha' + (1-\theta)\alpha''$ is feasible for $\varepsilon = \theta\varepsilon' + (1-\theta)\varepsilon''$. The convexity of Ψ now follows since the objective function is also linear in α.

The fact that $\Psi(\varepsilon)$ converges to $\Psi(0)$ as $\varepsilon \to 0$ follows from the Berge Maximum Theorem. To see that $\Psi(0) > 0$, note that for $\varepsilon = 0$ the only feasible solution to the optimization problem defining Ψ is (p_M, p_M). Since this is not a Nash equilibrium, $\Psi(0) > 0$. ■

A.1.1 Non-responsive strategies

The induced ex ante distribution over P_j in period t induced by a strategy profile s is

$$\alpha_j^t(s) = E_s\left[s_j^t\left(h_j^{t-1}\right)\right] \in \Delta(P_i) \tag{11}$$

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

The following lemma shows that the function Ψ, which determines the incentives versus efficiency trade-off in the one-shot game, embodies the same trade-off in a *repeated* setting if the non-deviating player follows a non-responsive 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 A.2 (Smoothing) *For any strategy profile s whose profits are greater than $2\pi_M - \varepsilon$,*

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

where $\overline{s}_i$ denote the strategy of firm i in which it plays $\overline{p}$ with probability one following any history.

Proof. Define

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

as the difference between the sum of efficient 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,

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

where the first equality follows from the fact that the induced distribution over prices p_j^t is the same under $(\overline{s}_i, \alpha_j)$ as it is under s. The second inequality follows from the definition of Ψ.

Now note that Lemma A.1 guarantees that a solution to the problem

$$\min_{\{\varepsilon(t)\}} (1-\delta) \sum_{t=1}^{\infty} \delta^t \Psi\left(\varepsilon\left(t\right)\right)$$

subject to

$$(1-\delta) \sum_{t=1}^{\infty} \delta^t \varepsilon\left(t\right) \leq \varepsilon$$

is to set $\varepsilon\left(t\right) = \varepsilon$ for all t. Thus, we have that

$$(1-\delta) \sum_{t=1}^{\infty} \delta^t \Psi\left(\varepsilon\left(t\right)\right) \geq \Psi\left(\varepsilon\right)$$

■

A.1.2 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 $h_j^t \in H_j^t = (P_j \times Y_j)^t \subset R^{2t}$. Similarly, let $\overline{\lambda}_j^t$ be the probability distribution over j's private histories that results from the strategy pair $(\overline{s}_i, s_j)$.[16] We wish to determine the total variation distance between λ_j^t and $\overline{\lambda}_j^t$.

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

$$\left\|G - \overline{G}\right\|_{TV} = \tfrac{1}{2} \sup_{\|\varphi\|_\infty \leq 1} \left|E\left[\varphi\right] - \overline{E}\left[\varphi\right]\right| \tag{12}$$

where E and $\overline{E}$ denote the expectations with respect to the distribution G and $\overline{G}$, respectively and the supremum is taken over all measurable functions φ with sup norm $\|\varphi\|_\infty \leq 1$. Note that the definition in (12) 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 A.3 *Given two distributions G and $\overline{G}$ over $R^m \times R^n$,*

$$\left\|G - \overline{G}\right\|_{TV} \leq \left\|G_X - \overline{G}_X\right\|_{TV} + \sup_x \left\|G_{Y|X} - \overline{G}_{Y|X}\right\|_{TV}$$

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

[16] Recall that $\overline{s}_i$ denotes the strategy of firm i in which it sets $\overline{p}$ with probability one following any history.

Proof. In what follows we denote by E all expectations with respect to G and by $\overline{E}$, all expectations with respect to $\overline{G}$.

Given any function $\varphi : R^m \times R^n \to [-1, 1]$, we have

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

and so

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

■

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

Lemma A.4 *For any* h_j^{t-1},

$$
\left\|\lambda_j^t\left(\cdot \mid h_j^{t-1}\right)-\overline{\lambda}_j^t\left(\cdot \mid h_j^{t-1}\right)\right\|_{TV} \leq \eta
$$

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

Now, for any $\varphi : P_j \times Y_j \to [-1, 1]$

$$
\begin{aligned}
& \left| \int_{P_j \times Y_j} \varphi d\lambda_j^t \left(\cdot \mid h_j^{t-1} \right) - \int_{P_j \times Y_j} \varphi d\overline{\lambda}_j^t \left(\cdot \mid h_j^{t-1} \right) \right| \\
= & \left| \int_{P_i} \int_{P_j} \int_{Y_j} \varphi dF_j \left(y_j \mid p \right) dS_j \left(p_j \mid h_j^{t-1} \right) d\widehat{S}_i \left(p_i \mid h_j^{t-1} \right) \right. \\
& \left. - \int_{P_j} \int_{Y_j} \varphi dF_j \left(y_j \mid \overline{p}, p_j \right) dS_j \left(p_j \mid h_j^{t-1} \right) \right| \\
= & \left| \int_{P} \int_{Y_j} \varphi dF_j \left(y_j \mid p \right) d\widehat{S} \left(p \mid h_j^{t-1} \right) - \int_{P} \int_{Y_j} \varphi dF_j \left(y_j \mid \overline{p}, p_j \right) d\widehat{S} \left(p \mid h_j^{t-1} \right) \right| \\
= & \left| \int_{P} \int_{Y_j} \varphi \left[f_j \left(y_j \mid p \right) - f_j \left(y_j \mid \overline{p}, p_j \right) \right] dy_j d\widehat{S} \left(p \mid h_j^{t-1} \right) \right| \\
\leq & \int_{P} \left[\int_{Y_j} \left| \varphi f_j \left(y_j \mid p \right) - \varphi f_j \left(y_j \mid \overline{p}, p_j \right) \right| dy_j \right] d\widehat{S} \left(p \mid h_j^{t-1} \right) \\
\leq & \ 2 \int_{P} \eta d\widehat{S} \left(p \mid h_j^{t-1} \right) \\
= & \ 2\eta
\end{aligned}
$$

where the second equality follows from the fact that

$$
\int_{P_j} \int_{Y_j} \varphi dF_j \left(y_j \mid \overline{p}, p_j \right) dS_j \left(p_j \mid h_j^{t-1} \right) = \int_{P_i} \int_{P_j} \int_{Y_j} \varphi dF_j \left(y_j \mid \overline{p}, p_j \right) dS_j \left(p_j \mid h_j^{t-1} \right) d\widehat{S}_i \left(p_i \mid h_j^{t-1} \right)
$$

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

Thus,

$$
\begin{aligned}
\left\| \lambda_j^t \left(\cdot \mid h_j^{t-1} \right) - \overline{\lambda}_j^t \left(\cdot \mid h_j^{t-1} \right) \right\|_{TV} &= \frac{1}{2} \sup_{\|\varphi\|_\infty \leq 1} \left| \int_{P_j \times Y_j} \varphi d\lambda_j^t \left(\cdot \mid h_j^{t-1} \right) - \int_{P_j \times Y_j} \varphi d\overline{\lambda}_j^t \left(\cdot \mid h_j^{t-1} \right) \right| \\
&\leq \eta
\end{aligned}
$$

■

Combining the preceding two results we obtain

Lemma A.5 *For all t,*

$$
\left\| \lambda_j^t - \overline{\lambda}_j^t \right\|_{TV} \leq t\eta
$$

Proof. The proof is by induction. For $t = 1$, there is no history and Lemma A.4 implies the result directly.

Now suppose that the result holds for $t-1$. Using Lemma A.3, we have

$$\begin{aligned}\left\|\lambda_j^t - \overline{\lambda}_j^t\right\|_{TV} &\leq \left\|\lambda_j^{t-1} - \overline{\lambda}_j^{t-1}\right\|_{TV} + \sup_{h_j^{t-1}} \left\|\lambda_j^t\left(\cdot \mid h_j^{t-1}\right) - \overline{\lambda}_j^t\left(\cdot \mid h_j^{t-1}\right)\right\|_{TV} \\ &\leq (t-1)\,\eta + \eta\end{aligned}$$

by the induction hypothesis and Lemma A.4. ■

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 non-responsive strategy. The importance of the lemma is in quantifying this difference.

Lemma A.6 *Let α be the non-responsive strategy as defined in (11). Then,*

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

where $\overline{\pi} = \max_{p_j} \pi_i\left(\overline{p}, p_j\right)$.

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

$$\pi_i\left(\overline{s}_i, s_j\right) = (1-\delta)\sum_{t=1}^{\infty}\delta^t \int_{H_j^{t-1}} E\left[\pi_i\left(\overline{p}, s_j\right) \mid h_j^{t-1}\right] d\overline{\lambda}_j\left(h_j^{t-1}\right)$$

Also, if $S_j\left(\cdot \mid h_j^{t-1}\right)$ denotes the distribution over j's prices induced by the strategy s_j following the history h_j^{t-1}, then

$$\begin{aligned}\pi_i\left(\overline{s}_i, \alpha_j\right) &= (1-\delta)\sum_{t=1}^{\infty}\delta^t \int_{P_j} \pi_i\left(\overline{p}, p_j\right) d\alpha_j^t\left(p_j\right) \\ &= (1-\delta)\sum_{t=1}^{\infty}\delta^t \int_{P_j} \pi_i\left(\overline{p}, p_j\right)\left(\int_{H_j^{t-1}} dS_j\left(p_j \mid h_j^{t-1}\right) d\lambda_j\left(h_j^{t-1}\right)\right) \\ &= (1-\delta)\sum_{t=1}^{\infty}\delta^t \int_{H_j^{t-1}} E\left[\pi_i\left(\overline{p}, s_j\right) \mid h_j^{t-1}\right] d\lambda_j\left(h_j^{t-1}\right)\end{aligned}$$

since by definition

$$\alpha_j^t\left(p_j\right) = \int_{H_j^{t-1}} S_j\left(p_j \mid h_j^{t-1}\right) d\lambda_j\left(h_j^{t-1}\right)$$

Thus,

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

where the second inequality is a consequence of Lemma A.5 and the fact that, as in (12), given any two distributions λ and $\overline{\lambda}$, $\left|E[\varphi]-\overline{E}[\varphi]\right| \leq 2\|\varphi\|_{\infty} \times\left\|\lambda-\overline{\lambda}\right\|_{TV}$ for any bounded measurable function φ. ■

A.2 Explicit collusion

Lemma A.7 *Suppose firm 2 follows the strategy $\left(s_2^*, r_2^*\right)$. Following a price of p_1, the optimal reporting threshold for firm 1 is*

$$\lambda\left(p_1\right)=\mu_1+\frac{1}{\rho}\left(\mu_M-\mu_2\right)$$

Proof. If firm 1 sets a price of p_1, and uses a reporting threshold of λ, then the probability that the sales reports agree (see (5)) can be rewritten (after standardizing the variables) as

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

where ϕ is a standard bivariate normal density with correlation coefficient $\rho \in(0,1)$ of the form in Assumption 1.

Maximizing this with respect to λ results in the first-order condition

$$\int_{-\infty}^{\frac{\mu_M-\mu_2}{\sigma}} \phi\left(\tfrac{\lambda-\mu_1}{\sigma}, z_2 ; \rho\right) d z_2=\int_{\frac{\mu_M-\mu_2}{\sigma}}^{\infty} \phi\left(\tfrac{\lambda-\mu}{\sigma}, z_2 ; \rho\right) d z_2$$

Dividing by the marginal density of Z_1 at $\frac{\lambda-\mu_1}{\sigma}$, and writing in terms of the cumulative distribution, we obtain

$$\Phi_{Z_2 \mid Z_1}\left(\tfrac{\mu_M-\mu_2}{\sigma} \mid \tfrac{\lambda-\mu_1}{\sigma}\right)=1-\Phi_{Z_2 \mid Z_1}\left(\tfrac{\mu_M-\mu_2}{\sigma} \mid \tfrac{\lambda-\mu_1}{\sigma}\right) \tag{13}$$

This says that the optimal λ is such that $\frac{\mu_M-\mu_2}{\sigma}$ is the *median* of the distribution $\Phi_{Z_2|Z_1}$ of Z_2 conditional on $Z_1 = \frac{\lambda-\mu_1}{\sigma}$. But since $\phi_{Z_2|Z_1}$ is also normal, its median is the same as its mean, $\rho\frac{\lambda-\mu_1}{\sigma}$. Thus the optimal strategy is to choose λ such that

$$\tfrac{\mu_M-\mu_2}{\sigma} = \rho\tfrac{\lambda-\mu_1}{\sigma}$$

from which the result follows. ■

Lemma A.8 *For any* $p_1 \geq 1$,

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

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

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

where

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

using Sheppard's formula[17] (see Tihansky, 1972) and Φ is the cumulative distribution function of a standard univariate normal. Thus,

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

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

[17] This employs a change of variables to the original formula.

Finally,

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

in a manner analogous to (13). This converges to 0 as $\sigma \to \infty$.

This completes the proof. ■

References

[1] Abreu, D., D. Pearce and E. Stacchetti (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58 (5), 1041–1063.

[2] Aoyagi, M. (2002): "Collusion in Dynamic Bertrand Oligopoly with Correlated Private Signals and Communication," *Journal of Economic Theory*, 102, 229–248.

[3] Athey, S. and K. Bagwell (2001): "Optimal Collusion with Private Information," *Rand Journal of Economics*, 32 (3), 428–465

[4] Awaya, Y. (2014a): "Private Monitoring and Communication in Repeated Prisoners' Dilemma," Working Paper, Penn State University, http://www.personal.psu.edu/yxa120/research.html

[5] Awaya, Y (2014b): "Cooperation without Monitoring," Working Paper, Penn State University, http://www.personal.psu.edu/yxa120/research.html

[6] Cherry, J. and L. Smith (2011): "Unattainable Payoffs for Repeated Games of Private Monitoring," Working Paper, University of Wisconsin, http://lonessmith.com/sites/default/files/unattainable_new.pdf

[7] Clark, R. and J-F. Houde (2014): "The Effect of Explicit Communication on Pricing: Evidence from the Collapse of a Gasoline Cartel," *Journal of Industrial Economics*, 62 (2), 191–227.

[8] Compte, O. (1998): "Communication in Repeated Games with Imperfect Private Monitoring," *Econometrica*, 66 (3), 597–626.

[9] Cooper, D. and K-U. Kühn (2014): "Communication, Renegotiation, and the Scope for Collusion," *American Economic Journal: Microeconomics*, 6(2): 247–278.

[10] Escobar, J. and J. Toikka (2013): "Efficiency in Games with Markovian Private Information," *Econometrica*, 81 (5), 1887–1934.

[11] Fonseca, M. and H-T. Normann (2012): "Explicit vs. Tacit Collusion—The Impact of Communication in Oligopoly Experiments," *European Economic Review*, 56 (8), 1759–1772.

[12] Fudenberg, D. and D. Levine (2007): "The Nash-Threats Folk Theorem with Communication and Approximate Common Knowledge in Two Player Games," *Journal of Economic Theory*, 132, 461–473.

[13] Genesove, D. and W. Mullin (2001): "Rules, Communication, and Collusion: Narrative Evidence from the Sugar Institute Case," *American Economic Review*, 91 (3), 379–398.

[14] Green, E., R. Marshall and L. Marx (2013): "Tacit Collusion in Oligopoly," *Handbook of International Antitrust Economics*, R. Blair and D. Sokol (eds.), Oxford University Press (forthcoming).

[15] Harrington, J. (2005): "The Collusion Chasm," Lecture presented at the Csef-Igier Symposium on Economics and Institutions, http://assets.wharton.upenn.edu/~harrij/pdf/Harrington_Csef-Igier%202005.pdf‘

[16] Harrington, J. and A. Skrzypacz (2011): "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices," *American Economic Review*, 101 (6), 2425–2449.

[17] Hörner, J., and W. Olszewski (2006): "The Folk Theorem for Games with Private Almost-Perfect Monitoring," *Econometrica*, 74 (6), 1499–1544.

[18] Kandori, M. and H. Matsushima (1998): "Private Observation, Communication and Collusion," *Econometrica*, 66 (3), 627–652.

[19] Levin, D., Y. Peres and E. Wilmer (2009): *Markov Chains and Mixing Times*, American Mathematical Society.

[20] Mailath, G. and L. Samuelson (2006): *Repeated Games and Reputations*, Oxford University Press.

[21] Marshall, R. and L. Marx (2012): *The Economics of Collusion,* MIT Press.

[22] Obara, I. (2009): "Folk Theorem with Communication," *Journal of Economic Theory*, 144, 120–134.

[23] Pai, M., A. Roth and J. Ullman (2014): "An Anti-Folk Theorem for Large Repeated Games with Imperfect Monitoring," http://arxiv.org/abs/1402.2801v2 [cs.GT]

[24] Stigler, G. (1964): "A Theory of Oligopoly," *Journal of Political Economy*, 74, 44–61.

[25] Sugaya, T. (2013): "Folk Theorem in Repeated Games with Private Monitoring," Working Paper, Stanford University, https://sites.google.com/site/takuosugaya/home/research

[26] Tihansky, D. P. (1972): "Properties of the Bivariate Normal Cumulative Distribution," *Journal of the American Statistical Association*, 67 (340), 903–905.

[27] Zheng, B. (2008): "Approximate Efficiency in Repeated Games with Correlated Private Signals," *Games and Economic Behavior*, 63, 406–416.