Collusion Enforcement with Private Information and Private Monitoring*

Jimmy H. Chan Wenzhang Zhang [†]

Feburary 2012

Abstract

This paper describes how a cartel may enforce a collusive agreement even when it does not observe any of the prices, sales, and costs of the firms. The underlying mechanism applies to both price and quantity competition, allows a wide class of stochastic demand systems, and can be generalized to the case of multi-market collusion.

*An earlier version of this paper has been circulated under the title "Price Wars, Side-Payments, and Collusion Enforcement." We thank Joseph Harrington, Andrzej Skrzypacz, and participants in various seminars for comments and suggestions. Part of this research is funded by the Chinese National Science Foundation (Project No. 71171125). Chan is supported by the Shanghai Dongfang Xuezhe Program.

[†]School of Economics, Shanghai University of Finance and Economics, Shanghai 200433, China. E-mail: jimmy.hing.chan@gmail.com (Chan) and wenzhang.zhang@gmail.com (Zhang).

1 Introduction

Ever since Stigler (1964), a focus of cartel theory is to explain how cartels could enforce collusive agreements given their limited ability to monitor their members. There are two strands of literature. One strand assumes that prices are publicly observed and focuses on the issue of private cost information. Athey and Bagwell (2001) and Athey, Bagwell, and Sanchirico (2004) study the optimal collusion scheme in dynamic Bertrand games in which prices are publicly observed and each firm receives a privately observed cost shock in each period, and Athey and Bagwell (2008) study the same issue while allowing for each firm's cost shocks to follow a Markovian process. Hörner and Jamison (2007) consider a framework similar to Athey and Bagwell (2001) but without allowing for communication. Escobar and Toikka (2010) study general two-player Bayesian games in which the private information of each player also follows a Markovian process. Both show that first-best can be achieved asymptotically when players are sufficiently patient. Athey and Segal (2007) and Bergemann and Välimäki (2010) propose efficient mechanisms in a dynamic environment with private information.[1]

Another strand of the literature assumes costs are publicly observable and focuses on the issue of secret price cutting. Green and Porter (1984) show that in a repeated Cournot game with stochastic demand a cartel that observes the common price but not the sales of individual firms may use a price war as a collective punishment to deter the firms from secretly raising outputs. Harrington and Skrzypacz (2007) analyze a similar problem but in the context of repeated Bertrand competition. They show that a cartel that observes sales rather than prices can enforce a collusive agreement by requiring each colluding firm to pay other firms in the cartel a side-payment proportional to its output. These side-payments effectively serve as an output tax that discourages firms from secretly cutting price.[2] Harrington and Skrzypacz (2011) extends their 2007 model to the

[1]Other related works include Athey and Miller (2007) and Miller (Forthcomming), Aoyagi (2003), Skrzypacz and Hopenhayn (2004), and Blume and Heidhues (2006).

[2]Aoyagi (2002) analyzes collusion in a repeated Betrand game in which the demand shocks are positively correlated.

case where the cartel observes neither prices nor sales.[3]

While these two strands of literature has greatly deepened our understanding on how a cartel can overcome the private-cost-information problem and the secret-price-cutting problem separately, there has been little work on how these two problems can be tackled jointly. There is no economic reason why either costs, prices, or sales must be publicly observable. On the contrary, one would expect in many industries firms would have some private information about all three of them. Even firms that produce homogenous products may have significantly different costs due to technological and productivity differences. In a recent case study, Harrington (2006) notes that many cartels relied primarily on the sales and pricing information supplied by the firms. Given that collusion is illegal, it would be difficult for a cartel to devise a mechanism that would make it impossible for a firm to lie without arousing the suspicion of the antitrust authority.

In this paper we argue that a natural way to prevent a firm from lying about its costs *and* secretly cutting price is to punish it when the profits of the other firms are low. After all, the only reason a cartel would want to prevent these actions is because they hurt the profits of the other firms. Furthermore, if the original agreement is efficient, then the private gains a firm can obtain from any form of cheating must be outweighed by the total loss of the other firms; therefore, a firm would have no incentive to cheat if it is held responsible for the profits of the others. We formalize this idea in a model of repeated oligopoly and identify a set of conditions under which a cartel can enforce a collusive agreement without observing any of the prices, costs, and sales of the firms.

Our enforcement scheme is similar to the classic AGV mechanism in that it also use transfers to internalize the externalities of individual actions. However, there are several important differences between the two mechanisms. Whereas under the AGV mechanism an agent's transfer depends only on his reported type, under our scheme a firm's transfer depends both on its reported cost type and the

[3] More broadly, similar issues are also investigated in the literature of repeated games with communication under private monitoring. See, Aoyagi, 2002, Kandori and Matsushima, 1998, Compte, 1998, Fudenberg and Levine, 2007, Obara, 2009, and Zheng, 2008.

reported profits of the other firms. Furthermore, to limit the efficiency loss due to private monitoring, the transfer of each firm under our scheme is increasing in the profits of the other firms only up certain profit targets. Once a firm's profit is above its target, the transfers of the other firms are no longer tied to it. As a result, a firm which have secretly cut price may escape punishment when a positive demand shock masks the effect of the price cut. A key part of our argument is to show that, despite this truncation problem, firms could still be motivated to comply with the collusion agreement.

The greatest strength of our scheme is its versatility. It applies to both price and quantity competition and allows for a wide set of stochastic demand functions. The firms' products can be substitutes or complements, and there is no restriction on the correlations between the demand shocks of the firms. In his seminal work on collusion, Stigler (1964) points out there is no "single" collusive price. In many industries not only do firms produce many multiple products, they may also sell in different geographical markets or charge different customers different prices. In Section 5 we show that the scheme can be easily generalized to a setting where each firm is choosing a vector of prices. Furthermore, the main feature of our scheme—that each firm should be punished when the profits of the other firms in the cartel fall below certain targets—is broadly consistent with the practice of many real-life cartels. Firms in the citric acid cartel, for example, agreed to a set of sales quotas, and firms that sold above quota were required to purchase product from those that sold below.[4] Although our scheme may not fully maximize the cartel profit, we show that when the "size" of the demand shocks is small the cartel profit would be close to the monopoly profit so that any difference between the cartel profit under our scheme and that under the optimal scheme must be small.

Our model is related to Athey and Segal (2007) and Harrington and Skrzypacz (2011). Athey and Segal (2007) use a generalized AGV mechanism to support

[4]Harrington (2006) reports that some form of compensation schemes were used in the cartels in the market for choline chloride, citric acid, lysine, organic peroxides, sodium gluconate, sorbates, most vitamins, and zinc phosphate. Also see Harrington and Skrzypacz (2011), Levenstein and Suslow (2006).

efficient collusion when agents' private types are correlated over time. Their model allows an agent to choose a private action, but the action is allowed to affect only the agent's own payoff. As a result, their model does not apply to our oligopoly setting where the action of a firm affects the profits of the other firms. Harrington and Skrzypacz (2011) introduce a collusion scheme in a repeated Bertrand game where a cartel can observe the costs but not the sales and prices of the firms. Our model is more general in that it allows firms also to have private cost information. Another key difference between our enforcement scheme and that of Harrington and Skrzypacz (2011) lies in the way penalties are decided. Whereas the scheme of Harrington and Skrzypacz (2011) punishes a firm when it reports high sales, ours punishes it when other firms report low sales. Since under the scheme of Harrington and Skrzypacz (2011) firms are punished when they report high sales, they would have an incentive to under-report. They show that the firms can be induced to report their sales truthfully when the total industry demand is close to completely inelastic. Our scheme requires no such restriction and can be applied to industries that are subject to aggregate demand shocks.

2 Model

2.1 Demand

We consider an infinitely repeated oligopoly game. Let $\mathcal{N}$ denote a set of n firms, each with constant marginal cost. In each period $t = 1, ..., \infty$, each firm i chooses an "action" a_i from a compact interval $A_i \subset \Re_+$. Let $a \equiv (a_1, ..., a_n)$ denote an action profile and $A \equiv \Pi_{i=1}^n A_i$ the set of action profiles. The market "outcome" for each firm i is a function of a and ε, a vector of random shocks. Let a^t denote the action profile chosen in period t.[5] Firm i's "outcome" of the oligopoly game in period t is denoted by

$$y_i^t = y_i\left(a^t, \varepsilon^t\right) \in \Re_+,$$

[5]More generally, for any variable x we use x^t to denote the value of that variable in period t.

where y_i is the outcome function of firm i, and a^t and ε^t are the action profile and a real vector of random shocks, respectively, in period t. Our model encompasses both price and quantity competition. In the former case, a_i would be firm i's price and y_i its sales, and in the latter vice versa. We assume that the random shock ε^t is identically and independently distributed across t according to an atomless probability measure F on a compact support Ω. For each firm i, y_i is positive for all $a \in A$ and $\varepsilon \in \Omega$, and integrable in ε.[6] Let

$$\overline{y}_i(a) \equiv \int_{\varepsilon\in\Omega} y_i(a,\varepsilon)\, dF(\varepsilon)$$

denote firm i's expected outcome, and let

$$\widetilde{y}_i(a, y_i') \equiv \int_{\varepsilon\in\Omega} \min\left(y_i(a,\varepsilon), y_i'\right) dF(\varepsilon)$$

denote firm i's expected outcome truncated from above at y_i'. We assume that y_i is continuous in (a, ε), and, for any $a \in A$ and $j \in \mathcal{N}$, differentiable in each a_j except for a measure-zero set of ε. This implies that both $\overline{y}_i(a)$ and $\widetilde{y}_i(a, y_i')$ are differentiable in a_j, with[7]

$$\begin{aligned}
\frac{\partial \overline{y}_i(a)}{\partial a_j} &= \int_{\varepsilon\in\Omega} \frac{\partial y_i(a,\varepsilon)}{\partial a_j} dF(\varepsilon);\\
\frac{\partial \widetilde{y}_i(a)}{\partial a_j} &= \int_{\varepsilon\in\left\{\varepsilon'\in\Omega \mid y_i(a,\varepsilon')\le y_i'\right\}} \frac{\partial y_i(a,\varepsilon)}{\partial a_j} dF(\varepsilon).
\end{aligned}$$

In addition to the technical assumptions described above, we shall make three substantive assumptions about y_i.

Assumption 1. For any $i, j \in \mathcal{N}$, $i \neq j$, y_i is monotone in a_j, and whether y_i increases or decreases in a_j is independent of ε.

Assumption 2. For any $i, j \in \mathcal{N}$, $i \neq j$, $a_{-j} \in A_{-j}$, $a_j', a_j'' \in A_j$, and $y_i' \in \Re_+$,

$$\frac{\frac{\partial \widetilde{y}_i\left(a_j', a_{-j}, y_i'\right)}{\partial a_j}}{\frac{\partial \overline{y}_i\left(a_j', a_{-j}\right)}{\partial a_j}} \le \frac{\frac{\partial \widetilde{y}_i\left(a_j'', a_{-j}, y_i'\right)}{\partial a_j}}{\frac{\partial \overline{y}_i\left(a_j'', a_{-j}\right)}{\partial a_j}}, \tag{1}$$

[6] The assumption that y_i is positive means that a firm can dispose of its product freely in the case of quantity competition.

[7] See, e.g., Theorem 20.4 of Aliprantis and Burkinshaw (1990).

when $y_i\left(a'_j, a_{-j}, \varepsilon\right) \geq y_i\left(a''_j, a_{-j}, \varepsilon\right)$ for all ε and the denominators on both sides of the inequality are non-zero.

Assumption 3. For any $i, j \in \mathcal{N}$, $i \neq j$ there exists some finite $\eta > 0$ such that for all $a \in A$,

$$\left|\frac{\partial y_i\left(a, \varepsilon\right)}{\partial a_j}\right| \leq \eta \left|\frac{\partial \overline{y}_i\left(a\right)}{\partial a_j}\right| \quad \text{for almost all } \varepsilon \in \Omega.$$

Assumption 1 says that the products of any firms i and j are either substitutes for all ε or complements for all ε. By definition,

$$\frac{\partial \widetilde{y}_i\left(a'_j, a_{-j}, y'_i\right)}{\partial a_j} \leq \frac{\partial \overline{y}_i\left(a'_j, a_{-j}\right)}{\partial a_j}.$$

Because of the truncation output function $\widetilde{y}_i$ can only partially capture the marginal effect of a_j on $\overline{y}_i$, and the fraction captured varies with a_j. Assumption 2 says that the fraction captured is smaller when firm j is choosing an action that leads to a higher outcome for firm i.[8] Finally, Assumption 3 says the marginal effect of a_j on y_i is uniformly bounded across ε as a fraction of the average marginal effect. This rules out the pathological case where the marginal effect is concentrated on an arbitrarily small subset of ε.

Assumptions 1 to 3 are satisfied by a wide range of outcome functions. The following examples are commonly used in the industrial organization literature.

Example 1 (Linear Outcome)**.** The random shock is an $n \times n$ vector $\varepsilon \equiv (\varepsilon_{10}, ..., \varepsilon_{1n}, ..., \varepsilon_{n0}, ..., \varepsilon_{nn})$, and the outcome for each firm i is

$$y_i\left(a, \varepsilon\right) = \varepsilon_{i0} + \sum_{j=1}^{n} \varepsilon_{ij} a_j.$$

Example 2 (Log-Linear Outcome)**.** The random shock is a $2n$ vector $\varepsilon \equiv (\varepsilon_{10}, \varepsilon_{11}, ..., \varepsilon_{n0}, \varepsilon_{n1})$, and the outcome for each firm i is

$$y_i\left(a, \varepsilon\right) = \varepsilon_{i0} a_i^{\varepsilon_{i1}} \Pi_{j \neq i} a_j^{l_{ij}},$$

where l_{ij}, $i \neq j$, are constant.

[8] That is, the fraction captured is smaller if firm j chooses a higher price in the case of Bertrand competition and a lower quantity in the case of Cournot competition

Example 3 (Truncated Linear Outcome)**.** The random shock is an n vector $\varepsilon \equiv (\varepsilon_1, ..., \varepsilon_n)$, each ε_i has a compact support $[\underline{\varepsilon}_i, \overline{\varepsilon}_i]$, and the density of the marginal distribution of each ε_i is strictly positive, differentiable, and log-concave for any $\varepsilon_i \in [\underline{\varepsilon}_i, \overline{\varepsilon}_i]$. The outcome for each firm i is

$$y_i(a, \varepsilon) = \max\left(\varepsilon_i + \sum_{j=1}^{n} l_{ij} a_j, 0\right).$$

Example 4 (Logit Outcome)**.** The random shock ε has a compact support $[\underline{\varepsilon}, \overline{\varepsilon}]$, and the density is strictly positive, differentiable, and log-concave for any $\varepsilon \in [\underline{\varepsilon}, \overline{\varepsilon}]$. The outcome of each firm i is

$$y_i(a, \varepsilon) = \frac{\exp(-l_i a_i)}{\varepsilon + \sum_{j=1}^{n} \exp(-l_j a_j)}.$$

In Examples 1 and 2 the random shocks enter the demand function additively and multiplicatively. There is no restriction on the distribution of the random shocks, except that in Example 1, each ε_{ij} must be either always positive or always negative so that y_i is monotone in a_j, and Ω and A must be chosen so that y_i is always positive, and in Example 2 each ε_{i0} must be positive to ensure that y_i is positive. In Examples 3 and 4, firm i's outcome is non-linear in the random shocks. In these cases, we have to impose a stronger restriction on the distribution of the random shocks. Note that in Examples 1-3, there is no restriction on the correlation between the firms' outcomes. Thus, unlike Harrington and Skrzypacz (2011), our model applies to industries with aggregate demand shocks. It is clear that these examples satisfy Assumptions 1 and 3. We shall discuss why they also satisfy Assumption 2 in Section 4.

We now turn to the supply side of the market. We assume that each firm i's marginal cost, c_i, is subject to firm-specific shock in each period. Let c_i^t denote firm i's marginal cost in period t. We assume c_i^t is identically and independently distributed across i and t on a compact interval $C_i \subset \Re_+$ according to a distribution G_i. The cost shocks and the demand shocks are independent, both contemporaneously and over time. Write C for $\Pi_{i=1}^{n} C_i$. Let $c = (c_1, ..., c_n) \in C$ denote a cost profile and c_{-i} a cost profile minus c_i, and let G and G_{-i} denote the distribution of

c and c_{-i}, respectively.[9] Both F and $(G_1, \ldots, G_n)$ are common knowledge among firms. Firm i's profit is equal to

$$\pi_i (a, c_i, \varepsilon) \equiv \Phi_1 (c_i, a_i)\, y_i (a, \varepsilon) + \Phi_2 (c_i, a_i)$$

The functional forms of Φ_1 and Φ_2 depend on whether the firms compete in price or quantity. In the former case $\Phi_1 (c_i, a_i) \equiv (a_i - c_i)$ and $\Phi_2 (c_i, a) \equiv 0$; in the later $\Phi_1 (c_i, a_i) \equiv a_i$ and $\Phi_2 (c_i, a) \equiv -c_i a_i$.

A cost-action profile $\alpha : C \to A$ is a function that maps each cost profile c to an action profile $\alpha (c) = (\alpha_1 (c), ..., \alpha_n (c)) \in A$. Let $\mathcal{A}$ denote the set of all cost-action profiles. For any cost-action profile α, let

$$\overline{\pi}_i (\alpha) \equiv \int_c \int_\varepsilon \pi_i (\alpha (c), c_i, \varepsilon)\, dF (\varepsilon)\, dG (c)$$

and

$$\overline{\Pi} (\alpha) \equiv \sum_{i=1}^{n} \overline{\pi}_i (\alpha)$$

denote the expected profit of firm i and the expected total cartel profit when the cost-action profile α is chosen, respectively.

Let α^* denote a cost-action profile such that for each $c \in C$, $\alpha^* (c)$ maximizes

$$\sum_{i=1}^{n} \int_{\varepsilon \in \Omega} \pi_i (a, c_i, \varepsilon)\, dF (\varepsilon).$$

In general, for any particular $c \in C$, $\alpha^* (c)$ may not be unique or be in the interior of A. For example, it is possible that a particular high-cost firm i should be shut down to maximize the total cartel profit. If the firms are competing in quantity, then $\alpha_i^* (c)$ should be 0, which must be the lower bound of A_i. If the firms are competing in price, then $\alpha_i^* (c)$ could be any sufficiently large number so that $\overline{y}_i (a_i^*) = 0$.

Given our assumptions, the incomplete-information oligopoly game where each firm i chooses a_i independently after observing c_i has a pure-strategy Bayesian-Nash equilibrium.[10] We denote this equilibrium by the cost-action profile α^{NE}.

[9] In general, for any vector x, x_{-i} denote the vector minus the i-th element.

[10] See Fudenberg and Tirole (1991), Theorem 6.3.

Since the equilibrium action of firm i depends only on c_i, $\alpha_i^{NE}(c_i, c_{-i}) = \alpha_i^{NE}\left(c_i, c'_{-i}\right)$ for all $c_i \in C_i$ and $c_{-i}, c'_{-i} \in C_{-i}$. By definition, for any firm i and any $c \in C$,

$$\alpha_i^{NE}(c) \in \arg\max_{a_i \in A_i} \int_c \int_\varepsilon \pi_i\left(a_i, \alpha_{-i}(c), c_i, \varepsilon\right) dF(\varepsilon)\, dG_{-i}(c_{-i}).$$

2.2 Timing and Information

In each period t each firm i chooses a_i^t after observing c_i^t. Then each firm i observes its own outcome y_i^t after ε^t, which is not directly observable, is realized. We assume that c_i^t, a_i^t, and y_i^t are all private information. In every period the firms meet twice to exchange information, first after the marginal costs but before actions are chosen, and then again after outcomes are realized. Following Harrington and Skrzypacz (2011), we allow firms to exchange side-payments at the end of each period.

Specifically, the time-line of each period t is as follows: (1) each firm i first privately draws a cost parameter c_i^t from C_i; (2) it then sends a cost report $\widehat{c}_i^t \in C_i$ to the other firms; (3) after receiving the cost reports from the other firms, firm i chooses a_i^t; (4) the demand shock ε^t is realized; (5) firm i observes its outcome y_i^t and updates its belief about ε^t; (6) it sends an outcome report $\widehat{y}_i^t$ to the other firms; (7) after receiving the outcome reports from the other firms, it makes a side-payment $\tau_{ij}^t \geq 0$ to each firm j; these payments are publicly observed by all firms; (8) finally, it observes the outcome ω^t of a public randomization device that is uniformly distributed between 0 and 1.[11] See Figure 1.

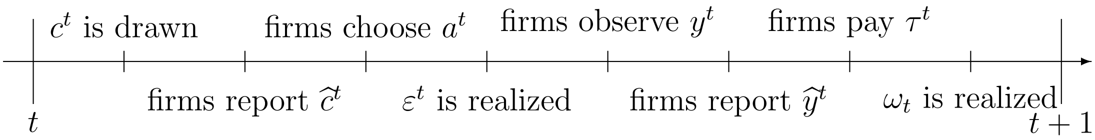

Figure 1: Timeline.

[11] The existence of a public randomization device, a standard assumption in the literature, allows firms to correlate their continuation strategies in future periods.

In period t each firm must decide what cost to report, what action to take, what outcome to report and, finally, how much to pay the other firms. Let $\widehat{c}^t = (\widehat{c}_1^t, ..., \widehat{c}_n^t)$ denote a profile of cost reports, and $\widehat{y}^t = (\widehat{y}_1^t, ..., \widehat{y}_n^t)$ a profile of outcome reports. A cost-reporting strategy for firm i is a function $\rho_i : C_i \to C_i$ that maps firm i's marginal cost c_i^t into a cost report $\widehat{c}_i^t \in C_i$. An action strategy is a function $\gamma_i : C \to A_i$ that determines firm i's period-t action, a_i^t, on the basis of c_i^t, firm i's own marginal cost and $\widehat{c}_{-i}^t$, the cost-report profile of all firms $j \neq i$. An outcome-reporting strategy is a function $r_i : C \times \Re_+ \to \Re_+$ that determines firm i's outcome report, $\widehat{y}_i^t$, on the basis of c_i^t, $\widehat{c}_{-i}^t$, and y_i^t. Finally, a transfer strategy is a function $b_i : C \times \Re_+^{n-1} \to \Re_+^n$ that maps c_i^t, $\widehat{c}_{-i}^t$, and $\widehat{y}_{-i}^t$ to a vector of side-payments $\tau_i^t = (\tau_{i1}^t, ..., \tau_{in}^t)$.[12] A reduced-normal-form stage-game strategy for firm i is a quadruple $(\rho_i, \gamma_i, r_i, b_i)$.[13]

A history of the game is an infinite sequence

$$\left\{\varepsilon^t, c^t, a^t, y^t, \widehat{c}^t, \widehat{y}^t, \tau^t, \omega^t\right\}_{t=1}^{\infty}.$$

At the beginning of period t (before observing the marginal cost of that period) each firm will have observed a public history h_{pub}^t that includes the cost reports, outcome reports, side-payments, outcomes of the randomization device in the first $t-1$ periods. In addition, firm i will have observed a private history h_i^t of its own costs, actions, and outcomes. Firm i's information at the beginning of period t includes both h_{pub}^t and h_i^t. A repeated-game strategy for firm i, denoted by σ_i, is a function that maps in each period t firm i's private information to a stage-game strategy. The firms discount future profits by a factor $\delta < 1$.The expected average discounted profit for firm i under strategy profile $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_n)$ is equal to

$$v_i(\sigma) = (1-\delta) \sum_{t=1}^{\infty} \delta^{t-1} E\left[\pi_i\left(a^t, c_i^t, \varepsilon^t\right) + \sum_{j \neq i}\left(\tau_{ji}^t - \tau_{ij}^t\right) \middle| \sigma\right],$$

where the expectation is taken over the distribution of histories induced jointly

[12]We assume firm i also pays itself a side-payment to simplify notations. Its value obviously has no significance.

[13]Allowing γ_i to depend also on $\widehat{c}_i$ will lead to more tedious notations without changing the result. The same is true for r_i and b_i.

by σ, F, and G. The solution concept we use is perfect public equilibrium.[14] A repeated-game strategy of a firm is a public strategy if the firm's stage-game strategy in any period depends only on the public history up to that period. A strategy profile is a perfect public equilibrium if the strategy of each firm is public and if the continuation strategy profile after any public history is a Nash equilibrium in the continuation game.

3 Result

The cartel ideally would like to implement the cost-action profile α^* in every period to maximize the total cartel profit. In general this would mean allowing firms with lower costs to produce more, and compelling each firm to choose an action that, given the actions of the other firms, does not maximize the firm's own profit. If actions were observable, it could implement α^* by applying an AGV mechanism in every period. If costs were observable, then the cartel problem would reduce to one of private monitoring. Harrington and Skrzypacz (2011) propose a mechanism to deal with this problem in a model of repeated price competition. But to apply their mechanism, a cartel must be able to calibrate the penalty for over-production so that each firm has the incentive to charge the exact right price. This becomes impossible when the cartel does not know the firms' marginal cost, as a firm's incentive to cut price depends on its cost. Because a firm can lie about its cost, deviate from its collusive action, and mis-report its outcome at the same time, a cartel cannot treat the private-cost-information problem and the private-monitoring-problem separately by first using an AGV mechanism to elicit the costs and then applying the Harrington-Skryzpacz mechanism to induce the right actions.

The contribution of this paper is to introduce a collusion enforcement scheme that allows for both private cost information and private monitoring. To introduce our main result, we need to introduce a measure of the "size" of the demand shock.

[14]See, e.g., Definition 5.3 of Fudenberg and Tirole (1991) for a formal description of a perfect public equilibrium.

Consider the following inequality:

$$F\left(\varepsilon \in \Omega \text{ s.t. } \left|\exists i \in \mathcal{N}, \left(y_i\left(a^*\left(c\right), \varepsilon\right) - \overline{y}_i\left(a^*\left(c\right)\right)\right) / \overline{y}_i\left(a^*\left(c\right)\right)\right| \geq \kappa\right) \leq \kappa. \qquad (2)$$

The left-hand side of (2) is the probability that some firm i's output is different from the mean by a factor of κ or more when the action profile $a^*(c)$ is chosen. When this probability is less than some small κ, the distribution of every firm i's output is concentrated around the mean. In the following we say that the size of the demand shocks is smaller than κ if (2) holds.

Proposition 1. Suppose the demand system satisfies Assumptions 1-3. Then given any $\zeta > 0$, there exists $\overline{\delta}$ and κ such that for any $\delta \geq \overline{\delta}$ there is a perfect public equilibrium σ with

$$\sum_{i=1}^{n} v_i(\sigma) \geq \overline{\Pi}\left(\alpha^*\right) - \zeta$$

if the size of the demand shocks is smaller than κ.

Proposition 1 says that there is a perfect public equilibrium in which the expected cartel profit is arbitrarily close to the monopoly profit if the firms are sufficiently patient, the demand shocks are sufficiently small, and the demand system satisfies Assumptions 1-3. The key step of the proof is to design incentives that implement α^*. In the following we provide a brief outline of the main issues. Details are left to the next section.

It is useful to focus on the incentives in a particular period. Consider a one-shot collusion game defined by steps 1-5 of the last section. But instead of making side-payments, each firm i directly receives a transfer w_i that is a function of $\widehat{c}$ and $\widehat{y}$. The total payoff for firm i in this game is

$$\pi_i\left(a, c_i, \varepsilon\right) + w_i\left(\widehat{c}, \widehat{y}\right).$$

Since there is no external source of funding in the original repeated game, the total transfer must be negative; that is,

$$\sum_{i=1}^{n} w_i\left(\widehat{c}, \widehat{q}\right) \leq 0, \quad \text{for all } \left(\widehat{c}, \widehat{y}\right) \in C \times \Re_+^n. \qquad (3)$$

As a cartel can destroy surplus through a costly price war, the total transfer can be strictly negative. But the discounted cartel profit will be low if the amount of surplus destroyed is large. The task, therefore, is to find a set of "not-too-costly" transfers such that it is a Nash equilibrium for each firm i to report c_i and y_i truthfully and chooses the action $\alpha_i^*(\widehat{c})$.

Our basic approach is to tie the transfer of a firm to the profits of the other firms. Recall that the profit of a firm depends solely on its own cost, action, and outcome.[15] Given the reports $\widehat{c}$ and $\widehat{y}$, we can calculate what the profit of firm j would be assuming that its cost and outcome are $\widehat{c}_i$ and $\widehat{y}_i$ and its action is $\alpha_i^*(\widehat{c})$. Let

$$\widetilde{\pi}_j(\widehat{y}_i, \widehat{c}) \equiv \Phi_1(\widehat{c}_i, \alpha_i^*(\widehat{c}))\,\widehat{y}_i + \Phi_2(\widehat{c}_i, \alpha_i^*(\widehat{c}))$$

denote the purported profit of firm j. We can implement α^* in the one-shot collusion game by setting for each firm i

$$w_i(\widehat{c}, \widehat{y}) = \sum_{j \neq i} \widetilde{\pi}_j(\widehat{y}_j, \widehat{c}) + L_i(\widehat{c}_{-i}),$$

where L_i is any function that depends only on the cost reports of firms $j \neq i$.[16] Since w_i is independent of $\widehat{y}_i$, no firm would gain by lying about its outcome. Furthermore, since α^* maximizes the expected cartel profit, any cheating, whether by lying about its cost, or by deviating from the equilibrium action, or both, must hurt the other firms more than it helps firm i. It follows that each firm i would have an incentive to report its cost truthfully and choose α_i^* when other firms are doing the same.

The problem with implementing α^* in this way is that the resulting efficiency loss is likely to be large. To conform with (3), we have to choose L_i, $i = 1, ..., n$, such that for any $c \in C$

$$\sum_{i=1}^{n} L_i(c_{-i}) \leq -(n-1) \sup_{\varepsilon \in \Omega} \sum_{i=1}^{n} \pi_i(\alpha_i^*(c), c_i, \varepsilon).$$

This means that the firms as a group will be punished whenever the realized cartel profit is lower than that under the most favorable demand shock. The total

[15] A firm's profit is affected by the actions of the other firms only indirectly through its outcome.

[16] Note that L_i cannot depend on $\widehat{y}_{-i}$ because a_i may affect $\widehat{y}_{-i}$ indirectly through y_{-i}.

expected transfer would be equal to

$$\sum_{i=1}^{n}\int_{c}\int_{\varepsilon} w_i\left(c, y\left(\alpha^*\left(c\right), \varepsilon\right)\right) dF\left(\varepsilon\right) dG\left(c\right)$$
$$\leq \left(n-1\right)\left(\overline{\Pi}\left(\alpha^*\right) - \int_{c}\sup_{\varepsilon\in\Omega}\sum_{i=1}^{n}\pi_i\left(\alpha_i^*\left(c\right), c_i, \varepsilon\right) dG\left(c\right)\right). \quad (4)$$

Note that, fixing C and Ω, the second bracketed term on the right-hand side of (4) would be bounded above from zero even when the distribution of ε converges to the mean, so long as the supremum of the total collusion profit is greater than the mean.[17]

To avoid this problem, we will tie the transfer of each firm to the *truncated* profit of the other firms. Specifically, for each i, $\widehat{c}$, and $\widehat{y}$, we set

$$w_i\left(\widehat{c}, \widehat{y}\right) = \sum_{j\neq i}\lambda_{ij}\left(\widehat{c}\right)\left(\min\left(K_j\left(\widehat{c}\right), \widetilde{\pi}_j\left(\widehat{y}_j, \widehat{c}\right)\right) - K_j\left(\widehat{c}\right)\right) + M_i\left(\widehat{c}\right),$$

where λ_{ij}, K_j, and M_i are all functions of $\widehat{c}$, and

$$\sum_{i=1}^{n} M_i\left(\widehat{c}\right) = 0 \quad \text{for all } \widehat{c}\in C.$$

Intuitively, under this set of transfers, each firm i is required to pay a fine only when the profit of firm $j\neq i$ is below a profit target $K_j\left(\widehat{c}\right)$. As the distribution of ε converges to the mean, the probability that a firm's profit is significantly below the mean goes to zero. We could, therefore, keep the total expected transfer close to zero by setting each $K_j\left(\widehat{c}\right)$ slightly above the mean profit of firm j.

Because the incentives are truncated, the transfers will no longer fully internalize the externalities of the firms' actions. But it turns out that it is not necessary for firm i's transfer to increase one to one with the profits of the other firms. For any cost profile $\widehat{c}$ and for any firm $j\neq i$, let $H_{ij}\left(\widehat{c}, a_i\right)$ denote the difference between the expected value of

$$\lambda_{ij}\left(\widehat{c}\right)\left(\min\left(K_j\left(\widehat{c}\right), \widetilde{\pi}_j\left(\alpha_j^*\left(\widehat{c}\right), y_j\left(a_i, \alpha_{-i}^*\left(\widehat{c}\right), \varepsilon\right), \widehat{c}_j\right)\right) - K_j\left(\widehat{c}\right)\right),$$

[17] Note that L_i could not be a function of $\widehat{y}_{-i}$ because doing so would allow firm i to influence L_i indirectly through a_i.

the part of w_i that ties to firm j's profit and the actual expected profit of firm j, as a function of firm i's action, assuming that firms other than i are choosing $\alpha^*_{-i}(\widehat{c})$. Since α^* maximize cartel profit, we only need to show that $\alpha_i(c)$ maximizes $H_{ij}(\widehat{c}, a_i)$ with respect to a_i for every $\widehat{c} \in C$ and every $j \neq i$. A key step of the proof of Proposition 1 is to show that, under Assumptions 1 and 2, H_{ij} is single-peaked in a_i. Hence, we can choose $\lambda_{ij}(\widehat{c})$ so that H_{ij} is maximized when $a_i = \alpha^*_i(\widehat{c})$ for every $\widehat{c} \in C$.

Assumption 1, which requires that each y_i is monotone in the action of each firm $j \neq i$, is natural in the oligopoly setting. To clarify the restrictions imposed by Assumption 2, we return to the examples in Section 2. In Examples 1 and 2, the marginal effect of a_j on y_i can be written as

$$\frac{\partial y_i}{\partial a_j} = \xi_{ij}(a)\,\zeta_{ij}(a_{-j}, \varepsilon), \tag{5}$$

where ξ_{ij} is independent of ε, and ζ_{ij} is positive and independent of a_j. Let

$$\Gamma(y'_i, a) \equiv \{\varepsilon \in \Omega | y_i(a, \varepsilon) \leq y'_i\}$$

denote the set of ε for which y_i is less than any cutoff y'_i. By factoring $\xi_{ij}(a)$ out of the integrals, we have

$$\frac{\frac{\partial \widetilde{y}_i\left(a_j, a_{-j}, y'_i\right)}{\partial a_j}}{\frac{\partial \overline{y}_i(a_j, a_{-j})}{\partial a_j}} = \frac{\int_{\varepsilon \in \Gamma\left(y'_i, a_j, a_{-j}\right)} \zeta_{ij}(a_{-j}, \varepsilon)\, dF(\varepsilon)}{\int_{\varepsilon \in \Omega} \zeta_{ij}(a_{-j}, \varepsilon)\, dF(\varepsilon)}. \tag{6}$$

Note that firm j's action a_j affects the right-hand side of (6) only through $\Gamma(y'_i, a_j, a_{-j})$. For any $a'_j, a''_j \in A_j$, we have $\Gamma\left(y'_i, a'_j, a_{-j}\right) \subset \Gamma\left(y'_i, a''_j, a_{-j}\right)$ if $y\left(a'_j, a_{-j}, \varepsilon\right) > y\left(a''_j, a_{-j}, \varepsilon\right)$ for all ε. Hence, the ratio on the left-hand side of (6) would be smaller under a'_j than a''_j. Intuitively, firm j is punished only firm i's outcome is less than the cutoff y'_i. Because firm i's output is more likely to be lower than y'_i when firm j is choosing an action that lowers firm i"s outcome, firm i's punishment captures a bigger fraction of of marginal effect of firm j's action when firm j's action reduces firm i's outcome.

Having demand shocks that enter the outcome function addictively or multiplicatively is not the only way Assumption 2 can be satisfied. In Examples 3 and

4, firm i's outcome is of the form

$$y_i(a,\varepsilon) = \nu_i(a_i)\, h_i(\chi_i(a) + \varepsilon_i),$$

with h_i increasing and χ_i monotone in each a_j. Since ε_i and $\chi_i(a)$ are not separable, we cannot factor out the term containing a_j as we do in Examples 1 and 2. However, in these examples the value of $\partial y_i/\partial a_j$ depend solely on the sum of $\chi_i(a)$ and ε_i. As χ_i becomes bigger due to firm j's action, the same value of $\partial y_i/\partial a_j$ will be associated with a smaller ε_i. We show in the appendix that when the density of ε_i is log-concave, such a change would reduce the marginal effect captured by the truncated profit function $\widetilde{y}_i$.

4 Proof of Proposition 1

4.1 Strategies

We will prove Proposition 1 by construction. In this section we first describe the trigger-strategy profile used in the construction. Our equilibrium trigger-strategy profile is characterized by two components: a probability function $\mu : \Re^n_+ \to [0,1]$, and an $n \times n$ compensation matrix β, where each component $\beta_{ij} : \Re^{2n}_+ \to [0,\infty)$. There are two states in trigger-strategy profile: collusive and non-collusive. Since the firms' continuation strategies depend only on the current state, we will drop the superscript t.

The equilibrium starts off in the collusive state in period 1. In the collusive state each firm i reports its cost c_i truthfully, chooses the action $\alpha_i^*(\widehat{c})$, and reports the outcome y_i truthfully. Let $\Lambda(c)$ denote the set of outcome profiles with positive densities when there is at most one firm i that does not choose $\alpha_i^*(c)$.[18] If $\widehat{y} \in \Lambda(\widehat{c})$, each firm pays $\beta_{ij}(\widehat{c},\widehat{y})$ to each firm j, and in this case the equilibrium will stay in the collusive state in the next period if $\omega > \mu(\widehat{c},\widehat{y})$ and switch to the non-collusive state if $\omega \leq \mu(\widehat{c},\widehat{y})$. Since ω is uniformly distributed between 0 and 1, the probability that the equilibrium will switch to the non-collusive state in

[18] Formally, $\Lambda(c) \equiv \left\{ y \in \Re^n_+ | \exists \varepsilon \in \Omega, i \in \mathcal{N}, a \in A \ s.t.\ y(a,\varepsilon) = y, a_{-i} = \alpha^*_{-i}(c) \right\}.$

the next period is $\mu(\widehat{c},\widehat{y})$. If $\widehat{y} \in \Lambda(\widehat{c})$, but some firms do not make the required payments, then the game will switch to the non-collusive state in the next period with probability one. Finally, if $\widehat{y} \notin \Lambda(\widehat{c})$, then the firms do not make any side-payments, and the game will switch to the non-collusive state in the next period with probability one. In the non-collusive state firm i sends the same cost report $\widetilde{c}_i$, chooses the action profile $\alpha_i^{NE}(\widehat{c})$, sends the same outcome report $\widetilde{y}_i$, and make no side-payments.[19] The non-collusive state is absorbing. Once the equilibrium enters the non-collusive state, it stays there forever.

In equilibrium firm i's ex ante average discounted profit in the non-collusive state is

$$v_i^N = \overline{\pi}_i\left(\alpha^{NE}\right),$$

while the ex ante average discounted profit in the collusive state is

$$v_i^* = (1-\delta)\left(\overline{\pi}_i(\alpha^*) + \sum_{j\neq i}\left(\overline{\beta}_{ji} - \overline{\beta}_{ij}\right)\right) + \delta\left(\overline{\mu}v_i^N + (1-\overline{\mu})\,v_i^*\right), \tag{7}$$

where

$$\overline{\beta}_{ij} \equiv \int_c\int_\varepsilon \beta_{ij}(c, y(\alpha^*(c),\varepsilon)\,dF(\varepsilon)\,dG(c),$$

and

$$\overline{\mu} \equiv \int_c\int_\varepsilon \mu(c, y(\alpha^*(c),\varepsilon))dF(\varepsilon)\,dG(c).$$

Let

$$w_i(\widehat{c},\widehat{y}) \equiv \sum_{j\neq i}(\beta_{ji}(\widehat{c},\widehat{y}) - \beta_{ij}(\widehat{c},\widehat{y})) + \delta(1-\delta)^{-1}\mu(\widehat{c},\widehat{y})(v_i^N - v_i^*) \tag{8}$$

denote the sum of the net side-payments received by firm i in the current period and the present value of the expected loss in future profit if the state switches to the non-collusive state in the next period. Henceforth, we shall call w_i firm i's "transfer" in the collusive state. Substituting (8) into (7) and rearranging terms, we have

$$v_i^* = \overline{\pi}_i(\alpha^*) + \int_c\int_\varepsilon w_i(c_i, y(\alpha^*(c),\varepsilon))\,dF(\varepsilon)\,dG(c).$$

Since the equilibrium starts in the collusive state, firm i's average equilibrium discounted profit is also v_i^*.

[19] Any $\widetilde{c}_i \in C_i$ and $\widetilde{y}_i \in \Re_+$ will suffice.

It is obvious that no firm can deviate profitably from the equilibrium strategies in the non-collusive state. Hence, we can focus on the firms' incentives in the collusive state. First, consider the decision on side-payments. Conditional on $(\widehat{c}, \widehat{y})$, firm i would receive a discounted profit equal to

$$w_i(\widehat{c}, \widehat{y}) + \delta (1-\delta)^{-1} v_i^*$$

in the continuation game if it makes the required side-payments at the end of the collusive state (and other firms follow the trigger-strategy profile). If it reneges, it would receive at most

$$\delta (1-\delta)^{-1} v_i^N + \sum_{j \neq i} \beta_{ji}(\widehat{c}, \widehat{y}).$$

Hence, it would be optimal for firm i to make the required side-payments at the end of the collusive state if

$$w_i(\widehat{c}, \widehat{y}) + \delta (1-\delta)^{-1} v_i^* \geq \delta (1-\delta)^{-1} v_i^N + \sum_{j \neq i} \beta_{ji}(\widehat{c}, \widehat{y}). \tag{9}$$

Now, turning to the reporting and action decisions. Let ρ_i^* denote the truth-telling cost-reporting strategy and r_i^* the truth-telling outcome-reporting strategy. For any ρ_i and γ_i, let α^{ρ_i,γ_i} denote the cost-action profile induced by firm i choosing (ρ_i, γ_i) and firms $j \neq i$ choosing $\left(\rho_{-i}^*, \alpha_{-i}^*\right)$. That is, for each $j \in \mathcal{N}$ and $c \in C$

$$\alpha_j^{\rho_i,\gamma_i}(c) \equiv \begin{cases} \gamma_i(c) & if \quad j = i \\ \alpha_j^*(\rho_i(c_i), c_{-j}) & if \quad j \neq i \end{cases}.$$

Suppose all firms $j \neq i$ are to follow the trigger-strategy profile. Firm i would receive an average discounted profit equal to

$$\begin{aligned} u_i(\rho_i, \gamma_i, r_i; w_i) \equiv\ & \overline{\pi}_i(\alpha^{\rho_i,\gamma_i}) + \delta (1-\delta)^{-1} v_i^* \\ & + \int_c \int_\varepsilon w_i(\rho_i(c_i), c_{-i}, r_i(c, y_i(\alpha^{\rho_i,\gamma_i}(c), \varepsilon)), y_{-i}(\alpha^{\rho_i,\gamma_i}(c), \varepsilon))\, dF(\varepsilon)\, dG(c) \end{aligned}$$

in the collusive state if it chooses cost-reporting strategy ρ_i, action strategy γ_i, outcome-reporting strategy r_i, sends side-payments β_i to the other firms, and follows the trigger-strategy profile from the next period onward. Let Σ_i denote the set of all cost-reporting, action, and outcome-reporting strategies. Assuming that firm i is to make the required side-payments after any $(\widehat{c}, \widehat{y}) \in C \times \Re^n_+$ and follow the trigger-strategy profile thereafter, it would be optimal for it to choose $(\rho^*_i, \alpha^*_i, r^*_i)$ in the collusive state if

$$(\rho^*_i, \alpha^*_i, r^*_i) \in \arg \max_{(\rho_i, \gamma_i, r_i) \in \Sigma_i} u_i (\rho_i, \gamma_i, r_i; w_i). \tag{10}$$

Lemma 1. *A trigger-strategy profile (μ, β) is a perfect public equilibrium if (10) holds and if (9) holds for each $(\widehat{c}, \widehat{y}) \in C \times \Re^n_+$.*

Lemma 1, which follows immediately from the standard one-step-deviation proof argument, provides a sufficient condition for a trigger-strategy profile (μ, β) to constitute a perfect public equilibrium. In the following we establish Proposition 1 in two steps. First, we introduce a set of "transfers" w^* that implements (ρ^*, α^*, r^*) in the collusive state in the sense of (10), and show that for any $\zeta > 0$ we can calibrate w^* so that the expected "cost" for implementing (ρ^*, α^*, r^*) is less than ζ when the size of the demand shocks is sufficiently small in the sense of (2). We then show that when the size of the demand shocks is sufficiently small we can generate the required w^* in the repeated game with a trigger-strategy profile (μ, β) such that (9) holds for each $\widehat{c} \in C$ and $\widehat{y} \in \Lambda (\widehat{c})$.

4.2 Stage-Game Mechanism

In this section we introduce a transfer scheme $w^* : C \times \Re^n_+ \to \Re^n$ that induces each firm i to choose $(\rho^*_i, \alpha^*_i, r^*_i)$ in the collusive state. Recall by Assumption 3 there exists η such that for all a'_i

$$\left| \frac{\partial y_j \left(a'_i, \alpha^*_{-i} (c), \varepsilon\right)}{\partial a_i} \right| \leq \eta \frac{\partial \overline{y}_j \left(a'_i, \alpha^*_{-i} (c)\right)}{\partial a_i} \quad \text{almost all } \varepsilon.$$

Throughout this section we assume that κ is sufficiently small that $\kappa\eta < 1$. The first step to define w^* is to define a vector of profit-target functions $K = (K_1, ..., K_n)$ and a set of "scaling factors" $\{\lambda_{ij}\}_{i,j \in \mathcal{N}, i \neq j}$.

For each firm j and each $c \in C$ we set firm j's profit target at a level that is greater than the expected profit of firm j when the firms choose $\alpha^*(c)$ by $\Phi_1\left(c_j, \alpha_j^*(c)\right)\kappa$; that is,

$$K_j(c) \equiv \Phi_1\left(c_j, \alpha_j^*(c)\right)\overline{y}_j\left(\alpha^*(c)\right)(1+\kappa) + \Phi_2\left(c_j, \alpha_j^*(c)\right).$$

When firm j chooses $\alpha_j^*(c)$, its profit will exceed $K_j(c)$ if its outcome is greater than

$$y_j'(c) \equiv \frac{K_j(c) - \Phi_2\left(c_j, \alpha_j^*(c)\right)}{\Phi_1\left(c_j, \alpha_j^*(c)\right)} = \overline{y}_j\left(\alpha^*(c)\right)(1+\kappa). \tag{11}$$

To define the scaling factor λ_{ij}, for each $c \in C$, let

$$\begin{aligned} B_{ij}^1(c) &\equiv \left\{a_i' \leq \alpha_i^*(c) \,|\, \partial\overline{y}_j\left(a_i', \alpha_{-i}^*(c)\right)/\partial a_i \neq 0\right\}; \\ B_{ij}^2(c) &\equiv \left\{a_i' \geq \alpha_i^*(c) \,|\, \partial\overline{y}_j\left(a_i', \alpha_{-i}^*(c)\right)/\partial a_i \neq 0\right\}. \end{aligned}$$

For $\overline{y}_j$ that increases in a_i, we define

$$\lambda_{ij}(c) \equiv \begin{cases} \sup_{a_i' \in B_{ij}^1(c)} \dfrac{\frac{\partial\overline{y}_j\left(a_i', \alpha_{-i}^*(c)\right)}{\partial a_i}}{\frac{\partial\widetilde{y}_j\left(a_i', \alpha_{-i}^*(c), y_j'(c)\right)}{\partial a_i}} & \text{if } B_{ij}^1(c) \text{ is non-empty;} \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

for $\overline{y}_j$ that decreases in a_i, we define

$$\lambda_{ij}(c) \equiv \begin{cases} \sup_{a_i' \in B_{ij}^2(c)} \dfrac{\frac{\partial\overline{y}_j\left(a_i', \alpha_{-i}^*(c)\right)}{\partial a_i}}{\frac{\partial\widetilde{y}_j\left(a_i', \alpha_{-i}^*(c), y_i'(c)\right)}{\partial a_i}} & \text{if } B_{ij}^2(c) \text{ is non-empty;} \\ 0 & \text{otherwise.} \end{cases} \tag{13}$$

Note that when $\partial\overline{y}_j\left(\alpha^*(c)\right)/\partial a_i$ is non-zero,

$$\lambda_{ij}(c) = \frac{\frac{\partial\overline{y}_j\left(\alpha^*(c)\right)}{\partial a_i}}{\frac{\partial\widetilde{y}_j\left(\alpha^*(c), y_i'(c)\right)}{\partial a_i}}.$$

In this case, $\lambda_{ij}(c)$ is simply the inverse of the fraction of the marginal effect of a_i on $\overline{y}_j$ captured by $\widetilde{y}_j$. We use the more general definition to allow for the possibility that $\partial\overline{y}_j\left(\alpha^*(c)\right)/\partial a_i = 0$. Under Bertrand competition the cartel may want an inefficient firm to sell zero unit by charging a very high price. In that case, a small change in the price of the inactive firm would not affect the sales of the active firms.

To see that $\lambda_{ij}(c)$ is well-defined, note that

$$\begin{aligned}
& \frac{\partial \widetilde{y}_j\left(a_i', \alpha_{-i}^*(c), y_i'(c)\right)}{\partial a_i} \\
= \ & \frac{\partial \overline{y}_j\left(a_i', \alpha_{-i}^*(c)\right)}{\partial a_i} - \int_{\varepsilon \notin \Gamma\left(y_j'(c), a_i', \alpha_{-i}^*(c)\right)} \frac{\partial y_j\left(a_i', \alpha_{-i}^*(c), \varepsilon\right)}{\partial a_i} dF(\varepsilon) \\
\geq \ & \frac{\partial \overline{y}_j\left(a_i', \alpha_{-i}^*(c)\right)}{\partial a_i} - \Pr\left(\varepsilon \notin \Gamma\left(y_j'(c), a_i', \alpha_{-i}^*(c)\right)\right)\eta \frac{\partial \overline{y}_j\left(a_i', \alpha_{-i}^*(c)\right)}{\partial a_i} \\
\geq \ & \left(1 - \Pr\left(\varepsilon \notin \Gamma\left(y_j'(c), \alpha^*(c)\right)\right)\eta\right) \frac{\partial \overline{y}_j\left(a_i', \alpha_{-i}^*(c)\right)}{\partial a_i} \\
\geq \ & (1 - \kappa\eta) \frac{\partial \overline{y}_j\left(a_i', \alpha_{-i}^*(c)\right)}{\partial a_i}.
\end{aligned} \tag{14}$$

Thus, $\partial \widetilde{y}_j\left(a_i', \alpha_{-i}^*(c), y_i'(c)\right)/\partial a_i > 0$ whenever $\partial \overline{y}_j\left(a_i', \alpha_{-i}^*(c)\right) > 0$. Intuitively, because firm j's profit target is set strictly above the mean profit and the marginal effect of a_i is uniformly bounded over ε, the truncated outcome function $\widetilde{y}_j$ must capture part of marginal effect of a_i when the distribution y_j is concentrated around the mean.

Given K and $\{\lambda_{ij}\}_{i,j \in \mathcal{N}, i \neq j}$, we define w^* as follows. For any $\widehat{c} \in C$ and $\widehat{y} \in \Lambda(\widehat{c})$, firm i's transfer after $(\widehat{c}, \widehat{y})$ is defined to be

$$\widehat{w}_i(\widehat{c}, \widehat{y}) \equiv \widehat{w}_i^1(\widehat{c}, \widehat{y}_{-i}) + \widehat{w}_i^2(\widehat{c}_i) - \frac{1}{n-1} \sum_{j \neq i} \widehat{w}_j^2(\widehat{c}_j), \tag{15}$$

where

$$\begin{aligned}
\widehat{w}_i^1(\widehat{c}, \widehat{y}_{-i}) & \equiv \sum_{j \neq i} \lambda_{ij}(\widehat{c}) \min\left(\widetilde{\pi}_j(\widehat{y}_j, \widehat{c}) - K_j(\widehat{c}), 0\right), \\
\widehat{w}_i^2(\widehat{c}_i) & \equiv \int_{c_{-i} \in C_{-i}} \int_{\varepsilon \in \Omega} \widehat{w}_i^1\left(\widehat{c}_i, c_{-i}, y_{-i}\left(\alpha^*(\widehat{c}_i, c_{-i}), \varepsilon\right)\right) dF(\varepsilon)\, dG_{-i}(c_{-i}) \\
& \quad - \sum_{j \neq i} \int_{c_{-i} \in C_{-i}} \int_{\varepsilon \in \Omega} \pi_j\left(\alpha^*(\widehat{c}_i, c_{-i}), c_j, \varepsilon\right) dF(\varepsilon)\, dG_{-i}(c_{-i}).
\end{aligned}$$

For any $(\widehat{c}, \widehat{y}) \in C \times \Re_+^n$, define firm i's transfer after $(\widehat{c}, \widehat{y})$ as

$$w_i^*(\widehat{c}, \widehat{y}) = \begin{cases} \widehat{w}_i(\widehat{c}, \widehat{y}) & if \quad \widehat{y} \in \Lambda(\widehat{c}) \\ \min_{\widehat{c} \in C, \widehat{y} \in \Lambda(\widehat{c})} \widehat{w}_i(\widehat{c}, \widehat{y}) & if \quad \widehat{y} \notin \Lambda(\widehat{c}) \end{cases}$$

To summarize, under the scheme w_i^* firm i will receive a transfer equal to $\widehat{w}_i(\widehat{c}, \widehat{y})$ if $\widehat{y}$ is consistent with at least $n-1$ firms choosing their actions according to $\alpha^*(\widehat{c})$; otherwise, it will receive a transfer lower than $\widehat{w}_i(c', y')$ for any $c' \in C$ and $y' \in \Lambda(c')$. The first component of $\widehat{w}_i$, $\widehat{w}_i^1$, punishes firm i whenever the reported profit of each firm $j \neq i$ is below $K_j(\widehat{c})$, and the second component, $\widehat{w}_i^2$ compensates firm i for the difference between the expected profits of the other firms and the expected value of $\widehat{w}_i^1$. The last component of firm i's transfer depends solely on the cost reports of firms $j \neq i$. Since each firm j reports $\widehat{c}_j$ before it $j \neq i$ observes $\widehat{c}_i$ and y_j, firm i cannot influence the value of $\widehat{w}_j^2$ through $\widehat{c}_i$ and a_i. This component is added so that it and $\widehat{w}_i^2$ sum to zero across firms. Hence, for all $\widehat{c} \in C$ and $\widehat{y} \in \Lambda(\widehat{c})$,

$$\sum_{i=1}^{n} \widehat{w}_i(\widehat{c}, \widehat{y}) = \left(\sum_{j \neq i} \lambda_{ji}(\widehat{c})\right) \min\left(\widetilde{\pi}_i(\widehat{y}_i, \widehat{c}) - K_i(\widehat{c}), 0\right) \leq 0.$$

The total expected transfer when each firm j choose $\left(\rho_j^*, \alpha_j^*, r_j^*\right)$ is, therefore, equal to

$$\overline{w}_i^* \equiv \sum_{j \neq i} \int_c \int_\varepsilon \lambda_{ij}(c) \min\left(\pi_j(\alpha^*(c), c_j, \varepsilon) - K_j(c), 0\right) dF(\varepsilon)\, dG(c).$$

We show below that given w^* it is a Nash equilibrium for each firm i to choose $(\rho_i^*, \alpha_i^*, r_i^*)$ and that $\overline{w}_i^*$ becomes arbitrarily small when κ converges to zero.

Lemma 2. *If Assumptions 1-3 hold, then, for each firm $i \in \mathcal{N}$, $(\rho_i^*, \alpha_i^*, r_i^*)$ maximizes $u_i(\rho_i, \gamma_i, r_i; w^*)$ with respect to all $(\rho_i, \gamma_i, r_i) \in \Sigma_i$.*

Proof of 2 . Note that for any $c \in C$ and $\varepsilon \in \Omega$, and for any cost-reporting strategy ρ_i and action strategy γ_i for firm i

$$y\left(\gamma_i(c_i), \alpha_{-i}^*(\rho_i(c_i), c_{-i}), \varepsilon\right) \in \Lambda(\gamma_i(c_i), c_{-i}).$$

Hence, if each firm $j \neq i$ chooses $\left(\rho_j^*, \gamma_j^*, r_j^*\right)$, then no matter what cost-reporting and action strategies firm i uses, the outcome report profile following any any cost-report profile $\widehat{c}$ will always belong to $\Lambda(\widehat{c})$ so long as firm i reports its outcome truthfully. Suppose firm i mis-reports its outcome after some cost report profile $\widehat{c}$.

If the resulting outcome report profile still belongs to $\Lambda(\widehat{c})$, then firm i's transfer will be the same as what it would have obtained if it didn't mis-report because $\widehat{w}_i$ does not depend on $\widehat{y}_i$. However, if the resulting outcome no longer belongs to $\Lambda(\widehat{c})$, then firm i will receive a transfer equal to $\min_{c\in C, y\in\Lambda(c)} w_i^*(c,y)$, which is lower than what it would have received if it did not mis-report. Hence, when each firm $j \neq i$ chooses $\left(\rho_j^*, \gamma_j^*, r_j^*\right)$, it is a best response for firm i to report its outcome truthfully.

It is therefore sufficient to prove Lemma 2 if we prove that (ρ_i^*, γ_i^*) maximizes $u_i(\rho_i, \gamma_i, r_i^*; w^*)$ with respect to all (ρ_i, γ_i). Define, for any firms i, $j \neq i$ and for any $c \in C$,

$$\begin{aligned} H_{ij}(c, a_i) \equiv\ & \lambda_{ij}(c) \int_\varepsilon \min\left(\pi_j\left(a_i, \alpha_{-i}^*(c), c_j, \varepsilon\right) - K_j(c), 0\right) dF(\varepsilon) \\ & - \int_\varepsilon \pi_j\left(a_i, \alpha_{-i}^*(c), c_j, \varepsilon\right) dF(\varepsilon). \end{aligned} \tag{16}$$

Substituting (15) and (16) into $u_i(\rho_i, \gamma_i, r_i^*; w^*)$, we have

$$\begin{aligned} u_i(\rho_i, \gamma_i, r_i^*; w^*) = & \sum_{j=1}^n \overline{\pi}_j\left(\alpha^{\rho_i, \gamma_i}\right) + \sum_{j\neq i} \int_c H_{ij}\left(\rho_i(c_i), c_{-i}, \gamma_i(c)\right) dG(c) \\ + \int_{c_i} \widehat{w}_i^2\left(\rho_i(c_i)\right) dG_i(c_i) & + \delta(1-\delta)^{-1} v_i^* - \frac{1}{n-1} \sum_{j\neq i} \int_{c_j} \widehat{w}_j^2(c_j)\, dG_j(c_j). \end{aligned} \tag{17}$$

The first term of (17) is the total expected cartel profit when firm i chooses (ρ_i, γ_i) and each firm $j \neq i$ chooses $\left(\rho_j^*, \alpha_j^*\right)$. Since α^* maximizes the total expected cartel profit, this term is maximized when firm i chooses (ρ_i^*, α_i^*).

Turning to the second term on the right-hand side of (17). For each $c \in C$ and $a_i \in A_i$, $H_{ij}(c, a_i)$ is differentiable in a_i, with

$$\frac{\partial H_{ij}}{\partial a_i} = \Phi_1\left(c_j, \alpha_j^*(c)\right)\left(\lambda_{ij}(c)\frac{\partial \widetilde{y}_j}{\partial a_i} - \frac{\partial \overline{y}_j}{\partial a_i}\right).$$

Note that by the definition of λ ((12) and (13)) and by Assumption 2, for any a_j' such that $\partial \overline{y}_j\left(a_i', \alpha_{-i}^*(c)\right)/\partial a_i \neq 0$,

$$\lambda_{ij}(c) \geq \frac{\frac{\partial \overline{y}_j\left(a_i', \alpha_{-i}^*(c)\right)}{\partial a_i}}{\frac{\partial \widetilde{y}_j\left(a_i', \alpha_{-i}^*(c), y_j'(c)\right)}{\partial a_i}} \quad \text{if and only if} \quad \overline{y}_j\left(\alpha^*(c)\right) \geq \overline{y}_j\left(a_i', \alpha_{-i}^*(c)\right). \tag{18}$$

Since y_j is monotone in a_i, $\partial \widetilde{y}_j / \partial a_i = 0$ if $\partial \overline{y}_j / \partial a_i = 0$. This, together with (18), implies that

$$\frac{\partial H_{ij}(c, a_i)}{\partial a_i} \geq 0 \quad \text{if and only if} \quad a_i \leq \alpha_i^*(c). \tag{19}$$

Since H_{ij} is continuous in a_i, $\alpha_i^*(c) \in \arg\max_{a_i \in A_i} H_{ij}(c, a_i)$ for any $c \in C$. Note that by definition, for any $\widehat{c}_i \in C_i$ and any $c_{-i} \in C_{-i}$

$$\int_{c_i \in C_i} \widehat{w}_i^2(\widehat{c}_i)\, dG_i(c_i) = -\sum_{j \neq i} \int_c H_{ij}(\widehat{c}_i, c_{-i}, \alpha_i^*(\widehat{c}_i, c_{-i}))\, dG(c).$$

Thus, for any ρ_i and γ_i,

$$\begin{aligned} \int_{c_i \in C_i} \widehat{w}_i^2(\rho(c_i))\, dG_i(c_i) &= -\sum_{j \neq i} \int_c H_{ij}(\rho(c_i), c_{-i}, \alpha_i^*(\rho(c_i), c_{-i}))\, dG(c), \\ &\leq -\sum_{j \neq i} \int_c H_{ij}(\rho_i(c_i), c_{-i}, \gamma_i(c_i, c_{-i}),)\, dG(c). \end{aligned}$$

Thus, the sum of the second and third terms on the right-hand side of (17) is maximized by any strategy (ρ_i, γ_i), including (ρ_i^*, α_i^*) in particular, that satisfies the condition that $\gamma_i(c_i, c_{-i}) = \alpha_i^*(\rho_i(c_i), c_{-i})$. Since (ρ_i^*, α_i^*) also maximizes the first term of (17), $u_i(\rho_i, \gamma_i, r_i^*; w^*)$ is maximized when $(\rho_i, \gamma_i) = (\rho_i^*, \gamma_i^*)$. □

Lemma 3. *Given any $\zeta > 0$, there exists $\kappa > 0$ such that*

$$\sum_{i=1}^{n} \overline{w}_i^* \geq -\zeta$$

if the size of the demand shocks is smaller than κ.

Proof of 3. From (14) we have for all $c \in C$

$$\lambda_{ij}(c) \leq \frac{1}{(1 - \kappa\eta)}.$$

By the definition of $K_j(c)$ we have

$$\pi_j(\alpha^*(c), c_j, \varepsilon) < K_j(c) - 2\kappa\Phi_1\left(c_j, \alpha_j^*(c)\right) \bar{y}_j(\alpha^*(c))$$

if and only if

$$y_j(\alpha^*(c), \varepsilon) < \bar{y}_j(\alpha^*(c))(1 - \kappa),$$

which by assumption occurs with probability less than or equal to κ.

Let

$$\begin{aligned} L_1 &= \max_{j,c,\varepsilon} \Phi_1 \left(c_j, \alpha_j^* (c)\right) y_j (\alpha^* (c), \varepsilon), \\ L_2 &= \min_{j,c,\varepsilon} (\pi_j (\alpha^* (c), c_j, \varepsilon) - K_j (c)), \end{aligned}$$

and

$$\begin{aligned} E_1 &= \left\{\varepsilon | \pi_j (\alpha^* (c), c_j, \varepsilon) \geq K_j (c) - 2\kappa \Phi_1 \left(c_j, \alpha_j^* (c)\right) \bar{y}_j (\alpha^* (c))\right\}, \\ E_2 &= \left\{\varepsilon \in \Omega | \pi_j (\alpha^* (c), c_j, \varepsilon) < K_j (c) - 2\kappa \Phi_1 \left(c_j, \alpha_j^* (c)\right) \bar{y}_j (\alpha^* (c))\right\}. \end{aligned}$$

Then, we have for each firm j and each $c \in C$,

$$\begin{aligned} & \int_\varepsilon \min (\pi_j (\alpha^* (c), c_j, \varepsilon) - K_j (c), 0)\, dF (\varepsilon) \\ \geq\ & \int_{\varepsilon \in E_1} -2\kappa L_1 dF (\varepsilon) + \int_{\varepsilon \in E_2} L_2 dF (\varepsilon) \\ \geq\ & -2\kappa L_1 + \Pr (y_j (\alpha^* (c), \varepsilon) < \bar{y}_j (\alpha^* (c)) (1 - \kappa))\, L_2 \\ \geq\ & -2\kappa L_1 + \kappa L_2. \end{aligned}$$

Hence

$$\begin{aligned} \sum_{i=1}^n \overline{w}_i^* &= \sum_{i=1}^n \sum_{j \neq i} \int_c \int_\varepsilon \lambda_{ij} (c) \min (\pi_j (\alpha^* (c), c_j, \varepsilon) - K_j (c), 0)\, dF (\varepsilon)\, dG (c) \\ &\geq \frac{n(n-1)}{1 - \eta\kappa} (-2\kappa L_1 + \kappa L_2), \end{aligned}$$

which tends to zero as κ tends to zero. □

4.3 Proof of Proposition 1

We are now in a position to prove Proposition 1. Note that the proposition is obviously true when $\sum_{i=1}^n \overline{\pi}_i \left(\alpha^{NE}\right) = \sum_{i=1}^n \overline{\pi}_i (\alpha^*)$. Suppose $\sum_{i=1}^n \overline{\pi}_i \left(\alpha^{NE}\right) < \sum_{i=1}^n \overline{\pi}_i (\alpha^*)$. Fix $\zeta > 0$. By Lemma 3, we can pick κ so that when the demand shock is smaller than κ,

$$\sum_i \overline{w}_i^* > \max \left(-\zeta, \sum_{i=1}^n \left(v_i^N - \overline{\pi}_i (\alpha^*)\right)\right). \tag{20}$$

Pick a vector $d = (d_1, ..., d_n)$, $\sum_{i=1}^{n} d_i = 0$, such that for each firm i

$$\overline{\pi}_i(\alpha^*) + \overline{w}_i^* + d_i > v_i^N.$$

Pick $\delta^* \in (0,1)$ such that for each firm i, and for any $\widehat{c} \in C$ and $\widehat{y} \in \Lambda(\widehat{c})$

$$\widehat{w}_i(\widehat{c}, \widehat{y}) + d_i \geq \delta^*(1-\delta^*)^{-1}\left(v_i^N - (\overline{\pi}_i(\alpha^*) + \overline{w}_i^* + d_i)\right). \tag{21}$$

Given any $\delta \geq \delta^*$, we define $\mu(\widehat{c}, \widehat{y})$ and $\beta(\widehat{c}, \widehat{y})$ for any $\widehat{c} \in C$, $\widehat{y} \in \Lambda(\widehat{c})$ as follows. First, set

$$\mu(\widehat{c}, \widehat{y}) \equiv \frac{\sum_{i=1}^{n}\left((1-\delta)\,\delta^{-1}\sum_{j\neq i}\lambda_{ij}(\widehat{c})\max\left(K_j(\widehat{c}) - \widetilde{\pi}_j(\widehat{y}_j, \widehat{c}), 0\right)\right)}{\sum_{i=1}^{n}\left(v_i^N - \overline{\pi}_i(\alpha^*) - \overline{w}_i^*\right)}. \tag{22}$$

By (20), $\mu(\widehat{c}, \widehat{y}) \in [0,1]$ for all $\widehat{c} \in C$ and $\widehat{y} \in \Lambda(\widehat{c})$. Then, set the transfer payment from firm i to firm j to be equal to

$$\beta_{ij}(\widehat{c}, \widehat{y}) \equiv \begin{cases} \beta_i^{net}(\widehat{c}, \widehat{y})\,\frac{\min\left(\beta_j^{net}(\widehat{c},\widehat{y}),0\right)}{\sum_k \min\left(\beta_k^{net}(\widehat{c},\widehat{y}),0\right)} & \text{if } \beta_i^{net}(\widehat{c}, \widehat{y}) > 0, \\ 0 & \text{otherwise,} \end{cases} \tag{23}$$

where

$$\beta_i^{net}(\widehat{c}, \widehat{y}) \equiv \widehat{w}_i(\widehat{c}, \widehat{y}) + d_i - \delta(1-\delta)^{-1}\mu(\widehat{c}, \widehat{y})\left(v_i^* - v_i^N\right). \tag{24}$$

Recall that we have defined $\mu(\widehat{c}, \widehat{y}) = 1$ and $\beta_{ij}(\widehat{c}, \widehat{y}) = 0$ for any $\widehat{c} \in C$ and $\widehat{y} \notin \Lambda(\widehat{c})$. It is straightforward to check that for all $\widehat{c} \in C$ the continuation profit of firm i as defined in (8) is equal to

$$w_i(\widehat{c}, \widehat{y}) = \begin{cases} \widehat{w}_i(\widehat{c}, \widehat{y}) + d_i & if \quad \widehat{y} \in \Lambda(\widehat{c}), \\ v_i^N & if \quad \widehat{y} \notin \Lambda(\widehat{c}). \end{cases}$$

Intuitively, for each $\widehat{c} \in C$, $\widehat{y} \in \Lambda(\widehat{c})$, we choose $\mu(\widehat{c}, \widehat{y})$ so that the expected value to be destroyed from non-collusion is equal to the total transfer, and then choose $\beta(\widehat{c}, \widehat{y})$ so that the individual transfer of each firm i is equal to $\widehat{w}_i(\widehat{c}, \widehat{y}) + d_i$.

Note that for $\widehat{c} \in C$, $\widehat{y} \in \Lambda(\widehat{c})$, $w_i(\widehat{c}, \widehat{y})$ is only different from $w_i^*(\widehat{c}, \widehat{y})$ by a constant, and like w_i^*, w_i satisfies the condition that for any $\widehat{c} \in C$, $\widehat{y} \in \Lambda(\widehat{c})$, and $\widehat{y}' \notin \Lambda(\widehat{c})$

$$w_i(\widehat{c}, \widehat{y}) \geq w_i(\widehat{c}, \widehat{y}').$$

Since w^* implements (ρ^*, α^*, r^*) by Lemma 2, so does w, meaning that (10) of Lemma 1 is satisfied. And (21) implies (9) of Lemma 1 is also satisfied. Hence, for any $\delta \geq \delta^*$, the trigger-strategy profile characterized by (μ, β) is a perfect public equilibrium.

Finally, by construction

$$\int_c \int_\varepsilon w_i\left(c_i, y\left(\alpha^*\left(c\right), \varepsilon\right)\right) dF\left(\varepsilon\right) dG\left(c\right) = \overline{w}_i^*.$$

Hence, firm i's discounted payoff in the collusive state is

$$v_i^* = \overline{\pi}_i\left(\alpha^*\right) + \overline{w}_i^* + d_i,$$

As $\sum_{i=1}^n \overline{w}_i^* > -\zeta$ (Lemma 3), the total discounted cartel payoff is equal to

$$\sum_{i=1}^n v_i^* \geq \overline{\Pi}\left(\alpha^*\right) - \zeta.$$

5 Extensions

5.1 Multiple-market Collusion

In this section we briefly describe how our enforcement scheme can be generalized to enforce a multi-market collusive agreement. Suppose that there are m markets, denoted by $l = 1, 2, \ldots, m$. Let $\mathcal{M}$ denote the set of markets. The stage-game action, outcome, and profit of each firm i are then represented by vectors

$$\begin{aligned} a_i &= (a_{i,1}, \ldots, a_{i,m}), \\ y_i &= (y_{i,1}, \ldots, y_{i,m}), \\ c_i &= (c_{i,1}, \ldots, c_{i,m}), \\ \pi_i &= (\pi_{i,1}, \ldots, \pi_{i,m}). \end{aligned}$$

Since the demands in different markets may be correlated, $y_{i,l}$, firm i's outcome in market l, could depend on the whole vector $a = (a_1, ..., a_n)$.

Recall that in the baseline model, given $K_j(c)$ and the corresponding $y'_j(c)$ (defined by (11)), the scaling factor $\lambda_{ij}(c)$ is characterized by the equation

$$\lambda_{ij}(c) = \frac{\frac{\partial \overline{y}_j(\alpha^*(c))}{\partial a_i}}{\frac{\partial \widetilde{y}_j\left(\alpha^*(c), y'_j(c)\right)}{\partial a_i}}$$

when $\partial \overline{y}_j(\alpha^*(c))/\partial a_i \neq 0$.

To extend our scheme to the multi-market case, we need to set a profit target, $K_{j,l}$, and a set of scaling factors, $\{\lambda_{ij,l}\}_{i \neq j}$, for each firm j and for each market l. But since each firm is producing more than one goods, it is possible that for some firm i and some markets l and k,

$$\frac{\frac{\partial \overline{y}_{j,l}(\alpha^*(c))}{\partial a_{i,l}}}{\frac{\partial \widetilde{y}_{j,l}\left(\alpha^*(c), y'_{j,l}(c)\right)}{\partial a_{i,l}}} \neq \frac{\frac{\partial \overline{y}_{j,l}(\alpha^*(c))}{\partial a_{i,k}}}{\frac{\partial \widetilde{y}_{j,l}\left(\alpha^*(c), y'_{j,l}(c)\right)}{\partial a_{i,k}}}.$$

In this case $\lambda_{ij,l}$ will not be uniquely defined.

To avoid this problem, we assume that a_i affects firm j's outcome indirectly through a real-valued function $\phi_{ij,l}$ that maps each $a_i \in A_i$ into a real number such that the outcome of firm j in market l can be written as

$$y_{j,l}\left(\phi_{1j,l}(a_1), ..., \phi_{nj,l}(a_n), \varepsilon\right).$$

We assume that $y_{j,l}$ is defined on $\Pi_{i=1}^n F_{ij,l} \times \Omega$, where $F_{ij,l}$ is a compact interval that contains the range of $\phi_{ij,l}$, continuous, and differentiable in each set of the values of $\phi_{1j,l}, ..., \phi_{nj,l}$ for almost all ε. Under this formulation we can think of each firm i choosing a scaler $\phi_{ij,l}$ instead of a vector a_i. Let $\phi_{j,l}$ be the vector $(\phi_{1j,l}, ..., \phi_{nj,l})$ and $\phi_{-ij,l}$ the vector $\phi_{j,l}$ minus $\phi_{ij,l}$. Let $F_{j,l}$ denote $\Pi_{i=1}^n F_{ij,l}$ and $F_{-ij,l}$ denote $\Pi_{k=1}^{i-1} F_{kj,l} \times \Pi_{k=i+1}^n F_{kj,l}$. We can then restate Assumptions 1-3 in terms of $\phi_{ij,l}$ instead of a_i.

Assumption 1A. *For any $i, j \in \mathcal{N}$, $i \neq j$, and any $l \in \mathcal{M}$, $y_{j,l}$ is monotone in $\phi_{ij,l}$, and whether $y_{j,l}$ increases or decreases in $\phi_{ij,l}$ is independent of ε.*

Assumption 2A. *For any $i, j \in \mathcal{N}$, $i \neq j$, $l \in \mathcal{M}$, $\phi'_{ij,l}, \phi''_{ij,l} \in F_{ij,l}$, $\phi_{-ij,l} \in F_{-ij,l}$, and $y'_{j,l} \in \Re_+$,*

$$\frac{\frac{\partial \widetilde{y}_{j,l}\left(\phi'_{ij,l}, \phi_{-ij,l}, y'_{j,l}\right)}{\partial \phi_{ij,l}}}{\frac{\partial \overline{y}_{j,l}\left(\phi'_{ij,l}, \phi_{-ij,l}\right)}{\partial \phi_{ij,l}}} \leq \frac{\frac{\partial \widetilde{y}_{j,l}\left(\phi''_{ij,l}, \phi_{-ij,l}, y'_{j,l}\right)}{\partial \phi_{ij,l}}}{\frac{\partial \overline{y}_{j,l}\left(\phi''_{ij,l}, \phi_{-ij,l}\right)}{\partial \phi_{ij,l}}}, \tag{25}$$

whenever $y_{j,l}\left(\phi_{j,l}(a_i', a_{-i}), \varepsilon\right) \geq y_{j,l}\left(\phi_{j,l}(a_i'', a_{-i}), \varepsilon\right)$ *for all* ε *and the denominators on both sides of the inequality are non-zero.*

Assumption 3A. *There exists some finite* $\eta > 0$ *such that for each* $\phi_{j,l} \in F_{j,l}$,

$$\left|\frac{\partial y_{j,l}\left(\phi_{j,l}, \varepsilon\right)}{\partial \phi_{ij,l}}\right| \leq \eta \left|\frac{\partial \overline{y}_{j,l}\left(\phi_{j,l}, \varepsilon\right)}{\partial \phi_{ij,l}}\right| \quad \text{for almost all } \varepsilon \in \Omega.$$

A multi-market version of Proposition 1 can be established following the same steps in Section 4 with Assumptions 1A-3A replacing Assumptions 1-3. Intuitively, although firm i's action is multi-dimensional—firm i can chooses higher actions in some markets and lower ones in others—its impact on $y_{j,l}$ remains one-dimensional through the function $\phi_{ji,l}$. Note that each firm i will now be punished when the profit of firm j in any market l falls below the profit target for that market. Hence, in this case we will have

$$\widehat{w}_i^1\left(\widehat{c}, \widehat{y}_{-i}\right) \equiv \sum_{j \neq i} \sum_{l=1}^{m} \lambda_{ij,l}\left(\widehat{c}\right) \min\left(\widetilde{\pi}_{j,l}\left(\widehat{y}_{j,l}, \widehat{c}\right) - K_{j,l}\left(\widehat{c}\right), 0\right).$$

The assumption that a_i can influence $y_{j,l}$ indirectly through $\phi_{ij,l}$ is naturally met in the multi-market version of Examples 3-4 because in those examples the demand shock is a scalar in each y_j. When there there are multiple demand shocks in the outcome of a firm, the assumption means that different components of firm i's action cannot be subject to different random shocks. For example, in the multi-market version of Example 1 we would need to have

$$y_i\left(a, \varepsilon\right) = \varepsilon_{i0} + \sum_{j=1}^{n} \varepsilon_{ij} \sum_{k=1}^{m} l_{j,k} a_{j,k}.$$

The coefficient of the individual components of a_j, $l_{j,k}$, $k = 1, ..., m$, cannot be random variables.

5.2 Supportable action profiles

Our enforcement mechanism exploits the fact that α^* maximizes the total cartel payoff. Obviously, by assigning different weights to different firms, we can implement other outcomes on the Pareto frontier. More interestingly, we can also apply a different set of weights to different firms. For example, when there are two firms, we can apply a weight of 2 to π_1 and a weight of 1 to π_2 when we calculate the transfer of firm 1 but a weight of 1 to π_1 and a weight of 2 to π_2 when we calculate the transfer of firm 2. Doing so allows us to implement outcomes within the Pareto frontier. This could be potentially important because when δ is low it could be that only inefficient outcomes can be maintained in equilibrium.

Let

$$\Psi_i \equiv \left\{ (\theta_1, \theta_2, \ldots, \theta_n) \middle| \theta_i > 0,\ \theta_j \geq 0 \text{ for each } j \neq i, \text{ and } \sum_{j=1}^{n} \theta_j = 1 \right\}$$

denote the set of weights that sum to one, and assign strictly positive weight to firm i and nonnegative weight to each firm $j \neq i$.

Definition 1. A cost-action profile $\alpha : C \to A$ is supportable if, for each i, there exists a set of weights $\theta^i = (\theta^i_1, \theta^i_2, \ldots, \theta^i_n)$ in Ψ_i such that for each $c \in C$, $\alpha_i(c)$ maximizes

$$\sum_{j=1}^{n} \theta^i_j \int_{\varepsilon \in \Omega} \pi_j \left(a_i, \alpha_{-i}(c), c_j, \varepsilon\right) dF\left(\varepsilon\right).$$

Following the steps in Section 4, we can construct a trigger-strategy perfect public equilibrium in which a cost-action profile $\alpha : C \to A$ supported by $(\theta^1, \theta^2, \ldots, \theta^n)$ is played in the collusive state. The only change is that we need to adjust each $\lambda_{ij}(c)$ to reflect the different weights assigned on firm i's and firm j's profits. For $\overline{y}_j$ that decreases in a_i, we define

$$\lambda_{ij}\left(c\right) = \begin{cases} \sup_{a'_i \in B^1_{ij}(c)} (\theta^i_i)^{-1} \theta^i_j \dfrac{\frac{\partial \overline{y}_j\left(a'_i, \alpha_{-i}(c)\right)}{\partial a_i}}{\frac{\partial \widetilde{y}_j\left(a'_i, \alpha_{-i}(c), y'_j(c)\right)}{\partial a_i}} & \text{if } B^1_{ij}\left(c\right) \text{ is non-empty;} \\ 0 & \text{otherwise.} \end{cases}$$

For $\overline{y}_j$ that decreases in a_i, we define

$$\lambda_{ij}\left(c\right) = \begin{cases} \sup_{a'_i \in B^2_{ij}(c)} (\theta^i_i)^{-1} \theta^i_j \dfrac{\frac{\partial \overline{y}_j\left(a'_i, \alpha_{-i}(c)\right)}{\partial a_i}}{\frac{\partial \widetilde{y}_j\left(a'_i, \alpha_{-i}(c), y'_j(c)\right)}{\partial a_i}} & \text{if } B^2_{ij}\left(c\right) \text{ is non-empty;} \\ 0 & \text{otherwise.} \end{cases}$$

To further clarify the idea of supportable action profiles, we characterize the set of supportable action profiles in a Bertrand duopoly example. Let p_i denote the price of firm i. For simplicity we assume that the firms always have zero production cost, and each firm i has an expected demand

$$\widehat{q}_i(p_i, p_j) = \max\left(1 - p_i + l_j p_j, 0\right),$$

where $l_j \in (0,1)$. Consider some (θ^1, θ^2) with $\theta^i = (\theta^i_1, \theta^i_2) \in \Psi_i$ for each i. Note that p_i maximizes the total weighted expected profits

$$\theta^i_i p_i \widehat{q}_i(p_i, p_j) + \theta^i_j p_j \widehat{q}_j(p_i, p_j)$$

if and only if

$$\theta^i_i - 2\theta^i_i p_i + l_j p_j = 0.$$

Solving for θ^i_i yields

$$\theta^i_i = \frac{l_j p_j}{2p_i - 1}.$$

The constraints $0 < \theta^1_1 \leq 1$ and $0 < \theta^2_2 \leq 1$ translate into

$$1 - 2p_1 + l_j p_2 \leq 0 \qquad \text{and} \qquad 1 - 2p_2 + l_j p_1 \leq 0. \tag{26}$$

The nonnegativity constraints $\widehat{q}_1 \geq 0$ and $\widehat{q}_2 \geq 0$ require that

$$1 - p_1 + l_j p_2 \geq 0 \qquad \text{and} \qquad 1 - p_2 + l_j p_1 \geq 0. \tag{27}$$

The set of supportable pairs (p_1, p_2) is the quadrangle bounded by the four inequalities in (26) and (27). Note that the left-hand sides of the inequalities in (26) are the marginal profits of the two firms. Hence, any pair of prices such that each firm's price is higher than its best response but lower than the zero-output price is supportable.

5.3 Non-monetary transfers

Instead of using direct monetary transfers a cartel may adjust future quotas to implement our enforcement mechanism. In this section we illustrate the idea with a simple example. Suppose there are two firms, each with zero production cost. We suppress the arguments c_i and $\widehat{c}_i$ throughout this section. Since production cost is constant, $\widehat{w}_i\left(\widehat{y}\right)$ reduces to

$$\widehat{w}_i\left(\widehat{y}\right) = -\lambda_{ij} \max\left(K_j - \widetilde{\pi}_j\left(\widehat{y}_j\right), 0\right).$$

Hence

$$v_i^* = \int_\varepsilon \pi_i\left(\alpha^*, \varepsilon\right) - \lambda_{ij} \max\left(K_j - \pi_j\left(\alpha^*, \varepsilon\right), 0\right) dF\left(\varepsilon\right)$$

for each firm i.

Suppose that, in addition, there are two supportable (degenerated) cost-action profiles α^1 and α^2. For each $k = 1, 2$, we use K_j^k and λ_{ij}^k to denote the corresponding profit target and scaling factor for action profile α^k, respectively. Let

$$v_i^k = \int_\varepsilon \pi_i\left(\alpha^k, \varepsilon\right) - \lambda_{ij}^k \max\left(K_j^k - \pi_j\left(\alpha^k, \varepsilon\right), 0\right) dF\left(\varepsilon\right)$$

for each firm i. Let Λ^k denote the set of outcome profiles with positive densities when there is at most one firm i that does not choose α_i^k. We assume further that, for each $k \in \{1, 2, *\}$ and each $\widehat{y} \in \Lambda \cup \Lambda^1 \cup \Lambda^2$, there exist four nonnegative real numbers $p_1^k\left(\widehat{y}\right)$, $p_2^k\left(\widehat{y}\right)$, $p_*^k\left(\widehat{y}\right)$, and $p_N^k\left(\widehat{y}\right)$ that sum to one and satisfy

$$\sum_{j\in\{1,2,*,N\}} p_j^k\left(\widehat{y}\right) v_1^j = v_1^* - \delta^{-1}(1-\delta)\widehat{w}_1\left(\widehat{y}\right), \tag{28}$$

$$\sum_{j\in\{1,2,*,N\}} p_j^k\left(\widehat{y}\right) v_2^j = v_2^* - \delta^{-1}(1-\delta)\widehat{w}_2\left(\widehat{y}\right), \tag{29}$$

This condition requires that for any $\widehat{y} \in \Lambda \cup \Lambda^1 \cup \Lambda^2$, the vector

$$\left(v_1^* - \delta^{-1}(1-\delta)\widehat{w}_1\left(\widehat{y}\right), v_2^* - \delta^{-1}(1-\delta)\widehat{w}_2\left(\widehat{y}\right)\right)$$

belongs to the convex hull of v^1, v^2, v^*, and v^N. Intuitively, this condition is more likely to be met if for each $i = 1, 2$,

$$v_i^i - v_i^* = \int_\varepsilon \pi_i\left(\alpha^i, \varepsilon\right) - \lambda_{ij}^i \max\left(K_j^i - \pi_j\left(\alpha^i, \varepsilon\right), 0\right) dF\left(\varepsilon\right) - v_i^*$$

is large. In that case a cartel can transfer profits from firm i to firm j by switching to α^j. Formally, consider the following repeated-game strategy profile. There are four states, denoted by $*$, 1, 2, and N, respectively. In the first period the firms start in state $*$. In each state $k \in \{*, 1, 2\}$, the firms play α^k in the current period, report their profits truthfully, and switch to state $j \in \{*, 1, 2, N\}$ with probability $p_j^k(\widehat{y})$. In state N, the firms play the state-game Nash equilibrium forever. The average discounted payoff for firm i in state k is v_i^k. Assuming that firm i, currently in state $k \in \{1, 2, *\}$, is going to follow its strategy from the next period onwards, its objective in the current period is to choose a_i to maximize

$$(1-\delta)\int_{\varepsilon} \pi_i\left(a_i, \alpha_j^k, \varepsilon\right) - \lambda_{ij}^k \max\left(K_j^k - \pi_j\left(a_i, \alpha_j^k, \varepsilon\right), 0\right) dF(\varepsilon) + \delta v_i^*.$$

By the discussion in Section 5.2, this expression is maximized by α_i^k. It then follows from the standard one-step-deviation-proof argument that the strategy profile is a perfect public equilibrium.

6 Concluding Remarks

In this paper we describe how a cartel can enforce a collusive agreement by constructing penalties that tie a firm's continuation profit with the current profits of the other firms' profits. We show that for this approach to work it is not necessary for a firm's continuation profit to increase one to one with the total profits of the other firms, as in a standard Clark-Grove mechanism. Instead, for a wide set of stochastic demand systems, it is sufficient that a firm's continuation profit is increasing in proportion with the upper-truncated profit functions of the other firms. The truncation imporves the cartel profit. We show that when demand shocks are small, the cartel profit will be closed to the monopoly profit.

This approach, we believe, has three nice features. First, because any form of cheating by a firm, whether by lying about cost or secretly lowering price, must inevitably harm the profits of the firms, it is particularly well-suited to an environment where neither costs nor prices are publicly observed. Second, because the penalty of a firm depends only on the profits of the other firms, a firm would have no incentive to lie about its sales under our approach. This avoids the complications that would have arisen if a firm's continuation profit is tied to its own sales report, as in Harrington and Skrzypacz (2011). Finally, the approach is very robust. It applies to both price and quantitiy competition, and can be easily generalized to allow the firms to collude in more than one price.

One common criticism against models of repeated games is that the strategies are too complex to be relevant in reality. In our case, while the actual equilibrium strategies are complicated, the idea of internalizing the externalities of a firm's action by holding it responisble for the profits of the other firms is simple and intuitive. Indeed, as Harrington (2006), Harrington and Skrzypacz (2011), Suslow (2005) have documented compensation schemes where firms that sell above quotas are required to compensate firms that sells below are common among cartels in the real world. Suslow (2005) finds that cartels that used some type of self-imposed penalty schemes were more stable. Our theoretical exercise shows that the underlying logic of our argument is sound in a very general setting. While in reality firms may not get the incentives exactly right, a reasonably calibrated compensation scheme, similar to the ones documented in the empirical literature, probably can go a long way to eliminate the incentives to cheat.

A Appendix

We will show that Examples 3 and 4 satisfy Assumption 2. In these examples, firm i's outcome can be written as

$$y_i\left(a,\varepsilon\right)=\nu_i\left(a_i\right)h_i\left(\chi_i\left(a\right)+\varepsilon_i\right),$$

where h_i is increasing and χ_i monotone. Specifically, in Example 3, $h_i\left(x\right)$ equals 0 when $x\leq 0$ and equals x when $x>0$, $\chi_i=\sum_{j=1}^{n}l_{ij}a_j$, and $\nu_i\left(a_i\right)=1$. In Example 4, $h_i\left(x\right)=-1/x$, $\chi_i=-\sum_{j=1}^{n}\exp(-l_ja_j)$, $\nu_i\left(a_i\right)=\exp(-l_ia_i)$, and $\varepsilon_i=-\varepsilon$.

Since a_{-j} and $\nu_i\left(a_i\right)$ are to be treated as constants throughout, we will suppress them in the following to simplify notations. With a slight abuse of notation, let F_i, f_i, and $[\underline{\varepsilon}_i,\overline{\varepsilon}_i]$ denote the distribution, density, and support of ε_i, respectively. Recall that we assume f_i is strictly positive, differentiable, and log-concave for any $\varepsilon_i\in[\underline{\varepsilon}_i,\overline{\varepsilon}_i]$. Let

$$m_i\left(y_i,a_j\right)\equiv h_i^{-1}\left(y_i\right)-\chi_i(a_j).$$

We can write the distribution of y_i as

$$G_i\left(y_i, a_j\right) = F_i\left(m_i\left(y_i, a_j\right)\right).$$

Let $y_i^{\max} \equiv \max_{a \in A} y_i\left(a, \overline{\varepsilon}_i\right)$. Through changing variables and integration by parts, we have

$$\begin{aligned}
\overline{y}_i\left(a_j\right) &= \int_0^{y_i^{\max}} y_i dG_i\left(y_i, a_j\right) = y_i^{\max} - \int_0^{y_i^{\max}} G_i\left(y_i, a_j\right) dy_i; \\
\widetilde{y}_i\left(a_j, y_i'\right) &= \int_0^{y_i^{\max}} \min\left(y_i, y_i'\right) dG_i\left(y_i, a_j\right) = y_i' - \int_0^{y_i'} G_i\left(y_i, a_j\right) dy_i.
\end{aligned}$$

Hence, for any a_j such that $\partial \overline{y}_i\left(a_j\right)/\partial a_j \neq 0$,

$$R\left(a_j\right) \equiv \frac{\frac{\partial \widetilde{y}_i\left(a_j, y_i'\right)}{\partial a_j}}{\frac{\partial \overline{y}_i\left(a_j\right)}{\partial a_j}} = \frac{\int_0^{y_i'} \frac{\partial G_i}{\partial a_j} dy_i}{\int_0^{y_i^{\max}} \frac{\partial G_i}{\partial a_j} dy_i} = \frac{\int_0^{y_i'} f_i\left(m_i\left(y_i, a_j\right)\right) dy_i}{\int_0^{y_i^{\max}} f_i\left(m_i\left(y_i, a_j\right)\right) dy_i} \geq 0.$$

Suppose $\partial \widetilde{y}_i\left(a_j, y_i'\right)/\partial a_j \neq 0$. In this case, we can write

$$R\left(a_j\right) = \frac{1}{1 + \widehat{R}\left(a_j\right)},$$

where

$$\widehat{R}\left(a_j\right) \equiv \frac{\int_{y_i'}^{y_i^{\max}} f_i\left(m_i\left(y_i, a_j\right)\right) dy_i}{\int_0^{y_i'} f_i\left(m_i\left(y_i, a_j\right)\right) dy_i}.$$

Differentiating $\widehat{R}$ with respect to a_j, we obtain

$$\frac{d\widehat{R}(a_j)}{da_j} = \chi_i'(a_j)\widehat{R}\left(a_j\right)\left(\frac{\int_0^{y_i'} f_i'\left(m_i\left(y_i, a_j\right)\right) dy_i}{\int_0^{y_i'} f_i\left(m_i\left(y_i, a_j\right)\right) dy_i} - \frac{\int_{y_i'}^{y_i^{\max}} f_i'\left(m_i\left(y_i, a_j\right)\right) dy_i}{\int_{y_i'}^{y_i^{\max}} f_i\left(m_i\left(y_i, a_j\right)\right) dy_i}\right). \tag{30}$$

Let

$$\begin{aligned}
\gamma_1(y_i, a_j) &\equiv \frac{f_i\left(m_i\left(y_i, a_j\right)\right)}{\int_{y_i'}^{y_i^{\max}} f_i\left(m_i\left(y_i, a_j\right)\right) dy_i}; \\
\gamma_2(y_i, a_j) &\equiv \frac{f_i\left(m_i\left(y_i, a_j\right)\right)}{\int_0^{y_i'} f_i\left(m_i\left(y_i, a_j\right)\right) dy_i}.
\end{aligned}$$

Since h is increasing and f is log-concave, $f_i'(m_i(y_i,a_j))/f_i(m_i(y_i,a_j))$ is decreasing in y_i. Hence the bracketed term on the right-hand side of (30) is equal to

$$\begin{aligned}
& \frac{\int_0^{y_i'} f_i'\left(m_i\left(y_i,a_j\right)\right) dy_i}{\int_0^{y_i'} f_i\left(m_i\left(y_i,a_j\right)\right) dy_i} - \frac{\int_{y_i'}^{y_i^{\max}} f_i'\left(m_i\left(y_i,a_j\right)\right) dy_i}{\int_{y_i'}^{y_i^{\max}} f_i\left(m_i\left(y_i,a_j\right)\right) dy_i} \\
= & \int_0^{y_i'} \gamma_2(y_i,a_j) \frac{f_i'\left(m_i\left(y_i,a_j\right)\right)}{f_i\left(m_i\left(y_i,a_j\right)\right)} dy_i - \int_{y_i'}^{y_i^{\max}} \gamma_1(y_i,a_j) \frac{f_i'\left(m_i\left(y_i,a_j\right)\right)}{f_i\left(m_i\left(y_i,a_j\right)\right)} dy_i \\
\geq & \frac{f_i'\left(m_i\left(y_i',a_j\right)\right)}{f_i\left(m_i\left(y_i',a_j\right)\right)} \left(\int_0^{y_i'} \gamma_2(y_i,a_j) dy_i - \int_{y_i'}^{y_i^{\max}} \gamma_1(y_i,a_j)\, dy_i \right) \\
= & \ 0.
\end{aligned}$$

Thus, $d\widehat{R}/da_j \geq 0$ if and only if $\chi_i'(a_j) \geq 0$. Since dR/da_j and $d\widehat{R}/da_j$ have opposite signs, $dR/da_j \leq 0$ if and only if $\chi_i'(a_j) \geq 0$.

We need to prove that for any a_j' and a_j such that $R\left(a_j'\right)$ and $R\left(a_j\right)$ are well-defined, $R\left(a_j'\right) \leq R\left(a_j\right)$ if and only if $\overline{y}_i\left(a_j'\right) \geq \overline{y}_i\left(a_j\right)$. We focus on the case where $\chi_i'(a_j) \geq 0$. The case for $\chi_i'(a_j) \leq 0$ is similar and omitted. Suppose $\chi_i'(a_j) \geq 0$ and $a_j' \geq a_j$. Suppose that $R\left(a_j\right) = 0$. Since $f_i\left(\varepsilon_i\right)$ has strictly positive for any $\varepsilon_i \in [\underline{\varepsilon}_i, \overline{\varepsilon}_i]$, and $y_i(a_j, \overline{\varepsilon}_i) \geq 0$ implies that $m_i\left(0, a_j\right) \leq \overline{\varepsilon}_i$, we must have $m_i\left(y_i', a_j\right) \leq \underline{\varepsilon}_i$. Since m_i is decreasing in a_j, it follows that $m_i\left(y_i', a_j'\right) \leq \underline{\varepsilon}_i$ and $R\left(a_j'\right) = 0$. Suppose that $R\left(a_j\right) > 0$ and $R\left(a_j'\right) = 0$. Then clearly $R\left(a_j'\right) < R\left(a_j\right)$. Finally, suppose that both $R\left(a_j\right) > 0$ and $R\left(a_j'\right) > 0$. Then $R\left(a_j''\right) > 0$ for all $a_j'' \in \left[a_j, a_j'\right]$. From the argument in the last paragraph, dR/da_j is defined and negative for all $a_j'' \in \left[a_j, a_j'\right]$. Hence, in this case $R\left(a_j'\right)$ must also be less than $R\left(a_j\right)$.

References

Aliprantis, Charalambos D., and Owen Burkinshaw, 1990, *Principles of Real Analysis* (Academic Press) 2nd edn.

Aoyagi, Masaki, 2002, Collusion in dynamic Bertrand oligopoly with correlated private signals and communication, *Journal of Economic Theory* 102, 229 – 248.

——— , 2003, Bid rotation and collusion in repeated auctions, *Journal of Economic Theory* 112, 79–105.

Athey, Susan, and Kyle Bagwell, 2001, Optimal collusion with private information, *The RAND Journal of Economics* 32, 428–465.

——— , 2008, Collusion with persistent cost shocks, *Econometrica* 76, 493–540.

——— , and Chris Sanchirico, 2004, Collusion and price rigidity, *The Review of Economic Studies* 71, 317–349.

Athey, Susan, and David A. Miller, 2007, Efficiency in repeated trade with hidden valuations, *Theoretical Economics* 2.

Athey, Susan, and Ilya Segal, 2007, An efficient dynamic mechanism, Discussion paper.

Bergemann, Dirk, and Juuso Välimäki, 2010, The dynamic pivot mechanism, *Econometrica* 78, 771–789.

Blume, Andreas, and Paul Heidhues, 2006, Private monitoring in auctions, *Journal of Economic Theory* 131, 179–211.

Compte, Olivier, 1998, Communication in repeated games with imperfect private monitoring, *Econometrica* 66, 597–626.

Escobar, Juan F., and Juuso Toikka, 2010, A folk theorem with markovian private information, mimeo, MIT.

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, 461 – 473.

Fudenberg, Drew, and Jean Tirole, 1991, *Game Theory* (The MIT Press).

Green, Edward J., and Robert H. Porter, 1984, Noncooperative collusion under imperfect price information, *Econometrica* 52, 87–100.

Harrington, Joseph E., and Andrzej Skrzypacz, 2011, Private monitoring and communication in cartels: Explaining recent collusive practices, *American Economic Review* 101, 2425–49.

Harrington, Joseph E. Jr., 2006, How do cartels operate?, *Foundations and Trends in Microeconomics* 2, 1–105.

——— , and Andrzej Skrzypacz, 2007, Collusion under monitoring of sales, *The RAND Journal of Economics* 38, 314–331.

Hörner, Johannes, and Julian Jamison, 2007, Collusion with (almost) no information, *The RAND Journal of Economics* 38, 804–822.

Kandori, Michihiro, and Hitoshi Matsushima, 1998, Private observation, communication and collusion, *Econometrica* 66, 627–652.

Levenstein, Margaret C., and Valerie Y. Suslow, 2006, What determines cartel success?, *Journal of Economic Literature* 44, 43–95.

Miller, David A., Forthcomming, Robust collusion with private information, *The Review of Economic Studies*.

Obara, Ichiro, 2009, Folk theorem with communication, *Journal of Economic Theory* 144, 120 – 134.

Skrzypacz, Andrzej, and Hugo Hopenhayn, 2004, Tacit collusion in repeated auctions, *Journal of Economic Theory* 114, 153–169.

Stigler, George J., 1964, A theory of oligopoly, *The Journal of Political Economy* 72, 44–61.

Suslow, Valerie Y., 2005, Cartel contract duration: Empirical evidence from interwar international cartels, *Industrial and Corporate Change* 14, 705 – 744.

Zheng, Bingyong, 2008, Approximate efficiency in repeated games with correlated private signals, *Games and Economic Behavior* 63, 406 – 416.