

Collusion under monitoring of sales

Joseph E. Harrington, Jr.*

and

Andrzej Skrzypacz**

Collusion under imperfect monitoring is explored when firms' prices are private information and their quantities are public information; such an information structure is consistent with several recent price-fixing cartels, such as those in lysine and vitamins. For a class of symmetric oligopoly games, it is shown that symmetric equilibrium punishments cannot sustain any collusion. An asymmetric punishment is characterized that does sustain collusion and it has firms whose sales exceed their quotas compensating those firms with sales below their quotas. In practice, cartels could have performed such transfers through sales among the cartel members.

... if I'm assured that I'm gonna get 67,000 tons [of lysine sales] by the year's end, we're gonna sell it at the prices we agreed to and I frankly don't care what you sell it for.

Terrance Wilson of Archer Daniels Midland from the March 10, 1994, meeting of the lysine cartel.

And that total for us for the year, calendar year is 68,000; 68,334. 68,334 and our target was 67,000 plus alpha. Almost on target.

Mark Whitacre of Archer Daniels Midland from the January 18, 1995, meeting of the lysine cartel.¹

1. Introduction

■ Many, if not most, price-fixing cartels involve firms selling to industrial buyers, with the lysine cartel being a notable example. As price can be settled through private negotiation, it is not typically observable. In such cases, compliance with the collusive agreement is often based on firms' sales. Indeed, cartels can go to great lengths to ensure that sales are public information among the cartel members. In the citric acid cartel, for example, firms hired an international accounting firm to independently audit sales reports (Connor, 2001). The objective of this article

* Johns Hopkins University; joe.harrington@jhu.edu.

** Stanford University; andy@gsb.stanford.edu.

We appreciate the comments of two anonymous referees, Marco Battaglini, Sadao Nagaoka, Rui Ota, Phil Reny, Mike Whinston and seminar or conference participants at the University of Tokyo, Hitotsubashi University, 2004 Duke-Northwestern-Texas IO Theory Conference, 2005 International Industrial Organization Conference and 2005 World Congress of the Econometric Society. The first author acknowledges the support of the National Science Foundation under grant nos. SES-0209486 and SES-0516943 and the hospitality of the Department of Economics at Harvard University. The second author acknowledges the financial support of the Center for Electronic Business and Commerce at the Stanford Graduate School of Business.

¹ These quotes are from the video transcript of "The International Lysine Cartel at Work, 3/28/00" provided by the U.S. Department of Justice, Antitrust Division.

is to explore collusion in an imperfect monitoring setting in which prices are private information and firms' quantities are public information.

In spite of such a monitoring environment being applicable to many market settings, there is relatively little work with such a structure even though, interestingly enough, it was the one described by Stigler (1964) when he originally raised the issue of imperfect monitoring. There are, of course, many articles using the classical monitoring setting of Green and Porter (1984), in which firms' quantities are private information and the market price is publicly observed. In the context of repeated auctions, Blume and Heidhues (2003) and Skrzypacz and Hopenhayn (2004) assume price is private information, while who won the auction is known. However, the assumption of one unit per period makes the model unsuitable for many markets and, pertinent to the issue at hand, constrains the monitoring of collusion through sales (a point we elaborate on later). Tirole (1988), Bagwell and Wolinsky (2002) and Campbell, Ray, and Muhanna (2005) allow for multiunit demand in the context of the infinitely repeated Bertrand price model with imperfect monitoring. The assumption that firm demand is discontinuous is obviously extreme and, furthermore, plays an important role in sustaining collusion. Our model is the first to consider collusion when prices are private information and monitoring occurs with respect to sales, while making standard and fairly general demand assumptions: demand is multiunit and expected firm demand is everywhere continuous.

Our first main finding is a surprising impossibility result. For a general class of symmetric demand structures with inelastic market demand, no collusion can be sustained by symmetric punishments. By way of example, one such demand structure is when the probability distribution of demand depends only on the difference in firms' prices, as is true with the discrete-choice model. The rough intuition for our result can be conveyed as follows for the duopoly case. To begin, one would expect punishment to occur when market shares are sufficiently skewed. Suppose, for example, punishment occurs when the market share of one of the firms exceeds 70%. A firm that considers charging a price below the collusive price raises the probability that its market share exceeds 70%—which makes punishment more likely—but lowers the probability that the other firm's market share exceeds 70%—which makes punishment less likely. What we show is that, for small price cuts, these two effects exactly offset each other, which implies that a firm's continuation payoff is unaffected by its price. Therefore, an equilibrium price for the infinite horizon game must be the same as that for the stage game. Though shown for the extreme case of fixed market demand, robustness prevails when market demand is stochastic and sensitive to firms' prices. Specifically, if market demand is very insensitive to firms' prices, then collusive prices are very close to noncollusive prices.

The conclusion we draw from this result is not that firms cannot collude but rather of the importance of treating apparent deviators differently from apparent nondeviators. The second main result is showing that collusion can be sustained with asymmetric punishments involving transfers in which firms that sold too much compensate those who sold too little. In fact, some price-fixing cartels, such as those in citric acid (Arbault et al., 2002) and sodium gluconate (European Commission, 2002), did indeed deploy asymmetric punishments through the use of interfirm sales, which can act as transfers. The main message of this article is that, if we are to understand the actual practices of some cartels, it is essential that we take account of imperfect monitoring with respect to prices and the role of asymmetric punishments that condition on sales.

After the model is described in Section 2, the inability of symmetric punishments to sustain collusion is established in Section 3. Some robustness issues are explored in Section 4, while a characterization of asymmetric equilibria that sustain collusion is provided in Section 5. We relate these results to the literature in Section 6 and briefly conclude in Section 7.

2. Model

■ Consider an infinitely repeated game in which $n \geq 2$ firms make simultaneous price decisions. Cost functions are common and linear and, without loss of generality, cost is zero. Demand

is fixed at m discrete units.² We often refer to there being m customers (with unit demands). Though total demand is fixed, firm demand is stochastic. Letting q_i denote the quantity of firm i , the set of feasible quantity vectors is

$$\Delta \equiv \left\{ (q_1, \dots, q_n) \in \{0, 1, \dots, m\}^n : \sum_{i=1}^n q_i = m \right\}.$$

Define $\psi(\underline{q}; \underline{p}) : \Delta \times \mathbb{R}^n \rightarrow [0, 1]$ as the probability of realizing quantity vector \underline{q} given price vector \underline{p} . As regards the stochastic nature of firm demand, one can imagine that products are differentiated or that they are homogeneous but buyer-specific shocks, which may be independent or correlated, that influence demand in each period. We describe some examples below.

We make three assumptions on the probability distribution on firm demand.

Assumption 1. ψ is continuously differentiable with respect to p_i for all i .

Assumption 2. $\psi(\underline{q}; \underline{p}) = \psi(\omega(\underline{q}; i, j); \omega(\underline{p}; i, j))$, for all i, j and all $(\underline{q}; \underline{p})$, where $\omega(\underline{q}; i, j)$ is the vector \underline{q} when elements i and j are exchanged.

Assumption 3. $\sum_{i=1}^n (\partial \psi(\underline{q}; (p, \dots, p)) / \partial p_i) = 0$, for all (\underline{q}, p) .

Assumption 1 is standard and Assumption 2 imposes symmetry in that permuting the price vector permutes the probability function. Assumption 3 is the key restriction, though it is satisfied in many models. Assumption 3 implies, that if we start at equal prices, then the distribution of demand remains unchanged if firms make small, identical price changes.

When $n = 2$, Assumption 3 holds if the demand distribution depends solely on the difference in prices; in that case, equal changes in price do not affect the difference. For general n , a sufficient condition for Assumption 3 to be true is that ψ depends only the price differences for all pairs of firms. To show this explicitly, consider $n = 3$ and suppose $\exists \xi : \Delta \times \mathbb{R}^3 \rightarrow [0, 1]$ such that

$$\psi(\underline{q}; \underline{p}) = \xi(\underline{q}; \Delta_{12}, \Delta_{23}), \quad \forall (\underline{q}; \underline{p}) \in \Delta \times \mathbb{R}^3,$$

where $\Delta_{ij} \equiv p_i - p_j$. Hence, the probability function depends only on the pairwise differences in firms' price. Then

$$\begin{aligned} \sum_{i=1}^3 \left(\frac{\partial \psi(\underline{q}; (p, \dots, p))}{\partial p_i} \right) &= \frac{\partial \xi(\underline{q}; 0)}{\partial \Delta_{12}} \left(\frac{\partial \Delta_{12}}{\partial p_1} + \frac{\partial \Delta_{12}}{\partial p_2} \right) + \frac{\partial \xi(\underline{q}; 0)}{\partial \Delta_{23}} \left(\frac{\partial \Delta_{23}}{\partial p_2} + \frac{\partial \Delta_{23}}{\partial p_3} \right) \\ &= \frac{\partial \xi(\underline{q}; 0)}{\partial \Delta_{12}} (1 - 1) + \frac{\partial \xi(\underline{q}; 0)}{\partial \Delta_{23}} (1 - 1) = 0, \end{aligned}$$

so that Assumption 3 holds.

An example from the literature that conforms to our demand specification is the following m -buyer generalization of the duopoly model of Cabral and Riordan (1994). Let the probability that firm 1 sells to a particular buyer equal $F(p_2 - p_1)$, where $F : \mathbb{R} \rightarrow [0, 1]$ is continuously differentiable and nondecreasing and F' is symmetric around zero. Assume also that buyers' decisions regarding from whom to buy are i.i.d. That implies that a firm's demand is binomially distributed,

$$\psi(b, m - b; p_1, p_2) = \frac{m!}{b!(m - b)!} F(p_2 - p_1)^b (1 - F(p_2 - p_1))^{m-b},$$

so only the price difference matters.

² See Section 4 for a generalization to when m is variable.

The discrete-choice model, in which consumer indirect utility is linear in price, is another common example satisfying Assumption 1–Assumption 3. In that case, the utility to consumer j from buying the product of firm i is $U_i^j - p_i$, so that customer j buys firm i 's product iff

$$U_i^j - p_i > U_h^j - p_h, \quad \forall h \neq i$$

or, equivalently,

$$U_i^j - U_h^j > p_i - p_h \quad \forall h \neq i.$$

If F_h^j is the c.d.f. on U_h^j , then the probability that customer j buys firm i 's product is

$$\int \prod_{h \neq i} F_h(U_i^j + p_h - p_i) dF_i(U_i^j),$$

so that, once again, only price differences matter. However, if there is an outside option, then the discrete-choice demand system does not satisfy Assumption 3.

More generally, note that we can represent $\psi(q; p)$ by $\hat{\psi}(q; f^q(p)) : \Delta \times \mathfrak{R} \rightarrow [0, 1]$, where $f^q(\cdot)$ is allowed to vary with q . It follows that Assumption 3 holds when

$$\sum_{i=1}^n \frac{\partial f^q(p, \dots, p)}{\partial p_i} = 0 \quad \forall q$$

and, furthermore, for any smooth transformation $g(f^q(\cdot))$ or $f^q(g(p_1), \dots, g(p_n))$. For example, start with $f^q(p_1, p_2) = p_1 - p_2$ and use the transformation: $g(p) = \ln(p)$. We then have $f^q(p_1, p_2) = \ln(p_1) - \ln(p_2) = \ln(p_1/p_2)$. Performing another transformation using $g(f^q) = \exp f^q$ gives us $f^q(p_1, p_2) = p_1/p_2$. Thus, if the probability distribution depends only on the ratio of prices, then it satisfies Assumption 3.

There is an infinite horizon and each firm's payoff is the expected present value of its profit stream where the common discount factor is $\delta \in (0, 1)$. The information structure is one of imperfect monitoring as firms' price decisions are private information though firms' quantities are public information. This conforms to the industrial buyer case in which price is negotiated between a seller and a buyer and thus is not publicly posted.³ It is sufficient to think of a public history at the start of period t , denoted h^{t-1} , to be a sequence of quantities sold by firm 1. Denote by H^{t-1} the set of all possible histories h^{t-1} . A firm's (public) strategy is then an infinite sequence of price functions, $\{p_i^t(\cdot)\}_{t=1}^{\infty}$, where $p_i^t : H^{t-1} \rightarrow \mathfrak{R}$. We restrict attention to perfect public equilibria so that firms do not condition their prices on their own past prices, just on the realized quantities.⁴ One final assumption is that first-order conditions are sufficient for defining an equilibrium.

The imperfect monitoring structure we consider obviously differs from the classical formulation of Green and Porter (1984), in which firms' quantities are private information and the market price is publicly observed. Assuming firms' prices are private information and monitoring is based on sales appears to better conform with many price-fixing cases. Nevertheless, there is a limited amount of work that considers such an informational structure. Blume and Heidhues (2003) and Skrzypacz and Hopenhayn (2004) consider collusion in repeated single-unit auctions. The limitation to one unit per period is restrictive and, as a result, their models are not applicable to many markets. Tirole (1988) and Bagwell and Wolinsky (2002) consider the Bertrand price model with uncertain aggregate demand in which firms' prices and quantities are private information. The standard Bertrand assumption of infinite elasticity of firm demand is clearly an extreme (though common) assumption, especially as even arbitrarily small deviations lead to a

³ Though list prices may be posted, they are often unrelated to transaction prices.

⁴ For equilibria in pure strategies, focusing on public strategies is without loss of generality.

discontinuous change in the distribution of the monitoring variable. Our model is then unique in the imperfect monitoring literature in allowing for the following compelling features: price is private information, monitoring occurs with respect to sales, multiunit demand, and expected firm demand is everywhere continuous. Though we do assume total market demand is fixed, robustness is established with respect to that assumption. All of these features—including highly inelastic market demand—fit well with many markets, including lysine and vitamins.

There are two assumptions that warrant discussion before moving on. First, it is assumed that firms' quantities are public information. Though this is not an issue when there are just two firms and market demand is perfectly inelastic—as a firm can infer the other firm's sales from its own—it is an issue more generally. In defense of this assumption, some cartels have gone to great lengths to make sales public information. The case of the citric acid cartel hiring an independent auditor was mentioned earlier. Furthermore, the documents on the lysine cartel suggest that cartel members acted "as if" the reported sales numbers were accurate. Though the validity of this assumption demands further investigation—both empirically and theoretically (what are the incentives to truthfully report sales?)—we feel the evidence provides some justification for making this assumption.

The second assumption is that, in some cases, we restrict a firm to charging the same price to all buyers, even though buyers are discrete. When it comes to deviating from a collusive arrangement, a firm may want to undercut its competitors' prices on only a subset of consumers so as to make detection less likely. (Of course, to a limited extent, it can do that by not undercutting as much.) For the impossibility result (Theorem 1), the assumption of a common price simplifies the analysis and, most important, the result is robust to allowing for customer-specific prices; if collusion is unsustainable when a deviating firm is constrained to setting the same price to all buyers, it is surely unsustainable when the set of options for a deviating firm is expanded. The assumption of a common price for all buyers is more of an issue when it comes to establishing positive results about collusion. Section 4 explores some robustness issues with respect to the impossibility theorem and has some results regarding when collusion is sustainable under that assumption. The main positive result is provided in Section 5, and there we do allow for customer-specific prices and show that one can construct asymmetric punishments to support collusion.⁵

Though we do not then need to assume a firm charges the same price to all buyers, let us conclude by arguing that, in some contexts, such an assumption may be reasonable. In many, if not most, price-fixing cartels, collusion is among high-level managers rather than sales representatives (that is, those who actually deal with customers); collusion then requires pricing authority to be centralized. For a manager to instruct sales representatives to charge different prices to different customers, it is necessary that they be identifiable *ex ante*, which means the manager needs to be able to predict who will be the buyers in the market in the current period. This, however, may not always be feasible. An individual customer's demand may be short-lived so that the market is continually replenished with new buyers. Even if the pool of customers is stable, the identity of the customers seeking to buy in that period may be random. Alternatively, a manager could mandate a distribution of prices without specifying the price for a specific customer but that may not be implementable if the number of customers is random and they arrive sequentially. There are then settings for which it may be reasonable to require a firm to set a common price for all buyers.

3. An impossibility result

■ With single-unit demand per period ($m = 1$), symmetric equilibria⁶ are trivially ineffective at supporting collusion if the players observe only sales and not actual prices.⁷ The reason is

⁵ Clearly, this result also holds when we constrain a firm to charge a common price to all buyers.

⁶ In the sense of strongly symmetric equilibria, as we define below.

⁷ This was first noted in Blume and Heidhues (2003) and Skrzypacz and Hopenhayn (2004), who explore collusion in repeated auctions. Their work will be discussed later.

prosaic: regardless of firms' prices, the customer buys from one of the sellers and this means that continuation play has to treat symmetrically the "winner" and the "loser." After either outcome, we would have to end up in a punishment or nonpunishment regime and, hence, there can be no symmetric punishment for secret price cutting.

However, with more than one customer per period, one might expect to be able to sustain collusion even with symmetric punishments. Considering the case of two customers and two sellers, two natural outcomes emerge: the sellers split the market or one of the sellers serves the whole market. If the collusive scheme recommends that they set a common high price, then a market split would seem less likely if one of the players deviates by charging a lower price. If so, then a punishment can be conditioned on market shares being skewed. This intuition is confirmed if we model the market as a continuum of independent customers as, by the law of large numbers, demand is then nonstochastic, which means deviations can be detected precisely. However, as we show in this section, that intuition is not correct in a large class of markets. If there is a finite number of customers, then no symmetric equilibrium can achieve prices above the competitive level.

In exploring collusion in a symmetric setting, it is common (and one might suppose natural) to first consider equilibria that take full advantage of this symmetry. For a particular strategy profile, let $v_i^t(\cdot) : H^{t-1} \rightarrow \mathfrak{R}$ denote the continuation payoff starting at t as a function of the public history. A set of symmetric histories consists of the initial null history, denoted h^0 , and, if m is a multiple of n , also of histories in which each firm had sales of m/n in every period. A symmetric Nash equilibrium is a Nash equilibrium in which the strategy profile calls for identical prices when the history is symmetric. This implies that continuation payoffs are identical across firms after such histories.

The next property, which we believe is compelling, is exchangeability. A Nash equilibrium is exchangeable if, when we permute the histories of firms i and j , the strategies of firms i and j are permuted while other firms' strategies are left unchanged. This produces the following property for continuation payoff functions:

$$\begin{aligned} v_i^t(\underline{q}^1, \dots, \underline{q}^{t-1}) &= v_j^t(\omega(\underline{q}^1; i, j), \dots, \omega(\underline{q}^{t-1}; i, j)) \\ v_j^t(\underline{q}^1, \dots, \underline{q}^{t-1}) &= v_i^t(\omega(\underline{q}^1; i, j), \dots, \omega(\underline{q}^{t-1}; i, j)) \\ v_k^t(\underline{q}^1, \dots, \underline{q}^{t-1}) &= v_k^t(\omega(\underline{q}^1; i, j), \dots, \omega(\underline{q}^{t-1}; i, j)), \quad k \neq i, j. \end{aligned}$$

A more restrictive but commonly imposed property is that of strong symmetry.⁸ A *strongly symmetric Nash equilibrium* is one in which strategies are symmetric for all histories. That implies the continuation payoffs are also symmetric after all histories,

$$v_i^t(h^{t-1}) = v_j^t(h^{t-1}) \quad \forall h^{t-1} \in H^{t-1}, \forall t, \forall i, j.$$

Let this common continuation payoff be denoted $v^t(\cdot)$.

The final restriction is designed to eliminate equilibria that condition on the history in nonmeaningful ways. Define $\tilde{p}^t(h^{t-1})$ to be the equilibrium price vector in period t after history h^{t-1} and let \hat{p} be a static Nash equilibrium price vector. An equilibrium is said to be history relevant when, if $\tilde{p}^t(h^{t-1}) = \hat{p}$, then $v_i^{t+1}(h^{t-1}, \underline{q}') = v_i^{t+1}(h^{t-1}, \underline{q}'')$ for all $\underline{q}', \underline{q}'' \in \Delta$, for all i . That is, if firms charge static Nash prices in period t , then the period $t + 1$ continuation payoff functions are independent of the period t quantities. This prevents equilibria—that start with static Nash prices—from circumventing the exchangeability restrictions through the use of an asymmetric history.⁹ Our view is that continuation payoff functions are made to depend on the current period's

⁸ This is assumed, for example, in Abreu (1986).

⁹ If firms wanted to break symmetry so as to get around the exchangeability restrictions, they surely have more direct means, such as having the managers introduce themselves to each other by name!

outcome in a repeated game in order to influence that period's behavior; this is the mechanism by which good behavior is induced. Thus, if current behavior is independent of those continuation payoff functions, then there is not an economic rationale for those continuation payoff functions to depend on the current period's outcome. As our objective is to characterize limits to collusion in symmetric equilibria, we think that restricting equilibria to be history-relevant is the natural assumption.¹⁰

Our first main finding is an impossibility result. Equilibria satisfying these properties cannot sustain collusion regardless of the discount factor. In stating this result, we define the set of equilibrium prices for the infinite horizon game to be those prices that arise with positive probability for some equilibrium.¹¹

Theorem 1. Assuming Assumptions 1–3, the set of strongly symmetric exchangeable history-relevant Nash equilibrium prices for the infinite horizon game coincides with the set of symmetric Nash equilibrium prices for the stage game.

Proof. Consider a strongly symmetric Nash equilibrium which, for period t , has all firms charge a price of $\tilde{p}^t(h^{t-1})$ and yields a continuation payoff of $v^{t+1}(h^{t-1}, \underline{q}^t)$. Firm i 's expected payoff from pricing at p_i^t is then

$$\sum_{\underline{q} \in \Delta} \psi(\underline{q}; \tilde{p}^t, \dots, p_i^t, \dots, \tilde{p}^t) [p_i^t \gamma_i(\underline{q}) + \delta v^{t+1}(h^{t-1}, \underline{q})],$$

where $\gamma_i(\underline{q})$ is defined to be the i th element of \underline{q} (in other words, firm i 's sales). By Assumption 1, a necessary condition for \tilde{p}^t to be an equilibrium price is

$$\begin{aligned} \sum_{\underline{q} \in \Delta} \left(\frac{\partial \psi(\underline{q}^t; \tilde{p}^t, \dots, \tilde{p}^t)}{\partial p_i^t} \right) \tilde{p}^t \gamma_i(\underline{q}) + \sum_{\underline{q}^t \in \Delta} \psi(\underline{q}^t; \tilde{p}^t, \dots, \tilde{p}^t) \gamma_i(\underline{q}) \\ + \sum_{\underline{q} \in \Delta} \left(\frac{\partial \psi(\underline{q}^t; \tilde{p}^t, \dots, \tilde{p}^t)}{\partial p_i^t} \right) \delta v^{t+1}(h^{t-1}, \underline{q}) = 0. \end{aligned}$$

Our method of proof is to show that the third term is zero, for, if that is the case, then \tilde{p}^t must satisfy

$$\sum_{\underline{q} \in \Delta} \left(\frac{\partial \psi(\underline{q}^t; \tilde{p}^t, \dots, \tilde{p}^t)}{\partial p_i^t} \right) \tilde{p}^t \gamma_i(\underline{q}) + \sum_{\underline{q} \in \Delta} \psi(\underline{q}^t; \tilde{p}^t, \dots, \tilde{p}^t) \gamma_i(\underline{q}) = 0,$$

which is the condition defining a symmetric Nash equilibrium for the stage game.¹² We then want to show that

$$\sum_{\underline{q} \in \Delta} \left(\frac{\partial \psi(\underline{q}; \tilde{p}^t, \dots, \tilde{p}^t)}{\partial p_i^t} \right) \delta v^{t+1}(h^{t-1}, \underline{q}) = 0. \quad (1)$$

¹⁰ Alternatively, Theorem 1 can be proven by replacing the requirement that equilibria be exchangeable and history relevant with that of recursive exchangeability. Define the state after period t to be $\underline{s}^t \equiv (v^t, \underline{q}^t)$, which is the vector of continuation payoffs at the beginning of period t and the vector of period t quantities. The equilibrium is recursively exchangeable if the continuation equilibrium strategies after a permutation of \underline{s}^t are an analogous permutation of the strategies after \underline{s}^t . This property is motivated by Abreu, Pearce, Stachetti (1990) in presuming that the vector of continuation payoffs act as sufficient state variables.

¹¹ For the case of a duopoly, Theorem 1 holds for strongly symmetric Nash equilibria, so that neither exchangeability nor history relevance is required. We have thus far been unable to prove that for the general case of n firms. The proof of the duopoly case is in the Appendix.

¹² Let us remind the reader that we are assuming the first-order condition is both necessary and sufficient for equilibrium. If it is not sufficient, then Theorem 1 as stated may not be true. Though the first-order conditions for the stage game and the infinitely repeated game still coincide, the second-order conditions need not. What is true, however, is that the set of strongly symmetric exchangeable history-relevant Nash equilibrium prices for the infinitely repeated game is a subset of the set of solutions to the first-order condition for the stage game.

The next step derives the key property used in proving Theorem 1. As Assumption 2 implies

$$\psi(\underline{q}; (p, \dots, p)) = \psi(\omega(\underline{q}; i, j); (p, \dots, p)), \quad \forall \underline{q}, p,$$

it follows that

$$\frac{\partial \psi(\underline{q}; (p, \dots, p))}{\partial p_i} = \frac{\partial \psi(\omega(\underline{q}; i, j); (p, \dots, p))}{\partial p_j}. \tag{2}$$

For an arbitrary quantity vector $\underline{q} \in \Delta$, define $\Gamma(\underline{q})$ to be the set of permutations of \underline{q} . Consider some arbitrary quantity vector \underline{q}^o . It then follows from (2) that

$$\sum_{\underline{q} \in \Gamma(\underline{q}^o)} \left(\frac{\partial \psi(\underline{q}; (p, \dots, p))}{\partial p_i} \right) = \sum_{\underline{q} \in \Gamma(\underline{q}^o)} \left(\frac{\partial \psi(\omega(\underline{q}; i, j); (p, \dots, p))}{\partial p_j} \right), \quad \forall i, j. \tag{3}$$

Because

$$\Gamma(\underline{q}^o) = \{\omega(\underline{q}; i, j) : \underline{q} \in \Gamma(\underline{q}^o)\},$$

(3) implies

$$\sum_{\underline{q} \in \Gamma(\underline{q}^o)} \left(\frac{\partial \psi(\underline{q}; (p, \dots, p))}{\partial p_i} \right) = \sum_{\underline{q} \in \Gamma(\underline{q}^o)} \left(\frac{\partial \psi(\underline{q}; (p, \dots, p))}{\partial p_j} \right), \quad \forall i, j. \tag{4}$$

By Assumption 3,

$$\sum_{\underline{q} \in \Gamma(\underline{q}^o)} \sum_{i=1}^n \left(\frac{\partial \psi(\underline{q}; (p, \dots, p))}{\partial p_i} \right) = 0.$$

Exchanging the summations, we have

$$\sum_{i=1}^n \sum_{\underline{q} \in \Gamma(\underline{q}^o)} \left(\frac{\partial \psi(\underline{q}; (p, \dots, p))}{\partial p_i} \right) = 0. \tag{5}$$

Using (4)–(5), we derive

$$\sum_{\underline{q} \in \Gamma(\underline{q}^o)} \left(\frac{\partial \psi(\underline{q}; (p, \dots, p))}{\partial p_i} \right) = 0, \quad \forall i. \tag{6}$$

It has then been shown that, given all firms charge a common price, the probability of $\Gamma(\underline{q}^o)$ —that is, all vectors that are permutable from \underline{q}^o —is locally independent of a firm’s price.

An equivalent representation of (1) is

$$\sum_{\underline{q}^o \in \Delta^*} \sum_{\underline{q} \in \Gamma(\underline{q}^o)} \left(\frac{\partial \psi(\underline{q}; \tilde{p}^t, \dots, \tilde{p}^t)}{\partial p_i} \right) \delta v^{t+1}(h^{t-1}, \underline{q}) = 0, \tag{7}$$

where Δ^* is the set of “basis vectors;” that is, no element of Δ^* is a permutation of any other element of Δ^* and any element of Δ is a permutation of some element of Δ^* . Consider $t = 1$ and let h^0 denote the null history. By exchangeability,

$$v^1(h^0, \underline{q}') = v^1(h^0, \underline{q}''), \quad \forall \underline{q}', \underline{q}'' \in \Gamma(\tilde{q}).$$

Therefore, (7) can be represented as

$$\sum_{\underline{q}^o \in \Delta^*} \delta v^2(h^0, \underline{q}^o) \sum_{\underline{q} \in \Gamma(\underline{q}^o)} \left(\frac{\partial \psi(\underline{q}; \tilde{p}^t, \dots, \tilde{p}^t)}{\partial p_i} \right) = 0,$$

which does indeed hold by (6). This implies $\tilde{p}^1(h^0) = \hat{p}$, where \hat{p} is a static Nash equilibrium price vector.

The proof is completed by strong induction. Suppose $\tilde{p}^\tau(h^0) = \hat{p}$, for all $\tau \leq t - 1$.¹³ Then, by the condition that equilibria be history relevant, $v^{t+1}(h^{t-1}, \underline{q}^t)$ is independent of h^{t-1} . Using exchangeability, (7) can be represented as

$$\sum_{\underline{q}^o \in \Delta^*} \delta v^{t+1}(h^{t-1}, \underline{q}^o) \sum_{\underline{q} \in \Gamma(\underline{q}^o)} \left(\frac{\partial \psi(\underline{q}; \tilde{p}^t, \dots, \tilde{p}^t)}{\partial p_i} \right) = 0,$$

which completes the proof. *Q.E.D.*

In explaining the intuition, it will be easiest to consider the duopoly setting. In thinking about punishment for perceived noncompliance in this setting, one would expect it to occur when market shares are sufficiently skewed: firm 1's sales are too high or too low. The former is consistent with firm 1 having undercut the collusive price and the latter with firm 2 having done so. Strong symmetry implies that the punishment entails identical behavior in the form of a price war. In such a situation, Theorem 1 shows that no collusion can be sustained.¹⁴ This result hinges on the fact that, when firm 1 sets a price marginally below the collusive price, it reduces the likelihood of having a low demand (say, $q'_1 < m/2$) and, at the same time, raises the probability of having a high demand (say, $m - q'_1$). The proof shows that the ensuing reduction in the probability of q'_1 is exactly equal to the rise in the probability of $m - q'_1$, so that the probability of q'_1 or $m - q'_1$ remains constant for a marginal change in price.¹⁵ Now suppose

$$v^{t+1}(h^{t-1}, (q_1, m - q_1)) = v^{t+1}(h^{t-1}, (m - q_1, q_1)), \quad \forall q_1,$$

so that the continuation payoffs depend only on the distribution of market share. It follows that the probability distribution over the continuation payoff is then unaffected by firm 1's price. Hence, an equilibrium price must maximize expected current profit because, at the margin, price has no effect on the expected continuation payoff. This implies the equilibrium price must be the same as that for a Nash equilibrium for the stage game.

A similar impossibility result can be obtained in environments such as are modeled in Blume and Heidhues (2003) and Skrzypacz and Hopenhayn (2004).¹⁶ These articles consider tacit collusion in auctions, where the bidders submit bids and the auctioneer chooses the best bid, announcing the winner but not the bids. One can think about those auctions as having one customer per period that performs closed-door price negotiations with the two potential sellers. In such a model, strongly symmetric equilibria also cannot support any collusion. The reason is more prosaic than in our model. At any point of the game, there are only two possible outcomes: firm 1 sells or firm 2 sells. That makes it impossible to detect a deviation if firms follow the same pricing strategy. In our model, however, asymmetric market shares can be used to detect deviations; it is just that the tests are too weak for small deviations. We will elaborate on this point later.

It is also interesting to ask why symmetric equilibria can be used to sustain collusion in Green and Porter (1984) but not in our model. As we show in the next section, it is not per se the difference in strategic variables (quantity versus price). It is instead the quality of information contained in the market signals. In Green and Porter, a deviation has a first-order effect on the probability of punishment. In our model, due to Assumption 1 and Assumption 3, a deviation to

¹³ It need not be the same static Nash equilibrium price vector in all periods.

¹⁴ As the proof of Theorem 1 used only the symmetry of the continuation payoffs and not their level, results would not change if we allowed for "money burning" activities that arbitrarily lowered $v(\cdot)$.

¹⁵ For the general n -firm case, a marginal change in price shifts probability mass over all quantity vectors that are permutations of each other but doesn't change the sum of their probabilities.

¹⁶ In Blume and Heidhues (2003) and Skrzypacz and Hopenhayn (2004), the bidders have private shocks that affect the efficient allocation of the object. This feature is not shared by our model, but it does not affect the result.

a lower price has no first-order effect on the probability of going to a punishment. That intuition becomes clearer in the next section, as we present an example which violates Assumption 1 and Assumption 3.

4. Robustness of the impossibility result

■ Here we explore the robustness of this impossibility result. In the first subsection, departures from Assumptions 1–3 are considered, while market demand is allowed to respond to firms' prices in the second subsection.

□ **Nondifferentiability of firm demand.** The proof of Theorem 1 is based on the property that the probability of a particular distribution of market shares is locally independent of a firm's price when firms charge a common price. Thus, skewness in market share is not made more likely when a firm undercuts the collusive price. One might imagine it is essential that ψ is continuously differentiable when firms charge a common price so that small price changes have small effects on the probability distribution over sales. That, however, proves not to be the case.¹⁷ Here, we present a simple model that assumes ψ is discontinuous at the point where firms have identical prices and show that collusion through symmetric punishments still need not be sustainable; we also show when it is sustainable.

Consider the following modest modification of a discrete version of the standard Bertrand price game with homogeneous goods. (By discrete, we mean that we are retaining our assumption of m units.) Assume that when firms do not charge the same price that all buyers go to the firm with the lowest price for sure. When instead, all prices are identical, then the probability distribution on firm demand is symmetric and let $y \in (0, 1)$ denote the probability at least two firms sell a positive amount (that is, no firm sells all m units).

This model is related to that in Tirole (1988) and Bagwell and Wolinsky (2002) in their specification of discontinuous expected residual demand when firms' prices are identical. Tirole (1988) assumes prices and quantities are private information. With homogeneous goods, market demand is stochastic and takes two possible states: it is positive (and nonincreasing in price) or it is zero for all prices. The inference problem is that, if a firm has zero sales, it isn't sure whether market demand was low (that is, zero) or market demand was positive and its rival cheated. Collusion is shown to be sustainable if the discount factor is sufficiently high. Whereas Tirole (1988) and Bagwell and Wolinsky (2002) allow market demand to be stochastic and firm demand to be deterministic (conditional on market demand), here we fix market demand and allow firm demand to be stochastic. This distinction proves unimportant as results are qualitatively similar.

Consider a strongly symmetric strategy profile that has all firms price at the collusive price p in period 1 and do so in period t as long as no firm had 100% of market sales in some past period. Otherwise, firms go to the static Nash equilibrium price of zero forever.¹⁸ Denoting the (rescaled) collusive payoff to be v , it is defined recursively by

$$v = (1 - \delta)(m/n)p + \delta yv.$$

From this, we get

$$v = \left(\frac{1 - \delta}{1 - \delta y} \right) p(m/n).$$

For the scheme with price p to be an equilibrium, the following incentive compatibility must hold:

$$v \geq (1 - \delta)mp \Leftrightarrow \left(\frac{1 - \delta}{1 - \delta y} \right) p(m/n) \geq (1 - \delta)mp \Leftrightarrow y \geq \frac{n - 1}{\delta n}.$$

¹⁷ We thank Phil Reny for conjecturing that differentiability is unnecessary.

¹⁸ It can be verified that, because the stage-game Nash equilibrium produces the minimax payoff and that only totally skewed market shares are possible when there is a deviation, the ensuing condition is both necessary and sufficient for collusion.

Thus, if $y \geq (n - 1)/\delta n$, then any collusive price (up to the maximum price that consumers are willing to pay) can be sustained by this strategy profile. If $y < (n - 1)/\delta n$, then only the static Nash equilibrium price is sustainable because the strategy profile contains the worst equilibrium punishment. Hence, regardless of the discount factor, collusion is not sustainable using symmetric punishments when $y < (n - 1)/n$.

The key issue here is whether firms can statistically detect deviations in the sense that the distribution of market shares under deviation and no deviation are different. Under the assumptions of Theorem 1, the likelihood of skewness in market share is unaffected when a firm marginally undercuts the collusive price. Thus, no statistical detection is possible. For the example of this section, this property doesn't hold, as the probability of a maximally skewed market share is one when a firm deviates and $1 - y$ when all firms are compliant. But that isn't sufficient for collusion. Though statistical detection follows from $y > 0$, collusion is sustainable only when $y \geq (n - 1)/\delta n$. The reason is that the probability of a false positive (that is, going to a punishment even though no firm deviated) is $1 - y$ and, if it is too high, the continuation payoff from colluding is too small, which makes it hard to provide incentives to collude. In order for collusion to be sustainable, the statistical test must be sufficiently precise so that punishment is sufficiently less likely when a firm colludes than when it cheats.

To see this more clearly, let us add some more structure by supposing that, when firms charge equal prices, each buyer randomly chooses between the n firms and their decisions are i.i.d. It follows that

$$y = 1 - \left(\frac{1}{n}\right)^{m-1}.$$

Because then $y \rightarrow 1$ as $m \rightarrow \infty$, collusion can be sustained with i.i.d. buyers as long as there are sufficiently many of them and $\delta > (n - 1)/n$. The probability of a false positive is $(1/n)^{m-1}$, so the statistical test is very precise when there are many buyers. This reduces the likelihood of wrongly triggering a punishment and thereby enhances the collusive payoff.¹⁹

□ **Elastic market demand.** Theorem 1 was proven under the extreme assumption that market demand is fixed and insensitive to firms' prices. Robustness is established by showing that very little collusion can be sustained when market demand is very inelastic.

Assume there is an upper bound on market demand of M units. A stochastic realization is comprised of total demand and an allocation of that demand, which can be represented as an element of

$$\Omega \equiv \{(m, \underline{q}) : m \in \{0, 1, \dots, M\}, \underline{q} \in \Delta(m)\},$$

where m is total sales, \underline{q} is the vector of firms' quantities and $\Delta(m)$ is as defined in Section 2, though now its dependence on m is made explicit. Letting $\xi : \Omega \times \mathbb{R}^n \rightarrow [0, 1]$, denote the probability function on (m, \underline{q}) given prices, the expected continuation payoff is

$$\sum_{m=0}^M \sum_{\underline{q} \in \Delta(m)} \xi(m, \underline{q}; \underline{p}) v(m, \underline{q}),$$

where the continuation payoff conditions on the most recent period's outcome, (m, \underline{q}) , and implicitly depends on the preceding history as well. Defining $\rho(m; \underline{p})$ as the marginal probability function on m and $\psi(\underline{q}; \underline{p}, m)$ as the conditional probability function on \underline{q} , then

$$\xi(m, \underline{q}; \underline{p}) = \rho(m; \underline{p}) \psi(\underline{q}; \underline{p}, m).$$

¹⁹ We also have a modified version of the Hotelling duopoly model, which makes similar points to those made in this subsection—see Harrington and Skrzypacz (2005). It entails a smooth expected demand function, but where the probability distribution on firm demand has a point of nondifferentiability (though is continuous everywhere). Symmetric punishments are still not capable of sustaining collusion when the kink is sufficiently small but collusion can be sustained when the kink is sufficiently large.

The expected continuation payoff can then be represented as

$$\sum_{m=0}^M \rho(m; \underline{p}) \sum_{q \in \Delta(m)} \psi(q; \underline{p}, m) v(m, \underline{q}).$$

Assume $\rho(\cdot; \underline{p})$ is differentiable in \underline{p} and is symmetric. Then

$$\frac{\partial \rho(m; p, \dots, p)}{\partial p_i} = \frac{\partial \rho(m; p, \dots, p)}{\partial p_j}, \quad \forall i, j, \forall p.$$

Finally, assume $\psi(q; \underline{p}, m)$ satisfies Assumptions 1–3, for all $m \in \{1, \dots, M\}$.

The maximization problem of firm i is

$$\max_{\underline{p}_i} \pi_i(\underline{p}) + \delta \left[\sum_{m=0}^M \rho(m; \underline{p}) \sum_{q \in \Delta(m)} \psi(q; \underline{p}, m) v(m, \underline{q}) \right],$$

where

$$\pi_i(\underline{p}) \equiv \sum_{m=0}^M \rho(m; \underline{p}) \sum_{q \in \Delta(m)} \psi(q; \underline{p}, m) p_i \gamma_i(q)$$

is the expected static profit of firm i (recall that $\gamma_i(q)$ is firm i 's sales according to vector q). The necessary first-order condition at a strongly symmetric Nash equilibrium is then

$$\begin{aligned} \frac{\partial \pi_i(p, \dots, p)}{\partial p_i} + \delta \left[\sum_{m=0}^M \left(\frac{\partial \rho(m; p, \dots, p)}{\partial p_i} \right) \sum_{q \in \Delta(m)} \psi(q; (p, \dots, p), m) v(m, \underline{q}) \right] \\ + \delta \left[\sum_{m=0}^M \rho(m; p, \dots, p) \sum_{q \in \Delta(m)} \left(\frac{\partial \psi(q; (p, \dots, p), m)}{\partial p_i} \right) v(m, \underline{q}) \right] = 0. \end{aligned} \tag{8}$$

By the method used in the proof of Theorem 1, one can show

$$\sum_{q \in \Delta(m)} \left(\frac{\partial \psi(q; (p, \dots, p), m)}{\partial p_i} \right) \delta v(q, m) = 0 \quad \forall m,$$

and, therefore, the third term in (8) is zero. Thus, (8) becomes

$$\frac{\partial \pi_i(p, \dots, p)}{\partial p_i} + \delta \left[\sum_{m=0}^M \left(\frac{\partial \rho(m; p, \dots, p)}{\partial p_i} \right) \sum_{q \in \Delta(m)} \psi(q; (p, \dots, p), m) v(m, \underline{q}) \right] = 0. \tag{9}$$

We conclude that a necessary condition to sustain collusion is that the second term in (9) is nonzero.

First, note that, if $\partial \rho(m; p, \dots, p) / \partial p_i = 0$, so that the distribution on market demand is independent of prices, (9) becomes the condition for a stage game equilibrium. Hence, collusion cannot be sustained as long as market demand is insensitive to prices, regardless of whether or not it is stochastic. When $\partial \rho(m; p, \dots, p) / \partial p_i$ is close to zero, then the set of values for p that satisfy (9) are, generically, close to the set of stage game symmetric equilibrium prices. We conclude that the collusive price is close to a stage game equilibrium price when market demand is sufficiently insensitive to firms' prices. In that sense, Theorem 1 is robust with respect to market demand being variable and sensitive to firms' prices.

An assumption of highly inelastic market demand is plausible for many of the price-fixing cartels mentioned, including those that arose in the markets for vitamins, lysine and citric acid. These products are largely being purchased by other firms as inputs—for example, vitamins and lysine are mixed with animal feed in the food-processing industry. As they make up a very small fraction of the unit cost of these purchasers, their demand is likely to be insensitive to price for a wide range of prices. Of course, the cartel members could set price high enough so as to induce a nonnegligible fall in market demand, but the size of the price increase required to make that happen may be of such a magnitude as to create suspicions among buyers that the input suppliers are colluding. As a result, the cartel might want to avoid such large price increases.²⁰ This may argue to the assumption that, over the relevant range of prices, market demand is highly inelastic.

Finally, note that, if we define

$$v(m) \equiv \sum_{q \in \Delta(m)} \psi(q; \underline{p}, m) v(m, q),$$

then (9) can be rewritten as

$$\frac{\partial \pi_i(p, \dots, p)}{\partial p_i} + \delta \sum_{m=0}^M \left(\frac{\partial \rho(m; p, \dots, p)}{\partial p_i} \right) v(m) = 0.$$

$v(m)$ is the expected continuation payoff conditional on the total market size and ignoring the division of market shares. This suggests that collusion may be supportable by conditioning on the size of market demand and not on the allocation of that demand across firms. An exploration of that conjecture we leave to future work. However, if $v(m)$ is constant, so firms expect the same continuation payoff regardless of the realized total market size, then it follows from $\sum_{m=0}^M (\partial \rho(m; p, \dots, p) / \partial p_i) = 0$ (which holds as the probabilities always sum up to 1) that again no collusion is sustainable.

5. Collusion with asymmetric punishments

■ We now show how collusion can be sustained with asymmetric punishments. In order to establish that collusion is sustainable against the most devious deviations, we allow a firm to set a unique price for each customer. Given n price vectors—each comprised of m prices (one for each buyer)—an outcome (which is generically denoted as \underline{r}) assigns a seller to each of the buyers and is an element of $\Theta \equiv \{1, 2, \dots, n\}^m$. Define $\varphi(\underline{r}; \underline{p}) : \Theta \times \mathbb{R}^{n \times m} \rightarrow [0, 1]$ to be the probability function over outcomes conditional on a price vector for each firm.

In designing a collusive strategy, we retain the focus on market-sharing schemes as have been used in a number of price-fixing cartels, such as the lysine cartel. If firms are in the collusive state, then, for each unit that a firm sells, it is to pay $z \geq 0$, which is then shared equally among the remaining members of the cartel. This implies that a firm's net transfer is increasing in its sales and its net transfer is positive (negative) when it sells more (less) than m/n units. The cartel starts in a collusive state and remains in it unless one of the firms fails to make the recommended transfer. If such a deviation is ever observed, the firms switch to static Nash equilibrium forever.

Given a realized outcome \underline{r} , let $\underline{\gamma}_i(\underline{r})$ be an $m \times 1$ vector of 0's and 1's specifying if a given customer buys from firm i . Let \underline{p}_i be an $m \times 1$ vector of prices set by firm i and $\underline{1}$ be an $m \times 1$ vector of ones. Assuming that future transfers will be honored, in the collusive state, firm i chooses its price vector to maximize

$$\sum_{\underline{r} \in \Theta} \varphi(\underline{r}; \underline{p}_1, \dots, \underline{p}_i, \dots, \underline{p}_n) \left[\underline{\gamma}_i(\underline{r})' \cdot (\underline{p}_i - z\underline{1}) + (1 - \underline{\gamma}_i(\underline{r}))' \left(\frac{1}{n-1} \right) z\underline{1} \right] + \delta v^{t+1}. \quad (10)$$

²⁰ For studies that model the effect of the prospect of detection on the cartel price path, see Harrington (2004, 2005) and Harrington and Chen (2006).

Note that, as long as all transfers are made, the continuation payoff (net of transfers) is independent of prices. Collecting the terms, we obtain

$$\sum_{r \in \Theta} \varphi(r; \underline{p}_1, \dots, \underline{p}_i, \dots, \underline{p}_n) \left[\gamma_i(r)' \cdot \left(\underline{p}_i - \frac{n}{n-1} z \underline{1} \right) \right] + \left(\frac{m}{n-1} \right) z + \delta v^{t+1}. \quad (11)$$

Therefore, the prices in such a scheme will be the same as prices in a one-shot game in which firms have a cost of $[n/(n-1)]z$ per unit. Assume that the equilibria of such games are well behaved, in the sense defined in Assumption 4, and that there is an equilibrium that is symmetric with respect to firms.

Assumption 4. The one-shot game with demand $\varphi(\cdot)$ and cost c per unit has, for all $c \geq 0$, a symmetric Nash equilibrium and there is a lower bound to the equilibrium price vector that is increasing and unbounded in c .

Let \hat{p} denote a stage-game symmetric equilibrium when $c = 0$. Verifying for a particular game if Assumption 4 is satisfied may be nontrivial because each firm has many ways to deviate with its price vector and, therefore, the sufficiency of first-order conditions will require careful investigation. This is relatively straightforward in some special cases, however. For example, if customer choices are independent—as in the generalization of the Cabral and Riordan (1994) model that we discussed in Section 2—one can treat the problem as m independent problems and hence focus on single-customer deviations. Finally, note that the condition of an increasing lower bound on the equilibrium price vector holds if all equilibrium prices weakly exceed unit cost.

Under Assumption 4 it is straightforward to prove that this collusive scheme can support arbitrarily high prices and profits if firms are sufficiently patient.

Theorem 2. If Assumption 4 holds, then, for any price vector $\underline{p} > \hat{p}$ (component-wise), there exists $\delta^* < 1$ such that for all $\delta \geq \delta^*$, there exists a subgame perfect equilibrium in which the cartel sets a price vector exceeding \underline{p} in every period.

Proof. First, assume that all future transfers will be paid in equilibrium. Firms then set prices to maximize (11). At $z = 0$, this results in the price vector \hat{p} . Now consider an arbitrary price vector, $\underline{p} > \hat{p}$. By the existence of an increasing unbounded lower-bound function, $\exists c^o > 0$ such that the lower bound exceeds \underline{p} . By choosing z so that $z[n/(n-1)] = c^o$, we have that an equilibrium price vector exceeds \underline{p} .

Second, after demand is realized, we must make sure that all firms have an incentive to pay its net transfer when it is positive. It is sufficient to verify the incentives of a firm that sells to all customers,

$$-mz + \delta v_z \geq \delta v_N,$$

where v_z is the sum of expected discounted payoffs along the proposed equilibrium path and v_N is analogously defined for the infinitely repeated static Nash equilibrium. As in the collusive state, all the firms set prices higher than \hat{p} , the difference in expected per-period profits in the collusive equilibrium and in the static Nash equilibrium is strictly positive regardless of δ . Therefore, as $\delta \rightarrow 1$, $\delta(v_z - v_N)$ increases to infinity. That establishes the existence of $\delta^* < 1$. *Q.E.D.*

We finish this section with two remarks. First, note that we have used monetary transfers after every period. If such transfers increase the probability of detection or are costly (so as to avoid detection by the antitrust authorities), one can use methods from Fudenberg, Levine, and Maskin (1994) to show that, for δ sufficiently close to one, similar collusive schemes are incentive compatible with monetary transfers replaced by transfers of continuation payoffs. Such transfers usually cause some loss of efficiency after asymmetric histories, but, for δ close to one, this loss is small.

Second, a class of collusive schemes was described that achieves very high prices and hence very high profits (unlike with symmetric price wars, in the above equilibrium, transfers

are balanced among the cartel members so the sum of profits across the firms depends only on the level of prices). Such high profits should not be taken too literally, of course, as the result follows from our extreme assumption that market demand is inelastic. While for some products this may be a reasonable approximation over some range of prices, the assumption becomes untenable for sufficiently high prices. Instead, we suggest interpreting the result as simply stating that asymmetric punishments can allow for a substantial degree of collusion.

In sum, collusion can be sustained by a punishment strategy in which firms with above-average sales compensate those firms with below-average sales. This is sustainable as long as firms are sufficiently patient and transfers can be made. In practice, transfers can be implemented by having a firm with excess sales buy output from a firm with insufficient sales at an inflated price. Several recent price-fixing cartels engaged in side payments of this sort, including the citric acid cartel of 1991–1995 (Connor, 2001), the graphite electrodes cartel of 1992–1997 (Levenstein, Suslow, and Oswald, 2004) and the vitamins cartel, in particular vitamins A and E over 1989–1999 (European Commission, 2003).

6. Related literature

■ Athey and Bagwell (2001) and Athey, Bagwell, and Sanchirico (2004) consider an infinitely repeated price game under perfect monitoring of both prices and sales but where each firm's cost is subject to i.i.d. shocks and, most important, is private information. They also show that asymmetric equilibria can achieve higher collusive payoffs than symmetric equilibria, though for fundamentally different reasons. In their model, sustaining high prices is not a problem and can be achieved with a sufficiently high discount factor. The problem lies in the inefficiency of a symmetric equilibrium in that buyers may not be efficiently allocated across sellers. Both types of problems—imperfect monitoring through sales and the efficient allocation of buyers—are relevant for real-life cartels and, furthermore, are intertwined. For example, splitting the market may be an effective method for circumventing the problem of monitoring secret price cuts, but it can increase the problem of allocating buyers efficiently across the sellers (and hence extracting higher surplus). We plan to study this trade-off in future research.²¹

The role of observability of prices has also been recently studied by Bergemann and Välimäki (2002). In their model, one buyer faces two sellers and all players are long lived and strategic. They show that, if prices are unobservable, the set of equilibrium payoffs is much smaller than if they are observable. By reacting strategically to the cartel, the buyer can “divide and conquer” the sellers. If one of them deviates to a lower price, the buyer has the option of not revealing it to the other seller and hence effectively coordinating at the same time with both sellers to the detriment of the cartel. The strategic behavior of buyers is largely neglected in the literature (for a general mechanism design approach, see Abdulkadiroglu and Chung, 2003) but is potentially important for cartels. In our model, the buyers are either short lived (so they don't internalize the effect their behavior has on future prices) or do not realize they are facing a cartel. However, the asymmetric equilibrium in Section 5 would seem robust to strategic deviations by the buyer, as future prices do not depend on the current market shares and incentives are provided directly through transfers. Such immunity to a buyer's strategic behavior may be another advantage to providing incentives via transfers rather than through the threat of a price war. This topic also deserves further investigation.

Cole and Kocherlakota (2005) explore conditions under which strongly symmetric equilibria cannot achieve payoffs higher than static Nash equilibrium payoffs. They show that this arises with strategies having finite history dependence, though infinite-memory equilibria can sustain higher payoffs. The economics are very different from our article. With Cole and Kocherlakota, the problem is one of insufficient incentives for the cartel to carry out punishments, while in our model, the problem lies in a deficiency of monitoring.

²¹ With many buyers per period, the problem of efficient allocation in a cartel with communication can be solved, as shown by Pesendorfer (2000).

Finally, the challenge of achieving efficient detection is a theme in Abreu, Milgrom, and Pearce (1991) and Sannikov and Skrzypacz (2005a). In these articles, as the frequency of moves increases, detection of deviations becomes more and more difficult. In fact, Sannikov and Skrzypacz (2005b) present a class of games in which, as the frequency of moves increases, the provision of incentives by symmetric punishments becomes increasingly costly and this results in the set of strongly symmetric equilibrium payoffs collapsing to the set of static Nash payoffs. This is in spite of the fact that deviation is, in principle, detectable (unlike in our model).

7. Concluding remarks

■ A common perception of collusive schemes is built around the idea of price wars: the cartel members are recommended to set high prices and if deviation is detected (actual or perceived through a noisy signal), then the firms punish each other by setting low prices. The actual practice of many well-documented price-fixing cartels tells a very different story. It is quite common to employ more complicated schemes involving history-dependent transfers among members. Our analysis suggests that imperfect monitoring in those markets may be the key reason why they did so. For a natural class of multiunit demand functions, symmetric price wars are incapable of sustaining any collusion regardless of how patient firms are. It may then be necessary for cartels to deploy punishments that discriminate between the firms that sold too much and those that sold too little. We put forth one punishment scheme that works and entails those who supplied over quota providing a transfer to those who sold under quota. This is a practice consistent with several recent price-fixing cartels.

There are a number of directions one can go from here. Though we have shown that collusion by virtue of symmetric punishments is minimal when market demand is highly inelastic, this leaves open the question of how effective price wars can be in sustaining collusion when market demand is responsive to price. Our initial investigation suggests conditioning punishments on total sales might be effective, though this is only a conjecture. It would be particularly interesting to characterize optimal equilibria that condition on both total sales and firms' market shares and see how that information is used.

This article has explored collusion when firms can only condition on market shares. Though this is one scheme that has been used in practice, another approach is to allocate customers rather than market share. In that case, each firm is entitled to sell to a set of customers and violation of that condition results in a punishment. Hence, the public information is not simply a firm's sales but also who bought from them. From a monitoring perspective, it would seem customer allocation is superior because a firm isn't even supposed to bid for another firm's customers, in which case a sale is a perfect signal of a deviation. This raises the question of why a market-share allocation scheme is used at all. More broadly, it is an open question as to what determines the type of allocation mechanism that cartels use.²²

Appendix

■ Theorem A1 and its proof follow.

Theorem A1. Assuming Assumptions 1–3 and $n = 2$, the set of strongly symmetric Nash equilibrium prices for the infinite horizon game coincides with the set of symmetric Nash equilibrium outcomes for the stage game.

Proof. First note that the derivation of (6) in the proof of Theorem 1 only requires strong symmetry. For $n = 2$, (6) takes the form

$$\frac{\partial \phi(b; p, p)}{\partial p_1} + \frac{\partial \phi(m - b; p, p)}{\partial p_1} = 0, \quad (\text{A1})$$

where ϕ is defined at the start of Section 4. If m is even, (A1) implies $\partial \phi(m/2; p, p)/\partial p_1 = 0$.

²² As mentioned in a previous section, allocating customers can be inefficient if there are firm-specific shocks, for example, like in Athey and Bagwell (2001).

Using (A1), (1) takes the following form when $n = 2$:

$$\sum_{b=0}^{(m/2)-1} \left(\frac{\partial \phi(b; p, p)}{\partial p_1} \right) [v(b) - v(m-b)] + \underbrace{\frac{\partial \phi(m/2; p, p)}{\partial p_1}}_0 v \left(\frac{m}{2} \right) = 0$$

when m is even, and takes the form:

$$\sum_{b=0}^{(m-1)/2} \left(\frac{\partial \phi(b; p, p)}{\partial p_1} \right) [v(b) - v(m-b)] = 0$$

when m is odd. To prove these conditions hold, subtract the first-order condition for firm 2 from the first-order condition for firm 1 to obtain

$$\begin{aligned} & \sum_{b=0}^m \left(\frac{\partial \phi(b; p, p)}{\partial p_1} - \frac{\partial \phi(m-b; p, p)}{\partial p_2} \right) pb + \delta \sum_{b=0}^m \left[\left(\frac{\partial \phi(b; p, p)}{\partial p_1} \right) v(b) - \left(\frac{\partial \phi(m-b; p, p)}{\partial p_2} \right) v(m-b) \right] \\ & + \sum_{b=0}^m [\phi(b; p, p) - \phi(m-b; p, p)]b = 0. \end{aligned}$$

The first summation is zero by (2) and the third summation is zero by Assumption 2 and that the probabilities add up to 1. Using (4) in the second summation, we derive

$$\sum_{b=0}^m \left(\frac{\partial \phi(b; p, p)}{\partial p_1} \right) [v(b) - v(m-b)] = 0.$$

This can be used to complete the proof; for example, if m is odd. Using again (6), it can be rewritten as

$$\sum_{b=0}^{(m-1)/2} \left(\frac{\partial \phi(b; p, p)}{\partial p_1} \right) [v(b) - v(m-b)] = 0,$$

which establishes the claim. *Q.E.D.*

References

- ABDULKADIROGLU, A. AND CHUNG, K.-S. "Auction Design with Tacit Collusion." Mineo, Columbia University, June 2003.
- ABREU, D. "Extremal Equilibria of Oligopolistic Supergames." *Journal of Economic Theory*, Vol. 39 (1986), pp. 191–225.
- , MILGROM, P., AND PEARCE, D. "Information and Timing in Repeated Partnerships." *Econometrica*, Vol. 59 (1991), pp. 1713–1733.
- , PEARCE, D., AND STACCHETTI, E. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica*, Vol. 58 (1990), pp. 1041–1063.
- ARBAULT, F. ET AL. "Commission Adopts Eight New Decisions Imposing Fines on Hard-Core Cartels." *EC Competitions Policy Newsletter*, Vol. 1 (February 2002), pp. 29–43.
- ATHEY, S. AND BAGWELL, K. "Optimal Collusion with Private Information." *RAND Journal of Economics*, Vol. 32 (2001), pp. 428–465.
- , ———, AND SANCHIRICO, C. "Collusion and Price Rigidity." *Review of Economic Studies*, Vol. 71 (2004), pp. 317–349.
- BAGWELL, K. AND WOLINSKY, A. "Game Theory and Industrial Organization." In R.J. Aumann and S. Hart, eds., *Handbook of Game Theory with Economics Applications*, Vol. 3. Amsterdam: Elsevier, 2002.
- BERGEMANN, D. AND VÄLIMÄKI, J. "Strategic Buyers and Privately Observed Prices." *Journal of Economic Theory*, Vol. 105 (2002), pp. 469–482.
- BLUME, A. AND HEIDHUES, P. "Modeling Tacit Collusion in Auctions." Mineo, University of Pittsburgh, February 2003.
- CABRAL, L. AND RIORDAN, M. "The Learning Curve, Market Dominance, and Predatory Pricing." *Econometrica*, Vol. 62 (1994), pp. 1115–1140.
- CAMPBELL, C., RAY, G., AND MUHANNA, W.A. "Search and Collusion in Electronic Markets." *Management Science*, Vol. 51 (2005), pp. 497–507.
- COLE, H.L. AND KOCHERLAKOTA, N.R. "Finite Memory and Imperfect Monitoring." *Games and Economic Behavior*, Vol. 53 (2005), pp. 59–72.
- CONNOR, J.M. *Global Price Fixing: Our Customers Are the Enemy*. Boston: Kluwer Academic Publishers, 2001.

- EUROPEAN COMMISSION. "Commission Fines Five Companies in Sodium Gluconate Cartel." Press release, March 19, 2002, http://europa.eu/index_en.htm.
- . "Case COMP/E-1/37.512—Vitamins." *Official Journal of the European Communities*, October 1, 2003.
- FUDENBERG, D., LEVINE, D., AND MASKIN, E. "The Folk Theorem with Imperfect Public Information." *Econometrica*, Vol. 62 (1994), pp. 997–1039.
- GREEN, E.J. AND PORTER, R.H. "Noncooperative Collusion Under Imperfect Price Information." *Econometrica*, Vol. 52 (1984), pp. 87–100.
- HARRINGTON, JR., J.E. "Cartel Pricing Dynamics in the Presence of an Antitrust Authority." *RAND Journal of Economics*, Vol. 35 (2004), pp. 652–674.
- . "Optimal Cartel Pricing in the Presence of an Antitrust Authority." *International Economic Review*, Vol. 46 (2005), 145–169.
- AND CHEN, J. "Cartel Pricing Dynamics with Cost Variability and Endogenous Buyer Detection." *International Journal of Industrial Organization*, Vol. 24 (2006), pp. 1185–1212.
- AND SKRZYPACZ, A. "Collusion Under Monitoring of Sales." Research Paper no. 1885, Stanford Graduate School of Business, March 2005.
- LEVENSTEIN, M.C., SUSLOW, V.Y., AND OSWALD, L.J. "Contemporary International Cartels and Developing Countries: Economic Effects and Implications for Competition Policy." *Antitrust Law Journal*, Vol. 71 (2004), pp. 801–852.
- PESENDORFER, M. "A Study of Collusion in First-Price Auctions." *Review of Economic Studies*, Vol. 67 (2000), pp. 381–411.
- SANNIKOV, Y. AND SKRZYPACZ, A. "Impossibility of Collusion Under Imperfect Monitoring with Flexible Production." Research Paper no. 1887, Stanford Graduate School of Business, May 2005a (*American Economic Review*, forthcoming).
- AND ———. "The Role of Information in Repeated Games with Frequent Actions." Mineo, Stanford University, 2005b.
- SKRZYPACZ, A. AND HOPENHAYN, H. "Tacit Collusion in Repeated Auctions." *Journal of Economic Theory*, Vol. 114 (2004), pp. 153–169.
- STIGLER, G.J. "A Theory of Oligopoly." *Journal of Political Economy*, Vol. 72 (1964), pp. 44–61.
- TIOLE, J. *The Theory of Industrial Organization*. Cambridge, Mass.: MIT Press, 1988.