Collusive Equilibrium in Cournot Oligopolies with Unknown Costs

Subir K. Chakrabarti*

Current Version: July 8, 2009

Department of Economics
Indiana University-Purdue University Indianapolis(IUPUI)
425 University Blvd.
Indianapolis, IN 46202

Abstract: This paper studies collusive equilibria in infinite horizon repeated games with discounting in which the firms play a Cournot game each period with unknown costs. We find that there are pooled Perfect Bayesian equilibria that are collusive in which a firm plays exactly the same strategy irrespective of its realized cost. There are also separating equilibria with signaling as well as with communication, in which the firms produce the optimal incentive compatible quantity vector. In the separating equilibria with signaling, the firms play the strictly separating Bayesian Nash equilibrium in period 1 and then from period 2 onwards produce the optimal incentive compatible collusive quantity vector. In the separating equilibrium with communication the firms produce the optimal incentive compatible quantity vector from period 1 onwards after an initial communication phase. We show that among these equilibria, the separating equilibrium with communication gives the highest expected discounted sum of joint profits.

Keywords: Oligopoly, Collusion, Repeated Games, Private Information, Folk Theorem, Pooling Equilibrium, Separating Equilibrium, Perfect Bayesian Equilibrium.

JEL Classification Numbers: Primary D2, D8, L1. Secondary L5

*The paper has benefited from presentations at the 2007 North American Summer Meetings of the Econometric Society, the Fall 2006 Midwest Theory Conference at Purdue University and at department seminars at Ohio State University, Virginia Polytechnic and State University, University of Iowa and Southern Methodist University. The author would like to thank Hans Haller, Dan Levin, Massimo Morelli, James Peck, B. Ravikumar, Jennifer Reinganum, Santanu Roy, Huanxing Yang, Lixin Ye, and Itzhak Zilcha for useful comments. I am also grateful to two anonymous referees for helpful comments and suggestions. The usual disclaimer applies.

1 Introduction

It is well known that firms may not fully know the costs of other firms. For instance, a firm may receive a privately observed technology shock that permanently changes its cost structure. Other firms may have some idea of some aspects of the technology shock and about the costs of their rivals, but may not have all the relevant information. While it is true that publicly held firms publish reports about their revenues and profits so that one might argue that the costs of the firms are never really private information as these have to be revealed on a periodic basis, this is certainly not true of partnerships and firms that are privately held, as they do not report costs or even profits in their annual statements[1]. Even with publicly held firms it can be argued that not everything about the costs of the firms are made explicit. Indeed, for such firms one may even pose the question a little differently. Since the firms report their costs, is there an incentive for the firms to misreport their true cost. Would this then work against some form of tacit collusion? Or would the nature of tacit collusion itself be modified in order to discourage the misrepresentation of the costs by firms. It thus seems that the question of whether collusion is possible in case the firms have private information about costs is an important one.

It is well known that in oligopoly markets firms can either compete by setting prices, as in the Bertrand model, or firms can compete in quantities, as in the Cournot model. With homogeneous products, competition in the Bertrand model can drive prices down to the point where firms earn zero profit. In the Bertrand model firms thus have very strong incentives to collude, that is, to set a common price that maximizes their joint profits. Even in the Cournot oligopoly, although the firms can usually expect to make positive profits (as in the Cournot equilibrium), the firms can make much higher profits when they collude and jointly set outputs to maximize joint profits. Given these strong incentives to collude, important questions arise as to whether collusion can be sustained as an equilibrium and the nature of the collusive behavior. A fairly extensive literature in industrial organization has dealt with this question and detailed discussions of issues related to this can be found for example in Tirole [1993], or Motta [2004]. Most

[1]For example law firms have traditionally never declared their profits or their cost. Barred from being public firms, they have no incentive to disclose their financial results (page 54 of the *Economist*, August 9, 2003). While one may observe that law firms are not quantity setting firms, it illustrates the fact that firms do not necessarily have to report their cost of production.

of this earlier literature has discussed collusion when there is complete information. We know that in the case of complete information, collusion can be sustained as part of an equilibrium behavior by using the threat of reverting back to the stage game equilibrium in both the Bertrand and Cournot setting. However, when the firms have private information about costs, differences emerge between the Bertrand and the Cournot models, as the stage game equilibrium is no longer well defined (the stage game equilibrium is sensitive to the costs of the firms). Recent works by Athey and Bagwell [2001] and Athey and Bagwell [2008] have studied collusive behavior in the Bertrand setting when the firms have private information about costs. What we do here is study the nature of collusive behavior in the Cournot setting when the cost of each firm is privately known only to the firm.

A major difference between the Bertrand model and the Cournot model arises from the fact that while the profit functions of the firms are continuous in quantities, they are not continuous in prices. In the Cournot model, the profit functions of the firms are continuous in quantities, which are the decision variables of the firms in the Cournot model, so that one can adjust the profit levels of the firms by appropriately adjusting the quantities produced by the firms. In the Bertrand model with homogeneous products this is not the case. A firm can increase its profits disproportionately by cutting prices. Therefore, in the intensely competitive setting of the Bertrand model with homogeneous products, collusive agreements may need to be agreements on both the price level as well as the market shares. As we shall see, this leads to important differences in the analysis of the Bertrand model and the Cournot model. It is worth noting, however, that in case there is product differentiation, the profit functions of the firms are continuous in prices even in the Bertrand model. The situation in this case is thus more like that of the Cournot model.

In what follows we discuss the nature of collusive equilibrium when the firms receive a privately observed technology shock that changes the cost of a firm for a sufficiently long duration of time. We do this in the context in which the firms play a Cournot game in each period. Since the costs of the firms are now private information, one can view this situation as one in which the firms play a Cournot quantity-setting game repeatedly over many periods without knowing the costs of the other firms. Since the firms do not know each other's costs, firms cannot play the Cournot-Nash quantities of the stage game, and therefore, the method of using trigger-price strategies with its reversion to Cournot-Nash equilibrium strategies cannot be used to deter deviations from the collusive output. This can be done only if there is some method of eliciting information about the true costs of the firms. However, playing the Cournot-Nash equilibrium quantities on the basis of what is reported will not work as firms with high costs will typically want to report low costs. Thus our setting is different from the one, for instance, in Green and Porter [1984], in which the firms do not observe each others output, but because they know each others payoff function, the firms can decide on the Cournot-Nash equilibrium output, and therefore, can respond to public signals by playing the Cournot-Nash output if the public signal indicates that with high probability there has been a deviation. Similarly, in Abreu, Pearce and Stachetti [1986] in which the firms again respond to public signals, the firms know each others payoff and thus can compute the continuation payoffs as a function of the public signals.

The works that are closest to the present work are the papers by Athey and Bagwell [2001] and [2008] as well as the paper by Athey, Bagwell and Sanchirico [2004][2]. All of these papers deal with the issue of private information and do so within an explicitly dynamic framework with multiple periods. In these papers the firms play an infinite-horizon version of the Bertrand price-setting game, in which the prices are perfectly observed, but the cost of the firms are subject to i.i.d. shocks every period. They show that the firms can collude at the efficient price by agreeing on appropriate splits of the market share. The high cost firm is willing to give up market share because it expects higher expected profit in the future. The result they obtain, however, depends crucially on the condition that the firm that is a high cost firm today could receive a technology "shock" in the future that would make it a low cost firm. In the case of the model presented here, because the high cost firm will remain a high cost firm, there seems to be little incentive for the high cost firm to reveal that information. The setting of Athey and Bagwell [2008] is closest to what we have here. In Athey and Bagwell [2008] there is persistence in the costs and in the limit there can be perfect persistence. This part of their analysis thus covers the case in which the firms receive a privately observed technology shock that persists for a long duration of time as in this paper. Their result is that if the distribution of costs is log concave and the firms are sufficiently patient then the optimal collusive scheme entails price rigidity; firms set the same price and share the

[2] A paper by LaCasse [1999] addresses a similar question as in this paper with two cost types but does not explicitly discuss collusive behavior.

market equally, regardless of their respective costs. Productive efficiency can be achieved under some circumstances, but such equilibria are not optimal. It should, however, be noted that the firms play a Bertrand price-setting game in each period. Further, and this could be the major element that drives their results, the demand side is given by a unit mass of identical consumers with a fixed reservation price r, such that $r > \bar{\theta}$, where $\bar{\theta}$ is the highest possible cost. This is different from having a downward sloping demand schedule as the firms know that the optimally collusive price is r irrespective of the privately observed costs of the firms. In case the demand is given by the usual downward sloping demand curve, the optimally collusive price depends on the realized costs of the firms, that is, if $\hat{p}(\theta)$ is the optimal collusive price when the realized costs of the firms are θ, then $\hat{p}(\bar{\theta}) > \hat{p}(\underline{\theta})$ if $\bar{\theta} > \underline{\theta}$.

The situation that we analyze here is one in which the firms not only do not know each others costs, but also face a downward sloping demand curve so the firms also do not know the optimal collusive quantity and price. As a result the firms may not know the kind of collusive agreements that can be implemented. Further, even if the firms agree on the collusive output vector it is not clear how firms can be deterred from deviating. One possible way to deter deviations would be to use punishment phases in which a deviating firm is punished for all possible cost configurations[3]. Some form of minimaxing may work in this case. But one needs to find minimaxing strategies that would work for all possible realizations of the costs of the firms. For the oligopoly games this difficulty is overcome by using a minimaxing strategy that punishes a deviator for all possible realizations of its cost.

Some recent work have analyzed the role that private information plays in an explicitly dynamic framework. Thus Cole and Kocherlakota [2001] analyze a class of games with hidden actions and hidden states. Kennan [2001] examines repeated bargaining in which the buyer's valuation is determined by a two-state Markov chain and this valuation is private information to the buyer. It should be noted that in both Cole and Kocherlakota [2001] as well as Kennan [2001], the private information is generated every period by a random shock to the state that is privately observed by some of the players but not by all. Among other works that are also closely related to the literature

[3]In Athey and Bagwell [2008], given that r is the optimal collusive price for all realized costs of the firms, an equilibrium can be played in which the firms either set the price p equal to r or in case of a deviation set the price equal to $\bar{\theta}$. In the case of the quantity setting firms in a market with the usual downward sloping demand curve, the optimal collusive quantity vector or price is not known.

on collusion in infinite horizon Oligopoly games with private information is the one by Hanazono and Yang [2007]. This work analyzes collusive behavior when the firms receive private signals about independently and identically distributed demand shocks. It thus analyzes situations in which the firms have private information about the demand side of the market. Another work that also looks at collusion within the framework of infinitely repeated games when there is private information about the demand side of the market is that of H. Gerlach [2007].

In this paper we show that the infinite horizon game in which the firms have private information about their costs, and play the oligopoly stage game repeatedly over an infinite horizon, has a fairly large equilibrium set. In these equilibria, the quantity choices could either be independent of the information about the costs, or be completely determined by them. In section 4 we analyze pooling equilibrium in which the quantity vector is independent of the realized costs as are the minimaxing strategies used to punish deviators. We call this kind of an equilibrium *pooling equilibrium* as the firms do not have to ever reveal their costs and produce the same output irrespective of their cost. In section 5 we analyze a completely different kind of equilibrium in which the strategies of the firms depend on the realized cost. We call this kind of collusive equilibrium *strictly separating*, as different realizations of cost lead to different strategies. In this type of equilibrium the quantity vector depends critically on the realized cost, as does the subsequent play of the game. We show that one can find strictly separating equilibrium in which the firms produce a quantity vector that maximizes the joint profit of the firms subject to some incentive constraints. These optimal incentive compatible collusive outputs are produced after the firms signal their cost in period 1 by playing a strictly separating Bayesian Nash equilibrium. We also show that if the firms communicate prior to producing their output, then there is a separating equilibrium in which the firms play their optimal incentive compatible quantity vector from period 1 onwards. In section 6 we compare the expected discounted sum of the joint profits of the firms from the optimal pooling equilibrium to those from the optimal strictly separating equilibrium with signaling and the optimal separating equilibrium with communication. We find that the expected joint profits from the separating equilibrium with communication is at least as large as those from either the optimal separating equilibrium with signaling or the optimal pooling equilibrium. One also notes that whether the firms are in a pooling equilibrium or in a separating equilibrium, the equilibrium prices and outputs are stable over time, indicating the classic sort of price rigidity that is common under collusion.

2 The Oligopoly with unknown costs

There are n firms. The marginal cost of firm i is some constant c_i. This is known only to firm i and is thus private information to the firm. The other firms know that c_i takes finitely many values $\{c_1, c_2, \cdots, c_{k_i}\}$. We will denote the set of possible costs of a firm by C. There is a common probability distribution over the set of possible marginal costs of the firms given by μ. Thus μ is a probability distribution over $\mathcal{C} = C \times C \times \cdots \times C$.[4] We will call an element $\mathbf{c}$ of $\mathcal{C}$ a *cost profile.* Given the private information of firm i that its marginal cost is c_i, the conditional distribution about the cost function of the other firms is given by $\mu(\mathbf{c}|c_i)$. Therefore the belief of firm i about the distribution of the costs of the other firms, given that its own cost is c_i, is $\mu(\mathbf{c}|c_i)$.

The firms all produce the same identical product and the inverse demand function $p(\mathbf{q})$ satisfies $p'(\mathbf{q}) < 0$. Each firm observes the output vector $\mathbf{q} = (q_1, \cdots, q_n)$ every period and the resulting market price $p(\mathbf{q}) = p(\sum_i q_i)$.

Firm i observes its own profit, which is a function of its cost c_i and this is given by

$$\pi_i(\mathbf{q}, c_i) = p(\mathbf{q})q_i - c_i q_i.$$

Firm i does not know the profit of the other firms and only knows that the cost of the other firms are distributed according to the conditional distribution $\mu(\mathbf{c}|c_i)$.

We make the following assumptions about the demand.

Assumption 1 *There is a $\bar{q}$ such that $p(\bar{q}) = 0$ and $p(0) < \infty$.*

and

Assumption 2 *There is a quantity vector $\mathring{\mathbf{q}}$ such that*

$$p(\mathring{\mathbf{q}}). \mathring{q}_i - \bar{c}\, \mathring{q}_i > 0$$

for all $i = 1, \cdots, n$,

[4]Even though we have denoted the set of possible costs of the firms as being the same, the set of actual costs of the firms can be different. If a particular firm's cost never take certain values in C, then that is reflected in the fact that the probability μ of those cost profiles is zero. Thus, if we have two firms, and one firm's cost can take the values $c_H > c_L > 0$ and the other firm's cost is given by c_H, then the probability distribution μ will satisfy the condition that $\mu(c_H, c_L) = \mu(c_L, c_L) = 0$.

where $\bar{c}$ is the highest possible marginal cost of a firm. This last assumption guarantees that there is sufficient demand in the market for a firm to operate at profit even if all the firms find that they have the highest possible marginal cost. With these assumptions the following holds.

Proposition 1 *Every firm's profit can be pushed down to* 0 *by the other firms independently of the firm's type.*

Proof: Consider the quantity vector $\mathbf{q}_{-i} = \{q_j\}_{j \neq i}$ such that $\sum_{j \neq i} q_j = \bar{q}$. Then $p(\mathbf{q}) \leq 0$ so that $\pi_i(\mathbf{q}, c_i) = 0 - c_i q_i \leq 0$ for any $c_i \in C$. ■

3 The infinite horizon game

The infinite horizon game is generated by repeating the incomplete information stage game of the oligopoly with unknown costs over an infinite horizon. Before the repeated game is played, the firms get to know their own marginal costs. The firms thus have private information about their costs which for each firm is randomly drawn from C according to the joint distribution μ on $\mathcal{C}$. After the firms get to know their marginal costs, the firms play the quantity setting Cournot game repeated over an infinite horizon. Thus, the infinite horizon sequential game is a game with imperfect information. The strategy of a firm i in this sequential game is a sequence $\{\sigma_{it}\}_{t=1}^{\infty}$ such that

$$\sigma_{it} : H_{t-1} \times C \to [0, \bar{q}]$$

where H_{t-1} is the set of histories of the game until period $t-1$ and an $h_{t-1} \in H_{t-1}$ is given by $h_{t-1} = \{\mathbf{q}^1, \mathbf{q}^2, \cdots, \mathbf{q}^{t-1}\}$, where $\mathbf{q}^t$ is the quantity vector in time period t. That is, h_{t-1} is a history that consists of quantity choices of the firms until period $t-1$. The action chosen by a firm i in period t thus depends on the past history of quantities chosen by the firms and the cost c_i of the firm. A strategy combination will be denoted by $\sigma = \{\sigma_t\}_{t=1}^{\infty}$, where $\sigma_t = \{\sigma_{1t}, \cdots, \sigma_{nt}\}$.

The payoff of a firm i when a strategy combination σ is used is given by

$$\pi_i^{\infty}(\sigma, c_i, \mathbf{c_{-i}}) = \sum_{t=1}^{\infty} \delta^{t-1} \pi_i(\sigma_{1t}(h_{t-1}, c_1) + \cdots + \sigma_{nt}(h_{t-1}, c_n), c_i)$$

Thus the payoff of firm i in the sequential game is the discounted sum of the single period profits. Since the profits depend on the true marginal costs and the quantity choices that are observed by the firms, the firms see their true payoffs, but do not know the payoffs of the other firms. Given a history h_t, the discounted sum of payoffs of player i from time period $t+1$ onwards, when the strategy combination σ is used, is

$$\pi_i^\infty(\sigma|_{h_t,c_i,\mathbf{c_{-i}}}) = \sum_{\ell=t}^{\infty} \delta^{\ell-1}\pi_{i\ell}(\sigma_{1\ell}(h_{\ell-1},c_1) + \cdots + \sigma_{n\ell}(h_{\ell-1},c_n),c_i).$$

The expected payoff, after a history h_t, given that the cost of firm i is c_i, is

$$\sum_{c\in C} \mu(\mathbf{c}|h_t,c_i)\pi_i^\infty(\sigma|_{h_t,c_i,\mathbf{c_{-i}}})$$

where $\mu(\mathbf{c}|h_t,c_i)$ is the conditional distribution over the costs given the history h_t and the private information c_i of firm i. It denotes the belief of firm i about the costs, after it has observed the history h_t and the cost c_i.

The equilibrium concept that we use here is that of a Perfect Bayesian equilibrium since we discuss equilibrium in a game with incomplete and hence imperfect information[5]. A **Perfect Bayesian equilibrium** is a strategy combination that continues to be an optimal strategy for every player given any history and the updated beliefs of the players given that history, when the beliefs are updated using Bayes' rule[6]. Thus, if σ is the strategy combination, then if $\mu(h_t|\sigma,\mathbf{c_{-i}},c_i)$ is the probability of the history h_t given σ and the cost profile $(\mathbf{c_{-i}},c_i)$ of the firms, then the probability of the cost profile being $\mathbf{c} = (\mathbf{c_{-i}},c_i)$ is

$$\mu(\mathbf{c}|h_t,c_i,\sigma) = \frac{\mu(h_t|\sigma,\mathbf{c_{-i}},c_i)\mu(\mathbf{c}|h_{t-1},\sigma,c_i)}{\sum_{\mathbf{c'}\in\mathcal{C}} \mu(h_t|\sigma,\mathbf{c'_{-i}},c_i)\mu(\mathbf{c'}|h_{t-1},\sigma,c_i)}. \quad (1)$$

Thus, in every period every firm updates its belief about the cost profile of the other firms using Bayes' rule, or equivalently, by the conditional probability of $\mathbf{c} \in \mathcal{C}$ given (h_t,σ,c_i), on the basis of the history it has observed.

Definition 1 *Given the strategy combination $\sigma^\star = (\sigma_1^\star,\cdots,\sigma_n^\star)$, the assessment $(\sigma^\star,\mu^\star)$ is a* **Perfect Bayesian** *equilibrium of the infinite horizon game if*
(i) $\mu^\star(.)$ is a system of beliefs that is determined by $\sigma^\star$ according to the rule given in (1),

[5] Strictly speaking, this is a game with incomplete information that can be viewed as a game with imperfect information, in which a chance move at the beginning of the game cannot be perfectly observed by all the players

[6] The marginal distributions of μ may not have full support as already observed in footnote 4. Therefore, beliefs are updated using Bayes' rule to the extent possible.

and

(ii) for every player i, for every time period t and for every history h_t of actions up to time period t,

$$\sum_{\mathbf{c}\in\mathcal{C}} \mu^\star(\mathbf{c}|h_t, c_i, \sigma^\star)\pi_i^\infty(\sigma^\star|_{h_t,c_i,\mathbf{c_{-i}}}) \geq \sum_{\mathbf{c}\in\mathcal{C}} \mu^\star(\mathbf{c}|h_t, c_i, (\sigma_i, \sigma^\star_{-i}))\pi_i^\infty((\sigma_i, \sigma^\star_{-i})|_{h_t,c_i,\mathbf{c_{-i}}})$$

for every $(\sigma_i|_{h_t,c_i})$.

Note that the strategies are conditioned on the private information of the firm and the history of the actions. For a detailed discussion of Perfect Bayesian equilibrium and Sequential Equilibrium one may refer to Fudenberg and Tirole [1991] and for a discussion of Sequential equilibrium see Kreps and Wilson [1982]. A Perfect Bayesian equilibrium will be called a **pooling equilibrium** if the strategies of the firms with different costs are the same. That is in a pooling equilibrium a firm plays the same strategy irrespective of its realized cost. A Perfect Bayesian equilibrium will be called a **separating equilibrium** if the strategy of a firm depends on its cost. It will be called a **strictly separating equilibrium** if the equilibrium strategy of a firm varies strictly with its cost.

4 Pooling Equilibrium, the Folk Theorem and Collusion

In this section we show that every quantity vector that gives positive profits is associated with a pooling perfect Bayesian equilibrium. This set of quantity vectors is the set of quantity vectors in which firms make positive profits even if their realized cost is the highest possible $\bar{c}$. We will show that given the cost profile of the firms, which will remain undisclosed in these pooling equilibria, any quantity vector in this set is the quantity vector of a perfect Bayesian equilibrium and a firm plays the same strategy irrespective of its realized cost. One can view this result as a version of the folk theorem for these repeated games with incomplete information.

The following gives the definition of a pooling perfect Bayesian equilibrium.

Definition 2 *A perfect Bayesian equilibrium $(\sigma^\star, \mu^\star)$ is a pooling perfect Bayesian equilibrium if for every firm i, and for any history h_{t-1} up to period t, and for every $c_i, c_i' \in C$, we have $\sigma^\star_{it}(h_{t-1}, c_i) = \sigma^\star_{it}(h_{t-1}, c_i')$.*

Thus in a pooling perfect Bayesian equilibrium, a firm will produce the same stream of output irrespective of its marginal cost. Let

$$Q = \{\hat{\mathbf{q}} = (\hat{q}_1, \cdots, \hat{q}_n) : \pi_i(\sum_i \hat{q}_i, c_i) > 0 \text{ for all } c_i \text{ and for all } i\}.$$

Thus Q is the set of quantity vectors that would allow every firm to make positive profits even if the realized marginal costs of the firms are the highest possible for each type. One can view Q as the set of quantity vectors that all firms could agree to produce if they did not have any idea of the cost of the other firms, since every firm would be able to make positive profits, whatever be their cost, at any of the quantity vectors in Q.

Proposition 2 *$Q \neq \emptyset$. Further, any quantity vector $\mathbf{q} = (q_1, \cdots, q_n)$ that satisfies the condition $p(\mathbf{q}) > \bar{c}$ is in Q.*

Proof: By assumption 2, there is a quantity vector $\mathring{\mathbf{q}}$ such that

$$p(\mathring{\mathbf{q}}).\, \mathring{q}_i - \bar{c}\, \mathring{q}_i > 0.$$

Since

$$p(\mathring{\mathbf{q}}).\, \mathring{q}_i - c_i.\, \mathring{q}_i \geq p(\mathring{\mathbf{q}}).\, \mathring{q}_i - \bar{c}\, \mathring{q}_i > 0$$

it follows that $\pi_i(\mathring{\mathbf{q}}, c_i) > 0$ for all c_i and for all $i = 1, \cdots, n$. Thus $\mathring{\mathbf{q}} \in Q$. Hence, $Q \neq \emptyset$.

The observation that any quantity vector $\mathbf{q} >> 0$ that satisfies the condition $p(\mathbf{q}) > \bar{c}$ is in Q, follows from noting that for all possible cost profiles, $\mathbf{q}$ will satisfy the condition

$$p(\mathbf{q})q_i - c_i q_i \geq p(\mathbf{q})q_i - \bar{c}q_i > 0.$$

This shows that $\mathbf{q}$ is in Q. ∎

At a quantity vector in the set Q all the firms will make positive profits, even if the marginal costs of all the firms are at the highest possible level $\bar{c}$, so that all the firms will make positive profits at all possible marginal costs.

Lemma 1 *For every $\hat{q}$ in Q there is a vector $\epsilon^i = (\epsilon^i_1, \cdots, \epsilon^i_n) >> 0$ and a quantity vector $\hat{\mathbf{q}}^{i,\epsilon}$ such that*

$$\pi_i^{i,\epsilon}(c_i) = \pi_i(\hat{\mathbf{q}}^{i,\epsilon}, c_i) \leq \pi_i(\hat{\mathbf{q}}, c_i) - \epsilon^i_i, \text{ and } \pi_j^{i,\epsilon}(c_j) = \pi_j(\hat{\mathbf{q}}^{i,\epsilon}, c_j) \geq \pi_j(\hat{\mathbf{q}}, c_j) + \epsilon^i_j$$

for all c_i, c_j, $j \neq i$.

Proof: Given a $\hat{\mathbf{q}}$ in Q, define the quantity vector $\hat{\mathbf{q}}^{i,\epsilon}$ as

$$\hat{q}_j^{i,\epsilon} = \hat{q}_j + \frac{\epsilon}{n-1}, \text{ and } \hat{q}_i^{i,\epsilon} = \hat{q}_i - \epsilon.$$

Then, clearly

$$\sum_{j=1}^{n} \hat{q}_j^{i,\epsilon} = \sum_{j=1}^{n} \hat{q}_j$$

and

$$\begin{aligned} \pi_i^{i,\epsilon}(c_i) &= [p(\hat{\mathbf{q}}) - c_i]\hat{q}_i^{i,\epsilon} \\ &= [p(\hat{\mathbf{q}}) - c_i](\hat{q}_i - \epsilon) = \pi_i(\hat{\mathbf{q}}, c_i) - \epsilon[p(\hat{\mathbf{q}}) - c_i]. \end{aligned} \quad (2)$$

Also

$$\begin{aligned} \pi_j^{i,\epsilon}(c_j) &= [p(\hat{\mathbf{q}}) - c_j]\hat{q}_j^{i,\epsilon} \\ &= [p(\hat{\mathbf{q}}) - c_j](\hat{q}_j + \frac{\epsilon}{n-1}) = \pi_j(\hat{\mathbf{q}}, c_j) + \frac{\epsilon}{n-1}[p(\hat{\mathbf{q}}) - c_j]. \end{aligned} \quad (3)$$

Define $\epsilon_i^i = \min_{c_i \in C}\{\epsilon[p(\hat{\mathbf{q}}) - c_i]\}$ and $\epsilon_j^i = \min_{c_j \in C}\{\frac{\epsilon}{n-1}[p(\hat{\mathbf{q}}) - c_j]\}$. The result then follows from equations (2) and (3). ∎

The next result shows that every quantity vector in Q is the quantity vector of a pooling equilibrium.

Theorem 1 (Pooling Equilibrium and the Folk Theorem) *For every output vector $\hat{\mathbf{q}}$ in Q, there is a $\delta(\hat{\mathbf{q}}) < 1$ such that for all $\delta \geq \delta(\hat{\mathbf{q}})$, $\hat{\mathbf{q}}$ is the output vector of a pooling Perfect Bayesian equilibrium.*

Proof: The claim is that the strategy combination $(\sigma_1^\star, \cdots, \sigma_n^\star)$ described below is a pooling equilibrium.

(i)$\sigma_{it}^\star(h_{t-1}) = \hat{q}_i$ if the past history has been the output vector $\hat{q}$.

(ii) If a firm i produces $q_i \neq \hat{q}_i$ in any period t and all the other firms had produced $\hat{q}_j$ in all previous periods, then all firms $j \neq i$ produce the output $\frac{1}{n-1}\bar{q}$ for a length of time T_i, from time $t+1$ onwards. This is a phase I punishment strategy.

(iii) If there are no deviations during a phase I punishment by any of the firms $j \neq i$, then after the length of time T_i, the firms produce the output vector $\{\hat{q}_j^{i,\epsilon}\}_{j=1}^n$ as defined in lemma 1.

(iv) If a firm $j \neq i$ deviates during a phase I punishment, then firms $\ell \neq j$ produce the

output $\frac{1}{n-1}\bar{q}$ for a length of time T_j. Such a punishment is a phase II punishment.
(v) After a phase II punishment the firms produce the quantity vector $\{\hat{q}_i^{j,\epsilon}\}_{i=1}^n$.
(vi) If a firm $\ell \neq j$ deviates from the quantity vector $\{\hat{q}_i^{j,\epsilon}\}_{i=1}^n$, then the other firms play the phase I punishment for a length of time T_ℓ, and then produce the quantity vector $\{\hat{q}_i^{\ell,\epsilon}\}_{i=1}^n$.
(vii) If firm j deviates from the quantity vector $\{\hat{q}_i^{j,\epsilon}\}_{i=1}^n$, then the other firms produce the output $\bar{q}$ as in the phase I punishment, but now for a period T_j^1, and then produce the quantity vector $\hat{\mathbf{q}}^{j,\epsilon}$ in the periods following that.

We now proceed to show that the strategy profile $\sigma^\star$ is an equilibrium irrespective of the realized cost of the firm.

Let M_i be the maximum "gain" a firm can make by deviating in any period irrespective of its cost. If firm i deviates in any period then its maximum payoff in the subsequent periods, if it has cost c_i, is

$$M_i + \delta^{T_i} \sum_{\nu=1}^{\infty} \delta^{\nu-1} \pi_i^{i,\epsilon}(\hat{\mathbf{q}}, c_i)$$

as for a length of time T_i firm $i's$ profit is zero or less every period (see proposition 1). If firm i does not deviate, its payoff in the subsequent periods is

$$\sum_{\nu=1}^{\infty} \delta^{\nu-1} \pi_i(\hat{\mathbf{q}}, c_i).$$

Therefore, from lemma 1 and the construction of the strategy profile, firm i does not gain from a deviation if

$$\sum_{\nu=1}^{\infty} \delta^{\nu-1} \pi_i(\hat{\mathbf{q}}, c_i) \geq M_i + \delta^{T_i} \sum_{\nu=1}^{\infty} \delta^{\nu-1} [\pi_i(\hat{\mathbf{q}}, c_i) - \epsilon_i^i]. \tag{4}$$

That is,

$$\frac{1-\delta^{T_i}}{1-\delta} \pi_i(\hat{\mathbf{q}}, c_i) \geq M_i - \frac{\delta^{T_i}}{1-\delta} \epsilon_i^i. \tag{5}$$

Now note that in (5) the expression $\frac{1-\delta^{T_i}}{1-\delta} \to T_i$ as $\delta \to 1$, therefore, there is a $\delta_1 : 0 < \delta_1 < 1$ and T_i sufficiently large for which equation (4) is satisfied for all $c_i \in C$. Choose T_i so that

$$\frac{1-\delta^{T_i}}{1-\delta} \min_{c_i} \pi_i(\hat{\mathbf{q}}, c_i) \geq M_i - \frac{\delta^{T_i}}{1-\delta} \epsilon_i^i. \tag{6}$$

Thus, phase I punishments can deter a firm from deviating irrespective of its cost.

Now consider a deviation made by a firm j during a phase I punishment. Let M_j be the maximum payoff firm j can get in a single period, and L_j the maximum loss every period that firm j sustains during a phase I punishment. Then firm $j's$ payoff after deviating is less than or equal to

$$M_j + \delta^{T_j} \sum_{\nu=1}^{\infty} \delta^{\nu-1}[\pi_j(\hat{\mathbf{q}}, c_j) - \epsilon_j^j],$$

and if a firm j does not deviate, the payoff in the subsequent periods is:

$$\delta^{T_i - t} \sum_{\nu=1}^{\infty} \delta^{\nu-1}[\pi_j(\hat{\mathbf{q}}, c_j) + \epsilon_j^i] - \sum_{\nu=1}^{T_i - t} \delta^{\nu-1} L_j.$$

Therefore, firm j does not gain by deviating during a phase I punishment if

$$\delta^{T_i - t} \sum_{\nu=1}^{\infty} \delta^{\nu-1}[\pi_j(\hat{\mathbf{q}}, c_j) + \epsilon_j^i] - \sum_{\nu=1}^{T_i - t} L_j \geq M_j + \delta^{T_j} \sum_{\nu=1}^{\infty} \delta^{\nu-1}[\pi_j(\hat{\mathbf{q}}, c_j) - \epsilon_j^j]. \tag{7}$$

This reduces to

$$\begin{aligned} \delta^{T_i - t} \frac{1 - \delta^{T_j - T_i + t}}{1 - \delta} \pi_j(\hat{\mathbf{q}}, c_j) &\geq M_j - \frac{\delta^{T_i - t}}{1 - \delta} \epsilon_j^i - \epsilon_j^j \frac{\delta^{T_j}}{1 - \delta} + L_j \frac{1 - \delta^{T_i - t}}{1 - \delta} \\ &\geq M_j - \frac{\delta^{T_i - t}}{1 - \delta} \epsilon_j^i - \epsilon_j^j \frac{\delta^{T_j}}{1 - \delta} + L_j \frac{1 - \delta^{T_i}}{1 - \delta}. \end{aligned} \tag{8}$$

In equation (8), for a given T_i and a given T_j, as $\delta \to 1$, the expression

$$\delta^{T_i - t} \frac{1 - \delta^{T_j - T_i + t}}{1 - \delta}$$

goes to $T_j + t - T_i$ and the expressions $\frac{\delta^{T_j}}{1-\delta}$ and $\frac{\delta^{T_i - t}}{1-\delta}$ both go to ∞. Further, the expression $\frac{1-\delta^{T_i}}{1-\delta}$ goes to T_i. Hence, there is a $\delta_{j2} : 0 < \delta_{j2} < 1$ such that equation (8) holds for all $\delta > \delta_{j2}$ and for all $c_j \in C$. Again choose T_j such that an expression like (8) holds for all c_j. That is, choose T_j such that

$$\delta^{T_i - t} \frac{1 - \delta^{T_j - T_i + t}}{1 - \delta} \min_{c_j} \pi_j(\hat{q}, c_j) \geq M_j - \frac{\delta^{T_i - t}}{1 - \delta} \epsilon_j^i - \epsilon_j^j \frac{\delta^{T_j}}{1 - \delta} + L_j \frac{1 - \delta^{T_i}}{1 - \delta}. \tag{9}$$

Thus, for all such discount factors $\delta > \delta_{j2}$ firm j does not gain from deviating during a phase II punishment irrespective of its cost.

Next, suppose that firm $\ell \neq j$ deviates during a phase II punishment, then the other firms play a phase II punishment for firm ℓ and then produce the output vector

$\{\hat{q}_i^{\ell,\epsilon}\}_{i=1}^n$. As in the case of the deviation during a phase II punishment by firm j, firm ℓ cannot gain from deviating if the discount factor $\delta_{\ell 2}$ is sufficiently high.

Finally, consider a deviation by firm j from the quantity vector $\mathbf{q}^{j,\epsilon}$. Then, if firm j deviates, firms $\ell \neq j$ each produce the output $\frac{1}{n-1}\bar{q}$ for a length of time T_j^1 and then again produce the quantity vector $\mathbf{q}^{j,\epsilon}$. Firm $j's$ discounted payoff after deviating is then less than or equal to

$$M_j + \delta^{T_j^1} \sum_{\nu=1}^{\infty} \delta^{\nu-1}[\pi_j(\hat{\mathbf{q}}, c_j) - \epsilon_j^j].$$

Therefore firm j does not gain by deviating from the output vector $\mathbf{q}^{j,\epsilon}$ if

$$\sum_{\nu=1}^{\infty} \delta^{\nu-1}[\pi_j(\hat{\mathbf{q}}, c_j) - \epsilon_j^j] \geq M_j + \delta^{T_j^1} \sum_{\nu=1}^{\infty} \delta^{\nu-1}[\pi_j(\hat{\mathbf{q}}, c_j) - \epsilon_j^j].$$

That is, if

$$\frac{1-\delta^{T_j^1}}{1-\delta}[\pi_j(\hat{\mathbf{q}}, c_j) - \epsilon_j^j] \geq M_j.$$

Choose T_j^1 to be large enough so that

$$T_j^1.[\pi_j(\hat{\mathbf{q}}, c_j) - \epsilon_j^j] > M_j.$$

Then since $\frac{1-\delta^{T_j^1}}{1-\delta} \to T_j^1$ as $\delta \to 1$, there is a δ_2 sufficiently large, such that for all $\delta \geq \delta_2$, we have

$$\frac{1-\delta^{T_j^1}}{1-\delta}[\pi_j(\hat{\mathbf{q}}, c_j) - \epsilon_j^j] \geq M_j.$$

Choose $\bar{\delta} = \max\{\delta_1, \{\delta_{j2}\}_{j\neq i}, \delta_2\}$, then if $\delta > \bar{\delta}$, firm i does not gain from deviating. Hence for $\delta > \bar{\delta}$, no firm can gain by deviating. This establishes the claim.

We now show that the strategy combination $\sigma^\star$ is a Perfect Bayesian equilibrium. We first note that since $(\sigma^\star|h_t, c_i) = (\sigma^\star|h_t, c_i')$ for every c_i, c_i', therefore from (1) we have $\mu(\mathbf{c}|h_t, c_i, \sigma^\star) = \mu(\mathbf{c}|h_{t-1}, c_i, \sigma^\star)$ for all $t \geq 1$. Hence, $\mu(\mathbf{c}|h_t, c_i, \sigma^\star) = \mu(\mathbf{c}|c_i)$ so that

$$\begin{aligned}\sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c}|h_t, c_i, \sigma^\star)\pi_i^\infty(\sigma^\star|h_t, c_i) &= \pi_i^\infty(\sigma^\star|h_t, c_i)\sum_{\mathbf{c}\in\mathcal{C}}\mu(\mathbf{c}|c_i) \\ &= \pi_i^\infty(\sigma^\star|h_t, c_i). \end{aligned} \quad (10)$$

Since we have already shown that for any h_t, c_i and for all strategy σ_i of firm i, $\pi_i^\infty(\sigma^\star|h_t, c_i) \geq \pi_i^\infty((\sigma_i, \sigma_{-i}^\star)|h_t, c_i)$, it now follows from (10) that

$$\begin{aligned}\sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c}|h_t, c_i, \sigma^\star)\pi_i^\infty(\sigma^\star|h_t, c_i) &= \pi_i^\infty(\sigma^\star|h_t, c_i) \\ &\geq \sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c}|h_t, c_i, (\sigma_i, \sigma_{-i}^\star))\pi_i^\infty((\sigma_i, \sigma_{-i}^\star)|h_t, c_i). \end{aligned} \quad (11)$$

But (11) then shows that $\sigma^\star$ is a Perfect Bayesian equilibrium. It is by construction a pooling equilibrium. This thus concludes the proof. ▪

In these pooling equilibrium, the firms interact without revealing their costs to their rivals, and there is simply no change in beliefs from the ones held initially. The result holds because of the way the players can be minimaxed without the other players' knowing their costs by simply working with the quantity vector that the firms have agreed to produce. The strength of the result in theorem 1 comes from the observation that one can get the result to work for any distribution of the marginal costs, as long as the support of the distribution has a reasonable upper bound, or alternatively, when the demand in the market remains sufficiently high.

While theorem 1 gives us some idea of the set of quantity vectors that would be produced in a pooling equilibrium it does not tell us anything about whether any of these quantity vectors are collusive. *Firms* **Collude** *when they agree either tacitly or explicitly to produce output levels that lead to high profits for the firms.* **Optimal Collusion** *occurs when the firms produce output levels that maximize joint profits.* The next result shows that under some mild conditions all possible optimal collusive quantity vectors are in Q. Let $\hat{\mathbf{q}}(\underline{\mathbf{c}})$ be a quantity vector that maximizes the joint profits of the firms when the marginal costs of all the firms are $\underline{c}$, the lowest possible, and $p(\hat{\mathbf{q}}(\underline{\mathbf{c}}))$ be the price that maximizes the joint profit in that case. The result that follows shows that if $p(\hat{\mathbf{q}}(\underline{\mathbf{c}})) > \bar{c}$ then all the collusive quantity vectors are in Q.

Proposition 3 *Let $\hat{\mathbf{q}}(\underline{\mathbf{c}})$ denote a quantity vector that maximizes the joint profit of the firms when the marginal costs of the firms are all given by $\underline{c}$. Then if*

$$p(\hat{\mathbf{q}}(\underline{\mathbf{c}})) - \bar{c} > 0 \tag{12}$$

the quantity vectors $\mathbf{q}$ that satisfy

$$\pi_i(\mathbf{q}, c_i) \geq \pi_i(\hat{\mathbf{q}}(\underline{\mathbf{c}}), c_i)$$

for every $i = 1, \cdots, n$, and for all $c_i \in C$, are in Q.

Proof: We observe that if the condition in (12) holds, then for the cost profile $\mathbf{c} = (c_1, \cdots, c_n)$ and the quantity vector $\mathbf{q}$,

$$\pi_i(\mathbf{q}|c_i) = p(\mathbf{q})q_i - c_i q_i \geq p(\hat{\mathbf{q}}(\underline{\mathbf{c}}))\hat{q}_i(\underline{\mathbf{c}}) - c_i\hat{q}_i(\underline{\mathbf{c}}) \geq p(\hat{\mathbf{q}}(\underline{\mathbf{c}}))\hat{q}_i(\underline{\mathbf{c}}) - \bar{c}\hat{q}_i(\underline{\mathbf{c}}) > 0$$

for all $i = 1, \cdots, n$. But this shows that the output vector $\mathbf{q}$ is in Q. ∎

Theorem 1 and proposition 3 show that the set of quantity vectors that would be produced in a Perfect Bayesian equilibrium is large and consists of almost any quantity vector at which the firms can make positive profits. It also seems that collusion is a distinct possibility. However, even though the set Q may contain collusive quantity vectors, an important issue here is whether the firms can agree on the optimal collusive quantities given their true marginal costs. Agreeing on which quantity vector is the optimal collusive quantity vector can be problematic in these cases. As the following example shows, the difficulty may lie in determining the true optimal collusive output vector of the firms when the marginal costs are not known.

Example 1 *The sub-optimality of pooling equilibrium relative to the case of complete information.*[7]

Consider an oligopoly game in which

$$p(Q) = 10 - Q$$

and the marginal cost of firm 1 is either $c_1 = 1$ or $c_1 = 2$ with

$$Prob(c_1 = 1) = Prob(c_1 = 2) = 0.5$$

and similarly, the marginal cost of firm 2 is either $c_2 = 1$ or $c_2 = 2$ with

$$Prob(c_2 = 1) = Prob(c_2 = 2) = 0.5.$$

It can be checked that the joint profit maximizing quantity vector when the marginal cost of the firms are $c_1 = 1$ and $c_2 = 1$ is $\hat{q}_1 = \hat{q}_2 = 2.25$[8]. The profit of each firm in this case is 10.125.

In case the marginal costs of the firms are $c_1 = c_2 = 2$, then the joint profit maximizing quantities are $\hat{q}_1 = \hat{q}_2 = 2$ and the profit of each firm is 8. In case the

[7] Notice that the example satisfies the condition of proposition 3.

[8] In case the marginal costs of the firms are $c_1 = c_2 = 1$, the price under optimal collusion is 5.5. If the marginal costs of the firms are $c_1 = c_2 = 2$, then the optimal collusive quantities are $\hat{q}_1 = \hat{q}_2 = 2$ and the optimal collusive price is 6. This shows that these oligopoly games are different from the class of oligopoly games analyzed in Athey and Bagwell [2008]. In their setting, in addition to the oligopoly game being a Bertrand price setting game, the optimal collusive price is independent of the marginal costs of the firms.

marginal cost of one firm is 2 and of the other firm is 1, then the joint profits would be maximized if the firm with the lower cost produced the entire output of 4.5; the output of the other firm is then zero.

Now note that the output that maximizes the joint expected profit, and which would be an optimal collusive output in a pooling equilibrium as described in theorem 1, is the output $Q = q_1 + q_2$ that maximizes

$$\begin{aligned} & 0.25[p(q_1+q_2)(q_1+q_2) - c_L q_1 - c_L q_2] + 0.25[p(q_1+q_2)(q_1+q_2) - c_L q_1 - c_H q_2] \\ + \; & 0.25[p(q_1+q_2)(q_1+q_2) - c_H q_1 - c_L q_2] + 0.25[p(q_1+q_2)(q_1+q_2) - c_H q_1 - c_H q_2] \\ = \; & p(Q)Q - 0.5c_L Q - 0.5c_H Q. \end{aligned}$$

The first order condition is

$$p'(Q)Q + p(Q) = 0.5c_L + 0.5c_H = 1.5$$

so that $\hat{Q} = 4.25$ and $\hat{q}_1 = \hat{q}_2 = 2.125$ are the outputs of the individual firms. In comparing this output level with the joint profit maximizing output levels for each of the four possible realizations of the marginal costs, we find that the quantity vector that maximizes the expected joint profit differs from the optimal collusive output levels for each of the four possible combinations of the realized costs. *Thus the quantity vector of the optimal pooling equilibrium would be uniformly sub-optimal after the firms receive their private information.* ∎

The observation made in the preceding example leads one to ask whether it would be possible for the firms to share the information about their marginal costs. However, notice that if the output is determined so as to maximize the joint profit according to the costs reported by the firms then the joint profit maximizing outputs of the firms when the marginal costs are reported as $c_1 = c_2 = 1$ are $\hat{q}_1 = \hat{q}_2 = 2.25$. However, when the marginal cost of firm 2 is $c_2 = 2$ instead, and firm 2 reveals that its marginal cost $c_2 = 2$, then firm 1 may want to produce the entire profit maximizing output of 4.5 units. Firm 2 would in this case be left with either a very small output or zero. If on the other hand firm 2 claimed that its marginal cost was indeed only 1, rather than 2, its profit would be $(10-4.5) \times 2.25 - 2 \times 2.25 = 7.875$, considerably more than what it could hope to get if it revealed its true marginal cost. This shows why firms may find it difficult to know what to make of the costs reported by the other firms; a high-cost firm would always want to claim that it is a low-cost firm. This incentive problem raises the question as to whether there is a way of inferring the true marginal costs of the firms, and whether this can be done as part of an equilibrium in which the firms collude optimally given some incentive constraints.

5 Separating Equilibrium

In the previous section we characterized the set of pooling equilibrium in the repeated Cournot Oligopoly. In such a pooling equilibrium the firms can produce the quantity vector that maximizes the expected joint profit of the firms. Example 1 shows that this quantity vector could be very different from the quantity vectors that maximize joint profits given the actual realized costs of the firms. Firms could therefore do better if they shared the information about the costs. But this too leads to problems as firms would not want to reveal their true cost. We take up this issue here and investigate whether there are equilibria in which firms reveal their true cost under some incentive constraints. We show that under some relatively mild conditions there is a strictly separating equilibrium in which the firms reveal their true costs and produce quantity vectors that maximize expected joint profits under incentive constraints. A concept that plays an important role in the construction of these strictly separating equilibrium points is the Bayesian Nash equilibrium of the single-period game.

Definition 3 *An n-vector of quantity choices* $\{q_1^\star(.), \cdots, q_n^\star(.)\}$ *is a* **Bayesian Nash Equilibrium** *of the game if for each firm i and cost* c_i *of firm i, we have*

$$q_i^\star(c_i) \in \operatorname{argmax} \sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c}|c_i)[\{p(\sum_{j\neq i} q_j^\star(c_j) + q_i) - c_i q_i\}]. \tag{13}$$

We will say a Bayesian Nash equilibrium is a **strictly separating** Bayesian Nash equilibrium if $q_i^\star(c_i) \neq q_i^\star(c_i')$, whenever $c_i \neq c_i'$. The next definition defines a strictly separating perfect Bayesian equilibrium of the infinite-horizon game.

Definition 4 *A Perfect Bayesian equilibrium* $(\sigma^\star, \mu^\star)$ *will be said to be a* **strictly separating** *Perfect Bayesian equilibrium if it is a Perfect Bayesian equilibrium and, in addition, satisfies the condition that for all i and any pair* (c_i, c_i') *from C*

$$\sigma_i^\star|_{h_t,c_i} \neq \sigma_i^\star|_{h_t,c_i'},$$

whenever $c_i \neq c_i'$.

Thus in a strictly separating equilibrium a firm's strategy is conditioned on its cost, and if the costs differ, then so does the strategy. In a strictly separating Bayesian equilibrium, firms with different costs will play differently, and thus will tend to reveal information about their costs, as the other firms would be able to infer the cost of a firm from observing its output choices. We will show that under some mild conditions, the infinite horizon game has a strictly separating equilibrium, and this equilibrium is generated by first playing a strictly separating Bayesian equilibrium of the single-period game of incomplete information. The first lemma shows that the Cournot quantity setting game with private information always has a Bayesian Nash Equilibrium.

Lemma 2 *If the inverse demand function $p : [0, \bar{q}] \to \mathbb{R}_+$ is concave, then the single-period Cournot quantity-setting game has a Bayesian Nash Equilibrium*[9].

Proof: For every firm i with cost c_i consider the correspondence[10]

$$B_i^{c_i} : \Pi_{(n-1)|C|}[0, \bar{q}] \to [0, \bar{q}]$$

defined as

$$B_i^{c_i}(\{\{q_j(c_j)\}_{c_j \in C}\}_{j \neq i}) = \{q_i^\star(c_i) | q_i^\star(c_i) \in \text{argmax} \textstyle\sum_{\mathbf{c} \in \mathcal{C}} \mu(\mathbf{c}|c_i)\pi_i(\mathbf{c}|c_i)\}$$

where $\pi_i(\mathbf{c}|c_i) = [p(\sum_{j \neq i} q_j(c_j)) + q_i) - c_i]q_i$ is the profit of firm i given a cost profile $\mathbf{c} \in \mathcal{C}$. Since the inverse demand function $p(.)$ is concave it follows that $\pi_i(\mathbf{c}|c_i)$ is concave, and hence, $\sum_{\mathbf{c} \in \mathcal{C}} \mu(\mathbf{c}|c_i)\pi_i(\mathbf{c}|c_i)$ is concave in q_i. Thus it can be checked that $B_i^{c_i}$ is a nonempty-valued, compact-valued and convex-valued correspondence that is upper semicontinuous.

Therefore, since C is finite, the correspondence

$$B : \Pi_{n|C|}[0, \bar{q}] \to \Pi_{n|C|}[0, \bar{q}]$$

defined as

$$B(\{q_i(c_i)\}_{c_i \in C}, \{\{q_j(c_j)\}_{c_j \in C}\}_{j \neq i}) = \Pi_{j=1}^n [\Pi_{c_j \in C} B_j^{c_j}(\{\{q_i(c_i)\}_{c_i \in C}\}_{i \neq j})]$$

[9]Even though the proof is fairly standard we give a version of the proof here, as the Bayesian Nash equilibrium plays an important role in the subsequent analysis. It also provides a useful background to the next result and explains the regulartity assumption of the concavity of $p(.)$.

[10]$|C|$ is the cardinality of the set C, the set from which the marginal costs of the firms are drawn.

is a nonempty-valued, compact-valued and convex-valued correspondence that is upper semicontinuous and defined on a compact subset of an Euclidean space. Thus the correspondence B has a fixed point $(\{\{q_i^\star(c_i)\}_{c_i \in C}\}_{i=1}^n)$. It is clear that this is a Bayesian Nash Equilibrium of the single-period game. ■

The next lemma shows that the Bayesian Nash equilibrium of the single-period game is a strictly separating equilibrium.

Lemma 3 *Assume that the costs of the firms are drawn independently of each other. Then if the inverse demand function $p(.)$ satisfies the conditions that $p'(.) < 0$ and $p''(.) \leq 0$, then a Bayesian Nash equilibrium of the quantity-setting game is a strictly separating Bayesian Nash Equilibrium.*

Proof: We first note that the condition on the inverse demand function shows that it is concave. Thus the profit-maximizing quantity choice, for every possible cost level $c_i \in C$, is given by the first order condition

$$\sum_{\mathbf{c} \in \mathcal{C}} \mu(\mathbf{c}|c_i)[p'(\sum_{j \neq i} q_j^\star(c_j) + q_i)q_i + p(\sum_{j \neq i} q_j^\star(c_j) + q_i)] = c_i.$$

Hence if $c_i' > c_i$, for every cost profile $\mathbf{c} \in \mathcal{C}$, we have

$$\begin{aligned} & \sum_{\mathbf{c} \in \mathcal{C}} \mu(\mathbf{c}|c_i)[p'(\sum_{j \neq i} q_j^\star(c_j) + q_i^\star(c_i))q_i^\star(c_i) + p(\sum_{j \neq i} q_j^\star(c_j) + q_i^\star(c_i))] = c_i \\ < \; & \sum_{\mathbf{c} \in \mathcal{C}} \mu(\mathbf{c}|c_i')[p'(\sum_{j \neq i} q_j^\star(c_j) + q_i^\star(c_i'))q_i^\star(c_i') + p(\sum_{j \neq i} q_j^\star(c_j) + q_i^\star(c_i'))] = c_i'. \end{aligned} \tag{14}$$

As the costs of the firms are realized independently of each other, we have

$$\sum_{\mathbf{c} \in \mathcal{C}} \mu(\mathbf{c}|c_i) = \sum_{\mathbf{c} \in \mathcal{C}} \mu(\mathbf{c}|c_i')$$

so that (14) implies that if $c_i' > c_i$, then

$$\begin{aligned} & p'(\sum_{j \neq i} q_j^\star(c_j) + q_i^\star(c_i))q_i^\star(c_i) + p(\sum_{j \neq i} q_j^\star(c_j) + q_i^\star(c_i)) \\ < \; & p'(\sum_{j \neq i} q_j^\star(c_j) + q_i^\star(c_i'))q_i^\star(c_i') + p(\sum_{j \neq i} q_j^\star(c_j) + q_i^\star(c_i')). \end{aligned}$$

For a given Q_{-i} consider the function

$$p'(Q_{-i} + q)q + p(Q_{-i} + q) : [0, \bar{q}] \to \mathbb{R}.$$

The first derivative of this function with respect to q is given by

$$p''(Q_{-i}+q)q + p'(Q_{-i}+q) + p'(Q_{-i}+q) = p''(Q_{-i}+q)q + 2p'(Q_{-i}+q) < 0,$$

since $p'(.) < 0$ and $p'' \leq 0$. Thus the function is a strictly decreasing function of q. Since this is true for every Q_{-i}, therefore the function

$$\sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c}|c_i)[p'(\sum_{j\neq i} q_j^\star(c_j) + q_i)q_i + p(\sum_{j\neq i} q_j^\star(c_j) + q_i)] : [0,\bar{q}] \to \mathbb{R}$$

is a strictly decreasing function of q_i. From this and from (14) it now follows that

$$q_i^\star(c_i') < q_i^\star(c_i). \tag{15}$$

But this shows that the Bayesian Equilibrium is a strictly separating equilibrium. ∎

In the next set of results we use the fact that there exists a strictly separating Bayesian Nash equilibrium to construct a strictly separating Perfect Bayesian equilibrium in the repeated game in which the private information about the costs of the firms is revealed early in the game. Example 1, however, shows quite clearly that the firms would reveal such information only under some incentive constraints. In what follows we describe the quantity vectors that satisfy the incentive constraints and maximize the joint profits of the firms subject to these incentive constraints.

Let

$$\mathbf{q} : C \times \cdots \times C \to \mathbb{R}^n_+$$

denote an assignment of quantity vectors as a function of the realized cost profile $\mathbf{c} = (c_1, \cdots, c_n)$. This assignment of quantity vectors will be said to be *incentive compatible* if it satisfies the following constraints. For all $i = 1, \cdots, n$, and for $c_i \in C$,

$$\pi_i(\mathbf{q}(\mathbf{c})|c_i) \geq \pi_i(\mathbf{q}(\mathbf{c_{-i}}, c_i')|c_i) \tag{16}$$

for all $c_i' \neq c_i$. That is, an assignment of quantity vector is incentive compatible if none of the firms have an incentive to claim that its cost is different from that of the true cost c_i. An assignment of quantity vectors

$$\tilde{\mathbf{q}} : C \times \cdots \times C \to \mathbb{R}^n_+$$

will be said to be an *optimal incentive compatible assignment of quantity vectors* if it is an incentive compatible assignment of quantity vectors that solves

$$\text{maximize} \sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c})\pi(\mathbf{q}(\mathbf{c})|\mathbf{c})$$

such that

$$\pi_i(\mathbf{q}(\mathbf{c})|c_i) \geq \pi_i(\mathbf{q}(\mathbf{c_{-i}}, c_i')|c_i) \tag{17}$$

for all $i = 1, \cdots, n$ and $c_i' \neq c_i$, where $\pi(\mathbf{q}(\mathbf{c}))$ denotes the joint profits (the sum of the profits) of the firms when the cost profile is $\mathbf{c}$, and the firms produce the quantity vector $\mathbf{q}(\mathbf{c})$. $\pi_i(.)$ as before denotes the profit of the individual firm i. Thus, an optimal incentive compatible quantity vector maximizes the expected joint profits of the firms, subject to the constraint that none of the firms would want to produce the output assigned to a firm with a different cost structure.

Lemma 4 *If $p(.)$ is continuous then there exists an optimal incentive compatible assignment*[11] *vector $\tilde{q} : C \times \cdots \times C \rightarrow \mathbb{R}^n_+$.*

Proof: Consider an assignment of quantity vectors $q : C \times \cdots \times C \rightarrow \mathbb{R}$ which is constant on $\mathcal{C} = C \times \cdots \times C$, that is $\mathbf{q}(\mathbf{c}) = \mathbf{q}(\mathbf{c}')$ for all $\mathbf{c}, \mathbf{c}' \in \mathcal{C}$. For such an assignment of quantity vectors we have

$$\pi_i(\mathbf{q}(\mathbf{c}|c_i)) = \pi_i(\mathbf{q}(\mathbf{c}_{-i}, c_i'|c_i)) \tag{18}$$

for all $c_i, c_i' \in C$.

Therefore, the set of quantity vectors that satisfy the constraints in (16) is a nonempty compact subset of $\mathbb{R}^{nk}_+$, where k is the total number of possible cost profiles. Since $p(.)$ is continuous, the function

$$\sum_{\mathbf{c} \in \mathcal{C}} \mu(\mathbf{c}) \pi(\mathbf{q}(\mathbf{c})|\mathbf{c})$$

is a continuous function of the the quantity vectors $\mathbf{q}$. Thus the problem given by (17) has a solution. This proves the existence of an optimal incentive compatible assignment of quantity vectors. ■

The optimal incentive compatible vector $\tilde{\mathbf{q}}(\mathbf{c})$ can be viewed as the most collusive output vector for the firms given that the incentive constraints have to hold. Example 1 shows that this cannot be avoided, because if an agreement does not satisfy the incentive constraints, a high-cost firm would always have an incentive to pretend to be a low-cost firm.

[11] The optimal incentive compatible assignment vector is interim efficient as, once the firms have decided on an assignment vector, a firm does not benefit by pretending to have a cost different from the actual realized cost, even after it receives the private information and gets to know the cost of the other firms. This is because the assignment vector satisfies the incentive constraints.

5.1 Strictly Separating Equilibrium with Signaling

We show here that the optimal incentive compatible assignment vector can be produced in a strictly separating Bayesian Nash equilibrium in which there would be a correct inference in equilibrium about the true costs of the firms. In this equilibrium the firms produce the strictly separating Bayesian equilibrium quantity vector in period 1 and then produce the optimal incentive compatible output vector from period 2 onwards. The first period in which the firms play the strictly separating Bayesian Nash equilibrium of the single-period game is played to signal the true costs of the firms.

Theorem 2 *If the costs of the firms are independently drawn and the inverse demand function $p(.)$ satisfies $p'(.) < 0$ and $p''(.) \leq 0$, then for every realized cost profile $\mathbf{c} = (c_1, \cdots, c_n)$, there is a $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$ there is a strictly separating Perfect Bayesian equilibrium in which, from period 2 onwards, the firms produce the optimal incentive compatible quantity vector $\tilde{\mathbf{q}}(c_1, \cdots, c_n)$.*

Proof: The claim is that the strategy combination $(\sigma_1^\star, \cdots, \sigma_n^\star)$ described below is the strategy combination of a strictly separating perfect Bayesian equilibrium.
(i) $\sigma_{i1}^\star(c_i) = q_i^\star(c_i)$ for all $i = 1, \cdots, n$.
(ii)$\sigma_{i2}^\star(h_1, c_i) = \tilde{q}_i(c_1, \cdots, c_n)$ where $(c_1, \cdots, c_n)$ satisfies the condition that $c_i = (q_i^\star)^{-1}(\sigma_{i1})$ for all $i = 1, \cdots, n$.
(iii) If a firm i produces $q_i \neq \sigma_{i2}$ in any period $t \geq 2$, and all the other firms had produced σ_{j2} in all previous periods, then all the firms $j \neq i$ produce the output $\frac{1}{n-1}\bar{q}$ for a length of time T_i. This is a phase I punishment strategy.
(iv) If there are no deviations during a phase I punishment by any of the firms $j \neq i$, then after the length of time T_i, the firms produce the output vector $\tilde{\mathbf{q}}^{i,\epsilon}(\mathbf{c}) = \tilde{\mathbf{q}}^{i,\epsilon}(c_1, \cdots, c_n) = \{\tilde{q}_j^{i,\epsilon}\}_{j=1}^n$(see lemma 1) such that

$$\pi_j(\tilde{\mathbf{q}}^{i,\epsilon}(\mathbf{c}|c_j)) = \pi_j(\tilde{\mathbf{q}}(\mathbf{c}|c_j)) + \frac{\epsilon_i}{n-1}$$

and

$$\pi_i(\tilde{\mathbf{q}}^{i,\epsilon}(\mathbf{c}|c_i)) = \pi_i(\tilde{\mathbf{q}}(\mathbf{c}|c_i)) - \epsilon_i.$$

(v) If a firm $j \neq i$ deviates during a phase I punishment, then firms $\ell \neq j$ produce the output $\frac{1}{n-1}\bar{q}$ for a length of time T_j. This is a Phase II punishment.
(vi) After a phase II punishment the firms produce the quantity vector $\tilde{\mathbf{q}}^{i,\epsilon}(\mathbf{c})$.

(vii) If a firm ℓ deviates from the quantity vector $\tilde{\mathbf{q}}^{j,\epsilon}(\mathbf{c})$, then the other firms play the phase I punishment for a length of time T_ℓ, and then produce the quantity vector $\tilde{\mathbf{q}}^{\ell,\epsilon}(\mathbf{c})$. (viii) Finally, if firm j deviates from the quantity vector $\tilde{\mathbf{q}}^{j,\epsilon}(\mathbf{c})$, then the other firms each produce $\frac{1}{n-1}\bar{q}$ for T_j^1 periods, and then all the firms produce the quantity vector $\tilde{\mathbf{q}}^{j,\epsilon}(\mathbf{c})$ after the T_j^1 periods.

We now show that $\sigma^\star$ is the strategy combination of a strictly separating equilibrium. We first show that if all the firms play their Bayesian Nash equilibrium output vector $q_i^\star(c_i)$ in the first period, then the strategy profile $\sigma^\star$ is an equilibrium. We then argue that the optimal strategy of the firms is to indeed produce the output vector $\{\{q_i^\star(c_i)\}_{c_i \in C}\}_{i=1}^n$ in the first period.

Let M_i be the maximum "gain" a player can make by deviating in any period irrespective of its cost. If firm i deviates in any period, then its payoff in the subsequent periods is given by at most

$$M_i + \delta^{T_i} \sum_{\nu=1}^{\infty} \delta^{\nu-1}[\pi_i(\tilde{\mathbf{q}}(\mathbf{c}|c_i)) - \epsilon_i],$$

as for a length of time T_i firm $i's$ profit is zero every period (see proposition 1). If firm i does not deviate its payoff in the subsequent periods is

$$\sum_{\nu=1}^{\infty} \delta^{\nu-1}[\pi_i(\tilde{\mathbf{q}}(\mathbf{c}|c_j))].$$

Therefore, firm i does not gain from a deviation if

$$\sum_{\nu=1}^{\infty} \delta^{\nu-1}[\pi_i(\tilde{\mathbf{q}}(\mathbf{c}|c_i))] \geq M_i + \delta^{T_i} \sum_{\nu=1}^{\infty} \delta^{\nu-1}[\pi_i(\tilde{\mathbf{q}}(\mathbf{c}|c_i)) - \epsilon_i]$$

or

$$\sum_{\nu=1}^{T_i} \delta^{\nu-1}\pi_i(\tilde{\mathbf{q}}(\mathbf{c}|c_i)) \geq M_i - \delta^{T_i}[\sum_{\nu=1}^{\infty} \delta^{\nu-1}\epsilon_i].$$

That is,

$$\frac{1-\delta^{T_i}}{1-\delta}\pi_i(\tilde{\mathbf{q}}(\mathbf{c}|c_i)) \geq M_i - \frac{\delta^{T_i}}{1-\delta}\epsilon_i. \tag{19}$$

Now note that in (19) the expression $\frac{1-\delta^{T_i}}{1-\delta} \to T_i$ as $\delta \to 1$, therefore, there is a $\delta_1 : 0 < \delta_1 < 1$ such that for T_i sufficiently large, equation (19) is satisfied for all cost profiles $(c_1, \cdots, c_n)$. Thus, phase I punishments can deter a firm from deviating.

Now consider a deviation made by a firm j during a phase I punishment. Let M_j be the maximum "gain" firm j can make irrespective of its cost and L_j the maximum "loss" firm j can sustain every period during a phase I punishment. Then firm $j's$ payoff after deviating is at most

$$M_j + \delta^{T_j} \sum_{\nu=1}^{\infty} \delta^{\nu-1}[\pi_j(\tilde{\mathbf{q}}(\mathbf{c}|c_j)) - \epsilon_j],$$

and if a firm j does not deviate, the payoff in the subsequent periods is

$$\delta^{T_i-t} \sum_{\ell=1}^{\infty} \delta^{\ell-1}[\pi_j(\tilde{\mathbf{q}}(\mathbf{c}|c_j)) + \frac{\epsilon_j}{n-1}] - \sum_{\nu=1}^{T_i-t} \delta^{\nu-1} L_j.$$

Therefore, firm j does not gain by deviating during a phase I punishment if

$$\begin{aligned} &\delta^{T_i-t} \sum_{\nu=1}^{\infty} \delta^{\nu-1}[\pi_j(\tilde{\mathbf{q}}(\mathbf{c}|c_j)) + \frac{\epsilon_j}{n-1}] - \sum_{\nu=1}^{T_i-t} \delta^{\nu-1} L_j \\ &\geq M_j + \delta^{T_j} \sum_{\nu=1}^{\infty} \delta^{\nu-1}[\pi_j(\tilde{\mathbf{q}}(\mathbf{c}|c_j)) - \epsilon_j]. \end{aligned}$$

This reduces to

$$\begin{aligned} \delta^{T_i-t}\frac{1-\delta^{T_j-T_i+t}}{1-\delta}\pi_j(\tilde{\mathbf{q}}(c_1,\cdots,c_n)|c_j) &\geq M_j - \frac{\delta^{T_i-t}}{1-\delta}\frac{\epsilon_j}{n-1} - \epsilon_j\frac{\delta^{T_j}}{1-\delta} + \sum_{\nu=1}^{T_i-t}\delta^{\nu-1}L_j \\ &\geq M_j - \frac{\delta^{T_i-t}}{1-\delta}\frac{\epsilon_j}{n-1} - \epsilon_j\frac{\delta^{T_j}}{1-\delta} + \sum_{\nu=1}^{T_i}\delta^{\nu-1}L_j. \end{aligned} \tag{20}$$

In equation (20), given T_i for a given T_j as $\delta \to 1$, the expression

$$\frac{1-\delta^{T_j-T_i+t}}{1-\delta}$$

goes to $T_j + t - T_i$ and the expressions $\frac{\delta^{T_j}}{1-\delta}$ and $\frac{\delta^{T_i-t}}{1-\delta}$ go to ∞. Hence, there is a $\delta_{j2} : 0 < \delta_{j2} < 1$ such that equation (20) holds for all $\delta > \delta_{j2}$ and for all $(\mathbf{c}|c_j)$. For all such discount factors $\delta > \delta_{j2}$, firm j does not gain from deviating irrespective of its cost.

Next, suppose that firm $\ell \neq j$ deviates during a phase II punishment, then the other firms play a phase II punishment for firm ℓ, and then produce the output vector $\mathbf{q}^{\ell,\epsilon}\mathbf{c}$. As in the case of firm j, firm ℓ cannot gain from deviating if the discount factor is $\delta > \delta_{\ell 2}$.

Finally, if firm j deviates when the quantity vector $\mathbf{q}^{j,\epsilon}(\mathbf{c})$ is being produced, then by choosing T_j^1 such that

$$T_j^1.\pi_j(\tilde{\mathbf{q}}^{j,\epsilon}(\mathbf{c}|c_j)) > M_j$$

it can be argued exactly as in the case of theorem 1(see the proof of theorem 1), that there is a $\delta_2 < 1$ sufficiently large such that for all $\delta \geq \delta_2$ we have

$$\frac{1-\delta^{T_j^1}}{1-\delta}\pi_j(\tilde{\mathbf{q}}^{j,\epsilon}(\mathbf{c}|c_j)) \geq M_j.$$

This shows that firm j cannot gain by deviating when the firms produce the quantity vector $\tilde{\mathbf{q}}^{j,\epsilon}(\mathbf{c})$.

Choose $\bar{\delta} = \max\{\delta_1, \{\delta_{j2},\}_{j\neq i}\}, \delta_2\}$, then firm i does not gain from deviating for $\delta > \bar{\delta}$. Thus no firm can gain by deviating from $\sigma^\star$ if $\delta > \bar{\delta}$.

We now show that the strategy combination $\sigma^\star$ is a Perfect Bayesian equilibrium strategy of the infinite horizon game given the beliefs $\mu(\mathbf{c}|h_t, c_i, \sigma^\star)$ generated by the $\sigma^\star$. In period 1, the expected payoff of firm i, given that it knows its marginal cost is c_i is given by

$$\sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c}|c_i)\pi_i^\infty(\sigma^\star, c_i) = \sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c}|c_i)[\pi_i(\{q_j^\star\}_{j=1}^n) + \frac{\delta}{1-\delta}(\pi_i(\tilde{\mathbf{q}}(\mathbf{c})|c_i)].$$

For $t \geq 2$ we have already shown that players cannot gain by deviating from $\sigma^\star$ in any period t, no matter what the beliefs are of the firms about the costs of the other firms. Further, from lemma 4 we know that since $\tilde{\mathbf{q}}(\mathbf{c}|c_i)$ is incentive compatible, playing according to the true costs of the firms in period 1 would be optimal in the game from period 2 onwards. Since the firms maximize expected payoffs when they play the strictly separating Bayesian equilibrium, the strategy $\sigma^\star$ together with the beliefs $\mu(\mathbf{c}|h_t, c_i, \sigma^\star)$ (see (1)) is a Perfect Bayesian equilibrium of the infinite horizon game.

The fact that it is a strictly separating equilibrium follows from observing that $\sigma^\star$ involves playing a single period Bayesian Nash equilibrium in period 1. Since the conditions of lemma 3 are satisfied it follows from that result that the single period Bayesian Nash equilibrium is a strictly separating Bayesian Nash equilibrium. Therefore, $\sigma^\star$ is a strictly separating equilibrium. ■

5.2 Equilibrium with Communication

In theorem 2 we showed that there is a strictly separating equilibrium in which the firms play a strictly separating Bayesian Nash equilibrium in period 1 to signal their costs.

However, the Bayesian Nash equilibrium of the single-period game frequently gives lower profits to the firms than the optimal incentive compatible outputs, especially in those cases in which collusion is most likely to be profitable for the firms. This, therefore, raises the question as to whether there is an equilibrium in which the firms can collude, without having to play the Bayesian equilibrium of the single-period game in period 1.

Consider the firms playing the infinite horizon game as described in section 3 but with an initial communication phase in which the firms report their cost to all the other firms. Thus, before choosing the output, firm i sends a report $r_i : C \to C$. Given the reported costs, the firms then choose their individual strategies for the infinite horizon game. The next result shows that there is a strictly separating equilibrium in which the firms report their costs truthfully and then produce the optimal incentive compatible quantity vector given the costs reported by the firms.

Theorem 3 *For every realized cost profile $\mathbf{c} = (c_1, \cdots, c_n)$ there is a $\underline{\delta} < 1$, such that for all $\delta > \underline{\delta}$ there is a Perfect Bayesian equilibrium in which the firms produce the optimal incentive-compatible collusive quantity vector $\tilde{\mathbf{q}}(\mathbf{c})$ from period 1 onwards, and in the initial communication period the firms truthfully report their realized costs.*

Proof: The claim is that the strategy combination $(\sigma_1^\star, \cdots, \sigma_n^\star)$ described below is the strategy combination of a strictly separating Perfect Bayesian equilibrium.
(i) $\sigma_{i0}^\star(c_i) = r_i^\star(c_i) = c_i$ for all $i = 1, \cdots, n$ and $c_i \in C$.
(ii) $\sigma_{i1}^\star(\mathbf{r}) = \tilde{q}_i(\mathbf{r})$, where $\mathbf{r} = (r_1, \cdots, r_n)$ is the vector of reported costs.
(iii) $\sigma_{it}^\star(h_{t-1}) = \tilde{q}_i(\mathbf{r})$, if the history up to period t has been the output vector $\tilde{\mathbf{q}}(\mathbf{r})$.
(iv) If a firm i produces $q_i \neq \tilde{q}_i(\mathbf{r})$ in any period $t \geq 2$, when the past history has been the output vector $\tilde{\mathbf{q}}(\mathbf{r})$, then all firms $j \neq i$ produce the output $\frac{1}{n-1}\bar{q}$ for a length of time T_i. This is a phase I punishment strategy.
(v) If there are no deviations during a phase I punishment by any of the firms $j \neq i$, then after the length of time T_i, the firms produce the output vector $\mathbf{q}^{i,\epsilon}(\mathbf{r}) = \{q_j^{i,\epsilon}(\mathbf{r})\}_{j=1}^n$(see lemma 1) such that

$$\pi_j(\mathbf{q}^{i,\epsilon}(\mathbf{r})|c_j) = \pi_j(\tilde{\mathbf{q}}(\mathbf{r})|c_j) + \frac{\epsilon_i}{n-1}$$

and

$$\pi_i(\mathbf{q}^{i,\epsilon}(\mathbf{r})|c_i) = \pi_i(\tilde{\mathbf{q}}(\mathbf{r})|c_i) - \epsilon_i.$$

(vi) If a firm $j \neq i$ deviates during a phase I punishment, then firms $\ell \neq j$ produce the output $\frac{1}{n-1}\bar{q}$ for a length of time T_j. This is a Phase II punishment.

(vii) After a phase II punishment the firms produce the quantity vector $\mathbf{q}^{j,\epsilon}(\mathbf{r})$.
(viii) If a firm ℓ deviates from the quantity vector $\mathbf{q}^{j,\epsilon}(\mathbf{r})$ then the other firms play the phase I punishment for a length of time T_ℓ and then produce the quantity vector $\mathbf{q}^{i,\epsilon}(\mathbf{r})$.
(ix) Finally, if a firm j deviates from the quantity vector $\tilde{\mathbf{q}}^{j,\epsilon}(\mathbf{r})$, then the firms $\ell \neq j$ each produce the output $\frac{1}{n-1}\bar{q}$ for a length of time T_j^1 and then the firms produce the quantity vector $\tilde{\mathbf{q}}^{j,\epsilon}(\mathbf{r})$.

Using arguments similar to those used to prove theorem 2 it can be shown that from period 1 onwards, the strategy profile $\sigma^\star$ is a Perfect Bayesian equilibrium. Further, since the assignment vector $\tilde{\mathbf{q}}(.)$ is incentive compatible, the firms do not gain from misreporting the realized cost. Thus, reporting the true cost is an optimal strategy for the firms. The strategy profile $\sigma^\star$ is, therefore, a Perfect Bayesian equilibrium. ■

6 Stationary Equilibria and Optimal Collusion

We have seen that the infinite horizon game has pooling equilibria as well as strictly separating equilibria with signaling as well as with communication. Here we investigate which among these different types of equilibrium gives the largest expected joint profit to the firms. We show that the expected joint profits from the separating equilibrium with communication is at least as large as the expected discounted profits from the separating equilibrium with signaling as well as from the expected joint profits from the pooling equilibrium. Therefore the separating equilibrium with communication in which the firms produce the optimal incentive compatible output vector from period one onwards is optimal among all these stationary equilibria, that is, among those equilibria in which the firms produce the same quantity vector every period[12].

Theorem 4 [13] *Assume that the realizations of the cost of the firms are independent of each other and that the inverse demand function satisfies $p' < 0$ and $p'' < 0$. Then the expected discounted sum of joint profits in a separating equilibrium with communication, in which the firms produce the optimal incentive compatible quantity vector, gives the maximum expected discounted joint profits among all pooling and separating equilibria.*

[12] Since the firms produce the same quantity vector in every period, these are all stationary equilibria and have stable prices and output.

[13] The result is also true when the inverse demand curve is linear and exactly the same proof goes through. The condition that the realizations of the cost are independent is only used to show that the expected joint profit from optimal incentive compatible quantities is at least as large as the expected joint profit from the Bayesian Nash equilibrium outputs.

Proof: We first show that the expected discounted sum of the joint profits when the firms produce the optimal incentive compatible quantity vector in a separating equilibrium with communication is at least as large as that from the optimal pooling equilibrium. Let $\hat{\mathbf{q}}$ be the quantity vector that maximizes the expected joint profit from a pooling equilibrium. Since $\hat{\mathbf{q}}$ is the quantity vector from a pooling equilibrium it follows that $\hat{\mathbf{q}}(\mathbf{c}) = \hat{\mathbf{q}}(\mathbf{c}')$ for all cost profiles $\mathbf{c}, \mathbf{c}' \in \mathcal{C}$. Therefore, $\hat{\mathbf{q}}$ trivially satisfies the incentive compatibility constraints given by (16). Therefore,

$$\sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c})\pi(\tilde{\mathbf{q}}(\mathbf{c})|\mathbf{c}) \geq \sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c})\pi(\hat{\mathbf{q}}(\mathbf{c})|\mathbf{c}). \tag{21}$$

From this it quickly follows that the present value of the expected joint profits from the separating equilibrium with communication, given by $\sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c})\frac{\pi(\tilde{\mathbf{q}}(\mathbf{c})|\mathbf{c})}{1-\delta}$ is greater than or equal to the expected joint profits from the optimal pooling equilibrium, given by $\sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c})\frac{\pi(\hat{\mathbf{q}}|\mathbf{c})}{1-\delta}$.

Next we show that the expected joint profits from the separating equilibrium with communication, in which the firms produce the optimal incentive compatible output vector from period one onwards, is at least as large as the expected joint profits from the strictly separating equilibrium with signaling. We will prove this by showing that the expected joint profit from the Bayesian Nash equilibrium cannot exceed the expected joint profit from the quantity vector that maximizes the expected joint profit.

We first note that since $p''(.) < 0$, the inverse demand function $p(.)$ is jointly concave in the quantity vector $\mathbf{q}$. This is verified by noting that $\frac{\partial^2 p}{\partial q_i^2} = \frac{\partial^2 p}{\partial q_i q_j} = p''$ for all $i \neq j$ so that the Hessian is given by

$$\begin{bmatrix} p'' \cdots p'' \\ p'' \cdots p'' \\ \cdots \\ p'' \cdots p'' \end{bmatrix}.$$

Since $p'' < 0$, the first principal minor of this Hessian is negative. Since it is also true that the higher order principal minors are all zero, it follows that $p(.)$ is concave.

Using the fact that the inverse demand function $p(.)$ is concave, the expected profit of firm i from the Bayesian Nash equilibrium, when the realized cost is c_i, is

$$E\pi_i^\star(c_i) \;=\; \sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c}|c_i)[p(\sum_{j\neq i} q_j^\star(c_j) + q_i^\star(c_i)) - c_i]q_i^\star(c_i)$$

$$\leq \quad [p(\sum_{\mathbf{c}\in\mathcal{C}}\mu(\mathbf{c}|c_i)\sum_{j\neq i}q_j^\star(c_j)+q_i^\star(c_i))-c_i]q_i^\star(c_i)$$

$$= \quad [p(\sum_{j\neq i}\sum_{c_j\in C}\mu(c_j|c_i)q_j^\star(c_j)+q_i^\star(c_i))-c_i]q_i^\star(c_i). \tag{22}$$

Therefore the expected profit of firm i from the Bayesian Nash equilibrium

$$\begin{aligned}\sum_{c_i\in C}\mu(c_i)E\pi_i^\star(c_i) &\leq \sum_{c_i\in C}\mu(c_i)[p(\sum_{j\neq i}\sum_{c_j\in C}\mu(c_j|c_i)q_j^\star(c_j)+q_i^\star(c_i))-c_i]q_i^\star(c_i)\\ &\leq [p(\sum_{j\neq i}\sum_{c_j\in C}\mu(c_j|c_i)q_j^\star(c_j)+\sum_{c_i\in C}\mu(c_i)q_i^\star(c_i))-\bar{c}_i]\sum_{c_i\in C}\mu(c_i)q_i^\star(c_i)\end{aligned} \tag{23}$$

where $\bar{c}_i = \sum_{c_i\in C}\mu(c_i)c_i$. The last inequality in (23) follows from the concavity of the inverse demand function. Since the realizations of the cost of the firms are independent, $\mu(c_j|c_i) = \mu(c_j)$ for all $j \neq i$. Let

$$\hat{\hat{q}}_i = \sum_{c_i\in C}\mu(c_i)q_i^\star(c_i),$$

and $\hat{\hat{\mathbf{q}}} = (\hat{\hat{q}}_1, \cdots, \hat{\hat{q}}_n)$, then (23) can be rewritten as

$$\begin{aligned}\sum_{c_i\in C}\mu(c_i)E\pi_i^\star(c_i) &\leq [p(\sum_{j\neq i}\sum_{c_j\in C}\mu(c_j|c_i)q_j^\star(c_j)+\sum_{c_i\in C}\mu(c_i)q_i^\star(c_i))-\bar{c}_i]\sum_{c_i\in C}\mu(c_i)q_i^\star(c_i)\\ &= [p(\hat{\hat{\mathbf{q}}})-\bar{c}_i]\hat{\hat{q}}_i.\end{aligned} \tag{24}$$

Since $\hat{\mathbf{q}}$ maximizes $E\pi(\mathbf{q})$, we have

$$\sum_{i=1}^{n}[p(\hat{\hat{\mathbf{q}}})-\bar{c}_i]\hat{\hat{q}}_i = E\pi(\hat{\hat{\mathbf{q}}}) \leq E\pi(\hat{\mathbf{q}}) \tag{25}$$

From (24) and (25) the expected joint profit from the Bayesian Nash equilibrium

$$E\pi^\star \leq E\pi(\hat{\mathbf{q}}). \tag{26}$$

From (21) and (26) it now follows that the present value of the expected joint profits from the separating equilibrium with communication given by

$$\sum_{\mathbf{c}\in\mathcal{C}}\mu(\mathbf{c})\frac{\pi(\tilde{\mathbf{q}}(\mathbf{c})|\mathbf{c})}{1-\delta}$$

is at least as large as the expected joint profits from the separating equilibrium with signaling given by

$$\sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c})\pi^{\star} + \delta \sum_{\mathbf{c}\in\mathcal{C}} \mu(\mathbf{c})\frac{\pi(\tilde{\mathbf{q}}|\mathbf{c})}{1-\delta}.$$

This completes the proof. ■

The following example illustrates some of these results.

Example 2 *The discounted sum of profits from the* **optimal separating equilibrium** *and the* **optimal pooling equilibria**.

Consider the oligopoly game of example 1 in which

$$p(Q) = 10 - Q$$

and the marginal cost of firm 1 is either $c_1 = 1$ or $c_1 = 2$ with

$$Prob(c_1 = 1) = Prob(c_1 = 2) = 0.5$$

and similarly, the marginal cost of firm 2 is either $c_2 = 1$ or $c_2 = 2$ with

$$Prob(c_1 = 1) = Prob(c_1 = 2) = 0.5.$$

In this case the Bayesian Nash equilibrium outputs of the firms are

$$q_1^{\star}(c_1 = 1) = 3.0833,\ q_1^{\star}(c_1 = 2) = 2.5833,\ q_2^{\star}(c_2 = 1) = 3.0833, \text{ and } q_2^{\star}(c_2 = 2) = 2.5833.$$

As a result the profit of the firms when they play the Bayesian Nash equilibrium are

$$\pi_1^{\star}(c_1 = 1, c_2 = 1) = \pi_2^{\star}(c_1 = 1, c_2 = 1) = 8.7362,$$

$$\pi_1^{\star}(c_1 = 1, c_2 = 2) = 10.2779, \quad \pi_2^{\star}(c_1 = 1, c_2 = 2) = 6.0279,$$

$$\pi_1^{\star}(c_1 = 2, c_2 = 1) = 6.0279, \quad \pi_2^{\star}(c_1 = 2, c_2 = 1) = 10.2779,$$

and

$$\pi_1^{\star}(c_1 = 2, c_2 = 2) = \pi_2^{\star}(c_1 = 2, c_2 = 2) = 7.3195.$$

The expected joint profit of the firms from the Bayesian Nash equilibrium quantities is therefore

$$E\pi^{\star} = 0.25 \times 17.4724 + 0.5 \times 16.3058 + 0.25 \times 14.639 = 16.1808. \tag{27}$$

The optimal incentive compatible collusive output of the firms are

$$\hat{q}_1(c_1 = 1, c_2 = 1) = 2.25,\ \ \hat{q}_2(c_1 = 1, c_2 = 1) = 2.25,$$

$$\hat{q}_1(c_1 = 2, c_2 = 2) = 2,\ \ \hat{q}_2(c_1 = 2, c_2 = 2) = 2,$$

$$\hat{q}_1(c_1 = 1, c_2 = 2) = 2.1869,\ \ \hat{q}_2(c_1 = 1, c_2 = 2) = 2.155,$$

and

$$\hat{q}_1(c_1 = 2, c_2 = 1) = 2.155,\ \ \hat{q}_2(c_1 = 2, c_2 = 1) = 2.1869.$$

The optimal incentive compatible profits are

$$\hat{\pi}_1(c_1 = 1, c_2 = 1) = \hat{\pi}_2(c_1 = 1, c_2 = 1) = 10.125,$$

$$\hat{\pi}_1(c_1 = 1, c_2 = 2) = 10.1867, \quad \hat{\pi}_2(c_1 = 1, c_2 = 2) = 7.8809,$$

$$\hat{\pi}_1(c_1 = 2, c_2 = 1) = 7.8809, \quad \hat{\pi}_2(c_1 = 2, c_2 = 1) = 10.1867,$$

and

$$\hat{\pi}_1(c_1 = 2, c_2 = 2) = \hat{\pi}_2(c_1 = 2, c_2 = 2) = 8.$$

The *expected joint profit from the* **optimal incentive compatible quantity vector** is

$$E\pi(\tilde{\mathbf{q}}) = 0.25 \times 20.25 + 0.5 \times 18.07 + 0.25 \times 16 = 18.10. \tag{28}$$

The output that maximizes the expected joint profit when the firms produce a pooled quantity vector is given by

$$p'(Q)Q + p(Q) = 0.5c_L + 0.5c_H = 1.5.$$

This gives $\hat{Q} = 4.25$ so that $\hat{q}_1 = \hat{q}_2 = 2.125$. The *expected joint profit of the firms in the case of the* **pooled quantity vector** is, therefore,

$$\begin{aligned} E\pi(\hat{\mathbf{q}}) &= 0.25 \times (5.75 - 1) \times 4.25 + 0.5 \times [(5.75 - 1) \times 2.125 + (5.75 - 2) \times 2.125] \\ &\quad +0.25 \times (5.75 - 2) \times 4.25 = 18.0625. \end{aligned} \tag{29}$$

From (27), (28) and (29) it follows that the expected discounted joint profits from the *separating equilibrium with communication* given by

$$\frac{18.10}{1 - \delta}$$

is strictly greater than the expected discounted joint profits from the *optimal separating equilibrium with signaling*, which is

$$16.1808 + \frac{\delta}{1-\delta}18.10,$$

as well as from the expected discounted joint profits from the *optimal pooling equilibrium*, given by

$$\frac{18.0625}{1-\delta}.$$

An interesting pooling equilibrium is one in which both the firms decide to play the pooling equilibrium assuming that each is low-cost firm.[14] Then if both are indeed low-cost firms, the profit of the firms are 10.125 each, whereas if a firm is a high-cost firm, its profit is 7.875. The expected joint profit of the firm is, therefore,

$$0.25 \times 20.25 + 0.5 \times (10.125 + 7.875) + 0.25 \times 15.75 = 18.0. \tag{30}$$

The expected joint profit is clearly not as high as from the optimal pooling equilibrium and, therefore, not as large as that from the optimal separating equilibrium with communication. ∎

The separating equilibrium with communication is clearly optimal for the firms in an ex ante sense. What is, however, also of interest is that in the optimal separating equilibrium with communication, the firms would produce the optimal collusive output of the complete information case if the realized costs are identical. In that case they would also be ex post optimal. It is also worth mentioning that the optimal separating equilibrium with communication is not only optimal ex ante but also optimal in the interim, that is, after the firms receive their private information. It is, however, not optimal ex post for every possible realizations of the costs. The present example is one in which there is a great deal of symmetry between the firms as each firm can be a high-cost or a low-cost firm with equal probability. We know from the complete information case that collusion is most likely in these symmetric cases, and this seems to be true here too.

In case the situation is not symmetric, theorem 4 of course will still hold, but optimal collusion of the sort that is possible in the symmetric cases may not work out

[14]This might in many cases be quite appealing to the firms because if a firm is indeed low-cost it would get the same profit as it would get if the other firm was low-cost and they produced the optimal collusive output. A high-cost firm may also be happy with such an arrangement as it avoids the possibility of having to share the market with a low-cost firm.

as smoothly. The following example illustrates what could happen in the non symmetric cases.

Example 3 *The discounted sum of profits from the* **optimal separating equilibrium** *and the* **optimal pooling equilibria** *in a non symmetric case.*

Consider the case in example 2, but with the probability distribution given by $\mu(c_L, c_L) = 0.05$, $\mu(c_L, c_H) = 0.9$, $\mu(c_H, c_H) = 0.05$, and $\mu(c_H, c_L) = 0$, so that firm 1 is very likely to be low-cost while firm 2 is almost certainly high-cost. In this case, the quantity vector that maximizes the expected joint profit is $\hat{q}_1 = 4.5$ and $\hat{q}_2 = 0$. The optimal incentive compatible quantity vector is

$$\tilde{q}_1(c_L, c_H) = 4.5,\ \ \tilde{q}_2(c_L, c_H) = 0,$$

$$\tilde{q}_1(c_H, c_H) = 3.94,\ \ \tilde{q}_2(c_H, c_H) = 0.06,$$

and

$$\tilde{q}_1(c_L, c_L) = 4.365,\ \ \tilde{q}_2(c_L, c_L) = 0.135.$$

Clearly, the expected joint profit from the separating equilibrium with communication is higher than the expected joint profit from the pooling equilibrium. However, firm 2 gets very little of the joint profit in either case[15] and the firms could revert to playing a non-collusive equilibrium in which the firms play the single period Bayesian Nash equilibrium in every period (see for example Lemma 4, Chakrabarti [2005]). ∎

One feature of the mostly stationary perfect Bayesian equilibrium points is that, even in the optimal separating equilibrium, the high cost firm produces a positive output starting from period 1. This as we all know is not the most efficient way to collude for the firms. If one firm is a low cost firm and the other firm is a high cost firm, then if the firms were to collude optimally, only the low cost firm would produce a positive output. In the next example we show that there are non-stationary perfect Bayesian equilibrium points in which, in the first period or the first few periods, the high cost firm produces zero output. However, as we shall see that in order to collude optimally in the first period the low cost firm has to commit to giving up substantial market share in the future.

[15]This is also what happens in the complete information case in which one firm is a low-cost firm and the other is a high-cost firm. Optimal collusion would mean that the high-cost firm produces zero output in which case firm 2 would never agree to collude, and the firms could end up playing the Nash equilibrium every period.

Example 4 *A Perfect Bayesian equilibrium in which the firms produce the optimal collusive output for each possible cost profile in period* 1.

Consider the situation as described in example 2 but with the following assignment of output for the different cost profiles.

$$\begin{aligned}
&\hat{q}_{t,1}(c_1 = 1, c_2 = 1) = 2.25,\ \hat{q}_{t,2}(c_1 = 1, c_2 = 1) = 2.25,\ \text{for all } t \geq 1 \\
&\hat{q}_{t,1}(c_1 = 2, c_2 = 2) = 2,\ \hat{q}_{t,2}(c_1 = 2, c_2 = 2) = 2,\ \text{for all } t \geq 1 \\
&\hat{q}_{1,1}(c_1 = 1, c_2 = 2) = 4.5,\ \hat{q}_{1,2}(c_1 = 1, c_2 = 2) = 0, \\
&\hat{q}_{t,1}(c_1 = 1, c_2 = 2) = 2.1869,\ \hat{q}_{t,2}(c_1 = 1, c_2 = 2) = 2.155,\ \text{for all } t \geq 2 \\
&\hat{q}_{1,1}(c_1 = 2, c_2 = 1) = 0,\ \hat{q}_{1,2}(c_1 = 2, c_2 = 1) = 4.5, \\
&\hat{q}_{t,1}(c_1 = 2, c_2 = 1) = 2.155,\ \hat{q}_{t,2}(c_1 = 2, c_2 = 1) = 2.1869,\ \text{for all } t \geq 2. \qquad (31)
\end{aligned}$$

Here $\hat{q}_{t,i}(c_1, c_2)$ denotes the output assigned to firm i in period t when the cost profile announced is (c_1, c_2). Note that according to this assignment of quantity vectors, in period 1, the firms produce the optimal collusive output vector for each of the four possible cost profiles. And from period 2 onwards the firms produce the optimal incentive compatible collusive output vector. It remains to be shown that this assignment of the quantity vectors is incentive compatible over time. As the situation is symmetric for the two firms, it is enough to check the cases for firm 1. There are two cases to be checked.

Case 1: $c_2 = 1$.

If firm $1's$ realized cost is $c_1 = 1$ and it reports $c_1 = 1$, then its discounted sum of profits is

$$\frac{10.125}{1-\delta}$$

whereas if it reports $c_1 = 2$, its discounted sum of profits is

$$0 + \frac{7.88\delta}{1-\delta}.$$

Therefore, firm 1 will always report $c_1 = 1$ if $c_1 = 1$.

If firm $1's$ realized cost is $c_1 = 2$ and it reports $c_1 = 2$, then its discounted sum of profits is

$$0 + \frac{7.88\delta}{1-\delta}$$

whereas if it reports $c_1 = 1$, then its discounted sum of profits is

$$7.875 + \frac{7.875}{1-\delta}.$$

Therefore, firm 1 will report $c_1 = 2$ if and only if

$$0 + \frac{7.88\delta}{1-\delta} \geq 7.875 + \frac{7.875}{1-\delta}$$

that is, if and only if

$$\delta \geq 0.9994. \tag{32}$$

Case 1: $c_2 = 2$.

If firm $1's$ realized cost is $c_1 = 1$ and it reports $c_1 = 1$, then its discounted sum of profits is

$$20.25 + \frac{10.1868\delta}{1-\delta}$$

whereas if it reports $c_1 = 2$, then its discounted sum of profits is

$$\frac{10}{1-\delta}.$$

Therefore, firm 1 will always report $c_1 = 1$ if $c_1 = 1$.

If firm $1's$ realized cost is $c_1 = 2$ and it reports $c_1 = 2$, then its discounted sum of profits is

$$\frac{8}{1-\delta}$$

whereas if it reports $c_1 = 1$, its discounted sum of profits is

$$15.75 + \frac{7.9999\delta}{1-\delta}.$$

Therefore, firm 1 will report $c_1 = 2$ if and only if

$$8 + \frac{8\delta}{1-\delta} \geq 15.75 + \frac{7.999\delta}{1-\delta}$$

that is, if and only if

$$\delta \geq 0.9999. \tag{33}$$

From (32) and (33) it follows that for a discount rate δ that satisfies

$$\delta \geq 0.9999$$

the quantity assignment given by (31) is incentive compatible and can be implemented as a perfect Bayesian equilibrium of the infinite-horizon game. ∎

The example highlights the fact that if firms are extremely patient, that is, the discount factor is close to 1 then it may be possible to find a perfect Bayesian equilibrium in which the firms produce the optimal collusive output for each possible cost profile in period 1. This separating equilibrium will then have an expected joint profit that will be larger than even the expected joint profit from the optimal separating equilibrium with communication. However, this would only be true of a discount factor δ that is very close to 1. This should be contrasted with what is required for the perfect Bayesian equilibria discussed in theorems 1, 2 and 3, in which the critical value of the discount factor δ is determined solely by the requirements of the equilibrium punishment strategies. Example 4 also suggests that if one changed the incentive compatible assignment of quantity vectors and allowed a larger margin of profit for reporting the true cost of a firm, it would be possible to lower the critical value of δ from 0.9999 to something a little lower. But this would then indicate that there is a trade off between higher expected joint profits in the first periods with lower expected joint profits in the later periods.

7 Conclusion

The results reported here draw quite heavily on the fact that the stage game is an oligopoly game. Proposition 1, for instance, is quite crucial as it shows that the firms can use a common minimaxing strategy to minimax deviating firms that is independent of the cost of the firm. This is, of course, not the case for games in general. Depending on the nature of the private information, it may not be possible to find a common minimaxing strategy.

Communication seems to be useful as long as there are the right incentives to report the costs correctly. This is consistent with observations made in other contexts in the literature on repeated games with private information as communication of some sort seems to be useful in many cases when there is private information. This is clearly the case in Compte [1998] and Kandori and Matsushima [1998]. Even when there is complete information it is generally understood that firms have to agree on the collusive output vector and some bargaining may take place prior to deciding on the collusive output vector. The reporting of a firm's cost can then be part of the same communication process. It is important, however, to note here that the communication here is quite limited and simply involves just a report of its cost by each firm and nothing more. Discussions among firms about setting output levels and/or prices can often be cited as evidence of firms trying to collude and firms can be prosecuted on the basis of such evidence when anti-trust laws are in effect. The very limited communication present here, which is a simple report of its cost by a firm, avoids any discussions about the output and/or the price levels. Such limited forms of communication could thus allow the firms to collude tacitly even in an environment in which antitrust laws are in effect. However, if the firms are reluctant to engage in even this very limited form of communication, it is possible that they may opt to collude tacitly by playing the optimal pooling equilibrium, especially if the potential loss is relatively small.

The other issue that one may raise here is whether the choice of a constant marginal cost is crucial for the results reported here. If one had U-shaped marginal costs, a version of proposition 1 would still be valid, as the price can again be driven down to low levels so that a deviator makes zero profits. Also, the non-deviators can be rewarded with a slightly larger market share in the future and the deviator can be punished by a slightly lower market share in the future after the minimaxing stage is over. Thus it seems that the assumption of constant marginal cost as opposed to U-shaped marginal costs is not that critical.

An assumption that we have retained throughout the paper is that there are a finite number of possible costs of the firms. One might therefore ask as to what would happen to the results presented here if the marginal costs of the firms were drawn from a continuum of possible values. It seems that this would have no major impact on the results of section 4. However, this would definitely affect the existence result on the Bayesian Nash equilibrium of section 5 and theorem 3; the incentive constraints would be hard to deal with as there would be infinitely many of them. It is thus unclear as to what would happen to the results of section 5. Clearly, the incentive compatibility conditions would have to be modified. In this case it may be that the result would hold only with a certain "noise," and the exact nature of the result may have to be rethought. This is thus left for future work.

The equilibria in the pooling case and the equilibria with communication are all stationary equilibria, namely equilibria in which the firms play the same strategy in each period if there is no deviation. The signaling equilibrium is also mostly stationary, as after period 1 the firms produce the same optimal incentive compatible output in each period. The signaling here is designed to minimize the cost that may arise from signaling and is distinct from the price wars that are often observed in markets that ultimately form Cartels; see for example the examination of the Lysine market in Nicola de Roos [2006], where it is shown that price wars occurred before collusion took place. Price wars of this kind may arise due to the inability of the firms to agree on the exact form of collusion.

Aside from the mostly stationary equilibria that we analyze here, there are non-stationary equilibria in these games, as shown in example 4 as well as in Chakrabarti [2005]. Would a non-stationary equilibrium lead to higher expected joint profits for the firms? The non-stationary equilibrium in Chakrabarti [2005] require the firms to play the Bayesian Nash equilibrium of the single-period game for several periods, until there is no incentive for a firm to play differently from what is dictated by its true cost. But this then leads to lower joint profits in the first periods, and unless the firms are very patient, the discounted sum of profits will be smaller. A different kind of non-stationary equilibrium point is examined in example 4. In that example the firms collude optimally in period 1 for each possible cost profile, and then produce the optimal incentive compatible quantity vector of example 2 from period 2 onwards. However, it is an equilibrium only for a discount factor very very close to 1. An alternative to this would be to give a firm the incentive to reveal its cost early; for instance by giving it a large market share, and then playing according to the information inferred. But the incentives have to be right, and if one gives away too much market share in the beginning, it would be hard to recoup the losses later, especially if future payoffs are discounted. The optimal incentive compatible quantity vectors may thus provide the best balance between the incentives to reveal costs truthfully and future profits.

References:

1. Abreu, D., D. Pearce, and E. Stachetti (1986): "Optimal cartel equilibria with imperfect monitoring," *Journal of Economic Theory*, 39, 251-269.

2. Athey, S. and K. Bagwell (2001): "Optimal Collusion with Private Information," *Rand Journal of Economics*, Autumn 32(3).

3. Athey, S. and K. Bagwell (2008): "Collusion with Persistent Shocks," *Econometrica*, 76, 493 - 540.

4. Athey, S., K. Bagwell and C. Sanchirico (2004): "Collusion and Price Rigidity," *Review of Economic Studies*, 71, 317-340.

5. Chakrabarti, S. K. (2005): "Equilibrium in Cournot Oligopolies with Unknown Costs," mimeo, Department of Economics, IUPUI.

6. Compte, O. (1998): "Communication in repeated games with imperfect private monitoring," *Econometrica*, 66, 597-626.

7. Cole, Harold L. and N. Kocherlakota, (2001): "Dynamic Games with Hidden Actions and Hidden States," *Journal of Economic Theory*, 98, 114-120.

8. The *Economist*, August 9, 2003, page 54.

9. Fudenberg, D. and E. Maskin (1986):"The folk theorem in repeated games with discounting or with incomplete information,"*Econometrica*, 62, 997-1039.

10. Fudenberg, D. and J. Tirole (1991): "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory*, 53, 236-260.

11. Geralch, H. (2007):"Stochastic Market Sharing, Partial Communication and Collusion," working paper no. 674, IESE Business School, University of Navarra.

12. Green, E. and R. Porter (1984): "Noncooperative collusion under imperfect price information," *Econometrica*, 52, 87-100.

13. Hanazono, H. and H. Yang (2007):"Collusion, Fluctuating Demand and Price Rigidity," *International Economic Review*, 48, 483-515 .

14. Kandori, M. and H. Matsushima (1998): "Private observation, communication and collusion," *Econometrica*, 66, 627-652.

15. Motta, M. (2004), *Competition Policy: Theory and Practice*, Cambridge University Press.

16. Tirole, J. (1993), *The Theory of Industrial Organization*, sixth edition, MIT Press.

17. Kennan, J. (2001): "Repeated Bargaining with Persistent Private Information," *Review of Economic Studies*, 68, 719-755.

18. Kreps, D. and R. Wilson (1982): "Sequential Equilibrium," *Econometrica*, 50, 863-894.

19. LaCasse, C. (1999): "Price-fixing with Incomplete Information," mimeo, University of Alberta.

20. Roos, Nicolas de (2006): "Examining Models of Collusion: the Market for Lysine," *International Journal of Industrial Organization*, 24(6), 1083-1107.