

Impossibility of Collusion under Imperfect Monitoring with Flexible Production

Yuliy Sannikov and Andrzej Skrzypacz*

October 2006

We show that it is impossible to achieve collusion in a duopoly when (1) goods are homogenous and firms compete in quantities, (2) new, noisy information arrives continuously, without sudden events and (3) firms are able to respond to new information quickly. The result holds even if we allow for asymmetric equilibria or monetary transfers. The intuition is that the flexibility to respond quickly to new information unravels any collusive scheme. Our result applies to both a simple stationary model and a more complicated one with prices following a mean-reverting Markov process, as well as to models of dynamic cooperation in many other settings. (JEL C73, D43, L13, L14)

Collusion is a major problem in many markets and has been an important topic of study in both applied and theoretical economics. From exposed collusive cases we know how numerous real-life cartels have been organized and what kinds of agreements (either implicit or explicit) are likely to be successful in obtaining collusive profits.¹ At least since George J. Stigler (1964), economists have recognized that imperfect monitoring may destabilize cartels. Nevertheless, the seminal paper of Edward J. Green and Robert H. Porter (1984) has shown that even with imperfect monitoring firms can create collusive incentives by allowing price wars to break out with positive probability.

We study the scope of collusion in a quantity-setting duopoly with homogenous goods and flexible production – that is, if firms can change output flow frequently.² As in Green and Porter (1984), in our model firms cannot observe each other's production decisions directly. They only observe noisy market prices/signals that depend on the total market supply. We show that collusion is impossible to achieve if

- (1) new, noisy information arrives continuously, without sudden events,
- (2) the firms have flexible production technologies and can thus react to new information quickly, and
- (3) public signals depend on total market supply only, and not individual decisions.

There are many markets that fit the general features of our model. For example, in markets with homogenous goods, e.g. chemicals, firms are selling both to a spot market and to clients, with client deals being private but affecting the spot market. The spot price can be used to monitor the success of collusion. One example that received a lot of attention was the cartel producing lysine, an amino acid.³ This cartel tried to collude by setting and monitoring a target price at first. However, those early attempts failed. Quoting from Cabral (2005), p. 201:

The topic of lysine prices came up at a dinner meeting in Chicago between ADM and European executives. The latter complained about low prices and accused ADM of being responsible for it. ADM's Whitacre responded that "one can point a lot of fingers," and that the best thing to do was to find a solution to the problem.

The cartel has encountered the difficulty related to condition (3): they could not identify the deviator.⁴ What solution did they agree upon? Along the lines of Green and Porter (1984), they could have agreed upon a new target price, committing to go to a price war if the price fell below the new target. However, that would repeat the old story: as our results suggest, that would not have worked. Instead, the solution was to let output figures "be collected every month by the trade association (...) If one company sold more than it was allotted, it would be forced to purchase lysine from companies lagging behind," (Eichenwald, 2000, p.205). Thus, the cartel began collecting individual company data and aggregating that data over monthly periods, breaking conditions (1) and (3).

The practice of collecting data on market shares has been especially common among cartels.⁵ Even the Joint Executive Committee railroad cartel, the motivation for the Green and Porter (1984) model of equilibrium price wars,⁶ collected data on individual members' market shares. Other cartels have also limited the flexibility of its members to respond to new information by setting strict rules regarding acceptable forms of contracts with customers and by collecting data about suspected deviations through secret investigations (e.g. see the sugar trust cartel described in David Genesove and Wallace P. Mullin (2001)).

The failure of the lysine cartel to collude by setting a target price at the beginning of its operation illustrates how the provision of incentives can break down under flexible production even when firms have very clear information about the success of collusion. Figure 1 illustrates this surprising fact using a theoretical model presented in Section VI. In this example spot prices are correlated over time and have an unconditional mean of $1-Q$, where Q is the total quantity. Figure 1 shows two sample paths of prices when the firms produce Cournot Nash quantities and when they split the monopoly quantity.⁷ Just by looking at the price level at any moment of time, it is obvious whether the firms are colluding or not. Yet, despite this apparent transparency, collusion is impossible when firms see prices continuously and act sufficiently frequently, and when prices depend only on the total supply.




Figure 1

Why? A first guess may be that fast arrival of information and the flexibility to respond to it facilitate collusion, as firms can punish potential deviators more quickly. However, although this intuition is true in games with perfect monitoring, it does not always hold in games with imperfect monitoring, as demonstrated for the first time in the classic paper of Dilip Abreu, Paul Milgrom and David Pearce (1991) (hereafter AMP).⁸ Let us see why collusion is impossible under the wide range of conditions we study.

First, let us consider the classic case of a stationary repeated game with a strongly symmetric collusive scheme, in which firms behave identically after all histories. Following Abreu, Pearce and Ennio Stacchetti (1986) (hereafter APS'86), an optimal symmetric equilibrium has two regimes: a collusive regime and a price war/punishment regime. In the collusive regime, firms produce less than the static Nash equilibrium quantities. If the price drops below a critical level, this arouses enough suspicion of cheating to cause a price war. In the price war regime, firms strongly overproduce, because the intensity of the price war (i.e. low prices) makes the return to the collusive regime more likely.

With shorter time periods between actions, firms must decide whether to trigger a price war by looking at noisier incremental information. As we show, this causes firms to make type I errors by triggering price wars on the equilibrium path disproportionately often, erasing all benefits from collusion. To see this intuitively, let us compare games in which the time period between actions is either Δ or 2Δ , where Δ is small. Suppose that information arrives continuously, so that the aggregate summary statistic of the information in each period is normally distributed. In Figure 2 the horizontal axis illustrates the summary statistic (e.g. average price) for one period of length Δ , and the vertical axis, the summary statistic in the next period of length Δ . In an optimal symmetric equilibrium, at the end of each period firms test the summary statistic against a cut-off level to decide whether to trigger a price war. Figure 2 illustrates the critical regions 1 and 2 that trigger a price war in a game with period length Δ .

If the time period between moves increases to 2Δ , two forces influence the scope of collusion. First, firms learn more information per period, which helps collusion. Second, the gain from deviation in a

given period increases, which hurts collusion.⁹ From Figure 2 we learn that the first effect is stronger when Δ is small, i.e. collusion is more difficult with smaller time periods between moves. Indeed, note that at the end of a period of length 2Δ we can perform the same tests to decide about a price war as in a game with periods of length Δ , and these tests provide approximately the right incentives against a deviation with periods 2Δ .¹⁰ However, if firms can change actions only once per time interval 2Δ , a more effective joint test (illustrated by the dashed diagonal line) can provide the same incentives more efficiently. Therefore collusion is more efficient for time periods of length 2Δ .¹¹ As we show in this paper, with information arriving continuously, the loss efficiency for small Δ is so large that collusion becomes impossible altogether.




Figure 2

In Section III we show that collusion is not possible in asymmetric equilibria even if players use monetary transfers. In many games players can enforce collusive schemes without price wars by using asymmetric continuation equilibria or monetary transfers. Namely, rather than destroying value, they can transfer payoffs among themselves, keeping the average profits high. As we show, our environment causes such collusive schemes to fail, a result that involves new insights. The main intuition is that with a one-dimensional signal and a continuum of quantities to choose from, transfers used to provide incentives for one player interfere with the incentives of the other player. As a result, collusion cannot be sustained by transfers alone. The provision of incentives necessarily involves the destruction of value.¹²

Figure 3 shows a stylized relationship between the length of the period Δ and the scope of collusion in that model (see Figure 6 in Section IV for a detailed example). As we see the relationship is not monotonic and the highest collusive payoffs are achieved at an interior value of Δ . Moreover, as Δ becomes small, the highest equilibrium profits decrease to the stage-game Nash equilibrium profits.

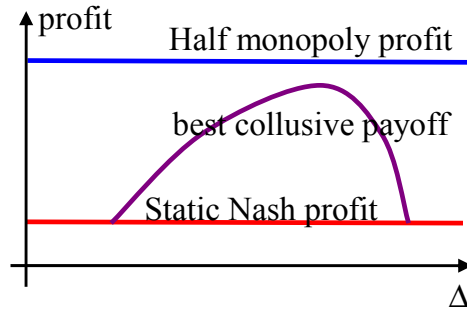


Figure 3

Besides looking at a classic stationary model of Cournot duopoly with constant marginal costs, we explore a number of extensions of applied interest. In Section V we look at the possibilities of increasing marginal costs and capacity constraints. In Section VI we extend the results to a more realistic non-stationary model with correlated prices. In Section VII we show that even if the monitoring technology is richer (so that deviations of the players affect signals differently), if there is a high correlation between the signals used to monitor the two players, the scope of collusion is substantially limited. Our results also extend to settings with more than two firms and multiple signals. Finally, the end of Section VII shows that our result holds even if we allow for a class of private strategies, which were shown in Michihiro Kandori and Ichiro Obara (2003) to greatly improve the players' payoffs in other environments.

Our analysis of cartels extends to cooperation in many other types of dynamic interaction, including for example moral hazard in teams. In particular, we present a model of a partnership in which efforts of team members are private while publicly observed outcomes depend (stochastically) on the sum of efforts of the partners. In that environment, cooperation over and above the static Nash equilibrium is not possible (as long as the marginal costs of effort are not too concave) if the partners observe the outcomes continuously and can react to information quickly.¹³

AMP have been the first to show that frequent actions can reduce the scope of collusion. They study symmetric equilibria in a Prisoner's Dilemma, in which just one type of signal arrives at a Poisson rate. This signal can be of a good type, indicative of cooperation, or a bad type, indicative of defection. When players act frequently, it is very unlikely that more than one signal arrives in a given period: effectively players observe at most one signal per period. If the signal is good, it cannot be used to trigger punishment, and cooperation becomes impossible with frequent actions. If the signal is bad, it can be used to trigger punishment, and limited cooperation is possible.

The nature of the signals is one of the fundamental differences between their model and ours. One is tempted to think that frequent actions should not hurt cooperation with Brownian signals. Tails of normally distributed signals are so informative about the players' actions that perfect collusion should be

possible if players are patient. Nevertheless, as $\Delta \rightarrow 0$ (i.e. as players become more patient and per-period information deteriorates) we show that surprisingly collusion is never possible (unlike with Poisson signals).¹⁴

Additionally, in comparison to AMP, we explore a number of other types of equilibria and many applied extensions. For example, we consider both symmetric and asymmetric equilibria, each possibly with monetary transfers. In fact, even though AMP do not explore this issue, in their setting asymmetric equilibria may achieve some cooperation even when no cooperation is possible in symmetric equilibria. We show that collusion becomes impossible even when one considers a non-stationary setting and compares collusive payoffs with those in Markov perfect equilibria. We also explore the issues of increasing marginal costs, multiple signals, mixed strategies, and even equilibria in private strategies.¹⁵

Another closely related paper is Radner, Myerson and Maskin (1986). They study a repeated partnership game in which in every period there are only two possible outcomes: a success or a failure. The probability of success depends on the sum of efforts. They show that the best equilibria in their setting are uniformly bounded away from efficiency for all discount rates. The intuition roughly consists of two parts: first, as in our paper, because it is difficult to distinguish between deviations of different players, asymmetric equilibria do not improve upon symmetric ones. Second, because the signals have bounded likelihood ratios, symmetric equilibria are necessarily costly. In contrast, in our setup the likelihood ratios are unbounded, so taking the interest rate to zero (while keeping Δ fixed) leads to asymptotically first-best cooperation. However, if we take Δ to zero, no cooperation is possible.

There is a large and growing theoretical literature trying to assess the impact of information on the scope of collusion in an environment with imperfect monitoring. Green and Porter (1984) were the first to propose a symmetric collusive equilibrium in which price wars are used on the equilibrium path to prevent deviations. Abreu, Pearce and Stacchetti (1986) characterize optimal symmetric equilibria in this setting and show that they involve two extreme regimes. The general analysis of games with hidden actions has been extended to asymmetric equilibria by Abreu, Pearce and Stacchetti (1990), Fudenberg, Levine and Maskin (1994) and Yuliy Sannikov (2004) among others.¹⁶

The paper is organized as follows. Section I presents a simple model of a repeated game with stationary prices. Section II proves the main result for symmetric public perfect equilibria. Section III proves the result for asymmetric equilibria with and without monetary transfers. Section IV presents a numerical example. Section V explores the possibility of nonlinear costs and capacity constraints. Section VI proves our impossibility result in a more complicated and realistic model in which market prices are correlated over time. Section VII discusses various modifications of the model. Section VIII concludes.

I. The Stationary Model

Two firms compete in a stationary market with homogenous products. The time horizon is infinite and firms discount future profits with a common interest rate r . Firms set quantities every period and the resulting prices depend on quantities and noise. We study how the scope of potential collusion depends on the firms' flexibility of production. We describe the flexibility of production in terms of Δ , the duration of time periods between which production decisions are made.

In a repeated game Γ^Δ , firms play the stage game at time points $t = 0, \Delta, 2\Delta, \dots$. The stage game is as follows: at time $t = n\Delta$ the firms choose (privately) supply rates $q_{it} \in [0, \bar{q}]$ for a time interval $[t, t+\Delta)$ (where \bar{q} is a large, exogenous capacity constraint). Denote by $Q_t = q_{1t} + q_{2t}$ the total supply rate of the two firms. During this time interval $[t, t+\Delta)$ firms supply ΔQ_t of the product. Both firms have the same constant marginal cost, which we normalize to zero.¹⁷ The profit of each firm is just the revenue, which depends on the price and the supply in a given period. Prices are publicly observable and depend on the total supply as well as a random shock

$$p_t = P(Q_t) + \varepsilon_t,$$

where the inverse demand function $P: [0, 2\bar{q}] \rightarrow \Re$ is strictly decreasing, twice continuously differentiable and $P(0) > 0$. Firm i 's revenue in time interval $[t, t+\Delta)$ is

$$\Delta q_{it} (P(Q_t) + \varepsilon_t)$$

To facilitate comparisons of equilibrium outcomes for different Δ , assume that random shocks come from a standard Brownian motion $\{Z_t; t \geq 0\}$. Shocks ε_t are formed in a way that makes total revenue of each firm depend not on Δ , but on the supply rates only. Specifically, the revenue that firm i gets from time 0 to t is

$$\sum_{s=0, \Delta, 2\Delta, \dots, t-\Delta} \Delta q_{is} (P(q_{1s} + q_{2s}) + \varepsilon_s) = \int_0^t q_{is} (P(q_{1s} + q_{2s}) ds + \sigma dZ_s),$$

where q_{is} with $s \in [n\Delta, n\Delta + \Delta)$ is the supply rate that firm i chooses at time $n\Delta$. and σ^2 is the variance of the noise. Then ε_t , the average price shock over time interval $[t, t+\Delta)$, is simply

$$\varepsilon_t = \frac{\sigma}{\Delta} (Z_{t+\Delta} - Z_t).$$

Therefore, $\varepsilon_t \sim N(0, \sigma^2/\Delta)$. Note that random shocks have greater variance over small time intervals, so that the information that firms learn by observing prices is proportional to the time interval. We can also interpret the noise structure in the following way. Suppose that market prices are quoted

every second and that these prices have a normal distribution with mean $P(Q)$ and variance σ^2 . If firms can change their supply rates every Δ seconds, then the average price over that interval has mean $P(Q)$ and variance σ^2/Δ (and as firms are risk neutral, these two interpretations of p_t are equivalent).

The fact that average prices can become unboundedly negative due to very high variance over short time periods is unattractive. However, it allows for a simple model in which the argument behind our main result is particularly clean.¹⁸ With the simple model we develop intuition that translates easily to more complex situations. See Section VI for a more realistic non-stationary model, with bounded instantaneous prices. Also, Section VII explains how our results extend to a stationary model in which the distribution of per-period prices is not necessarily normal. In particular, prices can be a nonnegative function of total supply and random shock, and, for example, follow a lognormal distribution.

A crucial assumption is that as in Green and Porter (1984) and Radner, Myerson and Maskin (1986) the price depends on the total supply only, so that the distribution of prices changes in the same way regardless of whether firm 1 or firm 2 increases its production by a unit, no matter what production levels they start with. That is, we assume that the products are homogenous. As demonstrated in Sections VI and VII, the assumption that the random price shocks are distributed normally and independently of past prices is not essential for the results.

In the repeated game firms choose supply rates after every history to maximize their expected discounted profit. Firm i 's (normalized) expected payoff is

$$E \left[\left(1 - e^{-r\Delta} \right) \sum_{t=0, \Delta, 2\Delta, \dots} e^{-rt} q_{it} (P(Q_t) + \varepsilon_t) \right],$$

where the expectation takes into account the dependence of future supply rates on past prices.

A history of firm i at time t is the sequence of price realizations and own supply decisions up to time t . A public history contains only the realizations of the past prices.¹⁹ We analyze the pure strategy public perfect equilibria (PPE) of the game. A strategy of a firm is public if it depends only on the public history of the game. Two public strategies form a PPE if after any public history the continuation strategies form a Nash equilibrium. Considering only pure strategies is admittedly restrictive. In Section VII we argue that the intuition holds also for public mixed strategies and for an important class of private mixed strategies that have been shown to improve collusive payoffs in other settings.²⁰

Throughout the paper we use the term PPE in the sense of pure strategy perfect public equilibria.

Remark 1. In the Abstract and the Introduction we have used extensively an informal phrase “information arrives continuously.” We can now explain this term formally. If prices are i.i.d. and the average price from time 0 to time t is a continuous function of t , then prices must take the form specified in our model (in particular, the noise has to be generated by a Brownian motion). Formally

speaking, if \bar{P}_t , the average price between times 0 and t is continuous, then $t\bar{P}_t$ is a Lévy process without jumps. Such a process can be represented as a sum of drift and diffusion terms, so average prices in each period are normally distributed.

Remark 2. It is important to note that the flexibility Δ plays two roles in our model: it affects the variance of the price distribution in a given period, and it affects the per-period discount rate $\delta \equiv e^{-r\Delta}$. The first effect makes collusion more difficult for smaller Δ since it results in less precise statistical inference, which makes deviations more difficult to detect. The second effect (standard in repeated games) makes it easier to collude because the single-period benefits to deviation become small relative to the continuation payoffs.

Remark 3. The methods and results for the repeated duopoly game are applicable far beyond the realm of collusion. To give an illustration, let us present a stylized model of moral hazard in teams, in which payoffs above the static Nash equilibrium are not achievable as $\Delta \rightarrow 0$ for the same reasons as in a repeated duopoly setting. Two partners choose effort rates (q_{1t}, q_{2t}) that affect the partnership's profit. The average profit per period (gross of private costs of effort) is

$$R_t = P(q_{1t} + q_{2t}) + \varepsilon_t$$

where $P(\cdot) > 0$ and $\varepsilon_t \sim N(0, \sigma^2 / \Delta)$. The payoffs to partner i are

$$(1 - e^{-r\Delta}) \sum_{t=0, \Delta, 2\Delta, \dots} e^{-rt} (R_{it} - c(q_{it}))$$

where R_{it} is the share of profit flow that player i receives at time t and $c(q)$ is the cost of effort (increasing and convex). Effort by player i benefits both players, but only player i pays its cost. This externality implies that the static Nash equilibrium does not yield joint profit maximization. Even though there is no direct mapping between this model and the duopoly model we study, the assumptions that information arrives *continuously* and that total profit depends only on *joint* effort drive the result that as $\Delta \rightarrow 0$, the highest payoffs the partnership can achieve are the static Nash equilibrium payoffs.²¹ The stage game is similar to Radner, Myerson and Maskin (1986), with the main differences being that we have normally distributed profits, while they have a binary distribution, and that we take $\Delta \rightarrow 0$ while they take $r \rightarrow 0$.

A. Structure of the Stage Game

We make the following two assumptions about the inverse demand function:

- **A1:** The marginal revenue (of total demand) is decreasing: $\frac{\partial^2}{\partial Q^2}(QP(Q)) < 0$.

- **A2:** The static best response to any $q \in [0, \bar{q}]$ is less than \bar{q} , i.e. the marginal revenue of

$$\text{residual demand is negative at } \bar{q} : \left. \frac{\partial(q'P(q'+q))}{\partial q'} \right|_{q'=\bar{q}} < 0.$$

Assumptions A1 and A2 are sufficient to guarantee that the best response in the stage game, $q^*(q)$, is unique and that the Nash equilibrium is symmetric and unique. Denote the Nash equilibrium of the stage game by (q_N, q_N) . We refer to it also as the ‘competitive equilibrium’ or the ‘static equilibrium’ especially when we talk about the repetition of (q_N, q_N) in each stage of the dynamic game. Define $v_N = P(2q_N)q_N$.

LEMMA 1. Assume A1 and A2. The static best response $q^*(q)$ is unique and less than \bar{q} . Also, the static Nash equilibrium is symmetric and unique.

The proof is standard and can be found in the Appendix.

II. Impossibility of Collusion: Symmetric Equilibria

In this section, we prove that collusion becomes impossible as $\Delta \rightarrow 0$ in *symmetric* PPE. From APS’86 we know that optimal symmetric PPE involve two regimes: a collusive regime and a price war regime. In each period, the decision whether to remain in the same regime or to switch is guided by the outcome of the price in that period alone. We show that as $\Delta \rightarrow 0$, transitions from the collusive regime to the price war regime must happen very frequently to provide the players with incentives. Because of that, price wars destroy all collusive gains, and payoffs above the static Nash become impossible in the limit. The methods we develop in this section are important for *asymmetric* PPE (possibly with monetary transfers) as well, but the analysis of those will require further insights. Denote by $[\underline{v}(\Delta), \bar{v}(\Delta)] \subset \mathfrak{R}$ the set of payoffs achievable in symmetric PPE in the game Γ^Δ . We would like to show that as $\Delta \rightarrow 0$, $\bar{v}(\Delta)$ converges to the ‘competitive payoff’ $v_N = P(2q_N)q_N$.

First, let us informally present the core of our argument. Using the results of APS’86, in the collusive regime players expect total payoffs $\bar{v}(\Delta)$ and choose some supply rates $(q, q) < (q_N, q_N)$ in the current period. In the next period, depending on the realized price, players either stay in the collusive regime with continuation payoff $\bar{v}(\Delta)$ or go to a price war with continuation payoff $\underline{v}(\Delta)$. The lower the punishment $\underline{v}(\Delta)$ is, the higher the collusive payoffs $\bar{v}(\Delta)$ can be: harsher punishments make providing incentives easier. We will show that $\bar{v}(\Delta)$ converges to v_N even if players can use a punishment of 0 (clearly $\underline{v}(\Delta) \geq 0$ because 0 is the minmax payoff – a firm can guarantee itself that payoff by producing 0). As $\Delta \rightarrow 0$, collusion becomes impossible because the statistical test to prevent deviations sends players to a price war with a disproportionately high probability.

To see this in greater detail, consider a deviation from q to a static best response $q^*(q) > q$ that reduces the mean of the observed average price from $\mu = P(2q)$ to $\mu' = P(q^*(q)+q) < \mu$. Lemma 3 below shows that the best statistical test to prevent this deviation is a tail test, which triggers a punishment when the price falls within a critical region $(-\infty, c]$. Given such a test, a deviation increases the probability of punishment by

$$\text{likelihood difference} = G'(c) - G(c),$$

where G' and G are normal cumulative distribution functions with variance σ^2 / Δ and means μ' and μ respectively. The probability of type I error, i.e. triggering punishment when no deviation has occurred, is given by the size of the test:

$$\text{size} = G(c)$$

Because the gain from a deviation in one period is on the order of Δ , we explore tail tests with a likelihood difference on the order of Δ . Lemma 2 shows that in such tests, as $\Delta \rightarrow 0$, the probability of making a type I error in each period blows up relative to Δ . Therefore, as illustrated in Figure 4, $\bar{v}(\Delta)$ cannot be sustained. The reasoning is as follows. The total expected payoff is a weighted sum of the current period payoff and the expected continuation payoff with weights equal to $(1-\delta)$ and δ . The continuation payoff is a weighted average of $\bar{v}(\Delta)$ and 0 (recall that we allow for punishments that are at least as harsh as the equilibrium ones, which relaxes the problem) with weights $1-G(c)$ and $G(c)$ respectively:

$$\bar{v}(\Delta) = (1-\delta)qP(2q) + \delta(\bar{v}(\Delta)(1-G(c)) + 0 \cdot G(c))$$

That is represented graphically as:




Figure 4

Now, as $\Delta \rightarrow 0$, we have two effects: the payoff gain from colluding in the current period $(1-\delta)(qP(2q) - q_N P(2q_N))$ becomes small, and the payoff loss due to punishment becomes small.

However, the payoff loss of $\delta \bar{v}(\Delta)G(c)$ becomes disproportionately large compared to the current-period payoff gain in the limit. As a result, any payoff above Nash cannot be sustained even if we both

- worry about only one deviation to a static best response and
- allow 0 as the harshest punishment, even though typically $\underline{v}(\Delta) > 0$.

When comparing games with different Δ one may be tempted to employ the following reasoning: consider two games with respective time periods Δ and $\Delta' = \Delta/2$. Suppose that with Δ we can construct a profitable collusive equilibrium. Now, moving to Δ' , instead of employing the factorization techniques of APS (namely, that it is sufficient to factorize the game into current period and continuation payoffs) consider the following strategies as a candidate for equilibrium. In the odd periods firms are recommended to follow the same strategies as in the game with Δ , using the average price over the last two periods to decide on continuation play. In the even periods the recommendation is to ignore the current prices and keep the quantities fixed from last period.

Clearly, if firms follow these recommendations, they will obtain the same payoffs as in the game with Δ . Unfortunately, these recommendations are not incentive compatible for several reasons. First, even if firm 2 follows the recommendation, firm 1 after seeing a high first period price has incentives to increase quantity in the second period and after seeing a low price to decrease quantity. That “option value” will also make firm 1 increase its quantity in the first period. Second, by that same reasoning firm 1 will expect firm 2 to react in the second period to the observed first-period price. That makes firm 1 increase its first-period quantity even more: firm 1 gets all the benefit, while the cost of reducing quantities in the second period if prices turn out to be low will be shared between the two firms.²²

A. Formal Argument

First, we prove that as $\Delta \rightarrow 0$, the probability of type I error under a tail test to prevent deviations becomes disproportionately high.

LEMMA 2. Fix $C_1 > 0$ and $\mu - \mu' > 0$. If $\Delta > 0$ is sufficiently small, then a tail test for a deviation with a likelihood difference of $C_1\Delta$ has a probability of type I error greater than $O(\Delta^{0.5+\varepsilon})$ for any $\varepsilon > 0$.

PROOF: See Appendix.

Next, we want to show that a tail test is an optimal way to deter any given deviation. Consider a deviation with instantaneous gain $D\Delta$, which reduces the mean of observed prices from μ to μ' . The optimal test maximizes the expected continuation payoff subject to providing incentives against this deviation. Lemma 3 shows that if we are worried about only one deviation, a tail test with a bang-bang property is best.

LEMMA 3. Suppose $D > 0$. Consider the problem

$$\begin{aligned} & \max_{v(x)} \int_{-\infty}^{\infty} v(x) g(x) dx \\ & s.t \\ & D\Delta \leq \delta \int_{-\infty}^{\infty} v(x) (g(x) - g'(x)) dx \quad \text{and} \quad \forall x \in \mathfrak{R} \quad v(x) \in [\underline{v}, \bar{v}] \end{aligned}$$

where g and g' are the densities of normal distributions with variance σ^2/Δ and means μ and $\mu' < \mu$ respectively. If this problem has a solution, then it takes the form

$$(1) \quad v(x) = \begin{cases} \bar{v} & \text{if } x > c \\ \underline{v} & \text{if } x \leq c \end{cases}$$

for some $c \in (-\infty, (\mu + \mu')/2]$.

PROOF: See Appendix.

With the help of Lemmas 2 and 3 we are ready to formulate our main result.

PROPOSITION I. As the time period between actions converges to 0, the maximal payoff achievable in symmetric PPE converges to the static Nash equilibrium payoff, i.e. $\bar{v}(\Delta) \rightarrow v_N$ as $\Delta \rightarrow 0$.

PROOF: See Appendix.

As $\Delta \rightarrow 0$, two effects influence the scope of collusion: more frequent moves lead to both weaker statistical tests (which makes collusion more difficult to sustain) and smaller immediate benefits to deviation (which makes collusion easier to sustain). The proof shows that as $\Delta \rightarrow 0$ the deterioration of information is so extreme that collusion becomes impossible altogether.

We finish this section by comparing the results to AMP, where information arrives discontinuously via Poisson jumps. Suppose that instead of continuous monitoring via prices the players can monitor deviations through public signals that arrive according to a Poisson arrival rate, which depends on the total supply Q . To facilitate comparison, focus on a single deviation to a static best response. This makes the setup very similar to the repeated Prisoners' Dilemma game in AMP. There are two cases to consider with such discontinuous monitoring: the *bad news* case (arrival rate increasing in Q , so that arrival rate is higher upon deviation) and the *good news* case (arrival rate decreasing in Q). In the best collusive symmetric equilibria, incentives are provided quite differently in these two cases: in the bad news case price wars are triggered when the signal arrives, while in the good news case they are triggered when the signal does not arrive.

It is very difficult to provide incentives in the good news case even for a fixed small Δ . In fact, for small Δ , the probability of type I error required to provide incentives is so high that even as $r \rightarrow 0$, no

collusion is possible in AMP (AMP Proposition 4). In contrast, in our model, for any given Δ , as $r \rightarrow 0$ the best equilibrium converges to first-best collusion. The reason is that with a smaller r , harsher punishments allow players to use more efficient statistical tests with cutoffs further in the tails of the price distribution. On the other hand, in the bad news case incentives can be provided much more efficiently than in our model. Conditioning punishment on exactly one signal arriving during the period²³ allows for statistical tests with likelihood difference *and* type I error probability both on the order of Δ . In fact the ratio of these two remains roughly constant as Δ changes, (unlike in our model) because the probability of one signal arriving is approximately Δ times the appropriate arrival rate, so that the optimal probability of going to a price war conditional on arrival is roughly constant for all Δ . That difference makes collusion possible in the AMP model even as $\Delta \rightarrow 0$ (AMP Proposition 5), while no collusion is possible in our model, i.e. when monitoring is based on information arriving continuously.

III. Asymmetric Equilibria and Monetary Transfers

As we discussed in the Introduction, in many games asymmetric equilibria achieve much higher payoffs than symmetric ones. The reason is that in symmetric equilibria incentives are provided via price wars, and hence they require destruction of value. In asymmetric equilibria incentives can often be provided by transferring payoffs between players, keeping the sum of profits high. Can firms sustain some collusion in our setting using asymmetric strategies? We claim the answer is no.

To see the intuition behind this claim, let us informally attempt to construct a natural asymmetric equilibrium and show why it fails. Consider a hypothetical collusive scheme in which one of the players always produces the monopoly quantity while the other produces nothing, and they change roles over time. One may hope that this collusive scheme can be made successful, because only the player who produces nothing is tempted to deviate in any period. We can then provide incentives against that deviation by letting the opponent, who is in the role of a monopolist, stay in that role longer if prices drops below a cut-off c . However, this arrangement fails because it interferes with the incentives of the monopolist: he becomes tempted to overproduce to stay in the role of the monopolist longer. Therefore, the scheme fails to keep the total supply low. The reason for such failure, which is formalized in the proof of Proposition II, is that the individual incentives created by transfers of future payoffs interfere with one other and cause the players to jointly produce the Cournot Nash sum of quantities.




Figure 5

The proof that collusion is impossible in asymmetric equilibria depends crucially on the assumption that deviations of different players cannot be statistically distinguished. The intuition is as follows. Consider a pair of quantities (q_1, q_2) with $Q = q_1 + q_2 \leq 2q_N$ and potential upward deviations of ε by each player. Both such deviations change the density of prices from g to g' and therefore are observationally equivalent. The incentive constraints against those deviations are

$$(2) \quad \begin{aligned} r\Delta((q_1 + \varepsilon)P(Q + \varepsilon) - q_1P(Q)) &\leq \int_{-\infty}^{\infty} v_1(x)(g(x) - g'(x))dx \\ r\Delta((q_2 + \varepsilon)P(Q + \varepsilon) - q_2P(Q)) &\leq \int_{-\infty}^{\infty} v_2(x)(g(x) - g'(x))dx \end{aligned}$$

where $v_1(x)$ and $v_2(x)$ are next-period continuation values as a function of price x . Adding up the two constraints and dividing by 2, we get a *joint constraint*

$$(3) \quad r\Delta\left(\left(\frac{Q}{2} + \varepsilon\right)P(Q + \varepsilon) - \frac{Q}{2}P(Q)\right) \leq \int_{-\infty}^{\infty} \frac{v_1(x) + v_2(x)}{2}(g(x) - g'(x))dx$$

This joint constraint is identical to the incentive constraint in a symmetric equilibrium with quantity pair $(Q/2, Q/2)$ that counters an upward deviation of ε . In order to satisfy this constraint (while sustaining current period collusive profits), the problem is that too much value has to be destroyed per period for $\Delta \rightarrow 0$, as we argued in the previous section.²⁴ Proposition II formalizes the preceding argument. Let $\bar{v}_a(\Delta)$ be the highest sum of payoffs that the players can achieve in equilibrium.

PROPOSITION II. As the time period between actions converges to 0, the maximal sum of payoffs achievable in any PPE converges to twice the static Nash equilibrium payoff, i.e. $\bar{v}_a(\Delta) \rightarrow 2v_N$ as $\Delta \rightarrow 0$.

PROOF: See Appendix.

A. Monetary transfers

So far we have assumed that the players cannot use any monetary transfers. This is a realistic assumption for some cartels (especially tacit ones). However, if we apply our analysis to other environments (for example, partnerships) in which transfers are legal and/or often used, we should relax that assumption. It turns out that the proof of Proposition II easily extends to show that collusion becomes impossible as $\Delta \rightarrow 0$ even if firms can use monetary transfers.

To allow for transfers, we assume that the players can write binding contracts that specify a transfer of money conditional on the observed realization of prices. Such a contract specifies an amount of money $t_i(x)$ that each firm is supposed to pay (or receive if t_i is negative) depending on the price realization. The

case of *balanced transfers*, i.e. $t_1(x) + t_2(x) = 0$, is particularly simple. The incentive constraints against an upward deviation of ε by each player separately are now

$$(4) \quad \begin{aligned} r\Delta((q_1 + \varepsilon)P(Q + \varepsilon) - q_1P(Q)) &\leq \int_{-\infty}^{\infty} (v_1(x) - t_1(x))(g(x) - g'(x))dx \\ r\Delta((q_2 + \varepsilon)P(Q + \varepsilon) - q_2P(Q)) &\leq \int_{-\infty}^{\infty} (v_2(x) - t_2(x))(g(x) - g'(x))dx \end{aligned}$$

Adding those two constraints and dividing by 2, we obtain the same joint constraint as (3):

$$r\Delta\left(\left(\frac{Q}{2} + \varepsilon\right)P(Q + \varepsilon) - \frac{Q}{2}P(Q)\right) \leq \int_{-\infty}^{\infty} \frac{v_1(x) + v_2(x)}{2}(g(x) - g'(x))dx$$

From this point, the proof of Proposition **II** applies to balanced transfers without any change. The intuition behind this result is the same: the provision of joint incentives destroys too much value, as measured by the sum of the players' payoffs.

If we allow *unbalanced transfers* $(t_1(x), t_2(x))$, we assume that they satisfy $t_1(x) + t_2(x) = b(x) \geq 0$, where $b(x)$ is the amount of money burning (e.g. payments made to a third party). Unbalanced transfers can improve somewhat the scope of collusion for a fixed Δ , as they allow for punishments stronger than the previous minmax payoff. However, our limit results still hold as long as the amount of money burning per period stays bounded. The reason is that as long as punishment is bounded, the likelihood difference for a statistical test to prevent a particular deviation must be still on the order of Δ . Such tests trigger punishment with a disproportionately high probability on the equilibrium path.

We summarize our findings in the following corollary.

Corollary. Even if monetary transfers are allowed (but the amount of money burning is bounded), as the time period between actions converges to 0, the maximal sum of payoffs achievable in any PPE converges to twice the static Nash equilibrium payoff, i.e. $\bar{v}_a(\Delta) \rightarrow 2v_N$ as $\Delta \rightarrow 0$.

PROOF: See Appendix.

We conclude this section with two observations related to bounds on money burning.

REMARK. Suppose that the players can sign binding contracts that specify a transfer $t_i(x)$ from player i to a third party when the realization of the signal is x , with $t_1(x) + t_2(x) \geq 0$. Then:

- i) If there is no bound on the transfers, then for any $\Delta > 0$ and $\varepsilon > 0$ there exists an equilibrium in which $\bar{v}_a(\Delta) \geq v_M - \varepsilon$ (where v_M is the monopoly profit).
- ii) If the amount of money burning is bounded by current wealth and the players start the game with finite wealth, then $\bar{v}_a(\Delta) \rightarrow 2v_N$ as $\Delta \rightarrow 0$ where $\bar{v}_a(\Delta)$ is calculated at the beginning of the game.

PROOF: See Appendix.

IV. A Numerical Example

We now provide a numerical example that illustrates the effects of changes in Δ on $\bar{v}(\Delta)$. The inverse demand is $P(Q) = 1 - Q$. It corresponds to a monopoly supply $Q = 1/2$ and expected price $P(Q) = 1/2$. In the static Nash (Cournot) equilibrium, quantities are $q_N = 1/3$, average price is $P(2q_N) = 1/3$ and payoffs are $v_N = 1/9$.

We take $r = 10\%$ to be the annual interest rate, so we interpret $\Delta=1$ as the firms making supply rate decisions every year. We take $\sigma = 0.03$, so that the daily standard deviation of average price is $\sigma / \sqrt{\Delta} = 0.03 / \sqrt{1/365} \approx 0.57$, which is on the order of the average daily monopoly price.

We do not compute the actual best collusive payoff, because it depends on whether we consider symmetric or asymmetric equilibria and various details of the model (e.g. the production capacity \bar{q} , which is directly important for the worst equilibria, and hence indirectly the highest equilibrium payoffs). Instead, we compute a robust lower bound by finding the best symmetric equilibrium with Nash reversion as a punishment, and a robust upper bound by finding the best symmetric equilibrium with the minmax payoff of 0 as a punishment. Both bounds are valid for both symmetric and asymmetric equilibria,²⁵ as well as collusion enforced by balanced monetary transfers. Figure 6 shows the two bounds. The horizontal axis is $\log_{10}(\Delta)$. The value of 0 corresponds to a move once a year, the value of -2, 100 times a year.

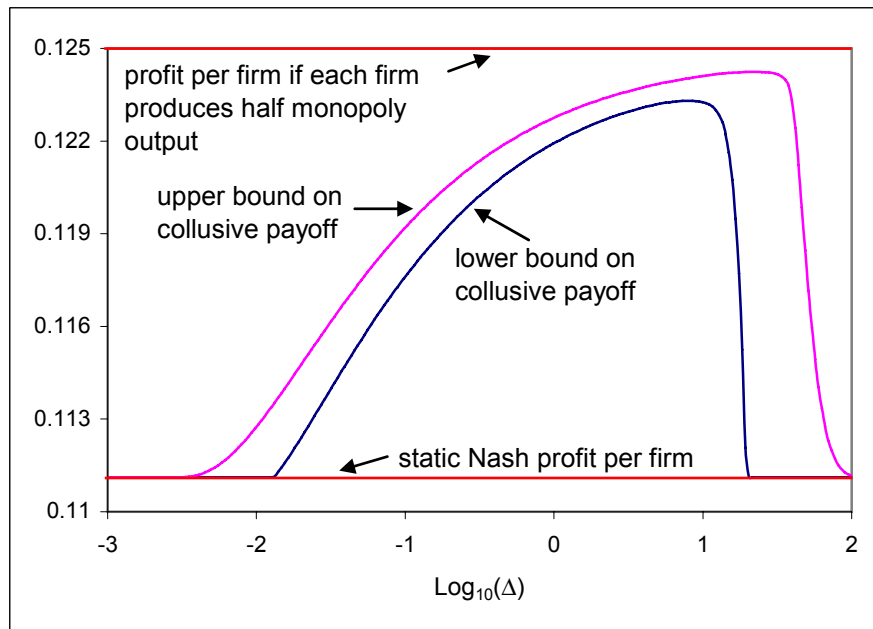


Figure 6

The behavior on the right hand side of the graph is standard: as the moves get less frequent the benefits from one-period deviation are large compared to continuation payoffs and collusion is difficult to sustain. The effect we focus on is the non-monotonicity of $\bar{v}(\Delta)$ in Δ : as we move to frequencies higher than once a year (we move to the left on the graph), $\bar{v}(\Delta)$ decreases. In fact if the players move 100 times a year, then $\bar{v}(\Delta) \leq 0.113$, which is at most a 1.5% improvement over the static Nash payoffs, while half of monopoly profit would represent a 12.5% improvement.

V. Convex Costs and Capacity Constraints.

In this section we extend our results to the case of increasing marginal costs. Recall that in our basic model, we assume constant marginal costs and normalize them to zero. Suppose that the costs of production accrue at rate $c(q)$ when the output flow is q , so that per-period profits are

$$\Delta(q_{it}(P(Q_t) + \varepsilon_t) - c(q_{it})).$$

Assume that costs are increasing and convex:

- **A3:** $c(0) = 0$, $c'(q) \geq 0$, $c''(q) \geq 0$

We show that with increasing marginal costs under A3 still no collusion can be sustained in symmetric equilibria. The scope of collusion in asymmetric equilibria depends on the convexity of marginal costs (i.e. on the third derivative of c). If they are not too convex, then even asymmetric equilibria with monetary transfers do not improve upon the static Nash payoffs. Otherwise, as we show in the extreme case of capacity constraints, collusion may be sustainable with transfers. The analysis of this section is brief as it directly applies the methodology of Sections II and III.

A. Symmetric Equilibria.

Let us summarize the argument why collusion is impossible in a symmetric equilibrium as $\Delta \rightarrow 0$. To improve upon the static Nash equilibrium payoffs, firms must use a pair of quantities (q, q) less than Nash. At any such pair, both firms can improve upon static payoffs by deviating upwards. A statistical test to prevent such deviations must have a likelihood difference on the order of Δ . As $\Delta \rightarrow 0$, such a test gives false positives with a disproportionately high probability, $O(\Delta^{0.5+\varepsilon})$, and destroys all the benefits of collusion.

We now argue that this argument still applies with increasing marginal costs. The symmetric static Nash equilibrium is unique. It is given by the first-order condition

$$(5) \quad P'(2q_N)q_N + P(q_N) - c'(q_N) = 0,$$

and note that by A1 and A3 the LHS of (5) is decreasing. When the firms produce (q, q) , the sum of profits is given by

$$\Pi(2q) = P(2q)q - 2c(q).$$

By A1 and A3 this is a strictly concave function. As $\Pi'(2q_N) < 0$, firms must choose quantities less than q_N to improve upon the static Nash equilibrium profits. But because the derivative of each firm's static profit with respect to quantity $P'(2q)q + P(q) - c'(q)$ is positive when $q < q_N$, each firm is tempted to overproduce. Moreover, the gain D from the most profitable static deviation is uniformly bounded away from 0 by some number $\varepsilon_\pi > 0$ for all $q < q_N - \varepsilon$. Now the argument in the proof of Proposition **I** applies directly. Quantities $q < q_N - \varepsilon$ cannot be used for collusion due to false positives as $\Delta \rightarrow 0$. Letting $\varepsilon \rightarrow 0$ shows that collusion becomes impossible. Therefore even with convex costs symmetric equilibria cannot improve upon the static Nash payoffs:

PROPOSITION III. Suppose total costs satisfy A3 (i.e. they are increasing and convex). As the time period between actions converges to 0, the maximal payoff achievable in symmetric PPE converges to the static Nash equilibrium payoff, i.e. $\bar{v}(\Delta) \rightarrow v_N$ as $\Delta \rightarrow 0$.

B. Asymmetric Equilibria with Monetary Transfers.

Can asymmetric equilibria help sustain collusion? With increasing marginal costs, asymmetric quantities are inefficient and may help collusion only if asymmetries significantly facilitate the provision of incentives. In most cases, the provision of asymmetric incentives is just as difficult as it is with constant marginal costs, and the provision of joint incentives destroys too much value. In a limited set of cases, e.g. with capacity constraints as discussed in the next subsection, collusion may be possible.

Let us go through the details of our argument. First, note that if firms move from a symmetric quantity pair $(Q/2, Q/2)$ to an asymmetric pair (q_1, q_2) with $Q = q_1 + q_2$, their total profits are hurt due to convex total costs (A3), i.e.

$$\Pi(q_1, q_2) = QP(Q) - c(q_1) - c(q_2) \leq P(Q)Q - 2c(Q/2).$$

Therefore, as in the case of symmetric quantity pairs, the sum of profits from (q_1, q_2) can be higher than twice the static Nash payoffs only if $Q = q_1 + q_2 < 2q_N$. Following our reasoning in Proposition **II**, we consider the sum of the firms' IC constraints when each considers increasing its quantity by some ε . The sum of derivatives of single firms' profits at (q_1, q_2) is

$$(6) \quad \frac{\partial \pi_1(q_1, q_2)}{\partial q_1} + \frac{\partial \pi_2(q_1, q_2)}{\partial q_2} = P'(Q)Q + 2P(Q) - c'(q_1) - c'(q_2).$$

If for any (q_1, q_2) such that $\Pi(q_1, q_2) - \Pi(q_N, q_N) > 0$ that sum is strictly positive, then asymmetric quantities do not facilitate incentive provision. The intuition is that because signals depend only on aggregate quantity, balanced transfers (of continuation values or money) between the firms can provide

incentives for one of the firms only at the cost of interfering with incentives of the other firm. Indeed, adding up two IC constraints for an upward deviation by each firm of ε , we obtain a *joint constraint* analogous to (3), where balanced transfers cancel out. Therefore, as in the proof of Proposition II, jointly providing incentives requires a disproportionately large destruction of value. Therefore if the value of (6) is positive whenever $\Pi(q_1, q_2) - \Pi(q_N, q_N) > 0$, collusion is impossible as $\Delta \rightarrow 0$.

A sufficient condition for this to be the case is

- **A4:** $c'''(q) \leq 0$ i.e. marginal costs are weakly concave.

A4 guarantees that asymmetric quantities are more difficult to support in equilibrium than symmetric ones (and, due to increasing marginal costs, they are also less efficient than symmetric quantity pairs). Indeed, if A4 holds, then for a fixed total quantity $Q = q_1 + q_2 \leq 2q_N$, the sum of derivatives at asymmetric quantities is higher than the sum at symmetric quantities:

$$P'(Q)Q + 2P(Q) - c'(q_1) - c'(q_2) \geq P'(Q)Q + 2P(Q) - 2c'(Q/2) \geq 0.$$

We summarize our conclusion in a proposition.

PROPOSITION IV Suppose total costs satisfy A3 and A4. Then as the time period between actions converges to 0, the maximal sum of payoffs achievable in any PPE converges to twice the static Nash equilibrium payoff, i.e. $\bar{v}_a(\Delta) \rightarrow 2v_N$ as $\Delta \rightarrow 0$.

Remark: Moral Hazard in Teams. In Section I we have presented a model of moral hazard in a repeated partnership and have said that our analysis can be applied to that model as well. In particular, Proposition IV holds for that model. In the partnership game sufficient assumptions that correspond to A1-A3 are that profits are increasing in total effort, weakly concave and twice differentiable, and that the private costs of effort are increasing and convex. Finally, an assumption analogous to A4 is that marginal costs of effort are weakly convex.²⁶

C. Capacity Constraints.

We finish this section by showing that if A4 is violated then sometimes collusion is possible with asymmetric equilibria that use balanced monetary transfers. We illustrate this fact in an extreme case of capacity constraints. In this case the concavity assumption A4 is violated because marginal costs jump to infinity at capacity. We show that although the result for symmetric equilibria is very robust, for asymmetric equilibria it is somewhat delicate. Technically, once one of the firms is at capacity, the sum in equation (6) is minus infinity, and therefore it is possible to provide incentives via transfers to one of the firms without disturbing incentives for the firm producing at capacity.

For example, suppose that both firms have zero marginal costs and that capacity constraints are at the monopoly level of production. Then it is still not possible to sustain profitable collusion with symmetric equilibria. In contrast, asymmetric equilibria with monetary transfers solve the monitoring problem. The way to build such equilibria is to allow the firms to alternate between being a monopolist and producing nothing. When a firm produces nothing, its incentives come from punishments when the price is low and rewards when the price is high. Unlike in Section III, these incentives can be provided only with monetary transfers between the firms, without money burning. Because the firm that acts as a monopolist is at capacity, it cannot overproduce to receive transfers by driving prices down. Through that mechanism, capacity constraints help us resolve the problem of separating deviators from non-deviators.

To see in greater detail how a collusive equilibrium may look, consider the following strategies. In odd periods firm 1 is recommended to produce at capacity (q_M) and firm 2 to produce nothing. In even periods the roles are reversed. After every period the firm that is supposed to produce nothing receives a transfer of $\Delta q_M (P(q_M) + \varepsilon_t)$, i.e. the realized profit (assuming no deviation). In such a scheme the current monopolist clearly has no incentives to deviate – lower quantity means lower profits and higher transfers. The other firm, which maximizes $\Delta(q_M + q_{it})(P(q_M + q_{it}) + \varepsilon_t)$, has no incentives to deviate either. Therefore the proposed strategies form an equilibrium (assuming the transfers can be enforced). That scheme achieves first-best.

This scheme requires the transfer to be very large after some extreme realizations; hence if transfers cannot be contractually enforced, players will have incentives to deviate. This can be resolved by truncating the realized price at 10 standard deviations away from the mean and slightly increasing the transfer within this range. This preserves incentives and keeps transfers on the order of $\sqrt{\Delta}$. Note that firms do not need to sign a binding contract to enforce such transfers: they can be enforced by a threat of Nash reversion. For small enough Δ this threat is credible and prevents deviations of players refusing to pay the transfer: since the equilibrium achieves first best, the expected continuation payoffs are almost half the monopoly profits forever. Not paying the transfer saves much less than the benefit of continued cooperation.

Remark. This setup has three ingredients that allow us to sustain collusion: monetary transfers, asymmetric equilibria and capacity constraints. We know that without any of the last two, collusion is not possible for small Δ . It can also be shown that the monetary transfers are necessary as well; in other words, using only transfers of continuation payoffs (say, promises to be a monopolist longer in the future) is not sufficient to provide incentives (for small Δ). This may be surprising since in repeated games when players are sufficiently patient, often transfers of continuation payoffs are asymptotically as good as monetary transfers.

The intuition behind the result is as follows. First, in order to sustain the highest sum of payoffs, in the current period one of the players must be at capacity and the other player must choose strictly less than his static best response. In order to provide incentives for this player not to overproduce, the transfers have to be at least on the order of $\sqrt{\Delta}$. The second part of the reasoning is about the largest difference in continuation payoffs that can be achieved in equilibrium for small Δ . Using arguments similar to the ones we have used so far, one can show that this upper bound on the continuation payoff transfer decreases to zero at a rate much faster than $O(\sqrt{\Delta})$.²⁷ To put it differently, the whole set of continuation payoffs rapidly collapses towards the 45 degree line as Δ gets small. Combining the two steps, for very flexible production, the range of feasible transfers via continuation payoffs alone is too small to provide incentives to reduce quantity below the static best response.

VI. Correlated Demand

We have presented the simplest duopoly model with homogenous goods to illustrate the idea that collusion becomes impossible if firms can act frequently. This standard model with i.i.d. prices is motivated by Green and Porter (1984). As firms act frequently, a few unrealistic features arise in this model if we keep it simple. For example, prices are unbounded (positive and negative) and the variance of per-period average prices grows to infinity as we take $\Delta \rightarrow 0$. Here we illustrate how the intuition that we have developed extends to a more complicated and realistic model.

In particular, assume that prices follow a mean-reverting process with mean that depends on the total supply:

$$p_{t+\Delta} = p_t + \Delta f(Q, p_t) + \varepsilon_{t+\Delta},$$

where $\varepsilon_{t+\Delta} \sim N(0, \Delta \sigma^2(p_t))$.²⁸ This price process is different from the previous model; in particular, it is non-stationary and the conditional variance of per-period prices decreases to zero as $\Delta \rightarrow 0$. However, the impact on prices of a one-period deviation is on the order of Δ and the standard deviation of noise is on the order of $\sqrt{\Delta}$. As the ratio of these two orders is the same as in the stationary model, similar issues arise: it is more difficult to detect deviations over smaller periods of time. Assume that $f_1 < 0$, $f_{11} < 0$, $f_2 < 0$ and that $f(Q, \cdot)$ has root $P(Q)$, the mean of prices when the total supply is Q .²⁹

The timing is as before: at periods $t \in \{0, \Delta, 2\Delta, \dots\}$ firms observe the current price p_t and decide on supply rates to be held constant for Δ units of time. The first price p_0 is given exogenously. Per-period profits are $(1 - \delta) q_t p_{t+\Delta}$.³⁰

To understand the potential for collusion in this model, denote by $F_M(p_t)$ the monopoly profit. A monopolist's Bellman equation when the current price is p_t is

$$F_M(p_t) = \max_Q E[(1-\delta)Qp_{t+\Delta} + \delta F_M(p_{t+\Delta})],$$

and the first-order condition is $E\left[(1-\delta)p_{t+\Delta} + \Delta f(Q, p_t)((1-\delta)Q + \delta F_M'(p_{t+\Delta}))\right] = 0$. Monopoly profit cannot be achieved in an equilibrium with two firms. Indeed, if a firm anticipated half the monopoly profit as its continuation value, currently it must solve

$$\max_q E\left[(1-\delta)qp_{t+\Delta} + \delta \frac{1}{2} F_M(p_{t+\Delta})\right].$$

Taking the first-order condition at half the monopoly quantity, we get

$$E\left[(1-\delta)p_{t+\Delta} + \Delta \underbrace{f_1(Q, p_t)}_{<0}((1-\delta)Q + \frac{1}{2} \delta \underbrace{F_M'(p_{t+\Delta})}_{>0})\right] > 0.$$

Therefore, each firm would be tempted to overproduce.

We argue that collusion is impossible in this setting as $\Delta \rightarrow 0$. First, we will discuss Markov Perfect Equilibria (MPE), which are analogous to Cournot competition in the simple model. Second, we will prove that it is impossible to achieve payoffs greater than those of an MPE.

A. Markov Perfect Equilibria

A (symmetric) MPE is a sequential equilibrium in which current supply rates $(\tilde{q}(p_t), \tilde{q}(p_t))$ depend only on current price. Denote by $F(p_t)$ the expected discounted payoff of each firm in a MPE. We will assume that $F(p_t)$ is a smooth function. Then $F(p_t)$ satisfies the following Bellman equation

$$F(p_t) = E[(1-\delta)qp_{t+\Delta} + \delta F(p_{t+\Delta})] = (1-\delta)qp_t + \delta \left(F(p_t) + \Delta f(2q, p_t)F'(p_t) + \frac{\Delta \sigma^2(p_t)}{2} F''(p_t) \right) + O(\Delta^2),$$

where to generate the second line we have used a Taylor expansion, the fact that $1-\delta = 1 - e^{-r\Delta} = O(\Delta)$ and the fact that $q = \tilde{q}(p_t)$ satisfies the first-order condition

$$(1-\delta)p_t + \delta \Delta f_1(2q, p_t)F'(p_t) + O(\Delta^2) = 0.$$

We are not concerned with existence, uniqueness or further characterization of MPE. Assuming that at least one MPE exists, let us show that firms cannot achieve higher payoffs in any PPE than in the best MPE as $\Delta \rightarrow 0$.

B. Impossibility of Collusion

Denote by $\bar{F}(p_t)$ the expected discounted payoff from the best symmetric PPE. We will assume that as $F(p_t)$ (which now denotes the payoffs in the best MPE) it is smooth and that there is an interior p_t at which the difference between these two functions is maximized.

PROPOSITION V. As the time period between actions shrinks to 0, the best collusive payoffs from symmetric PPE converge to Markov Perfect Equilibrium payoffs for all price levels, i.e. $\bar{F} \rightarrow F$ as $\Delta \rightarrow 0$.

For a sketch of the proof see the Appendix. The logic behind this proof is the same as in the simple model that we have analyzed in detail. In order to enforce collusion, firms must execute a statistical test of a normally distributed random variable that sends them to a punishment regime if there is enough evidence of cheating. The probability of punishment becomes disproportionately large as $\Delta \rightarrow 0$.

The unifying feature of the two models is that when Δ gets smaller, a deviation changes the mean of the signal by an amount that is $O(\sqrt{\Delta})$ relative to the standard deviation of this signal. In the nonstationary model presented in this section, the signal is the innovation in price; in the stationary model, the signal is the price itself.

VII. Discussion and Further Extensions

So far we have established the impossibility result and shown that the same reasoning can be applied in a model with monetary transfers or serially correlated demand. There are many additional ways in which one may modify the benchmark model, so we conclude the paper by informally discussing some of them.

Non-homogenous goods

In our model an important assumption that prevents asymmetric equilibria from helping collusion is that the goods are homogenous, so that the signal depends on the sum of quantities only. In fact, our results suggest that colluding firms may try to increase product differentiation to resolve this monitoring problem. However, as we argue below, product differentiation has to be sufficiently large in order to avoid the obstacles that prevent successful collusion in our setting.

Suppose that the firms can observe two signals

$$dX_{1t} = (\alpha q_{1t} + (1-\alpha)q_{2t})dt + dZ_{1t}$$

$$dX_{2t} = (\alpha q_{2t} + (1-\alpha)q_{1t})dt + dZ_{2t}$$

where $\alpha \in [1/2, 1]$ is a measure of how related the products are (and hence how difficult it is to distinguish between deviations by each player) and dZ_{it} are independent Brownian motions. For $\alpha=1/2$ this is our model (although with two signals - see below the section on Additional Signals). Analysis of such a game

is beyond the scope of this paper, but we can provide some intuition and conjectures (based on Sannikov (2004) and this paper). First, the best *symmetric* equilibrium payoff converges to v_N for the same reasons as in Section II. For any fixed r we expect that as $\alpha \rightarrow 1/2$, the best *asymmetric* equilibrium payoffs achievable for small Δ converge to Nash equilibrium payoffs. Also, the quantities involved in the best collusive equilibrium converge to q_N . The intuition is that even though asymmetric equilibria are helpful, the higher the correlation between the two signals, the more sensitive the continuation payoffs have to be to signal realizations to provide incentives. That leads to a higher volatility of continuation payoffs, which is costly since the set of continuation payoffs is convex. Thus, only quantity pairs close to q_N can be successfully enforced without destroying too much value. Summing up, we expect that for a given r and α close to $1/2$, the best collusive equilibrium payoff is non-monotonic in Δ , as in Figure 3, and as $\Delta \rightarrow 0$ it approaches v_N .

A stationary model with nonnegative prices

In our basic model of Section I the average price over any interval of time Δ is normally distributed. This allows us to divide the timeline into an arbitrarily fine grid in which at each point the firms adjust quantities, and to show that collusion is impossible in the limit. The basic model has an unattractive feature that prices can become unboundedly negative over short intervals of time. We can construct an alternative model with bounded prices as follows. Suppose that spot-market prices, observed every second, are given by

$$p_t = g(s_t), \quad t = 0, 1, \dots$$

where g is an increasing nonnegative function, $s_t = f(Q) + \varepsilon_t$ is a signal with $\varepsilon_t \sim N(0, \sigma^2)$ and f is a decreasing function. For example, if $g(s) = \exp(s)$, then prices would be lognormal. In this model, the expected inverse demand is $P(Q) = E[g(f(Q) + \varepsilon)]$. Assume that g and f are such that $P(Q)$ satisfies A1 and A2. Such a model separates the monitoring variable (s) from the prices.

Assume σ is sufficiently large. If firms have a production technology that allows them to change output flows every Δ seconds, then the expected per-period profits are $\Delta q_t P(Q)$. After observing one-second prices over the time period, an optimal way to test for deviations is to invert g and calculate the average s_t over that period. The average s_t has mean $f(Q)$ and variance σ^2/Δ . Applying the same reasoning as before, we can show that as Δ gets small, the highest average payoffs that can be sustained in equilibrium become very close to the static Nash payoffs.

Frequency of market data

In the game we have studied there are in fact two frequencies that determine the scope of collusion: the frequency of moves ($1/\Delta$) and the frequency at which the players observe prices. In Section I we have

assumed that these two frequencies are the same but remarked that nothing would change if firms observed prices at higher frequencies (in particular, observing the continuous price process). One could also imagine a market in which firms can move frequently, but the market data are collected by a third party and available only infrequently.

In general, what is important for the scope of collusion is the minimum of the two frequencies. The intuition is that if more prices are observed while quantities are fixed, then the sufficient statistic is the mean. On the other hand, if frequencies can be adjusted in between price observations, then given that there is no new information available to act on (and marginal revenue of residual demand is downward-sloping, making output smoothing optimal), such deviations are not profitable.

Endogenous selection of Δ

We have assumed that Δ is fixed exogenously. What would happen if it were chosen by the firms before any output game? Assume that the different choices of technology have the same costs (in the sense that lowering Δ is costless). It is easy to consider the following two scenarios.

First, suppose that choice of Δ by each firm is easily monitored. Then as a corollary to Propositions I and II, in the payoff-maximizing equilibrium firms would choose an interior Δ - neither too small nor too large.

Second, suppose that the firms can secretly choose their Δ and expect to play a collusive equilibrium afterwards. Then each will choose Δ as small as possible. The reason is that having a lower Δ gives the firm the option to react to the path of prices. As a result, firms will endogenously destroy any hope for collusion. The same reasoning applies if $\Delta=0$ but the firms can choose the frequency (or delay) at which they observe prices.

Additional signals

Assume that Δ is close to zero and that prices are observed continuously. Suppose there exists another source of information $y_t = G(Q_t) + \eta_t$ observed at intervals $\Delta_y > 0$, where η_t is normal noise independent of ε_t with variance inversely proportional to Δ_y .

If the time interval between observations Δ_y is fixed independently of Δ , then profitable collusion would be achievable in equilibrium: the players can disregard prices and focus on Y to provide incentives to restrict supply. This situation can arise if y_t is a signal from an industry survey released by an independent agency at regular time intervals. However if $\Delta_y = O(\Delta)$ then as $\Delta \rightarrow 0$ still no collusion is achievable. This can be the case if y_t is the price of an input commonly used in production by the two competing firms.

Delay of information

Finally, suppose $\Delta \equiv 0$ and that the information arrives continuously but with a delay of length d . As in AMP such delay allows for some information aggregation, and hence it can improve the scope of collusion. To see why collusion above Nash is possible, let us construct an equilibrium of this game based on the best equilibrium of a game without delay with $\Delta = d$ and $\delta = e^{-r2d}$. Divide time into time segments of length d . In the odd segments the players are recommended to play as in the best equilibrium with $\Delta = d$ and with $\delta = e^{-r2d}$. In the even segments they are recommended to play the static Nash. At the end of even segments players calculate the average price observed between times $t-d$ and t (which contains information about supply rates in the time interval $(t-2d, t-d)$), and then players decide on continuation payoffs based on that price.

Mixed strategy perfect public equilibria

Throughout the paper we have restricted our attention to pure strategy PPE. As we know from Fudenberg, Levine and Maskin (1994), in some games even if all pure strategy profiles are not identifiable, mixed strategy profiles may be. If some mixed strategy profiles have pairwise full rank, i.e. they allow the statistical distinction between deviations by different players and between different deviations by the same player, then those profiles can be sustained with balanced monetary transfers. In our game, the existence of such profiles depends on whether the action space (the range of quantities the firms can choose) is bounded. If in the mixed strategy profile both firms mix over quantities that can be locally increased, then a deviation by firm 1 to increase quantity by a unit (for all realizations of the mixing) is observationally equivalent to the analogous deviation by firm 2. Therefore if there are no bounds on quantity flows, then deviations from any mixed strategy profile are not identifiable and hence all our results still hold. However, if, as we have assumed in Section I, the firms have a capacity constraint \bar{q} , then a mixed strategy profile in which the two firms produce at capacity with some small probabilities could be both profitable and identifiable. Hence it may be possible to sustain collusion with balanced monetary transfers.

The use of private strategies

While PPE present a natural class of equilibria for games with public monitoring, recent theoretical literature has devoted a lot of attention to games with private monitoring and the use of private strategies. In a private strategy, a player's current action may depend not only on past public signals, but also on the private history of that player's actions.³¹ Equilibria in private strategies present a lot of theoretical challenges, as they are not connected with a convenient recursive structure of PPE. In fact, games with private monitoring are so difficult that even the Folk Theorem for those games has been proved only in a

limited class of settings.³² Nevertheless, equilibria in private strategies allow for a wider range of strategic interaction than PPE, and as shown by Kandori and Obara (2006), in some games they improve upon payoffs achievable in public perfect equilibria.

The idea of Kandori and Obara (2006) is that, by mixing, a player can sometimes detect deviations of her opponent more efficiently so as to minimize the likelihood of type I errors. This requires private strategies, because the decision to trigger punishment depends not only on the public signal but also on the action, which may have been used to test the opponent. Can such strategies help in our setting?

It turns out that such strategies may produce more efficient statistical tests for a fixed Δ , but as $\Delta \rightarrow 0$ type I errors still blow up disproportionately for all tests, as long as $P'(Q)$ is bounded. Consider a pair of collusive quantities (q, q) and let q' be an action used by one of the players to test her opponent. When one of the players deviates from q to $q + \varepsilon$ while the opponent is playing q , the mean of the public signal shifts from $P(2q)$ to $P(2q + \varepsilon)$ (by about $P'(2q)\varepsilon$). If the opponent is ‘checking’ the action of the former player by supplying the market at rate q' , the mean of the public signal shifts up by about $P'(q + q')\varepsilon$. While detection is improved when $|P'(q + q')\varepsilon| > |P'(2q)\varepsilon|$, there is still a uniform lower bound on type I errors from such statistical tests. That bound is on the order of about $\sqrt{\Delta}$, when $P'(Q)$ is bounded.

We conclude that the main known way by which private mixed strategies can be useful in other settings cannot be used to facilitate collusion when actions are frequent.³³ Unfortunately equilibria with private strategies are not understood well enough to show that our impossibility result holds for the entire class of such equilibria. We conjecture that the deterioration of information in our model is so severe that the whole set of PBE payoffs collapses to the static Nash payoffs as $\Delta \rightarrow 0$, but a proof of that conjecture remains an open question.

VIII. Conclusion

We have shown that scope of collusion can be affected very negatively by the ability of firms to react to information quickly if information about potential deviations comes to the market in a continuous fashion -- in particular, if the monitoring technology is such that it is difficult to tell apart deviations by different players. We have explored in great detail many forces that affect the scope of collusion: the use of asymmetric collusive schemes and monetary transfers, the existence of capacity constraints, less than homogenous goods and delay of information. We have also investigated how these forces affect collusion/cooperation in various other settings, e.g. in the case of non-stationary demand or moral hazard in teams. Our results explain many reasons why collusion may fail, and provide a rationale for some special arrangements that firms establish to achieve collusion. For example, as we discussed in the Introduction, firms often seek signals that go beyond the average market price in order to obtain collusion

at least in asymmetric schemes. Such asymmetric arrangements have been observed in many recent cases (see for example, François Arbault et al. (2002), European Commission (2002), or Harrington (2006)). Other arrangements that help collusion may include those that delay the release of information or take away firms' ability to adjust their production strategies quickly. We hope that our results will lead to a new, more complete understanding of how flexibility of actions affects dynamic collusion and cooperation, and that this paper will help and encourage further theoretical and empirical research in this area.

IX. Appendix

PROOF OF LEMMA 1: First, let us prove that the static best response is unique and less than q . The static best response problem is

$$\max_q P(q + q_j)q.$$

Let us show that this problem is concave. Differentiating twice with respect to q , we obtain:

$$\text{if } P''(q + q_j) < 0: \quad P''(q + q_j)q + 2P'(q + q_j) < 0$$

$$\text{else: } P''(q + q_j)q + 2P'(q + q_j) \leq P''(q + q_j)(q + q_j) + 2P'(q + q_j) < 0$$

where the last inequality follows from **A1**. So the problem is strictly concave. That establishes uniqueness. Directly from **A2** we have $q^*(q_j) < \bar{q}$.

The best response is defined by the first order condition

$$P'(q + q_j)q + P(q + q_j) = 0$$

whenever $P(q_j) > 0$ and it is equal to 0 otherwise. Now, if both firms choose positive quantities in equilibrium then the equilibrium satisfies the two FOCs:

$$P'(Q)q_i + P(Q) = 0$$

$$P'(Q)q_j + P(Q) = 0$$

Indeed, there does not exist an equilibrium with either of the quantities equal to 0. Suppose that $q_j = 0$. That would require $P(q_i) \leq 0$. But then the marginal profit of firm i is negative and hence it is not playing a best response. So in equilibrium both FOCs have to hold. Now, subtracting them from one another we get:

$$P'(Q)(q_i - q_j) = 0,$$

which can be satisfied only if $q_i = q_j$.

Finally, using this symmetry we can write $Q = 2q$ and then the equilibrium has to satisfy:

$$P'(Q)Q + 2P(Q) = 0 \quad (7)$$

The LHS of this equation is decreasing, as by **A1**:

$$P''(Q)Q + 3P'(Q) < P''(Q)Q + 2P'(Q) < 0$$

Furthermore, at $Q = 0$ the LHS of (7) is positive, and at $Q = 2\bar{q}$ it is negative by **A2**. Therefore, equation (7) has a unique interior solution $Q = 2q_N$.

PROOF OF LEMMA 2: There are two ways of representing the likelihood difference of a tail test on a graph, as shown on Figure 7.




Figure 7

Using the area to the right, we see that there exists $x^* \in (c, c + (\mu - \mu'))$ such that

$$\text{likelihood difference} = (\mu - \mu')g(x^*) = (\mu - \mu') \frac{\sqrt{\Delta}}{\sqrt{2\pi\sigma}} \exp\left(-\frac{(x^* - \mu)^2 \Delta}{2\sigma^2}\right).$$

Then if the likelihood difference is equal to $C_1\Delta$.³⁴

$$(\mu - \mu') \frac{\sqrt{\Delta}}{\sqrt{2\pi\sigma}} \exp\left(-\frac{(x^* - \mu)^2 \Delta}{2\sigma^2}\right) = C_1\Delta \Rightarrow \mu - x^* = \frac{\sigma}{\sqrt{\Delta}} \sqrt{-\log\left(2\pi\Delta \frac{C_1^2 \sigma^2}{(\mu - \mu')^2}\right)}.$$

Let $\alpha > 1$ be a number to be specified later. Let y^* satisfy $(\mu - y^*) = \alpha(\mu - x^*)$. Because $x^* - (\mu - \mu') < c$, the probability of type I error (i.e. the size of the test) is greater than the shaded area in Figure 8, i.e.

$$\begin{aligned} (x^* - (\mu' - \mu) - y^*)g(y^*) &= ((\alpha - 1)(\mu - x^*) - (\mu - \mu')) \frac{\sqrt{\Delta}}{\sqrt{2\pi\sigma}} \exp\left(-\frac{\alpha^2 (x^* - \mu)^2 \Delta}{2\sigma^2}\right) = \\ &\left((\alpha - 1) \frac{\sigma}{\sqrt{\Delta}} \sqrt{-\log\left(2\pi\Delta \frac{C_1^2 \sigma^2}{(\mu - \mu')^2}\right)} - (\mu - \mu') \right) \frac{\sqrt{\Delta}}{\sqrt{2\pi\sigma}} \left(\frac{C_1 \Delta \sqrt{2\pi\sigma}}{(\mu - \mu') \sqrt{\Delta}} \right)^{\alpha^2} > O(\Delta^{\alpha^2/2}) \end{aligned}$$




Figure 8

Taking α sufficiently close to 1 proves the Lemma.

PROOF OF LEMMA 3: Write the Lagrangian for the maximization problem

$$L = \int_{-\infty}^{\infty} v(x) g(x) dx + \lambda \left(\delta \int_{-\infty}^{\infty} v(x) (g(x) - g'(x)) dx - D\Delta \right) \\ + \int_{-\infty}^{\infty} \rho_1(x) (v(x) - \underline{v}) dx + \int_{-\infty}^{\infty} \rho_2(x) (\bar{v} - v(x)) dx,$$

where $\rho_1(x) > 0$ only if $v(x) = \underline{v}$, and $\rho_2(x) > 0$ only if $v(x) = \bar{v}$. Taking first-order conditions with respect to $v(x)$ gives

$$g(x) + \lambda \delta (g(x) - g'(x)) + (\rho_1(x) - \rho_2(x)) = 0.$$

If $g(x) + \lambda \delta (g(x) - g'(x)) < 0$ then $\rho_1(x) > 0$ and $v(x) = \underline{v}$. If $g(x) + \lambda \delta (g(x) - g'(x)) > 0$ then $\rho_2(x) > 0$ and $v(x) = \bar{v}$. We have

$$g(x) + \lambda \delta (g(x) - g'(x)) < 0 \Leftrightarrow \frac{g'(x)}{g(x)} > \frac{\lambda \delta + 1}{\lambda \delta}.$$

Because $g'(x)/g(x)$ is decreasing in x , the last inequality holds when $x < c$ for some constant c . Moreover, since $g'(x)/g(x) = 1 < (\lambda \delta + 1)/\lambda \delta$, when $x = (\mu + \mu')/2$, $c < (\mu + \mu')/2$. We conclude that the solution to the maximization problem satisfies (1).

Lemma 4 is technical. It investigates profitable deviations.

LEMMA 4. Consider symmetric supply rates (q, q) for $q \in [0, q_N - \varepsilon]$. For any $\varepsilon > 0$, the gain from the most profitable deviation is bounded away from 0 i.e. there exists $\varepsilon_\pi > 0$ such that

$$\forall q \in [0, q_N - \varepsilon], \quad P(q^*(q) + q)q^*(q) - P(2q)q > \varepsilon_\pi,$$

where $q^*(q)$ is a static best response to q .

Also, the price reduction in the most profitable deviation is bounded i.e. there exists M such that

$$\forall q \in [0, q_N - \varepsilon], \quad P(2q) - P(q^*(q) + q) < M.$$

PROOF OF LEMMA 4: First, the gains from deviations are bounded away from 0. Indeed, for all $q \in [0, q_N - \varepsilon]$, $q^*(q) \neq q$, so

$$P(q^*(q) + q)q^*(q) - P(2q)q > 0.$$

Also, $P(q^*(q) + q)q^*(q) - P(2q)q$ is continuous. (In fact, it is differentiable by the Envelope theorem, with derivative $P'(q^*(q) + q)q^*(q)$.) A positive continuous function on a compact set $[0, q_N - \varepsilon]$ must be bounded away from 0.

Second, the price reduction from the most profitable deviation is bounded because $q \in [0, q_N - \varepsilon]$, $q^*(q) \in [0, \bar{q}]$, and prices are continuous over this range.

PROOF OF PROPOSITION I. We can characterize the best symmetric equilibrium (using methods from APS'86) as a solution to the following problem:

$$(8) \quad \begin{aligned} \bar{v}(\Delta) &= \max_{q, v(x)} (1 - e^{-r\Delta})qP(2q) + e^{-r\Delta} \int_{-\infty}^{\infty} v(x)g(x)dx \\ \text{s.t.} \\ \forall q' \in [0, \bar{q}]: & \quad (1 - e^{-r\Delta})(q'P(q + q') - qP(2q)) \leq e^{-r\Delta} \int_{-\infty}^{\infty} v(x)(g(x) - g_{q'}(x))dx \\ v(x) &\in [\underline{v}(\Delta), \bar{v}(\Delta)] \end{aligned}$$

where (q, q') is the quantity pair chosen in the current period, g denotes the density of prices for total supply $2q$, $g_{q'}$ denotes the density for total supply $q + q'$, and $v(x)$ denotes the continuation payoff when the realized price is x . The first constraint is a standard one-period incentive compatibility constraint and the second constraint states that continuation payoffs must be achievable in equilibrium.

Let us relax the problem. First, consider only one incentive constraint against a deviation to a static best response. Second, allow a wider range of continuation payoffs. Third, weaken slightly the incentive constraint by noting that $r\Delta < (1 - e^{-r\Delta})/e^{-r\Delta}$. In this way we obtain the following relaxed problem:

$$\begin{aligned}
\bar{v}_R(\Delta) &= \max_{q, v(x)} (1 - e^{-r\Delta}) q P(2q) + e^{-r\Delta} \int_{-\infty}^{\infty} v(x) g(x) dx \\
(9) \quad &\text{s.t.} \\
q' = q^*(q) &: \quad r\Delta (q' P(q + q') - q P(2q)) \leq \int_{-\infty}^{\infty} v(x) (g(x) - g'(x)) dx \\
v(x) &\in [0, \bar{v}_R(\Delta)]
\end{aligned}$$

where g' denotes the density of prices for total supply $q + q^*(q)$. Clearly, $\bar{v}_R(\Delta) \geq \bar{v}(\Delta) \geq v_N$ so it will be sufficient to show that $\bar{v}_R(\Delta) \rightarrow v_N$ as $\Delta \rightarrow 0$. Consider any $\varepsilon > 0$. Let us show that if Δ is sufficiently small, then no supply rates $q \in [0, q_N - \varepsilon]$ can achieve $\bar{v}_R(\Delta)$. We will follow the reasoning presented in Figure 4.

Suppose the solution to (9) is $q \in [0, q_N - \varepsilon]$. Let $\varepsilon_\pi > 0$ be a lower bound on the static profit gain from any deviation $q^*(q)$ for $q \in [0, q_N - \varepsilon]$, and let M be an upper bound on the drop in prices from such deviations (see Lemma 4 above). Denote the static gain from $q^*(q)$ by

$$D = P(q^*(q) + q) q^*(q) - P(2q) q \geq \varepsilon_\pi.$$

This deviation makes the expected price drop by $P(2q) - P(q^*(q) + q) \leq M$.

Keeping the q that solves (9) fixed, $v(x)$ solves:

$$\begin{aligned}
(10) \quad &\max_{v(x)} \int_{-\infty}^{\infty} v(x) g(x) dx \\
&\text{s.t.} \\
q' = q^*(q) &: \quad r\Delta D \leq \int_{-\infty}^{\infty} v(x) (g(x) - g'(x)) dx \\
v(x) &\in [0, \bar{v}_R(\Delta)]
\end{aligned}$$

By Lemma 3, $v(x)$ takes the form:

$$(11) \quad v(x) = \begin{cases} \bar{v}_R(\Delta) & \text{if } x > c \\ 0 & \text{if } x \leq c \end{cases}$$

This form of v is associated with a tail test with critical region $(-\infty, c]$. The likelihood difference in this test is

$$\text{likelihood difference} = \int_{-\infty}^c (g'(x) - g(x)) dx$$

Denote by $v_M \geq \bar{v}_R(\Delta)$ half the monopoly profit. From the incentive compatibility constraint in (10), we have

$$(12) \quad r\Delta D \leq \int_c^\infty \bar{v}_R(\Delta)(g(x) - g'(x))dx \leq v_M \int_c^\infty (g(x) - g'(x))dx = v_M \underbrace{\int_{-\infty}^c (g'(x) - g(x))dx}_{\text{likelihood difference}}$$

That implies:

$$(13) \quad \text{likelihood difference} \geq r\Delta D / v_M \geq r\Delta \varepsilon_\pi / v_M$$

so it is of order Δ .

Referring to Lemma 2, let $\Delta^* > 0$ be such that for all $\Delta \leq \Delta^*$, a tail test for a difference in mean of M with likelihood difference $\varepsilon_\pi \Delta r / v_M$ has a probability of type I error greater than $C\Delta^\alpha$ for some $C > 0$ and $\alpha < 1$. Then any test for a difference in mean of less than M with the same likelihood difference has an even greater probability of type I error.

Consider $\Delta < \Delta^*$. The probability of making a type I error is

$$(14) \quad \int_{-\infty}^c g(x)dx > C\Delta^\alpha.$$

This implies:

$$(15) \quad \int_{-\infty}^\infty v(x)g(x)dx = \int_c^\infty \bar{v}_R(\Delta)g(x)dx < \bar{v}_R(\Delta)(1 - C\Delta^\alpha).$$

Finally, returning to (9), for any $\alpha < 1$

$$(16) \quad \underbrace{(1 - e^{-r\Delta})P(2q)q}_{O(\Delta)} + \underbrace{e^{-r\Delta} \int_{-\infty}^\infty v(x)g(x)dx}_{1 - O(\Delta)} < \bar{v}_R(\Delta)$$

$< \bar{v}_R(\Delta) - O(\Delta^\alpha)$

for sufficiently small Δ (from Lemma 2 we know that we can pick any $\alpha > 1/2$). This leads to a contradiction because, according to (9), we must have equality in (16). Therefore if q is the current-period supply rate that achieves value $\bar{v}_R(\Delta)$, then $q > q_N - \varepsilon$. But then

$$\bar{v}_R(\Delta) \leq (1 - e^{-r\Delta})P(2q)q + e^{-r\Delta}\bar{v}_R(\Delta) \Rightarrow \bar{v}_R(\Delta) \leq P(2q)q$$

Note that $P(2q)q$ is continuous and (by **A1**) decreasing for q greater than a half of the monopoly quantity. Letting $\varepsilon \rightarrow 0$, we conclude that $\bar{v}_R(\Delta) \rightarrow v_N$ as $\Delta \rightarrow 0$.

PROOF OF PROPOSITION II AND ITS COROLLARY. Denote by $E_a(\Delta)$ the set of payoff pairs achievable in asymmetric PPE (without monetary transfers) of game Γ^Δ .

Letting $Q = q_1 + q_2$, $\bar{v}_a(\Delta)$ solves

$$(17) \quad \bar{v}_a(\Delta) = \max_{q_1, q_2, v_1(x), v_2(x)} (1 - e^{-r\Delta})QP(Q) + e^{-r\Delta} \int_{-\infty}^{\infty} (v_1(x) + v_2(x))g(x)dx$$

subject to the players' incentive constraints and the constraint that future continuation values $(v_1(x), v_2(x))$ must be in the set $E_a(\Delta)$, i.e. $(v_1(x), v_2(x)) \in E_a(\Delta)$. We relax the problem in several steps. First, as in the proof of Proposition I, we relax the problem by keeping only the incentive constraints against upward deviations of ε by each firm, where ε will be specified later. Each deviation has the same effect on the distribution of prices, decreasing the mean of the observed price to $P(Q + \varepsilon)$. The incentive constraints against those deviations are given above by (2). We further relax the problem by replacing constraints (2) with their sum (see (3)) and by replacing the constraint $(v_1(x), v_2(x)) \in E_a(\Delta)$ with a weaker constraint $v_1(x) + v_2(x) \in [0, \bar{v}_a(\Delta)]$. As a result, we arrive at the problem

$$(18) \quad \begin{aligned} \bar{v}_{a,R}(\Delta) &= \max_{q_1, q_2, v_1(x), v_2(x)} (1 - e^{-r\Delta})QP(Q) + e^{-r\Delta} \int_{-\infty}^{\infty} (v_1(x) + v_2(x))g(x)dx \\ \text{s.t.} \\ r\Delta \left(\left(\frac{Q}{2} + \varepsilon \right) P(Q + \varepsilon) - \frac{Q}{2} P(Q) \right) &\leq \int_{-\infty}^{\infty} \frac{v_1(x) + v_2(x)}{2} (g(x) - g'(x)) dx \\ v_1(x) + v_2(x) &\in [0, \bar{v}_{a,R}(\Delta)] \end{aligned}$$

Clearly, $\bar{v}_{a,R}(\Delta) \geq \bar{v}_a(\Delta) \geq 2v_N$. Problem (18) is the same as problem (9) multiplied by 2 (here $v_1(x) + v_2(x)$ plays the same role as $2v(x)$ in problem (9)) so $\bar{v}_{a,R}(\Delta) = 2\bar{v}_R(\Delta)$. It was shown in the proof of Proposition I that $\bar{v}_R(\Delta) \rightarrow v_N$ as $\Delta \rightarrow 0$. Therefore, $\bar{v}_{a,R}(\Delta) \rightarrow 2v_N$ and $\bar{v}_a(\Delta) \rightarrow 2v_N$.

With monetary transfers and money burning bounded by \bar{b} in each period, we conclude similarly that the best collusive profits are bounded by the solution of the relaxed problem

$$\bar{v}_{a,R'}(\Delta) = \max_{q_1, q_2, v_1(x), v_2(x)} (1 - e^{-r\Delta}) QP(Q) + e^{-r\Delta} \int_{-\infty}^{\infty} (v_1(x) + v_2(x) - b(x))g(x) dx$$

s.t.

$$r\Delta \left(\left(\frac{Q}{2} + \varepsilon \right) P(Q + \varepsilon) - \frac{Q}{2} P(Q) \right) \leq \int_{-\infty}^{\infty} \frac{v_1(x) + v_2(x) - b(x)}{2} (g(x) - g'(x)) dx$$

$$v_1(x) + v_2(x) - b(x) \in [-\bar{b}, \bar{v}_{a,R'}(\Delta)]$$

As in the proof of Proposition **I**, a solution to this problem is based on a tail test with a likelihood difference on the order of Δ . We know that such tests destroy too much value as $\Delta \rightarrow 0$.

PROOF OF REMARK IN SECTION III: i) Fix $\Delta > 0$ and $\varepsilon > 0$. Denote by q_M the monopoly output. Suppose that the transfers are unbounded: in particular that the players can each commit to burning any amount of money T (in the sense that they each pay T to the third party). Then we can provide incentives statically, period-by-period. Suppose that the players sign a contract

$$t_i(x) = \begin{cases} 0 & \text{if } x \geq c \\ T & \text{if } x < c \end{cases}$$

which acts as a tail test that prescribes the players to burn T if the average price in the given period drops below c . The continuation payoff is independent of current prices.

Let us fix c and show that firms will choose quantity pair $(q_M/2, q_M/2)$ for an appropriate choice of T . We have to make sure that $q' = q_M/2$ maximizes

$$(19) \quad \Delta q' P(q' + q_M/2) - T \int_{-\infty}^{c - P(q' + q_M/2)} \phi(x) dx,$$

where ϕ is the density of a normal distribution with mean 0 and variance σ^2/Δ . The choice of q' affects the current-period profit and the probability of punishment. Taking the first-order condition, we get

$$T = -\Delta \frac{P(q_M) + P'(q_M)q_M/2}{\phi(c - P(q_M))P'(q_M)} > 0$$

Since there is no bound on T , for any small c we can find T that provides the players with incentives to choose $(\frac{q_M}{2}, \frac{q_M}{2})$. Then the per-period expected payoffs are

$$\frac{1}{2} \Delta v_M - T \Phi(c - P(q_M)) = \Delta \left(\frac{1}{2} v_M - K \frac{\Phi(c - P(q_M))}{\phi(c - P(q_M))} \right),$$

where $K = -\frac{P(q_M) + P'(q_M)q_M/2}{P'(q_M)}$ does not depend on c (Φ is the *cdf* of the standard normal

distribution). As $\lim_{z \rightarrow -\infty} \frac{\Phi(z)}{\phi(z)} = 0$, the firms can achieve arbitrarily close to the monopoly profit by

choosing a sufficiently small c .

ii) We sketch the reasoning behind this proof: Suppose that we want to obtain collusion with

$\bar{v}_a(\Delta) > 2v_N + \varepsilon$. From the previous part we know that to obtain payoffs $\bar{v}_a(\Delta) > 2v_N + \frac{1}{n}\varepsilon$ we need

to be able to burn at least $T(\Delta)$ in a given period with $\lim_{\Delta \rightarrow 0} T(\Delta) = \infty$. Now, we pick Δ small enough so

that the initial wealth is smaller than $T(\Delta)$ and that with probability $1 - \varepsilon'$ the wealth of the players at

any time $t \in [0, \tau]$ is smaller than $T(\Delta)$ as well. Then the expected payoffs at time 0 are bounded by

$$(1 - e^{-r\tau}) \left(2v_N + \frac{1}{n}\varepsilon \right) + \varepsilon' v_M + (1 - \varepsilon') e^{-r\tau} v_M,$$

which, as $n, \tau \rightarrow \infty$ and $\varepsilon' \rightarrow 0$, converges to $2v_N$.

SKETCH OF PROOF OF PROPOSITION V. Let p^* be the price that maximizes $\bar{F}(p) - F(p)$, which implies that $\bar{F}'(p^*) = F'(p^*)$ and $\bar{F}''(p^*) \leq F''(p^*)$. Let $\tilde{q} = \tilde{q}(p^*)$. Let us show that players cannot achieve payoffs higher than $F(p^*)$ by choosing a quantity pair $(q, q) \geq (\tilde{q}, \tilde{q})$ in the current period. Intuitively, if players wanted to collude, they would need to cut quantities, not to increase them. Indeed, if (q, q) is used currently, then the Bellman equation implies

$$\begin{aligned} \bar{F}(p_t) &= (1 - \delta)\tilde{q}p_t + \delta \left(\bar{F}(p_t) + \Delta f(2q, p_t)\bar{F}'(p_t) + \frac{\Delta\sigma^2(p_t)}{2}\bar{F}''(p_t) - \text{exp. punishment} \right) + O(\Delta^2) \Rightarrow \\ &(1 - \delta)\bar{F}(p_t) \stackrel{\text{1}}{\leq} (1 - \delta)qp_t + \delta \left(\Delta f(2q, p_t)\bar{F}'(p_t) + \frac{\Delta\sigma^2(p_t)}{2}\bar{F}''(p_t) \right) + O(\Delta^2) \stackrel{\text{2}}{\leq} \\ &(1 - \delta)qp_t + \delta \left(\Delta f(2q, p_t)F'(p_t) + \frac{\Delta\sigma^2(p_t)}{2}F''(p_t) \right) + O(\Delta^2) \stackrel{\text{3}}{\leq} \\ &(1 - \delta)\tilde{q}p_t + \delta \left(\Delta f(2\tilde{q}, p_t)F'(p_t) + \frac{\Delta\sigma^2(p_t)}{2}F''(p_t) \right) + O(\Delta^2) = (1 - \delta)F(p_t), \end{aligned}$$

where the inequality 1 follows from ignoring punishment and inequality 2 follows because $\bar{F}'(p^*) = F'(p^*)$ and $\bar{F}''(p^*) \leq F''(p^*)$. For inequality 3, the first-order condition in a MPE at price p^* implies that

$$(1-\delta)p_t + \delta\Delta f_1(2\tilde{q}, p_t)F'(p_t) + O(\Delta^2) = 0 \Rightarrow$$

$$\forall q \geq \tilde{q}, (1-\delta)p_t + 2\delta\Delta f_1(2q, p_t)F'(p_t) + O(\Delta^2) < 0,$$

so $(1-\delta)qp_t + \delta\Delta f(2q, p_t)F'(p_t)$ must be decreasing in q for all $q \geq \tilde{q}$.

We claim that it is impossible to achieve payoffs larger than $F(p_t)$ even by using a quantity pair $(q, q) < (\tilde{q}, \tilde{q})$. Intuitively, this happens because a statistical test to detect potential deviations is too costly. Let M be an upper bound on the size of punishment possible in an equilibrium. If firms use quantity pair $(\hat{q}, \hat{q}) < (\tilde{q}, \tilde{q})$ to achieve $\bar{F}(p_t)$, they must perform a statistical test on $p_{t+\Delta}$ to decide when to go to punishment. Let $q^*(\hat{q})$ be a supply rate that maximizes

$$(1-\delta)qp_t + \delta\Delta f(q + \hat{q}, p_t)\bar{F}'(p_t),$$

i.e. that maximizes expected payoffs (modulo constants and terms of order Δ^2) without taking into account the possibility of switching to the punishment phase. Let T be the best tail test on

$$\frac{p_{t+\Delta} - p_t}{\Delta} \sim N\left(f(q + \hat{q}, p_t), \frac{\sigma^2(p_t)}{\Delta}\right)$$

that prevents deviation towards $q^*(\hat{q})$ through a punishment of

magnitude M . Then we have

$$(1-\delta)\bar{F}(p_t) \leq (1-\delta)\hat{q}p_t + \delta\left(\Delta f(2\hat{q}, p_t)F'(p_t) + \frac{\Delta\sigma^2(p_t)}{2}F''(p_t) - M \cdot \text{size}(T)\right) + O(\Delta^2)$$

and the incentive compatibility constraint

$$(1-\delta)q^*(\hat{q})p_t + \delta\Delta f(q^*(\hat{q}) + \hat{q}, p_t)\bar{F}'(p_t) - (1-\delta)\hat{q}p_t - \delta\Delta f(2\hat{q}, p_t)\bar{F}'(p_t) \leq M \cdot LD(T).$$

Unless $\hat{q} \approx \tilde{q}$, the likelihood difference $LD(T)$ must be on the order of Δ . But then $\text{size}(T)$ is on the order of $\Delta^{0.5+\varepsilon}$ and it becomes impossible to sustain $\bar{F}(p_t) > F(p_t)$.

References

- Abreu, Dilip, Paul R. Milgrom, and David Pearce.** 1991. "Information and Timing in Repeated Partnerships." *Econometrica*, 59: 1713-1733.
- Abreu, Dilip, David Pearce, and Ennio Stacchetti.** 1986. "Optimal Cartel Equilibria with Imperfect Monitoring." *Journal of Economic Theory*, 39: 251-269.
- Abreu, Dilip, David Pearce, and Ennio Stacchetti.** 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica*, 58: 1041-1063

- Aoyagi, Masaki.** 2003. "Bid Rotation and Collusion in Repeated Auctions." *Journal of Economic Theory*, 112: 79-105.
- Arbault, François, Sari Suurnäki, Francisco Peiró, Barbara Nijs, Paul Bridgeland, Gerald Berger, Maarit Lindroos, and Erwan Marteil.** 2002. "Commission Adopts Eight New Decisions Imposing Fines on Hard-Core Cartels." *EC Competitions Policy Newsletter*, February(1): 29-43.
- Athey, Susan and Kyle Bagwell.** 2001. "Optimal Collusion with Private Information." *RAND Journal of Economics*, 32: 428-465.
- Athey, Susan, Kyle Bagwell, and Chris Sanchirico.** 2004. "Collusion and price rigidity." *Review of Economic Studies*, 71: 317–350.
- Blume, Andreas and Paul Heidhues.** 2003. "Modeling Tacit Collusion in Auctions." working paper, University of Pittsburgh.
- Blume, Andreas and Paul Heidhues.** Forthcoming. "Private Monitoring in Auctions." *Journal of Economic Theory*.
- Cabral, Luis M. B.** 2005. "Collusion Theory: Where to Go Next?" *Journal of Industry, Competition and Trade*, 5: 199-206.
- Eichenwald, Kurt.** 2000. *The Informant*, New York: Broadway Books.
- European Commission.** 2002. "Commission Fines Five Companies in Sodium Gluconate Cartel." Press Release, March 19.
- Fudenberg, Drew, David K. Levine, and Eric Maskin.** 1994. "The Folk Theorem with Imperfect Public Information." *Econometrica*, 62: 997-1040.
- Fudenberg, Drew and David K. Levine.** 2006. "Continuous Time Models of Repeated Games with Imperfect Public Monitoring," working paper, Harvard University.
- Genesove, David and Wallace P. Mullin.** 2001. "Rules, Communication, and Collusion: Narrative Evidence from the Sugar Institute Case." *American Economic Review*, 91: 379-398.
- Green, Edward J. and Robert H. Porter.** 1984. "Noncooperative Collusion under Imperfect Price Information." *Econometrica*, 52: 87-100.
- Harrington, Joseph.** 2006. "How Do Cartels Operate?" *Foundations and Trends in Microeconomics*, 2, Issue 1.
- Harrington Joseph and Andrzej Skrzypacz.** Forthcoming. "Collusion under Monitoring of Sales." *RAND Journal of Economics*.
- Hay, George A. and Daniel Kelley.** 1974. "An Empirical Survey of Price Fixing Conspiracies." *Journal of Law and Economics*, 17: 13-38.

- Hörner, Johannes and Wojciech Olszewski.** Forthcoming. “The Folk Theorem for Games with Private Almost-Perfect Monitoring”, *Econometrica*.
- Kandori, Michihiro and Ichiro Obara.** 2006. “Efficiency in Repeated Games Revisited: the Role of Private Strategies.” *Econometrica*, 74: 499-519.
- Levenstein, Margaret C. and Valerie Y. Suslow.** 2004. “What Determines Cartel Success?” working paper, University of Michigan.
- Matsushima, Hitoshi.** 2004. "Repeated Games with Private Monitoring: Two Players." *Econometrica*, 72: 823-852
- McAfee, R. Preston and John McMillan.** 1992. “Bidding Rings.” *American Economic Review*, 82: 579-599.
- Miller, David.** 2004. “The Dynamic Cost of Ex Post Incentive Compatibility in Repeated Games of Private Information.” working paper, UC San Diego.
- Porter Robert.** 1983. “A Study of Cartel Stability: The Joint Executive Committee, 1880-1886.” *Bell Journal of Economics*, 14: 301-314.
- Radner, Roy, Roger Myerson and Eric Maskin.** 1986. “An Example of a Repeated Partnership Game with Discounting and with Uniformly Inefficient Equilibria.” *Review of Economic Studies*, 53: 59-69.
- Sannikov, Yuliy.** 2006. “Games with Imperfectly Observable Actions in Continuous Time.” working paper, UC Berkeley.
- Skrzypacz, Andrzej and Hugo Hopenhayn.** 2004. “Tacit Collusion in Repeated Auctions.” *Journal of Economic Theory*, 114: 153-169.
- Stigler, George J.** 1964. “A Theory of Oligopoly.” *Journal of Political Economy*, 72: 44-61.
- Ulen, Thomas.** 1983. “Railroad Cartels before 1887: The Effectiveness of Private Enforcement of Collusion.” *Research in Economic History*, 8, 125-144.

FOOTNOTES

* Sannikov: 549 Evans Hall, Berkeley, CA 94720, sannikov@econ.berkeley.edu; Skrzypacz: Stanford Graduate School of Business, 518 Memorial Way, Stanford, CA 94305, andy@gsb.stanford.edu. We thank Jeremy Bulow, Kyna Fong, Drew Fudenberg, Joseph Harrington, Ichiro Obara, Robert Wilson and seminar participants at the SED 2004, IIOC 2005, Midwest Macro 2005, SITE 2005, Arizona State University, Berkeley, Columbia, University of Chicago, University of Illinois in Urbana Champaign, University of Iowa, New York University, Penn State, University of Pennsylvania, Princeton, Rutgers, Stanford, University of Texas in Austin, Washington University for useful comments and suggestions.

¹ There are many papers describing explicit and tacit collusion among firms. For comprehensive studies, see for example George A. Hay and Daniel Kelley (1974), Margaret C. Levenstein and Valerie Y. Suslow (2004) or Joseph Harrington (2006).

² The term *flexible production* is usually understood as describing low costs of changing the amount and the variety of output. In this paper we use this term in a narrower sense, having firms face low costs of changing flow but restricting them to produce only one type of output. Flexibility on the variety dimension has separate effects on the scope of collusion: increased product differentiation may improve collusion, but increased complexity of monitoring may destabilize cartels.

³ The lysine cartel has received a lot of attention thanks to the abundance of detailed information available on its inner workings, including FBI videotapes of cartel members' meetings. The cartel was described in detail by Kurt Eichenwald (2000) in a 600-page book, and discussed further by Luis M. B. Cabral (2005) and Harrington (2006).

⁴ The problems were real – ADM was indeed overproducing.

⁵ For a comprehensive list of such arrangements see Harrington (2006).

⁶ See Thomas Ulen (1983) and Porter (1983) for a detailed study of this cartel.

⁷ In this non-stationary setting the absence of collusion is characterized by a Markov Perfect Equilibrium, and the first-best collusion is similarly a state-dependent strategy. However, the Nash equilibrium and monopoly quantity of a one-shot analogue of this game (with constant demand $1 - Q$) serve well for illustration.

⁸ AMP pioneered the theory of games with frequent actions. We discuss the relationship to their paper at the end of the introduction, and, in more detail, at the end of Section 3.1.

⁹ Also, punishment is delayed when firms keep quantities fixed for a period of time of 2Δ , which hurts collusion, but this effect is negligible for small Δ .

¹⁰ Strictly speaking, if we perform tests with critical regions 1 and 2 in a game with period 2Δ , they provide a bit weaker incentives than in a game with period Δ , because when the first signal falls in region 1, punishment is delayed by time Δ . However, for small Δ this difference is negligible.

¹¹ We know from APS'86 that the most efficient equilibrium does not use review strategies, and so the firms cannot use the more efficient two-period tests when the period between actions is Δ . Indeed, incentives would break down if the firms tried to do that. First, after seeing the first-period prices the firms will likely react, either by reducing production to avoid a price war, or by increasing production if the first-period outcome makes it unlikely that the price war will be triggered after the second-period review.. Second, anticipating that (own) reaction, each firm will overproduce in the first period. Third, anticipating the reaction of the *other* firm, a firm expects its opponent to share the cost of preventing punishment in the second period and hence will overproduce in the first period even more. Since a single decision maker benefits from more flexibility (which is captured by the first two effects), it is the third, strategic effect that prevents collusion as $\Delta \rightarrow 0$.

¹² This intuition is related to the concept of identifiability from Drew Fudenberg, David K. Levine and Eric Maskin (1994); however, lack of identifiability is not sufficient to establish the result.

¹³ An example of such a partnership game is found in Roy Radner, Roger Myerson and Maskin (1986), which we discuss below.

¹⁴ The differences between Brownian and Poisson information have been recently further explored by Fudenberg and Levine (2006).

¹⁵ We also analyze a different game - the Green and Porter duopoly with a continuum of action choices per period.

¹⁶ Many papers look at the role of asymmetric equilibria and monetary transfers in achieving collusion. In a static setting with private information, R. Preston McAfee and John McMillan (1992) show the necessity of transfers for any collusive scheme utilizing the private information. Susan Athey and Kyle Bagwell (2001), Masaki Aoyagi (2003), Andreas Blume and Paul Heidhues (2003) and Andzej Skrzypacz and Hugo Hopenhayn (2004) among others extend that intuition by studying collusion in repeated games without monetary transfers and emphasizing the need for transfers of continuation payoffs. Harrington and Skrzypacz (2004) study repeated competition with hidden prices and observed stochastic market shares, and show that symmetric equilibria (with price wars after skewed market shares) do not improve

upon competitive outcomes but asymmetric equilibria do. Similarly, Athey and Bagwell (2001) show that asymmetric strategies greatly improve the scope of collusion (which is shown to be quite limited in the symmetric equilibria of Athey, Bagwell and Sanchirico (2004)). In contrast, our impossibility result holds even for asymmetric equilibria and with any monetary transfers.

¹⁷ See Section V for the extension to increasing marginal cost.

¹⁸ It is important to point out that our results do not follow simply from the per-period variance increasing to infinity as $\Delta \rightarrow 0$ but rather from the rate at which it increases.

¹⁹ One can also allow for public randomizations, so that before each period the players observe a realization of a public random variable (which becomes a part of the public history). It does not change any of the results - since prices have full support, any randomization can be supported using only them.

²⁰ When we consider pure strategies, it is not restrictive to focus on public strategies. Indeed, for every private strategy, one can find a public strategy that induces the same probability measure over private histories.

²¹ This result holds under assumptions similar to the ones we make in the duopoly model. See Remark in Section V.B for further discussion.

²² Although the second of these two effects is the main reason cooperation is impossible for small Δ , clearly the two effects are related.

²³ For small Δ it is sufficient since the probability of multiple jumps arriving is on the order of Δ^2 and hence has such events have negligible effect on incentives.

²⁴ A similar method can be used to show that in the partnership game of Radner, Myerson and Maskin (1986) asymmetric equilibria cannot improve upon symmetric ones.

²⁵ According to the proof of Proposition II, asymmetric equilibria cannot improve upon symmetric equilibria if punishment 0 is allowed.

²⁶ The difference from A4 in the duopoly model arises because in the duopoly case colluding firms want to overproduce while in the partnership cooperating agents want to under-supply effort.

²⁷ We subtract the local incentive compatibility constraints to arrive at a relaxed problem. This relaxed problem has a solution given by a cut-off rule. Then, bounding the type-I errors, we can prove that the maximal difference in equilibrium continuation payoffs drops to 0 at a rate much faster than $O(\sqrt{\Delta})$.

²⁸ To guarantee that prices are bounded, we could truncate the distribution of ε_t in far tails. This would not affect results.

²⁹ For example, we can have that $p_{t+\Delta} = (1 - \alpha\Delta)p_t + \alpha\Delta\phi(Q) + \varepsilon_{t+\Delta}$, which yields $f(Q, p_t) \equiv \alpha(\phi(Q) - p_t)$.

³⁰ We can model correlated demand in many other ways. For example, we can imagine that the noise is correlated over time. Such a model is much less tractable, because a deviating firm would have private information about future distribution of prices. That introduces the problem of private monitoring, making the analysis very difficult.

³¹ Conditioning on past actions can be advantageous only if the player mixed in the past, see footnote 19.

³² For example, see Hitoshi Matsushima (2004) and Johannes Hörner and Wojciech Olszewski (2006).

³³ Blume and Heidhues (2003) show that private strategies can improve upon PPE in a game with private signals, but their methods do not apply to our setting, since all signals in our model are public.

³⁴ From now on we take $x^* < \mu'$. Otherwise, the size is strictly positive as $\Delta \rightarrow 0$ and the Lemma is trivially satisfied.