

Online Appendix:

The Power of Communication

By David Rahman

I. Infrequent Coordination

According to [Sannikov and Skrzypacz \(2007\)](#), collusion breaks down in public Nash equilibrium. Here, though, firms collude in public communication equilibrium as long as equation (5) in the paper holds. If it fails, this agreement falls apart for the same reason as in [Sannikov and Skrzypacz's](#) result: obedient firms cannot be identified, so discouraging overproduction requires punishing everyone whenever prices are low. By Lemma 1 of the paper, this punishment cost is of order $r\sqrt{\Delta t}$, but the benefit, in terms of current collusion, is of order $r\Delta t$. As $\Delta t \rightarrow 0$, the costs overwhelm any benefits. Nevertheless, firms can still collude in private communication equilibrium even if (5) fails, as I show next. Intuitively, firms coordinate infrequently: they agree that it takes more than one low price to trigger punishment, thus tempering punishment costs. Just which low prices trigger punishment is temporarily kept secret from firms to recycle incentives.

For simplicity, $P(q) = q^{-e}$ with $0 < e < 1$ and firms face a constant marginal cost of $\gamma > 0$,¹ although the equilibrium construction clearly extends beyond this case. I also simplify the previous process to a random walk: firms observe the stock of prices, \hat{p}_t , whose innovations $\Delta\hat{p}_t = \hat{p}_t - \hat{p}_{t-\Delta t}$ equal $\sigma\sqrt{\Delta t}$ (resp. $-\sigma\sqrt{\Delta t}$) with probability $\Pi(q_t) = \frac{1}{2}[1 + (P(q_t)/\sigma)\sqrt{\Delta t}]$ (resp. $1 - \Pi(q_t)$). This change is innocuous, since both processes tolerate the equilibrium below, but helps to distill the argument.²

Fix a block of calendar time $c > 0$. In every period of this block, the mediator fulfills two functions. As before, he secretly recommends an output profile to each firm, IID across periods, to be publicly revealed at the end of the block. He also constructs a latent score $S_i = \{S_{it}\}$ for each firm i . S_i follows a random walk whose drift depends on both realized prices and his secret recommendations. At the end of the block, if

¹I only introduce these marginal costs so that the static Cournot outcome is easily well defined.

²The process above is a random walk representation of—hence converges as $\Delta t \rightarrow 0$ to—the previous model's process for the stock of prices. This limit process is a Brownian motion with law $d\hat{p}_t = P(q_{1t} + q_{2t})dt + \sigma dZ_t$, where Z stands for Wiener process. Moreover, being binomial, the random walk above withstands [Fudenberg and Levine's \(2007; 2009\)](#) hemi-continuity concerns.

S_{ic} exceeds a given threshold $\bar{z}c$ (linear in c) then firm i is punished by losing a given amount of continuation value w ; otherwise its continuation value remains the same. The score is constructed so that its drift equals zero in equilibrium, but any deviation raises the drift. This way, firms learn nothing about their score while they follow the mediator’s recommendations, so they are never confident of eluding punishment. By construction, every deviation raises the probability of such punishment significantly relative to any deviation gains. As a result, the threat of punishment always looms—so much so that it discourages firms from deviating altogether.

A. Overview of the Literature on Block Strategies

Block strategies of some sort or another have been pointed out in the literature, but all previous attempts to sustain collusion in this way fail under either public monitoring, frequent actions or both. Radner (1985, 1986) introduced “review strategies,” which punish players for statistical deviations from prescribed behavior. However, since players play pure—hence public—strategies, they remain subject to the basic problems of value-burning. This is epitomized in Radner, Myerson and Maskin (1986), which offers a classic example of a game with unsurmountable inefficiency because of value-burning and despite access to review strategies.

Abreu, Milgrom and Pearce (1991) assume that public signals (such as prices) arrive in blocks of fixed length of time, rather than every period, and exogenously manipulated the block length T . They constructed equilibria that economized on punishment as follows. When the signal arrives every period, public equilibria punish every time the public signal is “bad news.” With delay and lumping of the signals, it is possible to avoid punishing every such time and only punish sometimes. Indeed, Abreu, Milgrom and Pearce (1991) construct equilibria where players are punished only after T “bad news” signals. As Levin (2003, p. 847) suggested, Fuchs (2007) applied their construction to relational contracting. This approach suffers from two basic problems. First, it breaks down if public signals are not delayed: once a firm knows of at least one “good news” signal the fear of punishment disappears. Second, it is too lenient: incentives break down with frequent actions and imperfect monitoring (for instance, Brownian motion), even with arbitrary delay of exogenous information. See Rahman (2013b) for details.

Kandori and Matsushima (1998, motivated by Matsushima, 1995) and Compte (1998) overcame some of these problems by considering games with private monitoring and

conditionally independent signals. They replaced exogenous information delay with endogenous delay of other players’ reported signals, on which continuation values were allowed to depend. Players monitored and scored others directly through their signals over blocks of time. Relying on conditional independence, the same signal latency is possible as with exogenous delay. This restores collusion as in [Abreu, Milgrom and Pearce \(1991\)](#). Of course, though, the approach breaks down completely with public monitoring, which suggests that public monitoring can be a greater challenge to collusion than private monitoring. Additionally, [Kandori and Matsushima \(1998\)](#) considered a less lenient punishment trigger, but still lenient enough to temper value-burning. However, this trigger offers limited success at best with frequent actions, rendering full collusion generally unattainable. Again, see [Rahman \(2013b\)](#) for details.

[Obara \(2009\)](#) and [Obara and Rahman \(2006\)](#) extend [Compte \(1998\)](#) to allow for correlated signals, but fail with frequent actions. [Fong et al. \(2011\)](#) build on the work of [Matsushima \(2004\)](#) and [Ely, Hörner and Olszewski \(2005\)](#) in a two-player Prisoners’ Dilemma with private monitoring, no communication, and a small amount of correlation amongst players’ signals. There are several basic problems with their construction which mine overcomes. First, they require “sufficiently private monitoring,” which completely eliminates the more challenging case of public signals. Secondly, their construction breaks down with frequent actions, even under private monitoring, since it relies on sufficiently accurate hypothesis testing and this fails with imperfect monitoring in the limit. Lastly, their construction is implicit and does not apply to general games.

Below, I overcome all these challenges by deriving scoring rules and thresholds that obtain virtually full collusion in the frequent actions limit regardless of whether signals are public or private, and using other players’ recommendations to make these scores latent. Since recommendations are independent of actual behavior, their purpose is precisely to encrypt a player’s score.

B. Construction and Incentive Properties of Latent Scores

Usually, firms are asked to produce the collusive duopoly output, $q_d = \frac{1}{2}q_m$, but are occasionally asked to underproduce, q_u , and overproduce, q_o , for some given output levels q_u and q_o such that $0 < q_u < q_d < q_o$. Every period, the mediator recommends an output profile $(q_1, q_2) \in \{q_u, q_d, q_o\} \times \{q_u, q_d, q_o\}$ with probability $\mu_{12} > 0$, where μ_{dd} is close to 1. The rest of μ is chosen such that μ is symmetric ($\mu_{12} = \mu_{21}$) and—to

establish Lemma 1 below—so that, given $q_1 \in \{q_u, q_d, q_o\}$,

$$(\mu_{1u} + \mu_{1o})(q_1 + q_d)^{-(1+e)} = \mu_{1u}(q_1 + q_u)^{-(1+e)} + \mu_{1o}(q_1 + q_o)^{-(1+e)}.^3 \quad (1)$$

To construct firm i 's latent score, start with $S_{i0} = 0$. For all $t \in \Delta t \cdot \{1, \dots, \lfloor c/\Delta t \rfloor\}$, let $S_{it} = S_{it-\Delta t} \pm \sqrt{\Delta t}$ with probability $\frac{1}{2}$ if $\Delta \hat{p}_t < 0$ and $\frac{1}{2} \pm \zeta_i(q_{1t}, q_{2t})$ if $\Delta \hat{p}_t > 0$, where (q_{1t}, q_{2t}) is the mediator's recommendation profile at time t ,

$$\zeta_1(q_1, q_2) = \frac{\frac{1}{2}\alpha_1(q_1)}{\Pi(q_1 + q_2)} \quad \text{if } q_2 \neq q_d, \quad \zeta_1(q_1, q_d) = -\frac{\frac{1}{2}\alpha_1(q_1)}{\Pi(q_1 + q_d)} \frac{\mu_{1u} + \mu_{1o}}{\mu_{1d}},$$

ζ_2 is defined symmetrically to ζ_1 and $\alpha_i(q_i) > 0$ is a constant such that $\frac{1}{2} \pm \zeta_i(q_1, q_2)$ is a probability for all $\Delta t > 0$ sufficiently small. Intuitively, ζ determines stochastically when price increases raise or lower a firm's score. It tends to raise a firm's score if the other firm was asked to under- or over-produce, and otherwise tends to lower it.

Lemma 1. *In equilibrium, every firm i has a driftless score S_i , even conditional on i 's information. If i ever deviates then S_i has a positive drift for all small $\Delta t > 0$.*

Proof. Without loss, focus on firm 1. By construction, S_{1t} has no drift if \hat{p}_t drops. Otherwise, if firm 1 deviates from q_1 by h and \hat{p}_t rises, S_{1t} increases with probability

$$\Pr(\Delta S_{it} > 0 | q_1, h, \Delta \hat{p}_t > 0) = \frac{1}{2} + \frac{\sum_{q_2} \zeta_1(q_1, q_2) \mu_{12} \Pi(q_1 + q_2 + h)}{\sum_{q_2} \mu_{12} \Pi(q_1 + q_2 + h)}.$$

In equilibrium, $h = 0$. Substituting for ζ_1 , it follows that the numerator above equals $\frac{1}{2}\alpha_1(q_1)[\mu_{1u} - (\mu_{1u} + \mu_{1o}) + \mu_{1u}] = 0$. That is, S_1 has no drift given firm 1's information, hence unconditionally, too. If $h \neq 0$ then, after rearrangement, the numerator equals

$$\frac{1}{2}\alpha_1(q_1) \left[\mu_{1u} \frac{\Pi(q_1 + q_u + h)}{\Pi(q_1 + q_u)} - (\mu_{1u} + \mu_{1o}) \frac{\Pi(q_1 + q_d + h)}{\Pi(q_1 + q_d)} + \mu_{1o} \frac{\Pi(q_1 + q_o + h)}{\Pi(q_1 + q_o)} \right].$$

A first-order Taylor series expansion around $\sqrt{\Delta t} = 0$ yields $\frac{1}{4}\mu_1 z(h|q_1) \sqrt{\Delta t}$, where $z(h|q_1) = \frac{2\alpha_1(q_1)}{\sigma\mu_1} [\mu_{1u} \Delta P(h|q_1 + q_u) - (\mu_{1u} + \mu_{1o}) \Delta P(h|q_1 + q_d) + \mu_{1o} \Delta P(h|q_1 + q_o)]$, $\Delta P(h|q) = P(q + h) - P(q)$ and $\mu_1 = \sum_{q_2} \mu_{12}$ is the marginal probability of q_1 .

I will now use (1) to show that $z(h|q_1) > 0$ if $h \neq 0$, therefore S_i has positive drift. Let $f(h) = \mu_{1u} \Delta P(h|q_1 + q_u) - (\mu_{1u} + \mu_{1o}) \Delta P(h|q_1 + q_d) + \mu_{1o} \Delta P(h|q_1 + q_o)$. Since clearly $f(0) = 0$ and f is differentiable, I will equivalently show that $f'(h) > 0$ if $h > 0$, $f'(h) < 0$ if $h < 0$, and $f'(0) = 0$. Let $\rho = 1 + e > 0$. Substituting for $P(q)$,

$$f'(h) = -e[\mu_{1u}(q_1 + q_u + h)^{-\rho} - (\mu_{1u} + \mu_{1o})(q_1 + q_d + h)^{-\rho} + \mu_{1o}(q_1 + q_o + h)^{-\rho}].$$

³It is easy to see that these 3 linear equations (one for each q_1) in μ have infinitely many solutions: by symmetry, probabilities adding up to one and μ_{dd} being given, there are 4 unknowns.

Define $\tilde{q}_1(h)$ to satisfy

$$\tilde{q}_1(h|\rho) = [\gamma_u(q_1 + q_u + h)^{-\rho} + \gamma_o(q_1 + q_o + h)^{-\rho}]^{-1/\rho},$$

where $\gamma_u = \mu_{1u}/(\mu_{1u} + \mu_{1o})$ and $\gamma_o = 1 - \gamma_u$.

Claim 1. $d\tilde{q}_1/dh > 1$.

Proof of Claim 1. Let $q_{1uh} = q_1 + q_u + h$ and $q_{1oh} = q_1 + q_o + h$. By routine calculations,

$$\frac{d\tilde{q}_1}{dh} = \frac{\gamma_u q_{1uh}^{-\rho-1} + \gamma_o q_{1oh}^{-\rho-1}}{\tilde{q}_1(h|\rho)^{-\rho-1}} = \left[\frac{\tilde{q}_1(h|\rho + 1)}{\tilde{q}_1(h|\rho)} \right]^{-\rho-1},$$

so $d\tilde{q}_1/dh > 1$ if and only if $\tilde{q}_1(h|\rho + 1) < \tilde{q}_1(h|\rho)$. Moreover,

$$\begin{aligned} \frac{d\tilde{q}_1}{d\rho} &= \frac{d}{d\rho} \exp \left\{ -\frac{1}{\rho} \ln(\gamma_u q_{1uh}^{-\rho} + \gamma_o q_{1oh}^{-\rho}) \right\} \\ &= \tilde{q}_1(h|\rho) \left[\frac{1}{\rho^2} \ln(\gamma_u q_{1uh}^{-\rho} + \gamma_o q_{1oh}^{-\rho}) + \frac{1}{\rho} \frac{\gamma_u q_{1uh}^{-\rho} \ln q_{1uh} + \gamma_o q_{1oh}^{-\rho} \ln q_{1oh}}{\gamma_u q_{1uh}^{-\rho} + \gamma_o q_{1oh}^{-\rho}} \right] \\ &= \tilde{q}_1(h|\rho)^{\rho+1} \frac{1}{\rho^2} \left[(\gamma_u q_{1uh}^{-\rho} + \gamma_o q_{1oh}^{-\rho}) \ln(\gamma_u q_{1uh}^{-\rho} + \gamma_o q_{1oh}^{-\rho}) \right. \\ &\quad \left. - \gamma_u q_{1uh}^{-\rho} \ln q_{1uh}^{-\rho} - \gamma_o q_{1oh}^{-\rho} \ln q_{1oh}^{-\rho} \right] < 0, \end{aligned}$$

where the last inequality follows because $x \ln x$ is clearly a convex function of x . Since this is true for every ρ , it follows that $\tilde{q}_1(h|\rho + 1) < \tilde{q}_1(h|\rho)$, hence $d\tilde{q}_1/dh > 1$. \square

By construction, (1) implies that $f'(0) = 0$. Therefore, since, by Claim 1, $h > 0$ implies that $q_1 + q_d + h < \tilde{q}_1(h|\rho)$, hence $(q_1 + q_d + h)^{-\rho} > \tilde{q}_1(h|\rho)^{-\rho}$, it follows that $f'(h) > 0$. On the other hand, if $h < 0$ then, again by Claim 1, $q_1 + q_d + h > \tilde{q}_1(h|\rho)$, so $(q_1 + q_d + h)^{-\rho} < \tilde{q}_1(h|\rho)^{-\rho}$ and $f'(h) < 0$. This finally establishes Lemma 1. \square

In the proof of Lemma 1, $z(h|q_1)$ was defined to be the drift of S_i given i 's information, since $\Pr(\Delta S_{it} > 0 | q_1, h, \Delta \hat{p}_t > 0) \approx \frac{1}{2}[1 + z(h|q_1)\sqrt{\Delta t}]$ by a similar Taylor series expansion to the one there. The unconditional drift of S_i equals $\frac{1}{2}z(h|q_1)$: it is easy to see that $\Pr(\Delta S_{it} > 0 | q_1, h) \approx \frac{1}{2}[1 + \frac{1}{2}z(h|q_1)\sqrt{\Delta t}]$ before observing prices.

The graph on the left of Figure 1 illustrates Lemma 1 by showing how every deviation increases the drift of i 's score S_i . As $\Delta t \rightarrow 0$, S_i converges to a Brownian motion with $S_{ic} \sim N(0, c)$ in equilibrium, so the probability of punishment is approximately $1 - \Phi(\bar{z}\sqrt{c})$ when the punishment cutoff is $\bar{z}c$. Many cutoffs give the right incentives; I use the largest drift from any deviation, assuming that a firm produces at least 0:

$$\bar{z} = \sup_{h, q_1} \{z(h|q_1) : q_1 \in \{q_u, q_d, q_o\}, q_1 + h \geq 0\}.$$




Figure 1: Drift of i 's score as a function of i 's deviation (left) and a sample path of i 's score in blue against a time-linear cutoff in red (right).

It is easy to see that $0 < \bar{z} < \infty$.⁴ I will now use the cutoff \bar{z} to find a punishment w that discourages firms from disobeying the mediator. I apply the approach of Rahman (2013a) for general repeated games to this Cournot oligopoly and derive w intuitively; see Rahman (2013a) for formal details. To apply this approach, I assume that the magnitude of deviations is bounded below—as I show, the agreement above cannot discourage profitable infinitesimal deviations in this environment.

Definition 1. Let $\eta > 0$. A deviation h from q is *feasible* if $|h| \geq \eta$ and $h + q \geq 0$.

This assumption can be easily dropped if firms face a kinked demand curve $P(q)$, can only produce in discrete amounts, have finitely many output choices, or incur fixed output adjustment costs. Otherwise, I need it to sustain virtually full collusion.

Let $\Delta R(h|q_1) = \sum_{q_2} [(q_1 + h)P(q_1 + q_2 + h) - q_1P(q_1 + q_2) - \gamma h]\mu_{12}/\mu_1$ be a firm's expected profit change from deviating by h after being asked to produce q_1 . Define

$$\frac{dR}{dz} = \sup_{h, q_1} \{ \Delta R(h|q_1)/z(h|q_1) : q_1 \in \{q_u, q_d, q_o\}, |h| \geq \eta, q_1 + h \geq 0 \}.$$

dR/dz is a firm's maximal incentive cost—the highest ratio of profit from a deviation per unit increase in the drift of i 's score. If deviations are feasible, $z(h|q_1)$ is bounded below by some $\underline{z} > 0$, so $dR/dz < \infty$. Let φ be the standard normal PDF.

Lemma 2. *Every feasible deviation is discouraged for all small enough $\Delta t > 0$ if*

$$w > \frac{rce^{rc}}{\varphi(\bar{z}\sqrt{c})^{\frac{1}{2}}\sqrt{c}} \frac{dR}{dz}. \quad (2)$$

⁴By Lemma 1, $\bar{z} > 0$, $z(h|q_1)$ is decreasing in $h < 0$, so $z(h|q_1) \leq z(-q_1|q_1) < \infty$ for $h \in [-q_1, 0]$, and $z(h|q_1)$ is increasing in $h > 0$. Now $\bar{z} < \infty$ follows because z is bounded, since clearly $z(h|q_1)$ converges to $\alpha_1(q_1)[(\mu_{1u} + \mu_{1o})P(q_1 + q_d) - \mu_{1u}P(q_1 + q_u) - \mu_{1o}P(q_1 + q_o)] < \infty$ as $h \rightarrow \infty$.

Proof Sketch. If firm 1 deviates by h after q_1 was recommended in the first period, the drift of S_1 will increase by about $\frac{1}{2}z(h|q_1)$ before observing the first period's price. With no more deviations, the punishment probability increases from about $1 - \Phi(\bar{z}\sqrt{c})$ to about $1 - \Phi([\bar{z}c - \frac{1}{2}z(h|q_1)\Delta t]/\sqrt{c})$. To discourage this deviation, future costs $e^{-rc}w[1 - \Phi([\bar{z}c - \frac{1}{2}z(h|q_1)\Delta t]/\sqrt{c}) - (1 - \Phi(\bar{z}\sqrt{c}))]$ must outweigh current gains $(1 - e^{-r\Delta t})\Delta R(h|q_1) \approx r\Delta t\Delta R(h|q_1)$. As $\Delta t \rightarrow 0$, this rearranges in flows to

$$r\Delta R(h|q_1) < e^{-rc}w\varphi(\bar{z}\sqrt{c})\frac{1}{2}z(h|q_1)/\sqrt{c}. \quad (3)$$

Divide both sides by $z(h|q_1)$, take the supremum and rearrange to obtain (2). To discourage every other deviation, start at the last period of the block. Ignoring discounting, current gain flows are bounded above by the LHS of (3). The flow of punishment costs is just like the RHS of (3) except that $\varphi(\bar{z}\sqrt{c})$, the punishment probability flow, is replaced with $\varphi([\bar{z} - \theta]\sqrt{c})$ for some $\theta \geq 0$, since past deviations cannot lower the drift of S_1 . Since \bar{z} is as large as any drift increase, $\theta \leq \bar{z}$. Hence, $\varphi([\bar{z} - \theta]\sqrt{c}) \geq \varphi(\bar{z}\sqrt{c})$, so (3) still holds with the probability flow adjusted by θ . Last-period deviations are now discouraged (see Figure 2 below). But this argument applies also to previous periods, so by induction every deviation is discouraged. \square

By Lemma 2, $w \rightarrow 0$ as $r \rightarrow 0$ for any $c > 0$, so w is feasible for sufficiently patient firms. Let $R = \sum_{(q_1, q_2)} q_1 P(q_1 + q_2) \mu_{12}$ be a firm's expected revenue when everyone follows the mediator. Using (3), lifetime equilibrium payoffs u are given by

$$u \approx (1 - \delta)R + \delta[\Phi(\bar{z}\sqrt{c})u + (1 - \Phi(\bar{z}\sqrt{c}))(u - w)] \approx R - \frac{rc}{1 - e^{-rc}} \frac{1 - \Phi(\bar{z}\sqrt{c})}{\varphi(\bar{z}\sqrt{c})\frac{1}{2}\sqrt{c}} \frac{dR}{dz}.$$

Sufficiently patient firms can therefore sustain virtually full collusion by coordinating infrequently to rarely punish each other, since normal hazard rates explode linearly:

$$\lim_{c \rightarrow \infty} \lim_{r \rightarrow 0} u = R.$$

Given a smooth set W of payoff profiles as in Figure 2 above, Rahman (2013a) uses punishments and rewards to self-decompose W in “block-public” communication equilibrium. To apply this result and complete the equilibrium construction above, I verify below a necessary bound on deviation gains. I also offer general sufficient conditions on an arbitrary demand curve P for such a Folk Theorem.

As in the previous section, firms can dispense with the mediator by communicating their intentions directly, although this time with a delay of c instead of just Δt . Notice also that this construction generalizes easily to other environments, such as




Figure 2: Discouraging multiple deviations (left) and feasible punishments (right).

ones with private monitoring and information. For instance, private monitoring is arguably easier because firms can use each other's private observations to keep scores secret, which is the key to recycling incentives that facilitates collusion.

C. Folk Theorem with Bounded Deviation Gains

Rahman (2013a, Theorem 2) establishes a Folk Theorem for repeated games with frequent actions that satisfy a conditional identifiability condition and a bound on deviation gains. The former requirement is equivalent to the existence of a score that has no drift after obeying the mediator but whose drift increases with every deviation. Since such a score exists, the demand curve $P(q) = q^{-e}$ satisfies conditional identifiability. For some demand curves, such as a linear one, this condition fails. By definition, P exhibits conditional identifiability if

$$P(\cdot|q) \notin \text{conv}\{P(\cdot|q+h) : q+h \geq 0\} + L_1 \quad \forall q \geq 0, \quad (4)$$

where $P(x|q) = P(q+x)$ and $L_1 = \{\lambda \mathbf{1} : \lambda \in \mathbb{R}\}$ is the line generated by the constant function. It is easy to see that both $P(q) = \alpha - \beta q$ and $P(q) = e^{-\gamma q}$ fail conditional identifiability, but many other demand curves satisfy it.

Write monopoly profit as $R_m = q_m[q_m^{-e} - \gamma] = e[(1-e)/\gamma]^{(1-e)/e}$ and a firm's profit from best-responding to monopoly output as $R_m^d = q^*(q_m)[(q^*(q_m) + q_m)^{-e} - \eta]$, where $q^*(q)$ is a firm's best response to q . To apply the Folk Theorem, I also assume that

$$R_m^d < \frac{1}{2}R_m. \quad (5)$$

Claim 2. (5) is not pathological: for every $\gamma > 0$ there exists a nonempty open set of parameter values $E \subset (0, 1)$ such that $e \in E$ is consistent with (5).

Proof. The Cournot output when i 's marginal cost is γ_i equals $q_i = \bar{q}[\sigma_i + (\sigma_j - \sigma_i)/e]$, where $\sigma_i = \gamma_i/\bar{\gamma}$, $\bar{\gamma} = \gamma_1 + \gamma_2$ and $\bar{q} = q_1 + q_2 = [(2 - e)/\bar{\gamma}]^{1/e}$. Therefore, profit equals $R_i = e[\sigma_j + (\sigma_i - \sigma_j)/e]^2 e [(2 - e)/\bar{\gamma}]^{(1-e)/e}$. Firm 2 producing q_m corresponds to an equilibrium with $\gamma_1 = \gamma$ and γ_2 solving $\bar{q}[\sigma_2 + (\sigma_1 - \sigma_2)/e] = [(1 - e)/\gamma]^{(1-e)/e}$. If $e = 1/2$, for instance, then σ_1 solves $9\sigma_1^2(3\sigma_1 - 1) = 1$. The unique solution to this cubic equation is $\sigma_1 \approx 0.489$, so $\sigma_2 = 0.511$. By direct calculation, (5) holds if

$$\frac{\sigma_1 + (\sigma_2 - \sigma_1)/e}{\sigma_1} \frac{\sigma_1 + (\sigma_2 - \sigma_1)/e}{\sigma_2 + (\sigma_1 - \sigma_2)/e} \frac{1 - e}{2 - e} < \frac{1}{2}. \quad (6)$$

When $\sigma_1 \approx 0.489$, (6) holds. By continuity, it still holds near $e = 1/2$. \square

Intuitively, (5) means that a firm is better off with its share of monopoly rents than best-responding to the other firm monopolizing the market. I use this assumption as follows. Broadly, the Folk Theorem in [Rahman \(2013a\)](#) gives individual punishments and rewards to firms independently of each other, in contrast with [Fudenberg, Levine and Maskin \(1994\)](#), say, who correlate continuation values across firms to avoid value-burning. As a result, a firm that is rewarded must be compensated for every deviation with its reward. Of course, any such deviation ought to lower the likelihood of reward, yet in equilibrium the expected reward must still compensate for every deviation. Hence, the feasible lifetime average payoff for a firm being motivated to produce a temporarily disadvantageous amount with rewards must accommodate at least the best deviation from its temporary disadvantage. Punishments, on the other hand, can be rare in equilibrium as long as deviations increase the chance of punishments. This argument originates in [Compte \(1998\)](#) in a discrete game with private monitoring. For the monopoly profit line to be approachable for patient players, at least part of it must remain after restricting reward-driven payoffs to compensate every deviation. This is what (5) achieves. See [Figure 2](#) above and [Figure 3](#) below.

Formally, let W be a smooth set of payoff profiles in the interior of U_n , the set of feasible payoffs in excess of the static Nash equilibrium (q_n, q_n) , where

$$U_n = \{R(q_1, q_2) \geq R(q_n, q_n) : (q_1, q_2) \geq 0\}$$

and $R(q_1, q_2) = (q_1, q_2)[(q_1 + q_2)^{-e} - \gamma]$ is the profit profile when firms produce (q_1, q_2) . For W to be a subset of equilibrium payoffs, it must be locally self-decomposable with continuation values at the end of a given block of length c . Just as with secret monitoring, local self-decomposability is necessary at each point of ∂W . Consider a point whose outward normal vector gives positive weight to one firm and negative weight to the other, as in the bottom right-hand corner of [Figure 2](#) above. At this




Figure 3: Providing incentives with rewards.

point, the firm with negative weight must be given incentives with rewards, therefore its expected continuation value must increase by at least its best deviation gains, as was argued in the previous paragraph.

Formally, this is expressed as follows. Let \mathcal{M}_ε be the set of probability measures over output profiles in \mathbb{R}_+^2 with finite support and at least ε probability on each element. Let

$$U_{\varepsilon_1, \varepsilon_2}^J = \{u \in U_n : \exists \mu \in \mathcal{M}_k^{\varepsilon_1} \text{ s.t. } u_i \geq R_i^d(\mu) + \varepsilon_2 \forall i \notin J, u_i \leq R_i(\mu) - \varepsilon_2 \forall i \in J\},$$

where $R_i(\mu) = \sum_{(q_1, q_2)} R_i(q_i, q_j) \mu(q_i, q_j)$ is firm i 's profit from everyone following μ and

$$R_i^d(\mu) = \max_h \left\{ \sum_{(q_1, q_2)} R_i(q_i + h(q_i), q_j) \mu(q_i, q_j) : q_i + h(q_i) \geq 0 \right\}$$

is firm i 's profit after deviating according to $h(q_i)$ when q_i was recommended. Moreover, define

$$U_{\varepsilon_1, \varepsilon_2}^* = \bigcap_{J \subset \{1, 2\}} U_{\varepsilon_1, \varepsilon_2}^J \quad \text{and} \quad U^* = \bigcup_{\varepsilon_1, \varepsilon_2 \gg 0} U_{\varepsilon_1, \varepsilon_2}^*.$$

Claim 3. $R(q_d, q_d) \in \overline{U^*}$ if (5) holds, so by [Rahman \(2013a, Theorem 2\)](#) sufficiently patient firms can sustain virtually full collusion.

Proof. First, notice that if $u \in \text{int} U_n$ and $J = \{1, 2\}$ or $J = \emptyset$ then there exist $\varepsilon_1 > 0$ and $\varepsilon_2 > 0$ arbitrarily small such that $u \in U_{\varepsilon_1, \varepsilon_2}^J$. If $J = \{1, 2\}$, this is clear: choose μ to be a convex combination of $(0, q_m)$, $(q_m, 0)$ and (q_n, q_n) close enough to u that $u_i \leq R_i(\mu) - \varepsilon_2$. If $J = \emptyset$ then again the result is clear: let μ put all its probability mass on (q_n, q_n) . For sufficiently small $\varepsilon_2 > 0$, every $u \in \text{int} U_n$ satisfies

$u_i \geq R_i^d(\mu) + \varepsilon_2 = R_i(\mu) + \varepsilon_2$, where $R_i^d(\mu) = R_i(\mu)$ because μ is a Nash equilibrium of the stage game. Now let $J = \{1\}$. If $\mu = (q_m, 0)$ then $R_m^d < \frac{1}{2}R_m$, so for $\varepsilon_1 > 0$ and $\varepsilon_2 > 0$ small enough $R(q_d, q_d) - \varepsilon_2 \in U_{\varepsilon_1, \varepsilon_2}^J$. By symmetry, the same argument applies to $J = \{2\}$.

Theorem 2 of [Rahman \(2013a\)](#) defines U slightly differently with μ having a fixed support, but this is just to ensure conditional identifiability at the given support and doesn't change the result: by continuity, replace every μ here with one that places arbitrarily small probability on some other appropriate output levels q_u and q_o so that (4) is satisfied with this new support and the payoff bounds of the previous paragraph are still respected within an arbitrarily small margin. \square

References

- Abreu, Dilip, Paul Milgrom, and David G. Pearce.** 1991. "Information and Timing in Repeated Partnerships." *Econometrica*, 59(6): 1713–33.
- Compte, Olivier.** 1998. "Communication in Repeated Games with Imperfect Private Monitoring." *Econometrica*, 66(3): 597–626.
- Ely, Jeffrey C., Johannes Hörner, and Wojciech Olszewski.** 2005. "Belief-Free Equilibria in Repeated Games." *Econometrica*, 73(2): 377–415.
- Fong, Kyna, Olivier Gossner, Johannes Hörner, and Yuliy Sannikov.** 2011. "Efficiency in the Repeated Prisoners' Dilemma with Private Monitoring." mimeo.
- Fuchs, William.** 2007. "Contracting with Repeated Moral Hazard and Private Evaluations." *American Economic Review*, 97(4): 1432–1448.
- Fudenberg, D., and D.K. Levine.** 2007. "Continuous time limits of repeated games with imperfect public monitoring." *Review of Economic Dynamics*, 10(2): 173–192.
- Fudenberg, D., and D.K. Levine.** 2009. "Repeated Games with Frequent Signals." *The Quarterly Journal of Economics*, 124(1): 233–265.
- Fudenberg, Drew, David Levine, and Eric Maskin.** 1994. "The Folk Theorem with Imperfect Public Information." *Econometrica*, 62(5): 997–1039.
- Kandori, Michihiro, and Hitoshi Matsushima.** 1998. "Private Observation, Communication, and Collusion." *Econometrica*, 66(3): 627–652.
- Levin, Jonathan.** 2003. "Relational Incentive Contracts." *American Economic Review*, 93(3): 835–847.
- Matsushima, Hitoshi.** 1995. "Multimarket contact, imperfect monitoring, and implicit collusion." working paper.
- Matsushima, Hitoshi.** 2004. "Repeated games with private monitoring: Two players." *Econometrica*, 72(3): 823–852.
- Obara, Ichiro.** 2009. "Folk theorem with communication." *Journal of Economic Theory*, 144(1): 120–134.

- Obara, Ichiro, and David Rahman.** 2006. “Secret Contracts for Efficient Partnerships.” working paper.
- Radner, Roy.** 1985. “Repeated principal-agent games with discounting.” *Econometrica: Journal of the Econometric Society*, 1173–1198.
- Radner, Roy.** 1986. “Repeated partnership games with imperfect monitoring and no discounting.” *The Review of Economic Studies*, 53(1): 43–57.
- 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(1): 59–69.
- Rahman, David.** 2013*a*. “Frequent Actions with Infrequent Coordination.” working paper.
- Rahman, David.** 2013*b*. “Information Delay in Games with Frequent Actions.” working paper.
- Sannikov, Y., and A. Skrzypacz.** 2007. “Impossibility of collusion under imperfect monitoring with flexible production.” *The American Economic Review*, 97(5): 1794–1823.