Dynamic Information Management in Repeated Games with Frequent Actions

David Rahman*
University of Minnesota

First Draft: August 27, 2012
This Version: November 7, 2012

Abstract

I study repeated games with mediated communication and frequent actions and obtain a Folk Theorem in case side transfers fail to provide the right incentives but otherwise a strong form of individual identifiability holds. Even in the limit, when noise is driven by Brownian motion and actions are arbitrarily frequent, as players become increasingly patient it is possible for players to attain equilibrium payoffs far beyond static outcomes. This brings together the work on repeated games in discrete and continuous time. As an application, I suggest how firms may mediate dynamic collusive agreements in oligopoly.

JEL Classification: D21, D23, D82.
Keywords: folk theorem, private strategies, mediated communication.

*I thank the National Science Foundation for financial support through Grant No. SES 09-22253.

1 Introduction

Recently, the topic of repeated games in continuous time/with frequent actions has received considerable attention in leading economics journals (Sannikov, 2007; Sannikov and Skrzypacz, 2007, 2010; Faingold and Sannikov, 2011; Fudenberg and Levine, 2007, 2009). Arguably, one of the main results of this burgeoning literature is that when players' noisy observations are driven by Brownian motion, it is impossible to obtain (symmetric) outcomes beyond the repetition of static equilibria, i.e., equilibria of the stage game being repeated. All these papers share one key assumption: they restrict attention to equilibria in public strategies.

In this paper I argue that removing their restriction to equilibria in public strategies restores the possibility of truly dynamic equilibria, well beyond the repetition of static equilibria. In doing so, I reconcile the theory of repeated games in discrete time with its continuous time counterpart. To make this point as clearly as possible, as well as to deliver a basic interpretation of how information is managed dynamically to obtain the possibility result below, I employ *mediated strategies* (Forges, 1986; Myerson, 1986), which are a plausible generalization of private strategies. In many circumstances, mediated strategies are not strictly necessary for dynamic equilibria, but clarify exposition and understanding of the forces at play. See Section 6 for further discussion. A mediated strategy is the following information management institution: players are given confidential, non-binding recommendations by a disinterested party (or machine, or not-necessarily-public randomization device), called a "mediator," that depend on the mediator's own past observations. In general, these observations may include all kinds of previous communication with the players.

Information management institutions are useful generally, but to make their usefulness more difficult, I focus the model below on public monitoring. With private monitoring, it is even easier to implement the kinds of dynamic equilibria that I construct below. (See Section 6 for discussion.) But just how are these institutions useful? In the language of mediation, the answer is explained easily. It is well known (Abreu et al., 1990) that delaying the arrival of information to players helps to economize on incentives because they can be effectively "recycled." Yet, in repeated games with public monitoring, the public signal cannot be delayed without changing the game's underlying information structure. However, the mediator can delay the arrival to a player of others' recommendations, and use that delay to economize on incentives—so much so that truly dynamic strategies can be equilibria with Brownian noise.

It is perhaps best to explain the basic argument in two parts. In the first part, I make infeasible changes to an information structure, and in the second, I apply these changes in a feasible way. I begin with the first part, but let me say right away that although the general approach below is motivated by the insight of Abreu et al. (1990), their construction does not work to overcome the impossibility results for Brownian motion. Furthermore, although a version of the construction by Kandori and Matsushima (1998, Theorem 3) overturns the implausibility results for Brownian motion, in the sense that something may be possible beyond static Nash, it is not enough to obtain a Folk Theorem. I propose an intermediate approach that not only overturns the impossibility results, but also delivers a Folk Theorem. See Section 2.

For intuition, consider first the Prisoners' Dilemma with imperfect public monitoring, whose binary signals carry probabilities that are consistent with the random walk representation of an arithmetic Brownian motion as the time interval between interactions vanishes. Of course, we assume that the volatility parameter of the Brownian motion does not depend on the action profiles, although the drift is lower when a player defects than when he cooperates. One way of summarizing the impossibility results in the literature is as follows. The best (public) symmetric equilibrium payoff must necessarily involve some punishment after a down jump, and indeed may be calculated as the cooperative payoff net of incentive costs: deviation payoffs divided by the likelihood ratio test of whether or not someone deviated. As the time between interactions, $\Delta > 0$, diminishes, this likelihood ratio converges to zero, therefore the incentive cost becomes arbitrarily large, and all (symmetric) incentives break down as $\Delta \to 0$, for any fixed interest rate $r > 0$.

Consider a T-period block of length $T = c/\Delta$, where c is a constant. Such T-period block has the same "calendar time" length, regardless of Δ. Suppose that the arrival of all the public signals in the block can be delayed until the Tth period. Begin recommending that everyone cooperates. If p is the probability of a down-jump when both players cooperate and $q > p$ when one player defects, construct a score based on the empirical likelihood of a unilateral deviation: if there are more than qT down jumps in a T-period block then mutual defection ensues with some probability, call it Π. Otherwise, "rinse and repeat:" the score is reset to zero and a new T-period block begins with the same rules as before.[1]

[1] Kandori and Matsushima (1998) used pT instead of qT, which is not enough for a Folk Theorem in this paper's model. See Section 2 for details.

If the punishment is large enough to deter a single deviation then it will deter all deviations, because, by delaying the arrival of information (the public signal), it is possible to recycle incentives. The size of expected punishment necessary to deter a single deviation is approximately a one-period gain divided by the probability change due to a single deviation. By the Central Limit Theorem, this converges to a constant that depends on c and $r > 0$ as $\Delta \to 0$. Finally, as $r \to 0$ we may choose $c \to \infty$ so that in the limit as $r \to 0$, there is a symmetric equilibrium where payoffs are arbitrarily close to those from full cooperation. But here we had already set $\Delta \to 0$, so in the limit of repeated games with frequent actions, sufficiently patient players who are able to delay the arrival of some information can obtain a Folk Theorem.

For the second part of the argument, I apply this idea to the delay of the arrival of others' recommendations rather than the arrival of some public signal, but otherwise apply the overall construction just described. I make assumptions on payoffs and signal probabilities to make my point simply and generally. In this sense, I broadly follow Compte (1998), although there are important differences in my approach. Mostly I make technical assumptions so that I can easily extend and apply the algorithm of Fudenberg and Levine (1994) to T-public communication equilibria. With respect to payoffs, I make technical assumptions so that I can achieve a Folk Theorem without requiring pairwise identifiability. In terms of probabilities, I make a technical restriction that in the public monitoring case is stronger than individual identifiability (but not in the private monitoring case) but close in spirit to it, and much weaker than pairwise identifiability in that it is consistent with it being impossible to statistically identify a deviator after a unilateral deviation.

Identifying deviators is the spirit of pairwise identifiability; it is used extensively not only by Fudenberg et al. (1994) to prove their Folk Theorem, but also by Sannikov and Skrzypacz (2010). The latter argue that there is no way to sustain outcomes beyond repetition of static Nash equilibria if they involve value-burning as $\Delta \to 0$, for any $r > 0$, because incentive costs explode. However, they also show that it is possible to attain truly dynamic outcomes with "budget-balanced" continuation values, i.e., with transfers of value amongst players that do not burn any value. For this to render desirable outcomes enforceable, it is necessary to distinguish between the innocent and the guilty, to be able to reward some while punishing others and maintain budget balance (Rahman and Obara, 2010). When this is possible, one obtains a Folk Theorem, after taking the limit as $\Delta \to 0$, in a way that value-burning altogether. Information management institutions can add a significant amount of value in this context, too. For instance, in repeated oligopoly with capacity constraints, it is possible to sustain (almost) perfect collusion, because a sufficiently weak version of pairwise identifiability holds. What is possible, with a weaker version of pairwise identifiability that accommodates mediated communication, is a generalization of the Folk Theorem with public communication equilibria. See Rahman (2012b) for a detailed discussion of this issue, as well as Section 5 below. In this paper, however, I show that it is possible to provide incentives to players in repeated games with frequent actions without the use of such value transfers.

The main result of the paper is a "Nash-threats" Folk Theorem (although the threats are not Nash equilibria, but correlated equilibria), which argues that information management can lead to "mutual cooperation" even if value-burning is required in repeated games with frequent actions, as players becomes unboundedly patient. The order of limits matters in this statement: first $\Delta \to 0$ and then $r \to 0$. Taking limits in the opposite order is the standard approach of discrete-time repeated games, whereas taking limits in this way is the approach followed by, e.g., Sannikov (2007); Sannikov and Skrzypacz (2007, 2010); Faingold and Sannikov (2011). All these papers argue that it is not possible to have a Folk Theorem in continuous time when value-burning is necessary. However, the argument above suggests that this is indeed possible. The rest of the paper involves applying this idea to general repeated games, where the information that is delayed is the mediator's recommendations.

Other observations easily follow from the analysis of the paper. For instance, it is not difficult to see that when signals converge to a Poission arrival process, the construction above shows that full efficiency is achievable in the "bad news" case, but not in the "good news" case—see also Abreu et al. (1990). This is because in the bad news case the likelihood ratio does not vanish as $\Delta \to 0$, so delaying information allows for "reusing" incentive costs to deter more deviations than just one. On the other hand, in the good news case, the likelihood ratio decreases at rate Δ. Changing T does not help to prevent losses due to incentive costs from becoming unbounded.

The rest of the paper is organized as follows. In Section 2, I lay out a basic argument that changes the underlying information structure. In Section 3, I present the main model and assumptions. In Section 4, I try to apply the argument of Section 2 and obtain a Folk Theorem by delaying the arrival of recommendations instead of changing the information structure. In Section 5, I apply it to the Prisoners' Dilemma and repeated oligopoly. Finally, Sections 6 and Section 7 discuss the model and conclude.

2 Basic Argument

In this section, I will present a streamlined version of the basic argument that is generalized later. To get there, first I will argue that neither the approach due to Abreu et al. (1990) nor that due to Kandori and Matsushima (1998), despite being able to deliver a Folk Theorem in discrete time, is not able to deliver it in the continuous time limit with Brownian information. Then I argue that an intermediate approach, somewhere between the two approaches above, does work. All this discussion will take place around arguably the simplest possible context: the Prisoners' Dilemma.

2.1 Prisoners' Dilemma

Consider the Prisoners' Dilemma with imperfect public monitoring, repeated in discrete time, with common discount factor δ.

	C	D
C	u, u	$-b, u+g$
D	$u+g, -b$	$0, 0$

Payoffs

Here, $u > g - b$. There are two signals, γ and β, with $0 < \Pr(\beta|C,C) = p < q = \Pr(\beta|C,D) = \Pr(\beta|D,C) < 1$.

Public equilibrium payoffs with symmetric strategies, v, solve:

$$v = (1-\delta)u + \delta[(1-p)v + p(1-\alpha)v] \quad \Leftrightarrow \quad v = u - \frac{\delta}{1-\delta}p\alpha v,$$

$$v \geq (1-\delta)(u+g) + \delta[(1-q)v + q(1-\alpha)v] \quad \Leftrightarrow \quad \alpha v \geq \frac{1-\delta}{\delta}\frac{g}{q-p},$$

where $\alpha \in [0,1]$ is the probability of mutual defection henceforth. Since the incentive constraint above will bind, writing $\ell = q/p$ yields

$$v = u - \frac{g}{\ell - 1}.$$

The feasibility constraint that $0 \leq \alpha v \leq v$ implies that

$$u \geq \frac{g}{\ell-1}\left[1 + \frac{1-\delta}{\delta}p\right],$$

therefore, to attain the value $v = u - g/(\ell-1)$ it is necessary for δ to be large enough for this inequality to be satisfied. Clearly, such a δ exists strictly between 0 and 1.

2.2 Abreu et al. (1990) Fails

Recall first the argument due to Abreu et al. (1990). Their insight was to delay the arrival of information, so instead of the public signal arriving every period, no information arrives until the end of a T-period block, where all the information arrives at once. Their equilibrium construction relies on the following strategies: Start with mutual cooperation for the entire T-period block. If at the end of the block the vector β^T is observed then mutual defection henceforth will occur with probability α. Otherwise, the cooperative phase just described continues for another block. Now, T-public equilibrium symmetric payoffs can be found similarly as above:

$$v = (1-\delta^T)u + \delta^T[(1-p^T)v + p^T(1-\alpha)v] \quad \Leftrightarrow \quad v = u - \frac{\delta^T}{1-\delta^T}p^T\alpha v.$$

Incentive compatibility requires one more argument. First, assume that a player deviates only in the first period. The gain from such a deviation is $(1-\delta)g$. The cost of the deviation is $\delta^T(\ell-1)p^T\alpha v$. So α must satisfy $(1-\delta)g \leq \delta^T(\ell-1)p^T\alpha v$. It is easy to see that if this constraint is satisfied then all the incentive constraints in the T-period block are satisfied. Indeed, for τ deviations, the gain is at most $(1-\delta)\tau g$, whereas the cost is $\delta^T(\ell^\tau-1)p^T\alpha v$. Since $\ell > 1$, it follows that $\ell^\tau - 1 \geq \tau(\ell-1)$, which implies incentive compatibility for τ deviations from a single one.

Recognizing that this constraint will bind at an optimum and rearranging yields

$$v = u - \frac{1-\delta}{1-\delta^T}\frac{g}{\ell-1}.$$

As $\delta \to 1$, $v \to u - \frac{1}{T}\frac{g}{\ell-1}$, and as $T \to \infty$, $v \to u$, which delivers a Folk Theorem in discrete time, as long as it is possible to delay information.

This argument fails in repeated games with frequent actions, i.e., when the limit is taken with respect to Δ instead of δ, because ℓ depends on Δ, too. To see this, let $\delta = e^{-r\Delta}$. In order to obtain Brownian motion in the limit, define p and q in terms of the random walk representation of Brownian motion: $p = \frac{1}{2}[1-(x/\eta)\sqrt{\Delta}]$, and $q = \frac{1}{2}[1-(y/\eta)\sqrt{\Delta}]$, where $x > y$ are the drifts of the Brownian motion that players observe and $\eta > 0$ is the volatility parameter.

As $\Delta \to 0$, the efficiency loss from punishing bad news can be written in terms of Δ as follows, for some constant $C > 0$:

$$\frac{1-\delta}{1-\delta^T}\frac{g}{\ell-1} \approx \frac{r\Delta}{r\Delta T}\frac{g}{C\Delta^{1/2}} = \frac{1}{T}\frac{g}{C\Delta^{1/2}}.$$

It might seem that, as long as $T \to \infty$ faster than $\Delta^{-1/2}$, the efficiency loss from punishing bad news can be made arbitrarily small, but this will violate the feasibility constraint that $\alpha \leq 1$. Indeed, the feasibility constraint can be rewritten as

$$\frac{1-\delta}{\delta}\frac{g}{\delta^T p^T(\ell-1)} \leq \alpha v \leq v \approx u - \frac{1}{T}\frac{g}{C\Delta^{1/2}} \quad \Leftrightarrow \quad u \geq \frac{g}{C\Delta^{1/2}}\left[\frac{1}{T} + \frac{r\Delta}{e^{-r\Delta T}p^T}\right].$$

But clearly $p^T \to 0$ much faster than Δ if $T \to \infty$ faster than $\Delta^{-1/2}$, so the second term on the right-hand side inequality above is violated as $T \to \infty$. This finally implies that the argument due to Abreu et al. (1990) fails with frequent actions and any fixed discount rate $r > 0$.

2.3 Kandori and Matsushima (1998) Doesn't Go All the Way

There is another way of managing information that "partially survives" the frequent action limit as long as players are sufficiently patient. This is in stark contrast with the previous subsection, where nothing survived regardless of how patient players were. It is related to an example in Kandori and Matsushima (1998). As $\Delta \to 0$, let $T = c/\Delta$, for some constant c to be specified later. Count the number of γ signals and β signals over a T-period block, and let the equilibrium strategies be as follows. Players plan mutual cooperation over the entire block. At the end of the block, if the number of β-signals is greater than pT, then mutual punishment is triggered with probability α. Otherwise, there is no punishment, the signal realizations of the block are forgotten, and the next block begins afresh.

The (symmetric) lifetime average payoff to a player is now

$$v = u - \frac{\delta^T}{1-\delta^T}\alpha v x,$$

where

$$x = \sum_{t=0}^{\lfloor (1-p)T \rfloor} \binom{T}{t} p^{T-t}(1-p)^t$$

and $\lfloor z \rfloor$ is the greatest integer less than or equal to z.

Let $n^* = \lfloor (1-p)T \rfloor$. Discouraging a single deviation requires

$$(1-\delta)g \leq \delta^T \alpha v (q-p)\binom{T-1}{n^*} p^{T-1-n^*}(1-p)^{n^*}.$$

Indeed, punishment probability after one deviation equals

$$q \sum_{n=0}^{\lfloor (1-p)T \rfloor} \binom{T-1}{n} p^{T-1-n}(1-p)^n + (1-q) \sum_{n=0}^{\lfloor (1-p)T \rfloor -1} \binom{T-1}{n} p^{T-1-n}(1-p)^n.$$

With no deviation, just replace q with p, so the probability change is

$$(q-p)\left[\sum_{n=0}^{\lfloor (1-p)T \rfloor} \binom{T-1}{n} p^{T-1-n}(1-p)^n - \sum_{n=0}^{\lfloor (1-p)T \rfloor -1} \binom{T-1}{n} p^{T-1-n}(1-p)^n\right]$$
$$= (q-p)\binom{T-1}{n^*} p^{T-1-n^*}(1-p)^{n^*}.$$

For τ deviations, $(1-\delta^\tau)g \leq \delta^T \alpha v \Delta\Pi_\tau$, where $\Delta\Pi_\tau = \Pi_\tau - \Pi$ and Π_τ is the probability of at least pT failures given τ deviations. Hoeffding (1956) shows that

$$\Pi_\tau \geq \sum_{n=0}^{\lfloor (1-p)T \rfloor} \binom{T}{n} \overline{p}^{T-n}(1-\overline{p})^n,$$

where $\overline{p} = [(q\tau + p(T-\tau)]/T$ is the arithmetic mean of each of the Bernoulli trial probabilities under consideration. By the Central Limit Theorem, if T is large enough, $\Delta\Pi_\tau$ is approximated arbitrarily closely by

$$\frac{1}{\sqrt{2\pi T\overline{p}(1-\overline{p})}} \int_0^{(\overline{p}-p)T} \exp\left\{-\frac{z^2}{2T\overline{p}(1-\overline{p})}\right\} dz$$
$$\approx \Delta\Phi\left(\frac{\tau\frac{x-y}{\eta}\sqrt{c}}{T}\right) = \Phi\left(\frac{\tau\frac{x-y}{\eta}\sqrt{c}}{T}\right) - \Phi(0),$$

where Φ is the standard normal cumulative distribution function. Without loss, by taking a subsequence if necessary, we may assume that $\tau/T \to \rho \in [0,1]$ as $T \to \infty$. The incentive constraint for τ deviations becomes

$$g \leq \alpha v \frac{\Phi(\rho\frac{x-y}{\eta}\sqrt{c}) - \Phi(0)}{(1-e^{-rc\rho})/e^{-rc}}.$$

If rc is sufficiently small, the inequality becomes approximately

$$g \leq \alpha v \frac{\Phi(\rho\frac{x-y}{\eta}\sqrt{c}) - \Phi(0)}{\rho rc}.$$

It is easy to see that the right-hand side of this inequality is minimized with respect to ρ at $\rho = 1$. Graphically, this is illustrated in Figure 1 below. The right-hand side is (proportional to) the gradient of a line from the mid-point of the standard normal

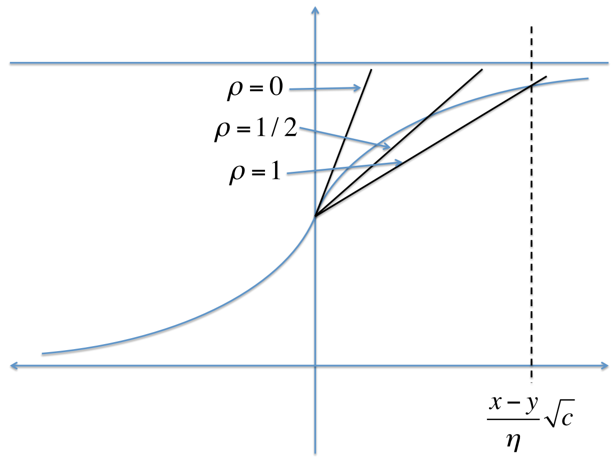

Figure 1: Cumulative normal has lowest slope between 0 and $\rho\frac{x-y}{\eta}\sqrt{c}$ at $\rho = 1$.

cdf and some point to its right depending on ρ. The lowest gradient is at the furthest possible point, at $\frac{x-y}{\eta}\sqrt{c}$, where $\rho = 1$.

Therefore, it is not the case, as in the previous subsection, that discouraging one deviation implies that all other deviations are discouraged. In fact, the most tightly binding incentive constraint is that meant to discourage T deviations, corresponding to $\rho = 1$. Nevertheless, if T deviations are discouraged then it is clearly the case that τ deviations are also discouraged for all τ less than T, so satisfying the previous inequality guarantees incentive compatibility.

In trying to maximize v, players' symmetric payoffs, this constraint will bind. Substituting for αv then yields

$$v = u - \frac{g/2}{\Phi(\frac{x-y}{\eta}\sqrt{c}) - \Phi(0)},$$

where $\Pi_0 \to 1/2$ as $\Delta \to 0$. This is is feasible only if

$$u \geq \frac{g}{\Phi(\frac{x-y}{\eta}\sqrt{c}) - \Phi(0)}\left[\frac{1}{2} + \frac{rc}{e^{-rc}}\right].$$

Therefore, c must be large and rc small. Even assuming feasibility, and having taken the limit $\Delta \to 0$, consider the limit as $c \to \infty$, $r \to 0$, and $rc \to 0$. (E.g., let $r \to 0$ and $c = r^{-1/2}$.) If u is big enough relative to g, it is possible to obtain a symmetric equilibrium outcome that is strictly better than static Nash, but we still don't get a Folk Theorem as players become patient. From the previous two expressions,

$$v \to u - g \quad \text{as} \quad r \to 0.$$

This is reminiscent of Compte's (1998) comparison of rewards and punishments.

2.4 Intermediate Approach

We will now consider an approach to providing incentives that lies somewhere in between Abreu et al. (1990) and Kandori and Matsushima (1998). I will amend the previous subsection's construction as follows. Just like the previous subsection, as $\Delta \to 0$, let $T = c/\Delta$, for some constant c to be specified later. Count the number of γ signals and β signals over a T-period block, and let the equilibrium strategies be as follows. Players plan mutual cooperation over the entire block. At the end of the block, if the number of β-signals is greater than qT (instead of pT), then mutual punishment is triggered with probability α. Otherwise, there is no punishment, the signal realizations of the block are forgotten, and the next block begins afresh. The key difference here is that punishment takes place after more failures. However, as $\Delta \to 0$, $q - p \to 0$, so the difference between the number of failures required in this subsection versus the previous subsection vanishes slowly (at rate $\sqrt{\Delta}$).

I will now show that discouraging just one deviation is enough to provide incentives, unlike the previous subsection. Moreover, I will prove a Folk Theorem: as players become (unboundedly) patient, the efficient outcome becomes an equilibrium.

The lifetime average payoff to a player is now

$$v = u - \frac{\delta^T}{1-\delta^T}\alpha v \Pi_0, \qquad \text{where} \qquad \Pi_0 = \sum_{n=0}^{\lfloor (1-q)T \rfloor} \binom{T}{n} p^{T-n}(1-p)^n.$$

By similar calculations to those in the previous subsection,

$$\Pi_0 \approx \tfrac{1}{2} - \frac{1}{\sqrt{2\pi T/4}} \int_0^{(q-p)T} \exp\left\{ \frac{-z^2}{2T/4} \right\} dz$$

for large T. Consider a player's plan to deviate τ times. Let Π_τ be the probability of at least qT failures given τ deviations. Boland et al. (2004) shows that

$$\Pi_\tau \geq \sum_{n=0}^{\lfloor (1-q)T \rfloor} \binom{T}{n} \tilde{p}^{T-n}(1-\tilde{p})^n,$$

where $\tilde{p} = q^{\tau/T} p^{(T-\tau)/T}$ is the geometric mean of each of the Bernoulli trial probabilities under consideration. Just as with Π_0, for large T this is approximately

$$\tfrac{1}{2} - \frac{1}{\sqrt{2\pi T/4}} \int_0^{(q-\tilde{p})T} \exp\left\{ \frac{-z^2}{2T/4} \right\} dz$$

Therefore,

$$\Delta\Pi_\tau \geq \frac{1}{\sqrt{2\pi T/4}} \int_{(q-\tilde{p})T}^{(q-p)T} \exp\left\{\frac{-z^2}{2T/4}\right\} dz = \Phi(2(q-p)\sqrt{T}) - \Phi(2(q-\tilde{p})\sqrt{T})$$
$$\approx \Phi(\tfrac{x-y}{\eta}\sqrt{c}) - \Phi((1-\rho)\tfrac{x-y}{\eta}\sqrt{c}),$$

where the last approximating equation follows in two steps. First, assume without loss of generality that $\tau/T \to \rho \in [0,1]$. Next, notice that $\tilde{p} \approx \overline{p}$ as $\Delta t \to 0$. To see this, consider the Taylor series expansion of $\tilde{p}$ with respect to $\sqrt{\Delta}$ around $\sqrt{\Delta} = 0$. Simple calculations yield that $\tilde{p}$ equals $\overline{p}$ plus terms of order higher than $\sqrt{\Delta}$.

This time, it is easy to see that the right-hand side of the inequality above divided by ρ is minimized by $\rho = 0$. For graphical intuition, see Figure 2 below.

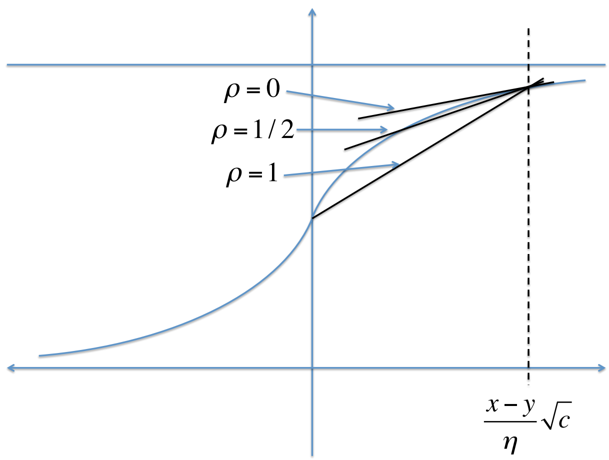

Figure 2: Cumulative normal has lowest slope from $(1-\rho)\frac{x-y}{\eta}\sqrt{c}$ to $\frac{x-y}{\eta}\sqrt{c}$ at $\rho = 0$.

In the same spirit as the previous subsection, if rc is small then the most binding incentive constraint is that associated with $\rho \to 0$, which yields

$$rcg \leq e^{-rc}\alpha v \varphi(\tfrac{x-y}{\eta}\sqrt{c})\tfrac{x-y}{\eta}\sqrt{c},$$

where φ is the standard normal density function. Substituting for αv yields

$$v = u - g\frac{1-\Phi(\frac{x-y}{\eta}\sqrt{c})}{\varphi(\frac{x-y}{\eta}\sqrt{c})\frac{x-y}{\eta}\sqrt{c}}.$$

Therefore, feasibility requires that

$$u \geq \frac{rcg}{\varphi(\frac{x-y}{\eta}\sqrt{c})\frac{x-y}{\eta}\sqrt{c}}\left[\frac{1}{e^{-rc}} + \frac{1-\Phi(\frac{x-y}{\eta}\sqrt{c})}{rc}\right]$$

Having already taken the limit $\Delta \to 0$, let $c \to \infty$ and choose $r(c) = \varphi(\frac{x-y}{\eta}\sqrt{c})\frac{x-y}{\eta}/c$, so that $r \to 0$ and $rc \to 0$ as $c \to \infty$. Using the fact that the normal distribution has a "linearly" exploding hazard rate, i.e., $(1 - \Phi(z))/(\varphi(z)z) \to 0$ as $z \to \infty$, we finally obtain a Folk Theorem:

$$v \to u \quad \text{as} \quad r \to 0.$$

To summarize: *information management yields asymptotic efficiency in repeated games with frequent actions as players become unboundedly patient.* The order of limits matters in this statement: first $\Delta \to 0$ and then $r \to 0$. Taking limits in the opposite order is the standard approach of discrete-time repeated games, whereas taking limits the way I just did is consistent with, e.g., Sannikov (2007); Sannikov and Skrzypacz (2007, 2010); Faingold and Sannikov (2011). All these papers argue that a Folk Theorem is not possible in continuous time when value-burning is necessary. However, the argument above suggests that this is indeed possible. The rest of the paper involves applying this idea to communication equilibria of general repeated games, where the information that is delayed is the mediator's recommendations.

As a final comment, it is not difficult to see that when signals converge to a Poission arrival process, the construction above shows that full efficiency is achievable in the "bad news" case, but not in the "good news" case—see also Abreu et al. (1990). This is because in the bad news case $\ell - 1$ does not vanish as $\Delta \to 0$, so delaying information allows for "reusing" incentive costs to deter more deviations than just one. On the other hand, in the good news case, $\ell - 1$ decreases at rate Δ. Changing T does not help to prevent losses due to incentive costs from becoming unbounded.

3 Model

Consider a repeated game with public monitoring in discrete time, with very similar notation as the previous section. The stage game consists of a finite set $I = \{1, \dots, n\}$ of players, a finite set A_i of actions for each player $i \in I$, where $A = \prod_i A_i$, utility profile $u : I \times A \to \mathbb{R}$, where $u_i(a)$ is the utility to player i from action profile a. Players have a common discount factor δ, where $\delta \in [0, 1)$. Given a sequence of action profiles $a^\infty = (a_1, a_2, \dots)$, the utility to player i is given by

$$U_i(a^\infty) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(a_t).$$

I now make two substantial restrictions on payoffs and the information structure.

Let $U = \mathrm{conv}\{u(a) \in \mathbb{R}^n : a \in A\}$ be the convex hull of stage-game payoff vectors. For every player i and correlated strategy μ, let

$$\underline{u}_i(\mu) = \max_{\beta_i : A_i \to A_i} \sum_{a \in A} \mu(a) u_i(\beta_i(a_i), a_{-i}),$$

and write $\underline{u}(\mu) = (\underline{u}_1(\mu), \ldots, \underline{u}_n(\mu))$. Let $\underline{u} = (\underline{u}_1, \ldots, \underline{u}_n)$ be the vector of correlated minmax values, i.e., $\underline{u}_i = \min_\mu \underline{u}_i(\mu)$ for all i. Let $\underline{U} = \{u \in U : u \geq \underline{u}\}$ be the set of *feasible, individually rational payoffs.* Let $\Delta^\circ(A)$ be the set of completely mixed correlated strategies, and for each coalition J, write

$$\underline{U}_J = \{u \in U : \exists \mu \in \Delta^\circ(A) \text{ s.t. } u_i \geq \underline{u}_i(\mu)\ \forall i \in J \text{ and } u_i \leq u_i(\mu)\ \forall i \notin J\} \text{ and } \underline{U}^* = \bigcap_{J \subset I} \underline{U}_J.$$

Finally, let $\underline{U}^* = \bigcap_{J \subset I} \underline{U}_J$. This construction of $\underline{U}^*$ follows closely (though not exactly) the one by Compte (1998); see specifically his Theorem 2 on page 609.

Assumption 1 (Payoffs)**.** I restrict u to resemble the Prisoners' Dilemma as follows:

(a) The set $\underline{U}$ has dimension n.

(b) The stage game has a strictly Pareto inefficient correlated equilibrium, with payoff profile denoted by u_0. Write $U_0 = \{u \in \underline{U} : u \geq u_0\}$.

(c) $\underline{U}^* \supset \mathrm{ri}\, U_0$, where "ri" stands for relative interior.

Assumption 1 makes the game similar to a Prisoners' Dilemma in some key aspects. This assumption is satisfied by a lot of important games in economics, such as the Prisoners' Dilemma, of course, and Cournot oligopoly. The point of making these assumptions is to simplify the derivation of a Folk Theorem in repeated games with frequent actions when transfers (of continuation values) across agents are ineffective. Otherwise, the Folk Theorem would apply directly from a combination of Tomala (2009) and Rahman and Obara (2010) with Sannikov and Skrzypacz (2010, Proposition 1). These assumptions can be relaxed without jeopardizing the main results of the paper, but no new strategic insights would emerge. The point of the paper is simply that whenever delaying information can substitute for transferring continuation values across players, value-burning becomes a useful tool for providing incentives, in contrast with the results of Sannikov and Skrzypacz (2010), where value-burning was deemed ineffective under Brownian information as a result of their restriction to public strategies. This restriction, which arguably amounts to no strategic management of information, can be problematic in games with frequent actions.

Assumption 1(a) says that in principle it is possible to punish players individually, which is standard, and 1(b) that there is room for everyone's improvement beyond static equilibrium. Assumption 1(c) is a technical shortcut: it simplifies the proof of the Folk Theorem, in line with (Compte, 1998, Theorem 2). I will explain this in more detail later, but it basically says that when one must give incentives to a coalition J through rewards, the necessary compensation associated with those rewards above and beyond payoffs from correlated equilibrium is arbitrarily small.

The plan now is to develop a discrete-time model of imperfect monitoring whose noise converges to Brownian motion as the time interval between interactions vanishes. To this end, we shall fix the simplest possible consistent expression of this idea, the (binomial) random walk representation of Brownian motion. Although this is the simplest approach, the results below easily apply to more general representations.

For every action profile $a \in A$, let $p(a) = \frac{1}{2}[1 - (x(a)/\eta)\sqrt{\Delta}]$ be the probability that a Bernoulli trial yields a failure.[2] As the time interval length $\Delta \to 0$, the binomial process defined by $p(a)$, where success translates into a jump up of length $\eta\sqrt{\Delta}$ and failure a jump down by the same amount, yields a Brownian motion X with law

$$dX = x(a)dt + \eta dW,$$

where W corresponds to the Wiener process.

The following strong detectability condition is necessary for the results of the paper. It comes from its version in the discrete-time model of Rahman (2011).

Definition 1. The vector p is said to exhibit *conic independence* if

$$p(a_i, \cdot) \notin \text{cone}\{p(b_i, \cdot) : b_i \neq a_i\} \text{ and } 1 - p(a_i, \cdot) \notin \text{cone}\{1 - p(b_i, \cdot) : b_i \neq a_i\} \quad \forall (i, a_i).$$

Conic independence is stronger than individual identifiability; it may be interpreted as saying that even signal-contingent deviations are statistically detectable. However, this condition is much weaker than pairwise identifiability, which requires that, for every profile of unilateral deviations, it is possible to statistically identify the culprit. Since we will be varying Δ below and p depends on Δ, it might be more convenient to derive a condition on x that guarantees conic independence.[3] Conic independence can easily be expressed in terms of the drift x as follows.

[2]Therefore, a success also has positive probability when Δ is small. This assumption means that every signal is possible, regardless of what players play. It is common in the literature, largely made for simplicity: no signal is ever off the equilibrium path.

[3]Actually, we only require conic independence for a sequence $\{\Delta_m > 0\}$ converging to zero.

Lemma 1. *There exists $\underline{\Delta} > 0$ such that the vector of probabilities p_Δ exhibits conic independence for all $\Delta \in (0, \underline{\Delta})$ if the vector of drifts x satisfies*

$$x(a_i, \cdot) \notin \text{conv}\{x(b_i, \cdot) : b_i \neq a_i\} + L_{\mathbf{1}} \quad \forall (i, a_i), \tag{2}$$

*where "*conv*" stands for convex hull and $L_{\mathbf{1}} = \{\alpha(1, \ldots, 1) : \alpha \in \mathbb{R}\}$.*

Proof. Let us argue the contrapositive. Suppose that for every $\underline{\Delta} > 0$ there exists $\Delta \in (0, \underline{\Delta})$ such that p_Δ fails to exhibit conic independence, i.e., there are positive linear combinations y_Δ^+ and y_Δ^- (for the up-jumps and down-jumps, respectively) such that, using the random walk representation of Brownian motion,

$$x(a) \;=\; \sum_{b_i} y_{i\Delta}^{\pm}(a_i, b_i) x(b_i, a_{-i}) \pm (1 - \sum_{b_i} y_{i\Delta}^{\pm}(a_i, b_i))/\sqrt{\Delta} \quad \forall (i, a).$$

Take a sequence $\{\Delta_m > 0\}$, strictly decreasing to 0, and consider the corresponding sequence $\{y_m^\pm = y_{\Delta_m}^\pm\}$. Clearly, $\sum_{b_i} y_{im}^\pm(a_i, b_i) \to 1$ as $m \to \infty$, since otherwise the equation above will fail. Since $y \geq 0$, it follows that each $y_{im}^\pm$ is a bounded sequence, hence has a convergent subsequence. Taking subsequences of subsequences if necessary, there is a subsequence such that all $y_{im}^\pm$ converge together to some limit $y_i^\pm$. Depending on the rate at which $\sum_{b_i} y_{im}^\pm(a_i, b_i) \to 1$ relative to $\Delta_m \to 0$, the term $(1 - \sum_{b_i} y_{im}^\pm(a_i, b_i))/\sqrt{\Delta_m}$ can converge to any real number, independent of a_{-i}. $\square$

Lemma 1 gives a condition on the drifts of the limiting Brownian process such that the "canonical" random walk representation described above exhibits conic independence sufficiently close to it. As a result, the Folk Theorem of Section 4 applies. However, it should be noted that the condition on drifts in Lemma 1 is sufficient but not necessary for the analysis below to apply. What is necessary is the implication in Lemma 2, for which conic independence is used. The advantage of conic independence is that it's a simple condition to describe. However, it depends on Δ, which is a disadvantage. The condition of Lemma 1 solves this problem.

For an example of this condition on drifts, see the Prisoners' Dilemma of Section 5, illustrated in Figure 3 below. There, (2) above is satisfied, therefore conic independence holds for all small Δ. For further discussion of (2), see Section 6.

We will restrict attention to signaling structures that satisfy the condition on drifts in Lemma 1 above. Although stronger than individual identifiability, this condition is still substantially weaker than pairwise identifiability.

Assumption 2 (Probabilities)**.** The vector x of drifts satisfies (2).

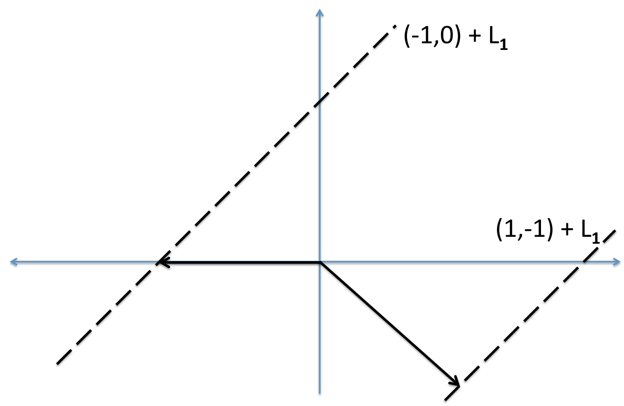

Figure 3: Conic independence holds for the Prisoners' Dilemma of Section 5.

4 Folk Theorem

I break up this section into punishments, rewards and the paper's main result.

4.1 Punishments

The next key lemma is an easy consequence of conic independence.

Lemma 2. *If p exhibits conic independence then $V(\mu) > 0$ for every completely mixed correlated strategy μ, where V is defined pointwise by*

$$\begin{aligned} V(\mu) = \max_{\gamma,\xi,\zeta\geq 0,\pi} \gamma \quad s.t. \quad \xi_i(a), \zeta_i(a) \leq 1 \quad \forall (i,a), \\ \gamma \left\|\Delta u_i(\mu)\right\| \leq \sum_{a_{-i}} (\xi_i(a) - \zeta_i(a))\Delta p(a, b_i)\mu(a) \quad \forall (i, a_i, b_i \neq a_i), \\ \sum_{a_{-i}} p(a_{-i}, b_i)\mu(a)\pi_i \leq \sum_{a_{-i}} \xi_i(a) p(a_{-i}, b_i)\mu(a) \quad \forall (i, a_i, b_i), \\ \sum_{a_{-i}} (1 - p(a_{-i}, b_i))\mu(a)\pi_i \leq \sum_{a_{-i}} \zeta_i(a)(1 - p(a_{-i}, b_i))\mu(a) \quad \forall (i, a_i, b_i), \\ \pi_i = \sum_a [\xi_i(a)p(a) + \zeta_i(a)(1 - p(a))]\mu(a) \quad \forall i, \end{aligned}$$

where $\left\|\Delta u_i(\mu)\right\| = \max_{(a_i,b_i)} \Delta u_i(\mu, a_i, b_i)$.

Proof. The proof is a simple application of duality. The dual of this problem is

$$V(\mu) = \max_{\eta,\sigma,y\geq 0,\kappa} \sum_{(i,a)} \eta_i^+(a) + \eta_i^-(a) \quad \text{s.t.} \quad \sum_i \Delta u_i(\mu,\sigma_i) \geq 1,$$

$$\eta_i^+(a) \geq \Delta p(a,\sigma_i)\mu(a) + \sum_{b_i} y_i^+(a_i,b_i)p(a_{-i},b_i)\mu(a) - \kappa_i p(a)\mu(a) \quad \forall (i,a),$$

$$\eta_i^-(a) \geq -\Delta p(a,\sigma_i)\mu(a) + \sum_{b_i} y_i^-(a_i,b_i)(1-p(a_{-i},b_i))\mu(a) - \kappa_i(1-p(a))\mu(a) \quad \forall (i,a),$$

$$\kappa_i = \sum_{(a,b_i)} y_i^+(a_i,b_i)p(a_{-i},b_i)\mu(a) + y_i^-(a_i,b_i)(1-p(a_{-i},b_i))\mu(a) \quad \forall i.$$

If $V(\mu) = 0$ then necessarily $\Delta p(a,\sigma_i)\mu(a) + \sum_{b_i} y_i^+(a_i,b_i)p(a_{-i},b_i)\mu(a) - \kappa_i p(a)\mu(a) = -\Delta p(a,\sigma_i)\mu(a) + \sum_{b_i} y_i^-(a_i,b_i)p(a_{-i},b_i)\mu(a) - \kappa_i p(a)\mu(a) = 0$ for all (i,a). Otherwise, if one of these terms is positive then one of the η's will also be positive, a contradiction. If one of these terms is negative then, since $\eta \geq 0$, it follows that $\sum_{(i,a)} \eta_i^+(a) + \eta_i(a)^-$ is strictly greater than the sum of the left-hand sides of the corresponding inequalities, which clearly add up to 0.

Now, conic independence implies that every deviation is detectable, and clearly, the first constraint implies that σ_i is a deviation for some i, so σ_i is detectable. But conic independence further implies that then there is no y_i such that the right-hand side of every dual inequality equals zero. □

The key condition needed below is that $V(\mu) > 0$. However, conic independence is reminiscent of individual identifiability, and in this sense intuitive. Call the first family of constraints above γ-*incentive constraints*, and the next two families of constraints *dynamic stability constraints*. The last family of constraints is just notation.

Next, I will use the fact that $V(\mu) > 0$ to obtain a Folk Theorem in the following sense. I will construct a "pseudo-equilibrium" of the repeated game as $\Delta \to 0$, which will be a communication equilibrium of an auxiliary T-period game with transfers, and then show that these transfers become self-generated as $r \to 0$.

Given a completely mixed correlated strategy μ, let μ^T be its T-period iid repetition. Suppose that $(\gamma, \xi, \zeta, \pi)$ solves the problem of Lemma 2 with value $V(\mu) > 0$, and, for each player i, let σ_i^* solve

$$\max_{\sigma_i} \frac{\|\Delta u_i(\mu)\| \, \sigma_i(\{b_i \neq a_i\})}{\sum_{(a,b_i)} \sigma_i(b_i|a_i)(\xi_i(a) - \zeta_i(a))\Delta p(a,b_i)\mu(a)}.$$

Hence, σ_i^* is a deviation that maximizes the incentive cost for player i, i.e., the utility gained from the deviation divided by the associated change in probability.

Define the score of the block by the following secret process. First, pick a player i at random with probability $1/n$. For any history (a^T, s^T), where s^T is a vector of up-jumps and down-jumps of the random walk, the mediator performs T independent Bernoulli trials. The tth trial has success probability $\xi_i(a_t)$ if s_t was a down-jump and $\zeta_i(a_t)$ if s_t was an up-jump. Following the intermediate approach of Section 2, punishment for player i ensues if the sum of successes in the T trials exceeds

$$\pi_i^* = \pi_i(\sigma_i^*) = \sum_{a \in A} [\xi_i(a) p(\sigma_i^*(a_i), a_{-i}) + \zeta_i(a)(1 - p(\sigma_i^*(a_i), a_{-i}))]\mu(a),$$

where $p(\sigma_i^*(a_i), a_{-i}) = \sum_{b_i} p(b_i, a_{-i})\sigma_i^*(b_i|a_i)$ is the probability of a down-jump if player i plays according to σ_i^* instead of obeying the mediator. Punishment to player i entails a transfer of some fixed amount w_i of what I will call "money" for now.

Assuming that the construction above is an equilibrium, the average lifetime utility to player i is given by

$$v_i = (1 - \delta^T) u_i(\mu) + \delta^T [(1 - \Pi_i/n) v_i + (\Pi_i/n)(v_i - w_i)],$$

where Π_i, the probability that i is punished if everyone abides by the candidate equilibrium and i was selected, equals

$$\Pi_i = \sum_{t=0}^{\lfloor (1-\pi_i^*)T \rfloor} \binom{T}{t} \pi_i^{T-t} (1 - \pi_i)^t$$

and $\pi_i = \sum_a [\xi_i(a) p(a) + \zeta_i(a)(1 - p(a))]\mu(a)$. Rearranging yields

$$v_i = u_i(\mu) - \frac{\delta^T}{1 - \delta^T} \Pi_i w_i / n.$$

4.2 Incentive Compatibility

Now I will argue that the contract defined above is incentive compatible for all large enough T, and that $w_i \to 0$ as $r \to 0$. To begin, assume that player i has not deviated yet, and considers deviating to σ_i^* for τ periods. Following the analysis of Section 2, it is clear that w_i feasibly exists to dissuade player i from any such deviation. On the other hand, if player i considers deviating to some other σ_i then the associated probability of failure in that period cannot decrease relative to σ_i^* by definition of σ_i^*, so the incentive constraint for σ_i is implied by the incentive constraint for σ_i^*. This logic applies to τ deviations, too.

As in Section 2, the above contract discourages non-history-contingent deviations, so it remains to show that history-contingent deviations are discouraged. I will use a modification of the following discrete-time argument. Suppose that the score of a block was constructed, instead of as described above, according to a variation on Abreu et al. (1990), where a player is punished if there are T failures (instead of $\pi_i^* T$).

It is easy to see that the dynamic stability condition of Lemma 2 guarantees incentive compatibility for each one-step deviation after every history, which implies overall incentive compatibility. Indeed, for any partial history h_i^t of recommendations and signal realizations for player i, and actual actions b_i^t, let $\Pr(F^{t-1}, h_i^t | b_i^t)$ be the joint probability of $t-1$ failures in the previous $t-1$ periods and the partial history h_i^t conditional on b_i^t, and let $\Pr(F^{T-t})$ denote the unconditional probability that from periods $t+1$ to T there will be only failures. Finally, for a given deviation in a single period, $\Delta \Pr(F_t)$ corresponds to the change in probability of contemporaneous failure from the deviation. With this notation, dynamic stability implies that $\Pr(F^{t-1}, h_i^t | b_i^t) \geq \pi_i^{t-1} \Pr(h_i^t | b_i^t)$ and $\Pr(F^{T-t}) = \pi_i^{T-t}$. Therefore, since the probability of failure after obedience satisfies $\Pr(F^{t-1}, h_i^t | a_i^t) = \pi_i^{t-1} \Pr(h_i^t | a_i^t)$, it follows that discouraging a one-step deviation after a history of obedience discourages all one-step deviations, since

$$\begin{aligned} \Pr(h_i^t | b_i^t)(1-\delta) \left\| \Delta u_i \right\| &\leq \delta^T \pi_i^{T-1} \Pr(h_i^t | b_i^t) \Delta \Pr(F_t) w_i \\ &\leq \delta^T \underbrace{\Pr(F^{t-1}, h_i^t | b_i^t)}_{\geq \pi_i^{t-1} \Pr(h_i^t | b_i^t)} \Delta \Pr(F_t) \underbrace{\Pr(F^{T \setminus t})}_{= \pi_i^{T-t}} w_i. \end{aligned}$$

As a result, incentive compatibility is achievable by the Abreu et al. (1990) version of the contract above in the strong sense that for every history, a one-step deviation can always be discouraged. Unfortunately, this argument does not quite hold with the actual contract proposed above for frequent actions. The problem is that after sufficiently many deviations and "bad" signals, it is possible that the probability of punishment does not really change by a large enough amount to discourage a one-step deviation. However, it can be shown that, in this case, even though a one-step deviation cannot be discouraged, the deviations required to get to that point can.

I will make this argument next in a few steps. But first, let us focus on deviations that are not history-contingent. Taking a subsequence if necessary, without loss ξ and ζ converge to $\overline{\xi}$ and $\overline{\zeta}$, say, as $\Delta \to 0$. Since the random walks driven by p converge to Brownian motion, it follows that π_i and π_i^* each converge at rate $\sqrt{\Delta}$ to

$$\overline{\pi}_i = \sum_{a \in A} \tfrac{1}{2} (\overline{\xi}_i(a) + \overline{\zeta}_i(a)) \mu(a).$$

Lemma 3. *Assumption 2 implies that* $\overline{\pi}_i \in (0,1)$.

Proof. Let $W(\mu,\Delta) = V(\mu,\Delta)\sqrt{\Delta}$, where $V(\mu,\Delta)$ is the value function from Lemma 2, and replace γ in the problem that defines it with $\lambda\sqrt{\Delta}$. By conic independence, it follows that $W(\mu,\Delta) > 0$ for all small $\Delta > 0$. It remains to show that $W(\mu,0) > 0$. This clearly yields $\overline{\pi}_i \in (0,1)$, since $\overline{\pi}_i \in \{0,1\}$ and the incentive compatibility constraint imply that $W(\mu;0) = 0$. The dual of the problem defining W is

$$\min_{\eta,y\geq 0} \sum_{(i,a)} \eta_i^+(a) + \eta_i^-(a) \quad \text{s.t.} \quad \|\Delta u_i\| \sum_{b_i \neq a_i} \sigma_i(b_i|a_i) \geq 1,$$

$$\eta_i^{\pm}(a) \geq \pm\Delta x_i(a,\sigma_i) + \tfrac{1}{2}\mu(a)[\sum_{b_i} y_i^{\pm}(a_i,b_i) - \overline{y}_i] \quad \forall(i,a),$$

$$\overline{y}_i = \tfrac{1}{2}\mu(a_{-i}) \sum_{b_i} y_i^+(a_i,b_i) + y_i^-(a_i,b_i) \quad \forall i.$$

If $W(\mu,0) = 0$ then, since $\eta \geq 0$, it follows that $\pm\Delta x_i(a,\sigma_i) + \frac{1}{2}\mu(a)[\sum_{b_i} y_i^{\pm}(a_i,b_i) - \overline{y}_i] \leq 0$ for all (i,a). But since (using the equation in the dual that pins down $\overline{y}$) the sum of all these left-hand sides equals zero, it follows that each of them equals zero individually. But this contradicts Assumption 2. □

Without loss, σ_i^* also converges to $\overline{\sigma}_i$, say. It is easy to see that, moreover, $\overline{\sigma}_i$ is a deviation. Indeed, by monotonicity σ_i^* maximizes

$$C_i(\sigma_i) = \frac{\|\Delta u_i(\mu)\|\, \sigma_i(\{b_i \neq a_i\})\sqrt{\Delta}}{\sum_a(\xi_i(a) - \zeta_i(a))\Delta p(a,\sigma_i)\mu(a)} = \frac{\|\Delta u_i(\mu)\|\, \sigma_i(\{b_i \neq a_i\})}{-\sum_a \frac{1}{2}(\xi_i(a) + \zeta_i(a))(\Delta x(a,\sigma_i)/\eta)\mu(a)}$$

This converges to $\|\Delta u_i(\mu)\|\, \sigma_i(\{b_i \neq a_i\})/-\sum_a \frac{1}{2}(\overline{\xi}_i(a) + \overline{\zeta}_i(a))(\Delta x(a,\sigma_i)/\eta)\mu(a)$, which is maximized by $\overline{\sigma}_i$. Let

$$\overline{z}_i = z(\overline{\sigma}_i) = \frac{-\sum_a \frac{1}{2}(\overline{\xi}_i(a) + \overline{\zeta}_i(a))(\Delta x(a,\overline{\sigma}_i)/\eta)\mu(a)}{\sqrt{\overline{\pi}_i(1-\overline{\pi}_i)}}.$$

Following the derivations of Section 2, the amount of punishment required to discourage any number of uncontingent deviations is bounded by

$$\overline{w}_i(r) = \frac{nrc}{\varphi(\overline{z}_i\sqrt{c})\overline{z}_i\sqrt{c}} \|\Delta u_i(\mu)\|.$$

The value to player i from this arrangement is given by

$$\overline{v}_i = u_i(\mu) - \frac{\|\Delta u_i(\mu)\|}{h(\overline{z}_i\sqrt{c})\overline{z}_i\sqrt{c}},$$

where $h(z)$ is the hazard rate of a standard normal random variable. Notice that $\overline{v}_i$ is independent of r, even though $\overline{w}_i(r)$ does depend on the interest rate—linearly, in fact. Of course, $\overline{v}_i$ does depend on c, which may be viewed as a function of r.

This is summarized in the last lemma of this subsection.

Lemma 4. *Fix $r > 0$, let $T = \lfloor c/\Delta \rfloor$, and define $C_i^* = \max_{\sigma_i} C_i(\sigma_i)$. By conic independence, $C_i^* \in (0, \infty)$. Moreover,*

$$\begin{aligned} v_i &\to \overline{v}_i = u_i(\mu) - \frac{\|\Delta u_i(\mu)\|}{h(\overline{z}_i\sqrt{c})\overline{z}_i\sqrt{c}} \qquad \textit{and} \\ w_i &\to \overline{w}_i(r) = \frac{nrc}{e^{-rc}\varphi(\overline{z}_i\sqrt{c})\overline{z}_i\sqrt{c}}\|\Delta u_i(\mu)\| \qquad \textit{as } \Delta \to 0. \end{aligned}$$

Lemma 4 follows immediately from the previous discussions. Given μ, r and c, the above values for v_i and w_i may or may not be self-generated, i.e., they may or may not describe a communication equilibrium of the repeated game. However, we will see that there exists $\overline{r} > 0$ such that for all $r < \overline{r}$, the payments above can be made self-generated by letting $c \to \infty$ at the appropriate rate. Moreover, as $r \to 0$, it will be seen that the set of communication equilibrium values converges to U_0, yielding a "Nash threats" Folk Theorem.

Now, let us consider history-contingent deviations. I will use dynamic stability below to show that the arrangement above discourages all history-contingent deviations for sufficiently large T. Fix any history-contingent deviation σ_i^T, and let $\Delta\Pi$ denote the change in the probability that player i is punished by the amount w_i above. Decompose $\Delta\Pi$ with respect to each partial history as follows. For each partial history $h_i^t = (b_i^{t-1}, s^{t-1}, a_i^t)$ of recommendations a_i^t, actions b_i^{t-1} and observations s^{t-1} for player i during the first t periods of a T-period block, let $\Pr(h_i^t)$ denote its probability and $\Delta\Pi(h_i^t)$ denote the change in probability of punishment for player i between deviating according to σ_i^T up to history h_i^t, and then obeying the mediator for the rest of the T-period block, and deviating up to history h_i^{t-1} (again, obeying henceforth). Thus, for example, $\Delta\Pi(h_i^1) = \Pi(h_i^1) - \Pi(h_i^0)$, where $\Pi(h_i^0)$ is the probability of punishment after deviating according to σ_{i1} in the first period and obeying henceforth, and $\Pi(h_i^1)$ is the probability of punishment given that player i deviated in period 1 according to σ_{i1} and deviated in period 2, after partial history h_i^2, according to $\sigma_{i2}(h_i^2)$. It will be useful to decompose $\Delta\Pi$ as follows.

$$\Delta\Pi = \Delta\Pi(h_i^0) + \sum_{h_i^t} \Pr(h_i^t)\Delta\Pi(h_i^t)$$

Since $\Delta\Pi(h_i^t) \geq 0$ assumes no more deviations after t, we may rewrite it as follows. Let F_1 denote the event that the first Bernoulli trial in the scoring rule described above leads to a failure, and let $\Delta\Pi^1(h_i^t)$ be the change in probability of punishment between deviating according to σ_i^T for the first t periods and deviating for the first $t-1$ periods given the partial history h_i^t, given that the first Bernoulli trial was a success. Similarly, $\Delta\Pi^0(h_i^t)$ corresponds to the same change in probability except that now it is assumed that the first Bernoulli trial was a failure. With this notation,

$$\begin{aligned}\Delta\Pi(h_i^t) &= \Pr(F_1, h_i^t)\Delta\Pi^0(h_i^t) + (\Pr(h_i^t) - \Pr(F_1, h_i^t))\Delta\Pi^1(h_i^t) \\ &= \Pr(h_i^t)\Delta\Pi^1(h_i^t) + \underbrace{\Pr(F_1, h_i^t)}_{\geq \Pr(h_i^t)\pi_i}\underbrace{[\Delta\Pi^0(h_i^t) - \Delta\Pi^1(h_i^t)]}_{=(*)}.\end{aligned}$$

By dynamic stability, it immediately follows that $\Pr(F_1, h_i^t) \geq \Pr(h_i^t)\pi_i$. Therefore, as long as $(*)$ above is not negative, we can apply a similar trick to the discrete-time case and collapse $\Delta\Pi(h_i^t)$ to a count of the number of deviations. Indeed, if $(*)$ is not negative then $\Delta\Pi(h_i^t) \geq \Pr(h_i^t)[\pi_i\Delta\Pi^0(h_i^t) + (1-\pi_i)\Delta\Pi^1(h_i^t)]$. However, $[\pi_i\Delta\Pi^0(h_i^t) + (1-\pi_i)\Delta\Pi^1(h_i^t)]$ equals the difference in probability of punishment assuming that player i obeyed the mediator in period 1.[4]

Given the history h_i^t, the probability change $\Delta\Pi^0(h_i^t)$ is just the difference in probability of n^* successes or fewer in the $T-1$ heterogeneous Bernoulli trials that correspond to periods 2 through T. Similarly, $\Delta\Pi^1(h_i^t)$ is the difference in probability of n^*-1 successes or fewer. Let $\Delta\Pi_1(h_i^t) = \pi_i\Delta\Pi^0(h_i^t) + (1-\pi_i)\Delta\Pi^1(h_i^t)$ denote the probability difference of n^* successes or fewer in the $T-1$ heterogeneous Bernoulli trials that correspond to periods 2 through T together with the Bernoulli trial corresponding to period 1 with success probability $1-\pi_i$.

Now we can iterate this process. Let F_2 denote the event that the second Bernoulli trial in the scoring rule described above leads to a failure. Just as with F_1, dynamic stability implies that $\Pr(F_2, h_i^t) \geq \Pr(h_i^t)\pi_i$. Define $\Delta\Pi_1^0(h_i^t)$ and $\Delta\Pi_1^1(h_i^t)$ similarly to $\Delta\Pi^0(h_i^t)$ and $\Delta\Pi^1(h_i^t)$, but with reference to the second period instead of the first. By the same argument as before,

$$\begin{aligned}\Delta\Pi(h_i^t) &\geq \Pr(F_2, h_i^t)\Delta\Pi_1^0(h_i^t) + (\Pr(h_i^t) - \Pr(F_2, h_i^t))\Delta\Pi_1^1(h_i^t) \\ &= \Pr(h_i^t)\Delta\Pi_1^1(h_i^t) + \Pr(F_2, h_i^t)[\Delta\Pi_1^0(h_i^t) - \Delta\Pi_1^1(h_i^t)] \\ &\geq \Pr(h_i^t)[\pi_i\Delta\Pi_1^0(h_i^t) + (1-\pi_i)\Delta\Pi_1^1(h_i^t)].\end{aligned}$$

[4]There is one minor caveat: player i may correlate subsequent deviations with his first-period deviation. This is captured already in the probability of h_i^t, whose dependence on σ_i^T is maintained.

Again, label $\Delta\Pi_2(h_i^t) = \pi_i \Delta\Pi_1^0(h_i^t) + (1-\pi_i)\Delta\Pi_1^1(h_i^t)$. Iterating this process repeatedly until period $t-1$ yields $\Delta\Pi_{t-1}(h_i^t)$, which consists of the difference in the probability of punishment between deviating in period t and not deviating in period t, but obeying the mediator during all other periods. By the analysis of Section 2, at history h_i^t, player i has no incentive to deviate in period t, since the utility gain $(1-\delta)\delta^{t-1}\Pr(h_i^t)\,\|\Delta u_i\|$ is outweighed by the loss $\delta^T \Delta\Pi_{t-1}(h_i^t) w_i$. If $(*)$ were non-negative for all h_i^t, we could add these gains and losses with respect to h_i^t, weighted by their probability, and obtain overall incentive compatibility with respect to σ_i^T.

However, it is not necessarily the case that $(*)$ is non-negative for all h_i^t, therefore the previous argument fails. To see this, write $n^* = (1-\pi_i^*)T$, and notice that $\Delta\Pi^0(h_i^t) - \Delta\Pi^1(h_i^t)$ equals the probability of exactly n^* successes after deviating according to σ_i^T during periods 2 through t minus the probability after deviating during periods 2 through $t-1$. If the history h_i^t is such that the probability of failure during the periods of deviation is higher than σ_i^* and if there are sufficiently many such periods, it is possible that n^* successes are too many, in other words, that the failure probabilities are so high that n^* successes are above the expected number of successes given history h_i^t. As a result, the probability of n^* successes is higher when the tth trial has higher success probability, which is of course the case when player i does not deviate in period t.

I will show that although this implies that some one-step deviations may be profitable after some histories, the behavior required to reach such histories is unprofitable as $T \to \infty$. I will bound incentives as follows. Whenever, the difference in probability of n^* successes above is not positive, I shall simply fix the difference in probability of punishment after history h_i^t to equal zero. Clearly, this probability difference is always non-negative, so such fix provides a lower bound on deviation costs.

To see the effect of such fix, for any history h_i^t such that $\Delta\Pi^0(h_i^t) - \Delta\Pi^1(h_i^t) \le 0$, set $\Delta\Pi(h_i^t) = 0$, as well as for all the successors of h_i^t. For any T-period history h_i^T subsequent to h_i^t, the difference between the probability of punishment given h_i^T and the probability of punishment given no deviation, $\Pi(h_i^T) - \Pi_0$, satisfies

$$\Pi(h_i^T) - \Pi_0 \ge \Pi(h_i^t) - \Pi_0 \ge \Phi(\overline{z}_i\sqrt{c}) - \Phi((1-\rho^*)\overline{z}_i\sqrt{c})$$

for large T, where $\rho^* = \pi_i^*/\pi_i^{**}$ and $\pi_i^{**} = \max_{\sigma_i} \max\{\Pr(F|\sigma_i, \text{up}), \Pr(F|\sigma_i, \text{down})\}$ is the largest possible failure probability from a one-period deviation σ_i given (ζ_i, ξ_i), regardless of whether the realized public signal was "up" or "down." To see this, notice first that Π_0 converges to $\Phi(\overline{z}_i\sqrt{c})$ as $T \to \infty$. Moreover, the probability $\Pi(h_i^t)$ of punishment after h_i^t must include enough previous deviations and signal realizations for which the conditional probability of punishment is strictly greater than π_i^* that $\Delta\Pi^0(h_i^t) - \Delta\Pi^1(h_i^t) \leq 0$. These deviations cannot induce a failure probability greater than π_i^{**}, since π_i^{**} is clearly an upper bound on the failure rate. Therefore, at a minimum, at least $\rho^* T$ deviations must have taken place for the average success probability of the Bernoulli trials consisting of $\rho^* T$ trials with failure rate π_i^{**} and $(1-\rho^*)T$ trials with failure rate π_i be equal less than or equal to $1-\pi_i^*$. Repeating the steps of Section 2, this leads to the estimate $\Phi((1-\rho^*)\overline{z}_i\sqrt{c})$ for the probability of punishment when T is large.

This argument bounds the cost of deviating conditional on history h_i^T. As regards the deviation gains, they are clearly bounded by $(1-\delta^T)\Delta u_i$, hence the following "tough" incentive constraint implies incentive compatibility given history h_i^T:

$$(1-\delta^T)\Delta u_i \leq \delta^T w_i[\Phi(\overline{z}_i\sqrt{c}) - \Phi((1-\rho^*)\overline{z}_i\sqrt{c})].$$

Finally, as $T \to \infty$, it is easy to see that $\rho^* \to 1$, therefore, replacing our previous expression for w_i into the incentive constraint above yields the limiting relationship

$$\Delta u_i \leq \|\Delta u_i(\mu)\| \frac{\Phi(\overline{z}_i\sqrt{c}) - \Phi(0)}{\varphi(\overline{z}_i\sqrt{c})\overline{z}_i\sqrt{c}}.$$

As $r \to 0$, $c \to \infty$ and $rc \to 0$, the inequality is satisfied for large enough T. Indeed, it is implied by incentive compatibility with respect to one-period deviations.

4.3 Rewards

The previous argument deals with sustaining equilibrium via individual punishments. However, sometimes it will be convenient to implement rewards. We take care of these next, to then derive a Folk Theorem. Next we will use a similar scoring approach as above to give rewards, i.e., in case some score exceeds a given amount, a player's utility is increased. This leads to only slightly different results. Let us begin with the corresponding lemma for rewards.

Lemma 5. *If p exhibits conic independence then $W(\mu) > 0$ for every completely* *mixed correlated strategy μ, where W is defined pointwise by*

$$W(\mu) = \max_{\gamma,\xi,\zeta\geq 0,\pi} \gamma \quad s.t. \quad \xi_i(a), \zeta_i(a) \leq 1 \quad \forall (i,a),$$

$$\gamma \Delta u_i(\mu, a_i, b_i) \leq \sum_{a_{-i}} (\zeta_i(a) - \xi_i(a)) \Delta p(a, b_i) \mu(a) \quad \forall (i, a_i, b_i),$$

$$\sum_{a_{-i}} p(a_{-i}, b_i) \mu(a) \pi_i \geq \sum_{a_{-i}} \xi_i(a) p(a_{-i}, b_i) \mu(a) \quad \forall (i, a_i, b_i),$$

$$\sum_{a_{-i}} (1 - p(a_{-i}, b_i)) \mu(a) \pi_i \geq \sum_{a_{-i}} \zeta_i(a)(1 - p(a_{-i}, b_i)) \mu(a) \quad \forall (i, a_i, b_i),$$

$$\pi_i = \sum_a [\xi_i(a) p(a) + \zeta_i(a)(1 - p(a))] \mu(a) \quad \forall (i,a).$$

The proof of this result is identical to that of Lemma 2, so omitted. For simplicity, although we could use the same scoring rule here as for punishments, we will adapt the idea of Abreu et al. (1990) to rewards instead. Given a solution to the linear program of Lemma 5 and a history (a^T, s^T) for the mediator, after having chosen player i with probability $1/n$, she will decide not to reward i with probability $\pi_i(a^T, s^T) = \prod_t [\xi_i(a)(1 - \mathbf{1}(s_t)) + \zeta_i(a) \mathbf{1}(s_t)]$, whose expectation is, of course, π_i^T. Assuming that this construction is an equilibrium, the average lifetime utility to player i is given by

$$v_i = (1 - \delta^T) u_i(\mu) + \delta^T [(1 - (1 - \pi_i^T)/n) v_i + (1 - \pi_i^T)/n(v_i + w_i)],$$

where now w_i is player i's reward in terms of continuation value. Rearranging yields

$$v_i = u_i(\mu) + \frac{\delta^T}{1 - \delta^T}(1 - \pi_i^T) w_i/n,$$

of course. To establish incentive compatibility, a slightly different value of w_i applies from the one for punishments. The reason is that for rewards, deterring one deviation no longer suffices to deter more of them. On the other hand, it is possible to deter all of them relatively easily, since the change in likelihood ratio from so many deviations is non-negligible as $\Delta \to 0$. Specifically, the benefit from deviating in the first period is given by $(1 - \delta) \Delta u_i(a_{i1}, b_{i1})$, whereas the cost is $\delta^T (w_i/n)(1 - \ell) y_{i1}$, where $\ell = \pi_i(a_{i1}, b_{i1})/\pi_i(a_{i1})$, $\pi_i(a_i, b_i) = \sum_{a_{-i}} [\xi_i(a) p(a_{-i}, b_i) + \zeta_i(a)(1 - p(a_{-i}, b_i))] \mu(a)$, $\pi_i(a_i) = \pi_i(a_i, a_i)$, and $y_{i1} = \mu(a_{i1}) - \pi_i^{T-1} \pi_i(a_{i1})$. By Lemma 5, $y_{i1} \geq \mu(a_{i1})(1 - \pi_i^T)$. Iterating this observation further, the utility gain from deviating every period equals $(1 - \delta^T) \Delta u_i(\mu, \sigma_i^*)$, for some deviation σ_i^*. On the other hand, the monetary loss equals $\delta^T (w_i/n)(1 - \ell(\sigma_i^*)^T)(1 - \pi_i^T)$. The associated inequality implies that all deviations are unprofitable, since $1 - \ell^T \leq (T/\tau)(1 - \ell^\tau)$ for $0 \leq \tau \leq T$ and $0 \leq \ell < 1$.

If $T = c/\Delta$ then, as $\Delta \to 0$, the incentive constraint for T deviations becomes simply this: $(1-e^{-rc})\Delta u_i(\mu, \sigma_i^*) \leq e^{-rc}(w_i/n)$. Similarly, the limit lifetime utility to player i as $\Delta \to 0$ becomes $v_i = u_i(\mu) + [e^{-rc}/(1-e^{-rc})]w_i/n = u_i(\mu) + \Delta u_i(\mu, \sigma_i^*) = u_i(\mu, \sigma_i^*)$. Define $\underline{w}_i(r) = n[(1-e^{-rc})/e^{-rc}]\Delta u_i(\mu, \sigma_i^*)$ and $\underline{v}_i = u_i(\mu, \sigma_i^*)$. Just as in the case of punishments, the continuation rewards $\underline{w}_i$ vanish as $r \to 0$, which permits self-generation for sufficiently small $r > 0$, as will be argued next.

As a final comment on rewards, notice that $\underline{v}_i$, which again does not depend on $r > 0$, corresponds to the scheme by Kandori and Matsushima (1998) described in Section 2, since the tightest incentive constraint is the one associated with always deviating.

4.4 Main Result

Now we can state the main result of the paper, quite simply that the "equilibrium threats" Folk Theorem holds as players' patience increases without bound. We could also prove the arguably stronger "minmax" Folk Theorem, but we would require stronger assumptions. This issue is discussed later.

Theorem 1 (Folk Theorem)**.** *Under Assumptions 1 and 2, the Folk Theorem holds in repeated games with frequent actions: for any payoff profile u in the interior of U_0, there is an interest rate $\underline{r}$ and a time step $\underline{\Delta}$ such that for all $r \leq \underline{r}$ and all $\Delta \leq \underline{\Delta}$ there exists a communication equilibrium whose payoff profile equals u.*

Proof. Given the previous arguments, all that remains is to establish self-generation for sufficiently small $r > 0$ and Δ. To this end, I will apply a simple extension of the algorithm by Fudenberg and Levine (1994). For each Δ, the proof of Fudenberg and Levine (1994) relied on establishing self-decomposability of smooth subsets. Let W be such a smooth subset of int U_0. We will show that, with punishments and rewards, W is self-generated, therefore W is a subset of the set of equilibria of the repeated game for some discount rate $\underline{r} > 0$. We must show that W is self-decomposable on tangent hyperplanes. For W in the interior, this is easy by using the correlated equilibrium of the stage game in Assumption 1. Now suppose that v is on the boundary of W.

Let $\overline{\delta} = \delta^T = e^{-r\Delta T} = e^{-rc}$, where $T = c/\Delta$. Pick any $\varepsilon, \varepsilon' > 0$ small. Let λ be the outward unit normal vector to W at v. If $\lambda_i \geq \varepsilon'$ for all i then we can self-generate v with unilateral punishments for every player, as in Figure 4 below, with a strictly Pareto-improving μ. We can now self-generate v with such μ. Indeed, let c be large enough that $\overline{v}_i = v_i$, and let $r < r(c) = \varphi(\overline{z}_i\sqrt{c})\overline{z}_i/c$. Choose $\underline{\Delta} > 0$ small enough that $\overline{w}_i(r)+\varepsilon$ satisfies the incentive constraint for punishments and $\underline{w}_i(r)+\varepsilon$ satisfies the incentive constraint for rewards for all $\Delta \leq \underline{\Delta}$. The welfare loss associated with these punishments clearly converges to zero as $r \to 0$, and for sufficiently small $r > 0$, continuation punishments clearly belong to W.

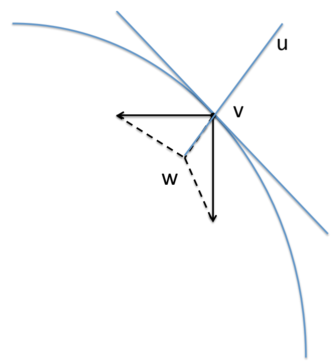

Figure 4: Local self-decomposability for positive welfare weights: $\lambda_i \geq \varepsilon'$ for all i.

Now, if all welfare weights are bounded above by ε' then we can just use the static equilibrium μ_0 with payoffs u_0 to self-generate v, as depicted in Figure 5 below.

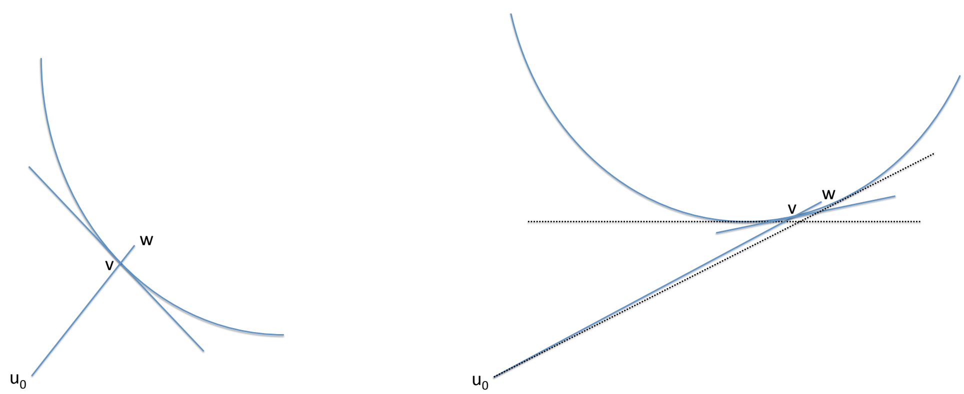

Figure 5: Local self-decomposability for weights of opposite signs, $|\lambda_i| \geq \varepsilon'$ for all i.

The left panel of Figure 5 shows how to self-generate v from u_0 given interest rate r when $|\lambda_i|$ is large for all i. The right panel shows how u_0 can be used to self-generate v even when λ_i is positive for some i, as long as it is not too large. Any given $\varepsilon' > 0$ sufficiently small implies a bound $r > 0$ such that self generation applies to $\lambda_i \leq \varepsilon'$.

If some players have positive weights and others negative weights under λ, with $|\lambda_i| \geq \varepsilon'$ for each i, by Assumption 1(c) we may choose μ with full support such that $\Delta u_i(\mu, \sigma_i^*) = v_i$ for all i with negative weight. Therefore, we may apply rewards consistently: as long as $v \gg u_0$, it is possible to self-generate v with some correlated strategy μ together with punishments and rewards. For intuition, see Figure 6.

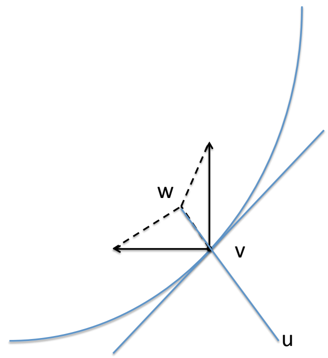

Figure 6: Local self-decomposability for weights of opposite signs, $|\lambda_i| \geq \varepsilon'$ for all i.

The last remaining case to self-generate is when $\lambda_i > \varepsilon'$ for some i and $|\lambda_i| \leq \varepsilon'$ for other i. Intuitively, Figure 7 shows how to self-generate such v. To show formally that this works, suppose that $\lambda_i = 1$ for some fixed i and $\lambda_j = 0$ for all $j \neq i$. With the equilibrium construction above, all that needs to be argued is that for sufficiently small $r > 0$ the payments required to provide incentives eventually lie in the interior of W. (This is not an issue for player i, of course.) Since W is smooth, so its boundary is twice differentiable, it suffices to show that z (see Figure 7) goes to 0 faster than $\sqrt{(1 - e^{-rc})/e^{-rc}} \approx \sqrt{rc}$ as $r \to 0$. But this is trivial: $\overline{w}_j(r) = Kr$ for some constant $K > 0$, which tends to zero linearly in r. This, of course, is much faster than $r^{1/2}$. The other cases are established either similarly or more easily, and hence omitted.

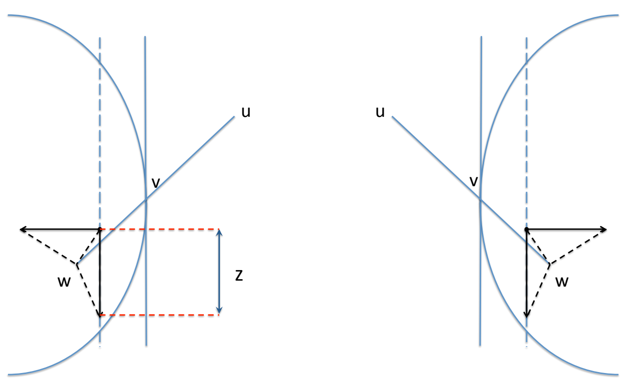

Figure 7: Local self-decomposability when $\phi_i > \varepsilon'$ and $|\lambda_j| \leq \varepsilon'$ for some i and j.

Finally, since W is compact, and we established self-generation for a neighborhood of each $v \in W$ and $r > 0$, i.e., an open cover of W,[5] it follows that there is a finite sub-cover, hence an $\underline{r} > 0$ for which every $v \in W$ is self-generated. $\square$

5 Examples

Below, I present two examples. The first is a Prisoners' Dilemma that requires the construction above for efficient outcomes to be attainable as players become patient. The second example shows how it is possible for firms to collude in repeated Cournot oligopoly, even when they have no capacity constraints.

5.1 Prisoners' Dilemma without Monotonicity

Consider a version of the Prisoners' Dilemma with imperfect public monitoring:

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0$

Payoffs

	C	D
C	1	-1
D	-1	0

Drifts

Volatility parameters are constant, say 1. To illustrate, if $\Delta = 1/9$ then $p(C, C) = \frac{1}{3}$, $p(C, D) = p(D, C) = \frac{2}{3}$ and $p(D, D) = \frac{1}{2}$. The key feature of this game is that the probability of good news is not increasing in the number of cooperators, unlike all other invocations of the Prisoners' Dilemma with imperfect monitoring.[6] Rahman and Obara (2010) show how such assumption is conceptually comparable to pairwise identifiability by arguing that monotonicity makes it possible to identify obedient agents. Also, in a sense public monitoring is especially challenging. Fong et al. (2007) require "sufficiently private monitoring" in addition to every player being able to detect everyone else's deviation, so their model does not apply here.

Rahman and Obara (2010) cannot provide a Folk Theorem for this example either because the signal structure above fails to identify obedient agents. Intuitively, there

[5] Although the statements above were made for each v rather than a neighborhood of each v, the additional step is standard and immediate.

[6] This monotonicity is comparable to the first-order stochastic dominance assumptions usually assumed in the literature on oligopoly with noisy prices, starting from Green and Porter (1984) and Abreu et al. (1986) all the way to Harrington and Skrzypacz (2010).

is a profile of deviations (e.g., for $\Delta = 1/9$, defecting if asked by the mediator to cooperate and cooperating with probability $\frac{1}{2}$ if asked to defect) that are profitable "close to" mutual cooperation such that for any action profile, the change in signal probabilities from a unilateral deviation is the same for every player: it could have been anyone. Hence, to discourage a deviation in this profile, every player must be punished on the path of play, leading to inefficiency.

Nevertheless, this inefficiency can be overcome with dynamic information management. By Theorem 1, conic independence is all we need (since Assumption 1 clearly holds), but this is easy to see. By Lemma 1, the vector x of drifts clearly satisfies (2), as shown graphically in Figure 1. Indeed, for $x(D, \cdot) = (-1, 0)$ and $x(C, \cdot) = (1, -1)$, there do not exist constants (α_1, α_2) such that $x(D, \cdot) = x(C, \cdot) + \alpha_1(1, 1)$ and $x(C, \cdot) = x(D, \cdot) + \alpha_2(1, 1)$. Therefore, the Folk Theorem holds in this Prisoners' Dilemma, using punishments and rewards as in the proof of Theorem 1.

Interestingly, the Folk Theorem holds even in some circumstances where (2) fails, although using a different incentive scheme from the one in Theorem 1. Indeed, suppose that the vector x has $x(D, D) = -3$ instead of 0. Now, it is clear that (2) fails. However, it is possible to identify obedient agents: if players are recommended to play an asymmetric profile, such as (C, D), one player (player 2) can only increase the drift of the Brownian motion, whereas the other player can only decrease it (player 1). Therefore, it is possible to identify obedient agents, and give the right incentives without value-burning with incentive schemes such as those proposed by Rahman and Obara (2010, Example 1) and Rahman (2012b). Moreover, these incentives lead to a Folk Theorem of the limiting game as $\Delta \to 0$, i.e., every feasible, (strictly) individually rational payoff profile is an equilibrium in the limit as $r \to 0$.

5.2 Collusion with Flexible Production

I will streamline the model of Sannikov and Skrzypacz (2007) quite a bit, but will otherwise stick to their notation and main assumptions (i.e., their stationary model). Their model consisted of a standard repeated Cournot oligopoly, where firms were able to produce any amount of a homogeneous good $q \in [0, \overline{q}]$, for some exogenous capacity constraint $\overline{q}$. Market prices are given by p_t, where

$$p_t = P(q_1 + q_2) + \varepsilon_t,$$

P is the instantaneous demand curve, $\varepsilon_t \sim N(0, \sigma^2/\Delta)$ is a random shock to prices, σ^2 is a variance parameter, and Δ is the duration of time periods between which production decisions are made. I assume, in line with Sannikov and Skrzypacz (2007), that P has bounded steepness: there exists $K > 0$ with $0 \geq P'(Q) \geq -K$ for all Q.

Let q_m denote the monopoly quantity associated with P, and consider the following mediated strategy. Fix $\mu > 0$ small, to be specified later. With probability $\frac{1}{2}$, a mediator selects one of the firms. If the selected firm is firm 1 then the mediator recommends to firm 2 to shut down, i.e., produce 0. With conditional probability $1 - \mu$, the mediator recommends to firm 1 to produce the monopoly output, q_m, and with conditional probability μ, the mediator recommends to firm 1 to produce at capacity, i.e., $\overline{q}$. If firm 2 was selected then the roles above are switched.

Rahman (2012b) shows that for this mediated strategy there exist continuation values that describe a public communication equilibrium as $\Delta \to 0$ without "burning value." Therefore, collusion can be mediated easily when there are capacity constraints, even with flexible production. The key to the argument there is that sometimes one firm shuts down while the other firm produces at capacity. This way, for any deviation profile, it is possible to identify someone who did not deviate. If the price rises stochastically relative to what the price should have been had nobody deviated, then it must have likely been that a firm which was meant to produce at capacity decided to decrease its output. Similarly, if the price drops then a firm scheduled to shut down must have likely decided to sell some output.

Without capacity constraints, however, this argument breaks down, because both firms could deviate by increasing their production, and then the identity of the deviator would not be inferred. Yet, using the construction of Section 4, there is a way to still sustain some amount of collusion as long as firms are sufficiently patient. Moreover, this amount of collusion increases unboundedly as $r \to 0$.

Consider the following mediated strategy. Given $\mu > 0$ small, with probability $\frac{1}{2}(1-\mu)$ the mediator tells firm 1 to shut down and firm 2 to produce the monopoly output. With probability $\frac{1}{2}\mu(1-\mu)$, the mediator tells firm 1 to shut down and firm 2 to produce q_μ, specified momentarily, and with probability $\frac{1}{2}\mu^2$ the mediator tells both players to play q_μ. For every asymmetric recommendation, all the roles above are reversed with the same probabilities. Letting q^* be each firm's best response correspondence, define q_μ as the fixed point that solves $q_\mu \in q^*((1-\mu)[0] + \mu[q_\mu])$.[7] It is

[7] The notation $[q]$ means Dirac measure at q. Here, I assume that such fixed point exists and be-

easy to see that $q_\mu < q_m$ and $q_\mu \to q_m$ as $\mu \to 0$.

Punishments take the following form. Suppose that firm 1 was asked to shut down (the case for firm 2 is symmetric). Firm 2's continuation values will now not depend on what happened previously. This is clearly incentive compatible because when a firm is asked to produce either q_μ or q_m, there is no better "static" response. Therefore, without loss we can construct the following punishment scheme that is only slightly different from that of Theorem 1. Fix a price cutoff $\overline{p} \in \mathbb{R}$.[8] If firm 2 was asked to produce q_μ then firm 1's punishment score will increase by 1 unit if prices are low, i.e., below $\overline{p}$, whereas if firm 2 was asked to produce q_m then firm 1's punishment score will increase by μ if prices are high, i.e., above $\overline{p}$. Punishments take a slightly different form to those of Theorem 1, but easily seen to apply to Theorem 1 because when anything but shutting down is recommended to a firm, the firm is willing to obey without any additional incentives. During a T-period block, a punishment score accumulates over time as just described. At the beginning of the T-period block, the mediator draws all recommendations and tells each firm how many times it will be asked to shut down, call it τ_i. This way, we may ignore dynamic consistency for not shutting down. At the end of the T-period block, the mediator chooses firm i with probability 1/2, and firm i is punished if the chosen firm's actual score ends up exceeding $\pi^*\tau_i$, where

$$\pi^* = \tfrac{1}{2}\mu(1-\mu)[\pi_{d^*\mu} + (1 - \pi_{d^*m})] < \tfrac{1}{2}\mu(1-\mu),$$

and π_{ij} is the probability of a low price (i.e., below $\overline{p}$) if firm 1 produces q_i units of output and firm 2 produces q_j units (and $q_0 = 0$, of course), and q_{d^*} maximizes

$$\frac{\Delta R(0,d)}{\frac{1}{2}\mu(1-\mu)[(\pi_{d\mu} - \pi_{0\mu}) + (\pi_{0m} - \pi_{dm})]},$$

where $\Delta R(0,d)$ stands for the change in (expected) revenue between producing q_d and q_0, and the denominator corresponds to the probabilistic change from producing q_0 to q_d.

I sketch an argument that this is an equilibrium below. Let us first argue that the scoring function above is consistent with Lemma 2 (modulo the minor caveat of the previous page), that is, it satisfies γ-incentive compatibility and dynamic stability for a positive weight γ given the recommendation q_0. Thus, γ-incentive compatibility

haves well. This holds, for instance, if the profit function is concave and continuously differentiable.

[8] This cutoff renders firms' relevant public information binary, in line with Theorem 1.

amounts to the requirement that if firm 1 produces $q_d > 0$ instead of q_0 then

$$0 < \gamma \Delta R(0,d) \;\leq\; \tfrac{1}{2}\mu(1-\mu)[(\pi_{d\mu} - \pi_{0\mu}) + (\pi_{0m} - \pi_{dm})].$$

To see that the right-hand side is positive, notice that, even though $\pi_{dj} > \pi_{0j}$ for all q_j, there exists $\overline{p}$ such that $\pi_{d\mu} - \pi_{0\mu} > \pi_{dm} - \pi_{0m}$, so $\gamma > 0$ exists to satisfy the above constraints. To see this, just pick $\overline{p}$ to the right of $P(q_\mu)$ such that the slope of the normal density only diminishes to the right of $\overline{p}$, i.e., $\overline{p} = P(q_\mu) + \sqrt{2\sigma^2}$. For intuition, see Figure 8 below.

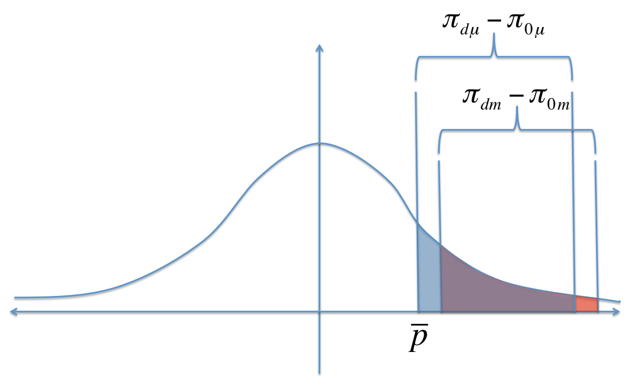

Figure 8: Differences in cutoff probability with respect to q_μ and q_m.

It is easy to see that the right hand side converges to a positive derivative times $q_d - q_0$ as $q_d \to q_0$, and converges to a positive number as $q_d \to \infty$. As for the left-hand side, by the bounded steepness assumption, it is clear that $\Delta R(0,d)$ converges to a derivative times times $q_d - q_0$ as $q_d \to q_0$, and converges to 0 as $q_d \to \infty$. Therefore, there exists $\gamma > 0$ to satisfy the inequality above for every q_d.

Next, dynamic stability requires that

$$\begin{aligned} \hat{\pi}[\tfrac{1}{2}\mu(1-\mu)(1-\pi_{d\mu}) + \tfrac{1}{2}(1-\mu)(1-\pi_{dm})] \;&\leq\; \tfrac{1}{2}\mu(1-\mu)(1-\pi_{dm}), \text{ and} \\ \hat{\pi}[\tfrac{1}{2}\mu(1-\mu)\pi_{d\mu} + \tfrac{1}{2}(1-\mu)\pi_{dm}] \;&\leq\; \tfrac{1}{2}\mu(1-\mu)\pi_{d\mu}. \end{aligned}$$

Each of these clearly holds whenever $\mu > 0$ is sufficiently small, which justifies the punishment scheme. For the construction of Section 4 to apply, we must still argue that a suitable rewards scheme delivers local self-decomposability when one firm has positive weight and another has negative weight. A similar approach to the one above is easily seen to work, so I omit this step, as it brings nothing really new.

Finally, for Theorem 1 to apply, it remains to establish the following final details. First, Assumption 1(a) and (b) immediate. To obtain (c), pick any coalition $J \subset I$ and let all players not in I produce so much that the best response to any player in J is to almost shut down. Now make the strategy of players not in J have full support but the high mean suggested above. The best-response profit of a player in J is close to zero, significantly less than the Nash equilibrium profit. Therefore, (c) holds. As for Assumption 2, it is not strictly necessary. What is necessary for the logic of Theorem 1 to work is the existence of a punishment function like the one above. This leads to the last detail. To apply the construction of Theorem 1, it is easy to see that the binary signal structure with an up-jump every time the price exceeds $\overline{p}$ and a down-jump every time it does not converges to a diffusion process for which the results above apply.

Therefore, as in Theorem 1, at the limit of $\Delta \to 0$, the collusive agreement described above converges to perfect collusion as $r \to 0$.

6 Discussion

In this section I comment on the model's limitations and possible extensions. I begin by discussing private monitoring, and then comment on the strengths and weaknesses of relying on a mediator.

6.1 Private Monitoring and Private Information

The model above assumed public monitoring, partly to make the results more difficult to obtain. With private monitoring it is even easier to obtain the model's results. This is because with private monitoring the mediator has more information at his disposal to manage. Let us assume, as usual, full support. Players observe a private signal every period that may be imperfectly correlated with others' signals. The information management protocol becomes slightly more sophisticated than before. Now, the mediator not only makes confidential (still non-binding) recommendations to players, but also asks players to confidentially report their observations. The mediator then makes recommendations contingent not only on what he told players, but also on what players told him. Notice that now the mediator is able to delay the arrival of two pieces of information: (i) others' recommendations, and (ii) others' reports. Therefore, as Rahman (2011) shows, the detectability requirements for the construction of this paper to work become substantially weaker. In fact, they boil down to precisely individual identifiability: every deviation being detectable suffices for the construction of this paper to work in games with private monitoring and full support.

Indeed, to see this, notice that the key step is to obtain a version of Lemma 2. With private monitoring, the linear programming problem of Lemma 2 must include incentive constraints that ensure players are willing to report their information honestly, and the dynamic stability constraints can integrate out others' signals. It is easy to see that the corresponding "cone condition" that suffices for such a private-monitoring version of Lemma 2 is the following:

$$\Pr(s_i, \cdot | a_i, \cdot) \notin \text{cone}\,\{\Pr(r_i, \cdot | b_i, \cdot) : (r_i, b_i) \neq (s_i, a_i)\} \quad \forall (i, a_i, s_i),$$

where $\Pr(s|a)$ is the probability of a profile $(s_1, \ldots, s_n)$ of up- and down-jumps for each player's private random walk given that everyone played the action profile $(a_1, \ldots, a_n)$.

It is also easy to see that this condition is equivalent to every deviation being detectable. Now, from the private-monitoring version of Lemma 2, the rest of the construction of this paper follows. If players have payoff-relevant private information, as is common in reputation models, a similar logic applies to the one just described, with a similar, although slightly more general version of conic independence. For details, see Rahman (2011).

6.2 Justifying the Mediator

There are several ways to justify studying mediated communication in dynamic games. Below I discuss what in my opinion are the more compelling justifications.

Firstly, at worst mediated communication provides an important benchmark for the study of repeated games in at least two ways. The first is in terms of communication. Lots of work on repeated games, along the lines of belief-free equilibria by Ely et al. (2005), has focused on the Folk Theorem without communication. Assuming that players have unlimited communication possibilities is no more extreme than that they have none whatsoever, and provides a benchmark for attainable outcomes. Second, and perhaps more interestingly, is in terms of private strategies. The use of private strategies has been highlighted as a way for players to correlate their play (Lehrer, 1991; Mailath et al., 2002), so communication equilibrium provides an upper bound for all Nash equilibria of a repeated game, including those with private strategies. This is important because there is no known characterization of Nash equilibria with private strategies. Communication equilibria can shed light on this question with a sometimes tight, often tractable upper bound. At the same time, games where this bound is not tight are ones with clear benefits to improved communication channels. In other words, if players had access to specific, perhaps private communication protocols, they might all be better off.

In some contexts, mediated communication is a realistic assumption. By Aumann's (1987) revelation principle, mediated communication can be considered shorthand for a not-necessarily-public randomization device, or simply unlimited scope for possibly secret communication (cheap talk), which is more than just plausible in many applications. Moreover, according to Forges (1990), any mediated communication protocol can be decentralized with (i) at least four players and (ii) plain conversation—these assumptions are also often reasonable in real-world applications.

Mediated communication is also arguably intrinsically interesting. Although some authors eschew any communication other than through actions, others embrace it.[9] Mediated communication specifically has not been studied in full generality, though,[10] yet it is interesting in and of itself. Characterizing the set of communication equilibrium payoffs in a repeated game, even as $\delta \to 1$, is still an open problem that is drawing increased interest in the game theory community.[11]

Mediation can also sometimes deliver simple answers. Mediation permits simple equilibria without arguably contrived constructions used to overcome having to communicate only through actions (e.g., Sekiguchi, 1997; Yamamoto, 2007; Sugaya, 2010; Fong et al., 2007), which may be unrealistic, or at least undesirable if a much simpler alternative is available.

Moreover, positive results are possible with a mediator in environments where without one they would be impossible. For instance, much of the literature has focused on

[9] See Kandori and Matsushima (1998), Compte (1998), Kandori (2003), Obara (2009) and Tomala (2009).

[10] Aoyagi (2005) looks at ε-perfect monitoring, Tomala (2009) focuses on public communication equilibria, and Cherry and Smith (2010) do not fix the game's information structure for their results.

[11] Renault and Tomala (2004) assume what could be called "$\delta = 1$," so are subject to the critiques of Rahman (2012a).

some form of mutual monitoring (e.g., Ben-Porath and Kahneman, 1996, 2003; Miyagawa et al., 2008, 2009; Fong et al., 2007; Sugaya, 2010), leading again to arguably contrived constructions for "monitoring monitors" in games with private monitoring. With a mediator, Rahman (2012a) shows that nobody needs to explicitly monitor a single monitor, and efficiency is possible even with one-sided monitoring.

Finally, there is also the issue of purifiability. Equilibria with mediated communication can in some circumstances deliver strict incentives, unlike belief-free and other equilibria. Therefore, they can be robust to Bhaskar's (2000) critique (see also Bhaskar et al., 2008), unlike the rest of the literature, such as Kandori (2003), Fong et al. (2007) or Sugaya (2010).

On the other hand, the use of a mediator is not without weaknesses. To assume that communication is completely costless, and that every kind of secret communication is possible, although a useful benchmark, may be unrealistic in some circumstances. Also, the fact that players can communicate freely might help them to coordinate their deviations, which raises the issue of coalition proofness.

6.3 Dispensing with the Mediator

In many circumstances, the mediator can be dispensed with. For instance, players, instead of taking recommendations, can simply report what they played, in line with Kandori (2003). However, this approach is subject to Bhaskar's (2000) critique. On the other hand, there are games where detectability requires perfectly correlated behavior by some players. In such games, it is unlikely that a mediator can be dispensed with. For a detailed discussion of these and other related issues, as well as an example of such a game, see Rahman (2012a).

Moreover, as was already mentioned, even if a mediator cannot be dispensed with, at least it may be possible to decentralize it with plain conversation, as argued by Forges (1990, 1986).

6.4 Conic Independence

Conic independence was shown above to be a sufficient condition on the probabilities p for the constructions of Theorem 1 to apply. Sometimes, other stronger conditions may apply and more sophisticated incentives schemes may be used, such as ones where punishments and rewards are correlated across players. Other times, the conic independence condition may not hold, so the construction above may not apply. However, conic independence is not strictly necessary for the above construction to hold. Rather, it is the conclusion of Lemma 2 that is required for the above construction. Even if the conclusion of Lemma 2 fails, it may be possible to provide incentives to players in other ways. For instance, as was illustrated in the discussion of the Prisoners' Dilemma at the end of Section 5.1, it is possible that p fails conic independence yet it can identify obedient agents, which allows for incentives to be correlated across players.

Lemma 1 provided a sufficient condition on the drifts of Brownian motions that imply conic independence for sufficiently small time intervals between interactions. The condition of Lemma 1 is arguably quite strong. Indeed, it is not difficult to see that it requires $|A_i| \leq |A_{-i}|$ for all i for it to hold generically. On the other hand, assumptions such as this one are not that uncommon in the literature.

6.5 Other Signal Structures

Relatedly, part of the reason why the condition of Lemma 1 is so strong is that the information structure used to reach the continuous time limit is arguably sparse. As shown by Fudenberg and Levine (2007, 2009), the kind of information structure used to reach the continuous time limit can have crucial, result-changing consequences. In this paper we decided to take the extreme case and make it as difficult as possible to obtain our desired result. Specifically, to this end we assumed a binary signal-driven random walk, where actions affect the drift but not the volatility of the limit diffusion process. This is important because it is well known that volatility is effectively observable, since it can be immediately inferred from a Brownian motion (Fudenberg and Levine, 2007, 2009, and references therein).

At this point, I should report once again that if the limiting signal structure is Poisson, rather than Brownian, it is easy to see that the impossibility results remain in the "good news" case (i.e., the arrival of a signal has lower probability after a given deviation) because there is no way of curtailing the associated implosion in likelihood ratio. Indeed, in this case the likelihood ratio $\ell - 1 \to 0$ at rate Δ, for which there is no function T that survives the basic argument of Section 2. On the other hand, in the "bad news" case, the moderate inefficiency reported in Abreu et al. (1990) can be easily eliminated, since there the likelihood ratio does not implode.

6.6 More Sophisticated Incentive Schemes

As a final comment, it is possible to provide more sophisticated incentives for fixed $r > 0$. For instance, the timing of up and down jumps can be added to the incentive scheme in order to make all the incentive constraints of Section 2 bind, not just the first-period constraint. This allows for more efficient incentive provision. Also, depending on the richness of the information structure, it may be possible to exploit additional information to improve incentives further. Finding the best incentive schemes for fixed interest rate is an interesting topic for further research.

7 Conclusion

This paper derives a "possibility" result for repeated games with frequent actions, showing how to sustain truly dynamic equilibria even with imperfect monitoring that converges to Brownian motion. The approach developed above relies on the use of a plausible generalization of private strategies, mediated strategies, which crucially allow for the delay of the arrival of endogenous strategic information. These mediated strategies may be thought of as dynamic information management institutions. The dynamic equilibria derived in this paper reconcile the literature on repeated games in discrete time and that in continuous time with a Folk Theorem that crucially depends on players' patience. The punishment schemes employed in this paper rely on carefully designed empirical likelihood tests of obedience that not only apply to discrete-time problems, but also to continuous-time problems. The results, despite having been derived for the case of public monitoring, apply much more generally, including to environments with private monitoring and private payoff-relevant information.

However, it is possible that continuous-time games may be intrinsically useful for pointed analysis of games with fixed interest rates. It would be interesting in the future to use the continuous-time limit to characterize the best forms of dynamic incentives for such fixed interest rates.

References

Abreu, D., P. Milgrom, and D. G. Pearce (1990): "Information and Timing in Repeated Partnerships," *Econometrica*, 59, 1713–33. 1, 2, 4, 5, 6, 7, 10, 12, 19, 25, 39

Abreu, D., D. Pearce, and E. Stacchetti (1986): "Optimal cartel equilibria with imperfect monitoring," *Journal of Economic Theory*, 39, 251–269. 29

Aoyagi, M. (2005): "Collusion Through Mediated Communication in Repeated Games with Imperfect Private Monitoring," *Economic Theory*, 25, 455–475. 36

Aumann, R. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1–18. 36

Ben-Porath, E. and M. Kahneman (1996): "Communication in Repeated Games with Private Monitoring," *Journal of Economic Theory*, 70, 281–297. 37

——— (2003): "Communication in repeated games with costly monitoring," *Games and Economic Behavior*, 44, 227–250. 37

Bhaskar, V. (2000): "The robustness of repeated game equilibria to incomplete payoff information," University of Essex Working Paper. 37

Bhaskar, V., G. J. Mailath, and S. Morris (2008): "Purification in the infinitely-repeated prisoners' dilemma," *Review of Economic Dynamics*, 11, 515–528. 37

Boland, P., H. Singh, and B. Cukic (2004): "The stochastic precedence ordering with applications in sampling and testing," *Journal of applied probability*, 41, 73–82. 10

Cherry, J. and L. Smith (2010): "Unattainable Payoffs for Repeated Games of Private Monitoring," Working paper. 36

Compte, O. (1998): "Communication in Repeated Games with Imperfect Private Monitoring," *Econometrica*, 66, 597–626. 3, 9, 13, 14, 36

Ely, J. C., J. Hörner, and W. Olszewski (2005): "Belief-Free Equilibria in Repeated Games," *Econometrica*, 73, 377–415. 35

FAINGOLD, E. AND Y. SANNIKOV (2011): "Reputation in Continuous-Time Games," *Econometrica*, 79, 773–876. 1, 4, 12

FONG, K., O. GOSSNER, J. HÖRNER, AND Y. SANNIKOV (2007): "Efficiency in the Repeated Prisoners' Dilemma with Private Monitoring," Mimeo. 29, 36, 37

FORGES, F. (1986): "An approach to communication equilibria," *Econometrica*, 54, 1375–1385. 1, 37

——— (1990): "Universal mechanisms," *Econometrica: Journal of the Econometric Society*, 1341–1364. 36, 37

FUDENBERG, D. AND D. LEVINE (1994): "Efficiency and observability with long-run and short-run players," *Journal of Economic Theory*, 62, 103–135. 3, 26

——— (2007): "Continuous time limits of repeated games with imperfect public monitoring," *Review of Economic Dynamics*, 10, 173–192. 1, 38

——— (2009): "Repeated Games with Frequent Signals," *The Quarterly Journal of Economics*, 124, 233–265. 1, 38

FUDENBERG, D., D. LEVINE, AND E. MASKIN (1994): "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62, 997–1039. 3

GREEN, E. AND R. PORTER (1984): "Noncooperative collusion under imperfect price information," *Econometrica: Journal of the Econometric Society*, 87–100. 29

HARRINGTON, JR., J. E. AND A. SKRZYPACZ (2010): "Private monitoring and communication in cartels: Explaining recent collusive practices," *American Economic Review*, forthcoming. 29

HOEFFDING, W. (1956): "On the distribution of the number of successes in independent trials," *The Annals of Mathematical Statistics*, 27, 713–721. 8

KANDORI, M. (2003): "Randomization, Communication, and Efficiency in Repeated Games with Imperfect Public Monitoring," *Econometrica*, 71, 345–353. 36, 37

KANDORI, M. AND H. MATSUSHIMA (1998): "Private Observation, Communication, and Collusion," *Econometrica*, 66, 627–652. 2, 5, 7, 10, 26, 36

LEHRER, E. (1991): "Internal correlation in repeated games," *International Journal of Game Theory*, 19, 431–456. 36

Mailath, G., S. Matthews, and T. Sekiguchi (2002): "Private strategies in finitely repeated games with imperfect public monitoring," *The BE Journal of Theoretical Economics*, 2, 2. 36

Miyagawa, E., Y. Miyahara, and T. Sekiguchi (2008): "The folk theorem for repeated games with observation costs," *Journal of Economic Theory*, 139, 192–221. 37

——— (2009): "Repeated Games with Costly Imperfect Monitoring," Kyoto University Discussion Paper No. 043. 37

Myerson, R. (1986): "Multistage games with communication," *Econometrica*, 54, 323–358. 1

Obara, I. (2009): "Folk theorem with communication," *Journal of Economic Theory*, 144, 120–134. 36

Rahman, D. (2011): "A Folk Theorem with Private Strategies," Mimeo. 14, 35

——— (2012a): "But Who Will Monitor the Monitor?" *American Economic Review*, forthcoming. 36, 37

——— (2012b): "Mediating Collusion with Flexible Production," Mimeo. 4, 30, 31

Rahman, D. and I. Obara (2010): "Mediated Partnerships," *Econometrica*, 78, 285–308. 3, 13, 29, 30

Renault, J. and T. Tomala (2004): "Communication equilibrium payoffs in repeated games with imperfect monitoring," *Games and Economic Behavior*, 49, 313–344. 36

Sannikov, Y. (2007): "Games with imperfectly observable actions in continuous time," *Econometrica*, 75, 1285–1329. 1, 4, 12

Sannikov, Y. and A. Skrzypacz (2007): "Impossibility of collusion under imperfect monitoring with flexible production," *The American Economic Review*, 97, 1794–1823. 1, 4, 12, 30, 31

——— (2010): "The role of information in repeated games with frequent actions," *Econometrica*, 78, 847–882. 1, 3, 4, 12, 13

Sekiguchi, T. (1997): "Efficiency in repeated prisoner's dilemma with private monitoring," *Journal of Economic Theory*, 76, 345–361. 36

Sugaya, T. (2010): "Folk Theorem in Repeated Games with Private Monitoring," Tech. rep., mimeo. 36, 37

Tomala, T. (2009): "Perfect Communication Equilibria in Repeated Games with Imperfect Monitoring," *Games and Economic Behavior*, 67, 682–694. 13, 36

Yamamoto, Y. (2007): "Efficiency results in N player games with imperfect private monitoring," *Journal of Economic Theory*, 135, 382–413. 36