

CARESS Working Paper #99-09
Repeated Games with Almost-Public Monitoring*

George J. Mailath
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104 USA
gmailath@econ.sas.upenn.edu

Stephen Morris
Cowles Foundation
Yale University
30 Hillhouse Avenue
New Haven, CT 06520 USA
stephen.morris@yale.edu

August 1999

Abstract

In repeated games with imperfect public monitoring, players can use public signals to coordinate their behavior perfectly, and thus support cooperative outcomes with the threat of punishments. But with even a small amount of private monitoring, players' private histories may lead them to have sufficiently different views of the world that such coordination on punishments is no longer possible (we describe a simple strategy profile that is a perfect public equilibrium of a repeated prisoner's dilemma with imperfect public monitoring, and yet is not an equilibrium for arbitrarily close games with private monitoring). If a perfect public equilibrium has players' behavior conditioned only on finite histories, then it induces an equilibrium in all close-by games with private monitoring. This implies a folk theorem for repeated games with almost-public almost-perfect monitoring.

*This is a substantial revision of "Repeated Games with Imperfect Private Monitoring: Notes on a Coordination Perspective." That paper was written while Mailath was visiting the Australian National University and Morris was visiting Northwestern University: we thank both institutions for their hospitality. We thank an Associate Editor and two referees for helpful comments. Mailath gratefully acknowledges financial support from NSF Grant #SBR-9810693 and Morris gratefully acknowledges financial support from the Alfred P. Sloan Foundation and NSF Grant #SBR-9896328.

Repeated Games with Almost-Public Monitoring

by George J. Mailath and Stephen Morris

1. Introduction

Perfect public equilibria of repeated games with imperfect *public* monitoring are well-understood.¹ When public signals provide information about past actions, punishments contingent on public signals provide dynamic incentives to choose actions that are not static best responses (see Green and Porter [16] and Abreu, Pearce, and Stacchetti [2]). Moreover, if the public signals satisfy an identifiability condition, a folk theorem holds: if the discount rate is sufficiently close to one, any individually rational payoff can be supported as the average payoff of an equilibrium of the repeated game (Fudenberg, Levine, and Maskin [15]). Perfect public equilibria of games with public monitoring have a recursive structure that greatly simplifies their analysis (and plays a central role in Abreu, Pearce, and Stacchetti [2] and Fudenberg, Levine, and Maskin [15]). In particular, any perfect public equilibrium can be described by an action profile for the current period and continuation values that are necessarily equilibrium values of the repeated game. However, for this recursive structure to hold, all players must be able to coordinate their behavior after any history (i.e., play an equilibrium after any history). If the relevant histories are public, then this coordination is clearly feasible.

Repeated games with *private* monitoring have proved less tractable. Since the relevant histories are typically private, equilibria do not have a simple recursive structure.² Consider the following apparently ideal setting for supporting non-static Nash behavior. There exist “punishment” strategies with the property that all players have a best response to punish if they know that others are punishing; and private signals provide extremely accurate information about past play, so that punishment strategies contingent on those signals provide the requisite dynamic incentives to support action profiles that are not static Nash. Even in

¹A strategy is public if it only depends on the public history, and a perfect public equilibrium is a profile of public strategies that induces a Nash equilibrium after every public history. Recent work (Kandori [18], Mailath, Matthews, and Sekiguchi [21], and Obara [25]) exploring private strategy profiles in games with imperfect public monitoring suggests that there is more to learn about games with public monitoring.

²Amarante [4] provides a very large state space recursive characterization of the equilibrium set of repeated games with private monitoring.

these circumstances, there is no guarantee that non-static Nash behavior can be supported in equilibrium. Even when one player is almost sure that another has deviated and would want to punish if he believed that others were punishing, he is not sure that others are almost sure that someone has deviated. With private signals, unlike public signals, there is not common knowledge of the histories that trigger punishments. If there is approximate common knowledge of the history of play, it should be possible to support non-static Nash behavior with the type of punishment strategies that we are familiar with from the perfect and imperfect public monitoring cases. But in what sense must there be approximate common knowledge of past play, and what kind of strategies will generate approximate common knowledge of past play?

We approach these questions as follows: Fix a repeated game with imperfect public monitoring and a strict pure strategy (perfect public) equilibrium of that game. Consider first the simplest perturbation of the game to allow private monitoring. Fix the set of signals in the public monitoring technology. Let each player observe a (perhaps different) signal from that set under the private monitoring technology. The private monitoring technology is said to be close to the public monitoring technology if the probability that all players observe the same signal, under the private technology, is close to the probability of that signal under the public technology. In this case, we say that there is *almost-public monitoring*. Now suppose players follow the original strategy profile, behaving as if the private signals they observe were in fact public. When is this an equilibrium of the perturbed game with private monitoring?

An important representation trick helps us answer this question. All perfect public equilibria of a repeated game with public monitoring can be represented in a recursive way by specifying a state space, a transition function mapping public signals and states into new states, and decision rules for the players, specifying behavior in each state (Abreu, Pearce, and Stacchetti [2]). We can use the same state space, transition function and decision rules to summarize behavior in the private monitoring game. Now each player will now have a *private state*, but we can use the transition function and decision rules to define a Markov process on vectors of private states. While this representation is sufficient for describing behavior under the given strategies and is invaluable in our analysis, it is not sufficient to check if the strategies are optimal. For this it is necessary to also know how each player's beliefs over the private states of other players evolve. A sufficient condition for a strict strategy profile to remain an equilibrium with private monitoring is that after every history each player assigns probability uniformly close to one to other players being in the same private state (Lemma 3). Thus, approximate common knowledge of histories throughout the game is sufficient

for equilibria with public monitoring to be robust to private monitoring. But for which strategy profiles will this approximate common knowledge condition be satisfied, for nearby private monitoring? A necessary condition is that the public strategy profile is *connected*: there is always a sequence of public signals that leads to the same final state independent of the initial state. However, we show by example that connectedness is not sufficient. One sufficient condition is that strategies only depend on a finite history of play (Theorem 2).

These results concern the robustness to private monitoring of perfect public equilibria of a fixed repeated game with imperfect public monitoring, with a given discount rate. We consider these results, and the examples illustrating them, to be the main results of the paper. The importance of finite histories is particularly striking, given that many of the standard strategies studied, while simple, do depend on infinite histories (e.g., trigger strategies). Our results convey a negative message for the recursive approach to analyzing repeated games with imperfect public monitoring. This approach is powerful precisely because it allows for the characterization of feasible equilibrium payoffs without undertaking the difficult task of exhibiting the strategy profiles supporting those payoffs. Our results suggest that if one is concerned about the robustness of perfect public equilibria to even the most benign form of private monitoring, fine details of those strategy profiles matter.

Our main results hold for a fixed discount rate. How close the private monitoring technology must be to the public monitoring technology depends, in general, on the discount rate. However, we also provide results that hold uniformly over discount rates sufficiently close to one. A connected finite public strategy profile is said to be *patiently strict* if it is a uniformly strict public equilibrium for all discount rates sufficiently close to one. In this case, approximate common knowledge of histories is enough to show that there exists $\varepsilon > 0$, such that for all discount rates sufficiently close to one, the strategy profile is an equilibrium of the private monitoring game, if the private monitoring technology is ε -close to the public monitoring technology (Theorem 3). This result is used to prove a pure-action folk theorem for repeated games with *almost-public*, *almost-perfect* monitoring (Theorem 5). Public monitoring is said to be *almost perfect* if the set of signals is the set of action profiles and, with probability close to one, the signal is the true action profile. There is almost-public almost-perfect monitoring if the private-monitoring technology is close to some almost-perfect public-monitoring technology. The folk theorem in this case follows from our earlier results, since it is possible to prove almost perfect monitoring folk theorems by constructing patiently strict finite history strategy profiles.

Thus far, our analysis has focused on a simple way of perturbing the public

monitoring technology where we held fixed the set of signals. We briefly also analyze private monitoring technologies with arbitrary signal sets. This analysis builds on the insights developed in the two period example of Section 2 and brings out more formally the role of approximate common knowledge assumptions.

There is now a growing and important literature looking at the role of mixed strategies in repeated games with private monitoring. Sekiguchi [28] showed that it is possible to achieve efficiency in a version of the repeated prisoner's dilemma, even if the private monitoring technology is independent, as long as it is sufficiently accurate; this result and technique have been significantly generalized (Bhaskar [8], Ely and Välimäki [13], Piccione [27], and Sekiguchi [29]). We discuss Sekiguchi's [28] construction (that is based on the grim trigger strategy profile) in Section 5 after the grim trigger example (Example 5). This literature has focused on the case of almost-perfect monitoring, but does not, as we do, require that signals be almost public. The technique of Ely and Välimäki [13] parallels ours in that they look at fully mixed strategy equilibria of a repeated game with public monitoring and characterize when they continue to be equilibria in a related class of private monitoring.³

An important negative result is proved by Compte [10], who considers trigger-strategy equilibria of the infinitely-repeated prisoner's dilemma. A trigger-strategy equilibrium has the property that if, in equilibrium, a player defects, he then defects thereafter with probability one. Compte shows that, for some class of full-support independent-signal private-monitoring technologies and discount rates close to one, the average expected payoff is close to the payoff from defection.⁴

We consider only the case of *full support* private monitoring with *no communication*. Thus, we exclude private monitoring environments where a subset of players perfectly observe the behavior of some player (Ben-Porath and Kahneman [6] and Ahn [3]); and we exclude the possibility of using cheap talk among the players to generate common belief of histories (Compte [11], Kandori and Matsushima [19], and Aoyagi [5]). In both approaches, the coordination problem that is the focus of our analysis can be solved, although of course new and interesting incentive problems arise. We also always analyze equilibria (not ε -equilibria) and assume strictly positive discounting. If players are allowed to take sub-optimal actions at some small set of histories, either because we are examining ε -equilibria of a discounted game or equilibria of a game with no discounting, then it is possible

³On mixed strategies in repeated games with private monitoring, see also Kandori [17] and Bhaskar and van Damme [9].

⁴Matsushima [23] also shows an anti-folk theorem. In particular, suppose that signals are independent and that players are restricted to pure strategies which depend on payoff-irrelevant histories *only* if that payoff-irrelevant history is correlated with other players' future play. These restrictions are enough to prevent coordination.

	D	C
D	$0, 0$	$x + 1, -x$
C	$-x, x + 1$	$1, 1$

Figure 1: A Prisoner's Dilemma

to prove stronger results (Fudenberg and Levine [14] and Lehrer [20]).

The paper is organized as follows. Section 2 discusses a two-period example that develops intuition for the results that follow, as well as providing one setting where it is possible to fully characterize pure strategy equilibria using approximate common knowledge properties of the private monitoring technology. Section 3 introduces repeated games with public monitoring and the same-signal-set class of nearby private monitoring technologies that we focus on. Sections 4 and 5 contain our main results on approximating arbitrary strict public equilibria, and examples illustrating the results. Section 6 presents the high discounting version of our finite history result and Section 7 applies this result to derive a folk theorem for repeated games with almost perfect monitoring. Section 8 extends these results to more general private monitoring technologies.

2. A Two Stage Example

We start by analyzing a two stage game. This example serves to introduce the issues of approximate common knowledge of past play, as well as illustrating two crucial assumptions. Section 2.2 shows that coordination is possible even if monitoring is not almost public, once we allow for mixed strategies;⁵ thus the restriction to pure strategies is a substantive one. Our main positive results concern strict pure strategy profiles of infinitely repeated finite-action stage games with private monitoring. Section 2.3 describes an example which indicates the necessity of the finite action assumption. Although we are interested in analyzing repeated games, it is useful to start with a case where the game varies across periods. We could illustrate the same points in a more complicated twice repeated game.

Two players are involved in a two stage game. In the first period, they play the prisoner's dilemma in Figure 1, where $x > 0$. In the second period, they play the coordination game in Figure 2, where $\alpha > \beta > 0$. There is no discounting.

⁵This analysis builds on Bhaskar and van Damme [9].

	G	B
G	α, α	$0, 0$
B	$0, 0$	β, β

Figure 2: A Coordination Game

2.1. Pure Strategy Equilibria

Consider first the case with perfect monitoring. It is possible to support (C, C) in the first period with pure strategies only if $\alpha - \beta \geq x$. In this case, there is a pure strategy equilibrium where each player chooses C in the first period, G in the second period if both chose C in the first period, and B in the second period if either player chose D in the first period.

Now consider the case with imperfect public monitoring. Let $\rho(y|a)$ be the probability that public signal $y \in Y$ is observed if action profile a is chosen; assume that $\rho(\cdot|a)$ has full support for each a and is symmetric (i.e., $\rho(\cdot|CD) = \rho(\cdot|DC)$). In any pure strategy equilibrium, either GG or BB must be played following each signal. Thus we are interested in symmetric strategy profiles of the form: each player chooses C in the first period, chooses G if the public signal ω is an element of Y^G and chooses B if a signal $y \in Y^B = Y \setminus Y^G$. Such a strategy profile is an equilibrium if and only if

$$\begin{aligned} \alpha [\rho(Y^G|CC) - \rho(Y^G|DC)] + \beta [\rho(Y^B|CC) - \rho(Y^B|DC)] &= \\ (\alpha - \beta) [\rho(Y^G|CC) - \rho(Y^G|DC)] &\geq x. \quad (1) \end{aligned}$$

Thus there exists a pure strategy equilibrium with first period cooperation if and only if there exists $Y^G \subseteq Y$ such that (1) holds. The equilibrium is strict if the inequality in (1) is strict.

Finally, consider the case with private monitoring. The simplest case of private monitoring (and the focus of Sections 3-7) is a private perturbation of the public monitoring technology (Y, ρ) , in which each player privately observes a signal from the space Y , and for each action profile a , there is a probability distribution $\pi(\cdot|a)$ on Y^2 . We say that π is ε -close to ρ if $|\pi(yy|a) - \rho(y|a)| < \varepsilon$ for all a and y . Consider the profile that specifies for player i , C in the first period, and after $y \in Y^G$, play G , and after $y \in Y^B$, play B . Suppose moreover that (1) holds strictly. Consider first the incentives in the second period, after observing $y \in Y^G$. For ε small, player i assigns a probability larger than $\beta/(\alpha + \beta)$ to the other player also having observed the same y (in fact, it is enough for the other player to have observed some signal in the set Y^G), and so player i finds

it optimal to play G . A similar argument applies for $y \in Y^B$. The last incentive constraint to check is for the first period. Cooperating yields an expected payoff

$$1 + \alpha\pi(Y^G \times Y^G | CC) + \beta\pi(Y^B \times Y^B | CC)$$

while deviating yields the payoff

$$x + 1 + \alpha\pi(Y^G \times Y^G | DC) + \beta\pi(Y^B \times Y^B | DC),$$

and so deviating is not profitable if

$$\alpha [\pi(Y^G \times Y^G | CC) - \pi(Y^G \times Y^G | DC)] + \beta [\pi(Y^B \times Y^B | CC) - \pi(Y^B \times Y^B | DC)] \geq x. \quad (2)$$

But since (1) holds strictly, for ε small, the above inequality is satisfied.

We now consider more general private monitoring technologies (the case of Section 8). Player i observes a signal ω_i from a finite set Ω_i concerning the first period action profile; write $\pi((\omega_1, \omega_2) | a)$ for the positive probability that signals (ω_1, ω_2) are observed when first period actions are $a \in \{C, D\}^2$; again assume that $\pi(\cdot | a)$ has full support for each a and is symmetric (i.e., $\pi((\omega_j, \omega_i) | (a_j, a_i)) = \pi((\omega_i, \omega_j) | (a_i, a_j))$).

When the players play according to a Nash equilibrium of this two stage game with private monitoring, the second stage can be viewed as a game of incomplete information, with player i having $T_i \equiv \{C, D\} \times \Omega_i$ as his type space. Moreover, there is a joint distribution over this type space induced by first period behavior and π .

In order to characterize the critical properties of π , we introduce belief operators (Monderer and Samet [24]). For any $E \subset \Omega_1 \times \Omega_2$, say that i p -believes E at $\omega = (\omega_1, \omega_2)$ if $\pi(E | \omega_i) \geq p$. The *belief operator* for player i identifies the signals at which i p -believes E , i.e.,

$$B_i^p(E; a) = \{\omega \in \Omega_1 \times \Omega_2 : \pi(E | \omega_i, a) \geq p\}.$$

The event E is p -evident (given a) if $E \subset B_i^p(E; a)$ for $i = 1, 2$. Note that if E is p -evident, then (since belief operators are monotonic, in the sense that $E' \subset E$ implies $B_i^p(E'; a) \subset B_i^p(E; a)$) $E \subset B_i^p(E; a) \subset B_i^p(B_j^p(E; a); a)$, i.e., i also assigns a probability of at least p to j assigning probability of at least p to E .

Lemma 1. *There exists a pure strategy equilibrium with cooperation in the first period if and only if each Ω_i can be partitioned into sets $\{\Omega_i^G, \Omega_i^B\}$ such that*

1. $\Omega_1^G \times \Omega_2^G$ is $\frac{\beta}{\alpha+\beta}$ -evident (given CC),

2. $\Omega_1^B \times \Omega_2^B$ is $\frac{\alpha}{\alpha+\beta}$ -evident (given CC), and
3. $\alpha [\pi (\Omega_1^G \times \Omega_2^G | CC) - \pi (\Omega_1^G \times \Omega_2^G | DC)]$
 $+ \beta [\pi (\Omega_1^B \times \Omega_2^B | CC) - \pi (\Omega_1^B \times \Omega_2^B | DC)] \geq x$.

This follows almost immediately from the definition of equilibrium. Suppose there is a pure strategy equilibrium (\hat{s}_1, \hat{s}_2) with

$$\hat{s}_i^1 = C \text{ and } \hat{s}_i^2(a_i, \omega_i) = \begin{cases} G, & \text{if } \omega_i \in \Omega_i^G, \\ B, & \text{if } \omega_i \in \Omega_i^B. \end{cases}$$

Properties [1] and [2] in the Lemma follow from the requirement that second period strategies constitute a (Bayesian) Nash equilibrium of the second period game following CC (this is a degenerate incomplete information game). Property [3] of the lemma is the counterpart of (2) and ensures that it is optimal to play C in the first period.

The Lemma provides an exact characterization of cooperative pure strategy equilibria. The characterization places restrictions on all conditional probabilities. However, if we just require a sufficient condition for cooperative pure strategy equilibrium, we can allow for a set of signals that we do not allocate to one action or another.

Corollary 1. *There exists a pure strategy equilibrium with cooperation in the first period if each Ω_i can be partitioned into sets $\{\tilde{\Omega}_i^G, \tilde{\Omega}_i^U, \tilde{\Omega}_i^B\}$ such that;*

1. $\tilde{\Omega}_1^G \times \tilde{\Omega}_2^G$ is $\frac{\beta}{\alpha+\beta}$ -evident (given CC)
2. $\tilde{\Omega}_1^B \times \tilde{\Omega}_2^B$ is $\frac{\alpha}{\alpha+\beta}$ -evident (given CC), and
3. $\alpha \pi (\tilde{\Omega}_1^G \times \tilde{\Omega}_2^G | CC) - \alpha \pi ((\tilde{\Omega}_1^G \cup \tilde{\Omega}_1^U) \times (\tilde{\Omega}_2^G \cup \tilde{\Omega}_2^U) | DC)$
 $+ \beta \pi (\tilde{\Omega}_1^B \times \tilde{\Omega}_2^B | CC) - \beta \pi ((\tilde{\Omega}_1^B \cup \tilde{\Omega}_1^U) \times (\tilde{\Omega}_2^B \cup \tilde{\Omega}_2^U) | DC) \geq x$.

To see why, suppose the condition of the corollary is satisfied. Carry out the following iterative procedure. Let $\hat{\Omega}_i^B(0) = \tilde{\Omega}_i^B$ and let $\hat{\Omega}_i^B(t+1) = \hat{\Omega}_i^B(t) \cup \{\omega_i \in \tilde{\Omega}_i^U : \pi (\hat{\Omega}_1^B(t) \times \hat{\Omega}_2^B(t) | \omega_i, CC) \geq \frac{\alpha}{\alpha+\beta}\}$. Observe that $\hat{\Omega}_i^B(t)$ is an increasing sequence of sets. Let $\Omega_i^B = \bigcup_{t \geq 1} \hat{\Omega}_i^B(t)$ and let Ω_i^G be the complement of Ω_i^B . By construction, $\{\Omega_i^G, \Omega_i^B\}$ partition Ω_i , $\Omega_1^G \times \Omega_2^G$ is $\frac{\beta}{\alpha+\beta}$ -evident (given

CC) and $\Omega_1^B \times \Omega_2^B$ is $\frac{\alpha}{\alpha+\beta}$ -evident (given CC). Also, observe that $\tilde{\Omega}_i^G \subseteq \Omega_i^G$ and $\tilde{\Omega}_i^B \subseteq \Omega_i^B$ imply that

$$\begin{aligned} & \alpha [\pi (\Omega_1^G \times \Omega_2^G | CC) - \pi (\Omega_1^G \times \Omega_2^G | DC)] + \beta [\pi (\Omega_1^B \times \Omega_2^B | CC) - \pi (\Omega_1^B \times \Omega_2^B | DC)] \\ & \geq \alpha \pi (\tilde{\Omega}_1^G \times \tilde{\Omega}_2^G | CC) - \alpha \pi ((\tilde{\Omega}_1^G \cup \tilde{\Omega}_1^U) \times (\tilde{\Omega}_2^G \cup \tilde{\Omega}_2^U) | DC) \\ & + \beta \pi (\tilde{\Omega}_1^B \times \tilde{\Omega}_2^B | CC) - \beta \pi ((\tilde{\Omega}_1^B \cup \tilde{\Omega}_1^U) \times (\tilde{\Omega}_2^B \cup \tilde{\Omega}_2^U) | DC) \geq x. \end{aligned}$$

So the premises of Lemma 1 are satisfied.

Our results on general private monitoring technologies in Section 8 build on the sufficient condition of the Corollary.

2.2. Mixed Strategy Equilibria

The pure strategy restriction is key to the above analysis. The role of mixed strategies in this context has been explored by Bhaskar and van Damme [9]. They consider a once repeated game with private monitoring. Our two stage game is essentially a stripped down version of their repeated game. They note that cooperation (i.e., an efficient but dominated action) is impossible in pure strategy equilibria with independent signals, but that cooperation is possible with correlated signals. Our Lemma 1 provides an exact description of how great the deviation from independence must be to allow cooperation.

Bhaskar and van Damme go on to show that mixed strategies allow cooperation in the first period, even with independent signals.⁶ We now illustrate this point in our example. In doing so, we describe how mixed strategies generate the requisite approximate common knowledge even when the private monitoring technology generates no correlation and thus no approximate common knowledge of the signals.

Consider an independent monitoring technology, with $\Omega_i = \{c, d\}$, in which each player observes his opponent's action correctly with independent probability $1 - \varepsilon$, and incorrectly with probability ε . We denote the signal of the opponent's

⁶The probability of cooperation is bounded away from one as the noise goes to zero. But the additional use of public sunspots allows cooperation with probability approaching one using mixed strategies, as the noise goes to zero.

	Cc	Cd	Dc	Dd
Cc	$\mu^2(1-\varepsilon)^2$	$\mu^2\varepsilon(1-\varepsilon)$	$\mu(1-\mu)\varepsilon(1-\varepsilon)$	$\mu(1-\mu)\varepsilon^2$
Cd	$\mu^2\varepsilon(1-\varepsilon)$	$\mu^2\varepsilon^2$	$\mu(1-\mu)(1-\varepsilon)^2$	$\mu(1-\mu)\varepsilon(1-\varepsilon)$
Dc	$\mu(1-\mu)\varepsilon(1-\varepsilon)$	$\mu(1-\mu)(1-\varepsilon)^2$	$(1-\mu)^2\varepsilon^2$	$(1-\mu)^2\varepsilon(1-\varepsilon)$
Dd	$\mu(1-\mu)\varepsilon^2$	$\mu(1-\mu)\varepsilon(1-\varepsilon)$	$(1-\mu)^2\varepsilon(1-\varepsilon)$	$(1-\mu)^2(1-\varepsilon)^2$

Figure 3: The distribution over $T_1 \times T_2$ generated by mixing with probability α on C .

	Cc	Cd	Dc	Dd
Cc	$(1-\varepsilon)^2$	$\varepsilon(1-\varepsilon)$	0	0
Cd	$\varepsilon(1-\varepsilon)$	ε^2	0	0
Dc	0	0	0	0
Dd	0	0	0	0

Figure 4: The distribution if $\alpha = 1$.

action by a lower case letter. Thus,

$$\pi((\omega_1, \omega_2) | (a_1, a_2)) = \begin{cases} (1-\varepsilon)^2, & \text{if } \omega_1 = a_2 \text{ and } \omega_2 = a_1, \\ \varepsilon(1-\varepsilon), & \text{if } \omega_1 = a_2 \text{ and } \omega_2 \neq a_1, \\ \varepsilon(1-\varepsilon), & \text{if } \omega_1 \neq a_2 \text{ and } \omega_2 = a_1, \\ \varepsilon^2, & \text{if } \omega_1 \neq a_2 \text{ and } \omega_2 \neq a_1. \end{cases}$$

Suppose that each player cooperates in the first period with probability μ and defects with probability $1-\mu$. The induced distribution over the type space $T_1 \times T_2$ is given in Figure 3. Setting $\mu = 1$ gives the pure strategy outcome of first period cooperation, with an induced distribution over types given in Figure 4, which is not consistent with punishment strategies in the second period (by the argument of Lemma 1). In particular, the event $\{(Cc, Cc)\}$ is $(1-\varepsilon)$ -evident, while the event $\{(Cd, Cd)\}$ is only ε -evident. The difficulty, of course, is that specifying B after a realization of c (as would be required by condition [2]) removes any incentive to choose C in the first stage. From the induced probability distribution on the type space $T_1 \times T_2$ we also see that for $\mu = 1$, the types are independent.

On the other hand, mixed strategies *generate* correlated types. As $\varepsilon \rightarrow 0$ (holding μ constant), the distribution over types tends to the distribution given in Figure 5. In other words, for *any* $p < 1$, $\{Cc\} \times \{Cc\}$ and $\{Cd, Dc, Dd\} \times \{Cd, Dc, Dd\}$ are both p -evident sets for ε small. Thus, for small ε and mixed first period strategies, there is no problem coordinating punishments in the second

	Cc	Cd	Dc	Dd
Cc	μ^2	0	0	0
Cd	0	0	$\mu(1-\mu)$	0
Dc	0	$\mu(1-\mu)$	0	0
Dd	0	0	0	$(1-\mu)^2$

Figure 5: The distribution if $\varepsilon = 0$.

period.

It is straightforward to construct symmetric mixed strategy equilibria using the above insight. Consider the following mixed strategies:

$$\sigma_i^1 [a_i^1] = \begin{cases} \mu, & \text{if } a_i^1 = C, \\ 1 - \mu, & \text{if } a_i^1 = D, \end{cases}$$

and

$$\sigma_i^2 (a_i^1, \omega_i) [a_i^2] = \begin{cases} 1, & \text{if } a_i^2 = G \text{ and } a_i^1 = C \text{ and } \omega_i = c, \\ 1, & \text{if } a_i^2 = B \text{ and either } a_i^1 = D \text{ or } \omega_i = d, \\ 0, & \text{otherwise.} \end{cases}$$

Second period optimality requires that the event $\{Cc\} \times \{Cc\}$ be $\frac{\beta}{\alpha+\beta}$ -evident while the event $\{Cd, Dc, Dd\} \times \{Cd, Dc, Dd\}$ is $\frac{\alpha}{\alpha+\beta}$. Thus is equivalent to requiring

$$\begin{aligned} \frac{\mu(1-\varepsilon)^2}{\mu(1-\varepsilon) + (1-\mu)\varepsilon} &\geq \frac{\beta}{\alpha+\beta}; \\ \frac{\mu\varepsilon^2 + (1-\mu)(1-\varepsilon)}{\mu\varepsilon + (1-\mu)(1-\varepsilon)} &\geq \frac{\alpha}{\alpha+\beta}; \\ \frac{\mu(1-\varepsilon)^2 + (1-\mu)\varepsilon}{\mu(1-\varepsilon) + (1-\mu)\varepsilon} &\geq \frac{\alpha}{\alpha+\beta}; \\ \text{and } \frac{\mu\varepsilon(1-\varepsilon) + (1-\mu)(1-\varepsilon)}{\mu\varepsilon + (1-\mu)(1-\varepsilon)} &\geq \frac{\alpha}{\alpha+\beta}. \end{aligned} \tag{3}$$

These inequalities will all hold as $\varepsilon \rightarrow 0$ as long as μ is bounded away from 1. For the mixed strategy to be optimal in the first period, we must have the payoff to cooperating,

$$\mu \{1 + (1-\varepsilon)^2 \cdot \alpha + \varepsilon^2 \cdot \beta\} + (1-\mu) \{-x + (1-\varepsilon) \cdot \beta\}$$

equal to the payoff from defecting,

$$\mu \{1 + x + (1 - \varepsilon) \cdot \beta\} + (1 - \mu) \{\beta\}.$$

Thus we must have

$$\mu = \mu(\varepsilon) = \frac{x + \varepsilon\beta}{(1 - \varepsilon)^2(\alpha - \beta) + 2\varepsilon^2\beta}.$$

As $\varepsilon \rightarrow 0$, $\mu(\varepsilon) \rightarrow \frac{x}{\alpha - \beta}$. Thus if $\alpha - \beta > x$, there will be an equilibrium with first period cooperation, for sufficiently small ε .

As a final comment on mixed strategies, note that an earlier period with private monitoring may allow the mixed strategy equilibrium to be purified.

2.3. An Anti-Folk Example

The pure strategy results in our two stage example relied on the discrete action space. Within the class of pure strategy equilibria, non-strict pure strategy equilibria only arise as non-generic cases. But strictness combined with the finite action assumption implies uniform strictness, i.e., for some strictly positive ε , the payoff to an equilibrium action is ε larger than the payoff of any other action. We use this property heavily in our repeated game arguments. We show here how replacing the second stage game with a continuum action coordination game destroys the possibility of first period cooperation under any private monitoring technology. This is even though first period cooperation is possible under imperfect public monitoring.

Let two individuals play the prisoner's dilemma in Figure 1 (with $x > 0$) in the first period. But in the second period, they play the following *convention game*. Player i chooses an action $a_i \in [0, 1]$. Payoff functions are

$$g_1(a_1, a_2) = a_2 - \gamma(a_1 - a_2)^2$$

and

$$g_2(a_1, a_2) = 1 - a_1 - \gamma(a_1 - a_2)^2,$$

where $\gamma > 0$. This is a special case of a class of games analyzed by Shin and Williamson [30]. There is a continuum of Nash equilibria: since each player's best response is to copy his opponent's action, (a_1, a_2) is a Nash equilibrium if and only if $a_1 = a_2$. Thus the *sum* of players' payoffs is 1 in any equilibrium. But player 1 prefers equilibria with high actions, while player 2 prefers equilibria with low actions. Note that for small γ , the symmetric efficient outcome has player 1

choosing action 0, player 2 choosing action 1, and both players obtaining a payoff of $1 - \gamma$.

With perfect monitoring, it is possible to support cooperation in the first period exactly if $\frac{1}{2} \geq x$. In this case, there is a pure strategy equilibrium where both players choose C in the first period. If both chose C in the first period, then both choose action $\frac{1}{2}$ in the second period. If player 1 chose action C and player 2 chose action D , then they both choose action 1 in period 2. If player 2 chose action C and player 1 chose action D , then they both choose action 0 in period 2.

One can similarly support first period cooperation with imperfect public monitoring. We will establish that first period cooperation is impossible with private monitoring. We first analyze what happens when the convention game is played once, but each player has access to some payoff irrelevant signal; specifically, each player observes a payoff irrelevant signal $\omega_i \in \Omega_i$, where each Ω_i is finite and (ω_1, ω_2) is drawn according to some full support distribution $\pi \in \Delta(\Omega_1 \times \Omega_2)$. For any $z \in [0, 1]$, this game has a *constant* equilibrium where each type of each player chooses action z . Shin and Williamson [30] showed that there are no other equilibria. The argument is elementary. Let \bar{a} be the largest action chosen by any type in an equilibrium. This is a best response only if every type of the other player chooses action \bar{a} . Thus play contingent on (full support) payoff irrelevant signals is inconsistent with equilibrium in this example. But this in turn implies that players' second period strategies in the two stage game with private monitoring are independent of their type, i.e., their first period action and observed private signal. Thus each player must defect in the first period.

If a game has multiple Nash equilibria, the finite folk theorem of Benoit and Krishna [7] shows that one may obtain any feasible and individually rational payoff in perfect equilibrium if the game is repeated often enough. Analogous imperfect public monitoring results also hold. However, it is a straightforward corollary of our earlier arguments (and backward induction) that if the convention game is repeated any finite number of times, with full support private monitoring, any Nash equilibrium of the repeated game must consist of a sequence of static Nash equilibria. A formal proof of this claim is presented in Mailath and Morris [22, Section 3].

3. Games with Almost-Public Monitoring

In this section, we begin our investigation of the extent to which games with public monitoring can be approximated by games with private monitoring. We first describe the game with public monitoring. The finite action set for player

$i \in \{1, \dots, N\}$ is A_i . The public signal is denoted y and is drawn from a finite set Y . The probability that the signal y occurs when the action profile $a \in A \equiv \prod_i A_i$ is chosen is denoted $\rho(y|a)$. We assume $\rho(y|a) > 0$ for all $y \in Y$ and all $a \in A$. We are thus restricting attention to *full support* public monitoring (this plays an important role, see Lemma 2 below). Since y is the only signal a player observes about opponents' play, it is common to assume that player i 's payoff after the realization (y, a) is given by $u_i^*(y, a_i)$. Stage game payoffs are then given by $u_i(a) \equiv \sum_y u_i^*(y, a_i) \rho(y|a)$.⁷ The infinitely repeated game with public monitoring is the infinite repetition of this stage game in which at the end of the period each player learns only the realized value of the signal y . Players do not receive any other information about the behavior of the other players. All players use the same discount factor, δ .

Following Abreu, Pearce, and Stacchetti [2] and Fudenberg, Levine, and Maskin [15], we restrict attention to perfect public equilibria of the game with public monitoring. A strategy for player i is *public* if, in every period t , it only depends on the public history $h^t \in Y^{t-1}$, and not on i 's private history. Henceforth, by the term *public profile*, we will always mean a strategy profile for the game with public monitoring that is itself public. A *perfect public equilibrium* is a profile of public strategies that, after observing any public history h^t , specifies a Nash equilibrium for the repeated game. Under imperfect full-support public monitoring, every public history arises with positive probability, and so every Nash equilibrium in public strategies is a perfect public equilibrium. Henceforth, equilibrium for the game with public monitoring means Nash equilibrium in public strategies (or, equivalently, perfect public equilibrium).

Any pure public strategy profile can be described as an automaton as follows: There is a set of states, W , an initial state, $w^1 \in W$, a transition function $\sigma : Y \times W \rightarrow W$, and a collection of decision rules, $d_i : W \rightarrow A_i$.⁸ In the first period, player i chooses action $a_i^1 = d_i(w^1)$. The vector of actions, a^1 , then generates a signal y^1 according to the distribution $\rho(\cdot|a^1)$. In the second period, player i chooses the action $a_i^2 = d_i(w^2)$, where $w^2 = \sigma(y^1, w^1)$, and so on. Since we can take W to be the set of all histories of the public signal, $\cup_{k \geq 0} Y^k$, W is at most countably infinite. A public profile is *finite* if W is a finite set.

If the profile is an equilibrium, each state has a continuation value, described by the mapping $\phi : W \rightarrow \mathfrak{R}^N$, so that the following is true (Abreu,

⁷While interpreting u_i as the expected value of u_i^* yields the most common interpretation of the game, the analysis that follows does not require it.

⁸Since we are restricting attention to pure strategies, the restriction to public strategies is without loss of generality: any pure strategy is realization equivalent to some pure public strategy (see Abreu, Pearce, and Stacchetti [2]).

Pearce, and Stacchetti [2]): Define a function $g : A \times W \rightarrow W$ by $g(a; w) \equiv (1 - \delta)u(a) + \delta \sum_y \phi(\sigma(y; w)) \rho(y|a)$. Then, for all $w \in W$, the action profile $(d_1(w), \dots, d_N(w)) \equiv d(w)$ is a pure strategy equilibrium of the static game with strategy spaces A_i and payoffs $g_i(\cdot; w)$ and, moreover, $\phi(w) = g(d(w), w)$. Conversely, if $(W, w^1, \sigma, d, \phi)$ describes an equilibrium of the static game with payoffs $g(\cdot; w)$ for all $w \in W$, then the induced pure strategy profile in the infinitely repeated game with public monitoring is an equilibrium.⁹

Here we consider private monitoring where the space of potential signals is also Y ; Section 8 extends the analysis to a broader class of private signals. Each player has action set A_i (as in the public monitoring game) and the set of private signals is Y^N . The underlying payoff structure is unchanged, being described by $u_i^*(y_i, a_i)$. For example, if y is aggregate output in the original public monitoring game, then y_i is i 's perception of output (and i never learns the true output). In a partnership game, y may be the division of an output (with output a stochastic function of actions), and private monitoring means that player i is not certain of the final payment to the other partners.

The probability that the vector of private signals $\mathbf{y} \equiv (y_1, \dots, y_N) \in Y^N$ is realized is denoted $\pi(\mathbf{y}|a)$. We say that the *private monitoring distribution* π is ε -close to the *public monitoring distribution* ρ if $|\pi(y, \dots, y|a) - \rho(y|a)| < \varepsilon$ for all y and a . If π is ε -close to ρ , then $\sum_y \pi(y, \dots, y|a) > 1 - \varepsilon|Y|$ for all a , where $|Y|$ is the cardinality of Y . We denote the vector $(1, \dots, 1)$ by $\mathbf{1}$, whose dimension will be obvious from context. Thus, $\pi(y, \dots, y|a)$ is written as $\pi(y\mathbf{1}|a)$. Let $\pi_i(\mathbf{y}_{-i}|a, y_i)$ denote the implied conditional probability of $\mathbf{y}_{-i} \in Y^{N-1}$. Moreover, for $w \in W$, $w\mathbf{1}$ denotes the vector (w, \dots, w) . Note that for all $\eta > 0$, there is an $\bar{\varepsilon} > 0$ such that if π is ε -close to ρ for $\varepsilon \in (0, \bar{\varepsilon})$, then

$$\left| \sum_{\mathbf{y}} u_i^*(y_i, a_i) \pi(\mathbf{y}|a) - \sum_{y_i} u_i^*(y_i, a_i) \rho(y_i|a) \right| < \eta.$$

An important implication of the assumption that the public monitoring is full support is that when a player observes a private signal y , then (for ε small) that player assigns high probability to all other players also observing the same signal, irrespective of the actions taken:

Lemma 2. *Fix $\eta > 0$. There exists $\bar{\varepsilon} > 0$ such that if $\varepsilon \in (0, \bar{\varepsilon})$ and π is ε -close to ρ , then for all $a \in A$,*

$$\pi_i(y\mathbf{1}|a, y) > 1 - \eta.$$

⁹We have introduced a distinction between W and the set of continuation payoffs for convenience. Any pure strategy equilibrium *payoff* can be supported by an equilibrium where $W \subset \mathbb{R}^I$ and $\phi(w) = w$ (again, see Abreu, Pearce, and Stacchetti [2]).

Proof. Fix the action profile $a \in A$. The probability that player i observes the private signal y is $\sum_{y_{-i}} \pi(y, \mathbf{y}_{-i}|a)$ and this probability is smaller than

$$\begin{aligned} (\rho(y|a) + \varepsilon) + \sum_{\mathbf{y}_{-i} \neq \mathbf{y}^1} \pi(y, \mathbf{y}_{-i}|a) &< \rho(y|a) + \varepsilon + \sum_{y \in Y} \sum_{\mathbf{y}_{-i} \neq \mathbf{y}^1} \pi(y, \mathbf{y}_{-i}|a) \\ &< \rho(y|a) + \varepsilon + \left(1 - \sum_{y \in Y} \pi(y \mathbf{1}|a)\right) \\ &< \rho(y|a) + \varepsilon(1 + |Y|). \end{aligned}$$

Thus, the probability that player i assigns to the other players observing the same signal y , $\pi_i(y \mathbf{1}|a, y)$, is at least as large as $\pi(y \mathbf{1}|a) \{\rho(y|a) + \varepsilon(1 + |Y|)\}^{-1} > (\rho(y|a) - \varepsilon) \{\rho(y|a) + \varepsilon(1 + |Y|)\}^{-1}$. Thus, by choosing

$$\varepsilon < \min_{a \in A} \frac{\eta \rho(y|a)}{2 + |Y| - \eta(1 + |Y|)},$$

we have $\pi_i(y \mathbf{1}|a, y) > 1 - \eta$ for all a . ■

Every public profile induces a private strategy profile (i.e., a profile for the game with private monitoring) in the obvious way:

Definition 1. *The public strategy profile described by the collection (W, w^1, σ, d) induces the private strategy profile $s \equiv (s_1, \dots, s_N)$ given by:*

$$\begin{aligned} s_i^1 &= d_i(w^1), \\ s_i^2(a_i^1, y_i^1) &= d_i(\sigma(y_i^1, w^1)) \equiv d_i(w_i^2), \end{aligned}$$

and defining states recursively by $w_i^{t+1} \equiv \sigma(y_i^t, w_i^t)$, for $h_i^t \equiv (a_i^1, y_i^1; a_i^2, y_i^2; \dots; a_i^{t-1}, y_i^{t-1}) \in (A \times Y)^{t-1}$,

$$s_i^t(h_i^t) = d_i(w_i^t).$$

That is, we are considering the public strategy translated to the private context. Note that these strategies ignore a player's actions, depending only on the realized signals. If W is finite, each player can be viewed as following a finite state automaton. Hopefully without confusion, we will abuse notation and write $w_i^t = \sigma(h_i^t; w^1) = \sigma(h_i^t)$, assuming that the initial state is taken as given. We describe w_i^t as player i 's *private state* in period t . It is important to note that while all players are in the same private state in the first period, since the signals

are private, after the first period, different players may be in different private states. The *private profile* is the translation to the game with private monitoring of the public profile of the game with public monitoring.

If player i believes that the other players are following a strategy that was induced by a public profile, then a sufficient statistic for h_i^t is player i 's private state and i 's beliefs over the other players' private states, i.e., (w_i^t, β_i^t) , where $\beta_i^t \in \Delta(W^{N-1})$. In principle, W may be quite large. For example, if the public strategy profile is nonstationary, it may be necessary to take W to be the set of all histories of the public signal, $\cup_{k \geq 0} Y^k$. On the other hand, the strategy profiles typically studied can be described with a significantly more parsimonious collection of states, often finite. When W is finite, the need to only keep track of each player's private state and that player's beliefs over the other players' private states is a tremendous simplification.

Example 1. Consider the prisoner's dilemma with payoffs, u , given by:

		Player 2	
		C	D
Player 1	C	2, 2	-1, 3
	D	3, -1	0, 0

The game with public monitoring has signals $y \in \{\underline{y}, \bar{y}\}$, where

$$\rho\{\bar{y}|a_1 a_2\} = \begin{cases} p, & \text{if } a_i = C, i = 1, 2, \\ q, & \text{otherwise,} \end{cases}$$

with $p > q$. The first profile we consider is described as follows: $W = \{\underline{w}, \bar{w}\}$, $w^1 = \bar{w}$, $d_i(\bar{w}) = C$, $d_i(\underline{w}) = D$, and

$$\sigma(yw) = \begin{cases} \bar{w}, & \text{if } y = \bar{y}, \\ \underline{w}, & \text{if } y = \underline{y}. \end{cases}$$

This is an equilibrium of the game with public monitoring if $1 > \delta > [3(p - q)]^{-1}$. A notable feature of this profile is that σ has finite (in fact, one period) memory. The actions of the players only depend upon the realization of the signal in the previous period. Thus, if player 1 (say) observes \bar{y} and assigns a probability sufficiently close to 1 that player 2 had also observed \bar{y} , then it seems reasonable that player 1 will find it optimal to play C.

Consider now a private monitoring distribution π , where $\pi(y_1 y_2 | CD) = \pi(y_1 y_2 | DC) = \pi(y_1 y_2 | DD) \equiv \pi_{y_1 y_2}^D$ and $\pi(y_1 y_2 | CC) \equiv \pi_{y_1 y_2}^C$. Note that the private monitoring distribution is identical if at least one player chooses D. Suppose player 1

observes the private signal \bar{y} and assigns a probability β to her opponent having also observed \bar{y} . We can write the incentive constraint for player 1 to follow the profile's specification of C as follows:

$$\begin{aligned} & \beta \left\{ (1 - \delta) 2 + \delta \left[\pi_{\bar{y}\bar{y}}^C V_1^{\bar{y}\bar{y}} + \pi_{\bar{y}y}^C V_1^{\bar{y}y} + \pi_{y\bar{y}}^C V_1^{y\bar{y}} + \pi_{yy}^C V_1^{yy} \right] \right\} \\ & + (1 - \beta) \left\{ (1 - \delta) (-1) + \delta \left[\pi_{\bar{y}\bar{y}}^D V_1^{\bar{y}\bar{y}} + \pi_{\bar{y}y}^D V_1^{\bar{y}y} + \pi_{y\bar{y}}^D V_1^{y\bar{y}} + \pi_{yy}^D V_1^{yy} \right] \right\} \\ & \geq \beta \left\{ (1 - \delta) 3 + \delta \left[\pi_{\bar{y}\bar{y}}^D V_1^{\bar{y}\bar{y}} + \pi_{\bar{y}y}^D V_1^{\bar{y}y} + \pi_{y\bar{y}}^D V_1^{y\bar{y}} + \pi_{yy}^D V_1^{yy} \right] \right\} \\ & + (1 - \beta) \left\{ \delta \left[\pi_{\bar{y}\bar{y}}^D V_1^{\bar{y}\bar{y}} + \pi_{\bar{y}y}^D V_1^{\bar{y}y} + \pi_{y\bar{y}}^D V_1^{y\bar{y}} + \pi_{yy}^D V_1^{yy} \right] \right\}, \end{aligned}$$

where $V_1^{y_1 y_2}$ is the continuation to player 1 under the profile when player 1 is in state (has just observed) y_1 and player 2 is in state y_2 . This expression simplifies to

$$\begin{aligned} & \beta \delta \left\{ (\pi_{\bar{y}\bar{y}}^C - \pi_{\bar{y}\bar{y}}^D) V_1^{\bar{y}\bar{y}} + (\pi_{\bar{y}y}^C - \pi_{\bar{y}y}^D) V_1^{\bar{y}y} + (\pi_{y\bar{y}}^C - \pi_{y\bar{y}}^D) V_1^{y\bar{y}} + (\pi_{yy}^C - \pi_{yy}^D) V_1^{yy} \right\} \\ & \geq (1 - \delta). \end{aligned} \quad (4)$$

The continuation values satisfy:

$$\begin{aligned} V_1^{\bar{y}\bar{y}} &= (1 - \delta) 2 + \delta \left\{ \pi_{\bar{y}\bar{y}}^C V_1^{\bar{y}\bar{y}} + \pi_{\bar{y}y}^C V_1^{\bar{y}y} + \pi_{y\bar{y}}^C V_1^{y\bar{y}} + \pi_{yy}^C V_1^{yy} \right\}, \\ V_1^{\bar{y}y} &= -(1 - \delta) + \delta \left\{ \pi_{\bar{y}\bar{y}}^D V_1^{\bar{y}\bar{y}} + \pi_{\bar{y}y}^D V_1^{\bar{y}y} + \pi_{y\bar{y}}^D V_1^{y\bar{y}} + \pi_{yy}^D V_1^{yy} \right\}, \\ V_1^{y\bar{y}} &= (1 - \delta) 3 + \delta \left\{ \pi_{\bar{y}\bar{y}}^D V_1^{\bar{y}\bar{y}} + \pi_{\bar{y}y}^D V_1^{\bar{y}y} + \pi_{y\bar{y}}^D V_1^{y\bar{y}} + \pi_{yy}^D V_1^{yy} \right\}, \text{ and} \\ V_1^{yy} &= \delta \left\{ \pi_{\bar{y}\bar{y}}^D V_1^{\bar{y}\bar{y}} + \pi_{\bar{y}y}^D V_1^{\bar{y}y} + \pi_{y\bar{y}}^D V_1^{y\bar{y}} + \pi_{yy}^D V_1^{yy} \right\}. \end{aligned}$$

Thus,

$$V_1^{\bar{y}\bar{y}} - V_1^{yy} = \frac{(1 - \delta) \left\{ 2 + \delta \left(\pi_{\bar{y}y}^D - \pi_{\bar{y}y}^C \right) + 3\delta \left(\pi_{\bar{y}\bar{y}}^C - \pi_{\bar{y}\bar{y}}^D \right) \right\}}{1 - \delta \left(\pi_{\bar{y}\bar{y}}^C - \pi_{\bar{y}\bar{y}}^D \right)}. \quad (5)$$

Suppose the private monitoring distribution has the particularly simple form $\pi_{\bar{y}\bar{y}}^C = p(1 - 2\varepsilon)$, $\pi_{\bar{y}y}^D = q(1 - 2\varepsilon)$, and $\pi_{\bar{y}y}^D = \pi_{y\bar{y}}^C = \pi_{y\bar{y}}^D = \pi_{yy}^C = \varepsilon$.¹⁰ Then (4) simplifies to

$$\beta \delta (p - q) (1 - 2\varepsilon) \left(V_1^{\bar{y}\bar{y}} - V_1^{yy} \right) \geq (1 - \delta),$$

¹⁰This is the case analyzed in Mailath and Morris [22, Section 4.1].

while (5) simplifies to

$$V_1^{\bar{y}\bar{y}} - V_1^{yy} = \frac{2(1-\delta)}{1-\delta(p-q)(1-2\varepsilon)},$$

so that the incentive constraint is satisfied if

$$(1+2\beta)(1-2\varepsilon) \geq \frac{1}{\delta(p-q)}. \quad (6)$$

Recall that the public profile is an equilibrium of the game with public monitoring if $3 \geq [\delta(p-q)]^{-1}$, and fix some $\underline{\delta} > [3(p-q)]^{-1}$. We claim that the same upper bound on ε suffices for the private profile to be an equilibrium of the game with private monitoring for any $\delta \geq \underline{\delta}$. For the incentive constraint describing behavior after \bar{y} , it suffices to have the inequality (6) hold for $\delta = \underline{\delta}$, which can be guaranteed by β close to 1 and ε close to 0. A similar calculation for the incentive constraint describing behavior after y yields the inequality

$$\delta(p-q)(1-2\varepsilon)(3-2\beta) \leq 1,$$

which can be guaranteed by appropriate bounds on β and ε , independent of δ . Finally, Lemma 2 guarantees that β can be made uniformly close to 1 by choosing ε small (independent of history).

Example 1 has the strong property that bound on the private monitoring can be chosen independent of δ . We return to this point in Section 6.

4. Approximating Arbitrary Strict Public Equilibria - The Fixed Discount Factor Case

We now formalize the idea that if players are *always* sufficiently confident that the other players are in the same private state as themselves, then the private profile induced by a strict public equilibrium is an equilibrium. In this section and the next, we focus on the case of fixed discount factors. More specifically, we ask: If a public profile is a strict equilibrium for some discount factor δ , then, for close by private monitoring distributions, is the same profile an equilibrium in the game with private monitoring for the *same* discount factor δ ?

Denote by $V_i(h_i^t)$ player i 's expected average discounted payoff under the private strategy profile after observing the private history h_i^t . Let $\beta_i(\cdot|h_i^t) \in \Delta(W^{N-1})$ denote i 's posterior over the other players' private states after observing the private history h_i^t . We denote the vector of private states by $\mathbf{w} =$

(w_1, \dots, w_N) , and \mathbf{w}_{-i} has the obvious interpretation. We also write $d(\mathbf{w}) \equiv (d_1(w_1), \dots, d_N(w_N))$. Finally, $\beta_i(\sigma(h_i^t) \mathbf{1} | h_i^t)$ is the posterior probability that player i assigns to all the other players being in the same private state as player i after the private history h_i^t , i.e., in the private state $\sigma(h_i^t)$.

Lemma 3. *Fix δ and a public profile. For all $\nu > 0$, there exists $\eta > 0$ and $\varepsilon > 0$, such that if the posterior beliefs induced by the private profile satisfy $\beta_i(\sigma(h_i^t) \mathbf{1} | h_i^t) > 1 - \eta$ for all h_i^t , and if π is ε -close to ρ , then for all \hat{h}_i^t ,*

$$\left| V_i(\hat{h}_i^t) - \phi_i(\sigma(\hat{h}_i^t)) \right| < \nu.$$

Proof. Consider first the case where the stage game payoff is independent of π and given by $u_i(a)$.

Under the strategy profile s , the vector of actions chosen in each period is only a function of the vector of private states. Fix a private history for player i , h_i^t , and let $w_i^t \equiv \sigma(h_i^t)$. Note that since V_i is in the convex hull of stage game payoffs, $|V_i(h_i^t)| \leq B \equiv \max_{i,a} |u_i(a)|$ for all h_i^t . Since V_i and ϕ_i are the equilibrium continuation values in the games with private and public monitoring, respectively, we have

$$\begin{aligned} V_i(h_i^t) = & \sum_{\mathbf{w}_{-i}} \left\{ (1 - \delta) u_i(d_i(w_i^t), d_{-i}(\mathbf{w}_{-i})) \right. \\ & \left. + \delta \sum_{y_i} \sum_{\mathbf{y}_{-i} \in Y^{N-1}} V_i(h_i^t; d_i(w_i^t), y_i) \pi(y_i, \mathbf{y}_{-i} | d_i(w_i^t), d_{-i}(\mathbf{w}_{-i})) \right\} \beta_i(\mathbf{w}_{-i} | h_i^t) \end{aligned}$$

and

$$\phi_i(w_i^t) = (1 - \delta) u_i(d(w_i^t)) + \delta \sum_{y_i} \phi_i(\sigma(y_i; w_i^t)) \rho(y_i | d(w_i^t)).$$

Then, since $\beta_i(w_i^t \mathbf{1} | h_i^t) > 1 - \eta$, $|V_i(h_i^t) - \phi_i(w_i^t)| \leq$

$$\begin{aligned} & (1 - \eta) \delta \sum_{y_i} \left| V_i(h_i^t; d_i(w_i^t), y_i) - \phi_i(\sigma(y_i; w_i^t)) \rho(y_i | d(w_i^t)) \right| + 2\eta B \\ & \leq \delta(1 - \eta) \sum_{y_i} |V_i(h_i^t; d_i(w_i^t), y_i) - \phi_i(\sigma(y_i; w_i^t))| \rho(y_i | d(w_i^t)) + 2\eta B \end{aligned}$$

$$\begin{aligned}
& +\delta(1-\eta) \sum_{y_i} |V_i(h_i^t; d_i(w_i^t), y_i)| \times \left| \sum_{\mathbf{y}_{-i} \in Y^{N-1}} \pi(y_i, \mathbf{y}_{-i} | d(w_i^t)) - \rho(y_i | d(w_i^t)) \right| \\
& \leq \delta(1-\eta) \sum_{y_i} |V_i(h_i^t; d_i(w_i^t), y_i) - \phi_i(\sigma(y_i; w_i^t))| \rho(y_i | d(w_i^t)) + 2\eta B \\
& +\delta(1-\eta) B \sum_{y_i} \left| \pi(y_i \mathbf{1} | d(w_i^t)) - \rho(y_i | d(w_i^t)) + \sum_{\mathbf{y}_{-i} \neq y_i \mathbf{1}} \pi(y_i, \mathbf{y}_{-i} | d(w_i^t)) \right| \\
& \leq \delta(1-\eta) \sum_{y_i} |V_i(h_i^t; d_i(w_i^t), y_i) - \phi_i(\sigma(y_i; w_i^t))| \rho(y_i | d(w_i^t)) + 2\eta B \\
& \quad +\delta(1-\eta) B \sum_{y_i} \left\{ \varepsilon + \sum_{\mathbf{y}_{-i} \neq y_i \mathbf{1}} \pi(y_i, \mathbf{y}_{-i} | d(w_i^t)) \right\} \\
& \leq \delta(1-\eta) \sum_{y_i} |V_i(h_i^t; d_i(w_i^t), y_i) - \phi_i(\sigma(y_i; w_i^t))| \rho(y_i | d(w_i^t)) + 2\eta B + \delta(1-\eta) B \varepsilon 2|Y|.
\end{aligned}$$

Taking the supremum on both sides with respect to private histories gives

$$\sup_{h_i^t} |V_i(h_i^t) - \phi_i(\sigma(h_i^t))| \leq \delta(1-\eta) \sup_{\hat{h}_i^t} |V_i(\hat{h}_i^t) - \phi_i(\sigma(\hat{h}_i^t))| + 2B(\eta + (1-\eta)\delta\varepsilon|Y|).$$

Rearranging,

$$\sup_{h_i^t} |V_i(h_i^t) - \phi_i(\sigma(h_i^t))| \leq \frac{2B(\eta + (1-\eta)\delta\varepsilon|Y|)}{1 - \delta(1-\eta)},$$

and, for fixed δ , the right hand side can be made arbitrarily small by choosing η and ε small.

Now consider $u_i(a) = \sum_{\mathbf{y}} u_i^*(y_i, a_i) \pi(\mathbf{y} | a)$. Since $\left| \sum_{\mathbf{y}} u_i^*(y_i, a_i) \pi(\mathbf{y} | a) - \sum_{y_i} u_i^*(y_i, a_i) \rho(y_i | a) \right|$ can be made arbitrarily small by choosing ε small, the above argument applies, *mutatis mutandis*. \blacksquare

Given this lemma, the next result is almost immediate. It is worth emphasizing again that the induced profile in the game with private monitoring is an equilibrium (for sufficiently small ε) for the *same* discount factor; in contrast to Example 1, the bound on ε depends on δ , and this bound becomes tighter as $\delta \rightarrow 1$.

Theorem 1. *Suppose the public profile is a strict equilibrium of the game with public monitoring for some δ and that it has a finite state automaton description. For all $\kappa > 0$, there exists η and ε such that if the posterior beliefs induced by the private profile satisfy $\beta_i(\sigma(h_i^t) \mathbf{1} | h_i^t) > 1 - \eta$ for all h_i^t , and if π is ε -close to ρ , then the private profile is an equilibrium of the game with private monitoring for the same δ , and the expected payoff in that equilibrium is within κ of the public equilibrium payoff.*

Proof. Let (W, w^1, σ, d) be the automaton description of the public profile. The one-stage deviation principle applies, so it is enough to show that for all \hat{h}_i^t , no player has an incentive in period t to choose an action different from $\hat{a}_i = d_i(\sigma(\hat{h}_i^t))$. Fix a history \hat{h}_i^t . Let $w_i^t = \sigma(\hat{h}_i^t)$. Since the public monitoring equilibrium is strict, there exists θ such that, for all $a_i \neq \hat{a}_i = d_i(w_i^t)$,

$$\phi_i(w_i^t) - \theta \geq (1 - \delta) u_i(a_i, d_{-i}(w_i^t)) + \delta \sum_y \phi_i(\sigma(y; w_i^t)) \rho(y | a_i, d_{-i}(w_i^t)). \quad (7)$$

Using Lemma 3, by choosing η and $\bar{\varepsilon}$ sufficiently small, we have for all $a_i \neq \hat{a}_i$,

$$V_i(\hat{h}_i^t) \geq \sum_{\mathbf{w}_{-i}} \left\{ (1 - \delta) u_i(a_i, d_{-i}(\mathbf{w}_{-i})) + \delta \sum_{y_i} \sum_{\mathbf{y}_{-i} \in Y^{N-1}} V_i(h_i^t; d_i(\sigma(h_i^t)), y_i^{t+1}) \pi(y_i, \mathbf{y}_{-i} | d_i(\sigma(h_i^t)), d_{-i}(\mathbf{w}_{-i})) \right\} \beta_i(\mathbf{w}_{-i} | h_i^t).$$

Since A_i and W are finite, θ can be chosen independent of \hat{h}_i^t , and so we are done. ■

A similar result holds for strict public profiles that have an infinite state automaton description, as long as the incentive constraints are “uniformly strict,” i.e., θ in (7) can be chosen independently of the state w_i^t .

Thus, the key question is the behavior of $\beta_i(\sigma(h_i^t) \mathbf{1} | h_i^t)$. In particular, can $\beta_i(\sigma(h_i^t) \mathbf{1} | h_i^t)$ be made arbitrarily close to 1 uniformly in h_i^t , by choosing ε sufficiently small? It is straightforward to show that for any integer T and $\eta > 0$, there is an $\bar{\varepsilon}$ such that for any private monitoring distribution that is ε -close to the public monitoring distribution, with $\varepsilon \in (0, \bar{\varepsilon})$, the beliefs for player i satisfy $\beta_i(\sigma(h_i^t) \mathbf{1} | h_i^t) > 1 - \eta$ for $t \leq T$. The difficulty is in extending this to arbitrarily large T . As T becomes large, the bound on ε becomes tighter.

There is one important case where $\beta_i(\sigma(h_i^t) \mathbf{1} | h_i^t)$ can be made arbitrarily close to 1 uniformly in h_i^t , and that is when the public strategy profile only requires finite memory of the public signals.

Definition 2. A public profile has finite memory (of public signals) if there is an integer L such that W can be taken to be the set $(Y \cup \{*\})^L$, and $\sigma(y, (y^L, \dots, y^2, y^1)) = (y, y^L, \dots, y^2)$ for all $y \in Y$. The initial state is $w^1 = (*, \dots, *)$.

We have introduced the “dummy” signal $*$ to account for the first L periods. This allows for finite memory profiles in which behavior in the first L periods is different from that when L periods have elapsed. An example of a profile that does not have finite memory is the grim trigger profile in the infinitely repeated prisoner’s dilemma. In this profile, player i plays C in the first period, and continues to play C as long as \bar{y} is observed, and plays D forever if ever \underline{y} is observed. While it is true that this profile only requires one-period memory if that memory includes both last period signal *and* action, if only signals can be remembered, then the entire history of signals is required: the strategy requires player i to play C after $(\bar{y}, \bar{y}, \dots, \bar{y})$, and to play D after $(\underline{y}, \bar{y}, \dots, \bar{y})$. Moreover, this dependence on arbitrarily long histories of signals implies that in some situations, even though the public trigger profile is an equilibrium, the private trigger profile is not an equilibrium in close-by games (Example 5). On the other hand, other significant strategies from the literature on repeated games with imperfect public monitoring do have finite memory: the two-phase, “stick and carrot,” strategies of Abreu [1] have finite memory and are optimal punishments, within the class of symmetric strategies, in a repeated Cournot game.

Theorem 2. Given a finite memory public profile, for all $\eta > 0$, there exists $\varepsilon > 0$ such that if π is ε -close to ρ , the posterior beliefs induced by the private profile satisfy $\beta_i(\sigma(h_i^t) \mathbf{1} | h_i^t) > 1 - \eta$ for all h_i^t .

Proof. Denote by L the length of the memory of the public profile. Each player’s private state is determined by the last L observations of his/her private signal. Suppose $t + 1 > L$ and denote player i ’s last L observations by $w \equiv (y_i^1, \dots, y_i^L)$

(this is just player i 's private state w_i^{t+1}). In period τ , $t+1-L \leq \tau \leq t$, player i chooses action $a_i^\tau = d_i(w_i^\tau)$, where w_i^τ is player i 's private state in period τ , given the private state w_i^{t+1-L} and the sequence of private observations y_i^1, \dots, y_i^ℓ , where $\ell = \tau - (t-L)$. Note that the index ℓ runs from 1 to L . For notational simplicity, we write a_i^ℓ for $a_i^{t-L+\ell}$. We need to show that by making ε sufficiently small, the probability that player i assigns to all the other players observing the same sequence of private signals in the last L periods can be made arbitrarily close to 1. Let $a^{(L)} \in A^L$ denote a sequence of L action profiles, where $a_{-i}^\ell \in A_{-i}$ is arbitrary. Then,

$$\Pr \left\{ \mathbf{w} = w\mathbf{1} | a^{(L)} \right\} = \prod_{\ell=1}^L \pi \left(y_i^\ell \mathbf{1} | a^\ell \right)$$

and

$$\Pr \left\{ w_i = w | a^{(L)} \right\} = \sum_{(\mathbf{y}_{-i}^1, \dots, \mathbf{y}_{-i}^L) \in Y^{(N-1)L}} \prod_{\ell=1}^L \pi \left(\mathbf{y}_{-i}^\ell, y_i^\ell | a^\ell \right).$$

Since these probabilities are conditional on the actions taken in the last L periods, they do not depend upon player i 's private state in period $t+1-L$. Then for any $\eta > 0$, there exists $\bar{\varepsilon} > 0$ such that for all $a^{(L)} \in A^L$ and $\varepsilon \in (0, \bar{\varepsilon})$,

$$\begin{aligned} \Pr \left\{ \mathbf{w}_{-i} = w\mathbf{1} | w_i = w, a^{(L)} \right\} &= \frac{\Pr \left\{ \mathbf{w} = w\mathbf{1} | a^{(L)} \right\}}{\Pr \left\{ w_i = w | a^{(L)} \right\}} \\ &= \frac{\prod_{\ell=1}^L \pi \left(y_i^\ell \mathbf{1} | a^\ell \right)}{\sum_{(\mathbf{y}_{-i}^1, \dots, \mathbf{y}_{-i}^L) \in Y^{(N-1)L}} \prod_{\ell=1}^L \pi \left(\mathbf{y}_{-i}^\ell, y_i^\ell | a^\ell \right)} \\ &> 1 - \eta. \end{aligned}$$

[For $\varepsilon = 0$, $\Pr \left\{ \mathbf{w}_{-i} = w\mathbf{1} | w_i = w, a^{(L)} \right\} = 1$. Moreover, the denominator is bounded away from zero, for all $\varepsilon \geq 0$ and all $a^{(L)} \in A^L$, and so continuity implies the result. The details are almost identical to the proof of Lemma 2.]

Let $\lambda \in \Delta(A^L)$ denote the beliefs for player i over the last L actions taken by the other players after observing the private signals w . Then,

$$\begin{aligned} \Pr \left\{ \mathbf{w}_{-i} = w\mathbf{1} | w_i = w \right\} &= \sum_{a^{(L)} \in A^L} \Pr \left\{ \mathbf{w}_{-i} = w\mathbf{1} | w_i = w, a^{(L)} \right\} \lambda \left(a^{(L)} \right) \\ &> (1 - \eta) \sum_{a^{(L)} \in A^L} \lambda \left(a^{(L)} \right) = 1 - \eta. \end{aligned}$$

■




Figure 6: A profile supporting some cooperation when p and q are both close to 1.

5. Examples

To illustrate the general problem, we analyze a series of examples in some depth.

Example 2. (*Prisoner’s Dilemma again.*) Consider the environment described in Example 1. For $3p > 2 + q$ and $\delta \geq 1/(3p - q - 1)$, the profile described as follows is a public perfect equilibrium: $W = \{\underline{w}, \bar{w}\}$, $w^1 = \bar{w}$, $d_i(\bar{w}) = C$, $d_i(\underline{w}) = D$, and

$$\sigma(yw) = \begin{cases} \bar{w}, & \text{if } w = \bar{w} \text{ and } y = \bar{y}, \\ & \text{or } w = \underline{w} \text{ and } y = \underline{y}, \\ \underline{w}, & \text{if } w = \bar{w} \text{ and } y = \underline{y}, \\ & \text{or } w = \underline{w} \text{ and } y = \bar{y}. \end{cases}$$

The idea of the profile is that behavior starts at CC , and continues there as long as the “good” signal \bar{y} is observed, and then switches to DD after the “bad” signal \underline{y} . In order to generate sufficient punishment, the expected duration in the punishment state \underline{w} cannot be too short, and so play only leaves DD after the less likely signal \underline{y} is realized. The profile is illustrated in Figure 6. This profile (unlike the previous one) does not require $q < 2/3$, although it does require $p > 2/3$.

This profile cannot be approximated in some close-by games with private monitoring. Consider the private monitoring technology π obtained by the compound randomization in which in the first stage a value of y is determined according to ρ , and then in the second stage, that value is reported to player i with probability

$(1 - \varepsilon)$ and the other value with probability ε ; conditional on the realization of the first stage, the second stage randomizations are independent across players.

Suppose that the associated private profile is an equilibrium. The profile induces a Markov chain on the state space $W^2 \equiv \{\bar{w}\bar{w}, \underline{w}\underline{w}, \bar{w}\underline{w}, \underline{w}\bar{w}\}$, with transition matrix:

$$\begin{pmatrix} p(1-\varepsilon)^2 + (1-p)\varepsilon^2 & (1-q)(1-\varepsilon)^2 + q\varepsilon^2 & \varepsilon(1-\varepsilon) & \varepsilon(1-\varepsilon) \\ (1-p)(1-\varepsilon)^2 + p\varepsilon^2 & q(1-\varepsilon)^2 + (1-q)\varepsilon^2 & \varepsilon(1-\varepsilon) & \varepsilon(1-\varepsilon) \\ \varepsilon(1-\varepsilon) & \varepsilon(1-\varepsilon) & q(1-\varepsilon)^2 + (1-q)\varepsilon^2 & (1-q)(1-\varepsilon)^2 + q\varepsilon^2 \\ \varepsilon(1-\varepsilon) & \varepsilon(1-\varepsilon) & (1-q)(1-\varepsilon)^2 + q\varepsilon^2 & q(1-\varepsilon)^2 + (1-q)\varepsilon^2 \end{pmatrix}.$$

This Markov chain has a unique invariant distribution α^ε , where

$$\Pr(\bar{w}\bar{w}) \equiv \alpha_1^\varepsilon = \frac{(1-q)(1-\varepsilon)^2 + q\varepsilon^2 + \varepsilon(1-\varepsilon)}{2\left((p+q)\varepsilon^2 + (2-p-q)(1-\varepsilon)^2 + 2\varepsilon(1-\varepsilon)\right)},$$

$$\Pr(\underline{w}\underline{w}) \equiv \alpha_2^\varepsilon = \frac{(1-p)(1-\varepsilon)^2 + p\varepsilon^2 + \varepsilon(1-\varepsilon)}{2\left((p+q)\varepsilon^2 + (2-p-q)(1-\varepsilon)^2 + 2\varepsilon(1-\varepsilon)\right)},$$

and

$$\Pr(\bar{w}\underline{w}) \equiv \alpha_3^\varepsilon = \frac{1}{4} = \Pr(\underline{w}\bar{w}) \equiv \alpha_4^\varepsilon.$$

For ε small, this distribution is close to α^0 , where

$$\alpha_1^0 = \frac{(1-q)}{2(2-p-q)}, \quad \alpha_2^0 = \frac{(1-p)}{2(2-p-q)}, \quad \text{and} \quad \alpha_3^0 = \alpha_4^0 = \frac{1}{4}.$$

Consider now the question of what beliefs a player should have over the opponents' private state after a very long history. Observe first that the probability that 1 assigns to 2 being in state \bar{w} , conditional on 1 being in state \bar{w} , is close to (for ε small)

$$\Pr\{\bar{w}|\bar{w}\} = \frac{2(1-q)}{2(1-q) + (2-p-q)} = \frac{2(1-q)}{4-p-3q},$$

while the probability that 1 assigns to 2 being state in \bar{w} , conditional on 1 being in state \underline{w} , is close to

$$\Pr\{\bar{w}|\underline{w}\} = \frac{2-p-q}{4-3p-q}.$$

Then, for ε small,

$$\Pr \{\bar{w}|\bar{w}\} < \Pr \{\bar{w}|\underline{w}\}.$$

Since this probability is the (asymptotic) expected value of the player's beliefs, where expectations are taken over the private histories, there are two histories h_i^t and \hat{h}_i^t such that $\underline{w} = \sigma(h_i^t)$ and $\bar{w} = \sigma(\hat{h}_i^t)$ and $\Pr \{\bar{w}|\hat{h}_i^t\} < \Pr \{\bar{w}|h_i^t\}$. But if C is optimal after \hat{h}_i^t , then it must be optimal after h_i^t , and so the profile is not an equilibrium of the game with private monitoring.

In this example, once there is disagreement in private states, the public profile maintains disagreement. Moreover, when there is private monitoring, disagreement arises almost surely, and so players must place substantial probability on disagreement. Thus, a necessary condition for beliefs to be asymptotically well behaved is that the public profile at least sometimes moves a vector of private states in disagreement into agreement. The following property of connectedness also plays a critical role in obtaining on ε that are uniform in δ , for δ close to 1.

Definition 3. A public profile is connected if, for all $w, w' \in W$, there exists m and y^1, \dots, y^m and $\bar{w} \in W$ such that

$$\sigma(y^m, \sigma(y^{m-1}, \dots, \sigma(y^1, w))) = \bar{w} = \sigma(y^m, \sigma(y^{m-1}, \dots, \sigma(y^1, w'))).$$

While stated for any two states, connectedness implies that any disagreement over all the players is removed after some sequence of public signals, and so, is removed eventually with probability one.

Lemma 4. For any connected finite public profile, there is a finite sequence of signals y^1, \dots, y^n and a state \bar{w} such that

$$\sigma(y^n, \sigma(y^{n-1}, \dots, \sigma(y^1, w))) = \bar{w}, \quad \forall w \in W.$$

Proof. We prove this for $|W| = 3$, the extension to an arbitrary finite number of states is straightforward. Fix w^1, w^2 , and w^3 . Let y^1, \dots, y^m be a sequence that satisfies $\sigma(y^m, \sigma(y^{m-1}, \dots, \sigma(y^1, w^1))) = \sigma(y^m, \sigma(y^{m-1}, \dots, \sigma(y^1, w^2))) \equiv w$. Since the profile is connected, there is a sequence of signals $y^{m+1}, \dots, y^{m+m'}$ such that $\sigma(y^{m+m'}, \sigma(y^{m+m'-1}, \dots, \sigma(y^{m+1}, w))) = \sigma(y^{m+m'}, \sigma(y^{m+m'-1}, \dots, \sigma(y^{m+1}, w')))$, where $w' \equiv \sigma(y^m, \sigma(y^{m-1}, \dots, \sigma(y^1, w^3)))$. The desired sequence of signals is then $y^1, \dots, y^{m+m'}$. ■

The profile in Example 2 is not connected. However, connectedness is not sufficient, as the following example illustrates. Connectedness only ensures that any disagreement in private states is eliminated eventually, with probability one. On the other hand, it may take a long time before all players simultaneously observe the sequence of signals that leads to agreement. Moreover, before this “agreement” sequence is observed, each player will be updating her beliefs over the private states of her opponents. Even if initial beliefs place small probability on disagreement (over private states), the updates eventually may place large probability on disagreement. Since monitoring is almost public, a player’s updating of her beliefs over the private states of her opponents given the observation of a public signal y (determined by the public monitoring distribution ρ) is close to her update after observing the private signal y (determined by the private monitoring distribution π). Thus, in the next example, we examine the updating of a player with nondegenerate beliefs over the private states of her opponent in the game with public monitoring.

Example 3. *This profile is a modification of Example 2 which is connected: $W = \{\underline{w}, \tilde{w}, \bar{w}\}$, $w^1 = \bar{w}$, $d_i(\bar{w}) = C$, $d_i(\tilde{w}) = d_i(\underline{w}) = D$, and*

$$\sigma(yw) = \begin{cases} \bar{w}, & \text{if } w = \bar{w} \text{ and } y = \bar{y}, \\ & \text{or } w = \underline{w} \text{ and } y = \underline{y}, \\ \tilde{w}, & \text{if } w = \bar{w} \text{ and } y = \underline{y}, \\ & \text{or } w = \tilde{w} \text{ and } y = \underline{y}, \\ \underline{w}, & \text{if } w = \tilde{w} \text{ and } y = \bar{y}, \\ & \text{or } w = \underline{w} \text{ and } y = \bar{y}. \end{cases}$$

It is illustrated in Figure 7. This profile is an equilibrium if $3p > 1 + 3q - q^2$ (which is implied by $3p > 2 + q$ and $3(p - q) > 1$) and δ is sufficiently close to 1. Note that the sequence $\underline{y}, \underline{y}$ always leads to the state \tilde{w} . However, this is not sufficient for this profile to be an equilibrium of the game with private monitoring.

Before considering the game with private monitoring, consider the game with public monitoring when a player does not assign probability one to the other player being in the same private state (as would occur under private monitoring). More specifically, suppose player 1’s is initially in private state $w_1^1 = \bar{w}$ and has beliefs $\beta = (1 - \eta^1, 0, \eta^1)$ on $(\bar{w}, \tilde{w}, \underline{w})$. Now consider the impact on player i ’s beliefs of the sequence $\underline{y} \bar{y} \underline{y} \bar{y} \underline{y} \bar{y} \dots$. The first observation of \underline{y} causes player 1’s private state to change to \tilde{w} . In addition, player 1 will update her beliefs over




Figure 7: The sequence $\underline{y}, \underline{y}$ always leads to the state \tilde{w} .

the private state of player 2 according to:

$$\begin{aligned}
 \Pr \{w_2^2 = \tilde{w} | w_1^2 = \tilde{w}\} &\equiv 1 - \eta^2 = \Pr \{w_2^2 = \tilde{w} | \underline{y}, w_1^1 = \bar{w}\} \\
 &= \Pr \{w_2^1 = \bar{w} | \underline{y}, w_1^1 = \bar{w}\} + \Pr \{w_2^1 = \tilde{w} | \underline{y}, w_1^1 = \bar{w}\} \\
 &= \frac{\Pr \{w_2^1 = \bar{w}, \underline{y} | w_1^1 = \bar{w}\} + \Pr \{w_2^1 = \tilde{w}, \underline{y} | w_1^1 = \bar{w}\}}{\Pr \{\underline{y} | w_1^1 = \bar{w}\}} \\
 &= \frac{\Pr \{\underline{y} | w_1^1 = \bar{w}, w_2^1 = \bar{w}\} \Pr \{w_2^1 = \bar{w}\} + \Pr \{\underline{y} | w_1^1 = \bar{w}, w_2^1 = \tilde{w}\} \Pr \{w_2^1 = \tilde{w}\}}{\Pr \{\underline{y} | w_1^1 = \bar{w}\}} \\
 &= \frac{(1-p)(1-\eta^1)}{(1-p)(1-\eta^1) + (1-q)\eta^1}
 \end{aligned}$$

which is strictly smaller than $1 - \eta^1$, since $q < p$. Thus, after observing \underline{y} , $\eta^2 > \eta^1$ and so player 1 assigns a lower probability to the other player being in the same private state. The next two observations, of \bar{y} and \underline{y} , do not result in any updating

signal	1's private state	Beliefs over $(\bar{w}, \tilde{w}, \underline{w})$
	\bar{w}	$(1 - \eta^1, 0, \eta^1)$
\underline{y}	\tilde{w}	$(\eta^2, 1 - \eta^2, 0)$
\bar{y}	\underline{w}	$(\eta^2, 0, 1 - \eta^2)$
\underline{y}	\bar{w}	$(1 - \eta^2, \eta^2, 0)$
\bar{y}	\bar{w}	$(1 - \eta^4, 0, \eta^4)$

Figure 8: The transitions on beliefs induced by a sequence of public signals.

of the posterior. The only changes in beliefs reflect the transition function on private states (see Figure 8). After the second \underline{y} , player 1 now observes \bar{y} and again nontrivial updating can occur:

$$\begin{aligned} & \Pr \{w_2^4 = \bar{w} | w_2^4 = \bar{w}\} \equiv 1 - \eta^4 = \Pr \{w_2^3 = \bar{w} | \bar{y}, w_2^3 = \bar{w}\} \\ &= \frac{\Pr \{\bar{y} | w_2^3 = \bar{w}, w_2^3 = \bar{w}\} \Pr \{w_2^3 = \bar{w}\}}{\Pr \{\bar{y} | w_2^3 = \bar{w}\}} = \frac{p(1 - \eta^2)}{p(1 - \eta^2) + q\eta^2} \end{aligned}$$

which is strictly larger than $1 - \eta^2$, since $q < p$. The issue then is whether η^1 is larger or smaller than η^4 . In order for player 1 to assign a probability close to one that player 2 is in private state \bar{w} when 1 is in private state \bar{w} after an arbitrarily long history of the form $\underline{y} \bar{y} \underline{y} \bar{y} \underline{y} \bar{y} \cdots$, it is necessary and sufficient that for η^1 in a neighborhood of 0, $\eta^4 < \eta^1$. Since when $\eta^1 = 0$, $\eta^4 = 0$, it is thus enough to investigate when $d\eta^4/d\eta^1|_{\eta^1=0} < 1$. Now,

$$\left. \frac{d\eta^4}{d\eta^1} \right|_{\eta^1=0} = \frac{d\eta^4}{d\eta^2} \times \left. \frac{d\eta^2}{d\eta^1} \right|_{\eta^1=0} = \frac{q}{p} \times \frac{(1 - q)}{(1 - p)}$$

and so $d\eta^4/d\eta^1|_{\eta^1=0} < 1$ if and only if $q(1 - q) < p(1 - p)$ (i.e., $p < 1 - q$). In fact, if $p > 1 - q$, after a sufficiently long history of the form $\underline{y} \bar{y} \underline{y} \bar{y} \underline{y} \bar{y} \cdots$, player 1 assigns a probability very close to 0 that player 2 is also in private state \bar{w} when 1 is. Note that while the condition $q < p < 1 - q$ requires $q < 1/2$, it is not inconsistent with existence of equilibrium for large δ , since the condition on p and q is $3p > 1 + 3q - q^2$, which can be satisfied by small q . In fact, there are values of p and q for which the profile in Example 1 is not an equilibrium, and this profile is (since Example 1 requires $3p > 1 + 3q$).

The intuition for the condition $q(1 - q) < p(1 - p)$ is as follows. Note first that on this history, nontrivial updating can only occur when player 1 is in private state \bar{w} (since in the other private states, player 1 is choosing D , and so the distribution of the public signal is independent of the behavior of player 2). On the subsequence where player 1 is in private state \bar{w} , half the observations are of the public signal \bar{y} while the other half are of the signal \underline{y} . Finally, $\Pr\{\bar{y}\underline{y}|CD, CD\} = q(1 - q)$ and $\Pr\{\bar{y}\underline{y}|CC, CC\} = p(1 - p)$, so $q(1 - q) < p(1 - p)$ is equivalent to $\Pr\{\bar{y}\underline{y}|CD\} < \Pr\{\bar{y}\underline{y}|CC\}$.

If player 1 starts with beliefs that assign positive probability to player 2 being in each of the three private states, after the sequence of public signals $\underline{y}\bar{y}\underline{y}\bar{y}$, player 1 assigns zero probability to player 2 being in the state \tilde{w} , and so considering beliefs $(1 - \eta^1, 0, \eta^1)$ is without loss of generality. Moreover, if player 1's beliefs after the history $\underline{y}\bar{y}\underline{y}\bar{y}\underline{y}\bar{y}\dots$ are well behaved (in the sense that the probability assigned to agreement does not move away from neighborhoods of 1), then they are well behaved after all histories. This is for three reasons: First, any history that contains two consecutive \underline{y} leads to immediate agreement on \tilde{w} . Second, other histories contain consecutive \bar{y} signals. And, third, signals are only informative about the private state of the other player when player 1 is in private state \bar{w} , in which case observing \bar{y} raises the posterior that player 2 is also in private state \bar{w} .

Consider now games with private monitoring. For ε sufficiently small, the private profile will be an equilibrium if and only if $p < 1 - q$ (the bound on ε depends on p and q): Again, the critical private histories we need to consider are of the form $\underline{y}\bar{y}\underline{y}\bar{y}\underline{y}\bar{y}\dots$. Let $\beta^\varepsilon(\eta^1) \equiv \beta^\varepsilon(\bar{w}|\underline{y}\bar{y}\underline{y}\bar{y})(\eta^1)$ denote the posterior probability assigned by player 1 to player 2 being in private state \bar{w} after observing the private history $\underline{y}\bar{y}\underline{y}\bar{y}$ and with prior probability $(1 - \eta^1, 0, \eta^1)$. We have seen that if $p < 1 - q$, then $\beta^0(\eta^1) > 1 - \eta^1$, so that $\lim_{t \rightarrow \infty} (\beta^0)^t(\eta^1) = 1$. By choosing ε small, we can make $\sup_{\eta^1} |\beta^\varepsilon(\eta^1) - \beta^0(\eta^1)|$ as small as we like. In particular, for any $\bar{\eta} > 0$, there is an $\bar{\varepsilon} > 0$ such that for all $\varepsilon \in (0, \bar{\varepsilon})$, $\beta^\varepsilon(\eta^1) > 1 - \eta^1$ for all $\eta^1 \in (\bar{\eta}, 1 - \bar{\eta})$. But this implies $\lim_{t \rightarrow \infty} (\beta^\varepsilon)^t(\eta^1) \geq 1 - \bar{\eta}$ for all $\eta^1 > \bar{\eta}$.

Example 4. The assumed signaling structure of the previous example simplified calculations, since some signals were uninformative. In this example, we show that the phenomenon displayed in that example did not require uninformative signals. Consider the same strategy profile as in the previous example, but with a different signaling distribution. In particular, suppose

$$\rho\{\bar{y}|a_1a_2\} = \begin{cases} p, & \text{if } a_1a_2 = CC, \\ q, & \text{if } a_1a_2 = CD \text{ or } DC, \\ r, & \text{if } a_1a_2 = DD, \end{cases}$$

with $p > q > r$. The public profile described in the previous example remains an equilibrium with this new ρ , as long as r is close to q . The important difference is that the signal in every private state is potentially informative. This implies that even if $p < 1 - q$, the private profile is not an equilibrium. To see this, consider the public monitoring game after the history $\underline{y}\bar{y}\bar{y}\bar{y}\bar{y}\cdots$. After such a history, player 1 is in private state \underline{w} . Suppose that in period 1 player 1 assigned positive probability to player 2 starting in private state \underline{w} . If player 2 started in private state \underline{w} , then after the sequence $\underline{y}\bar{y}\bar{y}\bar{y}\bar{y}\cdots$, player 2 will be in private state \bar{w} , and since $q > r$, player 1's posterior that player 2 is in state \bar{w} converges to 1 as the number of consecutive observations of \bar{y} goes to infinity.

Suppose now that $p > r > q$ (an admittedly unusual, but generic, configuration). Observe first that after very long histories of the form $\underline{y}\bar{y}\bar{y}\bar{y}\bar{y}\cdots$, player 1 is in private state \underline{w} and assigns very high probability that player 2 is in the same private state ($\Pr\{\bar{y}|DD\} > \Pr\{\bar{y}|DC\}$). The critical history turns out again to be those of the form $\underline{y}\bar{y}\bar{y}\bar{y}\bar{y}\cdots$. The probability of observing the sequence $\underline{y}\bar{y}\bar{y}\bar{y}$ when both players start in private state \bar{w} is $(1-p)r(1-r)p$, while the probability of observing the same sequence when 1 starts in state \bar{w} , while 2 starts in state \underline{w} , is $(1-q)q(1-q)q$. Thus, a sufficient condition for the private profile to be an equilibrium for ε small is that $p > r > 1 - q$.

These examples illustrate the difficulty that arises in checking whether a given public equilibrium induces a private equilibrium: there may be some histories that yield beliefs that assign substantial probability to players being in different private states. Sometimes, it is possible to show that such histories occur with very low probability: this is true in Examples 3 and 4 and can be shown to be true for all connected and finite public profiles.¹¹ Unfortunately, this is not enough to have an equilibrium under private monitoring, and there is no simple way to adjust such strategy profiles to make them equilibria.

We conclude this section with a brief discussion of grim trigger. We include this example because of the central role grim trigger has played in the earlier literature, in particular, Sekiguchi [28] and Compte [10]. After the example, we discuss the result of Sekiguchi [28].

Example 5. Consider again the environment described in Example 1. For $3p > 1 + 2q$ and $\delta \geq 1/(3p - 2q)$, the profile described as follows is a public perfect

¹¹Under some additional technical conditions (that are satisfied if ρ is almost uniform), the following is true: For all $\xi > 0$ and t , there exists $\bar{\varepsilon} > 0$ such that if π is ε -close to ρ , then the unconditional probability that any player assigns a probability of less $1 - \xi$ to all the players being in the same private state as herself is less than ξ .

equilibrium: $W = \{\underline{w}, \bar{w}\}$, $w^1 = \bar{w}$, $d_i(\bar{w}) = C$, $d_i(\underline{w}) = D$, and

$$\sigma(yw) = \begin{cases} \bar{w}, & \text{if } w = \bar{w} \text{ and } y = \bar{y}, \\ \underline{w}, & \text{if } w = \bar{w} \text{ and } y = \underline{y}, \text{ or } w = \underline{w}. \end{cases}$$

In this profile, behavior starts at CC , and continues there as long as the “good” signal \bar{y} is observed, and then switches to DD permanently after the “bad” signal \underline{y} . As $\delta \rightarrow 1$, the average payoff in this profile converges to 0, since the expected waiting time till a bad signal is finite and independent of δ . In order for the payoff in this profile to be close to 2 as $\delta \rightarrow 1$, we need to take $p \rightarrow 1$.

This trigger profile is an equilibrium of the first game with private monitoring described in Example 1.¹² Intuitively, if this profile is not an equilibrium, it is because, after a very long sequence of $\bar{y} \cdots \bar{y}$, player i is no longer sufficiently confident that the opponent is still in private state \bar{w} . But, note that since player i has been in state \bar{w} throughout this sequence, and so has always been choosing C , continually observing \bar{y} is confirming evidence that player j is also in \bar{w} . Thus, a similar argument to that in Example 3 shows that i 's posterior that j is in private state \bar{w} cannot get too low. The other case is more direct, since once a player is in private state \underline{w} , that player chooses D thereafter and the private signals are uninformative about the play of the opponent.

If instead we had the private monitoring of Example 4 with $q > r$, then after a sufficiently long history $\underline{y} \bar{y} \cdots \bar{y}$, a player must assign probability close to 1 that the opponent is in private state \bar{w} and so the trigger profile is not an equilibrium.

We now compare this result with the impressive result of Sekiguchi [28]. Sekiguchi [28] showed that, for some repeated prisoner's dilemmas, there exists a nearly efficient sequential equilibrium, when private monitoring is arbitrarily accurate and players are patient. There are three features we draw the reader's attention to. First, Sekiguchi's result makes few assumptions on the nature of the private monitoring (it includes both independent and correlated, though not almost-public, signals¹³). Second, his equilibrium is in mixed strategies, while ours is in pure. Finally, while his equilibrium builds on grim trigger, the final equilibrium is not grim trigger.

Crudely summarizing Sekiguchi's argument, he begins by considering a strategy profile that randomizes between always defection, and grim trigger. This profile is an equilibrium (given a payoff restriction) for moderate discount factors and sufficiently accurate private monitoring. Crudely, there are two things

¹²This is the case analyzed in Mailath and Morris [22, Section 4.2].

¹³It is crucial in Sekiguchi [28] that (in his notation) $\Pr\{dc|DC\}$ be close to 1, see footnote 15. This implies that $\Pr\{y_2 = d|a_1 = D, y_1 = c\}$ is also close to 1, which cannot occur if the private signals are almost public.

to worry about. First, if a player has been cooperating for a long time and has always received a cooperative signal, will the player continue to cooperate? The answer here is yes, given sufficiently accurate private monitoring (this is the analog to considering histories of the form $\bar{y} \cdots \bar{y}$ in Example 5).

Second, will a player defect as soon as a defect signal is received? This is where the randomization and upper bound on the discount factor comes in. For illustrative purposes, suppose the players are playing the pure strategy profile of grim trigger. After the initial period, if player i observes the defect signal, then the highest order probability events are that player j did not defect but i received an erroneous signal (in which case j is still cooperative), and that player j in the previous period had received an erroneous signal (in which case j now defects forever). These two events have equal probability, and if players are not too patient (so that they are not willing to experiment), player i will defect. If players are patient, then even a large probability that the opponent is already defecting may not be enough to ensure that the player defects: One more observation before the player commits himself may be quite valuable. Of course, in the initial period player j is not responding to any signal, so in order for player i to assign positive probability to the signal reflecting j 's behavior, j must defect in the initial period with positive probability.¹⁴ Our assumption that monitoring is almost public implies that these latter considerations are irrelevant. As soon as a player receives a defect signal, he assigns very high probability to his opponents having received the same signal, and so will defect. This is why we do not need randomization, nor an upper bound on the discount factor.¹⁵

Sekiguchi then removes the upper bound on the discount factor by observing (following Ellison [12]) that the repeated game can be divided into N distinct games, with the k th game played in period $k + tN$, where $t \in \mathbf{N}$. This gives

¹⁴In other words, grim trigger is not a Nash equilibrium because players have an incentive to ignore defect signals when first received (players believe their opponents are still cooperating and do not want to initiate the defect phase) and so players have no incentive to cooperate in the initial period.

Of course, the players must be indifferent between cooperation and defection in the initial period, and this determines the randomization probability. Moreover, as long as the discount is close to the value at which a player is indifferent between cooperation and defection against grim trigger in a game with perfect monitoring, then for sufficiently accurate monitoring, this randomization probability assigns small weight to initial defection.

¹⁵There is one last subtlety that is worth mentioning, and this concerns the incentives of a player to defect, when the player is already in the defect private state and the player observes a cooperative signal (suggesting that in fact the opponent is not in the defect private state). In Sekiguchi [28], the player still finds it optimal to defect in the future, because he believes that his choice of defect *this* period triggers a transition to the defect private state by his opponent with high probability (another implication of almost perfect monitoring). This is where it is important that the signaling is not almost public (see footnote 13).

an effective discount rate of δ^N on each game. Of course, the resulting strategy profile does not look like grim trigger.

6. The Case of Arbitrarily Patient Players

In this section, we obtain bounds on ε that are uniform in δ , for δ close to 1. We first rewrite the incentive constraints of the public monitoring game so that, at least for connected finite public profiles (see Definition 3), they make sense when evaluated at $\delta = 1$. A public profile is a strict equilibrium if, for all $i \in N$, $w \in W$, and all $a_i \neq d_i(w)$,

$$\phi_i(w) > (1 - \delta) u_i(d_{-i}(w), a_i) + \delta \sum_y \phi_i(\sigma(y; w)) \rho(y|d_{-i}(w), a_i),$$

where

$$\phi_i(w) = (1 - \delta) u_i(d(w)) + \delta \sum_y \phi_i(\sigma(y; w)) \rho(y|d(w)).$$

For simplicity, write $\hat{u}_i(w)$ for $u_i(d(w))$, and $\hat{u}_i(w, a_i)$ for $u_i(d_{-i}(w), a_i)$. For a fixed state $w \in W$, the mapping σ induces a partition on Y :

$$y_w(w') = \{y \in Y : \sigma(y; w) = w'\}.$$

The incentive constraints at w can be written more transparently, focusing on the transitions between states, as

$$\phi_i(w) > (1 - \delta) \hat{u}_i(w, a_i) + \delta \sum_{w'} \phi_i(w') \theta_{ww'}(d_{-i}(w), a_i), \quad (8)$$

where $\theta_{ww'}(a)$ is the probability of transiting from state w to state w' under the action profile a , i.e.,

$$\theta_{ww'}(a) = \begin{cases} \sum_{y \in y_w(w')} \rho(y|a), & \text{if } y_w(w') \neq \emptyset, \\ 0, & \text{if } y_w(w') = \emptyset. \end{cases}$$

Substituting for $\phi_i(w)$ in (8) and rearranging yields (writing $\hat{\theta}_{ww'}$ for $\theta_{ww'}(d(w))$ and $\hat{\theta}_{ww'}(a_i)$ for $\theta_{ww'}(d_{-i}(w), a_i)$),

$$\delta \sum_{w'} \phi_i(w') \left(\hat{\theta}_{ww'} - \hat{\theta}_{ww'}(a_i) \right) > (1 - \delta) (\hat{u}_i(w, a_i) - \hat{u}_i(w)). \quad (9)$$

For any $\bar{w} \in W$, (9) is equivalent to

$$\delta \sum_{w' \neq \bar{w}} (\phi_i(w') - \phi_i(\bar{w})) \left(\hat{\theta}_{ww'} - \hat{\theta}_{ww'}(a_i) \right) > (1 - \delta) (\hat{u}_i(w, a_i) - \hat{u}_i(w)). \quad (10)$$

If the profile is finite and connected, the Markov chain on W implied by the profile is ergodic, and so has a unique stationary distribution. As a consequence, $\lim_{\delta \rightarrow 1} \phi_i(w)$ is independent of $w \in W$, and so simply taking $\delta \rightarrow 1$ in (10) yields $0 \geq 0$. On the other hand, if we can divide by $(1 - \delta)$ and take limits as $\delta \rightarrow 1$, we have a version of the incentive constraint that is independent of δ . The next lemma assures us that this actually makes sense (the proofs of the Lemmas in this section are in the Appendix).

Lemma 5. *Suppose the public profile is finite and connected. For any two states $w, \bar{w} \in W$,*

$$\Delta_{w\bar{w}}\phi_i \equiv \lim_{\delta \rightarrow 1} (\phi_i(w) - \phi_i(\bar{w})) / (1 - \delta)$$

exists and is finite.

So, if a connected finite public profile is a strict equilibrium for discount factors arbitrarily close to 1, we have

$$\sum_{w' \neq \bar{w}} \Delta_{w'\bar{w}}\phi_i \times \left(\hat{\theta}_{ww'} - \hat{\theta}_{ww'}(a_i) \right) \geq \hat{u}_i(w, a_i) - \hat{u}_i(w).$$

Strengthening the weak inequality to a strict one gives a condition that implies (10) for δ large.

Definition 4. *A connected finite public profile is patiently strict if for all players i , states $w \in W$, and actions $a_i \neq d_i(w)$,*

$$\sum_{w' \neq \bar{w}} \Delta_{w'\bar{w}}\phi_i \times \left(\hat{\theta}_{ww'} - \hat{\theta}_{ww'}(a_i) \right) > \hat{u}_i(w, a_i) - \hat{u}_i(w), \quad (11)$$

where \bar{w} is any state.

The particular choice of \bar{w} in (11) is irrelevant: if (11) holds for one \bar{w} such that $\hat{\theta}_{w\bar{w}} \in (0, 1)$, then it holds for all such \bar{w} . The next lemma is obvious.

Lemma 6. *For any patiently-strict connected finite public profile, there exists $\underline{\delta} < 1$ such that, for all $\delta \in (\underline{\delta}, 1)$, the public profile is a strict public equilibrium of the game with public monitoring.*

The remainder of this section proves the following theorem. It is worth remembering that every finite memory public profile is both a connected finite public profile and, by Theorem 2, induces posterior beliefs that assign uniformly large probability to agreement in private states.

Theorem 3. *Suppose a connected finite public profile is patiently strict. There exist $\underline{\delta} < 1$, $\eta > 0$, and $\varepsilon > 0$ such that, if the posterior beliefs induced by the private profile satisfy $\beta_i(\sigma(h_i^t) \mathbf{1}|h_i^t) > 1 - \eta$ for all h_i^t , π is ε -close to ρ , and $\delta \in (\underline{\delta}, 1)$, then the private profile is an equilibrium of the game with private monitoring.*

The finite public profile induces in the game with private monitoring a finite state Markov chain (Z, Q^π) , where $Z \equiv W^N$ and, for $\mathbf{w} = (w_1, \dots, w_N)$ and $\mathbf{w}' = (w'_1, \dots, w'_N)$,

$$q_{\mathbf{w}\mathbf{w}'}^\pi(a) = \begin{cases} \sum_{y_1 \in y_{w_1}(w'_1)} \cdots \sum_{y_N \in y_{w_N}(w'_N)} \pi(\mathbf{y}|a), & \text{if } y_{w_i}(w'_i) \neq \emptyset \text{ for all } i, \\ 0, & \text{if } y_{w_i}(w'_i) = \emptyset \text{ for some } i. \end{cases}$$

The value to player i at the vector of private states \mathbf{w} is

$$\begin{aligned} \psi_i^\pi(\mathbf{w}) &= (1 - \delta) u_i(d(\mathbf{w})) + \delta \sum_{\mathbf{y}} \psi_i^\pi(\sigma(y_1; w_1), \dots, \sigma(y_N; w_N)) \pi(\mathbf{y}|d(\mathbf{w})) \\ &= (1 - \delta) u_i(d(\mathbf{w})) + \delta \sum_{\mathbf{w}'} \psi_i^\pi(\mathbf{w}') q_{\mathbf{w}\mathbf{w}'}^\pi(d(\mathbf{w})) \\ &= (1 - \delta) \tilde{u}_i(\mathbf{w}) + \delta \sum_{\mathbf{w}'} \psi_i^\pi(\mathbf{w}') \tilde{q}_{\mathbf{w}\mathbf{w}'}^\pi, \end{aligned}$$

where $\tilde{u}_i(\mathbf{w}) = u_i(d(\mathbf{w}))$ and $\tilde{q}_{\mathbf{w}\mathbf{w}'}^\pi = q_{\mathbf{w}\mathbf{w}'}^\pi(d(\mathbf{w}))$.

Analogous to Lemma 5, we have the following:

Lemma 7. *Suppose the public profile is finite and connected.*

1. For any two vectors of private states $\mathbf{w}, \bar{\mathbf{w}} \in W^N$,

$$\Delta_{\mathbf{w}\bar{\mathbf{w}}} \psi_i^\pi \equiv \lim_{\delta \rightarrow 1} (\psi_i^\pi(\mathbf{w}) - \psi_i^\pi(\bar{\mathbf{w}})) / (1 - \delta)$$

exists and is finite;

2. $\Delta_{\mathbf{w}\bar{\mathbf{w}}} \psi_i^\pi$ has an upper bound independent of π ; and

3. for any two states $w, \bar{w} \in W$, and any $\zeta > 0$, there exists $\varepsilon > 0$ such that, for all π ε -close to ρ , $|\Delta_{w\mathbf{1}, \bar{w}\mathbf{1}} \psi_i^\pi - \Delta_{w\bar{w}} \phi_i| < \zeta$.

This lemma implies that an inequality similar to (11) holds.

Lemma 8. *If a connected finite public profile is patiently strict, then for ε small, and π ε -close to ρ ,*

$$\sum_{\mathbf{w}' \neq \bar{w}\mathbf{1}} \Delta_{\mathbf{w}', \bar{w}\mathbf{1}} \psi_i^\pi \times (\tilde{q}_{w\mathbf{1}, \mathbf{w}'}^\pi - \tilde{q}_{w\mathbf{1}, \mathbf{w}'}^\pi(a_i)) > \hat{u}_i(w, a_i) - \hat{u}_i(w), \quad (12)$$

where \bar{w} is any state.

Now, the value player i assigns to being in state w , when she has beliefs β_i over the private states of her opponents, is

$$V_i^\pi(w; \beta_i) = \sum_{\mathbf{w}_{-i}} \psi_i^\pi(w, \mathbf{w}_{-i}) \beta_i(\mathbf{w}_{-i}),$$

and her incentive constraint in private state w is given by, for all $a_i \neq d_i(w)$,

$$V_i^\pi(w; \beta_i) \geq \sum_{\mathbf{w}_{-i}} \left\{ (1 - \delta) \tilde{u}_i(\mathbf{w}_{-i}, a_i) + \delta \sum_{\mathbf{w}'} \psi_i^\pi(\mathbf{w}') \tilde{q}_{w\mathbf{w}_{-i}, \mathbf{w}'}^\pi(a_i) \right\} \beta_i(\mathbf{w}_{-i}).$$

If β_i assigns probability close to 1 to the vector $w\mathbf{1}$, this inequality is implied by

$$\psi_i^\pi(w\mathbf{1}) > (1 - \delta) \tilde{u}_i(w\mathbf{1}, a_i) + \delta \sum_{\mathbf{w}'} \psi_i^\pi(\mathbf{w}') \tilde{q}_{w\mathbf{1}, \mathbf{w}'}^\pi(a_i). \quad (13)$$

Substituting for $\psi_i^\pi(w\mathbf{1})$ yields

$$\begin{aligned} \delta \sum_{\mathbf{w}'} \psi_i^\pi(\mathbf{w}') (\tilde{q}_{w\mathbf{1}, \mathbf{w}'}^\pi - \tilde{q}_{w\mathbf{1}, \mathbf{w}'}^\pi(a_i)) &> (1 - \delta) (\tilde{u}_i(w\mathbf{1}, a_i) - \tilde{u}_i(w\mathbf{1})) \\ &= (1 - \delta) (\hat{u}_i(w, a_i) - \hat{u}_i(w)). \end{aligned}$$

For any state $\bar{w} \in W$, (13) is equivalent to

$$\delta \sum_{\mathbf{w}' \neq \bar{w}\mathbf{1}} (\psi_i^\pi(\mathbf{w}') - \psi_i^\pi(\bar{w}\mathbf{1})) (\tilde{q}_{w\mathbf{1}, \mathbf{w}'}^\pi - \tilde{q}_{w\mathbf{1}, \mathbf{w}'}^\pi(a_i)) > (1 - \delta) (\hat{u}_i(w, a_i) - \hat{u}_i(w)).$$

Dividing by $(1 - \delta)$ and taking limits then yields (12). Thus, (12) implies that, if player i assigns a probability close to 1 to all her opponents being in the same private state as herself, the incentive constraint for i at that private state holds, and so (since there are only a finite number of incentive constraints) the theorem is proved.

7. An Application to Folk Theorems

A natural question is whether some form of the folk theorem holds for games with almost-public monitoring. As a corollary of our earlier results, we find that, if the monitoring is *also* sufficiently accurate, then a pure action version of the folk theorem holds *in general*.

Fix a pure action profile $a \in A$ that is individually rational in the stage game $g : A \rightarrow \mathfrak{R}^N$. In repeated games with perfect monitoring, players observe the action profile of all previous periods. The folk theorem asserts that, under a dimensionality condition, there is a discount factor δ' such that for each $\delta \geq \delta'$, there is a subgame perfect equilibrium of the repeated game with perfect monitoring in which a is played in every period. Since this equilibrium can be chosen to have finite memory (see, for example, the profile in Osborne and Rubinstein [26, Proposition 151.1]), we immediately have the following result:

Fix a discount factor $\delta > \delta'$ such that, for every history $h^t \in A^t$, the continuation value to any player from following the profile is strictly larger than that from deviating in period t and then following the profile thereafter. Say that a public monitoring technology (Y, ρ) is η -perfect if $Y = A$ and $\rho(a|a) > 1 - \eta$. There then exists $\eta' > 0$ such that if (Y, ρ) is η' -perfect, then the profile is a strict public equilibrium of the game with public monitoring (the arguments are almost the same as the proofs of Lemma 3 and Theorem 1). Since the public profile has finite memory, we then have (from Theorems 1 and 2) a bound on ε (that depends on ρ and δ) so that if the private monitoring technology is ε -close to ρ , the public profile induces an equilibrium in the game with private monitoring and the equilibrium payoffs are close to $g(a)$.

This is a weak result in the sense that, first, η depends on the discount factor and $\eta \rightarrow 0$ as $\delta \rightarrow 1$, and second, even if η was independent of δ , $\varepsilon \rightarrow 0$ as $\delta \rightarrow 1$.¹⁶ The remainder of this section is concerned with obtaining bounds on η and ε that are independent of δ . Since it is crucial that when a player observes the private signal y , she assigns a sufficiently high probability to her opponents all having observed the same signal, the bound on ε must depend on ρ (see Lemma 2).

In order to apply the techniques of Section 6, we first modify the profile used for the perfect monitoring folk theorem so that in η -perfect public monitoring games, it is patiently strict.¹⁷ Let \underline{g}_i denote player i 's pure strategy minmax

¹⁶Sekiguchi [29] proved such a result for efficient outcomes in the repeated prisoner's dilemma.

¹⁷While the profile from Osborne and Rubinstein [26, Proposition 151.1] in η -perfect games of public monitoring is finite and connected, the Markov chain on W is *not* ergodic for games with perfect monitoring. Since the Markov chain is ergodic for games with public monitoring, the incentive properties of the profile (in terms of strict patience) may differ between perfect and public monitoring. Property 4 in Theorem 4 implies that, even when the monitoring is perfect,

payoff, \underline{a}_i^i the action profile that minmaxes player i , and \underline{a}_i^i a myopic best response for i to \underline{a}_{-i}^i (so that $\underline{g}_i = g_i(\underline{a}^i)$). An action profile a is strictly individually rational if $g_i(a) > \underline{g}_i$. We use the version of the folk theorem given in Osborne and Rubinstein [26, Proposition 151.1]. We do not know if a similar result holds for the mixed action version, with unobservable mixing. The proofs of Theorem 4 and its corollary are in the Appendix.

Definition 5. *The action $a^* \in A$ satisfies the perfect folk theorem condition if it is strictly individually rational, and there exists N strictly individually rational action profiles, $\{a(i) : i \in N\}$, such that for all $i \in N$, $g_i(a^*) > g_i(a(i))$ and $g_i(a(j)) > g_i(a(i))$ for all $j \neq i$.*

Theorem 4. *(Perfect Monitoring) Suppose a^* satisfies the perfect folk theorem condition. Then there exists $L < \infty$ and $\underline{\delta} < 1$, such that for all $\delta \in (\underline{\delta}, 1)$, there is a subgame perfect equilibrium of the δ -discounted infinitely repeated game with perfect monitoring such that*

1. on the equilibrium path, a^* is played in every period;
2. for every history $h^t \in A^t$, the continuation value from following the profile is strictly larger than that from deviating in period t and then following the profile thereafter;
3. behavior in period t only depends on the action profiles of the last $\min\{L, t\}$ periods; and
4. after any history $h^t \in A^t$, under the profile, play returns to a^* in every period after L periods.

Moreover, the equilibrium strategy profile can be chosen independent of $\delta \in (\underline{\delta}, 1)$.

If the public monitoring technology has as signal space $Y = A$, then any profile of the repeated game with perfect monitoring also describes a public profile of the repeated game with public monitoring. As a corollary to Theorem 4, we have:

Corollary 2. *(Imperfect Public Monitoring) Fix a stage game $g : A \rightarrow \mathbb{R}^N$. Suppose a^* satisfies the perfect folk theorem condition. Let s denote the subgame perfect equilibrium profile described in Theorem 4. There exists $\underline{\delta} < 1$ and $\eta > 0$ such that if the public monitoring technology is η -perfect, then for any $\delta \in (\underline{\delta}, 1)$, the profile s is a public equilibrium of the δ -repeated game with public monitoring. Moreover, the profile is strictly patient.*

the Markov chain on W is ergodic.

This corollary, with Theorems 2 and 3, yields:

Theorem 5. (*Private Monitoring*) Fix a stage game $g : A \rightarrow \mathfrak{R}^N$. Suppose a^* satisfies the perfect folk theorem condition. For all $\nu > 0$, there exists $\delta' < 1$ and $\eta > 0$ such that for all η -perfect public monitoring technologies (Y, ρ) , there exists $\varepsilon > 0$ such that for all private monitoring distributions, π , ε -close to ρ , for all $\delta \in (\delta', 1)$, there is an equilibrium of the repeated game with private monitoring whose equilibrium payoff is within ν of $g(a^*)$.

8. Expanding the Set of Private Signals

In this section, we allow for a broader set of signals. The game with public monitoring is as described in Section 3. In a game with private monitoring in this section, each player i observes a signal $\omega_i \in \Omega_i$ (write ω and Ω for the vector of signals and set of signal profiles). As before, $\pi(\omega|a)$ is the probability of signal profile ω given action profile a . Write $\pi(\omega_{-i}|(a, \omega_i))$ for the implied conditional probability.

Definition 6. The private monitoring technology (Ω, π) is ε -close to the public monitoring technology (Y, ρ) if there exist functions $f_i : \Omega_i \rightarrow Y \cup \{\emptyset\}$ such that the following two properties hold:

1. for each $a \in A$ and $y \in Y$,

$$|\pi(\{\omega : f_i(\omega_i) = y \text{ for each } i\} | a) - \rho(y|a)| \leq \varepsilon,$$

and

2. for each $a \in A$, $y \in Y$, and $\omega_i \in f_i^{-1}(y)$,

$$\pi(\{\omega_{-i} : f_j(\omega_j) = y \text{ for each } j \neq i\} | (a, \omega_i)) \geq 1 - \varepsilon.$$

Note that some private signals may not be associated with any public signals: It is possible that there is a signal ω_i satisfying $f_i(\omega_i) = \emptyset$; this signal then may contain no information about the signals observed by the other players. It is worth noting that for the case $\Omega_i = Y$, the first condition implies the second (Lemma 2), and coincides with the notion of closeness used in Section 3.

The condition of ε -closeness can be restated as follows. Recall that an event is p -evident if, whenever it is true, everyone believes it with probability at least p . We say that (Ω, π) is ε -close to (Y, ρ) if for each public signal y , there is a set of private signal profiles, $f^{-1}(y)$, such that $f^{-1}(y)$ is $(1 - \varepsilon)$ -evident (contingent

on any action profile) and the probability of $f^{-1}(y)$ is within ε of the probability of the public signal y (contingent on any action profile).

An obvious way to proceed to the case of general private signals is to apply the argument from the previous section to the case of close monitoring technologies. As we discuss below, approximating arbitrary public equilibria is difficult, because the public state in period t is determined, in principle, by the entire history h^t . Approximating finite memory public profiles is significantly simpler.

Fix a strict public equilibrium with finite memory, (W, w^1, σ, d) . Since the profile has memory L , for some L , recall that we can take $W = (Y \cup \{*\})^L$, $w^1 = (*, \dots, *)$, and $\sigma(y, (y^L, \dots, y^2, y^1)) = (y, y^L, \dots, y^2)$. Fix a private monitoring technology (Ω, π) with associated signaling functions f_i that is ε -close to (Y, ρ) . Following Monderer and Samet [24], we first consider a *constrained game* where behavior after “exceptional signals” is arbitrarily fixed. Define the set of “exceptional” private histories, $H_i^e = \{h_i^t : \omega_i^\tau \in f_i^{-1}(\emptyset), \text{ some } \tau \text{ satisfying } t - L \leq \tau \leq t - 1\}$. This is the set of private histories for which in any of the last L periods, a private signal ω_i^τ satisfying $f_i(\omega_i^\tau) = \emptyset$ is observed. We fix *arbitrarily* player i 's action after any private history $h_i^t \in H_i^e$. For any private history that is not exceptional, each of the last L observations of the private signal can be associated with a public signal by the function f_i . Denote by $w_i(h_i^t)$ the private state so obtained. That is,

$$w_i(h_i^t) = \left(f_i(\omega_i^{t-1}), \dots, f_i(\omega_i^{t-L}) \right),$$

for all $h_i^t \notin H_i^e$. We are then left with a game in which in period $t \geq 2$ player i only chooses an action after a signal ω_i^{t-1} yields a private history not in H_i^e . We claim that for ε sufficiently small, the profile $(\hat{s}_1, \dots, \hat{s}_N)$ is an equilibrium of this constrained game, where \hat{s}_i is the strategy for player i :

$$\hat{s}_i^t(h_i^t) = \begin{cases} d_i(w_i^1), & \text{if } t = 1, \\ d_i(w_i(h_i^t)), & \text{if } t > 1 \text{ and } h_i^t \notin H_i^e. \end{cases}$$

But this follows from arguments almost identical to that in the proof of Theorem 2: since a player's behavior depends only on the last L signals, for small ε , after observing a history $h_i^t \notin H_i^e$, player i assigns a high probability to player j observing a signal that leads to the same private state. The crucial point is that for ε small, the specification of behavior after signals ω_i satisfying $f_i(\omega_i) = \emptyset$ is irrelevant for behavior at signals ω_i satisfying $f_i(\omega_i) \in Y$. It remains to specify optimal behavior after signals ω_i satisfying $f_i(\omega_i) = \emptyset$. So, consider a new constrained game where player i is required to follow \hat{s}_i where possible. This constrained game has an equilibrium, and so by construction, we thus have an equilibrium of the unconstrained game. We have thus proved:

Theorem 6. *Suppose a finite memory public profile is a strict equilibrium. For all $\kappa > 0$, there exists $\bar{\varepsilon}$ such that, if (Ω, π) is ε -close to (Y, ρ) for $\varepsilon \in (0, \bar{\varepsilon})$, then the induced private profile is an equilibrium of the game with private monitoring technology (Ω, π) , and the expected payoff in that equilibrium is within κ of the public equilibrium payoff.*

We could similarly extend our earlier results on patiently strict connected finite public profiles (Theorem 3) and on the almost public almost perfect folk theorem (Theorem 5) to this more general notion of nearby private monitoring technologies.

We now describe an example that suggests that even under a significant strengthening of ε -closeness, general positive results will be difficult to obtain. Say that (Ω, π) is *strongly* ε -close to (Y, ρ) if (Ω, π) is ε -close to (Y, ρ) and the signaling functions map into Y . That is, *every* private signal can be associated with some public signal.

Example 6. *We return to the public monitoring technology and profile of Example 3. The private monitoring technology is described as follows: The signal spaces are $\Omega_1 = \{\underline{y}, \bar{y}', \bar{y}''\}$ and $\Omega_2 = \{\underline{y}, \bar{y}\}$. With probability ε' , there is a uniform draw from $\Omega_1 \times \Omega_2$, and with probability $1 - \varepsilon'$, there is draw from the joint distributions*

$$\begin{array}{c}
 CC \\
 \begin{array}{c}
 \underline{y} \\
 \bar{y}' \\
 \bar{y}''
 \end{array}
 \end{array}
 \begin{array}{cc}
 & \underline{y} & \bar{y} \\
 \begin{array}{c}
 \underline{y} \\
 \bar{y}' \\
 \bar{y}''
 \end{array}
 & \begin{array}{|c|c|}
 \hline
 1 - p' - p'' & 0 \\
 \hline
 0 & p' \\
 \hline
 0 & p'' \\
 \hline
 \end{array}
 \end{array}$$

and

$$\begin{array}{c}
 a_1 a_2 \\
 \begin{array}{c}
 \underline{y} \\
 \bar{y}' \\
 \bar{y}''
 \end{array}
 \end{array}
 \begin{array}{cc}
 & \underline{y} & \bar{y} \\
 \begin{array}{c}
 \underline{y} \\
 \bar{y}' \\
 \bar{y}''
 \end{array}
 & \begin{array}{|c|c|}
 \hline
 1 - q' - q'' & 0 \\
 \hline
 0 & q' \\
 \hline
 0 & q'' \\
 \hline
 \end{array}
 \end{array}$$

for $a_1 a_2 \neq CC$. The private monitoring technology is clearly strongly ε -close to the public monitoring technology (Y, ρ) of Example 1 for some ε , with $p = p' + p''$ and $q = q' + q''$ (with the obvious signaling functions $f_1(\bar{y}') = f_1(\bar{y}'') = f_2(\bar{y}) = \bar{y}$ and $f_1(\underline{y}) = f_2(\underline{y}) = \underline{y}$). Moreover, ε can be made arbitrarily small by making ε' small. We now argue that by choosing the probabilities appropriately, the private profile implied by the signaling functions is not an equilibrium. Consider first the private monitoring technology $(Y^2, \hat{\pi})$, where player 1 cannot distinguish between

\bar{y}' and \bar{y}'' . As we saw Example 3, the private profile will be an equilibrium in that case for ε small, as long as $p < 1 - q$. A conjecture is that this is a sufficient condition when the private monitoring technology is in fact (Ω, π) . The problem, of course, is that the information content of \bar{y}' in the private monitoring technology (Ω, π) may be quite different from \bar{y} in the private monitoring technology $(Y^2, \hat{\pi})$. For example, consider the impact of the sequence $\underline{y}\bar{y}'\underline{y}\bar{y}'$ on player 1's beliefs over player 2's private state. As in Example 3, the critical issue is whether $\Pr\{\underline{y}\bar{y}'|CC\}$ is larger or smaller than $\Pr\{\underline{y}\bar{y}'|CC\}$. Since we are considering ε small, it is enough to compare these when $\varepsilon' = 0$. Player 1 will eventually assign large probability to player 2 being in private state \underline{w} when 1 is in private state \bar{w} if $p'(1 - p' - p'') < q'(1 - q' - q'')$. For example, if $p' = 1/12$, $p'' = 1/2$, $q' = q'' = 1/8$, we have $p = p' + p'' = 7/12$ and $q = 1/4$. Thus, under the private monitoring technology $(Y^2, \hat{\pi})$, for sufficiently small ε , the three state private profile will be an equilibrium. On the other hand, with the private monitoring technology (Ω, π) , the private profile is not an equilibrium.

A. Proofs for Section 6. The Case of Arbitrarily Patient Players

We need the following standard result (see, for example, Stokey and Lucas [31, Theorem 11.4]). If (Z, R) is a finite-state Markov chain with state space Z and transition matrix R , R^n is the matrix of n -step transition probabilities and $r_{ij}^{(n)}$ is the ij -th element of R^n . For a vector $x \in \mathfrak{R}^\ell$, define $\|x\|_\Delta \equiv \sum_j |x_j|$.

Lemma A. *Suppose (Z, R) is a finite state Markov chain. Let $\eta_j^{(n)} = \min_i r_{ij}^{(n)}$ and $\eta^{(n)} = \sum_j \eta_j^{(n)}$. Suppose that there exists ℓ such that $\eta^{(\ell)} > 0$. Then, (Z, R) has a unique stationary distribution p^* and, for all $p \in \Delta(Z)$,*

$$\left\| pR^{k\ell} - p^* \right\|_\Delta \leq 2 \left(1 - \eta^{(\ell)} \right)^k.$$

Proof of Lemma 5. Let Θ denote the matrix of transition probabilities on W induced by the public profile (W is a finite set by assumption). The ww' -th element is $\theta_{ww'}(d(w)) = \hat{\theta}_{ww'}$. If $\hat{u}_i \in \mathfrak{R}^W$ and $\phi_i \in \mathfrak{R}^W$ are the vectors of stage payoffs and continuation values for player i associated with the states, then

$$\phi_i = (1 - \delta) \hat{u}_i + \delta \Theta \phi_i.$$

Solving for ϕ_i yields

$$\begin{aligned}\phi_i &= (1 - \delta) (I_W - \delta\Theta)^{-1} \hat{u}_i \\ &= (1 - \delta) \sum_{t=0}^{\infty} (\delta\Theta)^t \hat{u}_i,\end{aligned}$$

where I_W is the $|W|$ -dimensional identity matrix. Let e_w denote the w -th unit vector (i.e., the vector with 1 in the w -th coordinate and 0 elsewhere). Then,

$$\phi_i(w) - \phi_i(\bar{w}) = (1 - \delta) \sum_{t=0}^{\infty} (e_w - e_{\bar{w}}) (\delta\Theta)^t \hat{u}_i = (1 - \delta) \sum_{t=0}^{\infty} \delta^t (e_w \Theta^t - e_{\bar{w}} \Theta^t) \hat{u}_i.$$

Because the public profile is connected, for any two distributions on W , α and α' , $\|\alpha\Theta^t - \alpha'\Theta^t\| \rightarrow 0$ at an exponential rate (Lemmas 4 and A). This implies that $\sum_{t=0}^{\infty} (e_w \Theta^t - e_{\bar{w}} \Theta^t) \hat{u}_i$ is absolutely convergent, and so $(\phi_i(w) - \phi_i(\bar{w})) / (1 - \delta)$ has a finite limit as $\delta \rightarrow 1$. ■

Proof of Lemma 7. The proof of the first assertion is identical to that of Lemma 5.

Since the public profile is finite and connected, for the purposes of applying Lemma A, we can take $\ell = n$, independent of π , where (y^1, \dots, y^n) is the finite sequence of signals from Lemma 4. Moreover, there exists $\varepsilon > 0$ such that for all π ε -close to ρ ,

$$\sum_{\mathbf{w}'} \min_{\mathbf{w}} q_{\mathbf{w}\mathbf{w}'}^{\pi, (n)} > \frac{1}{2} \times \sum_{w'} \min_w \theta_{ww'}^{(n)} \equiv \eta^*.$$

This gives a bound on the rate at which $\alpha(Q^\pi)^t$ converges to α^π , the stationary distribution of (Z, Q^π) , independent of π and $\alpha \in \Delta(Z)$. This then implies the second assertion.

Now,

$$\Delta_{w\bar{w}}\phi_i = \sum_{t=0}^{\infty} (e_w \Theta^t - e_{\bar{w}} \Theta^t) \hat{u}_i$$

and

$$\Delta_{w\mathbf{1}, \bar{w}\mathbf{1}}\psi_i^\pi = \sum_{t=0}^{\infty} (e_w \mathbf{1} (Q^\pi)^t - e_{\bar{w}} \mathbf{1} (Q^\pi)^t) \tilde{u}_i.$$

Fix $\zeta > 0$. There exists T such that

$$\left| \sum_{t=0}^T (e_w \Theta^t - e_{\bar{w}} \Theta^t) \hat{u}_i - \Delta_{w\bar{w}} \phi_i \right| < \zeta/3$$

and, for all π ε -close to ρ ,

$$\left| \sum_{t=0}^T (e_{w\mathbf{1}} (Q^\pi)^t - e_{\bar{w}\mathbf{1}} (Q^\pi)^t) \tilde{u}_i - \Delta_{w\mathbf{1}, \bar{w}\mathbf{1}} \psi_i^\pi \right| < \zeta/3.$$

Order the states in (Z, Q^π) so that the first $|W|$ states are the states in which all players' private states are in agreement. Then, we can write the transition matrix as

$$Q^\pi = \begin{bmatrix} Q_{11}^\pi & Q_{12}^\pi \\ Q_{21}^\pi & Q_{22}^\pi \end{bmatrix},$$

and so $[I_W : 0] \tilde{u}_i = \hat{u}_i$. As π approaches ρ , $Q_{11}^\pi \rightarrow \Theta$, $Q_{12}^\pi \rightarrow 0$, and $Q_{22}^\pi \rightarrow 0$.

Now,

$$[(Q^\pi)^2]_{11} = (Q_{11}^\pi)^2 + Q_{12}^\pi Q_{21}^\pi$$

and

$$[(Q^\pi)^2]_{12} = Q_{11}^\pi Q_{12}^\pi + Q_{12}^\pi Q_{22}^\pi,$$

and, in general,

$$[(Q^\pi)^t]_{11} = (Q_{11}^\pi)^t + Q_{12}^\pi [(Q^\pi)^{t-1}]_{21}$$

and

$$[(Q^\pi)^t]_{12} = Q_{11}^\pi [(Q^\pi)^{t-1}]_{12} + Q_{12}^\pi [(Q^\pi)^{t-1}]_{22}.$$

Thus, for all t , $[(Q^\pi)^t]_{11} \rightarrow \Theta^t$ and $[(Q^\pi)^t]_{12} \rightarrow 0$, as π approaches ρ . Hence, there exists $\varepsilon' > 0$ such that for all $t \leq T$, if π is ε' -close to ρ ,

$$\left| \sum_{t=0}^T (e_w \Theta^t - e_{\bar{w}} \Theta^t) \hat{u}_i - \sum_{t=0}^T (e_{w\mathbf{1}} (Q^\pi)^t - e_{\bar{w}\mathbf{1}} (Q^\pi)^t) \tilde{u}_i \right| < \zeta/3.$$

So, for $\varepsilon'' = \min \{\varepsilon, \varepsilon'\}$, if π is ε'' -close to ρ ,

$$|\Delta_{w\bar{w}}\phi_i - \Delta_{w\mathbf{1},\bar{w}\mathbf{1}}\psi_i^\pi| < \zeta.$$

■

Proof of Lemma 8. Let

$$\zeta = \frac{1}{2} \left\{ \sum_{w \neq \bar{w}} \Delta_{w\bar{w}}\phi_i \times \left(\hat{\theta}_{ww} - \hat{\theta}_{ww}(a_i) \right) - \left(\hat{u}_i(w, a_i) - \hat{u}_i(w) \right) \right\}.$$

Since the public profile is patiently strict, $\zeta > 0$.

The left hand side of (12) is

$$\sum_{w \neq \bar{w}} \Delta_{w\mathbf{1},\bar{w}\mathbf{1}}\psi_i^\pi \times \left(\tilde{q}_{w\mathbf{1},w\mathbf{1}}^\pi - \tilde{q}_{w\mathbf{1},w\mathbf{1}}^\pi(a_i) \right) + \sum_{\substack{w' \neq w\mathbf{1}, \\ w \in W}} \Delta_{w',\bar{w}\mathbf{1}}\psi_i^\pi \times \left(\tilde{q}_{w\mathbf{1},w'}^\pi - \tilde{q}_{w\mathbf{1},w'}^\pi(a_i) \right)$$

and, by Lemma 7, there exists $\varepsilon'' > 0$ such that for π ε'' -close to ρ ,

$$\left| \sum_{\substack{w' \neq w\mathbf{1}, \\ w \in W}} \Delta_{w',\bar{w}\mathbf{1}}\psi_i^\pi \times \left(\tilde{q}_{w\mathbf{1},w'}^\pi - \tilde{q}_{w\mathbf{1},w'}^\pi(a_i) \right) \right| < \zeta/2.$$

Moreover, again by Lemma 7, by choosing ε small, for π ε -close to ρ ,

$$\left| \sum_{w \neq \bar{w}} \Delta_{w\bar{w}}\phi_i \times \left(\hat{\theta}_{ww} - \hat{\theta}_{ww}(a_i) \right) - \sum_{w \neq \bar{w}} \Delta_{w\mathbf{1},\bar{w}\mathbf{1}}\psi_i^\pi \times \left(\tilde{q}_{w\mathbf{1},w\mathbf{1}}^\pi - \tilde{q}_{w\mathbf{1},w\mathbf{1}}^\pi(a_i) \right) \right| < \zeta/2,$$

and so

$$\begin{aligned} \sum_{w' \neq \bar{w}\mathbf{1}} \Delta_{w',\bar{w}\mathbf{1}}\psi_i^\pi \times \left(\tilde{q}_{w\mathbf{1},w'}^\pi - \tilde{q}_{w\mathbf{1},w'}^\pi(a_i) \right) &> \sum_{w \neq \bar{w}} \Delta_{w\mathbf{1},\bar{w}\mathbf{1}}\psi_i^\pi \times \left(\tilde{q}_{w\mathbf{1},w\mathbf{1}}^\pi - \tilde{q}_{w\mathbf{1},w\mathbf{1}}^\pi(a_i) \right) - \zeta/2 \\ &> \sum_{w \neq \bar{w}} \Delta_{w\bar{w}}\phi_i \times \left(\hat{\theta}_{ww} - \hat{\theta}_{ww}(a_i) \right) - \zeta \\ &> \hat{u}_i(w, a_i) - \hat{u}_i(w), \end{aligned}$$

which is the desired inequality (12).

■

B. Proofs for Section 7. An Application to Folk Theorems

Proof of Theorem 4. While the profile specified in the proof of Proposition 151.1 of Osborne and Rubinstein [26] satisfies the first three properties, it does not satisfy the requirement that play eventually return to a^* . The following modification does. We first describe the profile presented in Osborne and Rubinstein [26]. The profile has three types of phases, $C(0)$, $C(j)$, and $P(j)$. Player i chooses a_i^* in phase $C(0)$, $a_i(j)$ in phase $C(j)$, and \underline{a}_i^j in phase $P(j)$. Play starts in phase $C(0)$, and remains there unless there is a unilateral deviation, by player j , say. After such a deviation, the profile switches to phase $P(j)$ for L^* periods, after which play switches to $C(j)$, and remains there. If there is a unilateral deviation in $P(j)$ or $C(j)$ by player k , say, the profile switches to $P(k)$ for L^* periods, and then to $C(k)$, and remains there. Now modify the profile so that once the profile switches to $C(j)$, it stays in $C(j)$ for L^{**} periods, for an L^{**} to be determined, after which it reverts to $C(0)$.

For notational simplicity, set $a(0) = a^*$. First choose L^* large enough so that, for all $j \in N \cup \{0\}$ (where $M \equiv \max_{i,a} |g_i(a)|$),¹⁸

$$M - g_i(a(j)) < L^* \left(g_i(a^*) - \underline{g}_i \right). \quad (\text{B.1})$$

Second, choose L^{**} sufficiently large so that, for all i ,

$$M - g_i(\underline{a}^j) + L^* \left(\underline{g}_i - \min \{ g_i(a^*), g_i(\underline{a}^j) \} \right) < L^{**} (g_i(a(j)) - g_i(a(i))). \quad (\text{B.2})$$

Each player has a strict incentive to follow the prescribed path when in phase $C(j)$ if, for all $\ell \in \{1, \dots, L^{**}\}$ (where ℓ is the number of periods remaining in phase $C(j)$),¹⁹

$$M + \sum_{k=2}^{L^*+1} \delta^{k-1} \underline{g}_i + \sum_{k=L^*+2}^{L^*+L^{**}+1} \delta^{k-1} g_i(a(i)) < \sum_{k=1}^{\ell} \delta^{k-1} g_i(a(j)) + \sum_{k=\ell+1}^{L^*+L^{**}+1} \delta^{k-1} g_i(a^*). \quad (\text{B.3})$$

Evaluating this inequality at $\delta = 1$ and rearranging yields

$$\begin{aligned} M - g_i(a(j)) &< (\ell - 1) (g_i(a(j)) - g_i(a(i))) + L^* \left(g_i(a^*) - \underline{g}_i \right) \\ &\quad + (L^{**} - (\ell - 1)) (g_i(a^*) - g_i(a(i))), \end{aligned}$$

¹⁸Osborne and Rubinstein [26, Proposition 151.1] fix L^* large enough so that $M - g_i(a(j)) < L^* \left(g_i(a(j)) - \underline{g}_i \right)$, rather than as in (B.1), because in their profile, after a deviation play never returns to a^* .

¹⁹If $j = 0$, then the value of ℓ is irrelevant.

which is implied by (B.1), since $g_i(a(j)) > g_i(a(i))$ and $g_i(a^*) > g_i(a(i))$. Thus, there exists δ' such that for $\delta \in (\delta', 1)$, and any $\ell \in \{1, \dots, L^{**}\}$, (B.3) holds.

Each player has a strict incentive to follow the prescribed path when in phase $P(j)$ if, for all $\ell \in \{1, \dots, L^*\}$ (where ℓ is now the number of periods remaining in phase $P(j)$),

$$\begin{aligned} M + \sum_{k=2}^{L^*+1} \delta^{k-1} \underline{g}_i + \sum_{k=L^*+2}^{L^*+L^{**}+1} \delta^{k-1} g_i(a(i)) & \quad (B.4) \\ < \sum_{k=1}^{\ell} \delta^{k-1} g_i(\underline{a}^j) + \sum_{k=\ell+1}^{\ell+L^{**}} \delta^{k-1} g_i(a(j)) + \sum_{k=\ell+L^{**}+1}^{L^*+L^{**}+1} \delta^{k-1} g_i(a^*). \end{aligned}$$

Evaluating this inequality at $\delta = 1$ and rearranging yields

$$M - g_i(\underline{a}^j) + L^* \underline{g}_i - (\ell - 1) g_i(\underline{a}^j) - (L^* + 1 - \ell) g_i(a^*) < L^{**} (g_i(a(j)) - g_i(a(i))),$$

which is implied by (B.2). Thus, there exists δ'' such that for $\delta \in (\delta'', 1)$, and any $\ell \in \{1, \dots, L^*\}$, (B.4) holds.

The proof is completed by setting $L = L^* + L^{**}$ and $\underline{\delta} = \max\{\delta', \delta''\}$. By construction, all one-shot deviations are strictly suboptimal (the incentive constraints (B.3) and (B.4) hold strictly). ■

Proof of Corollary. Let (W, w^1, σ, d) be a finite state automaton description of the strategy profile from Theorem 4. Let $v_i : W \rightarrow \Re$ describe player i 's continuation values under this profile. Observe that, for all $w \in W$, $v_i(w) \rightarrow g_i(a^*)$ as $\delta \rightarrow 1$. Not surprisingly,

$$v_i = (1 - \delta) \hat{g}_i + \delta D v_i,$$

where $\hat{g}_i \in \Re^W$ is given by $\hat{g}_i(w) = g_i(d(w))$ and D is the transition matrix with ww' -th element given by

$$D_{ww'} = \begin{cases} 1, & \text{if } w' = \sigma(d(w); w), \\ 0, & \text{otherwise.} \end{cases}$$

We can view D as a degenerate stochastic matrix for the Markov chain (W, D) .

By construction, this Markov chain is ergodic. Now,

$$\begin{aligned} v_i(w) - v_i(w') &= (1 - \delta) \sum_{t=0}^{\infty} \delta^t (e_w D^t - e_{w'} D^t) \hat{g}_i \\ &= (1 - \delta) \sum_{t=0}^L \delta^t (e_w D^t - e_{w'} D^t) \hat{g}_i, \end{aligned}$$

and so

$$\Delta_{w'\bar{w}} v_i \equiv \lim_{\delta \rightarrow 1} (v_i(w') - v_i(\bar{w})) / (1 - \delta)$$

is well-defined and finite. Moreover, (B.1) and (B.2) imply that the profile is patiently strict: for all players i and states $w \in W$, for all $a_i \neq d_i(w)$,

$$\sum_{w' \neq \bar{w}} \Delta_{w'\bar{w}} v_i \times (D_{ww'} - D_{ww'}(a_i)) > \hat{u}_i(w, a_i) - \hat{u}_i(w),$$

where \bar{w} is any state, and

$$D_{ww'}(a_i) = \begin{cases} 1, & \text{if } w' = \sigma(d_{-i}(w), a_i; w), \\ 0, & \text{otherwise.} \end{cases}$$

The proof of Lemma 8 can then be used to show that the public profile in the η -perfect game of public monitoring is patiently strict for η small. ■

References

- [1] Dilip Abreu. Extremal equilibria of oligopolistic supergames. *Journal of Economic Theory*, 39:191–225, 1986.
- [2] Dilip Abreu, David Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58:1041–1063, 1990.
- [3] Illtae Ahn. Imperfect information repeated games with a single perfect observer. Mimeo, 1996.
- [4] Massimiliano Amarante. Discounted repeated games with imperfect private monitoring, 1998. M.I.T.

- [5] Masaki Aoyagi. A collusion mechanism in dynamic bertrand oligopoly under private observation of correlated demand shocks. Mimeo, University of Pittsburgh, 1997.
- [6] Elchanan Ben-Porath and Michael Kahneman. Communication in repeated games with private monitoring. *Journal of Economic Theory*, 70:281–297, 1996.
- [7] Jean-Pierre Benoit and Vijay Krishna. Finitely repeated games. *Econometrica*, 53:905–922, 1985.
- [8] V. Bhaskar. Sequential equilibria in the repeated prisoner’s dilemma with private monitoring. Mimeo, University of Essex, 1999.
- [9] V. Bhaskar and Eric van Damme. Moral hazard and private monitoring. Mimeo, University of St. Andrews and CentER Tilburg, 1997.
- [10] Olivier Compte. On failing to cooperate when monitoring is private. Mimeo, C.E.R.A.S.-E.N.P.C. and CEPRAMAP, 1996.
- [11] Olivier Compte. Communication in repeated games with imperfect private monitoring. *Econometrica*, 66:597–626, 1998.
- [12] Glenn Ellison. Cooperation in the prisoner’s dilemma with anonymous random matching. *Journal of Economic Theory*, 61:567–588, 1994.
- [13] Jeffrey C. Ely and Juuso Välimäki. A robust folk theorem for the prisoner’s dilemma. Mimeo, Northwestern University, 1999.
- [14] Drew Fudenberg and David K. Levine. An approximate folk theorem with imperfect private information. *Journal of Economic Theory*, 54:26–47, 1991.
- [15] Drew Fudenberg, David K. Levine, and Eric Maskin. The folk theorem with imperfect public information. *Econometrica*, 62:997–1040, 1994.
- [16] Edward Green and Robert Porter. Noncooperative collusion under imperfect price formation. *Econometrica*, 52:87–100, 1984.
- [17] Michihiro Kandori. Cooperation in finitely repeated games with imperfect private information. Mimeo, Princeton University, 1991.
- [18] Michihiro Kandori. Check your partners’ behavior by randomization: New efficiency results on repeated games with imperfect monitoring. Technical Report CIRJE-F-49, University of Tokyo, 1999.

- [19] Michihiro Kandori and Hitoshi Matsushima. Private observation, communication and collusion. *Econometrica*, 66:627–652, 1998.
- [20] Ehud Lehrer. On the equilibrium payoffs set of two player repeated games with imperfect monitoring. *International Journal of Game Theory*, 20:211–226, 1992.
- [21] George J. Mailath, Steven A. Matthews, and Tadashi Sekiguchi. Private strategies in repeated games with imperfect public monitoring. University of Pennsylvania, 1999.
- [22] George J. Mailath and Stephen Morris. Repeated games with imperfect private monitoring: Notes on a coordination perspective. CARESS Working Paper #98-07, University of Pennsylvania, 1998.
- [23] Hitoshi Matsushima. On the theory of repeated games with private information. Part I: Anti-folk theorem without communication. *Economic Letters*, 35:253–256, 1991.
- [24] Dov Monderer and Dov Samet. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1:170–190, 1989.
- [25] Ichiro Obara. Private strategy and efficiency: Repeated partnership game revisited. University of Pennsylvania, 1999.
- [26] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. The MIT Press, Cambridge, MA, 1994.
- [27] Michele Piccione. The repeated prisoner’s dilemma with imperfect private monitoring. Mimeo, University of Southampton, 1999.
- [28] Tadashi Sekiguchi. Efficiency in repeated prisoner’s dilemma with private monitoring. *Journal of Economic Theory*, 76:345–361, 1997.
- [29] Tadashi Sekiguchi. Robustness of efficient equilibria in repeated games with imperfect private monitoring. Mimeo, 1998.
- [30] Hyun Shin and Timothy Williamson. How much common belief is necessary for a convention? *Games and Economic Behavior*, 13:252–268, 1996.
- [31] Nancy Stokey and Robert E. Lucas. *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, MA, 1989.