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
July 13, 2001

Running Head: Almost-Public Monitoring

Corresponding author:
George J. Mailath
phone:(215) 898-7908
fax: (215) 573-2057

*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 Drew Fudenberg, Michihiro Kandori, 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

Abstract

In repeated games with imperfect public monitoring, players can use public signals to coordinate their behavior, and thus support cooperative outcomes. But with private monitoring, such coordination may no longer be possible. Even though grim trigger is a perfect public equilibrium (PPE) in games with public monitoring, it often fails to be an equilibrium in arbitrarily close games with private monitoring. If a PPE 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. *Journal of Economic Literature* Classification Numbers: C72, C73.

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 [17] 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 [16]). Perfect public equilibria (PPE) 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 [16]). In particular, any PPE can be described by an action profile for the current period and continuation values that are necessarily PPE 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 need 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 these circumstances, there is no guarantee that non-static Nash behavior can be supported in equilibrium. While a player may be almost sure another has deviated and would want to punish if he believed that others were punishing, he cannot be sure that others are also almost sure that someone has deviated. With private signals, unlike public signals, there is a lack of 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?[3]

We approach these questions as follows: Fix a repeated game with imperfect public monitoring and a strict pure strategy PPE of that game. Consider the simplest perturbation of the game to allow private monitoring. Fix the set of public monitoring signals. Let each player privately observe a (perhaps different) signal from that set. The private-monitoring distribution is said to be close to the public-monitoring distribution if the probability that all players observe the same signal, under the private-monitoring distribution, is close to the probability of that signal under the public-monitoring distribution. 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. Recall that all PPE 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 use the same state space, transition function and decision rules to summarize behavior in the private monitoring game. Each player will now have a *private state*, and the transition function and decision rules define a Markov process on vectors of private states. This representation is sufficient to describe behavior under the given strategies, but (with private monitoring) it is *not* sufficient to check if the strategies are optimal. It is also necessary to know how each player's beliefs over the private states of other players evolve. A sufficient condition for a strict equilibrium 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 (Theorem 4.1). 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 be *connected*: there is always a sequence of public signals that leads to the same final state independent of the initial state (Section 3.2 analyzes an example failing connectedness). However, connectedness is not sufficient, since the grim-trigger strategy profile in the repeated prisoners' dilemma with public monitoring is connected and yet is never a sequential equilibrium in any close-by game with private monitoring (Section 3.3).[4] In fact, for "intuitive" public monitoring distributions (where cooperation always increases the probability of observing the "good" signal), grim trigger is not even Nash. One sufficient condition is that strategies depend only on a finite history of play (Theorem 4.3).

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. 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.

In the results discussed so far, the bound on the distance between the private-monitoring and the public-monitoring distribution depends, in general, on the discount rate, with the bound converging to zero as the discount factor approaches one. We also provide results that hold uniformly over discount rates close to one. A connected finite public strategy profile is said to be *patiently strict* if it is a uniformly strict PPE for all discount rates close to one. In this case, approximate common knowledge of histories is enough for there to exist $\varepsilon > 0$, such that for all discount rates close to one, the strategy profile is an equilibrium of any ε-close private monitoring game (Theorem 5.1). This result implies a pure-action folk theorem for repeated games with *almost-public almost-perfect* monitoring (Theorem 6.2). 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 distribution is close to some almost-perfect public-monitoring distribution. 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.

This paper describes techniques and results for determining whether a fixed pure strategy public profile remains an equilibrium when translated to a private (but almost-public) monitoring setting. These techniques do not directly answer the traditional question of interest: what payoffs are achievable in equilibrium (with perhaps very complicated equilibrium strategies). A growing and important literature does examine this question, with a focus on the infinitely-repeated prisoners' dilemma with almost-perfect but not necessarily almost-public monitoring.[5] One branch of this literature (Sekiguchi [27], Compte [12]) shares with our work a focus on the evolution of beliefs; we discuss this literature in Section 3.4. The other branch (initiated by Piccione [26] and extended by Ely and Välimäki [14]) follows a completely different approach. In this approach, it is not necessary to study the evolution of beliefs, because each player is indifferent over his own actions after all private histories (he is not indifferent about the behavior of his opponent). This indifference is obtained by requiring players to randomize after every history (and since they are indifferent, they will be willing to do so).[6] In these equilibria, there is no need for players to coordinate their behavior at histories where they lack common knowledge. It is worth emphasizing that the role of randomization in these equilibria is very different from that in Sekiguchi [27], which we describe in Section 3.4.

We consider only the case of private monitoring close to full-support public 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]) as well as 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 to prove stronger results (Fudenberg and Levine [15] and Lehrer [21]).

The paper is organized as follows. Section 2 introduces repeated games with public monitoring and close-by private-monitoring distributions. Section 3 focuses on the repeated prisoners' dilemma, describing some strict public equilibria that are robust to private monitoring and some that are not, including grim trigger. Section 4 contains the results on approximating arbitrary strict public equilibria for fixed discount factors. Section 5 presents the high discounting version of our results and Section 6 applies this result to derive a folk theorem for repeated games with almost-public almost perfect monitoring.

2. Almost-Public Monitoring

We begin our investigation of the extent to which games with public monitoring can be approximated by games with private monitoring by describing the game with public monitoring. The finite action set for player $i \in \{1 \ldots, N\}$ is A_i. The public signal, denoted y, 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 restrict attention to *full-support* public monitoring (this plays an important role, see Lemma 1 below):

Assumption. *$\rho(y|a) > 0$ for all $y \in Y$ and all $a \in A$.*

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 [16], 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* (*PPE*) is a profile of public strategies that, conditional on any public history h^t, specifies a Nash equilibrium for the repeated game. Under full-support public monitoring, every public history arises with positive probability, and so every Nash equilibrium in public strategies is a PPE. Henceforth, equilibrium for the game with public monitoring means Nash equilibrium in public strategies (or, equivalently, PPE).

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 $\sigma : Y \times W \to W$, and a collection of decision rules, $d_i : W \to A_i$.[8] In the first period, player i chooses action $a_i^1 = d_i\left(w^1\right)$. The vector of actions, a^1, then generates a signal y^1 according to the distribution $\rho\left(\cdot|a^1\right)$. In the second period, player i chooses the action $a_i^2 = d_i\left(w^2\right)$, where $w^2 = \sigma\left(y^1, w^1\right)$, and so on.[9] 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 a mapping $\phi : W \to \Re^N$, so that the following is true (Abreu, Pearce, and Stacchetti [2]): Define a function $g : A \times W \to \Re^N$ by

$$g(a; w) \equiv (1-\delta)u(a) + \delta \sum_y \phi\left(\sigma\left(y; w\right)\right) \rho(y|a).$$

Then, for all $w \in W$, the action profile $(d_1(w), \ldots, d_N(w)) \equiv d\left(w\right)$ is a pure strategy equilibrium of the static game with strategy spaces A_i and payoffs $g_i(\cdot; w)$ and, moreover, $\phi\left(w\right) = g\left(d\left(w\right), w\right)$. Conversely, if $\left(W, w^1, \sigma, d, \phi\right)$ describes an equilibrium of the static game with payoffs $g\left(\cdot; w\right)$ for all $w \in W$, then the induced pure strategy profile in the infinitely repeated game with public monitoring is an equilibrium.[10] We say that the equilibrium is *strict* if $d\left(w\right)$ is a strict equilibrium of the static game $g\left(\cdot; w\right)$ for all $w \in W$.

In this paper, we consider private monitoring where the space of potential signals is also Y; Mailath and Morris [23] 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 $u_i^*\left(y_i, a_i\right)$. 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 $\mathbf{y} \equiv (y_1, \ldots, y_N) \in Y^N$ 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 $|\pi(y, \ldots, y|a) - \rho(y|a)| < \varepsilon$ for all y and a. If π is ε-close to ρ, then $\sum_y \pi(y, \ldots, y|a) > 1 - \varepsilon |Y|$ for all a, where $|Y|$ is the cardinality of Y. We denote the vector $(1, \ldots, 1)$ by $\mathbf{1}$, whose dimension will be obvious from context. Thus, $\pi(y, \ldots, y|a)$ is written as $\pi(y\mathbf{1}|a)$. Let $\pi_i(\mathbf{y}_{-i}|a, y_i)$ denote the implied conditional probability of $\mathbf{y}_{-i} \in Y^{N-1}$. Note that for all $\eta > 0$, there is an $\varepsilon > 0$ such that if π is ε-close to ρ, then $\left|\sum_{\mathbf{y}} u_i^*(y_i, a_i)\, \pi(\mathbf{y}|a) - \sum_{y_i} u_i^*(y_i, a_i)\, \rho(y_i|a)\right| < \eta$. Consequently, we can treat the stage game payoffs in the game with private monitoring as if they are given by $u_i(a)$ (i.e., the stage game payoffs of the game with public monitoring, $\sum_{y_i} u_i^*(y_i, a_i)\, \rho(y_i|a)$).

An immediate but important implication of the assumption that the public monitoring distribution has full support is that, for ε small, a player observing the private signal y assigns high probability to all other players also observing the same signal, irrespective of the actions taken:

Lemma 1. *Fix a full support public monitoring distribution* ρ *and* $\eta > 0$*. There exists* $\varepsilon > 0$ *such that if* π *is* ε*-close to* ρ*, then for all* $a \in A$*,*

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

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* $s \equiv (s_1, \ldots, s_N)$ *given by:*

$$s_i^1 = d_i(w^1),$$
$$s_i^2(a_i^1, y_i^1) = d_i(\sigma(y_i^1, w^1)) \equiv d_i(w_i^2),$$

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

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

We are thus 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 $w_i^t = \sigma\left(h_i^t; w^1\right) = \sigma\left(h_i^t\right)$, taking the initial state 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 strategies induced by a public profile, then player i's belief over the other players' private states, $\beta_i^t \in \Delta\left(W^{N-1}\right)$, is a sufficient statistic for h_i^t, in that i's expected payoff from any continuation strategy only depends on h_i^t through the induced belief β_i^t. Thus, after any private history h_i^t, player i is completely described by the induced private state and belief, $\left(w_i^t, \beta_i^t\right)$. In principle, W may be quite large. For example, if the public strategy profile is nonstationary, it may be necessary to take W as 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 keep track of only each player's private state and that player's beliefs over the other players' private states is a tremendous simplification.

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 $\beta_i\left(\cdot|h_i^t\right) \in \Delta\left(W^{N-1}\right)$ 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 $\mathbf{w} = (w_1, \ldots, w_N)$, with $\mathbf{w}_{-i}$ having the obvious interpretation. We also write $d(\mathbf{w}) \equiv (d_1(w_1), \ldots, d_N(w_N))$.

The following version of the no one-shot deviation principle plays an important role in our analysis.

Lemma 2. *The private profile induced by* $\left(W, w^1, \sigma, d\right)$ *is a sequential equilibrium if for all private histories* h_i^t, $d_i\left(\sigma\left(h_i^t\right)\right)$ *maximizes*

$$\sum_{\mathbf{w}_{-i}} \left\{(1-\delta)u_i(a_i, d_{-i}(\mathbf{w}_{-i})) + \delta \sum_{\mathbf{y}} V_i\left(h_i^t, a_i, y_i\right)\pi(y_i, \mathbf{y}_{-i}|a_i, d_{-i}(\mathbf{w}_{-i}))\right\} \beta_i\left(\mathbf{w}_{-i}|h_i^t\right). \tag{2.1}$$

This lemma follows from the observation that, under any private profile induced by a public profile, future behavior depends only on private states, and private histories determine private states only through private signals. The action chosen is relevant only in forming the posterior beliefs over the other players' private states from the realized private signal. This is true even for actions off-the-equilibrium-path. We emphasize that the lemma requires equation (2.1) to be satisfied at *all* histories, including histories that a player only reaches following his own deviations. Weaker conditions are sufficient for Nash equilibrium. We return to this issue briefly in Example 3.3 and Theorem 4.2 below.

3. The Prisoners' Dilemma

In this section, we study a repeated prisoners' dilemma with public monitoring and examine the robustness of three alternative strict PPE of the public monitoring game. The first and second profiles provide clean illustrations of robust and non-robust profiles. Our third example, the classic grim trigger, illustrates a number of more subtle issues.

The stage game is given by

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

There are two signals, $\underline{y}$ and $\bar{y}$, and the public monitoring distribution ρ is given by

$$\rho\left\{\bar{y}|a_1a_2\right\} = \begin{cases} p, & \text{if } a_1a_2 = CC, \\ q, & \text{if } a_1a_2 = CD \text{ or } DC, \\ r, & \text{if } a_1a_2 = DD. \end{cases} \tag{3.1}$$

Each public profile we consider can be described by a two-state automaton, with state space $W = \{w_C, w_D\}$, initial state $w^1 = w_C$, and decision rules $d_i(w_C) = C$, $d_i(w_D) = D$, for $i = 1, 2$; the profiles differ only in the specification of the transition function.

3.1. A Robust Profile.

The transition function is

$$\sigma(yw) = \begin{cases} w_C, & \text{if } y = \bar{y}, \\ w_D, & \text{if } y = \underline{y}. \end{cases} \tag{3.2}$$

The resulting profile is a strict equilibrium of the public monitoring game if $r \geq q$ and $\delta > [3p - 2q - r]^{-1}$ (and thus $p > \frac{2}{3}q + \frac{1}{3}r + \frac{1}{3}$).[11]

Under this profile, play starts at CC, and continues as long as the good signal $\overline{y}$ is observed, and then switches to DD after the "bad" signal $\underline{y}$. Play returns to CC once the good signal $\overline{y}$ is observed again. 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 (and, in fact, true) that player 1 will find it optimal to play C. Theorem 4.3 below implies that the public profile is therefore robust to public monitoring.

To illustrate this conclusion, let the public monitoring distribution satisfy $r = q$ and the private monitoring distribution π satisfy $\pi(y_1y_2|CD) = \pi(y_1y_2|DC) =$

$\pi(y_1y_2|DD) \equiv \pi^D_{y_1y_2}$ and $\pi(y_1y_2|CC) \equiv \pi^C_{y_1y_2}$. Note that the private monitoring distribution is identical if at least one player chooses D. To explicitly compute the equilibrium conditions, we assume $\pi^C_{\bar{y}\bar{y}} = p(1-2\varepsilon)$, $\pi^C_{\underline{y}\underline{y}} = (1-p)(1-2\varepsilon)$, $\pi^D_{\bar{y}\bar{y}} = q(1-2\varepsilon)$, $\pi^D_{\underline{y}\underline{y}} = (1-q)(1-2\varepsilon)$, and $\pi^a_{\bar{y}\underline{y}} = \pi^a_{\underline{y}\bar{y}} = \varepsilon$ for $a \in \{C, D\}$.[12]

Let $V_{aa'}$ be the continuation value to a player under the profile when he is in state w_a and his opponent is in state $w_{a'}$. The continuation values satisfy

$$\begin{aligned}
V_{CC} &= (1-\delta)\,2 + \delta\left\{\pi^C_{\bar{y}\bar{y}}V_{CC} + \varepsilon V_{CD} + \varepsilon V_{DC} + \pi^C_{\underline{y}\underline{y}}V_{DD}\right\}, \\
V_{CD} &= -(1-\delta) + \delta\left\{\pi^D_{\bar{y}\bar{y}}V_{CC} + \varepsilon V_{CD} + \varepsilon V_{DC} + \pi^D_{\underline{y}\underline{y}}V_{DD}\right\}, \\
V_{DC} &= (1-\delta)\,3 + \delta\left\{\pi^D_{\bar{y}\bar{y}}V_{CC} + \varepsilon V_{CD} + \varepsilon V_{DC} + \pi^D_{\underline{y}\underline{y}}V_{DD}\right\}, \text{ and} \\
V_{DD} &= \delta\left\{\pi^D_{\bar{y}\bar{y}}V_{CC} + \varepsilon V_{CD} + \varepsilon V_{DC} + \pi^D_{\underline{y}\underline{y}}V_{DD}\right\}.
\end{aligned}$$

Thus,

$$V_{CC} - V_{DD} = \frac{2(1-\delta)}{1-\delta(p-q)(1-2\varepsilon)}. \tag{3.3}$$

Suppose a player is in state w_C and assigns a probability ζ to his opponent also being in state w_C. His incentive constraint to follow the profile's specification of C is

$$\begin{aligned}
&\zeta\left\{(1-\delta)\,2 + \delta\left[\pi^C_{\bar{y}\bar{y}}V_{CC} + \varepsilon V_{CD} + \varepsilon V_{DC} + \pi^C_{\underline{y}\underline{y}}V_{DD}\right]\right\} \\
&+(1-\zeta)\left\{(1-\delta)(-1) + \delta\left[\pi^D_{\bar{y}\bar{y}}V_{CC} + \varepsilon V_{CD} + \varepsilon V_{DC} + \pi^D_{\underline{y}\underline{y}}V_{DD}\right]\right\} \\
&\geq \zeta\left\{(1-\delta)\,3 + \delta\left[\pi^D_{\bar{y}\bar{y}}V_{CC} + \varepsilon V_{CD} + \varepsilon V_{DC} + \pi^D_{\underline{y}\underline{y}}V_{DD}\right]\right\} \\
&+(1-\zeta)\left\{\delta\left[\pi^D_{\bar{y}\bar{y}}V_{CC} + \varepsilon V_{CD} + \varepsilon V_{DC} + \pi^D_{\underline{y}\underline{y}}V_{DD}\right]\right\};
\end{aligned}$$

this expression simplifies to

$$\zeta\delta(p-q)(1-2\varepsilon)(V_{CC} - V_{DD}) \geq 1-\delta. \tag{3.4}$$

Substituting from (3.3) into (3.4), we have

$$(1+2\zeta)(1-2\varepsilon) \geq \frac{1}{\delta(p-q)}. \tag{3.5}$$

Recall that if $r = q$ 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 (3.5) hold for $\delta = \underline{\delta}$, which can be guaranteed by ζ close to 1 and ε close to 0. A similar calculation for the incentive constraint describing behavior after $\underline{y}$ yields the inequality

$$(3-2\zeta)(1-2\varepsilon) \leq \frac{1}{\delta(p-q)},$$

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

This example has the strong property that the bound on the private monitoring can be chosen independent of δ. We return to this point in Section 5.

3.2. A Nonrobust Profile.

We now describe a profile that *cannot* be approximated in some arbitrarily close games with private monitoring. The transition function is given by

$$\sigma(yw) = \begin{cases} w_C, & \text{if } w = w_C \text{ and } y = \bar{y}, \\ & \text{or } w = w_D \text{ and } y = \underline{y}, \\ w_D, & \text{if } w = w_C \text{ and } y = \underline{y}, \\ & \text{or } w = w_D \text{ and } y = \bar{y}. \end{cases} \tag{3.6}$$

The resulting profile is a strict equilibrium of the game with public monitoring if $q \geq r$ and $\delta > [3p - 2q + r - 1]^{-1}$ (and thus $p > \frac{2}{3}q - \frac{1}{3}r + \frac{2}{3}$).[13]

Under this profile, 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 w_D cannot be too short, and so play only leaves DD after the less likely signal $\underline{y}$ is realized. This profile (unlike the previous one) is consistent with q and r being close to 1, as long as they are less than p.

The private monitoring distribution π is 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 $\varepsilon > 0$; conditional on the realization of the first stage value of y, the second-stage randomizations are independent across players.

Rather than attempting to directly calculate the beliefs of a player after different private histories, we proceed as follows. Note first that the profile induces a Markov chain on the state space $W^2 \equiv \{w_C w_C, w_D w_D, w_C w_D, w_D w_C\}$, with transition matrix

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

Let γ_t^ε be the unconditional distribution over $W \times W$ in period t. While player 1's belief about player 2's private state, $\beta_1^t\left(w|h_1^t\right)$, depends on 1's private history, the "average" belief of player 1 (conditioning only on her private state) is described by γ_t^ε:

$$\mathcal{E}\left\{\beta_1^t\left(w|h_1^t\right)\middle|\,\sigma\left(h_1^t\right) = w'\right\} = \frac{\gamma_t^\varepsilon(w'w)}{\gamma_t^\varepsilon(w'w_D) + \gamma_t^\varepsilon(w'w_C)}.$$

Since the Markov chain on W^2 is irreducible, it has a unique invariant distribution, α^ε. Moreover, $\gamma_t^\varepsilon \to \alpha^\varepsilon$ as $t \to \infty$. The invariant distribution α^ε is given by

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

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

and

$$\Pr(w_C w_D) \equiv \alpha_3^\varepsilon = \frac{1}{4} = \Pr(w_D w_C) \equiv \alpha_4^\varepsilon.$$

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

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

Consider now the question of what beliefs a player should have over his opponent's private state after a very long history. Observe first that the probability that 1 assigns to 2 being in state w_C, conditional only on 1 being in state w_C, is close to (for ε small)

$$\Pr\{w_C|w_C\} = \frac{2(1-r)}{2(1-r) + (2-p-r)} = \frac{2(1-r)}{4-p-3r},$$

while the probability that 1 assigns to 2 being state in w_C, now conditional only on 1 being in state w_D, is close to

$$\Pr\{w_C|w_D\} = \frac{2-p-r}{4-3p-r}.$$

Then, for ε small,

$$\Pr\{w_C|w_C\} < \Pr\{w_C|w_D\}.$$

Since this probability is the asymptotic expected value of the player's beliefs, there are two histories h_i^t and $\hat{h}_i^t$ such that $w_D = \sigma(h_i^t)$ and $w_C = \sigma(\hat{h}_i^t)$ and $\beta_i(w_C|\hat{h}_i^t) < \beta_i(w_C|h_i^t)$. 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. Notice that this argument relies only on behavior on the equilibrium path. Thus the induced private profile is not a Nash equilibrium.

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 (we call this property connectedness in Section 5).[14]

3.3. Grim Trigger.

The above two examples provide the simplest possible illustration of why some public profiles are robust and some public profiles are not. In our final example, we examine a grim-trigger strategy profile. Because the strategy depends on infinite history, there is no easy technique for proving robustness. In particular, since the profile has an absorbing state (both players will eventually end up defecting), we cannot use the asymptotic behavior of the profile to prove non-robustness, and so focus on the evolution of beliefs along different play paths. The analysis also illustrates the distinction between Nash and sequential equilibrium.

The grim trigger profile is described by the transition function

$$\sigma(yw) = \begin{cases} w_C, & \text{if } w = w_C \text{ and } y = \bar{y}, \\ w_D, & \text{if } w = w_C \text{ and } y = \underline{y}, \text{ or } w = w_D. \end{cases} \tag{3.7}$$

The resulting profile is a strict equilibrium of the public monitoring game if $\delta > [3p - 2q]^{-1}$ (and thus $p > \frac{2}{3}q + \frac{1}{3}$). 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 first "bad" signal $\underline{y}$.[15]

We will show that if $q > r$, the implied private profile is not a Nash equilibrium in *any* close-by game with full-support private monitoring.[16] On the other hand, if $r > q$, while the implied profile is not a sequential equilibrium in *any* close-by game with full-support private monitoring, it is a Nash equilibrium in *every* close-by game with private monitoring.[17] When it is Nash, since the game has no observable deviations, there is a realization-equivalent sequential equilibrium (Sekiguchi [27, Proposition 3]). Since the game and strategy profile are both symmetric, we focus on the incentives of player 1, assuming player 2 follows grim trigger.

In the case $q > r$, in close-by games with private monitoring, it is not optimal for player 1 to play D following a private history of the form $\left(C, \underline{y}; D, \overline{y}; D, \overline{y}; D, \overline{y}; \ldots\right)$, as required by grim trigger. The intuition is as follows. Immediately following the signal $\underline{y}$, player 1 assigns a probability very close to 0 to player 2 being in the private state w_C (because with probability close to 1, player 2 also observed the signal $\underline{y}$). Thus, playing D in the subsequent period is optimal. However, since π has full support, player 1 is *not sure* that player 2 is in state w_D, and observing the signal $\bar{y}$ after playing D is an indication that player 2 had played C (recall that $\rho\left(\bar{y}|DC\right) = q > r = \rho\left(\bar{y}|DD\right)$). This makes player 1 *less* sure that 2 was in state w_D and, if player 2 was in state w_C and observes $\bar{y}$,[18] then 2 will still be in state w_C. Eventually, player 1 believes that player 2 is almost certainly in state w_C, and so will have an incentive to cooperate.

We now formalize this intuition. Suppose player 1 initially assigns prior probability η to player 2 being in state w_C. Write $\zeta^{\pi}\left(\eta, a, y\right)$ for the posterior probability that he assigns to player 2 being in state w_C one period later, if he chooses action a and observes the private signal y, believing that his opponent is following grim trigger. Then,

$$\zeta^{\pi}\left(\eta, D, \overline{y}\right) = \frac{\pi\left(\bar{y}\bar{y}|DC\right)\eta}{\left\{\pi\left(\bar{y}\bar{y}|DC\right) + \pi\left(\bar{y}\underline{y}|DC\right)\right\}\eta + \left\{\pi\left(\bar{y}\bar{y}|DD\right) + \pi\left(\bar{y}\underline{y}|DD\right)\right\}\left(1-\eta\right)}.$$

If π is ε-close to ρ,

$$\zeta^{\pi}\left(\eta, D, \overline{y}\right) > \frac{\left(q-\varepsilon\right)\eta}{\left(q+2\varepsilon\right)\eta + \left(r+2\varepsilon\right)\left(1-\eta\right)}.$$

A simple calculation shows that for any $\overline{\eta} > 0$, there exists $\bar{\varepsilon} > 0$, such that for all private-monitoring distributions π $\bar{\varepsilon}$-close to ρ, $\zeta^{\pi}\left(\eta, D, \overline{y}\right) > \eta$ for all $\eta \leq \overline{\eta}$ [set $\bar{\varepsilon} = \frac{1}{4}\left(1-\overline{\eta}\right)\left(q-r\right)$]. Hence, for any $\eta \in \left(0, \overline{\eta}\right)$, there exists an n such that after observing n signals $\bar{y}$, the posterior must exceed $\overline{\eta}$ (n goes to infinity as η becomes small). This implies that after a sufficiently long history $\left(C, \underline{y}; D, \overline{y}; D, \overline{y}; D, \overline{y}; \ldots\right)$, player 1 eventually becomes very confident that player 2 is in fact in state w_C, and so no longer has an incentive to play D.[19] Since $\left(C, \underline{y}; D, \overline{y}; D, \overline{y}; D, \overline{y}; \ldots\right)$ is a history that occurs with positive probability under grim trigger, grim trigger is not a Nash equilibrium. This, of course, implies that grim trigger is also not sequential.

We now argue that grim trigger is not a sequential equilibrium for any configuration of q and r. In particular, it is not sequentially rational to follow grim trigger's specification of D after long private histories of the form $\left(C, \underline{y}; C, \overline{y}; C, \overline{y}; C, \overline{y}; \ldots\right)$. For π ε-close to ρ, player 1's posterior satisfies

$$\zeta^{\pi}\left(\eta, C, \bar{y}\right) > \frac{\left(p-\varepsilon\right)\eta}{\left(p+2\varepsilon\right)\eta + \left(q+2\varepsilon\right)\left(1-\eta\right)}.$$

Similar to the previous paragraph, for any $\bar{\eta} > 0$, there exists $\bar{\varepsilon} > 0$, such that for all π $\bar{\varepsilon}$-close to ρ, $\zeta^{\pi}(\eta, C, \overline{y}) > \eta$ for all $\eta \leq \bar{\eta}$, and so after a sufficiently long history $\left(C, \underline{y}; C, \overline{y}; C, \overline{y}; C, \overline{y}; \ldots\right)$, player 1 eventually becomes very confident that player 2 is in fact still in state w_C, and so 1 will find it profitable to deviate and play C. Since grim trigger specifies D after such a history, grim trigger is not a sequential equilibrium.

However, histories of the form $\left(C, \underline{y}; C, \overline{y}; C, \overline{y}; C, \overline{y}; \ldots\right)$ occur with zero probability under grim trigger. We now show that, when $r > q$, given any $\xi > 0$, we can choose $\varepsilon > 0$ such that for every history reached with positive probability under grim trigger, if a player has observed $\underline{y}$ at least once, he assigns probability less than ξ to his opponent being in state w_C; and if he has always observed $\overline{y}$, he assigns probability at least $1-\xi$ to his opponent being in state w_C. This implies that grim trigger is Nash (and the induced outcome is sequential) in any close-by game with almost-public monitoring (see Theorem 4.2).

Consider histories in which at least one $\underline{y}$ has been observed. The easy case is a history in which $\underline{y}$ was observed in the last period. Then, for ε small, immediately following such a signal, player 1 assigns a probability of at least $1-\xi$ that 2 also observed $\underline{y}$ (and so will be in state w_D).[20] We now turn to histories in which $\bar{y}$ was observed in the last period. Observing $\bar{y}$ after playing D is an indication that player 2 had played D (this requires $q < r$), and[21]

$$\zeta^{\pi}(\eta, D, \overline{y}) < \frac{(q+\varepsilon)\,\eta}{(q+\varepsilon)\,\eta + (r-\varepsilon)\,(1-\eta)}.$$

There thus exists $\varepsilon > 0$ such that for all π that are ε-close to ρ, $\zeta^{\pi}(\eta, D, \overline{y}) < \eta$ for all $\eta \in (0,1)$ [set $\varepsilon < \frac{1}{2}(r-q)$].[22] So, irrespective of the private signal a player observes, along every play path, the player becomes increasingly confident that his opponent is in state w_D.

Finally, consider histories of the form $(C, \overline{y}; C, \overline{y}; C, \overline{y}; C, \overline{y} \ldots)$. The posterior on such histories satisfies

$$\zeta^{\pi}(\eta, C, \overline{y}) > \frac{(p-\varepsilon)\,\eta}{p\eta + (q+2\varepsilon)\,(1-\eta)}.$$

For any $\bar{\eta} > 0$, there exists $\varepsilon > 0$ such that for all π that are ε-close to ρ, $\zeta^{\pi}(\eta, C, \overline{y}) > \bar{\eta}$ for all $\eta > \bar{\eta}$ [set $\varepsilon = \frac{1}{3}(1-\bar{\eta})(p-q)$]. Thus the player's belief that his opponent is in state w_C will never fall below $\bar{\eta}$ given private histories of the form $(C, \overline{y}; C, \overline{y}; C, \overline{y}; C, \overline{y} \ldots)$. It is worth noting that for η close to 1, $\zeta^{\pi}(\eta, C, \overline{y}) < \eta$. For example, even if $\eta = 1$ (such as in period 1), and player 1 observed $\bar{y}$, there is a positive probability that 2 observed $\underline{y}$, and so switched state to w_D. As a consequence, even after initially observing a sequence of $\bar{y}$'s, player 1's posterior probability that 2 is in state w_C must fall. However, it cannot fall below the fixed point η_{ε} of $\zeta^{\pi}(\eta, C, \bar{y})$, and $\eta_{\varepsilon} \to 1$ as $\varepsilon \to 0$.

3.4. Discussion.

Our analysis of the grim trigger strategy profile has a similar flavor to that of Compte [12] and Sekiguchi [27]: we all study strategies with the property that a player's belief about a binary state of the other player is a sufficient statistic for their private histories; thus checking for optimality involves understanding the evolution of beliefs over long histories.

We restricted attention to a very simple strategy (grim trigger) and exploited the closeness of the evolution of beliefs under nearby almost-public monitoring distributions to that of the beliefs under public monitoring. This allowed us to prove positive and negative results about the pure grim trigger strategy.

Compte [12], for the case of conditionally independent monitoring (in the sense that, conditional on the action profile, players' signals are independent[23]), studies a much larger class of *trigger-strategy* equilibria: those with the property that once a player defects, he thereafter defects with probability one. Compte [12] shows that, for sufficiently patient players, the average expected payoff in *any* trigger-strategy equilibrium is close to the payoff from defection.[24] Compte's task is complicated by the need for a result that holds uniformly across trigger-strategy profiles, i.e., for different triggering events.[25] Essentially, when players are patient, if a triggering event is sufficiently forgiving to imply payoffs under the profile significantly better than the payoff from defection, it cannot provide sufficient incentives to support cooperation, since a player could defect occasionally without significantly changing the beliefs of his opponent.

Sekiguchi [27] showed that it was possible to achieve efficiency in the repeated prisoners' dilemma, even if the private monitoring is conditionally independent, as long as the monitoring is almost perfect.[26,27] Since always play D is the only pure strategy equilibrium with conditionally-independent signals, the equilibrium Sekiguchi constructs is in mixed strategies. Consider first a profile in which players randomize at the beginning of the game between the strategy of always play D, and grim trigger strategy. This profile is an equilibrium (given a payoff restriction) for moderate discount factors and sufficiently accurate private monitoring. Crudely, there are two things 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 argument is the analog to our argument that a player who always observes signal $\overline{y}$ will think it more likely that his opponent is in w_C.

Second, will a player defect as soon as a defect signal is received? Suppose first that the first defect signal occurs after a long history. That is, suppose player i is playing grim trigger, has observed a long sequence of cooperative signals (so that i is reasonable confident that player j is also playing grim trigger), and then observes the defect signal. The two highest order probability events are that i received an erroneous signal (in which case j is still cooperative) and that player j had received an erroneous signal in the previous period (in which case j now defects forever). These two events have equal probability, and if players are not too patient (so they are unwilling to experiment), player i will defect. (If players are patient, 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.) Suppose now that the first defect signal occurs early in the game. The most extreme possibility is the defect signal is the first signal observed. Clearly, in the initial period player j is not responding to any signal, and so for player i to assign positive probability to j playing D in the initial period, j must in fact defect in the initial period with positive probability (and this is why the profile is in mixed strategies).[28] 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.

Sekiguchi thus shows that the existence of an always defect/grim-trigger strategy mixture that is a Nash equilibrium for moderate discounting. The upper bound on the discount factor can be removed by observing (following Ellison [13]) that the repeated game can be divided into N distinct games, with the kth game played in period $k+tN$, where $t \in \mathbf{N}$. This gives an effective discount rate of δ^N on each game. Of course, the resulting strategy profile does not look like grim trigger.

4. Approximating Arbitrary Strict Public Equilibria - Fixed Discount Factors

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 PPE is an equilibrium. In this section, we focus on the case of fixed discount factors. More specifically, we ask: If a public profile is a strict PPE 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?

Recall that $\beta_i\left(\sigma\left(h_i^t\right)\mathbf{1}|h_i^t\right)$ 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 $\sigma\left(h_i^t\right)$. Our first observation is that if all players assign uniformly high probability to all players always being in the same private state and if the private monitoring is sufficiently close to the public monitoring distribution, then $V_i\left(h_i^t\right)$ is close to the continuation value of the state $\sigma\left(h_i^t\right)$ in the game with public monitoring:

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

We sketch a proof: For fixed δ, there is a T such that the continuation value from $T+1$ is less than $\upsilon/3$. For π sufficiently close to ρ, the probability that all players observe the same T period history of private signals is close to 1. Thus, for ε small, for any private history $\hat{h}_i^t$, conditional on all players starting in the common private state $\sigma\left(\hat{h}_i^t\right)$, the value over T periods is within $\upsilon/3$ of the value over T periods of the game with public monitoring starting in the state $\phi_i\left(\sigma\left(\hat{h}_i^t\right)\right)$. Finally, for η small, the impact on the continuation value of players being in disagreement is less than $\upsilon/3$.

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 3.1, the bound on ε depends on δ, and this bound becomes tighter as $\delta \to 1$.

Theorem 4.1. *Suppose the public profile $\left(W, w^1, \sigma, d\right)$ is a strict equilibrium of the game with public monitoring for some δ and $|W| < \infty$. For all $\kappa > 0$, there exists η and ε such that if the posterior beliefs induced by the private profile satisfy $\beta_i\left(\sigma\left(h_i^t\right)\mathbf{1}|h_i^t\right) > 1-\eta$ for all h_i^t, and if π is ε-close to ρ, then the private profile is a sequential 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 $\left(W, w^1, \sigma, d\right)$ be the automaton description of the public profile. From Lemma 2, 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 $\hat{a}_i = d_i\left(\sigma\left(\hat{h}_i^t\right)\right)$. Fix a history $\hat{h}_i^t$. Let $w_i^t = \sigma\left(\hat{h}_i^t\right)$. Since the public monitoring equilibrium is strict, there exists θ such that, for all $a_i \neq \hat{a}_i = d_i\left(w_i^t\right)$,

$$\phi_i\left(w_i^t\right) - \theta \geq (1-\delta)\, u_i\left(a_i, d_{-i}\left(w_i^t\right)\right) + \delta \sum_y \phi_i\left(\sigma\left(y; w_i^t\right)\right)\rho\left(y|a_i, d_{-i}\left(w_i^t\right)\right). \quad (4.1)$$

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

$$V_i\left(\hat{h}_i^t\right) \geq \sum_{w_{-i}} \left\{ (1-\delta)\, u_i\left(a_i, d_{-i}\left(\mathbf{w}_{-i}\right)\right) \right.$$

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

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 (4.1) can be chosen independently of the state w_i^t.

While all our later results concern sequential equilibria, we note that if we only require the restriction on posterior beliefs to hold on private histories consistent with the strategy profile, we obtain a Nash equilibrium. The proof is in the Appendix.

Theorem 4.2. *Suppose the public profile $\left(W, w^1, \sigma, d\right)$ is a strict equilibrium of the game with public monitoring for some δ and $|W| < \infty$. For all $\kappa > 0$, there exists η and ε such that if the posterior beliefs induced by the private profile satisfy $\beta_i\left(\sigma\left(h_i^t\right)\mathbf{1}|h_i^t\right) > 1-\eta$ for all $h_i^t = \left(d_i\left(w^1\right), y_i^1; d_i\left(w_i^2\right), y_i^2; \ldots; d_i\left(w_i^{t-1}\right), y_i^{t-1}\right)$, where $w_i^{\tau+1} \equiv \sigma\left(y_i^\tau, w_i^\tau\right)$, and if π is ε-close to ρ, then the private profile is a Nash equilibrium of the game with private monitoring for the same δ, there is a realization-equivalent sequential equilibrium, and the expected payoff in that equilibrium is within κ of the public equilibrium payoff.*

Returning to sequential equilibria, the key issue is the behavior of $\beta_i\left(\sigma\left(h_i^t\right)\mathbf{1}|h_i^t\right)$. In particular, can $\beta_i\left(\sigma\left(h_i^t\right)\mathbf{1}|h_i^t\right)$ 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 $\beta_i\left(\sigma\left(h_i^t\right)\mathbf{1}|h_i^t\right) > 1-\eta$ 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 $\beta_i\left(\sigma\left(h_i^t\right)\mathbf{1}|h_i^t\right)$ 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 $\sigma\left(y, \left(y^L, \ldots, y^2, y^1\right)\right) = \left(y, y^L, \ldots, y^2\right)$ for all $y \in Y$. The initial state is $w^1 = (*, \ldots, *)$.*

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, studied in Section 3.3. While the grim trigger profile requires only one-period memory when 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 $(\bar{y}, \bar{y}, \cdots, \bar{y})$, and to play

D after $\left(\underline{y}, \bar{y}, \cdots, \bar{y}\right)$. This difference is crucial, since this profile is never a sequential equilibrium, and is often not a Nash equilibrium in close-by private monitoring games. Cole and Kocherlakota [9] show that, for some imperfect public monitoring distributions that support a folk theorem with infinite histories, all symmetric public profiles with finite memory must have defection in every period. On the other hand, other 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 strongly symmetric strategies, in a repeated Cournot game. The proof of the following theorem is in the Appendix.

Theorem 4.3. *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 $\beta_i\left(\sigma\left(h_i^t\right)\mathbf{1}|h_i^t\right) > 1-\eta$ for all h_i^t.*

5. 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 the class of finite public profiles defined below, 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 $a_i \neq d_i(w)$,

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

where

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

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

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

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

$$\phi_i(w) > (1-\delta)\,\hat{u}_i(w, a_i) + \delta \sum_{w'} \phi_i\left(w'\right)\theta_{ww'}\left(d_{-i}(w), a_i\right), \tag{5.1}$$

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

$$\delta \sum_{w'} \phi_i(w') \left(\hat{\theta}_{ww'} - \hat{\theta}_{ww'}(a_i)\right) > (1-\delta)(\hat{u}_i(w, a_i) - \hat{u}_i(w)). \tag{5.2}$$

For any $\bar{w} \in W$, (5.2) is equivalent to

$$\delta \sum_{w' \neq \bar{w}} (\phi_i(w') - \phi_i(\bar{w})) \left(\hat{\theta}_{ww'} - \hat{\theta}_{ww'}(a_i)\right) > (1-\delta)(\hat{u}_i(w, a_i) - \hat{u}_i(w)). \tag{5.3}$$

The following property of connectedness plays a critical role in obtaining a bound 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 $y^1, \ldots, y^m$ and $\bar{w} \in W$ such that*

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

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 (the proofs of the Lemmas in this section are in the Appendix).

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

$$\sigma\left(y^n, \sigma\left(y^{n-1}, \ldots, \sigma\left(y^1, w\right)\right)\right) = \bar{w}, \qquad \forall w \in W.$$

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, $\lim_{\delta \to 1} \phi_i(w)$ is independent of $w \in W$, and so simply taking $\delta \to 1$ in (5.3) yields $0 \geq 0$. On the other hand, if we can divide by $(1-\delta)$ and take limits as $\delta \to 1$, we have a version of the incentive constraint that is independent of δ. The next lemma assures us that this actually makes sense.

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

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

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 (5.3) for δ large.

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

$$\sum_{w' \neq \bar{w}} \Delta_{w'\bar{w}} \phi_i \times \left(\hat{\theta}_{ww'} - \hat{\theta}_{ww'}(a_i) \right) > \hat{u}_i(w, a_i) - \hat{u}_i(w), \tag{5.4}$$

where $\bar{w}$ *is any state.*

The particular choice of $\bar{w}$ in (5.4) is irrelevant: if (5.4) 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 4.3, induces posterior beliefs that assign uniformly large probability to agreement in private states. Note that if we only imposed the belief requirement after histories consistent with the strategy profile, the profile would still be a Nash equilibrium (and the induced outcome sequential) in the game with private monitoring (by Theorem 4.2).

Theorem 5.1. *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* $\beta_i\left(\sigma\left(h_i^t\right)\mathbf{1}|h_i^t\right) > 1 - \eta$ *for all* h_i^t, π *is* ε*-close to* ρ*, and* $\delta \in (\underline{\delta}, 1)$*, then the private profile is a sequential 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 $\mathbf{w} = (w_1, \ldots, w_N)$ and $\mathbf{w}' = (w'_1, \ldots, w'_N)$,

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

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

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

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

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* $\mathbf{w}$, $\bar{\mathbf{w}} \in W^N$,

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

   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 (5.4) holds.

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

$$\sum_{\mathbf{w}'\neq\bar{w}\mathbf{1}} \Delta_{\mathbf{w}',\bar{w}\mathbf{1}}\psi_i^\pi \times \left(\tilde{q}^\pi_{w\mathbf{1},\mathbf{w}'} - \tilde{q}^\pi_{w\mathbf{1},\mathbf{w}'}(a_i)\right) > \hat{u}_i(w,a_i) - \hat{u}_i(w), \tag{5.5}$$

where $\bar{w}$ *is any state.*

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 $a_i \neq d_i(w)$,

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

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

$$\psi_i^\pi (w\mathbf{1}) > (1-\delta)\, \tilde{u}_i (w\mathbf{1}, a_i) + \delta \sum_{\mathbf{w}'} \psi_i^\pi \left(\mathbf{w}'\right) \tilde{q}^\pi_{w\mathbf{1},\mathbf{w}'} (a_i) \,. \tag{5.6}$$

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

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

For any state $\bar{w} \in W$, (5.6) is equivalent to

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

Dividing by $(1-\delta)$ and taking limits then yields (5.5). Thus, (5.5) 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.

6. 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 $g : A \to \Re^N$. 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 [25, 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 distribution (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 PPE of the game with public monitoring (the arguments are almost the same as the proofs of Lemma 3 and Theorem

4.1). Since the public profile has finite memory, we then have (from Theorems 4.1 and 4.3) a bound on ε (that depends on ρ and δ) so that if the private monitoring distribution is ε-close to ρ,[29] 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 $\eta \to 0$ as $\delta \to 1$, and second, even if η was independent of δ, $\varepsilon \to 0$ as $\delta \to 1$.[30] 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 1).

In order to apply the techniques of Section 5, we first modify the profile used for the perfect monitoring folk theorem so that in η-perfect public monitoring games, it is patiently strict.[31] Let $\underline{g}_i$ denote player i's pure strategy minmax 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 [25, Proposition 151.1]. We do not know if a similar result holds for the mixed action version, with unobservable mixing. The proofs of Theorem 6.1 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 6.1. *(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 distribution 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 6.1, we have:

Corollary 1. *(Imperfect Public Monitoring) Fix a stage game $g : A \to \Re^N$. Suppose a^* satisfies the perfect folk theorem condition. Let s denote the subgame perfect equilibrium profile described in Theorem 6.1. There exists $\underline{\delta} < 1$ and $\eta > 0$ such that if the public monitoring distribution 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 patiently strict.*

This corollary, with Theorems 4.3 and 5.1, yields:

Theorem 6.2. *(Private Monitoring) Fix a stage game $g : A \to \Re^N$. 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 distributions (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 a sequential equilibrium of the repeated game with private monitoring whose equilibrium payoff is within ν of $g(a^*)$.*

A. Proofs for Section 4. Fixed Discount Factors

Proof of Theorem 4.2. A private history h_i^t is consistent with the strategy profile if it is of the form $\left(d_i\left(w^1\right), y_i^1;\ d_i\left(w_i^2\right), y_i^2; \ldots; d_i\left(w_i^{t-1}\right), y_i^{t-1}\right)$, where $w_i^{\tau+1} \equiv \sigma\left(y_i^\tau, w_i^\tau\right)$. We need the following (which can be proved along the lines of the proof of Lemma 3): Suppose the strategies for players $j \neq i$ are described by $\left(W, w^1, \sigma, d\right)$. If s_i is an arbitrary (continuation) pure strategy for player i and h_i^t is an arbitrary history consistent with the strategy profile, denote by $V_i\left(s_i | h_i^t\right)$ player i's continuation value of s_i in the game with private monitoring, conditional on the private history h_i^t. Denote by $\phi_i\left(s_i | w\right)$ player i's continuation value of s_i in the game with public monitoring, when the other players are in the common state w. (Note that if s_i is the strategy corresponding to σ and d_i, starting at state $\sigma\left(h_i^t\right)$, then $V_i\left(s_i | h_i^t\right) = V_i\left(h_i^t\right)$ and $\phi_i\left(s_i | h_i^t\right) = \phi_i\left(h_i^t\right)$.) Then, for all $\upsilon > 0$, there exists ε and $\eta > 0$ such that for all s_i and all histories on the equilibrium path, if $\beta_i\left(\sigma\left(h_i^t\right)\mathbf{1} | h_i^t\right) > 1 - \eta$ and π is ε-close to ρ , then $\left|V_i\left(s_i | h_i^t\right) - \phi_i\left(s_i | \sigma\left(h_i^t\right)\right)\right| < \upsilon$.

Fix a history h_i^t consistent with the strategy profile. Let $\tilde{s}_i$ be a deviation continuation strategy for player i with $\tilde{s}_i^1 \neq d_i\left(w_i^t\right)$, and let $\hat{s}_i$ denote the continuation strategy that agrees with $\tilde{s}_i$ in period t and then follows d_i and σ ($\hat{s}_i$ is a one-shot deviation). Then,

$$\phi_i\left(\tilde{s}_i | \sigma\left(h_i^t\right)\right) \leq \phi_i\left(\hat{s}_i | \sigma\left(h_i^t\right)\right),$$

and since the public profile is strict, there exists $\theta > 0$ (as above, θ can be chosen independent of h_i^t) such that

$$\phi_i\left(\hat{s}_i|\sigma\left(h_i^t\right)\right) < \phi_i\left(w_i^t\right) - \theta.$$

Finally, by choosing $\upsilon < \theta/3$, we have

$$\begin{aligned} V_i\left(\tilde{s}_i|h_i^t\right) &< \phi_i\left(\tilde{s}_i|\sigma\left(h_i^t\right)\right) + \theta/3 \\ &< \phi_i\left(w_i^t\right) - 2\theta/3 \\ &< V_i\left(h_i^t\right) - \theta/3, \end{aligned}$$

so that $\tilde{s}_i$ is not a profitable deviation. Thus, $\left(W, w^1, \sigma, d\right)$ describes a Nash equilibrium of the game with private monitoring.

Finally, since games with almost-public monitoring have no observable deviations, any Nash equilibrium outcome can be supported by a sequential equilibrium (Sekiguchi [27, Proposition 3]).

■

Proof of Theorem 4.3. 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 $w \equiv \left(y_i^1, \ldots, y_i^L\right)$ (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 $a_i^\tau = d_i\left(w_i^\tau\right)$, 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 $y_i^1, \ldots, y_i^\ell$, 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

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

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 $\varepsilon \in (0, \bar{\varepsilon})$,

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

[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.]

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

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

■

B. Proofs for Section 5. Arbitrarily Patient Players

Proof of Lemma 4. 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 $y^1, \ldots, y^m$ be a sequence that satisfies $\sigma\left(y^m, \sigma\left(y^{m-1}, \ldots, \sigma\left(y^1, w^1\right)\right)\right) = \sigma\left(y^m, \sigma\left(y^{m-1}, \ldots, \sigma\left(y^1, w^2\right)\right)\right) \equiv w$. Since the profile is connected, there is a sequence of signals $y^{m+1}, \ldots, y^{m+m'}$ such that $\sigma(y^{m+m'}, \sigma(y^{m+m'-1}, \ldots, \sigma\left(y^{m+1}, w\right))) = \sigma(y^{m+m'}, \sigma(y^{m+m'-1}, \ldots, \sigma\left(y^{m+1}, w'\right)))$, where $w' \equiv \sigma\left(y^m, \sigma\left(y^{m-1}, \ldots, \sigma\left(y^1, w^3\right)\right)\right)$. The desired sequence of signals is then $y^1, \ldots, y^{m+m'}$. ■

We need the following standard result (see, for example, Stokey and Lucas [29, 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 $x \in \Re^{\ell}$, 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 $\hat{u}_i \in \Re^W$ and $\phi_i \in \Re^W$ 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,

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

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

■

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 $\left(y^1, \ldots, y^n\right)$ is the finite sequence of signals from Lemma 4. Moreover, there exists $\varepsilon > 0$ such that for all π ε-close to ρ,

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

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,

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

and

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

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

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

and, for all π ε-close to ρ,

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

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, $Q_{11}^\pi \to \Theta$, $Q_{12}^\pi \to 0$, and $Q_{22}^\pi \to 0$.

Now,

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

and

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

and, in general,

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

and

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

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

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

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

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

■

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 (5.5) is

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

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

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

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

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

and so

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

which is the desired inequality (5.5).

■

C. Proofs for Section 6. An Application to Folk Theorems

Proof of Theorem 6.1. While the profile specified in the proof of Proposition 151.1 of Osborne and Rubinstein [25] 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 [25]. 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 $M \equiv \max_{i,a} |g_i(a)|$),[32]

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

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

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

Each player has a strict incentive to follow the prescribed path when in phase $C(j)$ if, for all $\ell \in \{1, \ldots, L^{**}\}$ (where ℓ is the number of periods remaining in phase $C(j)$),[33]

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

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

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

which is implied by (C.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 $\ell \in \{1, \ldots, L^{**}\}$, (C.3) holds.

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

$P(j)$),

$$M + \sum_{k=2}^{L^*+1} \delta^{k-1} \underline{g}_i + \sum_{k=L^*+2}^{L^*+L^{**}+1} \delta^{k-1} g_i(a(i)) \tag{C.4}$$
$$< \sum_{k=1}^{\ell} \delta^{k-1} g_i\left(\underline{a}^j\right) + \sum_{k=\ell+1}^{\ell+L^{**}} \delta^{k-1} g_i(a(j)) + \sum_{k=\ell+L^{**}+1}^{L^*+L^{**}+1} \delta^{k-1} g_i(a^*).$$

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

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

which is implied by (C.2). Thus, there exists δ'' such that for $\delta \in (\delta'', 1)$, and any $\ell \in \{1, \ldots, L^*\}$, (C.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 (C.3) and (C.4) hold strictly).

■

Proof of Corollary. Let (W, w^1, σ, d) be a finite state automaton description of the strategy profile from Theorem 6.1. Let $v_i : W \to \Re$ describe player i's continuation values under this profile. Observe that, for all $w \in W$, $v_i(w) \to g_i(a^*)$ as $\delta \to 1$. 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,

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

and so

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

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

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

where $\bar{w}$ is any state, and

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

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] D. Abreu, Extremal equilibria of oligopolistic supergames, *J. Econ. Theory* **39** (1986), 191–225.

[2] D. Abreu, D. Pearce, and E. Stacchetti, Toward a theory of discounted repeated games with imperfect monitoring, *Econometrica* **58** (1990), 1041–1063.

[3] I. Ahn, Imperfect information repeated games with a single perfect observer. Mimeo, 1996.

[4] M. Amarante, Recursive structure and equilibria in games with private monitoring. Columbia University, April 1999.

[5] M. Aoyagi, A collusion mechanism in dynamic bertrand oligopoly under private observation of correlated demand shocks. Mimeo, University of Pittsburgh, 1997.

[6] E. Ben-Porath and M. Kahneman, Communication in repeated games with private monitoring, *J. Econ. Theory* **70** (1996), 281–297.

[7] V. Bhaskar and I. Obara, Belief-based equilibria in the prisoners' dilemma with private monitoring. University of Essex and University of Pennsylvania, 2000.

[8] V. Bhaskar and E. van Damme, Moral hazard and private monitoring. Mimeo, University of St. Andrews and CentER Tilburg, 1997.

[9] H. L. Cole and N. Kocherlakota, Finite memory and imperfect monitoring. Working Paper 604, Federal Reserve Bank of Minneapolis, 2000.

[10] O. Compte, Coordination in repeated games with almost perfect private monitoring. CERAS, 1994.

[11] O. Compte, Communication in repeated games with imperfect private monitoring, *Econometrica* **66** (1998), 597–626.

[12] O. Compte, On failing to cooperate when monitoring is private. Mimeo, C.E.R.A.S.-E.N.P.C., 2000.

[13] G. Ellison, Cooperation in the prisoner's dilemma with anonymous random matching, *Rev. Econ. Stud.* **61** (1994), 567–588.

[14] J. C. Ely and J. Välimäki, A robust folk theorem for the prisoner's dilemma. Mimeo, Northwestern University, 1999.

[15] D. Fudenberg and D. K. Levine, An approximate folk theorem with imperfect private information, *J. Econ. Theory* **54** (1991), 26–47.

[16] D. Fudenberg, D. K. Levine, and E. Maskin, The folk theorem with imperfect public information, *Econometrica* **62** (1994), 997–1040.

[17] E. Green and R. Porter, Noncooperative collusion under imperfect price formation, *Econometrica* **52** (1984), 87–100.

[18] M. Kandori, Cooperation in finitely repeated games with imperfect private information. Mimeo, Princeton University, 1991.

[19] M. Kandori and H. Matsushima, Private observation, communication and collusion, *Econometrica* **66** (1998), 627–652.

[20] M. Kandori and I. Obara, Efficiency in repeated games revisited: The role of private strategies. Mimeo, University of Tokyo and University of Pennsylvania, 2000.

[21] E. Lehrer, On the equilibrium payoffs set of two player repeated games with imperfect monitoring, *Int. J. Game Theory* **20** (1992), 211–226.

[22] G. J. Mailath, S. A. Matthews, and T. Sekiguchi, Private strategies in finitely repeated games with imperfect public monitoring. CARESS Working Paper #01-10, University of Pennsylvania, 2001.

[23] G. J. Mailath and S. Morris, Finite state strategies and coordination in repeated games with private monitoring. University of Pennsylvania and Yale University, 2001.

[24] H. Matsushima, On the theory of repeated games with private information. Part I: Anti-folk theorem without communication, *Econ. Letters* **35** (1991), 253–256.

[25] M. J. Osborne and A. Rubinstein, "A Course in Game Theory," The MIT Press, Cambridge, MA, 1994.

[26] M. Piccione, The repeated prisoner's dilemma with imperfect private monitoring. Mimeo, University of Southampton, 1999.

[27] T. Sekiguchi, Efficiency in repeated prisoner's dilemma with private monitoring, *J. Econ. Theory* **76** (1997), 345–361.

[28] T. Sekiguchi, Robustness of efficient equilibria in repeated games with imperfect private monitoring. Mimeo, 1998.

[29] N. Stokey and R. E. Lucas, Jr., "Recursive Methods in Economic Dynamics," Harvard University Press, Cambridge, MA, 1989.

Notes

[1]A 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 and Obara [20] and Mailath, Matthews, and Sekiguchi [22]) exploring private strategy profiles in games with imperfect public monitoring suggests that there is more to learn about games with public monitoring.

[2]Amarante [4] provides a *nonstationary* recursive characterization of the equilibrium set of repeated games with private monitoring.

[3]Compte [10] was concerned with similar robustness issues, with perfect monitoring, rather than imperfect public monitoring, as the benchmark.

[4]For some public monitoring distributions, the grim-trigger strategy profile is a Nash equilibrium, and so there is a realization-equivalent sequential equilibrium (Theorem 4.2). However, the specification of the off-the-equilibrium-path behavior in the *original* strategy profile (reinterpreted as a profile for the game with private monitoring) is not sequentially rational.

[5]In much of this literature, each player observes a private noisy (but reasonably accurate) signal of the opponent's play.

[6]This work builds on an idea in Kandori [18].

[7]While 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.

[8]Since 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]).

[9]We assume W contains no irrelevant states, i.e., every $w \in W$ is reached from w^1 by some sequence of public signals.

[10]We 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 $W \subset \Re^N$ and $\phi(w) = w$ (again, see Abreu, Pearce, and Stacchetti [2]).

[11]If $q > r$, this profile will not be an equilibrium for δ close to 1, as players would have an incentive to cooperate in state w_D.

[12]In this section only, we will use a to denote an arbitrary action C or D, rather than

an action profile.

[13]If $q < r$, this profile will not be an equilibrium for δ close to 1, as players would have an incentive to cooperate in state w_D.

[14]A similar point was made in Compte [10].

[15]As $\delta \to 1$, the average payoff in this profile converges to 0, since the expected waiting time till a bad signal is finite and independent of δ. This profile is an example of a 1-period grim trigger. In a K-period grim trigger, behavior begins in CC, and switches to DD permanently after the first K "bad" signals. By choosing K sufficiently large, for fixed δ, the average payoff can be made arbitrarily close to 1. We analyze general K-period grim trigger strategy profiles in Mailath and Morris [23].

[16]A private monitoring distribution has *full-support* if $\pi(\mathbf{y}|a) > 0$ for all $\mathbf{y} \in Y^N$ and $a \in A$.

[17]It will be clear from the argument that the specification $\rho\{\bar{y}|CD\} = \rho\{\bar{y}|DC\}$ is purely for notational convenience, since nothing depends upon the relative values of $\rho\{\bar{y}|CD\}$ and $\rho\{\bar{y}|DC\}$. Thus, we characterize the robustness of grim trigger under generic two signal imperfect public monitoring technologies.

[18]And it is very likely that player 2 has observed $\bar{y}$ when player 1 has observed $\bar{y}$, because the monitoring is almost public.

[19]We are implicitly appealing to an argument by contradiction. Suppose grim trigger is a Nash equilibrium. After private histories that arise with positive probability under the strategy profile, it is optimal to play C when in state w_C and optimal to play D when in state w_D. Moreover, these incentives are strict when a player assigns probability one to the opponent being in the same private state as himself. Thus, whenever a player assigns arbitrarily large probability to the opponent being in state w_D, playing C is suboptimal.

[20]After observing $\underline{y}$ in period t, a player's posterior that his opponent's private state in period t was w_C increases. However, if the opponent's signal were also $\underline{y}$, then the opponent's private state in period $t+1$ is necessarily w_D. For ε-close private monitoring distributions, the second consideration clearly dominates the first.

[21]Recall that $x/(x+a)$ is increasing in x for $a > 0$.

[22]Note that this implies another failure of sequentiality. Suppose player 1 is in state w_C. After long histories of the form $(D, \bar{y}; D, \bar{y}; D, \bar{y}; \ldots)$, player 1 assigns high probability to player 2 being in state w_D, and even though he is still in state w_C, will prefer

to play D.

[23]This immediately rules out almost-public monitoring.

[24]Matsushima [24] 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.

[25]These triggering events include the triggering events of footnote 15, where K becomes arbitrarily large as $\delta \to 1$.

[26]This result and the associated technique, which builds on Bhaskar and van Damme [8], has been significantly generalized by Bhaskar and Obara [7] and Sekiguchi [28].

[27]The class studied includes both independent and correlated signal distributions. Denote by lower case letters the private signals that a player observes, so that c_i is the signal that player i observes about player j. Almost-perfect monitoring requires not only that $\Pr\{c_1c_2|C_1C_2\}$ and $\Pr\{d_1d_2|D_1D_2\}$ be close to one, but that $\Pr\{c_1d_2|D_1C_2\}$ be close to one. If we interpret $\bar{y}$ as c_i and $\underline{y}$ as d_i, then this is inconsistent with almost-public monitoring, for any choice of p, q, and r.

There is an interpretation, however, under which almost-perfect monitoring is consistent with almost-public monitoring in our example. Suppose p and r are close to 1, and q is close to 0. Then, if i has chosen C, $\bar{y}$ is an almost sure indication that j had also chosen C, while $\underline{y}$ is an almost sure indication that j had chosen D. On the other hand, if i had chosen D, $\bar{y}$ is an almost sure indication that j had also chosen D, while $\underline{y}$ is an almost sure indication that j had chosen C. The interpretation of $\bar{y}$ (as a signal of j's choice) now depends upon i's action. Observe, however, that since grim trigger only depends upon the signal while the player is choosing C, the differences in interpretation have no impact on play. The differences are relevant in belief formation. Note that when player i in state w_D observes c_i ($\underline{y}$), indicating that player j is still in state w_C, player j will observe (with high probability) d_j ($\underline{y}$) and so switch to state w_D (cf. footnote 20).

[28]In other words, grim trigger is not a Nash equilibrium because players have an incentive to ignore defect signals received in the first period (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.

The randomization probability is chosen to make the players indifferent between cooperation and defection in the initial period. 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.

[29]It is worth emphasizing we are imposing a particular structure on the private monitoring distribution. There are almost-perfect private monitoring distributions (with support A^N) that are conditionally independent, and so not almost public. Also note that for private monitoring distributions that are ε-close to an η-perfect public monitoring distribution, observation errors are almost perfectly correlated.

[30]Sekiguchi [28] proved such a result for efficient outcomes in the repeated prisoners' dilemma.

[31]While the profile from Osborne and Rubinstein [25, 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 6.1 implies that, even when the monitoring is perfect, the Markov chain on W is ergodic.

[32]Osborne and Rubinstein [25, 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 (C.1), because in their profile, after a deviation play never returns to a^*.

[33]If $j = 0$, then the value of ℓ is irrelevant.