

**COORDINATION FAILURE IN REPEATED GAMES
WITH ALMOST-PUBLIC MONITORING**

By

George J. Mailath and Stephen Morris

March 2005

COWLES FOUNDATION DISCUSSION PAPER NO. 1479R



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY

Box 208281

New Haven, Connecticut 06520-8281

<http://cowles.econ.yale.edu/>

Coordination Failure in Repeated Games with Almost-Public Monitoring*

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

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

March 23, 2005
first version August 20, 2004

Abstract

Some private-monitoring games, that is, games with no public histories, can have histories that are *almost* public. These games are the natural result of perturbing public-monitoring games towards private monitoring. We explore the extent to which it is possible to coordinate continuation play in such games. It is always possible to coordinate continuation play by requiring behavior to have *bounded recall* (i.e., there is a bound L such that in any period, the last L signals are sufficient to determine behavior). We show that, in games with general almost-public private monitoring, this is essentially the only behavior that can coordinate continuation play.

Keywords: repeated games, private monitoring, almost-public monitoring, coordination, bounded recall.

JEL Classification Numbers: C72, C73, D82.

*Earlier versions of this material have appeared under the title “Finite State Strategies and Coordination in Repeated Games with Private Monitoring.” Some parts of Sections 4 and 5 first appeared in “Repeated Games with Imperfect Private Monitoring: Notes on a Coordination Perspective.” That paper is subsumed by Mailath and Morris (2002) and this paper. We thank Andrew Postlewaite for helpful conversations.

Coordination Failure in Repeated Games with Almost-Public Monitoring

by George J. Mailath and Stephen Morris

1. Introduction

Intertemporal incentives often allow players to achieve payoffs that are inconsistent with myopic incentives. For games with public histories, the construction of sequentially rational equilibria with nontrivial intertemporal incentives is straightforward. Since continuation play in a public strategy profile is a function of public histories only, the requirement that continuation play induced by any public history constitute a Nash equilibrium of the original game is both the natural notion of sequential rationality and relatively easy to check (Abreu, Pearce, and Stacchetti (1990)). These *perfect public equilibria* (or *PPE*) use public histories to coordinate continuation play.

While games with private monitoring (where actions and signals are private) have no public histories to coordinate continuation play, some do have histories that are *almost* public. We explore the extent to which it is possible to coordinate continuation play for such games. It is always possible to coordinate continuation play by requiring behavior to have *bounded recall* (i.e., there is a bound L such that in any period, the last L signals are sufficient to determine behavior). We show that, in games with general almost-public private monitoring, this is essentially the only behavior that can coordinate continuation play. To make this precise, we must describe what it means for a game to have “general but almost-public private monitoring” and “essentially.”

Since the coordination-of-continuation-play interpretation depends on the structure of the strategy profile, we focus on equilibrium strategy profiles, rather than on the equilibrium payoff set, of private-monitoring games. Very little is known about the general structure of the equilibrium payoff set for general private-monitoring games. We return to this issue at the end of the Introduction.

Fix a game with full support public monitoring (so that every signal arises with strictly positive probability under every action profile). In the minimal perturbation of the public-monitoring game towards private monitoring, each player observes a *private* signal drawn from the space of public signals, and the other specifications of the game are unchanged. In this private-monitoring game, at the end of each period, there is a profile of private signals, and we say the game has *minimally-private almost-public monitoring* if the probability of any profile in which all players observe the same value of the signal is close to the probability of that signal in the public-monitoring game (there is also positive probability that different players observe different values of the public signal).

Any strategy profile of a public-monitoring game naturally induces behavior in minimally-private almost-public-monitoring games.¹ Mailath and Morris (2002) introduced a useful representation device for these profiles. Recall that all PPE of a public-monitoring game 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 (1990)). 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) is *not* sufficient to verify that the strategies are optimal. It is also necessary to know how each player's beliefs over the private states of other players evolve. This is at the heart of the question of whether histories can coordinate continuation play, since, given a strategy profile, a player's private state determines that player's continuation play. 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 all other players being in the *same* private state (Mailath and Morris (2002, Theorem 4.1)). PPE with bounded recall satisfy this sufficient condition, since for sufficiently close-by games with minimally-private almost-public monitoring, the probability that all players observed the same last L signals can be made arbitrarily close to one. However, under other strategy profiles, the condition may fail. The grim trigger PPE in some parameterizations of the repeated prisoners' dilemma, for example, does not induce an equilibrium in *any* close-by minimally-private almost-public-monitoring game (Example 2 in Section 3.1).

The restriction to minimally-private almost-public monitoring is substantive, since all players' private signals are drawn from a common signal space. In this paper, we allow for the most general private monitoring consistent with the game being "close-to" a public-monitoring game. We assume there is a signalling function for each player that assigns to each private signal either some value of the public signal or a dummy signal (with the interpretation that that private signal cannot be related to any public signal). Using these signalling functions (one for each player), there is a natural sense in which the private monitoring distribution can be said to be close to the public monitoring distribution, even when the sets of private signals differ, and may have significantly larger cardinality than that of the set of public signals. We say such games have *almost-public monitoring*. If every private signal is mapped to a public signal, we say the almost-public-monitoring game is *strongly close* to the public-monitoring game.

Using the signalling functions, any strategy profile of the public-monitoring game induces behavior in strongly-close-by almost-public-monitoring games. As in minimally-private almost-public-monitoring games, a player's private state determines that player's continuation play.

¹Since player i 's set of histories in the public-monitoring game and in the minimally-private almost-public-monitoring game agree, the domains for player i 's strategy in the two games also agree.

Given a sequence of private signals for a player, that player's private state is determined by the induced sequence of public signals that are the result of applying his signalling function. Consequently, it might appear that the richness of the private signals does not alter the situation from the case of minimally-private almost-public monitoring. However, the richness of the private signals is important for the formation of that player's beliefs about the other players' private states. It turns out that the requirement that the private-monitoring distribution be close to the public-monitoring distribution places essentially no restriction on the manner in which private signals enter into the formation of posterior beliefs. Nonetheless, if the profile has bounded recall, the richness of the private signals is irrelevant. Indeed, even if the private-monitoring games are not strongly close to the public-monitoring game, there is still a natural sense in which every strict PPE with bounded recall induces equilibrium behavior in every close-by almost-public-monitoring game (Theorem 2).

When a strategy profile of the public-monitoring game does not have bounded recall, realizations of the signal in early periods can have long-run implications for behavior. Subject to some technical caveats, we call such a profile *separating*. While the properties of bounded recall and separation do not exhaust possible behavior, they do appear to cover most behaviors of interest.² When the space of private signals is sufficiently *rich* in the values of posterior-odds ratios (this is what we mean by “general almost public”), and the profile is separating, it is possible to manipulate a player's updating over other players' private states through an appropriate choice of private history. This suggests that it should be possible to choose a private history with the property that a player (say, player i) is in one private state and assigns arbitrarily high probability to all the other players being in a different common private state.

There is a significant difficulty that needs to be addressed in order to make this argument: The history needs to have the property that player i is very confident of the other players' state transitions for any given initial state. This, of course, requires the monitoring be almost-public. At the same time, monitoring must be sufficiently imprecise that player i , after an appropriate initial segment of the history, assigns positive probability to the other players being in a common state different from i 's private state. This is the source of the difficulty: for any T -length history, there is an ε (decreasing in T) such that for private monitoring ε -close to the public monitoring, player i is sufficiently confident of the period T private states of players $j \neq i$ as a function of their period $t < T$ private states (and the history). However, this ε puts an upper bound on the prior probability that player i can assign in period t to the players $j \neq i$ being in a common state different from i 's private state. Since the choice of T is decreasing in this prior (i.e., larger T required for smaller priors), there is a tension in the determination of T and ε .

We show, however, that any separating profile implementable using a finite number of states has enough structure that we can choose the history so that not only do relevant states cycle, but that every other state transits under the cycle to a cycling state. The cycle allows us to

²We provide one example of a non-separating profile without bounded recall in Section 5 (Example 6). This profile is not robust to the introduction of even minimally-private monitoring.

effectively choose the T above independently of the prior, and gives us our main result (Theorem 4): Separating strict PPE profiles of public-monitoring games implementable using a finite number of states do not induce Nash equilibria in *any* strongly-close-by games with rich private monitoring.

Thus, separating strict PPE of public-monitoring games are not robust to the introduction of even a minimal amount of private monitoring. Consequently, separating behavior in private-monitoring games typically cannot coordinate continuation play (Corollary 1). On the other hand, bounded recall profiles are robust to the introduction of private monitoring. The extent to which bounded recall is a substantive restriction on the set of payoffs is unknown.³ Our results do suggest, even for public-monitoring games, that bounded recall profiles are particularly attractive (since they are robust to the introduction of private monitoring). Moreover, other apparently simple strategy profiles are problematic.

Our focus on equilibrium strategy profiles is in contrast with much of the literature in repeated games with private monitoring.⁴ For the repeated prisoners' dilemma with almost-perfect private monitoring, folk theorems have been proved using both equilibria with a coordination interpretation (for example, Sekiguchi (1997), which we discuss in Example 1, and Bhaskar and Obara (2002)) and those that are "belief-free" (for example, Piccione (2002), Ely and Välimäki (2002), and Matsushima (2004)⁵). Loosely, belief-free equilibria are constructed so that after relevant histories, players are indifferent between different choices. In games with finite signal spaces, this requires a significant amount of randomization (randomization is not required with a continuum of signals, but only because behavior can be purified using signals). Not only is the generality of this approach unclear (Ely, Hörner, and Olszewski (2005)), these equilibria do not have a clean coordination interpretation due to the extent of player indifference.

Finally, we view our findings as underlining the importance of communication in private-monitoring games as a mechanism to facilitate coordination. For some recent work on communication in private-monitoring games, see Compte (1998), Kandori and Matsushima (1998), Fudenberg and Levine (2004), and McLean, Obara, and Postlewaite (2002).

³Cole and Kocherlakota (forthcoming) show that for some parameterizations of the repeated prisoners' dilemma, the restriction to strongly symmetric bounded recall PPE results in a dramatic collapse of the set of equilibrium payoffs.

⁴See Kandori (2002) for a brief survey of this literature, as well as the accompanying symposium issue of the *Journal of Economic Theory* on "Repeated Games with Private Monitoring."

⁵Matsushima (2004) covers some two player games with private monitoring that need not be almost perfect, with signals that are either conditionally independent or have a particular correlation structure. His analysis does not cover almost-public-monitoring games.

2. Games with Imperfect Monitoring

2.1. Private-Monitoring Games

The infinitely-repeated game with private monitoring is the infinite repetition of a stage game in which at the end of the period, each player learns only the realized value of a private signal. There are n players, with a finite stage-game action set for player $i \in N \equiv \{1, \dots, n\}$ denoted A_i . At the end of each period, each player i observes a private signal, denoted ω_i drawn from a finite set Ω_i . The signal vector $\omega \equiv (\omega_1, \dots, \omega_n) \in \Omega \equiv \Omega_1 \times \dots \times \Omega_n$ occurs with probability $\pi(\omega|a)$ when the action profile $a \in A \equiv \prod_i A_i$ is chosen. Player i does not receive any information other than ω_i about the behavior of the other players. All players use the same discount factor, δ .

Since ω_i is the only signal a player observes about opponents' play, we assume (as usual) that player i 's payoff after the realization (ω, a) is given by $u_i^*(\omega_i, a_i)$. Stage game payoffs are then given by $u_i(a) \equiv \sum_{\omega} u_i^*(\omega_i, a_i) \pi(\omega|a)$. It will be convenient to index games by the monitoring technology (Ω, π) , fixing the set of players and action sets.

A pure strategy for player i in the private-monitoring game is a function $s_i : \mathcal{H}_i \rightarrow A_i$, where

$$\mathcal{H}_i \equiv \cup_{t=1}^{\infty} (A_i \times \Omega_i)^{t-1}$$

is the set of private histories for player i .

Definition 1 A pure strategy is action-free if, for all $h_i^t, \hat{h}_i^t \in \mathcal{H}_i$ satisfying $\omega_i^\tau = \hat{\omega}_i^\tau$ for all $\tau \leq t$,

$$s_i(h_i^t) = s_i(\hat{h}_i^t).$$

Since action-free strategies play a central role in our analysis, it is useful to note the following immediate result, which does not require full-support monitoring (its proof is omitted):

Lemma 1 Every pure strategy in a private-monitoring game is realization equivalent to an action-free strategy. Every mixed strategy is realization equivalent to a mixture over action-free strategies.

Remark 1 Behavior strategies realization equivalent to a mixed strategy will typically not be action-free. For example, consider the once repeated prisoners' dilemma, with action spaces $A_i = \{e_i, n_i\}$,⁶ $\Omega_i = \{g_i, b_i\}$, and the mixed strategy assigning equal probability to the two action-free strategies \bar{s}_1 and \tilde{s}_1 , where

$$\bar{s}_1(\emptyset) = e_1; \bar{s}_1(g_1) = e_1, \bar{s}_1(b_1) = n_1,$$

and

$$\tilde{s}_1(\emptyset) = n_1; \tilde{s}_1(g_1) = n_1, \tilde{s}_1(b_1) = n_1.$$

⁶Interpreting the prisoners' dilemma as a partnership game, e_i is "exert effort," while n_i is "no effort."

	e_2	n_2
e_1	2, 2	-1, 3
n_1	3, -1	0, 0

Figure 1: The prisoners' dilemma.

A behavior strategy realization equivalent to this mixed strategy must specify in the second period behavior that depends nontrivially on player 1's first period action. (A similar observation applies to public-monitoring games: every pure strategy is realization equivalent to a public strategy, every mixed strategy is realization equivalent to a mixture over public strategies, and yet all behavior strategies that are realization equivalent to a mixed strategy may not be public.)

Every pure action-free strategy can be represented by a set of states W_i , an initial state w_i^1 , a decision rule $d_i : W_i \rightarrow A_i$ specifying an action choice for each state, and a transition function $\sigma_i : W_i \times \Omega_i \rightarrow W_i$. In the first period, player i chooses action $a_i^1 = d_i(w_i^1)$. At the end of the first period, the vector of actions, a^1 , then generates a vector of private signals ω^1 according to the distribution $\pi(\cdot | a^1)$, and player i observes the signal ω_i^1 . In the second period, player i chooses the action $a_i^2 = d_i(w_i^2)$, where $w_i^2 = \sigma_i(w_i^1, \omega_i^1)$, and so on. Any action-free strategy requires at most the countable set $W_i = \cup_{t=1}^{\infty} \Omega_i^{t-1}$.

Any collection of pure action-free strategies can be represented by a set of states W_i , a decision rule d_i , and a transition function σ_i (the initial state indexes the pure strategies). One class of mixed strategies is described by $(W_i, \mu_i, d_i, \sigma_i)$, where μ_i is a probability distribution over the initial state w_i^1 , and W_i is countable. Not all mixed strategies can be described in this way, since the set of all pure strategies is uncountable (which would require W_i to be uncountable).

Remark 2 A consequence of Remark 1 is that action-free strategy profiles, and profiles of mixtures over action-free strategies, are often not sequentially rational. However, when the monitoring has full support, every Nash equilibrium has a realization-equivalent sequentially rational strategy profile (see Sekiguchi (1997, Proposition 3) and Kandori and Matsushima (1998, p. 648)). Consequently, we focus on Nash equilibria of games with private monitoring.

Example 1 We will often use the repeated prisoners' dilemma under various monitoring assumptions. The ex ante stage game is given by the normal form in Figure 1.⁷ Much of the literature has studied *almost-perfect conditionally-independent private monitoring*: player i 's signals are given by $\Omega_i = \{\hat{e}_i, \hat{n}_i\}$, with $\hat{a}_i \in \Omega_i$ a signal of $a_j \in A_j \equiv \{e_j, n_j\}$. Players 1 and 2's

⁷Here (and in other examples) we follow the literature in assuming the ex ante payoff matrix is independent of the monitoring distribution. This simplifies the discussion and is without loss of generality: Ex ante payoffs are close when the monitoring distributions are close (Lemma 4) and all relevant incentive constraints are strict.

signal are, conditional on the action profile, independently distributed, with

$$\pi(\hat{a}_1 \hat{a}_2 | a_1 a_2) = \pi_1(\hat{a}_1 | a_2) \pi_2(\hat{a}_2 | a_1)$$

and

$$\pi_i(\hat{a}_i | a_j) = \begin{cases} 1 - \varepsilon, & \text{if } \hat{a}_i = a_j, \\ \varepsilon, & \text{if } \hat{a}_i \neq a_j, \end{cases}$$

where $\varepsilon > 0$ is a small constant. As will be clear, we focus on a different class of private monitoring distributions.

In an important article, Sekiguchi (1997) constructed an efficient equilibrium for the almost-perfect conditionally-independent case (as well as for correlated but almost-perfect monitoring). Let $W_i = \{w^e, w^n\}$, $\sigma_i(w^n, \hat{a}_i) = w^n$ for all \hat{a}_i , $\sigma_i(w^e, \hat{e}_i) = w^e$, and $\sigma_i(w^e, \hat{n}_i) = w^n$. The pure strategy of grim trigger (begin playing e_i , and continue to play e_i as long as \hat{e}_i is observed, switch to n_i after \hat{n}_i and always play n_i thereafter) is induced by the initial state $w_i^1 = w^e$. The pure strategy of always play n_i is induced by the initial state $w_i^1 = w^n$. The critical insight in Sekiguchi (1997) is that while grim trigger is not a Nash equilibrium of this game, the symmetric mixed strategy profile where each player independently randomizes over initial states w^e and w^n is an equilibrium (as long as δ is not too close to 1). Sekiguchi (1997) then constructs an equilibrium for larger δ by treating the game as M distinct games, with the k^{th} game played in periods $k + tM$, for $t \in \mathbb{N}$. The mixed equilibrium for $M = 3$ is constructed from the machine in Figure 2. The state w^{ene} for example corresponds to grim trigger in “games” 1 and 3, and always n_i in game 2.

2.2. Public-Monitoring Games

We turn now to the benchmark public-monitoring game for our games with private monitoring. The finite action set for player $i \in N$ is again A_i . The public signal is denoted y and is drawn from a finite set Y . The probability that the signal y occurs when the action profile $a \in A \equiv \prod_i A_i$ is chosen is denoted $\rho(y|a)$. We refer to (Y, ρ) as the public-monitoring distribution. Player i 's payoff after the realization (y, a) is given by $\tilde{u}_i^*(y, a_i)$. Stage game payoffs are then given by $\tilde{u}_i(a) \equiv \sum_y \tilde{u}_i^*(y, a_i) \rho(y|a)$. 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 .

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.⁸ Henceforth, by the term *public profile*, we will always mean a strategy profile for the public-monitoring game that is itself public. A *perfect public equilibrium (PPE)* is a profile of public strategies that, after observing any public history h^t ,

⁸Note that strategies of public-monitoring games are public if and only if they are action-free when we view the public-monitoring game as a game with (trivial) private monitoring.




Figure 2: The automaton described the pure strategies in Sekiguchi (1997) for $M = 3$. The decision rules are $d_i(w^{abc}) = a_i$. Unlabeled arrows are unconditional transitions.

specifies a Nash equilibrium for the repeated game. Under imperfect full-support public monitoring, every public history arises with positive probability, and so every Nash equilibrium in public strategies is a 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 $\sigma : W \times Y \rightarrow W$, and a collection of decision rules, $d_i : W \rightarrow A_i$. In the first period, player i chooses action $a_i^1 = d_i(w^1)$. The vector of actions, a^1 , then generates a signal y^1 according to the distribution $\rho(\cdot|a^1)$. In the second period, player i chooses the action $a_i^2 = d_i(w^2)$, where $w^2 = \sigma(w^1, y^1)$, and so on. Since we can take W to be the set of all histories of the public signal, $\cup_{k \geq 0} Y^k$, W is at most countably infinite. A public profile is *finite* if W is a finite set. Note that, given a pure strategy profile (and the associated automaton), continuation play after any history is determined by the *public* state reached by that history. In games with private monitoring, by contrast, given an action-free strategy profile (and the associated automaton), a sufficient statistic for continuation play after any history is the vector of current *private* states, one for each player.

Denote the vector of average discounted expected values of following the public profile (W, w, σ, d) (i.e., the initial state is w) by $\phi(w)$. Define a function $g : A \times W \rightarrow W$ by $g(a; w) \equiv (1 - \delta)u(a) + \delta \sum_y \phi(\sigma(w, y)) \rho(y|a)$. We have (from Abreu, Pearce, and Stacchetti (1990)), that if the profile is an equilibrium, then, for all $w \in W$, the action profile $(d_1(w), \dots, d_N(w)) \equiv d(w)$ is a pure strategy equilibrium of the static game with strategy spaces A_i and payoffs $g_i(\cdot; w)$ for each i and, moreover, $\phi(w) = g(d(w), w)$. Conversely, if (W, w^1, σ, d) describes an equilibrium of the static game with payoffs $g(\cdot; w)$ for all $w \in W$, then the induced pure

strategy profile in the infinitely repeated game with public monitoring is an equilibrium.⁹ A PPE (W, w^1, σ, d) is *strict* if, for all $w \in W$, $d(w)$ is a strict Nash equilibrium of the static game $g(\cdot, \cdot; w)$.¹⁰

A maintained assumption throughout our analysis is that public monitoring has full support.

Assumption 1 $\rho(y|a) > 0$ for all $y \in Y$ and all $a \in A$.

Definition 2 An automaton (W, w^1, σ, d) is minimal if for every state $\hat{w} \in W$, there exists a sequence of signals $\hat{y}^1, \dots, \hat{y}^\ell$ such that $\hat{w} = \sigma(\hat{y}^1, \dots, \hat{y}^\ell; w^1)$, where $\sigma(\hat{y}^1, \dots, \hat{y}^\ell; w^1) \equiv \sigma(\hat{y}^\ell, \sigma(\dots, \sigma(\hat{y}^1, w^1)))$, and for every pair of states $w, \hat{w} \in W$, there exists a sequence of signals y^1, \dots, y^L such that for some i , $d_i(\sigma(y^1, \dots, y^L; w)) \neq d_i(\sigma(y^1, \dots, y^L; \hat{w}))$.

The restriction to minimal automata is without loss of generality: every profile has a minimal representing automaton. Moreover, this automaton is essentially unique.¹¹ Accordingly, we treat a public strategy profile and its minimal representing automaton interchangeably.

3. Almost-Public Monitoring

3.1. Minimally-private almost-public monitoring

Games with public monitoring (Y, ρ) are nested within games with private monitoring, since public monitoring simply means that all players always observe the same signal, i.e., $\Omega_i = \Omega_j = Y$, and $\pi(y, \dots, y|a) = \rho(y|a)$ for all a . Mailath and Morris (2002) discussed the case of minimally-private monitoring, in the sense that there is a public monitoring distribution (Y, ρ) with $\Omega_i = Y$ and π close to ρ :

Definition 3 A private-monitoring game $(u^*, (Y^n, \pi))$ is ε -close to a public-monitoring game $(\tilde{u}^*, (Y, \rho))$, if $|\tilde{u}_i^*(y, a_i) - u_i^*(y, a_i)| < \varepsilon$ and $|\pi((y, \dots, y)|a) - \rho(y|a)| < \varepsilon$ for all $i \in N$, $y \in Y$ and all $a \in A$. We also say that such a private-monitoring game has minimally-private almost-public monitoring.

⁹We have introduced a distinction between W and the set of continuation payoffs for convenience. Any pure strategy equilibrium *payoff* can be supported by an equilibrium where $W \subset \mathfrak{R}^I$ and $\phi(w) = w$ (again, see Abreu, Pearce, and Stacchetti (1990)).

¹⁰Equivalently, a PPE is strict if each player strictly prefers his equilibrium strategy to every other *public* strategy. For a large class of public-monitoring games, strictness is without loss of generality, in that a folk theorem holds for strict PPE (Fudenberg, Levine, and Maskin (1994, Theorem 6.4 and remark)).

¹¹Suppose (W, w^1, σ, d) and $(\tilde{W}, \tilde{w}^1, \tilde{\sigma}, \tilde{d})$ are two minimal automata representing the same public strategy profile. Define a mapping $\varphi: W \rightarrow \tilde{W}$ as follows: Set $\varphi(w^1) = \tilde{w}^1$. For $\hat{w} \in W \setminus \{w^1\}$, let $\hat{y}^1, \dots, \hat{y}^\ell$ be a public history reaching \hat{w} (i.e., $\hat{w} = \sigma(\hat{y}^1, \dots, \hat{y}^\ell; w^1)$), and set $\varphi(\hat{w}) = \tilde{\sigma}(\hat{y}^1, \dots, \hat{y}^\ell; \tilde{w}^1)$. Since both automata are minimal and represent the same profile, φ does not depend on the choice of public history reaching \hat{w} . It is straightforward to verify that φ is one-to-one and onto. Moreover, $\tilde{\sigma}(y, \tilde{w}) = \varphi(\sigma(y, \varphi^{-1}(\tilde{w})))$, and $\tilde{d}(\tilde{w}) = d(\varphi(w))$.

For $\eta > 0$ there is $\varepsilon > 0$ such that if $(u^*, (Y^n, \pi))$ is ε -close to $(\tilde{u}^*, (Y, \rho))$, then

$$\left| \sum_{y_1, \dots, y_n} u_i^*(y_i, a_i) \pi(y_1, \dots, y_n | a) - \sum_y \tilde{u}_i^*(y, a_i) \rho(y | a) \right| < \eta.$$

In other words, the ex ante stage payoffs of any minimally-private almost-public-monitoring game are close to the ex ante stage payoffs of the benchmark public-monitoring game.

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

Lemma 2 *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 $a \in A$ and $y \in Y$,*

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

A public strategy profile (W, w^1, σ, d) in the public-monitoring game induces a strategy profile (s_1, \dots, s_n) in minimally-private almost-public-monitoring games in the obvious way: $s_i^1 = d_i(w^1)$, $s_i^2(a_i^1, y_i^1) = d_i(\sigma(w^1, y_i^1)) \equiv d_i(w_i^2)$, and defining states recursively by $w_i^{t+1} \equiv \sigma(w_i^t, y_i^t)$, for $h_i^t \equiv (a_i^1, y_i^1; a_i^2, y_i^2; \dots; a_i^{t-1}, y_i^{t-1}) \in (A \times Y)^{t-1}$, $s_i^t(h_i^t) = d_i(w_i^t)$. This private strategy is, of course, action-free.

If W is finite, each player can be viewed as following a finite state automaton. Hopefully without confusion, when we can take the initial state as given, we abuse notation and write $w_i^t = \sigma(w^1, h_i^t) = \sigma(h_i^t)$. 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 private-monitoring game of the public profile (of the public-monitoring game).

If player i believes that the other players are following a strategy that was induced by a public profile, then a sufficient statistic of h_i^t for the purposes of evaluating continuation strategies is player i 's private state and i 's beliefs over the other players' private states, i.e., (w_i^t, β_i^t) , where $\beta_i^t \in \Delta(W^{N-1})$. In principle, W may be quite large. For example, if the public strategy profile is nonstationary, it may be necessary to take W to be the set of all histories of the public signal, $\cup_{k \geq 0} Y^k$. On the other hand, the strategy profiles typically studied can be described with a significantly more parsimonious collection of states, often finite. When W is finite, the need to only keep track of each player's private state and that player's beliefs over the other players' private states is a considerable simplification, as the following result (Mailath and Morris (2002, Theorem 4.2)) demonstrates.

Theorem 1 *Suppose the public profile (W, w^1, σ, d) is a strict equilibrium of the full-support public-monitoring game for some δ and $|W| < \infty$. For all $\kappa > 0$, there exists η and ε such that in any game with minimally-private almost-public monitoring, if the posterior beliefs induced by the private profile satisfy $\beta_i(\sigma(h_i^t) \mathbf{1} | h_i^t) > 1 - \eta$ for all $h_i^t = (d_i(w^1), y_i^1; d_i(w_i^2), y_i^2; \dots; d_i(w_i^{t-1}), y_i^{t-1})$, where $w_i^{t+1} \equiv \sigma(w_i^t, y_i^t)$, 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 and the expected payoff in that equilibrium is within κ of the public equilibrium payoff.*

Example 2 We return to the repeated prisoners' dilemma, with ex ante stage game given by Figure 1 (recall footnote 7). In the benchmark public-monitoring game, the set of public signals is $Y = \{y, \bar{y}\}$ and public monitoring distribution is

$$\rho(\bar{y} | a_1 a_2) = \begin{cases} p, & \text{if } a_1 a_2 = e_1 e_2 \\ q, & \text{if } a_1 a_2 = e_1 n_2 \text{ or } n_1 e_2, \\ r, & \text{if } a_1 a_2 = n_1 n_2. \end{cases}$$

The grim trigger strategy profile for the public-monitoring game is described by the automaton $W = \{w^e, w^n\}$, initial state w^e , decision rules $d_i(w^a) = a_i$, and transition rule

$$\sigma(w, y) = \begin{cases} w^e, & \text{if } y = \bar{y} \text{ and } w = w^e, \\ w^n, & \text{otherwise.} \end{cases}$$

Grim trigger is a strict PPE if $\delta > (3p - 2q)^{-1} > 0$ (a condition we maintain throughout this example). We turn now to minimally-private-monitoring games that are ε -close to this public-monitoring game. It turns out that, for ε small, grim trigger induces a Nash equilibrium in such games if $q < r$, but not if $q > r$. Consider first the case $q > r$ and the private history $(e_1 \underline{y}_1, n_1 \bar{y}_1, n_1 \bar{y}_1, \dots, n_1 \bar{y}_1)$. We now argue that, after a sufficiently long such history, the grim trigger specification of n_1 is not optimal. Intuitively, while player 1 has transited to the private state w_1^n , player 1 always puts strictly positive (but perhaps small) probability on his opponent being in private state w_2^e . Since $q > r$ (and ε is small), the private signal \bar{y}_1 after playing n_1 is an indication that player 2 had played e_2 (rather than n_2), and so player 1's posterior that player 2 is still in w_2^e increases. Eventually, player 1 is sufficiently confident of player 2 still being in w_2^e that he finds n_1 suboptimal. On the other hand, when $q \leq r$, such a history is not problematic because it reinforces 1's belief that 2 is also in w_2^n . Two other histories are worthy of mention: $(e_1 \underline{y}_1, n_1 \underline{y}_1, n_1 \underline{y}_1, \dots, n_1 \underline{y}_1)$ and $(e_1 \bar{y}_1, e_1 \bar{y}_1, e_1 \bar{y}_1, \dots, e_1 \bar{y}_1)$. Under the first history, while the signal \underline{y}_1 is now a signal that 2 had chosen e_2 in the previous period, for ε small, 1 is confident that 2 also observed \underline{y}_2 and so will transit to w_2^n . For the final history, the signal \bar{y}_1 continually reassures 1 that 2 is still playing e_2 , and so e_1 remains optimal. (See Mailath and Morris (2002, Section 3.3) for the calculations underlying this discussion.)

Example 3 As the players become patient, the payoffs from grim trigger converge to $(0, 0)$. A grim trigger profile (i.e., a profile in which the specification of n_i is absorbing) can only achieve




Figure 3: Forgiving grim trigger where any two realizations of \underline{y} lead to w^n .

significant payoffs for patient players by being forgiving.¹² Such a profile provides a different example of how a strict PPE can fail to induce a Nash equilibrium in close-by minimally-private-monitoring games. The simplest forgiving profile requires two realizations of \underline{y} to switch to $n_1 n_2$. The automaton for this profile has a set of states $W = \{w^e, \hat{w}^e, w^n\}$, initial state w^e , decision rules $d_i(w^a) = a_i$ and $d_i(\hat{w}^e) = e_i$, and transition function

$$\sigma(w, y) = \begin{cases} w^e, & \text{if } y = \bar{y} \text{ and } w = w^e, \\ \hat{w}^e, & \text{if } y = \underline{y} \text{ and } w = w^e \text{ or } y = \bar{y} \text{ and } w = \hat{w}^e, \\ w^n, & \text{otherwise.} \end{cases}$$

The profile is illustrated in Figure 3. This PPE never induces a Nash equilibrium in close-by minimally-private-monitoring games: consider a private history in which player 1 plays e_1 and observes \bar{y}_1 for T periods, and then observes \underline{y}_1 . Under the forgiving profile, player 1 is supposed to switch to the private state \hat{w}_1^e and continue to play e_1 (until another \underline{y}_1 is observed). But, for large T , it is more likely that player 2 has observed \underline{y}_2 in exactly one of the first T periods than having observed \bar{y}_2 in every period.¹³ Consequently, for large T , player 1 will not find e_1 optimal. Clearly, the same analysis applies to forgiving grim triggers that require more realizations of \underline{y} to switch to w^n .

Another class of forgiving grim trigger profiles requires successive realizations of \underline{y} to switch to w^n . In the three state version, the automaton is identical to that above except $\sigma(\hat{w}^e, \bar{y}) = w^e$ (see Figure 4). The analysis of this profile is similar to that of Example 2. The profile does not induce a Nash equilibrium in close-by minimally-private-monitoring games if $q > r$ for similar reasons. There are now two possibilities for the case $q \leq r$, since isolated observations of \underline{y}_1 do not lead to w_2^n . For the histories considered in Example 2, the same argument applies once we note that, conditional on players being in one of w^e or \hat{w}^e , a player assigns very high probability to the other player being in the same state, since this is determined by the last signal.

¹²This is the class of profiles studied by Compte (2002) for the conditionally-independent private-monitoring prisoners' dilemma.

¹³This type of drift of beliefs is a general phenomenon when players choose the same action in adjacent states (see also Example 6).




Figure 4: Forgiving grim trigger where two successive realizations of \underline{y} lead to w^n .

The remaining histories are those with isolated observations of \underline{y}_1 . The critical history (since it contains the largest fraction of \underline{y}_1 's consistent with e_1) is $(e_1 \bar{y}_1, e_1 \underline{y}_1, e_1 \bar{y}_1, e_1 \underline{y}_1, \dots, e_1 \bar{y}_1)$, that is, alternating \underline{y}_1 and \bar{y}_1 . If $p(1-p) \geq q(1-q)$, then such a history (weakly) indicates that player 2 is still playing e_2 , while the reverse strict inequality indicates that player 2 is playing n_2 . Summarizing, the profile induces a Nash equilibrium in close-by minimally-private-monitoring games if and only if $q \leq r$ and $p(1-p) \geq q(1-q)$.

3.2. General almost-public monitoring

We now turn to the most general private monitoring structure that nonetheless preserves the essential characteristics of both Definition 3 and Lemma 2.¹⁴

Definition 4 *The private monitoring distribution (Ω, π) is ε -close under f to the public monitoring distribution (Y, ρ) , where $f = (f_1, \dots, f_n)$ is a vector of signaling functions $f_i : \Omega_i \rightarrow Y \cup \{\emptyset\}$, if*

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

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

and

¹⁴While there is a connection to informational smallness (see, for example, McLean and Postlewaite (2004)), these are distinct notions. For concreteness, suppose ω_i is a noisy signal of y . Then, (Ω, π) is ε -close to (Y, ρ) if and only if the private signal is a sufficiently accurate signal of y . A player is *informationally small* if the posterior on y , conditional on the other players' private signals, on average does not vary too much with that player's private signal. Even if each player's private signal is very accurate, the posterior can vary dramatically in a player's signal if that player's signal is sufficiently accurate relative to the other players. Moreover, if there are many players, even when signals are very noisy, each player will be informationally small.

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

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

The private monitoring distribution (Ω, π) is strongly ε -close under f to the public monitoring distribution (Y, ρ) if it is ε -close under f , and in addition, all the signaling functions map into Y .

A private monitoring distribution (Ω, π) is (strongly) ε -close to the public monitoring distribution (Y, ρ) if there exists a vector of signaling functions f such that (Ω, π) is (strongly) ε -close under f to (Y, ρ) .

If the private monitoring is ε -close under f , but not strongly ε -close under f , then some private signals are not associated with any public signal: there is a signal ω_i satisfying $f_i(\omega_i) = \emptyset$. Such an “uninterpretable” signal may contain *no* information about the signals observed by the other players.

Note that the second condition implies that every player has at least one private signal mapped to each public signal. Moreover, for the case $\Omega_i = Y$, the first condition implies the second (Lemma 2).

The condition of ε -closeness in Definition 4 can be restated as follows. Recall from Monderer and Samet (1989) that an event is *p-evident* if, whenever it is true, everyone assigns probability at least p to it being true. The following Lemma is a straightforward application of the definitions, and so we omit the proof.

Lemma 3 Suppose $f_i : \Omega_i \rightarrow Y \cup \{\emptyset\}$, $i = 1, \dots, n$, is a collection of signaling functions. The private monitoring distribution (Ω, π) is ε -close under f to the public monitoring distribution (Y, ρ) if and only if for each public signal y , the set of private signal profiles $\{\omega : f_i(\omega_i) = y \text{ for all } i\}$ is $(1 - \varepsilon)$ -evident (conditional on any action profile) and has probability within ε of the probability of y (conditional on that action profile).

Example 4 We now allow player 1 to have a richer set of private signals, $\Omega_1 = \{y_1, \bar{y}'_1, \bar{y}''_1\}$, keeping player 2 signals unchanged, $\Omega_2 = \{y_2, \bar{y}_2\}$. The probability distribution of the signals is given in Figure 5. This private-monitoring distribution is $\sqrt{\varepsilon}$ -close to the public-monitoring distribution of Example 2 under the signaling functions $f_i(y_i) = \underline{y}$ and $f_2(\bar{y}_2) = f_1(\bar{y}'_1) = f_1(\bar{y}''_1) = \bar{y}$, as long as ε is sufficiently small, relative to $\min\{\alpha', \alpha - \alpha'\}$.

Definition 5 A private-monitoring game $(u^*, (\Omega, \pi))$ is ε -close (under f) to the public-monitoring game $(\tilde{u}^*, (Y, \rho))$, if (Ω, π) is ε -close under f to (Y, ρ) and

$$|\tilde{u}_i^*(f_i(\omega_i), a_i) - u_i^*(\omega_i, a_i)| < \varepsilon$$

for all $i \in N$, $a_i \in A_i$, and $\omega_i \in f_i^{-1}(Y)$. We will also say that such a private-monitoring game has almost-public monitoring.

$a_1 a_2$	\underline{y}_2	\bar{y}_2
\underline{y}_1	$(1 - \alpha)(1 - 3\varepsilon)$	ε
\bar{y}'_1	ε	$\alpha'(1 - 3\varepsilon)$
\bar{y}''_1	ε	$(\alpha - \alpha')(1 - 3\varepsilon)$

Figure 5: The probability distribution of the private signals for Example 4. The distribution is given as a function of the action profile $a_1 a_2$, where $\alpha = p$ if $a_1 a_2 = e_1 e_2$, q if $a_1 a_2 = e_1 n_2$ or $n_1 e_2$, and r if $a_1 a_2 = n_1 n_2$ (analogously, α' is given by p' , q' , or r' as a function of $a_1 a_2$). All probabilities are strictly positive.

As above, the ex ante stage payoffs of any almost-public-monitoring game are close to the ex ante stage payoffs of the benchmark public-monitoring game (the proof is in the Appendix).

Lemma 4 *For all $\eta > 0$, there is $\varepsilon > 0$ such that if $(u^*, (\Omega, \pi))$ is ε -close to $(\tilde{u}^*, (Y, \rho))$, then*

$$\left| \sum_{\omega_1, \dots, \omega_n} u_i^*(\omega_i, a_i) \pi(\omega_1, \dots, \omega_n | a) - \sum_y \tilde{u}_i^*(y, a_i) \rho(y | a) \right| < \eta.$$

Fix a public profile (W, w^1, σ, d) of a full-support public-monitoring game $(\tilde{u}^*, (Y, \rho))$, and, under f , a strongly ε -close private-monitoring game $(u^*, (\Omega, \pi))$. The public profile induces a private profile in the private-monitoring game in a natural way: Player i 's strategy is described by the automaton (W, w^1, σ_i, d_i) , where $\sigma_i(w, \omega_i) = \sigma(w, f_i(\omega_i))$ for all $\omega_i \in \Omega_i$ and $w \in W$. The set of states, initial state, and decision function are from the public profile. The transition function σ_i is well-defined, because the signaling functions all map into Y , rather than $Y \cup \{\emptyset\}$. As for games with minimally-private almost-public monitoring, if player i believes that the other players are following a strategy induced by a public profile, a sufficient statistic of h_i^t for the purposes of evaluating continuation strategies is player i 's private state and i 's beliefs over the other players' private states, i.e., (w_i^t, β_i^t) , where $\beta_i^t \in \Delta(W^{N-1})$. Finally, we can recursively calculate the private states of player i as $w_i^2 = \sigma(w^1, f_i(\omega_i^1)) = \sigma_i(w^1, \omega_i^1)$, $w_i^3 = \sigma_i(w_i^2, \omega_i^2)$, and so on. Thus, for any private history h_i^t , we can write $w_i^t = \sigma_i(h_i^t)$.

Example 5 In Example 2, we argued that if $q < r$, grim trigger induces Nash equilibrium behavior in close-by minimally-private-monitoring games. We now argue that under the private monitoring distribution of Example 4, even if $q < r$, grim trigger will not induce a Nash equilibrium behavior in some close-by games. In particular, suppose $0 < r' < q' < q < r$. Under this parameter restriction, the signal \bar{y}''_1 after n_1 is indeed a signal that player 2 had also played n_2 . However, the signal \bar{y}'_1 after n_1 is a signal that player 2 had played e_2 and so a sufficiently long private history of the form $(e_1 \underline{y}_1, n_1 \bar{y}'_1, n_1 \bar{y}''_1, \dots, n_1 \bar{y}'_1)$ will lead to a posterior for player 1 at which n_1 is not optimal.

4. PPE with bounded recall

As we saw in Example 5, arbitrary public equilibria need not induce equilibria of almost-public-monitoring games, because the public state in period t is determined, in principle, by the entire history h^t . For profiles that have bounded recall, the entire history is not needed, and equilibria in bounded recall strategies will induce equilibria in almost-public-monitoring games.¹⁵

Definition 6 *A public profile s has bounded recall if there exists L such that for all $h^t = (y^1, \dots, y^{t-1})$ and $\hat{h}^t = (\hat{y}^1, \dots, \hat{y}^{t-1})$, if $t > L$ and $y^\tau = \hat{y}^\tau$ for $\tau = t-L, \dots, t-1$, then*

$$s(h^t) = s(\hat{h}^t).$$

Let W_t be the set of states reachable in period t , $W_t \equiv \{w \in W : w = \sigma(w^1, y^1, y^2, \dots, y^{t-1}) \text{ for some } (y^1, y^2, \dots, y^{t-1}), \text{ where } w^1 \text{ is the initial state}\}$. The following characterization of bounded recall (proved in the Appendix) is useful.

Lemma 5 *The public profile induced by the minimal automaton (W, w^1, σ, d) has bounded recall if and only if there exists L such that, for all t and $w, w' \in W_t$ and for all $h \in Y^\infty$,*

$$\sigma(w, h^L) = \sigma(w', h^L).$$

Fix a strict public equilibrium with bounded recall, (W, w^1, σ, d) . Fix a private monitoring technology (Ω, π) ε -close under f to (Y, ρ) . Following Monderer and Samet (1989), we first consider a *constrained game* where behavior after “uninterpretable signals” is arbitrarily fixed. Define the set of “uninterpretable” private histories, $H_i^u = \{h_i^t : \omega_i^\tau \in f_i^{-1}(\emptyset), \text{ some } \tau \text{ satisfying } t-L \leq \tau \leq t-1\}$. This is the set of private histories for which in any of the last L periods, a private signal ω_i^τ satisfying $f_i(\omega_i^\tau) = \emptyset$ is observed. We fix *arbitrarily* player i 's action after any private history $h_i^t \in H_i^u$. For any private history that is not uninterpretable, each of the last L observations of the private signal can be associated with a public signal by the function f_i . Denote by $w_i(h_i^t)$ the private state so obtained. That is,

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

for all $h_i^t \notin H_i^u$. We are then left with a game in which in period $t \geq 2$ player i only chooses an action after a signal ω_i^{t-1} yields a private history not in H_i^u . We claim that for ε sufficiently

¹⁵Denote a dummy signal by $*$. Mailath and Morris (2002) used the term *bounded memory* for public profiles with the property that there is an integer L such that a representing automaton is give by $W = (Y \cup \{*\})^L$, $\sigma(y, (y^L, \dots, y^2, y^1)) = (y, y^L, \dots, y^2)$ for all $y \in Y$, and $w^1 = (*, \dots, *)$. Our earlier notion implicitly imposes a time homogeneity condition, since the caveat in Lemma 5 that the two states should be reachable in the same period is missing. The strategy profile in which play alternates between the same two action profiles in odd and even periods has bounded recall, but not bounded memory.

small, the profile $(\hat{s}_1, \dots, \hat{s}_N)$ is an equilibrium of this constrained game, where \hat{s}_i is the strategy for player i :

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

But this follows from arguments almost identical to that in the proofs of Mailath and Morris (2002, Theorems 4.2 and 4.3): since a player's behavior depends only on the last L signals, for small ε , after observing a history $h_i^t \notin H_i^u$, player i assigns a high probability to player j observing a signal that leads to the same private state. The crucial point is that for ε small, the specification of behavior after signals ω_i satisfying $f_i(\omega_i) = \emptyset$ is irrelevant for behavior at signals ω_i satisfying $f_i(\omega_i) \in Y$. It remains to specify optimal behavior after signals ω_i satisfying $f_i(\omega_i) = \emptyset$. So, consider a new constrained game where player i is required to follow \hat{s}_i where possible. This constrained game has an equilibrium, and so by construction, we thus have an equilibrium of the unconstrained game. We have thus proved:

Theorem 2 *Fix a full-support public-monitoring game $(\tilde{u}^*, (Y, \rho))$ and a strict public perfect equilibrium, \tilde{s} , with bounded recall L . There exists $\varepsilon > 0$ such that for all private-monitoring games $(u^*, (\Omega, \pi))$ ε -close under f to $(\tilde{u}^*, (Y, \rho))$,*

1. *if $f_i(\Omega_i) = Y$ for all i , the induced private profile is a Nash equilibrium; and*
2. *if $f_i(\Omega_i) \neq Y$ for some i , there is a Nash equilibrium of the private-monitoring game, s , such that, for all $h^t = (y^1, \dots, y^{t-1})$ and $h_j^t = (\omega_j^1, \dots, \omega_j^{t-1})$, if $t > L$ and $y^\tau = f_j(\omega_j^\tau)$ for $\tau = t - L, \dots, t - 1$, then*

$$s_j(h_j^t) = \tilde{s}_j(h^t)$$

for all j . Moreover, for all $\kappa > 0$, ε can be chosen sufficiently small that the expected payoff to each player under s is within κ of their public equilibrium payoff.

We could similarly extend our results on patiently-strict, connected, finite public profiles (Mailath and Morris (2002, Theorem 5.1)) and on the almost-public almost-perfect folk theorem (Mailath and Morris (2002, Theorem 6.1)) to this more general notion of nearby private-monitoring distributions.

5. Failure of Coordination

Example 5 illustrates that updating in almost-public-monitoring games can be very different than would be expected from the underlying public-monitoring game. In this section, we build on that example to show that when the set of signals is sufficiently rich (in a sense to be defined), many profiles fail to induce equilibrium behavior in almost-public-monitoring games.

Our negative results are based on the following converse to Theorem 1 (the proof is in the Appendix). Since the theorem is negative, the assumption of strong ε -closeness enhances the usefulness of the result.¹⁶

Theorem 3 *Suppose the public profile (W, w^1, σ, d) is a strict equilibrium of the full-support public-monitoring game $(\bar{u}^*, (Y, \rho))$ for some δ and $|W| < \infty$. There exists $\eta > 0$ and $\varepsilon > 0$ such that for any game with private monitoring $(u^*, (\Omega, \pi))$ strongly ε -close to $(\bar{u}^*, (Y, \rho))$, if there exists a player i , a private history for that player h_i^t , and a state w such that $d_i(w) \neq d_i(\sigma_i(h_i^t))$ and $\beta_i(w | h_i^t) > 1 - \eta$, then the induced private profile is not a Nash equilibrium of the game with private monitoring for the same δ .*

We implicitly used this result in our discussions of the repeated prisoners' dilemma. For example, in Example 5, we argued that there was a private history for player 1 that leaves him in the private state w_1^n , but his posterior after that history assigns probability close to 1 that player 2's private state is w_2^e .

Our approach is to ask when it is possible to so “manipulate” a player's beliefs through selection of private history that the hypotheses of Theorem 3 are satisfied. In particular, we are interested in the weakest independent conditions on the private-monitoring distributions and on the strategy profiles that would allow such manipulation.

Fix a PPE of the public-monitoring game and a close-by almost-public-monitoring game. The logic of Example 5 runs as follows: Consider a player i in a private state \hat{w} who assigns strictly positive (albeit small) probability to all the other players being in some other common private state $\bar{w} \neq \hat{w}$ (full-support private monitoring ensures that such an occurrence arises with positive probability). Let $\tilde{a} = (d_i(\hat{w}), d_{-i}(\bar{w}))$ be the action profile that results when i is in state \hat{w} and all the other players are in state \bar{w} . Suppose that if any other player is in a different private state $w \neq \bar{w}$, then the resulting action profile differs from \tilde{a} . Suppose, moreover, there is a signal y such that $\hat{w} = \sigma(\hat{w}, y)$ and $\bar{w} = \sigma(\bar{w}, y)$, that is, any player in the state \hat{w} or \bar{w} observing a private signal consistent with y stays in that private state (and so the profile cannot have bounded recall, see Lemma 5). Suppose finally there is a private signal ω_i for player i consistent with y that is more likely to have come from \tilde{a} than any other action profile, i.e., $\omega_i \in f_i^{-1}(y)$ and (where $\pi_i(\omega_i | a)$ is the probability that player i observes the signal ω_i under a)

$$\pi_i(\omega_i | \tilde{a}) > \pi_i(\omega_i | (d_i(\hat{w}), a'_{-i})) \quad \forall a'_{-i} \neq d_{-i}(\bar{w}). \quad (1)$$

¹⁶While we have stated this theorem, and Theorem 4 below, for pure strategies, they also hold for some mixed strategy profiles. Recall from Section 2.1 that given an automaton (W, d_i, σ_i) describing a collection of pure strategies for player i (taking any state $w \in W$ as the initial state gives a pure strategy), a probability distribution over W gives a mixed strategy. Consider now a mixed strategy PPE of the game with public-monitoring. Clearly, such a profile cannot be strict. However, there may exist a period T , such that all the incentive constraints after period T are strict (the equilibria in Sekiguchi (1997) are important examples). In that case, Theorem 3 holds if the hypotheses are satisfied for $t \geq T$.

Then, after observing the private signal ω_i , player i 's posterior probability that all the other players are in \bar{w} should increase (this is not immediate, however, since the monitoring is private). Moreover, since players in \hat{w} and \bar{w} do not change their private states, we can eventually make player i 's posterior probability that all the other players are in \bar{w} as close to one as we like. If $d_i(\hat{w}) \neq d_i(\bar{w})$, an application of Theorem 3 shows that the induced private profile is not an equilibrium.

The suppositions in the above logic can be weakened in two ways. First, it is not necessary that the *same* private signal ω_i be more likely to have come from \tilde{a} than *any* other action profile. It should be enough if for each action profile different from \tilde{a} , there is a private signal more likely to have come from \tilde{a} than that profile, as long as the signal not mess up the other inferences too badly. In that case, realizations of the other signals could undo any damage done without negatively impacting on the overall inferences. For example, suppose there are two players, with player 1 the player whose beliefs we are “manipulating,” and in addition to state \bar{w} , player 2 could be in state \hat{w} or w . Suppose also $A_2 = \{\tilde{a}_2, a'_2, a''_2\}$. As before, suppose there is a signal y such that $w = \sigma(w, y)$, $\hat{w} = \sigma(\hat{w}, y)$, and $\bar{w} = \sigma(\bar{w}, y)$, that is, any player in the state w , \hat{w} , or \bar{w} , observing a private signal consistent with y stays in that private state. We would like the odds ratio $\Pr(w_2 \neq \bar{w} | h_1^t) / \Pr(w_2 = \bar{w} | h_1^t)$ to converge to zero as $t \rightarrow \infty$, for appropriate private histories. Let $\tilde{a}_1 = d_1(\hat{w})$, $\tilde{a}_2 = d_2(\bar{w})$, $a'_2 = d_2(\hat{w})$, and $a''_2 = d_2(w)$, and suppose there are two private signals, ω'_1 and ω''_1 consistent with y , satisfying

$$\pi_1(\omega'_1 | \tilde{a}_1, a''_2) > \pi_1(\omega'_1 | \tilde{a}) > \pi_1(\omega'_1 | \tilde{a}_1, a'_2)$$

and

$$\pi_1(\omega''_1 | \tilde{a}_1, a'_2) > \pi_1(\omega''_1 | \tilde{a}) > \pi_1(\omega''_1 | \tilde{a}_1, a''_2).$$

Then, after observing the private signal ω'_1 , we have

$$\frac{\Pr(w_2 = \hat{w} | h_1^t, \omega'_1)}{\Pr(w_2 = \bar{w} | h_1^t, \omega'_1)} = \frac{\pi_1(\omega'_1 | \tilde{a}_1, a'_2)}{\pi_1(\omega'_1 | \tilde{a})} \frac{\Pr(w_2 = \hat{w} | h_1^t)}{\Pr(w_2 = \bar{w} | h_1^t)} < \frac{\Pr(w_2 = \hat{w} | h_1^t)}{\Pr(w_2 = \bar{w} | h_1^t)}$$

as desired, but $\Pr(w_2 = w | h_1^t, \omega'_1) / \Pr(w_2 = \bar{w} | h_1^t, \omega'_1)$ increased. On the other hand, after observing another private signal ω''_1 , also consistent with y , while the odds ratio $\Pr(w_2 = w | h_1^t, \omega''_1) / \Pr(w_2 = \bar{w} | h_1^t, \omega''_1)$ falls, $\Pr(w_2 = \hat{w} | h_1^t, \omega''_1) / \Pr(w_2 = \bar{w} | h_1^t, \omega''_1)$ increases. However, it may be that the increases can be offset by appropriate decreases, so that, for example, ω'_1 followed by two realizations of ω''_1 results in a decrease in *both* odds ratios. If so, a sufficiently high number of realizations of $\omega'_1 \omega''_1 \omega''_1$ result in $\Pr(w_2 \neq \bar{w} | h_1^t) / \Pr(w_2 = \bar{w} | h_1^t)$ being close to zero.

In terms of the odds ratios, the sequence of signals $\omega'_1 \omega''_1 \omega''_1$ lowers both odds ratios if, and only if,

$$\frac{\pi_1(\omega'_1 | \tilde{a}_1, a'_2)}{\pi_1(\omega'_1 | \tilde{a})} \left(\frac{\pi_1(\omega''_1 | \tilde{a}_1, a''_2)}{\pi_1(\omega''_1 | \tilde{a})} \right)^2 < 1$$

and

$$\frac{\pi_1(\omega'_1|\tilde{a}_1, a''_2)}{\pi_1(\omega'_1|\tilde{a})} \left(\frac{\pi_1(\omega''_1|\tilde{a}_1, a''_2)}{\pi_1(\omega''_1|\tilde{a})} \right)^2 < 1.$$

Our richness condition on private monitoring distributions captures this idea. For a private monitoring distribution, (Ω, π) , define $\gamma_{ad'_{-i}}(\omega_i) \equiv \log \pi_i(\omega_i|a_i, a_{-i}) - \log \pi_i(\omega_i|a_i, a'_{-i})$, and let $\gamma_a(\omega_i) = \left(\gamma_{ad'_{-i}}(\omega_i) \right)_{a'_{-i} \in A_{-i}, a'_{-i} \neq a_{-i}}$ denote the vector in $\mathfrak{R}^{|A_{-i}|-1}$ of the log odds ratios of the signal ω_i associated with different action profiles. The last two displayed equations can then be written as $\frac{1}{3}\gamma_{\tilde{a}}(\omega'_1) + \frac{2}{3}\gamma_{\tilde{a}}(\omega''_1) > \mathbf{0}$, where $\mathbf{0}$ is the 2×1 zero vector.¹⁷

Definition 7 A private-monitoring distribution (Ω, π) is rich, given signaling functions f , if for all $a \in A$ and all $y \in Y$, the convex hull of the set of vectors $\{\gamma_a(\omega_i) : \omega_i \in f_i^{-1}(y) \text{ and } \pi_i(\omega_i|a_i, a'_{-i}) > 0 \text{ for all } a'_{-i} \in A_{-i}\}$ has a nonempty intersection with $\mathfrak{R}_{++}^{|A_{-i}|-1}$.

It will be useful to quantify the extent to which the conditions of Definition 7 are satisfied. Since the space of signals and actions are finite, there are a finite number of constraints in Definition 7, and so for any rich private monitoring distribution, the set of ζ over which the supremum is taken in the next definition is non-empty.¹⁸

Definition 8 Given f , the richness of a rich private-monitoring distribution (Ω, π) is the supremum of all $\zeta > 0$ satisfying: for all $y \in Y$, the convex hull of the set of vectors $\{\gamma_a(\omega_i) : \omega_i \in f_i^{-1}(y) \text{ and } \pi_i(\omega_i|a_i, a'_{-i}) \geq \zeta \text{ for all } a'_{-i} \in A_{-i}\}$ has a nonempty intersection with $\mathfrak{R}_{\zeta}^{|A_{-i}|-1} \equiv \{x \in \mathfrak{R}_{++}^{|A_{-i}|-1} : x_k \geq \zeta \text{ for } k = 1, \dots, |A_{-i}| - 1\}$.

The second weakening concerns the nature of the strategy profile. The logic assumed that there is a signal y such that $\hat{w} = \sigma(\hat{w}, y)$ and $\bar{w} = \sigma(\bar{w}, y)$. If there were only two states, \hat{w} and \bar{w} , it would clearly be enough that there be a finite sequence of signals such that both \hat{w} and \bar{w} cycle. When there are more states, we also need to worry about what happens to the other states. In addition, we need to allow for time-dependent profiles, and profiles that use some states for only a finite time. Define $R(\tilde{w})$ as the set of states that are repeatedly reachable in the same period as \tilde{w} (i.e., $R(\tilde{w}) = \{w \in W : \{w, \tilde{w}\} \subset W_t \text{ infinitely often}\}$).

We generalize the cycling idea to the notion that there be a path that allows some distinguished state to be separated from every other state that could ever be reached. Given an outcome path $h \equiv (y^1, y^2, \dots) \in Y^\infty$, let ${}^\tau h \equiv (y^\tau, y^{\tau+1}, \dots) \in Y^\infty$ denote the outcome path from period τ , so that $h = (h^\tau, {}^\tau h)$ and ${}^\tau h^{\tau+t} = (y^\tau, y^{\tau+1}, \dots, y^{\tau+t-1})$.

¹⁷The convex combination is strictly positive (rather than negative) because the definition of $\gamma_{ad'_{-i}}$ inverts the odds ratios from the displayed equations.

¹⁸The bound ζ appears twice in the definition. Its first appearance ensures that for all $\zeta > 0$, there is uniform upper bound on the number of private signals satisfying $\pi_i(\omega_i|a_i, a'_{-i}) \geq \zeta$ in any private-monitoring distribution with a richness of at least ζ .

	A	B	C
A	3,3	0,0	0,0
B	0,0	3,3	0,0
C	0,0	0,0	2,2

Figure 6: The normal form for Example 6.

Definition 9 *The public strategy profile is separating if there is some state \tilde{w} and an outcome path $h \in Y^\infty$ such that there is another state $w \in R(\tilde{w})$ that satisfies $\sigma(w, h^t) \neq \sigma(\tilde{w}, h^t)$ for all t , and for all τ and $w \in R(\sigma(\tilde{w}, h^\tau))$, if $\sigma(w, {}^\tau h^{\tau+t}) \neq \sigma(\tilde{w}, h^{\tau+t})$ for all $t \geq 0$, then*

$$d_i(\sigma(w, {}^\tau h^{\tau+t})) \neq d_i(\sigma(\tilde{w}, h^{\tau+t})) \text{ infinitely often, for all } i.$$

When the set of states is finite, our next Lemma (proved in the Appendix) shows that separation is *equivalent* to a simpler and seemingly stronger cycling condition.

Lemma 6 *A finite public strategy profile of the public-monitoring game is separating if, and only if, there is a finite sequence of signals $\bar{y}^1, \dots, \bar{y}^m$, a collection of states W_c , and a state $\bar{w} \in W_c$ such that*

1. $\sigma(w, \bar{y}^1, \dots, \bar{y}^m) = w$ for all $w \in W_c$,
2. $\sigma(w, \bar{y}^1, \dots, \bar{y}^m) \in W_c$ for all $w \in R(\bar{w})$,
3. $\forall w \in W_c \setminus \{\bar{w}\}, \forall i \exists k, 1 \leq k \leq m$, such that

$$d_i(\sigma(w, \bar{y}^1, \dots, \bar{y}^k)) \neq d_i(\sigma(\bar{w}, \bar{y}^1, \dots, \bar{y}^k)),$$

and

4. for some i and $\hat{w} \in W_c \setminus \{\bar{w}\}$, $d_i(\hat{w}) \neq d_i(\bar{w})$.

Note that part 4 implies that $|W_c| \geq 2$. We emphasize that each state in the set of states W_c cycles under the given finite sequence of signals *and* every state reachable (infinitely often) in the same period as \bar{w} is taken into W_c by one round of the cycle.

Clearly, a separating profile cannot have bounded recall. Moreover, it is easy to construct PPE that neither have bounded recall nor are separating (Example 6). Nonetheless, we are not aware of any strict PPE of substantive interest that neither have bounded recall nor are separating.




Figure 7: The strategy profile for Example 6. In states w^A and \hat{w}^A , the action A is played, while in w^B the action B and in w^C , the action C is played.

Example 6 The stage game is given in Figure 6. In the public-monitoring game, there are two public signals, y' and y'' , with distribution ($0 < q < p < 1$)

$$\rho(y''|a_1a_2) = \begin{cases} p, & \text{if } a_1 = a_2, \\ q, & \text{otherwise.} \end{cases}$$

Finally, the public profile is illustrated in Figure 7. This profile is not separating: Under any path in which $y'y$, $y = y'$ or y'' , appears, all states transit to the same state. If only y'' appears, only w^A and \hat{w}^A arise. The definition of separation fails because play is the same at states w^A and \hat{w}^A .

The profile is also not robust: After enough realizations of private signals corresponding to y'' , beliefs must assign roughly equal probability to w^A and \hat{w}^A ,¹⁹ and so after the first realization of a private signal corresponding to y' , B is the only best reply (even if the current state is w^C). This example (like the second forgiving grim trigger of Example 3) illustrates the possibility that beliefs over private states can drift to a stationary distribution when play is identical in different states.

¹⁹This is most easily seen by considering the Markov chain describing player 2's private state transitions conditional on player 1 always playing A and always observing the same private signal consistent with y'' (a Markov chain is associated with each $\omega_1 \in f_1(y'')$). Each such Markov chain is ergodic, and so has a unique stationary distribution. A straightforward calculation shows that, in the limit (as the private-monitoring distributions become arbitrarily close), the probability assigned to w_2^A equals $\frac{1}{2}$.

It remains to ensure that, under private monitoring, players may transit to different states. It suffices to assume the following, weaker than full-support, condition:²⁰

Definition 10 *A private monitoring distribution (Ω, π) that is ε -close to a public monitoring distribution (Y, ρ) has essentially full support if for all $(y_1, \dots, y_n) \in Y^n$,*

$$\pi\{(\omega_1, \dots, \omega_n) \in \Omega : f_i(\omega_i) = y_i, i = 1, \dots, n\} > 0.$$

Theorem 4 *Fix a separating strict finite PPE of a full-support public-monitoring game $(\tilde{u}^*, (Y, \rho))$. For all $\zeta > 0$, there exists $\varepsilon' > 0$ such that for all $\varepsilon < \varepsilon'$, if $(u, (\Omega, \pi))$ is a private-monitoring game strongly ε -close under some signalling function f to $(\tilde{u}^*, (Y, \rho))$ with (Ω, π) having richness, given f , at least ζ and essentially full support, then the induced private profile is not a Nash equilibrium of the private monitoring game.*

It is worth noting that the bound on ε is only a function of the richness of the private monitoring. It is *independent* of the probability that a disagreement in private states arises. By considering finite state profiles that are separating, not only is the difficulty identified in the Introduction dealt with (as we discuss at the end of the next Section), but we can accommodate arbitrarily small probabilities of disagreement.

Thus, separating strict PPE of public-monitoring games are not robust to the introduction of private monitoring.²¹ It, of course, also implies that separating behavior in the private-monitoring game typically cannot coordinate continuation play in the following sense. Say a profile is ε -strict if all the incentive constraints are satisfied by at least ε . (The result follows immediately from upperhemicontinuity and Theorem 4.)

Corollary 1 *Fix a vector of signaling functions f , $f_i : \Omega_i \rightarrow Y$. Suppose $\{(u^k, (\Omega, \pi^k))\}$ is a sequence of private-monitoring games, with $(u^k, (\Omega, \pi^k))$ $1/k$ -close to some public-monitoring game $(\tilde{u}^*, (Y, \rho))$ and $\{(\Omega, \pi^k)\}$ a rich sequence of distributions. Fix a pure strategy profile of the private monitoring game in which each player's strategy respects his signaling function f_i (i.e., $\sigma_i(h_i, a_i, \omega_i) = \sigma_i(h_i, a_i, \hat{\omega}_i)$ if $f_i(\omega_i) = f_i(\hat{\omega}_i) \neq \emptyset$). Suppose this profile is separating (when interpreted as a public profile). For all $\varepsilon > 0$, there exists k' such that for $k > k'$, this profile is not an ε -strict Nash equilibrium.*

Since the equilibrium failure of separating profiles seem to arise after private histories that have low probability, an attractive conjecture is that equilibrium can be restored by appropriately modifying the profile at only the problematic histories. Unfortunately, such a modification appears to require additional modifications to the profile, destroying the connection to the public-monitoring game.

²⁰If an essentially-full-support private monitoring distribution does not have full support, Nash equilibria of the private-monitoring game may not have realization-equivalent sequentially-rational strategy profiles (recall Remark 2).

²¹The extension to mixed strategies described in footnote 16 also holds for Theorem 4.

6. The Proof of Theorem 4

The proof of Theorem 4 is by contradiction. So, suppose there exists $\zeta > 0$ such that for all k there exists a private monitoring game $(u, (\Omega^k, \pi^k))$ strongly $1/k$ -close under some f to $(\tilde{u}^*, (Y, \rho))$ with (Ω^k, π^k) having richness at least ζ , with the induced private profile a Nash equilibrium of the private-monitoring game.

To develop intuition, suppose the space of signals for each player were independent of k , so that $\Omega_i^k = \Omega_i$. Then, we can assume π^k converges to a limit distribution π^∞ on Ω (by choosing a subsequence if necessary). The behavior of beliefs of player i over the private states of the other players under the limit private monitoring distribution (Ω, π^∞) is significantly easier to describe. Since (Ω, π^k) is strongly $1/k$ -close to (Y, ρ) and $\pi^k \rightarrow \pi^\infty$, for each $y \in Y$ the event $\{(\omega_1, \dots, \omega_n) : \omega_i \in f_i^{-1}(y)\}$ is common belief under π^∞ . Moreover, if the other players start in the same state (such as \bar{w}) then they stay in the same state thereafter. We can thus initially focus on finding the appropriate sequence of signals to manipulate i 's updating about the current private states of the other players, without being concerned about the possibility that subsequent realizations will derail the process (we will deal with that issue subsequently). The difficulty, of course, is that Ω_i^k depends on k , and moreover, that in principle as k gets large, so may Ω_i^k .

We can however, proceed as follows: For each k and $a_i \in A_i$, let

$$\Omega_i^{k, a_i} = \left\{ \omega_i \in \Omega_i^k : \pi_i^k(\omega_i | a_i, a'_{-i}) > \zeta \text{ for all } a'_{-i} \in A_{-i} \right\}.$$

Since (Ω^k, π^k) is strongly close to (Y, ρ) , every signal in Ω_i^k is associated with some public signal, and so we can partition Ω_i^{k, a_i} into subsets of private signals associated with the same public signal, $\Omega_i^{k, a_i}(y)$. Order arbitrarily the signals in $\cup_{a_i} \Omega_i^{k, a_i}(y)$, and give the ℓ -th signal in the order the label (y, ℓ) . Let $k_{i, y} \equiv \left| \cup_{a_i} \Omega_i^{k, a_i}(y) \right|$; note that $k_{i, y}$ is (crudely) bounded above by $|A_i|/\zeta$ for all k . With this relabeling, and defining $\Omega_i \equiv \cup_{y \in Y} \{(y, 1), (y, 2), \dots, (y, k_{i, y})\}$, a finite set, we have, for all i and k ,

$$\Omega_i^k \subset \Omega_i \cup \left(\Omega_i^k \setminus \left(\cup_{a_i \in A_i} \Omega_i^{k, a_i} \right) \right) \quad (2)$$

and

$$\Omega_i^k \cap \Omega_i \neq \emptyset.$$

Without loss of generality, we can assume (2) holds with equality (simply include any signal $\omega_i \in \Omega_i \setminus \Omega_i^k$ in Ω_i^k , so that $\pi_i^k(\omega_i | a) = 0$).

We augment Ω_i , for each $y \in Y$, by a new signal denoted ω_i^y , and define $\Omega_i^\infty \equiv \Omega_i \cup (\cup_y \{\omega_i^y\})$. We interpret ω_i^y as the set of i 's private signals associated with y that are not in Ω_i . For each k , we can interpret Ω_i^∞ as a partition of Ω_i^k (each $\omega_i \in \Omega_i$ appears as a singleton, while $\omega_i^y \equiv \left\{ \omega_i \in \Omega_i^k \setminus \left(\cup_{a_i \in A_i} \Omega_i^{k, a_i} \right) : f_i(\omega_i) = y \right\}$ may be empty). For each $a \in A$, denote by $\hat{\pi}^k(\cdot | a)$ the probability distribution on $\prod_i \Omega_i^\infty$ induced by $\pi^k(\cdot | a)$. Note that we now have a

sequence of probability distributions $\{\hat{\pi}^k(\cdot|a)\}_k$ for each $a \in A$ on a common finite signal space $\prod_i \Omega_i^\infty$.

By passing to a subsequence if necessary, we can assume $\{\hat{\pi}^k(\omega|a)\}_k$ is a convergent sequence with limit $\pi^\infty(\omega|a)$ for all $a \in A$, $\omega \in \prod_i \Omega_i^\infty$. Note that $(\Omega^\infty, \pi^\infty)$ is 0-close to (Y, ρ) . Moreover, by passing to a further subsequence if necessary, we can also assume that, for each i , $a_i \in A_i$, and $y \in Y$, the convex hull of the set of vectors $\{\gamma_a^\infty(\omega_i) : \omega_i \in f_i^{-1}(y), \pi_i^\infty(\omega_i|a_i, a'_{-i}) > \zeta \text{ for all } a'_{-i} \in A_{-i}\}$ has a nonempty intersection with $\mathfrak{R}_\zeta^{|A_{-i}|-1}$, where $\gamma_{aa'_{-i}}^\infty(\omega_i) \equiv \log \pi_i^\infty(\omega_i|a_i, a'_{-i}) - \log \pi_i^\infty(\omega_i|a_i, a'_{-i})$ and $\gamma_a^\infty(\omega_i) = \left(\gamma_{aa'_{-i}}^\infty(\omega_i) \right)_{a'_{-i} \in A_{-i}, a'_{-i} \neq a_{-i}}$.

In the following lemma, a private signal ω_j for player j is *consistent* with the private signal ω_i for player i if $f_j(\omega_j) = f_i(\omega_i)$, where f_i and f_j are the signaling functions from Definition 4. It is an implication of this lemma that if player i assigns strictly positive probability to all the other players being in the state \bar{w} , then after sufficient repetitions of the cycle $\vec{\omega}_i^L$ (defined in Lemma 7), player i eventually assigns probability arbitrarily close to 1 that at the end of a cycle, all the other players are in the state \bar{w} .

Lemma 7 *Fix a finite separating public profile of the public-monitoring game, and let \bar{w} , \hat{w} , W_c , and i be the states, set of states, and player identified in Lemma 6. Then, there exists a finite sequence of private signals for player i , $\vec{\omega}_i^L \equiv (\omega_i^1, \omega_i^2, \dots, \omega_i^L)$, such that*

1. $\sigma_i(\hat{w}, \vec{\omega}_i^L) = \hat{w}$,
2. for all sequences of private signals, $\vec{\omega}_j^L$, for player $j \neq i$ consistent with $\vec{\omega}_i^L$, $\sigma_j(w, \vec{\omega}_j^L) = w$ for all $w \in W_c$, and
3. for all $\mathbf{w} \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}$,

$$A(\vec{\omega}_i^L; \mathbf{w}) \equiv \frac{\Pr_\infty(\vec{\omega}_i^L | w_{-i} = \mathbf{w}, w_i = \hat{w})}{\Pr_\infty(\vec{\omega}_i^L | w_{-i} = \bar{w}\mathbf{1}, w_i = \hat{w})} < 1, \quad (3)$$

where \Pr_∞ denotes probabilities calculated under π^∞ and the assumption that all players follow the private profile.

Proof. The cycle $\bar{y}^1, \dots, \bar{y}^m$ from Lemma 6 induces a cycle in the states $\bar{w} = \bar{w}^1, \dots, \bar{w}^{m+1} = \bar{w}^1$ and $\hat{w} = \hat{w}^1, \dots, \hat{w}^{m+1} = \hat{w}$. We index the cycle by ℓ and write $\bar{a}^\ell = d(\bar{w}^\ell)$ and $\hat{a}_i^\ell = d_i(\hat{w}^\ell)$. Let $\vec{a}^\ell \equiv (\hat{a}_i^\ell, \bar{a}_{-i}^\ell)$. Richness implies that for each ℓ , there exists a vector of nonnegative integers, $(n_{\omega_i})_{\omega_i \in f_i^{-1}(\bar{y}^\ell)}$, so that for all $a'_{-i} \neq \bar{a}_{-i}^\ell$,

$$\sum_{\omega_i \in f_i^{-1}(\bar{y}^\ell)} \gamma_{\hat{a}_i^\ell, a'_{-i}}^\infty(\omega_i) n_{\omega_i} > 0.$$

Since

$$\gamma_{\hat{a}_i^\ell, a'_{-i}}^\infty(\omega_i) = \log \pi_i^\infty(\omega_i | \vec{a}^\ell) / \pi_i^\infty(\omega_i | \hat{a}_i^\ell, a'_{-i}),$$

we have, for all $a'_{-i} \neq \bar{a}_{-i}^\ell$,

$$\prod_{\omega_i \in f_i^{-1}(\bar{y}^\ell)} \left(\frac{\pi_i^\infty(\omega_i | \bar{a}^\ell)}{\pi_i^\infty(\omega_i | \hat{a}_i^\ell, a'_{-i})} \right)^{n_{\omega_i}} > 1. \quad (4)$$

Letting $n_\ell = \sum_{\omega_i \in f_i^{-1}(\bar{y}^\ell)} n_{\omega_i}$ for each ℓ , denote by N' the lowest common multiple of $\{n_1, \dots, n_m\}$.

Let $\vec{\omega}_i^L$ denote the cycle of private signals for player i consistent with cycling N times through the public signals $\bar{y}^1, \bar{y}^2, \dots, \bar{y}^m$ and in which for each ℓ , the private signal $\omega_i \in f_i^{-1}(\bar{y}^\ell)$ appears $(N'/n_\ell)n_{\omega_i}$ times. This cycle is of length $L \equiv mN'$.

Given a private state profile $\mathbf{w} \in W_c^{n-1}$, let \check{a}_{-i}^ℓ denote the action profile taken in period ℓ of the cycle. Then,

$$\begin{aligned} A(\vec{\omega}_i^L; \mathbf{w}) &\equiv \frac{\Pr_\infty(\vec{\omega}_i^L | w_{-i}^t = \mathbf{w}, w_i = \hat{w})}{\Pr_\infty(\vec{\omega}_i^L | w_{-i}^t = \bar{w}\mathbf{1}, w_i = \hat{w})} \\ &= \left(\prod_{\ell=1}^m \left(\prod_{\omega_i \in f_i^{-1}(\bar{y}^\ell)} \left(\frac{\pi_i^\infty(\omega_i | \hat{a}_i^\ell, \check{a}_{-i}^\ell)}{\pi_i^\infty(\omega_i | \bar{a}^\ell)} \right)^{n_{\omega_i}} \right)^{N/n_\ell} \right). \end{aligned}$$

For $\mathbf{w} \neq \bar{w}\mathbf{1}$, then in each period at least one player is in a private state different from \bar{w} . From Lemma 6.2, $\check{a}_{-i}^\ell \neq \bar{a}_{-i}^\ell$ for at least one ℓ , and so $A(\vec{\omega}_i^L; \mathbf{w})$ must be strictly less than 1. \blacksquare

We are, of course, primarily concerned with private monitoring under the distribution (Ω^k, π^k) . In this situation, one must deal with the possibility that player j 's private signals may be inconsistent with player i 's observations. However, by choosing k sufficiently large, one can ensure that this possibility does not arise with large probability *along* the cycle $\vec{\omega}_i^L$. The subsequent lemma implies that this possibility never arises with large probability.

Lemma 8 *Assume the hypotheses of Lemma 7, and let h_i^t be a private history for player i satisfying $\hat{w} = \sigma_i(h_i^t)$. For all $\eta > 0$, there exists $\xi > 0$ and k' (independent of h_i^t) such that, for all $k > k'$, if $\eta < \Pr_k(w_{-i}^t \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\} | h_i^t) < 1$ and $\Pr_k(w_{-i}^t \notin W_c^{n-1} | h_i^t) < \xi$, then*

$$\frac{\Pr_k(w_{-i}^{t+L} \neq \bar{w}\mathbf{1} | \vec{\omega}_i^L, h_i^t)}{\Pr_k(w_{-i}^{t+L} = \bar{w}\mathbf{1} | \vec{\omega}_i^L, h_i^t)} < (1 - \xi) \frac{\Pr_k(w_{-i}^t \neq \bar{w}\mathbf{1} | h_i^t)}{\Pr_k(w_{-i}^t = \bar{w}\mathbf{1} | h_i^t)}, \quad (5)$$

where \Pr_k denotes probabilities calculated under π^k and the assumption that all players follow the private profile, and $\vec{\omega}_i^L$ is the sequence identified in Lemma 7.

Proof. For clarity, we suppress the conditioning on h_i^t . Denote the event that players other than i observe some sequence of private signals consistent with the cycle $(\bar{y}^1, \dots, \bar{y}^m)^N$ by \vec{y}_{-i} , and the complementary event by $\neg\vec{y}_{-i}$. Then,

$$\Pr_k(w_{-i}^{t+L} \neq \bar{w}\mathbf{1}, \vec{\omega}_i^L) = \Pr_k(w_{-i}^{t+L} \neq \bar{w}\mathbf{1}, \vec{\omega}_i^L, \vec{y}_{-i}) + \Pr_k(w_{-i}^{t+L} \neq \bar{w}\mathbf{1}, \vec{\omega}_i^L, \neg\vec{y}_{-i})$$

and

$$\begin{aligned}
& \Pr_k(w_{-i}^{t+L} \neq \bar{w}\mathbf{1}, \vec{\omega}_i^L, \vec{y}_{-i}) \\
& \leq \Pr_k(w_{-i}^t \neq \bar{w}\mathbf{1}, \vec{\omega}_i^L, \vec{y}_{-i}) \\
& = \Pr_k(w_{-i}^t \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}, \vec{\omega}_i^L, \vec{y}_{-i}) + \Pr_k(w_{-i}^t \notin W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}, \vec{\omega}_i^L, \vec{y}_{-i}),
\end{aligned}$$

where the inequality arises because a player $j \neq i$ may be in a private state not in W_c . Now,

$$\begin{aligned}
& \Pr_k(w_{-i}^t \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}, \vec{\omega}_i^L, \vec{y}_{-i}) \\
& = \Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} | w_{-i}^t \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}) \Pr_k(w_{-i}^t \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}) \\
& \leq \Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} | w_{-i}^t \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}) \Pr_k(w_{-i}^t \neq \bar{w}\mathbf{1}),
\end{aligned}$$

and if $\Pr_k(w_{-i}^t \notin W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}) < \xi$ (where ξ is to be determined),

$$\begin{aligned}
& \Pr_k(w_{-i}^t \notin W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}, \vec{\omega}_i^L, \vec{y}_{-i}) + \Pr_k(w_{-i}^{t+L} \neq \bar{w}\mathbf{1}, \vec{\omega}_i^L, \neg\vec{y}_{-i}) \\
& < \xi + \Pr_k(w_{-i}^{t+L} \neq \bar{w}\mathbf{1}, \vec{\omega}_i^L, \neg\vec{y}_{-i}) \\
& \leq \xi + \Pr_k(\vec{\omega}_i^L, \neg\vec{y}_{-i}) \\
& = \xi + \Pr_k(\neg\vec{y}_{-i} | \vec{\omega}_i^L) \Pr_k(\vec{\omega}_i^L).
\end{aligned}$$

Moreover,

$$\begin{aligned}
& \Pr_k(w_{-i}^{t+L} = \bar{w}\mathbf{1}, \vec{\omega}_i^L) \geq \Pr_k(w_{-i}^t = \bar{w}\mathbf{1}, \vec{\omega}_i^L, \vec{y}_{-i}) \\
& = \Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} | w_{-i}^t = \bar{w}\mathbf{1}) \Pr_k(w_{-i}^t = \bar{w}\mathbf{1}).
\end{aligned}$$

Defining

$$x^t(k) \equiv \frac{1}{\Pr_k(w_{-i}^t \neq \bar{w}\mathbf{1})} (\xi + \Pr_k(\neg\vec{y}_{-i} | \vec{\omega}_i^L) \Pr_k(\vec{\omega}_i^L)),$$

we have,

$$\begin{aligned}
& \frac{\Pr_k(w_{-i}^{t+L} \neq \bar{w}\mathbf{1} | \vec{\omega}_i^L)}{\Pr_k(w_{-i}^{t+L} = \bar{w}\mathbf{1} | \vec{\omega}_i^L)} \\
& < \frac{\Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} | w_{-i}^t \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}) + x^t(k)}{\Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} | w_{-i}^t = \bar{w}\mathbf{1})} \times \frac{\Pr_k(w_{-i}^t \neq \bar{w}\mathbf{1})}{\Pr_k(w_{-i}^t = \bar{w}\mathbf{1})} \\
& \leq \frac{\max_{\mathbf{w} \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}} \Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} | w_{-i}^t = \mathbf{w}) + x^t(k)}{\Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} | w_{-i}^t = \bar{w}\mathbf{1})} \times \frac{\Pr_k(w_{-i}^t \neq \bar{w}\mathbf{1})}{\Pr_k(w_{-i}^t = \bar{w}\mathbf{1})}. \tag{6}
\end{aligned}$$

From Lemma 7,

$$\max_{\mathbf{w} \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}} A(\vec{\omega}_i^L; \mathbf{w}) = \max_{\mathbf{w} \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}} \lim_{k \rightarrow \infty} \frac{\Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} | w_{-i}^t = \mathbf{w})}{\Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} | w_{-i}^t = \bar{w}\mathbf{1})} < 1,$$

and so there is an $\xi' > 0$ sufficiently small so that (recall that the denominator has a strictly positive limit)

$$\max_{\mathbf{w} \in W_c^{n-1} \setminus \{\bar{\mathbf{w}}\mathbf{1}\}} \lim_{k \rightarrow \infty} \frac{\Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} | w_{-i}^t = \mathbf{w}) + \xi'}{\Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} | w_{-i}^t = \bar{\mathbf{w}}\mathbf{1})} < 1 - \xi'.$$

The finiteness of the state space and the number of players allows us to interchange the max and lim operations. Consequently, there exists k'' such that for all $k \geq k''$,

$$\frac{\max_{\mathbf{w} \in W_c^{n-1} \setminus \{\bar{\mathbf{w}}\mathbf{1}\}} \Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} | w_{-i}^t = \mathbf{w}) + \xi'}{\Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} | w_{-i}^t = \bar{\mathbf{w}}\mathbf{1})} < 1 - \xi'. \quad (7)$$

Since (Ω, π^k) is strongly $1/k$ -close to (Y, ρ) , $\lim_{k \rightarrow \infty} \Pr_k(\neg \vec{y}_{-i} | \vec{\omega}_i^L) = 0$, and so there exists k''' such that $\Pr_k(\neg \vec{y}_{-i} | \vec{\omega}_i^L) < \xi' \eta / 2$ for all $k \geq k'''$. Suppose $\xi = \xi' \eta / 2$ and $k' = \max\{k'', k'''\}$. Since $\eta < \Pr_k(w_{-i}^t \in W_c^{n-1} \setminus \{\bar{\mathbf{w}}\mathbf{1}\}) \leq \Pr_k(w_{-i}^t \neq \bar{\mathbf{w}}\mathbf{1})$, $x^t(k) \leq \xi'$. Consequently (7), with (6), implies (5) (since $\xi < \xi'$). \blacksquare

Lemma 6 guarantees that one round of the cycle of signals will always take a state not in W_c into W_c , ensuring that the probability on states in $W \setminus W_c$ can be controlled.

Lemma 9 *Assume the hypotheses of Lemma 7, and let h_i^t be a private history for player i satisfying $\hat{w} = \sigma_i(h_i^t)$. Fix $\eta > 0$ and let ξ and k' be the constants identified in Lemma 8. There exists T such that if $t \geq T$, then for all $k > k'$,*

$$\Pr_k(w_{-i}^{t+L} \notin W_c^{n-1} | \vec{\omega}_i^L, h_i^t) < \xi.$$

Proof. Fix T large enough, so that if $\bar{w} \in W_t$ (the set of states reachable in period t) for $t \geq T$, then $W_t \subset R(\bar{w})$. Separation then implies $\Pr_k(w_{-i}^{t+L} \notin W_c^{n-1}, \vec{y}_{-i}) = 0$, and so

$$\begin{aligned} & \Pr_k(w_{-i}^{t+L} \notin W_c^{n-1} | \vec{\omega}_i^L) \\ &= \Pr_k(w_{-i}^{t+L} \notin W_c^{n-1}, \vec{y}_{-i} | \vec{\omega}_i^L) + \Pr_k(w_{-i}^{t+L} \notin W_c^{n-1}, \neg \vec{y}_{-i} | \vec{\omega}_i^L) \\ &= \Pr_k(w_{-i}^{t+L} \notin W_c^{n-1}, \neg \vec{y}_{-i} | \vec{\omega}_i^L) \\ &\leq \Pr_k(\neg \vec{y}_{-i} | \vec{\omega}_i^L), \end{aligned}$$

which is less than ξ for $k \geq k'$. \blacksquare

We are now in a position to complete the proof. Suppose \hat{h}_i^t is a private history for player i that leads to the private state \hat{w} with $t \geq T$, and let η be the constant required by Theorem 3. Since \hat{w} and \bar{w} are both reachable in the same period, with positive probability player i observes a private history \hat{h}_i^t that leads to the private state \hat{w} . Moreover, at \hat{h}_i^t his posterior believes that all the other players are in the private state \bar{w} , $\Pr_k(w_{-i}^t = \bar{\mathbf{w}}\mathbf{1} | \hat{h}_i^t)$, is strictly positive for all k , though converging to 0 as $k \rightarrow \infty$ (where \Pr_k denotes probabilities under π^k). If

$\Pr_k(w_{-i}^t \neq \bar{w}\mathbf{1}|\hat{h}_i^t) \leq \eta$, then $\Pr_k(w_{-i}^t = \bar{w}\mathbf{1}|\hat{h}_i^t) > 1 - \eta$, and since $d_i(\hat{w}) \neq d_i(\bar{w})$, Theorem 3 yields the desired conclusion.

Suppose then that $\Pr_k(w_{-i}^t \neq \bar{w}\mathbf{1}|\hat{h}_i^t) > \eta$, and $k > k'$, where k' is from Lemma 8. Lemmas 8 and 9 immediately imply that, as long as $\Pr_k(w_{-i}^{t+\kappa L} \neq \bar{w}\mathbf{1}|\hat{h}_i^t, (\bar{\omega}_i^L)^\kappa) > \eta$, after the first cycle, the odds ratio falls until eventually, $\Pr_k(w_{-i}^{t'} \neq \bar{w}\mathbf{1}|\hat{h}_i^{t'}) \leq \eta$, at which point we are in the first case (since \hat{w} cycles under $\bar{\omega}_i^L$, i 's private state continually returns to \hat{w}).

We conclude by explaining how the difficulty identified in the Introduction is dealt with. In the above argument, the length of the cycle was determined by Lemma 7 from the limit distribution $(\Omega^\infty, \pi^\infty)$, independently of $\Pr_k(w_{-i}^t = \bar{w}\mathbf{1}|\hat{h}_i^t)$. Separation is critical here, since it allows us to focus on a cycle, rather than an entire outcome path. We then considered private-monitoring games sufficiently far out in the sequence, such that along the cycle, state transitions occur as expected with high probability (Lemmas 8 and 9). Since we can use a cycle to manipulate beliefs, the magnitude of the prior is irrelevant; all we need is that $\Pr_k(w_{-i}^t = \bar{w}\mathbf{1}|\hat{h}_i^t) > 0$.

A. Omitted Proofs

Proof of Lemma 4. Suppose $(u^*, (\Omega, \pi))$ is ε -close to $(\tilde{u}^*, (Y, \rho))$ with associated signaling functions (f_1, \dots, f_n) . Then, for all a ,

$$\begin{aligned} & \left| \sum_{\omega_1, \dots, \omega_n} u_i^*(\omega_i, a_i) \pi(\omega_1, \dots, \omega_n | a) - \sum_{y_i} \tilde{u}_i^*(y_i, a_i) \rho(y | a) \right| \\ & \leq \left| \sum_y \sum_{\omega_1 \in f_1^{-1}(y), \dots, \omega_n \in f_n^{-1}(y)} u_i^*(\omega_i, a_i) \pi(\omega_1, \dots, \omega_n | a) - \tilde{u}_i^*(y, a_i) \rho(y | a) \right| \\ & \quad + |Y| \varepsilon \max_{\omega_i, a_i} |u_i^*(\omega_i, a_i)| \\ & \leq \left| \sum_y \tilde{u}_i^*(y, a_i) \left\{ \sum_{\omega_1 \in f_1^{-1}(y), \dots, \omega_n \in f_n^{-1}(y)} \pi(\omega_1, \dots, \omega_n | a) - \rho(y | a) \right\} \right| \\ & \quad + \varepsilon + |Y| \varepsilon \max_{\omega_i, a_i} |u_i^*(\omega_i, a_i)| \\ & \leq 2|Y| \varepsilon \max_{\omega_i, a_i} |u_i^*(\omega_i, a_i)| + \varepsilon + \varepsilon^2 |Y|, \end{aligned}$$

where the first inequality follows from $\sum_y \pi(\{\omega : f_i(\omega_i) = y \text{ for each } i\} | a) > 1 - \varepsilon |Y|$ (an implication of part 1 of Definition 4), the second equality follows from $|\tilde{u}_i^*(y, a_i) - u_i^*(\omega_i, a_i)| < \varepsilon$ for all $i \in N$, $a_i \in A_i$, and $\omega_i \in f_i^{-1}(y)$, and the third inequality from part 1 of Definition 4 and $\max_{y, a_i} |\tilde{u}_i^*(y, a_i)| \leq \max_{\omega_i, a_i} |u_i^*(\omega_i, a_i)| + \varepsilon$. The last term can clearly be made smaller than η by appropriate choice of ε . ■

Proof of Lemma 5. Suppose there exists L such that for all $w, w' \in W$ reachable in the same period and for all $h \in Y^\infty$,

$$\sigma(w, h^L) = \sigma(w', h^L).$$

Then, for all $w, w' \in W$ reachable in the same period and for all $h \in Y^\infty$,

$$d(\sigma(w, h^t)) = d(\sigma(w', h^t)) \quad \forall t \geq L + 1.$$

If $w = \sigma(w^1, y^1, \dots, y^{t-L-1})$ and $w' = \sigma(w^1, \hat{y}^1, \dots, \hat{y}^{t-L-1})$, then

$$\begin{aligned} s(h^t) &= d(\sigma(w, y^{t-L}, \dots, y^{t-1})) \\ &= d(\sigma(w', y^{t-L}, \dots, y^{t-1})) \\ &= d(\sigma(w', \hat{y}^{t-L}, \dots, \hat{y}^{t-1})) = s(\hat{h}^t). \end{aligned}$$

Suppose now the profile s has bounded recall. Let (W, w^1, σ, d) be a representation of s . Suppose w and w' are two states reachable in the same period. Then there exists h^τ and \hat{h}^τ such that $w = \sigma(w^1, h^\tau)$ and $w' = \sigma(w^1, \hat{h}^\tau)$. Then, for all $h \in Y^\infty$, (h^τ, h^t) and (\hat{h}^τ, h^t) agree for the last $t - 1$ periods, and so if $t \geq L + 1$, they agree for at least the last L periods, and so

$$\begin{aligned} d(\sigma(w, h^t)) &= s(h^\tau, h^t) \\ &= s(\hat{h}^\tau, h^t) = d(\sigma(w', h^t)). \end{aligned}$$

Minimality of the representing automaton then implies that for all $h \in Y^\infty$ and $w, w' \in W$ reachable in the same period, $\sigma(w, h^L) = \sigma(w', h^L)$. \blacksquare

Proof of Theorem 3. Let $\phi_i(w)$ be player i 's continuation value from the strategy profile (W, w, σ, d) in the game with public monitoring (i.e., $\phi_i(w)$ is the continuation value of state w under the profile (W, w^1, σ, d)), and let $\phi_i(s_i|w)$ be the continuation value to player i from following the strategy s_i when all the other players follow the strategy profile (W, w, σ, d) . Since the public profile is a strict equilibrium and $|W| < \infty$, there exists $\theta > 0$ such that for all $i, w \in W$, and \tilde{s}_i , a deviation continuation strategy for player i with $\tilde{s}_i^1 \neq d_i(w)$,

$$\phi_i(\tilde{s}_i|w) < \phi_i(w) - \theta.$$

Every strategy \tilde{s}_i in the game with public monitoring induces a strategy s_i in the games with private monitoring that are strongly ε -close in the natural manner:

$$s_i(a_i^1, \omega_i^1; a_i^2, \omega_i^2; \dots, a_i^{t-1}, \omega_i^{t-1}) = \tilde{s}_i(a_i^1, f_i(\omega_i^1); a_i^2, f_i(\omega_i^2); \dots, a_i^{t-1}, f_i(\omega_i^{t-1})).$$

Denote by $V_i^\pi(w)$ the expected value to player i in the game with private monitoring $(u^*, (\Omega, \pi))$ from the private profile induced by (W, w, σ, d) . Let $V_i^\pi(s_i|h_i^t)$ denote player i 's continuation value of a strategy s_i in the game with private monitoring, conditional on the private history h_i^t .

There exists ε and $\eta > 0$ such that for all strategies \tilde{s}_i for player i in the game with public monitoring, and all histories h_i^t for i in the game with private monitoring, if the game with private monitoring is strongly ε -close to the game with public monitoring and $\beta_i(w\mathbf{1}|h_i^t) > 1 - \eta$, then $|V_i^\pi(s_i|h_i^t) - \phi_i(\tilde{s}_i|w)| < \theta/3$, where s_i is the induced strategy in the game with private monitoring. (The argument is essentially the same as that of Mailath and Morris (2002, Lemma 3).)

Suppose there exists a player i , a private history h_i^t , and a state w such that $d_i(w) \neq d_i(\sigma_i(h_i^t))$ and $\beta_i(w\mathbf{1}|h_i^t) > 1 - \eta$. Denote by s_i' the private strategy described by (W, w, σ_i, d_i) , \tilde{s}_i' the public strategy described by (W, w, σ, d_i) , s_i the private strategy described by $(W, \sigma_i(h_i^t), \sigma_i, d_i)$, and \tilde{s}_i the public strategy described by $(W, \sigma_i(h_i^t), \sigma, d_i)$. Then,

$$\begin{aligned} V_i^\pi(s_i'|h_i^t) &> \phi_i(\tilde{s}_i'|w) - \theta/3 = \phi_i(w) - \theta/3 \\ &> \phi_i(\tilde{s}_i|w) + 2\theta/3 \\ &> V_i^\pi(s_i|h_i^t) + \theta/3 \\ &= V_i^\pi(\sigma_i(h_i^t)) + \theta/3, \end{aligned}$$

so that s_i' is a profitable deviation. ■

Proof of Lemma 6. It is immediate that if the profile satisfies the conditions in the lemma, then it is separating. Suppose, then, that the profile is separating. Given the outcome path $h \in Y^\infty$ and state \tilde{w} from the definition of separation, $\sigma(w, h^t)$ denotes the state reached after the first $t - 1$ signals in h from the state w .

The idea is to construct the set W_c by iteratively adding the states necessary to satisfy parts 1 and 2; parts 3 and 4 will then be implications of separation. We start by considering all states reached infinitely often from states in $R(\tilde{w})$ along h . While this implies a cycle of those states, there is no guarantee that other states reachable in the same period will be mapped into the cycle. Accordingly, we include states that are reached infinitely often from states that are reachable under any history in the same period as the states just identified, and so on. Proceeding in this way, we will construct a set of states and a finite sequence of signals with the properties that the states cycle under the sequence, and every state that could arise is mapped under the finite sequence of signals to a cycling state.

We begin by denoting by $\mathbf{w}^1(t)$ the vector of states $(\sigma(w, h^t))_{w \in R(\tilde{w})} \in W^{R(\tilde{w})}$. Since W is finite, so is $W^{R(\tilde{w})}$, and there exists T_1^1 such that for all $\tau \geq T_1^1$, $\mathbf{w}^1(\tau)$ appears infinitely often in the sequence $\{\mathbf{w}^1(t)\}_t$. Let $W^1 \equiv \left\{ \sigma(w, h^{T_1^1}) : w \in R(\tilde{w}) \right\}$, i.e., W^1 is the collection of states that can be reached in period T_1^1 under h , starting from any state in $R(\tilde{w})$. Separation implies $|W^1| \geq 2$. By the definition of T_1^1 , there exists an increasing sequence $\{T_1^k\}_{k=2}^\infty$, with $T_1^k \rightarrow \infty$ as $k \rightarrow \infty$, satisfying, for all $k \geq 2$,

$$\mathbf{w}^1(T_1^k) = \mathbf{w}^1(T_1^1),$$

and for all $t \geq T_1^1$ and $k \geq 1$, there exists a period τ with $T_1^k < \tau \leq T_1^{k+1}$ such that

$$\mathbf{w}^1(t) = \mathbf{w}^1(\tau).$$

The first displayed equation implies that for all $w \in W^1$, $\sigma(w, T_1^1 h^{T_1^k}) = w$ for all k . The second implies that for any state w in $R(\tilde{w})$ and any $t \geq T_1^1$, the state $w' = \sigma(w, h^t)$ appears at least once between each pair of dates T_1^k and T_1^{k+1} , for all k . For $t \geq T_1^1$, $\mathbf{w}^1(t)$ has $|W^1|$ distinct states, and so is equivalent to $\left(\sigma(w, T_1^1 h^t)\right)_{w \in W^1} \in W^{W^1}$.

The recursion is as follows: For a set of states W^k and a period T_k^1 , let $\mathbf{w}^k(t) = \left(\sigma(w, T_k^1 h^t)\right)_{w \in W^k}$ for $t \geq T_k^1$. The recursive step begins with a set of states W^k and an increasing sequence $\{T_k^j\}_{j=1}^\infty$, with $T_k^k \rightarrow \infty$ as $k \rightarrow \infty$, satisfying, for all $k \geq 2$,

$$\mathbf{w}^k(T_k^k) = \mathbf{w}^k(T_k^1),$$

and for all $t \geq T_k^1$ and $k \geq 1$, there exists a period τ with $T_k^k < \tau \leq T_k^{k+1}$ such that

$$\mathbf{w}^k(t) = \mathbf{w}^k(\tau).$$

Define $R(W^k) \equiv \cup_{w \in W^k} R(w)$; note that $W^k \subset R(W^k)$. Let $\mathbf{w}^{k+1}(t)$ denote the vector of states $\left(\sigma(w, T_k^1 h^{T_k^1+t})\right)_{w \in R(W^k)} \in W^{R(W^k)}$. There exists $\hat{t} \geq 1$ such that for all $\tau \geq \hat{t}$, $\mathbf{w}^{k+1}(\tau)$ appears infinitely often in the sequence $\{\mathbf{w}^{k+1}(t)\}_t$. Moreover, there exists $T_{k+1}^1 \geq T_k^1 + \hat{t}$ such that

$$\sigma(w, T_k^1 h^{T_{k+1}^1}) = w \quad \forall w \in W^k.$$

Now, define $W^{k+1} = \{\sigma(w, T_k^1 h^{T_{k+1}^1}) : w \in R(W^k)\}$. By the definition of T_{k+1}^1 , $W^k \subset W^{k+1}$. Just as in the initial step, there is an increasing sequence $\{T_{k+1}^k\}_{k=2}^\infty$, with $T_{k+1}^k \rightarrow \infty$ as $k \rightarrow \infty$, satisfying, for all $k \geq 2$

$$\mathbf{w}^{k+1}(T_{k+1}^k) = \mathbf{w}^{k+1}(T_{k+1}^1),$$

and for all $t \geq T_{k+1}^1$ and $k \geq 1$, there exists a period τ with $T_{k+1}^k < \tau \leq T_{k+1}^{k+1}$ such that

$$\mathbf{w}^{k+1}(t) = \mathbf{w}^{k+1}(\tau),$$

concluding the recursive step.

Since W is finite, this process must eventually reach a point where $W^{k+1} = W^k$. We have thus identified a set of states W^k and two dates T_k^1 and T_k^2 , such that letting $(\bar{y}^1, \dots, \bar{y}^m) \equiv T_k^1 h^{T_k^2}$ and setting $\bar{w} = \sigma(\tilde{w}, h^{T_k^1})$ yields parts 1 and 2 of the Lemma.

Separation implies that under h , for any state $w \in R(\tilde{w}) \setminus \{\tilde{w}\}$ and for all players i , there is some state reached infinitely often from w under h at which i plays differently from the state reached in that period from \tilde{w} . The dates T_k^1 and T_k^2 have been chosen so that any state reached

infinitely often under h from a state $w \in R(\tilde{w})$ appears at least once between T_k^1 and T_k^2 on the path starting in period T_k^1 from the state $\sigma(w, h^{T_k^1})$. Consequently, we have part 3.

Finally, since $|W^1| \geq 2$, $|W_c| \geq 2$. If part 4 does not hold for the current choice of cycle and states, by part 3, it will hold in some period of the cycle $(\bar{y}^1, \dots, \bar{y}^m)$, say period ℓ . Part 4 then holds as well for the cycle beginning in period ℓ , $(\bar{y}^\ell, \dots, \bar{y}^m, \bar{y}^1, \dots, \bar{y}^{\ell-1})$, the state $\bar{w} = \sigma(\tilde{w}, h^{T_k^1}, \bar{y}_1, \dots, \bar{y}^{\ell-1})$, and the set of cycling states is given by $\{\sigma(w, \bar{y}^1, \dots, \bar{y}^{\ell-1}) : w \in W_c\}$. ■

References

- ABREU, D., D. PEARCE, AND E. STACCHETTI (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58(5), 1041–1063.
- BHASKAR, V., AND I. OBARA (2002): "Belief-Based Equilibria in the Repeated Prisoners' Dilemma with Private Monitoring," *Journal of Economic Theory*, 102(1), 40–69.
- COLE, H. L., AND N. R. KOCHERLAKOTA (forthcoming): "Finite Memory and Imperfect Monitoring," *Games and Economic Behavior*.
- COMPTE, O. (1998): "Communication in Repeated Games with Imperfect Private Monitoring," *Econometrica*, 66, 597–626.
- (2002): "On Failing to Cooperate When Monitoring Is Private," *Journal of Economic Theory*, 102(1), 151–188.
- ELY, J. C., J. HÖRNER, AND W. OLSZEWSKI (2005): "Belief-Free Equilibria in Repeated Games," *Econometrica*, 73(2), 377–415.
- ELY, J. C., AND J. VÄLIMÄKI (2002): "A Robust Folk Theorem for the Prisoner's Dilemma," *Journal of Economic Theory*, 102(1), 84–105.
- FUDENBERG, D., AND D. K. LEVINE (2004): "The Nash Threats Folk Theorem with Communication and Approximate Common Knowledge in Two Player Games," unpublished.
- FUDENBERG, D., D. K. LEVINE, AND E. MASKIN (1994): "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62, 997–1040.
- KANDORI, M. (2002): "Introduction to Repeated Games with Private Monitoring," *Journal of Economic Theory*, 102(1), 1–15.
- KANDORI, M., AND H. MATSUSHIMA (1998): "Private Observation, Communication and Collusion," *Econometrica*, 66, 627–652.

- MAILATH, G. J., AND S. MORRIS (2002): "Repeated Games with Almost-Public Monitoring," *Journal of Economic Theory*, 102(1), 189–228.
- MATSUSHIMA, H. (2004): "Repeated Games with Private Monitoring: Two Players," *Econometrica*, 72(3), 823–852.
- MCLEAN, R., I. OBARA, AND A. POSTLEWAITE (2002): "Informational Smallness and Private Monitoring," University of Pennsylvania.
- MCLEAN, R., AND A. POSTLEWAITE (2004): "Informational Size and Efficient Auctions," *Review of Economic Studies*, 71(3), 809–827.
- MONDERER, D., AND D. SAMET (1989): "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior*, 1, 170–190.
- PICCIONE, M. (2002): "The Repeated Prisoner's Dilemma with Imperfect Private Monitoring," *Journal of Economic Theory*, 102(1), 70–83.
- SEKIGUCHI, T. (1997): "Efficiency in Repeated Prisoner's Dilemma with Private Monitoring," *Journal of Economic Theory*, 76, 345–361.