

Coordination failure in repeated games with almost-public monitoring

GEORGE J. MAILATH

Department of Economics, University of Pennsylvania

STEPHEN MORRIS

Department of Economics, Princeton University

Some private-monitoring games, that is, games with no public histories, 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. C72, C73, D82.

1. INTRODUCTION

Intertemporal incentives often allow players to achieve payoffs that are inconsistent with myopic incentives. For repeated 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 et al. 1990). These *perfect public equilibria* (or *PPE*) use public histories to coordinate continuation play.

George J. Mailath: gmailath@econ.upenn.edu

Stephen Morris: smorris@princeton.edu

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 and an anonymous referee and especially the editor, Jeffrey Ely, for valuable suggestions. Mailath is grateful for support from the National Science Foundation under grants #SES-0095768 and #SES-0350969. Morris is grateful for support from the John Simon Guggenheim Foundation, the Center for Advanced Studies in the Behavioral Sciences, and National Science Foundation Grant #SES-0518929.

Copyright © 2006 George J. Mailath and Stephen Morris. Licensed under the [Creative Commons Attribution-NonCommercial License 2.5](http://creativecommons.org/licenses/by-nc/2.5/). Available at <http://econtheory.org>.

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 perfect public equilibrium strategies continue to be equilibria when histories are only almost public. We show that 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).¹ But we also show a partial converse: in games with general almost-public private monitoring, this is the only behavior that can coordinate continuation play under an apparently mild restriction on strategies. To make this precise, we must describe “general but almost-public private monitoring” and characterize the restriction on strategies

When is a general private-monitoring technology close to some public monitoring technology? To be close, there must be a signaling 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 signaling functions (one for each player), the private monitoring is close to the public monitoring if the probability of private signals mapping to a given public signal, under the private-monitoring technology, is close to the probability of that public signal under the public monitoring (for any given action profile). If there exist such signaling functions satisfying this condition, we say there is *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 signaling functions, any strategy profile of the public-monitoring game induces behavior in strongly-close-by almost-public-monitoring games. 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 signaling function. We show that every strict PPE with bounded recall induces equilibrium in every strongly-close-by almost-public-monitoring game; and 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 1**). The idea is that with bounded recall we can always restrict posterior beliefs to be sufficiently close to the public monitoring by requiring the private-monitoring technology to be sufficiently close to the public-monitoring technology. This result generalizes the main result in **Mailath and Morris (2002)**, where the private signal set was assumed to equal the public signal set.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. We call profiles with this property *separating*. While the properties of bounded

¹Thus when we refer to strategy profiles that coordinate continuation play in games with private monitoring, we mean strategy profiles where players' choices are best responses if histories are sufficiently close to being public.

²The extension is nontrivial because 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.

recall and separation do not exhaust possible behavior, they do appear to cover most behaviors of interest.³ When the space of private signals is sufficiently *rich* for some player i 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 that 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 player i is in one private state and assigns arbitrarily high probability to all the other players being in a different common private state.

A significant difficulty 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 to 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: Fix a period t . For any T -length history ($T > t$), 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 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 is required for smaller priors), there is a tension in the determination of T and ε .

We show, however, that this tension can be resolved for separating profiles implementable using a finite number of states. For such profiles the history can be chosen so that not only do the relevant states cycle, but every other state transits under the cycle to a cycling state. The cycle allows us to effectively choose the T above independently of the prior, and gives us our main result (**Theorem 3**): 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.⁴ 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.

³We provide one example of a non-separating profile without bounded recall in **Section 4 (Example 3)**. This profile is not robust to the introduction of private monitoring. We do not know if there exist non-separating profiles without bounded recall that are robust to private monitoring.

⁴Cole and Kocherlakota (2005) 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.

We have analyzed the robustness of fixed strategy profiles to private monitoring. Our results do not say anything about the set of all equilibrium payoffs in private-monitoring games.⁵ While this classic question is important, we believe that there are at least three reasons why it is nonetheless also interesting to focus on a fixed strategy profile. First, researchers using repeated game theory to understand economic phenomena are interested in hypothesizing and testing particular strategy profiles.⁶ Second, understanding properties of particular strategy profiles may turn out to be an important step in characterizing the set of all equilibrium payoffs. Finally, one of our findings is that fine details of strategy profiles, such as history dependence, that are irrelevant for the classic recursive characterization of the PPE payoff set are very important for the robustness question we consider, and such fine details might turn out to be significant for other questions as well.

Both our positive and negative results restrict attention to *strict* PPE, and the assumption is important for both kinds of results. In such equilibria, players are not indifferent between alternative actions and are thus coordinated in their continuation play. Such strategy profiles capture basic intuitions about how cooperation can be sustained in repeated games by the threat of coordinated deviation to punishment paths; they form the basis of empirical applications of repeated game theory (see the references in footnote 6); and we believe they are interesting objects of study. However, as noted in footnote 5, the most permissive results in the private-monitoring literature have used strategies with a significant amount of randomization and indifference. The results in this paper do not have anything to say about the robustness of such strategies.⁷

This paper introduces a useful representation of finite state strategies for private-monitoring games. Each player has a finite set of private states, a transition function mapping private signals and states into new states, and decision rules for the players, specifying behavior in each state. 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 is *not* sufficient to verify that the strategies are

⁵Mailath and Samuelson (2006, Chapter 12) introduces the main issues and concepts. See 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.” 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 and Bhaskar and Obara 2002) and those that are “belief-free” (for example, Piccione 2002, Ely and Välimäki 2002, and Matsushima 2004), where equilibrium strategies are constructed using randomization to ensure that players are indifferent between some actions at all histories. While folk theorems cannot be proved using belief-free strategies for general payoff matrices (Ely et al. 2005), variations on belief-free strategy profiles have been used to prove general folk theorems (Hörner and Olszewski 2005).

⁶See, for example, Axelrod (1984), Ellison (1994), and Greif (2005).

⁷Bhaskar and van Damme (2002) and Ely (2002) show that trigger strategy profiles, which are strict PPE of a repeated prisoners’ dilemma with imperfect public monitoring, can be approximated in nearby games with private monitoring with a strategy profile with strict mixing. This possibility suggests that allowing non-strict equilibria may greatly assist in establishing robustness results. On the other hand, the equilibria with mixing require players to randomize differently at different payoff-equivalent histories, which is arguably implausible. Bhaskar (1998) and Bhaskar and van Damme (2002) suggest that such strategies often do not survive extensive form purification perturbations.

optimal. It is also necessary to know how each player's beliefs over the other players' private states 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. The crux of our analysis concerns how to track the evolution of beliefs over other players' private states during the course of play. In this paper, we use this representation to analyze private-monitoring profiles constructed from a PPE. However, the method is more general and we believe that it may be of more general use. Examples can be found in [Mailath and Samuelson \(2006, Section 12.4 and Chapter 14\)](#), where the method is used to analyze the mixed strategy employed in the classic analysis of [Sekiguchi \(1997\)](#) and define belief-free equilibria.

Finally, we note that we have not allowed any communication beyond that contained in the equilibrium strategies. We view our findings as underlining the importance of public 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 et al. \(2002\)](#).

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 each period, each player learns only the realized value of a private signal. There are n players, with the 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) depends only on (ω_i, a_i) . We denote this payoff 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 is 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

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

is the set of private histories for player i .

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 each 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 , the action it prescribes depends only on the public history $h^t \in Y^{t-1}$, and not on i 's private history. Henceforth, by the term *public profile*, we 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 any public history h^t , specifies a Nash equilibrium for the repeated game. Under imperfect full-support public monitoring, every public history arises with positive probability, and so every Nash equilibrium in public strategies is a 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, each 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, each 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, $\cup_{t \geq 1} Y^t$, 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.

Denote the vector of average discounted expected values of following the public profile (W, w, σ, d) (so that 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 et al. 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.⁸ 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 $g(\cdot; w)$.⁹

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 $a \in A$.

⁸We have introduced a distinction between W and the set of continuation payoffs for convenience. Any pure strategy equilibrium *payoff* can be supported by an equilibrium where $W \subset \mathbb{R}^I$ and $\phi(w) = w$ (again, see [Abreu et al. 1990](#)).

⁹Equivalently, 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 et al. 1994](#), Theorem 6.4 and remark).

We extend the domain of σ from $W \times Y$ to $W \times \cup_{t=1}^{\infty} Y^t$ by recursively defining $\sigma(w^1, h^t) = \sigma(\sigma(w^1, h^{t-1}), y^t)$ for all $h^t \in Y^{t-1}$, where $h^t = (h^{t-1}, y^t)$.

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

The restriction to minimal automata is without loss of generality: every profile has a minimal representing automaton. Moreover, this automaton is essentially unique.¹⁰ Accordingly, we treat a public strategy profile and its minimal representing automaton interchangeably.

2.3 Almost-public monitoring

We now define what it means for a private-monitoring distribution to be close to a public-monitoring distribution.

DEFINITION 2. 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 $y \in Y$,

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

and

2. for all $y \in Y$, $\omega_i \in f_i^{-1}(y)$, and all $a \in A$, if $\pi(\{\omega_i\} \mid a) > 0$, then

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

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.

¹⁰Suppose (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 $\hat{h}^\ell$ be a public history reaching $\hat{w}$ (i.e., $\hat{w} = \sigma(w^1, \hat{h}^\ell)$), and set $\varphi(\hat{w}) = \tilde{\sigma}(\tilde{w}^1, \hat{h}^\ell)$. 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, $\tilde{\sigma}(\tilde{w}, y) = \varphi(\sigma(\varphi^{-1}(\tilde{w}), y))$, and $d(w) = \tilde{d}(\varphi(w))$.

The condition of ε -closeness in [Definition 2](#) 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 1. 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).

DEFINITION 3. 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 say also that such a private-monitoring game has *almost-public monitoring*.

Note that because of our maintained assumption that public-monitoring games have full support monitoring, a private-monitoring game that has almost-public monitoring relative to a fixed ρ does not have “almost perfect” monitoring in the sense usually assumed in the literature.¹¹

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 2. 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 $Y \cup \{\emptyset\}$. Note that by construction, each player's strategy is “action-free,” i.e., it depends only on past signals and not on past actions of that player. (See [Mailath and Samuelson 2006](#), Chapter 12 for more discussion of “action-free.”)

¹¹The order of quantifiers is important: We can construct almost-perfect almost-public monitoring distributions by considering full-support public-monitoring distributions arbitrarily close to perfect monitoring—see [Mailath and Morris \(2002, Section 6\)](#).

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

FIGURE 1. The prisoners' dilemma.

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 $\beta_i^t \in \Delta(W^{N-1})$. With a slight abuse of notation, we write $\beta_i(w_{-i} | h_i^t)$ for the probability that player i assigns to his opponents being in private states w_{-i} at history h_i^t . 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. For any private history h_i^t , we write $w_i^t = \sigma_i(h_i^t)$ for the private state of the player in period t .

REMARK 1. In private-monitoring games that are ε -close, but not strongly so, a public profile induces only that part of the private profile determined by histories of signals $\omega_i \in f_i^{-1}(Y)$, with the remaining specification of behavior not determined by the public profile. For an example, see part (ii) of [Theorem 1](#); see also [footnote 15](#).

2.4 Prisoners' dilemma examples

We illustrate our definitions and results using the repeated prisoners' dilemma under various monitoring assumptions. The ex ante stage game is given by the normal form in [Figure 1](#).¹²

First, consider the leading example from [Mailath and Morris \(2002, Section 3.3\)](#). The example illustrates that without bounded recall, beliefs may vary in extreme ways to prevent a strict PPE from being an equilibrium in nearby private monitoring games.

EXAMPLE 1. In the benchmark public-monitoring game, the set of public signals is $Y = \{\underline{y}, \bar{y}\}$ and the 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

¹²Here (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 2](#)) and all relevant incentive constraints are strict.

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

FIGURE 2. The probability distribution of the private signals for **Example 2**. 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.

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 consider the ε -close private-monitoring technology where $\Omega_i = Y$ and the signaling functions are the identity functions. 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 has 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 transits 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.) $\diamond$

We now consider a richer case where the private signal set is not equal to the public signal set. The example illustrates that allowing richer signal sets may be important.

EXAMPLE 2. Let $\Omega_1 = \{\underline{y}_1, \bar{y}_1', \bar{y}_1''\}$ and $\Omega_2 = \{\underline{y}_2, \bar{y}_2\}$. The probability distribution of the signals is given in **Figure 2**. This private-monitoring distribution is $\sqrt{\varepsilon}$ -close to the public-monitoring distribution of **Example 1** 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 $\min\{\alpha', \alpha - \alpha'\}$. In **Example 1**, we argued that if $q < r$, grim trigger induces Nash equilibrium behavior in close-by private-monitoring games with $\Omega_i = Y$. We now argue that under the richer private-monitoring distribution of this example, even if $q < r$, grim trigger does

not induce Nash equilibrium behavior in some close-by games. In particular, suppose $0 < r' < q' < q < r$. Under this parameter restriction, the signal $\tilde{y}_1''$ after n_1 is indeed a signal that player 2 has also played n_2 . However, the signal $\tilde{y}_1'$ after n_1 is a signal that player 2 has played e_2 and so a sufficiently long private history of the form $(e_1 y_1, n_1 \tilde{y}_1', n_1 \tilde{y}_1', \dots, n_1 \tilde{y}_1')$ leads to a posterior for player 1 at which n_1 is not optimal. $\diamond$

3. PPE WITH BOUNDED RECALL

As we saw in Examples 1 and 2, 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 induce equilibria in almost-public-monitoring games.¹³

DEFINITION 4. A public profile s has *L-bounded recall* if 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 , $W_t \equiv \{w \in W : w = \sigma(w^1, h^t) \text{ for some } h^t, \text{ where } w^1 \text{ is the initial state}\}$. The following characterization of bounded recall is useful.

LEMMA 3. *The public profile induced by the minimal automaton (W, w^1, σ, d) has L-bounded recall if and only if for all t , all $w, w' \in W_t$, and all $h \in Y^\infty$,*

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

If a public profile induced by a finite automaton (W, w^1, σ, d) , where W has K elements, does not have $K(K-1)$ -bounded recall, then it has unbounded recall.

PROOF. The first claim is proved in the [Appendix](#).

For the second claim, suppose that the profile induced by the finite automaton (W, w^1, σ, d) , where W has K elements, does not have $K(K-1)$ -bounded recall. From the first claim, for some t , there exist $w, w' \in W_t$ and history $h \in Y^\infty$ such that

$$\sigma(w, h^\tau) \neq \sigma(w', h^\tau),$$

for $\tau = 1, \dots, K(K-1)$. The sequence $\{(\sigma(w, h^\tau), \sigma(w', h^\tau))\}_{\tau=0}^{K(K-1)}$, where $(\sigma(w, h^0), \sigma(w', h^0)) = (w, w')$, consists of $K(K-1)+1$ terms of pairs of states. Since pairs of identical states cannot arise, some pair of nonidentical states must be repeated. That is, there

¹³Denote a dummy signal by $*$. [Mailath and Morris \(2002\)](#) use the term *bounded memory* for public profiles with the property that there is an integer L such that a representing automaton is given by $W = (Y \cup \{*\})^L$, $\sigma((y^2, \dots, y^L, y), y) = (y^2, \dots, y^L, y)$ for all $y \in Y$, and $w^1 = (*, \dots, *)$. Our earlier notion implicitly imposes a time homogeneity condition, since the caveat in [Lemma 3](#) 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.

exist $w'' \neq w'''$ and periods $0 \leq \tau_1 < \tau_2 \leq K(K-1)$ such that

$$(w'', w''') = (\sigma(w, h^{\tau_1}), \sigma(w', h^{\tau_1})) = (\sigma(w, h^{\tau_2}), \sigma(w', h^{\tau_2})).$$

Now we have w'' , reachable in the same period as w''' infinitely often, such that letting $\tilde{h}$ be the infinite repetition of the cycle of outcomes ${}^{\tau_1}h^{\tau_2}$, we have $\sigma(w'', \tilde{h}^t) \neq \sigma(w''', \tilde{h}^t)$ for all t . $\square$

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, $H_i^u = \{h_i^t : \omega_i^t \in f_i^{-1}(\emptyset), \text{ some } \tau \text{ satisfying } t-L \leq \tau \leq t-1\}$. This is the set of private histories for which in any of the last L periods, a private signal ω_i^t satisfying $f_i(\omega_i^t) = \emptyset$ 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,

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

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

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

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 (recall [Lemma 1](#)). 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 1. *Fix a full-support public-monitoring game $(\tilde{u}^*, (Y, \rho))$ and a strict perfect public 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 $(\tilde{u}^*, (Y, \rho))$,*

- (i) *if $f_i(\Omega_i) = Y$ for all i , the induced private profile is a Nash equilibrium; and*

- (ii) 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 mutual minmax folk theorem to this more general notion of nearby private-monitoring distributions.¹⁴

4. FAILURE OF COORDINATION

Examples 1 and 2 illustrate 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 (rather than ε -closeness) does not limit its usefulness. The assumption clarifies the source of the failure of the induced profile to be a Nash equilibrium, which is not due to a difficulty with interpreting “uninterpretable” signals. Moreover, this failure arises in *any* strongly ε -close game in which the belief hypothesis holds. Recall also that the public profile completely determines a strategy profile in a private-monitoring game only when the private-monitoring game is strongly ε -close (Section 2.3).¹⁵

THEOREM 2. *Suppose the public profile (W, w^1, σ, d) is a strict equilibrium of the full-support public-monitoring game $(\tilde{u}^*, (Y, \rho))$ 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 $(\tilde{u}^*, (Y, \rho))$, 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 $\beta_i(w \mid h_i^t) > 1 - \eta$, then the induced private profile is not a Nash equilibrium of the game with private monitoring for the same δ .*

¹⁴We incorrectly claimed that the profile described in the “proof” of the almost-public almost-perfect folk theorem (Mailath and Morris 2002, Theorem 6.1) has bounded recall. See Mailath and Samuelson (2006, Proposition 13.6.1) for a proof of the weaker result reported in the text.

¹⁵The result does extend to private-monitoring games that are ε -close, but not strongly so. Any pure private strategy for i can be represented as an automaton $(\tilde{W}_i, \tilde{w}^1, \tilde{\sigma}_i, \tilde{d}_i)$, where $\tilde{\sigma}_i : \tilde{W}_i \times A_i \times \Omega_i \rightarrow \tilde{W}_i$ and (as usual) $\tilde{d}_i : \tilde{W}_i \rightarrow A_i$. Say a private profile $(\tilde{W}_i, \tilde{w}^1, \tilde{\sigma}_i, \tilde{d}_i)_i$ reflects the public profile (W, w^1, σ, d) if for all i (perhaps after relabeling states, see footnote 10) $W \subset \tilde{W}_i$, $\tilde{w}^1 = w^1$, $\tilde{\sigma}_i(w, a_i, \omega_i) = \sigma(w, f_i(\omega_i))$ for all $(w, a_i, \omega_i) \in W \times A_i \times f_i^{-1}(Y)$, and finally, $\tilde{d}_i(w) = d_i(w)$ for all $w \in W$.

Then, there exists $\eta > 0$ and $\varepsilon > 0$ such that for any close-by private-monitoring game and any private profile reflecting the public profile, if there is a player and a private history, and a state $w \in W \subset \tilde{W}_i$ with the specified properties, then the private profile is not a Nash equilibrium.

We implicitly used this result in our discussions of the repeated prisoners' dilemma. For example, in [Example 1](#), 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 the selection of a private history that the hypotheses of [Theorem 2](#) 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 1](#) 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 3](#)). 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

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

(where $\pi_i(\omega_i | a)$ is the probability that player i observes the signal ω_i under a). 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 2](#) 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 that signal does 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 $w^\dagger$. Suppose also $A_2 = \{\tilde{a}_2, \hat{a}_2, a_2^\dagger\}$. As before, suppose there is a signal y such that $\hat{w} = \sigma(\hat{w}, y)$, $\bar{w} = \sigma(\bar{w}, y)$, and $w^\dagger = \sigma(w^\dagger, y)$, that is, any player in the state $\hat{w}$, $\bar{w}$, or $w^\dagger$ 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 $\tilde{a}_1 = d_1(\hat{w})$, $\tilde{a}_2 = d_2(\bar{w})$, $\hat{a}_2 = d_2(\hat{w})$, and $a_2^\dagger = d_2(w^\dagger)$, and suppose there are two private signals, ω'_1 and ω''_1 consistent with y , satisfying

$$\pi_1(\omega'_1 | \tilde{a}_1, a_2^\dagger) > \pi_1(\omega'_1 | \tilde{a}) > \pi_1(\omega'_1 | \tilde{a}_1, \hat{a}_2)$$

and

$$\pi_1(\omega''_1 | \tilde{a}_1, \hat{a}_2) > \pi_1(\omega''_1 | \tilde{a}) > \pi_1(\omega''_1 | \tilde{a}_1, a_2^\dagger).$$

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

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

as desired, but $\Pr(w_2 = w^\dagger | h_1^t, \omega'_1) / \Pr(w_2 = \bar{w} | h_1^t, \omega'_1)$ increases. On the other hand, after observing another private signal ω''_1 , also consistent with y , while the odds ratio $\Pr(w_2 = w^\dagger | h_1^t, \omega''_1) / \Pr(w_2 = \bar{w} | h_1^t, \omega''_1)$ 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,

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

and

$$\frac{\pi_1(\omega'_1 | \tilde{a}_1, a_2^\dagger)}{\pi_1(\omega'_1 | \tilde{a})} \left(\frac{\pi_1(\omega''_1 | \tilde{a}_1, a_2^\dagger)}{\pi_1(\omega''_1 | \tilde{a})} \right)^2 < 1.$$

Our richness condition on private-monitoring distributions captures this idea. For a private-monitoring distribution (Ω, π) , define

$$\gamma_{aa'_i}(\omega_i) \equiv \log \pi_i(\omega_i | a_i, a_{-i}) - \log \pi_i(\omega_i | a_i, a'_{-i})$$

and let $\gamma_a(\omega_i) = (\gamma_{aa'_i}(\omega_i))_{a'_i \in A_{-i}, a'_i \neq a_{-i}}$ denote the vector in $\mathbb{R}^{|A_{-i}|-1}$ 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.¹⁶

DEFINITION 5. A private-monitoring distribution (Ω, π) is *rich* for player i , given his signaling function f_i , 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 $\mathbb{R}_{++}^{|A_{-i}|-1}$.

¹⁶The convex combination is strictly positive (rather than negative) because the definition of $\gamma_{aa'_i}$ inverts the odds ratios from the displayed equations.

Note that we require only that private monitoring be rich for one player.

It is useful to quantify the extent to which the conditions of [Definition 5](#) are satisfied. Since the spaces of signals and actions are finite, the number of constraints in [Definition 5](#) is finite, 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.¹⁷

DEFINITION 6. Given f , the *richness* of a rich private-monitoring distribution (Ω, π) for i is the supremum of all $\zeta > 0$ satisfying: for all $a \in A$ and all $y \in Y$, the convex hull of the set of vectors $\{\gamma_a(\omega_i) : \omega_i \in f_i^{-1}(y) \text{ and } \pi_i(\omega_i | a_i, a'_{-i}) \geq \zeta \text{ for all } a'_{-i} \in A_{-i}\}$ has a nonempty intersection with $\mathbb{R}_\zeta^{|A_{-i}|-1} \equiv \{x \in \mathbb{R}_+^{|A_{-i}|-1} : x_k \geq \zeta \text{ for } k = 1, \dots, |A_{-i}| - 1\}$.

The second weakening of the logic of [Example 1](#) described above 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 $\bar{w} = \sigma(\bar{w}, y)$. Thus along the history $(y, y, \dots)$, if the player started out in distinct states $\hat{w}$ or $\bar{w}$, he would remain in those distinct states and would continue to play in distinct ways. But the logic continues to hold if there exists an arbitrary history h such that some distinct initial states lead to distinct states forever and if, from such distinct states, play is distinct along that particular history infinitely often. This is the idea behind the following definition of a separating strategy profile.

Define $R(\tilde{w})$ as the set of states that are repeatedly reachable in the same period as $\tilde{w}$ (i.e., $R(\tilde{w}) = \{w \in W : \{w, \tilde{w}\} \subset W_t \text{ infinitely often}\}$). 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 ${}^\tau h^{\tau+t} = (y^\tau, y^{\tau+1}, \dots, y^{\tau+t-1})$. Consider a *continuation path* $(\tilde{w}, h)$ consisting of an initial state $\tilde{w}$ followed by an outcome path h . The continuation path $(\tilde{w}, h)$ satisfies *state-separation* if there is another state $w \in R(\tilde{w})$ such that starting in state w instead of $\tilde{w}$ would lead to distinct states into the infinite future: formally, there exists another state $w \in R(\tilde{w})$ that satisfies $\sigma(w, h^t) \neq \sigma(\tilde{w}, h^t)$ for all t . In this case, state w is separated from $\tilde{w}$ along history h . Recall from the proof of the second claim in [Lemma 3](#) that every unbounded recall profile induced by a *finite* automaton has a continuation path $(\tilde{w}, h)$ satisfying state-separation.

The logic of our proof requires not only state-separation, but in addition distinct behavior on the continuation path satisfying state-separation. The continuation path $(\tilde{w}, h)$ satisfies *behavior-separation* if whenever state $w \in R(\sigma(\tilde{w}, h^\tau))$ is separated from $\sigma(\tilde{w}, h^\tau)$, then all players choose different actions along the outcome path ${}^\tau h$ infinitely often. Formally, 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

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

Notice that every continuation path satisfies behavior-separation if, for each player, distinct states always lead to distinct actions. The need to behavior-separate the state $\tilde{w}$ from every other state that can be reached infinitely often is illustrated by our earlier discussion: because private monitoring implies all such states are assigned positive

¹⁷The bound ζ appears twice in the definition. Its first appearance ensures that for all $\zeta > 0$, there is a 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 3. The normal form for **Example 3**.

probability by a player's beliefs, we need to have signals that are informative about these states relative to $\tilde{w}$. Now we have:

DEFINITION 7. The public strategy profile is *separating* if there is a state $\tilde{w}$ and an outcome path $h \in Y^\infty$ such that $(\tilde{w}, h)$ satisfies state-separation and behavior-separation.

Clearly, a separating profile cannot have bounded recall. The key question is how much stronger is this property than having unbounded recall under the restriction to finite state strategies. Since every finite unbounded recall profile has a state-separating path, the *only* way a finite state strategy profile with unbounded recall can fail separation is if every continuation path satisfying state-separation fails behavior-separation. The following example illustrates this possibility.

EXAMPLE 3. The stage game is given in **Figure 3**. In the public-monitoring game, there are two public signals, y' and y'' , with distribution ($0 < q < p < 1$)

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

Finally, the public profile is illustrated in **Figure 4**. Under any outcome path in which the sequence $y'y'$ or $y'y''$ occurs, all states transit to the same state. Under any outcome path in which only y'' appears, the state eventually cycles between w^A and $\hat{w}^A$. Thus continuation path (w, h) is state-separating only if $h = (y'', y'', \dots)$. But this continuation path is not behavior separating, since action A is then played forever. $\diamond$

We think of this failure as pathological. In this example, it is easy to see that the profile is not robust. After enough realizations of private signals corresponding to y'' , beliefs must assign roughly equal probability to w^A and $\hat{w}^A$,¹⁸ 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).

We do not have an example of a finite state strategy profile with unbounded recall that fails separation but is robust. **Example 3** suggests an intuition why such an example might be hard to find: a strategy profile with unbounded recall can fail separation only if all state-separated states give rise to identical behavior most of the time. With the

¹⁸The details of this calculation can be found in **Mailath and Samuelson (2006, Example 13.4.6)**.




FIGURE 4. The strategy profile for **Example 3**. 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.

possibility of belief drift, as in the example, it seems hard to make this consistent with equilibrium. Moreover, this possibility of drift implies also that showing that these unbounded recall strategy profiles are not robust requires a quite different proof strategy than that pursued in this paper.

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:¹⁹

DEFINITION 8. 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 3. 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 signaling function f to $(\tilde{u}^*, (Y, \rho))$ with (Ω, π) having richness, given f , for some player i of 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 a function only 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.

¹⁹If 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.

Thus, separating strict PPE of public-monitoring games are not robust to the introduction of private monitoring. It, of course, implies also 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 3](#).)

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 $(u^k, (\Omega, \pi^k))$ strongly $1/k$ -close to some public-monitoring game $(\tilde{u}^*, (Y, \rho))$ and $\{(\Omega, \pi^k)\}$ a rich (for some player i) sequence of distributions. Fix a pure strategy profile of the private-monitoring game in which each player's strategy respects his signaling function f_j (i.e., $\sigma_j(h_j, a_j, \omega_j) = \sigma_j(h_j, a_j, \hat{\omega}_j)$ if $f_j(\omega_j) = f_j(\hat{\omega}_j) \neq \emptyset$). 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 seems 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.

5. THE PROOF OF [THEOREM 3](#)

Our proof exploits an alternative characterization of separation that holds for finite state strategies, reported in the next lemma and corollary (proved in the [Appendix](#)).

LEMMA 4. *A finite public strategy profile of the public-monitoring game is separating if, and only if, there is a finite sequence of signals h^m , a collection of states W_c , and a state $\bar{w} \in W_c$ such that*

- (i) $\sigma(w, h^m) = w$ for all $w \in W_c$,
- (ii) $\sigma(w, h^m) \in W_c$ for all $w \in R(\bar{w})$,
- (iii) $\forall w \in W_c \setminus \{\bar{w}\}, \forall i \exists \ell, 2 \leq \ell \leq m$, such that

$$d_i(\sigma(w, h^\ell) \neq d_i(\sigma(\bar{w}, h^\ell),$$

and

- (iv) $|W_c| \geq 2$.

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.

COROLLARY 2. *Suppose (W, w^1, σ, d) is the minimal automaton of a separating finite public strategy profile. For any player i , the history h^m , set of states W_c , and state $\bar{w} \in W_c$ can be chosen so that, in addition, $d_i(\hat{w}) \neq d_i(\bar{w})$ for some $\hat{w} \in W_c \setminus \{\bar{w}\}$.*

The proof of **Theorem 3** is by contradiction. 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.

The basic argument is most easily seen if the space of signals for each player is 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 the 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 derail the process (we 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

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

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 $\cup_{a_i \in A_i} \Omega_i^{k, a_i}(y)$, and give the ℓ -th signal in the order the label (y, ℓ) . Let $\lambda_{i, y} \equiv |\cup_{a_i \in A_i} \Omega_i^{k, a_i}(y)|$; note that $\lambda_{i, y}$ is (crudely) bounded above by $\lambda^* \equiv |A_i|/\zeta$ for all k (recall **footnote 17**). With this relabeling, and defining $\Omega_i \equiv \cup_{y \in Y} \{(y, 1), (y, 2), \dots, (y, \lambda^*)\}$, a finite set, we have, for all i and k ,

$$\Omega_i^k \subset \Omega_i \cup \left(\Omega_i^k \setminus \left(\cup_{a_i \in A_i} \Omega_i^{k, a_i} \right) \right) \quad (1)$$

and

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

Without loss of generality, we can assume (1) 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 $\omega_i^y \equiv \{\omega_i \in \Omega_i^k \setminus (\cup_{a_i \in A_i} \Omega_i^{k, a_i}) : f_i(\omega_i) = y\}$ 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) .

Because there are only a finite number of players, by passing to a further subsequence if necessary, we can assume that the private-monitoring distribution is rich for the *same* player i ; we call this player the *rich* player. Moreover, by passing to yet a further subsequence if necessary, we can assume also that, for the rich player i , $a_i \in A_i$, and $y \in Y$, the convex hull of the set of vectors $\{\gamma_a^\infty(\omega_i) : \omega_i \in f_i^{-1}(y), \pi_i^\infty(\omega_i | a_i, a'_{-i}) > \zeta\}$ for all $a'_{-i} \in A_{-i}$ has a nonempty intersection with $\mathbb{R}_\zeta^{|A_{-i}|-1}$, 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 $\gamma_a^\infty(\omega_i) = (\gamma_{aa'_{-i}}^\infty(\omega_i))_{a'_{-i} \in A_{-i}, a'_{-i} \neq a_{-i}}$.

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 2**. 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 5**), 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 5. *Fix a finite separating public profile of the public-monitoring game, and let $\bar{w}$, $\bar{w}$, W_c , be the states and set of states identified in **Corollary 2** for the rich player i . 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*

- (i) $\sigma_i(\bar{w}, \vec{\omega}_i^L) = \bar{w}$,
- (ii) for all sequences of private signals, $\vec{\omega}_j^L$, for any 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
- (iii) for all $\mathbf{w} \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}$,

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

where $\Pr_\infty$ 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 4** induces a cycle in the states $\bar{w} = \bar{w}^1, \dots, \bar{w}^{m+1} = \bar{w}^1$ and $\hat{w} = \hat{w}^1, \dots, \hat{w}^{m+1} = \hat{w}^1$. We index the cycle by ℓ and write $\bar{a}^\ell = d(\bar{w}^\ell)$ and $\hat{a}^\ell = d(\hat{w}^\ell)$. Let $\tilde{a}^\ell \equiv (\hat{a}_i^\ell, \hat{a}_{-i}^\ell)$. 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 $a'_{-i} \neq \bar{a}_{-i}^\ell$,

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

Since

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

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

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

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 $\check{a}_{-i}^\ell$ denote the action profile taken in period ℓ of the cycle. Then,

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

For $\mathbf{w} \neq \bar{w}\mathbf{1}$, in each period at least one player is in a private state different from $\bar{w}$. From **Lemma 4.2**, $\check{a}_{-i}^\ell \neq \tilde{a}_{-i}^\ell$ for at least one ℓ , and so $A(\vec{\omega}_i^L; \mathbf{w})$ must be strictly less than 1. $\square$

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 6. *Assume the hypotheses of **Lemma 5**, and let h_i^t be a private history for player i satisfying $\bar{w} = \sigma_i(h_i^t)$. For all $\eta > 0$, there exist $\xi > 0$ 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*

$$\frac{\Pr_k(w_{-i}^{t+L} \neq \bar{w}\mathbf{1} | \vec{\omega}_i^L, h_i^t)}{\Pr_k(w_{-i}^{t+L} = \bar{w}\mathbf{1} | \vec{\omega}_i^L, h_i^t)} < (1 - \xi) \frac{\Pr_k(w_{-i}^t \neq \bar{w}\mathbf{1} | h_i^t)}{\Pr_k(w_{-i}^t = \bar{w}\mathbf{1} | h_i^t)}, \quad (2)$$

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 5**.

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,

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

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

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

Moreover,

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

Defining

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

we have,

$$\begin{aligned} \frac{\Pr_k(w_{-i}^{t+L} \neq \bar{w}\mathbf{1} \mid \vec{\omega}_i^L)}{\Pr_k(w_{-i}^{t+L} = \bar{w}\mathbf{1} \mid \vec{\omega}_i^L)} \\ < \frac{\Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} \mid w_{-i}^t \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}) + x^t(k)}{\Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} \mid w_{-i}^t = \bar{w}\mathbf{1})} \times \frac{\Pr_k(w_{-i}^t \neq \bar{w}\mathbf{1})}{\Pr_k(w_{-i}^t = \bar{w}\mathbf{1})} \\ \leq \frac{\max_{\mathbf{w} \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}} \Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} \mid w_{-i}^t = \mathbf{w}) + x^t(k)}{\Pr_k(\vec{\omega}_i^L, \vec{y}_{-i} \mid w_{-i}^t = \bar{w}\mathbf{1})} \times \frac{\Pr_k(w_{-i}^t \neq \bar{w}\mathbf{1})}{\Pr_k(w_{-i}^t = \bar{w}\mathbf{1})}. \quad (3) \end{aligned}$$

From Lemma 5,

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

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

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

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 $k \geq k''$,

$$\frac{\max_{\mathbf{w} \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}} \Pr_k(\bar{\omega}_i^L, \bar{y}_{-i} \mid w_{-i}^t = \mathbf{w}) + \xi'}{\Pr_k(\bar{\omega}_i^L, \bar{y}_{-i} \mid w_{-i}^t = \bar{w}\mathbf{1})} < 1 - \xi'. \quad (4)$$

Since (Ω, π^k) is strongly $1/k$ -close to (Y, ρ) , $\lim_{k \rightarrow \infty} \Pr_k(\neg \bar{y}_{-i} \mid \bar{\omega}_i^L) = 0$, and so there exists k''' such that $\Pr_k(\neg \bar{y}_{-i} \mid \bar{\omega}_i^L) < \xi' \eta / 2$ for all $k \geq k'''$. Suppose $\xi = \xi' \eta / 2$ and $k' = \max\{k'', k'''\}$. Since $\eta < \Pr_k(w_{-i}^t \in W_c^{n-1} \setminus \{\bar{w}\mathbf{1}\}) \leq \Pr_k(w_{-i}^t \neq \bar{w}\mathbf{1})$, $x^t(k) \leq \xi'$. Consequently (4), with (3), implies (2) (since $\xi < \xi'$). $\square$

Lemma 4 guarantees that one round of the cycle of signals always takes a state not in W_c into W_c , ensuring that the probability on states in $W \setminus W_c$ can be controlled.

LEMMA 7. Assume the hypotheses of **Lemma 5**, 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 6**. There exists T such that if $t \geq T$, then for all $k > k'$,

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

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 $\Pr_k(w_{-i}^{t+L} \notin W_c^{n-1}, \bar{y}_{-i}) = 0$, and so

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

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

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 2**. 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 belief that all the other players are in the private state $\bar{w}$, $\Pr_k(w_{-i}^t = \bar{w}\mathbf{1} \mid \hat{h}_i^t)$, is strictly positive for all k , though converging to 0 as $k \rightarrow \infty$ (where $\Pr_k$ denotes probabilities under π^k). If $\Pr_k(w_{-i}^t \neq \bar{w}\mathbf{1} \mid \hat{h}_i^t) \leq \eta$, then $\Pr_k(w_{-i}^t = \bar{w}\mathbf{1} \mid \hat{h}_i^t) > 1 - \eta$, and since $d_i(\hat{w}) \neq d_i(\bar{w})$, **Theorem 2** yields the desired conclusion.

Suppose then that $\Pr_k(w_{-i}^t \neq \bar{w}\mathbf{1} \mid \hat{h}_i^t) > \eta$ and $k > k'$, where k' is from **Lemma 6**. Lemmas 6 and 7 immediately imply that, as long as $\Pr_k(w_{-i}^{t+L} \neq \bar{w}\mathbf{1} \mid h_i^t, (\bar{\omega}_i^L)^k) > \eta$, after the first cycle, the odds ratio falls until eventually $\Pr_k(w_{-i}^{t'} \neq \bar{w}\mathbf{1} \mid h_i^{t'}) \leq \eta$, 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 $\hat{w}$).

REMARK 2. How is the difficulty identified in the Introduction dealt with? In the above argument, the length of the cycle was determined by Lemma 5 from the limit distribution $(\Omega^\infty, \pi^\infty)$, independently of $\Pr_k(w_{-i}^t = \bar{w}\mathbf{1} \mid \hat{h}_i^t)$. 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 6 and 7). Since we can use a cycle to manipulate beliefs, the magnitude of the prior is irrelevant; all we need is that $\Pr_k(w_{-i}^t = \bar{w}\mathbf{1} \mid \hat{h}_i^t) > 0$.

REMARK 3. The difficulty with extending Theorem 3 to private-monitoring games that are ε -close, but not strongly so, to a public-monitoring game is that the public profile does not uniquely determine a private profile (see Remark 1). Without further restrictions on the private profile, it is difficult to determine the evolution of beliefs. However, the proof of Theorem 3 (with essentially no changes) shows the following (see footnote 15 for automaton representations of private strategies and the notion of *reflects*):

THEOREM 4. Fix (W, w^1, σ, d) , 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 ε -close under some signaling function f to $(\tilde{u}^*, (Y, \rho))$ with (Ω, π) having richness, given f , for some player i of at least ζ and essentially full support, and if the private profile $(\tilde{W}_i, \tilde{w}^1, \tilde{\sigma}_i, \tilde{d}_i)_i$ reflects the public profile (W, w^1, σ, d) with $\tilde{W}_i = W$ for all i , then the private profile is not a Nash equilibrium of the private-monitoring game.

The key observation is that under the assumption that the private profile not introduce any new states (i.e., $\tilde{W}_i = W$), the continuation play even after uninterpretable signals is still described by the public profile (conditional on the state). The logic described in Remark 2 hence applies to this case as well.

The restriction that the private profile not introduce any new states is substantive. It is easy to construct counterexamples to Part (ii) of Theorem 1 under this additional restriction.

APPENDIX: REMAINING PROOFS

PROOF OF LEMMA 2. 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 ,

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

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

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 2**), 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 follows from part 1 of **Definition 2** 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 . $\square$

PROOF OF THE FIRST CLAIM IN LEMMA 3. 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 for h^t and $\hat{h}^t$ as specified in **Definition 4**,

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

Suppose now the profile s has L -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 exist h^τ and $\hat{h}^\tau$ 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 $(\hat{h}^\tau, h^t)$ 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, $\sigma(w, h^L) = \sigma(w', h^L)$. $\square$

PROOF OF THEOREM 2. 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 $\tilde{s}_i$, for any 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 exist 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 4. 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 (i) and (ii); parts (iii) and (iv) are then 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 are 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 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 $W^1 \equiv \{\sigma(w, h^{T_1^1}) : w \in R(\tilde{w})\}$, 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 $(\sigma(w, T_1^1 h^t))_{w \in W^1} \in W^{W^1}$.

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

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

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

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

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

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

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

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

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

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

concluding the recursive step.

Since W is finite, this process must eventually reach a point where $W^{\kappa+1} = W^\kappa$. We have thus identified a set of states W^κ and two dates T_κ^1 and T_κ^2 , such that letting $(\bar{y}^1, \dots, \bar{y}^m) \equiv T_\kappa^1 h^{T_\kappa^2}$ and setting $\bar{w} = \sigma(\tilde{w}, h^{T_\kappa^1})$ yields parts (i) and (ii) 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_κ^1 and T_κ^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_κ^1 and T_κ^2 on the path starting in period T_κ^1 from the state $\sigma(w, h^{T_\kappa^1})$. Consequently, we have part (iii).

Finally, since $|W^1| \geq 2$, $|W_c| \geq 2$. □

PROOF OF COROLLARY 2. If $d_i(\hat{w}) \neq d_i(\bar{w})$ for some $\hat{w} \in W_c \setminus \{\bar{w}\}$ does not hold for the current choice of cycle and states, by part (iii), it holds in some period of the cycle $h^m = (\bar{y}^1, \dots, \bar{y}^{m-1})$, say period ℓ . Start the cycle in period ℓ , $(\bar{y}^\ell, \dots, \bar{y}^{m-1}, \bar{y}^1, \dots, \bar{y}^{\ell-1})$, and define the new $\bar{w}$ by $\sigma(\bar{w}, \bar{y}_1, \dots, \bar{y}^{\ell-1})$. Finally, the set of cycling states is given by $\{\sigma(w, \bar{y}^1, \dots, \bar{y}^{\ell-1}) : w \in W_c\}$. □

REFERENCES

- Abreu, Dilip, David Pearce, and Ennio Stacchetti (1990), "Toward a theory of discounted repeated games with imperfect monitoring." *Econometrica*, 58, 1041–1063. [311, 316]
- Axelrod, Robert M. (1984), *The Evolution of Cooperation*. Basic Books, New York. [314]
- Bhaskar, V. (1998), "Informational constraints and the overlapping generations model: Folk and anti-folk theorems." *Review of Economic Studies*, 65, 135–149. [314]
- Bhaskar, V. and Ichiro Obara (2002), "Belief-based equilibria in the repeated Prisoners' Dilemma with private monitoring." *Journal of Economic Theory*, 102, 40–69. [314]
- Bhaskar, V. and Eric van Damme (2002), "Moral hazard and private monitoring." *Journal of Economic Theory*, 102, 16–39. [314]
- Cole, Harold L. and Narayana R. Kocherlakota (2005), "Finite memory and imperfect monitoring." *Games and Economic Behavior*, 53, 59–72. [313]
- Compte, Olivier (1998), "Communication in repeated games with imperfect private monitoring." *Econometrica*, 66, 597–626. [315]
- Ellison, Glenn (1994), "Theories of cartel stability and the joint executive committee." *RAND Journal of Economics*, 25, 37–57. [314]
- Ely, Jeffrey C. (2002), "Correlated equilibrium and trigger strategies with private monitoring." Unpublished paper, Northwestern University. [314]
- Ely, Jeffrey C., Johannes Hörner, and Wojciech Olszewski (2005), "Belief-free equilibria in repeated games." *Econometrica*, 73, 377–416. [314]
- Ely, Jeffrey C. and Juuso Välimäki (2002), "A robust folk theorem for the Prisoner's Dilemma." *Journal of Economic Theory*, 102, 84–105. [314]

Fudenberg, Drew and David K. Levine (2004), “The Nash-threats folk theorem with communication and approximate common knowledge in two player games.” *Journal of Economic Theory*, forthcoming. [315]

Fudenberg, Drew, David K. Levine, and Eric Maskin (1994), “The folk theorem with imperfect public information.” *Econometrica*, 62, 997–1039. [316]

Greif, Avner (2005), *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*. Cambridge University Press, Cambridge. [314]

Hörner, Johannes and Wojciech Olszewski (2005), “The folk theorem for games with private almost-perfect monitoring.” *Econometrica*, forthcoming. [314]

Kandori, Michihiro (2002), “Introduction to repeated games with private monitoring.” *Journal of Economic Theory*, 102, 1–15. [314]

Kandori, Michihiro and Hitoshi Matsushima (1998), “Private observation, communication and collusion.” *Econometrica*, 66, 627–652. [315]

Mailath, George J. and Stephen Morris (2002), “Repeated games with almost-public monitoring.” *Journal of Economic Theory*, 102, 189–228. [312, 318, 319, 320, 321, 322, 323, 337]

Mailath, George J. and Larry Samuelson (2006), *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, New York. [314, 315, 318, 323, 327]

Matsushima, Hitoshi (2004), “Repeated games with private monitoring: Two players.” *Econometrica*, 72, 823–852. [314]

McLean, Richard, Ichiro Obara, and Andrew Postlewaite (2002), “Informational smallness and private monitoring.” Unpublished paper, University of Pennsylvania. [315]

Monderer, Dov and Dov Samet (1989), “Approximating common knowledge with common beliefs.” *Games and Economic Behavior*, 1, 170–190. [318, 322]

Piccione, Michele (2002), “The repeated Prisoner’s Dilemma with imperfect private monitoring.” *Journal of Economic Theory*, 102, 70–83. [314]

Sekiguchi, Tadashi (1997), “Efficiency in repeated Prisoner’s Dilemma with private monitoring.” *Journal of Economic Theory*, 76, 345–361. [314, 315]