Private Strategies in Finitely Repeated Games with Imperfect Public Monitoring*

George J. Mailath and Steven A. Matthews
University of Pennsylvania

Tadashi Sekiguchi
Kobe University

May 20, 2002

Abstract

We present three examples of finitely repeated games with public monitoring that have sequential equilibria in *private* strategies, i.e., strategies that depend on own past actions as well as public signals. Such *private sequential equilibria* can have features quite unlike those of the more familiar perfect public equilibria: (*i*) making a public signal less informative can create Pareto superior equilibrium outcomes; (*ii*) the equilibrium final-period action profile need not be a stage game equilibrium; and (*iii*) even if the stage game has a unique correlated (and hence Nash) equilibrium, the first-period action profile need not be a stage game equilibrium.

Journal of Economic Literature classification numbers: C72, C73.

Keywords: Private strategies, repeated games, public perfect equilibria.

*The research reported here was undertaken while Sekiguchi was visiting the University of Pennsylvania. Sekiguchi is grateful to the University of Pennsylvania for its kind hospitality. We thank Dilip Abreu, Michihiro Kandori, Leslie Marx, Ichiro Obara, and two referees for helpful comments. The first two authors thank the National Science Foundation and the third author thanks the Japan Economic Research Foundation for financial support. First version March 5, 2001.

1. Introduction

A repeated game has public monitoring if, after each period, each player observes the same signal of that period's action profile. The most easily studied sequential equilibria of these games are those in which each player's strategy is public, i.e., depends only on the history of public signals. These "perfect public equilibria" have been studied prominently by Abreu, Pearce, and Stacchetti (1990) and Fudenberg, Levine, and Maskin (1994). However, if the monitoring is imperfect, sequential equilibria may exist in which players use "private strategies" that depend on privately known past actions in addition to publicly known past signals. We refer to these equilibria as *private sequential equilibria.*

In this paper we provide three examples to illustrate features that distinguish private sequential equilibria from perfect public equilibria. These distinctions appear even in finitely repeated games. In our examples a stage game is played twice, with the players observing a public signal of the first-period actions.

In each example a private sequential equilibrium exists that Pareto dominates every perfect public equilibrium and, indeed, every subgame perfect equilibrium of the corresponding game with perfect monitoring. The examples thus show that if a repeated game's public signal is made less informative, new sequential equilibrium payoffs can arise that are not in the convex hull of the original game's set of sequential equilibrium payoffs. This is in contrast to the opposite result obtained by Kandori (1992) for perfect public equilibria.[1]

The third example is also of interest because it explores a novel aspect of equilibrium enforcement: after a deviation, nondeviating agents may have inconsistent beliefs over continuation play.

Our first two examples show that the following well-known result also does not extend to private sequential equilibria: If the stage game of a finitely repeated game has a unique Nash equilibrium, then the unique perfect public equilibrium of the finitely repeated game consists of playing the stage-game equilibrium in every period after any history. In contrast, our first two examples have unique stage game equilibria, and nonetheless exhibit private sequential equilibria in which the stage-game equilibrium is not played in each period. In the second example, the stage game's Nash equilibrium is also its only correlated equilibrium.

In the first example, the unique stage-game equilibrium is played in the first period but not in the second. Instead, a non-Nash correlated equilibrium of the

[1] However, an example in Kandori (1991) shows that the convex hull of the set of sequential equilibrium payoffs can expand if *private* signals are made less informative.

stage game is played in the second period. This is possible because the stage game's Nash equilibrium is in mixed strategies. When it is played in the first period, each player's realized first-period action becomes his private information. In the second period, after also observing the realization of the public signal, each player has a strict best response that depends on both his first-period action and the public signal. The random first-period actions and the public signal together constitute an appropriate correlating device for the second-period play. This is like the "internal correlation" of Lehrer (1991), as we discuss in Section 6.[2]

In the second example, unlike in the first, the unique stage-game correlated equilibrium is not played in the first period. Instead, player 2 plays another mixed strategy in the first period. Her second-period strategy is pure, and it determines her action as a function of her realized first-period action and the public signal.[3] These two random variables are independent if player 1 does not deviate in the first period. If player 1 does deviate in the first period, player 2's first-period action and the public signal become correlated. Moreover, since player 2's second-period action depends on her first-period action and the signal, this correlation causes player 2 to play (with high probability) an action that is particularly bad for player 1. This "unwitting punishment" deters player 1 from deviating in the first period.

The third example, like the second, is of a private sequential equilibrium in which first-period play is not a stage-game equilibrium. A deviation in the first period is deterred by subsequent play that is not a stage-game correlated equilibrium. The example appears to be more robust than the second example to perturbations of signals and payoffs. Rather than being inspired by purification, it uses arguments that rely on there being at least three players. The presentation starts with a signal that has the property that while any unilateral deviation is observed, the deviator's identity is not. The resulting game has a pure-strategy private sequential equilibrium that Pareto dominates every perfect public equilibrium. The equilibrium is in pure strategies, and does not entail the play of a stage-game equilibrium in the first period—even though the stage game has a unique pure-strategy equilibrium payoff vector. The relevant deviation (by player 3) from the private sequential equilibrium is deterred because when he deviates, the other players' beliefs differ as to the probable identity of the deviator, and so of continuation play. While this equilibrium is not robust to perturbing the

[2]It is not true that private strategies allow *anything* to happen in the last period: on the equilibrium path of any sequential equilibrium of a finitely repeated game with public monitoring, conditional on the signal a correlated equilibrium is played in the last period (Proposition 2 below).

[3]On the equilibrium path, player 2's pure second-period strategy purifies her mixed stage-game equilibrium strategy.

probabilities so that the signal has non-moving support, there is a similar private sequential equilibrium in mixed strategies that is.[4]

The next section presents preliminary material. The examples are presented in the three subsequent sections. The final section discusses related literature, and the appendix considers the robustness of the second example.

2. Preliminaries

The set of players is $N = \{1, \ldots, n\}$. The stage game is a strategic form game $G = (u, A)$, where $A = \prod_{i=1}^{n} A_i$ is a finite set of action profiles and $u : A \to \Re^n$ is the payoff function. The repeated game has two periods, $t = 1, 2$. In period t the players simultaneously choose their actions a_i^t. The resulting action profiles are a^1 and a^2, and the resulting payoffs are $u(a^1) + u(a^2)$. The first-period profile determines the probability distribution, $\pi(\cdot \mid a^1)$, of a signal y that is publicly observed by the players between the periods. Payoffs are received at the end of the second period.[5] The set of possible signals is a finite set Y. For each $y \in Y$, some $a \in A$ exists such that $\pi(y \mid a) > 0$. Denote this once-repeated game with imperfect public monitoring as $G(\pi, Y)$.

A (behavior) strategy profile for $G(\pi, Y)$ is denoted $\sigma = (\sigma_1, \ldots, \sigma_n)$, where $\sigma_i = (\sigma_i^1, \sigma_i^2)$, $\sigma_i^1 \in \Delta(A_i)$, and $\sigma_i^2 : A_i \times Y \to \Delta(A_i)$. If σ_i^2 is a pure strategy, $\sigma_i^2(a_i^1, y)$ is the action taken; otherwise $\sigma_i^2(a_i^2 \mid a_i^1, y)$ is a probability. The *outcome* of a profile σ is the distribution it induces on $A \times Y \times A$. Its *payoff-relevant outcome* is the corresponding marginal distribution on $A \times A$. We restrict attention to sequential equilibrium strategy profiles (Kreps and Wilson (1982)).

A strategy σ_i is *public* if $\sigma_i^2(a_i, y) = \sigma_i^2(\hat{a}_i, y)$ for all $a_i, \hat{a}_i \in A_i$ and $y \in Y$. A profile σ is a *perfect public equilibrium* (PPE) if it is a sequential equilibrium and each σ_i is public. A sequential equilibrium in which some strategy is not public is a *private sequential equilibrium* (PSE).

It is well known that in repeated games with imperfect public monitoring, any outcome of a pure strategy sequential equilibrium is also the outcome of a PPE.

[4]Some of the logic of this example appears in Marx and Matthews (2000) for a dynamic, non-repeated game with noiseless imperfect public monitoring. The similar repeated-game example in Matthews (1998) is subsumed by this paper.

[5]This ensures that first period payoffs do not provide additional information to players about the opponents' first period play. Equivalently, we can proceed as follows: Player i receives at the end of each period an ex post payoff $v_i(a_i, y)$ that only depends on i's action and the realized public signal. The payoff u_i is the expected value of v_i: $u_i(a) = \sum_y v_i(a_i, y)\pi(y|a)$.

It is worth emphasizing that this result requires the support of the signal to not vary with the action profile (the first example in Section 5 illustrates the necessity of the full support assumption). The precise statement of this result in our setting is the following proposition.

Proposition 1. *If $\pi(y \mid a) > 0$ for all $(y, a) \in Y \times A$, the outcome of any sequential equilibrium of $G(\pi, Y)$ with pure first-period strategies is a PPE outcome.*

Proof. Let σ be a sequential equilibrium in which, say, $a \in A$ is surely played in the first period. Since any signal is realized with positive probability after any action profile, even if player i plays $\hat{a}_i \neq a_i$ in the first period, he must still believe the others will play the profile $(\sigma_j^2(a_j, y))_{j \neq i}$ in the second period, given any realization $y \in Y$. His best action after he plays $\hat{a}_i$ in the first period is hence a best reply to $(\sigma_j^2(a_j, y))_{j \neq i}$. One best reply to this profile is his equilibrium-path strategy, $\sigma_i^2(a_i, y)$. Thus, replacing player i's second-period action $\sigma_i^2(\hat{a}_i, y)$ by $\sigma_i^2(a_i, y)$ for every $\hat{a}_i \in A_i$ yields a *public* strategy $\hat{\sigma}_i$ from which player i has no incentive to deviate. This shows that $(\hat{\sigma}_i, \sigma_{-i})$ is a sequential equilibrium that has the same outcome as does σ. Continuing in this fashion for all the players yields a PPE $\hat{\sigma}$ with the same outcome as σ. ■

Perfect public equilibria are relatively tractable because they have a recursive formulation (Abreu, Pearce, and Stacchetti (1990)). For a once-repeated stage game, this recursive formulation takes the following form: given any PPE σ and signal $y \in Y$, the distribution $\sigma^2(y)$ of the second-period action profile conditional on y is a Nash equilibrium of the stage game. Thus, even after a one-shot deviation in the first period, second-period play conditional on the signal is a stage-game equilibrium.

There is no analogous simple recursive formulation for private sequential equilibria. Conditional on a signal realization, a PSE need not yield a stage-game equilibrium in the second period, even on the equilibrium path. In fact, if a player makes a one-shot deviation in the first period, second-period play, conditional on the signal realization, need not even be a correlated equilibrium. Our second and third examples rely on this property. The only recursiveness necessarily exhibited by a PSE is that exhibited by any Nash equilibrium. This recursiveness, informally observed by many authors, is given in the following proposition.

Proposition 2. *Let σ be a Nash equilibrium of $G(\pi, Y)$. If $y \in Y$ is realized with positive probability in this equilibrium, then the equilibrium distribution of second-period actions, conditional on y, is a correlated equilibrium of G.*

Proof. [6] Let $p \in \Delta(A \times Y \times A)$ be the probability distribution (outcome) induced by σ. Fix $y \in Y$ such that $p(y) > 0$. The conditional distribution over second-period action profiles, $p(a^2 \,|\, y)$, is a correlated equilibrium of G if for all $i \in N$ and $\hat{a}_i^2 \in A_i$,

$$\hat{a}_i^2 \in \operatorname*{argmax}_{a_i \in A_i} \sum_{a_{-i}^2 \in A_{-i}} p(\hat{a}_i^2, a_{-i}^2 \,|\, y) u_i(a_i, a_{-i}^2). \tag{2.1}$$

The maximand in (2.1) can be written as

$$\sum_{a_{-i}^2 \in A_{-i}} \sum_{a_i^1 \in A_i} p(a_i^1, \hat{a}_i^2, a_{-i}^2 \,|\, y) u_i(a_i, a_{-i}^2) = \sum_{a_i^1 \in A_i} \sum_{a_{-i}^2 \in A_{-i}} p(a_i^1, \hat{a}_i^2, a_{-i}^2 \,|\, y) u_i(a_i, a_{-i}^2).$$

Hence (2.1) holds if, for all $a_i^1 \in A_i$,

$$\hat{a}_i^2 \in \operatorname*{argmax}_{a_i \in A_i} \sum_{a_{-i}^2 \in A_{-i}} p(a_i^1, \hat{a}_i^2, a_{-i}^2 \,|\, y) u_i(a_i, a_{-i}^2). \tag{2.2}$$

Fix $a_i^1 \in A_i$. If $\sigma_i^1(a_i^1) = 0$ or $\sigma_i^2(\hat{a}_i^2 \,|\, a_i^1, y) = 0$, then $p(a_i^1, \hat{a}_i^2, a_{-i}^2 \,|\, y) = 0$ for all $a_{-i}^2 \in A_{-i}$. In this case (2.2) holds trivially. So we can assume $\sigma_i^1(a_i^1) > 0$ and $\sigma_i^2(\hat{a}_i^2 \,|\, a_i^1, y) > 0$, i.e., that $(a_i^1, \hat{a}_i^2, y)$ is on the equilibrium path. Since σ is a Nash equilibrium,

$$\hat{a}_i^2 \in \operatorname*{argmax}_{a_i \in A_i} \sum_{a_{-i}^2 \in A_{-i}} p(a_{-i}^2 \,|\, a_i^1, y) u_i(a_i, a_{-i}^2). \tag{2.3}$$

Since

$$\begin{aligned} p(\hat{a}_i^2, a_{-i}^2 \,|\, a_i^1, y) &= p(a_{-i}^2 \,|\, a_i^1, y) p(\hat{a}_i^2 \,|\, a_i^1, y, a_{-i}^2) \\ &= p(a_{-i}^2 \,|\, a_i^1, y) \sigma_i^2(\hat{a}_i^2 \,|\, a_i^1, y), \end{aligned}$$

and $\sigma_i^2(\hat{a}_i \,|\, a_i^1, y) > 0$, (2.3) implies

$$\hat{a}_i^2 \in \operatorname*{argmax}_{a_i \in A_i} \sum_{a_{-i}^2 \in A_{-i}} p(\hat{a}_i^2, a_{-i}^2 \,|\, a_i^1, y) u_i(a_i, a_{-i}^2). \tag{2.4}$$

Multiplying the maximand in (2.4) by the positive term $p(a_i^1 \,|\, y)$ yields that of (2.2), and we are done. ■

[6] The essential idea is that if an action is a best reply after two different histories, and so potentially to two different beliefs over the opponents' play, then that action is a best reply to any average of those beliefs. This logic is much the same as the "obedience" part of the general revelation principle of Myerson (1982).

It is worth emphasizing that if we were only interested in Nash (rather than sequential) equilibria, the distinction between public and private strategies is not significant. Recall that every pure strategy is realization equivalent to a pure public strategy. This implies, of course, that every mixed strategy is realization equivalent to a public mixed strategy (a public mixed strategy has only public strategies in its support). It is not true, however, that every public mixed strategy is realization equivalent to a public behavior strategy. (Examples are easy to construct.) Our interest in sequentiality forces us to restrict attention to public behavior strategies and, as we will see, this is a real restriction.[7]

3. Internal Correlation

The stage game of this section's example is shown below.

	s_1	s_2	s_3
r_1	$0,0$	$1,2$	$2,1$
r_2	$2,1$	$0,0$	$1,2$
r_3	$1,2$	$2,1$	$0,0$

(3.1)

This game has a unique Nash equilibrium, in which each player plays each action with probability 1/3. The resulting payoff vector is $(1,1)$. Non-Nash correlated equilibria also exist, one of which is the following distribution on A :

	s_1	s_2	s_3
r_1	0	$\frac{1}{6}$	$\frac{1}{6}$
r_2	$\frac{1}{6}$	0	$\frac{1}{6}$
r_3	$\frac{1}{6}$	$\frac{1}{6}$	0

. (3.2)

This correlated equilibrium gives rise to the payoff vector $(3/2, 3/2)$.

The public signal (π, Y) is defined in terms of a partition, $\{A^1, A^2, A^3\}$, of A, where

$$A^1 = \{(r_1, s_1), (r_2, s_2), (r_3, s_3)\},$$

$$A^2 = \{(r_1, s_2), (r_2, s_3), (r_3, s_1)\},$$

[7]We thank a referee for encouraging us to discuss this issue here. Related issues arise in repeated games with private monitoring. An extended discussion can be found in Mailath and Morris (2002b) and Mailath and Morris (2002a).

$$A^3 = \{(r_1, s_3), (r_2, s_1), (r_3, s_2)\}.$$

The set of possible signals is $Y = \{y_1, y_2, y_3\}$, and the signal distribution is given by

$$\pi(y_k \mid a) = \begin{cases} 1/2, & \text{if } a \notin A^k, \\ 0, & \text{if } a \in A^k. \end{cases} \tag{3.3}$$

Thus, conditional on the signal and his own action, a player can rule out one action of the opponent. For example, if player 1 observes y_2 after choosing r_i, he knows player 2 did not choose s_{i+1}.[8]

We now describe a profile $\sigma = (\sigma^1, \sigma^2)$ of private strategies in which the stage-game Nash equilibrium is played in the first period, and the correlated equilibrium shown in (3.2) is played in the second period. The profile is given by $\sigma_1^1 = \sigma_2^1 = (1/3, 1/3, 1/3)$,

$$\sigma_1^2(r_i, y^1) = \begin{cases} r_i, & \text{if } y^1 = y_1, \\ r_{i-1}, & \text{if } y^1 = y_2, \\ r_{i+1}, & \text{if } y^1 = y_3, \end{cases} \tag{3.4}$$

and

$$\sigma_2^2(s_j, y^1) = \begin{cases} s_j, & \text{if } y^1 = y_1, \\ s_{j+1}, & \text{if } y^1 = y_2, \text{ and} \\ s_{j-1}, & \text{if } y^1 = y_3. \end{cases} \tag{3.5}$$

We now argue that this profile is a sequential equilibrium. Consider first the sequential rationality of play in the second period. If $y^1 = y_1$, then player 1, having played r_i in the first period, has the *ex post* belief that player 2 played s_{i+1} with probability 1/2 and s_{i-1} with probability 1/2. From (3.5), player 1 thus believes that 2 will play s_{i+1} with probability 1/2 and s_{i-1} with probability 1/2 in the second period. Given this belief, it is sequentially rational to play r_i again in the second period, as specified by σ_1^2. A similar argument applies to player 2 after the signal $y^1 = y_1$.

If $y^1 = y_2$, player 1, having played r_i in the first period, believes that 2 played s_i with probability 1/2 and s_{i-1} with probability 1/2. Thus by (3.5), he assigns equal probability to player 2, in the second period, playing s_{i+1} and s_i, and so playing r_{i-1} is optimal. On the other hand, player 2 who played s_j believes that 1 played r_j with probability 1/2 and r_{j+1} with probability 1/2. By (3.4), player 2 therefore believes that 1 plays each of r_{j-1} and r_j with probability 1/2. It is

[8] We use the convention that the first action follows the third.

thus sequentially rational to play s_{j+1} given this belief. The case where $y^1 = y_3$ is established through a similar argument.

Finally, no player has an incentive to deviate in the first period. Each player's continuation payoff under the profile in the second period is independent of history, being equal to 3/2. Since the choice of action in the first period has no payoff implications in the second period, it is a best response if it maximizes first-period payoffs, which it does.

This example is robust in the sense that a similar non-trivial equilibrium is obtained even if we perturb the initial monitoring structure π. To see this, note that play in the second period given σ and any history remains sequentially rational under any monitoring structure sufficiently close to π, because each $\sigma_i^2(a_i^1, y^1)$ is a strict best response given the beliefs. This is so even if each player's first-period strategy is not equal to $(1/3, 1/3, 1/3)$, but is sufficiently close to it.

Suppose now that the players are restricted to play σ^2 in the second period. This yields a one-shot game in which only the first-period actions are chosen. Its unique equilibrium is $\sigma^1 = ((1/3, 1/3, 1/3), (1/3, 1/3, 1/3))$. Therefore, under any monitoring structure sufficiently close to π, there exists an equilibrium σ^* of this game that is close to σ^1. By the above argument, the strategy profile in which the players play σ^* in the first period and then σ^2 in the second period is an equilibrium if the monitoring structure is sufficiently close to π.

4. Unwitting Punishment

The stage game of this section's example is shown below.[9]

	s_1	s_2	s_3	s_4
r_1	6, 0	0, 1	0, 0	0, 0
r_2	5, 6	1, 5	11, 0	11, 1
r_3	0, 0	0, 0	10, 10	10, 10

In this game r_3 is strictly dominated, and its removal causes s_3 and s_4 to become strictly dominated. The unique correlated (and hence Nash) equilibrium is for each player to play each of his first two actions with equal probability. The

[9]This game is based on the example in Kandori (1991). Both games have the property that in the unique correlated equilibrium, each player has a preference over the strategy choice of the other player independent of his own behavior. We discuss the relationship to Kandori (1991) at the end of this Section.

equilibrium payoff vector is $(3,3)$. Profiles (r_3, s_3) and (r_3, s_4) are desirable, but are not equilibria because player 1 has an incentive to deviate to r_2. We now present a signal structure such that the two-period game has a PSE in which, in the first period, player 1 plays r_3 and player 2 plays s_3 and s_4 with equal probability.

The public signal structure is given by $Y = \{y', y''\}$, and the following table of conditional probabilities $\pi(y' \mid r_i, s_j)$:

	s_1	s_2	s_3	s_4
r_1	.5	.5	.5	.5
r_2	.5	.5	1	0
r_3	.5	.5	.5	.5

(4.1)

Thus, if player 1 plays r_2 and player 2 randomizes between s_3 and s_4, the public signal is perfectly correlated with player 2's action—player 1 will then surely learn whether s_3 or s_4 was played. But if player 1 plays r_3 (or r_1), the signal and player 2's action are independent, and player 1 will learn nothing from the signal about player 2's action.

The candidate strategy for player 1, σ_1, requires that he play r_3 in the first period:

$$\sigma_1^1(r_3) = 1.$$

In the second period he plays r_1 and r_2 with equal probability, *provided* he played r_3 or r_1 in the first period. If he deviated to r_2 in the first period, he plays r_2 in the second:

$$\sigma_1^2(r_1 \mid r_i^1, y) = \begin{cases} 0, & \text{if } r_i^1 = r_2 \\ 1/2, & \text{otherwise,} \end{cases}$$

and $\sigma_1^2(r_2 \mid r_i^1, y) = 1 - \sigma_1^2(r_1 \mid r_i^1, y)$. The candidate strategy for player 2 requires her to play s_3 and s_4 with equal probability in the first period:

$$\sigma_2^1(s_j) = \begin{cases} 1/2, & \text{if } s_j = s_3, s_4 \\ 0, & \text{if } s_j = s_1, s_2. \end{cases}$$

In the second period she plays s_2 if her private history is (y', s_3) or (y'', s_4), and otherwise she plays s_1:

$$\sigma_2^2(s_1 \mid s_j^1, y) = \begin{cases} 0, & \text{if } y = y' \text{ and } s_j^1 = s_3, \\ & \text{or } y = y'' \text{ and } s_j^1 = s_4, \\ 1, & \text{otherwise,} \end{cases}$$

and $\sigma_2^2(s_2 \mid s_j^1, y) = 1 - \sigma_2^2(s_1 \mid s_j^1, y)$.

On the equilibrium path of σ, conditional on either signal realization, the stage-game equilibrium is played in the second period (as required by Proposition 2). In particular, player 1 is content to play r_1 and r_2 with equal probability because the signal is uninformative about player 2's past action, and hence about her second-period action, even though she is actually using a pure strategy in the second period. Her fifty-fifty mixture of s_1 and s_2 is purified by her random first-period action and the signal.

However, if player 1 deviates in the first period to his myopic best reply r_2, the continuation play conditional on either signal realization is not even a correlated equilibrium of the stage game: (r_2, s_2) is played after either y' or y''. If, for example, y' is realized after player 1 deviates to r_2, he will know that player 2's first-period action was s_3, and hence that she will surely play s_2, and so r_2 is his only best reply. Player 2 is nonetheless still content to play s_2 after (s_3, y'); observing y' does not reveal to her that player 1 deviated, and so she still believes he is playing the fifty-fifty mixture of r_1 and r_2.

It is now easy to see why player 1 does not deviate to r_2 in the first period, and hence that σ is a PSE. This deviation results in the play of (r_2, s_2) instead of the stage-game equilibrium in the second period, and so it costs him $3 - 1 = 2$ utiles. This is more than his myopic gain from the deviation, $11 - 1 = 1$. He thus has a strict incentive not to deviate from σ in the first period.

The example is robust in that any game obtained by slightly perturbing the signal probabilities has a PSE close to σ. We focussed on player 1 in this example only for convenience; a two-player example in which both players have an incentive to myopically deviate can be obtained from the authors. Moreover, that example is robust in the stronger sense that any game obtained by slightly perturbing either the payoffs or the signal probabilities has a PSE with desired properties.[10] However, this is not true if *both* structures are perturbed. In the Appendix we show that almost all small perturbations of the payoffs and probabilities together yield a game that has a unique sequential equilibrium: the perfect public equilibrium in which the stage-game equilibrium is played after any history. We conjecture that if a *two*-player finitely-repeated game with public monitoring has a stage game with generic payoff and signal probability structures, and if the stage game has a unique correlated equilibrium, then that correlated equilibrium is played after any history in the unique sequential equilibrium of the repeated game. We have no conjecture if the number of players is greater than two.

[10]The example described here is not robust to payoff perturbations because player 2 must be indifferent between $s_3$and s_4 when player 1 chooses r_3.

As mentioned in footnote 9, this example and the example of Kandori (1991) share the feature that in the unique correlated equilibrium, each player has a preference over the strategy choice of the other player independent of his own behavior. The idea in both is to provide incentives through disadvantageous changes in opponent's behavior that result from deviations. However, because signals are public in our case, while private in Kandori (1991), the mechanisms through which these disadvantageous changes arise are different.

5. Punishment by Disparate Beliefs

There are three players in this section's example. Each player has two actions, c and d. The stage game is shown below, with player 3 choosing the left or right matrix by his choice of c or d, respectively.

P1 \ P2	c	d
c	$15, 15, 17$	$-85, 5, 3$
d	$5, -85, 3$	$5, 5, -9$

c

P1 \ P2	c	d
c	$-1, -1, 26$	$1, 1, 12$
d	$1, 1, 12$	$0, 0, 0$

d

Action d is strictly dominant for player 3; his payoff increases by 9 if he plays d instead of c. The pure strategy equilibria are (c, d, d) and (d, c, d), and both generate the payoffs $(1, 1, 12)$. In the remaining equilibrium, players 1 and 2 each play c with probability $\frac{1}{3}$; the resulting payoffs are $(\frac{1}{3}, \frac{1}{3}, 8\frac{2}{9})$.

In any once-repeated game $G(\pi, Y)$ based on this stage game, when player 3 plays d in the second period, each player's payoff is maximized by the outcome

$$\overrightarrow{a} \equiv ((c, c, c), (c, d, d)).$$

The outcome $\overrightarrow{a}$ is not a PPE outcome. Recall that PPE requires, upon the realization of *any* signal, even off-the-equilibrium path, that a stage-game equilibrium be played in the second period. Player 3 can thus be punished by at most the maximal difference in his stage-game equilibrium payoffs, $12 - 8\frac{2}{9} = 3\frac{7}{9}$, which is less than his gain of 9 obtained by playing d rather than c in the first period. It follows that $((c, d, d), (c, d, d))$ is a Pareto dominant PPE outcome.[11] Since this outcome is Pareto dominated by $\overrightarrow{a}$, every PPE outcome of any of the games $G(\pi, Y)$ is Pareto dominated by $\overrightarrow{a}$.

[11] A payoff-equivalent PPE outcome is $((d, c, d), (d, c, d))$.

On the other hand, player 3 would not deviate to d in the first period if doing so caused the second-period profile to switch from (c, d, d) to (d, d, d).[12] While this is impossible in a PPE, it is possible in a PSE given an appropriate signaling structure, as we now show.

Consider the noiseless signal consisting of the number of players who play c in the first period. Letting $C(a) \equiv \#\{i \mid a_i = c\}$, the possible signal realizations are $\bar{Y} = \{0, 1, 2, 3\}$ and the probability function is

$$\bar{\pi}(y \,|\, a^1) \equiv \begin{cases} 1, & \text{if } y = C(a^1), \\ 0, & \text{if } y \neq C(a^1). \end{cases}$$

The game $G(\bar{\pi}, \bar{Y})$ has a PSE with outcome $\overrightarrow{a}$. The equilibrium profile, $\bar{s}$, is in pure strategies and defined as follows. In the first period, (c, c, c) is played: $\bar{s}_i^1 = c$ for $i = 1, 2, 3$. In the second period, player 3 plays d regardless of history: $\bar{s}_3^2(\cdot, \cdot) = d$. Players 1 and 2 play (c, d) in the second period if the number of players who played c in the first period was not two:

$$\bar{s}_1^2(\cdot, y) = c \text{ and } \bar{s}_2^2(\cdot, y) = d \text{ for } y \neq 2.$$

Otherwise, each of players 1 and 2 takes the action he did not take previously:

$$\bar{s}_i^2(c, 2) = d \text{ and } \bar{s}_i^2(d, 2) = c \text{ for } i = 1, 2. \tag{5.1}$$

This completes the definition of $\bar{s}$. Observe that its outcome is $\overrightarrow{a}$, and that (d, d, d) is played in the second period if player 3 deviates to d in the first. It follows that $\bar{s}$ is a Nash equilibrium.

To show that $\bar{s}$ is a PSE profile, suppose each of players 1 and 2 believes, when he takes action c in the first period and observes $y = 2$, that player 3 was the player who chose the other c. A deviation by player 3 then causes players 1 and 2 to each believe the other was the deviator (when in fact it was player 3). Given (5.1), each of players 1 and 2 then believes the other will play c, and so his best reply is d. These beliefs thus make (d, d, d) sequentially rational in the second period if player 3 unilaterally deviates in the first.

To complete the argument that $\bar{s}$ is a sequential equilibrium profile, we now show that the assessment obtained by pairing $\bar{s}$ with the specified beliefs is consistent. For $\varepsilon > 0$, define the following profile of mixed first-period strategies:

$$\sigma^{1,\varepsilon} \equiv ((1 - \varepsilon) \circ c + \varepsilon \circ d,\ (1 - \varepsilon) \circ c + \varepsilon \circ d,\ \left(1 - \varepsilon^2\right) \circ c + \varepsilon^2 \circ d).$$

[12] Indeed, if the monitoring is perfect, the path $\overrightarrow{a}$, together with (d, d, d) after any deviation, constitutes a Nash equilibrium.

As $\varepsilon \to 0$, $\sigma^{1,\varepsilon} \to (c, c, c)$. Let $\sigma^{2,\varepsilon}$ be any completely mixed second-period strategy profile for which $\sigma^{2,\varepsilon} \to \bar{s}^2$. Then $\sigma^\varepsilon \equiv (\sigma^{1,\varepsilon}, \sigma^{2,\varepsilon}) \to \bar{s}$. To prove consistency, we need only verify that the limiting belief, as $\varepsilon \to 0$, of both players 1 and 2 after a history $(c, 2)$ is that player 3 chose the other c. When σ^ε is played, each of players 1 and 2 has the following belief after history $(c, 2)$:

$$\Pr\{3 \text{ chose } c \,|\, c, 2\} = \frac{(1-\varepsilon^2)\,\varepsilon}{(1-\varepsilon^2)\,\varepsilon + \varepsilon^2\,(1-\varepsilon)} = \frac{1-\varepsilon^2}{1-\varepsilon^2+\varepsilon\,(1-\varepsilon)}.$$

As this converges to 1, $\bar{s}$ paired with the specified beliefs is indeed consistent.[13,14]

Remark. The punishment of player 3 in this equilibrium is sequentially rational only because, if he deviates, his identity as the deviator is not detectable. Neither other player knows it was player 3 who deviated, and so neither knows (d, d, d) will be played – each thinks the other will play c. This confusion is a *lack* of coordination that allows (d, d, d) to be sequentially rational. In contrast, most, if not all, of the previous literature is motivated by the observation that imperfect detection of deviators can make it difficult to coordinate on punishing profiles. Assumptions are thus made to insure that monitoring is not too imperfect, so that a deviation and the deviator's identity are statistically detectable, and punishment can be coordinated.[15] The logic of our example shows instead that imperfect monitoring can sometimes increase punishment levels *because* it impedes coordination.

The profile $\bar{s}$ remains a sequential equilibrium if the payoffs are perturbed, but not if the signaling structure is perturbed. It is easy to show that if π is any "full support" perturbation of $\bar{\pi}$, in the sense that $\pi(y \,|\, \cdot) > 0$ for all $y \in \bar{Y}$, then no sequential equilibrium outcome of $G(\bar{\pi}, \bar{Y})$ is close to $\overrightarrow{a}$. (This is in accordance with Proposition 1, since $\bar{s}$ has pure first-period strategies, but a non-PPE outcome.)

However, another PSE is robust to perturbations of the signal distribution. Denoting it as $\bar{\sigma}$, it is the same as $\bar{s}$ in the second period, and in the first period

[13]Less extreme beliefs can also be paired with $\bar{s}$ to obtain a consistent assessment. It is sufficient for each of players 1 and 2, after history $(c, 2)$, to believe player 3 played c with probability $p \geq 1/3$. To obtain the appropriate trembles, replace $\sigma_3^{1,\varepsilon}$ with $(1 - \varepsilon(1-p)/p) \circ c + \varepsilon(1-p)/p \circ d$.

[14]In fact, $\bar{s}$ is even an (extensive-form trembling-hand) perfect equilibrium, since $\bar{s}$ is a best reply to σ^ε for small ε.

[15]For example, Fudenberg, Levine, and Maskin (1994) assume "individual full rank" and "pairwise identifiability", and Mailath and Morris (2002b) assume "almost-public monitoring."

it is given by

$$\bar{\sigma}_1^1(c) = \bar{\sigma}_2^1(c) = \frac{9}{10}, \text{ and } \bar{\sigma}_3^1(c) = 1.$$

The nature of $\bar{\sigma}$ is similar to that of $\bar{s}$. It too has the property that a first-period deviation by player 3 causes (d, d, d) to be subsequently played with high probability, now .81 instead of 1. Also like $\bar{s}$, $\bar{\sigma}$ Pareto dominates any PPE of any of the games $G(\pi, Y)$. The fact that $\bar{\sigma}$ is a Nash equilibrium is obvious. It is thus a sequential equilibrium, since now the only problematic signal realization, $y = 2$, occurs with positive probability on the equilibrium path. Moreover, $\bar{\sigma}$ is robust to any small perturbation of $\bar{\pi}$, and it is a perfect equilibrium of $G(\bar{\pi}, \bar{Y})$:

Proposition 3. *For any sequence of distributions* $\{\pi^k\}$, $\pi^k : \bar{Y} \times A \to [0, 1]$, *converging to* $\bar{\pi}$, *there is a sequence* $\{\sigma^k\}$ *of strategy profiles converging to* $\bar{\sigma}$ *such that each* σ^k *is a perfect equilibrium of* $G(\pi^k, \bar{Y})$.

Proof. For large k, the first-period part of player 3's strategy, $\bar{\sigma}_3^1$, is a strict best reply to $\bar{\sigma}$ in $G(\pi^k, \bar{Y})$. Since $\bar{\sigma}_3^2$ always puts all probability on d and is hence a second-period best reply to any profile, $\bar{\sigma}_3$ is thus a best reply in $G(\pi^k, \bar{Y})$ to any profile in a neighborhood of $\bar{\sigma}$. We can thus fix player 3's strategy at $\bar{\sigma}_3$ in $G(\pi^k, \bar{Y})$ and consider the resulting game between players 1 and 2, denoted G^k.

Let G^{k1} be the one-shot game obtained from G^k by fixing the second-period strategies at $\bar{\sigma}_{-3}^2 = (\bar{\sigma}_1^2, \bar{\sigma}_2^2)$. For large k, G^{k1} is the same as the original stage game, with player 3's action fixed at c, except that a number near 1 is added to each of the payoffs of players 1 and 2. The profile $\bar{\sigma}_{-3}^1 = (\bar{\sigma}_1^1, \bar{\sigma}_2^1)$ is an equilibrium of this one-shot game if precisely 1 is added to each of their payoffs. Since a completely mixed equilibrium of a 2×2 game that has no dominant strategies is continuous in payoffs, G^{k1} has an equilibrium σ_{-3}^{k1} close to $\bar{\sigma}_{-3}^1$.

Let $\sigma_{-3}^k = (\sigma_{-3}^{k1}, \bar{\sigma}_{-3}^2)$. By the argument just given, each first-period strategy in this profile is a best reply to σ_{-3}^k in the two-period game G^k. For large k, since $(\pi^k, \sigma_{-3}^{k1})$ is close to $(\bar{\pi}, \bar{\sigma}_{-3}^1)$, when σ_{-3}^{k1} is played in G^k, Bayes' rule implies that for $i \neq j = 1, 2$, $\Pr(a_j^1 = d \mid a_i^1 = c, y = 2)$ is close to 1, and so $\bar{\sigma}_i^2(\cdot \mid c, 2)$ is a strict best reply to $\bar{\sigma}_i^2$ in the second period. Similarly, for large k the strategy $\bar{\sigma}_i^2(\cdot \mid d, 2)$ is a strict best reply to $\bar{\sigma}_i^2$. Following any $y \neq 2$, $\bar{\sigma}_1^2$ and $\bar{\sigma}_2^2$ are strict best replies to each other regardless of the first-period actions. Thus, the strategies in $\bar{\sigma}_{-3}^2$ are strict best replies to each other in the second period if k is large enough. It follows, since σ_{-3}^{k1} is completely mixed, that σ_{-3}^k is a perfect equilibrium of G^k. Finally, it is clear that $\sigma^k = (\sigma_{-3}^k, \bar{\sigma}_3)$ converges to $\bar{\sigma}$. ■

6. Related Literature

In a series of papers, Lehrer (1990, 1991, 1992) introduces the idea of internal correlation, in the setting of undiscounted infinitely-repeated games with deterministic imperfect public monitoring. These papers consider equilibria in which players use private histories of past play to correlate future choices, in effect constructing "internal" correlating devices. A similar idea underlies our first example, where the first-period actions, together with the signal, correlate second-period play. Since the first-period equilibrium payoffs of the players are independent of their first-period actions in the example, the first-period actions have some of the features of cheap talk announcements. As such, the example is also related to the literature on mediated cheap talk. For instance, Lehrer and Sorin (1997) show that almost any correlated equilibrium outcome is a Nash equilibrium of the extended game obtained by adding a mediator to whom the players send private messages, and who then replies with a public deterministic message, before the original game is played. The mediator in Lehrer and Sorin (1997) has a similar role as the public signal in our example, although its first-period actions are not literally cheap talk.

Like our examples, the examples of Kandori and Obara (2000) show that private sequential outcomes need not be in the convex hull of the set of PPE outcomes. Their examples, unlike our's, are equilibria of infinitely repeated games with discounting. While some of the equilibrium constructions in that paper are complicated and require an infinite horizon, their insight that private strategies can enhance the efficiency of monitoring also holds in finite horizon settings. Our examples illustrate other reasons for significant differences between PSE and PPE outcomes even with a finite horizon.

Tomala (1999) studies undiscounted infinitely repeated games with public monitoring in which observed deviations may be compatible with several potential deviators. As he points out, the inability to identify deviators can shrink the set of equilibrium payoffs. Our third example, on the other hand, shows the opposite can also occur, i.e., in some games, an inability to identity a deviator can expand the set of equilibrium payoffs.

Finally, the literature has noted another reason why PSE payoffs may differ from PPE payoffs. If there are more than two players, the correlated minmax payoff for a player (i.e., the minmax payoff when opponents can correlate their actions) may be less than the player's standard minmax payoff. Consequently, since correlation may be obtained via private strategies, some repeated games with imperfect public monitoring have PSE payoffs outside the set of feasible and individually rational payoffs, as usually defined. See Fudenberg and Tirole (1991, exercise 5.10) for a simple example and Tomala (1999) for a more detailed discussion.

A. Robustness of the Second Example

We show here that for almost all small perturbations of the payoffs and probabilities in the example of Section 4, the resulting game has a unique sequential equilibrium: the PPE in which the stage-game equilibrium is played after any history.

We first argue that since the stage game has a unique correlated equilibrium, each player must have identical beliefs over the future behavior of the other player after any private history.

Proposition A. *Let σ be a Nash equilibrium of a two-player game $G(\pi, Y)$, and suppose the stage game G has a unique correlated equilibrium, σ^*. Then, when σ is played, the equilibrium belief of player i about player j's second-period action, after any history (a_i^1, y) that has positive probability under σ, is given by σ_j^*.*

Proof. Let $y \in Y$ have positive probability under σ. When σ is played, the realization of y gives rise to a one-shot incomplete information game in the second period. The types of player i in this game are the actions a_i^1 that have positive probability under σ^1, conditional on y having been realized; denote this subset of A_i as T_i. The prior distribution, ρ, on the type space, $T \equiv \prod_i T_i$, is the conditional (on y) distribution of first-period action profiles under σ^1. This construction yields an incomplete information game, (G, T, ρ), in which the types are payoff-irrelevant. Since σ is a Nash equilibrium, and it puts positive probability on y, it induces a Bayes-Nash equilibrium on (G, T, ρ). Specifically, defining $\alpha_i(a_i \,|\, t_i) \equiv \sigma_i^2(a_i \,|\, t_i, y)$ for $i = 1, 2$ and $(a_i, t_i) \in A_i \times T_i$, α is a Bayes-Nash equilibrium of (G, T, ρ). The beliefs of player i, conditional on his type t_i, about the action player j will take in this equilibrium are given by $\sum_{t_j} \alpha_j \left(a_j \,|\, t_j\right) \rho_i \left(t_j \,|\, t_i\right)$. These are the same beliefs that player i holds when σ is played in $G(\pi, Y)$ about player j's second-period action, after the history $(a_i^1, y) = (t_i, y)$. The following Lemma now completes the proof. ■

Lemma A. *Let $G = (u, A)$ be a two-player game that has a unique correlated (and hence Nash) equilibrium, σ^*. Let (G, T, ρ) be a corresponding incomplete* *information game with payoff-irrelevant types. Then, for all* $(a_j, t_i) \in A_j \times T_i$ *and any Bayes-Nash equilibrium* α *of this game,*

$$\sum_{t_j} \alpha_j (a_j \mid t_j) \rho_i (t_j \mid t_i) = \sigma_j^* (a_j) \tag{A.1}$$

if $\rho_i (t_i) \equiv \sum_{t_j} \rho (t_i, t_j) > 0$.

Proof. Since every Bayesian-Nash equilibrium of (G, T, ρ) induces a correlated equilibrium of G, we have

$$\sum_{t_i} \alpha_i (a_i \mid t_i) \rho_i (t_i) \sum_{t_j} \alpha_j (a_j \mid t_j) \rho_i (t_j \mid t_i) = \sigma^* (a) = \sigma_i^* (a_i) \sigma_j^* (a_j).$$

For a in the support of σ_i^*, dividing by $\sigma_i^* (a_i)$ yields

$$\sum_{t_i} \frac{\alpha_i (a_i \mid t_i) \rho_i (t_i)}{\sigma_i^* (a_i)} \sum_{t_j} \alpha_j (a_j \mid t_j) \rho_i (t_j \mid t_i) = \sigma_j^* (a_j). \tag{A.2}$$

In other words, conditional on player i taking any action a_i in the support of σ_i^*, his average (over his type t_i) belief about the action of the opponent is $\sigma_j^* (a_j)$.

Let $A_i^* = \{a_i : \sigma_i^* (a_i) > 0\}$ be the support of σ_i^*. Fix $a_i \in A_i^*$. For $a_i' \in A_i \backslash \{a_i\}$, let $\beta_{a_i'} \in \Re^{|A_j^*|}$ be the vector given by $\left[\beta_{a_i'}\right]_{a_j} = u_i (a_i, a_j) - u_i (a_i', a_j)$. Define the set

$$C_{a_i'} \equiv \left\{\sigma_j \in \Delta(A_j^*) : \beta_{a_i'} \cdot \sigma_j \geq 0\right\}.$$

Then $\mathcal{C} \equiv \cap_{a_i' \in A_i \backslash \{a_i\}} C_{a_i'}$ is the set of beliefs over player j's action, that have support in A_j^*, for which a_i is a best reply. Let $T_i (a_i) \equiv \{t_i \in T_i : \alpha_i (a_i \mid t_i) > 0\}$ be the set of types that choose a_i with positive probability. Let $\sigma_j^{t_i} \in \Delta(A_j^*)$ denote the beliefs of type t_i. From (A.2),

$$\sigma_j^* (a_j) = \sum_{t_i \in T_i (a_i)} \frac{\alpha_i (a_i \mid t_i) \rho_i (t_i)}{\sigma_i^* (a_i)} \sigma_j^{t_i} (a_j), \tag{A.3}$$

so that σ_j^* is a convex combination of the beliefs $\sigma_j^{t_i}$ for $t_i \in T_i (a_i)$. Moreover, $\sigma_j^{t_i} \in \mathcal{C}$ for all $t_i \in T_i (a_i)$.

Since (σ_i^*, σ_j^*) is the unique Nash equilibrium, σ_j^* is the only distribution in $\mathcal{C}$ for which $\beta_{a_i'} \cdot \sigma_j = 0$ for $a_i' \in A_i^* \backslash \{a_i\}$. (If σ_j were another such distribution, then

(σ_i^*, σ_j) would also be a Nash equilibrium.) In other words, for any $\sigma_j \in \mathcal{C} \backslash \{\sigma_j^*\}$, $\beta_{a_i'} \cdot \sigma_j \geq 0$ for all $a_i' \neq a_i$ and for at least one $a_i' \in A_i^*$, $\beta_{a_i'} \cdot \sigma_j > 0$. Let

$$\bar{\beta} = \frac{1}{|A_i^*| - 1} \sum_{a_i' \in A_i^* \backslash \{a_i\}} \beta_{a_i'}.$$

Then, $\bar{\beta} \cdot \sigma_j > 0$ for all $\sigma_j \neq \sigma_j^*$. From (A.3), if $\sigma_j^{t_i} \neq \sigma_j^*$ for at least one t_i, then $\bar{\beta} \cdot \sigma_j^* > 0$, a contradiction. Hence, $\sigma_j^{t_i} = \sigma_j^*$ for all $t_i \in T_i(a_i)$. ■

Remark. The restriction to two players is needed for the proof, since this is what allows us to conclude that σ_j^* is the only distribution in $\mathcal{C}$ for which $\beta_{a_i'} \cdot \sigma_j = 0$ for $a_i' \in A_i^* \backslash \{a_i\}$. We do not know if Proposition A or Lemma A extend to three or more players.

We can now discuss the robustness of the example. Let G be a two-player stage game that has a unique correlated equilibrium, σ^*. Let (π, Y) be a monitoring structure with non-moving support: $\pi(y \,|\, a) > 0$ for all $(y, a) \in Y \times A$. Let $\hat{\sigma} = ((\hat{\sigma}_1^1, \hat{\sigma}_1^2), (\hat{\sigma}_2^1, \hat{\sigma}_2^2))$ be a sequential equilibrium of the associated repeated game $G(\pi, Y)$. Suppose, as in the example of Section4, that $\hat{\sigma}^1 \neq \sigma^*$. The outcome of $\hat{\sigma}$ is then not a PPE outcome. This implies that $\hat{\sigma}$ is a PSE, and by Proposition 1, at least one player, say 2, randomizes in the first period.

In order for player 2 to be willing to randomize in the first period, she must be *myopically* indifferent over the randomized actions: By Proposition A, every pair (a_2^1, y) on the equilibrium path causes her to have the same beliefs over player 1's past play, and so she has the same equilibrium expected payoff in period two. Her myopic indifference can be achieved in one of two ways. First, if $\hat{\sigma}_1^1$ is pure (as in our example), $a_1 \in A_1$ and $a_2 \neq a_2' \in A_2$ exist such that $u_2(a_1, a_2) = u_2(a_1, a_2')$. This is not a generic payoff function.

Alternatively, if $\hat{\sigma}_1^1$ is mixed, player 1 randomizes in such a way as to make player 2 indifferent. Then, by the previous paragraph, $(\hat{\sigma}_1^1, \hat{\sigma}_1^2)$ is a mixed strategy equilibrium of $(u, A_1^\sigma \times A_2^\sigma)$, where A_i^σ is the support of $\hat{\sigma}_i^1$ and satisfies $|A_i^\sigma| \geq 2$. For generic assignments of payoffs, $|A_1^\sigma| = |A_2^\sigma|$. Fix a value of the public signal, y. Then player 1 has $|A_1^\sigma|$ types that arise with positive probability in the game of incomplete information induced by σ (described in the proof of Proposition A). By Proposition A, each type of player 1 has the same beliefs over player 2's second-period action. That is, for all $a_1^1 \in A_1^\sigma$ and $a_2^2 \in A_2$,

$$\sum_{a_2^1 \in A_2^\sigma} \hat{\sigma}_2^2(a_2^2 \,|\, a_2^1; y) \Pr\left\{a_2^1 \,|\, a_1^1, y\right\} = \sigma^*(a_2^2).$$

For fixed a_2^2, there are $|A_1^\sigma|$ such equations, simultaneously determining $|A_2^\sigma|$ unknowns, i.e., $\hat{\sigma}_2^2(a_2^2 \,|\, a_2^1; y)$. One solution is given by $\hat{\sigma}_2^2(a_2^2 \,|\, a_2^1; y) = \sigma^*(a_2^2)$. But this solution yields a public strategy for player 2, and so $\hat{\sigma}$ would be a PPE. Thus, there must another solution, and hence the conditional probability matrix,

$$\left[\Pr\left\{a_2^1 \,|\, a_1^1, y\right\}\right]_{a_1^1 a_2^1},$$

must be singular. Since

$$\Pr\left\{a_2^1 \,|\, a_1^1, y\right\} = \frac{\pi\left(y \,|\, a^1\right)\hat{\sigma}_2^1(a_2^1)}{\sum_{a_2} \pi(y \,|\, a_1^1, a_2)\hat{\sigma}_2^1\left(a_2\right)},$$

and $\hat{\sigma}_2^1$ is player 2's strategy in a mixed strategy equilibrium of $(u, A_1^\sigma \times A_2^\sigma)$, the conditional probability matrix is singular only for a nongeneric choice of π.

We conclude that if G is a generic stage game with a unique correlated equilibrium, and (π, Y) is a generic monitoring structure, then that correlated stage-game equilibrium is played in the first period in any sequential equilibrium of the repeated game $G(\pi, Y)$. The example of Section 4 is not robust to perturbations of both payoffs and probabilities.

References

Abreu, D., D. Pearce, and E. Stacchetti (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58, 1041–1063.

Fudenberg, D., D. K. Levine, and E. Maskin (1994): "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62, 997–1040.

Fudenberg, D., and J. Tirole (1991): *Game Theory.* MIT Press, Cambridge, MA.

Kandori, M. (1991): "Cooperation in Finitely Repeated Games with Imperfect Private Information," mimeo, Princeton University.

——— (1992): "The Use of Information in Repeated Games with Imperfect Monitoring," *Review of Economic Studies*, 59, 581–594.

Kandori, M., and I. Obara (2000): "Efficiency in Repeated Games Revisited: The Role of Private Strategies," unpublished.

Kreps, D., and R. Wilson (1982): "Sequential Equilibria," *Econometrica*, 50, 863–894.

Lehrer, E. (1990): "Nash Equilibria of n-Player Repeated Games with Semi-Standard Information," *International Journal of Game Theory*, 19, 191–217.

——— (1991): "Internal Correlation in Rapeated Games," *International Journal of Game Theory*, 19, 431–456.

——— (1992): "Correlated Equilibria in Two-Player Repeated Games with Nonobservable Actions," *Mathematics of Operations Research*, 17(1), 175–199.

Lehrer, E., and S. Sorin (1997): "One-Shot Public Mediated Talk," *Games and Economic Behavior*, 20(2), 131–148.

Mailath, G. J., and S. Morris (2002a): "Finite State Strategies and Coordination in Repeated Games with Private Monitoring," University of Pennsylvania and Yale University.

——— (2002b): "Repeated Games with Almost-Public Monitoring," *Journal of Economic Theory*, 102(1), 189–228.

Marx, L. M., and S. A. Matthews (2000): "Dynamic Voluntary Contribution to a Public Project," *Review of Economic Studies*, 67(2), 327–358.

Matthews, S. A. (1998): "Beneficial Imperfect Public Signals in Multi-Stage Games," University of Pennsylvania.

Myerson, R. B. (1982): "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems," *Journal of Mathematical Economics*, 10, 67–81.

Tomala, T. (1999): "Nash Equilibria of Repeated Games with Observable Payoff Vectors," *Games and Economic Behavior*, 28(2), 310–324.