E C O N O M I C S   B U L L E T I N

A Simple Exposition of Belief-Free Equilibria in Repeated Games

Wojciech Olszewski
Northwestern University

Abstract

Recently, there has been made a substantial progress in the analysis of repeated games with private monitoring. This progress began with introducing a new class of sequential equilibrium strategies, called belief-free equilibria, that can be analyzed using recursive techniques. The purpose of this paper is to explain the general method of constructing belief-free equilibria, and the limit (or bound) on the set of payoff vectors that can be achieved in these strategies in a way that should be easily accessible, even for those who do not pretend to be experts in repeated games.

Citation: Olszewski, Wojciech, (2007) "A Simple Exposition of Belief-Free Equilibria in Repeated Games." *Economics Bulletin*, Vol. 3, No. 58 pp. 1-16
Submitted: October 18, 2007. **Accepted:** November 12, 2007.
URL: http://economicsbulletin.vanderbilt.edu/2007/volume3/EB-07C70038A.pdf

1. Introduction

Repeated games is the simplest setting for studying long-run interactions. Most real-life examples of long-run interactions involve imperfect private monitoring of actions, i.e. situations in which players observe possibly different signals, so they might not know the signals of their opponents. Repeated games with private monitoring have long been known for their intractability. This intractability comes from a complex statistical inference that one has to conduct to determine the optimal continuation strategies for any given player.[1]

Recently however, there has been made a substantial progress in the analysis of repeated games with private monitoring. This progress began with introducing a new class of sequential equilibrium strategies, *belief-free equilibria*, that can be analyzed using recursive techniques. The key characteristic of belief-free strategies is that when they are used, the set of optimal continuation strategies for any given player is independent of the prior history of play. That is, the problem of complex statistical inference is, for these class of strategies, assumed away.

Belief-free equilibrium was first used by Piccione (2002) to establish the folk theorem for the repeated, two-player Prisoner's Dilemma with private, almost-perfect monitoring. Ely and Välimäki (2002) proved the same result by using a substantially simpler construction (of belief-free equilibria). It should also be mentioned that Kandori and Obara (2006) used independently a construction similar to Ely and Välimäki in a paper on repeated games with public monitoring. For some time, there was a consensus among researchers that belief-free equilibrium was the most promising tool for the analysis of imperfect private monitoring. Ely et al. (2005) provide a simple and sharp characterization of equilibrium payoffs that can be achieved in belief-free equilibria; they conclude that while this set of payoffs is large, belief-free strategies are not rich enough to generate a folk theorem in most games besides the prisoner's dilemma, even under almost-perfect monitoring. Recently, their characterization was generalized to the case of more than two players in Yamamoto (2006b). Despite the impossibility of generating a folk theorem, belief-free equilibria remain an important tool in the studies on private monitoring (see, for example, Matsushima (2004), Hörner and Olszewski (2006), and Yamamoto (2006a)), as well as on some other topics in repeated games (see Takahashi (2007)).

Although belief-free strategies are the simplest among all known equilibria in repeated games with imperfect private monitoring, the construction of these strategies, especially the general methods developed in Ely et al., may not be so straightforward even for experts in repeated games. The purpose of this paper is to explain the general method of constructing belief-free equilibria, and the limit (or bound) on the set of payoff vectors that can be achieved in these strategies. The exposition should be fairly easily accessible, even for those who do not pretend to be experts in repeated games. This is achieved partially by restricting attention to specific examples. These

[1] We refer the reader to Kandori (2002), and Mailath and Samuelson (2006) for a detailed discussion of difficulties that arise in the analysis of repeated games with private monitoring.

examples require less tedious calculations and allow for more explicit description of strategies than the general case, yet basically identical ideas as those behind the specific examples apply in the general case.

We attempt to achieve our objective by presenting detailed (rather than verbal and informal) arguments. Also, we discuss only very briefly the properties of belief-free equilibria. We refer the reader to Ely et al., and Mailath and Samuelson (2006) for a much more complete discussion; Bhaskar, Mailath and Morris (2006) provide a criticism of usefulness of belief-free equilibria for studying long-run interactions on the grounds of Harsanyi's purifiability.

2. Basic Model of Repeated Games with Private Monitoring

We follow most of the notation from Fudenberg et al. (1994). In the stage game, player $i = 1, ..., n$ chooses action a_i from a finite set A_i. We call a vector $a \in A = A_1 \times ... \times A_n$ a profile of actions. A private monitoring structure is a collection of distributions $\{\pi(\cdot|a) : a \in A\}$ over $Y_1 \times ... \times Y_n$, where all sets Y_i are finite. The interpretation is that each player receives a signal $y_i \in Y_i$, and each signal profile $(y_1, \ldots, y_n)$ obtains with probability $\pi(y_1, \ldots, y_n|a)$. Player i's realized payoff $r_i(a_i, y_i)$ depends on the action a_i and the private signal y_i. Player i's expected payoff from action profile a is therefore

$$g_i(a) = \sum_{(y_1,\ldots,y_n)\in Y_1\times\ldots\times Y_n} \pi(y_1, \ldots, y_n|a) r_i(a_i, y_i).$$

Player i's private history in the repeated game is $h_i^t = (a_i^1, y_i^1, ..., a_i^{t-1}, y_i^{t-1})$. We call a vector $h^t = (h_1^t, ..., h_n^t)$ a profile of histories. A strategy σ_i for player i is a sequence of functions $(\sigma_i^t)_{t=1}^{\infty}$ where σ_i^t maps each h_i^t to a probability distribution over A_i. Players share a common discount factor $\delta < 1$. All repeated game payoffs are discounted and normalized by a factor $1 - \delta$. Thus, if $(g_i^t)_{t=1}^{\infty}$ is player i's sequence of stage-game payoffs, the repeated game payoff is

$$(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} g_i^t.$$

3. Definition of Belief-Free Equilibrium

The key difficulty in the analysis of repeated games with private monitoring is complex statistical inference. To determine the best-response continuation strategies of a player i at his/her private history h_i^t, one typically has to determine player i's expectation regarding his/her opponents' strategy; in practice, one thus has to compute the probability distribution, conditional on h_i^t, over private histories of player i's opponents. This task is very complicated even for very simple strategies and monitoring structures.

Belief-free equilibria is a class of equilibria in repeated games. They can be defined not only for private, but also for perfect and public monitoring structures. However, they became particularly useful in research on private monitoring. Belief-free equilibria are most tractable among all known classes of equilibria in repeated games with private monitoring. They are simple, because the problem of complex statistical inference is, by definition, assumed away for this class of equilibria.

Definition 1. *A strategy profile σ is belief-free if for every profile of histories $h^t = (h_1^t, ..., h_n^t)$, the continuation strategy $\sigma_i | h_i^t$ is a best response of player i to the continuation strategies $\sigma_j | h_j^t$ $(j \neq i)$ of player i's opponents.*

Notice the difference with the equilibrium condition, which requires that $\sigma_i | h_i^t$ is a best response contingent on player i's beliefs about the history profile h_{-i}^t at history h_i^t, not contingent on every single history profile h_{-i}^t.

Every belief-free equilibrium gives rise to a sequence of (nonempty) *regimes* $(\mathcal{A}^t)_{t=1}^{\infty}$, which are defined as the largest subsets of action profiles $\mathcal{A}^t = \mathcal{A}_1^t \times ... \times \mathcal{A}_n^t \subset A_1 \times ... \times A_n$ with the property that playing any sequence of actions $(a_i^t)_{t=t^*}^{\infty}$, $a_i^t \in \mathcal{A}_i^t$ for $t \geq t^*$, is a best response of player i to the continuation strategies of player i's opponents contingent on any history profile h_{-i}^t. By the definition of belief-free equilibrium, $\mathcal{A}_i^t$ contains all actions a_i such that for some history h_i^t mixed strategy $\sigma_i^t(h_i^t)$ assigns a positive probability to a_i, but $\mathcal{A}_i^t$ may contain also other actions.

At first, the definition of regime may look complicated, and the usefulness of this concept may not be immediately appreciated. We hope, however, that the examples from Section 4 will be turn out very helpful.

4. Folk Theorem for Prisoner's Dilemma

Consider the following two-player Prisoner's Dilemma:

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

Each player i obtains one of two possible signals: c or d; the probability distribution over signals depends only on the action of his/her opponent; the signal is correct, i.e. it is c when player $-i$'s action was C, or (respectively) it is d when player $-i$'s action was D, with probability $p > 1/2$.

This payoff matrix is purposely non-generic: A player's payoff to playing C is always by 1 smaller than the payoff to playing D, independently of the player opponent's action; also, a player's payoff is always by 2 larger when his/her opponent plays C (compared to playing D), independently of the player's own action. The two non-generic features enable us to present ideas without running into complicated calculations.

4.1. Full Cooperation

We shall first show that full cooperation, i.e. payoff vector $(2, 2)$, can be approximated by belief-free equilibria when $\delta \to 1$ and $p \to 1$, i.e. when discounting and monitoring imperfections vanish. Consider the following strategies with one-period memory:

$$\begin{aligned} \sigma_i^t &= C \text{ with prob. } 1 \text{ if } y_i^{t-1} = c; \\ \sigma_i^t &= D \text{ with prob. } \beta \text{ if } y_i^{t-1} = d. \end{aligned}$$

Player i begins with playing D with probability $(1-p)\beta$ in period 1. This pair of strategies is a belief-free equilibrium (with the constant sequence of regimes $\{C, D\} \times \{C, D\}$) if each player i is indifferent between playing C and playing D, independently of his/her opponent's history, i.e. when

$$1 = \delta(2p-1)\beta 2. \tag{1}$$

The left-hand side of this equality is equal to the difference in player i's flow payoffs to playing D and C; it is independent of player $-i$'s action, and so his/her history up to period t. The right-hand side is equal to the discounted expected difference in continuation payoffs that follow playing C and D. As strategies have one-period memory, this difference comes entirely from player $-i$'s action in period $t+1$. Playing D by player i increases (compared to playing C) the chance that player $-i$ will receive signal d by $2p-1$; receiving signal d by player $-i$ increases (compared to receiving signal c) the chance of playing D by β, and if the opponent plays D, it reduces player i's payoff by 2, independently of his/her own action. Again, the right-hand side is independent of player $-i$'s history up to period t.

The idea is that player $-i$ plays D with a slightly higher probability contingent on signal d. This gives player i an incentive to play C, and offsets his/her myopic incentive to play C. If the probability of playing D by player $-i$ contingent on signal d is appropriately chosen, player i is indifferent between playing C and D.

Notice that $\beta \to 1/2$ when $\delta \to 1$ and $p \to 1$.

Since the prescribed strategies are belief-free, player i's equilibrium payoff can be computed under the assumption that player i plays C in every period. Thus, player i receives in period 1 the flow payoff of $2[1-(1-p)\beta] = 2[p+(1-p)(1-\beta)]$ with probability 1; in periods $t > 1$, player i receives the flow payoff of 2 if player $-i$ plays C, which happens with probability $p + (1-p)(1-\beta)$, and the flow payoff of 0 if player $-i$ plays D, which happens with the remaining probability. This yields the total (normalized) payoff of

$$(1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} 2[p+(1-p)(1-\beta)] = [p+(1-p)(1-\beta)]2.$$

In the limit as p tends to 1, this yields payoff vector $(2, 2)$.

The same idea enabled Ely and Välimäki to approximate full cooperation in a "general" version of the Prisoner's Dilemma; however, condition (1) must then be replaced by a system of more complicated linear equations which must be satisfied by a number of probabilities which correspond to single probability β.

4.2. Other Payoffs

A repetition of stage-game Nash equilibrium (D, D) achieves the payoff vector $(1, 1)$. The set of feasible and IR payoffs (see Figure 1) is spanned by two more payoff vectors: $(2.5, 1)$ and $(1, 2.5)$; they can also be approximated in belief-free equilibria (when $\delta \to 1$ and $p \to 1$).

An obvious idea for constructing a belief-free equilibrium that approximates payoff vector $(2.5, 1)$ is to modify the strategies from the previous section so that: player 1 is allowed to play D in every second period with no effect on the continuation play, and (s)he is indifferent between playing C and D in all remaining periods, while player 2 is indifferent between playing C and D in all periods.

To examine this idea more formally, consider the following strategy (of player 1) with two-period memory:

$$\begin{aligned}
\sigma_1^t &= D \text{ with prob. } 1 \text{ if } t \text{ is even;} \\
\sigma_1^t &= C \text{ with prob. } 1 \text{ if } t \text{ is odd, and } y_1^{t-1} = y_1^{t-2} = c; \\
\sigma_1^t &= D \text{ with prob. } \beta \text{ if } t \text{ is odd, } y_1^{t-1} = c \text{ and } y_1^{t-2} = d; \\
\sigma_1^t &= D \text{ with prob. } \gamma \text{ if } t \text{ is odd, } y_1^{t-1} = d \text{ and } y_1^{t-2} = c; \\
\sigma_1^t &= D \text{ with prob. } \beta + \gamma \text{ if } t \text{ is odd, and } y_1^{t-1} = y_1^{t-2} = d,
\end{aligned}$$

where β is defined by (1), and γ is defined by

$$1 = \delta^2(2p - 1)\gamma 2. \tag{2}$$

Player 1 begins with playing C with probability q (which will be specified later) in period 1.

Consider also the following strategy of player 2:

$$\begin{aligned}
\sigma_2^t &= C \text{ with prob. } 1 \text{ if } t \text{ is even;} \\
\sigma_2^t &= C \text{ with prob. } 1 \text{ if } t \text{ is odd, and } y_2^{t-2} = c; \\
\sigma_2^t &= D \text{ with prob. } \gamma \text{ if } t \text{ is odd, and } y_2^{t-2} = d.
\end{aligned}$$

Player 2 begins with playin C with probability $p + (1 - p)(1 - \gamma)$.

Notice that the definition of the strategy of player 1 is incorrect, because

$$\beta + \gamma = \frac{1}{\delta(2p - 1)2} + \frac{1}{\delta^2(2p - 1)2} > 1$$

if δ or p is smaller than 1, but disregard this problem for a moment. We shall show that this is actually the only problem with this pair of strategies, and then we explain how they can be modified to obtain a well-defined belief-free equilibrium that approximates payoff vector $(2.5, 1)$.

Player 1 strictly prefers playing D in every even period, as it yields a higher flow payoff (compared to playing C) and the continuation payoff (from the following odd period) is independent of player 2's signal. By the argument used in the construction of equilibrium that approximates full cooperation, player 1 is indifferent between playing C and D in odd periods if condition (1) is satisfied. It is the case independently of player 2's history up to period t. Similarly, player 2 is indifferent (in every period) between playing C and D, independently of player 1's history. Notice that the indifference between playing C and D in even periods is guaranteed by condition condition (1), and in odd periods by condition (2).

Thus, the prescribed strategies would be a belief-free equilibrium with the sequence of regimes $\{C, D\} \times \{C, D\}$, $\{D\} \times \{C, D\}$, $\{C, D\} \times \{C, D\}$, $\{D\} \times \{C, D\}$, ..., if $\beta + \gamma$ were a number from interval $[0, 1]$. To compute the payoff of player 1 in this "equilibrium", one can assume that player 1 plays C in odd periods and D in even periods, while player 2 plays the prescribed strategy. This yields

$$(1-\delta)\left\{\sum_{t \text{ odd}} \delta^{t-1} 2[p + (1-p)(1-\gamma)] + \sum_{t \text{ even}} \delta^{t-1} 3\right\} =$$

$$= \frac{\delta 2[p + (1-p)(1-\gamma)]}{1+\delta} + \frac{3\delta}{1+\delta} \rightarrow_{p,\delta\to 1} 2.5.$$

Similarly, to compute the payoff of player 2 in this "equilibrium", one can assume that player 2 plays C in every period, while player 1 plays the prescribed strategy. This yields

$$(1-\delta)\sum_{t \text{ odd}} \delta^{t-1} 2[p^2 + p(1-p)(1-\beta) + p(1-p)(1-\gamma) + (1-p)^2(1-\beta-\gamma)] =$$

$$= \frac{2[p^2 + p(1-p)(1-\beta) + p(1-p)(1-\gamma) + (1-p)^2(1-\beta-\gamma)]}{1+\delta} \rightarrow_{p,\delta\to 1} 1,$$

if the probability that player 1 plays C in period 1 is defined by

$$q := p^2 + p(1-p)(1-\beta) + p(1-p)(1-\gamma) + (1-p)^2(1-\beta-\gamma).$$

Indeed, the flow payoff of player 2 is equal to 0 in even periods. In odd periods, the flow payoff of player 2 is equal to 0 if player 1 plays D, and it is equal to 2 if player 1 plays C. Player 1 plays C in odd periods $t > 1$ with positive probability in the following cases: (a) (s)he plays C with probability 1 if $y_1^{t-1} = y_1^{t-2} = c$, which happens with probability p^2; (b) with probability $(1-\beta)$ if $y_1^{t-1} = c$ and $y_1^{t-2} = d$ or with probability $(1-\gamma)$ if $y_1^{t-1} = d$ and $y_1^{t-2} = c$ (both happen with probability $p(1-p)$), and(c) with probability $(1-\beta-\gamma)$ if $y_1^{t-1} = y_1^{t-2} = d$, which happens with probability $(1-p)^2$.

Hence the only problem is that strategy σ_1 is not well-defined ($\beta + \gamma > 1$). Since player 1 is allowed to play D in every second period with no effect on the continuation play, his/her strategy in those periods cannot depend on signals. This leaves only "half of the periods" for giving player 2 incentives to play C, and it turns out not to be enough.

However, the problem does not seem too serious. Although, $\beta+\gamma > 1$, it converges to 1 as $\delta \to 1$ and $p \to 1$. This suggests that if there were a little more time (than half of the periods) for giving player 2 incentives to play C, strategies similar to σ_1 and σ_2 could be a well-defined belief-free equilibrium. One can therefore try to modify these strategies so that: player 1 is allowed to play D in every m out of $2m + 1$ periods with no effect on the continuation play, and (s)he is indifferent between playing C and D in the remaining $m + 1$ out of the $2m + 1$ periods, while player 2 is indifferent between playing C and D in all periods. In Appendix A, we show that this new pair of strategies is indeed a belief-free equilibrium (for sufficiently large discount factors), and it yields a payoff vector that converges to $(2.5, 1)$ as $\delta \to 1$.

5. A Bound on Belief-Free Equilibrium Payoff Vectors

We shall now provide a bound on the set of all payoff vectors that can be achieved in belief-free equilibria. Let $\mathcal{J}$ denote the set of all regimes. For all $i = 1, ..., n$ and $\mathcal{A} \in \mathcal{J}$ define

$$m_i^{\mathcal{A}} := \min_{\alpha_{-i} \in \Delta \mathcal{A}_{-i}} \max_{a_i \in A_i} g_i(\alpha) \text{ and } M_i^{\mathcal{A}} := \max_{\alpha_{-i} \in \Delta \mathcal{A}_{-i}} \min_{a_i \in \mathcal{A}_i} g_i(\alpha), \tag{3}$$

and let m_i and M_i denote vectors $\left(m_i^{\mathcal{A}}\right)_{\mathcal{A} \in \mathcal{J}}$ and $\left(M_i^{\mathcal{A}}\right)_{\mathcal{A} \in \mathcal{J}}$, respectively. Given a probability distribution over regimes $p \in \Delta \mathcal{J}$, let

$$p \circ m_i := \sum_{\mathcal{A} \in \mathcal{J}} p(\mathcal{A}) m_i^{\mathcal{A}} \text{ and } p \circ M_i := \sum_{\mathcal{A} \in \mathcal{J}} p(\mathcal{A}) M_i^{\mathcal{A}}.$$

Let V^{BF} be the set of all payoff vectors that can be achieved in belief-free equilibria

Theorem 1.

$$V^{BF} \subset \bigcup \left\{ \prod_{i=1}^{n} [p \circ m_i, p \circ M_i] : p \in \Delta \mathcal{J} \right\}$$

Notice that this theorem applies to any (finite) game and any (finite) monitoring structure. See Appendix B for the proof.

For a large class of stage games, inclusion can be replaced in Theorem 1 with equality, and for all games Theorem 1 is one of the key steps to a complete characterization of belief-free equilibrium payoffs. In the case of two players, such a characterization was obtained by Ely et al. (2005) by using a combination of the argument behind

Theorem 1 and a generalization of the construction from Section 4. In Appendix C, we apply Theorem 1 to showing that for some stage games, belief-free strategies are not rich enough to generate a folk theorem.

In the case of more than two players and almost-perfect monitoring, the characterization of belief-free equilibrium payoffs was obtained by Yamamoto (2006b). The construction from Section 4 does not extend easily to more than two players; instead, Yamamoto adapts a much more complicated construction from Hörner and Olszewski (2006). The difference in studying two and more than two players is that the self-generation techniques developed by Abreu et al. (1990) fairly easily generalize to private monitoring (for the class of belief-free equilibria) in the case of two players, but no generalization is known for more than two players. Nevertheless, some belief-free equilibria can be constructed in a manner similar to Section 4, or by techniques developed by Abreu et al. (1990); Yamamoto (2006a) shows, for example, that one can approximate in this way full cooperation in games similar to the n-player Prisoner's Dilemma.

6. Appendix A

We will construct a belief-free equilibrium that approximates payoff vector $(2.5, 1)$. Let the strategy of player 1 be given by

$$\begin{aligned}
\sigma_1^t &= D \text{ with prob. } 1 \text{ if } t = k(2m+1)+1, ..., k(2m+1)+m; \\
\sigma_1^t &= C \text{ with prob. } 1 \text{ if } t \in \{k(2m+1)+m+1, ..., k(2m+1)+2m\}, \\
y_1^{t-2m-1} &= y_1^{t-3m-1} = c; \\
\sigma_1^t &= D \text{ with prob. } \beta_1 \text{ if } t \in \{k(2m+1)+m+1, ..., k(2m+1)+2m\}, \\
y_1^{t-2m-1} &= c \text{ and } y_1^{t-3m-1} = d; \\
\sigma_1^t &= D \text{ with prob. } \gamma_1 \text{ if } t \in \{k(2m+1)+m+1, ..., k(2m+1)+2m\}, \\
y_1^{t-2m-1} &= d \text{ and } y_1^{t-3m-1} = c; \\
\sigma_1^t &= D \text{ with prob. } \beta_1 + \gamma_1 \text{ if } t \in \{k(2m+1)+m+1, ..., k(2m+1)+2m\}, \\
y_1^{t-2m-1} &= y_1^{t-3m-1} = d;
\end{aligned}$$

$$\begin{aligned}
\sigma_1^t &= D \text{ with prob. } \gamma_1 + \sum_{i=1}^{m} \alpha_i \cdot \varepsilon_i \text{ if } t = k(2m+1)+m+1, \text{ and } y_1^{t-2m-1} = d; \\
\sigma_1^t &= D \text{ with prob. } \sum_{i=1}^{m} \alpha_i \cdot \varepsilon_i \text{ if } t = k(2m+1)+m+1, \text{ and } y_1^{t-2m-1} = c,
\end{aligned}$$

where

$$\begin{aligned}
\varepsilon_i &= 1 \text{ if } y_1^{t-i-3m-1} = d; \\
\varepsilon_i &= 0 \text{ if } y_1^{t-i-3m-1} = c;
\end{aligned}$$

$$1 = \delta^m (2p-1)\gamma_1 2, \tag{4}$$

and

$$1 = \delta^{2m}(2p-1)\beta_1 2 + \delta^{2m+i}(2p-1)\alpha_i 2. \tag{5}$$

Of course, we have to pick numbers β_1 , γ_1, and α_i, $i = 1, ..., m$, such that all probabilities belong to interval $[0,1]$. For a fixed m, $\gamma_1 \to 1/2$ when $\delta \to 1$ and $p \to 1$; so, taking $\beta_1 = 1 - \gamma_1$, we have that also $\beta_1 \to 1/2$. Thus, we can take $\alpha_i \to 0$, and so all probabilities belong to interval $[0,1]$.

Let further the strategy of player 2 be given by

$$\begin{aligned}
\sigma_2^t &= C \text{ with prob. } 1 \text{ if } t = k(2m+1)+1, \text{ ..., } k(2m+1)+m; \\
\sigma_2^t &= C \text{ with prob. } 1 \text{ if } t \in \{k(2m+1)+m+1, \text{ ..., } (k+1)(2m+1)\}, \\
y_2^{t-2m-1} &= c; \\
\sigma_2^t &= D \text{ with prob. } \gamma \text{ if } t \in \{k(2m+1)+m+1, \text{ ..., } (k+1)(2m+1)\}, \\
y_2^{t-2m-1} &= d,
\end{aligned}$$

where

$$1 = \delta^{2m+1}(2p-1)\gamma_2 2. \tag{6}$$

Conditions (4)-(6) guarantee that player 2 is indifferent between playing C and D in all periods, and player 1 strictly prefers playing D in periods $k(2m+1)+1$, ..., $k(2m+1)+m$ and is indifferent between playing C and D in periods $k(2m+1)+m+1$, ..., $k(2m+1)+2m+1$. Thus, the prescribed strategy profile is a belief-free equilibrium with the sequence of regimes

$$\underbrace{\{D\} \times \{C,D\}}_{m \text{ times}}, \underbrace{\{C,D\} \times \{C,D\}}_{m+1 \text{ times}}, \underbrace{\{D\} \times \{C,D\}}_{m \text{ times}}, \underbrace{\{C,D\} \times \{C,D\}}_{m+1 \text{ times}}, \ldots$$

To compute the payoff of player 1, one can assume that player 1 plays D in periods $t = k(2m+1)+1$, ..., $k(2m+1)+m$ and C in periods $k(2m+1)+m+1$, ..., $(k+1)(2m+1)$, while player 2 plays the prescribed strategy; and to compute the payoff of player 2, one can assume that player 2 plays C in every period, while player 1 plays the prescribed strategy. The precise expression for this payoff vector is slightly tedious, but it is easy to see that in the limit (as $\delta \to 1$ and $p \to 1$) the payoff of player 1 is equal to the average of the payoff to playing (D,C) in m out of $(2m+1)$ periods and the payoff to playing (C,C) in the remaining $m+1$ out of $(2m+1)$ periods. Taking m sufficiently large, this average is close to 2.5. Similarly, the payoff of player 2 is equal to the average of the payoff to playing (D,C) in m out of $(2m+1)$ periods and the payoff to playing (C,C) in the remaining $m+1$ out of $(2m+1)$ periods; for m sufficiently large, this average is close to 1.

7. Appendix B

Proof of Theorem 1: Consider a belief-free equilibrium with the sequence of regimes $\mathcal{A}^1$, $\mathcal{A}^2$, ... Let $w_i^t(h_{-i}^t)$ denote the equilibrium continuation payoff of player

i (the payoff in the repeated game that begins in period t), contingent on private histories of player i's opponents at the beginning of period t. Because the equilibrium is belief-free, this continuation payoff is independent of player i's private history; in general, the equilibrium continuation payoff of player i would depend on his/her own private history, as the continuation strategy typically depends on player i's private history. Let further

$$W_i^t := \max_{h_{-i}^t} w_i^t(h_{-i}^t).$$

For every private history profile h_{-i}^t, we have

$$w_i^t(h_{-i}^t) \le (1-\delta) \min_{a_i \in \mathcal{A}_i^t} g_i(a_i, \sigma_{-i} \mid h_{-i}^t) + \delta W_i^{t+1} \le (1-\delta) M_i^{\mathcal{A}^t} + \delta W_i^{t+1};$$

the first inequality follows from the definition of regime and the definition of W_i^{t+1}, and the second inequality follows from the definition of $M_i^{\mathcal{A}^t}$. Thus,

$$W_i^t \le (1-\delta) M_i^{\mathcal{A}^t} + \delta W_i^{t+1};$$

applying this inequality iteratively, we obtain that

$$W_i^1 \le \sum_{t=1}^{\infty} (1-\delta)\delta^{t-1} M_i^{\mathcal{A}^t}.$$

Define a probability distribution over regimes $p \in \Delta\mathcal{J}$ by letting

$$p(\mathcal{A}) := \sum_{t:\mathcal{A}^t = \mathcal{A}} (1-\delta)\delta^{t-1}.$$

Then the last inequality can be written as

$$W_i^1 \le p \circ M_i.$$

Similarly, let

$$w_i^t := \min_{h_{-i}^t} w_i^t(h_{-i}^t).$$

For every private history profile h_{-i}^t, we have

$$w_i^t(h_{-i}^t) \ge (1-\delta) \max_{a_i \in A_i} g_i(a_i, \sigma_{-i} \mid h_{-i}^t) + \delta w_i^{t+1} \ge (1-\delta) m_i^{\mathcal{A}^t} + \delta w_i^{t+1},$$

so

$$w_i^t \ge (1-\delta) m_i^{\mathcal{A}^t} + \delta w_i^{t+1};$$

and applying this inequality iteratively,

$$w_i^1 \ge \sum_{t=1}^{\infty} (1-\delta)\delta^{t-1} m_i^{\mathcal{A}^t} = p \circ m_i.$$

It remains to notice that the equilibrium payoff of player i is equal to $w_i^1 = W_i^1$. ■

8. Appendix C

We shall now show that for some stage games, belief-free strategies are not rich enough to generate a folk theorem. Consider the following Battle of Sexes:

	B	O
B	$2,1$	$0,0$
O	$0,0$	$1,2$

This game has three stage-game Nash equilibria: $(B,B),(O,O)$, and an equilibrium in mixed strategies in which player 1 (row) plays B with probability 2/3 and player 2 (column) plays O with probability 2/3. The three stage-game equilibria yield payoffs: $(2,1)$, $(1,2)$ and $(2/3,2/3)$, respectively. We shall show that the convex hull the three payoff vectors contains V^{BF}, the set of all payoff vectors that can be achieved in belief-free equilibria of the repeated game. Thus, for the Battle of Sexes, the set V^{BF} is strictly smaller than the set of feasible and IR payoff vectors (see Figure 2).

It follows directly from formula (3) that

	$\{B\}\times\{B\}$	$\{B\}\times\{O\}$	$\{O\}\times\{B\}$	$\{O\}\times\{O\}$	$\{B,O\}\times\{B\}$
$M_1^{\mathcal{A}}$	2	0	0	1	0
$M_2^{\mathcal{A}}$	1	0	0	2	1
$m_1^{\mathcal{A}}$	2	1	2	1	2
$m_2^{\mathcal{A}}$	1	1	2	2	2/3

	$\{B,O\}\times\{O\}$	$\{B\}\times\{B,O\}$	$\{O\}\times\{B,O\}$	$\{B,O\}\times\{B,O\}$
$M_1^{\mathcal{A}}$	0	2	1	2/3
$M_2^{\mathcal{A}}$	2	0	0	2/3
$m_1^{\mathcal{A}}$	1	2/3	2/3	2/3
$m_2^{\mathcal{A}}$	2/3	1	2	2/3

For the purpose of computing the bound from Theorem 1, two regimes: $\{B\}\times\{O\}$ and $\{O\}\times\{B\}$ can be disregarded as $m_i^{\mathcal{A}} > M_i^{\mathcal{A}}$ for $i = 1,2$. Other two regimes $\{B,O\}\times\{B\}$ and $\{B\}\times\{B,O\}$ can be disregarded as well; $\{B,O\}\times\{B\}$ is dominated by $\{B,O\}\times\{O\}$, in the sense that $m_i^{\mathcal{A}}$ for the former is no lower than $m_i^{\mathcal{A}}$ for the latter and $M_i^{\mathcal{A}}$ for the former is no higher than $m_i^{\mathcal{A}}$ for the latter, and $\{B\}\times\{B,O\}$ is dominated (in the same sense) by $\{O\}\times\{B,O\}$.

Further, notice that

$$\left[m_1^{\mathcal{A}}, \times M_1^{\mathcal{A}}\right] \times \left[m_2^{\mathcal{A}}, \times M_2^{\mathcal{A}}\right] = \left[\frac{17}{21}, \frac{24}{21}\right] \times \left\{\frac{6}{7}\right\} \subset co\{(2,1),(1,2),(2/3,2/3)\},$$

for the regime $\mathcal{A} = \{B, O\} \times \{O\}$ (co stands for the convex-hull operator), and similarly

$$\left[m_1^{\mathcal{A}}, \times M_1^{\mathcal{A}}\right] \times \left[m_2^{\mathcal{A}}, \times M_2^{\mathcal{A}}\right] = \left\{\frac{6}{7}\right\} \times \left[\frac{17}{21}, \frac{24}{21}\right] \subset co\{(2,1),(1,2),(2/3,2/3)\}$$

for the regime $\mathcal{A} = \{O\} \times \{B, O\}$. Since

$$\begin{aligned} \left[m_1^{\{B\}\times\{B\}}, \times M_1^{\{B\}\times\{B\}}\right] \times \left[m_2^{\{B\}\times\{B\}}, \times M_2^{\{B\}\times\{B\}}\right] &= \{(2,1)\}, \\ \left[m_1^{\{O\}\times\{O\}}, \times M_1^{\{O\}\times\{O\}}\right] \times \left[m_2^{\{O\}\times\{O\}}, \times M_2^{\{O\}\times\{O\}}\right] &= \{(1,2)\}, \end{aligned}$$

and

$$\left[m_1^{\{B,O\}\times\{B,O\}}, \times M_1^{\{B,O\}\times\{B,O\}}\right] \times \left[m_2^{\{B,O\}\times\{B,O\}}, \times M_2^{\{B,O\}\times\{B,O\}}\right] = \{(2/3, 2/3)\},$$

regimes $\{B, O\} \times \{O\}$ and $\{O\} \times \{B, O\}$ can also be disregarded for the purpose of computing the bound from Theorem 1, and consequently, this bound is equal to $co\{(2,1),(1,2),(2/3,2/3)\}$.

9. References

ABREU, D., D. PEARCE, AND E. STACHETTI (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, **58**, 1041–1063.

BHASKAR, V., G. MAILATH, AND S. MORRIS (2006): "Purification in the Infinitely Repeated Prisoner's Dilemma," Mimeo, Cowles Foundation.

ELY, J. C., AND J. VÄLIMÄKI (2002): "A Robust Folk Theorem for the Prisoner's Dilemma," *Journal of Economic Theory*, **102**, 84–105.

ELY, J., J. HÖRNER, AND W. OLSZEWSKI (2005): "Belief-Free Equilibria in Repeated Games," *Econometrica*, **73**, 377–415.

FUDENBERG, D.,D. LEVINE, AND E. MASKIN (1994): "The Folk Theorem with Imperfect Public Information," *Econometrica*, **62**, 997–1039.

HÖRNER, J., AND W. OLSZEWSKI (2006): "The Folk Theorem for Games with Private Almost-Perfect Monitoring," *Econometrica*, **74**, 1499–1544.

KANDORI, M. (2002): "Introduction to Repeated Games with Private Monitoring," *Journal of Economic Theory*, **102**, 1–15.

KANDORI, M., AND I. OBARA (2003): "Efficiency in Repeated Games Revisited: The Role of Private Strategies," *Econometrica*, **74**, 499–519.

MAILATH, G. AND L. SAMUELSON (2006): "Repeated Games and Reputations: Long-Run Relationships," Oxford University Press.

MATSUSHIMA, H. (2004): "Repeated Games with Private Monitoring: Two Players," *Econometrica*, **72**, 823–852.

PICCIONE, M. (2002): “The Repeated Prisoners’ Dilemma with Imperfect Private Monitoring,” *Journal of Economic Theory*, **102**, 70–83.

TAKAHASHI, S. (2007): ”Community Enforcement when Players Observe Partners’ Past Play,” Mimeo, Harvard University.

YAMAMOTO, Y. (2006a): “Efficiency Results in N Player Games with Imperfect Private Monitoring,” forthcoming in *Journal of Economic Theory*.

YAMAMOTO, Y. (2006b): “A Characterization of Belief-Free Payoffs in Repeated Games with Almost-Perfect Monitoring,” Mimeo, Harvard University.

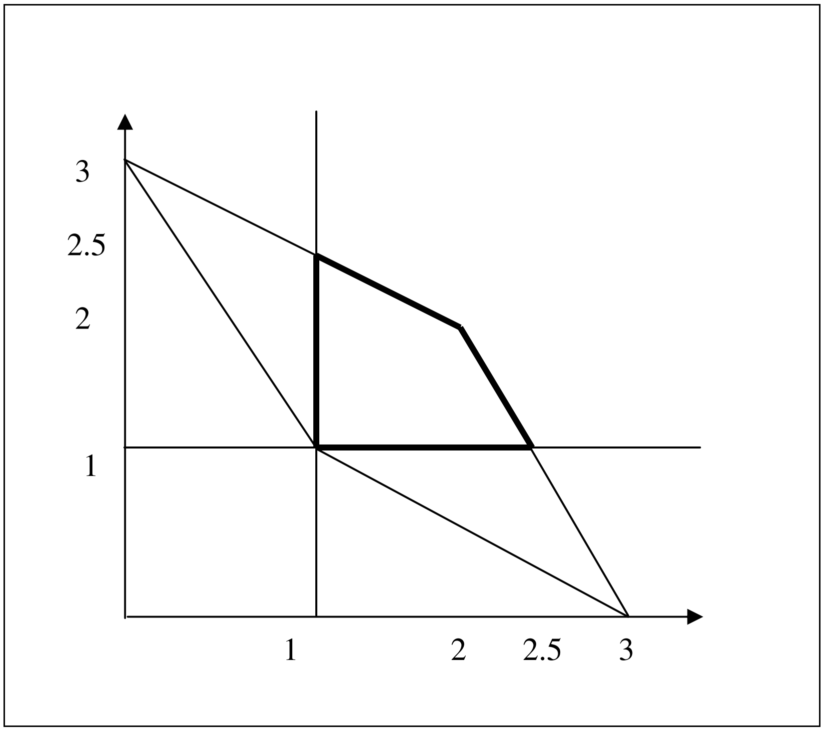

Figure 1 (The Prisoner's Dilemma): The set of feasible payoffs is spanned by vectors: (1,1), (2,2), (3,0) and (0,3), and the minmax payoff is equal to 1, both for player 1 and for player 2. Thus, the set of all feasible and IR payoffs (equal to the set of all belief-free payoff vectors) is spanned by payoff vectors: (1,1), (2,2), (2.5,1) and (1,2.5).

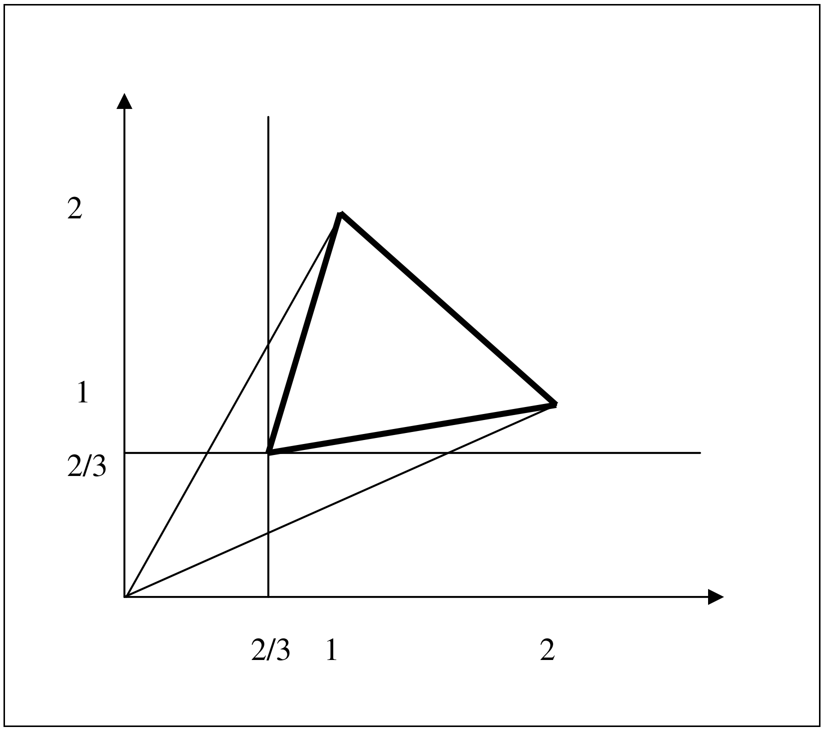

Figure 2 (The Battle of Sexes): The set of feasible payoffs is spanned by vectors: (2,1), (1,2), and (0,0), and the minmax payoff is equal to 2/3, both for player 1 and for player 2. The set of all belief-free payoff vectors is spanned by payoff vectors: (2,1), (1,2), and (2/3,2/3). It is equal to the convex hull of stage-game Nash equilibrium payoffs, and it is strictly smaller than the set of all feasible and IR payoffs.