CARESS Working Paper #00-16
Belief-Based Equilibria in the Repeated Prisoners' Dilemma with Private Monitoring[¤]

V. Bhaskar[y]
Department of Economics
University of Essex

Ichiro Obara[z]
Department of Economics
University of Pennsylvania

July 2000

Abstract

We analyze in…nitely repeated prisoners' dilemma games with imperfect private monitoring, and construct sequential equilibria where strategies are measurable with respect to players' beliefs regarding their opponents' continuation strategies. We show that, when monitoring is almost perfect, the symmetric e¢cient outcome can be approximated in any prisoners' dilemma game, while every individually rational feasible payo¤ can be approximated in a class of prisoner dilemma games. We also extend the approximate e¢ciency result to n-player prisoners' dilemma games and to prisoner's dilemma games with more general information structure. Our results require that monitoring be su¢ciently accurate but do not require very low discounting.

1 Introduction

We analyze a class of in…nitely repeated prisoners' dilemma games with imperfect private monitoring and discounting. The main contribution of this paper is to construct "belief-based" strategies, where a player's continuation strategy is a function only of her beliefs about her opponent's continuatioin strategy. This simpli…es the analysis considerably — in the two-player case, we explicitly construct sequential equilibria, enabling us to invoke the one-step deviation

[¤]This paper combines independent papers by each of the authors. We are grateful to Je¤ Ely, Johannes Horner, Michihiro Kandori, Stephen Morris, Nicolas Persico and Tadashi Sekiguchi for their comments on the earlier versions of this paper. The …rst author thanks the University of Tokyo for its hospitality while part of this work was done. The second author is grateful to George Mailath and Andrew Postlewaite for their advice and encouragement since the very early stage of this research.

[y]Email: vbhas@essex.ac.uk

[z]Email: obara@ssc.upenn.edu

principle of dynamic programming. By doing so, we prove that one can approximate the symmetric efficient payoff in any prisoners' dilemma game provided that the monitoring is sufficiently accurate. Furthermore, for a class of prisoners' dilemma games, one can approximate every individually rational feasible payoff. Our efficiency results also generalize to the n player case, where we show that the symmetric efficient payoff can similarly be approximated.

These results are closely related to an important paper by Sekiguchi [13], who shows that one can approximate the efficient payoff in two-player prisoners' dilemma games provided that the monitoring is sufficiently accurate. Sekiguchi's result applies for a class of prisoners' dilemma payoffs, and relied on the construction of a Nash equilibrium which achieves approximate efficiency. The results in this paper can be viewed as an extension and generalization of the approach taken in Sekiguchi's paper.

Our substantive results are also related to those obtained in recent papers by Piccione [12] and Ely and Välimäki [5], which adopt a very different approach. The current paper (and Sekiguchi's) utilizes initial randomization to ensure that a player's beliefs adjust so that she has the incentive to punish or reward her opponent(s) as is appropriate. The Piccione-Ely-Välimäki approach on the other hand relies on making each player indifferent between her different actions at most information sets, so that her beliefs do not matter[1]. We defer a more detailed discussion of the two approaches to the concluding section of this paper.

The rest of this paper is as follows. Section 2 constructs sequential equilibria which approximate the efficient outcome in the two-player case, while section 3 approximates the set of individually rational feasible payoffs in this case. Section 4 shows that the efficiency result can be generalized to n player prisoners' dilemma games. The final section concludes.

2 Approximating the Efficient Payoff

	C	D
C	1	$-l$
D	$1+g$	0

We consider the prisoners' dilemma with the stage game payoffs given above, where the row indicates the player's own action and the column indicates her opponent's action. Players only observe their own actions, and also observe a private signal which is informative about their opponent's action. This signal belongs to the set $\Omega = \{c, d\}$, where c (resp. d) is more likely when the opponent plays C (resp. D). The signalling structure is assumed symmetric, in the sense that the probability of errors does not depend on the action profile played. Given any action profile $a = (a_1, a_2)$; $a_i \in A = \{C, D\}$; the probability that exactly one player receives a wrong signal is $\varepsilon > 0$, and the probability that both receive

[1] Obara [11] found the same kind of strategy independently, but used it for repeated games with imperfect public monitoring to construct a sequential equilibrium which pareto-dominates perfect public equilibria in simple repeated partnership games.

wrong signals is » > 0: Players maximize the expected sum of stage game payo¤s discounted at rate ±: We also assume that at the end of each period, players observe the realization of a public randomization device uniformly distributed on the unit interval.

Our approach is closely related to Sekiguchi's [13]: we show that one can construct a mixed trigger strategy sequential equilibrium which achieves partial cooperation. With public randomization or by "dividing up the game" as in Ellison [4], one can modify the strategy appropriately in order to approximate full cooperation. Our approach involves the construction of a "belief-based" strategy, i.e. a strategy which is a function of the player's beliefs about his opponent's continuation strategy. This results in a major simpli…cation as compared to the more conventional notion of a strategy which is a function of the private information of the player.

We begin by de…ning partial continuation strategies. In any period t; de…ne the partial continuation strategy ¾$_D$ as follows: play D at period t; and at period t + 1 play ¾$_D$ if the realized outcomes in period t are (Dc) or (Dd): De…ne the partial continuation strategy ¾$_C$ as follows: in any period t play C; at period t + 1 play ¾$_C$ if the realized outcomes in period t is (Cc); and play ¾$_D$ if the realized outcome at t is (Cd): We call ¾$_C$ and ¾$_D$ a partial continuation strategy since each of these fully speci…es the player's actions in every subsequent period at every information set that arises given that he con…rms to the strategy: In consequence, the (random) path and payo¤s induced by any pair of partial continuation strategies is well de…ned. However, a partial continuation strategy does not specify the player's actions in the event that she deviates from the strategy at some information set. This is deliberate, since our purpose is to construct the full strategies that constitute a sequential equilibrium. Note also that for any player i; only the partial continuation strategy of player j is relevant when computing i's payo¤s in any equilibrium.

Let V_{ab}(±; "; »); a; b 2 fC; Dg denote the repeated game payo¤ of ¾$_a$ against ¾$_b$ — these payo¤s are well de…ned since the path induced by each pair is well de…ned. We have that V_{DD} > V_{CD}; for all parameter values. Furthermore, if ± > $\frac{g}{1+g}$; and (" + ») is su¢ciently small, then V_{CC} > V_{DC}: Suppose that player i believes that her opponent is playing either ¾$_C$ or ¾$_D$; and is playing ¾$_C$ with probability ¹: Then the di¤erence between the payo¤ from playing ¾$_C$ and the payo¤ from playing ¾$_D$ is given by

$$¢V(¹; ±; "; ») = ¹(V_{CC} \text{ ¡ } V_{DC}) \text{ ¡ } (1 \text{ ¡ } ¹)(V_{DD} \text{ ¡ } V_{CD}) \qquad (1)$$

Hence ¢V(¹) is increasing and linear in ¹ and there is a unique value, ¼(±; "; »); at which it is zero. Suppose now that at t = 1 both players are restricted to choosing between ¾$_C$ and ¾$_D$: There is a mixed equilibrium of the restricted game, where each player plays the strategy ¾ which plays ¾$_D$ with probability 1 ¡ ¼ and ¾$_C$ with probability ¼: Call this partial mixed strategy ¾: Note that ¼(±; "; ») increases to 1 as we decrease ± towards its lower bound $\frac{g}{1+g}$: Let ± be such that ¼ > $\frac{1}{2}$:

For future reference we emphasize that equation (1) applies to any period — if a player believes that her opponent's continuation strategy is σ_C with probability μ and σ_D with probability $1 - \mu$, then she prefers σ_C to σ_D if $\mu > \pi$ and prefers σ_D to σ_C if $\mu < \pi$. Note also that if a player's opponent begins at $t = 1$ with a strategy in $\{\sigma_C, \sigma_D\}$, her continuation strategy also belongs to this set, since σ_D induces only σ_D, while σ_C may induce either σ_C or σ_D, depending upon the private history that the opponent has observed.

We define the following four belief revision operators. Starting with any initial belief μ, we can define a player's new beliefs when she takes action a and receives signal ω. Her new belief (i.e. the probability that j's continuation strategy is σ_C) will be given by $\chi_{a\omega}(\mu)$. We have four belief operators, χ_{Cc}, χ_{Cd}, χ_{Dc}, χ_{Dd}, where each $\chi_{a\omega} : [0,1] \to [0,1]$ is defined, using Bayes rule, as follows

$$\chi_{Cc}(\mu) = \frac{\mu(1 - 2\varepsilon - \xi)}{\mu(1 - \varepsilon - \xi) + (\varepsilon + \xi)(1 - \mu)} \tag{2}$$

$$\chi_{Cd}(\mu) = \frac{\mu\varepsilon}{\mu(\varepsilon + \xi) + (1 - \varepsilon - \xi)(1 - \mu)} \tag{3}$$

$$\chi_{Dc}(\mu) = \frac{\mu\varepsilon}{\mu(1 - \varepsilon - \xi) + (\varepsilon + \xi)(1 - \mu)} \tag{4}$$

$$\chi_{Dd}(\mu) = \frac{\mu\xi}{\mu(\varepsilon + \xi) + (1 - \varepsilon - \xi)(1 - \mu)} \tag{5}$$

Starting with any initial belief p at the beginning of the game, a player's belief at any private history, i.e. after an arbitrary sequence $\langle (a\omega)_r \rangle_{r=1}^{t}$, can be computed by iterated application of the appropriate belief operators. Let $\Xi(p)$ be the set of possible beliefs, i.e. $\mu \in \Xi(p) \iff \exists \langle \mu_r \rangle_{r=1}^{t} : \mu_1 = p$, $\mu_t = \mu$ and $\mu_{r+1} = \chi_{(a\omega)_r}(\mu_r)$, $(a\omega)_r \in A \times \Omega$, $1 \leq r \leq t - 1$. Let τ be a (full) strategy, which is defined at every information set, i.e. after arbitrary private histories. Clearly, τ is a best response to σ after every private history if and only if it is optimal to play τ at every belief $\mu \in \Xi(p)$, i.e. at all possible beliefs given the initial belief p.

We now examine the properties of these belief operators. First, each is a strictly increasing function. The operator χ_{Cc} has an interior fixed point at θ, and $\chi_{Cc}(\mu) \gtrless \mu$ as $\mu \gtrless \theta$. The value of θ depends upon (ε, ξ) in the following way

$$\theta(\varepsilon, \xi) = \frac{1 - 3\varepsilon - 2\xi}{1 - 2\varepsilon - 2\xi} \tag{6}$$

We shall assume that $\max\{\varepsilon, \xi\} < \frac{1}{6}$, which in turn entails that $\theta > \frac{1}{2}$.

We now show $\chi_{a\omega}(\mu) < \mu$ for each of the other three belief operators, χ_{Cd}, χ_{Dc} and χ_{Dd}, provided that $\max\{\varepsilon, \xi\} < \frac{1}{6}$. To verify this, take any typical expression from (3)-(5), and divide by μ. This yields ε (or ξ) in the numerator, while the denominator is strictly larger since it is a convex combination of $(\varepsilon + \xi)$ and $(1 - \varepsilon - \xi)$.

Since $\mu < \theta \Rightarrow \chi_{a\omega}(\mu) < \theta$ for any belief operator, this immediately implies that if $\pi < \theta$, $\Xi(\pi) \subseteq [0, \theta]$. This follows from the fact that the initial belief p is strictly less than θ, and since we have demonstrated that no point $\mu' > \theta$ is the image of any $\mu \leq \theta$ under any belief operator.

Hence, provided that initial beliefs are given by $\pi < \theta$, it suffices to define our belief based strategy for beliefs in the set $[0, \theta]$. Let $\rho : [0, \theta] \to \{C, D, \pi\}$ be defined as follows: $\rho(\mu) = C$ if $\mu \in (\pi, \theta]$ and $\rho(\mu) = D$ if $\mu \in [0, \pi)$. If $\mu = \pi, \rho(\mu) = \pi$, i.e. ρ plays C with probability π and D with probability $1 - \pi$. Hence the pair (ρ, π), i.e. ρ in conjunction with an initial belief $\pi \leq \theta$, specifies an action at every possible belief, and hence a complete strategy.

The advantage of this specification is that a player's continuation strategy is specified even at information sets which arise due to a player's deviating from ρ in the past. The belief based strategy (π, ρ) is realization equivalent to the partial strategy σ if it induces the same probability distribution over actions at every private history. This reduces to the following condition:

Definition 1 (π, ρ) is realization equivalent to σ if $\mu \in [\pi, \theta] \Rightarrow [\chi_{Cc}(\mu) > \pi$ and $\chi_{Cd}(\mu) < \pi]$ and $\mu \in [0, \pi] \Rightarrow [\chi_{Dc}(\mu) < \pi$ and $\chi_{Dd}(\mu) < \pi]$.

Lemma 2 If $\frac{1}{2} < \pi < \theta(\varepsilon, \xi)$, (ρ, π) is realization equivalent to σ.

Proof. To verify that $\mu \in [\pi, \theta] \Rightarrow \chi_{Cc}(\mu) > \pi$, recall that $\chi_{Cc}(\mu) > \mu$ if $\mu < \theta$, so that $\chi^k_{Cc}(\pi) > \pi$ for any k. To verify $\mu \in [\pi, \theta] \Rightarrow \chi_{Cd}(\mu) < \pi$, it suffices to verify that $\chi_{Cd}(\theta) \leq \pi$, since χ_{Cd} is strictly increasing.

$$\chi_{Cd}(\theta) = \frac{\varepsilon(1 - 3\varepsilon - 2\xi)}{(\varepsilon + \xi)(1 - 3\varepsilon - 2\xi) + (1 - \varepsilon - \xi)\varepsilon} < \frac{\varepsilon}{2\varepsilon} = \frac{1}{2} < \pi \qquad (7)$$

$\mu \leq \pi \Rightarrow \chi_{Dc}(\mu) < \pi$ and $\mu \leq \pi \Rightarrow \chi_{Dd}(\mu) < \pi$ follow from the fact already established that χ_{Dc} and χ_{Dd} lie below the 45^0 line. ■

Note that if ε and ξ are sufficiently small, we can select $\delta > \frac{g}{1+g}$ so that $\pi(\delta, \varepsilon, \xi) \in (\frac{1}{2}, \theta)$ — this follows from the fact that $\pi(\delta, \varepsilon, \xi) \to 1$ as $\delta \to \frac{g}{1+g}$ and $(\varepsilon + \xi) \to 0$, while $\pi(\delta, \varepsilon, \xi) \to 0$ if $\delta \to 1$ and $(\varepsilon + \xi) \to 0$. Henceforth we shall assume that δ is such that $\pi \in (\frac{1}{2}, \theta)$ so that (ρ, π) is consistent.

We have therefore established that the pair (ρ, π) defines a full strategy which is behaviorally equivalent to σ.

Proposition 3 If $\frac{1}{2} < \pi < \theta(\varepsilon, \xi)$, the strategy profile where each player plays (ρ, π) is a sequential equilibrium.

Proof. Note first that if $\mu = \pi$, a player is indifferent between playing σ_C and σ_D, and hence a one-step deviation from playing ρ is not profitable. Since the payoffs from playing π, σ_C and σ_D are equal at belief π, one may also, for the purposes of computing payoffs, use σ_C or σ_D as is computationally convenient in the event of belief π.

Consider first the case when $\mu > \pi$. A one-step deviation from ρ is to play D, and to continue with ρ in the next period. The following sub-cases arise:

a) Suppose that $\psi_{Dc}(\mu) \leq \pi$ and $\psi_{Dd}(\mu) \leq \pi$. In this case, a one-step deviation from ρ is to play σ_D, whereas $\rho(\mu) = \sigma_C$. However, (1) establishes that in this case σ_C is preferable to σ_D, and hence a one-step deviation from ρ is unprofitable.

b) Suppose that $\psi_{Dc}(\mu) \leq \pi$ and $\psi_{Dd}(\mu) > \pi$, so that the one-step deviation is to play D today and continue with σ_D if Dc is reached, and to continue with σ_C if Dd is reached. Let $\Delta\bar{V}(\mu)$ be payoff difference between the equilibrium strategy and the one-step deviation. Note that the one step deviation differs from σ_D only at the information set Dd; at this information it continues by playing σ_C whereas σ_D continues with σ_D. Hence we can write $\Delta\bar{V}(\mu)$ as the payoff difference between σ_C and σ_D minus the payoff difference between σ_C and σ_D conditional on Dd being reached, as follows:

$$\Delta\bar{V}(\mu) = \Delta V(\mu) - \delta[\mu(\varepsilon + \lambda) + (1 - \mu)(1 - \varepsilon - \lambda)][\Delta V(\psi_{Dd}(\mu))] \qquad (8)$$

Note that $\psi_{Dd}(\mu) < \mu$. Equation (1) shows that this implies that $\Delta V(\mu) > \Delta V(\psi_{Dd}(\mu))$. Since the coefficient multiplying $\Delta V(\psi_{Dd}(\mu))$ is strictly less than one, this implies that $\Delta\bar{V}(\mu) > 0$. Hence if $\mu > \pi$, a one-step deviation is unprofitable.

c) Finally, we establish that $\psi_{Dc}(\mu) < \pi \; \forall \mu \leq \theta$, so that no other sub-case need be considered. Evaluating ψ_{Dc} at the upper bound θ, we have

$$\psi_{Dc}(\theta) = \frac{\varepsilon(1 - 3\varepsilon - 2\lambda)}{(1 - \varepsilon - \lambda)(1 - 3\varepsilon - 2\lambda) + (\varepsilon + \lambda)\varepsilon} < \frac{\varepsilon}{1 - \varepsilon - \lambda} < \frac{1}{2} \qquad (9)$$

where the last step follows from the assumption that $\max\{\varepsilon, \lambda\} < \frac{1}{6}$.

Consider now the case when $\mu < \pi$. In this case, a one-step deviation from ρ is to play C today, and to continue with σ_C if $\psi_{Cc}(\mu) \geq \pi$, but to continue with σ_D if $\psi_{Cc}(\mu) < \pi$. (Note that $\mu < \pi \Rightarrow \psi_{Cd}(\mu) < \pi$, so the continuation strategies do not differ in this event.) In the first sub-case, the one-step deviation from ρ corresponds to playing σ_C, and (1) establishes that in this case σ_D is preferable to σ_C, and hence a one-step deviation from ρ is unprofitable. In the second sub-case, the one-step deviation differs from σ_C only at the information set Cc — it plays σ_D at this information set rather than σ_C. Let $\Delta\hat{V}(\mu)$ denote the payoff difference between the one-step deviation and the equilibrium strategy σ_D. We have

$$\Delta\hat{V}(\mu) = \Delta V(\mu) - \delta[\mu(1 - \varepsilon - \lambda) + (1 - \mu)(\varepsilon + \lambda)][\Delta V(\psi_{Cc}(\mu))] \qquad (10)$$

Since $\pi > \chi_{Cc}(\mu) > \mu$, $\Delta V(\mu) < \Delta V(\chi_{Cc}(\mu)) < 0$. Also, the coefficient multiplying $\Delta V(\chi_{Cc}(\mu))$ is less than 1 which establishes that $\Delta \hat{V}(\mu) > 0$.

We have therefore established that if a player's opponent j plays the strategy σ (which randomizes between σ_C and σ_D), it is optimal for player i to play ρ, with initial belief π. However, (ρ, π) is consistent and behaviorally equivalent to the strategy σ. Hence the profile where both players play (ρ, π) is a sequential equilibrium. ■

Under what conditions is there a pure strategy sequential equilibrium where both players begin in period one by playing σ_C with probability one. The above analysis also permits an answer to this question, with the difference that the initial belief $\mu = 1$ rather than π. To ensure that it is always optimal to cooperate after receiving good signals, we require that $\chi^k_{Cc}(1) > \pi \ \forall k$, which will be satisfied as long as $\pi \leq \theta$. Additionally, it must be optimal to switch to playing D on receiving a bad signal, i.e. we must have $\chi_{Cd}(\mu) < \pi$ for $\mu = 1$ or $\mu = \chi^k_{Cc}(1)$ for some k. Hence it is necessary and sufficient that[2]

$$\chi_{Cd}(1) = \frac{\varepsilon}{\varepsilon + \xi} \leq \pi < \theta \tag{11}$$

This requires that the χ_{Cd} function always lies below θ, which requires the inequality

$$\varepsilon^2 < \xi(1 - 3\varepsilon - 2\xi) \tag{12}$$

This inequality will be satisfied if ε is sufficiently small relative to ξ, i.e. if signals are sufficiently "positively" correlated. It is easily verified that this inequality cannot be satisfied if signals are independent or "negatively" correlated so that the equilibrium must be in mixed strategies.

Note that π plays a dual role in the construction of the mixed strategy equilibrium. On the one hand it is the randomization probability in the first period, and on the other hand, it is simply a number which defines the threshold at which behavior changes. These roles are obviously distinct, as is apparent from our discussion of the pure strategy equilibrium. This distinction is particularly relevant when we discuss the folk theorem in the following section.

With the construction of the mixed equilibrium, one can approximate full cooperation by using one of two devices. If a public randomization device is available, then it is immediate that the equilibrium payoff set is monotonically increasing (in the sense of set inclusion) in δ — given any $\delta' > \delta$, players may simply re-start the game with probability $m = \frac{\delta}{\delta'}$. In the absence of such public randomization, one may use the construction introduced by Ellison [4] (see also, Sekiguchi [13]), of dividing the game into n separate repeated games, thereby reducing the discount factor.

[2] The conditions for the optimality of playing D once a player has played D are as before, and hence will also be satisfied.

Lemma 4 Let $\delta_0 < \delta_1 < 1$; and let there be a symmetric strategy profile which is a sequential equilibrium of the repeated game for any $\delta \in (\delta_0; \delta_1)$; yielding payoff $v(\delta) \geq v$ for any $\delta \in (\delta_0; \delta_1)$: There exists $\hat{\delta} < 1$ such that the repeated game has a symmetric sequential equilibrium with payoff greater than v for any $\delta \geq \hat{\delta}$. If a public randomization device is available and $(v_1; v_2)$ is a sequential equilibrium payoff for some $\delta \in (0; 1)$; it is also an equilibrium payoff for any $\delta' > \delta$:

Proof. For the proof of the first part of this lemma, see Ellison [4]. To prove the second part, let σ be the strategy profile giving the required payoff given δ: Given δ'; let $m = \frac{\delta}{\delta'}$. Players play a sequence of games: they begin with the strategy profile σ: If the sunspot in any period $\xi > m$; they play a new game and re-start with σ: ■

Proposition 5 For any $x < 1$; there exists a symmetric sequential equilibrium with payoff greater than x if ε and η are sufficiently small, provided that either (i) δ is sufficiently close to 1 or (ii) $\delta > \frac{g}{1+g}$ and a public randomization device is available.

Proof. Proposition 3 implies that if $(\varepsilon; \eta)$ are sufficiently small, so that $\theta(\varepsilon; \eta)$ is close to 1; we have an open interval of values of π such that $(\pi; \rho)$ is a sequential equilibrium. In this range, $\pi(\delta; \varepsilon; \eta)$ is a strictly decreasing function of δ; and hence if $(\varepsilon; \eta)$ are sufficiently small, there is an open interval of values of δ such that $(\pi(\delta; \varepsilon; \eta); \rho)$ is a sequential equilibrium. Since $(\varepsilon; \eta)$ are close to zero, we can select this interval of values of π close to 1; so that the payoff in any such equilibrium is greater than x: Part (i) of the proposition then follows from the first part of lemma 4. If a public randomization device is available, let $(\varepsilon; \eta) \to (0; 0)$ and $\delta(\varepsilon; \eta) \to \frac{g}{(1+g)}$; so that $\pi \to 1$: The equilibrium payoff tends to one. Lemma 4 ensures that this result holds for all $\delta > \frac{g}{1+g}$: ■

Although it is common to allow for vanishing discounting in proving folk theorems in repeated games, it is worth pointing out that in order to obtain approximate efficiency, such vanishing discounting is not required if we have a public randomization device. In the absence of such a randomization device, one does require vanishing discounting, essentially due to an "integer" problem.

2.1 Generalizing the information structure.

We now show that the above construction also extends for a more general information structure. Let Ω be a common finite set of signals observed by the players, and assume that the marginal distribution of p has full support, i.e. for any action profile a; $P(\omega_i = \omega_i'|a) = \sum_{\omega_j} p(\omega_i'; \omega_j|a) > 0$ for all $\omega_i \in \Omega$; $i = 1; 2$: Our assumptions on the signal structure are as follows:

1. Assume that the set Ω can be partitioned into the set of good signals Ω_c

and the set of bad signals Ω_d, where the likelihood ratios satisfy

$$\frac{P(\omega_1 = c \mid (a_1, D))}{P(\omega_1 = c \mid (a_1, C))} < 1 < \frac{P(\omega_1 = d \mid (a_1, D))}{P(\omega_1 = d \mid (a_1, C))} \tag{13}$$

for every $c \in \Omega_c, d \in \Omega_d$ and for any action $a_1 \in \{C, D\}$ taken by player 1. Although we focus on player 1, the same conditions and results also hold for player 2.

2. Assume that $\Omega_C \times \Omega_C$ is $1 - \varepsilon$ evident given the action profile (CC), where ε is a small number. This ensures that if player 1 receives any good signal, and her prior belief assigns high probability to the action profile (CC), then player 1 assigns high probability to the event that her opponent has also received a good signal. Note that this assumption is consistent with signals being independent conditional upon the action profile — under independence, if $P(\omega_2 \in \Omega_c \mid CC) > 1 - \varepsilon$, then $\Omega_C \times \Omega_C$ is $1 - \varepsilon$ evident given (CC).[3]

3. Let $\min_d P(\omega_1 = d \mid (CD))$ and $\frac{\min_c P(\omega_1 = c, \omega_2 \in \Omega_D \mid (CC))}{\min_d P(\omega_1 = d, \omega_2 \in \Omega_C \mid (CC))}$ be bounded below by some number $\eta > 0$, independent of ε. Assume also that

$\max_c P(\omega_2 \in \Omega_c \mid (DC), \omega_1 = c) \leq 1 - \eta'$ where $\eta' > 0$, again independent of ε.

Note that assumptions 1-3 above are consistent with the signals being independent conditional upon the strategy profile. Also note that this information structure can be quite different from almost perfect monitoring.

We now show that under these assumptions, beliefs evolve in such a way that they are always above the initial belief π as long as a player plays C and receives good signals in Ω_c, but they fall below π whenever a player receives a bad signal and they continue to stay below π when a player plays D.

When player 1 plays C and receives some signal $c \in \Omega_c$, the updated belief is given by

$$\psi_{Cc}(\mu) = \frac{\mu P(\omega_1 = c, \omega_2 \in \Omega_C \mid (CC))}{\mu P(\omega_1 = c \mid (CC)) + (1 - \mu) P(\omega_1 = c \mid (CD))} \tag{14}$$

The fixed point of this mapping will be close to 1 for every $c \in \Omega_C$ if $\frac{P(\omega_1 = c, \omega_2 \in \Omega_C \mid (CC))}{P(\omega_1 = c \mid (CC))} = P(\omega_2 \in \Omega_C \mid (CC), \omega_1 = c)$ is large enough for any $c \in \Omega_C$. This is greater than $1 - \varepsilon$ since $\Omega_c \times \Omega_c$ is $1 - \varepsilon$ evident.

The second condition is that a player should switch to playing σ_D when he receives a bad signal. For each $c \in \Omega_c$, we can define the associated fixed point of the mapping (14), just as in (6). Let $\bar{\mu}$ denote the largest such fixed point (in the set Ω_c) and let $\underline{\mu}$ denote the smallest such fixed point. A sufficient condition for our construction is that $\psi_{Cd}(\bar{\mu}) < \underline{\mu}$ for any $d \in \Omega_d$. In this case one can

[3] Mailath and Morris [8] introduce and use such conditions on the signal structure. Given their focus on almost public monitoring (where signals are correlated), they also assume a similar condition for $\Omega_d \times \Omega_d$.

select $\pi \in (\chi_{Cd}(\bar{\mu}); \underline{\mu})$ so that switching to σ_D is optimal: Since $\underline{\mu} \to 1$ as $\varepsilon \to 0$; and $\bar{\mu} \leq \max_{c \in \Omega_c} \frac{P(\omega_1 = c; \omega_2 \in \Omega_c | (CC))}{P(\omega_1 = c | (CC))}$; we need that $\chi_{C;d}(\max_{c \in \Omega_c} \frac{P(\omega_1 = c; \omega_2 \in \Omega_c | (CC))}{P(\omega_1 = c | (CC))})$ is bounded away from 1 independent of ε: A straightforward application of the belief operators shows that:

$$\chi_{C;d}(\bar{\mu}) \leq \frac{1}{1+\eta^2} \tag{15}$$

for all $d \in \Omega_D$; which is bounded away from 1 by assumption 2 above.

Finally, we need to ensure that $\chi_{D\omega_1}(\mu)$ remains low, so that a player continues with playing σ_D: If $\omega_1 \in \Omega_D$; it is easy to verify that $\chi_{D\omega_1}(\mu) < \mu$. If $\omega_1 \in \Omega_c$; then $p(\omega_2 \in \Omega_C | (DC); \omega_1 = c)$ is less than $1 - \eta'$; and bounded away from 1 for any Ω_c by assumption 3: Hence it is optimal to continue with σ_D once a player starts playing D:

Summarizing the above arguments and checking sequential rationality, we have the following theorem:

Proposition 6 Given $\eta; \eta' > 0$; there exists $\varepsilon^* > 0$ such that for any $\varepsilon < \varepsilon^*$; our mixed trigger strategy ($\rho(\omega) = C$ if $\omega \in \Omega_C$ and $\rho(\omega) = D$ if $\omega \in \Omega_D$::::) is a sequential equilibrium. If $\varepsilon \to 0$ we can approximate the efficiency outcome provided that either i) $\delta \to 1$ or ii) $\delta > \frac{g}{1+g}$ and a public randomization device is available.

3 Approximating Any Individually Rational Feasible Payoff

We now build on the construction of the previous section and show how to approximate any individually rational feasible payoff. We shall consider a prisoners' dilemma game where the symmetric efficient payoff is given by the profile $(C; C)$ (rather than by a convex combination of $(C; D)$ and $(D; C)$); and we also assume that there are only two signals, c and d: We also assume in this section that a public randomization device is available. The key step is to approximate the payoff $(\frac{1+g+\ell}{1+\ell}; 0)$, which is player 1's maximal payoff within the set of individually rational and feasible payoffs. Since the payoff $(1; 1)$ has already been approximated in the previous section, and $(0; 0)$ is a stage game equilibrium payoff, one can then use public randomization to approximate any individually rational feasible payoff.

It might be useful to outline the basic construction and to explain the complications that arise. The basic idea in approximating the extremal asymmetric payoff is that play begins in the asymmetric phase where player 1 plays D and player 2 randomizes, playing C with a high probability, υ. This asymmetric phase continues or ends, depending upon the realization of a public randomization device. Thus player 1's per-period payoff in the asymmetric phase is approximately $1 + g$ while player 2's per-period payoff is approximately $-\ell$:

Since the latter is less than the individually rational payoff for player 2, he must be rewarded for playing C. To ensure this, when the asymmetric phase ends, both player's continuation strategies depend upon their private information. Player 1 continues with σ_C if he has observed the signal c in the last period (i.e. if his information is Dc) and continues with σ_D if she has observed d (i.e. if his information is Dd). This ensures that player 2 is rewarded for playing C in the asymmetric phase. Similarly, player 2 continues with σ_C if his private information is Cd, the information set which is most likely when he plays C, and continues with σ_D if his private information is Dd. Hence, if the noise is small, player 2's continuation payoff when the asymmetric phase ends is approximately 1 if he has played C in the previous period and approximately zero if he has played D. Hence if δ is large relative to l ($\delta > \frac{l}{1+l}$), we can, by choosing the value of the sunspot appropriately, make player 2 indifferent between C and D in the asymmetric phase. The payoffs in this equilibrium converge to $(\frac{1+g+l}{1+l}, 0)$ as the noise vanishes.

However, one must also verify that the players find it optimal to play σ_C and σ_D, as appropriate, at each information set after the end of the asymmetric phase. A complication arises here, as compared to the previous section, since player 1 does not randomize in the asymmetric phase, i.e. she plays D for sure. (Indeed, she cannot play C with positive probability, since in that case her payoff in the asymmetric phase is bounded above by 1 and hence cannot approximate $1+g$.)[4] Hence when player 2 receives the signal c, he knows that there has been at least one error in signals, and his beliefs about player 1's continuation strategy depend upon the relative probability of one (ε) versus two errors (η). In other words, his continuation strategy at the information sets Cc and Dc depends upon the correlation structure of signals. Since player 2's continuation strategy depends upon this correlation structure, this implies that player 1's beliefs also depend upon the correlation structure.

We adopt two alternative approaches to handle this problem. First, we show that if signals are positively correlated, so that the probability of two errors is at least as large as the probability of one error, then one can approximate the asymmetric payoff, without any restriction upon payoffs. Second, we show that one does not need such positive correlation of signals provided that one can choose δ so that $\pi(\delta, \varepsilon, \eta)$ sufficiently close to one. This result applies to any prisoners' dilemma game where $g \geq l$ — in any such game one can approximate the asymmetric payoff arbitrarily closely. However, this second approach does not work if $l > g$, since in this case one cannot have $\pi(\delta, \varepsilon, \eta) \to 1$. The reason for this is the for π to be close to 1, we must have $\delta \to \frac{g}{1+g}$. However, in the asymmetric phase, player 2 incurs a loss of l by playing C, whereas his continuation payoff gain is no more than 1. Hence player 2 will be willing to play C in the asymmetric phase only if $\delta > \frac{l}{1+l}$. Hence if $l > g$, one cannot have

[4] This argument is more general and implies that one cannot have a folk theorem in completely mixed strategies for stage games with non-degenerate payoffs. Let $\hat{v}_1$ be the supremum payoff of player 1 in any equilibrium where player 1 randomizes in every period at every information set. Since $\hat{v}_1 \leq (1-\delta)\min_{a_1}\{\max_{a_2} u_1(a_1, a_2)\} + \delta\hat{v}_1$, this implies $\hat{v}_1 \leq \min_{a_1}\{\max_{a_2} u_1(a_1, a_2)\}$.

π close to 1 since δ is bounded away from $\frac{g}{1+g}$.

We make the following assumption for this section:

Assumption A: Either A1: $\xi \geq \varepsilon$ or A2: $g \geq l$ and $\xi(1-\xi)(1-3\varepsilon-2\xi) > \varepsilon^3$.

Note that A1 is a relatively strong assumption that signals are positively correlated, but does not require any assumption on payoffs. On the other hand, A2 requires an assumption on payoffs but is a mild assumption about the relative probability of errors. It is always satisfied if signals are positively correlated, or independent. In the independent signal case, the left hand side is a term of order ε whereas the right hand side is a term of order ε^3. Hence A2 is satisfied even if signals are negatively correlated provided that they are not too highly so.

We now define the players' strategies more precisely. In any period $t-1$ in the asymmetric phase, player 1 plays D for sure, while player 2 randomizes between C and D, choosing C with a constant probability υ which is close to 1. At the end of period, players observe the realization, ϕ_{t-1}, of a sunspot which is uniformly distributed on $[0,1]$. If $\phi_{t-1} > 1-\lambda$, both players continue in the asymmetric phase for the next period. If $\phi_{t-1} \leq \lambda$, the asymmetric phase ends for both players, and is never reached again. In this case, players' continuation strategies (i.e. their states) depend upon the realization of their private information at date $t-1$ (i.e. players ignore their private information from previous dates). Let ν_{t-1} denote the player's private information realization at date $t-1$. Player 1 continues with σ_C if $\nu_{t-1} = Dc$; if $\nu_{t-1} = Dd$, she continues in period t with σ_D.[5] Player 2's continues with σ_C if $\nu_{t-1} = Cd$, and continues with σ_D if $\nu_{t-1} = Dd$. If $\nu_{t-1} \in \{Cc, Dc\}$, player 2 continues with σ_C if $\mu_2(\nu_{t-1}) > \pi(\delta, \varepsilon, \xi)$ and with σ_D if $\mu_2(\nu_{t-1}) \leq \pi(\delta, \varepsilon, \xi)$.

Our analysis proceeds as follows. First, we show that player 2 is willing to randomize in the asymmetric phase provided that λ is appropriately chosen, and that the payoffs associated with this class of equilibria tend to $(\frac{1+g+l}{1+l}, 0)$ as the noise vanishes. Subsequently, we shall demonstrate that all players are choosing optimally at every information set.

Write $W_2(D)$ for the payoff of player 2 in the asymmetric phase given that he plays D, and $W_2(C)$ for the payoff in the asymmetric phase from playing C. Since $W_2(D) = W_2(C) = W_2$, we have

$$W_2(D) = \delta(1-\lambda)W_2 + \delta\lambda V_2(D) \tag{16}$$

where $V_2(D)$ is the expected payoff to player 2 conditional on the fact that the asymmetric phase has ended and that he has played D. Similarly, letting

[5] We show that any strategy which plays C in the asymmetric phase is dominated, and hence we need not define precisely the optimal continuation strategy after playing C. The existence of an optimal continuation strategy follows from the same argument as in Sekiguchi [13]. Since player 1 never plays C in the asymmetric phase, his continuation after his own deviation does not affect player 2's incentives.

$V_2(C)$ be the expected payoff to 2 conditional on the fact that the asymmetric phase has ended and that he has played C; we have

$$W_2(C) = (1 - \delta)(-l) + \delta(1 - \lambda)W_2 + \delta\lambda V_2(C) \tag{17}$$

Clearly, $V_2(D) \to 0$ as $(\varepsilon; \xi) \to (0; 0)$. We now show that $V_2(C) \to 1$ as $(\varepsilon; \xi) \to (0; 0)$. Let $V_2(Cd)$ (resp. $V_2(Cc)$) denote the continuation payoff at the end of the asymmetric phase, conditional on Cd (resp. Cc). Since player 1 plays D for sure in the asymmetric phase, we have

$$V_2(C) = (1 - \varepsilon - \xi)V_2(Cd) + (\varepsilon + \xi)V_2(Cc) \tag{18}$$

Hence it suffices to establish that $V_2(Cd) \to 1$ as $(\varepsilon; \xi) \to (0; 0)$. Write $\mu_2(Cd)$ for the probability that player 1's continuation strategy is σ_C, given that $\nu_{t-1} = Cd$. Since $\mu_2(Cd) \geq \frac{1-2\varepsilon-\xi}{1-\varepsilon-\xi}$, $\mu_2(Cd) \to 1$ as $\varepsilon \to 0$. Hence from equation (1) $V_2(Cd) \to V_{CC}$, where $V_{CC} \to 1$ as $(\varepsilon; \xi) \to (0; 0)$.

Hence if $\varepsilon + \xi$ is sufficiently small and $\delta > \frac{l}{1+l}$, there exists a value of λ which equates $W_2(C)$ and $W_2(D)$. Further, as $(\varepsilon + \xi) \to 0$, $\lambda \to \frac{(1-\delta)l}{\delta}$, and player 2's payoff converges to zero.

If $\upsilon \to 1$, player 1's per-period payoff tends to $(1 + g)$ in the asymmetric phase, and 1 in the cooperative phase. By substituting for the limiting value of λ which is $\frac{(1-\delta)l}{\delta}$, we see that player 1's payoff converges to $\frac{1+g+l}{1+l}$. (We shall establish later that $\upsilon \to 1$).

We now verify that each player plays optimally at each information set in this equilibrium. In the asymmetric phase, this is so for player 2 by construction, since he is indifferent between C and D. It is easy to see that player 1 also plays optimally in the asymmetric phase, since she is choosing her one shot best response.[6]

Consider now the transition to the cooperative phase, i.e. the player's actions in the first period after the sunspot signals at the end of the asymmetric phase. Since players only condition on their private information in the previous period, we may focus on this alone. Player 1 has two possible information sets, (Dc) and (Dd), whereas player 2 has four possible information sets. Let $\mu_i(\nu)$ denote the probability assigned by player i to her opponent's continuation strategy being σ_C, given that i is at information set ν.

As in the previous section, we shall assume that $\max\{\varepsilon; \xi\} < \frac{1}{6}$. Furthermore, as in the previous section, we assume that $\pi(\delta; \varepsilon; \xi) \in (\frac{1}{2}; \theta)$ — this assumption on π does not imply any restrictions upon g or l. However, if we invoke the assumption $g \geq l$ in A2, then we may also choose π to be arbitrarily close to

[6] It is possible that playing C in the asymmetric phase increases player 1's continuation payoff in the cooperative phase. However, it is easy to see that such an increase can never offset the loss from playing C. A simple proof is as follows. If playing C in the asymmetric phase is optimal for 1, then playing C in every period in the asymmetric phase is also optimal. The overall payoff of this strategy is approximately 1 if the noise is small, whereas the payoff of player 1 in the equilibrium tends to $\frac{1+g+l}{1+l}$, which is strictly greater.

its upper bound. We shall also assume that $\upsilon \in (\pi, \theta(\varepsilon, \xi))$. Since $\theta \to 1$ as $(\varepsilon, \xi) \to (0, 0)$ we can also have $\upsilon \to 1$.

Consider first the beliefs of player 2: Let $\mu_2(.)$ denote the probability assigned by 2 to the event that 1's continuation strategy is σ_C. Since player 1 plays σ_C at Dc and σ_D at Dd, and since player 1 does not play C in the asymmetric phase, we have

$$\mu_2(Cd) = \frac{1 - 2\varepsilon - \xi}{1 - \varepsilon - \xi} > \theta \tag{19}$$

Since $\pi < \theta$, it is optimal to continue with σ_C today at information set Cd. Further, we have

$$\chi_{Cd}(\mu_2(Cd)) = \frac{(1 - 2\varepsilon - \xi)\varepsilon}{(1 - 2\varepsilon - \xi)(\varepsilon + \xi) + (1 - \varepsilon - \xi)\varepsilon} < \frac{1}{2} \tag{20}$$

Hence it is optimal for player 2 to switch to the defection phase if he receives the signal Cd at any date in the future.

At Dd, we have

$$\mu_2(Dd) = \frac{\varepsilon}{1 - \varepsilon - \xi} \tag{21}$$

This is clearly less than $\frac{1}{2}$ since $\max\{\varepsilon, \xi\} < \frac{1}{6}$, so that it is optimal to continue with σ_D.

Consider now the beliefs of player 2 at (Cc) and (Dc), i.e. at the information sets where player 2 knows that there has been at least one error in the signals.

$$\mu_2(Dc) = \frac{\xi}{\xi + \varepsilon} \tag{22}$$

$$\mu_2(Cc) = \frac{\varepsilon}{\xi + \varepsilon} \tag{23}$$

Recall that player 2 plays σ_D at least at one of these information sets, since the above probabilities cannot be both greater than π, since this is greater than one-half. Hence there are three possibilities: either both $\mu_2(Dc)$ and $\mu_2(Cc)$ are less than π, or exactly one of these is greater than π. Now if $\mu_2(.) < \pi$ at any information set, it is optimal to continue with σ_D today, and at every future date. Hence it remains to verify the case when $\mu_2(.) \geq \pi$.

Suppose that $\frac{\xi}{\xi+\varepsilon} > \pi$, so that player 2 plays σ_C at Dc. If $\frac{\xi}{\xi+\varepsilon} \leq \theta$, lemma 2 verifies that it is optimal to continue with σ_C in this case. Hence consider the case where $\frac{\xi}{\xi+\varepsilon} > \theta$. We have that $\mu > \theta \Rightarrow \chi_{Cc}(\mu) < \mu$. Further, since χ_{Cd} is an increasing function, it suffices to verify that $\chi_{Cd}(\frac{\xi}{\xi+\varepsilon}) < \pi$, since this implies that $\chi_{Cd}(\mu) < \pi$ for $\mu = \chi^k_{Cc}(\frac{\xi}{\xi+\varepsilon})$ for any k.

$$\hat{A}_{Cd}\left(\frac{»}{" + »}\right) = \frac{"»}{»(" + ») + "(1 \text{ ¡ } " \text{ ¡ } »)} \qquad (24)$$

This is less than $\frac{1}{2}$ if max f"; »g is less than $\frac{1}{6}$: Hence player 2's continuation strategy is optimal at Dc:

Finally, we consider the case where that player 2 plays ¾$_C$ at Cc; i.e. when $\frac{"}{"+»}$ > ¼: Note that in this case A1 is violated. Hence we assume A2, which ensures that we can make ¼ arbitrarily close to its upper bound µ by selecting ± su¢ciently close to $\frac{g}{1+g}$: We can …nd a value of ¼ such that $\hat{A}_{Cd}(\frac{"}{"+»})$ < ¼ provided that $\hat{A}_{Cd}(\frac{"}{"+»})$ is less than the upper bound for ¼; i.e.

$$\hat{A}_{Cd}\left(\frac{"}{" + »}\right) = \frac{"^2}{"^2 + » \text{ ¡ } »^2} < µ \qquad (25)$$

It is easily veri…ed that the inequality above is ensured by condition A2.

Consider now the beliefs of player 1: Her beliefs will depend upon player 2's strategy, which in turn depends upon the parameters of the signal distribution, and as we have seen, there are three possible cases.

Consider …rst the case where 2 plays ¾$_C$ only at information set Cd:

$$¹_1(Dc) = \frac{À(1 \text{ ¡ } 2" \text{ ¡ } »)}{À(1 \text{ ¡ } " \text{ ¡ } ») + (1 \text{ ¡ } À)(" + »)} \qquad (26)$$

Note that the expression is such that ¹$_1$(Dc) = $\hat{A}_{Cc}$(À), where $\hat{A}_{Cc}$ is the belief revision operator de…ned in the previous section. Hence it follows that if À 2 [¼; µ); it follows that $\hat{A}^k_{Cc}$(À) 2 (¼; µ); 8k; and hence it is optimal for player 1 to continue with ¾$_C$ at every information set.

Consider 1's beliefs at (Dd): Once again, it is easy to verify that ¹$_1$(Dd) = $\hat{A}_{Cd}$(À); and since À < µ; it is optimal to continue with ¾$_D$ at this information set.

Consider next the case where ¾$_2$(Cd) = ¾$_2$(Cc) = ¾$_C$ and ¾$_2$(Dd) = ¾$_2$(Dc) = ¾$_D$: In this case, assumption A2 applies, so that we may choose ¼ close to its upper bound. We have

$$¹_1(Dc) = \frac{À(1 \text{ ¡ } " \text{ ¡ } »)}{À(1 \text{ ¡ } " \text{ ¡ } ») + (1 \text{ ¡ } À)(" + »)} \qquad (27)$$

If À > ¼; then ¹$_1$(Dc) > ¼ so that it is optimal to start by playing ¾$_C$ in this case. To see that player 1 will …nd it optimal to switch to ¾$_D$ on receiving a bad signal, note that requires

$$\hat{A}_{Cd}(¹_1(Dc)) = \frac{À"}{" + »} < ¼ \qquad (28)$$

Now if $\mu_1(Dc) \leq \theta$, lemma 2 has verified that a player who begins with σ_C will switch to σ_D on receiving signal Cd in any subsequent period. If $\mu_1(Dc) > \theta$, it suffices to verify that $\chi_{Cd}(\mu_1(Dc)) < \theta$, which is the upper bound for π. This yields the condition

$$\upsilon < \frac{(\varepsilon + \xi)\theta}{\varepsilon} \tag{29}$$

Since $\upsilon < \theta$, this condition is also satisfied.

Finally, we consider the case where $\sigma_2(Cd) = \sigma_2(Dc) = \sigma_C$ and $\sigma_2(Dd) = \sigma_2(Cc) = \sigma_D$.

$$\mu_1(Dc) = \frac{\upsilon(1 - 2\varepsilon - \xi) + (1 - \upsilon)\xi}{\upsilon(1 - \varepsilon - \xi) + (1 - \upsilon)(\varepsilon + \xi)} < \frac{1 - 2\varepsilon - \xi}{1 - \varepsilon - \xi} \tag{30}$$

Hence it suffices to evaluate χ_{Cd} at the upper bound, which yields

$$\chi_{Cd}\left(\frac{1 - 2\varepsilon - \xi}{1 - \varepsilon - \xi}\right) = \frac{(1 - 2\varepsilon - \xi)\varepsilon}{(1 - 2\varepsilon - \xi)(\varepsilon + \xi) + (1 - \varepsilon - \xi)\varepsilon} < \frac{1}{2} \tag{31}$$

Hence $\chi_{Cd}(\mu_1(Dc)) < \frac{1}{2}$ for every value of υ.

We have therefore proved that the payoff $(\frac{1+g+l}{1+l}, 0)$ (and obviously the payoff $(0, \frac{1+g+l}{1+l})$) can be approximated under assumption A provided that $\delta > \max\{\frac{g}{1+g}, \frac{l}{1+l}\}$ and provided that ε and ξ are sufficiently small. The payoff $(1, 1)$ has been approximated under a weaker set of assumptions ($\delta > \frac{g}{1+g}$ and ε and ξ sufficiently small), and the payoff $(0, 0)$ is a static Nash payoff. Since any payoff individually rational feasible payoff is a convex combination of these payoffs, and can be achieved via public randomization, we have proved the following theorem.

Theorem 7 Assume that Assumption A is satisfied, and players observe a public randomization device, then for any individually rational feasible payoff vector $u = (u_1, u_2)$ and any number $\zeta > 0$, there exist $\varepsilon(\zeta) > 0, \xi(\zeta) > 0$ such that there exists a sequential equilibrium with payoffs within ζ distance of u provided that $\varepsilon < \varepsilon(\zeta)$ and $\xi < \xi(\zeta)$ and $\delta > \max\{\frac{g}{1+g}, \frac{l}{1+l}\}$.

This result is most closely related to those obtained in a paper by Piccione [12], who also analyzes the prisoners' dilemma with imperfect private monitoring. Our results differ, both in terms of substance and in the techniques/strategies used. Piccione's substantive results are that full cooperation can always be approximated, and further, any individually rational feasible payoff can be approximated in a class of prisoners' dilemma games, i.e. for games where $l \geq g$. The "folk theorem" condition A in the present paper is, in a sense, the opposite of Piccione's condition. More recently, Ely and Välimäki [5] have considerably simplified the technique used in Piccione, and generalized the folk theorem obtained there. We shall discuss the differences between the approach of the present paper and the approach of Piccione and Ely-Välimäki in the concluding section.

4 The n-player case

In this section, we extend the approximate efficiency result to the n-player case. Let $N = \{1, 2, \ldots, n\}$ be the set of players and G be the stage game played by those players. The stage game G is as follows. Player i chooses an action a_i from the action set $A_i = \{C, D\}$. Actions are not observable to the other players and taken simultaneously. A n-tuple action profile is denoted by $a \in A = \prod_{i=1}^{n} A_i$. An action profile of all players but player i is $a_{-i} \in \prod_{j \neq i} A_j$.

Each player receives an $(n-1)$-tuple private signal profile about all the other players' actions. Let $\omega_i = (\omega_{i,1}, \ldots, \omega_{i,i-1}, \omega_{i,i+1}, \ldots, \omega_{i,n}) \in \{c, d\}^{n-1} = \Omega_i$ be a generic signal received by player i where $\omega_{i,j}$ stands for player i's signal about player j's action. A generic signal profile is denoted by $\omega = (\omega_1, \ldots, \omega_n) \in \Omega$. All players have the same payoff function u. Player i's payoff $u(a_i, \omega_i)$ depends on her own action a_i and private signal ω_i. Other players' actions affect a player i's payoff only through the distribution over the signal which player i receives. The distribution conditional on a is denoted by $p(\omega|a)$. It is assumed that $p(\omega|a)$ are full support distributions, that is, $p(\omega|a) > 0\ \forall a \forall \omega$. The space of a set of full support distributions $\{p(\omega|a)\}_{a \in A}$ is denoted by P.

We also introduce the perfectly informative signal distribution $P_0 = \{p_0(\omega|a)\}_{a \in A}$, where, for any $a \in A$, $p_0(\omega|a) = 1$ if $\omega_i = a_{-i}$ for all i. The whole space of the information structure $P \cup P_0$ is endowed with the Euclidean norm.

Since we are interested in the situation where information is almost perfect, we restrict attention mainly to a subset of P where information is almost perfect. Information is almost perfect when every person's signal profile is equal to the actual action profile with probability larger than $1 - \varepsilon$ for some small number ε.

To sum up, the space of the information structure we deal with in this section is the following subset of P:

$$P_\varepsilon = \left\{ \{p(\omega|a)\}_{a \in A} \in \mathbb{R}_{++}^{n \times (n-1) \times 2^n} \;\middle|\; \begin{array}{l} p(\omega|a) > 1 - \varepsilon \text{ if } \omega_{i,j} = c \Leftrightarrow a_j = C \text{ for all } i, j \\ \text{and } \forall a, \sum_{\omega} p(\omega|a) = 1 \end{array} \right\} \tag{32}$$

and we use p_ε for a generic element of P_ε.

A player's realized payoff only depends on the number of bad signals d that a player receives. Let $d(\omega_i)$ be the number of d contained in ω_i. Then, $u(a_i, \omega_i') = u(a_i, \omega_i'')$ if $d(\omega_i') = d(\omega_i'')$ for any a_i. Let $u\left(a_i, d^\theta\right)$ be the payoff of player i when $d(\omega_i) = \theta$. The deviation gain when θ defections are observed is $g(\theta) = u\left(D, d^\theta\right) - u\left(C, d^\theta\right)$, which is strictly positive for all θ. The largest deviation gain and the smallest deviation gain is $\overline{g}$ and $\underline{g}$ respectively, where $\overline{g} = \max_{0 \le k \le n-1} g(k)$ and $\underline{g} = \min_{0 \le k \le n-1} g(k)$.

We impose the following symmetry assumption on p_ε;

$$p(\omega|a) = p\left(\left(\omega_{\tau(i)\tau(j)}\right)_j \,\middle|\, \left(a_{\tau(1)},\ldots,a_{\tau(i)},\ldots,a_{\tau(N)}\right)\right) \quad (33)$$

for any permutation $\tau : N \to N$ and any $a \in A$.

This allows us to treat all the players in a symmetric way and to focus on only one player without loss of generality. Let $U_i\left(\left(a_i, D^\theta\right) : p\right)$ be the expected payoff of player i when θ players are playing D. The payoffs $U_i\left(\left(C, D^0\right) : p\right)$ and $U_i\left(\left(D, D^{n-1}\right) : p\right)$ are normalized to 1 and 0 respectively for all i. It is assumed that $(1,\ldots,1)$ is the symmetric efficient stage game payoff.

The stage game G is repeated infinitely by n players, who discount their payoffs with a common discount factor $\delta \in (0,1)$. Time is discrete and denoted by $t = 1,2,\ldots$. Player i's private history is $h_i^t = \left(\left(a_i^1, \omega_i^1\right),\ldots,\left(a_i^{t-1}, \omega_i^{t-1}\right)\right)$ for $t \geq 2$ and $h_i^1 = \emptyset$. Let H_i^t be the set of all such history h_i^t and $H_i = \bigcup_{t=1}^{\infty} H_i^t$. Player i's strategy is a sequence of mappings $\sigma_i = (\sigma_{i,1}, \sigma_{i,2}, \ldots)$, each $\sigma_{i,t}$ being a mapping from H_i^t to probability measures on A_i and denote $\sigma_i(h_i^t)$ for $\sigma_{i,t}(h_i^t)$. Discounted average payoff for player i is $V(\sigma : p, \delta) = (1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} E_p^{\sigma}\left[u\left(\left(a_i^t, \omega_i^t\right)\right)\right]$, where the probability measure on H_i^t is generated by (σ, p).

For this n-player repeated prisoner's dilemma, σ_C and σ_D are defined as the partial continuation strategies which are realization equivalent to the following grim trigger strategy and permanent defection respectively:

$$\sigma_C(h_i^t) = \begin{cases} C & \text{if } h_i^t = ((C,c),\ldots,(C,c)) \text{ or } t=1 \\ D & \text{otherwise} \end{cases}$$
$$\sigma_D(h_i^t) = D \text{ for all } h_i^t \in H_i$$

where $c = (C,\ldots,C)$.

This grim trigger strategy is the harshest one among all the variations of grim trigger strategies in the n player case. Players using σ_C switch to σ_D as soon as they observe any signal profile which is not fully cooperative. When player i is mixing σ_C and σ_D with probability $(\rho_i, 1-\rho_i)$, that strategy is denoted by $\rho_i\sigma_C + (1-\rho_i)\sigma_D$.

Suppose that either σ_C or σ_D is chosen in the first period by all players. Let $\theta \in \Theta$ be the number of players using σ_D as a continuation strategy among n players. Then a probability measure $\rho_{-i}(h_i^t; p)$ on the space $\Theta = \{0,1,\ldots,n-1\}$ is derived conditional on the realization of the private history h_i^t. Clearly, this measure also depends on the initial level of mixture between σ_C and σ_D by every player, but this dependence is not shown explicitly as it is obvious. Let U be the space of such probability measures, which is an $n-1$ dimensional simplex.

In the two player case, a player's strategy is represented as a function of belief, using the fact that the other player is always playing either σ_C or σ_D on and off the equilibrium path. Note that the space of the other players' "types" is much larger. However, there is a convenient way to summarize relevant information. We classify Θ into two sets; $\{0\}$ and $\{1,\ldots,n-1\}$; that is, the state no one have ever switched to σ_D and the state where there is at least one player who has already switched to σ_D. Player i's conditional subjective probability that no player has started using σ_D is denoted by $\phi\left(\mu_{-i}\right) = \mu_{-i}(0)$ given $\mu_{-i} \in U$. The reason why we just focus on this number is that the exact number of players who are playing σ_D does not make much difference to what will happen in the future given that everyone is playing σ_C. As soon as someone starts playing σ_D, every other player starts playing σ_D with very high probability from the very next period on by the assumption of almost perfect monitoring. What is important is not how many players have switched to σ_D, but whether anyone has switched to σ_D or not.

Finally, let $V(\sigma_i, \theta : p, \delta)$ be player i's discounted average payoff when θ other players are playing σ_D and $n - \theta - 1$ other players are playing σ_C. We need the following notations:

$$
\begin{aligned}
V\left(\sigma_i, \mu_{-i} : p, \delta\right) &= \sum_{\theta=0}^{n-1} V(\sigma_i, \theta : p, \delta)\, \mu_{-i}(\theta) && (34)\\
\Delta V(\theta : p_0, \delta) &= V(\sigma_C, \theta : p_0, \delta) - V(\sigma_D, \theta : p_0, \delta) && (35)\\
g\left(\mu_{-i}, p\right) &= \sum_{\theta=0}^{n-1} \left\{ u\left(\left(D, D^{\theta}\right) : p\right) - u\left(\left(C, D^{\theta}\right) : p\right)\right\} \mu_{-i}(\theta)
\end{aligned}
$$

4.1 Belief Dynamics and Best Response

For the two player case, the equilibrium strategy was described as a mapping from U to ΔA_i. The equilibrium strategy we will construct here has a similar structure except that belief lies in a larger space. It has the following expression:

$$
\rho\left(\mu_{-i}\right) = \begin{cases} C & \text{if } \mu_{-i} \in U^C \\ \pi & \text{if } \mu_{-i} \in U^I \\ D & \text{if } \mu_{-i} \in U^D \end{cases} \tag{36}
$$

where U^C, U^I, U^D are mutually exclusive and exhaustive sets in U, and π means playing C with probability π and playing D with probability $1 - \pi$.

In order to verify that ρ is a Nash equilibrium and achieves the approximate efficient payoff, we strengthen the path dominance argument used in Sekiguchi [13] instead of appealing to the one-shot deviation argument used in the previous sections. The argument is devided into several steps. First step is to give an almost complete characterization of the unique optimal action as a function of belief. As a next step, we analyze the dynamics of belief by introducing natural assumptions on the information structure when players are playing either σ_C or σ_D. The third step is to check consistency of this strategy profile, that is, to check if players are actually playing either σ_C or σ_D by following $\rho\left(\mu_{-i}\right)$.

Once it is established that ρ is a Nash equilibrium, then we can use the fact that there exists a sequential equilibrium which is realization equivalent to a

Nash equilibrium if the game is of non-observable deviation. Finally, as in the two player case, we can use a public randomization device or divide the original repeated game to component repeated games to implement the same payoff for large δ.

After we prove the existence of sequential equilibrium which is realization equivalent to $\rho\left(\mu_{-i}\right)$ and approximates the efficient outcome, we show, with one more assumption, that $\rho\left(\mu_{-i}\right)$ itself is actually a sequential equilibrium. The unique optimal action is indeed shown to have exactly the same form as $\rho\left(\mu_{-i}\right)$ for a certain range of δ if monitoring is almost perfect.

Before analyzing the unique optimal action, we first extend one property which holds in the two player case to the $n-$player case. In the two player case, the difference in payoffs by σ_C and σ_D is linear and there is a unique $\pi(\delta; \varepsilon; \xi)$ where a player is indifferent between σ_C and σ_D with perfect monitoring. When the number of players is more than two; the corresponding object $V\left(\sigma_C; \mu_{-i} : p_0; \delta\right) - V\left(\sigma_D; \mu_{-i} : p_0; \delta\right)$ is a slightly more complex. Even when players randomize independently and symmetrically playing σ_C with probability μ and σ_D with probability $(1-\mu)$, that is, $\mu_{-i}(\theta) = \sum_{\theta=0}^{n-1} (1-\mu)^{\theta} \mu^{n-1-\theta} \binom{n-1}{\theta}$ for $\theta = 0; \ldots; n-1$, it is a $n-1$ degree polynomial in $\mu \in [0;1]$. However, if $\delta > \frac{g(0)}{1+g(0)}$; it is possible to show that there exists a unique $\mu \in (0;1)$ such that $V\left(\sigma_C; \mu_{-i} : p_0; \delta\right) - V\left(\sigma_D; \mu_{-i} : p_0; \delta\right) = 0$.

Lemma 8 If $\delta > \frac{g(0)}{1+g(0)}$; there exists a unique $\mu^* \in (0;1)$ such that

$$\sum_{\theta=0}^{n-1} (1-\mu)^{\theta} \mu^{n-1-\theta} \binom{n-1}{\theta} \Delta V(\theta : p_0; \delta) = 0$$

Proof. Let $f(\mu) = \sum_{\theta=0}^{n-1} (1-\mu)^{\theta} \mu^{n-1-\theta} \binom{n-1}{\theta} \Delta V(\theta : p_0; \delta)$. Since $f(1) > 0$; $f(0) < 0$ and f is continuous, existence of such μ is guaranteed. To show uniqueness, we prove $f(\mu) = 0 \Rightarrow \frac{\partial f(\mu)}{\partial \mu} > 0$. Suppose that $f(\mu) = 0$. Then,

$$\frac{\partial f(\mu)}{\partial \mu} \tag{37}$$

$$= \sum_{\theta=0}^{n-1} \left\{ (1-\mu)^{\theta} (n-1-\theta) \mu^{n-2-\theta} - \theta (1-\mu)^{\theta-1} \mu^{n-1-\theta} \right\} \binom{n-1}{\theta} \Delta V(\theta : p_0; \delta) \tag{38}$$

$$= -\sum_{\theta=0}^{n-1} \theta (1-\mu)^{\theta} \mu^{n-1-\theta} \left(\frac{1}{\mu} + \frac{1}{1-\mu} \right) \binom{n-1}{\theta} \Delta V(\theta : p_0; \delta) \quad \text{(by } f(\mu) = 0\text{)}$$

Since $\Delta V(\theta : p_0; \delta) < 0$ for $\theta \geq 1$; $\frac{\partial f(\mu)}{\partial \mu} > 0$. ∎

Let $\pi(\delta; p_0)$ be this level of mixture where players are indifferent between σ_C and σ_D and monitoring is perfect, and denote the associated belief on Θ by μ^{π}_{-i}.

The following lemma extends a useful property in the two player case to the n player case.

Lemma 9 $\pi(\delta; p_0) \to 1$ as $\delta \downarrow \frac{g(0)}{1+g(0)}$

Proof. See Appendix

When monitoring is almost perfect, $V\left(\sigma_C; \mu_{-i} : p_\epsilon; \delta\right) - V\left(\sigma_D; \mu_{-i} : p_\epsilon; \delta\right)$ is very close to $V\left(\sigma_C; \mu_{-i} : p_0; \delta\right) - V\left(\sigma_D; \mu_{-i} : p_0; \delta\right)$. Actually, it is easy to see that the former converges to the latter uniformly in μ as $\epsilon \to 0$.[7] Hence, we can find $\pi(\delta; p_\epsilon)$ in the neighborfood of $\pi(\delta; p_0)$ when monitoring is almost perfect.

Now we derive the unique optimal action as a function of μ_{-i}. If monitoring is perfect, then $V\left(\sigma_C; \mu_{-i} : p_0; \delta\right) - V\left(\sigma_D; \mu_{-i} : p_0; \delta\right) > 0$ if and only if C is the unique optimal action for this belief μ_{-i}. We show that the unique optimal action with almost perfect monitoring is almost the same. So, this result is essentially the Maximum theorem in the sense that the optimal choice is "continuous". As a first step, the following lemma shows that σ_D is optimal if a player knows that someone has switched to the permanent defection and ϵ is small.

Lemma 10 There exists a $\bar{\epsilon} > 0$ such that $V_i\left(\sigma_i; \mu_{-i} : p_\epsilon; \delta\right)$ is maximized by σ_D for any p_ϵ, if $\mu_{-i}(\theta) = 1$ for any $\theta \neq 0$.

Proof. Take σ_D and any strategy which starts with C. The least deviation gain is $(1-\delta)\underline{g}$. The largest loss caused by the difference in continuation payoffs with σ_D and the latter strategy is $\delta\epsilon\overline{V}$. Setting $\bar{\epsilon}$ small enough guarantees that $(1-\delta)\underline{g} > \delta\epsilon\overline{V}$ for any $\epsilon \in (0, \bar{\epsilon})$. Then, D must be the optimal action for any such ϵ. Since players are using σ_D, $\mu_{-i}(\theta) = 1$ for some $\theta \neq 0$ in the next period. This implies that D is the unique optimal action in all the following periods. ■

Using $p(\cdot|\cdot)$ and given the fact that all players are playing either σ_C or σ_D, we define a transition probability of the number of players who have switched to σ_D. Let $q(l|m)$ be a probability that l players will play σ_D from the next period when m players are playing σ_D now. In other words, this $q(l|m)$ is a probability that $l-m$ players playing C receive the signal d when $n-m$ players play C and m players play D. Of course, $q(l|m) > 0$ if $l \geq m$ and $q(l|m) = 0$ if $l < m$.

The following lemma provides various informative and useful bounds on the variations of discounted average payoffs caused by introducing small imperfectness in private monitoring.

Lemma 11

[7] Also note that convergence of $V\left(\sigma_i; \mu_{-i} : p_\epsilon; \delta\right)$ to $V\left(\sigma_i; \mu_{-i} : p_0; \delta\right)$ is independent of the choice of the associated sequence $\{p_\epsilon\}$ because of the definition of P_ϵ.

1. $\inf_{p_\varepsilon \in P_\varepsilon} V(\sigma_C; 0 : p_\varepsilon; \delta) = \frac{(1-\delta)+\delta\varepsilon\underline{V}}{1-\delta(1-\varepsilon)}$

2. Given $\delta \in \left(\frac{g(0)}{1+g(0)}, 1\right)$; There exists a $\bar{\varepsilon} > 0$ such that for any $\varepsilon \in [0, \bar{\varepsilon}]$;

   $\sup_{\sigma_i, p_\varepsilon \in P_\varepsilon} V(\sigma_i; 0 : p_\varepsilon; \delta) \le \frac{1-\delta+\delta\varepsilon\overline{V}}{1-\delta(1-\varepsilon)}$

Proof. (1): For any $\varepsilon \in (0, 1)$ and $p_\varepsilon \in P_\varepsilon$;

$$V(\sigma_C; 0 : p_\varepsilon; \delta) = (1-\delta) + \delta q(0|0) V(\sigma_C; 0 : p_\varepsilon; \delta) + \delta(1 - q(0|0))\underline{V} \tag{39}$$

So,

$$V(\sigma_C; 0 : p_\varepsilon; \delta) = \frac{(1-\delta) + \delta(1 - q(0|0))\underline{V}}{1 - \delta q(0|0)} = \frac{(1-\delta) + \delta\varepsilon\underline{V}}{1-\delta(1-\varepsilon)} \tag{40}$$

(2): Given $\delta \in \left(\frac{g(0)}{1+g(0)}, 1\right)$; it is easy to check that $V(\sigma_C; 0 : p_0; \delta) > V(\sigma_D; 0 : p_0; \delta)$: Pick $\bar{\varepsilon}$ small enough such that (i) $V(\sigma_C; 0 : p_{\bar\varepsilon}; \delta) > V(\sigma_D; 0 : p_{\bar\varepsilon}; \delta)$ for any $p_{\bar\varepsilon}$ and (ii) $\bar{\varepsilon} < \hat{\varepsilon}$: Let σ_0^* be the optimal strategy given that everyone is using σ_C:[8] Suppose that σ_0^* assigns D for the first period. Then for any $\varepsilon \in [0, \bar{\varepsilon}]$;

$$\begin{aligned} V(\sigma_0^*; 0 : p_\varepsilon; \delta) \le & \; (1-\delta) U\left(\left(D, D^0\right) : p_\varepsilon\right) + \\ & \delta \left\{ q(1|1) V(\sigma_0^*; 0 : p_\varepsilon; \delta) + \sum_{\theta=2} q(\theta|1) V(\sigma_D; \theta - 1 : p_\varepsilon; \delta) \right\} \end{aligned} \tag{41}$$

In this inequality, the second component represents what player i could get if she knew the true continuation strategies of her opponents at each possible state. To see that this additional information is valuable, suppose that the continuation strategy of σ_0^* leads to a higher expected payoff than $V(\sigma_0^*; 0 : p_\varepsilon; \delta)$ or $V(\sigma_D; \theta - 1 : p_\varepsilon; \delta)$ at the corresponding states, then this contradicts the optimality of σ_0^* or σ_D by Lemma 10. Hence this inequality should hold.

Then, for any $\varepsilon \in [0, \bar{\varepsilon}]$;

$$\begin{aligned} V(\sigma_0^*; 0 : p_\varepsilon; \delta) & \le \frac{(1-\delta) U\left(\left(D, D^0\right) : p_\varepsilon\right) + \delta \sum_{\mu=2} q(k|1) V(\sigma_D; \theta - 1 : p_\varepsilon; \delta)}{1 - \delta q(1|1)} \\ & = V(\sigma_D; 0 : p_\varepsilon; \delta) \\ & < V(\sigma_C; 0 : p_\varepsilon; \delta) \end{aligned} \tag{42}$$

[8] Such σ^* exists because the strategy space is a compact space in product topology, on which discounted average payoff functions are continuous. Of course, this σ^* depends on the choice of p_ε:

Since this contradicts the optimality of σ_0^*; σ_0^* has to assign C for the first period.

Now,

$$V\left(\sigma_0^*; 0 : p_\varepsilon; \delta\right) \leq (1-\delta) + \delta q(0|0)\, V\left(\sigma_0^*; 0 : p_\varepsilon; \delta\right) + \delta\left(1 - q(0|0)\right)\overline{V} \tag{43}$$

So,

$$V\left(\sigma_0^*; 0 : p_\varepsilon; \delta\right) \leq \frac{(1-\delta) + \delta\left(1 - q(0|0)\right)\overline{V}}{1 - \delta q(0|0)} \leq \frac{(1-\delta) + \delta\varepsilon\overline{V}}{1 - \delta(1-\varepsilon)} \tag{44}$$

This implies that $\sup_{\sigma_i; p_\varepsilon \in P_\varepsilon} V_i\left(\sigma_i; 0 : p_\varepsilon; \delta\right) \leq \frac{1-\delta+\delta\varepsilon\overline{V}}{1-\delta(1-\varepsilon)}$ for any $\varepsilon \in [0; \overline{\varepsilon}]$: ■

(1) means that a small departure from the perfect monitoring does not reduce the payoff of σ_C much when all the other players are using σ_C. (2) means that there is not much to be exploited by using other strategies than σ_C with a small imperfection in the private signal as long as all the other players are using a σ_C.

The following result is an almost complete characterization of the optimal action as a function of μ_{-i}.

Proposition 12 Given δ; for any $\eta > 0$; there exists a $\overline{\overline{\varepsilon}} > 0$ such that for any $p_{\overline{\overline{\varepsilon}}}$;

- it is not optimal to play C for player i if μ_{-i} satisfies $\psi\left(\mu_{-i}\right) \leq \frac{1-\delta}{\delta} g\left(\mu_{-i}; p_0\right) - \eta$
- it is not optimal to play D for player i if μ_{-i} satisfies $\psi\left(\mu_{-i}\right) \geq \frac{1-\delta}{\delta} g\left(\mu_{-i}; p_0\right) + \eta$

Proof: (1): It is not optimal to play C if

$$\begin{aligned} &(1-\delta)\, g\left(\mu_{-i}; p_\varepsilon\right) \\ &> \delta\left[\psi\left(\mu_{-i}\right)\left\{(1-\varepsilon)\sup_{\sigma_i} V\left(\sigma_i; 0 : p_\varepsilon; \delta\right) + \varepsilon\overline{V}\right\} + \left(1 - \psi\left(\mu_{-i}\right)\right)\varepsilon\overline{V}\right] \end{aligned} \tag{45}$$

is satisfied because then any strategy which plays C now is dominated by σ_D:

By Lemma 11.2., this inequality is satisfied for any $\varepsilon \in [0; \overline{\varepsilon}]$ and any p_ε if

$$\begin{aligned} &(1-\delta)\, g\left(\mu_{-i}; p_\varepsilon\right) \\ &> \delta\left[\psi\left(\mu_{-i}\right)\left\{(1-\varepsilon)\frac{1-\delta+\delta\varepsilon\overline{V}}{1-\delta(1-\varepsilon)} + \varepsilon\overline{V}\right\} + \left(1 - \psi\left(\mu_{-i}\right)\right)\varepsilon\overline{V}\right] \end{aligned} \tag{46}$$

LHS converges to $(1-\delta)g\left(\mu_{-i}, p_0\right)$ and RHS converges to $\delta\psi\left(\mu_{-i}\right)$ as $\varepsilon \to 0$. So, if μ_{-i} satisfies $\psi\left(\mu_{-i}\right) \leq \frac{1-\delta}{\delta} g\left(\mu_{-i}, p_0\right) - \eta$ for any $\eta > 0$, then there exists a $\varepsilon'\left(\delta, \eta, \mu_{-i}\right) \in (0, \bar{\varepsilon})$ and a neighborhood $B\left(\mu_{-i}\right)$ of μ_{-i} such that C is not optimal for any $p_{\varepsilon'(\delta,\eta,\mu_{-i})}$ and any $\mu'_{-i} \in B\left(\mu_{-i}\right)$. This $\varepsilon'\left(\delta, \eta, \mu_{-i}\right) > 0$ can be set independent of μ_{-i} by standard arguments because $\left\{\mu_{-i} \mid \psi\left(\mu_{-i}\right) \leq \frac{1-\delta}{\delta} g\left(\mu_{-i}, p_0\right) - \eta\right\}$ is a compact subset in a $n-1$ dimensional simplex.

(2): It is not optimal to play D if

$$
\begin{aligned}
& (1-\delta) g\left(\mu_{-i}, p_\varepsilon\right) && (47)\\
< \; & \delta\left[\psi\left(\mu_{-i}\right)\left\{(1-\varepsilon) V\left(\sigma_C ; 0 : p_\varepsilon ; \delta\right) + \varepsilon \underline{V}\right\} + \left(1 - \psi\left(\mu_{-i}\right)\right)\left(\varepsilon \underline{V} - \varepsilon \bar{V}\right)\right]
\end{aligned}
$$

this inequality is satisfied for $\varepsilon \in (0, 1)$ and any p_ε if

$$
\begin{aligned}
& (1-\delta) g\left(\mu_{-i}, p_\varepsilon\right) && (48)\\
< \; & \delta\left[\psi\left(\mu_{-i}\right)\left\{(1-\varepsilon) \frac{1-\delta+\delta\varepsilon\underline{V}}{1-\delta(1-\varepsilon)} + \varepsilon \underline{V}\right\} + \left(1-\psi\left(\mu_{-i}\right)\right)\left(\varepsilon \underline{V} - \varepsilon \bar{V}\right)\right]
\end{aligned}
$$

This inequality converges to $\psi\left(\mu_{-i}\right) \geq \frac{1-\delta}{\delta} g\left(\mu_{-i}, p_0\right)$ as $\varepsilon \to 0$. So, if μ_{-i} satisfies $\psi\left(\mu_{-i}\right) \geq \frac{1-\delta}{\delta} g\left(\mu_{-i}, p_0\right) + \eta$ for any $\eta > 0$, there exists a $\varepsilon''\left(\delta, \eta, \mu_{-i}\right)$ such that D is not optimal for any $p_\varepsilon \in P_{\varepsilon''(\delta,\eta,\mu_{-i})}$ and any μ'_{-i} around μ_{-i}. Again, $\varepsilon''\left(\delta, \eta, \mu_{-i}\right)$ can be set independent of μ_{-i}.

Finally, setting $\bar{\varepsilon}(\delta, \eta) = \min\left\{\varepsilon'(\delta, \eta), \varepsilon''(\delta, \eta)\right\}$ completes the proof. $\blacksquare$

This proposition implies that the optimal action can be completely characterized except for an arbitrary small area around the manifold satisfying $\psi\left(\mu_{-i}\right) = \frac{1-\delta}{\delta} g\left(\mu_{-i}, p_0\right)$ in a $n-1$ dimensional simplex, where player i is indifferent between σ_C and σ_D with perfect monitoring, but we also characterize the optimal action for this region later.

Although this argument is essentially the path dominance argument used in Sekiguchi [13] for $n = 2$, it extends that argument to the n-player case and provides a sharper characterization even for $n = 2$.

An immediate corollary of this proposition is that C is the unique optimal action given that ψ is sufficiently close to 1, $\delta > \frac{g(0)}{1+g(0)}$, and ε is small.

Corollary 13 Given $\delta > \frac{g(0)}{1+g(0)}$, there exists $\bar{\psi} > 0$ and $\varepsilon > 0$ such that for any p_ε, it is not optimal for player i to play D if $\psi \in \left[\bar{\psi}, 1\right]$.

Since the optimal action is almost characterized as a function of μ_{-i}, now we need to know the dynamics of μ_{-i} associated with σ_C and σ_D. Given the optimal action shown above, what we need for consistency is that μ_{-i} stays in the "C area" described by Proposition 12 as long as player i has observed fully cooperative signals from the beginning and μ_{-i} stays in the "D area" once player i received a bad signal and started playing a defection for herself.

Let $W_k = \left\{\mu_{-i} \mid \psi\left(\mu_{-i}\right) \geq \frac{1-\delta}{\delta} g\left(\mu_{-i}; p_0\right) - k\right\}$, where $k \in \left(0, \frac{1-\delta}{\delta}\underline{g}\right)$. When k and ε is small, W_k just covers the region where C is the unique optimal action and the unique optimal action is indeterminate. We show that ψ_{t+n} is always above $\bar{\psi} \in (0,1)$ if a player plays C, observes c, $\psi_t \in W_k$, and monitoring is almost perfect.

Lemma 14 For any $\bar{\psi} > 0, k \in \left(0, \frac{1-\delta}{\delta}\underline{g}\right)$, there exists $\bar{\varepsilon}$ such that for any $\varepsilon \in (0, \bar{\varepsilon})$ $\psi\left(\mu_{-i}\left(h_i^{t+n}; p_\varepsilon\right)\right) \geq \bar{\psi}$ for $h_i^{t+n} = \left(h_i^t, (C, c), \ldots, (C, c)\right)$ when $\mu_{-i}\left(h_i^t\right) \in W_k$.

Proof. Let $h_i^{t+1} = \left(h_i^t, (C, c)\right)$ and $\psi_i^t = \psi\left(\mu_{-i}\left(h_i^t; p_\varepsilon\right)\right)$.
Applying Bayes' rule,

$$\begin{aligned} \psi_i^{t+1} &= \psi\left(\mu_{-i}\left(h_i^{t+1}; p_\varepsilon\right)\right) \\ &= \frac{\psi_i^t p(c|C)}{\psi_i^t P\left(\omega_i^t = c|C\right) + \left(1-\psi_i^t\right)\sum_{a_{-i}^t \neq c} P\left(\omega_i^t = c|a_{-i}^t\right) P\left(a_{-i}^t|h_i^t\right)} \end{aligned} \tag{49}$$

This function is increasing in ψ_i^t and crosses 45° line once. Note that this function is bounded below by $\varphi\left(\psi_i^t\right) = \frac{\psi_i^t(1-\varepsilon)}{\psi_i^t + \left(1-\psi_i^t\right)\varepsilon}$. Let $\hat{\psi}$ be the unique fixed point of this mapping. Given that $\psi_i^t = \psi\left(\mu_{-i}\left(h_i^t; p_\varepsilon\right)\right) \in W_k$, it is easy to see that $\varphi\left(\psi_i^t\right)$ and $\hat{\psi}$ can be made larger than $\bar{\psi} > 0$ by choosing ε small enough. If $\psi_i^t < \hat{\psi}$, then, as long as players continue to observe c, $\varphi^n\left(\psi_i^1\right)$ is going to increase monotonically to $\hat{\psi}$. On the other hand, since $\psi_i^{t+n} \geq \varphi\left(\psi_i^{t+n-1}\right)$ for $n = 1, 2, \ldots$ and φ is monotonically increasing, ψ_i^{t+n} is larger than $\varphi^n\left(\psi_i^t\right)$, hence larger than $\varphi\left(\psi_i^t\right)$ for any $n = 1, 2, \ldots$. On the other hand, if $\psi_i^t \geq \hat{\psi}$, then it is clear that $\psi_i^{t+n} > \varphi^n\left(\psi_i^t\right) \geq \hat{\psi} > \bar{\psi}$. These imply that $\left\{\psi_i^{t+n}\right\}_{n=1}^{\infty}$ is always above $\bar{\psi}$. ■

The above lemma guarantees that players are confident that the other players are cooperating after they observed a stream of cooperative signals and played C all the time. The next lemma is used to show that player i plays σ_D once a bad signal is observed or D has been played in the previous period. Let us define $\gamma_{p_\varepsilon} \in (0,1)$ as the smallest number such that

$$\gamma_{p^\varepsilon} = \frac{p(c|C)\,P\left(\omega_{-i}=\omega'_{-i}|C\right)}{\sum_{\omega_{-i}} p(c;\omega_{-i}|C)\,P\left(\omega_{-i}=\omega'_{-i}|c;f_i(\omega_{-i})\right)} \text{ for any } \omega'_{-i}\neq c \tag{50}$$

and

$$\gamma_{p^\varepsilon} = \frac{p\left(\omega'_{-i};c|D;a_{-i}=C\right)}{P\left(\omega_{-i}=\omega'_{-i}|D;a_{-i}=C\right)} \text{ for any } \omega'_{-i}$$

where $f_i:\omega_{-i}\mapsto\left(f_j(\omega_{-i})\right)_{j\neq i}\in\prod_{j\neq i}A_j$ is a mapping such that $f_j(\omega_j)=D$ if and only if $\omega_j\neq c$. The first condition implies that ϕ_i will be below γ_{p^ε} even if it is the first time for a player to observe anything other than c while C has been played. The second condition means that ϕ_i will be below γ_{p^ε} independent of the signal received when D is played. We impose the following regularity condition on the information structure we focus.

Assumption B: For some $\gamma<1$, $\gamma\geq\gamma_{p^\varepsilon}$.

Note that this assumption is not satisfied in the two player case in the previous sections. For example, $\psi_{Cd}(1)$ or $\psi_{Dd}(1)$ can be arbitrary close to 1 even if monitoring is almost perfect. However, we don't need this assumption when $n=2$. This assumption helps us to establish our result for $n\geq3$. The following is an example of information structure which satisfies this assumption when $n\geq3$ independent of ε[9]:

Example: Totally Decomposable Case

$p(\omega|a)=\prod_j\prod_{i\neq j}f(\omega_{i,j}|a_j)$ for all $a\in A$ and $\omega\in\Omega$

where $f(\omega|a)$ is a distribution function on $\{c,d\}$ such that $\omega=a$ with very high probability. Given the action by player j, the probability that player $i\neq j$ receives the right signal or the wrong signal about player j's action is the same across $i\neq j$. Also note that players' signals are conditionally independent over players.

The next lemma is an easy consequence of this assumption.

Lemma 15 $\phi\left(\mathbf{1}_{-i}\left(h_i^t;p^\varepsilon\right)\right)\leq\gamma$ after histories such as $h_i^t=\left(\ldots;(C;c);\left(C;\omega_i^{t-1}\right)\right)$ for $t\geq3$ where $\omega_i^{t-1}\neq c$ or $h_i^t=\left(\ldots;\left(D;\omega_i^{t-1}\right)\right)$

Proof. See appendix.

[9] An example of a more general class of p^ε which satisfies this assumption independent of the level of ε can be found in Obara [10].

4.2 Construction of Sequential Equilibrium

Let us introduce the following notations:

$$\begin{aligned}
U_\varepsilon^C &= \left\{\mu_{-i} \,|\, V\left(\sigma_C;\mu_{-i}:p_\varepsilon;\delta\right) > V\left(\sigma_D;\mu_{-i}:p_\varepsilon;\delta\right)\right\} \\
U_\varepsilon^\omega &= \left\{\mu_{-i} \,|\, V\left(\sigma_C;\mu_{-i}:p_\varepsilon;\delta\right) = V\left(\sigma_D;\mu_{-i}:p_\varepsilon;\delta\right)\right\} \\
U_\varepsilon^D &= \left\{\mu_{-i} \,|\, V\left(\sigma_C;\mu_{-i}:p_\varepsilon;\delta\right) < V\left(\sigma_D;\mu_{-i}:p_\varepsilon;\delta\right)\right\}
\end{aligned} \tag{51}$$

These are subsets of the belief space U. U_ε^ω is a manifold where player i is indifferent between σ_C and σ_D. In particular, $\mu^\pi_{-i}(\delta;p_\varepsilon) \in U_\varepsilon^\omega$ by definition. Note that U_ε^C converges to $U_0^C = \left\{\mu_{-i} \,|\, \phi\left(\mu_{-i}\right) > \frac{1-\delta}{\delta} g\left(\mu_{-i};p_0\right)\right\}$ and U_ε^D converges to $U_0^D = \left\{\mu_{-i} \,|\, \phi\left(\mu_{-i}\right) < \frac{1-\delta}{\delta} g\left(\mu_{-i};p_0\right)\right\}$ as $\varepsilon \to 0$.

Now define ρ as a mapping from $\mu_{-i} \in U$ to $\triangle\{C,D\}$ in the following way:

$$\rho_\varepsilon^*\left(\mu_{-i}\right) = \begin{cases} C & \text{if } \mu_{-i} \in U_\varepsilon^C \\ \pi(\delta;p_\varepsilon) & \text{if } \mu_{-i} \in U_\varepsilon^\omega \\ D & \text{if } \mu_{-i} \in U_\varepsilon^D \end{cases} \tag{52}$$

We know from Proposition 12 that this function assigns the best response action almost everywhere except for a neighborhood of U_0^ω when ε is small.

Now we use $\rho_\varepsilon^*\left(\mu_{-i}\right)$ to construct a Nash equilibrium approximating the efficient outcome, for which there exists a realization equivalent sequential equilibrium. All we have to do is to make sure that players are actually playing either σ_C or σ_D on the equilibrium path after they initially randomize between C and D.

Proposition 16 Suppose that Assumption B is satisfied. Then there exists a $\bar{\delta} \in \left(\frac{g(0)}{1+g(0)}, 1\right)$ such that for any $\delta \in \left(\frac{g(0)}{1+g(0)}, \bar{\delta}\right)$ there is a $\varepsilon(\delta) > 0$ where, for any p_ε, there exists a symmetric sequential equilibrium which is realization equivalent to $\rho_\varepsilon^*\left(\mu_{-i}\right)$, hence realization equivalent to $\pi(\delta;p_\varepsilon)\,\sigma_C + (1-\pi(\delta;p_\varepsilon))\,\sigma_D$.

Proof. Pick any $\bar{\delta} > \frac{g(0)}{1+g(0)}$ such that if $\delta \in \left(\frac{g(0)}{1+g(0)}, \bar{\delta}\right]$, then $\gamma < \mu^\pi_{-i}(p_0;\delta)(0) \left(< \bar{\chi}\right)$. First we prove that if $\phi_i \leq \gamma$, then D is the unique optimal action, hence the optimal continuation strategy is σ_D when ε is small enough.

Note that player i is indifferent between σ_C and σ_D with belief $\phi_i = \mu^\pi_{-i}(p_0;\delta)(0)$ with no noise, hence the following equality holds.

$$\sum_{\theta=0}^{n-1} \mu^\pi_{-i}(p_0;\delta)(\theta)\, \triangle V(\theta:p_0;\delta) = 0 \tag{53}$$

Suppose that $\phi_i = \mu_{-i}(p_0;\delta)(0) \leq \gamma < \mu^\pi_{-i}(p_0;\delta)(0)$. We show that $\sum_{\theta=0}^{n-1} \mu_{-i}(p_0;\delta)(\theta)\, \triangle V(\theta:p_0;\delta) < 0$ by comparing it to the above equality. If player i plays C here, then the positive payo¤ player i can get when µ = 0 decreases by at least ¡1 ¼$_{i}$ (p$_0$; ±) (0) ¡ ° ¢ 4V (µ : p$_0$; ±) compared to the case when Á$_i$ = 1¼$_{i}$ (p$_0$; ±) (0) : On the other hand, the additional gain from playing C when Á$_i$ **5** ° and µ 6= 0 is bounded above by ¡1 ¡ 1¼$_{i}$ (p$_0$; ±) (0) ¢ g compared to the case when Á$_i$ = 1¼$_{i}$ (p$_0$; ±) (0) : Since 1¼$_{i}$ (p$_0$; ±) (0) ! 1 as ± # $\frac{g(0)}{1+g(0)}$ by Lemma 9, playing C is strictly dominated when ± is chosen to be close to $\frac{g(0)}{1+g(0)}$. Any strategy playing C now continues to be dominated by ¾$_D$ even monitoring is almost perfect by the same argument as in Lemma 11[10].

Now, we check players' incentive on the equilibrium path to prove that ½¤" ¡¹ $_{i}$¢ is a symmetric Nash equilibrium. Players randomize between C and D with probabiity ¼ (±; p") ; 1 ¡ ¼ (±; p") respectively in the …rst period. First of all, Á$_i$ is strictly above Á as long as (C; c) has been observed by Lemma 14. C is the unique optimal action for such Á$_i$ by Lemma 13. Next, when a player …rst observes (C; ! (6= c)) after the second period, Á$_i$ gets below ° by Lemma 15. Hence, the unique optimal action is D by the above argument: If (C; ! (6= c)) is observed in the …rst period; then Á$_i$ is again clearly below ° for small " because ! is interpreted as a signal of ¾$_D$ being chosen in the …rst period rather than an error. So, D is always the unique optimal action after this kind of history which ends with (C; ! (6= c)) : Finally, when D is played, it is always the case that Á$_i$ **5** ° in the next period; hence the unique optimal action is again D: These imply that ¹ $_{i}$ 2 U$^C_"$ after (C; c) has been observed and ¹ $_{i}$ 2 U$^D_"$ after (C; ! (6= c)) is observed or, D is played in the previous period by de…nition of U$^C_"$ and U$^D_"$. So, ½¤" ¡¹ $_{i}$¢ is a symmetric Nash equilibrium, which is clearly realization equivalent to ¼ (±; p") ¾$_C$ + (1 ¡ ¼ (±; p")) ¾$_D$:

Existence of a sequential equilibrium which is realization equivalent to ½¤" ¡¹ $_{i}$¢ follows from the fact that this game is in a class of games with non-observable deviation. See Sekiguchi [13] for detail.¥

Since the probability that everyone chooses ¾$_C$ in this sequential equilibrium; ¼ (±; p")$^{n ¡ 1}$; gets closer to 1 as ± gets closer to $\frac{g(0)}{1+g(0)}$ by Lemma 9, an outcome arbitrary close to the e¢cient outcome can be achieved for ± arbitrary close to $\frac{g(0)}{1+g(0)}$. For high ±; we can use a public randomization device again to reduce ± e¤ectively or use Ellison's trick as in Ellison [4] to achieve an almost e¢cient outcome although the strategy is more complex and no longer a grim trigger. Hence, the following result is obtained.

Proposition 17 Suppose that Assumption B is satis…ed. Fix ± 2 ³$\frac{g(0)}{1+g(0)}$; 1´ : Then for any ¿ > 0; there is a " > 0 such that for any p"; there exists a sequential equilibrium whose symmetric equilibrium payo¤ is more than 1 ¡ ¿:

Finally we show that ½¤" ¡¹ $_{i}$¢ itself is actually sequential equilibrium with one more assumption:

Assumption C: If ¹ $_{i}$ 2 U$^D_"$; then Á$_i$ **5** ° after (C; ! (6= c)) is observed.

[10] This argument is not necessary when n = 2 because of the property; Â$_{D !}$ (¹) < ¹:

Proposition 18 Suppose that Assumption B and C is satisfied. For any $\delta \in \left(\frac{g(0)}{1+g(0)}, \bar{\delta}\right)$, if ε is small enough, then

- C is the unique optimal action if and only if $\mu_{-i} \in U_\varepsilon^C$
- D is the unique optimal action if and only if $\mu_{-i} \in U_\varepsilon^D$

Hence, $\sigma_\varepsilon^*\left(\mu_{-i}\right)$ itself is a sequential equilibrium.

Proof: Fix $\eta(\delta) > 0$ in Proposition 12 and set $k(\delta) > 0$ slightly larger than $\eta(\delta) > 0$ for each $\delta \in \left(\frac{g(0)}{1+g(0)}, \bar{\delta}\right)$. Take any μ_{-i} such that $\psi\left(\mu_{-i}\right) > \frac{1-\delta}{\delta} g\left(\mu_{-i}; p_0\right) - k$ and $\mu_{-i} \in U_\varepsilon^D$. We prove that D is the unique optimal action in this region if ε is sufficiently small. If a player plays C, then Lemma 14 and Assumption C imply that the continuation strategy is σ_C for small ε. This is because $\psi_i > \bar{\psi}(\delta)$ as long as (C, c) realizes and $\psi_i \leqq \gamma$ holds otherwise. Since σ_C is dominated by σ_D in this region, the unique optimal action should be D.

Similarly, take any μ_{-i} such that $\frac{1-\delta}{\delta} g\left(\mu_{-i}; p_0\right) + \eta > \psi\left(\mu_{-i}\right)$ and $\mu_{-i} \in U_\varepsilon^C$. If D is played, then again the continuation strategy is σ_D by Lemma 15. Since σ_D is dominated by σ_C in this region, the unique optimal action should be C.

Since any other $\mu_{-i} \in U_\varepsilon^D$ satisfies $\psi\left(\mu_{-i}\right) \leqq \frac{1-\delta}{\delta} g\left(\mu_{-i}; p_0\right) - k$ and any other $\mu_{-i} \in U_\varepsilon^C$ satisfies $\psi\left(\mu_{-i}\right) \geqq \frac{1-\delta}{\delta} g\left(\mu_{-i}; p_0\right) + \eta$ for small $\varepsilon > 0$, the proof is complete. ■

5 Concluding Comments

The main point of this paper has been to develop "belief-based" strategies as a way of constructing sequential equilibria in repeated games with private monitoring. This affords a major simplification as compared to the traditional method of analysis. While our construction has been restricted to the prisoners' dilemma, and to a strategy profile which consists only of two continuation strategies, the idea underlying this simplification is generalizable. If player i starts with a finitely complex (mixed) strategy which induces k possible continuation strategies, then the state space or the set of possible beliefs for player j for the entire repeated game is a $k-1$ dimensional simplex.

The approach of the present paper is based on generalizing "trigger strategy" equilibria to the private monitoring context. Under perfect or imperfect public monitoring, such trigger strategy can be constructed so as to provide strict incentives for players to continue with their equilibrium actions at each information set. Mailath and Morris [8] show that one can construct equilibria which provide similar strict incentives under private monitoring which is "almost-public". However, if private signals are not sufficiently correlated, pure trigger strategy profiles fail to be equilibria. The approach in the present paper, as in previous works such as Bhaskar and van Damme [3] and Sekiguchi [13], relies on approximating the grim trigger strategy with a mixed strategy. In the basic construction, a player is indifferent between cooperating and defecting in the initial period, but has strict incentives to play the equilibrium action at every subsequent information set. In particular, player i's strategy is measurable with respect to her beliefs about player j's continuation strategy. As Bhaskar [2] shows, such mixed strategies are robust to a small amount of incomplete payoff information as in Harsanyi's [6]. In particular, Bhaskar [2] shows in the context of the repeated prisoners' dilemma, where stage game payoffs are random and private, there exists a strict equilibrium with behavior corresponding to that of the mixed equilibrium of the present paper.[11] In the intial period, a player plays C for some realizations of his private payoff information, and D for other realizations, and continues with a trigger strategy in subsequent periods, independent of their payoff information.

The alternative approach to constructing non-trivial repeated game equilibria with private monitoring is due to Piccione [12] and Ely and Välimäki [5] [12]. This approach relies on using player j's mixed strategy to make a player i indifferent between playing C and D at every information set. Since player i is so indifferent, she is likewise willing to randomize so as to make j also indifferent between his actions at each information set. In this approach, beliefs are irrelevant, since a player's continuation payoff function does not depend upon her beliefs. Such equilibria seem to be less likely to survive if there is private payof information, and indeed this question is the subject of current research.

Appendix.

Proof of Lemma 9.

When $\delta = \frac{g(0)}{1+g(0)}$; $\pi(\delta; p_0) = 1$ is the solution of the equation in μ:

$$V\left(\sigma_C; \mu_{-i} : p_0; \delta\right) - V\left(\sigma_D; \mu_{-i} : p_0; \delta\right) = 0 \tag{54}$$

where $\mu_{-i}(\theta) = (1-\mu)^{\theta}\mu^{n-1-\theta}\binom{n-1}{\theta}$ for $\theta = 0, \ldots, n-1$.

We just need to show that $\frac{\partial \pi(\delta; p_0)}{\partial \delta}\Big|_{\delta = \frac{g(0)}{1+g(0)}} < 0$ using the implicit function theorem. Since,

$$\begin{aligned}
& V\left(\sigma_C; \mu_{-i} : p_0; \delta\right) - V\left(\sigma_D; \mu_{-i} : p_0; \delta\right) \\
= \; & (1-\delta)\sum_{\theta=0}^{n-1}(1-\mu)^{\theta}\mu^{n-1-\theta}\binom{n-1}{\theta}\Delta V(\theta : p_0; \delta) \qquad (55) \\
= \; & -(1-\delta)\sum_{\theta=0}^{n-1}(1-\mu)^{\theta}\mu^{n-1-\theta}\binom{n-1}{\theta} g(\theta) + \delta\mu^{n-1}
\end{aligned}$$

[11] This result is relevant since Bhaskar [1] considers a overlapping generations game with private monitoring and shows that incomplete payoff information as in Harsanyi implies an anti-folk theorem — players must play Nash equilibrium of the stage game in every period.

[12] See also the work of Kandori [7] in the context of a finitely repeated game.

a straightforward calculation gives the desired result as follows.

$$
\begin{aligned}
\frac{\partial \pi(\delta; p_0)}{\partial \delta}\Big|_{\delta=\frac{g(0)}{1+g(0)}} &= -\frac{\frac{\partial V_i(\sigma_C;\mu_{-i};p_0;\delta)-V_i(\sigma_D;\mu_{-i};p_0;\delta)}{\partial \delta}}{\frac{\partial V_i(\sigma_C;\mu_{-i};p_0;\delta)-V_i(\sigma_D;\mu_{-i};p_0;\delta)}{\partial \mu}}\Big|_{\delta=\frac{g(0)}{1+g(0)}} \qquad (56)\\
&= -\frac{1+g(0)}{(1-\delta)(n-1)(g(1)-g(0))+\delta(n-1)}\Big|_{\delta=\frac{g(0)}{1+g(0)}}\\
&= -\frac{1}{n-1}\frac{(1+g(0))^2}{g(1)} < 0
\end{aligned}
$$

■

Proof of Lemma 15

Suppose that $h_i^t = \left(\ldots;(C;c);\left(C;\omega_i^{t-1}\right)\right)$ with $\omega_i^{t-1} = \omega_i^0 \neq c$ and $t \geq 3$. ψ_i^t is bounded above by ψ^0 which is obtained by Bayes' rule after such an observation when $\psi_i^{t-2} = 1$. That ψ^0 is given by

$$
\begin{aligned}
&\frac{\Pr\left(\left(\omega^{t-2},\omega_i^{t-1},\omega_{-i}^{t-1}\right)=(c;\omega_i^0;c)\mid\theta_{t-2}=0\right)}{\Pr\left(\left(\omega^{t-2},\omega_i^{t-1}\right)=(c;\omega_i^0)\mid\theta_{t-2}=0\right)} \qquad (57)\\
\leq\ &\frac{p(c|C)\Pr(\omega_i=\omega_i^0|C)}{\sum_{\omega_{-i}} p(c;\omega_{-i}|C)\Pr(\omega_i=\omega_i^0|c;f(c;\omega_{-i}))}\\
\leq\ &\gamma
\end{aligned}
$$

Similarly, when $h_i^t = \left(\ldots;\left(D;\omega_i^{t-1}\right)\right)$ with $\omega_i^{t-1} = \omega_i^{00}$; ψ_i^t is bounded by

$$
\begin{aligned}
\psi_i^t &= \psi\left(\mathbf{1}_{-i}\left(h_i^t;p_\varepsilon\right)\right) \qquad (58)\\
&\leq \frac{p(\omega_i^{00};c|D;a_{-i}=C)}{\Pr(\omega_i=\omega_i^{00}|D;a_{-i}=C)}\\
&\leq \gamma \qquad (59)
\end{aligned}
$$

■

References

[1] V. Bhaskar, Informational Constraints and the Overlapping Generations Model: Folk and Anti-Folk Theorems, Review of Economic Studies, 65 (1998) , 135-149.

[2] V. Bhaskar, The Robustness of Repeated Game Equilibria to Incomplete Payo¤ Information, mimeo, 2000.

[3] Bhaskar, V., and E. van Damme, Moral Hazard and Private Monitoring, mimeo, 1997.

[4] G. Ellison, Cooperation in the prisoner's dilemma with anonymous random matching, Review of Economic Studies, 61 (1994), 567-588.

[5] J. C. Ely and J. Välimäki, A robust folk theorem for the prisoner's dilemma, mimeo, 1999.

[6] J. Harsanyi, Games with Randomly Disturbed Payo¤s: A New Rationale for Mixed-Strategy Equilibrium Points, Int. J. Game Theory 2 (1973), 1-23.

[7] M. Kandori,Cooperation in Finitely Repeated Games with Imperfect Private Information, mimeo, 1991.

[8] G.Mailath, and S. Morris, Repeated games with almost-public monitoring, mimeo, 1999.

[9] H. Matsushima, On the theory of repeated game with non-observable actions, part I: Anti-Folk Theorem without communication. Econ. Letters 35 (1990), 253-256.

[10] I. Obara, Repeated prisoners' dilemma with private monitoring: an N-player case, mimeo, 1999.

[11] I. Obara, Private strategy and e¢ciency: repeated partnership games revisited, mimeo, 1999.

[12] M. Piccione, The repeated prisoners' dilemma with imperfect private monitoring, mimeo, 1998.

[13] T. Sekiguchi, E¢ciency in the Prisoners' Dilemma with Private Monitoring, J. Econ. Theory,76 (1997), 345-361.