Centre for Development Economics

WORKING PAPER SERIES

Repeated Games With Almost Perfect Monitoring By Privately Observed Signals

V. Bhaskar

Delhi School of Economics

Working Paper No: 11

Centre for Development Economics
Delhi School of Economics
Delhi 110 007 INDIA

REPEATED GAMES WITH ALMOST PERFECT MONITORING

BY PRIVATELY OBSERVED SIGNALS

V. Bhaskar*
Delhi School of Economics
Delhi University
Delhi 110007, India

May 1994

* This paper was first presented at the Economic Theory Conference, Indian Statistical Institute, Delhi, in January 1994. I am grateful to conference participants and to Dilip Mookherjee and Debraj Ray for their comments. My special debt is to Eric van Damme, who pointed out the importance of considering mixed strategies. Indeed, the analysis of section 5 of the paper is essentially due to Eric.

I. INTRODUCTION

Following the work of Green and Porter (1984) and Abreu, Pearce and Stachetti (1986, 1990), there is now a large literature on repeated games with imperfect monitoring. This literature assumes that players observe a public signal, i.e. a signal which is common knowledge between the players. In this framework,Fudenberg, Levine and Maskin (1989) have shown that a version ofthe folk-theorem applies. This paper is concerned with a qualitatively different situation, where players monitor each other's actions by a privately observed signal. This is the situation in many important contexts, including the case of oligopolistic competition between price setting firms discussed by Stigler (1964) (see also the discussion in Fudenberg and Tirole (1991)). The critical difference with the public signals case is that the realization of signals is not mutual knowledge between the players. Our focus is on the situation where the signals are "almost perfect", i.e. where they almost perfectly correlated with actions. In other words, we are considering a game with imperfect monitoring which is arbitrarily close to a game with perfect monitoring. We find however a dramatic discontinuity - the equilibrium set of the game with almost perfect monitoring is qualitatively different from the equilibrium set of the game with perfect monitoring. If the errors in the observations made by players are uncorrelated, any pure strategy equilibrium of the repeated game requires players to play the Nash equilibrium of the stage game every period. This result appears to be robust to some correlation in the observation errors, as we show via an example. Although mixed strategy equilibria allow a wider range of behavior,we find that they do not allow us to approximate fully cooperative behavior. More precisely, we show that the equilibrium set and the set of equilibrium payoffs fails to be lower-hemicontinuos in the level of noise at the point of zero noise. Consequently, the fully cooperative outcome cannot be approximated even if monitoring is almost perfect. We must emphasize however that our analysis of mixed strategies is as yet preliminary, and much remains to be done.

The rest of this paper is as follows. Section 2 presents a simple example which demonstrates the basic idea of the paper. Section 3 discusses pure strategy equilibria in general repeated games when observational errors are independent. Section 4 returns to the example and analyzes the implications of correlated errors. Apart from demonstrating the robustness of our earlier result, we have another surprising finding – with correlated errors, the pure strategy equilibrium set is not monotonic in the degree of noise, and in some situations, increased noise may facilitate cooperation. Section 5 discusses mixed strategy equilibria and the final section concludes.

2. AN EXAMPLE

Consider the game G1, which is a variant on the prisoners' dilemma, and which we choose to interpret as a stylized model of bilateral trade. Two traders are exchanging fruit, an apple for an orange. They can send each other a good fruit (action C, for cooperate), or a spoiled fruit (action D, for defect). Fruit quality is unverifiable by a third party, so that the traders cannot write binding contracts to enforce the Pareto-optimal action pair (C,C). Players have, in addition, action E, which allows cooperation to be sustained when the game is finitely repeated –should one seek an interpretation, think of it as sending a poisoned apple! The payoff x can for the moment be taken to be 0.G1 has two pure strategy Nash equilibria – (D,D) and (E,E), with the latter being strictly worse for both players. This allows the players to cooperate when G is repeated, even if the repetition is only finite. For example when G is repeated twice, the action pair (C,C) in the first period can be supported; the players choose D in the second period if the first period actions are (C,C), and choose E otherwise. Extending this argument, if the game is repeated T times, the players can cooperate in the first T-1 periods, so that their average payoff is 8-5/T.

The standard way of adding noise to this model is to assume that players sometimes make mistakes, so that their actual actions are different from their intended actions. A player who intends to play C will sometimes make a mistake, so that with a small probability ε, his actual action will be D rather than C. Players can observe their opponent's actual actions, but are unable to observe their intentions. Consequently, a player who makes a mistake unintentionally must be punished, in order to prevent opportunistic behavior. This creates an inefficiency, but the inefficiency is small, being of order ε. When G is repeated twice,the expected payoff in the subgame perfect equilibrium described above is:

$$(1-\varepsilon)^2\, 11 + \varepsilon(1-\varepsilon)(10+3) + \varepsilon^2 4 \qquad (2.1)$$

This is less than the payoff of 11 in the game without noise, butas $\varepsilon \longrightarrow 0$, the payoff in the noisy game converges to 11. In otherwords, noise creates inefficiency, but this inefficiency vanishes as the noise vanishes, Further, as the number of repetitions gets large, and as ε tends to zero, the per-period payoff converges to 8.

The interpretation of the noise in the above model is that of Selten's trembling hand. I fully intend to play C, but my hand is jiggled, and I find myself choosing D, and sending you rotten fruit. When I do this I am aware that I have made a mistake, and that you will punish me for this. In other words, although my intended action is private information to me, my actual action is mutual knowledge (and common knowledge) between us at the end of the period. Put somewhat differently, the noisy game corresponds to one where actions (which we have called intended action here) are private information, but the signal (which we have called the actual action here) is publicly known.

Consider now an alternative formulation of noise which is the focus of this paper. I send you an good orange, but with a very small probability the orange deteriorates en route. I know my action, i.e. that I have sent you a good orange. However I do not know whether you have received a good orange or a bad one, i.e. I do not know what signal you have received. The signal is almost perfectly correlated with my action, but only almost. Neither the action nor the signal are mutual knowledge between us, although they are arbitrarily close to being so, in a probabilistic sense.

This formulation creates a drastic discontinuity. Any pure strategy equilibrium of the repeated game must consist of playing the Nash equilibrium of the stage game G in every stage. In other words, C cannot be played in such an equilibrium.

To see this consider the twice repeated game. Clearly, C can only be played in the first period, and to make the playing of C optimal, a player should punish a first-period deviation (to D) by playing E rather than D in the second period. However, this punishment, which depends on varying in second period behavior with the signal, is not optimal. Suppose that I believe that my opponent is playing the above strategy, i.e. playing C in the first period. If I now observe the signal D at the end of the period, I should believe that the signal arose by mistake - the application of Bayes' rule to my opponent's strategy implies that my opponent chose C, and that my observation of D is due to the noise in the signaling technology. I have played C, and with a very high probability $(1-\varepsilon)$ my opponent has observed C and is going to continue with D. Consequently, I should continue with D rather than E. Since varying second period behavior with the signal is not optimal, this makes it impossible to support the playing of C with probability one in the first period.

The inability to support cooperation is a robust feature given the information structure. In the game G, the only value of x for which playing C in the first period can be supported $x=3-\varepsilon/(1-\varepsilon)$. At this value of x, I am indifferent between C and D given my opponent's strategy, and this indifference holds regardless of the signal I receive. In other words, our negative result is generically true, as we shall see in the next section. Consider now the repetition of the stage game for a finite number of periods, say T. We shall see that our negative results apply no matter what the value of T. It might be surmised that this is due to a standard backward induction argument, using the analysis of the twice repeated game. However, this is not the case.In our model, the game that remains after T-2 periods is not analytically the same as a two-period game. After T-2 periods, the player's have imperfect information regarding the history of actions as well as of signals. Since the two players are not at the same information set, the continuation strategies at the end ofT-2 periods do not have to be equilibrium strategies of the two period game. Consequently, usual backward induction arguments cannot be used in this case, even though the game is only finitely repeated. As we shall see in the next section, the analysis has to begin in the first period and proceed by normal induction. As such,our method of proof does not distinguish between finitely and infinitely repeated games, and our results apply equally to both.

3. PURE STRATEGY EQUILIBRIA WITH INDEPENDENT SIGNALS: A GENERAL FRAMEWORK

Let $I = \{1,2,\ldots,n\}$ be the set of players, let A_i, $i \in I$ be the action set for each player. $A = \times_i A_i$, and U_i: $A \rightarrow R$ is the payoff function for player i. $U = (U_1,\ldots, U_n)$. The stage game G is the triple (I,A,U). G is repeated either for T periods or infinitely often. Players seek to maximize the expected discounted sum of payoffs, using a common discount rate δ. If G is repeated infinitely often, the discount rate, $\delta < 1$, incorporates the probability of termination in every period. If G is repeated finitely often, $\delta \leq 1$. Payoffs are received at the end of the game.

At the end of each period player i observes a (n-1) vector of signals, $b_i = (b_{i1}, b_{i2}, \ldots, b_{in})$. b_{ij} is player i's signal regarding the action taken by player j. Note that i and j are always distinct when we write b_{ij}. b_{ij} is drawn from a finite set B_{ij}, which is the set of all possible signals that i could receive regarding j's chosen action. To keep things simple we shall assume that $|B_{ij}| = |A_j|$, i.e. that the number of possible signals regarding j's action is the same as the number of possible actions that j could choose. $B_i = \times_j B_{ij}$ is the set of all possible signal combinations that i can receive.

Since signals are private information, histories are also player specific. A history upto to period t for player i, h_i^t, is a sequence of realizations of signals, $(b_i^1, b_i^2, \ldots, b_i^{t-1})$. The set of all possible histories upto period t for player i, H_i^t, is simply the product $(B_i)^{t-1}$. The history at period one for any player is the null history, h^1, i.e. H_i^1 is a singleton set for all players. We restrict attention to pure strategies; hence a strategy for player i, s_i, is a sequence $\langle s_i^t \rangle$, where $s_i^t: H_i^t \rightarrow A_i$. Given the restriction to pure strategies there is no loss of generality in not allowing s_i^t to depend upon player i's own past actions. A strategy profile is the n-tuple $s = (s_1, s_2, \ldots, s_n)$.

We turn now to the signalling technology. Write $p(b_{ij}/a_j)$ for the conditional probability that player i receives signal b_{ij} given that player j has chosen action a_j. We make the following assumptions.

Assumption 1. *Full Support.* $\forall\, i,j \in I$, $p(b_{ij}/a_j) > 0$ $\forall\, b_{ij} \in B_{ij}$, $\forall\, a_j \in A_j$.

Assumption 2. *Independence.* for all distinct $i,j,k \in I$,

$p(b_{ij}/a_j, b_{kj}) = p(b_{ij}/a_j)$, $\forall\, b_{ij} \in B_{ij}$, $\forall\, a_j \in A_j$, $\forall\, b_{kj} \in B_{kj}$.

Our concern in this paper is mainly with signalling technologies which are almost perfect. To make this precise, we define perfect and ε-perfect signalling technologies as follows.

Definition. The repeated game G^* has perfect monitoring if for every ordered pair of players (i,j), there exists a one-to-one correspondence $f_{ij}:A_j \longrightarrow B_{ij}$ such that $p(f_{ij}(a_j)/a_j) = 1$, $\forall\, a_j \in A_j$.

Definition. Let $\varepsilon>0$. The repeated game $G^*(\varepsilon)$ has ε-perfect monitoring if it satisfies the full support assumption and for every ordered pair of players (i,j), there exists a one-to-one correspondence $f_{ij}:A_j \longrightarrow B_{ij}$ such that $p(f_{ij}(a_j)/a_j) \geq 1-\varepsilon$, $\forall\, a_j \in A_j$.

If a repeated game has ε-perfect monitoring, it follows that $1 > p(f_{ij}(a_j)/a_j) \geq 1-\varepsilon$ for every ordered pair (i,j). Consider a strictly positive sequence $\langle\varepsilon_n\rangle$ converging to zero. At $\varepsilon=0$, the repeated game is one with perfect monitoring and therefore violates the full support assumption. Conversely, for any game G^* with perfect monitoring and a sequence $\langle\varepsilon_n\rangle$ converging to zero, one can find a sequence of games of imperfect monitoring, $G^*(\varepsilon_n)$. We shall identify G^* with the limit of the sequence $G^*(\varepsilon_n)$. Our concern is with equilibria of the game with perfect monitoring which are limits of sequences of games with ε-perfect monitoring where ε tends to zero.

Given any repeated game, let $V_i(s/h_i^t)$ denote the expected continuation payoff from period t to player i from strategy combination s conditional on i having observed the history h_i^t. A strategy combination $s = (s_i,s_{-i})$ is a Nash equilibrium if $V_i(s/h^1) \geq V_i(s_i',s_{-i}/h^1)$ every every other repeated game strategy s_i'.

Definition. Let G^* be a repeated game with perfect monitoring. A pure strategy profile $s = (s_1,s_2,\ldots, s_n)$ is a *robust equilibrium* of G^* if:

(i) s is a Nash equilibrium of G^* and

(ii) there exists a strictly positive sequence $\langle\varepsilon_n\rangle$ converging to zero, with an associated sequence $\langle G^*(\varepsilon_n)\rangle$ of games with ε_n-perfect monitoring, such that s is a Nash equilibrium of $G^*(\varepsilon_n)$ for all n.

Remark: Given our full support assumption it suffices to restrict attention to Nash equilibria – with ε-perfect monitoring, each signal is received with positive probability, so that there are no information sets that are ruled out due to the actions of other players.

To recall, the stage game is the triple (I,A,U). The repeated game with perfect monitoring is G^*, which may be either finitely or infinitely repeated. The associated repeated game with ε-perfect monitoring is $G^*(\varepsilon)$. In this section we restrict attention to signalling technologies which satisfy the independence assumption defined above. This implies that observational errors made by the player's (due to the imperfection of the signal) are independent.

Consider the stage game G, and fix I and A. There are $|A|$ possible action combinations and $n \times |A|$ possible payoff numbers. Let $m = n \times |A|$. Given I and A, the stage game is fixed by specifying a point in R^m. A property P is satisfied by almost all games if, given I and A, the set of points in R^m such that P is true is a closed set of Lebesgue measure zero in R^m.

Theorem 1. Consider a signalling technology which satisfies independence. For almost all stage games G, if s is a robust pure strategy equilibrium of the repeated game G^*, s requires the play of a Nash equilibrium of G at every stage and after any history.

To prove this theorem we need the following mathematical result which is set out as lemma 1. Before we present this result, note that δ is a fixed strictly positive real number strictly less than one.

Let $x \in R^m$, $x = (x_1, x_2, \ldots, x_k, \ldots, x_m)$.

Let $X = \{w \in R : w = e_k x,\ 1 \le k \le m\}$, i.e. $X = \{x_1, x_2, \ldots, x_k, \ldots, x_m\}$

Condition C1: x has the property that there exist distinct infinite sequences $\langle y_t \rangle$ and $\langle z_t \rangle$ with range X such that $\sum^{\infty} \delta^t (y_t - z_t) = 0$.

Condition C2: x has the property that there exist distinct finite sequences $\langle y_t \rangle$ and $\langle z_t \rangle$, $t=1,2,.,T$ with range X such that $\sum^{T} \delta^t (y_t - z_t) = 0$.

Condition C3: x has the property that there exist finite sequences $\langle y_t \rangle$ and $\langle z_t \rangle$, $t=1,2,\ldots,T$, with range X such that $\sum^{T} (y_t - z_t) = 0$, and $\langle y_t \rangle$ is not a permutation of $\langle z_t \rangle$.

Lemma 1.Let $S = \{x \in R^m$: x satisfies C1 or C2 or C3$\}$. S is a closed set of Lebesgue measure zero in R^m.

Proof: Each of the conditions C1,C2, and C3 define a single equation among the finite number of elements of X, so that the lemma follows. □

Lemma 2. If s is a robust Nash equilibrium, $s_i^t : H_i^t \rightarrow A_i$ is a constant function, for every player i, for almost all stage games G.

Proof: The proof is by induction.

For $t=1$, s_i^1 is obviously constant since H_i^1 is a singleton set.

Consider now arbitrary t, and consider a game with ε-perfect monitoring. Let G be generic, so that by the induction hypothesis, s_j^τ is constant on H_i^τ for every $\tau < t$, and for every player j. Let $h_i^t, h_i^{t'} \in H_i^t$ be distinct t-period histories.

We claim that :

$$V_i(s/h_i^t) = V_i(s/h_i^{t'}) \qquad 3.1$$

Since s is a pure strategy combination and (by the induction hypothesis) s_j^τ is a constant function for every $\tau<t$, for every player j, by Bayes' rule player i's beliefs regarding the actions chosen by j in periods 1,2,...,t-1, are degenerate: player i assigns probability one to the belief that j has chosen a deterministic action sequence, for every player j. Hence player i's beliefs regarding past actions are independent of the observed history. By the independence of signals (Assumption 2), player i's beliefs regarding the signals observed by other players, is independent of the realization of b_i^τ, for every $\tau<t$. Hence the probability assigned by player i to player j being at

history $h_j^t \in H_j^t$, is the same regardless of whether player i is at h_i^t or $h_i^{t'}$. Hence i's probability distribution over j's actions in period t and subsequently is the same, regardless of whether h_i^t or $h_i^{t'}$ is the observed history ofsignal. From the full support assumption, h_i^t and $h_i^{t'}$ are both observed with positive probability. Since s is a Nash equilibrium, player i must play a best response at h_i^t, and at $h_i^{t'}$, so that the equality of expected payoffs in (3.1) follows.

Consider now the game with perfect monitoring. In the game with perfect monitoring, at every information set, distinct actions generate distinct deterministic sequences of payoffs. Since s is a robust Nash equilibrium, (3.1) applies. We claim that $s_i(h_i^t) = s_i(h_i^{t'})$ for almost all stage games G. If $s_i(h_i^t) = s_i(h_i^{t'})$, so that player i takes two distinct actions, this generates two distinct real sequences, with range the set of pure action payoffs to i, which can call $\langle y_s \rangle$ and $\langle z_s \rangle$, $s = t, t+1, \ldots$.If G* is an infinitely repeated game, (3.1) implies that $\sum_{s=t}^{\infty} \delta^s (y_s - z_s) = 0$, so that $\langle y_s \rangle$ and $\langle z_s \rangle$ satisfy C1. By lemma 1, G is non-generic, i.e the set of payoffs of G such that this condition holds is a closed set of Lebesgue measure zero.

If G* is finitely repeated with discounting, (3.1) implies that $\sum_{s=t}^{T} \delta^s (y_s - z_s) = 0$, where $0 < \delta \leq 1$. Since $\langle y_s \rangle$ and $\langle z_s \rangle$ satisfy C2, lemma 1 shows that G is non generic.

Tf G* is finitely repeated without discounting, (3.1) implies $\sum_{s=t}^{T} (y_s - z_s) = 0$. Lemma 1 shows that if G is generic, $\langle y_s \rangle$ is a permutation of $\langle z_s \rangle$. However, if $\langle y_s \rangle$ is a permutation of $\langle z_s \rangle$, then for a generic game, the associated sequences of action n-tuples are permutations of each other. This is however impossible since there is no public signal on which to base the permutation. □

After lemma 2 the prooof of the theorem is straightforward- if it is G is not possible to vary in any period in response to signals, no player can be punished for deviating in any period. Hence in each period, the action combination must be a Nash equilibrium of the stage game. □

Remark: It seems possible to strengthen Theorem 1 so that it applies directly to the game with imperfect monitoring, and not just in the limit. This would require a stronger version of lemma 1, relating to sequences of random variables.

4 CORRELATED SIGNALS

In this section we examine the implications of relaxing Assumption 2, that the signals received by players are independent. We do not present a general analysis, but return to our two-period repetition of G1. For example, in our interpretation of this game as a model of bilateral trade, the probability of good fruit spoiling may be related to the weather, and weather shocks may affect the fruit that both players receive.

We adopt the following notation for this section: upper case letters denote actions, and the corresponding lower case letters denote signals.

Consider the following joint distribution of signals conditional on the action pair (C,C):

		2's signal c	2's signal d
1's signal	c	$(1-\varepsilon)^2 + \rho\varepsilon(1-\varepsilon)$	$(1-\rho)\varepsilon(1-\varepsilon)$
	d	$(1-\rho)\varepsilon(1-\varepsilon)$	$\varepsilon^2 + \rho\varepsilon(1-\varepsilon)$

If player i chooses D, the probability that j observes d is one. If i chooses C and j chooses D, the probability that j observes C is $(1-\varepsilon)$.

This probability distribution is parametrized by ε and ρ. As before, ε is the degree of noise in the signal, and is the (unconditional) probability that the signal takes the value d when action C is chosen. ρ is the degree of correlation between signals. ρ = 0 corresponds to case where the signals are independent as earlier, while ρ>0 correponds to positive correlation between signaling errors. ρ=1 is the case where the signals are perfectly correlated, which may be considered the case where signals are in fact publicly observed.

If signaling errors are positively correlated, this allows for the possibility of a player varying his actions depending on the signal he receives, along a pure strategy path. With positively correlated signaling errors, if I receive a bad signal, this makes it more likely that my opponent has also received a bad signals. Consequently an agreement to punish each other if we receive a bad signal could be made self enforcing. However, the degree of correlation must be large enough. In our game, if we consider the case where the noise vanishes, (i.e. $\varepsilon \rightarrow 0$), ρ must be greater than 0.75, if the payoff parameter x is zero, as we assume for most of this section.

Consider the sustainability of the cooperative strategy, i.e.:

1st period: C

2nd period: D if signal c

E if signal d

To check that this is a Nash equilibrium we need to see that second period behavior is optimal. It is optimal to play D if the probability of my opponent playing D is greater than 1/4, and optimal to play E if this probability is less than 1/4. In other words, along the equilibrium path where the action pair (CC) has been chosen, we require that if I have observed c, the conditional probability that my opponent has also observed c, p(c/c;CC), must satisfy:

$$p(c/c;CC) = (1-\varepsilon) + \rho\varepsilon \geq 1/4 \qquad (4.1)$$

Similarly, if I have observed d, the conditional probability that my opponent has observed c, must satisfy:

$$p(c/d;CC) = (1-\rho)(1-\varepsilon) \leq 1/4 \qquad (4.2)$$

These two inequalities are graphed in Fig 2. Recall that ε and ρ both lie in the unit interval (negative values for ρ are possible, but they only reinforce our earlier results). The shaded area shows values of ε and ρ such that the second period behavior of the above cooperative strategy is optimal. Fig 2 illustrates a number of interesting points regarding the sustainability of cooperation. First, as we have already seen, as $\varepsilon \rightarrow 0$, ρ must be greater than 3/4. Second, consider values of $\rho < 3/4$. The figure shows that cooperation is not possible if ε is small and close to zero, but may be possible for larger values of ε. Of course we still have to check that first period behavior is optimal. This requires:

$$8 + 3[(1-\varepsilon^2) + \rho\varepsilon(1-\varepsilon)] + 1\varepsilon^2 + \rho\varepsilon(1-\varepsilon) \geq 9 + 1 \tag{4.3}$$

The cooperative strategy is an equilibrium iff (ε,ρ) satisfy inequalities (4.1)-(4.3). It may be verified that if $\rho=2/3$, $\varepsilon=1/3$ satisfies these conditions. However, as we have already seen, the conditions are violated for sufficiently small values of ε; and further (4.3) is violated if ε is too large.

We summarize the results of this section in the following proposition:

Proposition 1. Consider the game G1, with correlated signals, parametrized by (ε,ρ).

i) As the noise parameter $\varepsilon \rightarrow$, cooperation can be supported only if the degree of correlation between signals, ρ, is sufficiently high, i.e. $\rho \geq 3/4$.

ii) The pure strategy equilibrium set need not be monotonic in ε, i.e. cooperation can be supported at intermediate levels of noise, but not either at very low or very noise.

Our result here contrasts sharply with Kandori's (1992) result. Kandori shows that if signals are publicly observed, the equilibrium set increases as the noise in the signal is reduced.

There is a distinct point which may be worth mentioning here, which arises when we vary the game G1, by varying parameter x. Let $\varepsilon \rightarrow 0$, and consider the values of ρ required to support cooperation as a function of x. $\rho(x)$ is decreasing in x, i.e. as x becomes larger, less correlation is required in order to support cooperation in the game with almost perfect monitoring. As x increases, the relative riskiness of the two stage game equilibria changes - (E,E) becomes less risky as compared to (D,D). This suggests that a conflict between Pareto-dominance and risk-dominance may help support cooperative behavior in repeated games.

5. MIXED STRATEGY EQUILIBRIA

In this section we discuss the role of mixed strategies. We consider again our example, G1. Further, we assume that signals are independent, i.e. we retain Assumption 2. The inability to support playing C in the first period can be seen as due to the following reason. In a pure strategy equilibrium, I can perfectly forecast my opponent's actions in the next period from his strategy alone. His signal consequently does not convey any additional information to me. Even though the signaling technology is almost perfect, it is only almost so, and my prior information takes precedence. To give the signal some bite, it must convey some information about my opponent's second period actions in equilibrium, i.e. even in the absence of any noise. This is possible if we allow for mixed strategies. Once again, the payoff parameter x in G1 is zero.

Consider the following strategies for the repeated game.

Strategy A: 1st period: C

2nd period: D if signal c

E otherwise

Strategy B: 1st period: D

2nd period: E for all contingencies

The payoff matrix for these two supergame strategies is:

	A	B
A	$8+3\varepsilon(1-\varepsilon)+\varepsilon^2$	3
B	10	4

Since A and B are best responses to each other in this matrix, the mixed strategy where A is played with probability approximately one-half $(p=1/(2-5\varepsilon+3\varepsilon^2))$ is a an equilibrium of the above matrix.

Proposition 2: The above mixed strategy equilibrium is a Nash equilibrium of the repeated game.

Proof: If my opponent plays p, it is sub-optimal to choose E in the first period, and I should start with C or D. If I start with D, my opponent receives signal d, and hence I should continue with E. If I start with C, I should continue with D or E, since C is strictly dominated. If I receive signal c, I know that he is playing the supergame strategy A, and that he has with very high probability$(1/1+\varepsilon)$ received signal c, and will play D in with probability close to 1, so that I should play D as well. If I receive signal d,the probability that he has played D is $(1/2)/(1/2\ +\varepsilon/2)\approx 1$. Hence he is most likely to continue with E, and I should play E as well. Hence I should conform to A if I start with A. Since only A and B are best responses against the mixed strategy p in the repeated game, p is an equilibrium. □

The payoff in the mixed strategy equilibrium is 4+6p, which converges to 7 as ε tends to zero. This is greater than the payoff of 6 which can be supported as a pure strategy equilibrium, but strictly less than the payoff of 11 which is achievable in the game without noise, i.e. when $\varepsilon=0$. It is also possible to show that a payoff of greater than 7 cannot be supported as long as ε is positive. This highlights a discontinuity in the transition from agame without noise to the game with noise. The set of sequential equilibria fails to be lower-hemicontinuous in ε at the point $\varepsilon=0$. This implies that the outcomes in the game with noise can be drastically different from the outcomes in the game without noise,even though the noise itself may be infinitesmal. This contrasts sharply with the standard way of analyzing a noisy game, where the signals are mutual knowledge between the players. As we saw in section 2, the payoff of 11 is the limit of the equilibrium payoffs in the noisy game as the noise vanishes.

6. CONCLUSIONS

The analysis of this paper is preliminary and much remains to be done. Repeated games with private signals are obviously a rich area for further research, and this paper has only probed the surface. To our knowledge, the only other paper on this area is the paper by Fudenberg and Levine (1990). They consider approximate equilibria, and prove a version of the folk theorem. Our results are dramatically different, at least insofar as pure strategies are concerned. However, it is well known (see Radner, 1980) that the set of approximate equilibria can be qualitatively different from the set of exact equilibria.

Reference

Abreu, D., D. Pearce and E. Stachetti, 1986, Optimal Cartel Equilibria with Imperfect Monitoring, *Journal of Economic Theory* 39:191-228.

Abreu, D., D. Pearce and E. Stachetti, 1990, Towards a theory of discounted repeated games with imperfect monitoring, *Econometrica* 56:1041-1064.

Fudenberg, D., and J. Tirole, 1991, *Game Theory*, Cambridge Mass: MIT Press.

Fudenberg, D., D. Levine and E. Maskin, 1989, The folk-theorem in repeated games with imperfect public information, mimeo, Massachusetts Institute of Technology.

Fudenberg, D.and D. Levine, 1990, Approximate equilibria in repeated games with imperfect private information, *Journal of Economic Theory* (forthcoming).

Green, E., and R.Porter, 1984, Noncooperative collusion under imperfect price information, *Econometrica* 52:87-100.

Kandori, M., 1992, Monotonicity of the equilibrium set with respect to the level of noise, *Review of Economic Studies.*

Radner, R., 1980, Collusive Epsilon equilibria in oligopolies with long but finite lives, *Journal of Economic Theory.*

Stigler, G., 1964, A theory of oligopoly, *Journal of Political Economy* 72:44-61.

I

Game G1

		II		
		C	D	E
I	C	8, 8	0, 9	0, 0
	D	9, 0	3, 3	0, x
	E	0, 0	x, 0	1, 1

$0 \leq x \leq 3$

Mostly, $x = 0$

Fig (2)

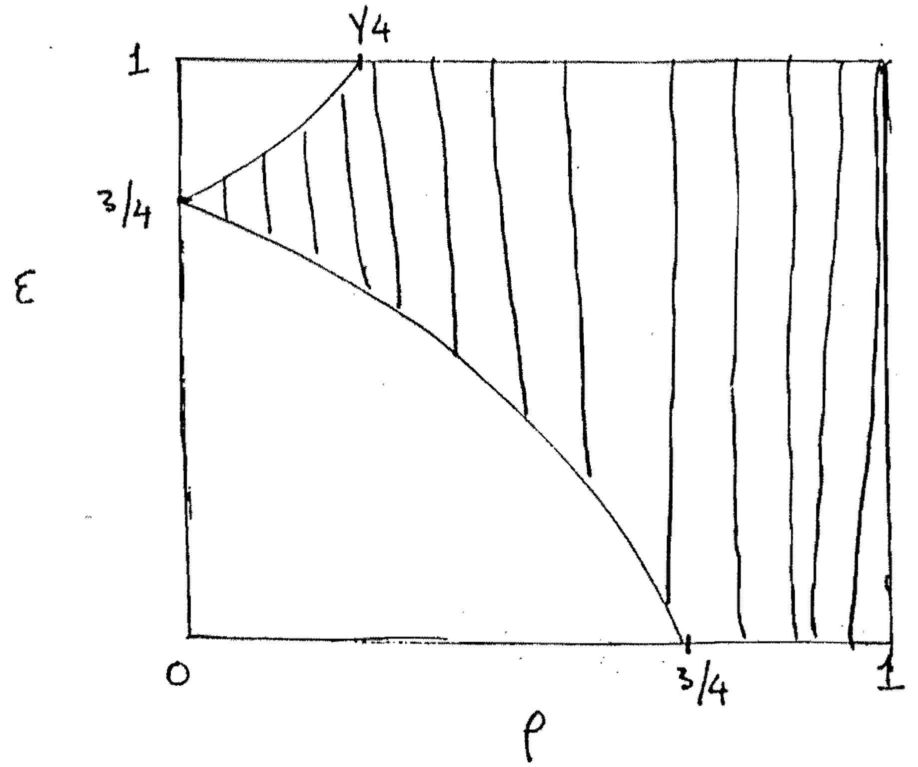

(ε, ρ) pairs such that cooperation is sustainable

CENTRE FOR DEVELOPMENT ECONOMICS WORKING PAPER SERIES

Number	Author(s)	Title
1	Kaushik Basu Arghya Ghosh Tridip Ray	The Babu and The Boxwallah : Managerial Incentives and Government Intervention
2	M. N. Murty Ranjan Ray	Optimal Taxation and Resource Transfers in a Federal Nation
3	V. Bhaskar Mushtaq Khan	Privatization and Employment : A Study of The Jute Industry in Bangladesh
4	V. Bhaskar	Distributive Justice and The Control of Global Warming
5	Bishnupriya Gupta	The Great Depression and Brazil's Capital Goods Sector : A Re-examination
6.	Kaushik Basu	Where There Is No Economist : Some Institutional and Legal Prerequisites of Economic Reform in India
7.	Partha Sen	An Example of Welfare Reducing Tariff Under Monopolistic Competition
8.	Partha Sen	Environmental Policies and North-South Trade : A Selected Survey of The Issues
9.	Partha Sen Arghya Ghosh Abheek Barman	The Possibility of Welfare Gains with Capital Inflows in A Small Tariff-Ridden Economy
10.	V. Bhaskar	Sustaining Inter-Generational Altruism when Social Memory is Bounded
11.	V. Bhaskar	Repeated Games with Almost Perfect Monitoring by Privately Observed Signals