Int. Journal of Game Theory. Vol. 9, Issue 3, page 157–167. ©*Physica-Verlag, Vienna.*

On Repeated Games with Incomplete Information Played by Non-Bayesian Players

By *N. Megiddo*, Tel Aviv[1])

Abstract: Unlike in the traditional theory of games of incomplete information, the players here are *not* Bayesian, i.e. a player does not necessarily have any prior probability distribution as to what game is being played. The game is infinitely repeated. A player may be absolutely uninformed, i.e. he may know only how many strategies he has. However, after each play the player is informed about his payoff and, moreover, he has perfect recall. A strategy is described, that with probability unity guarantees (in the sense of the liminf of the average payoff) in any game, whatever the player could guarantee if he had complete knowledge of the game.

1. Introduction

A game is said to be of incomplete information if at least one player does not know exactly which game is being played. *Harsanyi* [1967–68] proposed an embedding of the games of incomplete information within the class of games of complete information. The embedding is based on the assumption that the players are Bayesian. Specifically, the game is assumed to have been chosen by chance, with probability distribution which is itself public knowledge. Also, some information about chance's choice has been revealed to different players, according to rules that are themselves public knowledge. Essentially, different players have different prior probability distributions with respect to the game being played. As the game (i.e. the game that has once been chosen by chance) is repeated, these probabilities may be updated and, typically, a player has to consider future changes in other players' probability distributions that may be caused by his own decisions in the present.

Following *Harsanyi* [1967–68], contributions to this field have been made by (alphabetically) *Aumann/Maschler* [1968], *Kohlberg* [1975a, b], *Mertens* [1973], *Mertens/Zamir* [1971–72], *Ponssard/Zamir* [1973], *Stearns* [1967] and *Zamir* [1971–72, 1973]. All these papers deal with infinitely repeated two-person zero-sum games. They all assume that an uninformed player is also *not* informed about his payoff at the end of each stage; his payoffs are rather credited (or debited) somehow to his bank account, and he never receives any statements. On the other hand, he is informed about his opponent's choice according to prescribed rules.

[1]) Professor *N. Megiddo*, Statistics Department, Tel Aviv University, Tel Aviv, Israel.

In general, an informed player in a repeated game of incomplete information can take advantage of the fact that his opponent is uninformed about the payoffs. A known example is as follows. The game being repeatedly played is $\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$. However, while player I (the rows player) is informed about the game, player II starts with prior probabilities of .5 for the true game and .5 for $\begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$. If player I mixes his both pure strategies with the same probability, then player II never gains any additional information, and his optimal strategy under these circumstances is also to mix his both pure strategies with equal probabilities. Thus, the expected payoff is 1/4 per stage. On the other hand, the value of the same game with complete information is of course zero.

In this paper we consider a model which is quite different from the traditional one. In the first place, our players are not necessarily Bayesian, i.e. they do not necessarily have any prior probability distributions as to which game is being played. Secondly, the players are informed about their payoffs at the end of each stage, and they have perfect recall with respect to these payoffs. Our goal is to present a strategy for an absolutely uninformed player, that essentially guarantees him in any game, whatever he could guarantee if he were completely informed.

The game does not have to be two-person zero-sum. By playing our strategy, an uninformed player is guaranteed in any non-cooperative n-person game, to get as a long-run average, a payoff that is not less than his maximin expected payoff in the one-shot game with complete information.

After an earlier version of this paper had been issued [*Megiddo*, 1979], the author learned from Professor T. Ferguson about *Baños'* paper [1968]. Surprisingly, no reference to *Baños* [1968] is found in the literature on games of incomplete information. As a matter of fact, Baños proposed a strategy that essentially guarantees the same thing which is guaranteed by the strategy of the present paper. However, Baños' strategy is complicated and the present author believes that the strategy introduced here may shed more light on the problem, since its construction is simpler and intuitive.

2. The Strategy

The game that is being repeated infinitely many times is given in the normal form, i.e. a real r by c matrix G. Player I is the rows player and II is the columns player. The entries correspond to payoffs made by player II to player I. Even though we formulate everything in terms of two-person zero-sum games, the results can be interpreted in a more general setup, if G is the matrix of player I's payoff where columns correspond to joint strategies of all other players.

The strategy presented below is meant for player I. However, all player I needs to know at the start, is the number r. It is assumed that player I is informed about his payoff at the end of each play, and that he recalls all his payoffs from previous stages.

We start with an overview of the main ideas in the construction of our strategy. Roughly, the strategy may be viewed as a never-ending search for the optimal mixed strategy $s(G)$ of the real game G being played. During this search we discover strategies that yield increasing maximin payoffs. However, since we keep experimenting with new strategies, we frequently have to employ the previously discovered "good" strategies, in order to compensate for losses incurred during those "experiments."

The search for the optimal strategy of G relies on the following hierarchy of games. Let $C^n = \{G_1^n, \ldots, G_{m_n}^n\}$ be the set of all $r \times n$ game-matrices, whose entries are of the form k/n, where k is an integer such that $|k| \leq n^2$ (hence $m_n = (2n^2 + 1)^{rn}$) and $v(G_1^n) \geq \ldots \geq v(G_{m_n}^n)$. Let $|G'|$ denote the largest absolute value of an entry in G'. For $n \geq c$, by replicating the last column of G so as to form an $r \times n$ matrix, we can talk about the "distance" $|G_j^n - G|$ between G and G_j^n.

Our first observation is that an uninformed player can statistically test the following hypothesis. The real game G is "close" to a given game G'. Such an hypothesis is tested by repeatedly playing the optimal strategy of G' and comparing the average payoff so obtained with the value of G'. If that observed average is significantly less than the value of G', then the hypothesis is rejected. During our search we may accept games which are "far" from G, however that does not cause any loss in the long-run.

Obviously, for a sufficiently large n, we can find in C^n a game $G_{i_n}^n$ which is arbitrarily close to G. Now, suppose that we successively test (relative to the above hypothesis) the games $G_1^n, G_2^n, \ldots, G_{m_n}^n$ in this order, until the first acceptance. Obviously, by letting the number of plays in a single test be sufficiently large, we can guarantee (with probability arbitrarily close to one) that the average payoff observed in the last test will be greater than $v(G) - \epsilon$, for an arbitrarily small $\epsilon > 0$. Furthermore, we can design the test so that the conditional probabilities p_n, of accepting $G_{i_n}^n$ given that it is tested, converge sufficiently fast to one. That rate of convergence could be chosen so as to guarantee (with probability one) that rejections of games $G_{i_n}^n$ occur only finitely many times.

The principle now is quite simple. We keep testing games of C^n with increasing n, while frequently compensating for the losses by playing strategies of "accepted" games. The values of the accepted games constitute a sequence whose liminf is not less that $v(G)$. Even if a game G' is accepted by statistical error or by a "trick" of the opponent, it will be rejected later if it does not compensate as anticipated. In other words, the average payoff from any accepted game is continuously being watched and a test stops only with rejection.

The following sequence plays a key role in the construction of the strategy:

$$K_n = \left[\frac{16n^4}{1 - 2^{-1/2^n}}\right] \qquad (n = 1, 2, \ldots).$$

We are now ready to describe our strategy for player I in an infinitely repeated game. We describe the strategy in a form of an algorithm which includes the operation "play." Specifically, the algorithm is run for definite amounts of time between consecutive plays of the game G, and always provides player I with a strategy for the following stage. STACK is our pool of previously tested and accepted games.

Strategy S

0. Initialize with STACK $= \emptyset$ and $j_k = 0$ $(k = 1, 2, \ldots)$.
1. Set n to the least number such that both $j_n < m_n$ and for every STACK member G_i^k, $k < n$.
2. $j_n = j_n + 1$.
3. Repeat steps 31 and 32 K_n times:
   31. If STACK $= \emptyset$ then go to 32; otherwise let G_i^k be that STACK member whose upper index k is maximal. *Play* $s(G_i^k)$ during $n^3 m_n$ consecutive stages, subject to the following discipline: If at any time the average payoff for plays of $s(G_i^k)$ so far drops below $v(G_i^k) - 1/k$, then immediately remove G_i^k from STACK and go to 1.
   32. *Play* $s(G_{j_n}^n)$ once.
4. If the average payoff for plays of $s(G_{j_n}^n)$ so far is at least $v(G_{j_n}^n) - 1/n$, then place $G_{j_n}^n$ in STACK.
5. Go to 1.

The actual payoff at every stage depends of course only on the pure strategies chosen by the players. If mixed strategies are used, then the payoffs are random variables. Our main theorem is

Theorem: If player I plays strategy S and player II plays any strategy, then the payoff sequence $X_1, X_2, \ldots$ satisfies.

$$\text{Prob}\left\{\liminf_{q\to\infty} \frac{1}{q} \sum_{i=1}^{q} X_i \geqslant v(G)\right\} = 1.$$

3. Proof of the Theorem

Before proving the theorem we state several lemmas. Recall that by $|G'|$ we mean the maximum absolute value of an entry in G'.

Lemma 1: For every $n \geqslant \text{Max}(|G|, c)$ there is a game $G_{i_n}^n$ $(1 \leqslant i_n \leqslant m_n)$ such that if I plays $s(G_{i_n}^n)$ in G, then his expected payoff is at least $v(G_{i_n}^n) - 1/(2n)$.

Proof: Without loss of generality we assume that $n = c = |G|$, since columns of G may always be replicated without affecting the value. Obviously, there is a game $G_{i_n}^n$ $(1 \leqslant i \leqslant m_n)$, such that the absolute difference between any entry of $G_{i_n}^n$ and the corresponding entry of G is not less than $1/(2n)$. That implies our claim.

Corollary: $|v(G_{i_n}^n) - v(G)| \leqslant 1/(2n)$.

Lemma 2: Let $Y_1, Y_2, \ldots$ be a sequence of independent random (0,1)*-variables,* *such that* $p_k \equiv \text{Prob}\{Y_k = 1\} = 2^{-1/2^k}$ $(k = 1, 2, \ldots)$. *Under these conditions, with probability one, there is a number* K *such that for every* $k > K$, $Y_k = 1$.

Proof: For every $K \geqslant 0$, let A_K be the event in which $Y_k = 1$ for all $k > K$. Obviously,

$$\text{Prob}\,(A_K) = \prod_{k=K+1}^{\infty} 2^{-1/2^k} = 2^{-1/2^K} = p_K.$$

Thus,

$$\text{Prob}\{(\exists K)\,(\forall k > K)\,(Y_k = 1)\} = p\,(A_0) + \sum_{k=1}^{\infty} (1 - p_k)\,p\,(A_k)$$

$$= .5 + \sum_{k=1}^{\infty} (1 - 2^{-1/2^k})\,2^{-1/2^k} = \lim_{k \to \infty} 2^{-1/2^k} = 1.$$

For any $n \geqslant \text{Max}\,(|\,G\,|, c)$, let $G^n_{i_n}$ be the game whose existence is asserted in Lemma 1. Also, let

$$K_n = \left[\frac{16n^4}{1 - 2^{-1/2^n}}\right].$$

Lemma 3: Suppose that player I repeatedly plays the strategy $s\,(G^n_{i_n})$, *where* $n \geqslant \text{Max}\,(|\,G\,|, c)$, *in an infinitely repeated play of* G. *Under these conditions, with probability not less than* $2^{-1/2^n}$, *for every* $k \geqslant K_n$ *the average payoff for the first* k *stages is at least* $v\,(G^n_{i_n}) - 1/n$.

Proof: Without loss of generality, assume that player II repeatedly plays his best-reply strategy with respect to $s\,(G^n_{i_n})$ in G. Let X_i be the payoff for the i-th stage $(i = 1, 2, \ldots)$. $X_1, X_2, \ldots$ are mutually independent random variables with the same expectation $\mu \geqslant v\,(G^n_{i_n}) - 1/2n$ (Lemma 1) and the same variance $\sigma^2 \leqslant n^2$.

Kolmogorov's inequality [see *Feller*, p. 220] states: For every $\epsilon > 0$ and integer q,

$$\text{Prob}\left\{(\forall k \leqslant q)\left(\frac{1}{k}\sum_{i=1}^{k} X_i > \mu - \frac{\epsilon}{k}\right)\right\} \geqslant 1 - \frac{q\sigma^2}{\epsilon^2}.$$

It follows that

$$\text{Prob}\left\{(\forall k \geqslant K_n)\;\left(\frac{1}{k}\sum_{i=1}^{k} X_i > \mu - \frac{1}{2n}\right)\right\}$$

$$= \text{Prob}\bigcap_{j=1}^{\infty}\left\{K_n 2^{j-1} \leqslant k < K_n 2^j \Rightarrow \frac{1}{k}\sum_{i=1}^{k} X_i > \mu - \frac{1}{2n}\right\}$$

$$= 1 - \text{Prob}\bigcup_{j=1}^{\infty}\left\{(\exists k)\left(K_n 2^{j-1} \leqslant k < K_n 2^j, \frac{1}{k}\sum_{i=1}^{k} X_i \leqslant \mu - \frac{1}{2n}\right)\right\}$$

$$\geqslant 1 - \text{Prob} \bigcup_{j=1}^{\infty} \left\{ (\exists k) \left(K_n 2^{j-1} \leqslant k < K_n 2^j, \frac{1}{k} \sum_{i=1}^{k} X_i \leqslant \mu - \frac{1}{2n} \cdot \frac{K_n 2^{j-1}}{k} \right) \right\}$$

$$\geqslant 1 - \sum_{j=1}^{\infty} \text{Prob} \left\{ (\exists k < K_n 2^j) \left(\frac{1}{k} \sum_{i=1}^{k} X_i \leqslant \mu - \frac{1}{2n} \cdot \frac{K_n 2^{j-1}}{k} \right) \right\}$$

$$= 1 - \sum_{j=1}^{\infty} \left[1 - \text{Prob} \left\{ (\forall k < K_n 2^j) \left(\frac{1}{k} \sum_{j=1}^{k} X_i > \mu - \frac{1}{2n} \cdot \frac{K_n 2^{j-1}}{k} \right) \right\} \right]$$

$$\geqslant 1 - \sum_{j=1}^{\infty} \sigma^2 K_n 2^j \cdot \left(\frac{2n}{K_n 2^{j-1}} \right)^2 \geqslant 1 - \frac{8n^4}{K_n} \sum_{j=1}^{\infty} \frac{1}{2^{j-1}} = 1 - \frac{16n^4}{K_n} \geqslant 2^{-1/2^n}.$$

We now turn to the proof of the theorem. First, note that during a play of the infinitely repeated game, the variable n in strategy S exceeds any finite value. Also, the variables j_n change monotonically and hence each one of them finally attains some maximal value J_n $(1 \leqslant J_n \leqslant m_n)$. Given any sequence of strategy-choices by player II, J_n may be viewed as a random variable.

For every $n \geqslant \text{Max}(|G|, c)$, let i_n be the lower index of the game $G^n_{i_n}$ $(1 \leqslant i_n \leqslant m_n)$ whose existence is asserted in Lemma 1.

Assertion 4: For every $n \geqslant \text{Max}(|G|, c)$,

$$\text{Prob}\ \{J_n \leqslant i_n\} \geqslant 2^{-1/2^n}.$$

Proof: In order for J_n to exceed i_n, it is necessary that strategy $s(G^n_{i_n})$ is played k times, where $k \geqslant K_n$, and the average payoff per play for these plays is less than $\nu(G^n_{i_n}) - 1/n$. However, the probability of such an event is, by Lemma 3, less than $1 - 2^{-1/2^n}$.

Assertion 5: The probability that there is N, such that $J_n \leqslant i_n$ for all $n \geqslant N$, is one.

Proof: The proof follows from Lemma 2 and Assertion 4. However, note that the events $\{J_n \leqslant i_n\}$ are *not* independent, since the number of times a strategy $s(G^n_i)$ is played does depend on average payoffs obtained during plays of another strategy $s(G^{n+1}_j)$. Thus, we have to use the following argument. Suppose that the player extends his play of the strategy $s(G^n_{J_n})$ beyond strategy S. Suppose that he keeps playing it, against a fictitious player, as long as the average payoff remains above $\nu(G^n_{J_n}) - 1/n$. If it drops below that level, he switches to playing $s(G^n_{J_n+1})$ and so on. Let J^*_n be the final value of j_n in such a fictitious play. Obviously, the events $\{J^*_n \leqslant i_n\}$ are mutually independent and with probability one there is N such that $J_n \leqslant J^*_n \leqslant i_n$ for all $n \geqslant N$.

Henceforth, let N be fixed as that number which exists with probability one according to Assertion 5.

For any time t and every game G^k_i that has entered STACK at one time prior to t, we use the following notation. By $Z_i^k(t)$ we mean the average payoff per play, for plays performed prior to t, of either $s(G_i^k)$ or of some $s(G_{j_n}^n)$ immediately following a play of $s(G_i^k)$, but such that $G_{j_n}^n$ was not a member of STACK at any time prior to t.

Assertion 6: If G_i^k enters STACK at time t_0 and if $k \geqslant N$, then for all $t \geqslant t_0$, $Z_i^k(t) \geqslant v(G) - 5/k$.

Proof: First note that as long as G_i^k belongs to STACK, the average payoff for plays of $s(G_i^k)$ remains at least $v(G_i^k) - 1/k$. We now estimate the impact of plays of $s(G_{j_n}^n)$ on $Z_i^k(t)$. Since each play of $s(G_{j_n}^n)$ (which is included in $Z_i^k(t)$) has been preceded by at least $n^3 m_n$ plays of $s(G_i^k)$, it follows that the average payoff for plays either of $s(G_i^k)$ or of a specific $s(G_{j_n}^n)$ preceded by $s(G_i^k)$, (note that $n > k$) is not less than

$$\frac{n^3 m_n [v(G_i^k) - 1/k] - n}{n^3 m_n + 1} \geqslant v(G_i^k) - \frac{1}{k} - \frac{2}{n^2 m_n}.$$

Thus, if G_i^k is still in STACK at time t, then

$$Z_i^k(t) \geqslant v(G_i^k) - \frac{1}{k} - \sum_{n=k+1}^{\infty} m_n \cdot \frac{2}{n^2 m_n} \geqslant v(G_i^k) - \frac{3}{k}.$$

If G_i^k leaves STACK at time t, then since $s(G_i^k)$ has been played at least K_k times,

$$Z_i^k(t) \geqslant \frac{K_k [v(G_i^k) - 3/k] - k}{K_k + 1} \geqslant v(G_i^k) - \frac{4}{k}.$$

Finally, once a game G_i^k leaves STACK, $s(G_i^k)$ is never played again. Thus, since $k \geqslant N$, for all $t \geqslant t_0$

$$Z_i^k(t) \geqslant v(G_i^k) - \frac{4}{k} \geqslant v(G) - \frac{5}{k}.$$

Let $N_i^k(t)$ denote the number of plays accounted under $Z_i^k(t)$. It follows that the average payoff per play, for all plays prior to t, is

$$Z(t) = \frac{\sum_{k,i} N_i^k(t) Z_i^k(t)}{\sum_{k,i} N_i^k(t)},$$

(when $Z_i^k(t)$ may be set arbitrarily to $v(G)$ if $N_i^k(t) = 0$). Note that $N_i^k(t)$ is monotone non-decreasing and has some maximal value N_i^k. For any t,

$$Z(t) \geqslant v(G) - \frac{\sum_{k,i} N_i^k(t)(5/k)}{\sum_{k,i} N_i^k(t)}.$$

Given any number M, for any t,

$$Z(t) \geqslant v(G) - \frac{\sum_{k=1}^{M} \sum_{i=1}^{m_k} N_i^k (5/k)}{\sum_{k,i} N_i^k(t)} - \frac{5}{M}.$$

Since $\sum_{k,i} N_i^k(t) \to \infty$ as t increases, it finally follows that $\liminf_{t\to\infty} Z(t) \geqslant v(G)$. This completes the proof of the theorem.

4. Discussion

When a completely informed player plays an optimal mixed strategy in an infinitely repeated game, then he guarantees, in general, no more than that with probability one the liminf of the average actual payoff will not be less than $v(G)$. This is precisely what can be guaranteed by an absolutely uninformed player. However, one may argue that payoffs should be discounted, rather than averaged in the long run. It seems hard to analyze what precisely can be guaranteed in terms of a discount factor α $(0 < \alpha < 1)$. However, in the light of the lemma proved in the appendix the following is true. For every $\epsilon > 0$ there is α_0 $(0 < \alpha_0 < 1)$ such that for all $\alpha \geqslant \alpha_0$, $(\alpha < 1)$, strategy S guarantees the discounted payoff $\sum_{i=0}^{\infty} \alpha^i X_i$ to be at least $(v(G) - \epsilon)/(1-\alpha)$.

Another common approach in infinitely repeated games is to look at $\lim_{n\to\infty} V_n/n$, where V_n is the value of the finitely repeated game with n stages. A finitely repeated game here has no value, since the players are not assumed to be Bayesian. However, if V_n is defined to be the amount that player I can guarantee as his expected payoff then strategy S (followed up to n stages) guarantees expected payoffs such that $\lim_{n\to\infty} V_n/n = v(G)$.

We should mention that it is much easier for an uninformed player to achieve his maximin payoff relative to pure strategies. This is done as follows. The player always plays either a pure strategy he has not played before, or, if all have been played, the one whose worst case payoff so far is the greatest among all worst-case payoffs so far. Playing like that, the number of stages, in which the player may be paid less than his maximin (in pure strategies), will not be greater than $r - 1$ (where r is the number of his pure strategies).

It has been suggested by R.J. Aumann, J.F. Mertens, S. Sorin and S. Zamir that the latter idea can be generalized to the set of all mixed strategies for player 1. Thus, grids over the strategies simplex have to be selected, which have to get finer and finer, and the player has to "experiment" with mixed strategies belonging to these grids. "Compensation" may be provided by playing the best (in the sense of average payoff yielded) strategy so far. "Experiments" should however be performed only during a subsequence of stages whose density within the grand sequence is zero. J.F. Mertens has also pointed out that another strategy may be constructed by repeatedly applying the idea of [*Aumann/Maschler*] for the case when player 1 knows that G is in some finite set of games. Again, there has to be some "compensation" factor whenever a new round of considering a larger set of games begins.

Finally, one may wonder about the rate of convergence of the average payoff to the value of the game, while strategy S is being played. We were, of course, quite generous in selecting the different parameters of the strategy. It is conceivable that a more careful design of a learning strategy would lead to a better convergence rate, especially in situations where a player does have some partial information about the game at the start. However, our goal here was only to point out the feasibility of guaranteeing the value in the long-run under any circumstances.

Acknowledgments

The author is indebted to R.J. Aumann, T. Ferguson, E. Kalai, M. Maschler, J.F. Mertens, S. Sorin and S. Zamir.

Appendix

Lemma: If $\{a_n\}_{n=0}^{\infty}$ *is a bounded sequence then*

$$\liminf_{n\to\infty} \frac{1}{n+1} \sum_{i=0}^{n} a_i \leqslant \liminf_{\alpha\to 1-} (1-\alpha) \sum_{n=0}^{\infty} a_n \alpha^n .$$

This lemma is in fact a special case of an Abelian theorem stated in *Widder* [1941]. Specifically, apply Theorem 1 on p. 181 with $\gamma = 1, e^{-S} = \alpha$, and $\alpha(t)$ a step-function with jumps at the positive integers, whose sizes are a_n. For the sake of completeness we provide here a simple proof for our special case.

Proof: Without loss of generality we assume that

$$\liminf_{n\to\infty} \frac{1}{n+1} \sum_{i=0}^{n} a_i = 0 \text{ and } |a_i| \leqslant 1 \quad (i = 0, 1, \ldots).$$

Denote $S_n = \sum_{i=0}^{n} a_i$. It follows that for all α, $0 < \alpha < 1$,

$$\lim_{n\to\infty} S_n \alpha^n = 0.$$

That implies

$$\sum_{k=0}^{\infty} a_k \alpha^k = \lim_{n\to\infty} \sum_{k=0}^{n} a_k \alpha^k = \lim_{n\to\infty} [(1-\alpha) \sum_{k=0}^{n-1} S_k \alpha^k + S_n \alpha^n]$$

$$= (1-\alpha) \sum_{k=0}^{\infty} S_k \alpha^k .$$

Let $\epsilon > 0$ be given and let N be such that for every $n \geqslant N$, $S_n/(n+1) > -\epsilon/4$. Let α_1 be such that for all $\alpha, \alpha_1 < \alpha < 1, (1-\alpha)^2 \sum_{n=0}^{N-1} S_n \alpha^n > -\epsilon/2$.
Note that

$$f(\alpha) \equiv (1-\alpha)^2 \sum_{n=N}^{\infty} (n+1)\, \alpha^n = (1-\alpha)^2 \, [\sum_{n=N}^{\infty} \alpha^{n+1}]'$$

$$= (1-\alpha)^2 \left[\frac{\alpha^{N+1}}{1-\alpha}\right]' = (N+1)(1-\alpha)\,\alpha^N + \alpha^{N+1} \underset{(\alpha\to 1)}{\to} 1$$

and let α_2 be such that for all α $(\alpha_2 < \alpha < 1)\, f(\alpha) < 2$. It follows that for all α such that $\operatorname{Max}(\alpha_1, \alpha_2) < \alpha < 1$,

$$(1-\alpha) \sum_{n=0}^{\infty} a_n \alpha^n = (1-\alpha)^2 \sum_{n=0}^{\infty} S_n \alpha^n$$

$$= (1-\alpha)^2 \sum_{n=0}^{N-1} S_n \alpha^n + (1-\alpha)^2 \sum_{n=N}^{\infty} \frac{S_n}{n+1} (n+1)\, \alpha^n$$

$$> -\frac{\epsilon}{2} - \frac{\epsilon}{4} f(\alpha) > -\epsilon.$$

That implies that $\liminf_{\alpha\to 1-} (1-\alpha) \sum_{n=0}^{\infty} a_n \alpha^n \geqslant 0$ and hence completes the proof.

Remark: Under the same conditions it is known that $\liminf_{n\to\infty} \frac{1}{n+1} \sum_{i=0}^{n} a_i \geqslant$ $\liminf_{\alpha\to 1-} (1-\alpha) \sum_{n=0}^{\infty} a_n \alpha^n$.
[See *Titchmarsh*].

References

Aumann, R.J., and *M. Maschler*: Repeated Games with Incomplete Information: The Zero-Sum Extensive Case. Report to the U.S. Arms Control and Disarmament Agency, Washington, D.C. Final report on Contract ACDA/ST-143, prepared by Mathematica, Chapter III, Princeton 1968, 25–108.

Baños, A.: On Pseudo-Games. The Annals of Mathematical Statistics 39, 1968, 1932–1945.

Feller, W.: An Introduction to Probability Theory and its Applications. Vol. 1, New York 1957.

Harsanyi, J.C.: Games with Incomplete Information Played by 'Bayesian' Players. Management Science **14**, 1967–68, 159–182, 320–334, 486–502.

Kohlberg, E.: Optimal Strategies in Repeated Games with Incomplete Information. International J. Game Theory 4, 1975a, 7–24.

– : The Information Revealed in Infinitely Repeated Games of Incomplete Information. International J. Game Theory 4, 1975b, 57–59.

Megiddo, N.: On Repeated Games with Incomplete Information Played by Non-Bayesian Players. Discussion Paper No. 373, The Center for Mathematical Studies in Economics and Management Sciences, Northwestern University, March 1979.

Mertens, J.F.: The Value of Two-Person Zero-Sum Repeated Games: The Extensive Case. International J. Game Theory **1**, 1971–72, 217–227 continued in **2**, 1973, 231–234.
Mertens, J.F., and *S. Zamir*: The Value of Two-Person Repeated Games with Lack of Information on Both Sides. International J. Game Theory **1**, 1971–72, 39–64.
Ponssard, J.P., and *S. Zamir*: Zero-Sum Sequential Games with Incomplete Information. International J. Game Theory **2**, 1973, 99–107.
Stearns, R.E.: A Formal Concept of Games of Incomplete Information. Report to the U.S. Arms Control and Disarmament Agency, Washington, D.C. Final report on Contract ACDA/ST-116 prepared by Mathematica, Chapter IV, Princeton 1967, 405–433.
Titchmarsh, E.G.: The Theory of Functions. 2nd Edition. Oxford 1939, 226–229.
Widder, D.V.: The Laplace Transform. Princeton 1941.
Zamir, S.: On the Relation Between Finitely and Infinitely Repeated Games with Incomplete Information. International J. Game Theory **1**, 1971–72, 179–198.
– : On Repeated Games with General Information Function. International J. Game Theory **2**, 1973, 215–229.

Received April, 1979
(revised version April, 1980)