

Discussion Paper No. 686

A CHARACTERIZATION OF SEQUENTIAL
EQUILIBRIUM STRATEGIES IN INFINITELY REPEATED
INCOMPLETE INFORMATION GAMES

by

James Bergin*

April 1986

*Department of Managerial Economics and Decision Sciences, J. L. Kellogg
Graduate School of Management, Northwestern University, Evanston, Illinois
60201.

Thanks are due to Vijay Krishna, Harold Kuhn, J. F. Mertens, Hugo
Sonnenschein and Joe Stiglitz for many helpful conversations. All errors are
my own.

Abstract

For two-sided information infinitely repeated incomplete information games it is shown that any vector payoff can be sustained as a sequential equilibrium vector payoff if and only if there is a markov chain (with a state space of player type distributions and vector payoffs) starting at that vector payoff and satisfying a set of incentive compatibility constraints.

1. Introduction

This paper is concerned with providing a characterization of equilibrium strategies in infinitely repeated incomplete information games with discounting of payoffs. Attention is restricted to strategies which constitute sequential equilibria.

In general, in these games, equilibrium strategies may be complicated functions on the set of histories. A major problem is providing a good economic interpretation of these strategies. Is there a simple or intuitively appealing description that can be given to equilibrium strategies--a simple way of describing how players play? That question is the subject of this paper.

This question arises in any study of repeated games. A substantial part of the literature has focused on a class of equilibria sustained by threat of punishment in response to defection. For reasons to be described below, that approach is not adopted here. The major result of the paper is to show that any payoff vector arising from a sequential equilibrium may be sustained by strategies of a simple form. In the next few pages, some of the literature on repeated games is reviewed. Two types of information structure (complete and incomplete information) and two types of payoff criterion (limit of means and discounting) are considered. This review will illustrate some of the approaches taken and difficulties involved in characterizing equilibrium strategies. In section 2 the main result is given. Section 3 provides a discussion of the conditions under which the result can be strengthened and of the difficulties in doing so, in the context of zero-sum games. Section 4 discusses the extent to which the results of section 2 can be extended.

For games of complete information where payoff streams are evaluated

using the limit of means criterion, Rubinstein [9] has characterized the sets of equilibrium strategies and payoffs. He assumes that players do not randomize on pure strategies or if they do that mixed strategies are observable. This is a critical assumption for two reasons.

First, if the game has more than two players and mixed strategies are not observable (the normal assumption in game theory) then the characterization of perfect equilibria is incomplete. To see this, let the set of players be N and denote player i 's set of mixed strategies and payoff function in the one stage game by X_i and u_i . Also, let $X_{-i} = \prod_{j \neq i} X_j$. In a one stage game, player i 's payoff can be forced down to the level v_i where v_i is called the individually rational level and defined by

$$v_i = \min_{x_{-i} \in X_{-i}} \max_{x_i \in X_i} u_i(x_{-i}, x_i)$$

Call $x \in \prod X_i$ a weakly forced outcome if, for all $i \in N$, $\exists x^i \in \prod X_i$ such that

$$u_i(x_{-i}^i, \tilde{x}_i) \leq u_i(x), \forall \tilde{x}_i \in X_i.$$

The result of Rubinstein is that the set of perfect equilibrium payoffs are the payoffs that correspond to weakly forced outcomes. This result was also given by Aumann and Shapley [6] (see Fudenberg and Maskin [7] for a brief survey of this literature). The idea behind this result is that if a player can gain by defecting from a weakly forced outcome, other players have sufficient punishment power over finite periods of time to remove the gain of the defector with a credible threat. Similarly, a second layer of punishments applies to "punishers" who defect from implementing a punishment strategy. They also are credibly threatened with punishment over a finite period of time

and so on.

When attention is restricted to pure strategies or where mixed strategies are observable, this procedure works. However, if mixed strategies are not observable and punishment of player i requires randomization by j , say, then any defection by j from punishment of i within the support of j 's punishing strategy is undetectable. Since the approach depends on defections being detectable, it breaks down at the second layer of punishments. Thus, when player randomizations are not observable the idea of detection and punishment runs into serious difficulty.

The second reason why the exclusion of mixed strategies is an important assumption relates to communication. There is a communication aspect to strategies (e.g., a defection is communicated by an observation of a move that occurs in equilibrium with probability zero). If only pure strategies are allowed, the scope for communication is greatly reduced. Since the way communication occurs is itself of interest, this is an important omission. Moreover, in incomplete information games where very precise amounts of information must sometimes be transmitted through strategies, pure strategies are insufficient to achieve the degree of communication required.¹ Even in complete information games there is an interesting role for communication.²

Turning to infinitely repeated complete information games with discounting, Abreu [1] has provided a simple characterization of all pure strategy perfect equilibria. The central idea is that punishments which sustain equilibria are history independent. Furthermore, defection by some player at any point in the game is always met by the same punishment. This punishment corresponds to that players' payoff in the worst possible perfect equilibrium--that perfect equilibrium yielding the player the lowest payoff. Let v_i be lowest perfect equilibrium payoff to player i . A given equilibrium

is sustained by the following system of punishments. If player i defects, the defection leads to a subgame with an equilibrium yielding v_i to player i . If some player j defects from this equilibrium the defection leads to a subgame with the equilibrium on that subgame yielding v_j to player j . Layers of punishments are applied in this way and in this way all pure strategy perfect equilibria are sustained in a relatively simple way.

Furthermore, in the context of symmetric games,³ focusing on the set of symmetric strategy subgame perfect equilibria, he shows that the structure of punishment strategies is of a simple form. (In the context of symmetric strategy equilibria, punishment strategies are understood to be symmetric, so that a defection by any player leads to each player playing the same pure strategy on the induced subgame.) All symmetric strategy equilibria can be sustained by a two-phase punishment strategy. The first phase lasts for one period and in this phase each player receives a low payoff. The second and successive periods give each player a high payoff. Within the class of symmetric punishments, a punishment of this form yields the lowest possible equilibrium payoff on any subgame. The important insights given by Abreu have been developed in a different context by Abreu, Pearce and Stacchetti [2]. They focus on the characterization of symmetric strategy sequential equilibria using extreme point payoffs in the context of an oligopoly model. Their paper will be discussed (briefly) at the end of this section.

In the study of complete information games a central idea has been that of sustaining equilibrium by threat of off the equilibrium path punishment. This approach works because: (a) mixed strategies are avoided, (b) the games studied have subgames so that "punishing" equilibria can sensibly be attached to subgames, and (c) punishment is well defined. Turning to incomplete information games each of these points raises substantial difficulty. As has

been pointed out earlier, communication has a central role in incomplete information games. Without randomization it is impossible, for example, to have the type of signalling that arises naturally in incomplete information games. Restricting attention to pure strategy equilibria would clearly be inappropriate. With regard to point (b), incomplete information repeated games have no subgames. This raises difficulties in discussing off-the-equilibrium-path strategies. Finally, in incomplete information games, "punishment" may not be well defined. When a player may be one of many possible types, the exact identity unknown to other players, a strategy which would punish one type might not punish some other type. This leads naturally to the study of punishment in terms of vector payoffs, each element of the vector corresponding to a payoff to some player type. However, even utilizing the idea of vector payoffs, minmax type punishments may not be well defined in terms of the player type payoffs. This raises substantial difficulties, which are discussed at length in section 3. A simple example will illustrate the problem here. Consider a game of incomplete information with payoff matrices for player I (player I's player type payoff matrices).

$$A^1 = \begin{vmatrix} 3 & 1 & 0 \\ 2 & 1 & 0 \end{vmatrix} \quad A^2 = \begin{vmatrix} 2 & 0 & 1 \\ 3 & 0 & 1 \end{vmatrix}$$

Player I is the informed player and selects the row. To punish player I, player II should avoid playing the first column. If the prior distribution over player types is $(1/2, 1/2)$, any distribution over columns 2 and 3 played by player II minimizes the expected payoff to player I (giving an expected payoff to player I of $1/2$). However, if column 2 is played by player II with probability y , the vector payoff to player I is $(y, 1 - y)$. Thus, any vector payoff of the form $(y, 1 - y)$ may arise from a minmax strategy by player II.

It should be clear that a complete description of either equilibrium payoffs or punishment payoffs requires not only knowledge of the expected payoff to a player, but also the payoffs to each player type.

The study of infinitely repeated incomplete information games was begun by Aumann and Maschler [4]. Their results are summarized in section 3 in the discussion of zero sum games. Two papers dealing with the structure of equilibrium strategies will be discussed here. Both of these papers (Hart [6] and Ponsard and Sorin [8]) are quite involved. The discussion here is quite inadequate and intended simply to indicate the kind of approaches they have adopted.

In [6], Hart characterizes all the equilibrium points of two player incomplete information repeated games with lack of information on one side. The characterization is quite complex with the set of equilibrium points corresponding to a certain class of stochastic process. Periods in the game are divided into communication periods and payoff accumulation periods. At points where communication occurs (typically requiring randomization), payoffs are such that, the alternative paths from which the players can choose yield at most the same payoff as that given along an equilibrium path. In payoff accumulation phases, the strategies are pure. Thus, where undetectable derivations are possible, they are also worthless to the potential defector. This approach depends critically on the use of the limit of means criterion. For example, some equilibria may require an unbounded number of communications but with payoffs during periods of communication having no effect on the overall average payoff. This discussion does not at all describe Hart's characterization; it is intended merely to indicate aspects of his approach.

In [8], Ponsard and Sorin provide a description of equilibrium strategies for finitely repeated incomplete information zero sum games where

players move sequentially and where there is a two sided information structure (player I has K possible types and player II has R possible types). They provide a state variable characterization—strategies are determined by the state variables, vector payoffs and posterior distributions. Their procedure is to partition the simplices to which the player type distributions belong, into sets of convex polyhedra on which the value function is piecewise bilinear with respect to the player type distributions, p and q . The state variables are then defined using this property. The existence of such partitions and the bilinearity property typically do not hold in infinitely repeated games.

Abreu, Pearce and Stacchetti [2], have provided a simple characterization of all symmetric strategy sequential equilibria of an oligopoly game, in terms of extreme point payoffs. The game can briefly be described as follows: in the first period each player chooses a quantity, in successive periods each player chooses a quantity having observed the history of prices and the quantities supplied by that player only in previous periods. Thus, no player ever observes the action histories of other players and sufficient random disturbance is introduced into prices so that such information cannot be inferred from prices. Focusing on pure strategies, they show that all symmetric strategy sequential equilibrium payoffs can be sustained by strategies which take on only two possible values after the first period and these values are independent of the equilibrium being sustained.

In the discussion here on incomplete information games two of the themes in the literature discussed above will be developed. These are the state variable stochastic process approach of Hart and Ponsard and Sorin and the extreme point approach of Abreu and Abreu, Pearce and Stacchetti. In the following section it will be shown that the set of sequential equilibria

correspond to a class of stochastic process. To any point in the set of sequential equilibrium payoffs there exists a Markov chain starting at that point and satisfying a set of "incentive compatible" constraints. Conversely, any Markov chain satisfying these incentive compatibility constraints defines a sequential equilibrium. This may be compared with Hart's [6] characterization of equilibrium points as G-processes (see also Aumann and Hart [3]).

The latter part of the paper discusses extreme point transition functions for "equilibrium Markov chains." Unlike complete information games, restricting attention to symmetric games is unnatural as this imposes restrictions on the values of the prior distributions. In the general case (nonsymmetric games) an upper bound is obtained for the number of points in the range of the transition function that lie in the interior of the set of sequential equilibria.

II. A Characterization of Sequential Equilibria

The game to be discussed is the standard two sided incomplete information game. (The result to follow does not depend upon the assumption that there are two players. This is purely for notational convenience.)

A pair of matrices (A^{kr}, B^{kr}) is chosen from the set $\{A^{kr}, B^{kr} \mid k \in K, r \in R\}$ according to the distribution $(p, q) = ((p^1, \dots, p^K), (q^1, \dots, q^R))$. The pair (A^{kr}, B^{kr}) is chosen with probability $p^k q^r$. Player I is informed of the choice of k and player II is informed of the choice of r . The game chosen is played repeatedly, payoffs are discounted and at each stage both players are informed of the previous history of moves. A history of length $t - 1$, denoted h_t , is a sequence of moves $h_t = (i_1, j_1, \dots, i_{t-1}, j_{t-1})$. The set of all histories of length $t - 1$ is denoted H_t . Strategies in the game are defined as collections of functions:

For Player I

$$\sigma = (x_1, x_2, \dots, x_t, \dots)$$

with $x_t: H_t \times K \rightarrow \Delta^I$, $H_1 = \{\emptyset\}$.

For Player II:

$$\tau = (y_1, y_2, \dots, y_t, \dots)$$

with $y_t: H_t \times R \rightarrow \Delta^J$. (Δ^I , Δ^J , $I - 1$ and $J - 1$ dimensional simplexes.)

More concise notation may also be used by defining a strategy for I as a function σ with

$$\sigma: H \times K \rightarrow \Delta^I, H = \bigcup_{t=1}^{\infty} H_t$$

and a strategy for player II:

$$\tau: H \times R \rightarrow \Delta^J$$

Next, let $H_{\infty} = \prod_{t=1}^{\infty} (I \times J)$. Histories in H_{∞} are denoted h . Given a history $h \in H_{\infty}$ and a pair of types (k,r) , the payoff to player I is denoted

$$a^{kr}(h) = (1 - \delta_1) \sum_{t=1}^{\infty} \delta_1^{t-1} a_{i_t j_t}^{kr}, \quad A^{kr} = \{a_{ij}^{kr}\}_{i \in I, j \in J}$$

and the payoff to player II is denoted

$$b^{kr}(h) = (1 - \delta_2) \sum_{t=1}^{\infty} \delta_2^{t-1} b_{i_t j_t}^{kr}, \quad B^{kr} = \{b_{ij}^{kr}\}_{i \in I, j \in J}.$$

A pair of strategies and a pair of prior distributions (p, q) determine a probability distribution on $(H_{\infty} \otimes 2^{K \times R})$. Denote this distribution P and the corresponding expectation operator by E . Thus, given a pair of strategies and a pair of priors (σ, τ, p, q) , the expected payoff to player I is

$$E(a^{kr})$$

and the expected payoff to player II is

$$E(b^{kr})$$

Similarly, the expected payoff to player I, type k is

$$E\{a^{kr} | k\} = \xi^k$$

and the expected payoff to player II, type r is

$$E\{b^{kr} | r\} = \zeta^r$$

Let $\xi = \{\xi^k\}_{k \in K}$, $\zeta = \{\zeta^r\}_{r \in R}$ and note that

$$E(a^{kr}) = p \cdot \xi \quad \text{and} \quad E(b^{kr}) = q \cdot \zeta.$$

Since an understanding of the extensive form of the game is essential to the discussion to follow, a brief description of relevant aspects of the

extensive form is given here. The game has no subgames. Each history $h_t \in H$, leads to a minimal subform of the game. Each history $h \in H_\infty$ and pair (k,r) define a path on the extensive form. The set of paths in the minimal subform associated with h_t is

$$\bigcup_{\substack{k \in K \\ r \in R \\ h \in \prod_{s=1}^{\infty} (I \times J)}} (k,r,h_t,h) \subset K \times R \times H_\infty$$

Any information set reached by a history in H of length t periods or more and through which a path in the minimal subform passes is contained entirely in the minimal subform. For any path not in the minimal subform, no information set through which such a path passes, intersects the minimal subform. This is a heuristic description; further details are given in Selten [10] or Kreps-Wilson [4]. The intuitive idea is that after any history, h_t , the corresponding minimal subform is the smallest detachable set of branches of the extensive form (detachable in the sense that neither paths nor information sets are broken from there on). If to any minimal subform, a distribution is attached to the initial nodes of the subform, a replica of the original incomplete information game is defined.

The sequential equilibrium concept has been defined only for games with a finite number of pure strategies. The games considered here have a continuum of pure strategies but the extension to this situation is natural. Denote a strategy tuple, $\omega = (\sigma, \tau)$. We can write $(\omega_t, \{\omega(h_t)\}_{h_t \in H_t})$, with the understanding that ω_t defines the strategies on $\bigcup_{s=1}^{t-1} H_s$ and $\omega(h_t)$ is the "strategy" induced by the history h_t on the minimal subform defined by h_t . Adding type distributions to a minimal subform defines an infinitely repeated incomplete information game (denote from now on the minimal subform associated with h_t , $F(h_t)$). With posterior distributions $p(h_t), q(h_t)$ (defined for all

$h_t \in H$) denote such a subform game $G(h_t) = (F(h_t), p(h_t), q(h_t))$. With this notation, a sequential equilibrium may be defined. A sequential equilibrium is an equilibrium strategy tuple ω , such that for every $h_t \in H$, $\omega(h_t)$ is an equilibrium on the subform game $G(h_t)$. The posterior distributions $p(h_t)$, $q(h_t)$ are required to satisfy Bayesian consistency in the following way. If $\text{prob}(h_t) > 0$, then $p(h_t)$, $q(h_t)$ is defined by ω and the prior distributions. If $\text{prob}(h_t) = 0$, then given $p(h_t)$, $q(h_t)$, $p(h_t, i_t, j_t)$ is defined as

$$p^k(h_t, i_t, j_t) = \frac{x_{ti_t}^k(h_t) p^k(h_t)}{\text{Prob}(i_t | h_t)}, \quad \forall k \in K$$

if

$$\text{Prob}(i_t | h_t) = \sum p^k(h_t) x_{ti_t}^k(h_t) > 0.$$

If $\text{prob}(i_t | h_t) = 0$, then a sequence $x^n \in \text{int}(\Delta^I)^K$ is chosen with $x^n \rightarrow x(h_t)$. The sequence x^n defines a sequence of posteriors $p^n(h_t, i_t, j_t)$, the limit of which (taking subsequences if necessary) is defined to be $p(h_t, i_t, j_t)$. The posterior $q(h_t, i_t, j_t)$ is defined similarly. Define the set

$$\Delta = \prod_{t \geq 1} \prod_{h_t \in H_t} (\Delta^K \times \Delta^R)$$

Thus a sequential equilibrium is a pair $(\omega, \lambda) \in \Omega \times \Delta$ ($\Omega = \Sigma \times T$) inducing equilibria on every subform game and satisfying Bayesian consistency in the posterior sequence. Next, let

$$M = \max_{i, j, k, r} \{ |a_{ij}^{kr}|, |b_{ij}^{kr}| \}$$

and define the set of states of the game to be the set

$\bar{S} = \Delta^K \times \Delta^R \times [-M, M]^{K+R}$. A state is thus a vector $s = (p, q, \xi, \zeta)$ where (p, q)

are the priors and (ξ, ζ) are vector payoffs to players I and II.

The following propositions and theorems define a state space for the game and convert every sequential equilibrium to a first period payoff equivalent sequential equilibrium with representation as a Markov chain with stationary transition probabilities and a stationary transition function.

Define the correspondence $\phi(p, q)^4$ as

$$\phi(p, q) = \{\omega = (\sigma, \tau) \in \Omega \mid \omega \text{ is a equilibrium strategy pair}\}$$

ϕ is the equilibrium correspondence from priors to strategies.

Proposition 2.1: ϕ is an upper hemicontinuous (u.h.c.) correspondence from $\Delta^K \times \Delta^R$ to Ω .

Proof: Denote the probability distribution on $H_\omega \otimes 2^{K \times R}$ determined by (p, q, ω) as $P_{pq\omega}$, determining the expectation operator $E_{pq\omega}$. Observe that $E_{pq\omega}(a^{kr})$, $E_{pq\omega}(b^{kr})$, $E_{pq\omega}(a^{kr} \mid k)$ and $E_{pq\omega}(b^{kr} \mid r)$ are continuous functions of (p, q, ω) . It follows immediately that ϕ is u.h.c. \square

Define $\phi_S(p, q)$ the sequential equilibrium correspondence,

$$\phi_S: \Delta^K \times \Delta^R \rightarrow \Omega$$

$$\phi_S(p, q) = \{\omega \in \Omega \mid \forall h_t \in H, \omega(h_t) \in \phi(p(h_t), q(h_t))\}$$

Note $(p(h_t), q(h_t))$ is defined for every history (see the earlier discussion on this matter).

Proposition 2.2: ϕ_S is an u.h.c. correspondence from $\Delta^K \times \Delta^R$ to Ω .

Proof: Let $(p^n, q^n) \rightarrow (p, q)$ and let $\omega^n \in \phi_S(p^n, q^n)$ with corresponding sequence

$(P^n, Q^n) \in \Delta$. Let $\omega_n \rightarrow \omega$ and $(P^n, Q^n) \rightarrow (P, Q)$. For any history $h_t \in H$,

$$\omega^n(h_t) \in \phi(p^n(h_t), q^n(h_t)).$$

Since $\omega^n(h_t) \rightarrow \omega(h_t)$, $(p^n(h_t), q^n(h_t)) \rightarrow (p(h_t), q(h_t))$ and ϕ is u.h.c. it follows that for any $h_t \in H$

$$\omega(h_t) \in \phi(p(h_t), q(h_t)). \quad \square$$

Next, define a correspondence α , $\alpha: \Delta^K \times \Delta^R \rightarrow \bar{S}$, as

$$\alpha(p, q) = \{(\xi, \zeta) \in [-M, M]^{K+R} \mid \exists \omega \in \phi_S(p, q)$$

$$\text{with } \xi^k = E_{pq\omega}(a^{kr} \mid k), \zeta^r = E_{pq\omega}(b^{kr} \mid r) \text{ for all } (k, r) \in K \times R\}$$

Proposition 2.3: α is an u.h.c. correspondence from $\Delta^K \times \Delta^R$ to \bar{S} .

Proof: This follows immediately from the fact that ϕ_S is an u.h.c. correspondence and $E_{pq\omega}(a^{kr} \mid k)$, $E_{pq\omega}(b^{kr} \mid r)$ are continuous functions of (p, q, ω) . \square

Next, a set of sequential equilibrium states is defined. Let

$$S = \{(p, q, \xi, \zeta) \in \Delta^K \times \Delta^R \times [-M, M]^{K+R} \mid (\xi, \zeta) \in \alpha(p, q)\}.$$

The set S is the basic state space for the discussion to follow. Observe that since α is a u.h.c. correspondence, S is closed.

Proceeding, define random variables on the underlying space $H_\infty \otimes 2^{K \times R}$ as

follows. Let

$$a_t^{kr} = (1 - \delta_1) \sum_{s=t}^{\infty} \delta_1^s a_{i_s j_s}^{kr}$$

$$b_t^{kr} = (1 - \delta_2) \sum_{s=t}^{\infty} \delta_2^s b_{i_s j_s}^{kr}$$

and for fixed ω , let

$$\xi_t^k(h_t) = E_{pq\omega} \{a_t^{kr} | h_t, k\}, \forall k \in J$$

$$\zeta_t^r(h_t) = E_{pq\omega} \{b_t^{kr} | h_t, r\}, \forall r \in R.$$

Note that if ω is a sequential equilibrium, then for all $h_t \in H$,

$$\xi_t^k(h_t) = \sup_{\sigma \in \Sigma} E_{pq\sigma\tau} \{a_t^{kr} | h_t, k\}, \forall k \in K$$

and

$$\zeta_t^r(h_t) = \sup_{\tau \in \Gamma} E_{pq\sigma\tau} \{b_t^{kr} | h_t, r\}, \forall r \in R.$$

Thus, for example, $\xi_t^k(h_t)$ is the most that player I, type k can achieve from period t on, given the strategy of player II, if ω is a sequential equilibrium. (This is true for some version of the conditional distribution, if $\text{Prob}(h_t) = 0$ under (p, q, ω) .)

For a fixed sequential equilibrium ω , a state at time t following any history $h_t \in H$ is a vector of the form:

$$s_t(h_t) = (p(h_t), q(h_t), \xi_t(h_t), \zeta_t(h_t))$$

with

$$(\xi_t(h_t), \zeta_t(h_t)) \in \alpha(p(h_t), q(h_t)).$$

Define $S_t(p,q,\omega)$ as follows (ω a sequential equilibrium):

$$S_t(p,q,\omega) = \{s \in S \mid \exists h_t \in H_t, s_t(h_t) = s, \\ \text{with the function } s_t(\cdot) \text{ determined by } (p,q,\omega)\}$$

and define $S(p,q,\omega)$

$$S(p,q,\omega) = \bigcup_{t \geq 1} S_t(p,q,\omega) \\ = \bigcup_{t \geq 1} \bigcup_{h_t \in H_t} s_t(h_t)$$

For any fixed sequential equilibrium, ω , given priors (p,q) , $S(p,q,\omega)$ is the set of states that can be reached through any history $h_t \in H$ (with h_t possibly having zero probability of being reached) given the sequential equilibrium strategies $\omega = (\sigma,\tau)$. Observe that for any $\omega \in \Omega$, $S(p,q,\omega)$ has at most a countable number of states.

The content of Theorem 2.4 (see below) is that a sequential equilibrium may be replaced by another first period payoff equivalent sequential equilibrium with a simpler structure. The procedure is as follows: take a sequential equilibrium of the game with priors (p,q) which yield vector payoffs (ξ,ζ) and iteratively construct a sequence $\{\omega^n\}$ of sequential equilibria which yield the vector payoffs (ξ,ζ) . The limit of $\{\omega^n\}$, ω^* say, will be characterized with a state space interpretation which may be described in the following way. From the perspective of player I, type k , at period 1 player I_k computes the state s_1 and plays $x^k(s_1)$. History (i_1, j_1) occurs, the

player now computes $T(s_1, i_1, j_1) = s_2$ and plays $x^k(s_2)$. At time t , the player plays $x^k(s_t)$, computes $s_{t+1} = T(s_t, i_t, j_t)$ if (i_t, j_t) occurs and he plays $x^k(s_{t+1})$ in the following period.

Throughout the following discussion, the priors (p, q) are held fixed. To simplify notation, they will be suppressed and the set $S_t(p, q, \omega)$ will be denoted $S_t(\omega)$.

Theorem 2.4. Let ω be a sequential equilibrium, yielding expected payoffs (ξ, ζ) . There is a payoff equivalent sequential equilibrium (i.e., yielding the same expected payoffs), ω^* , where ω^* is characterized by a pair of functions (x, y) such that

$$(x, y): S(\omega^*) \rightarrow (\Delta^I)^K \times (\Delta^J)^R$$

and with the interpretation that if history h_t occurs at time t , player I type k plays the strategy $x^k(s_t(h_t))$; player II type r plays $y^r(s_t(h_t))$.

Furthermore, the process s_t , $t \geq 1$, determined by ω^* , is a Markov chain with stationary transition probabilities.

Proof: Let ω be a sequential equilibrium. A sequence of payoff equivalent sequential equilibria $\{\omega^l\}_{l \geq 1}$ will be constructed from ω with the property that $\omega^l \rightarrow \omega^*$, where ω^* has the desired property.

Set $\omega = \omega^1$, the first element of the sequence. Note that as ω^l changes, the state function $s_t(h_t)$ will change, as it is determined by ω^l . To indicate this dependency, write $s_{\lambda t}(h_t) \in S_t(\omega^l)$.

Given ω^1 , ω^2 is defined as follows.

At $t = 2$, for each $s \in S_2(\omega^1) \cap S_1(\omega^1)$,

put $\omega^2(h_2) = \omega^1$, $\forall h_2 \in s_{12}^{-1}(s)$.

For each $s \in S_2(\omega^1) \setminus S_1(\omega^1)$, pick $\bar{h}_2 \in s_{12}^{-1}(s)$ and

put $\omega^2(h_2) = \omega^1(\bar{h}_2)$, $\forall h_2 \in s_{12}^{-1}(s)$.

Define $\omega^2 = (\omega_1^1, \{\omega^2(h_2)\}_{h_2 \in H_2})$

At period t , for each $s \in S_t(\omega^{t-1}) \cap [\bigcup_{\tau=1}^{t-1} S_\tau(\omega^{t-1})]$

Pick some $\bar{h}_\tau \in s_{t-1,\tau}^{-1}(s)$, $\tau < t$, and

put $\omega^t(h_t) = \omega^{t-1}(\bar{h}_\tau)$, $\forall h_t \in s_{t-1,\tau}^{-1}(s)$

for each $s \in S_t(\omega^{t-1}) \setminus [\bigcup_{\tau=1}^{t-1} S_\tau(\omega^{t-1})]$

Pick some $\bar{h}_t \in s_{t-1,t}^{-1}(s)$

and put $\omega^t(h_t) = \omega^{t-1}(\bar{h}_t)$, $\forall h_t \in s_{t-1,t}^{-1}(s)$

Define $\omega^t = (\omega_{t-1}^{t-1}, \{\omega^t(h_t)\}_{h_t \in H_t})$

Two results will now be established, relating ω^t to ω^{t-1} .

Proposition 1. If ω^{t-1} is a sequential equilibrium, then ω^t is a sequential equilibrium.

Proof. If $\omega^t = \omega^{t-1}$ there is nothing to prove. Otherwise, $\omega^t \neq \omega^{t-1}$, so

$\omega^t(h_t) \neq \omega^{t-1}(h_t)$ for some $h_t \in H_t$. Therefore, $\exists \bar{h}_\tau$, $\tau \leq t$ and $\omega^t(h_t) = \omega^{t-1}(\bar{h}_\tau)$.

For this substitution to have occurred requires that $s_{t-1,t}(h_t) = s_{t-1,\tau}(\bar{h}_\tau)$. This implies that $(p_{t-1,t}(h_t), q_{t-1,t}(h_t)) = (p_{t-1,\tau}(h_\tau), q_{t-1,\tau}(h_\tau))$. Thus the subform games $G(h_t)$ and $G(\bar{h}_\tau)$ are identical. Consequently, if $\omega^{t-1}(\bar{h}_\tau)$ is a sequential equilibrium of $G(\bar{h}_\tau)$ then it is a sequential equilibrium of $G(h_t)$.

Thus, on minimal subforms $F(h_t)$, where strategy replacement occurs, the new strategies define sequential equilibria on those minimal subforms. On those minimal subforms where no replacement has occurred, the induced strategies are by definition sequential equilibria.

Thus, given the posterior distributions determined over H_t by ω_{t-1}^{t-1} , $\{\omega^t(h_t)\}_{h_t \in H_t}$ is optimal on each corresponding subform game. It remains to check that ω_{t-1}^{t-1} is optimal after replacement has occurred. Since $s_{t-1,t}(h_t) = s_{t,\tau}(h_t)$, $\forall h_t \in H_t$, the expected payoff to each player type is unchanged, so given the strategies of the others, a player type has no incentive to deviate at any history $h_{t-1} \in H_{t-1}$, given that he did not have such incentive before the replacement took place.

Proposition 2. Given ω^t , if $s_{tt'}(h_{t'}) = s_{t\tau}(h_\tau)$, $t', \tau < t$, then $\omega^t(h_{t'}) = \omega^t(h_\tau)$.

Proof. This follows from the construction of ω^t .

Continuing with the proof of the theorem, let ω^t or some subsequence converge: $\omega^t \rightarrow \omega^*$. Then ω^* is a sequential equilibrium. Define

$$S(\omega^*) = \bigcup_{t \geq 1} S_t(\omega^*).$$

Let s_t , $t \geq 1$ be the state function determined by ω^* . Note that

$$S(\omega^*) = \bigcup_{t \geq 1} \bigcup_{h_t \in H_t} s_t(h_t)$$

Suppose now, that at $h_t, h_\tau, s_t(h_t) = s_\tau(h_\tau)$. Then by the second proposition

$$\omega^*(h_t) = \omega^*(h_\tau)$$

In particular,

$$x_t^k(h_t) = x_\tau^k(h_\tau), \quad \forall k$$

$$y_t^r(h_t) = y_\tau^r(h_\tau), \quad \forall r$$

Define the functions $(x,y) = ((x^1, \dots, x^K), (y^1, \dots, y^R))$ on $S(\omega^*)$ by picking for each $s \in S(\omega^*)$ some h_t with $s_t(h_t) = s$. Then, put

$$x^k(s) = x^k(s_t(h_t)) = x_t^k(h_t), \quad h_t \in H$$

and

$$y^r(s) = y^r(s_t(h_t)) = y_t^r(h_t), \quad \text{for some } h_t \in H.$$

The functions, (x,y) , acting on the states, s_t , sustain the sequential equilibrium ω^* .

Turning to the characterization of states, suppose that $s_t(h_t) = s_\tau(h_\tau)$. Then $\omega^*(h_t) = \omega^*(h_\tau)$ so that

$$s_t(h_t, i, j) = s_\tau(h_\tau, i, j)$$

and

$$\text{prob}(i, j | h_t) = \text{prob}(i, j | h_\tau)$$

Thus, for any $s \in S(\omega^*)$, there are $I \times J$ possible successors. These are determined by picking some $h_t \in H$, with $s_t(h_t) = s$. The possible successors of s are $s_t(h_t, i, j)$, $i \in I$, $j \in J$. One can therefore define a transition function T , where

$$T: S(\omega^*) \times I \times J \rightarrow S(\omega^*)$$

If the current state is s and the "history" (i, j) then occurs, the new state is $T(s, i, j)$. If the current state at time t is s_t , then the probability of (i, j) is $(\sum_k p_t^k x_i^k(s_t))(\sum_r q_t^r y_j^r(s_t))$ which may be written as $\text{prob}(i, j | s_t)$. The probability that \bar{T} is reached next, given the current state s_t is equal to

$$\sum_{\{(i, j) | T(s_t, i, j) = \bar{T}\}} \text{prob}(i, j | s_t)$$

Finally, since $S(\omega^*)$ is countable, $\{s_t\}_{t \geq 1}$ is a Markov chain with stationary transition probabilities.

This completes the proof. \square

The function T obtained in Theorem 2.4 may be decomposed into two parts, $T = (u, v)$ with

$$u: S(\omega^*) \times I \times J \rightarrow \Delta^K \times \Delta^R$$

$$v: S(\omega^*) \times I \times J \rightarrow \alpha(u(S(\omega^*) \times I \times J))$$

u is the posterior updating function, $\forall (s, i, j) \in S(\omega^*) \times I \times J$,

$$u(s,i,j) = (\{p^k(s,i,j)\}_{k \in K}, \{q^r(s,i,j)\}_{r \in R})$$

$$s = (p,q,\xi,\zeta)$$

and

$$p^k(s,i,j) = \frac{x_i^k(s)p^k}{\sum_k p^k x_i^k(s)}, \text{ if } \sum_k p^k x_i^k(s) > 0$$

$$p^k(s,i,j) = \lim_{n \rightarrow \infty} \frac{x_i^{nk} p^k}{\sum_k p^k x_i^{nk}}, x^n \in \text{int}(\Delta^I)^K, x^n \rightarrow x(s)$$

(taking limits in some subsequence if necessary) for some sequence $\{x^n\} \in \text{int}(\Delta^I)^K$. A similar procedure applies to $q^r(s,i,j)$.

v assigns a vector payoff to each path (i,j) for each state s .

$$v(s,i,j) \in \mathbb{R}^{K+R}$$

Define functions f^k, g^r as follows:

$$f^k(s,x,y,v) = \{(1 - \delta_1) \sum_r q^r x^k A^{kr} y^r + \delta_1 \sum_{i,j} x_i^k y_j^k v_{ij}^k\}$$

$$g^r(s,x,y,v) = \{(1 - \delta_2) \sum_k p^k x^k A^{kr} y^r + \delta_2 \sum_{i,j} y_j^r x_i^r v_{ij}^r\}$$

$$\bar{x}_i^k = \sum_k p^k x_i^k, \quad \bar{y}_j^r = \sum_r q^r y_j^r$$

and for each (i,j) , v_{ij} is a vector in \mathbb{R}^{K+R} .

Note that f^k, g^r are written as functions of $s = (p,q,\xi,\zeta)$ when only p,q , enter as arguments of f^k, g^r . This is done for notational convenience later.

Define $f = \sum_k p^k f^k$ and $g = \sum_r q^r g^r$.

Observe now that the functions (x,y,T) defined in Theorem 2.4 satisfy

$$T = (u:v): S(\omega^*) \times I \times J \rightarrow S(\omega^*)$$

$$x: S(\omega^*) \rightarrow \Delta_I (= (\Delta^I)^K)$$

$$y: S(\omega^*) \rightarrow \Delta_{II} (= (\Delta^J)^R)$$

and $\forall s = (p, q, \xi, \zeta) \in S(\omega^*)$.

$$(i) \quad \xi^k = f^k(s, x(s), y(s), v(s)) \geq f^k(s, x, y(s), v(s)), \forall x \in \Delta_I, \forall k \in K$$

$$(ii) \quad \zeta^r = g^r(s, x(s), y(s), v(s)) \geq g^r(s, x(s), y, v(s)), \forall y \in \Delta_{II}, \forall r \in R$$

$$(iii) \quad v(s, i, j) \in \alpha(u(s, i, j)), \forall (i, j) \in I \times J$$

where $u(s, i, j) = (p(s, i, j), q(s, i, j))$

$$p(s, i, j) = \frac{x_i^k(s) p^k}{\sum_k p^k x_i^k(s)}, \text{ if } \sum_k p^k x_i^k(s) > 0$$

$$= \lim_{n \rightarrow \infty} \frac{x_i^{nk} p^k}{\sum_k p^k x_i^{nk}}$$

(in some subsequence if necessary) and

$$x^n \in \text{int} \Delta_I, x^n \rightarrow x(s).$$

$q(s, i, j)$ is defined similarly. This leads to a converse of Theorem 2.4.

Theorem 2.5: Given $S' \subseteq S$ (S is the set of sequential equilibrium states).

Let T satisfy

$$T = (u:v): S' \times I \times J \rightarrow S'$$

$$(u: S' \times I \times J \rightarrow \Delta^K \times \Delta^R)$$

if there exist functions x,y such that

$$x: S' \rightarrow \Delta_I$$

$$y: S' \rightarrow \Delta_{II}$$

and for all $s \in S'$ conditions (i)-(iii) are satisfied, then (T,x,y) defines a sequential equilibrium for each $s \in S'$.

Proof: Pick some $s \in S'$ and set $s_1 = s$. Define $x_1 = x(s_1)$, $y_1 = y(s_1)$. Let $s_2(i,j) = T(s_1,i,j)$ and define $x_2(i,j) = x(s_2(i,j))$, $y_2(i,j) = y(s_2(i,j))$. Define strategies recursively in this way so period t strategies are defined $s_t(h_t) = T(s_{t-1}(h_{t-1}), i_{t-1}, j_{t-1})$, $x_t(h_t) = x(s_t(h_t))$, $y_t(h_t) = y(s_t(h_t))$. This gives

$$\sigma = (x_1, x_2, \dots, x_t, \dots)$$

$$\tau = (y_1, y_2, \dots, y_t, \dots)$$

$$x_t: H_t \rightarrow \Delta_I \quad (x_t: H_t \times K \rightarrow \Delta^I)$$

$$y_t: H_t \rightarrow \Delta_{II} \quad (y_t: H_t \times R \rightarrow \Delta^J)$$

Let $s_t(h_t) = (p(h_t), q(h_t), \xi(h_t), \zeta(h_t))$. By construction, for each $k \in K$

$$\begin{aligned} \xi^k(h_t) &= \{(1 - \delta_1) \sum q^r(h_t) x_t^k A^{kr} y^r(h_t) \\ &+ \delta_1 \sum x_{ti_j}^k(h_t) \bar{y}_{tj_t}(h_t) \xi^k(h_t, i_t, j_t)\} \\ \bar{y}_{tj_t}(h_t) &= \sum_r q^r(h_t) y_{tj_t}^r(h_t) \end{aligned}$$

Expanding the expression gives $\xi^k(h_t) = E_{pq\omega} \{a_t^{kr} | h_t, k\}$ with $\omega = (\sigma, \tau)$. Also,

$$\begin{aligned} \xi^k(h_t) &= \max_{x \in \Delta_I} \{(1 - \delta_1) \sum q^r(h_t) x^k A^{kr} y^r(h_t) \\ &+ \delta_1 \sum x_i^k \bar{y}_{tj_t}(h_t) \xi^k(h_t, i_t, j_t)\} \end{aligned}$$

These properties are satisfied (by construction) $\forall h_t \in H$. Therefore, in particular,

$$\xi^k(h_t) = \sup_{\sigma \in \Sigma} E_{pq\sigma\tau} \{a_t^{kr} | h_t, k\}, \quad \forall h_t \in H.$$

A similar discussion applies to $\{\zeta(h_t)\}_{h_t \in H}$.

Finally, observe that for each $k \in K$, $h_t \in H$

$$\begin{aligned} p^k(h_t, i_t, j_t) &= \frac{x_{t1_t}^k(h_t) p^k(h_t)}{\sum p^k(h_t) x_{t1_t}^k(h_t)}, \text{ if } \sum p^k(h_t) x_{t1_t}^k(h_t) > 0 \\ &= \lim_{n \rightarrow \infty} \frac{x_{ti_t}^{nk} p^k(h_t)}{\sum p^k(h_t) x_{ti_t}^{nk}}, \text{ if } \sum p^k(h_t) x_{t1_t}^k(h_t) = 0 \end{aligned}$$

and with $x^n \in \text{int } \Delta_I$, $x^n \rightarrow x_t(h_t)$. Again, a similar discussion applies to

$q(h_t, i_t, j_t)$. Thus, (σ, τ) and the corresponding posterior distributions defined by $s_t(h_t)$, $h_t \in H$ define a sequential equilibrium. \square

To summarize the discussion of this section: every sequential equilibrium generates a (not necessarily unique) Markov chain with stationary transition probabilities; and any triple of functions satisfying conditions (i)-(iii) above defines a class of sequential equilibria.

III. Zero Sum Games

In the previous section, the state vector of variables contained both player type distributions and vector payoffs. Is it possible to simplify the state variable to contain only the player type distributions? A good point to begin the discussion is with a paper of Aumann and Maschler [4]. There, they showed that, for infinitely repeated incomplete information games with lack of information on one side and payoffs evaluated according to the limit of means criterion, the game has a value. However, the value of the game, given the prior distribution, is determined exclusively by a strategy for the uninformed player which is based on Blackwell's approachability theory for games with vector payoffs. In the present context, vector payoffs arise as payoffs to each player type. The ideas in the theory of approachability are important for the present discussion. A simple example will illustrate the theory. First, a brief review of notation is required.

For a zero sum game with two player types (so that there are two matrices A^1 and A^2) and prior, p , let $A(p) = pA^1 + (1 - p)A^2$. The function $u(p)$ defined:

$$u(p) = \max_x \min_y xA(p)y,$$

is the value of the game when the informed player (the maximizer, say) is not allowed to use type dependent information. Aumann and Maschler proved that the value of the infinitely repeated game is $\text{cav } u(p)$, the smallest concave function above $u(p)$. The proof uses Blackwell's approachability theory. A payoff in the game may be viewed as a vector, $a = (a^1, a^2)$ representing the payoff to each player type. Approachability relates to how well the uninformed player can simultaneously "push down" both elements of the vector payoff. A set (here a subset of \mathbb{R}^2) is approachable by the uninformed player if he can push the vector payoff into that set. With this definition, the Aumann-Maschler theorem may be stated (see, for example, Hart [6]): A necessary and sufficient condition for the set $Q = \{x \in \mathbb{R}^2 \mid x \leq a\}$ to be approachable by player 2 is that

$$q \cdot a \geq u(q), \forall q \in \Delta^2$$

In a game with prior p , any payoff vector, a , satisfying

$$q \cdot a \geq u(q), \forall q \in \Delta^2$$

and

$$p \cdot a = \text{Cav } u(p)$$

is an equilibrium payoff vector. The following example illustrates that, in equilibrium, the vector payoff may not be well defined.

Example 1.

$$A^1 = \begin{vmatrix} 2 & 0 \\ 2 & 1 \end{vmatrix} \quad \begin{vmatrix} 0 & 1 \\ 0 & 0 \end{vmatrix}$$

$$\begin{aligned} u(p) &= 2p, & p &\leq 1/3 \\ &= 1 - p, & 1/3 < p &\leq 1/2 \\ &= p, & 1/2 < p. \end{aligned}$$

So

$$\begin{aligned} \text{Cav } u(p) &= 2p, & p &\leq 1/3 \\ &= 1/2 + 1/2p, & p &> 1/3. \end{aligned}$$

These functions are plotted in figure 1. Next, consider the approachable sets defined by these functions. Figure 2 depicts these sets. Note that for $p < 1/3$, there is a unique equilibrium with vector payoff $a = (2,0)$. For $1/3 < p < 1/2$, $a = (2,0)$ is the tangent line to $\text{Cav } u(p)$. For $p > 1/2$ there is a unique vector payoff, $a = (1, 1/2)$, again the tangent line to $\text{Cav } u(p)$. However, at $p = 1/3$ any one of the rectangles in figure 2 is approachable (just to the corner) and whichever one is chosen to approach by player 2, the same expected payoff is obtained. Thus, at $p = 1/3$, any point of the form $\lambda(2,0) + (1 - \lambda)(1, 1/2)$ is an equilibrium payoff with $\lambda \in [0,1]$. The shaded area in figure 3 defines the set of lines corresponding to equilibrium vectors.

In example 2, the same problem occurs at $p = 1/3$. The equilibrium vector payoff here is any vector of the form

$$\lambda(2/3, -1/3) + (1 - \lambda)(4/3, -2/3), \lambda \in [0,1].$$

Example 2

$$A^1 = \begin{vmatrix} 2 & 0 \\ 0 & 1 \end{vmatrix} \quad A^2 = \begin{vmatrix} -1 & 0 \\ 0 & -2 \end{vmatrix}$$

For this example, regardless of the payoff criterion used in the infinitely repeated game, at $p = 1/3$, any vector of the form given above may arise as an equilibrium payoff vector for example, the kink in figure 4 at $p = 1/3$, defines the range of vector payoffs that may arise.

It should be clear from these examples that the "problem" of equilibrium vector payoffs not being well defined arises precisely at those points in the simplex where the value function is not differentiable. Since the focus of this paper is the structure of equilibrium strategies, does it matter that the vector payoffs are not well defined for some player type distributions? Could one not just "select" some vector payoff to be sustained by appropriate strategies whenever a player type distribution is reached, at which the equilibrium vector payoff is not well defined? This line of reasoning is reinforced by the observation that these zero sum games have a recursive structure.

In a two stage game (summing the payoffs), the value function satisfies:

$$v_2(p,q) = \max_x \min_y \{ \sum p^k q^k x^k A^{kr} y^r + \sum \bar{x}_i \bar{y}_j v_1(p(i),q(j)) \}$$

with $\bar{x}_i = \sum p^k x_i^k$

$$p^k(i) = p^k x_i^k / \bar{x}_i, \quad p(i) = p^1(i), \dots, p^K(i)$$

and with similar definitions for $\bar{y}_j, q(j)$.

Similarly, with an infinitely repeated game with discounting, the value function satisfies

$$v(p,q) = \text{Max}_x \text{Min}_y \{ (1 - \delta) \sum_p^k \sum_q^r x^k A^k y^r + \delta \sum_i \bar{x}_i \bar{y}_i v(p(i), q(j)) \}$$

In either of these cases, since the payoff from the second period on is fully determined by the posteriors, one might expect that second period strategies could be chosen optimally and such that those strategies would be the same on histories that lead to the same posteriors. This is clearly a necessary condition for posteriors to serve as state variables. Unfortunately, this is not the case, as the following example shows.

Example 3⁵

$$A^1 = \begin{vmatrix} 2 & 0 \\ 0 & -1 \end{vmatrix} \quad A^2 = \begin{vmatrix} -1 & 0 \\ 0 & 2 \end{vmatrix}$$

The value function of the one shot game is

$$v_1(p) = 2 \text{Min}(p, 1 - p)$$

Figure 5 depicts the set of equilibrium vector payoffs in the one stage game, at $p = 1/2$. The extensive form of the two stage game is given in figure 6. In this game, player I is the informed player and maximizer. Player I plays rows, player II plays columns. There are four possible histories that can occur following the first period. These are: $h_1 = (1,1), h_2 = (2,1), h_3 = (1,2)$ and $h_4 = (2,2)$. Strategies can be identified at each information

set just by giving the probability of playing the first row or column. Thus, for player II at the first period a strategy is given by y and for player I, a strategy is given by (x^1, x^2) . Corresponding to each of the four histories in the second period, player II has four strategies given by y_1, y_2, y_3, y_4 where $y_i = y(h_i)$. Note that in the second and last period, player I has a dominant strategy $(x^1(h_i), x^2(h_i)) = (1, 0), \forall i$, so that player I may be ignored in the second period. Figure 6 makes this description of the game clear. Thus any equilibrium in this game may be described by a vector of the form $(y, (x^1, x^2), (y_1, y_2, y_3, y_4))$. If $p = 1/2$, then the set of equilibrium points in this game is

$$\{(y, (x^1, x^2), (y_1, y_2, y_3, y_4)) \mid y = x^1 = x^2 = 1/2,$$

$$y_1 + y_3 = 1/4, y_2 + y_4 = 3/4\}.$$

Since $y_i \geq 0$, in particular observe that $y_1, y_3 \leq 1/4$ and $y_2, y_4 \geq 3/4$. Observe also that, since $x^1 = x^2 = 1/2$ the posterior on any history, $p(h_i)$ is equal to the prior $p (=1/2)$. In the one stage game with prior $p = 1/2$, the set of equilibrium vector payoffs was:

$$\{(2\alpha, 2(1 - \alpha)) \mid \alpha \in [0, 1]\}.$$

In the repeated game with histories leading from prior $p = 1/2$ in the first period and posteriors all equal to $1/2$ in the second period, the vector payoffs in the second period all lie in the set

$$\{(2\alpha, 2(1 - \alpha)) \mid \alpha \in [0, 1/4] \cup [3/4, 1]\}$$

Figure 7 depicts this set relative to the one stage payoff vector. Any one of the equilibria may be interpreted as follows. Player II plays $y = 1/2$ in the first period, ensuring neither player type of the opponent can get more than 1. Suppose player I type 1 tries to get the most possible in the first period by playing the 1st row. This leads to histories $h_1 = (1,1)$ or $h_2 = (1,2)$ occurring, if player I is type 1. Referring to figure 7, observe that any vector payoff following these histories is "bad" for player I type 1. From player II's perspective, the only player type of I that can gain from playing 1 is player type 1. Therefore, whenever a history occurs in which I plays 1, the payoff is "bad" for type 1 in the next period. A similar interpretation applies with player type 2. There are two points to note from this example: the posteriors are not sufficient as state variables to determine strategies and vector payoffs are central to the description of the equilibria.

Furthermore, it may be seen that the difficulties arise at precisely these points in the simplex where the value function is kinked. If the value function is everywhere differentiable on the simplex, then the vector payoffs associated with any distribution p (in games of one sided information) are uniquely defined. This observation leads to the following theorem.

Theorem 3.1. Let $v(p)$ be the value of an infinitely repeated game of incomplete information with discounted payoffs and prior distribution p . Then, if v is differentiable everywhere on the simplex the posterior distribution is a sufficient state variable to determine the strategies of both players. That is, there exists an equilibrium such that if at h_t and h_τ $p_t(h_t) = p_\tau(h_\tau)$ then $x_t^k(h_t) = x_\tau^k(h_\tau) \forall k \in K$ and $y_t(h_t) = y_\tau(h_\tau)$.

Proof. Let $a(p) = \frac{\partial v}{\partial p}(p)$.

Apply Theorem 2.1, the state variable satisfies

$$s_t = (p_t, a(p_t), p_t \cdot a(p_t))$$

if at t, t' , $p_t = p_{t'}$, then $a(p_t) = a(p_{t'})$ and $p_t a(p_t) = p_{t'} \cdot a(p_{t'})$. Thus, $s_t = s_{t'}$ if and only if $p_t = p_{t'}$. Therefore the state is uniquely determined by the posterior distribution.

Some additional insight is given to the role of posterior distributions in the following theorem.

Theorem 3.2. Suppose that behavioral strategies are observable, in addition to the histories. Then for any finitely repeated game, the posterior distributions are sufficient to determine strategies in each period. That is, if $p_t(h_t) = p_{t'}(h_{t'})$ then $x_t^k(h_t) = x_{t'}^k(h_{t'})$ and $y_t^r(h_t) = y_{t'}^r(h_{t'})$.

Proof. The value of the one period game is

$$v_1(p, q) = \text{Max}_x \text{Min}_y \{ \sum p^k q^k x^k A^{kr} y^r \}$$

Let $F_1(p, q) = \{(x, y) \mid (x, y) \text{ are equilibrium strategies in the one stage game with priors } (p, q)\}$. Select a function $f_1(p, q)$ such that

$$f_1(p, q) \in F_1(p, q) \text{ for all } (p, q) \in \Delta^K \times \Delta^R$$

(f_1 may be taken to be measurable, using the measurable selection theorem, and the fact that F_1 is a closed valued correspondence.) Write

$$f_1(p, q) = [(x_1^1(p, q), \dots, x_1^K(p, q)), (y_1^1(p, q), \dots, y_1^R(p, q))]$$

Proceed to the two period game, using f_1 . Call this game G_2 ,

$$v_2(p,q) = \text{Max}_x \text{Min}_y \{ (1 - \delta) \Sigma p^k q^r x^k A^{kr} y^r + \\ \delta \Sigma \bar{x}_i \bar{y}_j [\Sigma p_2^k(i) q_2^r(j) x_1^k(p(i), q(j)) A^{kr} y_1^r(p(i), q(j))] \}$$

$$\text{with } \bar{x}_i = \Sigma p^k x_i^k, \bar{y}_j = \Sigma q^r y_j^r$$

$$p^k(i) = p^k x_i^k / \bar{x}_i, q^r(j) = q^r y_j^r / \bar{y}_j$$

Denote the set of equilibrium strategies in this game by F_2

$$F_2(p,q) = \{x,y \mid (x,y) \text{ are equilibrium strategies in } G_2\}$$

Again, select a measurable function, f_2 , with

$$f_2(p,q) \in F_2(p,q) \text{ for all } (p,q) \in \Delta^K \times \Delta^R$$

$$f_2(p,q) = (x_2(p,q), y_2(p,q)).$$

One may proceed in this way to a n-stage game, obtaining functions $(x_t, y_t): \Delta^K \times \Delta^R \rightarrow \Delta^I \times \Delta^J$. This completes the proof. \square

Note again that the steps above depend upon the observability of behavioral strategies. For example, in the two period game, suppose player I considers changing x_1 in the first period, say to \tilde{x}_1 . This will generally change $p(i)$, the posterior, say to $\tilde{p}(i)$. Conditional on (i,j) occurring,

player I will evaluate the "new" expected payoff in the second period as $v(\tilde{p}(i), q(j))$. But this implicitly assumes that player II recomputes the posterior (due to the change from x_1 to \tilde{x}_1), and responds in the second period with a strategy optimal against $\tilde{p}(i)$.

IV. Extreme Transition Functions

Recall that in the earlier discussion (following Theorem 2.4), the transition function was decomposed into two parts, i.e., $T = (u,v)$, where u was the posterior updating function and v the vector of "add-on" payoffs. In many zero sum games, it is necessarily the case that the posterior sequence remain in the interior of the simplex or converges to an extreme point of the simplex at a slow rate in any equilibrium. Therefore any extreme point characterization should restrict attention to the function v . What properties will a transition function extreme in v possess? That question is the subject of the following discussion (see Theorem 4.1 below).

Recall the discussion between Theorems 2.4 and 2.5 and define the sets $A(s)$, $s \in S$ as follows:

$$A(s) = \{v: I \times J \rightarrow \mathbb{R}^{K+R} \mid \exists x(s) \in \Delta_I, y(s) \in \Delta_{II}\}$$

with $s = (\xi, \zeta, p, q)$ and

$$(i) \quad \xi^k = f^k(s, x(s), y(s), v) \geq f^k(s, x, y(s), v), \forall x \in \Delta_I, \forall k \in K$$

$$(ii) \quad \zeta^r = g^r(s, x(s), y(s), v) \geq g^r(s, x(s), y, v), \forall y \in \Delta_{II}, \forall r \in R$$

$$(iii) \quad v(i, j) \in \alpha(u(i, j))\}$$

Here $u(i,j) \in \Delta^K \times \Delta^R$, is the posterior distribution determine by $x(s)$, $y(s)$ and the occurrence of (i,j) . When $\text{prob}(i,j) = 0$ under $x(s)$, $y(s)$ it is defined as in the discussion between Theorems 2.4 and 2.5.

Given a state s , an element of $A(s)$ is a vector payoff attached to each history of the form (i,j) , such that these add-on payoffs are equilibrium vector payoffs on each subform and optimal strategies given these add-on payoffs yield vector payoffs (ξ, ζ) . Note that $A(s)$ is closed (since f^k , g^r are continuous and α is an upper hemicontinuous correspondence). Let

$$A = \prod_{s \in S} A(s)$$

A is nonempty and since it is a compact set in a locally convex linear topological space, A has extreme points. A point $v \in A$ is a function

$$v: S \times I \times J \rightarrow \mathbb{R}^{K+R}$$

$$(\text{or } v: S \rightarrow \mathbb{R}^{(I \times J)(K+R)})$$

Note that if $v \in A$ is an extreme point of A , then $v(s)$ is an extreme point of $A(s)$, for all $s \in S$. Corresponding to any extreme point of A , there are functions x, y , with

$$x: S \rightarrow \Delta_I; \quad y: S \rightarrow \Delta_{II}.$$

Theorem 4.1: Let $v \in \text{ext } A$. Then for all $s \in S$, $v(s, i, j) \notin \text{bd } \alpha(y(s, i, j))$ for at most $\min(I, J)$ of the points $(i, j) \in I \times J$.

Proof: For any $s \in S$ and $k \in K$ with $s = (\xi, \zeta, p, q)$

$$\xi^k = \sum_i x_i^k(s) \left[\sum_{j,r} q^r y_j^r(s) a_{ij}^{kr} + \delta \sum_j \bar{y}_j(s) v^k(s,i,j) \right]$$

$$> \sum_i x_i^k \left[\sum_{j,r} q^r y_j^r(s) a_{ij}^{kr} + \delta \sum_j \bar{y}_j(s) v^k(s,i,j) \right], \forall x^k \in \Delta^I.$$

Suppose that for some pair j, j' , that

$$v(s,i,j) \in \text{int } \alpha(u(s,i,j))$$

$$v(s,i,j') \in \text{int } \alpha(u(s,i,j'))$$

Fix some $k \in K$.

If $\bar{y}_j(s) = 0$, put $v_1^e(s,i,j) = v_2^e(s,i,j) = v^e(s,i,j)$, $\forall e \in (K \setminus \{k\}) \cup R$.

Put

$$v_1^k(s,i,j) = v^k(s,i,j) + \varepsilon$$

$$v_2^k(s,i,j) = v^k(s,i,j) - \varepsilon$$

and put

$$v_1(s,i,j) = v_2(s,i,j) = v(s,i,j) \text{ on } S \setminus \{s\} \times I \times J.$$

Then $v_1 \neq v$, $v_2 \neq v$, $v_1, v_2 \in A$ for ε sufficiently small and

$(1/2)(v_1 + v_2) = v$, contradicting the fact that v is extreme.

If $\bar{y}_j(s), \bar{y}_{j'}(s) > 0$, let $c = \bar{y}_j(s)v^k(s,i,j) + \bar{y}_{j'}(s)v^k(s,i,j)$. Let

$$v_1^k(s,i,j') = v^k(s,i,j') + \varepsilon$$

and choose $v_1^k(s,i,j)$ to satisfy

$$c = \bar{y}_j(s)v_1^k(s,i,j) + \bar{y}_{j'}(s)v_1^k(s,i,j')$$

Similarly, let $v_2^k(s,i,j') = v^k(s,i,j') - \epsilon$ and choose $v_2^k(s,i,j)$ to satisfy

$$c = \bar{y}_j(s)v_2^k(s,i,j) + \bar{y}_{j'}(s)v_2^k(s,i,j')$$

for

$$\tilde{j} \neq j, j', v_1^k(s,i,\tilde{j}) = v_2^k(s,i,\tilde{j}) = v^k(s,i,\tilde{j})$$

and let

$$v_1^e(s,i,j) = v_2^e(s,i,j) = v^e(s,i,j), \forall e \in (K \setminus \{k\}) \cup R, \forall s \in S, \forall (i,j) \in I \times J.$$

for ϵ sufficiently small $v_1, v_2 \in A$, $v_1 \neq v \neq v_2$ and $(1/2)(v_1 + v_2) = v$. This again contradicts the assumption that v is extreme.

Thus, for any $s \in S$ and any $i \in I$,

$$v(s,i,j) \notin \text{bd } \alpha(u(s,i,j)) \text{ for at most one element of } J.$$

Therefore, for any $s \in S$

$$v(s,i,j) \notin \text{bd } \alpha(u(s,i,j)) \text{ for at most } I \text{ elements of } I \times J.$$

The same reasoning implies that for that for any $s \in S$

$$v(s,i,j) \notin \text{bd } \alpha(u(s,i,j)) \text{ for at most } J \text{ elements of } I \times J.$$

Therefore, for any $s \in S$

$$v(s,i,j) \notin \text{bd } \alpha(u(s,i,j)) \text{ for at most } \min(I, J) \text{ elements of } I \times J. \quad \square$$

Notes

¹Consider a one sided incomplete information game where the informed player has two possible types and two pure strategies in the one stage game. For any prior $p, p \in (0,1)$, if players are restricted to pure strategies then the two player types can play the same at each stage in the repeated game, in which case all posteriors equal the prior; or differently at some stage leading to posteriors $\{0,1\}$. The set of possible values for the posterior in any stage and in any equilibrium is $\{0,p,1\}$.

²Suppose that (σ_1, τ_1) and (σ_2, τ_2) are two equilibrium strategy pairs in a complete information game with two players, each having two pure strategies in the one stage game. These pure strategies are $i \in \{1,2\}$ and $j \in \{1,2\}$. Suppose in the first period each player randomizes $(1/2, 1/2)$ on pure strategies. If in the second period $i + j$ is odd they play (σ_1, τ_1) and if $i + j$ is even they play (σ_2, τ_2) . Note that neither player can unilaterally affect the probability that the sum $i + j$ is odd or even. This gives an equilibrium point giving players the average of the payoffs from (σ_1, τ_1) and (σ_2, τ_2) . This is called a jointly controlled lottery, first introduced by Aumann, Maschler and Stearns [5].

³A two player game is symmetric if each player has the same strategy space, $X_1 = X_2 = X$, and $u_1(x_1, x_2) = u_2(x_2, x_1)$, $\forall x_1, x_2 \in X$.

⁴It is understood here that

$$E_{\omega pq}(a^{kr} | k) = \sup_{\sigma} E_{\sigma \tau pq}(a^{kr} | k)$$

and

$$E_{\omega_{pq}}(b^{kr} | r) = \sup_{\tau} E_{\sigma_{\tau pq}}(b^{kr} | r)$$

⁵This example was suggested to me by J. F. Mertens.

FIGURE 1




FIGURE 2




FIGURE 3





FIGURE 4




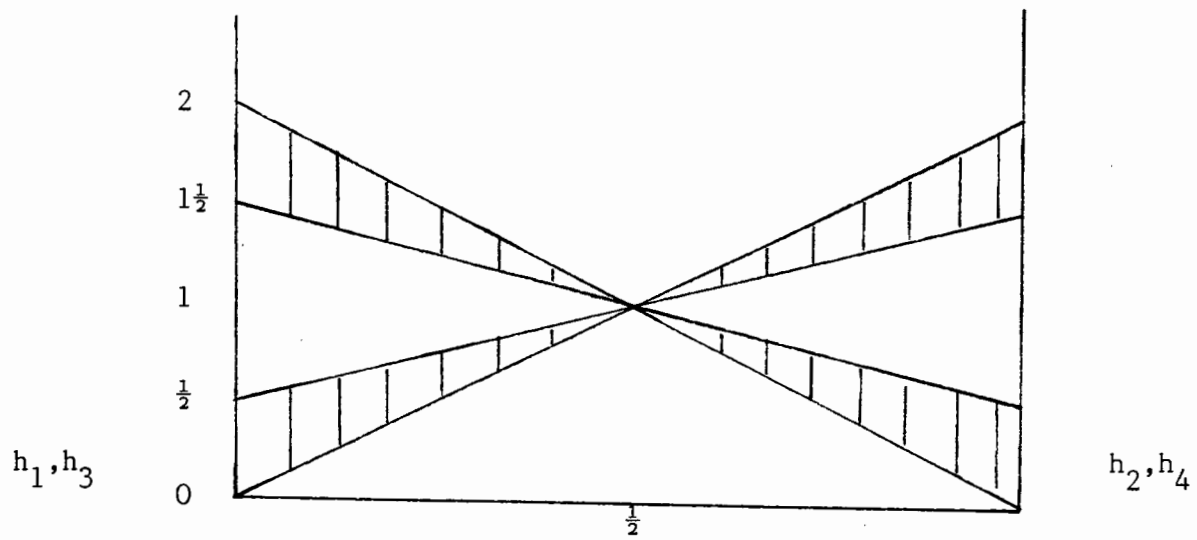



FIGURE 6 : OVERLEAF

FIGURE 7



$h_1=(1,1)$, $h_2=(2,1)$, $h_3=(1,2)$, $h_4=(2,2)$.

FIGURE 6



References

1. Abreu, D. (1982), "Repeated Games with Discounting: A General Theory and an Application to Oligopoly," mimeo, Princeton University.
2. Abreu, D., Pearce, D. and Stacchetti, E. (1985), "Optimal Cartel Equilibrium with Imperfect Monitoring."
3. Aumann, R. J. and Hart, S. (198-), "Bi-convexity and Bi-Martingales," Mathematics of Operations Research, forthcoming.
4. Aumann, R. J. and Maschler, M. (1966), "Repeated Games with Incomplete Information: A Survey of Recent Results," Mathematica, Report to the Arms Control and Disarmament Agency.
5. Aumann, R. J., Maschler, M. and Stearns, R. E. (1967), "Repeated Games of Incomplete Information: An Approach to the Non-Zero-Sum Case," Mathematica, Report to the Arms Control and Disarmament Agency.
6. Aumann, R. and Shapley, L. (1976), "Long Term Competition: A Game Theoretic Analysis," mimeo, Hebrew University.
7. Fudenberg, D. and Maskin, E. (1986), "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," Econometrica, Vol. 54, No. 3, May 1986.
8. Hart, S. (1985), "Non-Zero-Sum Two-Person Repeated Games with Incomplete Information," Mathematics of Operations Research, Vol. 10, No. 1, February 1985.
9. Kreps, D. and Wilson, R. (1982), "Sequential Equilibria," Econometrica, Vol. 50, No. 4.
10. Ponsard, J. P. and Sorin, S. (1982), "Optimal Behavioral Strategies in 0-Sum Games with Almost Perfect Information," Mathematics of Operations Research, Vol. 7, No. 1.
11. Rubinstein, A. (1977), "Equilibrium in Supergames," mimeo, The Hebrew University.
12. Selten, R. (1975), "Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games," International Journal of Game Theory, Vol. 4.