CHAPTER IV: REPEATED GAMES

1. Introduction

The two-person non-zero-sum game, "Prisoners' Dilemma", has the following payoff matrix:

1 \ 2	G	H
G	(1,1)	(5,0)
H	(0,5)	(4,4)

It is useful to think of G as a "greedy" strategy and H as a "helpful" strategy. The strategy G is dominating for both players and since each player wants to maximize his utility, G is the "rational" choice for both players. Moreover, (G,G) is the unique equilibrium point of this game. So if this game is played once, it is "reasonable" to single out (G,G) as the "solution" of the game.

Now suppose this game is to be played 100 times with the total payoff being the sum of the payoffs from the 100 plays of the game. At the 100-th play of the game, the game is to be played once, and the rational choices for the players are (G,G), since these are dominating strategies. At the 99-th play of the game, each player realizes that in the 100-th play the players will choose (G,G), so that the 100-th play is essentially determined, and the 99-th play is in strategic reality the last, so the players will choose their dominating strategies (G,G). By backwards induction we see that the rational choices for the players are (G,G) at each of the 100 plays of the game. These strategies are not only maxmin strategies, but also constitute an equilibrium point for the repeated game of 100 plays; moreover, the only equilibrium payoff is (100,100).

We now ask, what would be rational choices for the players if the game were to be played a countable number of times? One might expect that some sort of cooperation is induced: if a player deviates from a particular strategy at some play, in order to increase his own payoff, then the other player may be able to act in such a way that his opponent is penalized in every subsequent play of the game. In the sequel we study repeated games, consisting of a countable number of repeated plays of a single game.

2. Definition of a Repeated Game and Strategies

2.1 Notation. Let Γ denote the finite game with

(i) the set of players $N = \{1,\ldots,n\}$,

(ii) strategy sets $\Sigma^1,\ldots,\Sigma^n$, (let $\Sigma = \prod_{i\in N} \Sigma^i$),

(iii) payoff functions $h^i\colon \Sigma \to \mathbb{R}$, $\forall\, i \in N$.

Let X^i be the set of mixed strategies for player i (i.e., the set of probability distributions on Σ^i), let $X = \prod_{i\in N} X^i$ and let $H^i\colon X \to \mathbb{R}$ be the expected payoff function for player i.

2.2 Definition. The super-game, $\hat{\Gamma}$, consists of countably many repeated plays of the game Γ.

2.3 Remark. We should really speak of strategies for Γ and of "super strategies" for $\hat{\Gamma}$. However, to simplify terminology, we shall henceforth refer to $\sigma^i \in \Sigma^i$ (a strategy in Γ) as a "choice", and to a strategy in $\hat{\Gamma}$ simply as a strategy.

2.4 Definition. A (pure) strategy for player i is a sequence $F^i = \{f_k^i\}_{k=1}^{\infty}$, where f_k^i is a function, $f_k^i: (\Sigma)^{k-1} \to \Sigma^i$, (i.e., f_k^i dictates the choice of player i at the k-th play of the game Γ, as a function of the choices made by all the players in the preceding $(k - 1)$ plays of Γ). Let $\mathcal{F}^i = \{F^i\}$ = the set of all pure strategies for player i in $\hat{\Gamma}$.

2.5 Remark. Let $|\Sigma| = s$, $|\Sigma^i| = t_i$. Then, for each fixed $i \in \mathbb{N}$, $|\mathcal{F}^i| = \prod_{k=1}^{\infty} t_i^{(s^{k-1})} = t_i^{(\sum_{k=1}^{\infty} s^{k-1})} = t_i^{\aleph_0} = \aleph$. So each player has uncountably many pure strategies.

2.6 Remark. By considering the example where the game Γ has payoff matrix $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$, we can see that the $\underline{\lim}_{m\to\infty}$ and $\overline{\lim}_{m\to\infty}$ of the average payoff over m plays of Γ need not agree. So that, in general, the limit of the average payoff cannot be used as the payoff for $\hat{\Gamma}$, since this limit need not exist. In fact we shall see that there is no need to define a payoff function for $\hat{\Gamma}$, rather, in Section 3 we shall define equilibrium points in terms of preference relations on the strategy n-tuples.

In the following we define mixed strategies for the players.

2.7 Notation. Let $([0,1],\mathcal{B},\lambda)$ denote the probability space consisting of the unit interval, $[0,1]$, with the σ-field of its Borel subsets, $\mathcal{B}$, and the Lebesgue measure, λ. For each $i \in N$, player i has a probability space (Ω^i, A^i, P^i) which is isomorphic to $([0,1],\mathcal{B},\lambda)$. Let ω^i denote an element of Ω^i, let $\Omega = \prod_{i \in N} \Omega^i$, $A = \prod_{i \in N} A^i$ and let $P = \prod_{i \in N} P^i$ be the probability measure on (Ω, A).

2.8 Definition. A mixed strategy for player i in $\hat{\Gamma}$ is a sequence of functions $F^i = \{f_k^i\}_{k=1}^{\infty}$, where $f_k^i\colon (\Sigma)^{k-1} \times \Omega^i \to \Sigma^i$.

2.9 Remark. To implement a mixed strategy, F^i, in a play of $\hat{\Gamma}$, player i acts as follows. Before the play begins, player i performs a random experiment to select $\omega^i \in \Omega^i$, according to the probability distribution P^i on (Ω^i, A^i). Then at the k-th play of Γ, player i chooses $f_k^i(\sigma_1,\ldots,\sigma_{k-1},\omega^i) \in \Sigma^i$, where $\sigma_1,\ldots,\sigma_{k-1} \in \Sigma$ are the n-tuples of choices made by the players at the preceding $(k-1)$ plays of Γ. Note that if f_k^i is independent of ω^i for all k, then F^i is equivalent to a pure strategy.

2.10 Remark. If one needs (wants) independent randomizations at each play k of Γ, then (Ω^i, A^i, P^i) is sufficient, because it is isomorphic to $([0,1],\mathcal{B},\lambda)$, which in turn is isomorphic to the Cartesian product of denumerably many copies of itself; e.g., if $\omega^i \in [0,1]$ with decimal expansion $\omega^i = .\,\alpha_1\alpha_2\alpha_3\ldots$, then this generates realizations of countably many independent $\omega_k^i \in [0,1]$, $k = 1,2,\ldots,$ by the following diagonal construction:

$$\omega_1^i = .\alpha_1\,\alpha_2\,\alpha_4\,\ldots$$
$$\omega_2^i = .\alpha_3\,\alpha_5\,\ldots$$
$$\omega_3^i = .\alpha_6\,\ldots$$
$$\vdots$$

<u>2.11 Notation</u>. Let $\hat{X}^i = \{F^i\}$ = the set of all mixed strategies for player i in $\hat{\Gamma}$ and let $\hat{X} = \prod_{i \in N} \hat{X}^i$.

3. <u>Equilibrium Points</u>

<u>3.1 Notation</u>. For a collection of sets A^i, $i \in N$, let $A = \prod_{i \in N} A^i$ and $A^{-i} = \prod_{j \neq i} A^j$.

<u>3.2 Notation.</u> For $F \in \hat{X}$, $F = (F^i)_{i \in N}$, $F^i = \{f_k^i\}_{k=1}^{\infty}$ $\forall\, i \in N$, let $\{\underset{\sim}{\tau}_k^F\}_{k=1}^{\infty}$ be the sequence of random variables defined on Ω as follows. For $\omega = (\omega^1,\ldots,\omega^n) \in \Omega$, let $\underset{\sim}{\tau}_1^F(\omega) = (f_1^1(\omega^1), \ldots, f_1^n(\omega^n)) \in \Sigma$, the n-tuple of choices of the players at the first play of Γ, corresponding to ω,

$$\underset{\sim}{\tau}_2^F(\omega) = (f_2^1(\underset{\sim}{\tau}_1^F(\omega),\omega^1), \ \ldots, \ f_2^n(\underset{\sim}{\tau}_1^F(\omega),\omega^n))$$

$$\vdots$$

$$\underset{\sim}{\tau}_k^F(\omega) = (f_k^1(\underset{\sim}{\tau}_1^F(\omega),\ldots,\underset{\sim}{\tau}_{k-1}^F(\omega),\omega^1), \ \ldots, \ f_k^n(\underset{\sim}{\tau}_1^F(\omega),\ldots,\underset{\sim}{\tau}_{k-1}^F(\omega),\omega^n))$$

$$\vdots$$

Then $\underset{\sim}{\tau}_k^F$ is the n-tuple of choices made by the players in the k-th play of Γ (it is a random variable when the players are using mixed strategies).

We have not defined payoffs for $\hat{\Gamma}$ and so, to enable us to define an equilibrium point for $\hat{\Gamma}$, we introduce the following preference relations on the mixed strategies. (Also see Section 5 for a discussion of payoffs in $\hat{\Gamma}$).

3.3 Definition. Let $F, G \in \hat{X}$.

(a) Player i ℓ-prefers G to F if there exists an $\varepsilon > 0$ such that

$$P\{\omega \in \Omega: \frac{1}{m} \sum_{k=1}^{m} h^i(\underset{\sim}{\tau}_k^G(\omega)) > \frac{1}{m} \sum_{k=1}^{m} h^i(\underset{\sim}{\tau}_k^F(\omega)) + \varepsilon,$$

$$\text{for all but finitely many } m\text{'s}\} = 1$$

We denote this by $G \overset{\ell}{\underset{i}{>}} F$.

(b) Player i u-prefers G to F if there exists an $\varepsilon > 0$ such that

$$P\{\omega \in \Omega: \frac{1}{m} \sum_{k=1}^{m} h^i(\underset{\sim}{\tau}_k^G(\omega)) > \frac{1}{m} \sum_{k=1}^{m} h^i(\underset{\sim}{\tau}_k^F(\omega)) + \varepsilon$$

$$\text{for infinitely many } m\text{'s}\} > 0 \quad .$$

We denote this by $G \overset{u}{\underset{i}{>}} F$.

Remark: One can define other preference relations as well; the two above are the "extreme".

3.4 Definition. An n-tuple of mixed strategies, $F \in \hat{X}$, has a payoff (or: F is summable) if there is $a \in \mathbb{R}^n$ such that

$$\frac{1}{m}\sum_{k=1}^{m} h^i(\underset{\sim}{I}_k^F(\omega)) \xrightarrow[m\to\infty]{\text{a.s.}} a^i \quad , \quad \forall\, i \in N \quad .$$

3.5 Definition. An n-tuple of mixed strategies, $F \in \hat{X}$, is an $\left\{\begin{matrix}\text{upper}\\ \text{lower}\end{matrix}\right\}$ equilibrium point for $\hat{\Gamma}$ if:

(a) F has a payoff,

(b) for all $i \in N$ and for all $G^i \in \hat{X}^i$, $(F^{-i},G^i) \underset{i}{\overset{\{{u \atop \ell}\}}{\not\succ}} F$.

3.6 Notation. We shall use the abbreviation $\left\{\begin{matrix}\text{u.e.p.}\\ \text{l.e.p.}\end{matrix}\right\}$ for $\left\{\begin{matrix}\text{upper equilibrium point}\\ \text{lower equilibrium point}\end{matrix}\right\}$.

3.7 Remark. Since $G \underset{i}{\overset{\ell}{\succ}} F \Rightarrow G \underset{i}{\overset{u}{\succ}} F$, it follows that $G \underset{i}{\overset{u}{\not\succ}} F \Rightarrow G \underset{i}{\overset{\ell}{\not\succ}} F$ and hence $\{\text{u.e.p.}\} \subset \{\text{l.e.p.}\}$.

3.8 Example. Consider the repeated game, $\hat{\Gamma}$, in which Γ is the game of "Prisoners' Dilemma". The strategy pair in which both players choose G at each play of Γ is an (upper) equilibrium point for $\hat{\Gamma}$, with payoff $(1,1)$. The strategy pair in which each player plays "tit-for-tat", is also an (upper) equilibrium point for $\hat{\Gamma}$, with payoff $(4,4)$. To play his "tit-for-tat" strategy, a player chooses H in the first play of Γ and in each subsequent play of Γ, he chooses whatever his opponent chose in the immediately preceding play of Γ. We show that the "tit-for-tat" strategies do form an equilibrium point for $\hat{\Gamma}$ (it is even easier to show that the strategy pair where both players always choose G is an equililbrium point). Let $F = (F^1,F^2)$ denote the "tit-for-tat" strategy pair. When both players play according to F, they will choose H at each play of Γ and it follows that F has payoff (4,4). To prove that F satisfies condition (b) of Definition 3.5, it suffices, because of the symmetry between 1 and 2, to show that there is no $G^1 \in \hat{X}^1$ such that $(G^1,F^2) \overset{u}{\underset{1}{\gtrsim}} F$. Since F is a pure strategy we need only show that, for each $\varepsilon > 0$, no matter what pure strategy player 1 uses, as long as player 2 plays "tit-for-tat", then player 1's mean payoff over the first m plays of Γ is greater than $4 + \varepsilon$ for at most finitely many m. (In fact, we shall show that player 1's mean payoff is greater than 4 for at most a single value of m.)

Suppose player 2 plays "tit-for-tat" in a play of $\hat{\Gamma}$. If player 1 always chooses H then the mean payoff is always 4. On the other hand, if player 1 deviates from choosing H and k is the first play of Γ in which he chooses G, then his payoff for the k-th play is 5, but in the (k + 1)-st play of Γ, player 2 will play G and player 1's payoff for that play will be at most 1. So player 1's mean payoff for the first (k + 1) plays of Γ will be at most $(4(k - 1) + 5 + 1)/(k + 1) < 4$. For future plays of Γ, player 1's payoff will be at most 1 until the play following one in which he reverts to choosing H. It follows that player 1's mean payoff for the first $m > k + 2$ plays of Γ can be at most

$$\frac{4(k - 1) + 5 + 1 + 0 + 4(m - k - 3) + 5}{m} = 4 - \frac{5}{m} < 4 \quad .$$

So, player 1's mean payoff is greater than 4 only at the k-th play of Γ. Hence, whether or not player 1 deviates from playing "tit-for-tat", his mean payoff over the first m plays of Γ can be greater than 4 for at most one value of m, when player 2 plays "tit-for-tat".

We note that the payoff (0,5) cannot be an equilibrium payoff for $\hat{\Gamma}$, because, if there was a strategy pair $F = (F^1, F^2) \in X$ with payoff (0,5), then player 1 would ℓ-prefer the pair $(\hat{G}, F^2)$ to F, where $\hat{G}$ is the strategy for player 1 where he chooses G in each play of Γ. Player 1 ℓ-prefers $(\hat{G}, F^2)$ to F because by playing $\hat{G}$ against F^2 he obtains a mean payoff of at least 1, whereas playing F^1 against F^2, his mean payoff converges to 0 with probability 1.

We shall prove a theorem characterizing equilibrium payoffs, but before we can do so, we must introduce the notion of <u>individual rationality</u> for a player.

<u>3.9 Definition</u>. The <u>individual rationality level</u> for player i in Γ is

$$r^i = \min_{x^{-i} \in X^{-i}} \max_{x^i \in X^i} H^i(x) \quad .$$

<u>3.10 Remark</u>. The individual rationality level for player i is such that, when the players $j \neq i$ are not allowed to correlate their choices, r^i is the minimum payoff which those players can ensure player i cannot exceed (i.e., it is the maximum payoff which player i cannot be prevented from obtaining). The following example illustrates the fact that, in general, for an n-person game $(n > 2)$,

$$\min_{x^{-i}\in X^{-i}} \max_{x^i \in X^i} H^i(x) \neq \max_{x^i\in X^i} \min_{x^{-i}\in X^{-i}} H^i(x) .$$

3.11 Example. Consider the three-person game in which player 1's payoffs are as in the tables below.

2 \ 3	L	R
T	-1	0
B	0	0

Payoff matrix for player 1's first choice

2 \ 3	L	R
T	0	0
B	0	-1

Payoff matrix for player 1's second choice

Here $\min_{2,3} \max_{1} H^1 = \min_{q,r} \max \{-qr, -(1 - q)(1 - r)\} = -1/4$, and $\max_{1} \min_{2,3} H^1 = \max_{p} \min \{-p, 0, 0, -(1 - p)\} = -1/2$. The payoff $-1/4$ is the minimum which players 2 and 3 can ensure player 1 cannot exceed, when they make their choices independently. To ensure that player 1's expected payoff did not exceed $-1/2$, players 2 and 3 would have to correlate their choices to randomize (with equal probabilities) between the pairs (T,L) and (B,R).

3.12 Notation. Let $C = \text{conv } \{h(\sigma): \sigma \in \Sigma\}$ = the convex hull of the n-tuples of payoffs for Γ, (i.e., C is the set of payoffs which can be achieved by correlation (agreements)). Note that $\{H(x): x \in \Xi\} \subset C$; since Γ is a finite game, C is a compact set.

3.13 "Folk" Theorem. The following three sets are equal.

(a) U = {payoffs corresponding to u.e.p.'s in $\hat{\Gamma}$}.

(b) L = {payoffs corresponding to l.e.p.'s in $\hat{\Gamma}$}.

(c) $C_{IR} = \{a \in C: a^i \geq r^i \ \forall\ i \in N\}$.

3.14 Example. When Γ is the game of Prisoners' Dilemma,

$$r^1 = \min_q \max \{q + 5(1 - q),\ 4(1 - q)\} = \min_q \{5 - 4q\} = 1 \quad ,$$

and by symmetry, $r^2 = 1$. So, C_{IR} is the set of points, a, in C with $a = (a^1,a^2) \geq (1,1)$.

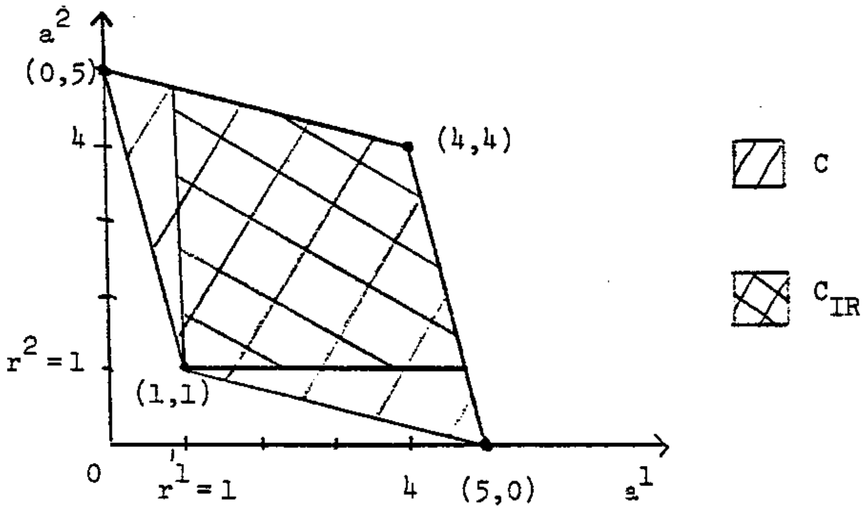

Proof of Theorem 3.13. Since {u.e.p.} $\subset$ {l.e.p.} it follows that (i) $U \subset L$. We shall prove that (ii) $L \subset C_{IR}$ and (iii) $C_{IR} \subset U$. Combining (i), (ii) and (iii) yields, $U \subset L \subset C_{IR} \subset U$, and hence $U = L = C_{IR}$. So we need only prove (ii) and (iii).

Since $\{h(\sigma): \sigma \in \Sigma\} \subset C$ and C is convex and closed, it follows from Definition 3.4 that U and L are subsets of C.

(ii) We shall prove that the complement of C_{IR} is contained in the complement of L, this being equivalent to $L \subset C_{IR}$. So suppose $a \in C \setminus C_{IR}$. Then there is $i \in N$ such that $a^i < r^i$. To prove that $a \notin L$, it suffices to prove that if $F \in \hat{X}$ has payoff a, then there is $G^i \in X^i$ such that $G = (F^{-i}, G^i) \overset{\ell}{\underset{i}{>}} F$.

So suppose $F \in \hat{X}$, $F = (F^i)_{i \in N}$, $F^i = \{f_k^i\}_{k=1}^{\infty}$ $\forall\, i \in N$, has payoff a. Since $r^i = \min_{x^{-i} \in X^{-i}} \max_{x^i \in \Sigma^i} H^i(x)$, then for each $x^{-i} \in X^{-i}$ there is $\beta^i(x^{-i}) \in \Sigma^i$ such that $H^i(x^{-i}, \beta^i(x^{-i})) \geq r^i$. The following example illustrates the first two steps in the inductive procedure for defining $G^i = \{g_k^i\}_{k=1}^{\infty}$, given F.

Example. Consider the case where Γ is the game of Prisoners' Dilemma. Suppose $\Omega^2 = [0,1]$, A^2 = Borel σ-field, P^2 = Lebesgue measure, (i.e., the uniform distribution on $[0,1]$, where each sub-interval has measure equal to its length). Suppose

$$f_1^2(\omega^2) = \begin{cases} G, & \text{if } \omega^2 \in [0, 1/2) \\ H, & \text{if } \omega^2 \in [1/2, 1] \end{cases}$$

$$f_2^2((G,G),\omega^2) = \begin{cases} G, & \text{if } \omega^2 \in [0, 1/3) \\ H, & \text{if } \omega^2 \in [1/3, 1] \end{cases}$$

$$f_2^2((G,H),\omega^2) = \begin{cases} G, & \text{if } \omega^2 \in [1/7, 5/7] \\ H, & \text{if } \omega^2 \in [0, 1/7) \cup (5/7, 1] \end{cases}$$

$$f_2^2((H,G),\omega^2) = \begin{cases} G, \text{ if } & \omega^2 \in [0,1/4) \\ H, \text{ if } & \omega^2 \in [1/4,1] \end{cases}$$

$$f_2^2((H,H),\omega^2) = \begin{cases} G, \text{ if } & \omega^2 \in [0,1/6) \\ H, \text{ if } & \omega^2 \in [1/6,1] \end{cases} .$$

Then g_1^1, g_2^1 are defined as follows. Since f_1^2 dictates that, in the first play of Γ, player 2 should choose each of G and H with probability 1/2, then $g_1^1 \equiv \beta^1(1/2\,G, 1/2\,H)$. Since $r^1 = 1$, the choice G for player 1 guarantees him a payoff of at least r^1 in Γ, and so we can let $\beta^1 \equiv G$.

Knowing f_1^1, f_2^1 and the choices of the players at the first play of Γ, g_2^1 is defined so as to depend on the <u>conditional</u> probability that, under f_2^2, player 2 will choose G or H, <u>given</u> the choices made by the players in the first play of Γ.

If the players chose (G,G) in the first play of Γ, then the probability that player 2 chooses G, (respectively H), in the second play of Γ, given (G,G) was chosen in the first play, is

$$P^2\{\omega^2 \in [0,1/3) \mid \omega^2 \in [0,1/2)\} = 2/3 \quad ,$$

$$(\text{respectively } P^2\{\omega^2 \in [1/3, 1] \mid \omega^2 \in [0, 1/2)\} = 1/3) \quad .$$

Similarly, if (G,H) was chosen in the first play, then the probability that player 2 chooses G, (respectively H), in the second play, given (G,H) was chosen in the first play, is

$$P^2\{\omega^2 \in [1/7,\ 5/7] \mid \omega^2 \in [1/2, 1]\} = 3/7 \quad ,$$

$$(\text{respectively} \quad P^2\{\omega^2 \in [0,\ 1/7) \cup (5/7,\ 1] \mid \omega^2 \in [1/2, 1]\} = 4/7) \quad .$$

We then define g_2^1 as follows.

$$g_2^1(G,G) = \beta^1(2/3\,G,\ 1/3\,H)$$

$$g_2^1(G,H) = \beta^1(3/7\,G,\ 4/7\,H) \quad .$$

The definitions of $g_2^1(H,G)$ and $g_2^1(H,H)$ are immaterial since $g_1^1 = G$.

We now give the general definition of the (pure) strategy $G^i = \{g_k^i\}_{k=1}^{\infty}$ and of $G = (F^{-i}, G^i)$.

For each $j \neq i$, let $x_1^j \in X^j$ be the mixture of choices in Γ, corresponding to f_1^j, i.e.,

$$x_1^j(\sigma^j) = P^j\{\omega^j \in \Omega^j : f_1^j(\omega^j) = \sigma^j\}$$

= the probability that under $F^j = \{f_k^j\}_{k=1}^{\infty}$, player j chooses $\sigma^j \in \Sigma^j$ at the first play of Γ .

Let $x_1^{-i} = (x_1^j)_{j \neq i}$ and let $g_1^i \equiv x_1^i \equiv \beta^i(x_1^{-i})$. For each $\omega \in \Omega$, let $\tau_1^G(\omega) = (f_1^{-i}(\omega), g_1^i)$, and $\underset{\sim}{y}_1(\omega) = h^i(\tau_1^G(\omega))$. Then

$$E(\underset{\sim}{y}_1) = \int_{\Omega} \underset{\sim}{y}_1(\omega)\,dP(\omega) = \text{expected value of } \underset{\sim}{y}_1$$

$$= H^i(x_1) \geq r^i \quad , \quad \text{by the choice of } g_1^i = x_1^i \quad .$$

Suppose $\sigma_1 = (\sigma_1^j)_{j=1}^n$ is an n-tuple of choices for the first play of Γ. For each $j \neq i$ and $\sigma^j \in \Sigma^j$, let

$$x_2^j(\sigma^j) = P^j\{\omega^j \in \Omega^j : f_2^j(\sigma_1, \omega^j) = \sigma^j \mid \underset{\sim}{\tau}_1^G(\omega) = \sigma_1\}$$

= the conditional probability that, under F^j, player j chooses σ^j in the second play of Γ, given that he chose σ_1^j and the other players chose σ_1^{-j} in the first play of Γ.

Let $x_2^{-i} = (x_2^j)_{j \neq i}$, let $g_2^i(\sigma_1) \equiv x_2^i \equiv \beta^i(x_2^{-i})$ and let $x_{2,\sigma_1} = (x_2^{-i}, x_2^i)$. For each $\omega \in \Omega$, let $\underset{\sim}{\tau}_2^G(\omega) = (f_2^{-i}(\underset{\sim}{\tau}_1^G(\omega),\omega),\ g_2^i(\underset{\sim}{\tau}_1^G(\omega)))$ and $\underset{\sim}{y}_2(\omega) = h^i(\underset{\sim}{\tau}_2^G(\omega))$. Then for each $\omega \in \Omega$,

$E(\underset{\sim}{y}_2 \mid \underset{\sim}{\tau}_1^G(\omega))$ = the expected value of $\underset{\sim}{y}_2$, given the n-tuple of choices, $\sigma_1 = \underset{\sim}{\tau}_1^G(\omega)$, made in the first play of Γ

$= H^i(x_{2,\sigma_1})$

$\geq r^i$.

Since the σ-algebra of Borel subsets of Ω generated by the random variable $\underset{\sim}{\tau}_1^G$ contains the σ-algebra generated by the random variable y_1 and since $E(\underset{\sim}{y}_2 \mid \underset{\sim}{\tau}_1^G(\omega)) \geq r^i \ \forall\ \omega \in \Omega$, implies that an average of such terms is $\geq r^i$, (Loève [1978], p. 16, Section 38.2, No. 4), then

$$E(\underset{\sim}{y}_2 \mid \underset{\sim}{y}_1(\omega)) = \text{the expected value of } \underset{\sim}{y}_2\text{, given the payoff, } \underset{\sim}{y}_1(\omega)\text{, for the first play of } \Gamma$$
$$= E(E(\underset{\sim}{y}_2 \mid \underset{\sim}{\tau}_1(\omega)) \mid \underset{\sim}{y}_1(\omega))$$
$$\geq r^i \quad .$$

The (pure) strategy $G^i = \{g_k^i\}_{k=1}^{\infty}$ is defined inductively as follows. Suppose that $\{g_\ell^i\}_{\ell=1}^{k-1}$ has been defined, some $k \geq 3$. Then for n-tuples of choices $\sigma_1,\ldots,\sigma_{k-1}$ made in the first $(k - 1)$ plays of Γ, let

$$x_k^j(\sigma^j) = P^j\{\omega^j \in \Omega^j\colon f_k^j(\sigma_1,\ldots,\sigma_{k-1},\omega^j) = \sigma^j \mid \underset{\sim}{\tau}_\ell^G(\omega) = \sigma_\ell \text{ for } \ell = 1,\ldots,k-1\} \quad ,$$

for each $j \neq i$ and $\sigma^j \in \Sigma^j$. Let $g_k^i(\sigma_1,\ldots,\sigma_{k-1}) \equiv x_k^i \equiv \beta^i(x_k^{-i})$. For each $\omega \in \Omega$, let

$$\underset{\sim}{\tau}_k^G(\omega) = (f_k^{-i}(\underset{\sim}{\tau}_1^G(\omega),\ldots,\underset{\sim}{\tau}_{k-1}^G(\omega),\omega),\ g_k^i(\underset{\sim}{\tau}_1^G(\omega),\ldots,\underset{\sim}{\tau}_{k-1}^G(\omega))) \quad ,$$

and let $\underset{\sim}{y}_k(\omega) = h^i(\underset{\sim}{\tau}_k^G(\omega))$. Then, in direct analogy with the case $k = 2$, we have that for each $\omega \in \Omega$,

(1) $$E(\underset{\sim}{y}_k \mid \underset{\sim}{y}_1(\omega),\ldots,\underset{\sim}{y}_{k-1}(\omega)) \geq r^i \quad .$$

Since the random variables, $\underset{\sim}{y}_k$, $k = 1,2,\ldots$, are uniformly bounded by M, say, we can apply the Strong Law of Large Numbers, see Loève [1978], p. 53, 32.1E with $b_n = M$), to deduce that:

$$\frac{1}{m}\sum_{k=1}^{m}[\underset{\sim}{y}_k(\omega) - E(\underset{\sim}{y}_k \mid \underset{\sim}{y}_1(\omega),\ldots,\underset{\sim}{y}_{k-1}(\omega))] \to 0 \quad \text{as} \quad m \to \infty$$

for almost all $\omega \in \Omega$. It follows, by (1), that

$$P\{\omega \in \Omega: \underline{\lim}_{m\to\infty} \frac{1}{m}\sum_{k=1}^{m}\underset{\sim}{y}_k(\omega) \geq r^i\} = 1 \quad ,$$

i.e.,

(2) $$P\{\omega \in \Omega: \underline{\lim}_{m\to\infty} \frac{1}{m}\sum_{k=1}^{m} h^i(\underset{\sim}{\tau}_k^G(\omega)) \geq r^i\} = 1 \quad .$$

But, by the assumption that F has payoff a, we have,

(3) $$P\{\omega \in \Omega: \lim_{m\to\infty} \frac{1}{m}\sum_{k=1}^{m} h^i(\underset{\sim}{\tau}_k^F(\omega)) = a^i\} = 1 \quad .$$

Since $r^i > a^i$, it follows from (2) and (3) that there is $\varepsilon > 0$ such that

$$P\{\omega \in \Omega: \frac{1}{m}\sum_{k=1}^{m} h^i(\underset{\sim}{\tau}_k^G(\omega)) > \frac{1}{m}\sum_{k=1}^{m} h^i(\underset{\sim}{\tau}_k^F(\omega)) + \varepsilon$$
$$\text{for all but finitely many } m\text{'s}\} = 1 \quad .$$

Hence, $G = (F^{-i},G^i) \underset{i}{\overset{\ell}{\succ}} F$, as required.

(iii) Let $a \in C_{IR}$. We shall show that there is an n-tuple of mixed strategies, $F \in \hat{X}$, such that F is an u.e.p. for $\hat{\Gamma}$, with payoff a.

Before we give the general construction for F, we shall illustrate this construction with the following example.

Example. Consider the case where Γ is the game of Prisoners' Dilemma. Then $(9/4, 7/2)$ is in C_{IR} (see Example 3.14), and it can be written as the following convex combination of the pure payoffs for Γ:

$$(9/4,\ 7/2) = \tfrac{1}{2}(4,4) + \tfrac{1}{4}(0,5) + \tfrac{1}{4}(1,1) \quad .$$

Consider the following sequences of choices A^1 and A^2, for players 1 and 2, respectively, in a play of $\hat{\Gamma}$:

$$A^1 = \{a_k^1\}_{k=1}^{\infty} = \{H,H,H,G,H,H,H,G,\ldots\}$$

$$A^2 = \{a_k^2\}_{k=1}^{\infty} = \{H,H,G,G,H,H,G,G,\ldots\} \quad ,$$

where the pair of choices for the players at the first and the second play of Γ is (H,H), (with payoff $(4,4)$), at the third play of Γ it is (H,G), (with payoff $(0,5)$), at the fourth play of Γ it is (G,G), (with payoff $(1,1)$), and then these four pairs of choices are repeated at each successive block of four plays of Γ. So, the proportion of the first m plays of Γ in which (H,H), (respectively (H,G) or (G,G)), is chosen, tends to $1/2$, (respectively $1/4$ or $1/4$), as $m \to \infty$, and hence the mean payoff vector, (over the first m plays), tends to $(9/4, 7/2)$ as $m \to \infty$.

We now illustrate the procedure for defining an u.e.p., $F = (F^1, F^2)$, with payoff $(9/4, 7/2)$. Firstly, note that, if player 1, (respectively player 2), chooses G in a play of Γ, then his opponent's payoff for that play can be no more than $r^2 = 1$ (respectively $r^1 = 1$). The strategy $F^1 = \{f_k^1\}_{k=1}^{\infty}$ is defined so that player 1 makes his choices according to A^1, so long as player 2 makes his choices according to A^2, but if player 2 deviates from the sequence of choices, A^2, then player 1 "punishes" him in all remaining plays of Γ, by choosing G in those remaining plays. The strategy $F^2 = \{f_k^2\}_{k=1}^{\infty}$ is similarly defined for player 2, so that player 2 chooses according to A^2 as long as player 1 chooses according to A^1, but if player 1 ever deviates from A^1, then player 2 punishes him by choosing G in all succeeding plays of Γ.

For example, if player 2 is to play according to F^2 and player 1's sequence of choices in a play of $\hat{\Gamma}$ is as follows:

$$\{H,H,H,G,G,G,G,H,G,G,\ldots\} ,$$

i.e., player 1 first deviates from A^1 at the fifth play of Γ, then player 2's sequence of choices will be:

$$\{H,H,G,G,H,G,G,G,G,G,\ldots\} ,$$

where player 2 punishes player 1 from the sixth play onwards by choosing G from thereon. Consequently, player 1's mean payoff over the first m plays of Γ, is greater than $9/4$ for at most finitely many m.

We now give the general definition of F, an u.e.p. with payoff a, where $a \in C_{IR}$.

Since $a \in C$, we can write a as a convex combination of the payoff vectors corresponding to n-tuples of choices in Γ, i.e., if $\Sigma = \{\sigma_{(j)}, j = 1,\ldots,s\}$, then there are real numbers $\alpha_{(j)} \geq 0$ such that $\sum_{j=1}^{s} \alpha_{(j)} = 1$ and $a = \sum_{j=1}^{s} \alpha_{(j)} h(\sigma_{(j)})$.

If the $\alpha_{(j)}$'s are all rational, then they can be written as fractions with a common denominator: $\alpha_{(j)} = p_{(j)}/q$, where $p_{(j)}$, q are positive integers, and q is independent of j. In this case, we define $A = (A^i)_{i \in N}$, where, for each $i \in N$, $A^i = \{a_k^i\}_{k=1}^{\infty}$ is a sequence of choices of player i (one for each play of Γ in $\hat{\Gamma}$), defined as follows. For the first $p_{(1)}$ plays of Γ, the choices of player i (i.e., $a_1^i,\ldots,a_{p_{(1)}}^i$) should all be $\sigma_{(1)}^i$, for the next $p_{(2)}$ plays of Γ, the choices of player i should all be $\sigma_{(2)}^i$, and so on up to the q-th play of Γ; then player i should repeat his first q choices at each successive block of q plays of Γ, (i.e., $a_k^i = a_{\bar{k}}^i$ where $\bar{k} \equiv k$ (modulo q)). It follows by construction, that when the players make the choices given by A, in a play of $\hat{\Gamma}$, the fraction of the first m plays of Γ in which $\sigma_{(j)}$ is chosen, tends to $\alpha_{(j)}$ as $m \to \infty$, and so, for each $i \in N$, player i's mean payoff over the first m plays of Γ converges to a^i as $m \to \infty$.

If not all the $\alpha_{(j)}$'s are rational, then, for each j, there is a sequence of rational numbers $(p_{(j)}^{\ell}/q^{\ell})_{\ell=1}^{\infty}$, convergent to $\alpha_{(j)}$ and such that when $\sigma_{(j)}$ is chosen for $p_{(j)}^{1}$ out of the first q^{1} plays of Γ, $p^2_{(j)}$ out of the next q^2 plays of Γ, and so on, then the fraction of the first m plays of Γ in which $\alpha_{(j)}$ is chosen, will tend to $\alpha_{(j)}$ as $m \to \infty$, and hence, for each $i \in N$, player i's mean payoff, over the first m plays of Γ, will converge to a^i as $m \to \infty$. Thus, even if the $\alpha_{(j)}$'s are not all rational, there is an n-tuple of sequences of choices,

$$A = (A^i)_{i \in N} \quad , \quad A^i = \{a^i_k\}_{k=1}^{\infty} , \ \forall \ i \in N \quad ,$$

such that when the players choose according to A in a play of $\hat{\Gamma}$, player i's mean payoff converges to a^i for each $i \in N$.

To construct an u.e.p., $F \in \hat{X}$, with payoff a, we modify the sequences of choices in A, so that if, in a play of $\hat{\Gamma}$, player i is the first player to deviate from the sequence of choices $\{a^i_k\}_{k=1}^{\infty}$, then for all plays of Γ following that in which he deviates, player i will be "punished" by the other players, in the sense that his expected payoff in each of these subsequent plays of Γ will be no more than r^i.

The formal definition of $F = (F^i)_{i \in N}$, $F^i = \{f^i_k\}_{k=1}^{\infty}$ $\forall$ $i \in N$, follows. For each $i \in N$, there is $\gamma^{-i} = x^{-i} \in X^{-i}$ such that

$$\max_{x^i \in \Sigma^i} H^i(\gamma^{-i}, x^i) = r^i \equiv \min_{x^{-i} \in X^{-i}} \ \max_{x^i \in \Sigma^i} H^i(x) \quad .$$

Thus, $H^i(\gamma^{-i}, x^i) \leq r^i$ for all $x^i \in \Sigma^i$, and hence for all $x^i \in X^i$. For each $i \in N$, $f^i_1(\omega^i)$ is defined to be equal to a^i_1 for all $\omega^i \in \Omega^i$, and for $k > 1$, f^i_k is defined on $(\Sigma)^{k-1} \times \Omega^i$ as follows. If

$(\sigma_1,\ldots,\sigma_{k-1}) \in (\Sigma)^{k-1}$ is equal to $(a_1,\ldots,a_{k-1})$, (i.e., in the first $(k-1)$ plays of Γ, the players have played according to A), then $f_k^i(\sigma_1,\ldots,\sigma_{k-1},\omega^i) = a_k^i$, for all $\omega^i \in \Omega^i$. If $(\sigma_1,\ldots\sigma_{k-1}) \neq (a_1,\ldots,a_{k-1})$, let $\bar{\ell} = \min\{\ell: \sigma_\ell \neq a_\ell\}$ and let $\bar{j} = \min\{j: \sigma_{\bar{\ell}}^j \neq a_{\bar{\ell}}^j\}$, (i.e., $\bar{j}$ is the "first" player to deviate from his A-sequence of choices). If $i \neq \bar{j}$, then $f_k^i(\sigma_1,\ldots,\sigma_{k-1},\cdot)$ is defined to be the random variable on Ω^i, corresponding to player i's mixed strategy component in $\gamma^{-\bar{j}}$. The random variable $f_k^{\bar{j}}(\sigma_1,\ldots,\sigma_{k-1},\cdot)$ is defined to be identically equal to $a_k^{\bar{j}}$.

If all of the players play according to F, in $\hat{\Gamma}$, then their sequences of choices are given by A and so it follows that F has pay-off a.

It remains to prove that F satisfies condition (b) of Definition 3.5. So let $i \in N$, let $G^i = \{g_k^i\}_{k=1}^{\infty} \in \hat{X}^i$ and let $G = (F^{-i},G^i)$. It suffices to prove that for each $\varepsilon > 0$,

$$(4) \qquad P\{\omega \in \Omega: \frac{1}{m}\sum_{k=1}^{m} h^i(\underset{\sim}{\tau}_k^G(\omega)) \leq \frac{1}{m}\sum_{k=1}^{m} h^i(\underset{\sim}{\tau}_k^F(\omega)) + \varepsilon$$

$$\text{for all but finitely many } m\text{'s}\} = 1 \quad .$$

For each $\omega \in \Omega$ and $k = 1,2,\ldots$, let

$$y_k(\omega) = h^i(\underset{\sim}{\tau}_k^G(\omega)) \quad , \text{ where } \underset{\sim}{\tau}_k^G \text{ is as in 3.2} \quad .$$

First consider the case where G^i is a pure strategy (independent of ω^i). Then, either $\underset{\sim}{\tau}_k^G(\omega) = f_k$ for all $\omega \in \Omega$ and $k = 1,2,\ldots$, in which case

$E(y_k \mid \tau_1^G(\omega),\ldots \tau_{k-1}^G(\omega)) = h^i(f_k)$ for all $\omega \in \Omega$ and $k = 1,2,\ldots,$ so that, (cf. part (ii) of this proof),

$$E(y_k \mid y_1(\omega),\ldots,y_{k-1}(\omega))$$
$$= E(E(y_k \mid \tau_1^G(\omega),\ldots,\tau_{k-1}^G(\omega)) \mid y_1(\omega),\ldots,y_{k-1}(\omega))$$
$$= h^i(f_k) \quad ,$$

for all $\omega \in \Omega$ and $k = 1,2,\ldots,$ and hence

(5) $$\frac{1}{m}\sum_{k=1}^{m} E(y_k \mid y_1(\omega),\ldots,y_{k-1}(\omega)) = \frac{1}{m}\sum_{k=1}^{m} r^i(f_k) \to a^i$$

as $m \to \infty$, for almost all $\omega \in \Omega$; or there is a smallest number ℓ such that $\tau_\ell^G(\omega) \neq f_\ell$, (where $\tau_\ell^G(\omega)$ is independent of ω), in which case, since player i is the first deviator (because everyone else plays according to F), we have, for all $\omega \in \Omega$,

$$E(y_k \mid \tau_1^G(\omega),\ldots,\tau_{k-1}^G(\omega)) \leq r^i$$

for all $k > \ell$, and hence, (cf. part (ii)), $E(y_k \mid y_1(\omega),\ldots,y_{k-1}(\omega)) \leq r^i$, so that

(6) $$\overline{\lim_{m\to\infty}} \frac{1}{m}\sum_{k=1}^{m} E(y_k \mid y_1(\omega),\ldots,y_{k-1}(\omega))$$
$$= \overline{\lim_{m\to\infty}} \frac{1}{m}\sum_{k=\ell+1}^{m} E(y_k \mid y_1(\omega),\ldots,y_{k-1}(\omega))$$
$$\leq r^i \leq a^i \quad ,$$

since $r^i \leq a^i$ (because $a \in C_{IR}$). Combining (5) and (6), we deduce that when G^i is a pure strategy,

(7) $$\overline{\lim_{m\to\infty}} \frac{1}{m} \sum_{k=1}^{m} E(y_k \mid y_1(\omega),\ldots,y_{k-1}(\omega)) \leq a^i$$

for almost all $\omega \in \Omega$.

Now, G^i may be a mixed strategy which allows player i to randomize amongst his pure strategies, but since (7) holds for all such pure strategies, and for a mixed strategy, the expression on the left of (7) will be an average of such expressions for pure strategies, it follows that (7) also holds when G^i is a mixed strategy.

We now apply the Strong Law of Large Numbers (cf. part (ii)) to deduce that:

(8) $$\overline{\lim_{m\to\infty}} \frac{1}{m} \sum_{k=1}^{m} y_k(\omega) \leq a^i$$

for almost all $\omega \in \Omega$. Also, since F has payoff a,

(9) $$\lim_{m\to\infty} \frac{1}{m} \sum_{k=1}^{m} h^i(\underset{\sim}{\tau}_k^F(\omega)) = a^i$$

for almost all $\omega \in \Omega$. It then follows from (8) and (9) that for each $\varepsilon > 0$,

$$P\{\omega \in \Omega: \frac{1}{m} \sum_{k=1}^{m} h^i(\underset{\sim}{\tau}_k^G(\omega)) \leq \frac{1}{m} \sum_{k=1}^{m} h^i(\underset{\sim}{\tau}_k^F(\omega)) + \varepsilon$$
$$\text{for all but infinitely many } m\text{'s}\} = 1 \quad . \qquad \blacksquare$$

4. Further Reading

Aumann ([1959] and [1975/76], Chapter 8, pp. 86-106) considers repeated games in which the players may form coalitions and he defines an associated notion of a strong equilibrium point. He then gives a characterization of the payoffs associated to strong equilibrium points. References to other work on repeated games can be found in Aumann's notes.

5. Payoff in a Repeated Game

Note that in Definitions 3.3 and 3.4, (preliminaries to the definition of equilibrium point in $\hat{\Gamma}$), we considered mean payoffs for each realization of $\omega \in \Omega$, rather than using expected mean payoffs: $1/m \sum_{k=1}^{m} E(h^i(\underline{\tau}_k^F(\omega)))$. The reason for this is that the supergame $\hat{\Gamma}$ is to be played once, so that a player should use the actual mean payoff for assessing the worth of his strategy, and not the expected mean payoff which is a measure of what he could expect, on average, if the supergame were to be played a number of times (which it is not).

Aumann ([1959], pp. 320-322) illustrates this with an example of a repeated game, in which the players may form coalitions, where he shows that it is misleading for the players to make judgements based on expected payoffs.

References

Aumann, R. [1959], "Acceptable Points in General Cooperative n-Person Games," Contributions to the Theory of Games, Vol. IV, Annals of Mathematics Studies, No. 40, Princeton: Princeton University Press, pp. 287-324.

Aumann, R. [1975/76], "Lectures on Game Theory," Institute for Mathematical Studies in the Social Sciences, Stanford University, Stanford.

Loève, M. [1978], Probability Theory, Volume II, 4th Edition, Springer-Verlag.

Luce, R.D. and H. Raiffa [1957], Games and Decisions, New York: John Wiley and Sons, Inc., Sections 5.4-5.7, pp. 94-102.