CHAPTER II: GAMES IN EXTENSIVE FORM

1. Extensive Form of an n-Player Game

An n-player game, Γ, in extensive form consists of the following.

(i) A set $N = \{1,2,\ldots,n\}$ of players.

(ii) A finite tree (i.e., a connected graph with no cycles), T, called the game tree.

(iii) A distinguished node of the tree (the root of the tree), referred to as the first move. A node of degree one (i.e., connected by one edge only) and different from the root, is called a terminal node. The set of all terminal nodes is denoted Ω.

(iv) A partition, $P^0, P^1, \ldots, P^n$ of the set of non-terminal nodes of the tree (where P^0 may be empty), called the player partition. The nodes in P^0 are called chance moves and the nodes in P^i are called moves of player i, for $i = 1,2,\ldots,n$. The set of non-terminal nodes, that is, the union of $P^0, P^1, \ldots P^n$, is called the set of moves for the game.

(v) For each node in P^0, a probability distribution on the branches out of it, which assigns a positive probability to each branch.

(vi) For each $i \in N$, a partition $U^i_1, U^i_2, \ldots, U^i_{k^i}$ of P^i (where U^i_j is called the j-th information set of player i), such that for each $j \in \{1,2,\ldots,k^i\}$,

  (a) all nodes in U^i_j have the same number of outgoing branches and there is a one-to-one correspondence between the sets of outgoing branches of different nodes in U^i_j, and

  (b) each path from the root to a terminal node can cross U^i_j at most once.

(vii) For each terminal node, t, an n-dimensional vector of real numbers, $(f^1(t),\ldots,f^n(t))$ called the <u>payoff</u> vector for t.

Each of the players of the game knows (i) - (vii).

One can imagine this n-player game as being played in the following manner. Each player has a number of <u>agents</u>, one for each of his information sets. The agents are isolated from each other and each knows the rules of the game, (i)-(vii).

A play of the game begins at the first move. Suppose the play has progressed to the move e. If e is a move of player i, then the agent whose information set contains e chooses one of the branches going out of e, knowing only that he is choosing a branch at one of the moves in his information set. If e is a chance move, then a branch out of e is chosen according to the probabilities specified for the node e, by (v). (Note that the probability distribution at e is independent of the distributions at the other chance moves.) In this manner a unique, finite path is constructed from the first move to some terminal node, t. At t, player i receives the payoff $f^i(t)$ for $i = 1,2,\ldots,n$.

<u>Examples</u>:

(i) <u>Matching pennies</u>.

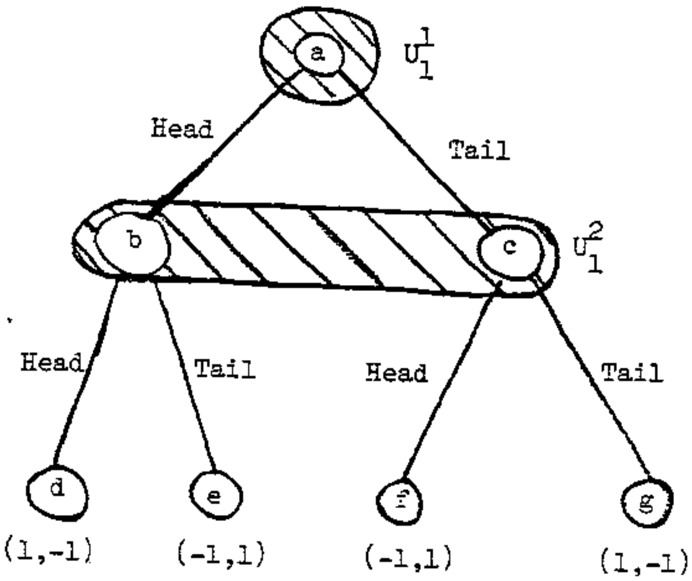

$N = \{1,2\}$

First move = a

{Moves} = {a,b,c}

$P^0 = \emptyset$

$P^1 = \{a\}$, $U_1^1 = \{a\}$

$P^2 = \{b,c\}$, $U_1^2 = \{b,c\}$

(ii)

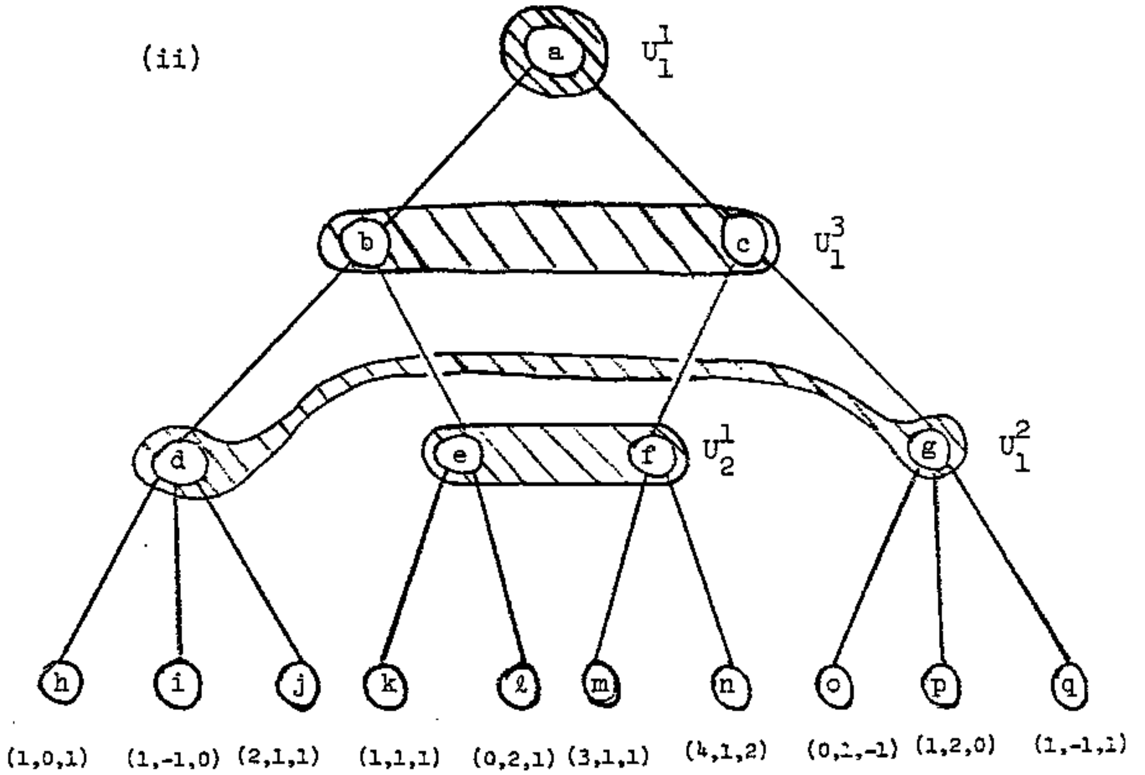

$N = \{1,2,3\}$, First move = a , {Moves} = {a,b,c,d,e,f,g} ,

$P^0 = \emptyset$, $P^1 = \{a,e,f\}$, $U_1^1 = \{a\}$, $U_2^1 = \{e,f\}$,

$P^2 = \{d,g\}$, $U_1^2 = \{d,g\}$, $P^3 = \{b,c\}$, $U_1^3 = \{b,c\}$.

At his second information set, U_2^1, player 1 does not recall what his choice was at U_1^1. So player 1 consists of two agents, one for each of his information sets, and these agents do not communicate.

2. Pure Strategies

2.1 Notation

For each information set U_j^i, let γ_j^i be the number of branches going out of each node in U_j^i, and number these branches from one through γ_j^i such that the one-to-one correspondence between the sets of branches for different nodes in U^i_j is preserved. Let $C(U^i_j) = \{1,2,\ldots,\gamma^i_j\}$, which is the set of choices available to player i at any move in U^i_j. Let $I^i = \{U^i_1,\ldots,U^i_{k^i}\}$, the set of information sets for player i; from now on, we will simplify notation by using U^i to denote a generic element of I^i.

2.2 Definition. A pure strategy for player i is a k^i-tuple, $\sigma^i = (\sigma^i(U^i))_{U^i \in I^i}$, where $\sigma^i(U^i) \in C(U^i)$ for all $U^i \in I^i$. That is, σ^i specifies, for each information set, U^i, of player i, a choice for player i at that information set.

2.3 Notation. Let $\Sigma^i = \{\sigma^i\} = \prod_{U^i \in I^i} C(U^i)$, the set of pure strategies for player i. Let $\Sigma = \Sigma^1 \times \ldots \times \Sigma^n$.

2.4 Definition. For an n-tuple of pure strategies, $\sigma = (\sigma^1,\ldots,\sigma^n) \in \Sigma$, the expected payoff to player i, $h^i(\sigma)$, is defined by

$$h^i(\sigma) = \sum_{t \in \Omega} p_\sigma(t) f^i(t)$$

where $p_\sigma(t)$ is the probability that a play of the game ends at the terminal node t, when the players use strategies $\sigma^1,\ldots,\sigma^n$. So that $p_\sigma(t) \neq 0$ only if, for each $i \in \{1,\ldots,n\}$ and for each node of player i on the path from the root to t, σ^i dictates that player i should choose the branch, at that node, which is along the path from the root to t. When $p_\sigma(t) \neq 0$ it is equal to the product of the probabilities, at each chance move along the path, of choosing the branch which is along the path. The function $h^i\colon \Sigma^1 \times \ldots \times \Sigma^n \to \mathbb{R}$ is called the payoff function for player i.

2.5 Remark. The reason why we consider expected payoffs is that all payoffs to player i are assumed to be in units of a von Neumann-Morgenstern utility for player i.

3. Normal Form

The normal form of Γ is specified as follows.

(i) There are n players.

(ii) The pure strategy sets are $\Sigma^1,\ldots,\Sigma^n$.

(iii) The payoff functions are $h^1,\ldots,h^n$.

3.1 Remark. In Section 2, we showed how to derive the normal form from the extensive form. Conversely, each n-person game given in normal form can be described in extensive form as follows. Starting with the root as the single node of player 1, there are as many branches out of this node as there are strategies for player 1. The end-points of these branches are the nodes for player 2 and these nodes are in the single information set of player 2. Each of these nodes has as many outgoing branches as there are strategies for player 2.

The tree is continued in this manner, with the nodes of player $k + 1$ being the end-points of the branches going out of the nodes of player k, and the nodes of each player are in one information set, with as many outgoing branches as there are strategies for the player. The branches out of the nodes of player n lead to terminal nodes, where the payoff at a terminal node is the payoff corresponding to the strategy choices, of the players, made along the unique path from the root to the terminal node.

3.2 Example. Normal form of a two player zero-sum game, Γ.

1 \ 2	γ	δ	ε
α	1	2	1
β	0	1	2

Entries in the table are payoffs to player 1.

A strategically equivalent extensive form of Γ.

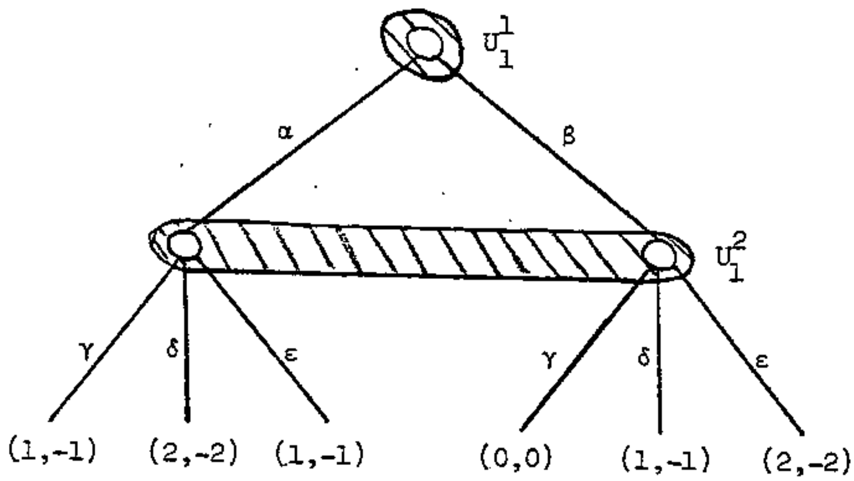

3.3 Remark. To each game in extensive form, there is essentially (up to one-to-one correspondence), a unique normal form representation, whereas there are many different games with the same normal form.

4. Mixed Strategies

4.1 Definition. The set of mixed strategies for player i is

$$X^i = \{x^i = (x^i(\sigma^i))_{\sigma^i \in \Sigma^i} \mid x^i(\sigma^i) \geq 0, \text{ for all } \sigma^i \in \Sigma^i, \sum_{\sigma^i \in \Sigma^i} x^i(\sigma^i) = 1\} ,$$

the set of probability distributions on player i's set of pure strategies. So that in playing a mixed strategy, $x^i = (x^i(\sigma^i))_{\sigma^i \in \Sigma^i}$, player i chooses the pure strategy σ^i with probability $x^i(\sigma^i)$. Let $X = X^1 \times \ldots \times X^n$.

4.2 Definition. The (expected) payoff to player i, corresponding to the mixed strategy n-tuple, $x = (x^1, \ldots, x^n) \in X$, is

$$H^i(x) = \sum_{\sigma \in \Sigma} x(\sigma) h^i(\sigma) \quad ,$$

where $x(\sigma) = \prod_{j=1}^{n} x^j(\sigma^j)$ is the probability that the pure strategy n-tuple $\sigma = (\sigma^1, \ldots, \sigma^n)$ is played under x.

4.3 Remark. Using Definition 2.4, an equivalent expression for the expected payoff can be given as:

$$H^i(x) = \sum_{t \in \Omega} p_x(t) f^i(t)$$

where $p_x(t) = \sum_{\sigma \in \Sigma} x(\sigma) p_\sigma(t)$ is the probability that a play of the game ends at t when the players play according to $x = (x^1, \ldots, x^n)$.

4.4 Remark. See Remark 2.5, which also applies here.

5. Equilibrium Point

5.1 Definition. A (mixed) strategy n-tuple $x = (x^1, \ldots, x^n) \in X$ is an equilibrium point for Γ if:

$$H^i(x|y^i) \leq H^i(x)$$

for all $y^i \in X^i$ and $i = 1, 2, \ldots, n$, where

$$x|y^i = (x^1, \ldots, x^{i-1}, y^i, x^{i+1}, \ldots, x^n) \quad .$$

5.2 <u>Theorem</u> (Nash [1950]). Every (finite) n-person game has an equilibrium point (in mixed strategies).

6. <u>Games with Perfect Information</u>

For a certain class of games, known as games with perfect information, there is an equilibrium point in <u>pure</u> strategies.

6.1 <u>Definition</u>. A game, Γ, has <u>perfect information</u> if, for each $i \in N$ and each information set $U^i \in \mathcal{I}^i$, we have $|U^i| = 1$, i.e., each information set consists of a single move.

The following special case of the more general result was proved by Zermelo.

6.2 <u>Theorem</u> (Zermelo [1912]): In the game of chess, either

(i) white can <u>force</u> a win, or

(ii) black can <u>force</u> a win, or

(iii) both players can <u>force</u> at least a draw.

<u>Proof</u>: See Aumann, Lectures on Game Theory [1975-76], pp. 1-4.

Further examples of "chess-like" games can be found in Aumann's lecture notes (pp. 5-7).

The general result for n-player games with perfect information is the following.

6.3 <u>Theorem</u> (von Neumann, Kuhn [1953]). A game, Γ, with perfect information has an equilibrium point in pure strategies.

Proof. The proof is by induction on the number of moves in the game. Suppose the number of moves is one. Then if the move is a chance move, the strategy sets are all empty and the theorem is vacuously true; whereas if the move is a move of player i, then to obtain an equilibrium strategy, player i should choose the branch which leads to a terminal node with the maximum payoff for him. Thus, the theorem is true when Γ has one move.

Suppose the theorem is true for any game with perfect information and with M moves or less (some $M \geq 1$). Let Γ be a game with perfect information and with $M + 1$ moves. Let T be the game tree for Γ, let r be the root of T, let k be the number of branches going out of r, and number these branches from 1 through k. The node at the end of the j-th branch from r is the root of a sub-tree, $T_{(j)}$, of T, where $T_{(j)}$ is the tree for a sub-game, $\Gamma_{(j)}$, of Γ, because Γ has perfect information (i.e., since each information set is a single node, it is included in one of the $T_{(j)}$'s). For each $\sigma^i \in \Sigma^i$, let $\sigma^i_{(j)}$ be the restriction of σ^i to the information sets in $\Gamma_{(j)}$, let $\Sigma^i_{(j)} = \{\sigma^i_{(j)}\}$ and let $h^i_{(j)}(\sigma_{(j)})$ be the expected payoff to player i in $\Gamma_{(j)}$, when the players play according to the strategy n-tuple $\sigma_{(j)} = (\sigma^1_{(j)},\ldots,\sigma^n_{(j)})$, for $j = 1,\ldots,k$. Each $\Gamma_{(j)}$ is a game with perfect information with M moves or less and so by the inductive assumption it has an equilibrium point $\bar{\sigma}_{(j)} = (\bar{\sigma}^1_{(j)},\ldots,\bar{\sigma}^n_{(j)})$, say, so that for $i = 1,\ldots,n$,

$$(1) \qquad h^i_{(j)}(\bar{\sigma}_{(j)}|\tau^i_{(j)}) \leq h^i_{(j)}(\bar{\sigma}_{(j)}) \quad \text{for all} \quad \tau^i_{(j)} \in \Sigma^i_{(j)} \quad .$$

The cases when r is a chance move and when r is a move of player i, are now considered separately.

(a) $\underline{r \in P^0}$. Let p_j be the probability assigned to the j-th branch going out of r, for $j = 1,\ldots,k$. For each $i \in N = \{1,2,\ldots,n\}$, let $\bar{\sigma}^i = \prod_{j=1}^{k} \bar{\sigma}^i_{(j)}$ be the strategy for player i in Γ, which corresponds to playing according to $\bar{\sigma}^i_{(j)}$ at information sets in $\Gamma_{(j)}$ for $j = 1,\ldots,k$, i.e., $\bar{\sigma}^i(U^i) = \bar{\sigma}^i_{(j)}(U^i)$ when U^i is an information set in $\Gamma_{(j)}$. We prove that $\bar{\sigma} = (\bar{\sigma}^1,\ldots,\bar{\sigma}^n)$ is an equilibrium point for Γ.

For each $i \in N$ and $\tau^i = \prod_{j=1}^{k} \tau^i_{(j)} \in \Sigma^i$, we have:

$$
\begin{aligned}
h^i(\bar{\sigma}|\tau^i) &= \sum_{j=1}^{k} p_j h^i_{(j)}(\bar{\sigma}_{(j)}|\tau^i_{(j)}) \\
&\leq \sum_{j=1}^{k} p_j h^i_{(j)}(\bar{\sigma}_{(j)}) \quad , \text{ by } (1) \\
&= h^i(\bar{\sigma}) \quad .
\end{aligned}
$$

Hence $\bar{\sigma}$ is a pure strategy equilibrium point for Γ.

(b) $\underline{r \in P^{i_0} \text{ for some } i_0 \in N}$. Let $\ell \in \{1,\ldots,k\}$ be such that

$$
h^{i_0}_{(\ell)}(\bar{\sigma}_{(\ell)}) = \max_{1\leq j\leq k} h^{i_0}_{(j)}(\bar{\sigma}_{(j)}) \quad ,
$$

then define $\bar{\sigma}^i \in \Sigma^i$ by:

$$\bar{\sigma}^i = \begin{cases} \prod_{j=1}^{k} \bar{\sigma}^i_{(j)} & , \text{ for } i \neq i_0 \\ \{\ell\} \times \prod_{j=1}^{k} \bar{\sigma}^{i_0}_{(j)} & , \text{ for } i = i_0 \end{cases}$$

i.e.,

$$\bar{\sigma}^i(U^i) = \begin{cases} \bar{\sigma}^i_{(j)}(U^i) & , \text{ if } U^i \text{ is an information set of player } i \text{ in } \Gamma_{(j)} \\ \ell & , \text{ if } U^i = \{r\} \ . \end{cases}$$

We prove that $\bar{\sigma} = (\bar{\sigma}^1,\ldots,\bar{\sigma}^n)$ is an equilibrium point for Γ.

For $i \neq i_0$ and $\tau^i = \prod_{j=1}^{k} \tau^i_{(j)} \in \Sigma^i$,

$$\begin{aligned} h^i(\bar{\sigma}|\tau^i) &= h^i_{(\ell)}(\bar{\sigma}_{(\ell)}|\tau^i_{(\ell)}) && , \text{ since player } i_0 \text{ chooses branch } \ell \text{ at } r, \text{ under } \bar{\sigma}^{i_0} \\ &\leq h^i_{(\ell)}(\bar{\sigma}_{(\ell)}) && , \text{ by } (1) \\ &= h^i(\bar{\sigma}) && , \text{ since player } i_0 \text{ chooses branch } \ell \text{ at } r, \text{ under } \bar{\sigma}^{i_0} \ . \end{aligned}$$

For $i = i_0$ and $\tau^{i_0} = \{m\} \times \prod_{j=1}^{k} \tau^i_{(j)} \in \Sigma^{i_0}$, $m \in \{1,\ldots,k\}$, we have:

$$h^{i_0}(\bar{\sigma}|\tau^{i_0}) = h^{i_0}_{(m)}(\bar{\sigma}_{(m)}|\tau^{i_0}_{(m)})$$, since player i_0 chooses branch m at r, under τ^{i_0}

$$\leq h^{i_0}_{(m)}(\bar{\sigma}_{(m)})$$, by (1)

$$\leq h^{i_0}_{(\ell)}(\bar{\sigma}_{(\ell)})$$, by choice of ℓ

$$= h^{i_0}(\bar{\sigma})$$, since player i_0 chooses branch ℓ at r, under $\bar{\sigma}^{i_0}$.

Hence, $\bar{\sigma}$ is a pure strategy equilibrium point for Γ.

Combining cases (a) and (b) we see that Γ has an equilibrium point in pure strategies. By induction the theorem is true for all Γ with perfect information. ∎

6.4 Remark. For any game in extensive form, with perfect information, we can find an equilibrium point by working backwards from the terminal nodes to the root. At each of his moves, a player chooses the branch which leads to the sub-tree having the highest equilibrium payoff for him. Working backwards from the terminal nodes this is a well-defined process. The following example illustrates this method.

6.5 Example.

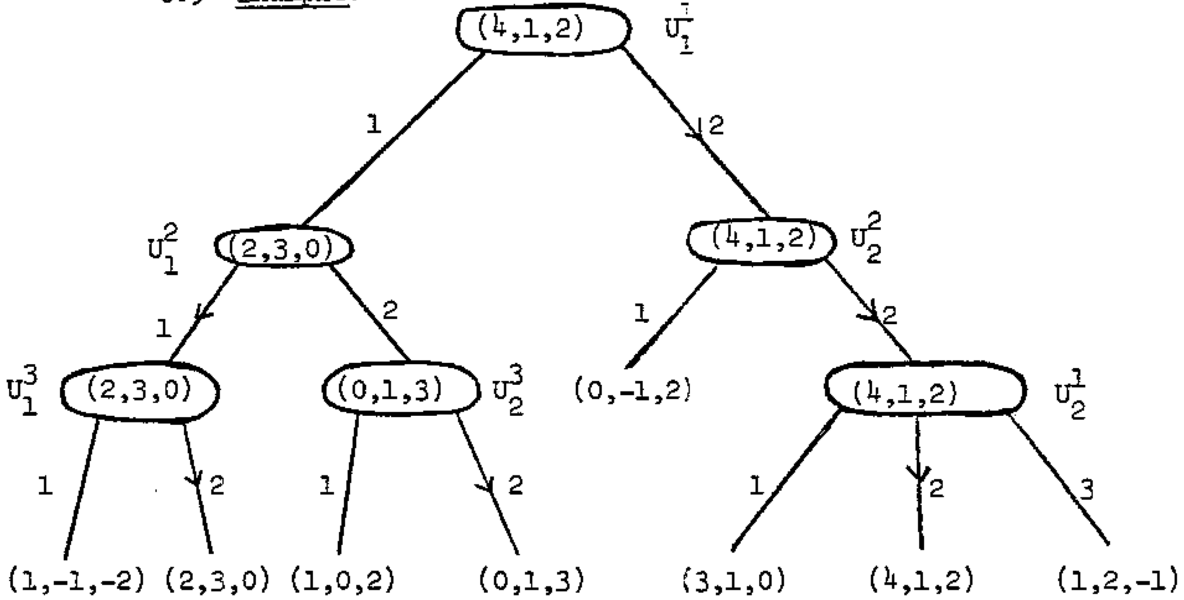

Arrows indicate choices for equilibrium strategies. Numbers in the nodes are equilibrium payoffs for the sub-trees rooted at those nodes.

$$\bar{\sigma}^1(U_1^1) = 2 \quad , \quad \bar{\sigma}^1(U_2^1) = 2$$

$$\bar{\sigma}^2(U_1^2) = 1 \quad , \quad \bar{\sigma}^2(U_2^2) = 2$$

$$\bar{\sigma}^3(U_1^3) = 2 \quad , \quad \bar{\sigma}^3(U_2^3) = 2$$

$$(h^1(\bar{\sigma}), h^2(\bar{\sigma}), h^3(\bar{\sigma})) = (4,1,2) \quad .$$

6.6 Problem. Two players take turns at removing one or two stones from a pile of five. Whoever takes the last stone is the winner and gets one unit of money from his opponent. Draw a tree to represent this game and find an equilibrium point.

7. <u>Behaviour Strategies</u>

A mixed strategy for a player is a probability distribution on his set of pure strategies. An alternative method of randomizing his choices would be for the player to specify, for each of his information sets, a probability distribution over the alternatives at that set. In this way, the player randomizes at each move and the choices at different information sets are made *independently*. A collection of such distributions, containing one distribution for each information set of the player, is called a *behaviour strategy* for the player. The nomenclature arises from the fact that these are the distributions one would measure in attempting to observe the behaviour of a player.

A useful way of viewing the difference between mixed and behaviour strategies is the following. One can think of each pure strategy of a player as a book of instructions, where for each of the player's information sets there is one page which states what choice he should make at that information set. The player's set of pure strategies is a library of such books. A mixed strategy for the player is a probability distribution on his library of books, so that in playing according to his mixed strategy, the player chooses one book from his library by means of a chance device having the probability distribution of his mixed strategy. A behaviour strategy is a single book of a different sort. Although each page still refers to a single information set, it specifies a probability distribution over the choices at that set, not a specific choice.

Given a behaviour strategy for a player, there is always a mixed strategy which, for all possible choices of mixed or behaviour strategies of the other players, will yield the same payoffs as the behaviour strategy. The converse is not in general true, but it is true when the player has what is called "perfect recall."

The following definitions formalize the notion of a behaviour strategy, and the correspondence between mixed strategies and behaviour strategies.

7.1 Definition. A behaviour strategy, b^i, for player i in Γ, is a collection of probability distributions on the information sets of player i, so that

$$b^i = (b^i(U^i,c))_{\substack{U^i \in I^i \\ c \in C(U^i)}}$$

where $b^i(U^i,c) \geq 0$, $\forall\, U^i \in I^i$, $\forall\, c \in C(U^i)$ and $\sum_{c \in C(U^i)} b^i(U^i,c) = 1$ $\forall\, U^i \in I^i$, i.e., $b^i(U^i,c)$ is the probability that player i chooses alternative c at the information set U^i.

7.2 Remark. To specify a mixed strategy for player i we need to choose a point on an $s^i \equiv (\prod_{j=1}^{k^i} \gamma^i_j - 1)$-dimensional simplex, whereas to specify a behaviour strategy for player i, we need to give $\delta^i \equiv \sum_{j=1}^{k^i} (\gamma^i_j - 1)$ real numbers, specifying one point on a $(\gamma^i_j - 1)$-dimensional simplex for $j = 1,\ldots,k^i$. In general, δ^i is much smaller than s^i.

7.3 Definition. The mixed strategy, x^i, corresponding to a behaviour strategy, b^i, is defined by:

$$x^i = (x^i(\sigma^i))_{\sigma^i \in \Sigma^i} \quad \text{where} \quad x^i(\sigma^i) = \prod_{U^i \in I^i} b^i(U^i, \sigma^i(U^i))$$

$$\text{for all} \quad \sigma^i \in \Sigma^i \quad .$$

Note that the probability that player i chooses $\sigma^i(U^i)$ at $U^i \in I^i$ under b^i is $b^i(U^i,\sigma^i(U^i))$, so the probability that he chooses $\sigma^i(U^i)$ at each $U^i \in I^i$ under b^i is $x^i(\sigma^i)$.

7.4 Remark. It is easily checked that x^i is indeed a mixed strategy and that the probability of making a choice, c, at an information set U^i, is the same under b^i as under the corresponding x^i. It follows that x^i leads to the same payoffs as b^i, regardless of the strategies chosen by the other players. We say that x^i is strategically equivalent to b^i.

Next, we define the behaviour strategy generated by a mixed strategy and then give an example to show that the behaviour strategy need not be strategically equivalent to its generating mixed strategy. Moreover, for this example, there is no behaviour strategy which is strategically equivalent to the given mixed strategy.

7.5 Definition. A move $e \in P^i$ is reachable under a mixed strategy $x^i \in X^i$, for player i, if there is a mixed strategy n-tuple, $x \in X$, containing x^i, such that the probability of reaching e under x is positive. An information set U^i for player i is reachable under $x^i \in X^i$ if some move $e \in U^i$ is reachable under x^i. Let $Rch(x^i)$ denote the collection of information sets of player i which are reachable under x^i.

7.6 Remark. Given a tree and a mixed strategy $x^i \in X^i$, we can find those moves of player i which are reachable under x^i by "pruning" the tree. To do this, we eliminate a choice $c \in C(U^i)$ and the sub-trees following c if there is no pure strategy $\sigma^i \in \Sigma^i$ such that σ^i chooses c at U^i and $x^i(\sigma^i) > 0$. The moves of player i which remain after this pruning process are the reachable moves, and an information set is reachable if there is at least one move in the information set remaining in the pruned tree.

7.7 Example.

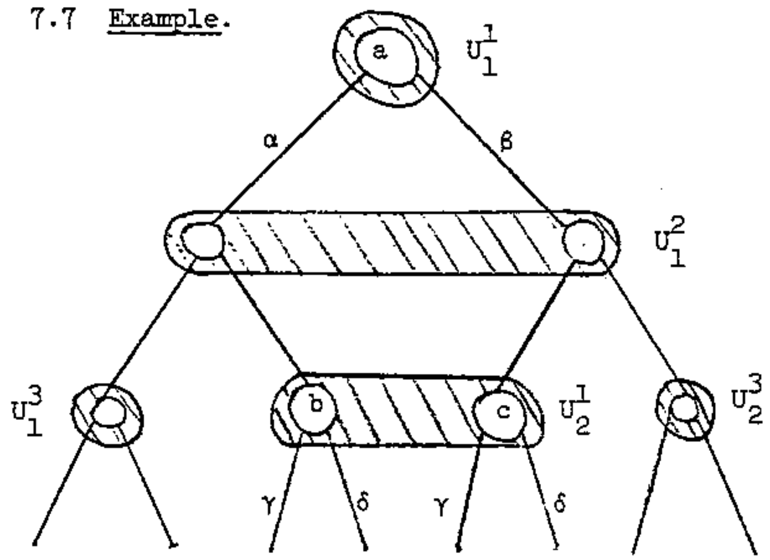

Player 1 has two information sets, with two choices at each set and so he has four pure strategies, $\sigma_1 = (\alpha,\gamma)$, $\sigma_2 = (\alpha,\delta)$, $\sigma_3 = (\beta,\gamma)$, $\sigma_4 = (\beta,\delta)$.

Suppose player 1 chooses the mixed strategy with $x^1(\sigma_1) = 1/2$, $x^1(\sigma_2) = 1/2$, $x^1(\sigma_3) = x^1(\sigma_4) = 0$.

In pruning the tree, we eliminate the choice β at U_1^1 and the sub-tree rooted at the end-node of the branch β. No other choices of player 1 can be eliminated. So the pruned tree is:

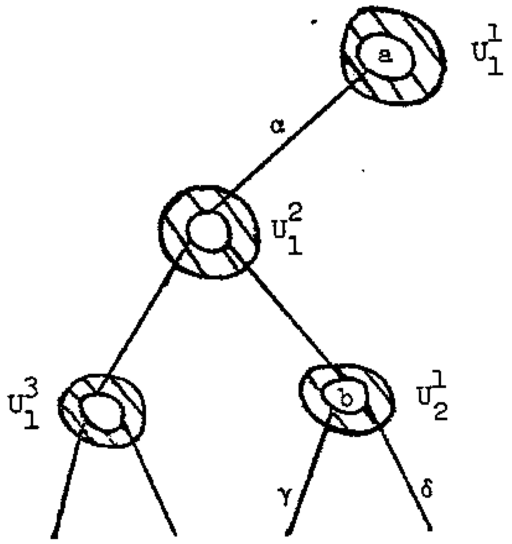

So moves a and b are reachable for player 1 under x^1, but move c is not reachable. However, the information sets U_1^1 and U_2^1 are both reachable, because there are nodes which are in the pruned tree and in these sets.

7.8 <u>Definition</u>. The behaviour strategy, b^i, <u>generated</u> by a mixed strategy $x^i \in X^i$ is defined by:

$$b^i(U^i,c) = \begin{cases} \left[\displaystyle\sum_{\substack{\sigma^i \in \Sigma^i:\ U^i \in Rch(\sigma^i) \\ \text{and } \sigma^i(U^i)=c}} x^i(\sigma^i) \right] \Bigg/ \left[\displaystyle\sum_{\sigma^i \in \Sigma^i:\ U^i \in Rch(\sigma^i)} x^i(\sigma^i) \right], & \text{if } U^i \in Rch(x^i) \\ \dfrac{1}{|C(U^i)|}, & \text{if } U^i \notin Rch(x^i) \end{cases}$$

$\forall\, U^i \in I^i$, $\forall\, c \in C(U^i)$. It is easily checked that b^i is a behaviour strategy.

(Note that when $U^i \notin Rch(x^i)$, the definition of $b^i(U^i,c)$ is immaterial and so we simply define it to be $1/|C(U^i)|$).

7.9 Example. Consider a two-player zero-sum game in which player 1 consists of two people, Jim and his wife, Mary, and player 2 is a single person, Paul. Two cards, one marked "High" and the other "Low", are dealt to Jim and Paul—the two possible deals occurring with equal probabilities. The person with the High card then receives one dollar from the person with the Low card and has the choice of stopping or continuing the play. If the play continues, Mary, not knowing the outcome of the deal, instructs Jim and Paul either to exchange or to keep their cards. Again, the holder of the High card receives a dollar from the holder of the Low card, and the game ends.

Discussion. The game tree for this game is sketched below. The outcome of the draw is indicated in the node at the end of the branch. The alternatives for the players are abbreviated as follows: S = stop, C = continue, K = keep cards, T = exchange (trade) cards. Payoffs at the terminal nodes are those of player 1.

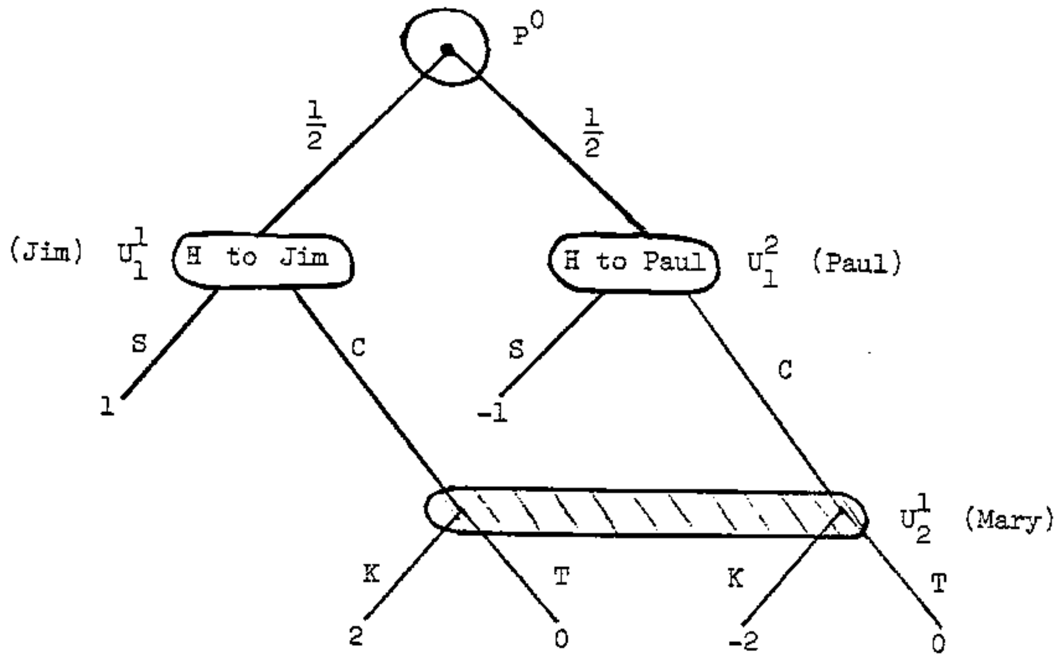

A pure strategy for player 1 is $(\sigma^1(U_1^1), \sigma^1(U_2^1))$, and a pure strategy for player 2 is $(\sigma^2(U_1^2))$. The normal form of this game is then:

1 \ 2	S	C
(S,K)	$\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot (-1) = 0$	$\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot (-2) = -\frac{1}{2}$
(S,T)	$\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot (-1) = 0$	$\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot (0) = \frac{1}{2}$
(C,K)	$\frac{1}{2} \cdot 2 + \frac{1}{2} \cdot (-1) = \frac{1}{2}$	$\frac{1}{2} \cdot 2 + \frac{1}{2} \cdot (-2) = 0$
(C,T)	$\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot (-1) = -\frac{1}{2}$	$\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot (0) = 0$

Strategies (S,K) and (C,T), for player 1, are strictly dominated. Eliminating these strategies, the reduced normal form is:

1 \ 2	S	C
(S,T)	0	$\frac{1}{2}$
(C,K)	$\frac{1}{2}$	0

The optimal mixed strategy for player 1 is to choose (S,T) with probability 1/2 and (C,K) with probability 1/2. The optimal mixed strategy for player 2 is to choose S with probability 1/2 and C with probability 1/2. The value is 1/4. If

$p = (p(S,K), p(S,T), p(C,K), p(C,T))$ and $q = (q(S), q(C))$ denote mixed strategies for players 1 and 2, respectively, where, for example, $p(S,T)$ is the probability that player 1 chooses the pure strategy (S,T), then the optimal mixed strategies are $\bar{p} = (0,(1/2),(1/2),0)$, $\bar{q} = ((1/2),(1/2))$. By playing his optimal mixed strategy, player 1 can guarantee that his expected payoff will be at least $1/4$.

On the other hand, suppose player 1 is restricted to using behaviour strategies. Let $b^1(U_1^1,\cdot) = (\alpha, 1-\alpha)$, $b^1(U_2^1,\cdot) = (\beta, 1-\beta)$, be the probability distributions at U_1^1 and U_2^1, respectively, where $\alpha, 1-\alpha, \beta, 1-\beta$ are the probabilities that he chooses S, C, K, T, respectively. Then player 1's expected payoff, when he uses the behaviour strategy b^1 is:

$$\frac{1}{2}[\alpha \cdot 1 + (1-\alpha)\{\beta \cdot 2 + (1-\beta) \cdot 0\}] + \frac{1}{2} \cdot (-1) = (1-\alpha)(\beta - \frac{1}{2}) \quad ,$$

if player 2 uses S ,

$$\frac{1}{2}[\alpha \cdot 1 + (1-\alpha)\{\beta \cdot 2 + (1-\beta) \cdot 0\}] + \frac{1}{2}[\beta \cdot (-2) + (1-\beta) \cdot 0] = \alpha(\frac{1}{2} - \beta) \quad ,$$

if player 2 uses C .

So the maximum expected payoff that player 1 can guarantee when he is restricted to behaviour strategies is

$$\max_{\alpha,\beta} [\min \{(1-\alpha)(\beta - \frac{1}{2}), \alpha(\frac{1}{2} - \beta)\}] \quad ,$$

which equals 0, since for each value of β either $\beta - 1/2$ or $1/2 - \beta$ is ≤ 0.

Thus, behaviour strategies for player 1 do a poorer job than mixed strategies. Indeed, there is no behaviour strategy for player 1 which has as corresponding mixed strategy, the optimal mixed strategy for player 1. This discrepancy is due to the uncorrelated nature of the probability distributions comprising a behaviour strategy.

Given a behaviour strategy with $b^1(U_1^1,\cdot) = (\alpha, 1-\alpha)$, $b^1(U_2^1,\cdot) = (\beta, 1-\beta)$, the corresponding mixed strategy is $p = (\alpha\beta, \alpha(1-\beta), (1-\alpha)\beta, (1-\alpha)(1-\beta))$. Given a mixed strategy $(p(S,K), p(S,T), p(C,K), p(C,T))$, the behaviour strategy generated by it is given by

$$b^1(U_1^1,\cdot) = (p(S,K) + p(S,T), p(C,K) + p(C,T))$$
$$b^1(U_2^1,\cdot) = (p(S,K) + p(C,K), p(S,T) + p(C,T)) \quad .$$

Note that, in this game, if player 1 reaches U_2^1 from U_1^1, then in going from U_1^1 to U_2^1, he "forgets" what the outcome of the draw was.

A player is said to have perfect recall, if, at each move, he remembers what he knew at previous moves and what choices he made at those moves. We shall prove that if a player has perfect recall then behaviour strategies are sufficient for his strategic purposes, in the sense that any mixed strategy chosen by him can be strategically matched, as far as payoffs are concerned, by its generated behaviour strategy, regardless of the strategies chosen by the other players.

In the sequel, we give a precise mathematical definition of perfect recall and prove the desired result.

7.10 Notation. Let T be the tree of a game Γ. If e is a node of T then $T(e)$ is the sub-tree of T with root at e. For a choice c at e, let $e * c$ be the node following e when c is chosen at e. If E is a set of nodes, let $T(E) = \bigcup_{e \in E} T(e)$, and if c is a choice at an information set, U, let $U * c = \{e * c \mid e \in U\}$.

7.11 <u>Definition</u>. In a game, Γ, player i has <u>perfect recall</u> if, whenever $e_1, e_2 \in P^i$, $e_1 \in U_1^i$, $e_2 \in U_2^i$ and $e_2 \in T(e_1)$, then there is a (unique) $c \in C(U_1^i)$ such that $U_2^i \subset T(U_1^i * c)$. A game Γ in which every player has perfect recall is called a game with perfect recall.

7.12 <u>Examples</u>.

(a) In the game of Example 7.9, player 1 does not have perfect recall.

(b) In the following game, we illustrate the fact that player I has perfect recall by drawing vertical dotted lines, partitioning the nodes of player I into sets which cannot be linked by information sets.

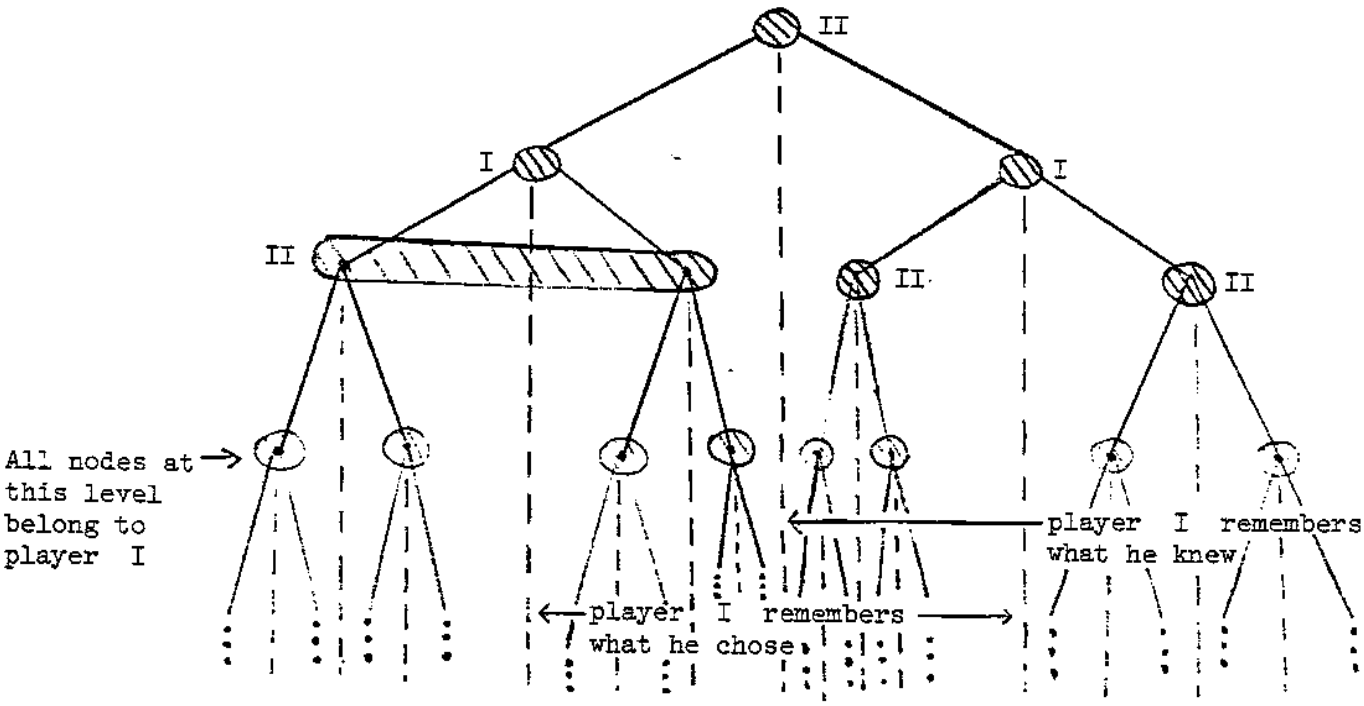

(c) The following example illustrates the coarsest partition of player I's nodes into information sets which is consistent with perfect recall.

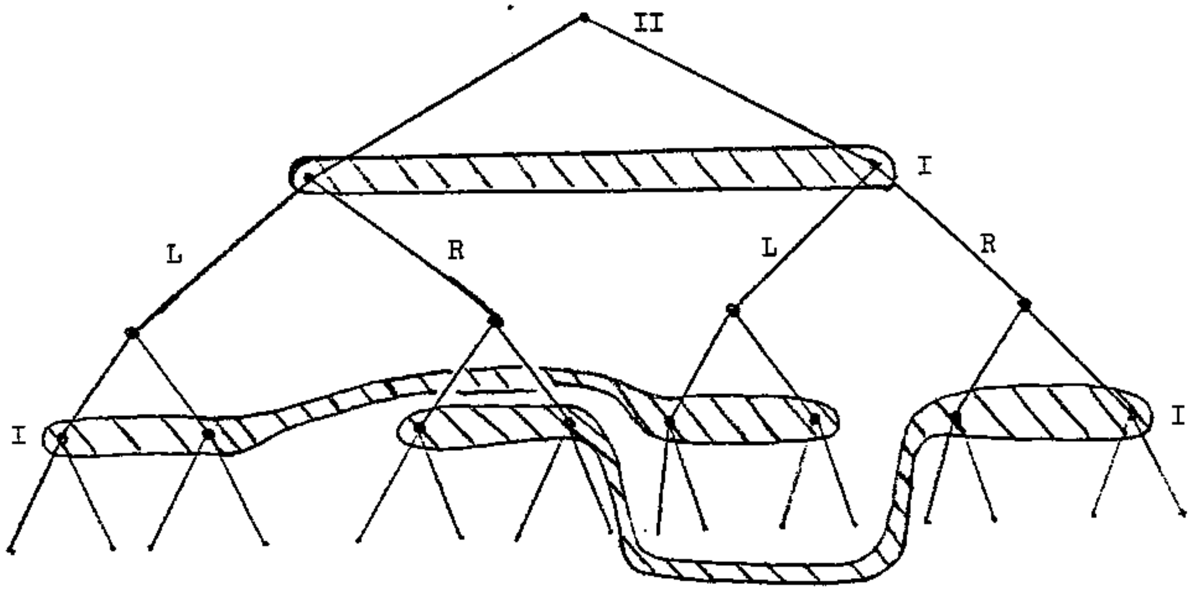

7.13 Theorem (von Neumann, Kuhn [1953]). Let Γ be a game (in extensive form), in which player i has perfect recall. Then, for each mixed strategy $x^i \in X^i$, the behaviour strategy, b^i, generated by x^i, is such that for all $x^k \in X^k$, $k \neq i$ and $j = 1,\ldots,n$,

$$H^j(x) = H^j(x|b^i)$$

where $x = (x^1,\ldots,x^n)$ and $x|b^i = (x^1,\ldots,x^{i-1},b^i,x^{i+1},\ldots,x^n)$. [Note that H^j is defined for $x = (x^1,\ldots,x^n) \in X$. By abuse of notation we also use it for the payoff when x^i is replaced by b^i.]

Proof. Without loss of generality, we assume that $i = 1$. Let $x^1 \in X^1$ and let b^1 be the behaviour strategy generated by x^1. Let $x^i \in X^i$ for each $i \neq 1$. Since $H^j(x) = \sum_{t\in\Omega} p_x(t)f^j(t)$ and $H^j(x|b^1) = \sum_{t\in\Omega} p_{x|b^1}(t)f^j(t)$, where $p_{x|b^1}(t)$ is the probability that the game ends at the terminal node t under the strategy choices $b^1, x^2, \ldots, x^n$, for the n players, then the theorem will be proved if we show that $p_x(t) = p_{x|b^1}(t)$ for all terminal nodes t.

Let t be a fixed terminal node. Then there is a unique path from the root, r, of the game tree, to t. Let $e_1, \ldots, e_\ell$ be the moves of player 1 on this path, in the order in which they occur, and let $c_1, \ldots, c_\ell$ be the choices required at those moves to keep on the path. Let $w_1, \ldots, w_m$ be the moves other than $e_1, \ldots, e_\ell$ on the path, and let $d_1, \ldots, d_m$ be the choices required at those moves, to keep on the path. Let q be the probability that $d_1, \ldots, d_m$ are chosen at $w_1, \ldots, w_m$, when player i plays x^i, for each $i \neq 1$. Since the choices at $w_1, \ldots, w_m$ do not depend on player 1's strategy, q is well-defined. Let $p_{x^1}(c_1, \ldots, c_\ell)$ (respectively $p_{b^1}(c_1, \ldots, c_\ell)$), be the probability that player 1 makes the choices $c_1, \ldots, c_\ell$ at $e_1, \ldots, e_\ell$ when he plays the mixed (respectively behaviour) strategy x^1 (respectively b^1). Then,

$$p_x(t) = p_{x^1}(c_1, \ldots, c_\ell) \cdot q\,, \quad p_{x|b^1}(t) = p_{b^1}(c_1, \ldots, c_\ell) \cdot q$$

If $q = 0$, then clearly $p_x(t) = p_{x|b^1}(t)$. So we assume henceforth that $q \neq 0$, that is, we assume that for each $j \in \{1, \ldots, m\}$, the choice d_j at w_j is made with positive probability, when player i uses x^i for each $i \neq 1$. Then we must prove that $p_{x^1}(c_1, \ldots, c_\ell) = p_{b^1}(c_1, \ldots, c_\ell)$. For each $j \in \{1, \ldots, \ell\}$, let U_j^1 be the information set of player 1 which contains e_j. If one of these information sets is not reachable under x^1, then there must be at least one $k \in \{1, \ldots, \ell\}$ such that the probability of choosing c_k at U_k^1 is zero under x^1, (otherwise the probability of reaching t would be positive and hence so would be the probability of reaching any move on the path from r to t, so that each move, e_j, and hence each information set, U_j^1, would be reachable under x^1, which is a contradiction). But then $b^1(U_k^1,c_k) = 0$ and so $p_{x^1}(c_1,\ldots c_\ell) = 0 = p_{b^1}(c_1,\ldots,c_\ell)$. Thus, we may assume henceforth that $U_j^1 \in \mathrm{Rch}(x^1)$ for $j = 1,\ldots,\ell$.

To prove that $p_{x^1}(c_1,\ldots,c_\ell) = p_{b^1}(c_1,\ldots,c_\ell)$, we need to show that

$$\sum_{\substack{\sigma^1\in\Sigma^1:\ \sigma^1(U_j^1)=c_j \\ \text{for } j=1,\ldots,\ell}} x^1(\sigma^1) = \prod_{j=1}^{\ell} b^1(U_j^1,c_j) \quad .$$

Now,

$$\prod_{j=1}^{\ell} b^1(U_j^1,c_j) = \prod_{j=1}^{\ell} \left[\sum_{\substack{\sigma^1\in\Sigma^1:\ U_j^1\in\mathrm{Rch}(\sigma^1) \\ \text{and } \sigma^1(U_j^1)=c_j}} x^1(\sigma^1)\right] \Bigg/ \left[\sum_{\sigma^1\in\Sigma^1:\ U_j^1\in\mathrm{Rch}(\sigma^1)} x^1(\sigma^1)\right]$$

For each $j \in \{1,\ldots,\ell-1\}$ we have $e_j,\ e_{j+1} \in P^1$, $e_j \in U_j^1$, $e_{j+1} \in U_{j+1}^1$ and $e_{j+1} \in T(e_j)$, so that since player 1 has perfect recall, $U_{j+1}^1 \subset T(U_j^1 * c_j)$. It follows by induction on j that for each $j \in \{1,\ldots,\ell-1\}$ and each $\sigma^1 \in \Sigma^1$, the following three statements are equivalent:

(a) $U^1_{j+1} \in \mathrm{Rch}(\sigma^1)$,

(b) $\sigma^1(U^1_k) = c_k$ for $k = 1,\ldots,j$,

(c) $U^1_j \in \mathrm{Rch}(\sigma^1)$ and $\sigma^1(U^1_j) = c_j$.

Thus, $\{\sigma^1 \in \Sigma^1 : U^1_j \in \mathrm{Rch}(\sigma^1), \sigma^1(U^1_j) = c_j\} = \{\sigma^1 \in \Sigma^1 : U^1_{j+1} \in \mathrm{Rch}(\sigma^1)\}$, for $j = 1,\ldots,\ell-1$. Hence $\prod_{j=1}^{\ell} b^1(U^1_j, c_j)$ is a "telescoping" product which equals:

$$\left[\sum_{\substack{\sigma^1\in\Sigma^1:\ U^1_\ell\in\mathrm{Rch}(\sigma^1)\\ \text{and } \sigma^1(U^1_\ell)=c_\ell}} x^1(\sigma^1)\right] \Bigg/ \left[\sum_{\sigma^1\in\Sigma^1:\ U^1_1\in\mathrm{Rch}(\sigma^1)} x^1(\sigma^1)\right]$$

But, $\sum_{\sigma^1\in\Sigma^1:\ U^1_1\in\mathrm{Rch}(\sigma^1)} x^1(\sigma^1) = 1$, since player 1 has made no choices prior to U^1_1 and so U^1_1 must be reachable under any $\sigma^1 \in \Sigma^1$. So,

$$\prod_{j=1}^{\ell} b^1(U^1_j, c_j) = \sum_{\substack{\sigma^1\in\Sigma^1:\ U^1_\ell\in\mathrm{Rch}(\sigma^1)\\ \text{and } \sigma^1(U^1_\ell)=c_\ell}} x^1(\sigma^1) = \sum_{\substack{\sigma^1\in\Sigma^1:\ \sigma^1(U^1_j)=c_j\\ \text{for } j=1,\ldots,\ell}} x^1(\sigma^1) \quad .$$

Hence, $p_{b^1}(c_1,\ldots,c_\ell) = p_{x^1}(c_1,\ldots,c_\ell)$. ∎

7.14 Corollary. Let Γ be a game with perfect recall. Then Γ has an equilibrium point in behaviour strategies.

<u>Proof</u>. By Theorem 7.13, any mixed strategy for Γ is strategically equivalent to a behaviour strategy. But, by Theorem 5.2 (Nash), Γ has an equilibrium point in mixed strategies and hence, by virtue of strategic equivalence, Γ has an equilibrium point in behaviour strategies. ∎

References

Aumann, R.J. [1975-76], Lectures on Game Theory, Stanford University.

Kuhn, H.W. [1953], "Extensive Games and the Problem of Information, Contributions to the Theory of Games," Annals of Mathematics Studies, Vol. II, No. 28, Princeton: Princeton University Press, pp. 193-216.

Luce, R.D. and H. Raiffa [1957], Games and Decisions, New York: John Wiley and Sons, Inc., Chapter 3, Section 7.4.

Nash, J.F. [1950], "Equilibrium Points in n-Person Games," Proceedings of the National Academy of Sciences, Vol. 36, pp. 48-49.

Zermelo, E. [1912], "Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels," Proc. Fifth Int. Cong. Math., Cambridge, Vol. II, pp. 501-504.

Further References

Aumann, R.J. [1964], "Mixed and Behaviour Strategies in Infinite Extensive Games, Advances in Game Theory," Annals of Mathematics Studies, No. 52, Princeton: Princeton University Press, pp. 627-650. (Aumann considers infinite games with perfect recall.)

Dalkey, N. [1953], "Equivalence of Information Patterns and Essentially Determinate Games, Contributions to the Theory of Games," Annals of Mathematics Studies, Vol. II, No. 28, Princeton: Princeton University Press, pp. 217-243.

Thompson, G.L. [1953], "Signalling Strategies in n-Person Games, Contributions to the Theory of Games," Annals of Mathematics Studies, Vol. II, No. 28, Princeton: Princeton University Press, pp. 267-277.