Chapter 2

GAMES IN EXTENSIVE AND STRATEGIC FORMS

SERGIU HART*

The Hebrew University of Jerusalem

Contents

0. Introduction 20
1. Games in extensive form 20
2. Pure strategies 25
3. Games in strategic form 26
4. Mixed strategies 28
5. Equilibrium points 29
6. Games of perfect information 29
7. Behavior strategies and perfect recall 32
References 40

*Based on notes written by Ruth J. Williams following lectures given by the author at Stanford University in Spring 1979. The author thanks Robert J. Aumann and Salvatore Modica for some useful suggestions.

Handbook of Game Theory, Volume 1, Edited by R.J. Aumann and S. Hart
© *Elsevier Science Publishers B.V., 1992. All rights reserved*

0. Introduction

This chapter serves as an introduction to some of the basic concepts that are used (mainly) in Part I ("Non-Cooperative") of this Handbook. It contains, first, formal definitions as well as a few illustrative examples, for the following notions: games in extensive form (Section 1), games in strategic form (Section 3), pure and mixed strategies (Sections 2 and 4, respectively), and equilibrium points (Section 5). Second, two classes of games that are of interest are presented: games of perfect information, which always possess equilibria in pure strategies (Section 6), and games with perfect recall, where mixed strategies may be replaced by behavior strategies (Section 7).

There is no attempt to cover the topics comprehensively. On the contrary, the purpose of this chapter is only to introduce the above basic concepts and results in as simple a form as possible. In particular, we deal throughout only with *finite* games. The reader is referred to the other chapters in this Handbook for applications, extensions, variations, and so on.

1. Games in extensive form

In this section we present a first basic way of describing a game, called the "extensive form". As the name suggests, this is a most detailed description of a game. It tells exactly which player should move, when, what are the choices, the outcomes, the information of the players at every stage, and so on.

We need to recall first the basic concept of a "tree" and a few related notions. The reader is referred to any book on Graph Theory for further details.

A (finite, undirected) *graph* consists of a finite set V together with a set A of unordered pairs of distinct members[1] of V. See Figure 1 for some examples of graphs. An element $v \in V$ is called a *vertex* or a *node*, and each $\{v_1, v_2\} \in A$ is an *arc*, a *branch* or an *edge* ("joining" or "connecting" the vertices v_1 and v_2). Note that A may be anything from the empty set (a graph with no arcs) to the set of all possible pairs (a "complete" graph[2]). An (open) *path* connecting the nodes v_1 and v_m is a sequence $v_1, v_2, \ldots, v_m$ of distinct vertices such that $\{v_1, v_2\}, \{v_2, v_3\}, \ldots, \{v_{m-1}, v_m\}$ are all arcs of the graph (i.e., belong to A). A *cycle* (or "closed path") is obtained when one allows $v_1 = v_m$ in the above definition.

[1]Note that neither multiple arcs (between the same two nodes) nor "loops" (arcs connecting a node with itself) are allowed.

[2]The complete graph with n nodes has $n(n-1)/2$ arcs.

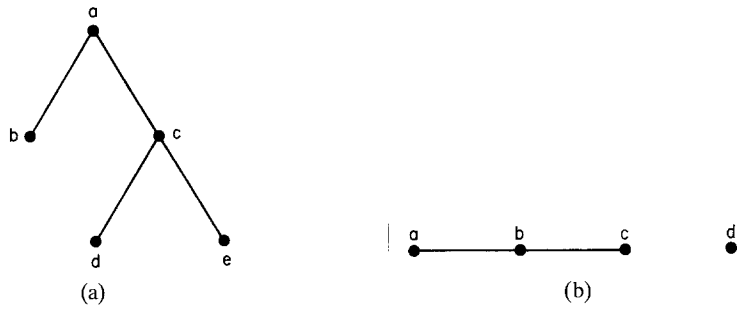

Figure 1. (a) A graph: $V = \{a, b, c, d, e\}$; $A = \{\{a, b\}, \{a, c\}, \{c, d\}, \{c, e\}\}$. (b) A graph: $V = \{a, b, c, d\}$; $A = \{\{a, b\}, \{b, c\}\}$.

A *tree* is a graph where any two nodes are connected by exactly one path. See Figures 2 and 3 for some examples of trees and "non-trees", respectively. It is easy to see that a tree with n nodes has $n - 1$ arcs, that it is a connected graph, and that it has no cycles.

Let T be a tree, and let r be a given distinguished node of T, called the *root* of the tree. One may then uniquely "direct" all arcs so they will point away from the root. Indeed, given an "undirected" arc $\{v_1, v_2\}$, either the unique path from r to v_2 goes through v_1 – in which case the arc becomes the ordered pair (v_1, v_2) – or the unique path from r to v_1 goes through v_2 – and then the arc is directed as (v_2, v_1). The root has only "outgoing" branches. All nodes having only "incoming" arcs are called *leaves* or *terminal nodes*; we will denote by $L \equiv L(T)$ the set of leaves of the tree T. See Figure 4 for an example of a "rooted tree".

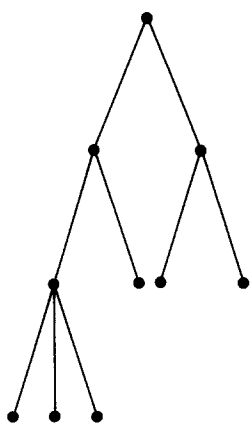

Figure 2. A tree.

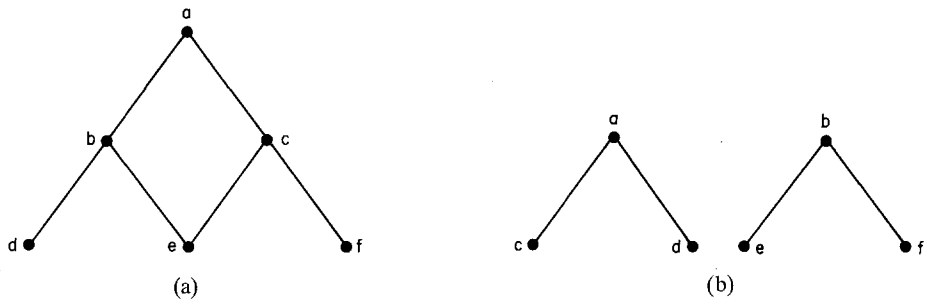

Figure 3. (a) Not a tree (two paths from a to e). (b) Not a tree (no path from a to b).

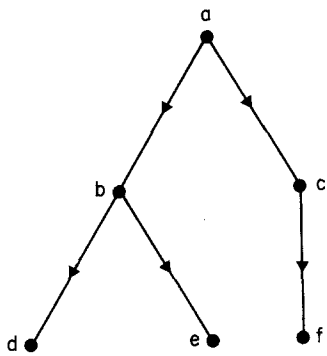

Figure 4. A rooted tree: root = a; leaves = d, e, f.

We can now formally define an *n-person game in extensive form*, Γ, as consisting of the following:[3]

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

(ii) A rooted tree, T, called the *game tree*.

(iii) A partition of the set of non-terminal nodes[4] of T into $n+1$ subsets denoted $P^0, P^1, P^2, \ldots, P^n$. The members of P^0 are called *chance* (or, *nature*) *nodes*; for each $i \in N$, the members of P^i are called the *nodes of player i*.

(iv) For each node in P^0, a probability distribution over its outgoing branches.

(v) For each $i \in N$, a partition of P^i into $k(i)$ *information sets*, $U^i_1, U^i_2, \ldots, U^i_{k(i)}$, such that, for each $j = 1, 2, \ldots, k(i)$:

[3]This definition is due to Kuhn (1953); it is more general than the earlier one of von Neumann (1928) [see Kuhn (1953, pp. 197–199) for a comparison between the two].

[4]A non-terminal node is sometimes called a "move".

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

(b) every (directed) path in the tree from the root to a terminal node can cross each U_j^i at most once.

(vi) For each terminal node $t \in L(T)$, an n-dimensional vector $g(t) = (g^1(t), g^2(t), \ldots, g^n(t))$ of *payoffs*.

(vii) The complete description (i)–(vi) is common knowledge among the players.[5]

One can imagine this game Γ as being played in the following manner. Each player has a number of *agents*, one for each of his information sets [thus i has $k(i)$ agents]. The agents are isolated from one another, and the rules of the game [i.e., (i)–(vii)] are common knowledge among them too. A *play*[6] of Γ starts at the root of the tree T. Suppose by induction that the play has progressed to a non-terminal node, v. If v is a node of player i (i.e., $v \in P^i$), then the agent corresponding to the information set U_j^i that contains v chooses one of the branches going out of v, knowing only that he is choosing an outgoing branch at one of the nodes in U_j^i [recall (v) (a)]. If v is a chance node (i.e., $v \in P^0$), then a branch out of v is chosen according to the probability distribution specified for v [recall (iv); note that the choices at the various chance nodes are independent]. In this manner a unique path is constructed from the root to some terminal node t, where the game ends with each player i receiving a payoff $g^i(t)$.

Remark 1.1. The payoff vectors $g(t)$ are obtained as follows: to each terminal node $t \in L$ there corresponds a certain "outcome" of the game, call it $a(t)$. The payoff $g^i(t)$ is then defined as $u^i(a(t))$, where u^i is a von Neumann–Morgenstern utility function of player i. As will be seen below, the role of this assumption is to be able to evaluate a *random* outcome by its *expected* utility.

Example 1.2 ("Matching pennies"). See Figure 5: $N = \{1, 2\}$; root $= a$; $P^0 = \emptyset$; $P^1 = U_1^1 = \{a\}$; $P^2 = U_1^2 = \{b, c\}$; payoff vectors $(g^1(t), g^2(t))$ are written below each terminal node t. Note that player 2, when he has to make his choice, does not know the choice of player 1.[7] Thus both players are in a

[5]That is, all players know it, each one knows that everyone else knows it, and so on; see the chapter on 'common knowledge' in a forthcoming volume of this Handbook for a formal treatment of this notion.

[6]One distinguishes between a *game* and a *play*: the former is a complete description of the rules (i.e., the whole tree); the latter is a specific instance of the game being played (i.e., just one path in the tree).

[7]This shows the role of information sets; the game changes dramatically if player 2 knows at which node he is (b or c) when he has to make his choice – he can then always "win" (i.e., obtain a payoff of 1).

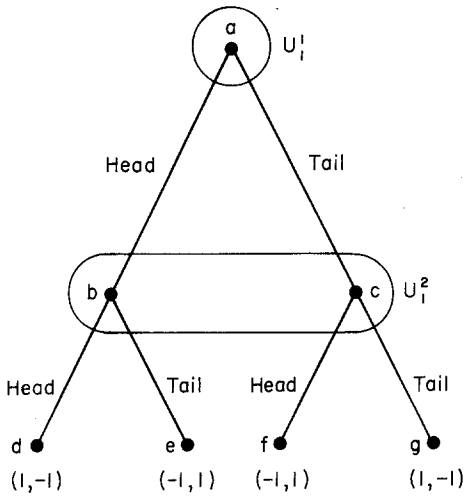

Figure 5. The game tree of Example 1.2.

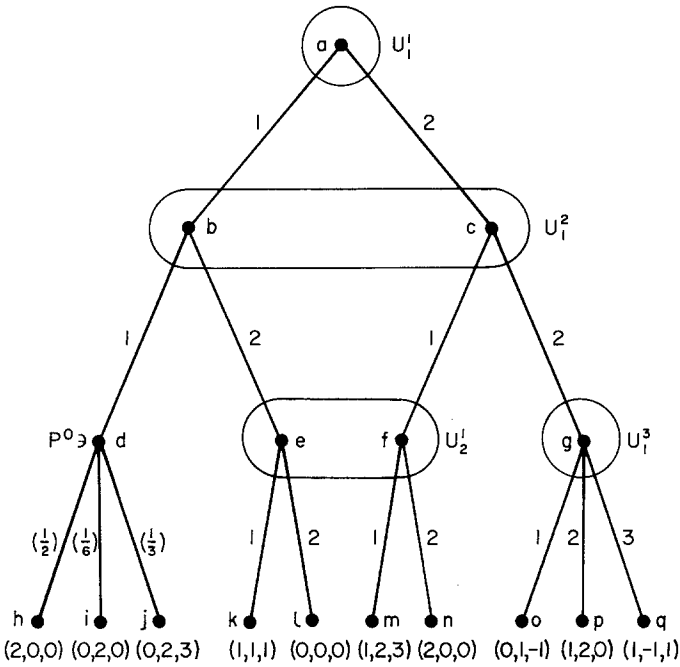

Figure 6. The game tree of Example 1.3.

similar situation: they do not know what is the choice of the other player; for instance, they may make their choices simultaneously. □

Example 1.3. See Figure 6: $N = \{1, 2, 3\}$; root $= a$; $P^0 = \{d\}$; $P^1 = \{a, e, f\}$; $U_1^1 = \{a\}$; $U_2^1 = \{e, f\}$; $P^2 = U_1^2 = \{b, c\}$; $P^3 = U_1^3 = \{g\}$; payoff vectors, the probability distribution at d, and the branches' correspondences [by (v) (a)] are all written on the tree. Note that 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 information set), who do not communicate during the play. □

2. Pure strategies

Let $\boldsymbol{I}^i := \{U_1^i, U_2^i, \ldots, U_{k(i)}^i\}$ be the set of information sets of player i; from now on we will simplify notation by using $U^i \in \boldsymbol{I}^i$ to denote a generic element of $\boldsymbol{I}^i$. For each information set U^i of player i, let $\nu \equiv \nu(U^i)$ be the number of branches going out of each node in U^i; number these branches from 1 through ν such that the one-to-one correspondence between the sets of outgoing branches of the different nodes of U^i is preserved. Thus, let $C(U^i) := \{1, 2, \ldots, \nu(U^i)\}$ be the set of choices available to player i at any node in U^i.

A *pure strategy* s^i *of player* i is a function

$$s^i : \boldsymbol{I}^i \to \{1, 2, \ldots\},$$

such that

$$s^i(U^i) \in C(U^i) \quad \text{for all } U^i \in \boldsymbol{I}^i .$$

That is, s^i specifies for every information set $U^i \in \boldsymbol{I}^i$ of player i, a choice $s^i(U^i)$ there. Let S^i denote the set of pure strategies of player i, i.e., $S^i := \Pi_{U^i \in \boldsymbol{I}^i} C(U^i)$. Let $S := S^1 \times S^2 \times \cdots \times S^n$ be the set of *n-tuples* (or *profiles*) *of pure strategies* of the players.

For an n-tuple $s = (s^1, s^2, \ldots, s^n) \in S$ of pure strategies, the (expected)[8] *payoff* $h^i(s)$ *to player* i is defined by

$$h^i(s) := \sum_{t \in L} p_s(t) g^i(t) , \tag{2.1}$$

where, for each terminal node $t \in L(T)$, we denote by $p_s(t)$ the probability

[8]Recall Remark 1.1.

that the play of the game ends at t when the players use the strategies $s^1, s^2, \ldots, s^n$. This probability is computed as follows. Let $\pi \equiv \pi(t)$ be the (unique) path from the root to the terminal node t. If there exists a player $i \in N$ and a node of i on π at which s^i specifies a branch different from the one along π, then $p_s(t) := 0$. Otherwise, $p_s(t)$ equals the product of the probabilities, at all chance nodes on the path π, of choosing the branch which is along π. The function[9] $h^i: S \to \Re$ defined by (2.1) is called the *payoff function of player i*.

Example 2.2. Consider again the game of Example 1.3. Player 1 has four pure strategies: $\langle 1,1\rangle$, $\langle 1,2\rangle$, $\langle 2,1\rangle$ and $\langle 2,2\rangle$, where $\langle j_1, j_2\rangle$ means that j_1 is chosen at U_1^1 and j_2 is chosen at U_2^1. Player 2 has two pure strategies: $\langle 1\rangle$ and $\langle 2\rangle$, and player 3 has three pure strategies: $\langle 1\rangle$, $\langle 2\rangle$ and $\langle 3\rangle$. To see how payoffs are computed, let $s = (\langle 2,1\rangle, \langle 2\rangle, \langle 3\rangle)$; then the terminal node q is reached, thus $h^1(s) = 1$, $h^2(s) = -1$, and $h^3(s) = 1$. Next, let $s' = (\langle 1,1\rangle, \langle 1\rangle, \langle 3\rangle)$; then $h(s') = (1/2)(2,0,0) + (1/6)(0,2,0) + (1/3)(0,2,3) = (1,1,1)$. □

3. Games in strategic form

A second basic way of describing a game is called the "strategic form" (also known as "normal form" or "matrix form").[10]

An *n-person game in strategic form* Γ consists of the following:

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

(ii) For each player $i \in N$, a finite set S^i of (pure) *strategies*. Let $S := S^1 \times S^2 \times \cdots \times S^n$ denote the set of n-tuples of pure strategies.

(iii) For each player $i \in N$, a function $h^i: S \to \Re$, called the *payoff function of player i*.

In the previous section we showed how the strategic form may be derived from the extensive form. Conversely, given a game in strategic form, one can always construct an extensive form as follows. Starting with the root as the single node of player 1, there are $|S^1|$ branches out of it, one for each strategy $s^1 \in S^1$ of player 1.[11] The $|S^1|$ end-nodes of these branches are the nodes of player 2, and they all form one information set. Each of these nodes has $|S^2|$ branches out of it, one for each strategy $s^2 \in S^2$ of player 2. All these $|S_1| \cdot |S_2|$ nodes form one information set of player 3. The construction of the tree is

[9]The real line is denoted $\Re$.

[10]We prefer "strategic form" since it is more suggestive.

[11]The number of elements of a finite set A is denoted by $|A|$.

Table 1

		s^2	
		Head	Tail
s^1	Head	1, −1	−1, 1
	Tail	−1, 1	1, −1

continued in this manner: there are $|S^i|$ branches – one for each strategy $s^i \in S^i$ of player i – going out of every node of player i; the end-points of these branches are the nodes of player $i+1$, and they all form one information set. The end-points of the branches out of the nodes of player n are the terminal nodes of the tree;[12] the payoff vector at such a terminal node t is defined as $(h^1(s), h^2(s), \ldots, h^n(s))$, where $s \equiv s(t)$ is the n-tuple of strategies of the players that correspond, by our construction, to the branches along the path from the root to t.

Example 3.1. Let $N = \{1, 2\}$; $S^1 = S^2 = \{\text{Head, Tail}\}$; the payoff functions are given in Table 1, where each entry is $h^1(s^1, s^2), h^2(s^1, s^2)$. Plainly, the above construction yields precisely the extensive form of Example 1.2 ("matching pennies"). □

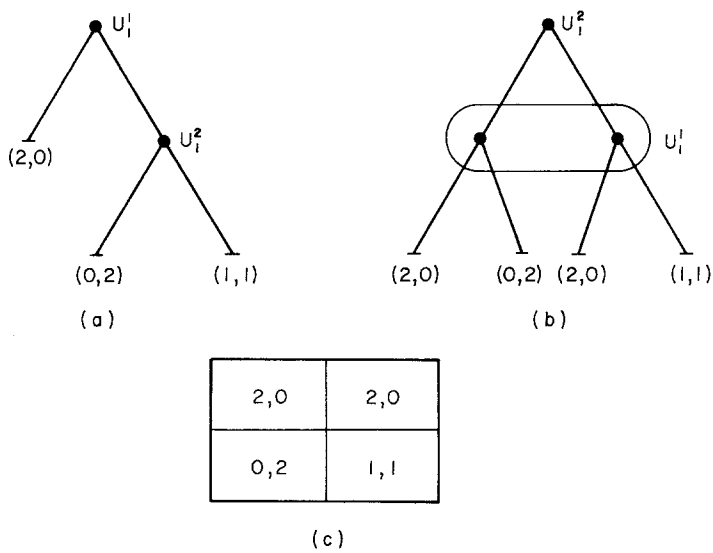

Figure 7. Two games in extensive form, (a) and (b), with the same strategic form, (c).

[12]There are $|S| = |S^1| \cdot |S^2| \cdot \ldots \cdot |S^n|$ terminal nodes.

It is clear that, in general, there may be many extensive forms with the same strategic form (up to "renaming" or "relabeling" of strategies). Such an example is presented in Figure 7. Thus, the extensive form contains more information about the game than the strategic form.

4. Mixed strategies

There are many situations in which a player's best behavior is to randomize when making his choice (recall, for instance, the game "matching pennies" of Examples 1.2 and 3.1). This leads to the concept of a "mixed strategy".

We need the following notation. Given a finite set A, the set of all probability distributions over A is denoted $\Delta(A)$. That is, $\Delta(A)$ is the $(|A|-1)$-dimensional unit simplex

$$\Delta(A) := \left\{ x = (x(a))_{a \in A} \colon x(a) \geqslant 0 \text{ for all } a \in A \text{ and } \sum_{a \in A} x(a) = 1 \right\}.$$

The set of *mixed strategies* X^i *of player* i is defined as $X^i := \Delta(S^i)$, where S^i is the set of pure strategies of player i. Thus, a mixed strategy $x^i = (x(s^i))_{s^i \in S^i} \in X^i$ of player i means that i chooses each pure strategy s^i with probability $x^i(s^i)$. From now on we will identify a pure strategy $s^i \in S^i$ with the corresponding unit vector in X^i.

Let $X := X^1 \times X^2 \times \cdots \times X^n$ denote the set of *n-tuples of mixed strategies*. For every $x = (x^1, x^2, \ldots, x^n) \in X$, the (expected)[13] *payoff of player* i is

$$H^i(x) := \sum_{s \in S} x(s) h^i(s),$$

where $x(s) := \Pi_{j \in N} x^j(s^j)$ is the probability, under x, that the pure strategy n-tuple $s = (s^1, s^2, \ldots, s^n)$ is played. We have thus defined a *payoff function* $H^i \colon X \to \Re$ for player i. Note that $\Gamma^* := (N; (X^i)_{i \in N}; (H^i)_{i \in N})$ is an n-player (infinite)[14] game in strategic form, called the *mixed extension* of the original game $\Gamma = (N; (S^i)_{i \in N}; (h^i)_{i \in N})$.

If the game is given in extensive form, one obtains [from (2.1)] an equivalent expression for H^i:

$$H^i(x) = \sum_{t \in L} p_x(t) g^i(t), \tag{4.1}$$

where, for each terminal node $t \in L(T)$, we let $p_x(t)$ be the probability that the terminal node t is reached under x; i.e., $p_x(t) := \Sigma_{s \in S} x(s) p_s(t)$.

[13]Again, recall Remark 1.1.
[14]The strategy spaces are infinite.

5. Equilibrium points

We come now to the basic solution concept for non-cooperative games.

A (mixed) n-tuple of strategies $x = (x^1, x^2, \ldots, x^n) \in X$ is an *equilibrium point*[15] if

$$H^i(x) \geqslant H^i(x^{-i}, y^i)$$

for all players $i \in N$ and all strategies $y^i \in X^i$ of player i, where $x^{-i} := (x^1, \ldots, x^{i-1}, x^{i+1}, \ldots, x^n)$ denotes the $(n-1)$-tuple of strategies, in x, of all the players except i. Thus $x \in X$ is an equilibrium whenever no player i can gain by changing his own strategy (from x^i to y^i), assuming that all the other players do not change their strategies. Note that the notion of equilibrium point is based only on the strategic form of the game; various "refinements" of it may however depend on the additional data of the extensive form (see the chapters on 'strategic equilibrium' and 'conceptual foundations of strategic equilibrium' in a forthcoming volume of this Handbook for a comprehensive coverage of this issue).

The main result is

Theorem 5.1 [Nash (1950, 1951)]. *Every (finite) n-person game has an equilibrium point (in mixed strategies).*

The proof of this theorem relies on a fixed-point theorem (e.g., Brouwer's or Kakutani's).

6. Games of perfect information

This section deals with an important class of games for which equilibrium points in *pure* strategies always exist.

An n-person game Γ (in extensive form) is a *game of perfect information* if all information sets are singletons, i.e., $|U^i| = 1$ for each player $i \in N$ and each information set $U^i \in \boldsymbol{I}^i$ of i. Thus, in a game of perfect information, every player, whenever called upon to make a choice, always knows exactly where he is in the game tree.

Examples of games of perfect information are Chess, Checkers, Backgammon (note that chance moves are allowed), Hex, Nim, and many others. In contrast, Poker, Bridge, Kriegsspiel (a variant of Chess where each player knows the position of his own pieces only) are games of *imperfect information*.

[15]Also referred to as "Nash equilibrium", "Cournot–Nash equilibrium", "non-cooperative equilibrium", and "strategic equilibrium".

(Another distinction, between *complete* and *incomplete information*, is presented and analyzed in the chapter on 'games of incomplete information' in a forthcoming volume of this Handbook.)

The historically first theorem of Game Theory deals with a game of perfect information.

Theorem 6.1 [Zermelo (1912)]. *In Chess, either*
(i) *White can force a win, or*
(ii) *Black can force a win, or*
(iii) *both players can force at least a draw.*

We say that a player can *force* an outcome if he has a strategy that makes the game terminate in that outcome, no matter what his opponent does. Zermelo's Theorem says that Chess is a so-called "determined" game: either there exists a pure strategy of one of the two players (White or Black) guaranteeing that he will always win, or each one of the two has a strategy guaranteeing at least a draw. Unfortunately, we do not know which of the three alternatives is the correct one (note that, in principle, this question is decidable in finite time, since the game tree of Chess is finite).[16]

The proof of Zermelo's Theorem is by induction, in a class of "Chess-like" games;[17,18] it is actually a special case of the following general result:

Theorem 6.2 [Kuhn (1953)]. *Every (finite) n-person game of perfect information has an equilibrium point in pure strategies.*

Proof. Assume by induction that the result is true for any game with less than m nodes. Consider a game Γ of perfect information with m nodes, and let r be the root of the game tree T. Let $v_1, v_2, \ldots, v_K$ denote the "sons" of r (i.e., those nodes that are connected to r by a branch), and let $T_1, T_2, \ldots, T_K$ (respectively), be the (disjoint) subtrees of T starting at these nodes. Each such T_k corresponds to a game Γ_k of perfect information (indeed, since Γ has perfect information, every information set is a singleton, thus completely included in one of the T_k's); Γ_k therefore possesses, by the induction hypothesis, an equilibrium point $s_k = (s_k^i)_{i \in N}$ in pure strategies. From these we construct a pure equilibrium point $s = (s^i)_{i \in N}$ for Γ, as follows. If r is a chance node, then s is just the "combination" (or "concatenation") of the s_k's i.e., $s^i(v) = s_k^i(v)$ for

[16]There are "Chess-like" games – for instance, Hex – where it can be proved that the first player can force a win, but nonetheless a winning strategy is not known. Other (simpler) games – e.g., Nim – have complete solutions (i.e., which player wins and by what strategy).

[17]For example, see Aumann (1989, pp. 1–4).

[18]Zermelo's Theorem 6.1 has been extended to two-person, zero-sum games by von Neumann and Morgenstern (1944, Section 15).

all nodes v of player i that belong to T_k. If r is a node of, say, player i, then, in addition to the above "combination", we put[19]

$$s^i(r) := \operatorname*{argmax}_{1 \leq k \leq K} h_k^i(s_k) ,$$

i.e., player i chooses at his first node r a branch k that leads to a subgame Γ_k where his equilibrium payoff is maximal. It is now straightforward to check that s is indeed a pure equilibrium point of Γ. □

Remark 6.3. The above proof yields a construction of equilibrium points in pure strategies by "backwards induction", from the terminal nodes to the root:[20] at each node of a player, choose a branch which leads to a subtree with the highest equilibrium payoff for that player;[21] at each chance node, average the equilibrium payoffs of the subtrees. Note that the equilibrium points constructed in this manner, when restricted to any subgame of the original game, yield equilibria in the subgame as well; such equilibria are called "(subgame) perfect". The reader is referred to the chapters on 'strategic equilibrium' and 'conceptual foundations of strategic equilibrium' in a forth-

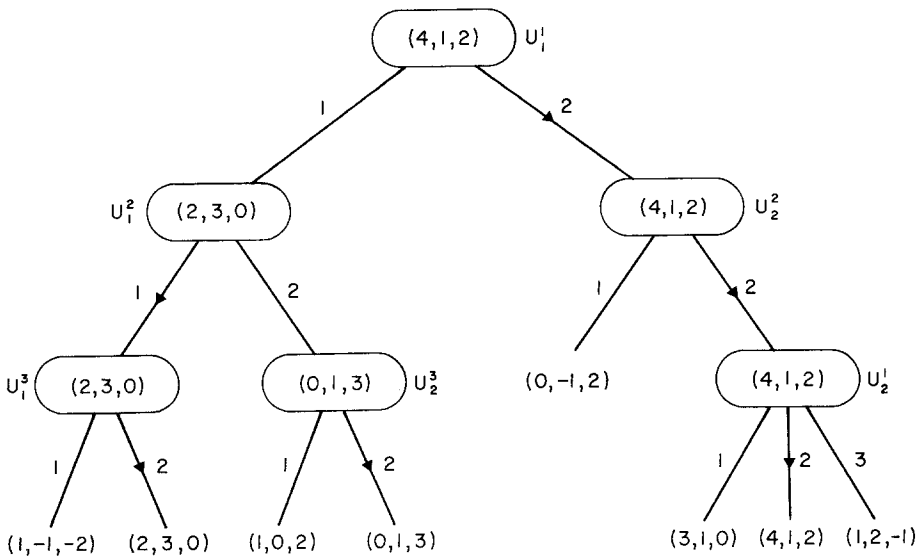

Figure 8. The game tree of Example 6.4 and the construction of the pure equilibrium point.

[19]We write h_k^i for the payoff function of player i in the subgame Γ_k, and "argmax" for a maximizer (if not unique, pick one arbitrarily).

[20]This is the standard procedure of "dynamic programming".

[21]Note that some of these choices need not be unique, in which case there is more than one such equilibrium.

coming volume of this Handbook for a discussion of these issues of backwards induction and perfection in relation to equilibrium points.

The following example illustrates the construction.

Example 6.4. See Figure 8: arrows indicate the choices forming the equilibrium strategies; the numbers in each node are the equilibrium payoffs for the subtree rooted at that node. The resulting equilibrium point is $s = (\langle 2, 2\rangle, \langle 1, 2\rangle, \langle 2, 2\rangle)$, with payoffs $h(s) = (4, 1, 2)$. □

The reader is referred to Chapter 3 in this volume for the development of the topic of games of perfect information, in particular in infinite games.

7. Behavior strategies and perfect recall

A pure strategy of a player is a complete plan for his choices in all possible contingencies in the game (i.e., at all his information sets). A mixed strategy means that the player chooses, *before* the beginning of the game, one such comprehensive plan at random (according to a certain probability distribution). An alternative method of randomization for the player is to make an independent random choice at each one of his information sets. That is, rather than selecting, for every information set, one definite choice – as in a pure strategy – he specifies instead a probability distribution over the set of choices there; moreover, the choices at different information sets are (stochastically) independent. These randomization procedures are called *behavior strategies*.

A useful way of viewing the difference between mixed and behavior strategies is as follows. One can think of each pure strategy 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 is a probability distribution on his library of books, so that, in playing according to a mixed strategy, the player chooses one book from his library by means of a chance device having the prescribed probability distribution. A behavior strategy is a single book of a different sort. Although each page still refers to a single information set of the player, it specifies a probability distribution over the choices at that set, not a specific choice.

We will see below that a behavior strategy is essentially a (special kind of) mixed strategy. Moreover, when a player has what is called "perfect recall", the converse also holds: every mixed strategy is fully "equivalent" to a behavior strategy.

We define a *behavior strategy* b^i *of player* i in the game Γ (in extensive form) as an element of

$$B^i := \prod_{U^i \in I^i} \Delta(C(U^i)), \tag{7.1}$$

that is, $b^i = (b^i(U^i))_{U^i \in I^i}$, where each $b^i(U^i)$ is a probability distribution over the set $C(U^i)$ of choices of player i at his information set U^i. We will write $b^i(U^i; c)$, rather than the more cumbersome $(b^i(U^i))(c)$, for the probability that the choice of player i at U^i is $c \in C(U^i)$; thus $\Sigma_{c \in C(U^i)} b^i(U^i; c) = 1$ and $b^i(U^i; c) \geqslant 0$.

Note that the linear dimension of the space of behavior strategies B^i of player i is $\Sigma_j (\nu^i_j - 1)$, whereas that of the space of mixed strategies X^i is $\Pi_j \nu^i_j - 1$, where j ranges from 1 to $k(i) = |I^i|$ and $\nu^i_j = |C(U^i_j)|$. Therefore B^i is a much smaller set than X^i.

Actually, the set B^i of behavior strategies of player i can be identified with a subset of the set X^i of mixed strategies of i. Indeed, given a behavior strategy, one may perform all the randomizations (for all information sets) before the game starts, which yields a (random) pure strategy – i.e., a mixed strategy. Formally, the *mixed strategy* x^i *corresponding to the behavior strategy* $b^i \in B^i$ is defined by $x^i = (x^i(s^i))_{s^i \in S^i}$, where

$$x^i(s^i) := \prod_{U^i \in I^i} b^i(U^i; s^i(U^i)) \tag{7.2}$$

for each pure strategy $s^i \in S^i$. Since $b^i(U^i; s^i(U^i))$ is the probability, under b^i, that player i chooses $s^i(U^i)$ at the information set U^i, it follows that $x^i(s^i)$ is precisely the probability that all his (realized) choices are according to the pure strategy s^i; in short, $x^i(s^i)$ is the probability, under b^i, of using s^i. The following lemma is thus immediate.

Lemma 7.3. *For any behavior strategy* $b^i \in B^i$ *of player* i, *the corresponding* x^i *given by* (7.2) *is a mixed strategy of* i *that is equivalent to* b^i.

We call the two strategies y^i and z^i of player i *equivalent* if they yield the same payoffs[22] (to everyone) for any strategies of the other players, i.e., $H^j(y^i, x^{-i}) = H^j(z^i, x^{-i})$ for all x^{-i} and all $j \in N$. Note that the argument given above shows that a stronger statement is actually true: for each terminal node $t \in L$, the probabilities that t is reached under (b^i, x^{-i}) and under (x^i, x^{-i}) are the same, for any x^{-i}.

[22]We have defined the (expected) payoff functions H^j for n-tuples of mixed strategies (see Section 4). The definition may be trivially extended to behavior strategies as well: use (4.1) with the probabilities $p_x(t)$ computed accordingly.

The difference between behavior and mixed strategies can thus be viewed as *independent* vs. (possibly) *correlated* randomizations (at the various information sets). This may be also seen by comparing directly the two definitions: B^i is a product of probability spaces [see (7.1)], whereas X^i is the probability space on the product [i.e., $\Delta(\Pi_{U^i \in I^i} C(U^i))$]. The following example is most illuminating.

Example 7.4 [Kuhn (1953)]. Consider a two-player, zero-sum[23] game in which player 1 consists of two people,[24] Alice and her husband Bill, and player 2 is a single person, Zeno. Two cards, one marked "High" and the other "Low", are dealt at random to Alice and Zeno. The person with the High card receives \$1 from the person with the Low card, and then has the choice of stopping or continuing the play. If the play continues, Bill, *not knowing the outcome of the deal*, instructs Alice and Zeno either to exchange or to keep their cards. Again, the holder of the High card receives \$1 from the holder of the Low card, and the game ends. See Figure 9 for the game tree (A = Alice,

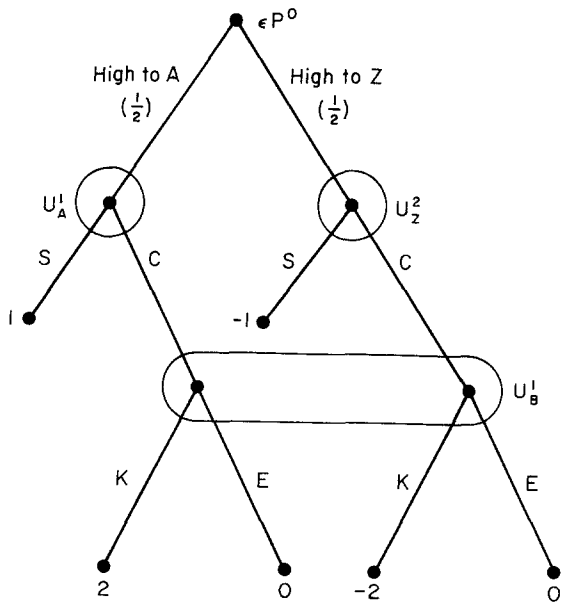

Figure 9. The game tree of Example 7.4.

[23]A two-player game is a *zero-sum* game if $h^1 + h^2 = 0$, i.e., what one player gains is what the other loses.

[24]With a joint bank account.

Table 2
Strategic form of Example 7.4

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, E⟩	$\frac{1}{2}\cdot 1+\frac{1}{2}\cdot(-1)=0$	$\frac{1}{2}\cdot 1+\frac{1}{2}\cdot(0)=\frac{1}{2}$
⟨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, E⟩	$\frac{1}{2}\cdot 0+\frac{1}{2}\cdot(-1)=-\frac{1}{2}$	$\frac{1}{2}\cdot 0+\frac{1}{2}\cdot(0)=0$

B = Bill, Z = Zeno; S = Stop, C = Continue, K = Keep and E = Exchange; payoffs at the terminal nodes are those paid by player 2 to player 1).

The strategic form of this game is given in Table 2. Note that the strategies ⟨S, K⟩ and ⟨C, E⟩ of player 1 are strictly dominated (by ⟨C, K⟩ and ⟨S, E⟩, respectively). Eliminating them yields the reduced strategic form of Table 3.

It is now easy to see that the unique optimal (mixed) strategies of the players are (0, 1/2, 1/2, 0) and (1/2, 1/2), respectively;[25] the value of the game is 1/4. Thus, in particular, player 1 can guarantee that his expected payoff will be at least 1/4, regardless of what player 2 will do.

Suppose now that player 1 uses only behavior strategies. Let $b^1 = (b^1(U^1_A), b^1(U^1_B)) = ((\alpha, 1-\alpha), (\beta, 1-\beta)) \in B^1$, i.e., Alice chooses S with probability α and C with probability $1-\alpha$, and Bill chooses K with probability β and E with probability $1-\beta$. [Note that the mixed strategy corresponding to b^1 is $(\alpha\beta,\ \alpha(1-\beta),\ (1-\alpha)\beta,\ (1-\alpha)(1-\beta))$ – see (7.2).] Then player 1's expected payoff is[26] $(1-\alpha)(\beta-1/2)$ if player 2 plays S, and $\alpha(1/2-\beta)$ if player 2 plays C. So the maximum payoff that player 1 can guarantee when restricted to behavior strategies is

Table 3
Reduced strategic form of Example 7.4

1 \ 2	⟨S⟩	⟨C⟩
⟨S, E⟩	0	$\frac{1}{2}$
⟨C, K⟩	$\frac{1}{2}$	0

[25]A mixed strategy of player 1 is written as the vector of probabilities for his pure strategies ⟨S, K⟩, ⟨S, E⟩, ⟨C, K⟩, ⟨C, E⟩, in that order; for player 2, the order is ⟨S⟩, ⟨C⟩.

[26]For example, the payoff if 2 plays S is computed as follows $(1/2)\cdot[\alpha\cdot 1+(1-\alpha)\cdot(\beta\cdot 2+(1-\beta)\cdot 0)]+(1/2)\cdot(-1)$.

$$\max_{0 \leq \alpha, \beta \leq 1} [\min\{(1-\alpha)(\beta - 1/2), \alpha(1/2 - \beta)\}],$$

which equals 0, since either $\beta - 1/2$ or $1/2 - \beta$ is always ≤ 0. □

Thus, behavior strategies for player 1 do a poorer job in this example than mixed strategies: player 1 can guarantee 1/4 with the latter, but only 0 with the former. Indeed, there is no behavior strategy corresponding to the unique optimal mixed strategy $x^1 = (0, 1/2, 1/2, 0)$ of player 1, since x^1 requires the randomizations at his two information sets to be fully correlated (rather than independent).

The reason that behavior strategies are inadequate in Example 7.4 is that player 1 consists of two agents who are *not allowed to communicate* during the play. This implies that, in going from U_A^1 to U_B^1, player 1 "forgets" what he knew, namely the outcome of the initial draw. Therefore the player needs to correlate, before the game starts, the random choices of his agents at his two information sets. Conversely, if a player always remembers what he knew as well as what he chose at all his previous nodes – in which case we say that the player has "perfect recall" – then he has no need to correlate the choices at his different information sets: indeed, being at any information set uniquely determines what happened at all the previous ones.

Formally, given a game tree T and a node v of T, we will denote by $T(v)$ the subtree of T with root at v. For an information set U and a choice there $c \in C(U)$, we will write $T(U; c)$ for the union of the trees $T(w)$, where w is connected to some node $v \in U$ by a branch labeled c [i.e., $T(U; c)$ is the "remainder" of the game after the information set U has been reached, and the choice c has been made there by the corresponding player]. We will say that *player i has perfect recall in the game Γ* (in extensive form) if the following condition is satisfied.[27] Let $v_1, v_2 \in P^i$ be nodes of player i, let $U_1^i \ni v_1$ and $U_2^i \ni v_2$ be the corresponding information sets of i, and assume that $v_2 \in T(v_1)$ (i.e., there exists a play of the game – a path – where v_2 "comes after" v_1); then there exists a unique choice $c \in C(U_1^i)$ 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 of perfect recall*. Note that a player who is a single person has perfect recall;[28] isolated agents are not needed to play the game for him. This is the case for most parlor games (but not for Bridge, when viewed as a two-player game, with each player consisting of two partners).

The condition in the definition of perfect recall can be separated into two parts:

[27]The original definition of Kuhn (1953) is different but equivalent to the one presented here; the advantage of the latter is that it is stated (and may be checked) directly on the structure of the tree.

[28]Provided he is not too absent-minded.

(i) "Player i recalls what he knew": $U_2^i \subset T(U_1^i)$, i.e., each node in U_2^i has to be reachable from some node in U_1^i. Otherwise, let $v_2' \in U_2^i \backslash T(U_1^i)$;[29] then, when reaching the information set U_2^i, player i does not recall whether the play went through U_1^i (as in v_2) or not (as in v_2') – he forgot what he knew (when he was at U_1^i). If player i had perfect recall, then, at U_2^i, he would be able to distinguish between v_2 and v_2', according to whether or not he has already been called upon to make a choice at U_1^i; hence v_2 and v_2' would have lied in different information sets. In Example 7.4 above, player 1 does not have perfect recall, since, if the play is: 'player 1 gets "High" and decides "Continue"', then at U_{B}^1 he no longer knows what he knew at U_{A}^1 (namely, the outcome of the draw).

(ii) "Player i recalls what he chose": $U_2^i \subset T(U_1^i; c)$ for a unique choice c at U_1^i. Otherwise, let $v_2' \in T(U_1^i; c') \cap U_2^i$ for some other choice $c' \neq c$; then, at U_2^i, player i does not recall whether his own choice at U_1^i was c (as is the case at v_2) or c' (as is the case at v_2') – he forgot what he chose (at U_1^i). If player i had perfect recall, then the nodes v_2 and v_2' would be distinguished by the choice he made at U_1^i, and would thus lie in different information sets. In Example 1.3 (in Section 1), for instance, condition (i) is satisfied but (ii) is not: U_2^1 contains two nodes (e and f) that follow different choices (1 and 2, respectively) of player 1 at U_1^1.

Example 7.5. In the game tree of Figure 10, assume that the root and the nodes at the third level do *not* belong to player 1 (it does not matter if they are chance nodes or personal nodes, and to which information sets they belong), and assume that all eight nodes at the fourth level are nodes of player 1. If player 1 has perfect recall, then condition (i) implies that $\{c, d, e, f\}$ is separated from $\{g, h, i, j\}$ (i.e., there can be no information set of player 1 containing nodes from both sets); condition (ii) separates $\{c, d\}$ from $\{e, f\}$, and $\{g, h\}$ from $\{i, j\}$. The dashed lines in Figure 10 show this partitioning. □

Example 7.6. Modify Example 7.5 by putting the two nodes a and b at the second level in one information set (of player 1). Then perfect recall for player 1 implies that $\{c, d, g, h\}$ must be separated from $\{e, f, i, j\}$ [condition (ii)]. □

We come now to the main result of this section.

Theorem 7.7 [Kuhn (1953)]. *Let Γ be a (finite) n-person game in which player i has perfect recall. Then for each mixed strategy $x^i \in X^i$ of player i there exists a corresponding behavior strategy $b^i \in B^i$ that is equivalent to x^i.*

[29]A "\" denotes set subtraction.

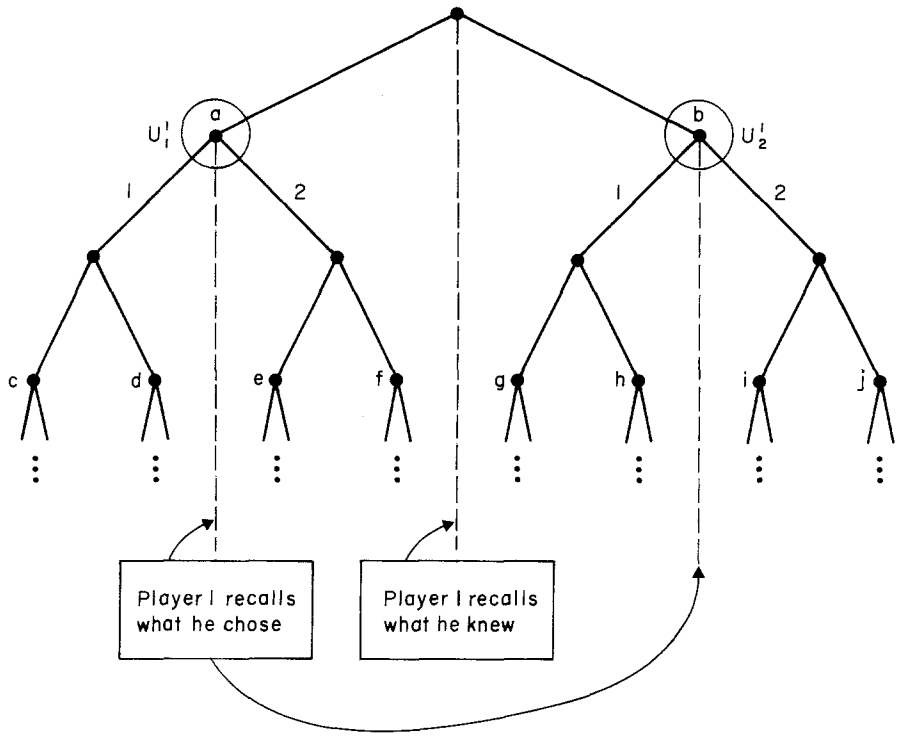

Figure 10. The game tree of Example 7.5.

Thus, having perfect recall is a sufficient[30] condition for restricting a player to behavior strategies instead of the (usually much) larger set of mixed strategies.

Proof (outline). Given a *pure* strategy $s^i \in S^i$ and an information set $U^i \in \boldsymbol{I}^i$, we will say that U^i is *reachable under* s^i if there exists a play of the game that goes through U^i and is consistent with s^i, i.e., there exists a path π in the tree that intersects U^i and, at every node of the player i on π, the path π follows the choice dictated by s^i. Given $x^i \in X^i$ we define the *corresponding behavior strategy* $b^i \in B^i$ as follows. For every information set $U^i \in \boldsymbol{I}^i$, let $\xi^i(U^i)$ be the probability that U^i is reachable under x^i, i.e., $\xi^i(U^i)$ is the sum of $x^i(s^i)$ over all $s^i \in S^i$ under which U^i is reachable. Similarly, for each choice $c \in C(U^i)$ at U^i, let $\xi^i(U^i; c)$ be the probability, under x^i, that U^i is reachable

[30]Kuhn (1953) shows that it is also *necessary*.

and the choice there is c [i.e., the sum of $x^i(s^i)$ over all s^i such that U^i is reachable and $s^i(U^i) = c$]. Finally, put

$$b^i(U^i; c) := \xi^i(U^i; c)/\xi^i(U^i) \tag{7.8}$$

if the denominator is positive, and define $b^i(U^i)$ arbitrarily when $\xi^i(U^i)$ vanishes (it will not matter, since then U^i is never reached when i plays x^i). One may interpret $b^i(U^i)$ as the "observed (random) behavior" of player i at U^i when he uses x^i.

Let x^{-i} be the strategies of the other players. We will show that

$$p_{(x^i, x^{-i})}(t) = p_{(b^i, x^{-i})}(t) \tag{7.9}$$

for each terminal node $t \in L$ (these are the probabilities that the play ends at t). Fix t, and denote by $\pi \equiv \pi(t)$ the path from the root to t. Let α be the probability that chance and all players except i always choose, at nodes on π, the branches along π; note that α depends on x^{-i} but not on the choices of player i. Similarly, let ξ (respectively, β) denote the probability under x^i (respectively, b^i) that all the choices of player i on π are along π. Then (7.9) becomes $\alpha \cdot \xi = \alpha \cdot \beta$, and it suffices to show that $\xi = \beta$ when $\alpha > 0$ (i.e., when t is reachable).

Let v_1 and v_2 be two consecutive nodes of i along π (i.e., there are no other nodes of i between them), and let $U_1^i \ni v_1$ and $U_2^i \ni v_2$ be the corresponding information sets; let c be the choice at v_1 along π. Player i has perfect recall, therefore $U_2^i \subset T(U_1^i; c)$, implying that U_2^i is reachable if and only if U_1^i is reachable and the choice there is c. Hence $\xi^i(U_2^i) = \xi^i(U_1^i; c)$; or, by (7.8), the denominator of $b^i(U_2^i; \cdot)$ equals the numerator of $b^i(U_1^i; c)$. To compute β we have to multiply the probabilities, under b^i, of all the choices of i along π. We thus obtain a telescoping product that simplifies to $\beta = \xi^i(U_m^i; c_m)$, where c_m is the choice of i along π at his last information set U_m^i on π. Now $\xi^i(U_m^i; c_m)$ is the probability, under x^i, that U_m^i is reachable and the choice there is c_m, or, equivalently (again by perfect recall, using induction), the probability that all the m choices of i on π are of the branches along π; but this probability is precisely ξ. Therefore $\beta = \xi$, and the proof is completed. □

Theorem 7.7 has been generalized to infinite games by Aumann (1964). As is pointed out there (p. 630), the extension of Kuhn's proof to the case where the length of the game as well as the number of choices at all information sets are at most countable poses no problems; the difficulties arise when there are uncountably many choices at some information set(s). In addition, games in which "time" is continuous pose special problems of their own, and do not easily fit into the framework of this chapter (see Chapter 3 in this volume, and

chapters on 'two-player games', 'differential games' and 'economic applications of differential games' in forthcoming volumes of this Handbook for some examples).

References

Aumann, R.J. (1964) 'Mixed and behavior strategies in infinite extensive games', in: M. Dresher, L.S. Shapley and A.W. Tucker, eds., *Advances in game theory*, Annals of Mathematics Studies 52. Princeton: Princeton University Press, pp. 627–650.

Aumann, R.J. (1989) *Lectures on game theory*. Boulder: Westview Press.

Kuhn, H.W. (1953) 'Extensive games and the problem of information', in: H.W. Kuhn and A.W. Tucker, eds., *Contributions to the theory of games*, Vol. II, Annals of Mathematics Studies 28. Princeton: Princeton University Press, pp. 193–216.

Nash, J.F. (1950) 'Equilibrium points in *n*-person games', *Proceedings of the National Academy of Sciences*, 36: 48–49.

Nash, J.F. (1951) 'Non-cooperative games', *Annals of Mathematics*, 54: 286–295

von Neumann, J. (1928) 'Zur Theorie der Gesellschaftsspiele', *Mathematische Annalen*, 100: 295–320. English translation: 'On the theory of games of strategy', in: A.W. Tucker and R.D. Luce, eds., *Contributions to the theory of games*, Vol. IV, Annals of Mathematics Studies 40. Princeton: Princeton University Press.

von Neumann, J. and O. Morgenstern (1944) *Theory of games and economic behavior*. Princeton: Princeton University Press.

Zermelo, E. (1912) 'Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels', *Proceedings of the Fifth International Congress of Mathematicians*, Vol. II, pp. 501–504.