Emerald

FOURTH EDITION

GAME THEORY

GUILLERMO OWEN

GAME THEORY

FOURTH EDITION

BY

GUILLERMO OWEN
Department of Mathematics
Naval Postgraduate School
Monterey, California

Emerald

United Kingdom – North America – Japan
India – Malaysia – China

Emerald Group Publishing Limited
Howard House, Wagon Lane, Bingley BD16 1WA, UK

First edition 1968
Second edition 1982
Third edition 1995
Fourth edition 2013

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

ISBN: 978-1-78190-507-4

Preface

Game theory once again proves to be a dynamic field. In less than 100 years since the name was coined, it has reached the status of a major branch of both mathematics and economics. It has also proved quite useful in several of the social sciences. Even biology and finance have been touched (and, I hope, improved) by game-theoretic analysis. The field continues to expand, and so a book that I last published in 1995 requires some modification. I am happy that Emerald Publishers is giving me a chance to bring out a new edition.

It would be very difficult to include all the new developments in the field. I have therefore merely added a chapter explaining some of the most important advances in game theory. In particular, I have tried to describe the developments that have been central to several of the Nobel Memorial Prizes in Economics.

I also wish to thank people who have helped me, by vigorous and stimulating discussion, to improve this book. These include Francesc Carreras, Hélène Ferrer, Gianfranco Gambarelli, Maurice Koster, Ines Lindner, Conrado Manuel, Gordon McCormick, Martha Saboyá, and Juan Tejada.

Finally, I am sad to mention the deaths of some members of the game theory community. Michael Maschler, my mentor and one of my best friends, was instrumental in the development of the bargaining set and related concepts. John Harsanyi, also a good friend, was one of the first to study games with incomplete information. I will always miss them both.

Guillermo Owen

Contents

Chapter 1

Definition of a Game

1.1. General Notions

The general idea of a game is that with which we are familiar in the context of parlor games. Starting from a given point, there is a sequence of personal moves, at each of which one of the players chooses from among several possibilities; interspersed among these there may also be chance, or random, moves such as throwing a die or shuffling a deck of cards.

Examples of this type of game are chess, in which there are no chance moves (except for the determination of who shall play first), bridge, in which chance plays a much greater part, but in which skill is still important, and roulette, which is entirely a game of chance in which skill plays no part.

The examples of bridge and chess help to point out another important element of a game. In fact, in a chess game each player knows every move that has been made so far, while in bridge a player's knowledge is usually very imperfect. Thus, in some games, a player is unable to determine which of several possible moves has actually been made, either by an opposing player, or by chance. The practical result of this is that, when a player makes a move, he does not know the exact position of the game, and must make his move remembering that there are several possible actual positions.

Finally, at the end of a game, there is normally some payoff to the players (in the form of money, prestige, or satisfaction) which depends on the progress of the game. We may think of this as a function which assigns a payoff to each "terminal position" of the game.

1.2. Games in Extensive Form

In our general idea of a game, therefore, three elements enter: (1) alternation of moves, which can be either personal or random (chance) moves, (2) a possible lack of knowledge, and (3) a payoff function.

We define, first, a *topological tree* or *game tree* as a finite collection of nodes, called *vertices*, connected by lines, called *arcs*, so as to form a connected figure which includes no simple closed curves. Thus it follows that, given any two vertices A and B, there is a unique sequence of arcs and nodes joining A to B.

From this we obtain

1.2.1 Definition. Let Γ be a topological tree with a distinguished vertex A. We say that a vertex C *follows* the vertex B if the sequence of arcs joining A to C passes through B. We say C follows B *immediately* if C follows B and, moreover, there is an arc joining B to C. A vertex X is said to be *terminal* if no vertex follows X.

1.2.2 Definition. By an *n-person game in extensive form* is meant

(α) a topological tree Γ with a distinguished vertex A called the *starting point* of Γ;
(β) a function, called the *payoff function*, which assigns an n-vector to each terminal vertex of Γ;
(γ) a partition of the nonterminal vertices of Γ into $n+1$ sets $S_0, S_1, \ldots, S_n$, called the *player sets*;
(δ) a probability distribution, defined at each vertex of S_0, among the immediate followers of this vertex;
(ε) for each $i = 1, \ldots, n$, a subpartition of S_i into subsets S_i^j, called *information sets*, such that two vertices in the same information set have the same number of immediate followers and no vertex can follow another vertex in the same information set;
(ζ) for each information set S_i^j, an index set I_i^j, together with a $1-1$ mapping of the set I_i^j onto the set of immediate followers of each vertex of S_i^j.

The elements of a game are seen here: condition α states that there is a starting point; β gives a payoff function; γ divides the moves into chance moves (S_0) and personal moves which correspond to the n players ($S_i, \ldots, S_n$); δ defines a randomization scheme at each chance move; ε divides a player's moves into "information sets": he knows in which information set he is, but not which vertex of the information set.

1.2.3 Example. In the game of *matching pennies* (Figure 1.1), player I chooses "heads" (H) or "tails" (T). Player II, not knowing player I's choice, also chooses "heads" or "tails." If the two choose alike, then player II wins a cent from player I; otherwise, player I wins a cent from II. In the game tree shown, the vectors at the

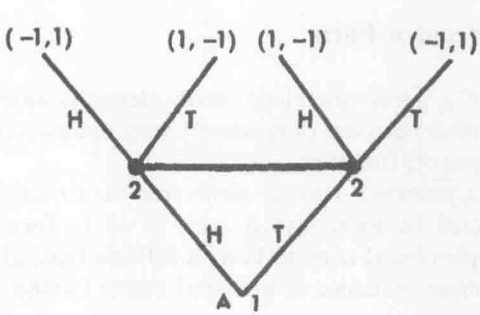

Figure 1.1.

terminal vertices represent the payoff function; the numbers near the other vertices denote the player to whom the move corresponds. The shaded area encloses moves in the same information set.

1.2.4 Example. The game of pure strategy, or GOPS, is played by giving each of two players an entire suit of cards (13 cards), A third suit is shuffled, and the cards of this third suit are then turned up, one by one. Each time one has been turned up, each player turns up one of his cards at will: the one who turns up the larger card "wins" the third card. (If both turn up a card of the same denomination, neither wins.) This continues until the three suits are exhausted. At this point, each player totals the number of spots on the cards he has "won"; the "score" is the difference between what the two players have.

With 13-card suits, this game's tree is too large to give here; however, we can give part of the tree of an analogous game using three-card suits (Figure 1.2).

There is a single chance move, the shuffle, which orders the cards in one of the six possible ways, each having a probability of $\frac{1}{6}$. After this the moves correspond to the two players, I and II, until the game ends. We have drawn parts of the game tree, including the initial point, several branches, and four of the terminal points. The remaining branches are similar to those we have already drawn. With respect to information, we have

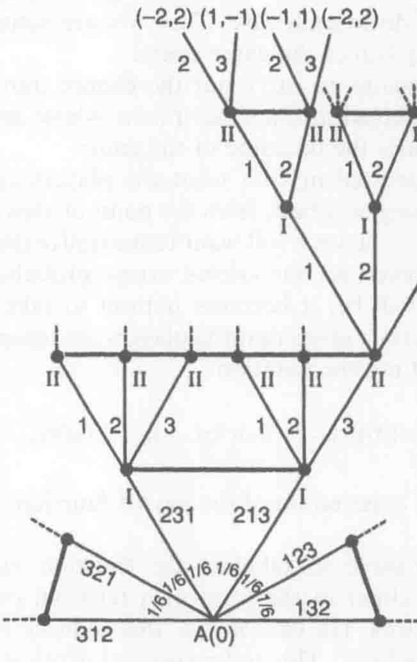

Figure 1.2.

1.2.5 Definition. Player i is said to have *perfect information* in Γ if his information sets S_i^j each consist of one element. The game Γ is said to have perfect information if each player has perfect information in Γ.

For example, chess and checkers have perfect information, whereas bridge and poker do not.

1.3. Strategies: The Normal Form

The intuitive meaning of a strategy is that of a plan for playing a game. We may think of a player as saying to himself, "If such and such happens, I'll act in such and such a manner." Thus, we have

1.3.1 Definition. By a *strategy* for player i is meant a function which assigns, to each of player i's information sets S_i^j, one of the arcs which follows a representative vertex of S_i^j.

The set of all strategies for player i will be called Σ_i.

In general, we are accustomed to the idea that a player decides his move in a game only a few moves, at best, in advance, and quite usually only at the moment he must make it. In practice this must be so, for in a game such as chess or poker the number of possible moves is so great that no one can plan for every contingency very much in advance. From a purely theoretic point of view, however, we can overlook this practical limitation, and assume that, even before the game starts, each player has already decided what he will do in each case. Thus, we are actually assuming that each player chooses a strategy before the game starts.

Since this is so, it only remains to carry out the chance moves. Moreover, the chance moves may all be combined into a single move, whose result, together with the strategies chosen, determines the outcome of the game.

Actually, what we are interested in, and what the players are interested in, is deciding which one of the strategies is best, from the point of view of maximizing the player's share of the payoff (i.e., player i will want to maximize the ith component of the payoff function). As, however, no one knows, except probabilistically, what the results of the chance moves will be, it becomes natural to take the mathematical expectation of the payoff function, given that the players are using a given n-tuple of strategies. Therefore, we shall use the notation

$$\pi(\sigma_1, \sigma_2, \ldots, \sigma_n) = (\pi_1(\sigma_1, \ldots, \sigma_n), \pi_2(\ldots), \ldots, \pi_n(\sigma_1, \ldots, \sigma_n))$$

to represent the mathematical expectation of the payoff function, given that player i is using strategy $\sigma_i \in \Sigma_i$.

From this, it becomes possible to tabulate the function $\pi(\sigma_1, \ldots, \sigma_n)$ for all possible values of $\sigma_1, \ldots, \sigma_n$, either in the form of a relation, or by setting up an n-dimensional array of n-vectors. (In case $n = 2$, this reduces to a matrix whose elements are pairs of real numbers.) This n-dimensional array is called the *normal form* of the game Γ.

1.3.2 Example. In the game of matching pennies (see Example 1.2.3) each player has the two strategies, "heads" and "tails." The normal form of this game is the matrix

	H	T
H	$(-1,1)$	$(1,-1)$
T	$(1,-1)$	$(-1,1)$

(where each row represents a strategy of player I, and each column a strategy of player II).

1.3.3 Example. Consider the following game: An integer z is chosen at random, with possible values 1, 2, 3, 4 (each with probability $\frac{1}{4}$). Player I, without knowing the results of this move, chooses an integer x. Player II, knowing neither the result of the chance move nor I's choice, chooses an integer y. The payoff is

$$(|y-z|-|x-z|, |x-z|-|y-z|)$$

i.e., the point is to guess close to z.

In this game each player has four strategies: 1, 2, 3, 4, since other integers are of little use. If, for instance, I chooses 1 and II chooses 3, then the payoff will be $(2,-2)$ with probability $\frac{1}{4}$, $(0,0)$ with probability $\frac{1}{4}$, and $(-2,2)$ with probability $\frac{1}{2}$. The expected payoff, then, is $\pi(1,3)=(-\frac{1}{2},\frac{1}{2})$. Calculating all the values of $\pi(\sigma_1,\sigma_2)$, we obtain

	1	2	3	4
1	$(0,0)$	$(-\frac{1}{2},\frac{1}{2})$	$(-\frac{1}{2},\frac{1}{2})$	$(0,0)$
2	$(\frac{1}{2},-\frac{1}{2})$	$(0,0)$	$(0,0)$	$(\frac{1}{2},-\frac{1}{2})$
3	$(\frac{1}{2},-\frac{1}{2})$	$(0,0)$	$(0,0)$	$(\frac{1}{2},-\frac{1}{2})$
4	$(0,0)$	$(-\frac{1}{2},\frac{1}{2})$	$(-\frac{1}{2},\frac{1}{2})$	$(0,0)$

1.3.4 Definition. A game is said to be finite if its tree contains only finitely many vertices.

Under this definition, most of our parlor games are finite. Chess, for instance, is finite, thanks to the laws which end the game after certain sequences of moves.

It should be seen that, in a finite game, each player has only a finite number of strategies.

1.4. Equilibrium *n*-Tuples

1.4.1 Definition. Given a game Γ, a strategy n-tuple $(\sigma_1^*,\sigma_2^*,\ldots,\sigma_n^*)$ is said to be *in equilibrium*, or an *equilibrium n-tuple*, if and only if, for any $i=1,\ldots,n$, and $\hat{\sigma}_i\in\Sigma_i$.

$$\pi_i(\sigma_1^*, \ldots, \sigma_{i-1}^*, \hat{\sigma}_i, \sigma_{i+1}^*, \ldots, \sigma_n^*) \leq \pi_i(\sigma_1^*, \ldots, \sigma_n^*)$$

In other words, an n-tuple of strategies is said to be in equilibrium if no player has any positive reason for changing his strategy, assuming that none of the other players is going to change strategies. If, in such a case, each player knows what the others will play, then he has reason to play the strategy which will give such an equilibrium n-tuple, and the game becomes very stable.

1.4.2 Example. In the game with normal form

	β_1	β_2
α_1	(2, 1)	(0, 0)
α_2	(0, 0)	(1, 2)

both (α_1, β_1) and (α_2, β_2) are equilibrium pairs.

Unfortunately, not every game has equilibrium n-tuples. As an example, the game of matching pennies (Example 1.3.2) has no equilibrium pairs.

In general, if a game has no equilibrium n-tuples, we usually see the several players trying to outguess each other, keeping their strategies secret. This suggests (and it is indeed true) that in games of perfect information, equilibrium n-tuples exist.

To prove this statement, we must study the question of decomposition of a game.

A game Γ will be said to *decompose at a vertex* X if there are no information sets which include vertices from both of (a) X, and all its followers, and (b) the remainder of the game tree. In this case, we can distinguish the subgame, Γ_X, consisting of X, and all its followers, and the quotient game, Γ/X, which consists of all the remaining vertices, plus X. For the quotient game, X will be a terminal vertex; the payoff here can be considered to be Γ_X: i.e., the payoff at this vertex is a play of the subgame Γ_X.

Now, as we have seen, a strategy for i is a function whose domain consists of the information sets of player i. If we decompose a game at X, then we can also decompose the strategy σ into two parts: $\sigma_{|\Gamma/X}$, obtained by restricting σ to information sets in Γ/X, and $\sigma_{|\Gamma_X}$, obtained by restricting σ to Γ_X. Conversely, a strategy for Γ/X and a strategy for Γ_X can be combined in the obvious way to yield a strategy for the larger game Γ.

1.4.3 Theorem. Let Γ decompose at X. For $\sigma_i \in \Sigma_i$, assign to X (considered as a terminal vertex of Γ/X) the payoff

$$\pi_X(\sigma_{1|\Gamma_X}, \sigma_{2|\Gamma_X}, \ldots, \sigma_{n|\Gamma_X})$$

In this case

$$\pi(\sigma_1, \ldots, \sigma_n) = \pi_{\Gamma/X}(\sigma_{1|\Gamma/X}, \ldots, \sigma_{n|\Gamma/X})$$

The proof of this theorem is clear and can be left as an exercise to the reader. Briefly, it is only necessary to verify that, for each possible outcome of the chance moves, the same terminal vertex is eventually reached either in the original or in the decomposed game.

With this, we can prove.

1.4.4 Theorem. Let Γ decompose at X, and let $\sigma_i \in \Sigma_i$ be such that (a) $(\sigma_{1|\Gamma_X}, \ldots, \sigma_{n|\Gamma_X})$ is an equilibrium n-tuple for Γ_X, and (b) $(\sigma_{1|\Gamma/X}, \ldots, \sigma_{n|\Gamma/X})$ is an equilibrium n-tuple for Γ/X, with the payoff $\pi(\sigma_{1|\Gamma_X}, \ldots, \sigma_{n|\Gamma_X})$ assigned to the terminal vector X. Then $(\sigma_1, \ldots, \sigma_n)$ is an equilibrium n-tuple for Γ.

Proof. Let $\hat{\sigma}_i \in \Sigma_i$. Because $(\sigma_{1|\Gamma_X}, \ldots, \sigma_{n|\Gamma_X})$ is an equilibrium n-tuple for Γ_X, it follows that

$$\pi_i(\sigma_{1|\Gamma_X}, \ldots, \hat{\sigma}_{i|\Gamma_X}, \ldots, \sigma_{n|\Gamma_X}) \le \pi_i(\sigma_{1|\Gamma_X}, \ldots, \sigma_{n|\Gamma_X})$$

On the other hand, by (b) we know that, if we assign the payoff $\pi(\sigma_{1|\Gamma_X}, \ldots, \sigma_{n|\Gamma_X})$ to the vector X, then

$$\pi_i(\sigma_{1|\Gamma/X}, \ldots, \hat{\sigma}_{i|\Gamma/X}, \ldots, \sigma_{n|\Gamma/X}) \le \pi_i(\sigma_{1|\Gamma/X}, \ldots, \sigma_{n|\Gamma/X})$$

Now, the payoff (for a given set of strategies) is a weighted average of the payoffs at some of the terminal vertices of a tree. Hence, if the payoff to player i at a given terminal vertex (X, in this case) is decreased, his expected payoff, for any choice of strategies, will either remain equal or be decreased. Thus, applying Theorem 1.4.3, we find that

$$\pi_i(\sigma_1, \ldots, \hat{\sigma}_i, \ldots, \sigma_n) \le \pi_i(\sigma_1, \ldots, \sigma_n)$$

and so $(\sigma_1, \ldots, \sigma_n)$ is an equilibrium n-tuple.

This is all we need to prove.

1.4.5 Theorem. Every finite n-person game with perfect information has an equilibrium n-tuple of strategies.

Proof. We shall define the length of a game as the largest possible number of edges that can be passed before reaching a terminal vertex, i.e., the largest possible number of moves before the game ends. Clearly a finite game has finite length. The proof is by induction on the length of the game.

If Γ has length 0, the theorem is trivially true. If it has length 1, then at most one player gets to move, and he obtains equilibrium by choosing his best alternative. If Γ has length m, then it decomposes (having perfect information) into several subgames of length less than m. By the induction hypothesis, each of these subgames has an equilibrium n-tuple; by Theorem 1.4.4 these form an equilibrium n-tuple for Γ.

1.5. The Monty Hall Game

We conclude this chapter with an example which, in the past, raised some controversy. Based on a well-known television show (*Let's Make a Deal*), the problem caused serious discussion among mathematicians and others who treated it as a problem in probability theory. (See items 3, 6, 8 in the bibliography for background on the problem and the controversy.) We will, rather, treat it as a game.

In this game, a Contestant (player I) is shown three screens. Behind one of these screens is an expensive, brand-new automobile. There is a goat behind each of the other two screens. It is understood that each screen has probability $\frac{1}{3}$ of being the "good" screen (i.e., the one with the automobile).

Once chance has chosen the good screen, player I, in ignorance of the results of the chance move, chooses one of the screens (call this screen a). At this point Monty Hall, the host of the show (player II), opens another one of the screens (call this screen b), to reveal a goat. In doing this, Monty knows the outcome of the chance move.

Player I may then *insist* on screen a, or opt to *switch* to screen c. In either case, he will receive whatever is behind his final choice (a or c). It is assumed that the automobile is much more valuable than the goat. What is the player's best strategy?

If we treat this as a game, the first thing we notice is that the game is not well defined. In fact, we are told how the Contestant may act, but we are told neither how Monty may act, nor what Monty's payoff will be. There would seem three possibilities.

Case 1. Monty the Dummy

Monty has no latitude in his actions. If the Contestant's choice, screen a, is the good screen, then Monty must open either b or c, each with probability $\frac{1}{2}$. If screen a is one of the two bad screens, then Monty must open the other bad screen (whether it be b or c).

Now, Contestant has two strategies: Insist (on his original choice) and Switch (from his original choice, a, to the remaining unopened screen, c).

Consider his first strategy, Insist. In this case, Contestant will win if (and only if) a was indeed the good screen. As mentioned above, this has probability $\frac{1}{3}$.

Consider next the strategy Switch. Then, if a is the good screen, Contestant will lose. If, on the other hand, a is a bad screen, then Monty will open the other bad screen. Contestant, by switching to the remaining screen, will win. This has probability $\frac{2}{3}$.

Thus, in this case, Contestant will do better by Switching.

Case 2. The Friendly Monty

Monty has some freedom. After player I has chosen screen a, Monty *may* (but need not) open a bad screen–either b or c. If Monty does not open a screen, then Contestant will be left with his original choice, a. Moreover, Monty wants Contestant to *win*.

In this case, Contestant, with Monty's help, is sure to win. In fact, if the original choice, a, is good, then Monty will not open any screen, and Contestant will have to open his choice, a. If, on the other hand, a is a bad screen, then Monty will open the other bad screen. Contestant, by switching, will win.

Thus, the strategy Switch will win with probability 1.

Case 3. The Unfriendly Monty

Monty has some freedom. After player I has chosen screen a, Monty *may* (but need not) open a bad screen–either b or c. If Monty does not open a screen, then Contestant will be left with his original choice, a. This time, however, Monty wants Contestant to *lose*.

In this case, Contestant, by Insisting, will win with probability $\frac{1}{3}$.

Suppose, however, Contestant uses his other strategy, Switch. If a is the good screen, then Monty's best strategy is to open one of the other screens–either b or c (both of which are bad). Contestant, by switching to the other bad screen, will lose. If a is a bad screen, then Monty's best strategy is not to open any screen, and thus Contestant will be left with his original choice. Thus Contestant will certainly lose.

In this case, then, Contestant will win with probability $\frac{1}{3}$ by Insisting, and probability 0 by Switching.

So what, then, is the best strategy for Contestant? Clearly it depends on the rules of the game, and also on Monty's objective: Contestant should Switch if Monty is a dummy, and also if Monty wants him to win; he should Insist if Monty wants him to lose. Unfortunately, Contestant is apparently not told the rules of the game, and probably does not know what Monty wants. Perhaps Monty's sponsors want Contestant to win–it will be good publicity for their products. Perhaps they want Contestant to lose–they have spent too much money already on prizes.

Clearly Contestant needs more information. This is an example of a *game with incomplete information*. We look at possible approaches to these games in Chapter 6 of this book.

Problems

1. An infinite game, even with perfect information, need not have an equilibrium n-tuple.
   (a) Consider a two-person game in which the two players alternate, and, at each move, each player chooses one of the two digits 0 and 1. If the digit x_i is chosen at the ith move, then each play of the game corresponds to a number

   $$x = \sum_{i=1}^{\infty} x_i 2^{-i}$$

   in the interval [0, 1]. Then, player I wins one unit from player II if $x \in S$, and loses one unit if $x \notin S$, where S is some subset of [0, 1].
   (b) Each player has exactly $2^{\aleph_0}$ strategies, which can therefore be indexed σ_β, τ_β respectively for $\beta < \alpha$, where α is the smallest ordinal preceded by at least $2^{\aleph_0}$ ordinals.
   (c) Let $\langle \sigma, \tau \rangle$ denote the play (or number x) obtained if the players choose strategies σ and τ, respectively. For each strategy σ of I, player II has $2^{\aleph_0}$ strategies τ which will give different values of $\langle \sigma, \tau \rangle$ (and similarly for each strategy τ of II).