MATHEMATICAL METHODS AND THEORY IN GAMES, PROGRAMMING, AND ECONOMICS

VOLUME II

S. KARLIN

Mathematical Methods and Theory in Games, Programming, and Economics

THIS BOOK IS IN THE
ADDISON-WESLEY SERIES IN STATISTICS

Z. W. Birnbaum
Consulting Editor

Mathematical

in Games, Programming

Volume II

THE THEORY OF INFINITE GAMES

Methods and Theory and Economics

by

SAMUEL KARLIN

Stanford University

PERGAMON PRESS
LONDON-PARIS
1959

ADDISON-WESLEY PUBLISHING COMPANY, INC.
READING, MASSACHUSETTS, U.S.A.
LONDON, ENGLAND

Printed in the United States of America

Library of Congress Catalog Card No. 60-5402

Sole distributors in the United Kingdom:

PERGAMON PRESS LTD.
4 and 5 Fitzroy Square, London W. 1

PERGAMON PRESS S.A.R.L.
*24 Rue des Écoles, Paris V^e^*

NOTE TO THE READER

This volume differs from Volume I in only one essential respect: the dimensionality of the strategy spaces available to the players of the game. In Volume I our strategy spaces were finite-dimensional; in this volume they are infinite-dimensional.

In order that this volume may be studied independently of Volume I, the essential background material from Volume I is reproduced here in its entirety: i.e., Chapter 1, which presents the underlying concepts of game theory in their simplest form and introduces the basic notation; and the appendixes, which review matrix theory, the properties of convex sets, and miscellaneous topics of function theory.

The organization of this volume is similar in all respects to that of Volume I.

SAMUEL KARLIN

Stanford, California
August 1959

LOGICAL INTERDEPENDENCE OF THE VARIOUS CHAPTERS

(applies only to unstarred sections of each chapter)

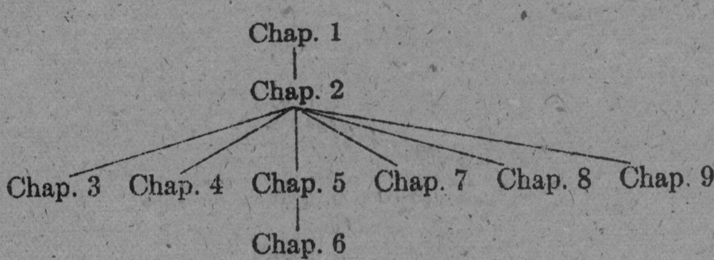

NOTATION

Vectors

A vector $\mathbf{x}$ with n components is denoted by

$$\mathbf{x} = \langle x_1, x_2, \ldots, x_n \rangle.$$

$\mathbf{x} \geq 0$ shall mean $x_i \geq 0$ $(i = 1, 2, \ldots, n)$.

$\mathbf{x} > 0$ shall signify $x_i \geq 0$ and at least one component of $\mathbf{x}$ strictly positive.

$\mathbf{x} \gg 0$ denotes that $x_i > 0$ $(i = 1, 2, \ldots, n)$ (all components are positive).

The inner product of two real vectors $\mathbf{x}$ and $\mathbf{y}$ in E^n is denoted by

$$(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^{n} x_i y_i.$$

In the complex case

$$(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^{n} x_i \bar{y}_i,$$

where $\bar{y}_i$ denotes conjugate complex.

The distance between two vectors $\mathbf{x}$ and $\mathbf{y}$ is denoted by $|x - y|$.

Matrices

A matrix $\mathbf{A}$ in terms of its components is denoted by $\|a_{ij}\|$.

A matrix $\mathbf{A}$ applied to vector $\mathbf{x}$ in the manner $\mathbf{xA}$ gives the vector

$$\langle \sum x_i a_{i1}, \sum x_i a_{i2}, \ldots, \sum x_i a_{in} \rangle.$$

Similarly,

$$\mathbf{Ay} = \langle \sum a_{1j} y_j, \sum a_{2j} y_j, \ldots, \sum a_{nj} y_j \rangle.$$

$(\mathbf{Ax})_j$ denotes the jth component of the vector $\mathbf{Ax}$.

The transpose of $\mathbf{A}$ is denoted by $\mathbf{A}'$.

The determinant of $\mathbf{A}$ is usually designated by $|\mathbf{A}|$ and alternatively by det $\mathbf{A}$.

Convex Sets

The convex set spanned by a set S is denoted alternatively by Co $(S) = [S]$.

The convex cone, spanned by S, is denoted by P_S or $\mathfrak{P}_S$.

Distributions

The symbols

$$\mathbf{x}_{t_0}, \mathbf{y}_{t_0} \quad \text{and} \quad \mathbf{I}_{t_0} \quad (0 \leq t_0 \leq 1)$$

are used interchangeably to represent a distribution defined on the unit interval which concentrates its full mass at the point t_0, that is,

$$\mathbf{x}_{t_0}(\xi) = \begin{cases} 0 & \xi < t_0, \\ 1 & \xi \geq t_0. \end{cases}$$

The symbol

$$\sum_{i=1}^{k} \lambda_i I_{\xi_i}$$

represents the probability distribution function with jumps λ_i located at ξ_i.

CONTENTS

CHAPTER 1

THE DEFINITION OF A GAME AND THE MIN-MAX THEOREM

1.1 Introduction. Games in normal form. The mathematical theory of games of strategy deals with situations involving two or more participants with conflicting interests. The outcome of such games is usually controlled partly by one side and partly by the opposing side or sides; it depends to some extent on chance, but primarily on the intelligence and skill employed by the participants. Aside from games proper, such as poker and chess, there are many conflicting situations to which the theory of games can be applied, notably in certain areas of operations research, economics, politics, and military science.

We shall first consider only two-person zero-sum games, i.e., games with only two participants (competing persons, teams, firms, nations) in which one participant wins what the other loses. It should be noted that the terms of this situation exclude the possibility of any bargaining between the participants. We shall also be concerned with two-person constant-sum games, in which the two players compete irreconcilably for the greatest possible share of the kitty. By suitable renormalization such a game may be converted into a zero-sum game.

A fundamental concept in game theory is that of *strategy*. A strategy for Player I is a complete enumeration of all actions Player I will take for every contingency that might arise, whether the contingency be one of chance or one created by a move of the opposing player. What a strategy is should not be interpreted too naively. It might seem that once Player I has chosen a strategy, every move he makes at any stage is determined, in the sense that he has in mind at the beginning of the game a sequence of moves which he will carry out no matter what his opponent does. However, what we mean by "strategy" is a rule which in determining Player I's ith moves takes into account everything that has happened before his ith turn. The reason a player does not change a strategy during a game is not that the strategy has committed him to a sequence of moves he must make no matter what his opponent does, but that it gives him a move to make in any circumstances that may arise.

It is usual, in describing a game, to regard all possible procedures, good or bad, as possible strategies. Even in simple games the number of possible strategies is often forbidding. Consider the game of ticktacktoe.* Suppose Player I makes the first move. There are nine possible positions for his first cross. Player II then has eight possible moves he can make, and what Player I does at his second opportunity to move will depend on the preceding move of Player II. In this situation there are seven possible moves he may make. Player I's third move will, of course, depend on all the preceding moves of both players; and so on.

A strategy might start off as follows: Player I's first cross is to be made in the upper right-hand square. If Player II marks either the square below or the square to the left, Player I makes his second cross in the center square; if Player II marks the center square, Player I makes his second cross in the lower left-hand square; if Player II marks any of the other five squares, Player I makes his second cross in the square immediately below the first; and thus the description goes on. Player I might even embody in his strategy the possibility of randomizing, according to a fixed probability distribution, among alternatives at a given move.

Clearly, an enormous number of possible strategies present themselves even for such a simple game as ticktacktoe. Although many of them intuitively seem to be poor strategies, we are obliged to include all the possibilities in order to give a complete description of the game. In the course of this book mathematical tools for manipulating and analyzing these large sets of strategies will be developed.

A second fundamental concept in game theory is that of the *pay-off*. The pay-off is the connecting link between the set of strategies open to Player I and the set open to Player II. Specifically, it is a rule that tells how much Player I may be expected to win from Player II if Player I chooses any particular strategy from his set of strategies and Player II chooses any particular strategy from his set. The pay-off function is always evaluated in terms of the appropriate utility units (see the notes to this section on p. 20).

We are now ready for a formal definition of a game.

A game is defined to be a *triplet* $\{X, Y, K\}$, where X denotes the space of strategies for Player I, Y signifies the space of strategies of Player II, and K is a real-valued function of X and Y. Player I chooses a strategy $\mathbf{x}$ from X and Player II chooses a strategy $\mathbf{y}$ from Y. For the pair $\{\mathbf{x}, \mathbf{y}\}$ the pay-off to Player I is $K(\mathbf{x}, \mathbf{y})$ and the pay-off to Player II is $-K(\mathbf{x}, \mathbf{y})$. We shall call K the pay-off kernel.

In the absence of a statement to the contrary, the following conditions are assumed to be satisfied throughout this chapter:

---

* Ticktacktoe is played on a 3×3 matrix grid. Players move alternately and on each turn are allowed to capture one of the remaining free squares. The first player who takes possession of three squares which are on a single horizontal, vertical, or diagonal line wins.

(a) X is a convex, closed, bounded set in Euclidean n-space E^n.

(b) Y is a convex, closed, bounded set in Euclidean m-space E^m.

(c) The pay-off kernel K is a convex linear function of each variable separately. Explicitly,

$$K[\lambda\mathbf{x}_1 + (1 - \lambda)\mathbf{x}_2, \mathbf{y}] = \lambda K(\mathbf{x}_1, \mathbf{y}) + (1 - \lambda)K(\mathbf{x}_2, \mathbf{y})$$

and

$$K[\mathbf{x}, \lambda\mathbf{y}_1 + (1 - \lambda)\mathbf{y}_2] = \lambda K(\mathbf{x}, \mathbf{y}_1) + (1 - \lambda)K(\mathbf{x}, \mathbf{y}_2),$$

where λ is a real number satisfying $0 \leq \lambda \leq 1$.

Several of these limitations will be relaxed in later chapters.

The property that is essential in Chapters 1–4 is that X and Y are convex sets and have the character of finite dimensionality. (A representation of a game as a triplet involving strategy spaces which are finite-dimensional is necessarily a restriction; its justification rests on the fact that numerous actual games are of this kind.) The identification of strategies with points in Euclidean n-space is a convenience that simplifies the mathematical analysis.

An important special class of games is obtained where X is taken as the simplex S^n in E^n, defined as the set of all $\mathbf{x} = \langle x_1, x_2, \ldots, x_n\rangle$ where $x_i \geq 0$ and

$$\sum_{i=1}^{n} x_i = 1,$$

and the space Y is the corresponding simplex T^m in E^m. The pay-off kernel then takes the form

$$K(\mathbf{x}, \mathbf{y}) = \sum_{j=1}^{m} \sum_{i=1}^{n} x_i a_{ij} y_j = (\mathbf{x}, \mathbf{A}\mathbf{y}),$$

where $\mathbf{A}$ is the matrix $\|a_{ij}\|$. In the case of such matrix games we shall often denote the pay-off corresponding to strategies $\mathbf{x}$ and $\mathbf{y}$ as $A(\mathbf{x}, \mathbf{y})$ in place of $K(\mathbf{x}, \mathbf{y})$ to suggest that these games are matrix games, i.e., games in which $K(\mathbf{x}, \mathbf{y}) = (\mathbf{x}, \mathbf{A}\mathbf{y})$.

Certain special strategies consisting of vertex points of X are denoted by $\boldsymbol{\alpha}_i = \langle 0, \ldots, 0, 1, 0, \ldots, 0\rangle$ $(i = 1, \ldots, n)$, where the 1 occurs in the ith component. These are Player I's *pure strategies*. Similarly, the strategies $\boldsymbol{\beta}_j = \langle 0, \ldots, 0, 1, 0, \ldots, 0\rangle$ of Y $(j = 1, \ldots, m)$ are referred to as Player II's pure strategies. Since $A(\boldsymbol{\alpha}_i, \boldsymbol{\beta}_j) = K(\boldsymbol{\alpha}_i, \boldsymbol{\beta}_j) = a_{ij}$, we see that the i, j element of the matrix array A expresses the yield to Player I when Player I uses the pure strategy $\boldsymbol{\alpha}_i$ and Player II employs the pure strategy $\boldsymbol{\beta}_j$.

A strategy $\mathbf{x} = \langle x_1, x_2, \ldots, x_n \rangle$ with no component equal to 1 is called a *mixed strategy*. In view of the relationship

$$\sum_{i=1}^{n} x_i K(\alpha_i, \mathbf{y}) = K(\mathbf{x}, \mathbf{y}), \tag{1.1.1}$$ *

the mixed strategy $\mathbf{x}$ can be effected as follows. An experiment is conducted with n possible outcomes such that the probability of the ith outcome is x_i. The ith pure strategy is used by Player I if and only if the ith outcome has resulted, and the yield to Player I becomes the expected yield from this experiment, namely

$$\sum_{i=1}^{n} x_i K(\alpha_i, \mathbf{y}).$$

This amounts to playing each pure strategy with a specified probability. We can therefore interpret a mixed strategy as a probability distribution defined on the space of pure strategies, and conversely. Similar interpretations apply to the strategy space T^m of Player II.

The usual introductory approach to game theory is to enumerate the spaces of pure strategies for both players and to specify the pay-off matrix A corresponding to these pure strategies. The concept of mixed strategy is introduced subsequently, and the pay-off function is then replaced in a natural way by the expected pay-off function. In this book we have started with the more general formulation of a game, in terms of a triplet specifying the complete strategy spaces for both players and the pay-off kernel. In this formulation, the distinction between pure and mixed strategies does not exist. Nevertheless, in many special classes of games it is natural to single out the pure strategies, which span convexly the space of all strategies. This will be done wherever it is advantageous, e.g., whenever pure strategies occur in a natural manner and possess special significance. Broadly speaking, however, we lose no generality by starting directly with two-person zero-sum games in normal form, where the strategy spaces are finite-dimensional.

Again, it has been common in the literature of game theory to make a distinction between games in extensive form and games in normal form and to take the first as a point of departure. A game formulated in extensive form is developed in terms of more primitive concepts such as "play," "chance move," "personal move," and "information structure." A strategy

* Although we have postulated this formula, it is possible to derive it by appealing to a suitable axiom system satisfied by a preference pattern which selects among alternative probability distributions over the space of outcomes resulting from the choice of a pure strategy (see the notes to this section).

is then defined within the framework of these notions, and the analysis of optimal strategies proceeds from this point. Finally, a theorem is proved to the effect that any game in extensive form may be in fact reduced to an equivalent game in normal form.

In contrast, our definition of a game as a triplet begins immediately with the concepts of strategy and pay-off, and is flexible and general enough to encompass all forms of finite game theory, including in particular the structure of games in extensive form. By carefully defining strategies and specifying completely X, Y and $K(\mathbf{x}, \mathbf{y})$ we are able to handle all forms of information patterns that arise. This will become clear as we study specific games.

It may happen that in special instances one can construct two apparent strategy spaces which are in fact equivalent in terms of pay-off. Whenever necessary we shall demonstrate this equivalence. For purposes of mathematical consistency, however, whenever any two given games differ in a specific component, the two games are taken to be distinct. For example, if we enlarge the X strategy space while the components Y and K remain unchanged, we create a new game. This is so even if the added strategies are obviously inferior and cannot affect either player's ultimate choice of an optimal strategy. In practice, as will be seen, the strategy spaces exhibited constitute an exhaustive class of procedures, i.e., a class of procedures that take into account all the fine structure of the model.

1.2 Examples. In dealing with finite matrix games it is sufficient to specify the pure strategies for both players and the corresponding pay-off matrix $\|a_{ij}\|$. The pay-off kernel for arbitrary mixed strategies $\mathbf{x}$ and $\mathbf{y}$ is given by the expression

$$K(\mathbf{x}, \mathbf{y}) = \sum_{j=1}^{m} \sum_{i=1}^{n} x_i a_{ij} y_j. \tag{1.2.1}$$

Example 1. Matching pennies. Players I and II each display simultaneously a single penny. If I matches II, i.e., if both are heads or both are tails, I takes II's penny. Otherwise, II takes I's penny. The pay-off kernel is represented in matrix form as follows:

		Player II H	Player II T
Player I	H	1	−1
	T	−1	1

The first pure strategy for I would be to display heads and the second tails. One possible mixed strategy would be to randomize equally between showing heads and tails (i.e., $\mathbf{x} = \langle \frac{1}{2}, \frac{1}{2} \rangle$).

Example 2. Two-finger Morra. Each player displays either one or two fingers and simultaneously guesses how many the opposing player will show. If both players guess correctly or both guess incorrectly, the game is a draw. If only one guesses correctly, he wins an amount equal to the total number of fingers shown by both players.

In this case each pure strategy will have two components: (a) the number of fingers to show, and (b) the number of fingers to guess. Thus, each strategy can be represented by a pair $\langle a, b \rangle$, where a denotes the first component and b the second. For example, the strategy $\langle 2, 1 \rangle$ for Player I is to show two fingers and guess one. There will be four such pure strategies for each player: $\langle 1, 1 \rangle$, $\langle 1, 2 \rangle$, $\langle 2, 1 \rangle$, and $\langle 2, 2 \rangle$. The pay-off matrix is shown below.

		Player II			
		$\langle 1, 1 \rangle$	$\langle 1, 2 \rangle$	$\langle 2, 1 \rangle$	$\langle 2, 2 \rangle$
Player I	$\langle 1, 1 \rangle$	0	2	−3	0
	$\langle 1, 2 \rangle$	−2	0	0	3
	$\langle 2, 1 \rangle$	3	0	0	−4
	$\langle 2, 2 \rangle$	0	−3	4	0

Example 3. Poker model. In this poker model there are only three possible hands, as compared with $\binom{52}{5}$ in regular poker, and all three are considered equally likely to occur. One hand is dealt to each player. There is a preference ordering among the hands: Hand 1 wins over hands 2 and 3, and hand 2 wins over hand 3. The ante is a units. Player I can either pass or bet b units. If Player I bets then Player II can either fold or call. If Player I passes and Player II bets, then Player I again has the option of folding or calling. The three possible courses of play can be diagrammed as follows:

$$
\begin{array}{lclcl}
\text{Player I} & & \text{Player II} & & \text{Player I} \\
\text{pass} & — & \text{pass} & & \\
\text{pass} & — & \text{bet} & — & \left\{\begin{array}{l}\text{fold}\\ \text{call}\end{array}\right\} \\
\text{bet} & — & \left\{\begin{array}{l}\text{fold}\\ \text{call}\end{array}\right\} & &
\end{array}
$$

If a pass follows a bet, the bet wins. If both contestants pass or if one contestant calls, the hands are compared and the player with the better hand wins the pot.