


The background of the book cover is a collage. It features several white dice with black pips. One die is in the top left, another in the top right, and a large one in the bottom right. A blue grid with white numbers is overlaid on the dice. The numbers include 108 , $13/16$, $91 \frac{3}{4}$, $-1 \frac{11}{16}$, $-2 \frac{1}{16}$, and 2.2 .

Decision Making using Game Theory

**An Introduction
for Managers**

Anthony Kelly

CAMBRIDGE

Decision Making using Game Theory

An introduction for managers

Anthony Kelly



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa

<http://www.cambridge.org>

© Cambridge University Press 2003

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2003

Printed in the United Kingdom at the University Press, Cambridge

Typeface Minion 11.5/14pt *System* Poltype® [vN]

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication data

Kelly, Anthony, 1957–

Decision making using game theory: An introduction for managers / Anthony Kelly
p. cm.

Includes bibliographical references and index.

ISBN 0 521 81462 6

1. Game theory I. Title.

HB144 .K45 2002

658.4'0353–dc21 2002019217

ISBN 0 521 81462 6 hardback

Preface

And, greatest dread of all, the dread of games!

John Betjeman 1906–1984 'Summoned by Bells'

Game theory is the science of strategic decision making. It is a powerful tool in understanding the relationships that are made and broken in the course of competition and cooperation. It is not a panacea for the shortcomings of bad management. For managers, or those who interact with management, it is simply an alternative perspective with which to view the process of problem solving. It is a tool, which, like all others, is best used by those who reflect on their own practice as a mechanism for improvement. Chance favours a prepared mind and this book is intended as much for those who are seeking effectiveness as for those who have already found it.

Game theory has been used to great effect in sciences as diverse as evolutionary biology and economics, so books on the subject abound. They vary from the esoteric to the populist; from the pedantic to the frivolous. This book is different in a number of ways. It is designed for both students and practitioners. It is theoretical insofar as it provides an introduction to the science and mathematics of game theory; and practical in that it offers a praxis of that theory to illustrate the resolution of problems common to management in both the commercial and the not-for-profit sectors.

The book is intended to help managers in a number of ways:

- To expand the conceptual framework within which managers operate and in doing so, encourage them to develop more powerful generic problem-solving skills.
- To resolve practical difficulties as and when they occur, more efficiently and with increased effectiveness.

- To find new solutions to familiar problems that have not been satisfactorily resolved, by giving practitioners a deeper understanding of the nature of incentives, conflict, bargaining, decision making and cooperation.
- To offer an alternative perspective on problems, both old and new, which may or may not yield solutions, but which at worst, will lead to an increased understanding of the objective nature of strategic decision making.
- To help managers understand the nature of power in multi-person systems and thereby reduce the perception of disenfranchisement among those who work in committee-like structures within organisations.

The book is a self-contained, though by no means exhaustive, study of game theory. It is primarily intended for those who work as managers, but not exclusively so. Students of politics, economics, management science, psychology and education may find the approach used here more accessible than the usual format of books on the subject. No great mathematical prowess is required beyond a familiarity with elementary calculus and algebra in two variables.

Game theory, by its very nature, offers a rational perspective and, in a society that has developed an aversion to such things, this will be sufficient reason for some to criticise it. This is as unfortunate as it is short-sighted. Research suggests that good managers are well informed, multi-skilled and flexible in their approach to problem solving. Organisations themselves are increasingly complex places, which can no longer afford to live in isolation from the expectations of their employees or the wider community. More than ever, they are workplaces where managers must continuously balance opposing forces. The resulting tensions are ever-changing, and know-how, mathematical or otherwise, is often what separates a failing manager from a successful one.

It has been said, by way of an excuse for curtailing knowledge, that a person with two watches never knows what time it is! Unfortunately, managers cannot afford such blinkered luxury. Game theory has clearly been successful in describing, at least in part, what it is to be a decision maker today and this book is for those who are willing to risk knowing more.

Contents

<i>Preface</i>	ix
----------------	----

1	Introduction	1
	Terminology	3
	Classifying games	6
	A brief history of game theory	8
	Layout	14

2	Games of skill	17
	Linear programming, optimisation and basic calculus	18
	The Lagrange method of partial derivatives	27

3	Games of chance	32
	An introduction to basic probability theory	33
	Games of chance involving risk	37
	Games of chance involving uncertainty	45

4	Sequential decision making and cooperative games of strategy	48
	Representing sequential decision making	49
	Sequential decision making in single-player games	52
	Sequential decision making in two-player and multi-player games	66
	Cooperative two-person games	72
5	Two-person zero-sum games of strategy	77
	Representing zero-sum games	78
	Games with saddle points	80
	Games with no saddle points	86
	Large matrices generally	90
	Interval and ordinal scales for pay-offs	93
6	Two-person mixed-motive games of strategy	98
	Representing mixed-motive games and the Nash equilibrium	99
	Mixed-motive games without single equilibrium points: archetype 1 – leadership games	102
	Mixed-motive games without single equilibrium points: archetype 2 – heroic games	104
	Mixed-motive games without single equilibrium points: archetype 3 – exploitation games	105
	Mixed-motive games without single equilibrium points: archetype 4 – martyrdom games	107
	Summary of features of mixed-motive prototypes	113
	The Cournot, von Stackelberg and Bertrand duopolies: an interesting application of mixed-motive games	115
	Solving games without Nash equilibrium points using mixed strategies	129

7	Repeated games	135
	Infinitely repeated games	135
	Finitely repeated games	139
8	Multi-person games, coalitions and power	149
	Non-cooperative multi-person games	150
	Mixed-motive multi-person games	151
	Partially cooperative multi-person games	153
	Indices of power: measuring influence	155
9	A critique of game theory	174
	Rationality	174
	Indeterminacy	177
	Inconsistency	178
	Conclusion	180
	Appendix A Proof of the minimax theorem	182
	Appendix B Proof of Bayes's theorem	190
	Bibliography	192
	Index	199

1

Introduction

Man is a gaming animal. He must always be trying to get the better in something or other.

Charles Lamb 1775–1834 'Essays of Elia'

Game theory is the theory of independent and interdependent decision making. It is concerned with decision making in organisations where the outcome depends on the decisions of two or more autonomous players, one of which may be nature itself, and where no single decision maker has full control over the outcomes. Obviously, games like chess and bridge fall within the ambit of game theory, but so do many other social situations which are not commonly regarded as games in the everyday sense of the word.

Classical models fail to deal with interdependent decision making because they treat players as inanimate subjects. They are cause and effect models that neglect the fact that people make decisions that are consciously influenced by what others decide. A game theory model, on the other hand, is constructed around the strategic choices available to players, where the preferred outcomes are clearly defined and known.

Consider the following situation. Two cyclists are going in opposite directions along a narrow path. They are due to collide and it is in both their interests to avoid such a collision. Each has three strategies: move to the right; move to the left; or maintain direction. Obviously, the outcome depends on the decisions of both cyclists and their interests coincide exactly. This is a fully *cooperative game* and the players need to signal their intentions to one other.

However, sometimes the interests of players can be completely opposed. Say, for example, that a number of retail outlets are each

ving for business from a common finite catchment area. Each has to decide whether or not to reduce prices, without knowing what the others have decided. Assuming that turnover increases when prices are dropped, various strategic combinations result in gains or losses for some of the retailers, but if one retailer gains customers, another must lose them. So this is a *zero-sum non-cooperative game* and unlike cooperative games, players need to conceal their intentions from each other.

A third category of game represents situations where the interests of players are partly opposed and partly coincident. Say, for example, the teachers' union at a school is threatening not to participate in parents' evenings unless management rescinds the redundancy notice of a long-serving colleague. Management refuses. The union now complicates the game by additionally threatening not to cooperate with preparations for government inspection, if their demands are not met. Management has a choice between conceding and refusing, and whichever option it selects, the union has four choices: to resume both normal work practices; to participate in parents' evenings only; to participate in preparations for the inspection only; or not to resume participation in either. Only one of the possible strategic combinations leads to a satisfactory outcome from the management's point of view – management refusing to meet the union's demands notwithstanding the resumption of normal work – although clearly some outcomes are worse than others. Both players (management and union) prefer some outcomes to others. For example, both would rather see a resumption of participation in parents' evenings – since staff live in the community and enrolment depends on it – than not to resume participation in either. So the players' interests are simultaneously opposed and coincident. This is an example of a *mixed-motive game*.

Game theory aims to find optimal solutions to situations of conflict and cooperation such as those outlined above, under the assumption that players are instrumentally rational and act in their own best interests. In some cases, solutions can be found. In others, although formal attempts at a solution may fail, the analytical synthesis itself can illuminate different facets of the problem. Either way, game theory offers an interesting perspective on the nature of strategic selection in both familiar and unusual circumstances.

The assumption of rationality can be justified on a number of levels.

At its most basic level, it can be argued that players behave rationally by instinct, although experience suggests that this is not always the case, since decision makers frequently adopt simplistic algorithms which lead to sub-optimal solutions.

Secondly, it can be argued that there is a kind of 'natural selection' at work which inclines a group of decisions towards the rational and optimal. In business, for example, organisations that select sub-optimal strategies eventually shut down in the face of competition from optimising organisations. Thus, successive generations of decisions are increasingly rational, though the extent to which this competitive evolution transfers to not-for-profit sectors like education and the public services, is unclear.

Finally, it has been suggested that the assumption of rationality that underpins game theory is not an attempt to describe how players actually make decisions, but merely that they behave *as if* they were not irrational (Friedman, 1953). All theories and models are, by definition, simplifications and should not be dismissed simply because they fail to represent all realistic possibilities. A model should only be discarded if its predictions are false or useless, and game theoretic models are neither. Indeed, as with scientific theories, minor departures from full realism can often lead to a greater understanding of the issues (Romp, 1997).

Terminology

Game theory represents an abstract model of decision making, not the social reality of decision making itself. Therefore, while game theory ensures that a result follows logically from a model, it cannot ensure that the result itself represents reality, except in so far as the model is an accurate one. To describe this model accurately requires practitioners to share a common language which, to the uninitiated, might seem excessively technical. This is unavoidable. Since game theory represents the interface of mathematics and management, it must of necessity adopt a terminology that is familiar to both.

The basic constituents of any game are its participating, autonomous decision makers, called *players*. Players may be individual persons, organisations or, in some cases, nature itself. When nature is desig-

nated as one of the players, it is assumed that it moves without favour and according to the laws of chance. In the terminology of game theory, nature is not 'counted' as one of the players. So, for example, when a deck of cards is shuffled prior to a game of solitaire, nature – the second player – is making the first move in what is a 'one-player' game. This is intrinsically different from chess for example, where nature takes no part initially or subsequently.

A game must have two or more players, one of which may be nature. The total number of players may be large, but must be finite and must be known. Each player must have more than one choice, because a player with only one way of selecting can have no strategy and therefore cannot alter the outcome of a game.

An *outcome* is the result of a complete set of strategic selections by all the players in a game and it is assumed that players have consistent preferences among the possibilities. Furthermore, it is assumed that individuals are capable of arranging these possible outcomes in some order of preference. If a player is indifferent to the difference between two or more outcomes, then those outcomes are assigned equal rank. Based on this order of preference, it is possible to assign numeric pay-offs to all possible outcomes. In some games, an ordinal scale is sufficient, but in others, it is necessary to have interval scales where preferences are set out in proportional terms. For example, a pay-off of six should be three times more desirable than a pay-off of two.

A *pure strategy* for a player is a campaign plan for the entire game, stipulating in advance what the player will do in response to every eventuality. If a player selects a strategy without knowing which strategies were chosen by the other players, then the player's pure strategies are simply equivalent to his or her choices. If, on the other hand, a player's strategy is selected subsequent to those of other players and knowing what they were, then there will be more pure strategies than choices. For example, in the case of the union dispute cited above, management has two choices and two pure strategies: concede or refuse. However, the union's strategic selection is made after management's strategic selection and in full knowledge of it, so their pure strategies are advance statements of what the union will select in response to each of management's selections. Consequently, although the union has only four choices (to resume both practices; to participate in parents' evenings only; to participate in preparations for gov-

Table 1.1 The union's pure strategies

If management chooses to . . .	Then the union will . . .	And if management chooses to . . .	Then the union will . . .
Concede	Resume both practices	Refuse	Resume both practices
Concede	Resume both practices	Refuse	Resume parents' evenings
Concede	Resume both practices	Refuse	Resume inspection preparations
Concede	Resume both practices	Refuse	Resume neither practice
Concede	Resume parents' evenings	Refuse	Resume both practices
Concede	Resume parents' evenings	Refuse	Resume parents' evenings
Concede	Resume parents' evenings	Refuse	Resume inspection preparations
Concede	Resume parents' evenings	Refuse	Resume neither practice
Concede	Resume Ofsted preparations	Refuse	Resume both practices
Concede	Resume Ofsted preparations	Refuse	Resume parents' evenings
Concede	Resume Ofsted preparations	Refuse	Resume inspection preparations
Concede	Resume Ofsted preparations	Refuse	Resume neither practice
Concede	Resume neither practice	Refuse	Resume both practices
Concede	Resume neither practice	Refuse	Resume parents' evenings
Concede	Resume neither practice	Refuse	Resume inspection preparations
Concede	Resume neither practice	Refuse	Resume neither practice

ernment inspection only; not to resume participation in either), they have 16 pure strategies, as set out in Table 1.1 above. Some of them may appear nonsensical, but that does not preclude them from consideration, as many managers have found to their cost!

In a game of *complete information*, players know their own strategies and pay-off functions and those of other players. In addition, each player knows that the other players have complete information. In games of *incomplete information*, players know the rules of the game and their own preferences of course, but not the pay-off functions of the other players.

A game of *perfect information* is one in which players select strategies sequentially and are aware of what other players have already chosen, like chess. A game of *imperfect information* is one in which players have to act in ignorance of one another's moves, merely anticipating what the other player will do.

Classifying games

There are three categories of games: games of *skill*; games of *chance*; and games of *strategy*. Games of skill are one-player games whose defining property is the existence of a single player who has complete control over all the outcomes. Sitting an examination is one example. Games of skill should not really be classified as games at all, since the ingredient of interdependence is missing. Nevertheless, they are discussed in the next chapter because they have many applications in management situations.

Games of chance are one-player games against nature. Unlike games of skill, the player does not control the outcomes completely and strategic selections do not lead inexorably to certain outcomes. The outcomes of a game of chance depend partly on the player's choices and partly on nature, who is a second player. Games of chance are further categorised as either involving risk or involving uncertainty. In the former, the player knows the probability of each of nature's responses and therefore knows the probability of success for each of his or her strategies. In games of chance involving uncertainty, probabilities cannot meaningfully be assigned to any of nature's responses (Colman, 1982), so the player's outcomes are uncertain and the probability of success unknown.

Games of strategy are games involving two or more players, not including nature, each of whom has partial control over the outcomes. In a way, since the players cannot assign probabilities to each other's choices, games of strategy are games involving uncertainty. They can be sub-divided into two-player games and multi-player games. Within each of these two sub-divisions, there are three further sub-categories depending on the way in which the pay-off functions are related to one another – whether the player's interests are completely coincident; completely conflicting; or partly coincident and partly conflicting:

- Games of strategy, whether two-player or multi-player, in which the players' interests coincide, are called *cooperative games of strategy*.
- Games in which the players' interests are conflicting (i.e. strictly competitive games) are known as *zero-sum games of strategy*, so called because the pay-offs always add up to zero for each outcome of a fair game, or to another constant if the game is biased.




Figure 1.1 A taxonomy of games.

- Games in which the interests of players are neither fully conflicting nor fully coincident are called *mixed-motive games of strategy*. Of the three categories, this last one represents most realistically the intricacies of social interaction and interdependent decision making and most game theory is concentrated on it.

A brief history of game theory

Game theory was conceived in the seventeenth century by mathematicians attempting to solve the gambling problems of the idle French nobility, evidenced for example by the correspondence of Pascal and Fermat (c. 1650) concerning the amusement of an aristocrat called de Mere (Colman, 1982; David, 1962). In these early days, largely as a result of its origins in parlour games such as chess, game theory was preoccupied with two-person zero-sum interactions. This rendered it less than useful as an application to fields like economics and politics, and the earliest record of such use is the 1881 work of Francis Edgeworth, rediscovered in 1959 by Martin Shubik.

Game theory in the modern era was ushered in with the publication in 1913, by the German mathematician Ernst Zermelo, of *Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels*, in which he proved that every competitive two-person game possesses a best strategy for both players, provided both players have complete information about each other's intentions and preferences. Zermelo's theorem was quickly followed by others, most notably by the minimax theorem, which states that there exists a strategy for each player in a competitive game, such that none of the players regret their choice of strategy when the game is over. The minimax theorem became the fundamental theorem of game theory, although its genesis predated Zermelo by two centuries. In 1713, an Englishman, James Waldegrave (whose mother was the daughter of James II) proposed a minimax-type solution to a popular two-person card game of the period, though he made no attempt to generalise his findings (Dimand & Dimand, 1992). The discovery did not attract any great attention, save for a mention in correspondence between Pierre de Montmort and Nicholas Bernouilli. It appears not to have unduly distracted Waldegrave either, for by 1721, he had become a career diplomat, serving as British ambassador to the Hapsburg court in Vienna. Nevertheless, by 1865, Waldegrave's solution was deemed significant enough to be included in Isaac Todhunter's *A History of the Mathematical Theory of Probability*, an authoritative, if somewhat dreary, tome. Waldegrave's contribution might have attracted more attention but for that dreariness and his minimax-type solution remained largely unknown at the start of the twentieth century.

In 1921, the eminent French academician Emile Borel began publishing on gaming strategies, building on the work of Zermelo and others. Over the course of the next six years, he published five papers on the subject, including the first modern formulation of a mixed-strategy game. He appears to have been unaware of Waldegrave's earlier work. Borel (1924) attempted, but failed, to prove the minimax theorem. He went so far as to suggest that it could never be proved, but as is so often the case with rash predictions, he was promptly proved wrong! The minimax theorem was proved for the general case in December 1926, by the Hungarian mathematician, John von Neumann. The complicated proof, published in 1928, was subsequently modified by von Neumann himself (1937), Jean Ville (1938), Hermann Weyl (1950) and others. Its predictions were later verified by experiment to be accurate to within one per cent and it remains a keystone in game theoretic constructions (O'Neill, 1987).

Borel claimed priority over von Neumann for the discovery of game theory. His claim was rejected, but not without some disagreement. Even as late as 1953, Maurice Frechet and von Neumann were engaged in a dispute on the relative importance of Borel's early contributions to the new science. Frechet maintained that due credit had not been paid to his colleague, while von Neumann maintained, somewhat testily, that until his minimax proof, what little had been done was of little significance anyway.

The verdict of history is probably that they did not give each other much credit. Von Neumann, tongue firmly in cheek, wrote that he considered it an honour 'to have labored on ground over which Borel had passed' (Frechet, 1953), but the natural competition that can sometimes exist between intellectuals of this stature, allied to some local Franco-German rivalry, seems to have got the better of common sense.

In addition to his prodigious academic achievements, Borel had a long and prominent career outside mathematics, winning the Croix de Guerre in the First World War, the Resistance Medal in the Second World War and serving his country as a member of parliament, Minister for the Navy and president of the prestigious Institut de France. He died in 1956.

Von Neumann found greatness too, but by a different route. He was thirty years younger than Borel, born in 1903 to a wealthy Jewish banking family in Hungary. Like Borel, he was a child prodigy. He