CHAPTER 1

Introduction

NATURE OF THE SUBJECT

It is all too clear at this moment that there are many ways for a book to begin; and most of those in plain sight are transparently bad. We are tantalized by the thought that somewhere among them *may* lie hidden a few having such noble qualities as these: The readers are informed—perhaps without suspecting it, though in the clearest prose—of what the writer intends to discuss; yet at the same time, it sounds like the Lorelei calling. Whereupon these readers resolve into two groups: The first, a large and happy family really, will stick to the book to the end, even though unimagined adversities impend. Further, this group will always think and speak kindly of it, and will doubtless have at least one copy in every room. The second group is most briefly described by stating that it differs from the first; but the book acts immediately as a soporific on all unpleasant passions, so, as it is sleepily laid aside, the sole lasting impression is that of a good gift suggestion.

If we could devise an opening strategy such as that, it would wonderfully exemplify the theme and aims of the book, for *our concern throughout will be with a method for selecting best strategies,* even in contexts where the word 'strategy' itself may not be in common use.

The contexts of interest to us are those in which people are at cross-purposes: in short, conflict situations. The problem of how to begin this book is recognizably of that type, for certainly you and the writer are at cross-purposes, as our interests are opposed—in a polite way, of course, but definitely opposed. For we hope to cozen you into a very difficult type of intellectual activity, while you, a reasonable person with enough troubles already, may crave only relaxation or satisfaction of curiosity. This conflict of interests is essential in the situations we shall study.

Another element is also essential and it is present here too: Each of us can exert *some* control over the situation. Many ways will occur to you: for one, you may throw the book at the cat, thus irritating both the writer and the cat,

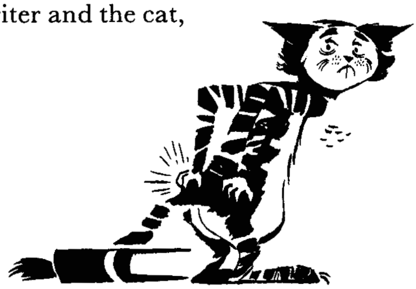

but at some cost in property, perhaps some in self-respect, and undoubtedly some in deteriorated relations with the cat. Or you may skim the hard parts, and so on. There are aspects within the control of the writer, too, such as the choice and treatment of content—but it is not necessary to labor the point. And a further characteristic element appears: Some aspects of the situation are not within the control of either of us; for example, a multitude of events in our pasts and extraneous influences during the writing and reading periods will play important roles. Of course this particular problem, of beginning the book in a really optimum way, has a further characteristic which we shall henceforth shun, namely, it is too hard—else we should have solved it.

The restrictions on the subject matter being so few and mild, it follows that the set of conflict situations we are willing to consider is most notable for its catholicity. There is no objection, in principle, to considering an H-bomb contest between Mars and Earth, or a love affair of the Barrett-Browning type. The contest may be economic in character, or it may be Musical Chairs. Or it may be almost any one of the myriad activities which take place during conventional war. It doesn't follow that we have a nostrum for strategic ills in all these fields, but there is a possibility that our offering may as a method, perform useful service in any of them.

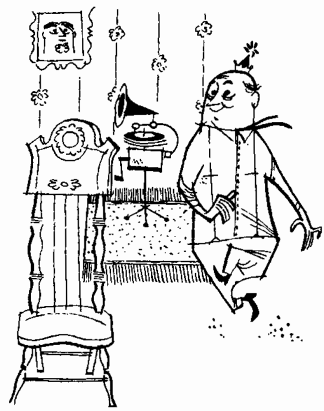

The method which will be presented is identified by the catch phrase *Game Theory* or, time permitting, the *Theory of Games of Strategy.* If this is your first encounter with that unlikely sequence of nouns, the sole reaction is probably: Why? Well, the idea takes its name from the circumstance that the study of games is a useful and usable starting point in the study of strategy. That does not really help, for again we hear: Why? Well, because games contain many of the ingredients common to all conflicts, and they are relatively amenable to description and to study. (Incidentally, having used the word 'game' to name the theory, we then call any conflict a game when we are considering it in the light of the theory.)

To illustrate the point, let us run our minds over a Poker game, keeping watch for items which are significant in, say, a military conflict. You and four others are thus studying human nature, under a system of rewards, you hope. We note at once that the players have opposing interests; each wants to win and, because the winnings of one are necessarily the losses of another, their interests are opposed. This provides the basis of conflict. We observe too that some elements of the action, being personal choices, are completely within your control. And the same being true for each player, there are elements which are not within your control; worse, they are controlled by minds having interests inimical to yours. Finally, there are elements of the game that are not, under the rules, within the control of any player, such as the order of the cards in the deck. These elements may be thought of as being controlled by Nature—who has a massively stable personality, a somewhat puckish attitude toward your important affairs, but who bears you no conscious malice. These are all surely familiar aspects of any conflict situation.

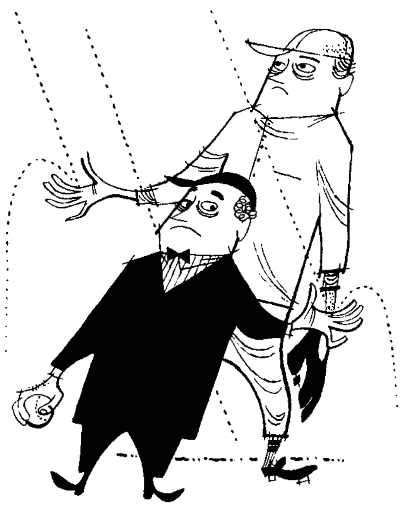

Another characteristic is that the state of information—intelligence, in the military sense—is a factor, and, as usual, is an imperfect and hence troublesome factor: We don't know what the other fellow's hole card is. There is also the bluff by which you, or the opposition, give false evidence regarding intentions or strength of forces. Other similarities will occur to you; people even get killed, occasionally.

But the analogy should not be pushed too far. You can think of many aspects of warfare which are not reflected in Poker. One tank will sometimes kill two tanks, in a showdown; whereas a pair of Jacks always wins over an Ace-high hand in the showdown. Of course Poker could be modified to make it contain showdown possibilities of this kind, say by ruling that an Ace is superior to any pair, up to Jacks, whenever anybody's wife phones during the play of a hand. But the fact is that games don't exhibit all the complexities of warfare and of other real-life conflict situations—which is precisely why they are usable starting points for a study of strategy. In the early stages of developing a theory it just is not possible simultaneously to handle very many interacting factors.

It is probably clear, then, that games do contain some of the basic elements that are present in almost any interesting conflict situation. Does it follow that we can learn useful things by beginning a study with them? Not necessarily. It may be that military, economic, and social situations are just basically too complicated to be approached through game concepts. This possibility gains credence from the fact that the body of Game doctrine now in existence is not even able to cope with full-blown real games; rather, we are restricted at present to very simple real games, and to watered-down versions of complicated ones, such as Poker.

It may be baffling then that someone devotes valuable energy to the study and development of Game Theory—and, moreover, expects you to participate! The reason it is done is in part an act of hope and of faith, stemming from past successes. For the invention of deliberately oversimplified theories is one of the major techniques of science, particularly of the 'exact' sciences, which make extensive use of mathematical analysis. If the biophysicist can usefully employ simplified models of the cell and the cosmologist simplified models of the universe, then we can reasonably expect that simplified games may prove to be useful models for more complicated conflicts.

Of course the mortality among such theories is higher than any military organization would tolerate in *its* activities, and those that are successful are not really immortal; the best that can be expected of one is that it be adequate for certain limited purposes, and for its day.

AN HISTORICAL THEORY

It may be useful to examine one successful scientific abstraction, to see what it is like and for the sake of the hints it may give us. We choose one which is surely an example of heroic oversimplification.

Let us assume that we may, in order to study their motions, replace each of the major bodies of the Solar System by a point; that each point has a mass equal to that of the body it replaces; that each pair of points experiences a mutual attraction; that we may estimate the attractive force by multiplying the mass of one point by the mass of the other, after which we divide that product by the square of the distance between the points; that we may neglect all else; and that it isn't patently stupid to consider this theory, else we would never get started.

The fact is that this theory, the Theory of Gravitation, has been adequate for predicting the motions of the planets for two and one-half centuries—and this in the face of constant checking by positional astronomers, who, it can fairly be said, carry precision to extremes. The worst strain has come from the orbit of Mercury, which unaccountably drifted from the predicted place by one-fifth of a mil (a foot, at a distance of a mile) *per century,* thus showing that the theory is rough after all, just as it looks. The improved theory, by Einstein, accounts for this discordance.

LESSONS AND PARALLELS

The elements of the theory stated above of course did not just float into a mind dazed by a blow from an apple. There was much information at hand regarding the actual behavior of the planets, thanks largely to Tycho Brahe, and a wearisome mess it was. Kepler finally grubbed out of it a few rules of thumb; with these, and with a lift from a new mathematical invention (the Calculus), Newton soon afterward hit upon the above abstraction. He had the misfortune to try it immediately on the Moon, which cost him years of happiness with his theory, for the data were seriously in error.

This example contains several lessons for us. One is that theories may be very simple, while the phenomena they model do not appear simple. Anybody who supposes that planetary motions are quite simple has never had the responsibility for predicting them; the ancients had good reason to name them the Wanderers. Another lesson is that a theory can be very general, being applicable to a wide variety of phenomena, without being sterile; the Theory of Gravitation is even more general than stated above, for it applies to *all* mass particles, not just to the major bodies of the Solar System. Another lesson is that theories often or usually are imperfect, though the one used as an example is embarrassingly good. Another—and this is a very important one—is that the theory covers only one of the interesting factors which may affect the motion of bodies; one, moreover, that is frequently negligible. For example, the gravitational attraction between two airplanes flying a tight formation is equivalent to the weight of a cigarette ash, perhaps a sixteenth of an inch long.

Still another lesson concerns the importance of having some relevant data. In this respect Newton was somewhat better off than we are—we who are trying to do abstraction in such a field as conflict. For most of the data we have on man relate to the individual—his physical and mental composition, health, ability, etc.—and, to a lesser extent, to the gross characteristics of the social group. The *interactions* between men, as individuals in a group or between groups, have not been studied on anything approaching the scale needed; and these interactions are the stuff which constitutes conflict.

Another lesson, or at least a suggestive note, is the fact that Newton almost simultaneously developed the Theory of Gravitation and a new branch of mathematics—the Calculus; and the theory would have been practically unusable without it. In fact the Calculus has played a dominant role in all physical science for a quarter of a millennium. It is provocative to speculate on whether Game Theory will develop a new mathematical discipline destined for a comparable role in analyzing the interactions of men. It is much too early to conjecture that it will; so far, there has been little that is recognizable as brand new, and much that is recognizable as borrowings from established branches. But it may happen, and perhaps even it must happen if the application of the method is to reach full flower. It is at least interesting that the original development of Game Theory is the work of one of the really great mathematicians and versatile minds of our day—John von Neumann.*

Game Theory is very similar in spirit to the Theory of Gravitation. Both attempt to treat broad classes of events according to abstract models. Neither tries to model all the complexities present in any situation. One of them, to the extent it is applicable to animal activity, concerns itself with some of the involuntary actions; thus the Theory of Gravitation can answer superbly all questions regarding the gross motions of a pilot, alone at 40,000 feet, who is unencumbered by aircraft, parachute, or other device. Game Theory, on the other hand, would be more interested in the strategy by which he achieved all this and with questions regarding its optimality among alternative strategies; it, therefore, enters the region of decisions and free will.

This comparison with Gravitation Theory will be unfortunate if it seems to imply comparable utility and (in a loose sense) validity—not to say social standing—for the two theories. The one is mature and comfortably established as a useful approximation to Nature, whereas the other is a lusty infant, which may be taken by a plague or which may grow up to great importance, but which is now capable only of scattered contributions. As an infant, it is proper for it to be a little noisy.

Having permited you to sense the galling bit of mathematics that will come (i.e., 'bit' as in the horse), we hasten to assure you that the approach we shall use is that of the primer, strictly, which means (you will recall) an elementary book for practice in spelling, and the like. We assume explicitly that you are not trained in mathematics beyond rudimentary arithmetic. In fact, if this is not true, simple charity requires that you close the book.

SECTARIAN REMARKS ON METHOD

It is sometimes felt that when phenomena include men, it is tremendously more difficult to theorize successfully; and our relative backwardness in these matters seems to confirm this. However, part

* Von Neumann's first paper on Game Theory was published in 1928, but the first extensive account appeared in 1944: *Theory of Games and Economic Behavior* by John von Neumann and Oskar Morgenstern (Princeton University Press, Princeton, N. J.). The challenging nature of this work was immediately appreciated by some reviewers, such as A. H. Copeland, who wrote "Posterity may regard this book as one of the major scientific achievements of the first half of the twentieth century" (*Bulletin of the American Mathematical Society*, vol. 51, 1945, pp. 498–504).

of the so-far minor effort made in this direction has been dissipated against hand-wringing protestations that it is too hard to do. Some of the impetus toward simple theory—simple theory being a few axioms and a few rules for operating on them, the whole being more or less quantitative—has come from amateurs; physical scientists, usually. These are often viewed by the professional students of man as precocious children who, not appreciating the true complexity of man and his works, wander in in wide-eyed innocence, expecting that their toy weapons will slay live dragons just as well as they did inanimate ones. Since Game Theorists are obviously children of this ilk, you doubtless anticipate that we shall now make some reassuring sounds, probably at the expense of the professionals, else we should not have raised the subject. If you do so anticipate, this shows how easy it really is, for it establishes you as a promising student of man, too!

The motive force that propels the Game Theorist isn't *necessarily* his ignorance of the true complexity of man-involved conflict situations; for he would almost surely try to theorize if he were not so ignorant. We believe, rather, that his confidence—better, his temerity—stems from the knowledge that he and his methods were completely outclassed by the problems of the inanimate world. He could not begin to comprehend them when he looked at them microscopically and, simultaneously, with a wide field of view; the quantity of detail and the complexity of its organization were overpowering. So, since he has had some success in that field, he suspects that sheer quantity and complexity cannot completely vitiate his techniques.

He is also aware that his successes occur spottily, so that his knowledge is much less complete than the uninitiated suspect—the uninitiated including of course those who believe the animate field must be vastly harder than the inanimate *because* the latter has been done so well (!). For example, modern physicists have only the foggiest notions about some atomic constituents—though they designed successful A-bombs. Their favorite particle, the electron, is shrouded in ignorance; such elementary information as where-is-it and, simultaneously, where-is-it-going is not known—worse, they have decided this information is in a strict sense forever unknowable. The mathematicians are likewise a puny breed. Item: after centuries of effort, they still don't know the minimum number of colors needed to paint a map (so that adjacent countries will not have the same color); it's fair to add that they suspect the number is four, but they haven't proved it.

Within the last hundred years, the physical scientists have added a

new arrow to their quiver, one which has played only the role of minor weapon in most of their campaigns so far, namely, mathematical statistics. They are now beginning to find more important uses for it, and there is a good prospect that it will become an increasingly important tool in the animate field; Game Theory has many points of contact with it. An early demonstration of its power, and a harbinger of its range of utility, was its success in accounting for the distribution (over the years) of deaths in Prussian Army Corps due to kicks from horses.* If you protest that horses are more predictable than men, we counter confidently with the assertion that the method is just as applicable to the distribution of horses kicked to death by Prussians. Of course the

* The reference is to these data, covering ten army corps over a twenty-year period (1875–1894). The deaths are per corps, per year.

Deaths	Occurrences Observed	Occurrences Computed
0	109	109
1	65	66
2	22	20
3	3	4
4	1	1

The computed values are derived from one bit of observed information, namely, that the fatalities average about one every twenty months, and from a statistical theory that is particularly applicable to rare events.

whole field of insurance is an example of statistical theory applied to some aspects of human affairs; the balance sheets of the insurance companies bear eloquent testimony to its success. Humans are not completely unpredictable.

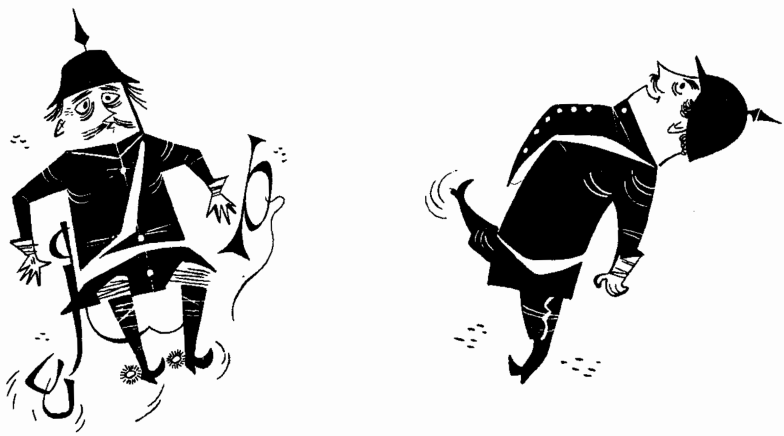

So what are reasonable expectations for us to hold regarding Game Theory? It is certainly much too simple a theory to blanket all aspects of interest in any military, economic, or social situation. On the other hand, it is sufficiently general to justify the expectation that it will illumine certain critical aspects of many interesting conflict situations.

There are at present some important things to be done. One is to develop further the theory itself, so that more difficult and more varied problems can be solved; this task falls to the scientists. Another is to find situations to which existing theory can profitably be applied; one purpose of this book is to increase the number of persons who, by knowing the rudiments of the theory, can suggest applications to problems selected from those they encounter. (Those who hang on far enough will be able to formulate and solve simple problems for themselves.) Another task is the collection of data in the field of human interaction, to improve the bases of abstraction.

PLAYERS AND PERSONS

Now to Game Theory itself. We shall begin by looking at Stud Poker, and we shall look just long enough to introduce some concepts and terminology that will be used throughout the book. You and four others are still sitting there, with a deck of cards, some money or other valuables, and an agreed-on set of rules that covers all contingencies. The rules govern how the cards are to be doled out, who may bet and when, how the various hands are to be judged in the showdown, and what happens to the pot.

One of the obvious things about this situation is that it is a five-person game. But this may be more obvious than true; for perhaps two of the players formed a coalition, in advance of the game, in which they agreed to pool their winnings or losses. If they did so, it is reasonable to suppose that they will play for their common good whenever circumstances permit it. Thus if one member of the coalition believes his partner has a good chance of winning a particular hand, he should take whatever action he can toward the common good. If only three hands are active, perhaps he should fold so that the burden of calling falls on the outsider, or perhaps he should raise the bet in order to increase the pot, even though he knows his cards cannot win the hand. In short, the members of the coalition will behave as much like a single individual, with two heads, as they can.

In the case where two players have formed a coalition, it is evident that it may be fruitful to consider it as a four-person, rather than as a five-person, game. Thus we come to believe *it is significant to count the number of sets of opposing interests* around the table, rather than the bodies. According to this principle, Bridge is classed as a two-person game, because there are only two sets of interests when the partners are permanent. You will note that the words 'person' and 'player,' as we use them, cover legal persons and organizations, as well as natural persons.

Again, you may prefer to regard the Poker game as a two-person game in which you are one of the players and the other four individuals are the other player. If they do not look at it the way you do, they will gain no advantage from the association you have imagined for them, and you will suffer no loss from it; it is as though they constitute a coalition with weak internal communications, or some other malady which makes it ineffective.

This is one of the fundamental distinctions in Game Theory, namely, the number of persons—distinct sets of interests—that are present in the game. The form of analysis and the entire character of the situation depend on this number. There are three values, for the number of persons, which have special significance: one, two, and more-than-two.

Solitaire is an example of a one-person game when played for recreation, for your interests are the only ones present. Even if you buy the deck for, say, $1 a card from somebody who is willing to pay you, perhaps, $5 a card for all cards transferred to the payoff piles, the case is the same: only chance events must be countered, and not the moves of a responsive human adversary. One-person games are uninteresting, from the Game Theory point of view, and therefore are not really studied here. Their solution is quite straightforward, conceptually: You simply select the course of action that yields the most and do it. If there are chance elements, you select the action which yields the most on the average, and do it. You may complain that we are glossing over an awful lot of practical difficulties; and that's right.

However, one-person games (including Solitaire) may be regarded as a special kind of two-person game in which you are one of the players and Nature is the other. This may be a useful viewpoint even if you don't believe that Nature is a malignant Being who seeks to undo you. For example, you may not know enough about Nature's habits to select the course which will yield the most on the average. Or it may happen that you know the kinds of behavior open to Nature, but know little about the frequency with which She uses them. In this case Game Theory does have something to say; it will lead you to conservative play, as we shall see later.

The true two-person game is very interesting. It occurs frequently and its solution is often within our present means, both conceptual and technological. This is the common conflict situation. You have an opponent who, you must assume, is intelligent and trying to undo you. If you choose a course of action which appears favorable, he may discover your plans and set a trap which capitalizes on the particular choice you have made. Many situations which are not strictly two-person games may be treated as if they were; the five-man Poker game was an example of this, where you could assign the interests present at the table to two 'persons,' yourself and everybody-not-you. Most of the work done to date in Game Theory deals with the two-person game.

When the number of distinct persons, i.e., sets of interests, exceeds two, qualitatively new things enter. The principal new factor is that the identities of the persons may change in the course of the game, as temporary coalitions are formed and broken; or certain players may form what is in effect a permanent partial coalition in some area of action where they conceive it to be beneficial. This could happen in the Poker game and would compromise our treatment of it as a two-person game, as proposed earlier. For example, you might wish to team up with others, informally but effectively, to act against a heavy winner; you might be motivated by fear that he would leave the game taking most of the cash with him, or you might prefer to see more of it in the hands of a weaker player. Our understanding of games that involve more than two persons is less complete at present than for two-person games, and the subject is rather complicated; in fact, it lies beyond the modest limits of this book.

THE PAYOFF

We have indicated that the number of persons involved is one of the important criteria for classifying and studying games, 'person' meaning a distinct set of interests. Another criterion has to do with the payoff: What happens at the end of the game? Say at the end of the hand in Poker? Well, in Poker there is usually just an exchange of assets. If there are two persons, say you (Blue) and we (Red), then if you should win \$10, we would lose \$10. In other words,

$$\text{Blue winnings} = \text{Red losses}$$

or, stated otherwise,

$$\text{Blue winnings} - \text{Red losses} = 0$$

We may also write it as

$$\text{Blue payoff} + \text{Red payoff} = \$10 - \$10 = 0$$

by adopting the convention that winnings are positive numbers and that losses are negative numbers.

It needn't have turned out just that way; i.e., that the sum of the payoffs is zero. For instance, if the person who wins the pot has to contribute 10 per cent toward the drinks and other incidentals, as to the cop on the corner, then the sum of the payoffs is not zero; in fact

$$\text{Blue payoff} + \text{Red payoff} = \$9 - \$10 = -\$1$$

The above two cases illustrate a fundamental distinction among games: It is important to know whether or not the sum of the payoffs, counting winnings as positive and losses as negative, to all players is zero. If it is, the game is known as a *zero-sum game.* If it is not, the game is known (mathematicians are not very imaginative at times) as a *non-zero-sum game.* The importance of the distinction is easy to see: In the zero-sum case, we are dealing with a good, clean, closed system; the two players and the valuables are locked in the room. It will require a certain effort to specify and to analyze such a game. On the other hand, the non-zero-sum game contains all the difficulties of the zero-sum game, plus additional troubles due to the need to incorporate new factors. This can be appreciated by noting that we can restore the situation by adding a fictitious player—Nature again, say, or the cop. Then we have

$$\begin{aligned} \text{Blue payoff} &= \$9 \\ \text{Red payoff} &= -\$10 \\ \text{Cop payoff} &= \$1 \end{aligned}$$

so now

$$\text{Blue payoff} + \text{Red payoff} + \text{Cop payoff} = \$9 - \$10 + \$1 = 0$$

which is a *three-person zero-sum* game, of sorts, where the third player has some of the characteristics of a millstone around the neck. But recall that we don't like three-person games so well as we do two-person games, because they contain the vagaries of coalitions. So non-zero-sum games offer real difficulties not present in zero-sum games, particularly if the latter are two-person games.

Parlor games, such as Poker, Bridge, and Chess, are usually zero-sum games, and many other conflict situations may be treated as if they were. Most of the development of Game Theory to date has been on this type of game. Some work on non-zero-sum games has been done, and more is in progress, but the subject is beyond our scope. A troublesome case of particular interest is the two-person game in which the nominally equal payoffs differ in utility to the players; this situation occurs often even in parlor games.

STRATEGIES

Just as the word 'person' has a meaning in Game Theory somewhat different from everyday usage, the word 'strategy' does too. This word, as used in its everyday sense, carries the connotation of a particularly skillful or adroit plan, whereas in Game Theory it designates any *complete* plan. *A strategy is a plan so complete that it cannot be upset by enemy action or Nature;* for everything that the enemy or Nature may choose to do, together with a set of possible actions for yourself, is just part of the description of the strategy.

So the strategy of Game Theory differs in two important respects from the conventional meaning: It must be utterly complete, and it may be utterly bad; for nothing is required of it except completeness. Thus, in Poker, all strategies must make provision for your being dealt a Royal Flush in Spades, and some of them will require that you fold instantly. The latter are not very glamorous strategies, but they are still strategies—after all, a Bridge player once bid 7 No-Trump while holding 13 Spades. In a game which is completely amenable to analysis, we are able—conceptually, if not actually—to foresee all eventualities and hence are able to catalogue all possible strategies.

We are now able to mention still another criterion according to which games may be classified for study, namely, the number of strategies available to each player. Thus, if Blue and Red are the players, Blue may have three strategies and Red may have five; this would be called a 3×5 game (read 'three-by-five game').

When the number of players was discussed, you will recall that certain numbers—namely, one, two, and more-than-two—were especially significant. Similarly, there are critical values in the number of strategies; and it turns out to be important to distinguish two major categories. In the first are games in which the player having the *greatest* number of strategies still has a finite number; this means that he can count them, and finish the task within some time limit. The second major category is that in which at least one player has infinitely many strategies, or, if the word 'infinitely' disturbs you, in which at least one player has a number of strategies which is larger than any definite number you can name. (This, incidentally, is just precisely what 'infinitely large' means to a mathematician.)

While infinite games (as the latter are called) cover many interesting and useful applications, the theory of such games is difficult. 'Difficult' here means that there are at least some problems the mathematician doesn't know how to solve, and further that we don't know how to present any of it within the friendly pedagogical limits of this book; such games require mathematics at the level of the Calculus and beyond—mostly beyond. Therefore we here resolve to confine our attention to finite games.

We shall find it convenient, in later chapters, to distinguish three cases among finite games: namely, those in which the player having the *least* number of strategies has exactly two, exactly three, or more-than-three. In addition, considerations of labor, fatigue, and the better life will cause us to develop a rather special attitude toward games having more than about ten strategies.

THE GAME MATRIX

We are now in a position to complete the description of games, i.e., conflict situations, in the form required for Game Theory analysis. We will freely invoke all the restrictions developed so far, so as to aim the description directly at the class of games which will be studied throughout the book. Hence our remarks will primarily apply to finite, zero-sum, two-person games.

The players are Blue and Red. Each has several potential strategies which we assume are known; let them be numbered just for identification. Blue's strategies will then bear names, such as Blue 1, Blue 2, and so on; perhaps, in a specific case, up to Blue 9; and Red's might range from Red 1 through Red 5. We would call this a nine-by-five game and write it as '9 × 5 game.' Just to demonstrate that it is possible to have a 9 × 5 game, we shall state one (or enough of it to make the point). Consider a game played on this road map:

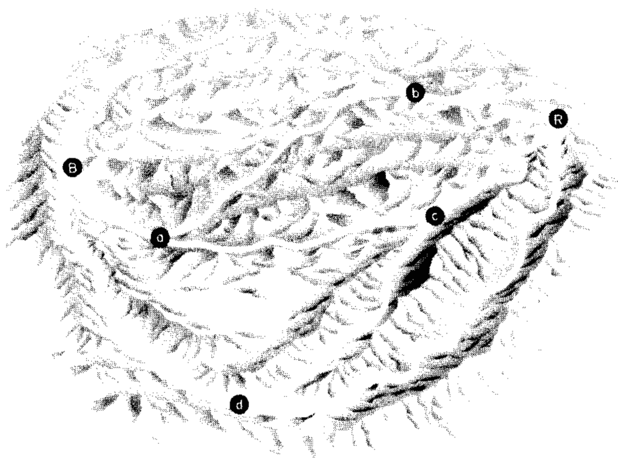

The rules require that Blue travel from *B* to *R*, along the above system of roads, without returning to *B* or using the same segment twice during the trip. The rules are different for Red, who must travel from *R* to *B*, always moving toward the west. Perhaps Blue doesn't want to meet Red, and has fewer inhibitions about behavior. You may verify that there are nine routes for Blue and five for Red.*

* To avoid even the possibility of frustrating you this early in the game, we itemize the routes. Blue may visit any of the following sets of road junctions (beginning with *B* and ending with *R* in each case):

b, bac, bacd, ab, ac, acd, dcab, dc, d

Red may visit

b, ba, ca, cd, d

The rules must also contain information from which we can determine what happens at the end of any play of the game: What is the payoff when, say, Blue uses the strategy Blue 7 (the northern route, perhaps) and Red uses Red 3 (the southern route, perhaps)? There will be $9 \times 5 = 45$ of these pairs and hence that number of possible values for the payoff; and these must be known. Whatever the values are, it is surely possible to arrange the information on this kind of bookkeeping form:

		Red				
		1	2	3	4	5
	1					
	2					
	3					
	4					
Blue	5					
	6					
	7					
	8					
	9					

Such an array of boxes, each containing a payoff number, is called a *game matrix*. We shall adopt the convention that a positive number in the matrix represents a gain for Blue and hence a loss for Red, and vice versa. Thus if two of the values in the game matrix are 3 and −8, as shown here,

		Red				
		1	2	3	4	5
	1					
	2		−8			
	3					
	4					
Blue	5					
	6				3	
	7					
	8					
	9					

the meaning is: When Blue uses Blue 6 and Red uses Red 4, Blue wins 3 units, whereas when Blue 2 is used vs. Red 2, Red wins 8 units.

When the original problem has been brought to this form, a Game Theory analysis may begin, for all the relevant information is represented in the descriptions of the strategies whose signatures border the matrix and in the payoff boxes. This is the Game Theory model of the conflict, and the applicability of the subsequent analysis will depend completely on the adequacy of this form of representation—a set of strategies and a payoff matrix.

IMPLICIT ASSUMPTIONS

Perhaps the last statement should be expanded. We narrow our attention for a moment to two complicated objects: One is the real conflict situation in which Blue and Red are involved. This includes the rules, regulations, taboos, or whatnots that are *really* operative; it includes the true motives of the players, the geography, and in fact everything that is significant to the actual game. The second object is also real, but it is much more simple: It is the rules we have *written,* the strategies we have enumerated and described *on paper,* and the game matrix we have *written.* There is a relationship—a significant one, we trust—between these two objects. The second object—the marks on paper—is an abstraction from the first. We can discover some non-obvious properties of this second object by making a Game Theory analysis, and these properties *may* have some validity in connection with the first object—the real world game. It will all depend on the adequacy of the abstraction.

The principal topic of this book will be discussion of how Game Theory operates on the second object, the abstract model. The difficulties and questions that will come up in that discussion will be, principally, technical ones, rather than conceptual ones. They will be questions of ingenuity in handling difficult mathematical problems or of devices to avoid outrageous labor; in general, just high-class crank turning. We should recognize, before passing to this relatively comfortable pastime, that all is now easy because we have already glided over many of the real difficulties, namely, the conceptual ones.

One of the conceptual problems, a critical point in Game Theory so far as its application to real-life conflict situations is concerned, is reached when we try to fill in the boxes with the values of the payoff. While there will be individual cases in which the requirements are less severe, in general we have to assume that the payoff can, in principle, be measured numerically; that we in fact know how to measure it; and that we do measure it, with sufficient accuracy. Further, the units of measurement must be the same in all boxes, and the units must be simple, dimensionally; that is to say, we are not prepared to cope with dollars in one box, grams of uranium in another, and lives in another—unless of course we happen to know exchange ratios between these items and can eliminate the heterogeneity of units of measurement. If the payoff in each box contains several numbers representing disparate items—which may be some dollars, some uranium, and some lives, say—we are still in trouble, except in special cases. This difficulty can arise in ordinary games; consider, for example, a two-person game between master players; the stakes may be sums of money and prestige. Unless we are prepared to adopt an exchange ratio between prestige and money, our analysis is likely to be in trouble.

Another conceptual difficulty in connection with real problems is that of defining the problem sufficiently crisply, so that the action alternatives available to the players may be completely itemized; and to do this without isolating the problem from the important influences of its original environment.

Other hazards will be pointed out from time to time in later chapters, as the discussion veers by an occasional rock. There will be some things to note on the other side of the question, too; for the model need not be an exact replica of the real-life situation in order to be useful. We shall see that there is sometimes considerable latitude in these matters.

THE CRITERION

A perennial difficulty in modelmaking of the analytical (as opposed to wooden) variety is the illness which might well be known as criterion-trouble. What is the criterion in terms of which the outcome of the game is judged? Or should be judged?

To illustrate the wealth of possible criteria in a homely example, consider a housewife who has $5 to spend on meat. What should she buy? If her criterion is simply quantity, she should buy the cheapest kind and measure the payoff in pounds. If it is variety, she should buy minimum, useful quantities of several kinds, beginning with the cheap-

est kinds; she measures the payoff by the number of kinds she buys. Or she may be interested in protein, fat, or calories. She may have to satisfy various side conditions, or work within certain constraints, such as allergies, tastes, or taboos. She may be interested in least total effort, in which case she may say, "I want five dollars worth of cooked meat—the nearest, of course—and deliver it sometime when you happen to be down our way."

Generally speaking, criterion-trouble is the problem of what to measure and how to base behavior on the measurements. Game Theory has nothing to say on the first topic, but it advocates a very explicit and definite behavior-pattern based on the measurements.

It takes the position that there is a definite way that rational people should behave, if they believe in the game matrix. The notion that there is some way people ought to behave does not refer to an obligation based on law or ethics. Rather it refers to a kind of mathematical morality, or at least frugality, which claims that the *sensible object of the player is to gain as much from the game as he can, safely, in the face of a skillful opponent who is pursuing an antithetical goal.* This is our model of rational behavior. As with all models, the shoe has to be tried on each time an application comes along to see whether the fit is tolerable; but it is well known in the Military Establishment, for instance, that a lot of ground can be covered in shoes that do not fit perfectly.

Let us follow up the consequences of this model in a zero-sum game, which, you will recall, is a closed system in which assets are merely passed back and forth between the players. It won't affect anything adversely (except Red), and it will simplify the discussion, if we assume for a moment that all payoffs in the game matrix are *positive;* this means that the strategy options available to the players only affect how many valuables Red must give to Blue at the end of a play of the game; this isn't a fair game for Red, but we will let him suffer for the common weal.

Now the viewpoint in Game Theory is that *Blue wishes to act in such a manner that the least number he can win is as great as possible, irrespective of what Red does;* this takes care of the safety requirement. *Red's comparable desire is to make the greatest number of valuables that he must relinquish as small as possible, irrespective of Blue's action.* This philosophy, if held by the players, is sufficient to specify their choices of strategy. If Blue departs from it, he does so at the risk of getting less than he might have received; and if Red departs from it, he may have to pay more than he could have settled for.

The above argument is the central one in Game Theory. There is a way to play every two-person game that will satisfy this criterion. However, as in the case of the housewife buying meat, it is not the only possible criterion; for example, by attributing to the enemy various degrees of ignorance or stupidity, one could devise many others. Since Game Theory does not attribute these attractive qualities to the enemy, it is a conservative theory.

You will note an apparent disparity in the aims of Blue and Red as stated above; Blue's aims are expressed in terms of winning and Red's in terms of losing. This difference is not a real one, as both have precisely the same philosophy. Rather, it is a consequence of our convention regarding the meaning of positive and negative numbers in the game matrix. The adoption of a uniform convention, to the effect that Blue is always the maximizing player and Red the minimizing player, will reduce technical confusion (once it becomes fixed in your mind); but let's not pay for this mnemonic by coming to believe that there is an essential lack of symmetry in the game treatment of Blue and Red.

EXAMPLE 1. THE CAMPERS

It may help to fix these ideas if we give a specific physical realization. When the payoffs are all positive, we may interpret them as the altitudes of points in a mountainous region. The various Blue and Red strategies then correspond to the latitudes and longitudes of these points.

To supply some actors and motivation for a game, let's suppose that a man and wife—being very specific always helps, so let's name them Ray and Dotty—are planning a camping trip, and that Ray likes high altitudes and Dotty likes low altitudes. The region of interest to them

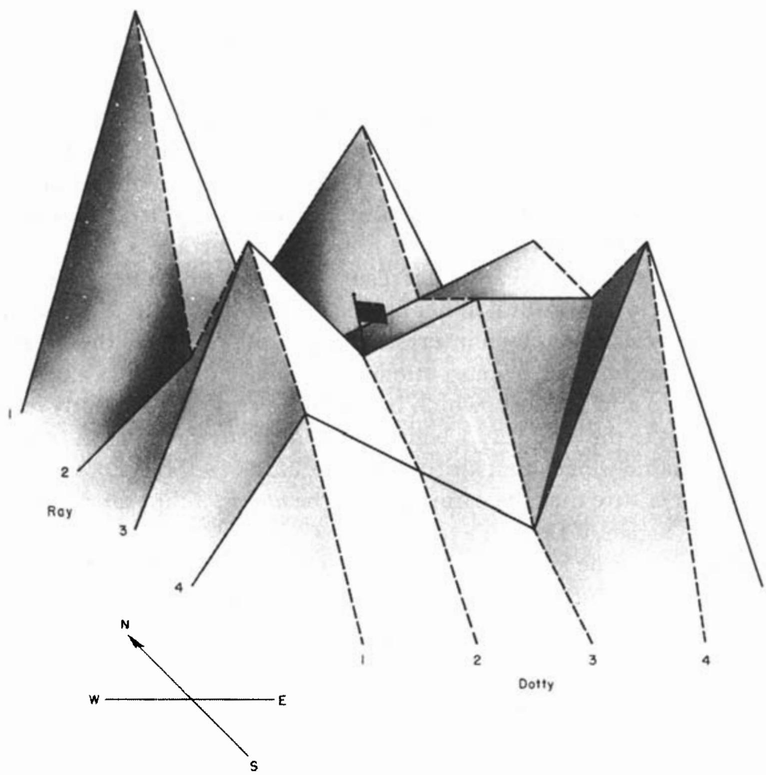

is crisscrossed by a network of fire divides, or roads, four running in each direction. The campers have agreed to camp at a road junction. They have further agreed that Ray will choose the east-west road and that Dotty will choose the north-south road, which jointly identify the junction. If Game Theory doesn't save them, frustration will kill them.

The junctions on the roads available to Ray have these altitudes (in thousands of feet):

Ray	1	7	2	5	1
	2	2	2	3	4
	3	5	3	4	4
	4	3	2	1	6

Being a reasonable person, who simply wants to make as much as possible out of this affair, he is naturally attracted to the road Ray 1—with junctions at altitudes of 7, 2, 5, and 1—for it alone can get him the 7-thousand-foot peak. However, he immediately recognizes this kind of thinking as dream stuff; he does not dare undertake a plan which would realize him a great deal if it succeeds, but which would lead to disaster if Dotty is skillful in her choice. Not anticipating that she will choose carelessly, his own interests compel him to ignore the breath-taking peaks; instead, he must attend particularly to the sinks and lows, of one kind and another, which blemish the region. This study leads him finally to the road Ray 3, which has as attractive a low as the region affords, namely, one at an altitude of 3 thousand feet. By choosing Ray 3, he can ensure that the camp site will be *at least* 3 thousand feet up; it will be higher, if Dotty is a little careless.

His wife—as he feared—is just as bright about these matters as he is. The critical altitudes on her roads are listed in the following table:

Dotty

1	2	3	4
7	2	5	1
2	2	3	4
5	3	4	4
3	2	1	6

As she examines these, she knows better than to waste time mooning over the deep valleys of Dotty 3 and Dotty 4, much as she would like to camp there. Being a realist, she examines the peaks which occur on her roads, determined to choose a road which contains only little ones. She is thus led, finally, to Dotty 2, where a 3-thousand-foot camp site is the worst that can be inflicted on her.

We now note that something in the nature of a coincidence has occurred. Ray has a strategy (Ray 3) which guarantees that the camp site will have an altitude of 3 thousand feet or more, and Dotty has one (Dotty 2) which ensures that it will be 3 thousand feet or less. In other words, either player can get a 3-thousand-foot camp site by his own efforts, in the face of a skillful opponent; and he will do somewhat better than this if his opponent is careless.

When the guaranteed minimum and maximum payoffs of Blue and Red are exactly equal, as they are here, the game is said to have a *saddle-point,* and the players should use the strategies which correspond to it. If either alone departs from the saddle-point strategy, he will suffer unnecessary loss. If both depart from it, the situation becomes completely fluid and someone will suffer.

Note too this consequence of having a saddle-point: security measures are not strictly necessary. Either Ray or Dotty can openly announce a choice (if it is the proper one), and the other will be unable to exploit the information and force the other beyond the 3-thousand-foot site.

We remarked that the existence of a Game Theory saddle-point is something of a coincidence. Yet it corresponds to a pass or saddle-point in the mountains, and almost any complicated arrangement of mountains will contain many passes. The trick is that a mountain pass must have special features to make it qualify as a Game Theory saddle-point. For one, the road through the pass must run north and south; i.e., this road must lie within the choice of the player who wants to keep the payoffs small. Another feature is that there must be no high ground north or south of the pass. Another is that there must be no low ground east or west of the pass. It is rather reasonable to find qualifications such as these; for after all a mountain pass is a *local* feature of the terrain, so some additional qualities are needed to ensure that it have the global properties of being best over the entire region.

While we shall always find it worth while to inspect games for saddle-points, the incidence of saddle-points is not very great, in general. In a 4×4 game, such as the present one, there is about one chance in ten that a matrix of random numbers will have a saddle-point.

In the present instance, the saddle-point can be eliminated by making an apparently minor change in the matrix, at any one of several points. For instance, if the altitude of the junction at the intersection of Ray 4 with Dotty 2 were changed from 2 to 6, the character of the game would become very different. In that game, i.e., in

Dotty

Ray	1	2	3	4
1	7	2	5	1
2	2	2	3	4
3	5	3	4	4
4	3	6	1	6

our elementary deductions regarding choice of strategies break down. If Ray argues as before, he will be led again to the road Ray 3, which ensures that the came site will be 3 thousand feet up, or higher; but Dotty will be led to Dotty 3 this time, which only guarantees that the camp site will be at 5 thousand feet, or less.

Thus there is a gap, between 3 and 5 thousand feet, in which the situation is out of control. Your intuition may suggest that there should be a way to play the game which will close this gap. In fact there is a way; but we must begin our study with simpler situations. In passing, we remark that good play will now require a more elaborate security system than was needed in the case of a saddle-point. In particular, the players will need to express their choices of stategy simultaneously, or in sealed ballots. What they should write on these ballots is quite a problem.