CHAPTER 3

Game Theory, Guess It, Foxholes

THE THEORY

GAME THEORY, one of the most useful branches of modern mathematics, was anticipated in the early 1920's by the French mathematician Emile Borel, but it was not until 1926 that John von Neumann gave his proof of the minimax theorem, the fundamental theorem of game theory. On this cornerstone he built almost single-handedly the beautiful basic structure of game theory. His classic 1944 work, *Theory of Games and Economic Behavior*, written with the economist Oskar Morgenstern, created a tremendous stir in economic circles (see "The Theory of Games," by Oskar Morgenstern, *Scientific American*, May 1949). Since then game theory has developed into a fantastic amalgam of algebra, geometry, set theory, and topology, with applications to competitive situations in business, warfare, and politics as well as economics.

Attempts have been made to apply game theory to all kinds of other conflict situations. What is the nation's optimal strategy in the Cold War Game? Is the Golden Rule, some philosophers have asked, the best strategy for maximizing happiness payoffs in the Great Game of Life? How can a scientist best play the Induction Game against his formidable opponent Nature? Even psychiatry has not been immune. Although Eric Berne's "transactional therapy" (popularized by his best-selling *Games People Play*) makes no use of game theory mathematics, it borrows many of its terms from, and obviously has been influenced by, the game theory approach.

Most game theory work has been on what are called two-person zero-sum games. This means that the conflict is between two players (if there are more, the theory gets muddied by coalitions) and whatever one player wins the other loses. (One reason game theory is difficult to apply to international conflicts is that they are not zero-sum; a loss for the U.S.S.R. is not necessarily a gain for the United States, for example.) The main purpose of this chapter is to present an interesting two-person zero-sum card game invented by Rufus Isaacs, a game theory expert who wrote *Differential Games* (John Wiley, 1965) and is professor of applied mathematics at Johns Hopkins University. But first a quick look at some elementary game theory.

Consider this trivial game. Players A and B simultaneously extend one or two fingers, then B gives A as many dollars as there are fingers showing. The game obviously is unfair since A always wins. How, though, should A play so as to make his wins as big as possible, and how should B play so as to lose as little as possible? Most games have numerous and complicated strategies, but here each player is limited to two: he can show one finger or he can show two. The "payoff matrix" can therefore be drawn on a 2-by-2 square as shown in Figure 4, left. By convention, A's two strategies are shown on the left and B's two strategies are shown above. The cells hold the payoffs for every combination of strategies. Thus if A shows one finger and B two, the intersection cell shows a \$3 payoff to A. (Payoffs are always given as payments from B to A even when the money actually goes the other way, in which case the payment to B is indicated by a minus sign.)

If A plays one finger, the least he can win is 2. If he plays two fingers, the least he can win is 3. The *largest* of these lows (the 3 at lower left) is called the maxmin (after maximum of

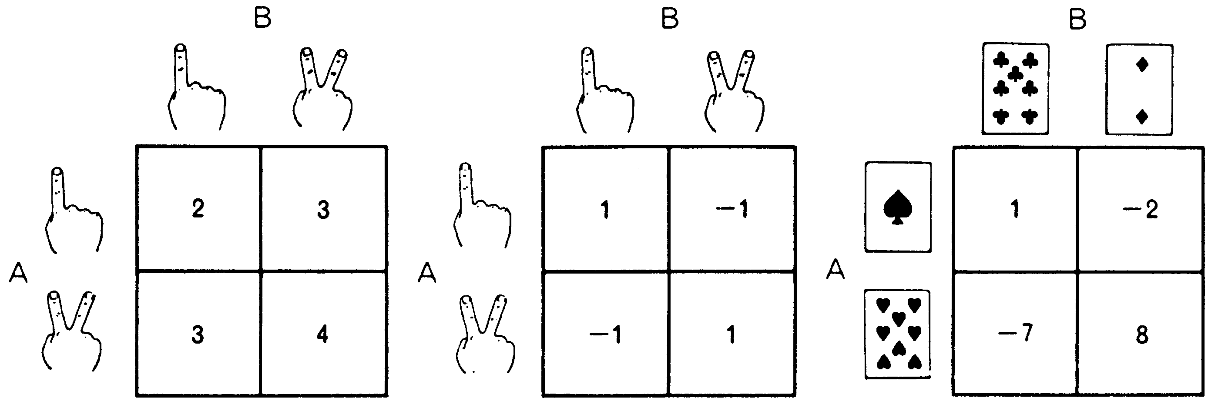

Payoff matrix for a trivial game

Matrix for odds-and-evens game

Matrix for card game

FIGURE 4

the minima). If B plays one finger, the most he can lose is 3. If he plays two fingers, the most he can lose is 4. The least of these highs (again the 3 at lower left) is called the minmax (minimum of the maxima). If the cell that holds the minmax is also the cell that holds the maxmin, as it is in this case, the cell is said to contain the game's "saddle point" and the game is "strictly determined."

Each player's best strategy is to play a strategy that includes the saddle point. A maximizes his gain by always showing two fingers; B minimizes his loss by always showing one. If both play their best, the payoff each time will be $3 to A. This is called the "value" of the game. As long as either player uses his optimal strategy he is sure to receive a payoff equal to or better than the game's value. If he plays a nonoptimal strategy, there is always an opposing strategy that will give him a poorer payoff than the value. In this case the game is of course so trivial that both optimal strategies are intuitively obvious.

Not all games are strictly determined. If we turn the finger game into "odds and evens" (equivalent to the game of matching pennies), the payoff matrix becomes the one shown in Figure 4, middle. When fingers match, A wins $1; when they do not match, B wins \$1. Since A's maxmin is -1 and B's minmax is 1, it is clear there is no saddle point. Consequently neither player finds one strategy better than the other. It would be foolish, for example, for A to adopt the strategy of always showing two fingers because B could win every time by showing one finger. To play optimally each player must mix his two strategies in certain proportions. Ascertaining the optimal proportions can be difficult, but here the symmetry of this simple game makes it obvious that they are 1 : 1.

This introduces an all-important aspect of game theory: to be effective the mixing must be done by a randomizing device. It is easy to see why nonrandom mixing is dangerous. Suppose A mixes by alternating one and two fingers. B catches on and plays to win every time. A can adopt a subtler mixing pattern but there is always the chance that B will discover it. If he tries to randomize in his head, unconscious biases creep in. When Claude E. Shannon, the founder of information theory, was at the Bell Telephone Laboratories, he and his colleague D. W. Hagelbarger each built a penny-matching computer that consistently won against human players when they made their own choices by pressing one of two buttons. The computer analyzed its opponent's plays, detected nonrandom patterns, and played accordingly. Because the two machines used different methods of analyzing plays, they were pitted against each other "to the accompaniment," Shannon disclosed, "of small side bets and loud cheering" (see "Science and the Citizens," *Scientific American*, July 1954). The only way someone playing against such a machine can keep his average payoff down to zero is to use a randomizer—for example, flipping a penny each time to decide which button to push.

The game matrix shown in Figure 4, right, provides an amusing instance of a game with a far from obvious mixed strategy. Player A holds a double-faced playing card made by pasting a black ace back to back to a red eight. Player B has a similar double card: a red two pasted to a black seven. Each chooses a side of his card and simultaneously shows it to the other. A wins if the colors match, B if they fail to match. In every case the payoff in dollars is equal to the value of the winner's card.

The game *looks* fair (has a value of zero) because the sum of what A can win ($8 + 1 = 9$) is the same as the sum of what B can win ($2 + 7 = 9$). Actually the game is biased in favor of B, who can win an average of $1 every three games if he mixes his two strategies properly. Since 8 and 1, in one diagonal, are each larger than either of the other two payoffs, we know at once that there is no saddle point. (A 2-by-2 game has a saddle point if and only if the two numbers of either diagonal are *not* both higher than either of the other two numbers.) Each player, therefore, must mix his strategies.

Without justifying the procedure, I shall describe one way to calculate the mixture for each player. Consider A's top-row strategy. Take the second number from the first: $1 - (-2) = 3$. Do the same with the second row: $-7 - 8 = -15$. Form a fraction (ignoring any minus signs) by putting the last number above the first: 15/3, which simplifies to 5/1. A's best strategy is to mix in the proportions 5 : 1, that is, to show his ace five times for every time he shows his seven. A die provides a convenient randomizer. He can show his ace when he rolls 1, 2, 3, 4, or 5, his seven when he rolls 6. The randomizer's advice must, of course, be concealed from his opponent, who otherwise would know how to respond.

B's best strategy is similarly obtained by taking the bottom numbers from the top. The first column yields 8, the second -10. Ignoring minus signs and putting the second above the first gives 10/8, or 5/4. B's best strategy is to show his seven five times to every four times for the two. As a randomizer he can use a table of random numbers, playing the seven when the digit is 1, 2, 3, 4, or 5 and the deuce when it is 6, 7, 8, or 9.

To calculate the game's value (the average payoff to A), assume that the cells are numbered left to right, top to bottom, a, b, c, d. The value is

$$\frac{ad - bc}{a + d - b - c}.$$

The formula in this case has a value of $-1/3$. As long as A plays his best strategy, the $5:1$ mixture, he holds his average loss per game to a third of a dollar. As long as B plays his best mixture, the $5:4$, he ensures an average win per game of a third of a dollar. The fact that every matrix game, regardless of size or whether it has a saddle point, has a value, and that the value can be achieved by at least one optimal strategy for each player, is the famous minimax theorem first proved by von Neumann. Readers may enjoy experimenting with 2-by-2 card games of this type but using different cards, and calculating each game's value and optimal strategies.

Most two-person board games, such as chess and checkers, are played in a sequence of alternating moves that continues until either one player wins or the game is drawn. Since the number of possible sequences is vast and the number of possible strategies is astronomically vaster, the matrix is much too enormous to draw. Even as simple a game as ticktacktoe would require a matrix with tens of thousands of cells, each labeled 1, -1, or 0. If the game is finite (each player has a finite number of moves and a finite number of choices at each move) and has "perfect information" (both players know the complete state of the game at every stage before the current move), it can be proved (von Neumann was the first to do it) that the game is strictly determined. This means that there is at least one best pure strategy that always wins for the first or for the second player, or that both of the players have pure strategies that can ensure a draw.

THE GAME OF GUESS IT

ALMOST ALL card games are of the sequential-move type but with incomplete information. Indeed, the purpose of making the backs of cards identical is to conceal information. In such games the optimal strategies are mixed. This means that a player's best decision on most or all of his moves can be given only probabilistically and that the value of the game is an average of what the maximizing player will win in the long run. Poker, for instance, has a best mixed strategy, although (as in chess and checkers) it is so complicated that only simplified forms of it have been solved.

Isaacs' card game, named Guess It by his daughter Ellen, is remarkable in that it is a two-person sequential-move game of incomplete information, sufficiently complicated by bluffing to make for stimulating play, yet simple enough to allow complete analysis.

The game uses eleven playing cards with values from ace to jack, the jack counting as 11. The packet is shuffled. A card is drawn at random and placed face down in the center of the table, neither player being aware of its value. The remaining ten cards are dealt, five to each player. The object of the game is to guess the hidden card. This is done by asking questions of the form "Do you have such-and-such a card?" The other player must answer truthfully. No card may be asked about twice.

At any time, instead of an "ask" a player may end the game by a "call." This consists of naming the hidden card. The card is then turned over. If it was correctly named, the caller wins; otherwise he loses. To play well, therefore, a player must try to get as much information as he can, at the same time revealing as little as possible, until he thinks he knows enough to call. The delightful feature of the game is that each player must resort to occasional bluffing, that is, asking about a card he himself holds. If he never bluffed, then whenever he asked about a card not in his opponent's hand, the opponent would immediately know that card must be the hidden one—and would call and win. Bluffing is therefore an essential part of strategy, both for defense and for tricking the opponent into a false call.

If player A asks about a card, say the jack, and the answer is yes, both players will then know B has that card. Since it will not be asked about again, nor will it be called, the jack plays no further role in the game. B places it face up on the table.

If B does not have the jack, he answers no. This places him in a quandary, although one that proves to be short-lived. If he thinks A is not bluffing, he calls the jack and ends the game, winning if his suspicion is correct. If he does not call it and the hidden card *is* the jack, then A (who originally asked about it) will surely call the jack on his next play, for he will know with certainty that it is the hidden card. Therefore, if A does *not* call the jack on his next play, it means he had previously bluffed and has the jack in his hand. Again, because the location of the card then becomes known to both players, it plays no further role. It is removed and placed face up on the table. In this way hands tend to grow smaller as the game progresses. After each elimination of a card the players are in effect starting a new game with fewer cards in hand.

It is impossible to give here the details of how Isaacs solved the game. The interested reader will find it explained in his article "A Card Game with Bluffing" in *The American Mathematical Monthly* (Vol. 62, February 1955, pages 99–108). I will do no more here than explain the optimal strategies and how they can be played with the aid of two spinners made with the dials shown in Figure 5. Readers are urged first to play the game without these randomizers, keeping a record of n games between players A and B. They should then play another n games with only A using the spinners, followed by a third set of n games with only B using the spinners. (If both players use randomizers, the game degenerates into a mere contest of chance.) In this way an empirical test can be made of the efficacy of the strategy.

The dials can be copied or mounted on a rectangle of stiff cardboard. Stick a pin in the center of each and over each pin put the loop end of a bobby pin. A flip of the finger sends the bobby pin spinning. The spinners must of course be kept out of your opponent's view when being used, either by turning your back when you spin them or keeping them on your lap below the edge of the table. After using them you must keep a "poker face" to avoid giving clues to what the randomizers tell you to do.

The top dial tells you when to bluff. The boldface numbers

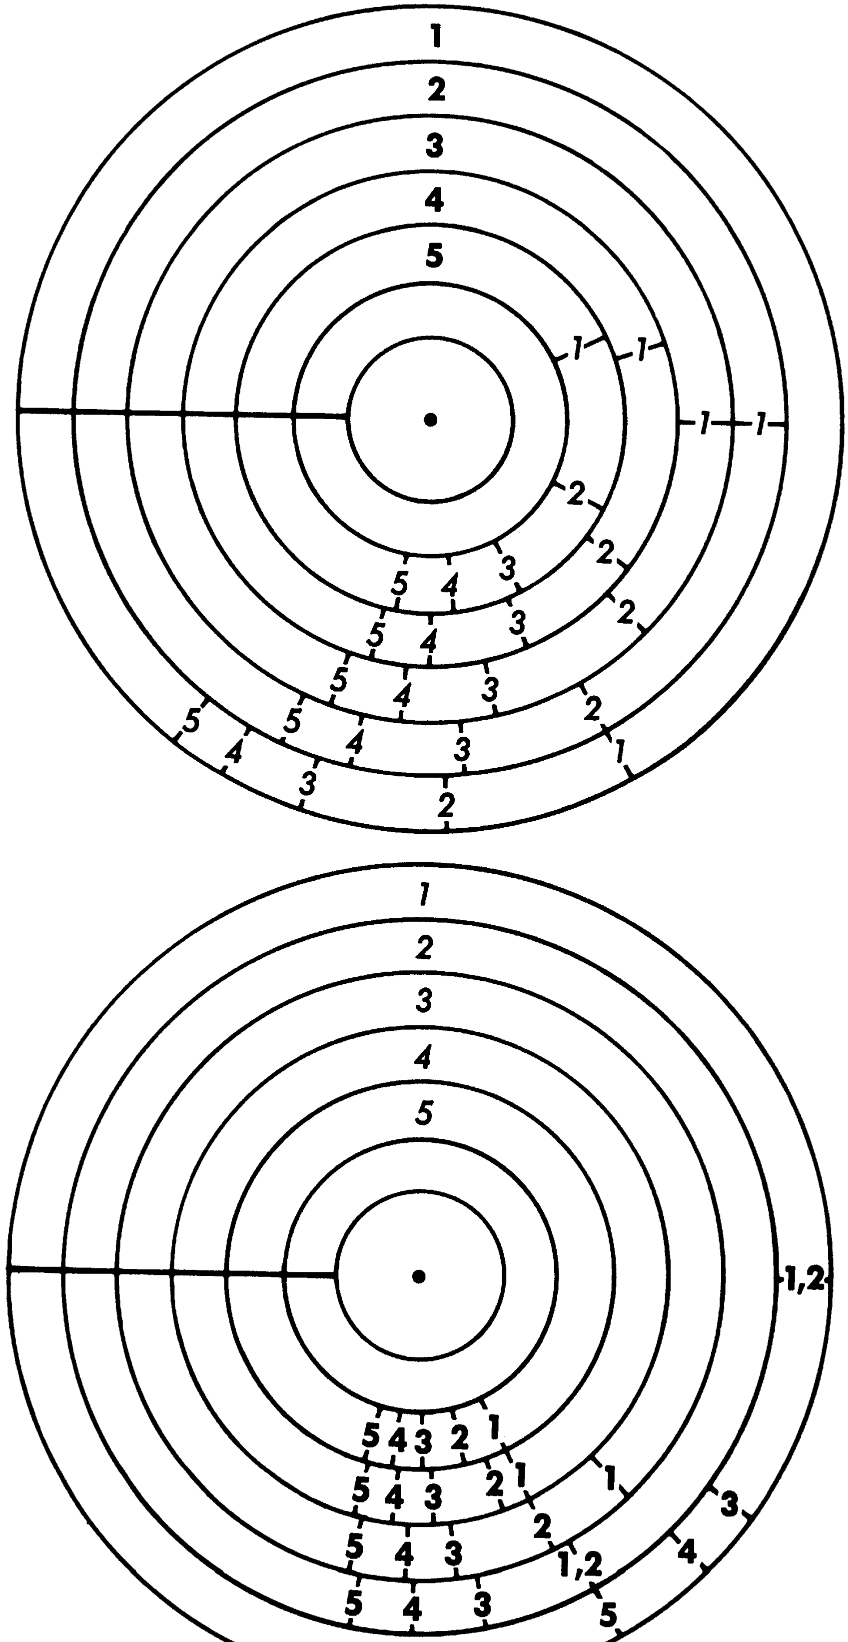

FIGURE 5

Randomizing dials for deciding when to bluff (top) and when to call (below) give the number of cards in your hand. The other numbers scattered over the dial and attached to marks stand for the number of cards in your opponent's hand. Assume that you have three cards and he has two. Confine your attention to the ring labeled with a boldface 3. Spin the bobby pin. If it stops in the portion of the ring that extends clockwise from mark 2 to the heavy horizontal line, you bluff. Otherwise you ask about a card that could be in your opponent's hand.

In either case, asking or bluffing, pick a card at random from the possibilities open to you. If a strict empirical test of strategy is to be made, you should use a randomizer for this selection. The simplest device would be a third spinner on a circle divided into 11 equal sectors and numbered 1 to 11. If the first spinner tells you to bluff, for example, and you have two, four, seven, and eight in your hand, you spin the third spinner repeatedly until it stops on one of those numbers. Without the aid of such a spinner, simply select at random one of the four cards in your hand. The danger of your opponent's profiting from an unconscious mental bias is so slight, however, that we shall assume a third spinner is not used.

The bottom dial is used whenever you have just answered no to an ask. On this dial the rings are labeled with italic numbers to indicate that they correspond to the number of cards in your opponent's hand. The boldface numbers near the marks give the number of cards you hold. As before, pick the appropriate ring and spin the bobby pin. If it stops in the portion of the ring that extends from the proper mark clockwise to the horizontal line, call the card previously asked. If it does not stop in this portion of the ring, your next action depends on whether your opponent has just one card or more than one. If he has only one, call the other unknown card. If he has more than one (and you have at least one card), you must ask. To decide whether to bluff or not, spin the first dial, but now you must pick your ring on the assumption that his hand is reduced by one card. The reason for this is that if he did not bluff on his last ask, your "no" answer will enable him to win on his next move. You therefore play as if he were bluffing and the game were to continue, in which case the card he asked about has been taken out of the game by your "no" answer even though it is not actually placed face up on the table until after his next move.

In addition to the circumstances just explained, you call only under the following circumstances: (1) When you know the hidden card. (This occurs when you have asked without bluffing and received a "no" reply, and he has not won the game by calling on his next turn; and it occurs of course when he holds no cards.) (2) When you have no cards and he has one or more, because if you do not call, he surely will call and win on his next play. If each of you holds just one card, it is immaterial whether you call or ask; the probability of winning is 1/2 and is obtained either way. (3) When instructed to call by the second dial, as explained before.

The table in Figure 6 shows the probability of winning for the player who has the move. The number of his cards appears at the top, those of the other player on the left. At the beginning, assuming that both players use randomizers for playing their best, the first player's probability of winning is .538, or

NUMBER OF CARDS IN OPPONENT'S HAND	NUMBER OF CARDS IN PLAYER'S HAND: 1	2	3	4	5
1	.5	.667	.688	.733	.75
2	.5	.556	.625	.648	.680
3	.4	.512	.548	.597	.619
4	.375	.450	.513	.543	.581
5	.333	.423	.467	.512	.538

FIGURE 6
Chart of probabilities of winning Guess It game

slightly better than 1/2. If the payoff to the first player is $1 for each win and zero for each loss, then $.538 is the value of the game. If after each game the loser pays the winner $1, the first player will win an average of 538 games out of every 1,000. Since he receives $538 and loses $462, his profit is $76, and his average win per game is $76/1,000, or $.076. With these payoffs the game's value is a bit less than eight cents per game. If the second player does not use randomizers, the first player's chance of winning increases substantially, as should appear in an empirical test of the game.

FOXHOLES

HERE IS a simple, idealized war game that Isaacs uses to explain mixed strategies to military personnel. One player, the soldier, has a choice of hiding in any one of the five foxholes shown in Figure 7. The other player, the gunner, has a choice of firing at one of the four spots *A*, *B*, *C*, *D*. A shot will kill the soldier if he is in either adjacent foxhole—shot *B*, for example, is fatal if he is in foxhole 2 or 3.

"We can see the need for mixing strategies," Isaacs writes, "for the soldier might reason: 'The end holes are vulnerable to only one shot, whereas the central holes can each be hit two ways. Therefore I'll hide in one of the end holes.' Unfortunately the gunner might foresee this reasoning and fire only at *A* or *D*. If the soldier suspects that the gunner will do this, he will hide in a central hole. But now the gunner may still be one-up by guessing that the soldier will think he will think this way, therefore he aims at the center. These attempts at outthinking the opponent lead only to chaos. The only way either player can be sure of deceiving his opponent is by mixing his strategies."

Assume that the payoff is 1 if the gunner kills the soldier, 0 if he does not. The value of the game is then the same as the probability of a hit. What are the optimal strategies for each player and what is the game's value?

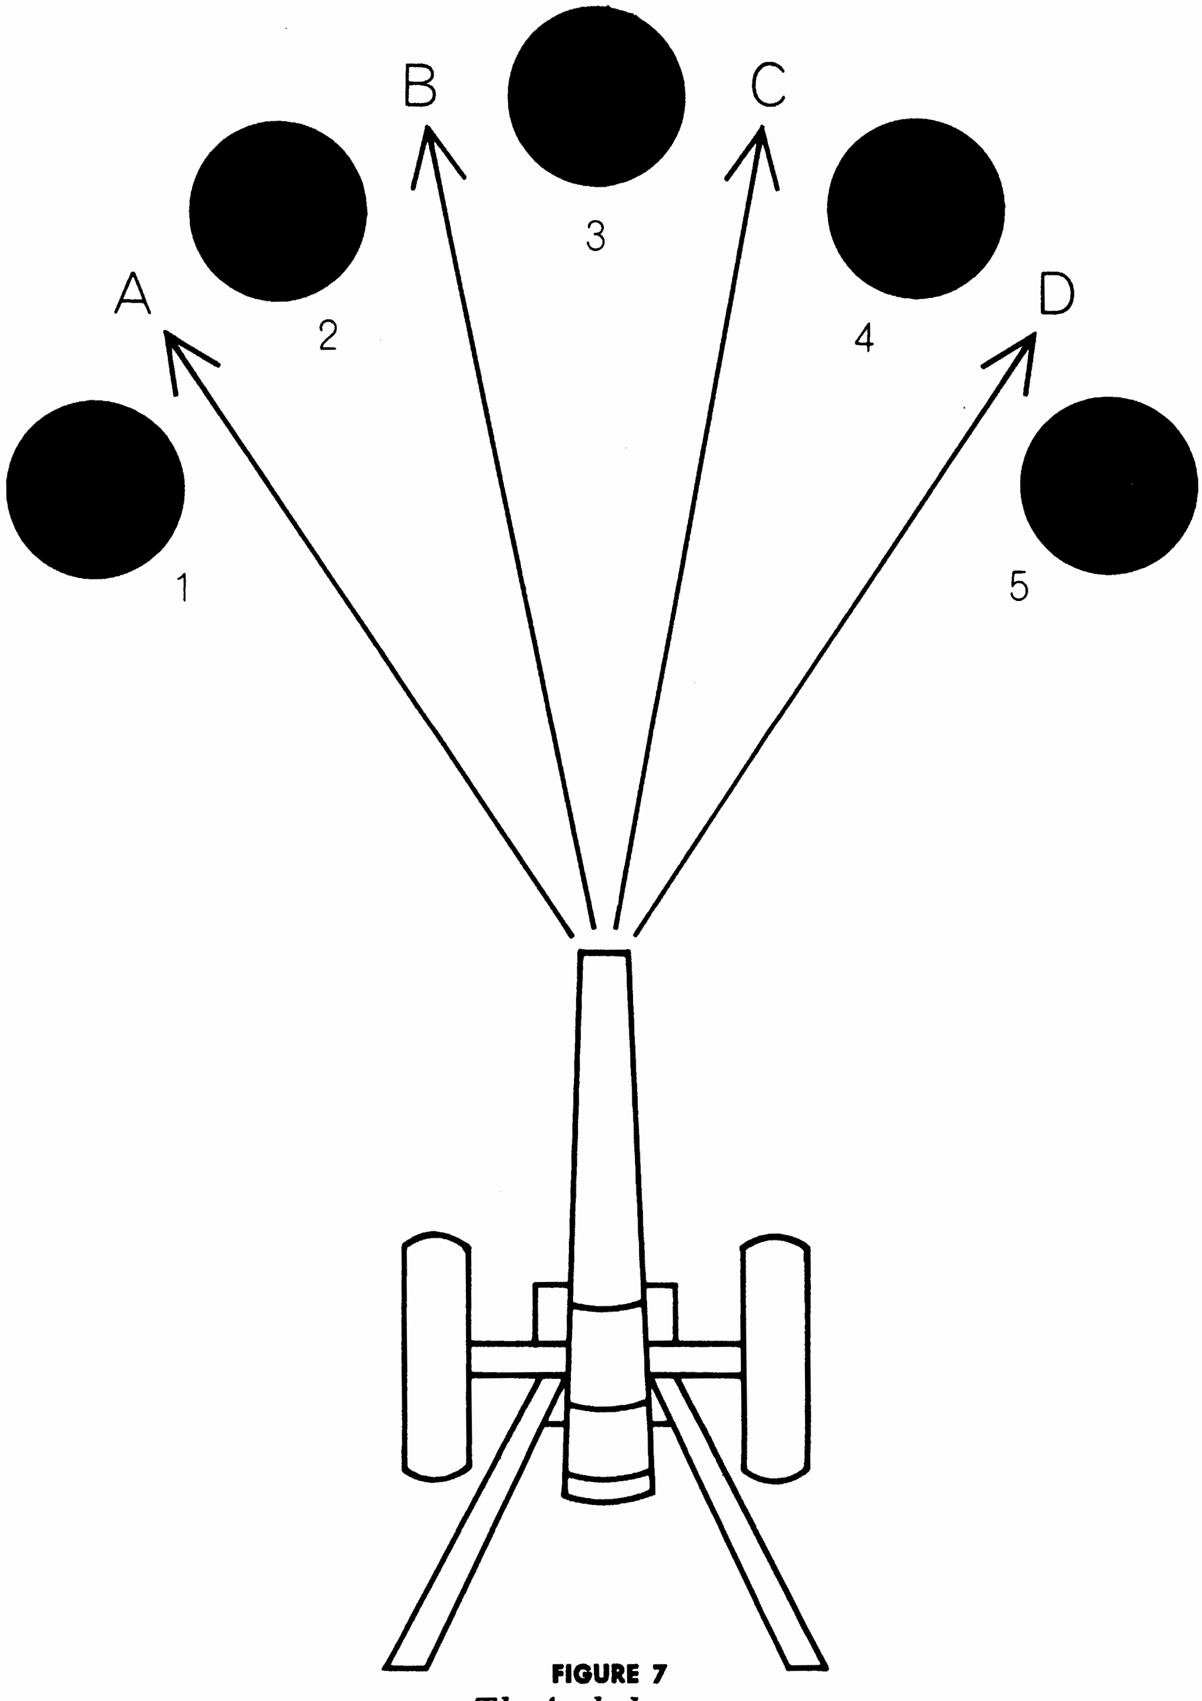

FIGURE 7
The foxhole game

ANSWERS

RUFUS ISAACS' foxhole game concerns a soldier who has a choice of hiding in one of five foxholes in a row and a gunner who has a choice of firing at one of four spots, A, B, C, D, between adjacent foxholes. An equivalent card game can be played with five cards, only one of which is an ace. One player puts the cards face down in a row. The other player picks two adjacent cards and wins if one of them is the ace.

"One can easily write a 4-by-5 matrix for this game and apply one of the general procedures described in the textbooks," Isaacs writes. "But, with a little experience, one learns in simple cases like this how to surmise the solution and then verify it."

The soldier's optimal mixed strategy is to hide only in holes 1, 3, and 5, selecting the hole with a probability of 1/3 for each. The gunner has a choice of any of an infinite number of optimal strategies. He assigns probabilities of 1/3 to A, 1/3 to D, and any pair of probabilities to B and C that add to 1/3. (For example, he could let B and C each have a probability of 1/6, or he could give one a probability of 1/3 and the other a probability of 0.)

To see that these strategies are optimal, consider first the soldier's probability of survival. If the gunner aims at A, the soldier has a 2/3 chance of escaping death. The same is true if the gunner aims at D. If he aims at B, he hits only if the soldier is in hole 3, so that again the probability of missing is 2/3. The same is true if he aims at C. Since each individual choice gives the soldier a 2/3 probability of survival, the probability remains 2/3 for any mixture of the gunner's choices. Thus the soldier's strategy ensures him a survival probability of at least 2/3.

Consider now the gunner's strategy. If the soldier is in hole 1, he has a hit probability of 1/3. If the soldier is in hole 2, he is hit only if the gunner fires at A or B, and consequently the probability of a hit is 1/3 plus whatever probability the gunner assigned to B. If the soldier is in hole 3, he is hit only if the gunner fires at B or C, to which are assigned probabilities adding to 1/3. Therefore the probability of a hit here is 1/3. If the soldier is in hole 4, the probability of a hit is 1/3 plus the probability assigned to *C*. If he is in hole 5, the probability is 1/3. Thus the gunner's strategy guarantees him a probability of at least 1/3.

Assuming a payoff of 1 to the gunner if he kills the soldier, 0 if he doesn't, the value of the game is 1/3. The gunner has an infinite number of strategies that guarantee him a hit probability of at least 1/3. It is possible he could do better against a stupid opponent, but against good opposition he can hope for no more because, as we have seen, the soldier has a strategy that keeps the probability of his death down to 1/3. A similar argument holds from the soldier's standpoint. By using his optimal strategy he keeps the payoff at 1/3 and cannot hope to do better because the gunner has a way of making it at least 1/3. As a further exercise, readers can try to prove there are no optimal strategies other than those explained here.

"The process of surmising the solution is not as hard as it looks," Isaacs adds. "The reader can so convince himself by generalizing this solution to the same game but with n foxholes. For odd n the preceding solution carries over in an almost obvious way, but with even n one encounters some modest novelty."