

1. Introduction

It is generally agreed that the minimax solution to a two-person zero-sum matrix game is intuitively satisfactory. Now in many applications of game theory, a game is not described a priori in matrix (or "normal" or "strategic") form, but rather in extensive form, i.e. by its rules. A game described in such a way may be reduced to a matrix game by means of the concept of "strategy". If, moreover, it is of perfect recall, then all mixed strategies, and in particular the optimal strategies of each player, are equivalent to behavior strategies [K]. The usual conclusion from these considerations is that for 2-person 0-sum games in extensive form, the minimax solution is intuitively satisfactory; and that in games of perfect recall, in particular, the players would do well to play in accordance with optimal¹ behavior strategies. In this paper we shall discuss some examples that, we believe, cast doubt on these conclusions.

In our main example, each player has only a single move, and each has a unique optimal behavior strategy. Nevertheless, when faced with his single move, there seems to be no good reason for player I to play in accordance with the dictates of his unique optimal behavior strategy. Specifically, the arguments for the minimax solution in zero-sum games have, in this example, no more force than those in favor of a Nash equilibrium point in nonzero-sum games; and the latter are known often to be less than convincing. Thus, it appears from this example that the passage from the extensive form to the matrix form is not without pitfalls even for zero-sum games. We shall study the nature of this passage and also indicate the conditions under which it remains valid.

One may be inclined to think that our criticism is relevant only for games in which the optimal strategies are mixed. This is not the case, and we shall show how our example can be modified so that all the phenomena occur in pure strategies.

We shall also show that even if the game is originally given in matrix form, similar phenomena may occur, provided that the matrix has no optimal pure strategies.

The authors would like to acknowledge some very helpful discussions with a number of people, including especially John Harsanyi and Guillermo Owen. Of course they bear no part of the responsibility for the opinions expressed herein; it is quite possible that one or both heartily disagree with some or all of these opinions.

¹ Throughout this paper "optimal" will refer to minimax and maximin strategies.

The paper contains a discussion of the implications of this example to applications of the minimax principle. Also, in the last section, we shall discuss the relevance of the example to Harsanyi's theory of games with incomplete information.

2. A Nonzero-Sum Game²

First, let us consider a nonzero-sum game, to support our statement that arguments in favor of Nash equilibrium points³ are sometimes less than convincing. Although arguments against the equilibrium point concept are well known, this particular game will make it easier to understand the main example below.

Consider the 2×2 nonzero-sum game in normal form whose payoff matrix is as follows (as usual, the row signifies the choice of player I, and the column the choice of player II; this convention will be maintained throughout the paper):

	L	R
L	1, 0	0, 1
R	0, 3	1, 0

This game has a unique equilibrium point, and it is in mixed strategies: Player I's strategy is $(\frac{3}{4}L, \frac{1}{4}R)$, and player II's strategy is $(\frac{1}{2}L, \frac{1}{2}R)$. If the equilibrium point concept is at all convincing, it should certainly be convincing here, where the equilibrium point is unique.

But is it? The payoff at the unique equilibrium point is $(\frac{1}{2}, \frac{3}{4})$. Now, $\frac{1}{2}$ and $\frac{3}{4}$ are also the maximin values of this game for players I and II respectively, i.e. these are the security levels of the players, the amounts that they can *guarantee* to themselves (in mixed strategies). Unfortunately, the equilibrium strategies *do not*, in fact, guarantee these values! In order to guarantee $\frac{1}{2}$, player I should play $(\frac{1}{2}L, \frac{1}{2}R)$; whereas, by playing $(\frac{3}{4}L, \frac{1}{4}R)$, he guarantees only $\frac{1}{4}$. Similarly, in order to guarantee $\frac{3}{4}$, player II should play $(\frac{1}{4}L, \frac{3}{4}R)$; by playing $(\frac{1}{2}L, \frac{1}{2}R)$, he guarantees only $\frac{1}{2}$. If the equilibrium strategies are played at all, they are presumably played with the *hope* that each player will obtain his equilibrium payoff. But why play "with a hope", when the players could each *guarantee* the same payoff by choosing their maximin rather than their equilibrium strategies? To be sure, we do not know what to recommend in this situation (because the maximin strategies are not in equilibrium); but the maximin strategies certainly seem preferable to the equilibrium strategies.

The reason for the curious phenomenon exhibited in this game is that to achieve equilibrium, each player must play *against* his opponent rather than *for* himself. In choosing his strategy $(pL, (1-p)R)$, player I must make certain that player II cannot improve his payoff, and in effect this means that p is chosen in such a way that player II's payoff is the same no matter which column he chooses; the choice of such a p depends only on player II's matrix. Player I's matrix could be changed drastically without changing player I's strategy at all; if say, his matrix would be

9	2
1	3

² This example was communicated to us by John Harsanyi.

³ An "equilibrium point" means an equilibrium pair of strategies. It should be distinguished from "equilibrium payoff".

rather than

1	0
0	1

player I's equilibrium strategy would not change at all. A similar situation holds for player II. Thus each player's equilibrium strategy depends chiefly⁴ on the magnitude of the entries in the *other* player's matrix; whereas his maximin strategy depends exclusively on his own matrix.

To sum up, we have exhibited a game with a unique equilibrium point, which has the following property: The equilibrium point yields to each player exactly his security level; and yet, neither one of the equilibrium strategies *guarantees* the security level. Under these conditions, the use of the equilibrium strategies does not seem reasonable.

3. The Main Example

Consider the 2-person 0-sum game G with the following rules: First chance chooses one of the two symbols L and R , with probabilities $\frac{1}{3}$ and $\frac{2}{3}$ respectively. If chance chose L , player I never gets to make a choice; the payoff (to I) is then dependent on II's choice only. Specifically, II may choose either L or R , and the payoff will be 9 or 0 respectively. If chance chose R , players I and II each choose one of the symbols L and R (without knowledge of the other's choice); the payoff is then in accordance with the following table:

	L	R
L	6	0
R	0	6

It is stipulated that II must choose without knowing chance's choice. Of course, if I is called upon to move at all, he will know that chance has chosen R .

The above table can also be viewed as the matrix of a 2-person 0-sum game in its own right; it will be denoted G_R . The relationship of G to G_R will be clarified presently. The extensive form of G is given in Figure 1. The figure has been drawn so that L and R correspond throughout to left and right respectively.

In order to avoid misunderstanding, let us stress that we are considering a 1-shot game. It is assumed that the players play the game just once and that they never expect to meet and play again.

Players I and II each have 2 pure strategies in G . The matrix of G is

	L	R
L	7	0
R	3	4

It has unique optimal strategies, namely, $(\frac{1}{3}L, \frac{2}{3}R)$ for I, and $(\frac{1}{2}L, \frac{1}{2}R)$ for II. These strategies may also be viewed as behavior strategies. The value of the game is $3\frac{1}{2}$.

⁴ We say "chiefly" because this continues to hold only as long as the "type" of the game remains unchanged; i.e., as long as the order relations between the entries in each payoff matrix are the same as in the original example.

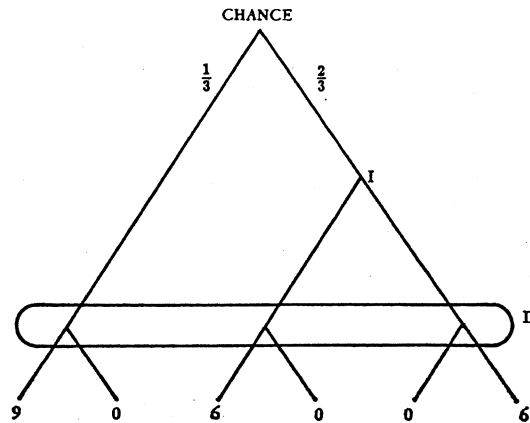


FIGURE 1

Note that if both players play optimally, then the payoff is $4\frac{1}{2}$ if chance chose L , and 3 if chance chose R .

The key to a proper understanding of G is the remark that II's optimal strategy in G is the same as in G_R . In other words, player II ignores the possibility that chance chooses L , and plays as if chance chooses R with certainty. Intuitively, it is not immediately clear why he should do so. One might have thought, for example, that he would hedge, i.e., choose some mixture of the minimax strategy in G_R , namely $(\frac{1}{2}L, \frac{1}{2}R)$, and the strategy which is best for him if chance chose L , namely R . But this is not the case, i.e. it is worthwhile for II to ignore the possibility of chance choosing L , and to accept the resulting penalty if in fact chance did choose L .

Let us now turn to player I's strategy. When chance chooses L , I has nothing to do. When chance chooses R , I's optimal strategy dictates $(\frac{1}{3}L, \frac{2}{3}R)$. To analyze this strategy, we distinguish between two cases: in the first, II plays optimally, and in the second, he does not. In the first case, since II's strategy is minimax in G_R , and since G_R is "totally mixed," the payoff will be the same no matter what I does. In the second case, by playing $(\frac{1}{3}L, \frac{2}{3}R)$ —a strategy that is not maximin in G_R —I may well get less than 3, the value of G_R . Why then does he not play $(\frac{1}{2}L, \frac{1}{2}R)$, which would guarantee the payoff of 3 for this "branch" of the game,⁵ and which is the most that can be guaranteed for this branch?

The answer is that if in G it were optimal for player I to play maximin in G_R , II would know this fact. He would therefore know that if chance chose R , it makes no difference what he does. Therefore he would concentrate on the possibility that chance chose L , and optimize on this assumption, i.e. choose R . This would lead to a loss for player I if chance chose L .

Now this argument may explain why $(\frac{1}{3}L, \frac{2}{3}R)$ is optimal in the mathematical sense of the word, but it does not convince us that player I should in fact play it. Note that the reason for $(\frac{1}{3}L, \frac{2}{3}R)$ when chance chose R is based entirely on the fact that chance might have chosen L . But if chance chose R , then player I *knows* this when he comes to make his choice. Why then should he take into account the possibility that chance might have chosen L , when he knows that in fact chance did not? If thereby he could somehow take advantage of II's lack of information, well and good. But we have seen that II should anyhow assume that chance chose R ; chance did in fact choose R ; player I knows that chance chose R ; yet we are told that player I should still play to protect his interests under the (false!) assumption that chance chose L !

⁵ I.e. the branch determined by chance's choice of R .

The situation is quite paradoxical: The player who does not know chance's choice boldly refuses to hedge; the player who knows that chance chose R , by playing maximin in G , still hedges on the possibility that chance chose L !

4. Discussion

Let us start by again making it clear that we are viewing a situation in which the game G is to be played once and only once. We are *not* considering repeated play. Under appropriate conditions of repeated play, it may be possible to justify some such solution as the above to G ; this is, however, a complicated subject, and we have no intention to get sidetracked into it here. Like von Neumann and Morgenstern, who emphatically rule out repeated play in their discussion (see [N-M, §17.3, pp. 146-147]), we will base our discussion here exclusively on the "one-shot" assumption.⁶

The arguments in favor of the minimax solution to a 2-person 0-sum game may be classified into two types, which we will call "equilibrium arguments" and "guaranteed value arguments".⁷ Let us examine these two types separately in the case of the game G before us.

The equilibrium arguments depend on the fact that a pair of optimal strategies is an equilibrium point; in case the minimax solution is unique, as in our game G , it is in fact the unique equilibrium point. A typical such argument runs as follows: If game theory is to "recommend" any specific pair of strategies, then each strategy in the pair must be "best possible" against the other strategy in the pair, i.e. the pair must be an equilibrium point. Otherwise, a knowledgeable player will know what the theory recommends for the other player, and so will want to choose a strategy that is better for him; that is, he will want to act contrary to what the theory recommends to him. Such an argument is especially convincing when, as in our case, the equilibrium point is unique.

We have no quarrel with these arguments; as we shall see, they retain their full force in the game G . We note, though, that these arguments also retain their full force in the case of certain nonzero-sum games with unique equilibrium points, such as the game described in §2; and in such games the equilibrium point is often rather unattractive, in spite of the arguments in its favor.

We now turn to the "guaranteed value" arguments; these run somewhat as follows: Two intelligent players are faced with a situation in which there is a number v such that I can guarantee that he can make II pay him at least v , and II can guarantee that he will not have to pay I more than v . Since both players are well aware of the situation as described, it is only reasonable to suppose that they will in fact act in order to guarantee an outcome that is $\geq v$ and $\leq v$ respectively, and therefore the outcome will be exactly v . This argument applies specifically to 2-person 0-sum games, and it is this argument that makes the minimax solution intuitively attractive in such games.

Let us now see how these two types of arguments apply to the game before us, and in particular, to the situation that arises if chance chose R . If game theory is to recom-

⁶ The chief reason for excluding repeated play in discussing the foundations of game theory is that the one-shot case is the more general one: A sequence of repeated plays can always be viewed as a single play of a large game (often called a "supergame" in the literature), but a single play of a game cannot in general be viewed as a sequence of repeated plays. Another reason is that if the payoff to a single play is expressed in utility units, then there are difficulties in relating the payoff to a sequence of plays to the payoffs to the individual plays; we might say that utilities are generally not "additive" in any natural sense.

⁷ In [N-M], they are called "indirect" and "direct" respectively; summaries of the two arguments may be found, respectively, in §17.3 (pp. 146-148) and in §17.8 (pp. 158-160).

mend any specific pair of strategies in this situation, then each strategy in the pair must be "best possible" against the other one. If we take into account the fact that player I knows that chance chose R and that player II does not, then it follows that such a strategy pair must be an equilibrium point of the nonzero-sum game with matrix⁸

6, -7	0, 0
0, -3	6, -4

There is one and only one such equilibrium point, and it coincides with the optimal behavior strategies in the game G . Thus we see that the equilibrium arguments in favor of the minimax solution retain their full force in the game we are considering.

Next, we consider the guaranteed value arguments. Unfortunately, the situation there is not so satisfactory. It is no longer true that there is a number v such that I can guarantee at least v and II can guarantee that he will not have to pay more than v . Indeed, after chance has chosen R , player I can guarantee a payoff of 3; but II can only guarantee that he will not have to pay more than $3\frac{1}{2}$, since he does not know what chance chose. What is even more disconcerting, though, is that I, though he could guarantee a payoff of 3, does not act to do so! This might be understandable if I could expect to get more than 3 by playing as he does; but this is not the case either, since he will get exactly 3 if II also plays optimally. In fact I's action is completely incomprehensible in the framework of a "guaranteed value" argument; it can be understood only in the context of an equilibrium argument (or in the context of repeated play, which we have ruled out).

Note the similarity between the present situation and the nonconstant-sum game discussed in §2.

What then went awry in the chain of reasoning with which we started this paper? The answer is as follows: In passing from the extensive to the matrix form of the game, we moved the point of time at which player I must reach his decision; in the extensive game he makes his decision after receiving the information as to chance's choice, and in the matrix game he makes his decision before receiving this information. It is quite true that passage from the extensive to the matrix form does not affect the possible courses of action at the disposal of I, and from this point of view the two situations are formally equivalent. But this passage does change the outlook of I. Indeed, if I decides on a strategy before chance makes its choice, there is considerable justification for taking into account the possibility that chance might choose L , and this is exactly what the maximin strategy does; but if I is informed before making his decision that chance has chosen R , there is no longer any justification for taking into account the possibility that chance might have chosen L . This change of outlook is ignored in the passage from the extensive to the normal form, and this fact invalidates the assertion in the introduction that the original extensive game may be "reduced" to a matrix game.

The argument may now be made that though by the rules of the extensive game I does not have to decide until after getting the information, he may of his own free will move the decision point to before chance's choice. Specifically, he could commit himself by taking the choice out of his own hands; he could for example give appropriate irreversible instructions to his agent, and then go home and sleep a just man's sleep, knowing that he has maximined.

⁸ We are of course not claiming that the situation is equivalent to this game.

We may note, though, that this is not the way in which the application of the minimax theorem to games of perfect recall has been generally understood. Moreover, it may be that an early self-commitment is physically impossible; player I may have no way of committing himself to taking a certain action before the time for this action comes. In this case it will be difficult to defend the minimax solution.

The fact that a passage from the extensive to the normal form may involve a shift in the time in which a choice must be made is known. Its impact, however, is usually discounted. It is claimed that a shift in time to the beginning of the play is immaterial, because as the beginning one can always choose that strategy which specifies how one would act when later reaching any specific stage. Our example points out the flaw in this reasoning.

5. A Pure Strategy Example

In this section we construct an example using pure strategies only, which has properties completely analogous to those of our previous example. For this purpose, it is only necessary to replace the mixed strategies that actually occur in the analysis of the previous example by explicit pure strategies. Specifically, player I has four pure strategies, which we will call A , B , C , and D (corresponding to L , $(\frac{1}{2}L, \frac{1}{2}R)$, $(\frac{1}{8}L, \frac{7}{8}R)$, and R respectively), and player II has three pure strategies, which we will call A , B , and C (corresponding to L , $(\frac{1}{2}L, \frac{1}{2}R)$, and R respectively). As before, chance chooses L or R with probabilities $\frac{1}{3}$ and $\frac{2}{3}$ respectively. If chance chose L , the payoffs are determined by II's choice, in accordance with the following table:

	A	B	C
L	9	$4\frac{1}{2}$	0

If chance chose R , the payoffs are determined by the choices of both players, in accordance with the following table:

	A	B	C
A	6	3	0
B	3	3	3
C	$\frac{3}{4}$	3	$5\frac{1}{4}$
D	0	3	6

Player I is informed as to chance's choice, but player II is not. The matrix of the game G is

	A	B	C
A	7	$3\frac{1}{2}$	0
B	5	$3\frac{1}{2}$	2
C	$3\frac{1}{2}$	$3\frac{1}{2}$	$3\frac{1}{2}$
D	3	$3\frac{1}{2}$	4

Player II has a unique pure optimal strategy, namely B . Player I also has a unique pure optimal strategy, namely to choose C when chance chooses R . If chance actually

does choose R , and if player II plays optimally, this yields a payoff of 3. But it does not *guarantee* a payoff of 3; player I could guarantee this amount by choosing B . However, B is not a maximin strategy in the original game.

In this case player I has a unique pure optimal strategy in G . But when it comes to actually making the only move that player I can make in the game, the arguments for playing according to this unique optimal pure strategy do not seem convincing.

6. A Game in Normal Form

We discovered an example similar to the example presented in §3 at the International Game Theory Workshop held in Jerusalem in 1965. At that time, Guillermo Owen pointed out that if our arguments are valid, then the minimax principle could be questioned even in the case of matrix games. Specifically, consider the 2-person zero-sum game H with matrix

	L	R
L	0	2
R	3	1

An optimal strategy for, say, player I, is $(\frac{1}{2}L, \frac{1}{2}R)$. So this player flips a coin; let us say that the coin flip dictates playing L . But now I still has to decide whether he should actually play L . *Before* the coin toss, the strategy $(\frac{1}{2}L, \frac{1}{2}R)$ maximized the utility payoff and, since this is a zero-sum game, it made good sense. *After* the toss, though, I has more information (namely, that the toss resulted in L), and playing L is no longer maximin. In fact, R yields a higher security level. We are not saying that we are recommending R ; we are simply saying that after the toss, the player still has to decide how to play, and he is essentially in the same quandary as before.⁹

As in the extensive games discussed in §§3, 4, and 5, this problem revolves around that of commitment. There are two possibilities. Either I can irrevocably commit himself, beforehand, to obeying the results of the toss, or he cannot. If he can—even if he doesn't have to—then the ordinary arguments in favor of the maximin strategy apply; the result is a recommendation to I to make a commitment to obey the results of a $\frac{1}{2}$ - $\frac{1}{2}$ toss. The question of whether then to obey the toss if it turns out L does not arise in this case, since the commitment to obey the toss is irrevocable. If, on the other hand, such an irrevocable prior commitment is impossible or forbidden, then the toss really becomes meaningless as soon as it is made, and it is difficult to know what to recommend.

At this point, it is perhaps in place to devote a few words to the phrase "irrevocable prior commitment". Such a commitment could involve giving instructions to a machine or an agent, or it could involve a contract to pay some large indemnity or otherwise incur a loss unless the toss is obeyed. Even a firm prior decision could come under the heading of "irrevocable prior commitment", if the player involved would suffer a sufficiently large loss of utility from the very act of changing his mind (i.e. not obeying the coin toss). Some people feel that if he wishes, a strong-willed player can practically always condition himself so that a violation of his own previous decision would involve a significant loss of utility. Others disagree, holding that there is no good reason for a player to take a prior decision into account after the toss, no matter how "firmly" it

⁹ See [L-R, p. 75] for another discussion of this problem.

was made beforehand; or at least that in important games, the disutility of reversing one's decision is not sufficiently great to deter such a reversal. Without necessarily taking a stand on this last question as far as a single person is concerned, we feel that there is surely no question in many situations in which a player consists of a group of persons (such as a government). Indeed, frequently there is no loss of face, and therefore no disutility, if such a group reverses its decision.

In any case, the precise forms that an "irrevocable prior commitment" can take do not really affect our discussion. What is important is only that when such a commitment is impossible, the use of the minimax strategy is questionable.

Note that this entire discussion applies almost word for word to the games described in §§3-5; it is only necessary to replace the toss made by the player by the choice made by chance. The upshot in both situations is that the minimax principle is unconvincing when irrevocable prior commitments cannot be made. Nevertheless, even in this case there are certain differences between the two situations, which make our attack still more convincing in the previous extensive games than it is in the matrix game considered in this section.

The first difference is as follows: Consider the situation S_R that arises if chance chooses R in one of the extensive games considered in the previous sections. If both players choose in S_R in accordance with strategies that are optimal in G , then I will receive 3. Such a choice by I, however, does not guarantee 3, *though he could guarantee 3 by playing a different strategy*. Playing in S_R in accordance with optimal strategies in G therefore seems particularly pointless. The game H considered in this section does not have the analogous property. Consider the situation U_L arising after the toss has dictated L . If both players choose in U_L in accordance with strategies that are optimal in H —i.e. if I plays L and II also tosses a coin (with appropriate probabilities) and obeys the toss, then I will receive a utility payoff $1\frac{1}{2}$ (because of the fact that II mixes). It is true that by playing L , I does not guarantee $1\frac{1}{2}$, but unlike the case with the extensive games, he now has *no choice*¹⁰ that will guarantee $1\frac{1}{2}$. The attack on the use of an optimal strategy is therefore not quite so telling here as in the extensive games.

The other difference concerns the notion of mixed strategies. This notion has often led to raised eyebrows; some people feel that it is a merely theoretical construct, and that in real life it is unthinkable to base an important decision on a coin flip. The example of this section might therefore be shrugged off by these people as being in any case irrelevant to real applications. The previous example is less easily shrugged off, because—as is shown in §5—it questions the minimax principle in a situation in which there is a pure-strategy solution to the game.¹¹

7. Concluding Remarks

We close this paper with a few remarks. First of all, there is nothing special or pathological about our example; we suspect that the phenomena exhibited therein are the rule rather than the exception.

Our second remark concerns the novelty of the ideas presented herein. Although they were new and quite disconcerting to us when we found them, it is quite possible that similar ideas have been "around" for some time. Specifically, people have thought about

¹⁰ That is, no pure choice. Note that since we are discussing the case in which a prior commitment is impossible, additional coin tosses will not help player I; eventually he has to choose a row, and at that moment he cannot guarantee $1\frac{1}{2}$.

¹¹ Note that a matrix game with a pure strategy saddle point cannot constitute such an example.

the problem of what to do when a player arrives, so to speak, "in the middle" of an extensive game. We do not know of any results in this direction, but we understand that this problem was considered during the 1950's. It is perhaps only a small step from these considerations to the remark that even if one arrives at the *beginning* of an extensive game, one eventually gets to the middle of it, and then one will have the same problems as if one had arrived in the middle. In any case, whether or not the difficulty that we have pointed out in this paper has been known, we feel fairly certain that it has never been satisfactorily resolved.

Our final remark is concerned with Harsanyi's Theory of *games with incomplete information* [H]. It should be recalled that many such games are converted in his theory to regular games similar to the game discussed in §3, where a chance move, so to speak, replaces the lack of information. Here, however, no commitment is possible before chance has made its choice, since in this theory the players are assumed to *enter* the game after chance has made its choice. In fact, the chance move is only a mathematical device, which never actually occurs in reality. Our criticism therefore applies to this case with full force, since an early commitment is totally impossible. Harsanyi himself already recognized that this time gap is crucial when cooperative games with incomplete information are being played. The present paper shows that difficulties due to the time gap exist even if the players are playing two-person zero-sum games with incomplete information.

References

- [H] HARSANYI, J. C., "Games with Incomplete Information Played by 'Bayesian' Players," Parts I, II, III, *Management Science* Vol. 14 (1967), pp. 159-182; Vol. 14 (1968), pp. 320-334; Vol. 14 (1968), pp. 486-502.
- [K] KUHN, H. W., "Extensive Games and the Problem of Information," in *Contributions to the Theory of Games II*, Princeton University Press, Princeton, N. J., 1953, pp. 193-216.
- [L-R] LUCE, R. D. AND RAIFFA, H., *Games and Decisions*, John Wiley and Sons, New York, 1957.
- [N-M] VON NEUMANN, J. AND MORGENSTERN, O., *Theory of Games and Economic Behavior*, Princeton University Press, Princeton, N. J., 1944, Third Edition, 1953.