Introduction to Game Theory

Peter Morris

Peter Morris Introduction to Game Theory Q

Peter Morris Introduction to Game Theory With 44 Illustrations Springer

Peter Morris Mathematics Department Penn State University University Park, PA 16802 USA Editorial Board (North America): S. Axler F.W. Gehring Mathematics Department Mathematics Department San Francisco State University East Hall San Francisco, C A 94132 University of Michigan USA Ann Arbor, MI 48109 USA K.A. Ribet Mathematics Department University of California at Berkeley Berkeley, CA 94720-3840 USA

To my mother, Opal Morris, and to the memory of my father, W. D. Morris

Preface The mathematical theory of games has as its purpose the analysis of a

wide range of competitive situations. These include most of the recreations which people usually call "games" such as chess, poker, bridge, backgammon, baseball, and so forth, but also contests between companies, military forces, and nations. For the purposes of developing the theory, all these competitive situations are called games. The analysis of games has two goals. First, there is the descriptive goal of understanding why the parties ("players") in competitive situations behave as they do. The second is the more practical goal of being able to advise the players of the game as to the best way to play. The first goal is especially relevant when the game is on a large scale, has many players, and has complicated rules. The economy and international politics are good examples. In the ideal, the pursuit of the second goal would allow us to describe to each player a strategy which guarantees that he or she does as well as possible. As we shall see, this goal is too ambitious. In many games, the phrase "as well as possible" is hard to define. In other games, it can be defined and there is a clear-cut "solution" (that is, best way of playing). Often, however, the computation involved in solving the game is impossible to carry out. (This is true of chess, for example.) Even when the game cannot be solved, however, game theory can often help players by yielding hints about how to play better. For example, poker is too difficult to solve,

viii Preface but analysis of various forms of simplified poker has cast light on how to play the real thing. Computer programs to play chess and other games can be written by considering a restricted version of the game in which a player can only see ahead a small number of moves. This book is intended as a text in a course in game theory at either the advanced undergraduate or graduate level. It is assumed that the students using it already know a little linear algebra and a little about finite probability theory. There are a few places where more advanced mathematics is needed in a proof. At these places, we have tried to make it clear that there is a gap, and point the reader who wishes to know more toward appropriate sources. The development of the subject is introductory in nature and there is no attempt whatsoever to be encyclopedic. Many interesting topics have been omitted, but it is hoped that what is here provides a foundation for further study. It is intended also that the subject be developed rigorously. The student is asked to understand some serious mathematical reasoning. There are many exercises which ask for proofs. It is also recognized, however, that this is an applied subject and so its computational aspects have not at all been ignored. There were a few

foreshadowings of game theory in the 1920's and 1930's in the research of von Neumann and Borel. Nevertheless, it is fair to say that the subject was born in 1944 with the publication of [vNM44]. The authors, John von Neumann and Oskar Morgenstern, were a mathematician and an economist, respectively. Their reason for writing that book was to analyze problems about how people behave in economic situations. In their words, these problems "have their origin in the attempts to find an exact description of the endeavor of the individual to obtain a maximum of utility, or, in the case of the entrepreneur, a maximum of profit." Since 1944, many other very talented researchers have contributed a great deal to game theory. Some of their work is mentioned at appropriate places in this book. In the years after its invention, game theory acquired a strange reputation among the general public. Many of the early researchers in the subject were supported in part or entirely by the U.S. Department of Defense. They worked on problems involving nuclear confrontation with the Soviet Union and wrote about these problems as if they were merely interesting complex games. The bloody realities of war were hardly mentioned. Thus, game theory was popularly identified with "war gaming" and was thought of as cold and inhuman. At the same time, the power of the theory was exaggerated. It was believed that it could solve problems which were, in fact, far too difficult for it (and for the computers of that time). Later, in reaction to this, there was a tendency to underestimate game theory. In truth, it is neither all-powerful nor a mere mathematician's toy without relevance to the real world. We hope that the usefulness of the theory will

Preface ix become apparent to the reader of this book. It is a pleasure to thank some people who helped in moving this book from a vague idea to a reality. Mary Cahill, who is a writer of books of a different sort, was unfailing in her interest and encouragement. Linda Letawa Schobert read part of the first draft and offered many suggestions which greatly improved the clarity of Chapter I. Ethel Wheland read the entire manuscript and corrected an amazing number of mistakes, inconsistencies, and infelicities of style. She is a great editor but, of course, all the errors that remain are the author's. Finally, thanks are due to several classes of students in Math 486 at Penn State who used early versions of this book as texts. Their reactions to the material had much influence on the final version.

Contents Preface vii List of Figures xv 1. Games in Extensive Form 1 1.1.

1 Games in Extensive Form

All the games that we consider in this book have certain things in common. These are: • There is a finite set of players (who may be people, groups of people, or more abstract entities like computer programs or "nature" or "the house"). • Each player has complete knowledge of the rules of the game. • At different points in the game, each player has a range of choices or moves. This set of choices is finite. • The game ends after a finite number of moves. • After the game ends, each player receives a numerical payoff. This number may be negative, in which case it is interpreted as a loss of the absolute value of the number. For example, in a game like chess the payoff for winning might be +1, for losing — 1, and for a draw 0. In addition, the following are properties which a game may or may not have: • There may

be chance moves. In a card game, the dealing of the hands is such a chance move. In chess there are no chance moves. • In some games, each player knows, at every point in the game, the entire previous history of the game. This is true of tic-tac- toe and backgammon but not of bridge (because the cards dealt to the other players are hidden). A game with this property is

2 1. Games in Extensive Form said to be of 'perfect information. Note that a game of perfect information may have chance moves. Backgammon is an example of this because a die is rolled at points in the game. We have just said that the players receive a numerical payoff at the end of the game. In real conflict situations, the payoff is often something non- quantitative like "happiness," "satisfaction," "prestige," or their opposites. In order to study games with such psychological payoffs, it is first necessary to replace these payoffs with numerical ones. For example, suppose that a player in a certain game can win one of three prizes: • A week in Paris. • A week in Hawaii. • Eight hours in a dentist's chair. Different people would assign different "happiness ratings" to these prizes. For a person with an interest in French culture, rating them as 100, 25, — 100, respectively, might be reasonable. To a surfer, the ratings might be 10, 100, —100. The point is that we are assuming that this conversion of nonquantitative payoffs to numerical ones can always be done in a sensible manner (at least in the games we consider). It is natural to represent a game by means of a "tree diagram." For example, consider the following simple game (called Matching Coins). There are two players (named Thelma and Louise); each conceals a coin (either a penny or a nickel) in her closed hand, without the other seeing it. They then open their hands and if the coins are the same, Thelma takes both of them. If they are different, Louise takes both. A tree diagram for Matching Coins is shown in Figure 1.1. The little circles in the diagram are called vertices, and the directed line segments between them are called edges. A play of the game starts at the top vertex (labeled "Thelma") and arrives, via one of the two vertices labeled "Louise," at one of the four vertices at the bottom. Each of these is labeled with an ordered pair of numbers which represent the payoffs (in cents) to Thelma and Louise, respectively. For example, if Thelma holds a penny, and Louise a nickel, then the game moves from Thelma's vertex by way of her left-hand edge, and leaves Louise's left-hand vertex by way of her right-hand edge. The bottom vertex reached is labeled (—1,1). This means that Thelma loses +1 ("wins" -1) and that Louise wins +1. In this chapter, we rigorously develop

this idea of tree diagrams. We also extend it to games with chance moves, and introduce a definition of strategy.

1.1. Trees 3 Thelma Louiseι/ \a •Γ "% Louise ∨ \5 ι / \5 (1,-1) (-U) (-5,5) (5^5) Figure 1.1. Matching Coins. 1.1. Trees A directed graph is a finite set of points, called vertices, together with a set of directed line segments, called edges, between some pairs of distinct vertices. We can draw a directed graph on the blackboard or on a piece of paper by drawing small circles for the vertices and lines with arrowheads for the edges. An example is given in Figure 1.2. It has five vertices and five edges. The theory of graphs (both directed and not directed) is a large and interesting subject. See [BM76] for more information. We will use uppercase letters like G or H or T to denote directed graphs. The vertices of such a graph will be denoted by lowercase letters like и от v. Subscripts will be used sometimes. Edges will be denoted by, for example, (u, v). Here, the edge goes from и to v. For a directed graph G, the set of all vertices is denoted V(G). A path from a vertex и to a vertex ν in a directed graph G is a finite sequence (^o, v\,..., vn) of vertices of G, where η > 1, vo = u, vn — v, and (vi-i,Vi) is an edge in G for г = 1, 2,..., η. For example, in the directed graph of Figure 1.2, (e,c,a) and (e,c, 6, e) are paths. A tree is a special kind of directed graph. We have the following: DEFINITION 1.1. A directed graph T is a tree if it has a distinguished vertex r, called the root, such that r has no edges going into it and such that for every other vertex ν of T there is a unique path from r to v. The example in Figure 1.2 is clearly not a tree. A directed graph consisting of a single vertex and no edges is a (trivial) tree.

4 1. Games in Extensive Form Figure 1.2. A directed graph. Usually, but not always, we draw trees in such a way that the root is at the top. An example is given in Figure 1.3; all the vertices have been labeled. This tree has eleven vertices and ten edges. Now let T be a tree. We make a few definitions. A vertex ν is a child of a vertex и if (tt, v) is an edge. Also, in this case, и is the parent of v. In the tree of Figure 1.3, vertex / is a child of 6, vertex g is a child of c, a is the parent of both d and e, and the root is the parent of a, 6, and c. In any tree, the set of children of и is denoted Ch(tt). Notice that a vertex may have many children. The root is the only vertex without a parent. A vertex without any children is called terminal. In our example, the terminal vertices are d, e, f, i, j, and ft. A vertex which is not the root and not terminal is called

intermediate. In our example, the intermediate vertices are a, 6, c, and g. A vertex v is a descendant of a vertex и if there is a path from и to v. In this case, и is an ancestor of v. Thus the root is an ancestor of every other vertex. In the example, c is an ancestor of j, and i is a descendant of c. The length of a path is the number of edges in it. Paths in trees do not cross themselves; that is, they consist of distinct vertices. This follows from Exercise (8). Thus, a path in a tree has length at most the number of vertices minus one. The depth of a tree is the length of the longest path in it. In Figure 1.3, (c, g, i) is a path of length 2, as is (root, 6, /). It is clear that the longest path in a tree starts at the root and ends at a terminal

1.1. Trees 5 Figure 1.3. A tree. vertex. If this were not true, we could extend the given path to a longer one. The depth of a tree T is denoted De(T). The depth of the tree in Figure 1.3 is three because (root, c, g, j) is a path of greatest length. Some general facts about trees are collected together in the following: Theorem 1.1. Lei T be a tree. Then we have: (1) No vertex has more than one parent. (2) If и and v are vertices of T and there is a path from и to v, then there is no path from v to u. (3) Every nonterminal vertex has a terminal descendant. Proof. (1) Suppose that w is a vertex with two parents и and v. By definition of a tree, there is a path from the root to u. Appending the edge (it, w) to this path produces a path from the root to w. Similarly, appending the edge (v, w) to the path from the root to v produces another path from the root to w. These two paths are not the same because the last edge in one of them is (u,w) and the last edge in the other is (v,w). This contradicts the definition of a tree. Thus, w cannot have two parents. (2) Suppose there is a path from v to u. Now start with the path from the root to и and append to it first the path from и to v and then the path from v back to u. This process produces a second path from the root to u. This contradicts the definition of a tree. Thus, there is no path from v to u. (3) Let и be a nonterminal vertex. Consider a path of greatest length starting at u. The vertex at the end of this path is terminal be-

6 1. Games in Extensive Form cause, if it were not, the path could be extended. This would contradict the fact that the path chosen has greatest length. П Now let T be a tree and let и be any vertex of T. The cutting of T determined by it, denoted TU1 is defined as follows: The vertices of Tu are и itself plus all of the descendants of u. The edges of Tu are all edges of T which start at a vertex

of Tu. Note that if и is the root of T, then Tu — T, and if и is terminal, then Tu is a trivial tree. The cutting Tc of the tree in Figure 1.3 has vertices c, g, ft, i, j, and edges (c,#), (#,г), (gj), (c, ft). Theorem 1.2. For any tree T and any vertex u, Tu is a tree with и as the root. Proof. First, и has no edges of Tu going into it because the starting vertex of such an edge could not be a descendant of u. Second, let v be a vertex of Tu. Then v is a descendant of и and so there is a path in Tu from и to v. The fact that there is only one such path follows from Exercise (14). П Let T be a tree and и a vertex of T. The quotient tree T/u is defined as follows: The vertices of T/u are the vertices of T with the descendants of и removed; the edges оïT/u are the edges of T which start and end at vertices of T/u. Thus, и is a terminal vertex оiT/u. For example, in the tree of Figure 1.3, the quotient tree T/a is obtained by erasing the edges (a, d) and (a, e), and the vertices d and e. Notice that if и is the root of T, then T/u is trivial, and if и is terminal, then T/u is T. The proof of the following should be obvious and is omitted. Theorem 1.3. If T is a tree and и is a vertex ofT, then T/u is a tree whose root is the root ofT. Finally, if T is a tree, then a subtree 5 of T is a tree whose vertices form a subset of the vertices of T; whose edges form a subset of the edges of T; whose root is the root of T; and whose terminal vertices form a subset of the terminal vertices ofT. In the tree of Figure 1.3, the tree consisting of the root together with vertices c, ft, g, ji, and edges (root,c), (c, ft), (c, #), (g,j) is a subtree. The following theorem gives us an alternative definition of subtree. The proof is left as an exercise. Theorem 1.4. Let S be a subtree of a tree T. Then S is the union of all the paths from the root to a terminal vertex of S. Conversely, if U is any nonempty subset of the set of all terminal vertices of the tree T, then the union of all paths from the root to a member of U is a subtree. The set of terminal vertices of this subtree is precisely U.

1.2. Game Trees 7 For example, in the tree of Figure 1.3, consider the set of terminal vertices U = {d,f,j}. Then the subtree determined by U is the union of the three paths (root, a, d), (root, 6,/), (root, c,$r,j). Exercises (1) Sketch all cuttings of the tree in Figure 1.4. (2) Sketch all quotient trees of the tree in Figure 1.4. (3) How many subtrees does the tree in Figure 1.4 have? (4) Sketch all cuttings of the tree in Figure 1.3. (5) Sketch all quotient trees of the tree in Figure 1.3. (6) How many subtrees does the tree in Figure 1.3 have? (7) Prove that a tree has only one root. (8) Define a. cycle in a. directed graph to be a path which begins and ends at the same vertex. Prove that trees do not contain

cycles. Then prove that paths in trees consist of distinct vertices. (9) Let G be a directed graph. For a vertex и of G, define p(u) to be the number of edges starting at и minus the number of edges ending at u. Prove that uev(G) (10) Let D be a directed graph. Prove that the number of vertices и of D with p(u) odd is even. [See Exercise (9) for the definition of p(u).] (11) Let T be a tree; let б be the number of edges in T; and let ν be the number of vertices. Prove that б = ν — 1. (12) LetTbea nontrivial tree. Prove that there is a nonterminal vertex such that all of its children are terminal. (13) Let T be a tree. Let W be a set of vertices of T such that every terminal vertex is in W and such that W contains a vertex whenever it contains all the children of that vertex. Prove that W contains all vertices. (14) Let T be a tree and it and ν vertices of T such that ν is a descendant of u. Prove that there is only one path from и to v. (15) Prove Theorem 1.4. (16) Define a vertex w in a directed graph G to have the unique path property if, for any other vertex и in G, there is a unique path from w to u. Give an example of a directed graph with two vertices with the unique path property. 1.2. Game Trees Let T be a nontrivial tree. We wish to use T to define an TV-player game without chance moves. First, let Pi, P2,..., Pn designate the players. Now

8 1. Games in Extensive Form Figure 1.4. Tree for Exercises (1), (2),and (3) of Section 1.1. label each nonterminal vertex with one of these designations. We will say that a vertex labeled with Pг· belongs to Pг· or that Pг· owns that vertex. Then label each terminal vertex ν with an TV-tuple of numbers p(v). The game is now defined. It is played as follows. The player who owns the root chooses one of the children of the root. If that child is intermediate, the player to whom it belongs chooses one of its children. The game continues in this way until a terminal vertex ν is reached. The players then receive payoffs in accordance with the TV-tuple labeling this final vertex. That is, player Pг· receives payoff pг(и), component number г of p(v). Figure 1.5 is an example with three players. In this game, there are never more than three edges out of a vertex. It is natural to designate these edges as L (for "left"), R (for "right"), and M (for "middle"). For example, if P^s first move is to vertex 6, then we would designate that edge (and move) by M. If P3 then responds by moving to g, the history of that play of the game could be designated (M, R). As a second example, another possible history is (L, L, R). In this case, the payoff vector is (0, —2, 2). A tree labeled in the way just described (using designations of players and payoff vectors) is a game tree, and the corresponding game is a

tree game. It is important to realize that the kind of game we have set up has no chance moves. In a later section, we will see how to include the idea of vertices "belonging to chance" in our game trees. When that has been done, we will be able to think of all games as tree games. Notice that every path from the root to a terminal vertex represents a different way in which the history of the game might evolve.

1.2. Game Trees 9 Figure 1.5. A three-player game tree. It is not hard to see how a game tree might, in theory, be set up even for such a complicated game as chess. The player moving first would own the root. Each of her legal opening moves would correspond to a child of the root. Each of these children would belong to the player moving second and would have a child corresponding to each legal move for that player, etc. A terminal vertex would correspond to a situation in which the game is over such as a checkmate or a draw. Clearly, such a tree is enormous. Let us consider a simpler example. Example 1.1 (Two-Finger Morra). This game is played by two people as follows. Both players simultaneously hold up either one finger or two and, at the same time, predict the number held up by the other player by saying "one" or "two." If one player is correct in his prediction while the other player is wrong in hers, then the one who is right wins from the other an amount of money equal to the total number of fingers held up by both players. If neither is right or both are right, neither wins anything. The difficulty in representing this game as a tree is the simultaneity of the players' moves. In a game defined by a tree the players obviously move consecutively. To remove this apparent difficulty, suppose that a neutral third party is present. Player Pi whispers her move (that is, the number of fingers she wishes to hold up and her prediction) to the neutral person. Player P2 then does the same. The neutral person then announces the result. Now the players' moves are consecutive. We see that Pi has a choice of four moves. Each can be represented by an ordered pair $(/, p)$ where / is the number of fingers held up and ρ is

10 1. Games in Extensive Form ($°>°$) P2 (1,1) PI (2,2) P2 Figure 1.6. Two-finger morra. the prediction. Thus each of / and ρ is either 1 or 2. We see from this that the root of our tree, which is labeled by Pi, has four children. These are all labeled by P2 and each of them has four children (one for each of P25s possible moves). These children are all terminal. The tree is shown in Figure 1.6. (Most of the payoff pairs have been left out.) This tree for two-finger

morra does not embody all the rules of the game. It appears from the tree that P2 could guarantee herself a positive payoff by making the correct move at any of her four vertices. This is certainly not true—the two players have symmetric roles in the game, and so if P2 could guarantee herself a positive payoff, then so could Pi. But they cannot both win. The problem is that P2 does not know Pi's move and thus does not know which of her vertices she is moving from. In other words, the game is not of perfect information. This concept was discussed briefly earlier, and will be discussed in more detail later. Notice that if T is a game tree, then any cutting of T is also a game tree. A quotient tree T/u of T becomes a game tree if an TV-tuple of numbers is assigned to the vertex u. For example, in Figure 1.5, the cutting Ta is a game in which P3 has no vertices but receives payoffs. Also, T/a is a game if we assign, arbitrarily, the payoff (0,1,1) to a. A subtree of T is also a game tree. These constructions of new game trees from old ones are often useful in studying parts of games (such as end-games in chess) and restricted versions of games.

1.2. Game Trees 11 1.2.1. Information Sets Let us go back and look at the game of Matching Coins (Figure 1.1). The rules of this game state that Louise does not know which coin Thelma is holding. This implies, of course, that the game is not of perfect information. In terms of the tree, it also means that Louise does not know which of her two vertices she is at. It is therefore impossible, under the rules, for her to plan her move (that is, decide which coin to hold) based on which of these two vertices she has reached. The set consisting of her two vertices is called an information set. More generally, an information set S for player P is a set of vertices, all belonging to P, such that, at a certain point in the game, P knows he is at one of the vertices in S but does not know which one. We mentioned earlier that the tree in Figure 1.6 does not embody all the rules of two-finger morra. The reason is that the tree does not take account of the information set for P2. This information set consists of all four of P2's vertices. Exercises (1) For the game in Figure 1.5, sketch the cutting Ta. What is P2 's best move in (2) For the game in Figure 1.5, sketch the quotient game T/α [with (0,1,1) assigned to u]. What is Pi's best move? (3) Concerning the game pictured in Figure 1.5, answer the following questions. • What is the amount player P3 is guaranteed to win, assuming that all players play rationally? • What choice would you advise P\ to make on his first move? • If the rules of the game allow P\ to offer a bribe to another player, how much should he offer to whom for doing what? (4) For the game tree in Figure 1.7,

the three vertices owned by player B form an information set (enclosed by a dashed line). How should A and B play? How much can they expect to win? (5) Suppose that, in the game of Figure 1.5, the rules are such that P2 does not know whether P\ has moved L or R. Sketch the tree with Pг'в information set indicated. How should P2 move? (6) A very simple version of the game of nim is played as follows: There are two players, and, at the start, two piles on the table in front of them, each containing two matches. In turn, the players take any (positive) number of matches from one of the piles. The player taking the last match loses. Sketch a game tree. Show that the second player has a sure win. (7) A slightly less simple version of nim is played as follows: There are two players, and, at the start, three piles on the table in front of them, each containing two matches. In turn, the players take any (positive) number of matches from one of the piles. The player taking the last match loses. Sketch a game tree. Show that the first player has a sure win.

12 1. Gaines in Extensive Form (2,2) (2,0) Figure 1.7. Game tree for Exercise (4) of Section 1.2. 1.3. Choice Functions and Strategies We now want to make a precise definition of a "strategy" for a player in a game. Webster's New Collegiate Dictionary defines this word as "a careful plan or method." This is the general idea, but we need to make the idea rigorous. There are three requirements which our definition of this concept must satisfy. The first is that it should be complete: A strategy should specify which move is to be made in every possible game situation. The second is that it should be definite: The move to be made in a given game situation should be determined by that game situation, and not by chance or by the whim of the player. It may help to think of a strategy as a description of how to play which can be implemented as a computer program. The third requirement is that information sets must be respected. This means that a strategy must call for the same move at every vertex in an information set. For example, in two-finger morra, a strategy for P^ must call for the same move from each of her four vertices. The following defines a concept which is our first approximation to what a strategy should be. Definition 1.2. Let T be a game tree and let P be one of the players. Define a choice function for P to be a function c, defined on the set of all vertices of T belonging to P, which is such that c(u) is a child of и for

1.3. Choice Functions and Strategies 13 every vertex и belonging to P. Thus, given that player P is playing according to a choice function c, P knows which

choice to make if the play of the game has reached vertex u (owned by P): She would choose $c(u)$. Now suppose that there are N players $P_1, P_2, \cdots, P_n$. Let us denote the set of all choice functions for player P_i by Γ_i. Given that each player P_i moves according to his choice function $c_i \in \Gamma_i$, for $1 < i < N$, a path through the game tree from the root to a terminal vertex is determined. Define $\pi_i(c_1, c_2, ..., c_N)$ to be component number i of the TV-tuple $p(w)$ which labels the terminal vertex w at which this path ends. Thus, this quantity is the payoff to P_i when the players play according to the N-tuple of choice functions $(c_1, c_2, ..., c_N)$. For example, in the game of Figure 1.5, suppose that P_1, P_2, and P_3 use the choice functions c_1, c_2, and c_3, respectively, where $c_1(root) = a$, $c_1(d) = j$, $c_1(e) — l$; $c_2(a) = e$, $c_2(c) — i$; $c_3(b) — f$. The path traced out is $(root, a, e, l)$, which terminates at the payoff vector $(0,1,1)$. We can probably agree that for any strategy there is a choice function which embodies it. However, there are reasons why choice function is not an appropriate definition of strategy. The first is that choice functions are usually defined at many vertices where they need not be defined, that is, at vertices which can never be reached in the course of the game, given earlier decisions by the player. For example, in the game shown in Figure 1.5, a choice function for player P_1 which calls for him to move to the middle or right-hand child from the root obviously need not specify moves from either of the other vertices belonging to P_1 since they cannot be reached. Nevertheless the definition of a choice function requires that its domain of definition should consist of all vertices belonging to a given player. 1.3.1. Choice Subtrees Let us try to eliminate the problem mentioned above. There are several ways in which this might be done. We could allow choice functions to be defined on subsets of the set of vertices belonging to a given player. This is unsatisfactory because the domain of definition would vary with the choice function. The following gives a better solution: Definition 1.3. Let T be a game tree and let P be one of the players. Let c be a choice function for P. Define a (P,c)-path to be a path from the root of T to a terminal vertex such that if u is a vertex of that path belonging to P, then the edge $(u, c(u))$ is an edge of the path. Thus, a (P, c)-path represents one of the possible histories of the game given that player P plays according to the choice function c. For example,

14 1. Gaines in Extensive Form in the game of Figure 1.5, suppose that P_1 uses the choice function c_1, where $c_1(root) = a$, $c_1(d) — k$, $c_1(e) = l$. Then the path $(root, a, e, l)$ is a (P_1, c_1)-path. So is $(root, a, d, k)$. In fact, these are the only

two. Then we have the following: Definition 1.4. Let T be a game tree, P a player, and c a choice function for P. Then the choice subtree determined by P and c is defined to be the union of all the (P, c)-path&. Thus, the choice subtree determined by P and c has as its set of vertices all the vertices of T which can be reached in the course of the game, given that P plays according to the choice function c. It is clear that a choice subtree is a subtree (by Theorem 1.4). The reader might now look at Figure 1.5 to see what the possible choice subtrees are. For example, one of them contains the vertices: root, a, e, d, fc, and /. Its edges are (root, a), (a,d), (a,e), (d, fc), and (e,/). It is interesting that if и is a vertex of a choice subtree which belongs to player P, then и has only one child in the subtree. On the other hand, a vertex v of a choice subtree which belongs to a player different from P has all its children in the subtree. (Both of these facts are proved in the next theorem.) We see that all the essential information about the choice function c is contained in the choice subtree. That is, if one knows the subtree one can play in accordance with the choice function. Suppose, for example, that player P is playing according to a choice subtree S. If the game has reached a vertex и belonging to P, then и is in S and only one of its children is in S. P's move is the edge from и to that child. On the other hand, the inessential information about the choice function (namely, how it is defined at vertices which can never be reached) is not contained in the choice subtree. For example, in the tree of Figure 1.5, consider two choice functions for P\: c_i(root) = 6, c_i(d) — fc, d(e) = /; and c'i(root) = 6, c[(d) = jf, c[(e) — I. These clearly embody the same strategy since vertex d is never reached. But the choice subtree determined by Pi and c\ is the same as the choice subtree determined by Pi and c[. Not all subtrees are choice subtrees. For example, in the game of Figure 1.5, the subtree whose terminal vertices are jf, fc, ft, and г is not a choice subtree for any of the players. (If it were a choice subtree for Pi, it would not contain both j and fc; if for P2, it could not contain both ft and i; and, if for P3, it would contain either / or g.) The two properties stated above for choice subtrees (concerning children of vertices) are enough to characterize them among all subtrees. This fact is contained in the following theorem. The advantage of this characterization is that we can recognize a subtree as a choice subtree without having to work with choice functions at all. A lemma is needed first.

1.3. Choice Functions and Strategies 15 Lemma 1.5. Let T be a game tree; let P be a player; and let c be a choice function for P. Suppose that Q is a path from

the root to a vertex ν and that Q satisfies the condition: (*) If w is a vertex on Q owned by P, then either w — ν or c(w) is on Q. Then Q can be extended to a (P, c)-path. Proof. Let Q' be a path of greatest length containing Q and satisfying (*). Let ζ be the last vertex on Q' (that is, the vertex furthest from the root). If ζ is not terminal, then either it belongs to P or it does not. If it does, let Q'' = Q'U(z,c(z)). If it does not, let y be any child of ζ and let Q'' = Q'U(z,y). Then Q'' satisfies (*) and is longer than Q'. This is a contradiction and so ζ is terminal. Hence, Q1 is a (P, c)-path. D Theorem 1.6. Let T be a game tree and let P be one of its players. A subtree S of T is a choice subtree determined by P and some choice function c if and only if both the following hold: (1) If и is a vertex in S and и belongs to P, then exactly one of the children of и is in S. (2) // и is a vertex in S and и does not belong to P, then all the children of и are in S. Proof. First, suppose 5 is a choice subtree determined by P and c. Let и be a vertex in S. If и belongs to P, then, by definition of choice subtree, there is a (P, c)-path containing и and contained in S. By definition of (P, c)-path, (it, c(u)) is an edge of this path and so c(u) is in S. Now if ν is a child of и different from c(w), then no (P, c)-path contains ν and thus ν is not in S. Hence (1) holds. Now if и does not belong to P, then let ν be any child of и and let Q be the path from the root to ν. The part of Q which ends at и is part of a (P, c)-path (since и is in S). Thus, Q satisfies (*) of Lemma 1.5 and so there is a (P, c)-path containing Q. Hence, ν is in S. For the other half of the proof, suppose that 5 is a subtree of T which satisfies both (1) and (2). We must define a choice function c for which S is the choice subtree determined by P and c. If ν is a vertex of S belonging to P, define c(v) to be that child of ν which is in S. If ν belongs to P but is not in 5, let c(v) be any child of ν. Now we prove that this works. Let Q be a (P, c)-path. To show that Q is entirely contained in 5, let и be the last vertex in Q which is in S. If и is terminal, then the entire path Q is in

16 1. Gaines in Extensive Form S. Suppose then that и is not terminal. There are two cases: If и belongs to P, then c(u) is in S and the edge (it, c(u)) is in Q. This contradicts the fact that и is the last vertex. This leaves only the case where и does not belong to P. But then if (tt, v) is the edge from и which is in <2, then we again reach a contradiction since, by (2), ν is in S. We have now shown that S contains the choice subtree determined by P and c. To show the opposite inclusion, let w be a vertex of S. Then, since 5 is a subtree, the path from the root to w is contained in S. By (1) and the way in which c was constructed, this path satisfies condition (*) of Lemma 1.5. Hence there is a (P,

c)-path containing w, and thus w is in S. ☐ We now replace choice functions by choice subtrees. We will modify our earlier notation and write $7Tj(Si,... , Sat)$ for the payoff to player Pi- resulting from the TV-tuple of choice subtrees $(Si,..., Sn)$- Now, not every choice subtree respects information sets. Thus, we cannot simply define a strategy to be a choice subtree. Instead, we define a strategy for a player P to be a member of a subset of the set of all choice subtrees determined by P and P's choice functions. In turn, the members of this subset are the choice subtrees which respect P's information sets. We have the following: Definition 1.5. Let T be a game tree with N players Pi,P2,...,JV. A game in extensive form based on T consists of T together with a nonempty set $\Sigma,-$ of choice subtrees for each player P,-. The set $\Sigma,-$ is called the strategy set for P,- and a member of $\Sigma,-$ is called a strategy for P,·. To summarize, it is this concept of strategy sets that allows us to impose the rules of the game which prohibit the players from having perfect information. For example, if P,· is not allowed to know which of a set of vertices he has reached at a certain point in the game, then Pj's strategy set would not contain choice subtrees which call for different moves depending on this unavailable information. We will use an uppercase Greek letter Γ, Δ, or Λ to denote a game in extensive form. Formally, we would write, for example, $\Gamma = (T,\{P\iota,P2,...,P,\nu\},\{\Sigma\iota,\Sigma2,...,\Sigma,\nu\})$ to describe the game Γ with game tree T, and with players Pi, P2,..., Pn having strategy sets $\Sigma\iota$, $\Sigma2$,..., $\Sigma\lambda\tau$, respectively. Each choice subtree (and thus each strategy) is a subtree. Thus a choice subtree is itself a game tree. It represents a smaller game in which one of the players has effectively been removed (since that player has no choices).

1.3. Choice Functions and Strategies 17 Finally, notice that if a choice subtree Si is chosen from each $E\Gamma$-, then the intersection of the Si's is a path from the root to a terminal vertex. This path represents the history of the game, given that each player moves in accordance with the appropriate Si. The TV-tuple of numbers labeling the terminal vertex at which this path ends gives us the payoffs to the N players. We write $7T;(Si, S2,..., Sn)$ for component number Γ of that N- tuple. This number is thus the payoff to player P,· resulting from all players playing according to the given choice subtrees. Let us consider an example. Its game tree is given in Figure 1.8 and it is a game of perfect information. There are two players, designated A and P. Player A has twelve choice functions and player P has four. To count the choice functions for A, note that there are three vertices belonging to A. One of them (the root) has

three children while each of the other two has two children. Thus, the number of choice functions is 2 χ 2 χ 3 = 12. There are six choice subtrees for player A. Three of them are shown in Figure 1.9. To enumerate them, let M denote the choice subtree for A corresponding to A's choice of her middle child (at the root), and let R denote the one corresponding to the right-hand child. If A chooses her left-hand child at the root, there are four possible choice subtrees. Each of these has a name of the form Lxy, where χ and y are each either L or R. In this notation, χ denotes the child chosen by A in case B moves left, and y is the child chosen by A in response to P's moving right. Thus these four choice subtrees are LLL, LLR, LRL, and LRR. For P, there are four choice subtrees. We denote them xy, where each of χ and y is either L or R. Here, χ is the response of P to A's choosing her left-hand child, and y is the response of P to A's choosing her middle child. One of these choice subtrees for P is shown in Figure 1.10. For an example, suppose that A plays according to LRR, while P plays according to LL. Then the payoff to A is 2, while the payoff to P is —1. Now suppose that the rules of the game are modified so that, after A's first move, P cannot distinguish between A's moving left and A's moving middle. Also suppose that if A moves left on her first move, then A cannot distinguish between P's moving left and P's moving right. Then the strategy set for A consists of LLL, LRR, M, and R. Also, the strategy set for P consists of the two choice subtrees, LL and RR. Exercises (1) Describe all the strategies for each player for the version of nim given in Exercise (6) of Section 1.2. (2) Describe all the strategies for each player in the version of nim given in Exercise (7) of Section 1.2.

18 1. Gaines in Extensive Form (1,0) t N (-5,2) (-6,3) (4Д) (2,-1) (3,-2) Figure 1.8. A game tree. A A /\ (-5,2) (-6,3) (3,4) (3,4) (1,0) (4,1) Figure 1.9. Some choice subtrees for A.

1.3. Choice Functions and Strategies 19 Figure 1.10. A choice subtree for B. (3) Consider the game of perfect information shown in Figure 1.8. What strategies would you advise A and B to adopt? (4) Suppose that, in the game of Figure 1.8, if A moves left at the root, she cannot distinguish between B's moving left and B's moving right. List the strategies for both players. (5) In the game shown in Figure 1.8, suppose that the rules are such that, after A's first move, B cannot distinguish between A's moving left and A's moving middle. Also suppose that if A moves left on her first move, then A cannot distinguish

between B's moving left and B's moving right. What strategies would you advise A and B to adopt? (6) Sketch the choice subtrees for both players for the game in Figure 1.7. Which of these are strategies? (That is, which respect the information set?) (7) In the game of Figure 1.5, how many choice subtrees are there for each player? (8) Let T be a game tree with N players. Let P be one of the players and let c be a choice function for P. Define a vertex v of T to have the (P, c)-property if the path R from the root to v satisfies condition (*) of Lemma 1.5. Prove that the choice subtree determined by P and c is equal to the set of all vertices with the (P, c)-property. (9) Let T be a game tree and let S be a choice subtree for player P. Let Tu be a cutting of T. Prove that either S Γ)Tи = 0 or S Π Tu is a choice subtree for P in Tu. (10) Let r = (T, {Plf...fPN}f{Elf...fEN}) be a game in extensive form. Define a vertex v of T to be reachable by a player Pi if there exists a strategy for Pt which contains v. A vertex v is simply reachable if it is reachable by every player. Prove the following statements, (a) The root is reachable.

20 1. Games in Extensive Form (b) If v is reachable, then its parent (if v is not the root) is reachable. (c) If v is not reachable, then none of its children is reachable. (d) If v is not terminal and no child of v is reachable then v is not reachable. (e) Define T* to be the set of all reachable vertices of T together with all edges of T which go from a reachable vertex to a reachable vertex. Then T* is a subtree of T. 1.4. Games with Chance Moves We now want to modify the definition given in the previous section so as to allow for chance moves in the game. This is done by allowing some of the nonterminal vertices to be labeled C for "chance" instead of by the designation of a player. The meaning of such a vertex is that, at that point in the game, there is a random choice of one of the children of that vertex. The probabilities with which the various children will be chosen are non- negative numbers summing up to 1. These probabilities are used as labels on the edges coming out of the vertex labeled "chance." We introduce the following notation: If и is a vertex belonging to chance and v is a child of it, then Pr(tt, v) denotes the probability that the edge (it, v) will be chosen. Thus, letting E(u) denote the set of all edges starting at it, $Pr(«, υ) > 0$ and $(u,v)£E(u)$ For example, suppose that the game begins with a roll of two fair dice, and that the course of the rest of the game then depends on the sum of the numbers on top of the dice. Thus, there are eleven possible outcomes: 2, 3,..., 12, and so the root of the tree has eleven children. The probabilities on the edges coming out of the root vary from 1/36

(for outcomes 2 and 12) to 1/6 (for outcome 7). The definition of choice function and of choice subtree is exactly the same as before. The definition of payoff $\pi, \cdot (S_i, S_2, \ldots, S_n)$ for player P_i no longer works, however, because the terminal vertex where the game ends is now dependent on the chance moves as well as on the choices made by the players. Suppose that each player P_i plays according to choice subtree S_i. In the case of games without chance moves, the intersection of the S;'s is a single path from the root to a terminal vertex. In the case where there are chance moves, this intersection is a subtree which may contain many terminal vertices. The correct picture is that every path in this subtree branches whenever it runs into a vertex belonging to chance. Thus, every terminal vertex и in this intersection is reached with a probability which is the product of probabilities at the vertices belonging to chance encountered

1.4. Gaines with Chance Moves 21 on the path from the root to u. The expected payoff to player P_i would then be a weighted average of ith components of the TV-tuples labeling the terminal vertices of the intersection. Our point of view here is that the game is played repeatedly, perhaps many thousands of times. It is clear that in a game which is only played once (or just a few times) the idea of expected payoff may have little meaning. We are making a rather important distinction since some games, by their nature, can only be played once. Consider, for example, games of nuclear confrontation in which payoffs may correspond to annihilation. Definition 1.6. Let T be a tree for a game with N players. For each г, let S_i G Ej. If w is a terminal vertex of H^S; and if R is the path from the root to w, then the probability of the game terminating at w is $Pr(S_i, \ldots, S_n \backslash w) = TT\{Pr(tt, v) : и$ belongs to chance and (it, v) G R$\}>$ Here, it is understood that if there are no chance vertices on the path R, then $Pr(S_i, \ldots, S_n]$ w) = 1. The payoff can then be defined as follows: Definition 1.7. In the setting of the previous definition, the expected Vay°ffi ^-$(S_i, S_2, \ldots, S_n)$ to player P_i resulting from the choice subtrees S_i G $\Sigma2\cdot$ is defined by $7T_i(S_i, \ldots, S_N) = {}^{\wedge}Pr(S_i, \ldots, S_N; w)p_i(w)$, where the sum is taken over all terminal vertices w in njLxSi. Let us do some calculating for the game shown in Figure 1.11. There are two players, designated A and Б, and the game is of perfect information. There are two vertices belonging to chance. The left-hand vertex has two children, and the probabilities attached to the two edges are 0.6 and 0.4, respectively. One might imagine that a biased coin is tossed at that point in the game. The right-hand chance vertex is similar except that the coin tossed is fair. Player A has only two strategies: He can either move left (denoted L) or

right (denoted R). The reader is invited to verify that player B has sixteen strategies. Now suppose that player A plays his strategy L, while player B plays her strategy LRLR. The meaning of this notation is that if A moves left and the chance move is then to the left child, B moves left. If A moves left and the chance move is then to the right, B moves right. Finally, if A moves right, then B moves left or right depending on whether the chance move is to the left or right. To compute the expected payoffs for both players, note first that the terminal vertices which can be reached, given that the

22 1. Gaines in Extensive Form Figure 1.11. A game with chance moves. players play according to these strategies, are the ones whose payoff pairs are (1, —2) and (—2,1). Thus, the expected payoff for player A is тгл(Ь, LRLR) = 0.6x1 + 0.4 χ -2 = -0.2, while the expected payoff for player B is 7гБ(1,LRLR) = 0.6 χ (-2) + 0.4x1 = -0.8. 1.4.1. A Theorem on Payoffs We now wish to prove a theorem which allows us to compute payoffs due to choice subtrees in terms of payoffs in smaller games. A lemma is needed first. Lemma 1.7. Lei T be a tree game with N players; let Si,...,Siv be choice subtrees for Pi,...,P/\r; respectively; and let r be the root of T. Then we have: (1) If r belongs to a player Pi and (r, u) is an edge of Si, then, for $1 < 3 < N$, Sj ПТи is a choice subtree ofTu for player Pj. (2) Ifr belongs to chance and и is any child ofr, then, for $1 < j < N$, Sj C\TU is a choice subtree ofTu for player Pj. Proof. By Exercise (9) of Section 1.3, it suffices for the proof of both (1) and (2) to show that Sj ПТи ф$. This is true in both cases because и £ Sj for l<j<N. П

1.4. Gaines with Chance Moves 23 Theorem 1.8. Let T be a tree game with N players Pι,...,P/v. Let S\,..., Sn be choice subtrees for P\,..., Pχ, respectively. Then, with r as the root ofT, we have: (1) Ifr belongs to the player P{ and (r, u) is in Si then, for $1 < j < N$, 7Tj(5i,52,... ,SV) = Kj(Si CiTUlS2 CiTU)... ,Sn C\Tu). (2) Ifr belongs to chance then, for $1 < j < N$, ^(5ι>52>...,5λγ)= Σ Pr(r,«)^(51ПГ‖>...>5лгПГ‖). u£Ch(r) Proof. To prove (1), note that each Sk C\TU is a choice subtree of Tu. This follows from Lemma 1.7. Then, notice that n%=1Sk is n%=1(Sk ПТи) together with the edge (r, u). Thus, the subtree C\^=1Sk of T and the subtree C\%=1(Sk f\Tu) of Tu have the same terminal vertices. Also the probability of the game terminating at each of these terminal vertices is the same in both game trees. The result follows from the definition of expected payoff. To prove (2), note that Lemma 1.7 implies that Sk (~}TU is a choice subtree of Tu for each к and each child и of r. Note also that Pr(5b...

>5iV;ti;) = Pr(r>ti)Pr(5inrtl>... ,SNnTu-w), for any terminal vertex w £ C\^=1(Sk ClTu). Then, writing Wj for π/(Si,. . .,Sat) and summing over all terminal vertices w £ Hj^Sfc, we have ттγ = Σ Pr(51' · · · ' Sn> w)Pj(w) Σ Σ Vt(Si,...,SN;w)pj(w) uGCh(r) wen(SknTu) = Σ Pr(r'w) Σ ^(S1nTUl...1SNnTu]w)pj(w) uGCh(r) wen(SkC\Tu) = Σ Pr(r>«)^(51ПГ‖>...>5лгПГ‖). D w€Ch(r) Exercises (1) For the game of perfect information presented in Figure 1.11, what is your advice to player A about how to play? What about player BⅠ

24 1. Gaines in Extensive Form (2) The game tree shown in Figure 1.12 has one chance move; it is not of perfect information (the information sets are enclosed by dashed lines.) List the strategies for each of the two players. (3) Suppose that the rules for the game shown in Figure 1.11 are changed so that B does not know the result of either one of the chance moves. What is your advice to player A about how to play? What about player BⅠ (4) Two players, call them Frankie and Johnny, play the following card game (which we will call One Card): First, each antes $1 (that is, puts it into the "pot" in the middle of the table). Then a card is dealt face-down to Frankie from a deck which contains only cards marked high and low (in equal numbers). Frankie looks at her card and bets either $1 or $5 (by putting the money into the pot). Whichever the bet, Johnny can either see (by matching the bet) or fold. If Johnny folds, Frankie takes all the money in the pot. If he sees, the card is revealed. Then, if it is high, Frankie takes all the money and, if low, Johnny takes it all. Write down a game tree for One Card, indicating the information sets. (5) The game of Sevens is played between two players (called A and B) as follows. Each player rolls a fair die in such a way that the other cannot see which number came up. Player A must then bet $1 that either: (i) The total on the two dice is less than seven, or (ii) the total is greater than seven. Then, B can either (i) accept the bet, or (ii) reject it. If the bet is rejected, the payoff to each is zero. Otherwise, both dice are revealed. If the total is exactly seven, then both payoffs are zero. Otherwise, one of the players wins the other's dollar. Describe a tree for Sevens, including information sets. (6) Let T be an ΛΓ-player game tree. Given a choice subtree Si for each i, let H = n^LjSi and prove that either (a) H is a single path from the root to a terminal vertex; or (b) every path in H from the root to a terminal vertex contains a vertex belonging to chance. 1.5. Equilibrium TV-tuples of Strategies Roughly speaking, an TV-tuple of strategies (one for each player) is in equilibrium if one player's

departing from it while the others remain faithful to it results in punishment for the straying player. The idea is that once the players start playing according to such an TV-tuple, then they all have good reason to stay with it. This gives us a "solution concept," one which is reasonably amenable to mathematical analysis. There is a large class of games (to be considered in the next chapter) for which everyone agrees that an equilibrium JV-tuple really is a "solution." For other games, this is not so clear. We will discuss this question later. The formal definition is the following:

1.5. Equilibrium N-tuples of Strategies 25 Figure 1.12. Game tree for Exercise (2) of Section 1.4. Definition 1.8. Let Γ be an TV-player game in extensive form and denote the players' strategy sets by $\Sigma_1, \Sigma_2,..., \Sigma\lambda\tau$. An TV-tuple (5*), where s? e ς,- for $1 < i < \eta$, is an equilibrium N-tuple (or is in equilibrium) if, for every i and for every Si £ Σ,·, we have щ(o1,... , Si,... , Stf) < TTi(S1,..., Si,... , SN). Thus, a single player who departs from the TV-tuple is hurt (or, at least, not helped). It is easy to check that the pair of strategies (L, LRLR) in the game shown in Figure 1.11 is in equilibrium. The following game is an important example in this context. Example 1.2 (The Prisoner's Dilemma). Two criminals, call them Bonnie and Clyde, are arrested for robbery. The police immediately separate them so that they are unable to communicate in any way. Each is offered the following deal: If you confess and implicate the other prisoner, then you will serve only one year in jail (if the other guy doesn't confess) or five years (if the other guy does confess). On the other hand, if you don't confess and the other prisoner does, you will serve ten years. Finally, if neither of you confesses, you will both serve two years since our case won't be very strong.

26 1. Games in Extensive Form Table 1.1. A game with two equilibrium TV-tuples. Strategy triples Payoff vectors (1.1.1) (0,0,0) (1.1.2) (0,-1,1) (1.2.1) (1,0,-1) (1,2,2)1 (0,0,0) (2,1,1)» (1,1,-2) (2.1.2) (2,0,-2) (2,2,1) (1,-1,0) (2,2,2) (-1,1,0) Viewing this as a game, we see that each player has only two strategies. Let's call them C (for "cooperate with the other criminal") and D (for "defect"). Adopting strategy C means to refuse to confess, that is, to act as though the two prisoners formed a team of people who could trust each other, while adopting D means to confess. The "payoffs" (namely, years in prison) are really punishments and so we should use their negatives as payoffs instead. We see that the payoff to each player arising from the pair of strategies (Д D)

is —5. For the strategy pair (Д C) (that is, Bonnie defects while Clyde cooperates), Bonnie "wins" —1 and Clyde "wins" —10. Now, it is easy to see that (D,D) is an equilibrium pair. Moreover, it is the only one. Nevertheless, it is a little difficult to think of (£), D) as the "solution" to Prisoner's Dilemma. They both do better if they play the strategy pair (C, C). Of course, they would have to trust each other in order to do so. This game will be discussed again in a later chapter. Another example is described in Table 1.1. In this game, there are three players. Each has only two strategies, denoted 1 and 2. The players simultaneously choose their strategies, so that no player knows what either of the other two chose. Thus, the game is not of perfect information. There are eight combinations of strategies. These combinations are listed in the left-hand column, and the corresponding 3-tuples of payoffs are in the right-hand column. The equilibrium 3-tuples are marked with jj. We see that there are two equilibrium 3-tuples, neither of which looks very stable. It is clear that the third player would strongly prefer the 3- tuple (1, 2, 2), while the other two players would prefer the 3-tuple (2,1,1). These two could force the issue by playing according to 2 and 1, respectively. However, it would be worth his while for player 3 to make a payment to player 2 in exchange for 2's playing strategy 2. Despite our reservations about whether knowing an equilibrium TV-tuple solves a game, the concept is still of value. It is therefore of interest to know conditions under which one exists. It is certainly true that there are many games for which there is no equilibrium TV-tuple. The reader may verify

1.5. Equilibrium N-tuples of Strategies 27 that this is true for two-finger morra. The most important class of games for which existence is guaranteed is that of games of perfect information. This type of game was defined earlier, and a definition can now be given in terms of our notation for games in extensive form. Definition 1.9. Let Γ be a game in extensive form. Then Γ is of perfect information if, for every player, each choice subtree is a strategy. We can now state the following: Theorem 1.9. Let T be a game in extensive form. If Γ is of perfect information, then there exists an equilibrium N-tuple of strategies. Proof. The proof is by induction on the depth of the game tree T. The smallest possible depth of a game tree is one, since we excluded trivial trees. Now, if De(T) = 1, then at most one player has a move (depending on whether the root belongs to a player or to chance). Assuming that the root belongs to player P, let the strategy for P be to choose a (terminal) child of the root for which that player's payoff is a maximum. The corresponding choice subtree for

P then consists of a single edge. Let the choice subtrees for the other players be all of T (in fact, no other possibilities exist). That the resulting TV-tuple is in equilibrium should be clear. In the case where the root belongs to chance, each player's choice subtree must be taken to be all of T (since no other possibilities exist). The proof is complete in case De(T) = 1. Now suppose that the depth of T is m > 1 and that the theorem holds for all games of perfect information for which the depth of the game tree is less than m. Let r be the root of T. For each child и of r, the cutting Tu is a game tree and Oe(Tu) < De(T). We regard Tu as the game tree for a game of perfect information by defining each player's strategy set to be the set of all choice subtrees determined by that player. By the inductive assumption, there is an equilibrium N-tuple, (5", S£ > · · · > ^n)j of strategies in Tu. We want to put these together to form an TV-tuple of strategies in T. There are two cases. First, if the root of T belongs to a player Pj, choose и to be a child of the root such that the payoff Kj(Si, S£ > · · · > &n) is a maximum over Ch(r). Now define Sj to be the union of S^ and the edge (r, u). For 1 < г < N and г ф j, define 5* to be U (O-.Ousn. v6Ch(r) Second, if r belongs to chance, define each S* according to the preceding formula. It should be clear that we have now defined an TV-tuple of choice subtrees. Since the game is of perfect information, all these choice subtrees are

28 1. Games in Extensive Form strategies. We must now show that this N-tuple of strategies is in equilibrium. To do so, suppose that one of the players, say Pг·, plays according to strategy Si instead of S*. We consider the following three cases: (1) Suppose the root belongs to player Pг·. Now if (r, w) is in 5г·, we have, by Theorem 1.8, ττ,·(5í,...,£·,...,S*N) = ττ,·(5í ПГЮ Sl-ПГ. S*N CiTw). (1.1) Then since (Sf,..., £%) is an equilibrium JV-tuple in Tw, we have Tw,..., Sjy nTw)<^(S^,...,S%). Now if the edge (r, u) belongs to 5*, we have, by the way in which S* was chosen, ^•(S™,..., Ξχ) < 7Tj(5i,..., 5]y·). By definition of the S*'s, we have 7Tj(5* ПТи,...,5^ ПГц) = тгг'(5",.. .,5]у·). Finally, by Theorem 1.8, 7Tj(S*,... ,Stf) — 7Ti(Sl C\TU1.. .,Stf Г\Ти). Combining all these, we get 7Гг'(5* , . . . , 5г', . . . , Stf) < K{(Sl, . . . , Stf). (2) Suppose that the root belongs to a player Pj different from Pг·. Let и be that child of the root so that (r, u) is in Sj . Then we have 7тг'(5*,.. .,5г·,... ,Sjv) = тгг'(5* ПТМ,.. .,S,· ПТМ,.. .,SJr ПТМ). But then Finally, 7Ti(5i,.. .,5]у·) = 7Гí(5í,.. .,Sjy). Combining these yields the desired inequality.

Table 1.2. Game for Exercise (4) of Section 1.5. Strategy triples Payoff vectors (1.1.1) (1,-1,1) (1.1.2) (0,0,0) (1.2.1) (-1,2,0) (1.2.2) (0,1,-1) (2.1.1) (1,1,-2) (2.1.2) (-2,1,0) (2,2,1) (1,0,1) (2,2,2) (0,0,1) (3) Suppose that the root belongs to chance. Then, for each child и of the root, we have Tu,..., SN Now apply part (2) of Theorem 1.8 to get that 7T;(S* , . . . , Si, . . . , S*N) < 7Ti(Sl, . . . , Stf). ☐ This theorem was essentially proved for the first time in [vNM44]. Exercises (1) Find all equilibrium pairs of strategies (if any) of the game in Figure 1.7. (2) Determine the equilibrium pairs of strategies for the game of One Card (see Exercise (4) of Section 1.4). (3) Verify that there are no equilibrium pairs of strategies for two-finger morra. (4) Table 1.2 describes a three-person game in the way that Table 1.1 does. Are there any equlibrium 3-tuples for this game? Sketch a tree for this game, indicating information sets. 1.6. Normal Forms Let r = (T,{P1,...,PiV}, {E1,...,Eiv}) be a game in extensive form. For every $\Lambda\Gamma$-tuple of strategies (Si,..., Sn) in the Cartesian product $\Sigma\iota\chi \cdots \chi \Sigma\lambda\gamma$, there is determined an TV-tuple 7r(Si, $\cdots$ j Sn) of payoffs to the players. This function from $\Sigma\iota\chi \cdots \chi \Sigma\nu$ to $tN (that is, the N-fold Cartesian product of the real line 3£ with itself) is called the normal form of the game Γ. We will often study games in their normal forms without ever mentioning their extensive forms.

It should be mentioned that the normal form of a game is not uniquely defined. If we permute the order of the players then we change the Cartesian product $\Sigma\iota\chi \cdots \chi \Sigma\nu$ and thus the function π. In practice, this lack of uniqueness presents no difficulty and we will ignore it. Thus when we say "the normal form" we will mean any of the possible ones. We abstract the idea of the normal form of a game in the following: Definition 1.10. Let $X\iota,...,Xn$ be finite nonempty sets and let π be a function from the Cartesian product $X\backslash x \cdots x Xn$ to $tN. Then π is called an N-player game in normal form with strategy sets $X\backslash$,..., Xn- Let us designate the players in an TV-player game π in normal form by Pi,..., Pn- They play as follows: Each player Pr· chooses Xi £ Xi. These choices are made simultaneously and independently. The payoffs are then given by the components of the TV-tuple 7г(#), where χ — (a?i,.. .,xn)- The definition of equilibrium TV-tuples carries over almost unchanged to the context of games in normal form. To be precise, we give Definition 1.11. Let π be a game in normal form with strategy sets $X\backslash$, ..., Xn - An iV-tuple x* e $X\iota\chi \cdots \chi Xn$ is an equilibrium N-tuple if, for all $1 < r < N$ and any X{ £ Xr·, Ki(xl,...,Xi,...,X*N) < TTi(x*). Certainly it is true that if a

game in extensive form has an equilibrium N-tuple of strategies, then so does its normal form. In the case where N — 2, a game in normal form can be pictured as a pair of matrices. In fact, let m be the number of strategies in Σι and let η be the number of strategies in Σ2. Form an m χ η matrix of payoffs for player Pi by labeling the rows with members of Σι and the columns with members of Σ2. The entry in the row labeled with a? £ Σι and in the column labeled with y £ Σ2 is defined to be 7Γι(#, y). Call this matrix M1. The payoff matrix M2 for player P2 is formed in a similar way. These matrices are not unique since a permutation of either or both of the sets Σι or Σ2 would also permute the rows or columns of the matrices. This lack of uniqueness presents no difficulty. Let us compute the normal form for the two-person game of perfect information whose tree is shown in Figure 1.13. The players are designated A and B. There are eight strategies for player A: LLL) LLR) LRL, LRR, RLL, RLR, RRL, RRR.

1.6. Normal Forms 31 8 / \ η IV \.ii 12./ \. ц x7/ \. 15 (-1.0) ^1,-1) (2,0) "(2,1) (-2,-1) (1,-1) H·"1) (1,0) Figure 1.13. A two-person game. Here again, L means "left" and R means "right." Thus, the strategy RLR corresponds to the choice function c a, where <*(2) = 5, <u(6) = 12, d(7) = 15. Player B also has eight strategies and they are designated in exactly the same way. For example, B's strategy LLR corresponds to the choice function cjg, where cв(3) = 6, cB(4) = 8, cB(5) = ll. The normal form of this game is contained in the following matrices. The order of the rows and columns corresponds to the order of strategies given above. The matrix M1 of payoffs for player Pi is given first. / -3/2 -3/2 0 0 0 0 3/2 ^ 3/2 -3/2 -3/2 0 0 0 0 3/2 3/2 -1/2 -1/2 1 1 0 0 3/2 3/2 -1/2 -1/2 1 1 0 0 3/2 3/2 -1 0 -1 0 1/2 3/2 1/2 3/2 -1 0 -1 0 1/2 3/2 1/2 3/2 0 1 0 1 1/2 3/2 1/2 3/2 0\ 1 0 1 1/2 3/2 1/2 3/2/ Then the payoff matrix M2 for player P^ is

32 1. Gaines in Extensive Form Table 1.3. Game for Exercise (10) of Section 1.6. Strategy triples (1,1,1 (1,1,2 (1,2,1 (1,2,2 (2,1,1 (2,1,2 (2,2,1 (2,2,2 (3,1,1 (3,1,2 (3,2,1 (3,2,2 Payoff vectors (0,-1,0) (0,-2,0) (3,0,-1) (1,-1,-1) (0,0,0) (0,0,-1) (-i>i>i) (2,1,-1) (0,0,2) (0,-1,1) (1--2.1) (1,1,-1) / -1/2 -1/2 ! -1/2 ! -1/2 -1/2 -1/2 -1/2 \ -1/2 -1/2 -1/2 -1/2 -1/2 0 0 0 0 -1 -1 -1 -1 -1/2 -1/2 -1/2 -1/2 -1 -1 -1 -1 0 0 0 0 -1/2 0 -1/2 0 -1/2 0 -1/2 0 -1/2 0 -1/2 0 0 1/2 0 1/2 -1 -1/2 -1 -1/2 -1/2 0 -1/2 0 -1 \ -1/2 -1 -1/2 0 1/2 0 1/2) From the two matrices, it is easy to pick out the equilibrium pairs of strategies. Indeed, it is clear that a

pair (S*,^) 1S an equilibrium pair if and only if the entry at coordinates (S*,^) in Mχ is a maximum in its column, while the entry in M2 at the same coordinates is a maximum in its row. It can be observed that there are six equilibrium pairs in this example. Both the extensive and normal forms of games have advantages. The normal form is perhaps simpler mathematically. For the kind of games to be considered in the next chapter, the normal form is the one which allows us to solve the games. On the other hand, it is simpler to go from a verbal description of a game to its extensive form than to its normal form. Also, if we are interested in considering small subgames of a game as a way of better understanding a game which is too large to analyze, then the extensive form is better. This is because it is fairly easy to determine which small pieces of the big tree to examine. It is difficult to break down a normal form into pieces in this way. This sort of analysis will come up in Chapter 7.

1.6. Normal Forms 33 Exercises (1) Choose five entries in each of the two 8x8 matrices presented in this section and verify them. (2) Give the normal form (as a pair of payoff matrices) of the game shown in Figure 1.7. (3) Write down the normal form for the game shown in Figure 1.12. Find the equilibrium pairs (if any). (4) Find all the equilibrium pairs of strategies for the game discussed in this section. (5) Write down the normal form for the game shown in Figure 1.5. (6) Write down the normal form for the game shown in Figure 1.8. (7) Write down the normal form for the game shown in Figure 1.11. (8) The game of Fingers is played as follows: The two players (Mary and Ned) simultaneously hold up either one or two fingers. Mary wins in case of a match (the same number), and Ned wins in case of a nonmatch. The amount won is the number of fingers held up by the winner. It is paid to the winner by the loser. (a) Describe the strategies for both players. (b) Write down the normal form (as a pair of payoff matrices). (c) Verify that there is no equilibrium pair of strategies. (9) Write down the normal form of two-finger morra (see Figure 1.6). (10) The normal form of a three-player game is given in Table 1.3. Player Pi has three strategies (1, 2, 3); players P2 and P3 each have two (1 and 2). Find the two equilibrium 3-tuples of strategies. (11) Refer to Exercise (10) of Section 1.3 for the terminology and notation of this exercise. Define $\Gamma^* = (T^*, \{P\text{Ь}...,P^*\},^*,...^\})$, where $\Sigma? = \{5, (I\ T^* : Si\ G\ \Sigma;\}$. Prove that Γ and Γ^* have the same normal form.

2 Two-Person Zero-Sum Games Let ĭ be a game in normal form with strategy

sets $X_1, \ldots, X_n$. We say that this game is zero-sum if $N \wedge\wedge\bullet(я?1,\ldots,я?лг) = 0$, $г = 1$ for every choice of $X_i \in X_i, 1 < г < N$. The corresponding definition for a game in extensive form states that the sum of the components of $p(v)$ is zero for each terminal vertex v. This condition is certainly true for ordinary recreational games. It says that one player cannot win an amount unless the other players jointly lose the same amount. Nonrecreational games, however, tend not to be zero-sum. Competitive situations in economics and international politics are often of the type where the players can jointly do better by playing appropriately, and jointly do worse by playing stupidly. The phrase "zero-sum game" has entered the language of politics and business. In this chapter we are concerned with games in normal form which are zero-sum and which have two players. In the previous chapter, we discussed the fact that a two-person game in normal form can be represented as a pair of payoff matrices (one for each player). If the game is also zero-sum, the two matrices are clearly just negatives of each other. For this reason, there is no need to write down both. From now on, we shall represent a two-person zero-sum game as a matrix game, that is, as a single $m \chi \eta$

36 2. Two-Person Zero-Sum Games matrix M. The two players are referred to as the row player and the column player, respectively. The row player has m strategies which are identified with the rows of M. The column player has η strategies which are identified with the columns of M. If the row player plays strategy $г$ and the column player plays strategy j, then the payoff to the row player is $rriij$ and the payoff to the column player is — $rriij$. It is important to make it clear from the beginning that larger numbers in M are favored by the row player and smaller ones by the column player. Thus, a negative entry is a loss to the row player but a gain (of the absolute value) to the column player. 2.1. Saddle Points The idea of an equilibrium pair of strategies can be easily translated into the context of matrix games. Indeed, if (p, q) is such a pair, then $TTliq < IJlpq$ for all 2, and $rripj > mpq$ for all j. Notice that the second inequality is reversed because the payoff to the column player is the negative of the matrix entry. These inequalities motivate the following definition. Definition 2.1. Let M be a matrix with real entries. An entry mpq of M is a saddle point of M if it is simultaneously a minimum in its row and a maximum in its column. Thus, if M is a matrix game, then mpq is a saddle point if and only if (p, q) is an equilibrium pair of strategies. In the following three examples, the entry $ГП21$ is the only saddle point of the first matrix, entry $m\backslash2$

and entry т₃₃ are both saddle points of the second matrix (there are two more), and the third matrix has no saddle points at all. сп) (-:н) (ι··) The first thing we prove about saddle points is the following: Theorem 2.1. If тьι and mpq are saddle points of the matrix M, then rrikq and mpi are also saddle points and ™>kl — ™>Pq — mkq — ™>Pl-

2.1. Saddle Points 37 Proof. Since saddle points are maxima in their columns and minima in their rows, we have while rripg < mpt < тьι. Thus, ™>kl — ™>pq — ™>kq — ™>pl- Also, rrikq is a saddle point since rrikq = m^i is a minimum in row fc, and rrikq = rripq is a maximum in column q. Similarly, mp\ is a saddle point. Π We want to establish that a saddle point provides an acceptable solution to a matrix game. To do so, let us consider how the row player should rationally play. First, it is reasonable to assume that, no matter how he plays, his opponent will respond by playing so as to maximize her (the column player's) payoff. After all, she knows the game as well as he does and it would be foolish to assume that she will, in the long run, go on being charitable or making mistakes. Now, the fact that the game is zero- sum means that maximization of the column player's payoff is precisely the same as minimizing the row player's payoff. Therefore, the row player should choose his strategy so that the minimum possible payoff due to this strategy is as large as possible. The column player should act in a similar way. These ideas lead to the following: Definition 2.2. Let M be an m χ η matrix game. The value to the row player and value to the column player are, respectively, ur(M) = maxminm,'j • j and uc(M) = minmaxm^·. Thus, ur(M) is an amount that the row player is guaranteed to win if he plays a strategy κ for which the maximum in the definition of ur(M) is attained, that is, so that mmmkj — ur(M). j The amount actually won by playing strategy κ might even be larger if the column player chooses a strategy unwisely. A similar interpretation holds for the column player. Her best guaranteed payoff is obtained by playing a column / such that maxrnj/ = vc(M).

38 2. Two-Person Zero-Sum Games In the three examples just given, the values to the row player are, respectively, — 1,1, and 1 and the values to the column player are, respectively, —1,1, and 2. The fact that the row player's value equals the column player's value in the two examples where a saddle point exists (and not in the other example) is no coincidence. Before proving this, we need a lemma. Lemma 2.2. For any matrix M, ur(M) < uc{M). Proof. For

any i and any /, we have minm2j < гпц. 3 Maximizing both sides over г gives ur(M) = maxminm2y < maxmj/, * j i for every /. Hence ur(M) < minmaxm27 = uc(M). D Theorem 2.3. If the matrix game M has a saddle point mpq, then ur(M) - ue(M) = mpq. Proof. We have that mm rripj =∎ mpq, 3 and so ur(M) > mpq. Also maxmjg = mpq, and so uc(M) < mpq. Combining these, we have uc(M) < mpq < ur(M). But, from the lemma, ur(M) < uc{M), and so the two values are equal. П The converse also holds. Theorem 2.4. Ifur(M) = uc{M}, then M has a saddle point.

2.1. Saddle Points 39 Proof. Choose к such that minmfcj = ur(M). 3 Then choose / such that mki = mmmkj = ur(M) — uc{M). 3 Now niki is a minimum in row к. There exists a column q such that maxmi? = ue(M). Thus ты = uc{M} = maxm!g > mkq. Since mfc/ is a minimum in its row, we have ты = rafcg, and so rrikq is also a minimum in its row. Finally, ™>kq = ?^fc/ = maxm,-g, г and so rrikq is a saddle point. П Let M be a matrix with a saddle point. Suppose that the row player and column player play row i and column j, respectively. If row г contains a saddle point, say m,·/, then rriij > mu = ur(M). On the other hand, if row г does not contain a saddle point, while column ,;' does contain one (say, m^·), then ™>ij < ™>kj = Ur(M). In other words, the row player, by playing a row containing a saddle point, can guarantee himself a payoff of at least ur(M). The consequence of not playing such a row is that the column player could (and would) play so as to make the row player's payoff at most ur(M). From these considerations, it is clear that the best play for the row player is a row containing a saddle point. By similar reasoning, the column player should play a column containing a saddle point. If they both follow our recommendation, then, by Theorem 2.1, the payoff entry rriij will be a saddle point. In summary, if a matrix game M has a saddle point, then its solution is: • The row player plays any row containing a saddle point. • The column player plays any column containing a saddle point. • The payoff to the row player is ur(M) = wc(M), while the payoff to the column player is, of course, — ur(M) = — uc(M).

40 2. Two-Person Zero-Sum Games Exercises (1) Find all saddle points of the matrix 2 -1 0 0 1 -2 -1 0 le matrix 1 -1 -2 0 -1 0 1 1 1 -1 0 1 1 -1 2 1 -1 1 0 0 2 0 1 2 0 -1 2 -1 1 1 2 1 0 2 0 1 1 (2) Find all saddle points of the matrix (3) For which values of o does the following matrix have a saddle point: -2 α λ o 1 J ' (4) For which values of a does the following matrix have a saddle point: 1 a 2

-1 (5) For the following matrix, compute ur(M) and uc(M): M = (6) For the following matrix, compute ur(M) and uc(M): M-. 1-1 2 0 1 2 0-11-1 -22001 0 1-10 2 2.2. Mixed Strategies The results of the previous section show that matrix games with saddle points can be easily solved. Let us consider a matrix without any saddle points. Here is a 2 x 2 example: M (-?"!)■

2.2. Mixed Strategies 41 It is easy to see that if the row player consistently plays strategy 1, then the column player, when she notices this, will play 2 consistently and win 1 every time. On the other hand, if the row player plays 2 consistently, then the column player will play 1 and again win 1. The same result would ensue if the row player varied his strategy, but in a predictable way. For example, if the row player decided to play 1 on odd-numbered days and major holidays, and 2 at other times, then the column player would eventually catch on to this and respond appropriately. In case it is the column player who adopts a fixed strategy or a predictable pattern of strategies, the row player could always respond so as to win either 1 or 2. A third possibility is that both players play flexibly, each responding to the other's previous moves. In this case, the row player might begin by playing 1. When the column player catches on, she starts playing 2. Then the row player switches to 2, whereupon the column player goes to 1. Then the row player goes back to 1. This rather mindless cycle could be repeated forever. A new idea is needed in order to get any closer to a satisfactory concept of a solution to a matrix game: Each player should choose, at each play of the game, a strategy at random. In this way, the other player has no way of predicting which strategy will be used. The 'probabilities with which the various strategies are chosen will probably be known to the opponent, but the particular strategy chosen at a particular play of the game will not be known. The problem for each player will then be to set these probabilities in an optimal way. Thus, we have the following: Definition 2.3. Let M be an m χ η matrix game. A mixed strategy for the row player is an m-tuple ρ of probabilities. That is, Pi > 0> 1 < i < m, and m Σ* = i· г=1 Similarly, a mixed strategy for the column player is an η-tuple q of probabilities. That is, Ц.2 > 0) 1 < 3 < я, and η

42 2. Two-Person Zero-Sum Games The idea of a mixed strategy for the row player is that he will, at each play of the game, choose his strategy at random, and that this choice will be made so that the probability of choosing strategy г is pi. Thus, there is no way for the column player to predict which particular

strategy will be played against her. In practice, this choice of strategy can be carried out by any convenient "chance device." For instance, suppose that the row player in the 2x2 game which we just discussed decides to play according to ρ — (1/2,1/2). Then he could simply flip a fair coin every time the game is played (without letting his opponent see how it comes up). He could then play row 1 if the toss is heads and row 2 if it is tails. For the same game, if the column player wants to use the mixed strategy q = (1/3,2/3), she could achieve this by glancing at a watch with a second hand just before playing. If it shows between 0 and 20 seconds, she would play column 1, otherwise column 2. Less homely chance devices would be needed in more complicated situations. A random-number generator (found on most calculators) can always be used. To emphasize the distinction between mixed strategies and ordinary strategies, we will refer to the latter as pure strategies. It should, however, be realized that pure strategies are really special cases of mixed strategies. The mixed strategy ρ for which pi — 0 for г ф к (and p^ = 1) is obviously the same as the pure strategy k. If one or both players adopt mixed strategies, the payoffs at each play of the game depend on which particular pure strategies happened to be chosen. The important quantity to use in studying mixed strategies is the expected payoff. It is an average over many plays of the game and is denoted $E(p > \theta)\iota$ where ρ and q are the mixed strategies in use. To see how to define this quantity, consider first the case where the row player plays according to the mixed strategy ρ and the column player plays the pure strategy j. Then, for an m χ η matrix M, the payoff to the row player is m^· with probability pi. The expected payoff is thus m Now, if the column player adopts the mixed strategy q, the payoff to the row player is E{p)j) with probability qj. Therefore η m η m $E(p,«) = \Sigma v \Sigma\lambda"*4 = \Sigma \Sigma\rho<3;™!'; \cdot$ j=l r=1 j=lr=1

2.2. Mixed Strategies 43 Note also that, interchanging the order of summation, we get m η $E(p > Я) = \Sigma$ Pi Σ «i mu · r = 1 j=l Let us compute some expected payoffs. With M the 2x2 matrix presented at the beginning of this section, suppose the row player plays the mixed strategy ρ = (1/2,1/2) and the column player plays the mixed strategy q = (2/3,1/3). Then £((1/2,1/2), (2/3,1/3)) = 2(2/3)(l/2) - (1/3)(1/2) - (2/3)(l/2) + (1/3)(1/2) = 1/3. Now suppose that the column player plays q — (2/3,1/3) as above. Given that information, how should the row player play so as to maximize his expected payoff? To compute his best strategy, note first that every mixed strategy for him is of the form (p, 1 — p), where $0 < \rho < 1$. Now Я((p,1-p),(2/3,1/3)) = 4p/3-1/3. From this we

see that the maximum expected payoff is 1 and is attained for $\rho = 1$. That is, the row player should play the pure strategy 1. 2.2.1. Row Values and Column Values We now make a pair of definitions which are analogous to the definitions of row player's and column player's values which we made earlier. The change is that mixed strategies are substituted for pure strategies. Definition 2.4. Let M be an m χ η matrix game. The row value is defined vr(M) — maxmini^]?, g), ρ я where ρ* and <f range over all mixed strategies for the row player and column player, respectively. Similarly, the column value is defined vc(M) = minmax E(p, q). я ρ Thus vr(M) is an amount that the row player is guaranteed to win on the average, assuming that he plays intelligently. He may win more if the column player makes mistakes, but he cannot count on it. It will be proved in a later chapter that the maximum in the definition of vr is actually

44 2. Two-Person Zero-Sum Gaines attained, that is, that there exists at least one mixed strategy r for the row player such that vr(M) = mm E(r, q). я Such a strategy is called an optimal mixed strategy for the row player. Similarly, there exists at least one optimal mixed strategy s for the column player such that vc(M) = maxi?(p, s). ρ There is an important point to be made about these optimal mixed strategies. If the row player plays such a strategy, he is guaranteed to gain at least vr(M) even if the column player plays as well as she can. On the other hand, if the row player knows that the column player is playing a stupid strategy, then there may well be a counter-strategy which gains him more than the "optimal" one does. For now, we will simply assume that we already know that these optimal mixed strategies exist. Another theorem (called the minimax theorem), which will be proved later, is that the row and column values are always equal. This is a surprising and vital piece of information. First, it is surprising because it is not at all clear from the definitions why it should be true. Second, it is vital because the whole theory we are developing in this chapter would not work without it. To gain some insight into this matter, we prove the following: Theorem 2.5. Let M be an m χ η game matrix and let r and s be optimal mixed strategies for the row player and column player, respectively. Then vr(M) <E(r,s) < ve(M). Proof. Since s*is a mixed strategy for the column player, we certainly have vr(M) - min E(r,q) < E(r,s). я Similarly, we have ve(M) - max^p, s) > E(r, s). ρ Combining these two inequalities gives the result. Π This theorem gives us the inequality vr(M) < vc(M), which is a weak version of the minimax theorem. More

importantly, if we assume the minimax theorem, vr(M) — vc(M), we get vr(M) = E(r,8) = vc(M).

2.2. Mixed Strategies 45 Now suppose, in addition, that r and s are optimal mixed strategies for the row player and column player, respectively. Then, if ρ and q are any mixed strategies for the players, E(p, 5) < vc(M) = E(r, 3) = vr(M) (2.1) and E(r, q) > vr(M) = E(r, 3) = vc(M). (2.2) Then (2.1) and (2.2) together say that the row player can do no worse than vr(M) by playing r, and may do worse than vr(M) playing p. Thus, the row player should play an optimal mixed strategy. Similarly, the column player should play an optimal mixed strategy. We have the following: Definition 2.5. Let M be a matrix game for which vr(M) A solution to M consists of three components: • An optimal mixed strategy for the row player. • An optimal mixed strategy for the column player. • The value of the game, v(M), defined by v(M) = vr(M) = ve(M). Inequalities (2.1) and (2.2) are reminiscent of the definition of a saddle point in a matrix. In fact, if we think of the mixed strategies as labeling the rows and columns of the infinite matrix whose entries are the expected payoffs E(p, q), then (2.1) and (2.2) say that the pair (F, s) of optimal mixed strategies is a saddle point. The following theorem is analogous to Theorem 2.1. Theorem 2.6. Let M be a matrix game for which vr(M) — vc(M). Suppose that r and t are both optimal mixed strategies for the row player, while s and и are both optimal mixed strategies for the column player. Then E(r, s) = E(r, u) = E(ï, и) - E(i, s). Proof. Using (2.1) and (2.2), we have E(r, s) > E(i, s) > E (ζ й), and E(t, и) > E(r, и) > E(r, s). Thus, all the inequalities above are really equalities, and the proof follows. П = vc(M).

46 2. Two-Person Zero-Sum Games It follows that, if there is a choice, it does not matter which optimal mixed strategies the players choose. Since saddle points of game matrices correspond to equilibrium pairs of pure strategies, we see that a pair of optimal mixed strategies is analogous to an equilibrium pair of mixed strategies. This concept will be formally defined and discussed in detail in Chapter 5. We emphasize that once the minimax theorem has been proved we will know that every matrix M satisfies the condition vr(M) — vc(M) and thus has a solution. The following theorem gives us a somewhat simpler way to compute the row and column values. It allows us to use only pure strategies in computing the inside minimum (or maximum). Theorem 2.7. Let M be an m χ η matrix. Then vr(M) = maxmin E(p,j) ρ J and vc(M) =

minmax^g). я * Here, j ranges over all columns and г ranges over all rows. Proof. Since pure strategies are special cases of mixed strategies, we certainly have minE(p, q) < minE(p, j), я J for any mixed strategy ρ for the row player. To prove the opposite inequality, let / be such that E(p,l) = mm E(p J). j Then, if g* is a mixed strategy for the column player, we have η E(p,q) = ^qjE(pJ)>E(p1l). j=i Since this holds for all <f, we have, by the choice of /, minE(p, q) > minE(p, j). я J Thus, for all p, we have min E(p, q) - min E(p, j). я j Maximizing over ρ on both sides of this equation completes the proof for the row value. The proof for the column value is similar. Π

2.2. Mixed Strategies 47 This proof also shows that a strategy r for the row player is optimal if and only if vr(M) - min E(f \j), 3 where ,; runs over all columns. Similarly, a strategy s*for the column player is optimal if and only if vc(M) = max^(f,s), i where i runs over all rows. This theorem will play an important role in the remaining sections of this chapter. If M has a saddle point, say m^√, then the pure strategies к and / for the row player and column player, respectively, together with the value, m^√, constitute a solution. This follows from: Corollary 2.8. If M has a saddle point rriki, then vr = vc = mk\, and so к and I are optimal mixed strategies. Proof. Since pure strategies are special cases of mixed strategies, we see from the theorem that vr(M) = maxmmE(p,j) > ur(M). ρ 2 Similarly, vc(M) < ue(M). Since ur(M) = uc(M) and vr(M) < vc(M), we see that vr(M) = vc(M) - ur(M) = uc(M) - ты. To show that к is optimal for the row player, notice that min E(k,j) = mmrrikj = ты = vr(M). 3 3 By the remark made after the proof of the theorem, к is optimal. Similarly, / is optimal. Π

48 2. Two Person Zero-Sum Games 2.2.2. Dominated Rows and Columns It is sometimes possible to assign probability zero to some of the pure strategies. In other words, certain pure strategies can sometimes be identified as ones which would never appear with positive probability in an optimal strategy. Definition 2.6. Let M be an m χ η matrix. Then row i dominates row At if ™>ij > ™>kj for all j. Also, column ,; dominates column / if rriij < mu for all i. Note that the inequality in the definition of domination of columns is reversed compared to the inequality in domination of rows. It should be obvious that a dominated row need never be used by the row player, and that a dominated column need never be used by the column player. This implies that if we are looking for optimal mixed strategies, we may as well assign probabilities of zero to any

such row or column. In fact, we may as well reduce the size of the matrix by erasing dominated rows and columns. Let us consider an example. Let (' -1 ~2\ M= 2-1 0 . V -1 1 1 / We see that row 1 is dominated by row 2. We erase the dominated row to get In this new matrix, column 3 is dominated by column 2 (there was no domination of columns in the original matrix). We arrive at (-? -!)■ in which there is no domination. We have the following: Definition 2.7. Let M be a matrix game and let ρ be a mixed strategy for the row player. Then row i of M is said to be active in ρ if pi > 0. Similarly, if q is a mixed strategy for the column player, then column j is active in q if qj > 0. A row or column which is not active is said to be inactive.

2.2. Mixed Strategies 49 Our discussion of domination can be summarized: A dominated row (or column) is inactive in an optimal strategy for the row player (or column player). [But see Exercise (13).] The following theorem will be useful. Theorem 2.9. Let M be an m χ η matrix game such thai vr(M) — vc(M). Let f and s be mixed strategies for the row player and column player, respectively. Then r and s are both optimal if and only if the following two conditions hold: (1) Row κ is inactive in r whenever E{k)s) < max^(i,s). (2) Column I is inactive in s whenever E(r,l) > min E(f\j). 3 Proof. Suppose first that r and s are optimal for the row and column player, respectively. To show that (1) holds, suppose that E(k,s) <maxE(i,s) - vc(M), but that rk > 0. Then m m 4r,s) = ^nE^s) < Σ'MM) = vc{M\ 2 = 1 « = 1 since E(i, s) < vc{M) for all i. This contradicts the optimality of r and s (see the remarks after Theorem 2.5). The proof that (2) holds is similar. Now suppose that both (1) and (2) hold. We must show that r and s are optimal. Using (1) and (2), together with Theorem 2.7, we have m E(r, 8) = ΣηE(ι, s) - τax^r, s) > ve(M), and η E{r,3) = Y\8jE{r,j) = minE(rJ) < vr(M). ■л J •7=1 Since vr(M) = ve(M), we have vc(M) — E(r, s) — τax£'(r, s), and vr(M) — E(r,s) — min E(r\j).

50 2. Two-Person Zero-Sum Games Thus, both r and s are optimal. Π This theorem is occasionally useful in computing optimal strategies. In fact, it will play such a role in the next section. It is very useful for verifying optimality of strategies. For example, let M = We use Theorem 2.9 to verify that r - (0,0,1) and s = (2/3,0,1/3) are optimal strategies for the row player and column player, respectively. We have Я(1, s) = 1 E(2, s) = -1/3 Я(3, s) = 1, 2 2 1 -1 0 2 -1 3 1 and Thus E(r, 1) = 1 E(r, 2) = 2 E(r, 3) = 1. max^(f,s) = 1 and min i? (r,j) — 1. • 3 Theorem 2.9 says that, in order for both r and s to be optimal, we

must have $\Gamma 2$ and $\ll 2$ equal zero. They both are, and so optimality is verified. Exercises (1) Let M = If the column player plays the mixed strategy $(2/5,1/3,4/15)$, what is the best way for the row player to play? (2) Let $/ * $ м= I -ι \ 2 (a) Compute £7$((l/5,2/5,2/5),(1/3,1/3,1/3))$. (b) On the assumption that the row player continues to play $(1/5,2/5,2/5)$, what is the best way for the column player to play? (3) Let M = Eliminate dominated rows and columns so as to reduce M to the smallest size possible. 3 1 0 2 1 -1 -1 -1 1 0 -1 2

2.2. Mixed Strategies 51 (4) Suppose you are the row player in a 3 X 4 game and want to play the mixed strategy $\rho = (5/12,1/4,1/3)$. Your calculator can produce 8-digit random numbers uniformly distributed in the interval $[0,1]$. Explain how you would use these random numbers to achieve ρ (as closely as possible). (5) For the matrix game / -1 2-2 0 1 -2-1320 2 1 0-1-2 0 0 2 11 \ 1 -1 0-2 1 verify that p = $(5/52,0,11/52,17/26,1/26)$, q = $(21/52,3/13,0,3/52,4/13)$,?;= 19/52, is a solution. (6) Given verify that the following is a solution: ?= $(1/2,0,1/2)$, <f = $(3/4,1/4,0)$, v(M) = 1/2. (7) Prove that if ρ and g*are mixed strategies for the row player and column player, respectively, such that min E(p, j) : ά then ρ and q are optimal. Given / 2 ι -1 0 1 \ -3 -3 2 1 2 4 = ma.xE(i,ql 4 -3 -2 -3 -5 "5 \ 4 3 4 6/ verify thatp = $(1/2,0,1/6,0,1/3)$ andg* = $(0,1/4,1/2,1/4)$ are optimal strategies. (9) Let M be a 2 X 2 matrix. Prove that M has a saddle point if and only if it has a dominated row or a dominated column. Is this true for larger matrices? (10) Suppose that both r and и are optimal strategies for the row player for the matrix game M. Prove that if $0 < t < 1$ then tf + (1 — t)u is also an optimal strategy for the row player. (11) Let M be a matrix and c a real number. Form the matrix M' by adding c to every entry of M. Prove that Vr(M') = vr(M)-j-c and Vc(M') = vc(M) + c

52 2. Two-Person Zero-Sum Gaines (12) Let M be a matrix and a a nonnegative real number. Prove that vr(aM) = avr(M) and vc(aM) = avc(M). What happens if a < 0? (13) The assertion was made that a dominated row or column is never active in an optimal strategy. Explain why this statement may be incorrect if two of the rows or columns are equal. 2.3. Small Games In this section, we discuss a graphical method for solving those matrix games in which at least one of the players has only two pure strategies. We start with the following 2x2 example: --(-? 1)· We know, from Theorem 2.7, that vr(M) — maxmin^(p,j). ρ J Now each ρ is of the form (p, 1 — p), where $0 < \rho < 1$. For

convenience of notation, we write $\pi_j(p) = E((p, 1 - p)J)$ for $j = 1,2$ and $0 < \rho < 1$. These functions of ρ are easy to compute. We get $\pi_1(p) = 2p-(1-p) = 3p-1$ and $\pi_2(p) = -3p + (1 - p) = -4p + 1$. Now these two functions are both linear. In Figure 2.1, we graph both of them and indicate, with a heavier line, their minimum. By definition, vr(M) is the maximum of this minimum. It is circled in the figure. It occurs at the point where $\pi_1(\rho)$ crosses $\pi_2(\rho)$· Setting $\pi_1(\rho) = \pi_2(\rho)$, we get that vr(M) — —1/7 and it is attained for ρ — 2/I. Thus ρ — (2/7, 5/7) is an optimal mixed strategy for the row player. To find an optimal strategy for the column player, let us define $\pi_i'(q) = E(i, (q, 1 - q))$ for $i = 1,2$ and $0 < q < 1$. Thus, vc(M) is the minimum over q of the maximum of the two linear functions $\pi_1(g)$ and $\pi_2(g)$. In Figure 2.2, we graph both of them. Their

2.3. Small Games 53 Figure 2.1. Finding the row value. maximum is indicated with a heavier line and the minimum of this maximum is circled. We have *1(q) = 2q-3(1-q) = bq-3 and *2(q) = -« + (i-«) = -2« + i. The minimum of the maximum occurs where these cross and so we easily compute vc(M) = -1/7, attained at q — (4/7,3/7). Since vr(M) — vc{M), we have indeed solved the game. To solve a 2 x 2 matrix game, it is not really necessary to draw the graph. This follows from: Theorem 2.10. Let —(: J) be α 2x2 matrix and suppose M has no saddle points. Then, (i) the straight lines $\pi_1(\rho)$ and $^(p)$ cross at a value p* of ρ in the open interval (0,1); moreover, $\pi_1(\rho^*) = \pi_2(\rho^*) = v_r(M)$. Also, (ii) the straight lines $\pi_1(g)$ and $\pi_2(g)$ cross at a value q* in the open interval (0,1); moreover, *V) = »V) = t;e(M).

54 2. Two-Person Zero-Sum Games Figure 2.2. Finding the column value. Proof. We prove (i); the proof of (ii) is similar. We must first prove that the two lines cross in the open interval (0,1). Suppose that this is not true. Then one of the two functions must lie above the other. That is, we must have either $\pi_1(\rho) > \pi_2(p)$, $0<p<1$, (2.3) or MP) > *i(p)i $0<p<1$. (2.4) If the first of these inequalities occurs, then we have that c - $\pi_1(0) > \pi_2(0) = d$ and a - $\pi_1(1) > \pi_2(1) = 6$. Thus, column 1 is dominated by column 2. By Exercise (9) of Section 2.2, M has a saddle point. This is a contradiction. If (2.4) holds, then column 2 is dominated by column 1, which is again a contradiction. We have verified that the two lines cross in (0,1). Let p* denote the value of ρ where this happens. Now we show that the maximum of the minimum occurs at p*. To the left of ρ^*, one of the two functions is strictly less than the other. Let us assume that $\pi_1(\rho) < \pi_2(\rho)$ for $0 < \rho < \rho^*$ (the other case is similar). Thus,

taking ρ — 0, we see that c < d. Consideration of the geometry of the situation shows that we only need to prove that $7\Gamma\iota(\rho)$ has positive slope and that $7r2(p)$ has negative slope. It is easily verified

2.3. Small Games 55 that the slope of $7\Gamma\iota(\rho)$ is a — c and that the slope of $^\wedge\{p)$ is 6 — d. First, suppose that a — c < 0. Then a < c. Thus, since c < d, c is a saddle point. This is a contradiction and so a — c > 0. Second, suppose that b — d > 0. Then, since α > c, row 1 dominates row 2. This means, as before, that M has a saddle point, which is a contradiction. Π Let us look at another example. Consider / 1 0 1\ M= -1 0 II. V-l 2 1) We see that there are no saddle points, but that row 2 is dominated. After removing it, we see that column 3 is dominated. Erasing it leaves us with (-!!)■ This has no saddle points and so Theorem 2.10 applies. Now $7\Gamma\iota(\rho) = 2p$ — 1 and $^\wedge(p)$ — —%P + 2. Setting these equal, we solve to get p* — 3/4. The row value is thus $7\Gamma\chi(3/4) = 7\Gamma2(3/4) = 1/2$. As for the column player, we see that $7\Gamma1(д) = q$ and n2(q) — — 3g + 2. Setting these equal, we get q* — 1/2. The column value is then 1/2. For the original 3x3 game, the probabilities of the erased row and column are both 0. The probabilities of the remaining rows and columns are the numbers just computed. We get that the optimal mixed strategy for the row player is (3/4,0,1/4) and the optimal strategy for the column player is (1/2,1/2,0). 2.3.1. 2 X η and m X 2 Games Now we consider 2 χ η matrix games. Here is an example: "■(-J-i-i)· There are no saddle points and no dominated rows or columns. We compute Kj(p) — E((p, 1 — p), j) for j — 1,2,3 and 0 < ρ < 1, and get $\pi\chi(\rho) = 8p - 4$, $\pi2(\rho) = -8p + 4$, $\pi\pi3(p) = 3p - 2$. Then we graph these three linear functions (in Figure 2.3), indicate their minimum by a heavier line, and circle the maximum of this minimum.

56 2. Two-Person Zero-Sum Games Figure 2.3. Row value for a 2 χ 3 game We see from the graph that the maximum of the minimum occurs where $7\Gamma2(p)$ crosses $7\Gamma3(p)$. Setting these equal, we get p* = 6/11 and so vr(M) = π2(6/11) = π3(6/11) = —4/11. To compute the optimal strategy for the column player, we note, from Figure 2.3, that $\pi\iota(\rho^*) > \min 7Tj(p^*)$. j It follows from Theorem 2.9 that column 1 is inactive in any optimal strategy for the column player. In other words, we can assign probability zero to column 1. This leaves us with a 2 x 2 problem. Solving for the column player's strategy gives q = (0,3/11,8/11) as optimal and vc(M) — —4/11. Finally, we consider m χ 2 matrix games. Here is an example: M = -i\ 1 -1 1 0 1/2 V 1/2 0 / There are no saddle points

and no dominated rows or columns. We

2.3. Small Games 57 1.0 0.5 0 -1.0 Figure 2.4. Column value for a 4 χ 2 game. compute 7Гг(д) = E(i, (g, 1 — q)) for г — 1, 2, 3,4 and 0 < q < 1 and get ^(q) = 2q - 1, тг2(д) = -2<j + 1, тг3(д) = -<j/2 + 1/2, тг4(д) = g/2. Then we graph these four linear functions (in Figure 2.4), indicate their maximum by a heavier line and circle the minimum of this maximum. We see that the minimum of the maximum occurs where 7r3(g) crosses 7r4(g). Setting these equal, we get q* = 1/2 and so vc(M) — 1/4. From the graph, we see that the first two rows will be inactive in an optimal strategy for the row player. This leaves us with a 2 χ 2 problem. Solving for the row player's strategy, we get ρ = (0,0,1/2,1/2) and vr(M) — 1/4. Exercises (1) First solve (-..-O- Then suppose that the row player believes that the column player is playing the mixed strategy (1/3,2/3). Is there a strategy for the row player which is a better response to this than the "optimal" one? What is it?

58 2. Two-Person Zero-Sum Gaines (2) Solve Then suppose that the column player has confidential but reliable information to the effect that the row player is playing the mixed strategy (1/2,1/2). How should he play in response? (3) Solve 2 0 λ 2 -1 ;· !) !)■ (6) For the matrix game (-1 1 1/2 \^ 1 -1 -1/2)' compute the value, the optimal strategy for the row player, and two optimal strategies for the column player. (7) How would you solve a 1 X η or m X 1 matrix game? (8) Define a matrix game to be fair if its value is zero. Consider the matrix game (;.;)■ For which values of the parameter o is the game fair? When does it favor the row player (positive value)? When does it favor the column player? (9) Let -(-; :)■ For what values of a is this game fair [see Exercise (8)]? For what values of a does the row player have an advantage (that is, when is the value positive)? For what values of a does the column player have an advantage? (4) Solve -1 1 1 -1 -1 4 -3 3 (5) Solve / 1 I 3 V -2

2.4. Symmetric Games 59 2.4. Symmetric Games A game is symmetric if the players are indistinguishable except for their names. Thus, if the two players interchange their roles in such a game, neither would find it necessary to modify his or her optimal strategy. Two- finger morra is an example of a symmetric game. Chess is not symmetric since one of the two players moves first and the other second. For a matrix game M to be symmetric, it is first necessary that it be square (since the two players have the same number of

pure strategies). Also, such a matrix must have the property that if the row player and column player interchange strategies, then they interchange payoffs. Thus the entry rr_{ij} must be the negative of rr_{ii}. This property is expressed in the following standard definition in linear algebra. Definition 2.8. A square matrix M of size $\eta \chi \eta$ is skew-symmetric if $rr_{ji} = -rr_{ij}$, $1 < i_{jj} < n$. Thus a game is symmetric if and only if its matrix is skew-symmetric. This collision in terminology between linear algebra and game theory is unfortunate but unavoidable. We will use the adjective "symmetric" to modify "game" and "skew-symmetric" to modify "matrix." An easy consequence of the definition is that every diagonal entry in a skew-symmetric matrix is zero (for, тц — —тц). It is intuitively clear that neither player in a symmetric game has an advantage over the other. It follows that the row value must be nonpositive (otherwise, there would be a mixed strategy for the row player such that he would always win). By symmetry, the column value must be nonnegative. We give a formal proof of these facts. THEOREM 2.11. Let M be a symmetric matrix game. Then $vr(M) = -vc(M)$ and $vr(M) < 0 < vc(M)$. Proof. First, if ρ and «fare any mixed strategies, $\eta \eta \eta \eta E(P,Я) = Y^{\wedge}Y^{\wedge}PiQjmj = -Y^{\wedge}Y^{\wedge}PiQjmji = \sim Щ{>}P)'$ i=li-1 j=li=l

60 2. Two-Person Zero-Sum Gaines Then $vr(M)$ — maxmini^]?, q) ρ я = maxmin($-E(q,p)$) ρ я = max($--$max i?($<$f, ρ)) ρ я = $-$mmmaxE(q,p) ρ я = -vc(M). Finally, by Theorem 2.5, $-vc(M) = vr(M) < vc(M)$. Thus, $vc(M) > 0$ and so $vr(M) < 0$. D Corollary 2.12. If the symmetric matrix game M has a solution, then its value is zero. Also, iff is an optimal strategy for the row player, then r is also an optimal strategy for the column player (and vice versa). Proof. If M has a solution, $vr(M) = vc(M)$. The fact that both these quantities are zero is immediate from the theorem. Now maxi?(]?, r) — —min £'(r,p) = $-vr(M) = 0$, ρ ρ and so f is optimal for the column player. П 2.4.1. Solving Symmetric Games We now present a method for solving symmetric matrix games. It has the advantage of being fairly easy to use if the game is small, and the disadvantage that it does not always work. The game of three-finger morra is an example for which the method fails. It is discussed in Chapter 4. Observe that if r is an optimal strategy (for either player, and hence both) for an $\eta \chi \eta$ symmetric matrix game M, then η Ж^) =][{>}rrгy > 0, (2.5) г = 1 for all j, and that some of these inequalities must be equalities [otherwise vr(M) would be positive]. Now, treating the n's as unknowns, we have the equation $n + ... + rn = 1$. (2.6) If we then choose $n-1$ of the inequalities (2.5)

and set them to be equalities, we have a total of η equations. The idea then is to solve, if possible, this

2.4. Symmetric Games

61 system of equations for the η unknowns to get a solution vector r. If each component r_i of this solution vector is nonnegative, and if the remaining one of the inequalities (2.5) is valid for the γ,-'s, then we have found an optimal strategy. Consider an example: M This is clearly a skew-symmetric matrix. The inequalities (2.5) are, for $j = 1,2,3$, respectively, $r_2 - 2r_3 > 0$, $-r_1 + 3r_3 > 0$, $2r_1 - 3r_2 > 0$. Arbitrarily set the first two to be equations. These, together with (2.6), give us the system of equations $r_2 - 2r_3 = 0$, $-r_1 + 3r_3 = 0$. This system is easily solved to give: $r_1 = 1/2$, $r_2 = 1/3$, $r_3 = 1/6$. These are all nonnegative, and the third inequality (the one not used to solve for the r_2-'s) is valid. Thus we have solved the game. The probability vector $r = (1/2, 1/3, 1/6)$ is an optimal strategy for both players and the value of the game is zero. The reader should verify that, in this case, setting any two of the three inequalities to zero gives this same solution. Here is a second example: M The inequalities (2.5) are $r_2 + r_4 > 0$, $-r_1 + r_3 - 2r_4 > 0$, $-r_2 + 2r_4 > 0$, $-r_1 + 2r_2 - 2r_3 > 0$. Before jumping into a calculation, let us think a bit. The first inequality has no negative coefficients. Thus, since a valid solution requires nonnegativity of the r's, setting it to be an equality would imply that $r_2 = r_4 = 0$.

62 2. Two-Person Zero-Sum Games $\begin{pmatrix} 0 & 1 & 0 & 1 \\ -1 & 0 & 2 & 0 \\ 0 & -2 & 0 & 2 \\ 1 & 0 & -2 & 0 \end{pmatrix}$ That, together with the fourth inequality, would imply that r_1 and r_3 are zero also. This is impossible since the r_i's must sum to one. Thus, the only possible way to proceed is to set the last three inequalities to equalities. These, together with (2.6), give us a system of four equations in four unknowns. Solving it gives $r = (0, 2/5, 2/5, 1/5)$. These are all nonnegative and satisfy the inequality not used in solving for them. For a third example, consider M = Using any three of the four inequalities leads to the parameterized set of solutions: $r_1 = 2/3 - 4r_4/3$, $r_2 = r_4$, $r_3 = 1/3 - 2r_4/3$, $r_4 = r_4$. Substituting values of r_4 leads to infinitely many valid optimal strategies. For example, $(2/3, 0, 1/3, 0)$ and $(0, 1/2, 0, 1/2)$ are both solutions. On the other hand, we also get $(2, -1, 1, -1)$, which is invalid. EXAMPLE 2.1. The game of scissors-paper-stone is played as follows. Each of the two players simultaneously makes a gesture indicating one of the three objects in the name of the game (a closed fist for "stone" etc.). If they choose the same object, the game is a draw. Otherwise,

the winner is decided by the rules: "Scissors cut paper, paper covers stone, stone breaks scissors." The payoff is +1 for a win and — 1 for a loss. If the three pure strategies are listed in the order given in the name of the game, then the payoff matrix is 0 -1 1 1 0 -1 -1 1 0 The game is certainly symmetric. The solution is left for Exercise (3). Exercises (1) Solve 0 1 -2 1 -1 0 1 1 2 -1 0 -1 -1 -1 1 0

2.4. Symmetric Gaines 63 (2) Solve (3) Solve scissors-paper-stone. (4) Solve / ° -1 1 0 -2 -1 3 0 1 0 , -i/ (5) Solve (6) Find at least two optimal mixed strategies for two-finger morra. (7) The game of two-finger morrette is played just like two-finger morra except that the payoff to the winner is always +1. Prove that it has infinitely many optimal strategies.

3 Linear Programming The subject of this chapter is the part of the field of linear programming which we will need later in the book. We will use it to solve matrix games. In addition, the theoretical aspects of linear programming developed here will, when applied to game theory, give us the minimax theorem and other theorems about matrix games. There are many good books which discuss linear programming in greater detail. See, for example, [Chv83], [Str89], or [Lue84]. 3.1. Primal and Dual Problems A linear programming problem consists of a linear real-valued objective function $w\{x)$ $= c\X\ + c2X2 H l\cdot cnxn + d$, of η variables, which is either to be maximized or minimized subject to a finite set of linear constraints. Each constraint has either the form $\alpha\chi\chi\iota + a2X2 H l\cdot anxn < 6$, or the form $\alpha\chi\chi\iota + a2X2 H l\cdot anxn >$ b. We will sometimes abbreviate the phrase linear programming problem by LP problem.

66 3. Linear Programming Thus, the problem is to maximize or minimize $w(x)$ over the set of all vectors χ which satisfy all the constraints. A solution consists of this maximum or minimum value, together with a vector χ at which it is attained. In applications, the objective function often represents a quantity like profit or cost, while the variables represent quantities of various commodities. Here is a simple example of a linear programming problem: maximize — $3\#i + 2x2 + x_3$ (3.1) subject to x\, a?2, #з > 0 Xl + %2 +#3 < 1. It is a maximization problem with three unknowns and a total of four constraints. The constraints $\#b\#2)\#3 > 0$ are called positivity constraints. In the great majority of all linear programming problems (including those arising from

game theory), all the unknowns are naturally nonnegative. For this reason, we restrict ourselves to such problems; positivity constraints will henceforth be understood and not written. Thus, the problem just stated becomes maximize — $3\#i + 2\#2 + \#_3$ (3-2) subject to $\chi_1 + \#2 + \#_3 < 1\cdot$ We mention also that there is no loss of generality in restricting ourselves to problems in which all unknowns are constrained to be nonnegative. If a problem contains an unknown which is allowed to be negative, then we can replace it throughout by the difference of two new unknowns which are constrained to be nonnegative. It should also be mentioned that equality constraints frequently arise. Such a constraint has the form a\X\ + o&2#2 + 1" anXn — b. This one constraint can be replaced by the equivalent pair of inequality constraints a\x\ H h anxn < 6, and a\x\ -\ V anxn > b. Thus, the class of linear programming problems we are considering is really very general. A vector χ which satisfies all the constraints of a given linear programming problem is called a feasible vector for that problem. An LP problem

3.1. Primal and Dual Problems 67 is said to be feasible if there exists at least one feasible vector for it. Problem (3.2) is feasible since χ — (0,0,0) is a feasible vector. On the other hand, the LP problem minimize $\text{Ъ}x\ + X2$ (3-3) subject to $x\ + X2 < — 1$ is infeasible because $x\, x^\wedge > 0$ and so their sum cannot be negative. Note that the objective function is irrelevant for feasibility. A feasible vector χ is optimal if the objective function attains its maximum (or minimum) at x. Some feasible problems have no solution because the objective function is unbounded. For example, the problem maximize $w(x) = —x\ + 2x2$ (3.4) subject to $x\ — X2 < 2$ is unbounded because the vector χ — (0, a?) is feasible for every $\chi > 0$ and $\lim w((Q,x))$ — -boo. X—ЮO It will be proved later that every bounded feasible problem has a solution. If w(x) is the objective function of a feasible maximization problem, we will use the notation maxu;(#) to denote its maximum value over the set of feasible vectors. If the problem is unbounded, we write maxu;(£) = -foo. For a minimization problem, we use the corresponding notation minw(x). If the problem is unbounded, we write minw(x) = —oo. 3.1.1. Primal Problems and Their Duals The following definition gives an important special type of problem. Definition 3.1. A linear programming problem is said to be primal if it has the form maximize f(x) — \\ cj xj + d η subject to 22,aiix3 ~ ^' 1 ~ * ~ m' where A — (a>ij) is an m χ η coefficient matrix, c is an η-tuple of numbers, d is a constant, and 6 is an m-tuple of numbers.

68 3. Linear Programming Problems (3.2) and (3.4) are primal but Problem (3.3) is not. Here is another example. maximize f(x\, #2> #з> %a) — x\ — x^ (3-5) subject to x\ + #2 + #з — #4 < 2 -a?i - 3a?2 ~ #з + 2a?4 < —1. It is feasible because, for example, χ — (0, 0,1, 0) is a feasible vector. It should be noticed that any linear programming problem can be converted into an equivalent problem which is primal. The method for doing so is briefly stated as follows: If the objective function is to be minimized, replace it with its negative; if there is a constraint in which the left-hand side is greater than or equal to the right-hand side, take negatives on both sides. A primal problem is really half of a pair of problems. The other half is called the dual problem. We have the following: Definition 3.2. Consider the primal problem η maximize f(x) — Yj cjXj + d η subject to /_.aijxj < &í> 1 < i < m. i=i The dual problem corresponding to this primal problem is m minimize g(y) — ^ &,·?/,· + d m subject to 2_] aijVi ^ cj > 1 5: 3 5: n· г = 1 Thus, the coefficients in the dual objective function are the right-hand sides of the primal constraints; the m-tuple of coefficients in the jth dual constraint is the jth column of the coefficient matrix A\ the right-hand sides of the dual constraints are the coefficients in the primal objective function. Note that we can construct the primal from the dual as easily as the dual from the primal. As stated before, the two problems form a dual/primal pair—mathematically, they are equally important. Their interpretations may make one of them more important to us than the other. For Problem (3.2), the coefficient matrix is A=(l 1 1).

3.1. Primal and Dual Problems 69 Also c=(-3 2 1), 6=(1). Thus, the corresponding dual problem is minimize 2/1 (3-6) subject to 2/1 > —3 Vi >2 У1>1. The dual problem corresponding to Problem (3.5) is minimize 2y\ — 2/2 (3-7) subject to 2/i — 2/2 > 1 2/1 - 32/2 > 0 2/1 ~ 2/2 > 0 -2/1 +22/2 >-l. This problem is feasible since y — (1,0) is a feasible vector. The solutions of the dual and primal problems are closely related. The following theorem reveals part of the relationship. THEOREM 3.1. If both the primal problem and its corresponding dual are feasible, then № < </(£>, for any feasible vector χ of the primal and any feasible vector y of the dual. Thus, max f(x) < min g(y). Proof. Compute η η m < ΣΣ i=lг=1 τη η í=1 j=l m < Σ6<№+<í 8=1 = 9{У)·

70 3. Linear Programming The second inequality is obtained by maximizing over χ on the left and then minimizing over y on the right. Π From this, we

obtain the following: Corollary 3.2. The following statements hold. (1) If both the 'primal and its corresponding dual are feasible, then both are bounded. (2) // the primal is unbounded, then the dual is infeasible. (3) // the dual is unbounded, then the primal is infeasible. Proof. For a fixed feasible vector yioi the dual, f(x) is bounded above by g(y), and thus the primal problem is bounded. Similarly, the dual problem is bounded. The other two statements follow easily from the first. Π As an example, we noted that both Problem (3.5) and its dual, Problem (3.7), are feasible. Thus, they are both bounded. For a second example, note that Problem (3.6) is easy to solve. It has only one unknown and the second constraint implies the other two. Hence, the problem calls for us to minimize y\ subject only to the constraint y\ > 2. Obviously, 2 is the minimum value of the objective function and it is attained for y\ = 2. Now, Problem (3.6) is the dual corresponding to Problem (3.2). From Theorem 3.1 and its corollary, we conclude that Problem (3.2) is bounded and that its objective function is bounded above by 2 on the set of feasible vectors. We can go one step further. The vector χ — (0,1,0) is a feasible vector for Problem (3.2) and the objective function has value 2 at this feasible vector. Since the objective function can be no larger than 2, we have solved Problem (3.2). Exercises (1) Consider the problem minimize — Ђx\ + 2x2 — xг subject to x\ + X2 + #3 = 3 xi ~ X2 > 0· Convert it into an equivalent primal problem. (2) Consider the problem maximize x\ — X2 — x3 + xa subject to x\ — xz < 0 X2 — a?4 < 1. Is it feasible? Is it bounded?

3.2. Basic Forms and Pivots 71 (3) Consider the dual problem minimize y\ - 2y2 + 2/3 subject to —2/1 —2/2+2/3 > -1. Is it feasible? Is it bounded? Write down the corresponding primal problem. Is the primal feasible? Is the primal bounded? (4) Solve the problem maximize x\ + X2 + x3 subject to x\ + 2x2 + 3#3 < 10. (5) Consider the problem maximize — x\ + 2#2 — 3#3 + Ax± subject to #3 — x± < 0 a?! — 2#3 < 1 2#2 + ^4 5; 3 —#1 + 3#2 < 5. Is it feasible? Is it bounded? (6) Consider the problem maximize x\ — 3x2 subject to x\ + X2 > 1 371 "~ ^2 5: 1 ^1 — #2 > —1· Is it feasible? Is it bounded? 3.2. Basic Forms and Pivots Consider a primal LP problem: maximize f(x) — V^ CjXj + d η subject to 2_^aiJxJ — ^»» 1 < г < m. i=i For each of the m constraints, define a slack variable by η %n+i — "г / v QijEj · i=i Thus, the m slack variables depend on the unknowns #1,..., xn, and each xn+i measures the extent to which the left-hand side of constraint number

72 3. Linear Programming i is less than the right-hand side. Notice also that $\chi =$ ($\chi_1,\ldots$, #n) is a feasible vector if and only if xk > 0, 1 < к < n + m. Using these slack variables, we rewrite the problem in a way which appears strange at first but which turns out later to be convenient. It is η maximize f(x) = Y^ CjXj + d (3.8) η subject to /Ja*j#7 ~ &* — —#n+*, 1 < i < m. i=i This form of the primal problem is called the (primal) basic form. We emphasize that it is mathematically equivalent to the original primal form in the sense that solving one is the same as solving the other. Let us write out a basic form for Problem (3.5). There are two constraints and so we have two slack variables, x$ and xq. The basic form is maximize x\ — #4 (3.9) subject to χ_1 + #2 + #3 — #4 — 2 = —x$ — X\ — 3#2 — χ_3 + 2#4 + 1 = — Xq. In the context of basic forms, the variables on the right-hand side of the constraints (that is, xn+i, $\cdot\,\cdot\,\cdot$, Xn+m) are called basic (or, sometimes, dependent) variables while the other variables are called nonbasic (or independent) variables. The set of basic variables is called the basis corresponding to the basic form. 3.2.1. Pivots With the use of a little simple algebra, we can compute a new basic form from an old one. To do so, choose, in Basic Form (3.8), any nonzero coefficient a^u where 1 < к < m and 1 < / < n. Thus aь\ is the coefficient of x\ in the equation for —xn+k. Since aьı Ф 0, we can solve the equation to get x\ in terms of xn+k and the remaining nonbasic variables (other than x\). Then we substitute for x\ in the other constraints and in the objective function. The result is another basic form, differing from the first in that a basic and nonbasic variable have traded places. The operation we have just carried out is called a pivot, and we say that we pivot on the coefficient a*/. It is important to realize that the new basic form is again mathematically equivalent to the old one, and thus to the original primal problem.

3.2. Basic Forms and Pivots 73 Let us pivot on the coefficient of #3 in the equation for — xq in Basic Form (3.9). Solving for #3 in that equation gives us #3 = -a?i - 3#2 + #6 + 2a?4 + 1. Substituting this expression for #3 in the equation for — #5 and combining terms give us -2a?2 + #6 + #4 - 1 = —a?5· The formula for the objective function in Basic Form (3.9) does not contain xз and so is not changed. The new basic form is maximize χ_1 — χ_4 (3.10) subject to — 2#2 + #6 + %4 — 1 = —#5 x\ + 3#2 — #6 — 2#4 — 1 = — xз- We carry out one more pivot, this time on the coefficient of x\ in the equation for — #3. The result is maximize — X3 — 3#2 + #6 + #4 + 1 (3.11) subject to — 2a?2 + #6 + #4 — 1 = —#5 X3 + 3#2 "~ #6 — 2x4 — 1 = — x\. Thus the

basis corresponding to Basic Form (3.9) is $\{x^{\wedge\wedge}xq\}$, the basis corresponding to Basic Form (3.10) is $\{\#5, \#3\}$, and the basis corresponding to Basic Form (3.11) is $\{\#5,\#i\}$. It is clear that there are, in general, many basic forms for each primal problem. Any finite sequence of pivots transforms one basic form into another. In the general case where there are η unknowns in the primal problem and m constraints, the total number of possible bases corresponding to basic forms is the binomial coefficient $fm + n\backslash \_ (m + n)! \backslash m J n\backslash m\backslash$ For example, if $\eta = 4$ and $m = 2$, this number is 15. The actual number of basic forms might be less than this quantity. Given a basic form, the basic variables and the objective function are expressed in terms of the nonbasic variables. If arbitrary values are assigned to the nonbasics, then the values of the basic variables are determined, and so we obtain an $(n + m)$-tuple ($\#1,...,$ xn+m). In case each $Xk > 0$, we say that this $(n + m)$-tuple is feasible. If $(a?i,.. .,\#n+m)$ is feasible, then the vector ($\#1,...,$ xn) is, as already noted, a feasible vector. The most important case of this computation occurs when the value zero is assigned to each nonbasic. We have

74 3. Linear Programming Definition 3.3. Given a basic form for a primal problem with η unknowns and m constraints, the corresponding basic solution is the $(m + n)$- tuple obtained by setting the η nonbasic variables to zero and then solving for the m basic variables from the constraints. For example, in Basic Form (3.9), the basic solution is $(0,0,0,0,2, —1)$. Note that the value of a basic variable in a basic solution is just the negative of the constant term in its constraint equation. In Basic Forms (3.10) and (3.11), the basic solutions are $(0,0,1,0,1,0)$ and $(1,0,0,0,1,0)$, respectively. Also note that the value of the objective function at the basic solution is the constant term in its equation. Definition 3.4. A basic form is feasible if its basic solution is feasible. A feasible basic solution is optimal if the value of the objective function at the basic solution is a maximum. In this case, the basic form is also called optimal. We know that a vector ($\#i,.. . ,\#n$) is a feasible vector for the primal problem if and only if $[x\backslash,..., $ xn+m) is feasible. It will be proved later that if a problem has an optimal feasible vector then it has an optimal feasible basic solution. We see, for example, that Basic Form (3.9) is infeasible while Basic Forms (3.10) and (3.11) are both feasible. In the next section, we shall describe the simplex algorithm, which is a numerical method for solving primal problems. The idea of it is to start with a feasible basic form and then to pivot in such a way that the new basic form is still feasible and gives a larger value of the

objective function. This process is repeated until the maximum is reached. The algorithm includes a simple method for deciding which coefficient to pivot on next. In order to recognize feasible and optimal basic forms, we need Theorem 3.3. The following hold: (1) A basic form is feasible if and only if the constant term in each constraint is nonpositive. (2) Suppose a basic form is feasible. If each coefficient (of a nonbasic variable) in the equation for the objective function is nonpositive, then the basic form is optimal. Proof. (1) This is clear since the value of a basic variable in the basic solution is the negative of the constant term in its equation. (2) We can write the objective function as $m+n$ $f\{£) = Y^CkXk + d, k = l$

3.2. Basic Forms and Pivots 75 where $Ck = 0$ if Xk is basic (since basic variables do not appear in the equation for the objective function). Our hypothesis is that $Ck < 0$ for each k. Now, if £is the basic solution, we have $f\{z) = d$. Then if χ is any feasible vector, we have, since each $Xk > 0$, $f(x) =$ $£c*z*+d<d=/(f)$. $k = l$ Thus, max $/(\#) = f(z)$. D We can illustrate the second part of the theorem by doing a pivot on Basic Form (3.11). In fact, if we pivot on the coefficient of #4 in the equation for — #5, we get maximize — #3 — #2 — #5 + 2 (3.12) subject to — 2#2 + #6 + #5 — 1 = — #4 x₃ — #2 + #6 + 2#5 — 3 = —x\. According to the theorem, this basic form is both feasible and optimal and thus we have solved Problem (3.5). To summarize, the solution is max $f(x) = 2$, $x\ — 3$, $x2 — x^ — 0$, $x± — 1$. 3.2.2. Dual Basic Forms Consider a dual problem minimize $Y^ \&,·?/,· + d *=i$ subject to $Y^a\#2/* > ty$, $1 < .; < n$. $2 = 1$ For each of the η constraints, define a surplus variable by m i=l Then $ym+y$ measures the difference between the left and right sides of the jth constraint. We see that $(2/1,.. . ,ym)$ is a feasible vector for the dual

76 3. Linear Programming problem if and only if $yk > 0$ for $1 < κ < m + n$. We rewrite the dual problem as m minimize $/]bjyj + d$ $i = 1$ m subject to $ym+j \sim ^2 aijVi - ch ^<3<n·$ $i = 1$ This is a dual basic form. The variables on the left-hand side of the equality signs are basic and the others are nonbasic. Just as in the primal case, we can pivot so as to interchange a basic and a nonbasic variable and produce a new dual basic form. Let us write down the dual basic form for Problem (3.7). It is minimize $2yi — y<i$ (3.13) subject to $y₃ = yi - 2/2 - 1$ $2/4 = 2/1 \sim 3y2$ $2/5 = 2/1- 2/2$ $2/6 = -2/1 +2j/2 +l$. If we pivot on the coefficient of y\ in the equation for y₃, we get minimize $2y3 + 2/2+2$ (3 14) subject to $2/1=2/3+2/2 + 1$ $2/4 = 2/3- 22/2 + 1$ $2/5 = 2/3 + 1$ $2/6 = -2/3 + 2/2-$ Given a

dual basic form, the definition of feasible $(n+m)$-tuple is the same as in the case of primal basic forms. Also, the basic solution corresponding to a dual basic form is obtained by setting the nonbasic variables equal to zero, and then solving for the basic variables. Clearly, the value of a basic variable in a basic solution is simply the constant term in its equation. We define a basic solution $(2/1,..., 2/m+n)$ (and its dual basic form) to be feasible if $yk > 0$ for $1 < \kappa < \tau + n$. The value of the objective function at the basic solution is clearly equal to the constant term in its equation. A feasible basic solution (and its dual basic form) is optimal if the value of the objective function is a minimum. In Dual Basic Form (3.13), the basic solution is $(0,0,-1,0,0,1)$; it is infeasible. In Dual Basic Form (3.14), the basic solution is $(1,0,0,1,1,0)$; it is feasible.

3.2. Basic Forms and Pivots 77 A theorem analogous to Theorem 3.3 holds for dual basic forms. The proof is omitted since it is very similar to the one for Theorem 3.3. Theorem 3.4. The following are true: (1) A dual basic form is feasible if and only the constant term in each constraint equation is nonnegative. (2) Suppose a dual basic form is feasible. If each coefficient (of a non- basic variable) in the equation for the objective function is nonnegative then the dual basic form is optimal. This theorem shows that Dual Basic Form (3.14) is both feasible and optimal. Thus the solution to Problem (3.7) is $mmg(y) = 2, */i = 1, y2 = 0$. Exercises (1) Consider the problem maximize $2\chi\chi + 3\#2$ subject to $x\backslash + X2 < 10$ $X\backslash - X2 < 2$. Write down all the basic forms. (2) Consider the primal problem maximize $x\backslash - 2x\%$ subject to $x\backslash < 2$ $-X\backslash + X2 < 3$. Write down all the basic forms. Which are feasible? Which is optimal? (3) Consider the dual problem minimize $y\backslash + 2/2$ subject to $2y\backslash + 3i/2 > 5$ $-2/1 > -6$. Write down all the dual basic forms. Which are feasible? Which is optimal? (4) Consider the problem minimize $- y\backslash + 2Y2 + 2/3$ subject to $2/i - 2/2 > 0$ $-2/1 + 2/3 > 1$ $-2/2 + 2j/3 > -1$. Carry out one pivot to reach an optimal basic form.

78 3. Linear Programming (5) Solve the problem maximize $x\backslash - X2 + 2\#_3$ subject to $x\backslash + X2 + \#3 < 3$. 3.3. The Simplex Algorithm In order to discuss the simplex algorithm, it is convenient to introduce a more compact way of writing a basic form. It is called a tableau and is discussed next. 3.3.1. Tableaus Let us consider an example, namely Basic Form (3.11). The tableau corresponding to it is $\#3$ $X2$ $\#6$ $\#4$ 0-211 1 3-1-2 -1-3 1 1 -1 1 1 -1 = $-x_b$ = $-X_\iota$ = f This tableau is essentially a 3×5 matrix with labels on the rows and columns. In the

general case, where the primal problem has η unknowns and m constraints, the size would be $(m+1) \times (n+1)$. In our example, the first four columns are labeled by the nonbasic variables, #3, #2> #6, $\chi 4$. The fifth column is labeled — 1. This label will be explained shortly. The first two rows are labeled on the right by the basic variables "= — #5,= — x\" The last row is labeled by the objective function "= /." Each row is to be read as an equation—the numbers in the row are multiplied by the labels at the tops of the corresponding columns, and these products are added. Thus, the first row represents the equation $(0)(*8) + (-2)(*2) + (1)(*\beta) + (1)(*4) + (1)(-1) =$ -хь- This simplifies to $-2x2 + \#6 + \#4 - 1 =$ -a?5i which is the first constraint in Basic Form (3.11). The second row gives the second constraint and the bottom row gives the objective function. The basic solution is easily read off the tableau. In fact, the values of the basic variables are simply the numbers in the right-hand column (excluding the bottom row); the value of the objective function at the basic solution is the negative of the number in the lower right-hand corner. Theorem 3.3 gives conditions on the coefficients in a basic form to ensure feasibility and optimality. This result can be easily translated into the context of tableaus.

3.3. The Simplex Algorithm 79 Theorem 3.5. Consider a tableau of size $(m+1) \chi (n+1)$. Then we have: (1) The tableau is feasible if and only if each entry in the right-hand column (not counting the bottom one) is nonnegative. (2) Suppose the tableau is feasible. If each entry in the bottom row (not counting the right-hand one) is nonpositive, then the tableau is optimal. The result of pivoting on the coefficient of $\chi 4$ in the equation for —#5 in Basic Form (3.11) is given in Basic Form (3.12). The corresponding tableau is X3 0 1 -1 X2 -2 -1 -1 xe 1 1 0 *5 1 2 -1 -1 1 3 -2 = -X4 = -Xι = f (3.16) The effect of a pivot on the labels is simply to interchange a basic and a nonbasic variable. In the example, basic variable x$ and nonbasic variable $\chi 4$ trade places. The next theorem shows how to compute the numbers in the new tableau from those in the old one. Theorem 3.6. Let T be a tableau of size $(m+1) \chi (n+1)$. Suppose a pivot is carried out on entry t^i. In other words, the pivot is on the coefficient of the nonbasic variable xp (labeling column I) in the constraint equation for the basic variable xq (labeling row k). Of course, $1 < \kappa < m$, $1 < / < n$, and tki φ 0. Let T' be the tableau resulting from the pivot. Then, if $1 < i < m+1$ and I $< j < n + I$, we have: (1) t'kl = VUi. (2) Ifj?l,thent'kj=tkj/tkl. (3) //!#*, <ien <{, =-<«/<«· (4) If г ф к, and j φ I, then '« = Proof. We prove only (4), since the other cases are similar but easier. To do this, let a be the label on column j.

Thus a is either one of the nonbasic variables other than xP) or the constant —1. Also, let β be the label on row i. Thus β is either the negative of a basic variable other than xq) or the objective function. Now, solving for xp from the constraint equation for -xq) we have hi J t_ hi + Γ, (3.17)

80 3. Linear Programming where Γ is a sum of terms not involving a. Then, from row г of T, we have β = Uixp+Uja + A, where Δ is a sum of terms not involving xp or a. Substituting from (3.17), we have β = Ui[- (γ1 J a + Γ] + t{ja + Δ. Thus, tjjtkl tjl*kj ,α + Θ, tkl where Θ consists of terms not involving a. Since t^· is the coefficient of a in the equation for /?, the proof is complete. Π Let us check the theorem on the example already discussed. The pivot carried out on Tableau (3.15) to produce Tableau (3.16) is on entry t\4 = I. From parts (1) and (2) of the theorem, we see that row 1 of Tableau (3.16) should be identical to row 1 of Tableau (3.15). Indeed, it is. The entries in column 4 of Tableau (3.16), other than ^14, should be the negatives of the entries in Tableau (3.15). This comes from part (3) of the theorem. Again, this is true. To check part (4) of the theorem, consider entry t'32 [of Tableau (3.16)]. It should be ,, _ *32*14 ~ *12*34 _ (-3)(1) - (-2)(1) _ /14 1 This is correct. We summarize the rules for pivoting in a tableau in the following algorithm. Algorithm 3.1 (The Pivoting Algorithm). We are given a tableau T of size (m + 1) χ (η + 1), and an entry tki Φ 0. We are to compute a new tableau T1 of the same size as T. (1) Interchange the labels on column I and row κ. The other labels on T' are the same as on T. (2) t'hl = l/tkl. (3) Iííφ I, ihent'kj=tkj/tki. (4) If г φ κ, then t'u - -*,·//**/· (5) If г φ κ and j φ l} then ,1 tij*kl til*kj

3.3. The Simplex Algorithm 81 3.3.2. The Simplex Algorithm Before stating the simplex algorithm, we need one more theorem about tableaus. Theorem 3.7. Consider a feasible tableau for a primal problem with η unknowns and m constraints. Suppose that one of the first η entries in the bottom row of the tableau is positive, but that all the other entries in the column containing that entry are nonpositive. Then the problem is unbounded. Proof. Suppose that it is column j which has a positive entry at the bottom and which has all nonpositive entries otherwise. Then column j is labeled by some nonbasic variable, say xq. Given any positive number #, define an (m + n)-tuple as follows: To xq) assign the value x; to a nonbasic χι, I φ g, assign the value 0; then solve for the basics from their constraint equations. Then each basic

variable in this $(m+n)$-tuple is nonnegative because of the condition on the entries in column j. Since the coefficient of xq in the objective function is positive, the objective function has limit $+00$ as $\chi \longrightarrow \blacksquare 00$. Thus, the objective function is unbounded. П For example, consider the problem maximize χι (3.18) subject to $\chi\iota — X2 < 1$. The initial tableau is X\ Xl 1 -1 1 0 -1 1 0 = - Ж3 = / This is feasible and nonoptimal. The theorem does not indicate that it is unbounded. We now pivot on the coefficient of $\chi \iota$ in the constraint equation for —X3. The new tableau is X3 #2 1 -1 -1 1 -1 1 -1 = -Xι = / This is feasible. The theorem tells us that the problem is unbounded. Algorithm 3.2 (The Simplex Algorithm). We are given a feasible tableau T of size $(m+1) \chi (n+1)$. (1) V Wij $< °$ for^<3<n, then STOP (the tableau is optimal).

82 3. Linear Programming (2) Choose any I with $1 < / < \eta$ such that tfm+i,/ > 0. (The nonbasic variable labeling column I is called the entering variable.) (3)// in < 0 for $1 < r < m$, STOP (the problem is unbounded). (4) Choose any к with $1 < к < m$ such that tki > 0 and h,n+i hi min < —— : $1 < rI < m$ and in > 0 > . I hi J (The basic variable labeling row к is called the leaving variable.) (5) Pivot on tki and replace T by the tableau resulting from this pivot. (Entry tki ^ called the pivot entry.j (6) Go to Step (1). In the next theorem, we verify that the algorithm produces a sequence of feasible tableaus such that the values of the objective function are non- decreasing. Theorem 3.8. The following hold for the simplex algorithm: (1) After each pivot [Step (5)], the new tableau is feasible. (2) After each pivot, the value of the objective function for the new tableau is greater than or equal to that for the previous tableau. Proof. (1) We are pivoting on entry tki of T to produce V. We must verify that t'in+1 > 0 for 1 $< r < τ$. This is clearly true for i = k, so assume г ф к. Then ,/ _ ij,n+ltkl — ^7^Jb,n+l Now, we know that 2г>+ъ tkh h,n+i are all nonnegative. There are two possibilities. If tu < 0 then t\ n+1 > 0. On the other hand, if tn > 0 then, by the choice of fc, tki ~~ tu Multiplying both sides by tkiUi, we get tiltk,n+l < ti}n+itkl. This implies that

3.3. The Simplex Algorithm 83 (2) The value of the objective function for T is —tfm+i,n+i; f°r Tf, it is-4+i,n+i· Now hi hi Taking negatives on both sides completes the proof. П There are two things about the algorithm which will be cleared up in the next section. The first is the problem of producing an initial feasible tableau (without which the algorithm cannot be started). The second is the problem of whether the algorithm eventually stops with the right answer. It

is clear that, if it stops, either the problem is unbounded or the solution has been reached. For an example, consider the primal problem, maximize $\chi\chi$ — $x2 + x3 \sim 2x^{\wedge}$ (3.19) subject to $x\backslash + a?2 + x3 + \#4 < 4$ $-X\chi + X3 < 0$ $x4 — \chi\alpha < 1$. The initial tableau is then X\ 1 -1 0 1 X2 1 0 1 -1 x3 1 *1 0 1 X4 1 0 -1 -2 -1 4 0 1 0 = -X5 = -X6 = -x7 = f This tableau is feasible but not optimal; unboundedness is not indicated. The algorithm allows two choices of entering variable: x\ or #3. We choose X3. In the column labeled by #3, there are two positive entries above the bottom row. The minimum ratio is zero and is achieved by the entry in the second row. Thus, X3 will enter the basis and xq will leave it. The pivot entry is marked with an "*." We pivot to get the tableau xi *2 -1 0 2 X2 1 0 1 -1 X6 -1 1 0 -1 x4 1 0 -1 -2 -1 4 0 1 0 = -xs = -x3 = -x7 = f This tableau is feasible (as it must be) but not optimal; unboundedness is not indicated. There is only one possible pivot entry: x\ must enter the basis and x$ must leave it. The pivot entry is again marked. We pivot to get

84 3. Linear Programming X5 X2 X6 Xa 1/2 1/2 -1/2 1/2 1/2 1/2 1/2 1/2 0 1 0 -1 -1 -2 0 -3 -1 2 2 1 -4 = —Xi = -x3 = -X7 = f This tableau is optimal. We get the solution max/(f) = 4, χ_1 — 2, #2 = 0, #3 = 2, #4 = 0. In fact, it has more than one solution. To see this, pivot in the last tableau so as to interchange xq and #3. Exercises (1) Solve maximize x\ subject to $2x\backslash — X2 < 2$ $x2 + {}^{\wedge}3 < 2$. (2) Solve maximize $\alpha?_1 + 2\#2 + 3\#3$ subject to $x\backslash + 072 + {}^{\wedge}3 < 3$ $—\#1 + 072 < 0$. (3) Solve maximize #1 + #2 + XZ subject to $x\backslash + \#3 < 3$ (4) Solve maximize $x\backslash + x\% — xz — 2x\pm$ subject to $x\backslash + 2a;2 + 2\#3 + x\pm < 1Q$ $xi — x3 < 0$ $*2 < 2$. (5) Verify that Problem (3.19) can be solved in one pivot. (6) Find a condition on an optimal tableau which guarantees that there are at least two solutions.

3.4. Avoiding Cycles and Achieving Feasibility 85 3.4. Avoiding Cycles and Achieving Feasibility As the simplex algorithm was stated in the previous section, it is possible for it to go into an infinite loop and never deliver a solution to the problem. There are several methods for avoiding this difficulty. We will present one of them. After that, we will discuss the question of how to find an initial feasible tableau so that the algorithm can be started. 3.4.1. Degeneracy and Cycles First, we have the following: Definition 3.5. Let T be an $(m + 1) \chi (\eta + 1)$ tableau. The pivot on entry iki is said to be degenerate if $i_{k,n+i} = 0$. The first pivot we carried out in solving Problem (3.19) was degenerate. The important fact about degenerate pivots is given in the following theorem. The proof should be obvious. Theorem 3.9. If the tableauV

is obtained from T by a degenerate pivot then $t_{i,n+1} = \hat{t}_{r',n+1})$ for all r with $1 < i < m + 1$. In other words, the last column of T' is identical to the last column of T. Corollary 3.10. // T' is obtained from T by a degenerate pivot, then the basic solution corresponding to T' is equal to the basic solution corresponding to T. Thus, the objective function has the same value for the two basic solutions. Proof. Since the pivot is degenerate, the value of the leaving variable in T is zero. This is also its value in T1 since it is nonbasic. The value of the entering variable remains zero; the values of the others clearly do not change. Π Thus, doing a degenerate pivot while carrying out the simplex algorithm causes no improvement in the objective function. This is usually harmless [as Problem (3.19) shows]. Degenerate pivots are eventually followed by nondegenerate ones, and the progression toward a solution continues. There is, however, a rare event which must be allowed for in the theory. It is called cycling and can put the algorithm into an infinite loop. A cycle consists of a sequence of feasible tableaus $T°, T1,..., Tk$, such that each is obtained from the previous one by a degenerate pivot chosen according to the algorithm, and such that Tk has the same basis as $T°$. Thus, Tk can differ from $T°$ only in the order of its rows and columns. Now, if the

86 3. Linear Programming pivot in Tk is chosen just as the one in $T°$ was, the result can be that the computation repeats the list of tableaus ad infinitum. As we have said, cycling is rare. There are examples involving reasonably small problems (an example is given in [Str89]), but only a few have have ever arisen which were not intentionally made up to illustrate the phenomenon. For the theory, however, cycling is a serious problem because it blocks a proof that the algorithm stops. Fortunately, there is a simple way to avoid the problem. It first appeared in [Bla77] and is an amendment to the simplex algorithm. It is the following: Bland's Anticycling Rule Modify the simplex algorithm so that, whenever there is a choice of entering or leaving variable, we always choose the one with smallest subscript. If we had applied Bland's rule in Problem (3.19), the first pivot would have interchanged χ_1 and x$ instead of xs and x$. Thus, the degenerate pivot would have been eliminated. In general, Bland's rule does not eliminate all degenerate pivots; it does, however, prevent cycles. Theorem 3.11. // the simplex algorithm is modified by Bland's rule, then no cycle occurs. Proof. The strategy of the proof is to assume that there is a cycle $T°,...,Tfc$, with pivots chosen according to Bland's rule, and derive a contradiction. Define t to be the largest member of the set C — {i : Xi enters

the basis during the cycle}. Thus C consists of all i such that Xi is basic for some tableaus in the cycle, and nonbasic for others. Since the bases for T° and Tk are identical, there is a tableau Tu in the cycle such that xt is basic in 3™ and nonbasic in Tw+1. Also, there is a tableau Tv such that x% is nonbasic in Tv and basic in Tv+1. Let xs be the variable which is nonbasic in Tu and basic in Tu+1. Thus xs and xt trade places in the pivot in Tu which produces Tw+1. Note that s < t since s is in C. Now, the objective function equation in Tu has the form n+m № = ά+Σ°*χχ> (3·2°) k = l where Ck — 0 whenever Xk is basic in Tu. Similarly, the objective function equation in Tv has the form n+m f(x) = d+J24*k, (3.21)

3.4. Avoiding Cycles and Achieving Feasibility 87 where cjj! = 0 whenever Xk is basic in Tv. The constant terms in (3.20) and (3.21) are the same because the value of the objective function does not change after a degenerate pivot. For any (n + m)-tuple χ satisfying the constraints, the two formulas (3.20) and (3.21) must give the same value. That is, subtracting, n+m £>* - c*)** = 0. (3.22) fc = l In particular, let us form χ by setting xs = 1, setting all the other variables nonbasic in Tu equal to 0, and solving for the variables basic in Tu from their constraint equations. We see that if χ ι is basic in Tw, then the expression for X{ has the form %i — bi — щ. In this equation, a; is the coefficient of xs in the constraint equation for —a?,·, and bi is the constant term in that equation. From (3.22), we get c.-cj+£(-<$)(*.·-«.·) = (). The sum here is over all variables basic in Tw, and we have used the fact that C{ = 0 when Xi is basic in Tu. Now, c* < 0 because, otherwise, xs would be a nonbasic variable in Tv eligible to enter the basis. However, xt actually enters the basis. Since s < t, Bland's rule is violated. Also, cs > 0 since xs enters the basis in the pivot on Tu. We conclude that cs -c*> 0, and so £c?(6,--a<)>0. We choose any r so that xr is basic in Tu and c*(br -ar) > 0. It follows that c* φ 0 and so xr is nonbasic in Tv. Thus br = 0. The last equation becomes c*ar < 0. Now, r φ t because at > 0 (it is the pivot entry in Tu) and c* > 0 (xt is the entering variable in Tv). Since xr is basic in Tu and nonbasic in Tv, r G C and so r < t. Finally, if ar > 0, xT is eligible to leave the basis in the pivot on Tu. Since r < t, the choice of x% violates Bland's rule. Thus ar < 0 and so c* > 0. But this means that xr is eligible to enter the basis in the pivot on Tv (instead of xt). Again, Bland's rule is violated. Π

88 3. Linear Programming To state this theorem differently, if we modify the

simplex algorithm by using Bland's rule, then no basis ever appears more than once. Since there are only finitely many bases, the algorithm must terminate after finitely many iterations. There is still the problem of finding an initial feasible tableau. 3.4.2. The Initial Feasible Tableau Let us consider the problem maximize $2x_1 — x_2 + 2x_3$ (3.23) subject to $x_1 + x_2 + x_3 < 6$ $-x_1 + x_2 < -1$ $-x_2 + x_3 < -1$. Introducing slack variables gives us the basic form maximize $2x_1 — x_2 + 2x_3$ subject to $x_1 + x_2 + x_3 — 6 = — X_4$ $~x_1 + x_2 + 1 = -x_5$ $—x_2 + x_3 + 1 = — x_6$· This basic form is infeasible, although the problem itself is feasible [(2,1,0) is a feasible vector]. We introduce an auxiliary variable u, and use it to construct the auxiliary problem. It is maximize — u subject to $x_1 + x_2 + x_3 — u — 6 = — x_4$ $-x_1 + x_2 - u + 1 = -x_5$ $-x_2 + x_3 - u + 1 = -x_6$. Notice that the auxiliary objective function has nothing whatsoever to do with the objective function for Problem (3.23). To understand why we are interested in the auxiliary problem, suppose that we can find a feasible tableau for it in which the auxiliary variable is nonbasic. The tableau obtained by erasing the column labeled by u is then a feasible tableau for Problem (3.23) (except that the original objective function must replace the auxiliary one). The reason for this is that $(x_1,..., x_q)$ satisfies the original constraints if and only if $(x_1,..., x_q, 0)$ satisfies the auxiliary constraints. (Here, the last entry in the second vector is u.) To continue with the example, we write the tableau

3.4. Avoiding Cycles and Achieving Feasibility 89 Xl 1 -1 0 0 *2 1 1 -1 0 X3 1 0 1 0 u -1 -1 -1 -1 -1 6 -1 -1 0 = = = = = -x4 -*s -я6 / This tableau is infeasible, but it is easy to pivot it into a feasible one. In this example, the entry in column 4 and row 2 (or row 3) works. Carrying out the first of these pivots gives Xl 2 1 1 1 X2 0 -1 -2 -1 X3 1 0 1 0 X& -1 -1 -1 -1 -1 7 1 0 1 = = = = = — X4 — tl -xe f This tableau is feasible. After two pivots chosen according to the simplex algorithm, we get X6 2 -1 -1 0 u -4 1 2 -1 xз 3 -1 -1 0 xь 1 0 -1 0 -1 3 1 2 0 = — #4 = -Я?2 = -Xι f This is optimal and the auxiliary variable is nonbasic. In order to write down an initial feasible tableau for the original Problem (3.23), it is only necessary to erase the column labeled by u and to replace the bottom row by the coefficients of the original objective function (written in terms of the nonbasics x_q, #3, and #5). Since the objective function is given in terms of x_1, #2) and #3, we only have to replace x_1 and #2 by the formulas for them given by the constraint equations. Thus $2x_1 ~ X_2 + 2x_3$ - 2(#6 + хз + хь + 2) - (я?6 + ж3 + 1) + 2x_3 = x_6 + 3жз + 2ж5 + 3. The initial

feasible tableau for Problem (3.23) is then xe 2 -1 -1 1 x₃ 3 -1 -1 3 хь 1 0 -1 2 -1 3 1 2 -3 = — #4 = -X2 = -Xι = f

90 3. Linear Programming The problem can then be solved in three pivots if we use Bland's rule, or in one if we do not. We formalize this method as follows. Algorithm 3.3 (The Feasibility Algorithm). We are given an infeasible primal basic form. (This is called the original problem.) (1) Introduce an auxiliary variable u. (2) Define the auxiliary objective function to be —u. (3) For each constraint in the original problem, form an auxiliary constraint by adding —u to the left-hand side. (4) Define the auxiliary problem to be that of maximizing the auxiliary objective function subject to the auxiliary constraints. (5) Set up the tableau for the auxiliary problem. (6) Choose a row (not the bottom one) so that the right-hand entry is smallest (the most negative); pivot on the entry in this row and in the column labeled by u. (The result is a feasible tableau.) (7) Apply the simplex algorithm to solve the auxiliary problem. (8) If the maximum value of the auxiliary objective function is negative, STOP (the original problem is infeasible). (9) Carry out a pivot so that the auxiliary variable is nonbasic (if necessary). (10) Set up a feasible tableau for the original problem by erasing the column labeled by the auxiliary variable, and then replacing the auxiliary objective function by the original one (written in terms of the nonbasics). The reader is asked to verify as exercises the assertions made in Steps (6) and (8). Note that the auxiliary problem always has a solution since its objective function is bounded above by zero. The only other point which seems to require comment is Step (9). It is conceivable that, after solving the auxiliary problem, we get a maximum of zero in a tableau in which the auxiliary variable is basic. Obviously, its value would be zero. It is, however, easy to see that a (degenerate) pivot makes it nonbasic. The following theorem summarizes the results of the previous sections. Theorem 3.12. If a primal linear programming problem is both feasible and bounded, then it has a solution. Moreover, a solution can be computed by using the feasibility algorithm (if necessary) together with the simplex algorithm, modified by Bland's rule.

3.5. Duality 91 Exercises (1) Solve maximize x\ + #3 subject to x\ + X2 + ^3 < 10 #1 — X2 — X2> < —1 XX > 2. (2) Solve maximize x\ + X2 — x₃ — x± subject to χχ + X2 + x₃ + x± < & X\ — X2 — X3 — X± < —1 X3 < 4. (3) Solve maximize 3x\ — 2x2 — x₃ + 5#4 subject to #i -f 2#2 + #3 + 2#4 < 10

X! <1 x± < 1 #2 — #3 < —1 X\ — 074 < —1. (4) Prove that if the pivot is chosen as in Step (6) of the feasibility algorithm, then the result is a feasible tableau. (5) Prove the assertion made in Step (8) of the feasibility algorithm. 3.5. Duality In this section, we describe the dual versions of the results already obtained for primal problems. It will then be easy to finish the job begun earlier (in Theorem 3.1 and its corollary) of revealing the connection between primal and dual problems.

92 3. Linear Programming 3.5.1. The Dual Simplex Algorithm For an example, let us consider Dual Basic Form (3.13). The dual tableau for this dual basic form is 2/1 2/2 -1 111-1 -1 -3 -1 2 10 0-1 = 2/3 = 2/4 =2/5 = 2/6 2 -1 0 = <? Here, the columns are read as equations. For example, column 2 is read (1)Ы + (-3)(й) + (0)(-1) = y4, which is the second constraint equation in the dual basic form. A pivot on the first entry in the first row produces Dual Basic Form (3.14), whose dual tableau is 2/3 2/2 -1 111-1 1-201 -1 -1 -1 0 = 2/1 =2/4 = 2/5 =2/6 2 1 -2 = <? A quick check shows that this second dual tableau could have been obtained by applying the pivoting algorithm (page 80). This is no coincidence. Indeed, the pivoting algorithm works for dual tableaus exactly as for primal ones. We omit the proof of this fact; it is similar to the proof of Theorem 3.6. There is also a theorem analogous to Theorem 3.5. It is the following: Theorem 3.13. Consider a dual tableau of size $(m+1) \chi (n+1)$. Then we have: (1) The tableau is feasible if and only if each entry in the bottom row (not counting the last one) is nonpositive. (2) Suppose the tableau is feasible. If each entry in the right-hand column (not counting the bottom one) is nonnegative7 then the tableau is optimal. Thus, the first dual tableau above is not feasible, but the second is feasible and optimal. There is, of course, a dual version of the simplex algorithm:

3.5. Duality 93 Algorithm 3.4 (The Dual Simplex Algorithm). We are given a feasible dual tableau T of size $(m+1) \chi (n+1)$. (i) // Чп+i ^ 0 for 1 < i < m, then STOP (the tableau is optimal). (2) Choose any к with 1 < к < m such that tk}n+i < 0. (The nonbasic variable labeling row к is called the entering variable.) (3) // tkj > 0 for 1 <j <n, STOP (the problem is unbounded.) (4) Choose any I with 1 < / < η such that tki < 0 and tm+l tki L· = min < m+ '3 Λ < j <n and tkj < 0 > I tkj J (The basic variable labeling column I is called the leaving variable.) (5) Pivot on tki and replace T by the dual tableau resulting from this pivot. (Entry tki ^ called the pivot entry.j (6) Go to Step (1). If this

algorithm is modified by Bland's rule, then cycles are impossible, and so the computation always terminates after finitely many iterations. There is, finally, a dual version of the feasibility algorithm which we illustrate by solving the following example: minimize $y_1 + y^\wedge$ (3.24) subject to $y_1 - 2y_2 + y_3 > 1$ $-2/1 \sim$ уз > -5. A dual basic form is minimize $y_1 + y_2$ subject to $y_4 = y_1 - 2y_2 + y_3 - 1$ Уъ - ~У\~ Уз + 5. This is infeasible. Construct an auxiliary problem by introducing an auxiliary variable v, an auxiliary objective function $+v$, and auxiliary constraints formed from the original constraints by adding v. The result is minimize v subject to $v_4 = y_1 - 2y_2 + y_3 + v - 1$ Уъ - ~У\~ Уз $+ v + 5$. The dual tableau for the auxiliary problem is

94 3. Linear Programming г/2 Уз v -1 1 -1 -2 0 1 -1 1 1 1 -5 = г/4 = уъ 0 0 0 1 0 = h We pivot on the entry in the row labeled by v and in the first column. The reason for this choice is that the entry at the bottom of the first column is positive. If there had been more than one column with positive bottom entry, we would have chosen the largest one. The result of the pivot is 2/1 2/2 2/3 2/4 -1 -1 -2 2 2 -1 -2 1 1 -1 -6 = V = J/5 -1 2 -1 1 -1 = h This is feasible. The dual simplex algorithm allows two choices of pivot. These are the entry in row 1, column 1, and the entry in row 3, column 1. We choose the first of these. The result is V 2/2 2/3 2/4 -1 -1 2 2 -2 -1 0 1 -1 -1 -4 = 2/i =2/5 1 0 0 0 0 = h This is optimal and the auxiliary variable v is nonbasic. To obtain the initial feasible dual tableau for the original problem, we erase the row labeled by υ and replace the auxiliary objective function by the original one (written in terms of the nonbasics 2/2, Уз, and y_4. The result is 2/2 2/3 2/4 -1 2 -2 -1 0 1 -1 -1 -4 = 2/1 =2/5 3 -1 1 -1 = 9

3.5. Duality 95 This is feasible but not optimal. The dual simplex algorithm calls for a pivot on the entry in row 2 and column 1. The result is 2/2 2/1 2/4 -1 2 -1 1 -1 = 2/3 -2 0 -1 -4 = 2/5 1 1 0 0 = 9 This is optimal and we read off the solution $\min g(y) = 0$, $y_1 = 0$, $y_2 = 0$, $y_3 - 1$. The reader is asked, in an exercise, to write down a formal statement of the dual feasibility algorithm. The following theorem summarizes the results about dual problems. Theorem 3.14. A dual linear programming problem which is both feasible and bounded has a solution. Moreover, a solution can be computed by using the dual feasibility algorithm (if necessary), together with the dual simplex algorithm modified by Bland's rule. 3.5.2. The Duality Theorem Consider a dual/primal pair of problems, and introduce slack and surplus variables so as to obtain a

primal and a dual basic form. The tableau and dual tableau corresponding to these basic forms consist of the same $(m+1)\times(n+1)$ matrix, and differ only in the labels. This suggests that we combine the two tableaus into one dual/primal tableau. This is done by putting both sets of labels on the matrix. For example, if the primal problem is maximize — $x_1 + 2\#_2 — 3\#_3 + 4x_4$ subject to $\#_3 — \#_4 < 0$ $\chi_1 — 2\#_3 < 1$ $2\#_2 + x_4 \wedge 3$ —$\chi_1 + 3a?_2 < 5$, then the dual/primal tableau is

96 3. Linear Programming Уч Уъ Уа -1 X\ X2 X3 XA 0 0 1-1 10-20 0 2 0 1 -13 0 0 -12-3 4 = У5 =Ув =У7 =2/8 -1 0 1 3 5 0 = 9 - -хъ - -хе - -X7 - -X8 = f Given this notation, we can simultaneously carry out pivots on both problems —it is simply a matter of keeping track of both sets of labels. For example, if we pivot on the second entry in the third row of the example, we get 2/1 2/2 2/6 2/4 -1 Xl 0 1 0 -1 -1 = 2/5 X7 0 0 1/2 -3/2 -1 = 2/3 хз 1 -2 0 0 -3 = 2/7 x4 -1 0 1/2 -3/2 3 = 2/8 -1 0 1 3/2 1/2 -3 = 9 = -хъ = -X6 = -a?2 = -X8 = f We are now ready for the most important theorem of this chapter. Theorem 3.15 (The Duality Theorem). If either of a dual/primal pair of problems has a solution, then so does the other; moreover, if they have solutions, the optimal values of the two objective functions are equal to each other. Proof. Suppose first that the primal problem has a solution. Set up the dual/primal tableau and carry out a sequence of pivots which produces an optimal feasible tableau for the primal. This is possible by Theorem 3.12. The final tableau in this sequence has all nonnegative entries in the last column (not counting the bottom row), and all nonpositive entries in the bottom row (not counting the last column). But, then, by Theorem 3.13, the dual tableau is both feasible and optimal. This proves that the dual problem has a solution. Also, the optimal values of both objective functions are equal to the negative of the bottom right-hand entry in the tableau, and thus are equal to each other. The proof that the primal problem has a solution if the dual does (and that the optimal values are equal) is identical to what we have just done, except for obvious changes. ∏

Exercises (1) Solve minimize $y_1 + У_2 + 2/3$ subject to — $y_1 + 2y_2 > 1$ -2/2+32/3 >1. (2) Solve the following primal problem and its dual maximize $bx_1 + X2 + 5\#_3$ subject to $x_1 + 2\#_2 + \#_3 < 4$ $a?_1 < 2$ $*_3 <2$. (3) Solve the following dual problem and its primal minimize $2/1 - 2/2 + 2/3$ subject to — $y_1 + У_2 > O$ -2/2 + 2/3 > -1. (4) Solve the following dual problem and its primal minimize $y_1 - У_2 + У_3$ subject to $2y_1 — y_2 > — 1$ -2/1 + 2/3 > -2. (5) Write down the dual feasibility algorithm.

4 Solving Matrix Games In this chapter, we apply linear programming to matrix games. The first task is to set up the problem of computing optimal strategies for the row player and column player as a dual/primal pair of linear programming problems. The minimax theorem will then be proved. The rest of the chapter consists of examples. The connection between game theory and linear programming was first discussed in [Dan51]. Our interest is in solving matrix games by using linear programming methods, that is, in stating the problem of solving a matrix game as a dual/primal pair. It is interesting that, conversely, linear programming problems can be stated as matrix games. See [LR57] for more information about this. 4.1. The Minimax Theorem Let M be an m χ η matrix game. The row player's problem is to compute the row value vr and an optimal mixed strategy P= (Pĩt-..iPm).

100 4. Solving Matrix Games That is, the problem is to find vr and a vector of probabilities ρ such that vr is as large as possible and m vr — min }]pimij. i - л 1 — \ We rephrase this problem in an equivalent way (and, also, drop the subscript r from v). The new formulation is: Find v and ρ such that v is as large as possible, subject to the following conditions: (1) pi > 0 for 1 < i < m. (2) E £iW = 1- (3) v < YlTzziPi171^ fᵒr 1 < J< n· To see that this really is equivalent to the first formulation, note that, since v is as large as possible, condition (3) implies that m v — mmy^j pjuiij. 1 — \ This new formulation is a linear programming problem. It has the disadvantage that the unknown v cannot be assumed to be nonnegative. (In many games, the column player has an advantage; in such games, v < 0.) It also is not in the form most convenient for solving. Now, the problem about v not being positive is easily taken care of. According to Exercise (11) on page 51, adding a constant to every entry in M results in adding the same constant to vr and vc. If we add a constant so large that the modified matrix has all entries positive, then vr and vc will also be positive. Clearly, too, adding a constant to every entry does not change the optimal strategies. Before going further, we assume that the following has been carried out. Preliminary Step Choose a constant c large enough so that rriij + c > 0 for all entries rriij in M. Replace M by the result of adding c to every entry. The plan is to solve the matrix game as modified by this preliminary step. The optimal strategies will be correct for the original matrix; the row value and column values of the original matrix will be those of the modified matrix minus the constant c. We return to the second formulation of the row player's problem. Let us make a change of variable yi = pi/v for 1 < г < т. (4.1)

4.1. The Minimax Theorem 101 The division by v is legal because of the preliminary step. Now, the condition that the p;'s sum to 1 [condition (2)] implies that ι m Γ = 1 Also, condition (3) implies that m У^ rriijyi > 1 for 1 < i < n. Γ = 1 Since maximizing v is equivalent to minimizing 1/v, we arrive at another (and final) formulation: minimize yi H h ym (4.2) m subject to /Jm^y; — 1' ^ОГ 1 — i — n- Γ = 1 Thus, the row player's problem is a dual linear programming problem. It should come as no surprise that the column player's problem is the corresponding primal. The derivation is very similar. First, we can formulate her problem as: Find v and <f such that v is as small as possible, subject to the following conditions: (1) qj > 0, for 1 < j < n. (2) Σ?=uí;=1- (3) v > Σ?=1 Я.3т13 1 f°r 1 < « < Wl. The appropriate change of variables is Xj — qj/v, for 1 < j < n. (4.3) The column player's problem becomes maximize x\ + l· xn (4-4) η subject to 2ZmiJxJ — 1' ^ОГ ^ — '— m* i=i Thus, as predicted, Problems (4.2) and (4.4) form a dual/primal pair. The method of solving a matrix game is now clear. It is: (1) Carry out the preliminary step (and remember the constant c; it will be needed later). (2) Set up and solve Problems (4.2) and (4.4).

102 4. Solving Matrix Games (3) Compute the column value of the modified (by the preliminary step) game by taking the reciprocal of the maximum value of the primal objective function. (4) Compute ρ and q, using the solutions of the linear programming problems and the change of variable equations [(4.1) and (4.3)]. (5) Compute the row value and column value of the original game by subtracting c from the corresponding quantities for the modified game. It is time to state the following: Theorem 4.1 (The Minimax Theorem). Let M be an m χ η matrix game. Then, both the row player and column player have optimal mixed strategies. Moreover, the row value and column values are equal. Proof. Problem (4.2) is feasible by Exercise (6). Its objective function is bounded below by zero. By Theorem 3.14, it has a solution. By Theorem 3.15 (the duality theorem), Problem (4.4) also has a solution. Therefore, both players have optimal mixed strategies. Finally, the optimal values of the two objective functions are equal; it follows immediately that the row value and column value are equal. Π Let us do an example to illustrate the method. Let M This game has no saddle point and there are no dominated rows or columns. We carry out the preliminary step with c = 3. The modified matrix is (-2 0 2 \-i 2 -1 1 3 1 \ 3 -1 o/ M The initial dual/primal tableau for the row player's and column player's problems is í/1 Уч Уз У4 -1 Xl 1 3 5 2 1 = УЪ X2 5 2 4 6 1 =

Уб хз 4 6 2 3 1 = У7 -1 0 = 9 — — — — — —#4 ~хъ -хе -Хη f

4.1. The Minimax Theorem 103

The primal tableau is already feasible, so we carry out the (primal) simplex algorithm. There are three pivots allowed by the algorithm. We choose the entry at row 3, column 1. The result of the pivot is 2/1 2/2 2/5 2/4 -1 Х6 -1/5 -3/5 1/5 -2/5 -1/5 = 2/3 Х2 21/5 -2/5 4/5 22/5 1/5 = 2/6 *3 18/5 24/5 2/5 11/5 3/5 = 2/7 -1 4/5 2/5 1/5 3/5 -1/5 = 9 = -x4 = -хъ = -Хι = -x7 f There are now two possible pivots; the one chosen is the entry at row 4, column 2. After this second pivot, one more pivot gives an optimal tableau 2/1 2/7 2/5 2/6 -1 X6 18/55 -7/55 3/11 -3/110 -13/110 = 2/з x7 -54/55 1/55 -2/11 12/55 -3/55 = 2/4 X5 -3/10 1/5 0 -1/10 -1/10 = 2/2 -1 1/11 1/11 1/11 1/11 -3/11 = 9 - -XA = -хз = -Хι = -χ2 = f From this, we read off the solutions to the linear programming problems: max/ = ming = 3/11, Xl — χ2 — χ3 = 1/11, У1=0, У2 = 1/Ю, уз = 13/110, 2/4 = 3/55. From the change of variable equations ((4.1) and (4.3)), we see that Pj =(ll/3)% and 9< = (11/3)* <. Thus, q = (1/3,1/3,1/3) and ρ = (0,11/30,13/30,1/5). Finally, the value of the original game is the value of the modified game (which is 11/3) minus the constant c from the preliminary step, v(M) = 11/3- 3 = 2/3.

104 4. Solving Matrix Games

Exercises (1) Solve the game (2) Solve the game (3) Solve the game (4) Solve the game -2 3 0 1 \ 1-3 2 0 0-2 13 -1 1-12 (5) Solve the game (6) Prove that Problem (4.2) is feasible.

4.2. Some Examples

In this section, we use linear programming to solve some interesting games. 4.2.1. Scissors-Paper-Stone This game was discussed in Chapter 2 (Example 2.1) and the reader was asked to solve it in Exercise (3) on page 63. The three pure strategies play indistinguishable roles in the game and so it is not surprising that the optimal strategy for each player is to play each of them with equal probability: ρ = (1/3,1/3,1/3). We now introduce a more complicated game based on scissors-paper- stone. This game, which we will call scissors-paper-stone-glass-water·, or

4.2. Some Examples 105

Figure 4.1. Scissors-paper-stone-glass-water. SPSGW, was first discussed in [Wil66]. In SPSGW, the two players simultaneously choose one of the five objects in the name of the game. The winner is determined by looking at Figure 4.1. In this directed graph, there is an edge between every pair of vertices. If the edge goes from vertex и to vertex v} and if one player chooses и while the other chooses v, then the

winner is the one who chooses u. Thus, glass beats stone; stone beats water; scissors beats glass etc. For comparison, the directed graph for ordinary scissors-paper-stone is given in Figure 4.2. The matrix for SPSGW is easily computed: /0-1 1 1 -1 \ 1 0 1—1—11 -1-1 0-1 1 . -1 1 1 0-1 \ 1 1-1 1 0 / Since this is a symmetric game, it could be solved using the method of Chapter 2. The linear programming method was used instead. The value of the game is zero, and the optimal strategy for each player is ρ =(1/9,1/9,1/3,1/9,1/3). This solution is surprising in at least one way. The third pure strategy seems to be the least favorable for each player. Three of the five outcomes

106 4. Solving Matrix Games scissors stone Figure 4.2. Scissors-paper-stone. paper of playing it are losses, and one of the others is a draw. Nevertheless, the optimal strategy calls for playing it one-third of the time. This illustrates the fact that, even for fairly simple games, the optimal solution may be difficult to arrive at intuitively. 4.2.2. Three-Finger Morra The rules are the same as for two-finger morra, except that each player can hold up one, two, or three fingers (and, at the same time, guess 1, 2, or 3). Thus, each player has 3x3 = 9 pure strategies. Each of these pure strategies can be designated by an ordered pair (/,p), where / is the number of fingers held up, and ρ is the prediction. In the matrix for this game, we have labeled the rows and columns by the appropriate designations: (1,1) (1,2) (1,3) (2,1) (2,2) (2,3) (3,1) (3,2) (3,3) (1,1) 0 -2 -2 3 0 0 4 0 0 (1,2) 2 0 0 0 -3 -3 4 0 0 (1,3) 2 0 0 3 0 0 0 -4 -4 (2,1) -3 0 -3 0 4 0 0 5 0 (2,2) 0 3 0 -4 0 -4 0 5 0 (2,3) 0 3 0 0 4 0 -5 0 -5 (3,1) -4 -4 0 0 0 5 0 0 6 (3,2) 0 0 4 -5 -5 0 0 0 6 (3,3) 01 0 4 0 0 5 -6 -6 0 This game is certainly symmetric. The matrix has no saddle points, and

4.2. Some Examples 107 Table 4.1. Solutions for three-finger morra. Solution 1 2 3 4 (1.3) 5/12 16/37 20/47 25/61 (2,2) 1/3 12/37 15/47 20/61 (3,1) 1/4 9/37 12/47 16/61 there are no dominated rows or columns. It is an example of a symmetric game for which the method of Chapter 2 fails. In fact, none of the nine possible systems of linear equations has any solutions at all. Solving the dual/primal pair of linear programming problems by hand is tedious, but possible. However, they were solved by computer, and four optimal basic solutions were found; these were converted to solutions of the game. It is interesting that they all involve only the three pure strategies (1,3), (2,2), and (3,1). The probabilities of playing these three pure strategies in each of the four solutions are given in Table 4.1. It is also interesting to note how close

together these four solutions are. 4.2.3. Colonel Blotto's Game The version of this game which we will discuss is found in [Dre81]. It is a military game and goes as follows. Colonel Blotto leads an infantry force consisting of four regiments. The enemy force, commanded by General Attila, consists of three regiments. There are two positions which both armies would like to capture (called San Juan Hill and Lookout Mountain). The problem for both commanding officers is to decide how many regiments to send to each position. We assume that a battle between two forces will result in a victory for the one with more regiments, and a draw if they are equal. The payoffs are computed as follows: If a force of r regiments defeats a force of s regiments, then the winner gains s+l (the +1 is there because the position is captured by the winner and is regarded as worth 1). Now, the five pure strategies for Colonel Blotto are: (4, 0), (0,4), (3,1), (1,3), and (2,2). Here, for example, (3,1) means that three regiments are sent to San Juan Hill and one to Lookout Mountain. Similarly, General Attila has four pure strategies: (3,0), (0,3), (2,1), and (1,2). The payoff matrix is

	(4,0)	(0,4)	(3,1)	(1,3)	(2,2)
(3,0)	4	0	1	-1	-2
(0,3)	0	4	-1	1	-2
(2,1)	2	1	3	0	2
(1,2)	1	1	2	0	3	2

108 4. Solving Matrix Games These entries are easy to compute. For example, if Blotto plays (3,1) and Attila plays (2,1), Blotto wins the battle of San Juan Hill and gains 2+1 = 3. The battle of Lookout Mountain is a draw. If Blotto plays (4,0) and Attila plays (2,1), the battle of San Juan Hill is a victory for Blotto, who gains 2+1 = 3. The battle of Lookout Mountain is won by Attila, who gains 0+1 = 1. Thus, Blotto gains a net payoff of 3 — 1 = 2. By the linear programming method, we get the following solution: value= 14/9, ρ = (4/9,4/9,0,0,1/9), g = (1/30,7/90,8/15,16/45). This says that Blotto should almost always concentrate all his regiments at one of the two positions, but that, one time in nine, he should divide them equally. As for Attila, he should divide his forces most of the time (with probability 8/15+ 16/45 = 8/9, to be precise), but should occasionally concentrate them all at one position. Thus, the stronger army and the weaker one use dramatically different strategies. There is another curious thing about this solution. Although the two positions are indistinguishable, Attila's pure strategies (3,0) and (0,3) are played with unequal probabilities. The same is true of the pure strategies (2,1) and (1, 2). This apparent paradox is removed by the observation that the probability vector obtained from <f by interchanging q\ and #2, and interchanging qs and q ± is also an optimal strategy for General Attila. If we average these two

optimal solutions, we get a symmetric optimal solution for Attila: q — (1/18,1/18,4/9,4/9). The fact that the value of the game is positive is to be expected since Blotto has one more regiment than Attila; thus, he should have an advantage. There is a reasonable philosophical objection to the solution above. It is certainly valid in case the "game" is played repeatedly, but this particular military campaign will only occur once. The consequences of an unfortunate combination of pure strategies might be devastating for one of the armies. For example, if Attila plays (2,1) while Blotto plays (4,0), two of Attila's regiments might be rendered unfit for further action for a long time. In the context of the entire war, this might be unacceptable. On the other hand, it is hard to see what a "safe" alternative is for Attila (except to retreat without fighting at all). 4.2.4. Simple Poker Real poker is usually played by more than two people and is difficult to analyze, partly because the variety of possible hands of cards is very great. For this reason, various forms of simplified poker have been studied in order to gain insight into the real game. Our version is played as follows. First,

4.2. Some Examples 109 there are two players: San Antonio Rose and Sioux City Sue. Each player antes $1. This means that that amount of money is placed in the middle of the table. The resulting pile of money (or chips representing money) is called the pot. Each player is then dealt a hand consisting of two cards, face down. Each player may look at her own cards, but cannot see the other player's. The deck of cards from which these hands are dealt is highly simplified. It consists of only two kinds of cards: High (abbreviated H), and Low (abbreviated L). Moreover, the deck is very large, so that the probability of dealing an H is always the same as dealing an L. Thus, $Pr(\#) = Pr(L) = 1/2$. There are three possible hands: HH,LL,HL. The probabilities of these are easily seen to be $Pr(\#\#) = 1/4$, $Pr(LL) = 1/4$, $Pr(\#L) = 1/2$. In accordance with the traditions of poker, the ranking of the hands is: $H H > LL > HL$. Thus, either pair beats a hand without a pair. Rose plays first. She has two choices: She can either bet by adding $2 to the pot, or check (that is, bet zero). It is now Sue's turn. Her options depend on what Rose did. If Rose bet, Sue can either see or fold . To see the bet means to match Rose's bet by putting her own $2 in the pot. To fold means to give up (without having to put any money in), and, in this case, Rose wins the pot. On the other hand, if Rose checks, Sue can either bet $2 or check. Finally, if Rose checks and Sue bets, then Rose has the option of either seeing the bet or folding (in which case, Sue wins the pot). If neither

folds, the two hands are turned over and compared. If they are the same, the pot is divided equally. Otherwise, the player with the higher hand takes the whole pot. The extensive form of our version of simple poker is large but not unmanageable. To see what Rose's pure strategies are, two things should be kept in mind. First, Rose knows her own hand but not Sue's. Second, if a given pure strategy calls for her to check, then that strategy must also tell her which action to take in case Sue's response is to bet. The first of these considerations tells us that the action to be taken according to a particular strategy can depend on her hand. Since there are three possible hands, it is natural to represent a pure strategy as a 3-tuple, where each entry represents an action for one of the possible hands. To designate these actions, we let "b," "s," "f" stand, respectively, for "bet," "check and see if Sue bets," and "check and fold if Sue bets." There are thus 33 = 27 pure strategies. A typical strategy is "bsf," which means: If Rose's hand is HH, she bets; if it is LL, she checks but will see if Sue bets; and, it is HL, she will check and fold if Sue bets.

110 4. Solving Matrix Games As for Sue's pure strategies, note that she knows not only her own hand but also the action taken by Rose. Therefore, there are six situations in which she could find herself (any of three hands times either of two actions taken by Rose). It is therefore natural to represent Sue's pure strategies as 6-tuples. The structure of these 6-tuples is indicated in the following table: HH, bet I HH, check | LL, bet | LL, check | HL, bet | HL, check We use uppercase letters in designating Sue's possible choices: "B" means "bet," "S" means "see," "C" means "check," and "F" means "fold." A typical pure strategy for Sue is: SBFBSC. The meaning of this is that if Sue is dealt HH and Rose bets, then she sees the bet; if Rose checks, she bets. If Sue is dealt LL and Rose bets, she folds; if Rose checks, she bets. Finally, if Sue is dealt HL and Rose bets, she sees; if Rose checks, she checks. Since each entry in the 6-tuple can be either of two possibilities (S/F or B/C, depending on what action Rose took), the total number of pure strategies for Sue is 26 = 64. Thus, it appears that the game matrix for simple poker will be 27 x 64 (taking Rose to be the row player). Common sense (or, rather, poker sense) allows us to eliminate some of the pure strategies for each player. For Rose, we can safely eliminate all strategies which call for her to fold HH. This hand cannot lose. This leaves 18 pure strategies for her. In Sue's case, any strategy which calls for her to fold HH can also be eliminated. With a little additional thought, we see that it is safe to eliminate those which call for her to check with HH.

This reduces the number of pure strategies for Sue to 16. All of them begin with SB. The other way of looking at this is that, if we wrote out the full 27 x 64 matrix, the rows and columns corresponding to pure strategies we have eliminated would be dominated. We now must explain how the matrix entries are computed. This is complicated because each entry is an average over all the 9 possible combinations of hands for the two players. To take an example, suppose Rose plays sbs, while Sue plays SBFBFB. We list the nine possible combinations of hands and analyze each: (1) Each is dealt HH [probability = (1/4) χ (1/4) = 1/16]; Rose checks, Sue bets, Rose sees; payoff is 0. (2) Rose is dealt HH, Sue gets LL [probability = (1/4) χ (1/4) = 1/16]; Rose checks, Sue bets, Rose sees; Rose wins $3 (the half of the pot contributed by Sue; the other half was already hers). (3) Rose is dealt HH, Sue gets HL [probability = (1/4) χ (1/2) = 1/8]; Rose checks, Sue bets, Rose sees; Rose wins $3. (4) Rose is dealt LL, Sue gets HH [probability = (1/4) χ (1/4) = 1/16]; Rose bets, Sue sees; Sue wins and so payoff = — $3.

4.2. Some Examples 111 (5) Each is dealt LL [probability = (1/4) χ (1/4) = 1/16]; Rose bets, Sue folds; Rose wins $1. (6) Rose is dealt LL, Sue gets HL [probability = (1/4) χ (1/2) = 1/8]; Rose bets, Sue folds; Rose wins $1. (7) Rose is dealt HL, Sue gets HH [probability = (1/2) χ (1/4) = 1/8]; Rose checks, Sue bets, Rose sees; the payoff is —$3. (8) Rose is dealt HL, Sue gets LL [probability = (1/2) χ (1/4) = 1/8]; Rose checks, Sue bets, Rose sees; the payoff is — $3. (9) Each is dealt HL [probability = (1/2) x (1/2) = 1/4]; Rose checks, Sue bets, Rose sees; the payoff is 0. The matrix entry in the row labeled sbs, and in the column labeled SBFBFB is the sum of the nine payoffs, each multiplied by the probability of that particular combination occurring: (1/16)(0) + (1/16)(3) + (1/8)(3) + (1/16X-3) + (1/16)(1) + (1/8)(1) + (1/8)(-3) + (1/8)(-3) + (1/4)(0) = -3/16. Computing all 288 entries is a very tedious job, if done by hand. In fact, however, it was done on a computer. Because of space considerations, the matrix will not be shown here. It turns out not to have a saddle point, but there are some dominated rows and columns. After these are eliminated, we are left with a 12 χ 12 matrix. Solving by the linear programming method gives us: value of simple poker is — 0.048. Thus, the game is slightly (about a nickel per hand) favorable to Sue. The dual/primal problem has many solutions. One of these gives the following optimal mixed strategies. Rose should play • bsb with probability 0.231 • bsf with probability 0.615 • bff with probability 0.077 • sff with probability 0.077. Sue should play

• SBSBFB with probability 0.154 • SBSCFC with probability 0.462 • SBFCFB with probability 0.077 • SBFCFC with probability 0.308. These mixed strategies can be restated in a way which is easier to grasp. Suppose, for example, that Rose is dealt HH. How should she play? Looking at her solution, we see that bsb, bsf, and bff call for her to bet on this hand.

112 4. Solving Matrix Games The fourth pure strategy, sff, calls for her to check (and see). The sum of the probabilities of those pure strategies calling for a bet is $0.231 + 0.615 + 0.077 = 0.923$. Our recommendation to her is, therefore, that she bet with probability 0.923, when dealt HH. Similar calculations are easily made in case she is dealt LL or HL. In summary, our recommendations to Rose are as follows: • On HH, bet with probability 0.923, check and see with probability 0.077. • On LL, check and see with probability 0.846, check and fold with probability 0.154. • On HL, bet with probability 0.231, check and fold with probability 0.769. Our recommendations to Sue are computed in a similar way. They are as follows: • On HH and Rose betting, always see. • On HH and Rose checking, always bet. • On LL and Rose betting, see with probability 0.616, fold with probability 0.384. • On LL and Rose checking, bet with probability 0.154, check with probability 0.846. • On HL and Rose betting, always fold. • On HL and Rose checking, bet with probability 0.231, check with probability 0.769. Some of these recommendations are intuitive but some are certainly not. It is to be expected that Rose should occasionally bluff with an HL, that is, bet it as if it were a stronger hand. She is told to do this almost one time in four. On the other hand, it would have been difficult to guess that Rose should check HH about one time in 13. We have mentioned that the dual/primal linear programming problem for this game has many solutions. It is a fascinating fact that if we convert any of the known solutions into recommendations as was done above, we get exactly the same probabilities! Exercises (1) Verify that the solution given for SPSGW is correct. (2) A surprising feature of the solution to SPSGW was discussed earlier in the text, namely, that the third pure strategy is played relatively frequently. Suppose that the row player has no faith in game theory and decides to vary the optimal strategy by playing row 3 with probability less than 1/3, and row 5 more than

4.2. Some Examples 113 Figure 4.3. Directed graph for Exercise (3). Figure 4.4. A game.

1/3 (the other probabilities being unchanged). How would the column player respond and what would the result be? (3) Define a two-person game based on the directed graph in Figure 4.3 in the same way in which SPSGW was based on its directed graph. Solve the game. (4) Define a two-person game based on the directed graph in Figure 4.4. Solve the game. (5) Verify the solutions given in the text for three-finger morra. (6) Suppose that one of the players in three-finger morra plays according to the first optimal strategy in the table on page 107, but that the other player plays each pure strategy with probability 1/9. Compute the expected payoff. Then, still assuming that one player plays each pure strategy with probability 1/9, how should the other player play? (7) In Colonel B lotto's game, suppose that the Colonel does not entirely trust game theory—instead of playing the strategy we computed, he decides to flip a fair coin and send all his regiments to either San Juan Hill or Lookout Mountain, depending on whether the coin comes up heads or tails. In other words, his mixed strategy is $\hat{} = (1/2,1/2,0,0,0)$. Assuming that Attila knows this, how would you advise him to respond, and what would be the result? Does Attila do better than the value of the game predicts? (8) Suppose that General Attila ignores our advice and plays the mixed strategy $(0,0,1/2,1/2)$. Assuming that Blotto knows this, how should he play, and what would be the result? (9) Verify that $v = 14/9$, $\rho = (4/9,4/9,0,0,1/9)$, $q = (1/18,1/18,4/9,4/9)$ constitutes a solution to Colonel Blotto's game. (10) Modify Colonel Blotto's game so that each army has three regiments. Solve it. (11) Compute the entry in the matrix for simple poker which is in the column labeled SBFBFB and in the row labeled bfb. (12) Compute the entry in the matrix for simple poker which is in the column labeled SBFBSC and in the row labeled bsb. (13) Suppose that Rose rejects our recommendations and decides to play a mixed strategy consisting of only bsb and bsf, each with probability 0.5. How would you advise Sue to respond, and what would be the result? (14) Suppose that, in simple poker, Sioux City Sue modifies our recommended strategy by always checking when she has HL and Rose checks. How should Rose respond and what would be the result?

5 Non-Zero-Sum Games In studying games which are not zero-sum, the distinction between cooperative and noncooperative games is crucial. There are two types of cooperation among players. The first is the making of binding agreements before play starts about how they will play (that is, coordination of strategies); the second is the making of agreements about sharing payoffs (or

about "side payments" from one player to another). The aim of the first type is to increase the payoffs to the cooperating players; the aim of the second is for one player (or group of players) to induce another player to coordinate strategies. It is important to understand that the agreements made concerning coordination of strategies must be binding. If "double crosses" are possible, then there might as well be no cooperation at all. An interesting example to keep in mind when thinking about these ideas is Prisoner's Dilemma (page 25). In this game, each of the two players (named Bonnie and Clyde) has a choice of two strategies: C (for "cooperate") or D (for "defect"). If they both play D, they both go to prison for five years. If they both play C, they each receive a two-year sentence. If one defects and the other cooperates, the defector gets one year, while the cooperator gets ten. It is clear that if they could play cooperatively (and make a binding agreement), they would both play C On the other hand, if the game is noncooperative, the best way for each to play is D. We will look at this game in another way in Section 5.2.4. Another interesting point about it is that the payoffs are

116 5. Non-Zero-Sum Gaines not transferable. That is, there is no possibility of pooling the payoffs and dividing them in accordance with a previous agreement. Zero-sum two-person games are noncooperative in nature because cooperation is never to either player's advantage. The total payoff is always zero and cannot be increased by any combination of strategies. In this chapter, we are interested in games with two players. However, a few theorems and definitions will be stated for an arbitrary number, TV, of players. We will study noncooperative games first and then go on to cooperative ones. 5.1. Noncooperative Games The normal form of a two-person game can be represented as a pair of $m \chi \eta$ matrices (see page 30). A more succinct way of expressing the same thing is as a ЬI-matrix^ that is, a matrix of pairs. For example, the bi-matrixfor Prisoner's Dilemma is / (_2,-2) (-10,-1) λ V (-1,-ю) (-5,-5) ;· In general, if C is an $m \chi \eta$ bi-matrix, then each entry Cij is an ordered pair of numbers. The members of this ordered pair are the payoffs to the two players Pχ, P2 (called, in analogy with the zero-sum case, the row player and column player, respectively), given that the row player plays row г, and the column player plays column j. We will use both the bi-matrix notation (especially for displaying games), and the two-matrix notation (especially for computations). Other examples follow. Example 5.1 (Battle of the Buddies). Two friends, named Norm and Cliff, have different tastes in evening

entertainment. Norm prefers professional wrestling matches, but Cliff likes roller derby. Neither likes to go to his choice of entertainment alone; in fact, each would rather go with the other to the other's choice than go alone to his own. On a numerical scale, each grades going alone to either event as a 0, going to his friend's choice with him as 1, and going with his friend to his own favorite event as 5. We think of these numbers as "happiness ratings." Regarding this situation as a game, we see that each player has two pure strategies: W (for "wrestling") and R (for "roller derby"). Let us suppose that they independently and simultaneously announce their decisions each evening. This makes the game noncooperative; we will discuss

5.1. Noncooperative Gaines 117 the cooperative variant later. The 2x2 bi-matrix is thus f (5,1) (0,0) \ V (0.0) (1,5) J We have designated Norm as the row player and Cliff as the column player. For this game, the advantages of cooperation are not so striking as in the Prisoner's Dilemma game. Of course, both Norm and Cliff are better off if they both choose the same activity—the problem is that they cannot agree on which one to choose. Rational people would resolve this problem by flipping a coin or alternating between the activities, but that would make the game cooperative. Example 5.2 (Chicken). Two teenage males with cars meet at a lonely stretch of straight road. They position the cars a mile apart, facing each other, and drive toward each other at a high rate of speed. The cars straddle the center line of the road. If one of the drivers swerves before the other, then the swerver is called "chicken" and loses the respect of his peers. The nonswerver, on the other hand, gains prestige. If both swerve, neither is considered very brave but neither really loses face. If neither swerves, they both die. We assign numerical values, somewhat arbitrarily, to the various outcomes. Death is valued at —10, being chicken is 0, not swerving when the other driver does is worth 5, swerving at the same time as the other is valued at 3. Our version of this game is taken from [Pou92]. For other versions, see the movies Rebel Without a Cause and American Graffiti. The bi-matrix for Chicken is (3,3) (0,5) \ (5,0) (-10,-10))· Cooperating players would surely agree that both should swerve. This conclusion is, however, highly dependent on our choice of payoff numbers. For example, if (3, 3) in the bi-matrix is replaced by (2, 2), then cooperating players would do better flipping a coin to decide who swerves and who does not. 5.1.1. Mixed Strategies The concept of an equilibrium pair of strategies was introduced in Chapter 1 (see page 24). The strategies to which this

definition refers are pure strategies (since mixed strategies had not yet been defined). There is no reason, however, why players in a non-zero-sum game cannot play mixed strategies, and no reason why we cannot extend the definition of an equilibrium pair to include them.

118 5. Non-Zero-Sum Gaines Definition 5.1. Let π be an TV-player game in normal form with strategy sets $X_1,..., X_n-$ A mixed strategy for player $P_r\cdot$ is a probability vector Pi — $(Pi(x))xeX_i-$ The entry $Pi(x)$ is interpreted as the probability that $P_r\cdot$ plays strategy $\chi \pounds X\{$. For a two-player game the notation is simpler. Let the $m \chi \eta$ matrices A and B be the payoff matrices for the row player and column player, respectively. Then a mixed strategy for the row player is an m-tuple ρ of probabilities; a mixed strategy for the column player is an n-tuple q of probabilities. We will use the symbol $M_r\cdot$ to denote the set of all mixed strategies for player $P;.$ The expected payoff to player $P_r\cdot$ due to the mixed strategies $\rho_i > \cdots > \text{Й}\backslash \text{т} >$ played by players $P_i,..., P_n$, respectively, is $*I\cdot(P_1,...,P\alpha\gamma) = \Sigma(P_1(\chi_1) \text{ x} \cdots \text{хрлг(жлг)})\text{тгг-(ж1,...,жлг)}$, where the sum is taken over all choices of $X\{ \pounds X_r\cdot, 1 < i < N$. Thus, the expected payoff is the sum of all payoffs due to N-tuples of pure strategies, each weighted according to the probability of that TV-tuple being played. If $TV = 2$ and A and B are the $m \chi \eta$ payoff matrices for the row player and column player, respectively, then the expected payoff to the row player due to the mixed strategies ρ and q is $m \eta$ $\pi_1\{P_1 < i) = \Sigma\Sigma P_i < \upsilon\alpha\upsilon'$ r=1j=l The expected payoff to the column player is $m \eta$ r=1i=l In the game of the Battle of the Buddies, a mixed strategy for Norm is an assignment of probabilities to the two pure strategies W and R. Suppose he decides to flip a fair coin each evening and choose W if it comes up heads, and R if tails. If, meanwhile, Cliff always chooses R, then the expected happiness of Norm is $(1/2)(0) + (1/2)(1) = 1/2$, while the expected happiness of Cliff is $(1/2)(0) + (1/2)(5) = 5/2$.

5.1. Noncooperative Gaines 119 On the other hand, if Cliff chooses R with probability 2/3, and W with probability 1/3, then Norm's expected happiness is 7/6, while Cliff's is 11/6. 5.1.2. Maximin Values Before formally defining equilibrium N-tuples, we discuss a quantity called the "maximin value," which a player can rather easily compute and which gives him or her a pessimistic estimate of how much payoff can be expected. To make the discussion definite, assume that we are dealing with a two-person game with players Pi and P2. The maximin value, $v\backslash$, for Pi is computed by making the assumption that the

other player will act so as to minimize Pi's payoff. Thus, v_1 — maxmin7Ti(p, g), ρ я where ρ and q range over all mixed strategies for Pi and P2, respectively. Now, unless the game is zero-sum, the assumption that P2 will play that way is probably false. In fact, P2 will act so as to maximize her payoff, and this is often not the same as minimizing his. Nevertheless, v_1 gives Pi a lower bound on his payoff (since he can play so as to guarantee at least that much), and it is easy to compute: We simply regard Pi's payoff matrix as a zero-sum game; its value is then V_1. This is true because, if P2 plays to minimize Pi's payoff, she is playing to maximize the negatives of those payoffs. This is another way of saying that she is acting as if she were the column player in the matrix game. Consider, for example, the following bi-matrix (1.1) (0,1) (2,0) \ (1.2) (-1,-1) (1,2) . (5.1) (2,-1) (1,0) (-1,-1)/ The payoff matrix for Pi, the row player, is C-·.:) As a matrix game, this is easily solved. After removing a dominated row and a dominated column, the remaining 2x2 matrix is solved to give ЬI = 1/2. The optimal strategy (for the matrix game) is to play rows 1 and 3 both with probability 1/2. The column player's maximin value is then computed by solving the matrix game obtained by transposing her payoff matrix. This transposition

120 5. Non-Zero-Sum Gaines is necessary because her payoffs are the actual matrix entries and not their negatives. The value of this transposed game will then be v2. Here is the matrix we need for the example: \0 2-1/ This is also easily solved to give v2 = -1/4. The mixed strategy for P^ which achieves this value is to play column 1 of the bi-matrix with probability 1/4, and column 2 with probability 3/4. The pessimistic nature of these maximin values is illustrated by the fact that if both players play according to the mixed strategies just computed (which are called the maximin solutions), then they both receive expected payoffs greater than their maximin values. In fact, Pi actually has an expected payoff of 3/4, while P2 has an expected payoff of 3/8. In the special case where the game is zero-sum, the minim ax theorem implies that the two maximin values are negatives of each other. 5.1.3. Equilibrium N-tuples of Mixed Strategies The definition of an equilibrium TV-tuple of mixed strategies is Definition 5.2. Let π be an TV-person game in normal form. An N- tuple of mixed strategies <fi,..., fa is an equilibrium N-tuple if ЭТ|(Й,...,Й,...,Й\г) < 7Гг(Й,...,длг), for all г and for all mixed strategies pi for player Pг·. Thus, if all players but one use the mixed strategies in the N-tuple, but the other departs from it, the one departing suffers (or, at least, does no better). The following

theorem is due to Nash (see [Nas51]; proofs are also to be found in [Owe82] and [Vor77]). We will omit the proof since it uses the Brouwer fixed-point theorem from topology. Theorem 5.1. Let π be any N-player game. Then there exists at least one equilibrium N-tuple of mixed strategies. This result is of fundamental importance. Note that it would be false if we replaced "mixed" by "pure." The reservations expressed in Chapter 1 about the usefulness of equilibrium TV-tuples in solving games are still valid for equilibrium TV-tuples of

5.1. Noncooperative Gaines 121 mixed strategies. Nevertheless, they are still useful in some cases. In general, computing equilibrium TV-tuples is difficult. We now explore a simple but interesting case where it can be done. 5.1.4. A Graphical Method for Computing Equilibrium Pairs We illustrate the method by working with the Battle of the Buddies. Since each player has only two pure strategies, the mixed strategies for Norm can all be written in the form (a?, 1 — a?), where $0 < \chi < 1$. Similarly, the mixed strategies for Cliff all have the form (y, 1 — y), where $0 < y < 1$. The expected payoffs, as functions of χ and y, are easy to compute. We have *i(x> У) - $XУ + í1 - x)(l -У)- §XУ ~ x ~ У + !> where we have abbreviated 7Ti((ar, 1 — ж), (y, 1 — y)) by 7Ti(ar, y). Similarly, $\pi 2$(«ι У) — XУ + 5(1 - x)(1 - y) — 6xy - Ьx - Ьy + 5. In order to find an equilibrium pair, we need to find a pair (a?*,y*) so that 7Ti(a?*, y*) is a maximum over 7Ti(ar, y*), and 7Г2(#*, y*) is a maximum over 7Г2(#*, y). The idea is to graph together the following two sets: A = {(ж, y) : 7Гι(#, y) is a maximum over ж, with y fixed}, and B = {(ж, y) : 7Г2(#, y) is a maximum over y, with a? fixed}. The points of intersection of these sets will be precisely the equilibrium pairs. First, we rewrite the payoff functions as tfifaii/) = {®y-i)x-y+1, and $\pi 2$(«ι 2/) = (6a? - 5)j/ - 5a? + 5. From the formula for 7Ti(ar,y), we see that if $y < 1/6$, then a? = 0 at the maximum (since the coefficient of χ is negative); if r/ > 1/6, a? = lat the maximum; and, if $y = 1/6$, any a? gives the maximum. For 7Г2(ar, y), we see that if $\chi < 5/6$, $y = 0$ at the maximum; if a? > 5/6, $y = 1$ at the maximum; and, if $\chi = 5/6$, any y gives the maximum. The sets A and B are graphed in Figure 5.1. The set A is indicated with a solid line, the set B by a dashed line, and the three equilibrium pairs are circled. They are ((0,1),(0,1)), ((5/6,1/6), (1/6,5/6)), and ((1,0), (1,0)). Norm naturally prefers the third of these (since his payoff is 5), and Cliff prefers the first. The second gives each of them 5/6, which is less than each receives at the other two. In fact, 5/6 is the maximin value for both players [by Exercise (3)].

АУ l/6_ 0 Θ- --Q 0 5/6 Figure 5.1. Equilibrium pairs for Norm and Cliff. Another example is (4,-4) (-1,-1) (0,1) (1,0) (5.2) For this game, $\tau\eta(я, y) = 4xy - x(1-y) + (1-x)(1-y) = (6y-2)x - y + 1$, while тг2(ж, у) - -4xy - x(1-y) + (I-x)y - (-4ж + \)y - x. The two sets are shown in Figure 5.2. We see that the only equilibrium pair is ((1/4,3/4), (1/3,2/3)). Playing in accordance with these mixed strategies, the row player has an expected payoff of 2/3, and the column player has an expected payoff of —1/4. Exercises (1) In Battle of the Buddies, suppose Cliff plays the mixed strategy (1/4,3/4). What is the best way for Norm to play in response?

5.1. Noncooperative Gaines

123 i4y- - - - Figure 5.2. Equilibrium pair for Bi-matrix (5.2). (2) For the game of Prisoner's Dilemma, verify that the maximin value for each player is -5, and that ((0,1),(0,1)) is the only equilibrium pair of mixed strategies. (3) For the game of Battle of the Buddies, verify that 5/6 is the maximin value for each player. (4) For the game given by Bi-matrix (5.2), verify that the maximin value for the row player is 2/3, and for the column player is -1. (5) Given the bi-matrix / (-1,3) (1,0) \ (2,-1) (0,1) , V (1,1) (-2,1) J compute the maximin values for both players. (6) Consider the bi-matrix / (-2,3) (-1,1) (1,-2) \ (0,1) (-1,-2) (1,1) . V (2,2) (2,-1) (0,0)) Compute the maximin values for both players.

124 5. Noη-Z его-Sum Gaines (7) For the bi-matrix / (2,-3) (-1,3) $\lambda\,^\wedge$ (0,1) (1,-2) J' compute the equilibrium pairs and the maximin values for both players. (8) For the bi-matrix (2,-1) (-1,1) λ (0,2) (1,-1))' find the maximin values and the equilibrium pairs of mixed strategies. (9) Compute the maximin values and equilibrium pairs for the game of Chicken. (10) Modify Battle of the Buddies as follows: Norm is not as interested in wrestling as Cliff is in roller derby; in fact his happiness rating from going to wrestling with Cliff is 2 instead of 5. Thus the bi-matrix for the modified game is (2,1) (0,0) \ (0,0) (1,5) J ■ Compute the maximin values and equilibrium pairs.

5.2. Solution Concepts for Noncooperative Games

The analysis of noncooperative games is difficult because it is so often dependent on nonmathematical considerations about the players. As a result, most conclusions are subject to honest disagreement. In this respect, these games may be closer to real life than are other parts of game theory! We illustrate this by discussing several examples. Before doing so, we will enlarge our stock of analytic tools by making some definitions. The first of these is the following: Definition 5.3. Let π be a two-

person game. The two-dimensional set $\Pi = \{(\tau\tau\iota(\rho, q), 7\Gamma2(p, q)):pe \text{ Muqe}$ M2] is called the noncooperative payoff region. The points in Π are called payoff pairs. The payoff region can be plotted in a Cartesian coordinate system, where the horizontal axis represents $\pi\backslash$ and the vertical axis represents $7\Gamma2$. The idea is to consider all possible pairs of mixed strategies and plot the corresponding pair of payoffs for each. One thing that can be read from a drawing of a payoff region is the relation of dominance between payoff pairs. This is defined by the following: Definition 5.4. Let (u, v) and (u',v') be two payoff pairs. Then (u, v) dominates (it', v') if $u > u'$ and $v > v'$.

5.2. Solution Concepts for Noncooperative Games 125 (0,0) Figure 5.3. Payoff region for Battle of the Buddies. In the payoff region, the dominating payoff pair is above and to the right ("northeast") of the dominated one. For Battle of the Buddies, the payoff region is shown in Figure 5.3. It is bounded by two line segments and the curve from (1,5) to (5,1). (This curve is a piece of a parabola.) The three (distinct) entries in the bi- matrix are the coordinates of points in the region resulting from playing pure strategies. The three equilibrium pairs of mixed strategies result in the points marked with small circles. Our second definition is the following: Definition 5.5. Let both (p, q) and (f, s) be equilibrium pairs of mixed strategies for a two-person game. Then (1) they are said to be interchangeable if (p, s) and (f, q) are also equilibrium pairs; (2) they are said to be equivalent if $*i(Pi\ll) = \wedge(r,s)$, for $i = 1,2$. Theorem 2.6 shows that in a zero-sum game, all equilibrium pairs of mixed strategies are interchangeable and equivalent. This is not always true of non-zero-sum games. It is sometimes said that a game in which any two equilibrium pairs are interchangeable and equivalent is solvable in the

126 5. Non-Zero-Sum Games Nash sense. On page 122, we gave an example [Bi-matrix (5.2)] of a game in which there is only one equilibrium pair. Such a game is certainly solvable in the Nash sense because the single equilibrium pair is interchangeable and equivalent with itself. 5.2.1. Battle of the Buddies Let us try to use the ideas developed so far in order to say something about how Norm and Cliff might actually play their noncooperative game. We will take Norm's viewpoint, but this is no limitation because the game is really symmetric. We observe first that the equilibrium pairs are neither interchangeable nor equivalent. Thus, Battle of the Buddies is not solvable in the Nash sense. From Figure 5.3, we observe that almost every payoff pair is

dominated by some other pair. The only exceptions are (1,5), (5,1), and the points on the bounding curve in the neighborhood of (3/2,3/2). Payoff pairs which are not dominated by any other pair are said to be Pareto optimal. This concept will be discussed later in connection with cooperative games. It is clear that rational cooperating players would never play so that their payoff pair is not Pareto optimal. We list some possible mixed strategies and comment on them. The first three listed are Norm's strategies in the equilibrium pairs of mixed strategies. These are always of great interest because it is natural to think that only an equilibrium pair has the stability to persist indefinitely. (1) Norm might play the pure strategy W if he thinks that Cliff will eventually decide that he will never play anything else; based on this decision, Cliff's only rational move is to play W himself. The success of this scenario requires a difference in degree of stubbornness between Norm and Cliff. Otherwise, by symmetry, Cliff might persist in playing R. This combination of W by Norm and R by Cliff results in the lowest possible payoff for both. (2) Norm should probably not play (5/6,1/6) (his strategy in the second equilibrium pair). Cliff could gain 5/6 playing either W or R; if he happens to choose R, Norm's payoff would drop to 1/6. (3) The strategy R is an interesting choice for Norm. If Cliff believed this pattern of play would continue, he could respond with R. The result would be excellent for Cliff, and not bad for Norm. Symmetry, however, comes into this. If R is a good idea for Norm, is not W a good idea for Cliff? That would be a rather comic turn of events, but the possibility of it does cast doubt on the idea. (4) The maximin value for each player is 5/6 [from Exercise (3) of the previous section]. Norm can assure himself of at least that by playing (1/6,5/6). In fact, if he plays that way, his payoff is the

same no matter how Cliff plays. Perhaps this is the best he can do. The conclusion that we just came to is not very satisfying; it is impossible to think of it as definitive. Suppose, for example, that Norm takes our advice and plays (1/6,5/6). If Cliff sees what his friend is doing, he can, by playing R, gain an expected payoff of 25/6, far greater than Norm's payoff. 5.2.2. Prisoner's Dilemma In Exercise (2) of the previous section, the reader was asked to verify that the maximin value for each player in Prisoner's Dilemma is —5, and that there is only one equilibrium pair of mixed strategies: ((0,1), (0,1)). The payoff to both players is then —5 for this equilibrium pair. As observed before, this game, like any

game with a single equilibrium pair, is solvable in the Nash sense. Now if the row player plays D (the second row), the column player has no reasonable choice; he must play D also. It is therefore highly likely that the payoff pair (D, D) would indeed be played. This is true even though it is not Pareto optimal [since (£), D) is dominated by (C, C)]. Note also that the maximin solutions are the same as the equilibrium pair. Our conclusion about Prisoner's Dilemma is hard to accept because we want the two prisoners to act more like decent and intelligent human beings. If they do so, the payoff pair will surely be (—2, —2). The trouble is that the game is presented to us as noncooperative. The cooperative variant of it is different and has a different solution. There is a third way to look at this which we discuss shortly. 5.2.3. Another Game The game given in Bi-matrix (5.2) was shown to have only one equilibrium pair of mixed strategies: ((1/4, 3/4), (1/3,2/3)). It is thus solvable in the Nash sense. In Exercise (4) of the previous section, the reader was asked to verify that the maximin value for the row player is 2/3, and for the column player, — 1. Now, is it reasonable for the row player to play (1/4,3/4) (his mixed strategy in the equilibrium pair)? If he does so, the column player gains —1/4 playing either column. However, the row player's expected payoff is 1 if the column player plays column 1, and only 1/2 if she plays column 2. This latter payoff is less than the row player can guarantee himself by playing according to the maximin solution. It is easy to verify that if both play according to the maximin solutions, the column player's actual payoff is —1/6, which is greater than the payoff of —1/4 from the equilibrium pair. It is therefore likely that if the row player does

128 5. Non-Zero-Sum Gaines play (1/4, 3/4), the column player would indeed play the second column in order to force him to change to the maximin solution. 5.2.4. Supergames Let us think about Prisoner's Dilemma in a different way. First of all, there are many conflict situations in which there is a dilemma similar to the one facing the prisoners. These range from ones involving personal relationships between two people to others involving great military questions which may determine the survival of civilization. The features common to all these games are: • Both players do well if they cooperate with each other. • If one player plays the cooperative strategy while the other betrays her by playing the defecting strategy, then the defector does very well and the cooperator does badly. • Neither player really trusts the other. We discuss two examples. First, there is the H-Bomb Game. In the period

immediately following World War II, both the United States and the Soviet Union had the scientific and industrial resources to develop the fusion bomb. It was known that such a weapon would be immensely more destructive than even the atomic bombs which the United States had just dropped on two Japanese cities. For each of the countries, the choice was between building the H-bomb (defecting from the wartime alliance between them) or not building it (continuing to cooperate). For each country (but more for the Soviet Union), the cost of defecting would be very great. The resources devoted to the project would be unavailable for improving the living standards of the people; also, the societies would continue to be distorted by the dominant role played by military considerations. Furthermore, if both countries defected, there would be a real possibility of a nuclear war which could throw the human race back to the Stone Age. On the other hand, if one country built the bomb and the other did not, the defector would have such an enormous superiority in military force that the very sovereignty of the cooperator would be at risk. Of course, both countries chose defection as the solution to the game. For a less momentous example, consider two people, Robert and Fran- cesca, who are romantically involved. Each has the choice of either remaining faithful to the other (cooperating), or going out with someone else on the sly (defecting). If one defects and the other cooperates, the defector is happy (for a while, at least). After all, variety is the spice of life. On the other hand, the cooperator feels betrayed and miserable. If both defect, they each feel fairly good for a while, but the level of tension between them detracts from their joy, and the probability of the relationship breaking up

5.2. Solution Concepts for Noncooperative Games 129 is high. This would make both of them very sad. If this Dating Game is viewed as noncooperative, they would probably both defect. On the other hand, if they trust each other enough, the game becomes cooperative and they will probably both remain faithful. Thus we see that game theory is relevant to romance. In general, we say that a bi-matrix game is of Prisoner's Dilemma type if it has the form ((a, a) (6,c) \ V M) (d,d) Y' where c> a > d > b and a > (6 + c)/2. These inequalities say that mutual cooperation (row 1, column 1) is better than mutual defection (row 2, column 2) and that, if you fear your opponent is going to defect, it is best to defect also. The last inequality says that, if the players can trust each other, the best way for both to play is to cooperate. Any bi-matrix game which satisfies these conditions will be referred to as Prisoner's

Dilemma. Another example is ((3,3) (0,5) \ V(5.0) (1,1) J' Some of the conflict situations modeled by Prisoner's Dilemma are such that they can only happen once; others can occur repeatedly. Our discussion of Prisoner's Dilemma and other noncooperative games is correct if we assume that the game is played only once, or that, if played repeatedly, both players decide on a mixed strategy to which they stick. Suppose, however, that we think about a supergame consisting of a certain large number of repetitions of a noncooperative game. Each of the separate plays of the game is called an iteration and a player may use different strategies in different iterations. In particular, a player's strategy may depend on previous moves. One way in which this idea could be used is for one player to estimate her opponent's mixed strategy by keeping track of the frequency with which he chooses his various pure strategies. She could then play so as to do the best possible against this (estimated) mixed strategy. For example, if the opponent in the first 100 iterations has played his first pure strategy 27 times, his second pure strategy 39 times, his third never, and his fourth 34 times, then she could estimate his mixed strategy as (.27, .39,0, .34). The supergame is just as noncooperative as the original one and so cooperation in the true sense is still not allowed, but a form of communication is now possible. To make this clearer, imagine 500 iterations of Prisoner's Dilemma. The extensive form of this supergame is enormous, but really (5.3)

130 5. Non-Zero-Sum Gaines rather simple in structure. Each nonterminal vertex has exactly two children. (A tree with this property is called binary.) We can picture it as follows. Start with the tree for Prisoner's Dilemma. Use each of its terminal vertices as the root of another copy of the tree for Prisoner's Dilemma. Repeat this process as many times as there are iterations in the supergame. There are some information sets in the supergame tree because a player does not know his opponent's move in the same iteration. However, he does know the moves made in previous iterations. Such information could be used in many different ways. For example, a player could use a strategy in which he "signals" his willingness to cooperate by playing C several times in a row. If the other player catches on, they would both play C and do well. If the other player refuses to begin cooperating, then the first player can begin playing D (perhaps sending the cooperation signal again later). Of course, if the opponent never plays anything but £), our player does poorly in those iterations in which he plays C. The point is, however, that there

are relatively few of these in comparison to the total number of iterations. In general, suppose we consider a supergame consisting of M iterations of a certain game (called the base game). There are two kinds of supergame strategies. There are those which take account of moves in previous iterations, and those which do not. We call the first kind adaptive and the second kind forgetful. For example, if the base game is Prisoner's Dilemma, the following are forgetful strategies: • Always cooperate. • Always defect. • Cooperate with probability .5, defect with probability .5. Here is an adaptive strategy for the same game: Cooperate on the first iteration; after ,; iterations have been played, compute pj to be the proportion of the times in these first .; iterations in which the opponent has cooperated; in iteration j +1, randomly choose to cooperate with probability pj, and to defect with probability 1— pj. In this strategy, the opponent is rewarded for frequently cooperating and punished for frequently defecting. Another way of looking at supergame strategies is as follows. If $1 < j < M$, the strategy to be used in the jth iteration can, in general, be a function of the moves made in the first j — 1 iterations. A supergame strategy is forgetful if the strategy used at the jth iteration is actually independent of the previous moves. Some fanciful names have been given to supergame strategies ([Sag93]). The Golden Rule is the strategy of always cooperating. The Iron Rule is the strategy of never cooperating. Tit-for-Tat is the strategy of doing whatever your opponent did on the previous iteration (and cooperating on the first one). In recent years, tournaments have been held in which various

5.2. Solution Concepts for Noncooperative Games 131 supergame strategies have competed against one another. These strategies were submitted by game theorists and the tournaments were organized in round-robin fashion so that each strategy was pitted against every other strategy in a supergame with Prisoner's Dilemma as the base game. The total payoff was added up for each strategy. The consistent winner was Tit- for-Tat. This is interesting partly because of the simplicity of the strategy, and also because it is one which people seem to use instinctively in many conflict situations. Even though Tit-for-Tat won the tournaments by accumulating the highest total payoff against the other strategies, it can never gain more payoff in any single supergame! [See Exercise (8).] The book [Axe84] has much discussion on this subject. There has been a great deal of interest in these supergames on the part of people interested in a variety of fields, including sociology and biology. Explanations

have been sought for the evolution of cooperative behavior in situations where it is difficult to explain in any other way. Some intriguing examples are discussed in [Axe84]. These include both biological evolution and evolution of behaviors in human societies. In particular, the problem is discussed of how a population of individuals acting completely selfishly can evolve into a state in which cooperation is the rule. Exercises (1) Analyze the game given in Exercise (8) of the previous section. (2) Sketch the payoff region for Prisoner's Dilemma. (3) For the game given in Bi-matrix (5.2), prove that (1,0) is a Pareto optimal payoff pair. (4) Analyze the modified Battle of the Buddies game described in Exercise (10) of the previous section. (5) Suppose that you have played 137 iterations of three-finger morra against the same opponent. You notice that he has played (1,1) 41 times, (2,2) 65 times, and (3,3) the rest of the times. Use this data to estimate his mixed strategy. How would you play in response? (6) Sketch the tree for the supergame consisting of two iterations of Prisoner's Dilemma. Be sure to include information sets and payoffs. (7) Consider a supergame consisting of M iterations of Prisoner's Dilemma. Prove that the only equilibrium pair for this supergame is the Iron Rule vs. the Iron Rule. (8) Suppose that a supergame with base game Prisoner's Dilemma (Bi-matrix (5.3)) is being played. One of the players is using Tit-for-Tat as her strategy. Prove that the total payoff accumulated by that player is less than or equal to the total accumulated by her opponent. Prove also that the difference in these total payoffs is at most c — 6.

132 5. Non-Zero-Sum Gaines 5.3. Cooperative Games In the theory discussed in this section, the players are allowed to make binding agreements about which strategies to play. They do not share payoffs, and there are no side payments from one to another. In many cases, the restriction against sharing payoffs is not due to the rules of the game, but arises from the nature of the payoffs. In Battle of the Buddies, it does not make sense for one of the players to give the other some of his "happiness ratings." In Prisoner's Dilemma, the law requires that each prisoner serve his or her own time. The other prisoner cannot take on the burden. In the next chapter, we will study games in which payoffs can be shared. Let us consider Battle of the Buddies as a cooperative game. There is an obvious way in which any two cooperating players would surely decide to play the game. Each evening, they would flip a fair coin and both go to the wrestling match if it comes up heads, and both go to the roller derby if tails. In this way, the expected payoff of each would be 3, far higher

than they could expect in the noncooperative variant. In Prisoner's Dilemma, the obvious way to play cooperatively is for the players to make a binding agreement with each other to cooperate. In Chicken, the players could agree that they both swerve, or could agree to flip a coin to decide who swerves and who does not. In each of these examples a joint strategy is used. The definition is as follows: Definition 5.6. Let π be a two-person game with mxn payoff matrices A and B. A joint strategy is an $m \chi \eta$ probability matrix P — (Pij). Thus, Pij > 0> fºr 1 < г < ?тг, 1 < J < η, and m η г=1 j=l Thus, a joint strategy assigns a probability to each pair of pure strategies. The expected payoff to the row player due to the joint strategy P is m η г = 1 j=l The expected payoff 7Г2(P) to the column player is the same except that dij is replaced by bij. For example, consider the bi-matrix ((2,0) (-1,1) (0,3) λ . U-2.-1) (3,-1) (0,2);· [ЬА)

5.3. Cooperative Games 133 (0,0) Figure 5.4. Cooperative payoff region for Battle of the Buddies. There are six pairs of pure strategies. Suppose that the players agree to play according to the joint strategy described in the matrix (1/8 0 1/3 λ ^ 1/4 5/24 1/12 J ' {°'0) Thus, they will play the pair of pure strategies (row 2, column 1) with probability 1/4, and the pair of pure strategies (row 1, column 3) with probability 1/3. Under this joint strategy, the expected payoff for the row player is (l/8)(2) + (0)(-l) + (l/3)(0) + (1/4X-2) + (5/24)(3) + (1/12)(0) = 3/8, and for the column player it is 17/24. The cooperative payoff region is the set $\{(\pi\chi(P), 7\Gamma 2(P)) : P$ is a joint strategy}. It is a larger set than its noncooperative counterpart [according to Exercise (10)]. Thus, there are more payoff pairs obtainable by the players if they cooperate. Figure 5.4 shows the cooperative payoff region for Battle of the Buddies. It should be compared with Figure 5.3.

134 5. Non-Zero-Sum Gaines Cooperative payoff regions are always convex sets (defined soon); they are also closed1 and bounded. Their vertices are points whose coordinate pairs are among the entries in the bi-matrix. The players in a cooperative game have the task of making an agreement about which joint strategy to adopt. There are two criteria which would surely be important to them. These are: • The payoff pair resulting from the joint strategy they have agreed on should be Pareto optimal. • For each player, the gain from the joint strategy should be at least as great as the maximin value. These considerations lead to a definition. Definition 5.7. The bargaining set for a

two-person cooperative game is the set of all Pareto optimal payoff pairs (ii, v) such that $u > v_1$, $v > v_2$, where v_1 and $V2$ are the maximin values. In Figure 5.4, the bargaining set is the line segment from (1,5) to (5,1). The problem for the players is then to agree on a payoff pair in the bargaining set. In Battle of the Buddies, the symmetry of the game strongly suggests the payoff pair (3,3), the midpoint of the line segment. Figure 5.5 shows the cooperative payoff region for Bi-matrix (5.4). The maximin values are $v_1 — 0$, $V2 = 2$. Thus the bargaining set is the line segment from (0, 3) to (3/4, 2). The row player prefers a payoff pair as far to the right as possible, while the column player prefers one as far up as possible. They negotiate in order to agree on one between the extremes. 5.3.1. Nash Bargaining Axioms The theory which we now present was developed by Nash (see [Nas50]). It is an attempt to establish a fair method of deciding which payoff pair in the bargaining set should be agreed on. The idea is to prove the existence of an arbitration procedure Φ which, when presented with a payoff region P and a status quo point $(\check{u}_0, \hat{o})$ £ -P» will produce a payoff pair (the arbitration pair) which is fair to both players. The status quo point is usually the pair of maximin values. (We will mention other choices later.) It is the pair of payoffs which the players must accept if they reject the payoff pair suggested by the arbitration procedure. The arbitration procedure Φ may be thought of as taking the place of a neutral human arbiter who is called in to settle a conflict. It is required to satisfy the six Nash axioms, Roughly speaking, this means they contain their bounding line segments. See [Apo74] or [PM91] for details.

5.3. Cooperative Games 135 (0,3) , . . e5----—.^ bargaining \7 set (-2,-D ' (3,-1) Figure 5.5. Payoff region for Bi-matrix (5.4). which may be thought of as the principles of fairness and consistency which might guide a human arbiter. It will not only be proved that an arbitration procedure exists but that it is unique. This fact lends much credibility to the idea. We denote the arbitration pair corresponding to payoff region P and status quo point (u0,v0) by *(P,(u0,«o)) = (u*>v*)· With this notation, the Nash axioms are: (1) [Individual Rationality] $u* > u_0$ and $v* > v0$. (2) [Pareto Optimality] (u*,v*) is Pareto optimal. (3) [Feasibility] (u*,v*)€P. (4) [Independence of Irrelevant Alternatives] If P' is a payoff region contained in P and both (u0, vq) and (u*,v*) are in P', then «(P',K^o)) = («*,«*)· (5) [Invariance Under Linear Transformations] Suppose P' is obtained from P by the linear transformation $u' = au + b$, $v' = cv + d$ where a,c>0. Then $\Phi(P;, (au_0 + b, cv_0 + d)) = (au* + b, cv* + d)$. (6)

[Symmetry] Suppose that P is symmetric (that is, (u, v) £ P if and only if (v, u) G P), and that uq — vq. Then u* = v*.

136 5. Non-Zero-Sum Gaines Figure 5.6. One nonconvex and two convex sets. Axioms (1) and (2) say that the arbitration pair is in the bargaining set. Axiom (3) simply says that the arbitration pair can be achieved. Axiom (4) says that if a different game has a smaller payoff region and the same status quo point and if this smaller payoff region contains the arbitration pair (r**, v*), then the arbitration procedure should suggest the same payoff pair for the other game. Axiom (5) says, in part, that if there is a change in the units in which the payoffs are computed, then there is no essential change in the arbitration pair. Axiom (6) says that if the players have symmetric roles as far as both the payoff region and the status quo point are concerned, then they should gain the same payoff. 5.3.2. Convex Sets Before stating and proving the theorem that there is a unique arbitration procedure satisfying Nash's axioms, we must discuss a small part of the theory of convexity. The book [Val64] treats this subject in detail and is recommended to the reader who wants to learn more about it. Definition 5.8. A subset S of 3£n is said to be convex if, for every χ and y in 5, and every number t with $0 < t < 1$, we have tx + (l-t)yeS. In words, this definition says that S is convex if every line segment whose end-points are in S lies entirely in S. In Figure 5.6, two of the sets are convex and the other is not.

5.3. Cooperative Games 137 In order to study how a convex set is built up from a smaller set, we make the following definitions: Definition 5.9. Let F — {a?i, ..., xk} be a finite subset of 3ïη. Then к r'=1 is a convex combination of F whenever tfi,... ,^ are nonnegative numbers whose sum is 1. We see, by induction, that if S is convex, then any convex combination of points of S is in S. Definition 5.10. Let A be any subset of $ftn. The convex hull of A, denoted co(A), is defined to be the set of all convex combinations of finite subsets of A. The reader is asked to prove [in Exercise (12)] that co(A) is actually convex. It follows that every convex set which contains A also contains co(A). For example, a triangle is the convex hull of its vertices; the second convex set in Figure 5.6 is the convex hull of its six vertices; the first convex set in Figure 5.6 is the convex hull of the ellipse which bounds it. The coefficients U in the definition of convex combination have the properties of probabilities. In fact, we have the following: Theorem 5.2. Lei if be a two-person game with mxn

payoff bi-matrix C. Then the cooperative payoff region is the convex hull of the set of points in $\Re^2$ whose coordinates are the entries in the bi-matrix. Proof. If P is a joint strategy, then the corresponding payoff pair is m η (τ1(Π^(p)) = 5]£рцщ, г'=1 j = 1 which is in the convex hull of {cij : 1 < i < m, 1 < j < n}. Conversely, any point of the convex hull is a payoff pair. Π Nash's Axiom (6) referred to symmetric payoff regions. We have Definition 5.11. A subset S of $\Re^2$ is said to be symmetric if (v,u) is in S whenever (u, v) is in 5. Thus, a set is symmetric if it is identical to its reflection in the diagonal χ — y. Figure 5.7 shows two symmetric sets, one convex and the other not.

138 5. Non-Zero-Sum Gaines X Figure 5.7. Two symmetric sets. Definition 5.12. Let A be a subset of $3\pounds2$. The symmetric convex hull is defined to be the convex hull of the set AU{(v,u) : (u,v) eA}. It is denoted sco(A). Thus, the symmetric convex hull is formed by reflecting A in the diagonal and then forming the convex hull of the union of A and its reflection. The reader is asked to prove [in Exercise (13)] that the result of this process is actually symmetric. The following is needed later: Lemma 5.3. Let A be a subset of№ and к be a number such that u + v < k, for every point (u, v) in A. Then the same inequality holds for every point in the symmetric convex hull of A. Proof. The inequality certainly holds for every point in the reflection of A in the diagonal. Thus, it holds for every point of the union of A and this reflection. The set of all points in $ft2 for which the inequality holds is convex. Thus, it contains sco(A). D 5.3.3. Nash's Theorem Theorem 5.4. There exists a unique arbitration procedure Φ satisfying Nash's axioms.

5.3. Cooperative Games 139 Proof. First, we describe how to compute Φ(P, (щ, vq))] then we prove that Nash's axioms are satisfied; then, finally, we show that there is only one such arbitration procedure. In computing Φ, there are two major cases, and three subcases of the second of these. Case (i). There exists (u,v) G P such that и > uq and v > vq. Let K be the set of all (u, v) satisfying these conditions. Then define g(u, v) = (u - u0)(v - v0), for (u, v) G K. Let (и*, v*) G K be the point at which </(u, v) attains its maximum value. [The point (ii*,v*) exists and is unique according to Lemma 5.5 proved later.] Define Φ(P, («o,vo)) = («*,«*)· Case (ii). No (u,v) G P exists such that и > щ and v > vq. Consider the following three subcases. Case (iia). There exists (г*о, v) G P with v > vq. Case (iib). There exists (u, vq) G P with и > щ. Case (iic). Neither (iia) nor (iib) is true. The first thing to notice about these

subcases is that (iia) and (iib) cannot both be true. For, suppose that they are and define $(«v) = (1/2)(«0)») + (1/2)(«,»o)$. Then (u',v') is in P (by convexity) and satisfies the condition of Case (i). Since Case (i) does not hold, this is a contradiction and thus (iia) and (iib) cannot both hold. Now we define $\Phi(P, (\Gamma^*o, \wedge o))$ in each of these subcases. In Case (iia), let v* be the largest v for which (щ,y) is in P2; then define $\Phi(P, («0, «o)) = («0i \ v^*)$· In Case (iib), let u* be the largest и for which (u, vq) is in P; then define $»(P,(tio,«o)) = (tt^*,«o)$. In Case (iic), let $\Phi(P, (\Gamma^*o, \wedge o)) = (\wedge O)\wedge o)$· Our arbitration procedure is now defined in all possible cases. We now prove that the Nash axioms hold. Axioms (1) and (3) are obvious in all the cases. Suppose that Axiom (2) is not true. Then there exists (u, v) G P which dominates (u*,v*) and is different from it. Now, in Case (i), we have $(u - no) > («^* - «o)i \ (v - vo) > (v^* - vo)$, and at least one of these inequalities is strict [since (u, v) φ (и*, v*)]. Thus, flf(ti,i;) $>g(u^*,v^*)$. This is a contradiction. In Case (iia), we must have u* — щ — w, since Case (iib) does not hold. Thus, $v > v^*$. But this contradicts the definition of v*. Case (iib) is dealt with in a similar way. In Case (iic), (Γ^{**}, v^*) — $(\Gamma^*o, \wedge o)$· If 2v* exists because P is closed and bounded.

140 5. Non- Zero- Sum Gaines и > ŭo, then Case (iib) holds; if $v > vq$, then Case (iia) holds. Since neither of these does hold, we have again reached a contradiction. We conclude that Axiom (2) holds. We now show that Axiom (4) holds. In Case (i), the maximum value of g over K Π P' is less than or equal to its maximum over K. But (u*, v*) belongs to P' and so the two maxima are equal. Thus $»(P/>(ti0,«o)) = »(P,K,«o))$. Cases (iia) and (iib) are similar; Case (iic) is easy. To verify Axiom (5), suppose first that Case (i) holds. Then Case (i) also holds for payoff region P' with status quo point (aщ + Ь, cvo + d). Also $(u' - (aщ + b))(v' - (cv0 + d))$ — $ac(u - u0)(v - v0)$. Since a, c > 0, the maximum of the left-hand side of the equation above is attained at (au* -f b,cv* + d). Thus the axiom holds in Case (i). The reasoning in the other cases is similar. Finally, we come to Axiom (6). Suppose that u* φ v*. By the symmetry condition, (v*,ii*) G P. In Case (i), we have $g(v^*,u^*) = g(u^*,v^*)$. By Lemma 5.5, g attains its maximum only at one point, and so this is a contradiction. Cases (iia) and (iib) cannot hold since, if one held, so would the other (by symmetry). As proved above, this is impossible. Case (iic) is obvious. It remains to prove that Φ is unique. Suppose that there is another arbitration procedure Φ which satisfies the Nash axioms. Since they are different, there is a payoff region P and a status quo point (ŭo,^o) G P such that $(«, v) = \Phi(P,$

(ti0> v0)) φ »(P, (ti0, v0)) = («*,«*). Suppose Case (i) holds. Then и* > щ and v* > vq. We define , U - Щ , V - V0 и = }v = . U* — Uq V* — Vq This linear change of variables takes (ĭĭo,^o) into (0,0) and (ti*,v*) into (1,1). Thus, by Axiom (5), *(P/I(0,0)) = (1,1). Also by Axiom (5), «(J", (0,0)) 54(1,1). We now verify that if (wr, v') G P', then v! + v' < 2.

5.3. Cooperative Games 141 Suppose this is not true. By the convexity of P', t{u\ v') + (1 - 0(1,1) G P', 0 < * < 1. Now define, for 1 < t < 1, h(t) = </(*(</, v') + (1 - <)(1,1)) - (<«' + (1 - t))(iv' + (1 - <)). Then ft(0) = 1 and the derivative of h(t) is h\t) = 2tu'v' + (1 - 2<)(ti; + i/) - 2(1 - 0, so that ft'(0) = ti' + i/ - 2 > 0. Thus, there exists a small positive t such ft(tf) > 1. But this contradicts the definition of (ti*, v*), since </(l, 1) = 1. Now let P be the symmetric convex hull of P1. Then, by Lemma 5.3, s +1 < 2 for all (s,0 G P, and, therefore, if (a, a) G P, a < 1. Since P is symmetric, it follows from Axiom (6) that Φ(P>(0,0)) = (1,1)) since, otherwise, there would exist a point (a, a) in P with a > 1. But, then, by Axiom (4), applied to Φ, Φ(P',(0,0)) = (1,1). This is a contradiction which proves uniqueness in Case (i). This leaves Cases (iia), (iib), and (iic). Now, (iic) is easy by Axiom (1). Cases (iia) and (iib) are similar so we only consider the first of these. Since we are not in Case (i), we see, from Axiom (1), that й = uq = u*. Since both (ti*,v*) and (TZ, v) are Pareto optimal, v = v*. This contradicts the assumption that (TZ, v) and (u*,v*) are different. Π Lemma 5.5. Let P be a payoff region and (τ/o,^o) G P. Suppose thai there exists a point (u,v) G P witt li > Uq, V > Vq, and let K be the set of all (u,v) G P satisfying these inequalities. Define, on K, g(u,v) = (u - u0)(v - v0). ТЛеп gf attains its maximum on K at one and only one point.

142 5. Non-Zero-Sum Gaines Proof. The set K' — {(u, v) £ P : и > щ, v > v0} is a closed bounded set containing K. By a theorem of mathematical analysis (see [Apo74] or [PM91], for example), the function </, being continuous, attains its maximum on K'. Clearly, the maximum of g over K' is the same as its maximum over K. It remains to prove that the maximum is attained only once. Suppose it is attained at two different points, (ĭĭι,vι) and (1*2,^2)· We let M - max0(u, v) = g(uuvi) = g{u2) v2). Now either ui > u2) ^i < v2l or Щ < ^2, ^1 > ^2· We carry out the proof assuming the first possibility. The proof with the second possibility assumed is similar. By the convexity of P, (ti3> v3) = (l/2) (tn, «1) + (l/2)(ti2> v2) e P. We compute 0(«з,«з) + + щ+и2 \ (v\ + v2 2 Uo) {— Γ 2 2 J V 2 2 y (l/4)[(«i - «o)(«i ~ vq) + («2 - «o)(«2 - «0) («1 - U0)(V2 -

Vq) + («2 - «θ)(«1 - «θ)] (1/4)[2M + 2M + (tii - <io)("2 - «0) («2 ~ «o)(i>i - «0) - («1 - «o)(«i - «0) («2 ~ U0)(V2 - V0)] M + (1/4)[(v2 - «o)(«i - «2) + («i - «o)(«2 - «1)] M+ (1/4)[(tn-ti2)(i;2-«i)] M. This is a contradiction because M is the maximum. Π

5.3. Cooperative Games 143 arbitration pair (1,-2) Figure 5.8. Cooperative payoff region for Bi-matrix (5.6). 5.3.4. Computing Arbitration Pairs The proof of Theorem 5.4 actually contains an algorithm for computing the arbitration pair. We illustrate the computation with some examples. First, let us look at Battle of the Buddies. The payoff region, which is shown in Figure 5.4, is a symmetric set. The pair of maximin values is (5/6,5/6). Taking this pair to be (r*o,^o)> we see that Axiom (6) applies. Thus the arbitration pair is of the form (a, a). Since (a, a) must be Pareto optimal [by Axiom (2)], we see that a = 3. Thus, we get the point which common sense led us to guess. For another example, consider the cooperative game given by the bi- matrix f (2,-1) V(-i.: (-2Д) (0,2) ,2) (0,2) (1,-2))· (5'6) The maximin values are easily computed to be v\— —2/5,^2 = 1· The payoff region is shown in Figure 5.8. The pair (— 2/5,1) is indicated with a small square. Taking (йo,^o) = (—2/5,1), we see that the arbitration pair is to be found among the points in the payoff region which dominate (—2/5,1),

144 5. Non-Zero-Sum Gaines and which are Pareto optimal. A glance at Figure 5.8 shows that these points constitute the line segment from (1,1) to (0,2) (indicated with a heavy line). We see also that, in the terminology of the proof, Case (i) holds. Thus, the arbitration pair can be computed by finding the point on this line segment at which the maximum of the function g(u, v) - (u - u0)(v - v0) occurs. Now, on this line segment, v = —ti + 2, and so the problem reduces to one of maximizing a function of one variable. We have flf(ti, v) = (u + 2/5)(v - 1) = (ti + 2/5)(-ti + 1) = -u2 + 3u/5 + 2/5. By the methods of calculus, we find that the maximum is attained for и = 3/10, v — 17/10. This point is indicated as the arbitration pair in Figure 5.8. For another example, consider the bi-matrix /(5,1) (7,4) (1,10) N \(\,\) (9,-2) (5,1))' ^ The payoff region is shown in Figure 5.9. The maximin values are easily computed to be v\ = 3 and V2 = 1. From the figure, we see that the bargaining set is the union of two line segments. These are the one from (3,8) to (7,4), and the one from (7,4) to (8,1). The bargaining set is indicated with a heavy line. With (uq, ^o) = (3,1), we are clearly in Case (i). Thus we must maximize g(u,v) = («-3)(v- 1)

over the bargaining set. At the three points (3,8), (7,4), and (8,1), we have $</(3,8) = 0, ^(7,4) = 12, 3(8,1) = 0$. On the upper line segment, $v = — u + 11$ and so $g(u, «) = («- 3)(-u + И - 1) = (u - 3)(—ti + 10) = -u2 + 13u - 30$. Setting the derivative with respect to и equal to zero gives $u = 13/2$ and $v = 9/2$. This point is in the bargaining set and gives us $0(13/2, 9/2) = 49/4$. On the lower line segment, $v = — 3u + 25$ and so $g(u, v) = (u- 3)(-3« + 25 - 1) = (u - 3)(-3u + 24) = -3u2 + 33u - 72$.

5.3. Cooperative Games 145 arbitration pair (9,-2) Figure 5.9. Payoff region for Bi-matrix (5.7). Setting the derivative with respect to и equal to zero gives $u = 11/2$ and $v = 17/2$. This point is, however, on the continuation of the line segment beyond (7,4) and is thus outside the bargaining set. Since $49/4 > 12$, the arbitration pair is (13/2, 9/2). 5.3.5. Remarks The method discussed just now, with the status quo point taken to be the pair of maximin values, is called the Shapley procedure. It is interesting to examine the effect of the status quo point on the arbitration pair. The arbitration pair for Bi-matrix (5.6) is (3/10, 17/10). This seems much more favorable to the column player than to the row player; for both, the maximum possible payoff is 2. The arbitration pair gives the row player only 15 percent of this maximum, but gives the column player 85 percent. This difference is, in part at least, a result of the fact that the status quo point is more favorable to the column player. A verbal description of the situation would be that the column player is in a stronger bargaining position. If the negotiations collapse, she still gets a good payoff, but the row player does not. Bi-matrix (5.7) also illustrates

146 5. Non-Zero-Sum Games this phenomenon. In this game, the row player is in a better position as far as the status quo point is concerned. The arbitration pair is (13/2, 9/2), which is also more favorable to the row player. The definition and properties of the arbitration procedure are valid for any choice of status quo point. This fact is made use of in the Nash procedure. In this method, the status quo point is determined by means of preliminary "threats" by the players. These threats are strategies (pure or mixed) and the goal of each player is to obtain the most favorable possible status quo point before arbitration begins. This preliminary phase can be viewed as a noncooperative game with infinitely many pure strategies. We do not discuss the Nash procedure because this sort of game is outside the scope of this book, but the interested reader should see [LR57], [Tho84], or [Owe82]. Exercises (1) For

the noncooperative game described by the bi-matrix in Exercise (7) on page 124, show that (2,-3), (—1,3), and (0,1) are Pareto optimal. (2) Find the arbitration pair for the cooperative game described by the bi-matrix in Exercise (7) on page 124. (3) Find the arbitration pair for the bi-matrix game (-1,-1) (4 (0,4) (4,0) ^ 1,-D J (4) Find the arbitration pair for the bi-matrix game (5) Find the arbitration pair for the modified Battle of the Buddies game described in Exercise (10) on page 124. (6) Find the arbitration pair for Chicken.. (7) Modify Chicken by changing the payoff to the nonswerver from 3 to 2. Compute the arbitration pair for the modified game and compare this answer with the arbitration pair for the original Chicken. How does this innocent-seeming change in the payoff numbers affect the "optimal" way of playing the game? (8) Find the arbitration pair for the game given in Bi-matrix (5.1). (9) Find the arbitration pair for the bi-matrix game / (-1/2,0) (-1/2,-4) \ (1,2) (-2,4) . V (4,-4) (-1/2,0) J (10) Prove that the noncooperative payoff region for a two-person game is a subset of the cooperative payoff region.

5.3. Cooperative Games 147 (11) Let P be the probability matrix (5.5) on page 133. Do there exist mixed strategies ρ and q for the two players such that Pi4j = P\j for all r, j? (12) Prove that the convex hull of a set is convex. (13) Prove that the symmetric convex hull of a set is symmetric.

6 N-Person Cooperative Games In the previous chapter, we studied cooperative games in which the players could coordinate their strategies but could not share payoffs. In the games considered in this chapter, the players are permitted to cooperate fully. Sets of players (called coalitions) cannot only make binding agreements about joint strategies, but can also agree to pool their individual payoffs and then redistribute the total in a specified way. In order for this latter kind of cooperation to be possible, we must assume that the payoffs are in some transferable quantity, such as money. The number of players is not explicitly restricted to be greater than two, but might as well be —the case of two players turns out either to be trivial or to have been covered in our earlier work. The theory developed in this chapter applies both to zero-sum and nonzero-sum games. A large part of it goes all the way back to [vNM44], that is, to the beginning of game theory. 6.1. Coalitions The idea of a coalition plays an essential role in the theory we are about to develop. As indicated above, a coalition is simply a subset of the set of players which forms in order to coordinate strategies and to agree on how the total payoff is

to be divided among its members. As in our earlier discussion of cooperative games, it must be understood that the agreements

150 6. N-Person Cooperative Games Table 6.1. A 3-player game. Strategy triples Payoff vectors (1.1.1) (-2,1,2) (1.1.2) (1,1,-1) (1.2.1) (0,-1,2) (1.2.2) (-1,2,0) (2.1.1) (1,-1,1) (2.1.2) (0,0,1) (2,2,1) (1,0,0) (2,2,2) (1,2,-2) which the members of the coalition make among themselves are absolutely binding. We introduce some notation. The set consisting of all N players is denoted V. A coalition is denoted by an uppercase script letter: <S, T, W, etc. Given a coalition 5CP, the counter-coalition to S is Sc=V-S = {PeV:P<£S}. Let us consider a 3-player example. Its normal form is given in Table 6.1. Each player has two strategies, denoted 1 and 2. Thus, there are eight combinations of strategies and eight corresponding 3-tuples of payoffs. In this game, V — {Pχ,Pι,Pζ}- There are eight coalitions. Three of them contain only one player each. These are {Pl},{P2},{P3h Then, there are three of two players each, namely, {P1,P2},{P1>P3},{P2>P₃}. Finally, V, consisting of all the players, is called the grand coalition, and the counter-coalition to V is 0, the empty coalition. In general, in a game with N players, there are 2^ coalitions. 6.1.1. The Characteristic Function If a coalition S has formed, it is natural to think of the game as a contest between two "players," these being the coalition S and the counter-coalition Sc. This two-"person" game is non-cooperative. Indeed, if the coalition and the counter-coalition cooperated, the grand coalition would form (in place of S). Suppose, then, that we are dealing with an N-player game for which V = {Pi,...tPN},

6.1. Coalitions 151 and for which the strategy set for player Pr· is denoted by X{. Let S C V be a coalition. We assume, for now, that S is not empty and does not contain all the players. The pure joint strategies available to this coalition are the members of the Cartesian product of those X^s for which P,- £ S. Similarly, the pure joint strategies available to the counter-coalition are the members of the Cartesian product of those X^s for which P; £ S. The game between the coalition S and the counter-coalition Sc is a bi-matrix game. The rows correspond to the pure joint strategies available to «S, and the columns to those available to Sc. An entry in the bi-matrix is a pair whose first member is the sum of the payoffs to the players in the coalition, and whose second member is the sum of the payoffs to the players in the counter-coalition (given that both "players" play according to the pure joint strategies labeling the row

and column in which the entry appears). In the game given by Table 6.1, let us consider the coalition S = {Pι,Pβ}. Then Sc — {P2}. The coalition has four pure joint strategies available to it. These may be designated: (1,1), (1,2), (2,1), and (2,2), where, for example, (2,1) means that Pi plays 2, and P3 plays 1. The counter-coalition has only two pure strategies: 1 and 2. The 4 χ 2 bi-matrix is (6.1) Here, for example, the pair in row 1 and column 2 is (2, —1) because the payoff 3-tuple corresponding to strategy combination (1,2,1) is (0,-1,2). Then the sum of payoffs to the players in the coalition is $0 + 2 = 2$, and the payoff to the counter-coalition is —1. The maximin value for the coalition (computed from the bi-matrix) is called the characteristic function of S and is denoted v(S). As explained in our previous discussion (page 119) of maximin values, the members of S are guaranteed to be able to choose a joint strategy with which they can gain a total payoff of at least v(S). The characteristic function has as its domain the set of all coalitions, and measures the values of these coalitions. Let us do some computing for the example in Table 6.1. We know that ι/({Pι, P3}) is the value of the matrix game whose entries are the first terms in the entries in Bi-matrix (6.1). This matrix is (1,1) (1,2) (2,1) (2,2) 1 (0,1) (0,1) (2,-1) (1,0) 2 (2,-1) (-1,2) (1,0) (-1,2) 0 0 2 1 2\ -1 1 -1/

152 6. N-Person Cooperative Games Two of the rows are dominated. The resulting 2x2 matrix game is easily solved to give "(№,ft}) = 4/3. The characteristic function of the counter-coalition, {P2}, is the value of (1 1 -1 o\ V -1 2 0 2/ Two columns are dominated. The resulting 2x2 matrix game is solved to give *({P2}) = -1/3. Computing in a similar way, we have u({PltP2}) = l, K{P₃}) = 0, ./({ft. ft}) = 3/4, i/({Pi}) = 1/4. The value of the characteristic function for the grand coalition is simply the largest total payoff which the set of all players can achieve. It is easily seen that v(V) = 1. Finally, by definition, the characteristic function of the empty coalition is ι/(β) = 0. By examining the values of the characteristic function, we can speculate about which coalitions are likely to form. Since Pi does better playing on his own than do the other two, it might be that P2 and P3 would bid against each other to try to entice Pi into a coalition. In exchange for his cooperation, Pi would demand a large share of the total payoff to the coalition he joins. He would certainly ask for more than 1/4, since he can get that much on his own. On the other hand, if he demands too much, P2 and P3 might join together, exclude Pi, and gain a total of 3/4. There is an interesting theorem about the characteristic

function. It says that "in union, there is strength." Theorem 6.1 (Superadditivity). Let S and T be disjoint coalitions. Then $v(S \cup T) > v(S) + v(T)$. Proof. By definition of the characteristic function, there is a joint strategy for S such that the total payoff to the members of S is at least $i/(\text{«}S)$. A similar statement holds for T. Since S and T are disjoint, it makes sense for each member of S and each member of T to play in accordance with these maximin solutions. The result is a joint strategy for the union of the two coalitions such that the total payoff is guaranteed to be at least $v(S) + v(T)$.

6.1. Coalitions 153 The maximin value for the union might be even larger. Π The example we studied earlier shows that the inequality in the theorem may be strict. For instance, $v(\{Pu\ P3\}) = 4/3 > 1/4 = 1/4 + 0 = i/(\{Pi\}) + u(\{Ps\})$. The first of the following two corollaries can be proved using the theorem and mathematical induction. The second corollary is a special case of the first. Corollary 6.2. If $Si,. ..,£\&$ are pairwise disjoint coalitions (that is, any two are disjoint), then $MU^* > i>(*) \cdot \ \text{«}=i / *=i$ Corollary 6.3. For any N-person game, $N \ ^)>\Sigma^P \cdot \text{»} \cdot \text{»}=1$ As far as the formation of coalitions is concerned, a game can be analyzed using only the characteristic function. This suggests the following definition. Definition 6.1. A game in characteristic function form consists of a set $V = \{Plt...,PN\}$ of players, together with a function $i/$, defined for all subsets of V, such that $\nu/(\beta) = 0$, and such that superadditivity holds. That is, $v(S \cup T) > v(S) + v(T)$, whenever S and T are disjoint coalitions of the players. As an abbreviation, we will use the single symbol ν to designate a game in characteristic function form.

154 6. N-Person Cooperative Games 6.1.2. Essential and Inessential Games We can single out a class of games which are trivial as far as coalitions are concerned. In fact, they have the property that there is no reason to prefer any one coalition over any other. Definition 6.2. An TV-person game ν in characteristic function form is said to be inessential if $\kappa?) = 1 > (\{я\}) \cdot *=i$ A game which is not inessential is said to be essential. In other words, a game is inessential if the inequality in Corollary 6.3 is actually an equality. In fact, for such a game, a stronger statement is true. Theorem 6.4. Let S be any coalition of the players in an inessential game. Then $\kappa^*) = \Sigma > (\{^\}) \cdot$ pes Proof. Suppose not. Then, by Corollary 6.2, $"(5) > \Sigma"(\{^\})$. pes and so, by superadditivity, $v(V) > v(S) + v(S<) > £>(\{*\}), *=1$ which contradicts the definition of an inessential game. Π Thus, in an inessential game, there is no reason for a

coalition actually to form—cooperation does not result in a greater total payoff. The fact that a game is inessential does not make it unimportant. It simply means that there is no reason for a coalition to form. The following theorem shows that there are many inessential games which are, nevertheless, important. Theorem 6.5. A two-person game which is zero-sum in its normal form is inessential in its characteristic function form. Proof. The minimax theorem says that $i>(\{Pi\})$ and $^\wedge(\{ft\})$ are negatives of each other, and thus their sum is zero. But, also, $v(V) — 0$, by the zero-sum property.

6.1. Coalitions 155 Thus, $v(V) = v(\{Pi\}) + v(\{P2\})$. D In the context of the previous chapter, an inessential two-person game is one for which there is no advantage to be gained by cooperating. Zero-sum games with more than two players can be essential. See Exercise (5). Exercises (1) Verify the values of the characteristic function for the game shown in Table 6.1. (2) Compute the characteristic function form of the game shown in Table 1.1, page 26. (3) Compute the characteristic function (for all coalitions) for the game shown in Table 1.2, page 29. (4) Show that the following (non-zero-sum) bi-matrix game is inessential: ((0,0) (1,-2) λ ^ (-1,1) (1,-1) J' (5) The 3-person game of Couples is played as follows. Each player chooses one of the other two players; these choices are made simultaneously. If a couple forms (for example, if Pi chooses P$, and P$ chooses P^), then each member of that couple receives a payoff of 1/2, while the person not in the couple receives — 1. If no couple forms (for example, if P\ chooses P2, P2 chooses P3, and P3 chooses Pχ), then each receives a payoff of zero. Prove that Couples is zero-sum and essential. (6) Let V — {P\, P2, Pz»Pa} be a set of players. Let o, 6, c, d, e, / be nonnegative numbers such that o-j-d < 1 6-j-e < 1 c+/ < 1. Define v by $1/(0) = 0$, $v(V) = l$, $i/(\text{any 3-person coalition}) = 1$, $i/(\text{any 1-person coalition}) = 0$, *№.«}) = ». »({РиРг}) = Ь, u({P1,Pi}) = c, v({P3,PA) = á, $1/(\{P2,P4\}) = e$, $v(\{P2,Pi\}) = f$. Prove that v is a game in characteristic function form.

156 6. N-Person Cooperative Games 6.2. Imputations Suppose that a coalition forms in an TV-person game. The problem we now wish to study is that of the final distribution of the payoffs. It is this which, presumably, the players themselves would be most interested in. Indeed, a player considering whether to join a given coalition would want to know how much she gains from doing so. Now, the amounts going to the various players form an TV-tuple χ of

numbers. We will argue that there are two conditions which such an TV-tuple must satisfy in order that it have any chance of actually occurring in the game. These are individual rationality and collective rationality. An TV-tuple of payments to the players which satisfies both these conditions is called an imputation. After the formal definition, we will try to justify the appropriateness of the two conditions. Definition 6.3. Let v be an TV-person game in characteristic function form with players $V = \{PU...,PN\}$. An TV-tuple χ of real numbers is said to be an imputation if both the following conditions hold: • {Individual Rationality) For all players Pг·, • {Collective Rationality) We have $i = 1$ The condition of individual rationality is easy to motivate. If for some i, then no coalition giving P,· only the amount X(would ever form—Pi would do better going on his own. As for collective rationality, we first argue that N i=1 Suppose that this inequality is false. We would then have N $\beta = v\{v\}$-$\Sigma\chi\acute{\imath}>°$- « = 1

6.2. Imputations 157 Thus, the players could form the grand coalition and distribute the total payoff, v(V)) in accordance with *{ = x{ + β/N, thus giving every player more. Hence, if χ is to have a chance of actually occurring, (6.2) must be true. We then argue that N 5><κη (6-3) *=1 To see this, suppose that χ actually occurs. That is, suppose that S is a coalition and that the members of it, and the members of its counter- coalition, agree to χ as their way of dividing the payoffs. Then, using superadditivity, N г=1 Pi £S Pi£Sc The combination of (6.2) and (6.3) gives the collective rationality condition. In the game shown in Table 6.1, any 3-tuple χ which satisfies the conditions #i + #2 + #з = 1, and *i > 1/4, X2 > -1/3, X3 > 0, is an imputation. It is easy to see that there are infinitely many 3-tuples which satisfy these conditions—for example, (1/3,1/3,1/3), (1/4,3/8,3/8), (1,0,0). The following theorem shows that essential games always have many imputations, while inessential ones do not. Theorem 6.6. Let v be an N-person game in characteristic function form. If v is inessential, then it has only one imputation, namely, S = (u({P1}),...M{Pn})). If v is essential, then it has infinitely many imputations. Proof. Suppose first that v is inessential, and that χ is an imputation. If, for some J,

158 6. N-Person Cooperative Games then $\Sigma>>\Sigma>(\{\rho<\}) = $ κη *=1 *=1 This contradicts collective rationality and shows that, for each i, Now suppose that v is essential and let N β=»iP)-Σ,vHρM>o· *=1 For any $\Lambda\Gamma$-tuple a of nonnegative numbers summing to /?, we have that Xi = v({Pi}) + (*i defines an

imputation. Since there are obviously infinitely many choices of a, there are also infinitely many imputations. Π For an essential game, there are too many imputations. The problem is to single out those which deserve to be called "solutions." For example, in the game shown in Table 6.1, none of the three imputations listed earlier for it seems likely to occur. For instance, consider the imputation (1/4,3/8, 3/8). It is unstable because players Pi and P2 could form a coalition, gain a total payoff of at least 1, and divide it between them so that each gains more than her entry in (1/4,3/8, 3/8) gives her. 6.2.1. Dominance of Imputations The following definition attempts to formalize the idea of one imputation being preferred over another by a given coalition. Definition 6.4. Let ν be a game in characteristic function form, let S be a coalition, and let #, y be imputations. Then we say that χ dominates y through the coalition S if the following two conditions hold: • #* > У% for all Pi £ <S. The notation for this relation is x >s У- The second condition in this definition says that χ is feasible, that is, that the players in S can attain enough payoff so that X{ can actually be paid out to each P,- G S. Since the inequality in the first condition is strict, every player in S does better under χ (compared to y).

6.2. Imputations 159 In the game whose normal form is shown in Table 6.1, we see that (1/3,1/3,1/3) dominates (1,0,0) through the coalition {P2, Лз}, and that (1/4,3/8,3/8) dominates (1/3,1/3,1/3) through the same coalition. Also, (1/2,1/2,0) dominates (1/3,1/3,1/3) through {Pu P2}· 6.2.2. The Core It seems intuitively clear that an imputation which is dominated through some coalition would never become permanently established as the way in which the total payoff is actually distributed. Instead, there would be a tendency for the existing coalition to break up and be replaced by one which gives its members a larger share. This idea motivates the following: Definition 6.5. Let ν be a game in characteristic function form. The core of ν consists of all imputations which are not dominated by any other imputation through any coalition. Thus, if an imputation χ is in the core, there is no group of players which has a reason to form a coalition and replace χ with a different imputation. The core is the first "solution concept" which we define for TV-person cooperative games. As we shall soon see, it has a serious flaw: The core may be empty! Deciding whether an imputation is in the core seems difficult if we use only the definition. The following theorem makes the job easier. Theorem 6.7. Let и be a game in characteristic function form with N 'players, and let χ be an imputation. Then χ is in the core of ν if and only if Σ *i > κ*), (6·4) Pi£S for

every coalition S. Proof. First, suppose that (6.4) holds for every coalition S. If some other imputation ζ dominates χ through a coalition «S, then Σ *·' > Σ *·' * "№> Pies p^s which violates the feasibility condition in Definition 6.4. Thus, χ is in the core. Now, suppose that χ is in the core, and suppose that S is a coalition such that Pi €S

160 6. N-Person Cooperative Games We must derive a contradiction. We observe that S φ V. If this were not true, the collective rationality condition in the definition of imputation would be violated. We next wish to show that there exists Pj £ Sc such that $x_j > v(\{P_j\})$. If this were not true, we would have, using superadditivity, N «=1 Pi£Sc which violates collective rationality. Thus, we can choose Pj £ Sc such that there exists a with $0 < a < X_j - v(\{P,·\})$, and a $< v(S) - \Sigma x_i$- Pi£S Now, with κ denoting the number of players in «S, we define a new imputation ζ by $Z_i = X_i + a/k$ for Pi £ «S, $^· = X_j - a$, and $Z_j = X_i$ for all other i. Then ζ dominates ж through «S, and so the assumption that χ is in the core is contradicted. Π The following corollary states a more convenient form of this result. The difference is that we do not have to check separately whether χ is an imputation. Corollary 6.8. Let и be a game in characteristic function form with N players, and let χ be an N-tuple of numbers. Then χ is an imputation in the core if and only if the following two conditions hold: • $\Sigma \rho i$€5 x_i ^ U(S) for everU coalition S. Proof. An imputation in the core certainly satisfies the two conditions. For the other half of the proof, let χ satisfy both conditions. Applying the second condition to one-player coalitions shows that individual rationality holds. The first condition is collective rationality, and thus χ is an imputation. It is in the core, by the theorem. Π

6.2. Imputations 161 Let us use this theorem to find the core of the game shown in Table 6.1. By the corollary, a 3-tuple (#1,#2,#з) is an imputation in the core if and only if (6.5) (6.6) (6.7) (6.8) (6.9) From (6.5), (6S)^ and (6.7), we see that xs = 0 and x\ + #2 = 1· From the first of these and (6.8) and (6.9), we have that x\ > 4/3 and #2 > 3/4. Adding these, we get that x\ + x^ > 25/12 > 1. This is a contradiction, and so we conclude that the core of this game is empty. As a second example, consider the 3-player game whose characteristic function is given by (6.10) X\ + #2 + #3 Xl X2 X3 Xl + X2 Xl + X3 X2 + #3 = > > > > > > 1, 1/4, "1/3, 0, 1, 4/3, 3/4. "(№) кт» KJ№) KíΛ, ft}) KiA.ft}) KJ№.ft}) Kift.ft.ft» — = = = = = = = "1/2: 0, -1/2: 1/4, 0, 1/2, 1. The reader should verify that superadditivity holds for this example. We see that a 3-tuple χ is an

imputation in the core of this game if and only if the following all hold: $X1 + X2 + X3$ $X1$ $X2$ $X3$ $X1 + X2$ $X1 + X3$ $X2 + X3$ $\geq \geq \geq \geq \geq \geq \geq$ 1, -1/2, 0, -1/2, 1/4, 0, 1/2. This system has many solutions. For example, (1/3,1/3,1/3) is in the core. As a third example, consider the following:

162 6. N-Person Cooperative Games Example 6.1 (The Used Car Game). A man named Nixon has an old car he wishes to sell. He no longer drives it, and it is worth nothing to him unless he can sell it. Two people are interested in buying it: Agnew and Mitchell. Agnew values the car at $500 and Mitchell thinks it is worth $700. The game consists of each of the prospective buyers bidding on the car, and Nixon either accepting one of the bids (presumably the higher one), or rejecting both of them. We can write down the characteristic function form of this game directly. Abbreviating the names of the players as N,A,M, we have $\#\}) = \varphi\})^{\wedge}(M) = o(\ i/(\{JV, A\}) = 500$, $v(\{N, M\}) = 700$, $v/(\{A, M\}) = 0$, $i/(\{JV,A,M\}) = 700$. The reasoning behind these numbers is as follows. Consider first the one- player coalition {N}. In the game between this coalition and its counter- coalition, N has only two reasonable pure strategies: (i) Accept the higher bid, or (ii) reject both if the higher bid is less than some lower bound. There exists a joint strategy for the counter-coalition {A, M} (namely, both bid zero) such that, if it plays that way, the maximum of TV's payoffs over his strategies is zero. By definition of the maximin value, $v(\{N\}) = 0$. The characteristic function values of the other two one-player coalitions are zero since the counter-coalition can always reject that player's bid. The coalition {JV, A} has many joint strategies which result in a payoff to it of $500, independently of what M does. For example, A could pay N $500 and take the car. The payoff to N is then $500, and the payoff to A is zero (the value of the car minus the money). On the other hand, they cannot get more than $500 without the cooperation of M. Similarly, $i/(\{TV, M\}) = 700$. Finally, the grand coalition has characteristic function value $700, since that is the largest possible sum of payoffs (attained, for example, if M pays N $700 for the car). An imputation $(\chi\chi, \chi\ a, xm)$ is in the core if and only if $\chi v, \chi \alpha, \chi \mu > 0$, $xn + xa + x_M - 700$, $xn + xa > 500$, $xn + x_M > 700$, $xa + xm > 0$. These are easily solved to give $500 < ж л г < 700$, $x_M = 7QQ-xN$, $xa = 0$. The interpretation of this solution is that M gets the car with a bid of between $500 and $700 ($xn$ is the bid). Agnew does not get the car, but his presence forces the price up over $500. This answer is fairly reasonable.

It is consistent with what actually happens in bidding situations, except for one thing. Since the game is cooperative, it is possible for A and M to conspire to bid low. For example, they might agree that Mitchell bids $300, and Agnew bids zero. In exchange, if Nixon accepts the higher bid, Mitchell would pay Agnew $200 for his cooperation. The imputation corresponding to this arrangement is (300,200,200). It is not in the core, and so our analysis ignores it. Another point is worth discussing. Suppose that Agnew and Mitchell play as above, but that Nixon rejects the bid. Thus, he keeps the car (which has no value to him), and the other two neither gain nor lose. The 3-tuple of payoffs is then (0,0,0). This is not an imputation —individual rationality holds, but collective rationality does not. Referring to our argument concerning collective rationality on page 156, we see that the reason this 3-tuple is ruled out as an imputation is that all three players could do better acting differently. Since the final decision is Nixon's, the conclusion is that he is being irrational in not accepting the bid of $300. Of course, the same reasoning would apply if Mitchell bid $1 instead of $300! The problem here is that, in real life, Nixon might very well reject a bid he thinks is too low. After all, he will have other chances to sell the car, either to Agnew, or to Mitchell, or to someone else. 6.2.3. Constant-Sum Games The following definition includes, as we shall show, games whose normal forms are zero-sum. Definition 6.6. Let v be a game in characteristic function form. We say that v is constant-sum if, for every coalition «S, we have $v(S) + v(Sc) = v(V)$. Further, v is zero-sum if it is constant-sum and if, in addition, $v(P) = 0$. The game shown in Table 6.1 is constant-sum. The Used Car Game is not, since, for example, $v(\{N, A\}) + v\{\{M\}) = 500 + 0 \; \varphi \; 700 = v(\{N, A, M\})$. We must be careful here. There is a concept of "zero-sum" both for games in normal form and for games in characteristic function form. They are almost, but not quite, the same. Furthermore, there is a natural definition of constant-sum for games in normal form. Again, the two concepts of "constant-sum" are not quite the same. In Exercise (2) of this section the reader is asked to prove that the characteristic function form of a game which is zero-sum in its normal form is zero-sum in its characteristic

function form, according to the definition above. We will prove a similar statement for constant-sum games. However, the converse is false—it is possible for a game which is not constant-sum in its normal form to be constant-sum in its characteristic function form. An example

of this phenomenon appears in Exercise (4). We make a formal definition of constant-sum for games in normal form. Definition 6.7. Let π be an TV-person game in normal form. Then we say that π is constant-sum if there is a constant c such that N £)я-<(a?1,...,a?Aг) = c, í=1 for all choices of strategies χ χ,..., χ ν for players Pχ,..., Pдг, respectively. If c = 0, this reduces to zero-sum. Theorem 6.9. // an N-person game π is constant-sum in its normal form, then its characteristic function is also constant-sum. Proof. Let c be the constant value of π appearing in the previous definition. We define a new game f by subtracting c/N from every payoff in π. Thus, тï(χι,...,χν) = *i(xi,...,XN)-c/N, for every choice of i, and for all choices of strategies. Then f is zero-sum. By Exercise (2), the characteristic function μ of f is zero-sum. But, it is easy to see that the characteristic function ν of π is related to μ by the formula i/(5) = /i(5) + kc/N, where κ is the number of players in the coalition S. From this, it is clear that ν is constant-sum. П The following theorem contains bad news about cores. Theorem 6.10. If ν is both essential and constant-sum, then its core is empty. Proof. Suppose ν has N players: V = {Pu...,PN}. The idea of the proof is to assume both that ν is constant-sum and that there is an imputation χ in its core, and then to prove that ν is inessential. For any player Pj, we have, by individual rationality, *i > "(№))- (6.11)

6.2. Imputations 165 Since χ is in the core, we have Adding these inequalities, we get, using collective rationality, v{v) = f> > Ц{РЛ) + <{PiY) = "(П « = 1 by the constant-sum property. It follows that (6.11) is actually an equality. Since it holds for every jf, we have г = 1 which says that ν is inessential. П 6.2.4. A Voting Game The theory of cooperative games has been applied to several problems involving voting. For example, the distribution of power in the United Nations Security Council has been studied in this way (see [Jon80]), as has the effect that the Electoral College method of electing U.S. presidents has had on the relative strengths of voters in different states (see [Owe82]). The game presented here is rather small, but shows some of the features of the lifesize ones. Example 6.2 (Lake Wobegon Game). The municipal government of Lake Wobegon, Minnesota, is run by a City Council and a Mayor. The Council consists of six Aldermen and a Chairman. A bill can become a law in Lake Wobegon in two ways. These are: • A majority of the Council (with the Chairman voting only in case of a tie among the Aldermen) approves it and the Mayor signs it. • The Council passes it, the Mayor vetoes it, but at least six of the seven members of the Council then vote to override the

veto. (In this situation, the Chairman always votes.) The game consists of all eight people involved signifying approval or disapproval of the given bill. In its normal form, the payoffs would be in units of "power" gained by being on the winning side. It is easier to set up the characteristic function form directly than to work with the normal form. Let us call a coalition among the eight players a winning coalition if it can pass a bill into law. For example, a coalition consisting of any three Aldermen, the Chairman,

166 6. N-Person Cooperative Games and the Mayor is winning. We call a coalition which is not winning a losing coalition. Thus, the coalition consisting only of four Aldermen is a losing one (since they do not have the votes to override a mayoral veto). Define $v(S) = 1$ if the coalition S is winning, and $i/(«S) = 0$ if S is losing. Since every one-player coalition is losing, and the grand coalition is winning, an 8-tuple is an imputation if and only if #м, #c, a? i,..., a?6 > 0 and xм + %c +x\-\ V #6 = 1. Here, M means Mayor, C means Chairman, and 1,..., 6 denote the Aldermen. In Exercise (9), the reader is asked to verify that Lake Wobegon is not constant-sum. Nevertheless, its core is empty. We have the following: Theorem 6.11. Lake Wobegon has an empty core. Proof. Suppose, on the contrary, that (#м,#c,#ъ · · · >#6) is in the core. Now, any coalition consisting of at least six members of the Council is winning. Thus xc + xi + h#6 > 1, and the same inequality holds if any one of the terms in it is dropped. Since all x's are nonnegative, and the sum of all eight is 1, this implies that all the x's in the inequality above are zero. This is a contradiction. П Lake Wobegon is an example of an important category of games. Definition 6.8. A game v in characteristic function form is called simple if all the following hold: • $i/(«S)$ is either 0 or 1, for every coalition S. • $i/($ the grand coalition$) = 1.$ • $^($ any one-player coalition$) = 0.$ In a simple game, a coalition S with $v(S) = 1$ is called a winning coalition, and one with $v(S) = 0$ is called losing. Exercises (1) In the game of The Used Car, suppose Agnew and Mitchell agree that Agnew is to bid $300, and Mitchell zero. Further, they agree that if Nixon accepts the bid, then Agnew will pay Mitchell $200 for his cooperation. If this happens, what is the 3-tuple of payoffs? Is it an imputation? If not, explain what is wrong with it. (2) Prove that if 7? is a zero-sum game in normal form, then its characteristic function form is also zero-sum.

6.3. Strategic Equivalence 167 (3) Referring to Exercise (6) on page 155, for

which values of the parameters is the game constant-sum? (4) Modify the game of Couples [see Exercise (5), page 155] in the following way: If P\ and P2 choose each other, and P3 chooses P\, then the payoff to P3 is —2. All other payoffs are the same as before. Prove that the modified game is not zero-sum in its normal form, but is zero-sum in its characteristic function form. (5) Does the game of Exercise (2) on page 155 have nonempty core? (6) A four-person game is given in characteristic function form as follows: "(W) = -i, "№» = o, Ц{ЬІ) = -1, "(№» = o, *({A,ft}) = o, *({Pi,ft}) = -i, »№,Pt}) = i, v({P2,P3}) = 0, u({P2,Pi}) = \, «/({ft,P4}) = 0, •'({A.ft.Ps}) = 1, •'({ft.ft.P*}) = 2, v{{P1,P3,Pi}) = 0, «/({Pa,P3,P4}) = 1, «/({Pt.P2,P3,Pι}) = 2, ι/(0) = 0. Verify that 1^ is a characteristic function. Is the core of this game nonempty? (7) Does the game of Exercise (3), page 155, have nonempty core? (8) Let ί/bea game in characteristic function form, (i) Prove that the set of all imputations is convex, (ii) Prove that the core is convex. (9) Verify that Lake Wobegon is not a constant-sum game. 6.3. Strategic Equivalence Consider two games ν and μ in characteristic function form. Suppose that the number of players is the same for both of them. The question we are now concerned with is this: When can we say that ν and μ are "essentially" the same? Suppose, for example, that we merely change the units in which the payoffs are computed. For instance, this would be the result of converting from U.S. dollars to Swiss francs. Such a change would not change the analysis of the game in any way. This change of units is equivalent to multiplying the characteristic function by a positive constant. Another modification of a game which should have no effect on the mathematical analysis of it is this: Suppose each player P,· receives a fixed amount cг·, independently of how she plays. (Of course, cг· could be negative, in which case, its absolute value represents a fixed payment for the privilege of playing.) Since the players can do nothing to change the c^'s, they would play as if these fixed amounts were not present. Combining the two modifications just discussed leads to the following:

168 6. N-Person Cooperative Games Definition 6.9. Let ν and μ be two games in characteristic function form with the same number N of players. Then μ is strategically equivalent to ν if there exist constants $\kappa > 0$, and $c_i,...,c_{jv}$ such that, for every coalition <S, /i(5) = M5)+ X)c<. (6.12) Note first that i/ and μ really play symmetric roles in this definition. That is, (6.12) can be solved for ν to give K5) = (i/t)M5)+j;(-ci/t)l p» es which has the same form as (6.12). For an example, the game whose normal form appears in Table 6.1 has

characteristic function v(V) = 1, i>(0) = 0, v({Pu P*}) = 1, v({Pu Ps}) = 4/3, v{{P2) P3}) = 3/4, Ц{P1}) = 1/4, *({Jb}) =-1/3, *({Ps}) = 0. Letting к — 2, and C1,сг,сз be —1,0,-1, respectively, we have that µ is strategically equivalent to i/, where /i(7>) = (2)1 + (-1 + o - 1) = 0, /i(0) = (2)0 = 0, µ({P1,P2}) = (2)1 + (-1 + 0)=1, µ({Pι, Рз}) = (2)(4/3) + (-1 - 1) = 2/3, µ({P2, P3}) = 1/2, A»({Pι}) = -1/2, µ({P2}) =-2/3, µ({P3}) =-1· In this example, µ is zero-sum. We emphasize that if v and µ are strategically equivalent, then the players would play the same in both games. That is, the relative likelihoods of the various coalitions forming would be the same, and the members of the coalitions would adopt the same joint strategies. We prove the following: Theorem 6.12. If и and µ are strategically equivalent, and v is inessential, then so is µ. Thus, if и is essential, so is µ.

6.3. Strategic Equivalence 169 Proof. Assuming v is inessential, compute N N i=l «=1 N N Емт» = D*"({fi})+«) « = 1 N N «=1 г = 1 N «=1 = МП This shows that µ is inessential. By symmetry, if µ is inessential, so is v. Thus, if v is essential, so is µ. П 6.3.1. Equivalence and Imputations Suppose that v and µ are strategically equivalent, that is, that the relation given by (6.12) holds. Then the following theorem relates the imputations for v to those for µ. Theorem 6.13. Let v and µ be strategically equivalent N-person games. Then we have: • An N-tuple χ is an imputation for v if and only if kx + c is an imputation for µ. • An imputation χ dominates an imputation y through a coalition S with respect to v if and only ifkx + c dominates ky + c with respect to µ through the same coalition. • An N-tuple χ is in the core of и if and only ifkx + c is in the core of µ. Proof. Suppose that χ is an imputation for v. Then, for 1 < г < TV, tiiPi}) = kv({Pi})+Ci < kxi+a,

170 6. N-Person Cooperative Games which verifies individual rationality since kxi + Ci is the ith component of kx + c. As for collective rationality, N µ(v) = *.,(?) +Σ« г = 1 N N = *Σ*<+Σ* г = 1 г = 1 N - Y^(kxi + a). * = 1 Thus, kx + c is an imputation for µ. The converse of this statement is true because of the symmetry of v and µ. The other two statements of the theorem are proved in a way similar to the proof of the first. П This theorem validates our belief that if we are studying a game in characteristic function form, then we are simultaneously studying all games which are strategically equivalent to it. In case v and µ are strategically equivalent, then we will use phrases like "i/ and µ are the same up to strategic equivalence" to emphasize this fact. 6.3.2. (0,1)-

Reduced Form Part of the usefulness of the observation we have just made is that we can replace a game by another one whose characteristic function is particularly easy to work with. Definition 6.10. A characteristic function μ is in $(0, 1)$-reduced form if both the following hold: • $\mu(\{P\}) — 0$ for every player P. • $\mu(T) = 1$. A game in $(0, 1)$-reduced form is obviously essential. Conversely, it is also true that, up to strategic equivalence, every essential game is in (unreduced form. We have the following: Theorem 6.14. If и is an essential game, then ν is strategically equivalent to a game μ in $(0,1)$-reduced form. Proof. Define к tf > 0,

6.3. Strategic Equivalence 171 and, for $1 < \text{г} < TV$, define a $= -kv(\{Pi\})$. Then μ is defined by (6.12). The easy verification that μ is in $(0, 1)$-reduced form is left to the reader. П A simple game is already in $(0, 1)$-reduced form. Let us do some computation for the game whose normal form is given in Table 6.1. From the formulas in the proof, we have and d $= -(12/13)(1/4) = -3/13$, $c2 = 4/13$, $c3 = 0$. Then μ is given by $\mu(T) = 1$, $\mu(0) = O$, /i(any one-player coalition) $= 0$, $\mu(\{Pι, Pi\}) = (12/13)(1) - 3/13 + 4/13 = 1$, and $\mu(\{P1,P3\}) = \mu(\{P2,P3\})=1$. For this game μ, we see three things immediately: • All three two-person coalitions are equally good. • If a two-person coalition forms, the players will probably divide the payoff equally (since the players have completely symmetric roles). • There is no advantage to a two-player coalition in bringing in the third player to form the grand coalition. We conclude that one of the two-player coalitions will form, the players in it will split the payoff, and the third player will be left out in the cold. Thus, the prevailing imputation will be either $(1/2,1/2,0)$, $(1/2,0,1/2)$, or $(0,1/2,1/2)$. Our analysis is unable to predict which of the three two-player coalitions will actually form. If we transform these conclusions back into terms of i/, we see that one of the two-player coalitions will form. The prevailing imputation will be one of three possibilities which can be computed using the relationship between μ and v. This computation will be carried out shortly. The reader may verify that the $(0, 1)$-reduced form μ of the game given in (6.10) on page 161 is such that $\mu(\{P1,P2\}) = 3/8$, /i($\{Pi,ft\}$) $= \mu([P2,P3]) = 1/2$. In this example, the nice symmetry of the previous game is lost. We can safely say that the two-player coalitions seem weak, and that the grand

172 6. N-Person Cooperative Games coalition is likely to form. Stated differently, any of the two-player coalitions would profit by recruiting the third

player. To make a guess about what the final imputation might be is hazardous. From the earlier discussion of the game, we know that the core is large. The $(0, 1)$-reduced form of a game is unique. This fact follows from the following theorem and from Exercise (1). Theorem 6.15. Suppose μ and ν are N-person games in $(0, \backslash)$-reduced form. If they are strategically equivalent, then they are equal. Proof. By definition of strategic equivalence, there exist constants $\kappa > 0$, and $c\backslash,..., слг$, such that Pi es for every coalition S. To prove that μ and ν are equal, we must show that $\kappa — \backslash$ and $C\{ = 0$ for all i. Since both characteristic functions are zero for all one-player coalitions, we see that $c^* = 0$ for all i. Since both characteristic functions are 1 for the grand coalition, we see that $\kappa = 1$. Π 6.3.3. Classification of Small Games Up to strategic equivalence, the number of games with two or three players is limited, as the following three theorems show. All three are easy to prove. Theorem 6.16. A two-player game in characteristic function form is either inessential or strategically equivalent to v, where $v\{$the grand coalition$) = 1$, $i/(0) = 0$, $i/($either one-player coalition$) — 0$. In the case of constant-sum games with three players, we have the following: Theorem 6.17. Every three-player constant-sum game in characteristic function form is either inessential or is strategically equivalent to v, where $i/($the grand coalition$) = 1$, $i/(0) = 0$, $i/($any two-player coalition$) — 1$, $i/($any one-player coalition$) — 0$.

6.3. Strategic Equivalence 173 Since the phrase "three-player, constant-sum, essential game in (unreduced form" is fairly clumsy, we will abbreviate it by speaking of the game THREE. The previous theorem says that every essential constant- sum game with three players is strategically equivalent to THREE. In particular, the game whose normal form is given in Table 6.1 is strategically equivalent to THREE. Up to strategic equivalence, the three-player games which are not necessarily constant-sum form a three-parameter family. We have the following: Theorem 6.18. Every three-player game in characteristic function form is either inessential or there exist constants a, 6, c satisfying $0 < a < 1$, $0 < 6 < 1$, $0 < c < 1$, such that the game is strategically equivalent to v, where $v\{$the grand coalition$) = 1$, $i/(0) = 0$. and $i/($any one-player coalition$) — 0$, $v(\{P1,P2\}) = a$, $ii\{PuP3\} = b$, $v(\{P2,P3\}) = c$. Exercises (1) Prove that if λ is strategically equivalent to μ, and if μ is strategically equivalent to ν, then λ is strategically equivalent to v. (2) Suppose that v and μ are strategically equivalent. Prove that if one of them is constant-sum, then so is the other. (3) Prove that a constant-sum game is strategically equivalent to a zero-sum game.

(4) The Zero Game is such that its characteristic function is equal to zero for every coalition. (It's not much fun to play.) Prove that every inessential game is strategically equivalent to the Zero Game. (5) Compute the (0, l)-reduced form for the Used Car Game. (6) Compute the (0, l)-reduced form of the game in Exercise (6) on page 167. (7) Compute the (0, l)-reduced form of the game referred to in Exercise (3), page 155. (8) Compute the (0, l)-reduced form of Couples. (9) State and prove a theorem which characterizes four-player constant-sum games.

174 6. N-Person Cooperative Games 6.4. Two Solution Concepts As a solution concept for the games we are studying, the core is flawed. Often, there are no imputations in the core; when there are, there are often so many that we have no reasonable way to decide which ones are actually likely to occur. Several attempts have been made over the years to define more acceptable solution concepts. We discuss two of them here. 6.4.1. Stable Sets of Imputations The definition of a stable set of imputations is fairly natural: Definition 6.11. Let X be a set of imputations for a game in characteristic function form. Then we say that X is stable if the following two conditions hold: • (Internal Stability) No imputation in X dominates any other imputation in X through any coalition. • (External Stability) If y is any imputation outside X, then it is dominated through some coalition by some imputation inside X. This idea was introduced by von Neumann and Morgenstern in [vNM44]. They went so far as to call a stable set a solution of the game. Before discussing their reasons for doing so, let us note that an imputation inside a stable set may be dominated by some imputation outside. Of course, that outside imputation is, in turn, dominated by some other imputation inside (by external stability). This seems wrong at first because we tend to assume that the relation of dominance through a coalition is "transitive," but this is not true. Now, let us consider what might happen in a game with a stable set X. By external stability, an imputation outside X seems unlikely to become established. There exists a coalition (possibly many) which definitely prefers one of the imputations inside X. Therefore, there would be a tendency toward a shift to such an inside imputation. By internal stability, all the imputations inside X are equal as far as domination through coalitions is concerned. Presumably, the one which actually prevails would be chosen in some way not amenable to mathematical analysis—pure chance, custom, precedent, etc. But, now, there is a problem. We just mentioned that there may be an imputation outside X which

dominates a given imputation inside. Why would not a new coalition form in order to take advantage of this outside imputation? If this happens, X is abandoned (until another coalition forms and the game moves back into X). Another way of looking at this problem is that, in fact, the stable set X is not unique. There may well be others (as later examples will show). Why should the game stay inside any given one of them? Of course, it could be that the history of the game might actually turn out to be a chaotic

6.4. Two Solution Concepts 175 series of formations and dissolutions of coalitions, of moving into and out of various stable sets. After all, real life often looks that way. But, if that is the way the game goes, why call a stable set a solution? The explanation given by von Neumann and Morgenstern for calling a stable set a solution is that stable sets correspond to "sound" or "accepted" standards of behavior. That is, they represent the principles of conduct which the community (that is, the set of players) accepts as ethically right. Certain imputations (ones outside a certain stable set) are condemned as wrong, and thus can never be established, even though coalitions exist which would benefit from them. For example, in civilized societies, the poorest people are not actually left to die on the street, even though richer people have to pay more taxes in order to prevent this happening. The von Neumann-Morgenstern concept of a stable set as a solution is thus based on extra-mathematical grounds. One can accept it as reasonable or unreasonable (depending, perhaps, on the game), but there is room for honest disagreement. Let us consider, for example, the game THREE. It has a nice stable set which we have already mentioned. Let X = {(0,1/2,1/2), (1/2,0,1/2), (1/2,1/2,0)}. These are the three imputations which we saw as likely to occur for THREE. We prove the following: Theorem 6.19. The set X defined above is a stable set for THREE. Proof. We denote the characteristic function for THREE by μ. First, we verify internal stability. By symmetry, it is enough to show that the imputation (0,1/2,1/2) does not dominate (1/2, 0,1/2) through any coalition. Now, we see that the only possible coalition through which this domination could occur is {P2}· But $\mu(\{P2\}) = 0 < 1/2$. and this violates the feasibility condition in the definition of dominance. Thus, internal stability is proved. To prove external stability, let y be an imputation for THREE outside X. We must show that one of the members of X dominates it through some coalition. Now note that there are at least two values of r for which Pi < 1/2. If this were not true, we would have yi > 1/2 for two values of i. But, since each y% is

nonnegative, and the j/f's sum to one, this implies that y is one of the imputations in X. This contradicts the assumption that y is not in X. By symmetry, we may assume that y\ and 2/2 are both less than 1/2. But, then, (1/2,1/2,0) dominates y through the coalition {Pι,P2}. Π

Note that there are imputations outside X which dominate members of X. For example, (2/3,1/3,0) dominates (1/2,0,1/2) through {Pu P2}. On the other hand, (0,1/2,1/2) (a member of X) dominates (2/3,1/3, 0) through {P2,P3}. Since every essential constant-sum game with three players is strategically equivalent to THREE, we can use X to obtain a stable set for any such game. For example, let v be the game whose normal form is shown in Table 6.1. Then /i(S) = M<S)+£>' Pi es for every coalition <S, where κ — 12/13, c\ — —3/13, c2 — 4/13, and c₃ = O. (See the computation on page 171.) Thus Pi£S Replacing each imputation χ in X by (l/k)x — (l/fc)c then gives us a stable set for i/, namely, {(19/24,5/24,0), (19/24, -1/3,13/24), (1/4,5/24,13/24)}. THREE has other stable sets. We have the following: Theorem 6.20. Let c be any constant such that $0 < c < 1/2$. Then the set of imputations Zc = {(c, x2) #₃) : Z2, x₃ > 0, x2 + x₃ - 1 - c} is a stable set for THREE. Proof. To verify internal stability, let us suppose that (c,x2)x₃) and (c,#2,#₃) are both in Zc, and that the first of these dominates the other through some coalition. Then this coalition is clearly either {P2} or {P3}· But this is not possible because the feasibility condition in the definition of domination would be violated. To prove external stability, let (2/1,2/2? 2/₃) be an imputation outside Zc. If 2/1 = c, then y£Zc. Therefore, there are two cases:
• Suppose 2/1 < c. Since the y^s are nonnegative and sum to one, either y2 < 1/2 or ι/₃ < 1/2. The two cases are practically the same, so let us assume that 2/2 < 1/2. Then choose a positive number a so small that c + a< 1/2.

This is possible because $c < 1/2$. Then the 3-tuple (c,l/2 + a,l/2- c- a) £ Ze, and dominates (2/1,2/2,2/₃) through the coalition {Pi,P2}. • Suppose 2/1 > c. Then 2/2 + 2/3 < 1 — c· Let /? =1"e~2ya"w>o, and let #2 = 2/2 + /?, Ж3 = 2/3 + β- Then (c, a?2,a?₃) is in Ze, and it dominates (2/1,2/2 > 2/3) through the coalition {P2, P3}. □ The stable set Zc is said to be discriminatory toward Pi. The idea is that P2 and P3 agree to give Pi the amount c, and to negotiate between themselves about the division of the rest. By symmetry, there are similar stable sets which are discriminatory toward P2 and P3, respectively. They are {(я?1,с,я?₃) :a?1,a?

з> 0,a?i +a?з = 1-e}, and {(a?i,a?2,c) :a?i,a?2 > 0, a?i + a?2 = 1 - c}, both defined for $0 < c < 1/2$. For simple games, there are some interesting stable sets of imputations. Let ν be such a game, and call a winning coalition S minimal if every coalition properly contained in S is losing. For example, in the Lake Wobegon Game, the coalition consisting only of the six Aldermen is a minimal winning coalition. Theorem 6.21. Let ν be an N-player simple game, and lei S be a minimal winning coalition. Then the set X of all imputations χ such that χι = 0 for P{ £ 5, is a stable set. Proof. To prove internal stability, suppose that χ and ζ are both in X, and that χ dominates ζ through some coalition T. Then, clearly, T is properly contained in S. But Pier which contradicts the feasibility condition in the definition of domination. Thus, internal stability holds.

178 6. N-Person Cooperative Games To prove external stability, let y be an imputation outside X. Define <* =l - Σyi > °- Pi es Then define, for Pi £ «S, Xi -УгЈI- a/κ, where κ is the number of players in «S; finally, define, for Pг· £ S, xi = 0. Then χ is in X and dominates y through S. Π In the Lake Wobegon Game, this type of stable set represents a situation in which a winning coalition refuses to share power with those who are not members. For some years it was hoped, but not proved, that every game had a stable set. However, in 1967 (see [Luc68]), an example of a 10-player game without a stable set was discovered. There is a discussion of an application of the theory of stable sets to economics (namely, the Edgeworth trading model) in [Jon80]. 6.4.2. Shapley Values The second solution concept discussed here is that of the Shapley value. The idea was introduced in [Sha53], and is an interesting attempt to define, in a fair way, an imputation which embodies what the players' final payoffs "should" be. It attempts to take into account a player's contribution to the success of the coalitions she belongs to. If the characteristic function of the game is i/, and if S is a coalition to which player Pг· belongs, then the number 6(PitS) = v(S)-v(S-{Pi}) is a measure of the amount that Pг· has contributed to S by joining it. These numbers will be used in defining the Shapley value, but, in themselves, are not very revealing. For example, consider the game THREE. For the grand coalition, we have S(Pi,V) = 0, for each player. That is, no one contributes anything. If S is a two-player coalition, then S(PitS) = l, for each player in S. That is, the sum of the contributions is greater than i/(«S). Both of these results are amusing, but unhelpful. To begin our derivation of the Shapley value (called <£г·, for player Pг), notice that once the players have

collectively agreed on an imputation, it

6.4. Two Solution Concepts 179 might as well be assumed that it is the grand coalition which forms. This is because the condition of collective rationality ensures that the total of all payments (via the imputation) is v(V). In any case, we make this assumption, and concentrate on the process by which the grand coalition comes into being. Our assumption is that this process starts with one player; she is joined by a second, and they are later joined by a third, etc. Thus, the process is characterized by an ordered list of the players, with the kth player in the list being the kth one to join. Let us consider a four-person example. Suppose that its characteristic function is given by i/(P) = 100, i/(0) = 0, (6.13) i/({Pi}) = 0, i/({ft}) =-10, i/({P3}) = 10, v({P4}) = 0, i/({P1,P2}) = 25, ν/({Pι,P3}) = 30, ν/αPι,P4}) = 10, i/({P2,P3}) = 10, i/({P2>P4})=10, i/({P3>P4}) = 30, |/({P1,P2,Pз}) = 50, 1/({P1,P2,P4}) = 30, |/({P1,Pз,P4}) = 50, 1/({P2,Pз,P4}) = 40. Here is one ordering of the players through which the grand coalition could form: P3,P2,Pi,P4. (6.14) There are several other possible orderings. In fact, there are 4x3x2 = 4! of them. In general, the number would be AT!. We think of the choice of the actual ordering by which the grand coalition comes about as a random event. Since there are, in general, N\ possibilities, it is reasonable to assign probability ∨N\ to each of them. Now, given that the grand coalition forms according to the ordering shown above in (6.14), the number «(Pi, {P3, P2> Pi}) = i/({P3, P2, Pi}) - i/({Ps, P2}) = 50 - 10 = 40 is a measure of the contribution Pi makes as she enters the growing coalition. In general, the definition of φ{ is this: Make the same sort of calculation for each of the N\ possible orderings of the players; weight each one by the probability 1/7V! of that ordering occurring; and add the results. We will use this definition in two ways. First, we will develop a formula which makes the computation of φι somewhat easier, and, second, we will prove that φ = (φ1}.. .,φκ) is an imputation. Thus, each player will be able to get back exactly the amount contributed. To derive a formula for φι, note that, among the N\ terms in the sum which defines <£;, there are many duplications. Indeed, suppose that we have an ordering of the players such that P(occurs at position k. Denote by S the set of κ players up to and

180 6. N-Person Cooperative Games including Pi in this ordering. Then if we permute the part of the ordering coming before Pj, and permute the part coming after it, we obtain a new ordering in which Pi again is in the kth position.

Moreover, for both the original and the permuted orderings, the term in the sum defining $<j>i$ is $6(Pi,S) = v(S)-v(S-\{Pt\})$. There are $(k—1)!$ permutations of the players coming before Pr·, and $(N—k)\backslash$ permutations of the players coming after Pr·. Thus, the term $\delta(P«, S)$ occurs $(JV\text{-}*)!(*\text{-}l)!$ times. Letting $|«S|$ denote the number of players in «S, we finally get Pi es The number $\varphi\iota$ is called the Shapley value for Pj, and φ is called the Shapley vector for the game. Before getting into examples, we mention that there is an alternative way of obtaining this formula. It can be done by listing three axioms which we would want the Shapley value to satisfy. Then, a theorem can be proved to the effect that (6.15) gives the unique value which satisfies the axioms. See [Sha53] or [Vor77] for more information. After considering some examples, we prove that φ is an imputation. Consider, first, the game whose normal form is shown in Table 6.1. Its characteristic function is given on page 168. To compute $\varphi\iota$, note that there are four coalitions containing Pχ, namely, $\{Pi\}$, $\{ft, P2\}$, $\{Pi, P3\}$, $\{ft, ft, P3]$. Therefore, (6.15) has four terms in this case. For each of the four coalitions containing Pχ, we compute $6(ft, \{ft\}) = 1/4\text{-} 0 = 1/4$, $5(ft, \{ft,P2\}) = 1 - (-1/3) = 4/3$, $«(ft, \{ft. P3\}) = 4/3 - 0 = 4/3$, $«(Pi, \{ft, Pit P3\}) = 1 - 3/4 = 1/4$. Then 2O1 lilt 1111 QW $\phi1 = \psi(1/4) + \psi(\varphi) + ^p(4/3) + ^\text{-}(1/4) = 11/18$. By similar calculations, we have $\varphi2 = 1/36$, $\varphi3 = 13/36$. Notice that φ is an imputation. The Shapley values can be interpreted as reflecting the bargaining power of the players. In this example, Pi's Shapley value is the largest of the three, indicating that he is the strongest; on the

6.4. Two Solution Concepts 181 other hand, Pr'в value is very small. Player P3 is in the middle. A glance at the characteristic function supports this interpretation. The Shapley vector for this game can be computed in another way. Note, first, that in the game of THREE, the Shapley vector is surely (1/3,1/3,1/3). This is because the players have absolutely symmetric roles in the game. (It is also easy to verify by computation that this is true.) In any case, any three-player constant-sum essential game v is strategically equivalent to THREE. We can use the transformation introduced in Theorem 6.13 to obtain the Shapley vector for v. For the game given in (6.10) on page 161, the reader may verify that the Shapley vector is (1/8,5/8,1/4). These numbers seem to reasonably reflect the advantage that player P2 has in the game. For more discussion of this game, see page 171. For the Used Car Game, the Shapley values are Φv — 433.33 ..., $\Phi\alpha$ — 83.33 ..., $\phi м \sim 183.33$ .... Thus, Mitchell gets the car for $433.33, but has to pay Agnew $83.33 as a bribe for not

bidding against him. The Shapley vector indicates that Nixon is in the most powerful position. Theorem 6.22. Let ν be a game in characteristic function form. Then the Shapley vector for ν is an imputation. Proof. To prove individual rationality, we must show that $\Phi_i >$ "$(\{P_i\})$- Now, by superadditivity, if Pi £ «S, S(P{,S) = u(S) - v(S - {Pi}) > и({P{}). Thus, \Pies ') The sum in this inequality is the sum of the probabilities of the different orderings of the players. As such, it must equal 1, and so To prove collective rationality, consider ç^çç'^T-'Nμ. г' = 1 г'=1 Pi£S

182 6. N-Person Cooperative Games In this double sum, let us fix our attention on the terms involving ^(T), where T is a fixed nonempty coalition which is not equal to V. Then there are two kinds of terms involving v(T)—those with a positive coefficient (when T = S): (N-\T\)\(\T-1)1 N\ and those with a negative coefficient (when T = S — {Pi})'· (N-l-\T\)\\T\\ The first kind occurs \T\ times (one for each member of T), and the second kind occurs N — \T\ times (once for each player outside T). Thus the coefficient of v(T) in the double sum is \T\ (N-\T\)\(\T\-1)\ _ (N - \T\)(N - 1 - |Г|)!|T|! = (N-\T\)\\T\\ _ (JV-|T1)!|T|! N\ N\ Therefore, the only terms left in the double sum are those involving the grand coalition, and those involving the empty coalition. We have, since v(β) = 0, V^, N(0\)(N - 1)! //m i—l This proves collective rationality. П For another example, we compute the Shapley values for the Lake Wobegon Game. We start with фм- The nonzero terms in (6.15) are those for which S — {M} is a losing coalition, but S is winning. That is, they are coalitions which, if the Mayor is removed, can pass a bill but cannot override a mayoral veto. A little thought shows that there are four types of these winning coalitions, namely, (1) S contains the Mayor, three of the Aldermen, and the Chairman. (2) S contains the Mayor and four Aldermen. (3) S contains the Mayor, four Aldermen, and the Chairman. (4) S contains the Mayor and five Aldermen. There are sets of the first type. Since |«S| = 5, (ДГ-|5|)!(|5|-1)! _ (8-5)1(5-1)! _ W\ " 8! - 1/m

6.4. Two Solution Concepts 183 Thus, the contribution to фм from these sets is 20/280 = 1/14. There are 15 sets of the second type, and the contribution to фм from them is (15)(s-5№-i)i=3/56 There are 15 sets of the third type, and the contribution to фм from them Finally, there are 6 sets of the fourth type, and the contribution to фм from them is (6)^ЬI = 1/2в. Adding these four numbers, we get фм = 1/14 + 3/56 + 5/56 + 1/28 = 1/4. To compute фc, note that there are only two types of sets which make a contribution to фc, namely, (1) S contains

the Chairman, three Aldermen, and the Mayor. (In this case, the vote among the Aldermen is a tie, the Chairman votes to approve, and the Mayor signs it.) (2) S contains the Chairman and five Aldermen. (In this case, the bill is vetoed, but with the Chairman's vote, the veto is overridden.) There are 20 sets of the first type, and 6 of the second. Then fe=M№!t(f)MM=i;il. Now, the sum of the </>'s is 1, and all the ф{$ are surely equal (by symmetry). Thus, for each г, φí = (1/6)(1 - 1/4 - 3/28) = 3/28. These results say that the Mayor has much more power than an Alderman or the Chairman. It turns out that the Chairman's power is exactly equal to that of an Alderman. [Also, see Exercise (9).] Exercises (1) Prove that the set {(ж, 0,700-ж) :0<ж<700} is a stable set of imputations for The Used Car Game.

184 6. N-Person Cooperative Games (2) Let X be a stable set of imputations for a game v. Prove that the core is a subset of X. (3) List all the minimal winning coalitions for the Lake Wobegon Game. (4) Compute the Shapley vector for the game in Exercise (6) on page 167. (5) Compute the Shapley vector for the game referred to in Exercise (3) on page 155. (6) Verify the Shapley values given in the text for the Used Car Game. (7) Compute the Shapley value for each player in the game given in (6.13) on page 179. (8) What is the Shapley vector for an inessential game? (9) Consider a committee with 2m ordinary members and a Chairman who only votes in case of a tie. Prove that the power of the Chairman (defined as the Shapley value) is equal to that of an ordinary member. (10) The town of Bedford Falls has a bicameral Council. The Upper House has three members, and the Lower House has seven. To pass a bill, a majority is needed in each House (there is no mayoral veto). Compute the power of each member of each House (defined as the Shapley value). (11) Prove that if a game has a stable set containing only one imputation, then it is inessential.

7 Game-Playing Programs Computer programs which are capable of playing chess and other games at a high level of skill have become familiar to most of us in recent years. In fact, there are chess-playing programs which can beat all but the best human players, and many experts think that the world champion will be a program within the next few years (see [Lev84]). Our aim in this chapter is to discuss the mathematical background involved in these programs. The actual writing of such a program involves more than this mathematical background, of course. There is a part of the process which requires an

intimate knowledge and "feel" for the game. The writing of a game-playing program is an interesting exercise in both game theory and computer programming. The moment when the programmer realizes that his or her creation plays the game better than its creator is both exhilarating and strangely unsettling. The famous Baron von Frankenstein must have had similar feelings. The extensive form of a game is used in this chapter. The first task is to discuss algorithms for computing a player's optimal move. Throughout this chapter, the games studied are two-person zero-sum games of perfect information. Thus each player has an optimal pure strategy (by Theorem 1.9). Iri addition, we assume that there are no chance moves. This last assumption could be removed but we make it so as to simplify the discussion.

186 7. Game-Playing Programs 7.1. Three Algorithms Let T be the tree for a game satisfying the conditions mentioned above. Let A and B be the two players. We assume that the players move alternately until a terminal vertex is reached. The algorithms of this section compute optimal first moves for the player who owns the root. In other words, the move produced by the algorithms is an edge from the root which belongs to an optimal pure strategy for that player. 7.1.1. The Naive Algorithm The following algorithm is simple enough to show the general idea, but is not as efficient as the other two algorithms of this section. The quantities +oo and —oo which appear in the algorithm would, in a program, be replaced by numbers which are greater than and less than, respectively, any possible payoff. Algorithm 7.1 (Naive). We are given a game treeT as just described. An optimal move from the root for the player owning it is computed. (1) We define two functions тг($и$) and cr(u); тг($и$) is defined on the set of all vertices and is real-valued; cr(u) is defined on the set of all nonterminal vertices, and takes its values in the set of vertices. They are defined recursively as follows. (a) If $и$ is terminal, let where pa(u) is the payoff to player A at terminal vertex u. (b) If $и$ belongs to A, let {vi,...,Vd} be the children of $и$ and compute as follows. (i) Set i = 1, s — —oo; w — v\. (ii) Set t = $\pi(v()$. (iii) Ift>s, set s = t and w = V{. (iv) Set i = i+ 1. (v) If i < d go to (ii). (vi) Set $\pi(ĭ()$ = s and a(u) = w. (c) If $и$ belongs to B, let {^i,...,^} be the children of $и$ and compute as follows. (i) Set i = 1, s = +oo; w — v\. (ii) Sett = $\pi(ví)$. (iii) Ift < s, set s — t and w = V{. (iv) Set i = i + 1. (v) If i <d go to (ii). (vi) Set $\pi(ĭ()$ = s and a(u) — w. (2) An optimal move for the player owning the root r is given by <r(r).

The computation in case и belongs to A results in finding the child ν of и for which π(ν) is a maximum. Then a(u) is that child, and тг(и) is the value of that maximum. In case и belongs to B, the computation is the same except that maximum is replaced by minimum. It is intuitively plausible that the restriction of the function a{u) to the vertices owned by A is an optimal choice function for A, that the restriction of a{u) to the vertices owned by B is an optimal choice function for Б, and that 7r(r) is the value of the game. Of course, the corresponding statements should hold if T is replaced by any of its cuttings. The formal proof of these facts is omitted; it is similar to the proof of Theorem 1.9. The theorem here is as follows: Theorem 7.1. Let T be a game tree as described above. Apply the naive algorithm to compute the functions <r(u) and тг(и). Then, if ν is a vertex ofT, we have: • π(ν) is the value of the game whose tree is the cutting Tv. • The restriction of cr(u) to the vertices owned by player A in the cutting Tv gives an optimal strategy for A in the game Tv. • The restriction of a(u) to the vertices owned by player B in the cutting Tv gives an optimal strategy for B in the game Tv. We now apply the algorithm to an example. The tree is shown in Figure 7.1. The root is in the center and the circled numbers labeling the terminal vertices are the payoffs to player A. The other numbers at the vertices are intended only to identify them. We assume that the algorithm scans children in increasing numerical order. Let us trace the computation of cr(r). In this computation, the algorithm first calls for the computation of 7г(1). In order to compute 7г(1), π(4) is needed. The children of vertex 4 are all terminal and vertex 4 belongs to A, and so we have that π(4) = 5 and σ(4) = 10. Similarly, π(5) = 7 and (j(5) = 13. Thus, since vertex 1 belongs to Б, 7г(1) = 5 and σ(1) = 4. Then the algorithm calls for 7г(2). To compute this, 7г(б) is needed. We see that тг(б) = 4 and σ(6) = 17. Then, тг(7) = 8 and σ(7) = 21. Thus, тг(2) = 4 and σ(2) = 6. Finally, тг(3) = 2 and σ(3) = 9. We get that a(r) = 1 and 7г(г) = 5. 7.1.2. The Branch and Bound Algorithm Let us examine further the calculation just carried out for the game of Figure 7.1. In the calculation of 7г(2), the value π(6) = 4 was computed first. At this point, we (but not the naive algorithm) would know that 7г(2) < 4. Since it is already known that 7г(1) = 5, it is clear that a(r) φ 2. This means that the calculation of 7г(7) is unnecessary. Stated differently, the payoffs labeling vertices 19, 20, and 21 could be changed to any other

Figure 7.1. An example for the algorithms. values without affecting the final result of the computation. A similar thing

happens in the computation of 7г(3). Неге, 7г(8) = 3 and so it is assured that 7г(3) < 3. Thus, a(r) φ 3 and therefore the computation of 7г(9) is unnecessary. The naive algorithm has no way of taking advantage of the observations we have just made. The second algorithm is an improvement in that it can stop itself from doing this sort of unneeded calculations. The idea is that it keeps track of a variable which represents a lower bound for the value of 7г(г). When a vertex belonging to B is under consideration, the computation is stopped when it is seen that no possibility exists of beating the lower bound. When a vertex belonging to A is under consideration, the value of this lower bound can be adjusted upward when it is seen that a larger value can be guaranteed. Algorithm 7.2 (Branch and Bound). We are given a game tree T as just described. An optimal move from the root for the player owning it is computed. (1) We define two functions $\pi(\breve{u},\alpha)$ and $\sigma(\breve{u},\alpha)$; 7г(и,a) is defined on the set of all vertices и and all real numbers a, and is real-valued;

7.1. Three Algorithms 189 $\sigma(\breve{u},\alpha)$ 25 defined on the set of all nonterminal vertices и and all real numbers a, and takes its values in the set of vertices. They are defined recursively as follows. (a) If и is terminal, let 7r(?i, a) = pa(u) fºr all &- (b) If и belongs to A, let $\{v\backslash,.. ,,Vd\}$ be the children of и and compute as follows. (i) Set i — 1, 7 = a, s = —oo; w = v\. (ii) Let t — $\pi(v,\cdot,7)$. (iii) Ift>s set s = t, w — V{. (iv) If s > j, set 7 = s. (v) Set i = 2 + 1. (vi) If i < d go to (ii). (vii) Set 7r(ii, a) = s and $\sigma(\breve{u}, a)$ — w. (c) If и belongs to B, let $\{vi,...,Vd\}$ be the children of и and compute as follows. (i) Set i — 1, s — +oo; w = v\. (ii) Let t — 7г(^г-,a). (iii) Ift<s set s — t, w — V{. (iv) If s < a, go to (vii). (v) Set i = i+ 1. (vi) If г < d go to (ii). (vii) Set 7r(ii, a) — s and $\sigma(\breve{u}, a)$ = w. (2) The optimal move from the root for the player owning it is $\sigma(\Gamma, -oo)$. The key difference between this algorithm and the naive algorithm is found at the point where the variable s is compared to a. This occurs when the vertex under consideration belongs to B and results in a "cutoff" of the computation if s < a. The new algorithm has to do a little more work at each vertex (in order to update the variable 7, and to decide whether to cut off computation) but has to consider fewer vertices. 7.1.3. The Alpha-Beta Pruning Algorithm An additional saving in computation can be made. In the example of Figure 7.1, consider the computation of 7г(1). After $\pi(4) = 5$ has been computed, we see that 7г(1) < 5. In the computation of 7г(5), 7г(13) = 7 is computed first and tells us that 7г(5) > 7. Thus, 7г(1) is known to be 5. Therefore, the rest of the calculation of 7г(5) can be cut off. The branch- and-

bound algorithm is not able to take advantage of such a "deep cutoff," but the third algorithm is. The idea is to keep track of a second variable

190 7. Game-Playing Programs which plays a role symmetric to that played by the variable used in the branch-and-bound algorithm. The new variable is an upper bound on the value of the game. It may be revised downward when a vertex belonging to player B is under consideration. When a vertex owned by A is under consideration, the computation is stopped if it is seen to be impossible to return a value less than this upper bound. Algorithm 7.3 (Alpha-Beta Pruning). We are given a game tree T as before. An optimal move from the root for the player owning it is computed. (1) We define two functions $7r(ii,a,/?)$ and $\sigma(\breve{u},\alpha,/?)$; $7\Gamma(u,a,/?)$ is defined on the set of all vertices u and all real numbers a and β, and is real-valued; $\sigma(\upsilon,\alpha,/?)$ is defined on the set of all nonterminal vertices u and all real numbers α and β, and takes its values in the set of vertices. They are defined recursively as follows. (a) If u is terminal, let $7r(ii, \alpha, \beta)$ — $Pa(u)$ for all a and β. (b) If u belongs to A, let $\{vi,...,fld\}$ be the children of u and compute as follows. (i) Set i = 1, 7 = a, s = —oo; w = v\. (ii) Leit = $\tau\varphi^{\wedge}$,γ,/?). (iii) Ift>s set s — t, w = V{. (iv) If s > β go to (viii). (v) If s > 7 set 7 = s. (vi) Set í = i+ 1. (vii) If i < d go to (ii). (viii) Set $7\gamma(\breve{v}, \alpha, \beta)$ = s and $\sigma(\breve{u}, \alpha, \beta)$ = w. (c) If u belongs to B, let $\{vi,...,Vd\}$ be the children of u and compute as follows. (i) Set i = 1, $\delta = \beta$, s = +oo; w — v\. (ii) Let t = $7\Gamma(^{\wedge}r,a,<\$)$. (iii) Ift<s set s = t, w = V{. (iv) If s < a go to (viii). (v) If s < δ set δ = s. (vi) Set $\iota = r + 1$. (vii) If i <d go to (ii). (viii) Set $7r(u, \alpha, \beta)$ — s and $a(u, \alpha, \beta)$ = W. (2) TЛe optimal move from the root for the player owning it is $\sigma(\Gamma, —oo,+oo)$.

7.2. Evaluation Functions 191 Figure 7.2. Game tree for exercises. These three algorithms are discussed in detail in [KM75]. Proofs are given there that they do what we have claimed they do. Also, some estimates are made of the computational efficiency of alpha-beta pruning in comparison to the naive algorithm. There are also discussions of alpha-beta pruning and its application to specific games in [Sol84] and [Lev84]. Exercises (1) Apply the naive algorithm to the game in Figure 7.2. Compute the values of tt(u) and <7(u) for every vertex u. (2) Apply the branch-and-bound algorithm to the game in Figure 7.2. Compute $\tau\tau(r, —oo)$ and $a(r, —oo)$. Indicate where cutoffs occur. (3) Apply alpha-beta pruning to the game in Figure 7.2. Compute $7r(r, —oo, +oo)$ and $a(r, —oo, +oo)$. Indicate where deep cutoffs occur. 7.2. Evaluation

Functions For serious games, the full game tree is far too large to be analyzed. In order to get an intuitive feeling for the sizes involved, let us consider an

192 7. Game-Playing Programs imaginary game in which each player has ten possible moves at each of her vertices, and in which the game terminates after each player has made ten moves. Then there are ten children of the root, 100 grandchildren, 1000 great-grandchildren, etc. The total number of vertices is 1 + 10 -f IO2 + 103 + h 1020 = 111111111111111111111. Even if the information about one vertex could be compressed into one byte of computer memory, no computer on earth could contain this tree. In fact, if all the computers in the world were wired together into one, it would still not be large enough. Furthermore, games like chess have trees which are much larger than this one. The average number of moves available to a player is greater than ten, and games usually last for more than 20 moves. It is true that an examination of the algorithms shows that the entire game tree would not have to be in memory at one time. If memory is reclaimed after it is no longer needed, the most that is required at one time is the amount necessary to hold the longest path from the root to a terminal vertex. This observation does not solve the problem. The time required to do the computation is proportional to the total number of vertices and, even if each one takes a tiny fraction of a second, the total builds up to an impractical amount. For example, in the imaginary game just discussed, suppose that the naive algorithm is used and that the time required to compute тг(и) is only 10"9 seconds (for each u). Then the total time required is more than 1011 seconds, which is over 3000 years. If we had a computer 10 times as fast, and if the alpha-beta pruning algorithm cut off 90 percent of the vertices, then the calculation could be carried out in only 30 years. If the algorithms just discussed are to be of any use, they will have to be applied to trees smaller than full game trees. 7.2.1. Depth-Limited Subgames Let T be a game tree and let и be a nonterminal vertex of T. Suppose, to be definite, that и belongs to player A. If one of the three algorithms could be applied to the cutting TU) then the result would be an optimal move for A from u. But if the entire tree T is enormous, then most of the cuttings Tu are too large to analyze. To obtain a suitably small tree, we use the following definition. Definition 7.1. Let T be a game tree and let ubea nonterminal vertex of T. Also let m be a positive integer. The depth-limited subgame S(u,m) with root и and depth m consists of those vertices v of Tu such that the path from и to v has length at most m. The edges of S(u, m) are the edges of T which begin and end

at vertices of S(u, m).

7.2. Evaluation Functions 193 Thus, a depth-limited subgame S(u, m) consists of all vertices of the game tree which can be reached in m or fewer moves, starting at u. In other words, its vertices are the states of the game which player A can "see" by looking ahead at most m moves. The terminal vertices of S(u, m) are, in general, not terminal in T. However, by assigning numbers to these vertices, we can make S(u, m) into a game tree. If m is not too large, 5(w, m) will be of reasonable size and we can apply the algorithms to it. If, further, the assigned numbers reasonably reflect the values of their vertices (to A), then the move produced by the algorithm should be a good one (although perhaps not the best one). Let us now make a definition. A procedure for assigning values to the vertices of the game tree (that is, to the possible states of the game) is called an evaluation function. Let e(u) denote such a function. We assume that the evaluation is from A's point of view, and that the evaluation of и from B's point of view is —e(u) [this ensures that S(u,m) is zero-sum]. Second, if и happens to be terminal in T, then it is reasonable to require that e(u) = pA(u). The idea of an evaluation function is that it assigns a real number to every state of the game. This number is a measure of the worth to player A of that state of the game. For example, if и is a state of the game from which A should (by playing correctly) be able to gain a large payoff, then e(u) will be relatively big. Observe that, in theory, there is a way to define a perfect evaluation function— simply let e(u) — ъ(и) be the value of Tw, as computed by, for example, the alpha-beta pruning algorithm! As we have been saying, this evaluation function is usually not computable and so we seek one which, though imperfect, is computable and is an approximation to w(u). The actual design of such a realistic evaluation function is a tricky business. It clearly requires a good deal of knowledge of the game being studied. For an interesting discussion of the evaluation of chess positions, [Lev84] is recommended. A game-playing program is likely to spend a great part of its time computing values of the evaluation function. Therefore, it is a good idea to make that function as simple as possible. The goals of simplicity and nearness of approximation to $\pi\{\mu)$ are antagonistic, and part of the problem in writing game-playing programs is to strike a balance between them. In summary, the general procedure which a game-playing program goes through is this. It is given a depth m (either fixed, or set by the user); when it is the program's turn to move (from vertex u, say), it applies the alpha-beta pruning

algorithm to the depth-limited subgame S(u, m) (which has been made into a game tree via an evaluation function). The result

194 7. Game-Playing Programs is a move which may not be optimal, but which, one hopes, is good. It is clear that the quality of the program's play depends on the quality of the evaluation function and that it improves as m increases. The task of writing a program to play a specific games involves solving many problems which depend on the game. The way in which the current game situation is represented internally, how moves are encoded, and so forth, are all highly game-dependent. Moreover, knowledge of the game often allows the programmer to introduce extra efficiency into the program—a move which a skilled player knows will probably lead to trouble can be avoided. In general, writing a game-playing program which plays as well as a mediocre player is not so hard, but writing one which plays at a championship level is a project which involves as much art as science. 7.2.2. Mancala We will discuss in some detail a game which is complicated enough to be nontrivial to play, but which is such that an evaluation function is easily written. It is one of a very large and very old family of board games which bears the generic name mancala. These games are played in all parts of Africa and in the Middle East and Asia. Their origins are lost in history. See [Mur52] and [Bel60] for more information about them. The game we discuss here may not be exactly like any of the versions of mancala actually played. It is very close to the Indonesian game congklak discussed in [BC88]. The board for our game consists of two rows of six holes, together with two larger bowls (called stores). There are also 72 small hard objects—seeds, marbles, or beans would all be suitable, but we refer to them as beads. At the start of play, each of the 12 holes contains 6 beads, and both stores are empty. The mancala board in its initial state is shown in Figure 7.3. The holes numbered 1 through 6 belong to player A, those numbered 7 through 12 belong to 5, and each player owns a store (as indicated). Player A moves first. Her move consists of choosing one of the holes on her side of the board, removing the beads contained in it, and "sowing" them in a clockwise direction (as indicated by the arrows). To sow the beads means to add one to each hole in turn (starting with the hole following the chosen one, and including the player's own store but not her opponent's). If the last bead is sown in the player's store, then the move is over and the other player moves. If the last bead is sown in a hole which is "loaded" (that is, nonempty), then its contents are sown (and so forth, as long as the last

bead is sown in a loaded hole). If the last bead is sown in an empty hole, then the move is over. However, if this empty hole belongs to the player whose move it is, then whatever beads are in the opposite hole (which belongs to the opponent) are moved to the player's store. The players alternate in this way until a player is unable to move

7.2. Evaluation Functions 195 7 8 9 10 11 12 Figure 7.3. Initial configuration of a mancala board. (because all his holes are empty). At this point, the other player puts all the beads remaining in her holes into her store. The winner is the player with more beads at the end. For example, suppose A's first move is to sow from hole number 1. Then one bead goes into each of holes 2 through 6, and the last bead goes into A's store. A's turn is then over, and B moves. For another example, suppose A's first sowing is from hole 2. Then holes 3 through 6 receive an additional bead, A's store receives one, and hole 7 receives the last one. Since this last hole is loaded, its contents are then sown. The last one goes into the empty hole 2. The move is over, except that the 7 beads in the opposite hole 11 go into A's store. Thus, at the end of this move, A's store contains 8 beads. Of course, 5's store contains none. Figure 7.4 shows the state of the board after A has made the move just described, and B has sown from hole 12. There is an obvious evaluation function for mancala. Let $и$ denote a vertex of the game tree. Then $и$ corresponds to a certain state of the game, that is, knowing that we are at $и$ tells us how many beads there are in each hole, how many there are in each player's store, and whose turn it is to move. We define $e(u)$ to be the number of beads in A's store minus the number in 5's. For example, if $и$ is the root, then $e(u)$ — 0; if $и$ is the vertex reached after A sows from hole 2, then $e(u)$ — 8; if $и$ is the vertex after A sows from hole 2 and B sows from hole 12, then $e(u)$ — —13.

196 7. Game-Playing Programs 7 8 9 10 11 12 Figure 7.4. Mancala board after two moves. A computer program was written to play mancala. It uses the evaluation function just described, and the depth can be set by the user. We now discuss what was learned about mancala and about the three algorithms from experimenting with this program. First of all, the program plays very well. Even with the depth set at 2, human players (who were, however, not experts at the game) found it almost impossible to win. It is an interesting question whether either player has an advantage in mancala [see, also, Exercise (3)]. The experiment was tried of having the program play against

itself for various choices of the depth. The result was always that player A wins. For example, the score was 42 to 30 when the depth was set at 6. These results make it extremely likely that A has the advantage. The comparisons among the three algorithms illustrate the striking superiority of branch-and-bound over naive, and of alpha-beta pruning over branch-and-bound. As just discussed, the program played against itself for various choices of the depth. The program was written in three different ways, implementing the three algorithms. Table 7.1 shows the total number of calls to the functions 7r(u), 7г(?/,a), and 7г(?/,a,/?), respectively, for three different values of the depth.

7.2. Evaluation Functions 197 Table 7.1. Comparison of the algorithms. Depth 2 4 6 Naive 577 10944 212165 Branch and Bound 489 5323 49523 Alpha-Beta 311 3695 25245 7.2.3. Nine-Men's Morris Our second example is a member of another ancient family of games which bears the generic name morris1. Besides the version discussed here (nine-men's morris), there are also three-men's, six-men's, and twelve- men's variants. The board for the game is shown in Figure 7.5. To illustrate the great antiquity of these games, we mention that this same board has been found incised on a roof slab of the temple of Kurna built at Thebes on the Nile in ancient Egypt. Construction of this temple was begun in the reign of Rameses I who ruled Egypt from 1400 to 1366 B.C. One pictures the workers on this great project playing the game during their lunch breaks. The games were played in many other parts of the world, and nine men's was popular in Europe through the Middle Ages. It is a shame that it is now almost extinct— the rules and board are simple, but the game is not at all trivial to play. See [Mur52] and [Bel60] for more details. As seen in Figure 7.5, the board consists of three concentric squares, together with line segments connecting the centers of the squares. The 24 points of the board are labeled αï,..., α8, 61,..., 68, and cl,..., c8. Each of the two players has nine small objects. The ones belonging to a given player can be identical, but those belonging to player A must be distinguishable from those belonging to player B. We refer to these objects as pieces. The game has two phases. In the first phase, the players take turns placing pieces at points on the board. If a player achieves a mill, that is, places three pieces in a row along one of the line segments (for example, 67 — 68 — 61 or a4 — 64 — c4), then she removes any one of the other player's pieces which is not in a mill. For example, suppose that player A has pieces at points a7, a8, al, and a2; that player B has pieces at points 64 and a4; and that it is f?'s turn to play. B can place a piece at

c4 and then take A's piece at a2. The pieces at a7, a8, and al cannot be taken because they are in a mill. After both players have placed their nine pieces, the first phase of the game is over. In the second phase, a player's turn consists of moving a piece from a point to an adjacent empty point. For example, if A has a piece at point c2 and point cl is empty, she may move her piece there. If such a move causes a mill to form, then any piece belonging to the This name comes from the Low Latin merellus, a coin or token.

198 7. Game-Playing Programs 7 Ы b6 c (a8 7 Ь8 c c8 c6 c4 ï5 Ь5 Ь4 ïĭ c2 c3 a] э1 Ь2 Ь3 a?, a3 Figure 7.5. Board for nine men's morris. other player which is not in a mill can be removed. The game is over when the player whose turn it is has only two pieces left, or is unable to move because all his pieces are blocked. At this point, the other player wins. For example, suppose the game has reached the situation shown in Figure 7.6. In this diagram, the round pieces belong to player A and the square ones to player B. It is A's turn to move. She has four possible moves, c8 — cl, c6 — c5, 62 — c2, aЪ — α4. Here, the notation c8 — cl, for example, means to move the piece at c8 to the empty point cl. Two of these moves seem poor because they would allow player B to form a mill on his next move. In this situation, player A's choice is 62 — c2. At a later move, she can choose either a2 - 62 or c2 - 62, thus forming a mill and gaining a piece. A computer program was written to play nine-men's morris. The first evaluation function tried was quite simple—it counted the number of A's pieces on the board, and subtracted the number of 5's pieces. With this function, the program played at a mediocre level. In fact, it was fairly easy to beat when it played first, even with the depth set at 6. (Setting the depth to a larger value made the program unacceptably slow.) This evaluation function was replaced by one which also counted pieces, but

7.2. Evaluation Functions 199 a7 a8 al Figure 7.6. Nine men's morris board after 18 moves. Table 7.2. A nine-men's morris game. 08 o7 68 o3 61 62 c2 c3 64 10) o7 - o6 11) o3-o4 o5 ol c8 c6 67 63 o2 cl 66 c6 - c7 : a6 o2-o3 12) c3-c4:63 c7 - c6 13) 64 - 63 : c6 o5 - 06 14) 63 - 64 :a6 66 - c6 15) 64 - 63 : c6 67 - 66 16) 63 - 64 : o3 66 - c6 17) 64 - 63 :c6 c8-c7 18) 63 - 64 : cl

200 7. Game-Playing Programs which gave extra weight to pieces in a mill and to pieces which are mobile (that is, not blocked from moving). With this function, the program played reasonably well, even with the depth set to 3 or 4.

Table 7.2 is a record of a game between the program (moving first, and with the depth set at 4) and a human player. In this table, the notation 64 — 63 : c6, for example, means that the piece at 64 was moved to the adjacent point 63 (completing a mill), and that the piece at point c6 was then removed. The computer won the game. The human's move 66 in line 9 was a crucial mistake. He was trying to be tricky in setting up future mills, but the more obvious move cl would have been better. Exercises (1) For the imaginary game discussed on page 191, how many vertices are there inS(r,4)? (2) Suppose Ays first move in mancala is to sow from hole 3. Draw the board after her move is complete. (3) The game of simple mancala is played just like mancala, except that each player has only two holes, and there are initially only two beads in each hole. Draw enough of the game tree of simple mancala to show that player B has a sure win.

Appendix Solutions In this appendix, we present solutions to some of the exercises. The notation (n.m.p) stands for exercise number ρ at the end of Section m of Chapter n. (1.1.1) There are 10 vertices and, therefore, 10 cuttings. One of them is the tree itself. Two others are shown in Figure A.l. (1.1.2) There are 10 quotient trees. Six of them (the ones corresponding to the terminal vertices) are the tree itself. One of the others is shown in Figure A.2. (1.1.3) There are six terminal vertices. Each nonempty set of terminal vertices uniquely determines a subtree (by Theorem 1.4). Thus, the number of subtrees is 26 - 1 = 63. (1.1.9) Each edge (u,v) contributes +1 and —1 to the sum [+1 to p(u) and —1 to p(v)]. Thus $\Sigma\, 'H = \pounds(\iota\text{-}\iota) = o$, u€V(G) where the second sum is over all edges. (1.1.10) From Exercise 1.1.9, we know that Let V\ = {u : p(u) is odd},

202 Solutions Figure A.l. Two cuttings. a S Figure A.2. A quotient tree.

Solutions 203 and Then +1 +1 Figure A.3. A tree for very simple nim. V2 — {u : p(u) is even}. Vi V2 ViUV2 The sum over V2 is even, and, therefore, so is the sum over V\. Since each term of this sum is odd, we must have that the number of terms is even. (1.2.4) B has only two strategies, L and R. They lead to the same payoff if A moves L or R. However, if A moves M, a payoff of 2 is possible from B's moving Я, and not from moving L. Thus, B prefers R. A's best initial move is M (and either second move). (1.2.6) Figure A.3 shows a game tree for this version of nim. The designations on the edges show the state

of the piles of matches after the move. For example, the designation (|,‖) means that, on completion of that move, there will be one pile with one match, and one with two. An empty pile is denoted "-." The terminal vertices are denoted either +1 (a win for player A) or —1 (a win for B). Player B has a sure win. If A moves left from the root, B moves middle; if A moves right, B moves right. (1.3.1) Player A has three strategies, Я, LL, LR. Strategy R means that he moves right from the root (and then has no further choices). LL and LR mean that he moves

204 Solutions (1,-1) (1,0) (1,0) (2,2) Figure A.4. Choice subtrees. left from the root and moves left and right, respectively, if he reaches the other vertex where a choice exists. Player B has six strategies, LL, LR, ML, MR, RL, and RR. For example, MR means that if A moves left from the root, B moves middle, while if A moves right from the root, B moves right. (1.3.6) Player A has four choice subtrees, L, R, ML, and MR. Here, MR, for example, means that A moves middle from the root, and moves right if B moves right. All four of these are strategies for A since A has perfect information. Player B has eight choice subtrees since she has two possible moves at each of her three vertices. Only two of these are strategies because she must make the same move at each of the three vertices. We denote these by L and R. One choice subtree for each player (L for B, and ML for A) is shown in Figure A.4. (1.4.2) Player A has four strategies, LL, LR, RL, RR. Here, for example, LR means that she moves left from either of the two vertices owned by her which are children of the root, and that she moves right from either of the two vertices in her other information set. Player B has eight strategies, xyz, where x, y, and ζ can be L or R. For example, RLL means that he moves right from the child of the root belonging to him, moves left from either of the two vertices in the left-hand one of the other two information sets, and moves left from either of the two vertices in his remaining information set. (1.5.1) Trying all 2 X 4 = 8 combinations, we find that there are four equilibrium pairs (L,L),(ML,R),(MR,L),(MR,R).

Solutions 205 (1.6.2) Player A's payoff matrix is μη 1 2 ' V 1 2) while B 's payoff matrix is μη. I 0 2 J \ ° ° / Here, the rows are in the order L, R, ML, MR; the columns are in the order L, R. (1.6.3) After tedious calculation, we get the payoff matrices. For player A, it is LL LR RL RR LLL 0 -1/2 1 1/2 LLR 1/4 -1/4 3/4 1/4 LRL 1/4 -1/4 5/4 3/4 RLL 0 0 1 1 LRR 1/2 0 1 1/2 RLR 1/4 1/4 3/4 3/4 RRL 1/4 1/4 5/4 5/4 ririri 1/2 1/2 1 1 For player B, the payoff matrix

is LL LR RL RR LLL -1/2 0 -1/4 1/4 LLR 0 1/2 0 1/2 LRL -1/2 1/2 -1/4 1/4 RLL 1 0 5/4 1/4 LRR 0 1/2 0 1/2 RLR 3/2 1/2 3/2 1/2 RRL 1 0 5/4 1/4 ririri 3/2 1/2 3/2 1/2 For example, the entries corresponding to strategies (RL, RLL) are computed as follows: If the chance move is to the left, the terminal vertex reached has payoffs (1,1); if the chance move is to the middle, the terminal vertex reached has payoffs (l, 0); if the chance move is to the right, the payoffs for the terminal vertex reached are (1,2). Thus, the pair of expected payoffs is (1/4)(1,1) + (1/4)(1,0) + (1/2)(1,2) = (1,5/4). Hence, 1 appears in the (RL,RLL) entry in the first matrix, and 5/4 in the (RL,RLL) entry in the second matrix. The equilibrium pairs can be found from the two matrices. They are (RL, RLR), (RL, RRR), (RR, RLR), (RR, RRR). (1.6.8) (a) Each player has two strategies: 1 (that is, hold up one finger), and 2 (that is, hold up two fingers). (b) The payoff matrices for Mary and Ned, respectively, are (-ι1 ~l)> {~\ -a)· (c) A pair of strategies is in equilibrium if the corresponding entry in Mary's matrix is a maximum in its column, while the corresponding entry in Ned's matrix is a maximum in its row. From the two matrices, we see that there are no pairs of strategies satisfying this condition. (1.6.10) The two equilibrium 3-tuples are (1,2,1) and (3,1,1).

206 Solutions (2.1.1) The entry 0 in row 2, column 2, is the only saddle point. (2.1.3) We have: • In order for —2 to be a saddle point, —2 < a and — 2 > o. Thus, a = -2. • In order for 1 to be a saddle point, 1 < a and 1 > o. Thus, a = 1. • In order for the upper right-hand corner to be a saddle point, a < — 2 and a > 1. These cannot both be true, and so the upper right-hand corner is never a saddle point. • In order for the lower left-hand corner to be a saddle point, a < 1 and a > —2. Thus, -2 < a < 1. In summary, there is a saddle point if — 2 < o < 1. If — 2 < o < 1, the only saddle point is the lower left-hand corner. If o = —2, then both entries in the first column (both equal to —2) are saddle points. If a = 1, then both entries in the second row (both equal to 1) are saddle points. (2.1.5) We have minmij = 0,minra2j = —l,minm3j = 0,minm4j = —1. 3 3 3 3 Thus, ur(M) = maxminmij = 0. * 3 Also, max mji = 3,maxm,2 = 2,тахг^з = 3,тахг^ = 2. i i i i Thus uc(M) = minmaxmij = 2. 3 * (2.2.1) We compute £7(1, (2/5,1/3,4/15)) =(2/5)(l)+(l/3)(2) + (4/15)(3) = 28/15, £7(2, (2/5,1/3,4/15)) = (2/5)(3)+(l/3)(0) + (4/15)(2)= 26/15, £7(3,(2/5,1/3,4/15)) =(2/5)(2)+(l/3)(l) + (4/l5)(0) = 17/15. Thus, the best pure strategy for the row player against the given strategy for the column player is to play row 1. There is no mixed strategy which does better than this because, if ρ is a mixed strategy for the

row player, and ξ — (2/5,1/3,4/15), then E(P,ξ) = pi£7(l,<?) + p2Я(2,<7) + рзЯ(3,<?) < E(l,fi. (2.2.5) We compute £7(1,5) = 19/52, £7(2,5) = -12/13, £7(3, q) = 19/52,

Solutions 207 and E(4,ξ) = 19/52, E(5,ξ) = 19/52. Also, E(p, 1) = 19/52, E(ρ, 2) = 19/52, E(ρ, 3) = 29/26, and E(p, 4) = 19/52, E(p, 5) = 19/52. Thus, max# (i,£) = 19/52 and min#(p, j) = 19/52. • j In order that both ρ and f be optimal, we must have pг = 0 and 53 = 0. They both are, and so both are optimal. Finally, 19/52 is the value of the game because max E(i j ξ) = min£7(p,j) = 19/52. * 3 (2.3.1) There are no saddle points. A theorem in the text applies. We compute *i(p) = -P, MP) = 2p - 2(1 - p) - Ap - 2. Setting these equal, we get ρ = 2/5. Then the value of the game is 714(2/5) = π2(2/5) =-2/5. To get the column player's value, compute π1 (q) = -q + 2(1 - q) = -3q + 2, ir2(q) = -2(1 - q) = 2q - 2. Setting these equal gives q = 4/5. As a check, compute ^(4/5) = -2/5 = π2 (4/5). In summary, the solution is ν =-2/5, ?= (2/5,3/5), <f= (4/5,1/5). Then, if the column player is playing (1/3,2/3), we compute £7(1, (1/3,2/3)) = 1, £7(2,(1/3,2/3)) = -4/3. If the row player plays the optimal strategy (2/5,3/5) against (1/3,2/3), his payoff is (2/5)(1)+(3/5)(-4/3)=-2/5. Thus, he should play row 1 instead. (2.3.3) We compute тг,(p) = E((p,l - p)J) for j = 1,2,3,4 and get 7Гι(ρ) =-2p + 1, π2(ρ) = 2ρ-1, π3(ρ) = -4ρ + 2, π4(ρ) = ρ-1. We graph these four functions in Figure A.5, indicate their minimum with a heavier line, and circle the maximum of this minimum. It occurs where π3(ρ) crosses π±(ρ). Setting these equal, we get p* = 3/5, vr = -2/5. We also see from Figure A.5 that columns 1 and 2 are inactive. This leaves us with a 2X2 problem This is easily solved to give q* = 1/5, vc = -2/5.

208 Solutions Ψ Figure A.5. Row value. In summary, the solution is p= (3/5,2/5), <f = (0,0,1/5,4/5), ν = -2/5. (2.3.4) We compute тг*(д) = E(i,(q,l - q)) fori = 1,2,3,4. We get π1 (q) = (-1)9 + (3)(1 - q) = -Aq + 3, тг2(9) = 5<? - 1, тг3(<г)=-8<г + 5, π4(9) = 29 + 1. These are graphed in Figure A.6; their maximum is shown in a heavier line; and the minimum of the maximum is circled. This minimum occurs where ir3(q) crosses 7r4(g). Setting these equal and solving for q gives q* = 2/5, vc = 9/5. We see from the figure that rows 1 and 2 are inactive. This leaves us with a 2 X 2 problem (i:)· This is easily solved to give p* = 1/5 and vr = 9/5. In summary, the solution is ?= (0,0,1/5,4/5), ξ= (2/5,3/5), « = 9/5.

Solutions 209 Figure A.6. Column value. (2.3.8) We first deal with the case where there is a saddle point. If $1 < a < 2$, then o is a saddle point. Since $\alpha > 0$, the value of the game is positive. If $\alpha > 2$, then 2 is a saddle point. Since $2 > 0$, the value is positive. If $a < 1$, there is no saddle point, and we compute optimal mixed strategies. To compute $(p^*, 1 — p^*)$, set and solve to get $P = \alpha\rho + 1$ -p = 2p- (1 -ρ), vr = 2p* - (1 - p*) = a + 2 4 — α 4 — a The row value is zero if and only if o = — 2. It is negative when o $< — 2$, and positive when o $> — 2$. Computing the column player's strategy gives 9' = 4-a q+2 4- o : vr. (2.4.1) The matrix is skew-symmetric. We notice that row 2 is dominated by row 4. Also, column 2 is dominated by column 4. Erasing the second row and second column, we have

210 Solutions This, of course, remains skew-symmetric. The three inequalities are -2p2+P3 > 0, 2pi - P3 > 0, -Pi +P2 > 0. We arbitrarily choose the first two to make equalities. Combining these with the condition that the pj's sum to one, we get the system of equations Pi + P2 + P3 = 1, -2p2 + P3 = 0, 2pi - P3 = 0. The solution to this system is Pi = 1/4, P2 = 1/4, P3 = 1/2. These are all nonnegative, and the third inequality holds when these values are plugged in. Thus, we have solved the 3x3 matrix game. Returning to the original matrix, we know that the probability of playing the second row is zero; the other probabilities are the numbers just computed. In summary, the solution is $v = 0$, $\rho = (1/4,0,1/4,1/2) = \xi$. (2.4.4) The bottom row is dominated (by row 4). Deleting it leaves the skew- symmetric matrix 0 12-3 -10 10 -2-1 0 1 3 0-10 Since the game is symmetric, we try the method of Section 2.4. The inequalities are -r2 - 2r3 + 3r4 > 0, 7*1 — Γ3 > 0, 2t*i + Γ2 — 7*4 > 0, -3n + r3 > 0. We first try setting the first three of these to equalities, and attempt to solve the system from those together with the equation 7*1 + 7*2 + 7*3 + 7*4 = 1· This system has no solution. Setting the first two inequalities together with the last to be equalities does work. We get (0,3/4,0,1/4). The third inequality (the one not set to an equality) is valid for these values of the unknowns. They are all nonnegative, and so we have a solution. As for the original matrix game, the probability of playing the bottom row must be zero. The solution is ? = (0,3/4,0,1/4,0), ?= (0,3/4,0,1/4), $v = 0$.

Solutions 211 If we set the first and the last two inequalities to equalities, we get f= (1/6,0,1/2,1/3). However, the omitted inequality is not valid for these values of the unknowns. Finally, omitting the first inequality gives a second

valid solution p= (0,1/2,0,1/2,0), f= (0,1/2,0,1/2), ν = 0. (3.1.2) The given problem is primal. It is clearly feasible since (0,0,0,0) satisfies the constraints. The dual problem is minimize 3/2 subject to 2/1 > 1 2/2 >-l -2/1 > -1 -2/2 > 1. The last constraint makes this dual problem infeasible (since 2/2 cannot be negative). The infeasibility of the dual does not tell us that the primal is unbounded. The implication in Corollary 3.2 only goes in the opposite direction. Returning to the primal, however, we see that if we let x± —▶ +00, while keeping the other variables equal to zero, we get that the objective function goes to infinity. Thus, the primal is unbounded. (3.2.1) Introducing two slack variables, #3 and #4, we get the basic form maximize 2x \ -J- 3#2 subject to x\ -f X2 — 10 = —xs xl ~~ x2 ~~ 2 = —X±. Since there are two unknowns in the primal form and two constraints, the total possible number of basic forms is (T>-· The basis for the basic form already given is {#3, x±}. To find the basic form whose basis is {^1,0:4}, solve the first constraint above for x\ to get xx = —x2 — xs + 10. (A.l) Substituting this into the equation for —x± gives us —2#2 — a?з + 8 = —074. Then, substituting (A.l) into the objective function gives X2 - 2*3 + 20. Finally, applying a little algebra to (A.l) gives *2 + *з - 10 = —x\. The new basic form, obtained from the first by interchanging x\ and #3, is then

212 Solutions maximize X2 — 2x^ + 20 subject to X2 + xз — 10 = —x\ —2x2 — xз + 8 = — #4. If we now interchange #4 and X2 in this second basic form, we get the basic form maximize \j2x± — Ъ/2xз + 24 subject to 1/2^4 + 1/2a;з — 6 = — x\ -1/2*4 + 1/2*3 - 4 = -a?2. Interchanging x\ and #3 in this gives maximize 3*4 + §x\ — 6 subject to #4 + 2x\ — 12 = —x$ —X± — X\ + 2 = —Д72 · Interchanging #2 and x\ now gives maximize — 2x± + 5a?2 + 4 subject to — #4 + 2a?2 — 8 = —0:3 #4 ~~ X2 ~~ 2 = — X\. An examination of the five basic forms computed so far shows that the only one missing is the one whose basis is {#2» #4}· This can be obtained from any one of the forms already computed, except the last. It is maximize — 3яз — x\ +30 subject to xз + 2x\ — 12 = —x± xз + x\ — 10 = —X2- (3.3.2) The initial tableau is (A.2) This tableau is feasible but nonoptimal. We choose to pivot in the first column. The top entry is the only positive one and so we pivot so as to interchange x\ and x±. The second tableau is X\ X2 X3 111 -110 12 3 -1 3 0 0 = — #4 = -xъ f X± X2 X3 111 12 1 -112 -1 CO CO -3 = -Xχ = -xъ f (A.3)

Solutions 213 This is still not optimal. We choose the second column to pivot

in. The minimum ratio is in the second row and so we pivot so as to interchange #2 and #5. The new tableau is (A.4) This is still not optimal. There is only one possible pivot column, the third. There is a tie for choice of pivot row. We arbitrarily choose the first. Thus the next pivot will interchange хз and x\. The result is x± 1/2 1/2 -3/2 X$ Xz -1/2 1/2 1/2 1/2 -1/2 3/2 -1 3/2 3/2 -9/2 = -Xχ = -X2 f X± Xs X\ 1-12 0 1-1 -3 1 -3 -1 3 0 -9 = -хз - -X2 f (A.5) Still not optimal! There is only one possible pivot—the one which interchanges x$ and X2. We get #4 X2 X\ 111 0 1-1 -3 -1 -2 -1 3 0 -9 = -хз - - xs f (A.6) This is, at last, optimal. The solution is thus $\chi\lambda = \chi 2 = 0$, xs = 3;max/ = 9. It took four pivots to do it. We note that a luckier choice of pivot at the beginning would have allowed us to get the answer in one pivot. (3.4.2) The initial basic form for this problem is maximize x\ + X2 — хз — x± (A.7) subject to x\ + X2 + хз + x± — 8 = —хь X\ ~~ X2 ~~ хз — x± ~b 1 = —X6 X3 — 4 = —xj. This basic form is infeasible because of the positive constant term in the constraint equation for — xq . We thus apply the feasibility algorithm. The initial basic form for the auxiliary problem is maximize — и subject to x\ + X2 + хз + x± — и — 8 = — хь $\chi\iota$ — X2 — хз — x± — и + 1 = —χ& xz — и — 4 = — χγ. (A.8)

214 Solutions The tableau for this basic form is Xl 1 1 0 0 X2 1 -1 0 0 X3 1 -1 1 0 x± 1 -1 0 0 и -1 -1 -1 -1 -1 8 -1 4 0 = = = = -хь -X6 -x7 h $\chi\iota$ 0 -1 -1 -1 X2 2 1 1 1 X3 2 1 2 1 #4 2 1 1 1 a* -1 -1 -1 -1 -1 9 1 5 1 = = = = -хь —и -x7 h Xl 2 -1 0 0 и -2 1 -1 -1 хз 0 1 1 0 з?4 0 1 0 0 X6 1 -1 0 0 -1 7 1 4 0 = = = = -хь - X2 -X7 h (A.9) We have used h to refer to the auxiliary objective function. According to the feasibility algorithm, we should pivot on the entry in row 2, column 5 (thus interchanging и and X6) · The new tableau is (A.10) As promised in the algorithm, this tableau is feasible. It is not optimal, however. We choose the second column as the pivot column. There is only one choice of pivot row, and so we pivot so as to interchange X2 and u. The resulting tableau is (A.U) This is optimal and the maximum value of h is zero. Thus the original problem is feasible. Following the algorithm, we now set up a (feasible) tableau for the original problem. It is (A.12) The bottom row of this tableau is obtained as follows. The objective function in the original problem is / == X\ + x2 — x3 — #4 · The nonbasic variables are now $\chi\iota$, X3, x±, X6, and so we must write / in terms of these. Thus, we want to get rid of X2 and replace it with X6> The second row of the tableau gives us — X2 in terms of the nonbasics. Substituting for X2 from this equation gives us the desired formula

for /. The tableau above is feasible but not optimal. One pivot (interchanging X6 and хъ) gives us a feasible tableau (A.13) Xl 2 -1 0 2 X3 0 1 1 -2 #4 0 1 0 -2 X6 1 -1 0 1 -1 7 1 4 -1 = = = = -хъ -X2 -X7 f Xl 2 1 0 0 X3 0 1 1 -2 274 0 1 0 -2 хъ 1 1 0 -1 -1 7 8 4 -8 = = = = -a* -X2 -X7 f

Solutions 215 From this, we read off the solution a?! = 07з = 274 = 0, о?2 = 8, max/ = 8. (3.5.3) The primal problem for this dual is maximize — 072 subject to — 07i < 1 X\ — X2 < —1 072 < 1. The dual/primal tableau is 2/1 2/2 2/3 -1 χι -1 1 0 0 = 2/4 X2 0 -1 1 -1 2/5 -1 1 -1 1 0 = 9 = = = = -073 —074 -075 / (A.14) This tableau is feasible as a dual tableau, but infeasible as a primal one. Therefore, it seems easier to use the dual simplex algorithm. It calls for us to pivot so as to interchange 2/2 and уъ (and, thus, 072 and 074). The result of this pivot is (A.15) This is feasible and optimal, both as a dual and a primal tableau. We read off the solutions 2/1=2/3= 0, 2/2 = 1; min^f = -1. 07i = 0,072 = 1, 073 = 1; max/ = —1. (4.1.1) First, the smallest entry in the matrix is —3. We therefore choose c = 4 in the preliminary step. Adding this quantity to each entry gives the modified matrix 2/1 2/5 2/3 -1 071 -1 -1 1 -1 = 2/4 074 0 -1 1 -1 = 2/2 -1 1 1 0 1 = 9 = = = = -XZ -X2 -хъ f The initial primal/dual tableau can then be set up. It is 2/1 2/2 2/3 2/4 -1 x\ 1 5 3 5 1 = 2/5 X2 6 2 4 5 1 = 2/6 xz 4 3 6 1 1 = 2/7 -1 0 = 9 = = = = = —074 -075 -076 -077 /

216 Solutions We choose to pivot in the first column. There is then a choice of two possible pivot rows (2 and 4). We choose row 2. That is, we will interchange x\ and x$ (and 2/2 and 2/5). After the pivot, the new tableau is 2/1 2/5 2/3 2/4 -1 Я5 -1/5 1/5 -3/5 -1 -1/5 = 2/2 X2 28/5 2/5 14/5 3 3/5 = 2/6 xз 17/5 3/5 21/5 -2 2/5 = 2/7 -1 4/5 1/5 2/5 0 -1/5 = 0 = —374 = -07i = -076 = -077 / We choose to pivot in the second column. Then, the only possible pivot row is 4. Thus, we interchange X2 and 077 (and 1/4 and ye). The new tableau is 2/1 2/5 2/3 2/6 -1 Я5 5/3 1/3 1/3 -1/3 0 = 2/2 077 -28/15 -2/15 -14/15 1/3 -1/5 = 2/4 S3 107/15 13/15 91/15 -2/3 4/5 = 2/7 -1 4/5 1/5 2/5 0 -1/5 = 9 = -074 = -a?i = -076 = -072 / This is still not optimal. There is only one possible pivot. It is the one which interchanges 073 and x$ (and 2/3 and 2/7). The new tableau is 2/1 2/5 2/7 2/6 -1 *5 07 07 07 07 -4/91 = 2/2 077 07 07 07 07 -1/13 = 2/4 076 07 07 07 07 -12/91 = 2/3 -1 30/91 1/7 6/91 4/91 -23/91 = 9 = = = = = —074 -071 -073 -072 / This tableau is optimal and so we have not bothered to compute the irrelevant entries. They are replaced by x. The solution to the primal and dual problems are thus 07i=l/7, 072=4/91, 073 = 6/91 2/i =0, 2/2 =

4/91, 2/3 = 12/91, 2/4 = 1/13 max/ = min^f = 23/91. According to the change-of-variable formulas, the probabilities for the column player are given by q_i = (l/7)/(23/91) = 13/23, q_2 = (4/91)/(23/91) = 4/23, q_3 = (6/91)/(23/91) = 6/23. Thus, the column player's optimal strategy is (13/23,4/23,6/23). Similarly, the row player's optimal strategy is computed to be (0,4/23,12/23,7/23). The value of the modified matrix is 91/23. Therefore, the value of the original game is v = 91/23 - c = 91/23 - 4 = -1/23. It is slightly favorable to the column player.

Solutions 217 (4.2.3) Each of the two players (call them Pi and P2) chooses one of the four vertices A, B, C, D. If they both choose the same vertex, the payoff is zero. If there is an edge from Pi's choice to P2 's choice, then Pi wins 1 and P2 loses 1. If there is an edge going the other way, then the payoff also goes the other way. If there is no edge between the choices, then the payoff is zero. The matrix for the game is 0 0 0 1 0 0 1-1 0-101 -11-10 Here, the order of the rows and columns corresponds to the alphabetical ordering of the vertices. This is, of course, a symmetric game and can be solved by the method of Chapter 2. However, an inspection of the matrix shows that the entry in row 1, column 1, is a saddle point. Thus, both players should always choose vertex A. (4.2.7) For each of the four columns of the matrix for Colonel Blotto's game, compute the expected payoff when that column is played against Colonel Blotto's mixed strategy. We find that these payoffs are 2 for both the first two columns, and 3/2 for the other two. It follows that Attila should put two regiments at one position and the other at the remaining position. The payoff, 3/2, is slightly better for Attila than the value of the game (which was computed to be 14/9). Thus Blotto is punished for not respecting game theory. It should also be noticed that any weighted average of the two pure strategies (2,1) and (1,2) would also give the same payoff. (4.2.11) There are nine combinations of hands: • Each is dealt HH(probability = 1/16)—Rose bets, Sue sees, payoff is zero. • Rose is dealt HH, Sue is dealt LL(probability = 1/16)—Rose bets, Sue folds, payoff is 1. • Rose is dealt HH, Sue is dealt HL(probability = 1/8)—Rose bets, Sue folds, payoff is 1. • Rose is dealt LL, Sue is dealt HH(probability = 1/16)—Rose checks, Sue bets, Rose folds, payoff is —1. • Both are dealt LL(probability = 1/16)—Rose checks, Sue bets, Rose folds, payoff is —1. • Rose is dealt LL, Sue is dealt HL(probability = 1/8)—Rose checks, Sue bets, Rose folds, payoff is —1. • Rose is dealt HL, Sue is dealt HH(probability = 1/8)—Rose bets, Sue sees, payoff is —3. • Rose is dealt HL, Sue is dealt LL(probability = 1/8)—Rose bets, Sue folds,

payoff is 1. • Both are dealt HL(probability = 1/4)—Rose bets, Sue folds, payoff is 1. Summing the products of these payoffs times the associated probabilities gives, for the expected payoff, the amount —1/16. (5.1.5) The row player's payoff matrix is The row player's maximin value is the value of this matrix game. It is easily computed to be 1/2 (after noting that the third row is dominated by the second).

218 Solutions The column player's payoff matrix is (-Π) To compute the column player's maximin value, we need the transpose of this matrix. It is The column player's maximin value is the value of this matrix. The third column is dominated by the second. The value is easily computed to be 3/5. (5.1.7) We compute $ni(x,y) = 2a?i/- a?(1 - y) + (1 - a?)(1 - y) = (4y-2)x-y+1$. Prom this, we see that the maximum over χ of $\pi\iota$ (χ, y) is attained for $\chi = 0$ when $y < 1/2$, for $\chi = 1$ when $y > 1/2$, and for all χ when $y = 1/2$. The set $A = \{(\#,2/) : \pi\iota$ (x, y) is a maximum over χ with y fixed$\}$ is shown in Figure A.7 with a solid line. Then we compute $ir2(x,y) = Sxy + 3a:(1 - y) + y(1 - x) - 2(1- x)(1 - y) = (\sim 9x + 3)y + 5x-2$. From this, we see that the maximum over y of $^\{^yУ)$ is attained for $y = 1$ when $\chi < 1/3$, for $y = 0$ when $\chi > 1/3$, and for all y when $\chi — 1/3$. The set $B = \{(я»2/) : я"2(\#»2/)$ is a maximum over y with a; fixed$\}$ is shown in Figure A.7. The intersection of the sets A and B is the set of equilibrium pairs. We see from Figure A.7 that there is only one, (1/3,1/2). We note that the payoffs to the two players corresponding to this equiUbrium pair are 7Γ! $(1/3,1/2) = 1/2, 7Γ2 (1/3,1/2) =-1/3$. The maximin values for the row player and column player are the values of the two matrices and (1-0· The values are then easily computed to be 1/2 (for the row player) and —1/3 (for the column player).

Solutions 219 X 1^ Figure A.7. Equilibrium pair. (5.3.1) The problem concerns the 7i07icooperative payoff region. An accurate drawing of this region would allow us to solve the problem, but such drawings are difficult to make. However, the following fact is easily verified from the definition of Pareto optimality: // a payoff pair (u,v) is in the noncooperative payoff region and is Pareto optimal in the cooperative region, then it is Pareto optimal in the noncooperative region. The cooperative payoff region is shown in Figure A.8. From it, we immediately see that the three payoff pairs are Pareto optimal. (5.3.2) From Exercise 5.1.7, the pair of maximin values is (1/2,-1/3). The arbitration pair is to be found in the set of payoff pairs which are Pareto

optimal and which dominate (1/2,-1/3). A glance at Figure A.8 shows that this set is the line segment from (1/2,0) to (2/3,-1/3). It is indicated with a heavier line in Figure A.8 and has an arrow pointing at it. In the terminology of the proof of Theorem 5.4, Case (i) holds. Thus, the arbitration pair occurs at the maximum of the function $0(u,t/) = (u-1/2)(t/+1/3)$, where the payoff pair (u, v) lies in the line segment just mentioned. On this line segment, and Thus $v = -2u+ 1$, $1/2 < и < 2/3$. $g(u,v) = (u-1/2)(-2u+ 1 + 1/3)= -2u2 + 7u/3 - 2/3$.

220 Solutions Figure A.8. Payoff region (cooperative). Setting the derivative of this function of $и$ equal to 0 gives $и = 7/12$ and so $v = -1/6$. This pair is on the line segment and $g(7/12,-1/6) = 1/72$. The values of $g(u, v)$ at both endpointsare zero. We conclude that $(7/12,-1/6)$ is the arbitration pair. (6.1.2) There are three players. Let us designate them as Pi, P2, P3. There are eight coalitions. Besides the grand coalition and the empty coalition, they are and w.w.w. {Pl,P2},{Pl,P*},{P2,P3}. The bi-matrix for the two-player game between {Pi} and its counter-coalition (that is, {P2,P3}) is / (0,0) (0,0) (1,-1) (0,0) ^ (1,-1) (2,-2) (1,-1) (-1,1) In this bi-matrix, the rows correspond to the two strategies 1 and 2 for {Pi}, and the four columns correspond to the joint strategies (1,1), (1,2), (2,1), and (2,2). Thus, the entry in the second row and second column is (2,-2) because the payoff triple

Solutions 221 corresponding to the strategy combination (2,1,2) is (2,0,-2). From this bi-matrix, the maximin values are quickly computed to give "(№}) = o, KN⍛,ft}) = o. For the coalition {P2}, the bi-matrix is ((0,0) (-1,1) (1,-1) (0,0) \^ (0,0) (0,0) (-1,1) (1,-1))■ The rows and columns again correspond to joint strategies in the same order as before. The entry in the first row, third column, is (1,-1) because the payoff triple corresponding to the strategy combination (2,1,1) is (1,1,-2). From the maximin values, we have "(№}) = - IA *({Pι,P3}) = 1/3. Finally, we get *({P₃}) = -2, и({PиP2}) = 2. Also, $v(\{Pi,P2,P3\}) = o$, *({0}) = o. We notice that the equation K5C) = -*(*) holds for every coalition S. The reason is that the game is zero-sum. The characteristic functions of zero-sum games are discussed in a later section. (6.2.5) The game is zero-sum and essential; thus, by Theorem 6.10, its core is empty. For the fun of it, however, let us verify this by a calculation. Using the values of the characteristic function computed in the previous problem, we see that $(#1, X2,x₃)$ is an imputation if $x\ > 0$, $X2 > —1/3$, $#3 > —2$, and $X\ + X2 + #3 = 0$. Also, an imputation $(x\, #2», #3)$ is in the core if and only if $Xl + X2$

> 2, X\ + X3 > 1/3, X2 + X3 > 0, since the other inequalities from Theorem 6.7 are automatically true for an imputation. Adding these last three inequalities, we get 2(x1+x2 + x₃)> 7/3. This contradicts Xl + X2 + ^3 = 0, and so the core is empty. (6.2.6) To verify that v is a characteristic function, we have to check that superad- ditivity, 1/(5 U T) >!/(£)+ i/(T), holds whenever S and T are disjoint coalitions. This is easily checked. For example, H{Pi<P2,Pi}) = 2> -1 +1 = K№}) + H{P2,Pi})- By Corollary 6.8, a 4-tuple (χι,X2ιχ3ιχï) is in the core if and only if both the following hold: XI + X2 + *3 + *4 = I/({Pï,P2,P3,P4» = 2,

222 Solutions and Pi€S for every coalition S. It is easily checked that, for example, (1,0,0,1) and (0,1,0,1) satisfy these conditions. Thus the core is not empty. (6.3.6) Using the formulas in the proof of Theorem 6.14, we compute κ = ^ = ï = 1/4. "W-EL^.}) 2-(-2) Then ci = -M{Pi}) = -(l/4)(-1)=l/4, c2 = 0, c3 = 1/4, c4 = 0. Then μ(5) = Ari/(5) + Y^ Ci Pi€S defines the (0, l)-reduced form of v. For example, μ({Λ,A»,P4}) = (l/4H{Pi,ft,P4}) + (1/4+ 1/4 + 0) = 1/2. (6.4.4) The Shapley value φï is computed from (6.15) on page 180. For player Pj, the sum is over the coalitions S for which i(Pi,$)*0. These are We have 01 = -1/4 + 1/12 + 1/12 + 1/12 + 1/4 = 1/4. The other three are computed in the same way to give the Shapley vector (1/4,13/12,-1/4,11/12). It is an imputation (as it is supposed to be). Players P^ and P4 are in powerful positions in this game.

Bibliography [Apo74] Tom W. Apostol. Mathematical Analysis. Addison-Wesley, second edition, 1974. [Axe84] Robert M. Axelrod. The Evolution of Cooperation. Basic Books, 1984. [BC88] Robbie Bell and Michael Cornelius. Board Games Round the World. Cambridge University Press, 1988. [Bel60] R.C. Bell. Board and Table Games. Oxford University Press, 1960. [Bla77] R.G. Bland. New finite pivoting rules for the simplex method. Mathematics of Operations Research, 2:103-107,1977. [BM76] J.A. Bondy and U.S.R. Murty. Graph Theory with Applications. North- Holland, 1976. [Chv83] Vasek Chvatal. Linear Programming. W.H. Freeman and Company, 1983. [CV89] James Calvert and William Voxman. Linear Programming. Harcourt Brace Jovanovich, 1989. [Dan5l] G.B. Dantzig. A proof of the equivalence of the programming problem and the game problem. In T.C. Koopmans, editor, Activity Analysis of Production and Allocation, number 13 in Cowles Commission Monograph, pages 330-338. John Wiley & Sons, 1951. [Dre8l]

Melvin Dresher. The Mathematics of Games of Strategy, Theory and Applications. Dover Publications, 1981. [Eps67] Richard A. Epstein. The Theory of Gambling and Statistical Logic. Academic Press, 1967. [Gas86] Saul I. Gass. Linear Programming, Methods and Applications. McGraw-Hill, 1986. [Jon80] A.J. Jones. Game Theory: Mathematical Models of Conflict. Ellis Horwood Limited, Chichester, 1980. [KM75] Donald E. Knuth and Ronald W. Moore. An analysis of alpha-beta pruning. Artificial Intelligence, 6:293-326,1975.

224 Bibliography [Lev84] David Levy. The Joy of Computer Chess. Prentice-Hall, Inc., 1984. [Lev88] David Levy, editor. Computer Games I. Springer-Verlag, 1988. [LR57] R. Duncan Luce and Howard Raiffa. Games and Decisions: Introduction and Critical Survey. John Wiley &; Sons, 1957. [Luc68] W.F. Lucas. A game with no solution. Bulletin of the American Mathematical Society, 74:237-239, 1968. [Lue84] David G. Luenberger. Linear and Nonlinear Programming. Addison-Wesley, 1984. [McK52] J.C.C. McKinsey. Introduction to the Theory of Games. McGraw-Hill, 1952. [MMS83] Albert H. Morehead and Geoffrey Mott-Smith. Hoyle's Rules of Games. Signet, 1983. [Mor68] Oskar Morgenstern. The theory of games. In Mathematics in the Modern World, chapter 39, pages 300-303. W. H. Freeman and Company, 1968. First appeared in Scientific American, May, 1949. [Mur52] H.J.R. Murray. A History of Board-Games Other Than Chess. Oxford University Press, 1952. Reprinted by Hacker Art Books, New York, 1978. [Nas50] J.F. Nash. The bargaining problem. Econometrica, 18:155-162, 1950. [Nas51] J.F. Nash. Non-cooperative games. Annals of Mathematics, 54(2):286-295, 1951. [NT93] Evar D. Nering and Albert W. Tucker. Linear Programs and Related Problems. Academic Press, 1993. [Owe82] Guillermo Owen. Game Theory. Academic Press, 1982. [PM91] M.H. Protter and C.B. Morrey. A First Course in Real Analysis. Springer- Verlag, second edition, 1991. [Pou92] William Poundstone. Prisoner's Dilemma. Doubleday, 1992. [Rap66] Anatol Rapoport. Two-Person Game Theory: The Essential Ideas. The University of Michigan Press, 1966. [Rap68] Anatol Rapoport. The use and misuse of game theory. In Mathematics in the Modern World, chapter 40, pages 304-312. W.H. Freeman and Company, 1968. First appeared in Scientific American, December, 1962. [Rap70] Anatol Rapoport. N-Person Game Theory. The University of Michigan Press, 1970. [Rus84] Laurence Russ. Mancala Games. The Folk Games Series, no. 1. Reference Publications, Algonac, Michigan,

1984. [Sag93] Carl Sagan. A new way to think about rules to live by. Parade Magazine, pages 12-14, November, 1993. [Sha53] L.S. Shapley. A value for τι-person games. In Contributions to the Theory of Games, II, Ann. Math. Studies No. 28, pages 307-317. Princeton University Press, 1953. [Sol84] Eric Solomon. Games Programming. Cambridge University Press, 1984. [Str89] James K. Strayer. Linear Programming and Its Applications. Springer-Verlag, 1989. [Tho66] Edward O. Thorp. Beat the Dealer: A Winning Strategy For the Game of Twenty One. Vintage Books, 1966. [Tho84] L.C. Thomas. Games, Theory and Applications. Ellis Horwood Limited, Chichester, 1984. [Val64] Frederick A. Valentine. Convex Sets. McGraw-Hill, 1964. [vNM44] John von Neumann and Oskar Morgenstern. Theory of Games and Economic Behavior. Princeton University Press, 1944. [Vor77] N.N. Vorob'ev. Game Theory, Lectures For Economists and Systems Scientists. Springer-Verlag, 1977. Translated and supplemented by S. Kotz.

Bibliography 225 [Wan88] Wang Jianhua. The Theory of Games. Tsinghua University Press, Beijing and Clarendon Press, Oxford, 1988. [Wil66] J.D. Williams. The Compleat Strategyst. McGraw-Hill, 1966.

Geometry Reisel: Elementary Theory of Metric Spaces Rickart: Natural Function Algebras Rotman: Galois Theory Rubel/Colliander: Entire and Meromorphic Functions Sagan: Space-Filling Curves Samelson: Notes on Lie Algebras Schiff: Normal Families Shapiro: Composition Operators and Classical Function Theory Simonnet: Measures and Probability Smith: Power Series From a Computational Point of View Smoryski: Self-Reference and Modal Logic Stillwell: Geometry of Surfaces Stroock: An Introduction to the Theory of Large Deviations Sunder: An Invitation to von Neumann Algebras Tondeur: Foliations on Riemannian Manifolds Wong: Weyl Transforms Zhang: Matrix Theory: Basic Results and Techniques Zong: Sphere Packings Zong: Strange Phenomena in Convex and Discrete Geometry