CHAPTER 11

Game Theory

In this chapter, we shall study if not the most practical then certainly an elegant application of linear programming. The subject is called game theory, and we shall focus on the simplest type of game, called the *finite two-person zero-sum game*, or just *matrix game* for short. Our primary goal shall be to prove the famous Minimax Theorem, which was first discovered and proved by John von Neumann in 1928. His original proof of this theorem was rather involved and depended on another beautiful theorem from mathematics, the Brouwer Fixed-Point Theorem. However, it eventually became clear that the solution of matrix games could be found by solving a certain linear programming problem and that the Minimax Theorem is just a fairly straightforward consequence of the Duality Theorem.

1. Matrix Games

A *matrix game* is a two-person game defined as follows. Each person first selects, independently of the other, an action from a finite set of choices (the two players in general will be confronted with different sets of actions from which to choose). Then both reveal to each other their choice. If we let i denote the first player's choice and j denote the second player's choice, then the rules of the game stipulate that the first player will pay the second player a_{ij} dollars. The array of possible payments

$$A = [a_{ij}]$$

is presumed known to both players before the game begins. Of course, if a_{ij} is negative for some pair (i, j), then the payment goes in the reverse direction—from the second player to the first. For obvious reasons, we shall refer to the first player as the *row player* and the second player as the *column player*. Since we have assumed that the row player has only a finite number of actions from which to choose, we can enumerate these actions and assume without loss of generality that i is simply an integer selected from 1 to m. Similarly, we can assume that j is simply an index ranging from 1 to n (in its real-world interpretation, row action 3 will generally have nothing to do with column action 3—the number 3 simply indicates that it is the third action in the enumerated list of choices).

Let us look at a specific familiar example. Namely, consider the game every child knows, called Paper–Scissors–Rock. To refresh the memory of older readers, this is a two-person game in which at the count of three each player declares either Paper, Scissors, or Rock. If both players declare the same object, then the round is a

R.J. Vanderbei, *Linear Programming*, International Series in Operations Research & Management Science 196, DOI 10.1007/978-1-4614-7630-6_11,

draw. But Paper loses to Scissors (since scissors can cut a piece of paper), Scissors loses to Rock (since a rock can dull scissors), and finally Rock loses to Paper (since a piece of paper can cover up a rock—it's a weak argument but that's the way the game is defined). Clearly, for this game, if we enumerate the actions of declaring Paper, Scissors, or Rock as 1, 2, 3, respectively, then the payoff matrix is

$$\begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix}.$$

With this matrix, neither player has an obvious (i.e., deterministic) winning strategy. If the column player were always to declare Paper (hoping that the row player will declare Rock), then the row player could counter by always declaring Scissors and guaranteeing herself a winning of one dollar in every round. In fact, if the column player were to stick to any specific declaration, then the row player would eventually get wise to it and respond appropriately to guarantee that she wins. Of course, the same logic applies to the row player. Hence, neither player should employ the same declaration over and over. Instead, they should randomize their declarations. In fact, due to the symmetry of this particular game, both players should make each of the three possible declarations with equal likelihood.

But what about less trivial games? For example, suppose that the payoffs in the Paper–Scissors–Rock game are altered so that the payoff matrix becomes

$$A = \begin{bmatrix} 0 & 1 & -2 \\ -3 & 0 & 4 \\ 5 & -6 & 0 \end{bmatrix}.$$

This new game still has the property that every deterministic strategy can be foiled by an intelligent opponent. Hence, randomized behavior remains appropriate. But the best probabilities are no longer uniformly $1/3$. Also, who has the edge in this game? Since the total of the payoffs that go from the row player to the column player is 10 whereas the total of the payoffs that go to the row player is 11, we suspect that the row player might have the edge. But this is just a guess. Is it correct? If it is correct, how much can the row player expect to win on average in each round? If the row player knows this number accurately and the column player does not, then the row player could offer to pay the column player a small fee for playing each round. If the fee is smaller than the expected winnings, then the row player can still be confident that over time she will make a nice profit. The purpose of this chapter is to answer these questions precisely.

Let us return now to the general setup. Consider the row player. By a *randomized strategy*, we mean that, at each play of the game, it appears (from the column player's viewpoint) that the row player is making her choices at random according to some fixed probability distribution. Let y_i denote the probability that the row player selects action i. The vector y composed of these probabilities is called a *stochastic vector*. Mathematically, a vector is a stochastic vector if it has nonnegative components that sum up to one:

$$y \geq 0 \qquad \text{and} \qquad e^T y = 1,$$

where e denotes the vector consisting of all ones. Of course, the column player must also adopt a randomized strategy. Let x_j denote the probability that the column player selects action j, and let x denote the stochastic vector composed of these probabilities.

The expected payoff to the column player is computed by summing over all possible outcomes the payoff associated with that outcome times the probability of the outcome. The set of possible outcomes is simply the set of pairs (i, j) as i ranges over the row indices $(1, 2, \ldots, m)$ and j ranges over the column indices $(1, 2, \ldots, n)$. For outcome (i, j) the payoff is a_{ij}, and, assuming that the row and column players behave independently, the probability of this outcome is simply $y_i x_j$. Hence, the expected payoff to the column player is

$$\sum_{i,j} y_i a_{ij} x_j = y^T A x.$$

2. Optimal Strategies

Suppose that the column player adopts strategy x (i.e., decides to play in accordance with the stochastic vector x). Then the row player's best defense is to use the strategy y^* that achieves the following minimum:

$$\begin{array}{ll} \text{minimize} & y^T A x \\ \text{subject to} & e^T y = 1 \\ & y \geq 0. \end{array} \tag{11.1}$$

From the fundamental theorem of linear programming, we know that this problem has a basic optimal solution. For this problem, the basic solutions are simply y vectors that are zero in every component except for one, which is one. That is, the basic optimal solutions correspond to deterministic strategies. This is fairly obvious if we look again at our example. Suppose that

$$x = \begin{bmatrix} 1/3 \\ 1/3 \\ 1/3 \end{bmatrix}.$$

Then

$$Ax = \begin{bmatrix} -1/3 \\ 1/3 \\ -1/3 \end{bmatrix},$$

and so the row player's best choice is to select either $i = 1$ (Paper) or $i = 3$ (Rock) or any combination thereof. That is, an optimal solution is $y^* = (1, 0, 0)$ (it is not unique).

Since for any given x the row player will adopt the strategy that achieves the minimum in (11.1), it follows that the column player should employ a strategy x^* that attains the following maximum:

$$\max_x \min_y y^T A x, \tag{11.2}$$

where the max and the min are over all stochastic vectors (of the appropriate dimension).

The question then becomes: how do we solve (11.2)? It turns out that this problem can be reformulated as a linear programming problem. Indeed, we have already seen that the inner optimization (the minimization) can be taken over just the deterministic strategies:

$$\min_y y^T Ax = \min_i e_i^T Ax,$$

where we have used e_i to denote the vector of all zeros except for a one in position i. Hence, the max-min problem given in (11.2) can be rewritten as

$$\begin{array}{ll} \text{maximize} & \left(\min_i e_i^T Ax\right) \\ \text{subject to} & \displaystyle\sum_{j=1}^n x_j = 1 \\ & \quad x_j \geq 0 \qquad j = 1, 2, \ldots, n. \end{array}$$

Now, if we introduce a new variable, v, representing a lower bound on the $e_i^T Ax$'s, then we see that the problem can be recast as a linear program:

$$\begin{array}{lll} \text{maximize} & v & \\ \text{subject to} & \qquad v \leq e_i^T Ax & \quad i = 1, 2, \ldots, m \\ & \displaystyle\sum_{j=1}^n x_j = 1 & \\ & \quad x_j \geq 0 & \quad j = 1, 2, \ldots, n. \end{array}$$

Switching back to vector notation, the problem can be written as

$$\begin{array}{ll} \text{maximize} & v \\ \text{subject to} & ve - Ax \leq 0 \\ & \qquad e^T x = 1 \\ & \qquad\quad x \geq 0. \end{array}$$

Finally, writing in block-matrix form, we get

$$\begin{array}{ll} \text{maximize} & \begin{bmatrix} 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ v \end{bmatrix} \\ \text{subject to} & \begin{bmatrix} -A & e \\ e^T & 0 \end{bmatrix} \begin{bmatrix} x \\ v \end{bmatrix} \begin{matrix} \leq \\ = \end{matrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} \\ & \qquad\qquad\qquad x \geq 0 \\ & \qquad\qquad\qquad v \ \text{free}. \end{array} \tag{11.3}$$

Now let's turn it around. By symmetry, the row player seeks a strategy y^* that attains optimality in the following min-max problem:

$$\min_y \max_x y^T Ax,$$

which can be reformulated as the following linear program:

$$\begin{array}{ll} \text{minimize} & u \\ \text{subject to} & ue - A^T y \geq 0 \\ & \qquad e^T y = 1 \\ & \qquad\quad y \geq 0. \end{array}$$

Writing in block-matrix form, we get

$$
\begin{array}{lll}
 & \text{minimize} & \begin{bmatrix} 0 & 1 \end{bmatrix} \begin{bmatrix} y \\ u \end{bmatrix} \\
(11.4) & \text{subject to} & \begin{bmatrix} -A^T & e \\ e^T & 0 \end{bmatrix} \begin{bmatrix} y \\ u \end{bmatrix} \begin{matrix} \geq \\ = \end{matrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} \\
 & & y \geq 0 \\
 & & u \ \text{free.}
\end{array}
$$

3. The Minimax Theorem

Having reduced the computation of the optimal strategies x^* and y^* to the solution of linear programs, it is now a simple matter to show that they are consistent with each other. The next theorem, which establishes this consistency, is called the Minimax Theorem:

THEOREM 11.1. *There exist stochastic vectors x^* and y^* for which*

$$\max_x \ {y^*}^T Ax = \min_y \ y^T Ax^*.$$

PROOF. The proof follows trivially from the observation that (11.4) is the dual of (11.3). Therefore, $v^* = u^*$. Furthermore,

$$v^* = \min_i e_i^T Ax^* = \min_y y^T Ax^*,$$

and similarly,

$$u^* = \max_j e_j^T A^T y^* = \max_x x^T A^T y^* = \max_x {y^*}^T Ax.$$

□

The common optimal value $v^* = u^*$ of the primal and dual linear programs is called the *value* of the game. From the Minimax Theorem, we see that, by adopting strategy y^*, the row player assures herself of losing no more than v units per round on the average. Similarly, the column player can assure himself of winning at least v units per round on the average by adopting strategy x^*. A game whose value is zero is therefore a *fair game*. Games where the roles of the two players are interchangeable are clearly fair. Such games are called *symmetric*. They are characterized by payoff matrices having the property that $a_{ij} = -a_{ji}$ for all i and j (in particular, m must equal n and the diagonal must vanish).

For the Paper–Scissors–Rock game, the linear programming problem that the column player needs to solve is

$$
\begin{array}{ll}
\text{maximize} & v \\
\text{subject to} & \begin{bmatrix} 0 & -1 & 2 & 1 \\ 3 & 0 & -4 & 1 \\ -5 & 6 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ v \end{bmatrix} \begin{matrix} \leq \\ \leq \\ \leq \\ = \end{matrix} \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \\
 & x_1, x_2, x_3 \geq 0 \\
 & v \ \text{free.}
\end{array}
$$

In nonmatrix notation, it looks like this:

$$
\begin{array}{lrcl}
\text{maximize} & & & v \\
\text{subject to} & -x_2 + 2x_3 + v & \leq & 0 \\
& 3x_1 \qquad - 4x_3 + v & \leq & 0 \\
& -5x_1 + 6x_2 \qquad + v & \leq & 0 \\
& x_1 + x_2 + x_3 & = & 1 \\
& x_1,\ x_2,\ x_3 & \geq & 0 .
\end{array}
$$

This linear programming problem deviates from our standard inequality form in two respects: (1) it has an equality constraint and (2) it has a free variable. There are several ways in which one can convert this problem into standard form. The most compact conversion is as follows. First, use the equality constraint to solve explicitly for one of the x_j's, say x_3:

$$x_3 = 1 - x_1 - x_2.$$

Then eliminate this variable from the remaining equations to get

$$
\begin{array}{lrcr}
\text{maximize} & v & & \\
\text{subject to} & -2x_1 - 3x_2 + v & \leq & -2 \\
& 7x_1 + 4x_2 + v & \leq & 4 \\
& -5x_1 + 6x_2 + v & \leq & 0 \\
& x_1 + x_2 & \leq & 1 \\
& x_1,\ x_2 & \geq & 0 .
\end{array}
$$

The elimination of x_3 has changed the last constraint from an equality into an inequality.

The next step is to write down a starting dictionary. To do this, we need to introduce slack variables for each of these constraints. It is natural (and desirable) to denote the slack variable for the last constraint by x_3. In fact, doing this, we get the following starting dictionary:

$$
\begin{array}{rcl}
\xi & = & \qquad\qquad\qquad v \\
\hline
x_4 & = & -2 + 2x_1 + 3x_2 - v \\
x_5 & = & 4 - 7x_1 - 4x_2 - v \\
x_6 & = & \qquad 5x_1 - 6x_2 - v \\
x_3 & = & 1 - x_1 - x_2 .
\end{array}
$$

The variable v is not constrained to be nonnegative. Therefore, there is no reason for it to be nonbasic. Let us do an arbitrary pivot with v as the entering variable and any basic variable as the leaving variable (well, not exactly any—we must make sure that it causes no division by 0, so therefore x_3 is not a candidate). Picking x_4 to leave, we get

$$
\begin{array}{rcl}
\xi & = & -2 + 2x_1 + 3x_2 - x_4 \\
\hline
v & = & -2 + 2x_1 + 3x_2 - x_4 \\
x_5 & = & 6 - 9x_1 - 7x_2 + x_4 \\
x_6 & = & 2 + 3x_1 - 9x_2 + x_4 \\
x_3 & = & 1 - x_1 - x_2 .
\end{array}
$$

Since v is free of sign constraints, it will never leave the basis (since a leaving variable is, by definition, a variable that hits its lower bound—v has no such bound). Therefore, we may as well remove it from the dictionary altogether; it can always be computed at the end. Hence, we note that

$$v = -2 + 2x_1 + 3x_2 - x_4,$$

or better yet that

$$v = \xi,$$

and the dictionary now becomes

$$\begin{array}{rcl} \xi &=& -2 + 2x_1 + 3x_2 - x_4 \\ \hline x_5 &=& 6 - 9x_1 - 7x_2 + x_4 \\ x_6 &=& 2 + 3x_1 - 9x_2 + x_4 \\ x_3 &=& 1 - x_1 - x_2 \ . \end{array}$$

At last, we are in a position to apply the simplex method. Two (tedious) iterations bring us to the optimal dictionary. Since it involves fractions, we multiply each equation by an integer to make each number in the dictionary an integer. Indeed, after multiplying by 102, the optimal dictionary is given by

$$\begin{array}{rcl} 102\xi &=& -16 - 27x_5 - 13x_6 - 62x_4 \\ \hline 102x_1 &=& 40 - 9x_5 + 7x_6 + 2x_4 \\ 102x_2 &=& 36 - 3x_5 - 9x_6 + 12x_4 \\ 102x_3 &=& 26 + 12x_5 + 2x_6 - 14x_4 \ . \end{array}$$

From this dictionary, it is easy to read off the optimal primal solution:

$$x^* = \begin{bmatrix} 40/102 \\ 36/102 \\ 26/102 \end{bmatrix}.$$

Also, since x_4, x_5, and x_6 are complementary to y_1, y_2, and y_3 in the dual problem, the optimal dual solution is

$$y^* = \begin{bmatrix} 62/102 \\ 27/102 \\ 13/102 \end{bmatrix}.$$

Finally, the value of the game is

$$v^* = \xi^* = -16/102 = -0.15686275,$$

which indicates that the row player does indeed have an advantage and can expect to make on the average close to 16 cents per round.

4. Poker

Some card games such as poker involve a round of bidding in which the players at times *bluff* by increasing their bid in an attempt to coerce their opponents into backing down, even though if the challenge is accepted they will surely lose. Similarly, they will sometimes *underbid* to give their opponents false hope. In this section, we shall study a simplified version of poker (the real game is too hard to analyze) to see if bluffing and underbidding are justified bidding strategies.

Simplified poker involves two players, A and B, and a deck having three cards, 1, 2, and 3. At the beginning of a round, each player "antes up" \$1 and is dealt one card from the deck. A bidding session follows in which each player in turn, starting with A, either (a) *bets* and adds \$1 to the "kitty" or (b) *passes*. Bidding terminates when

a bet is followed by a bet,
a pass is followed by a pass, or
a bet is followed by a pass.

In the first two cases, the winner of the round is decided by comparing cards, and the kitty goes to the player with the higher card. In the third case, bet followed by pass, the player who bet wins the round independently of who had the higher card (in real poker, the player who passes is said to *fold*).

With these simplified betting rules, there are only five possible betting scenarios:

A passes,	B passes:		\$1 to holder of higher card
A passes,	B bets,	A passes:	\$1 to B
A passes,	B bets,	A bets:	\$2 to holder of higher card
A bets,	B passes:		\$1 to A
A bets,	B bets:		\$2 to holder of higher card

After being dealt a card, player A will decide to bet along one of three lines:

1. Pass. If B bets, pass again.
2. Pass. If B bets, bet.
3. Bet.

Similarly, after being dealt a card, player B can bet along one of four lines:

1. Pass no matter what.
2. If A passes, pass, but if A bets, bet.
3. If A passes, bet, but if A bets, pass.
4. Bet no matter what.

To model the situation as a matrix game, we must identify each player's pure strategies. A pure strategy is a statement of what line of betting a player intends to follow for each possible card that the player is dealt. Hence, the players' pure strategies can be denoted by triples (y_1, y_2, y_3), where y_i is the line of betting that the player will use when holding card i. (For player A, the y_i's can take values 1, 2, and 3, whereas for player B, they can take values 1, 2, 3, and 4.)

Given a pure strategy for both players, one can compute the average payment from, say, A to B. For example, suppose that player A adopts strategy $(3, 1, 2)$ and player B adopts strategy $(3, 2, 4)$. There are six ways in which the cards can be dealt, and we can analyze each of them as follows:

card dealt A	card dealt B	betting session			payment A to B
1	2	A bets,	B bets		2
1	3	A bets,	B bets		2
2	1	A passes,	B bets,	A passes	1
2	3	A passes,	B bets,	A passes	1
3	1	A passes,	B bets,	A bets	-2
3	2	A passes,	B passes		-1

Since each of the six deals are equally likely, the average payment from A to B is $(2 + 2 + 1 + 1 - 2 - 1)/6 = 0.5$.

The calculation of the average payment must be carried out for every combination of pairs of strategies. How many are there? Player A has $3 \times 3 \times 3 = 27$ pure strategies and player B has $4 \times 4 \times 4 = 64$ pure strategies. Hence, there are $27 \times 64 = 1{,}728$ pairs. Calculating the average payment for all these pairs is a daunting task. Fortunately, we can reduce the number of pure strategies (and hence the number of pairs) that need to be considered by making a few simple observations.

The first observation is that a player holding a 1 should never answer a bet with a bet, since the player will lose regardless of the answering bet and will lose less by passing. This logic implies that, when holding a 1,

> player A should refrain from betting along line 2;
> player B should refrain from betting along lines 2 and 4.

More clearly improvable strategies can be ruled out when holding the highest card. For example, a player holding a 3 should never answer a bet with a pass, since by passing the player will lose, but by betting the player will win. Furthermore, when holding a 3, a player should always answer a pass with a bet, since in either case the player is going to win, but answering with a bet opens the possibility of the opponent betting again and thereby increasing the size of the win for the player holding the 3. Hence, when holding a 3,

> player A should refrain from betting along line 1;
> player B should refrain from betting along lines 1, 2, and 3.

Eliminating from consideration the above lines of betting, we see that player A now has $2 \times 3 \times 2 = 12$ pure strategies and player B has $2 \times 4 \times 1 = 8$ pure strategies. The number of pairs has therefore dropped to 96—a significant reduction. Not only do we eliminate these "bad" strategies from the mathematical model but also we assume that both players know that these bad strategies will not be used. That is, player A can assume that player B will play intelligently, and player B can assume the same of A. This knowledge then leads to further reductions. For example, when holding a 2, player A should refrain from betting along line 3. To reach this conclusion, we must carefully enumerate possibilities. Since player A holds the 2, player B holds either the 1 or the 3. But we've already determined what player B will do in both of those cases. Using this knowledge, it is not hard to see that player A would be unwise to bet along line 3. A similar analysis reveals that, when holding a 2, player B should refrain from lines 3 and 4. Therefore, player A now has only $2 \times 2 \times 2 = 8$ pure strategies and player B has only $2 \times 2 \times 1 = 4$ pure strategies.

At this point, no further reductions are possible. Computing the payoff matrix, we get

$$A = \begin{array}{c} \\ (1,1,2) \\ (1,1,3) \\ (1,2,2) \\ (1,2,3) \\ (3,1,2) \\ (3,1,3) \\ (3,2,2) \\ (3,2,3) \end{array} \begin{array}{c} \begin{array}{cccc} (1,1,4) & (1,2,4) & (3,1,4) & (3,2,4) \end{array} \\ \left[\begin{array}{cccc} & & \frac{1}{6} & \frac{1}{6} \\ & -\frac{1}{6} & \frac{1}{3} & \frac{1}{6} \\ \frac{1}{6} & \frac{1}{6} & -\frac{1}{6} & -\frac{1}{6} \\ \frac{1}{6} & & & -\frac{1}{6} \\ -\frac{1}{6} & \frac{1}{3} & & \frac{1}{2} \\ -\frac{1}{6} & \frac{1}{6} & \frac{1}{6} & \frac{1}{2} \\ & \frac{1}{2} & -\frac{1}{3} & \frac{1}{6} \\ & \frac{1}{3} & -\frac{1}{6} & \frac{1}{6} \end{array} \right] \end{array} .$$

Solving the matrix game, we find that

$$y^* = \left[\begin{array}{cccccccc} \frac{1}{2} & 0 & 0 & \frac{1}{3} & 0 & 0 & 0 & \frac{1}{6} \end{array} \right]^T$$

and

$$x^* = \left[\begin{array}{cccc} \frac{2}{3} & 0 & 0 & \frac{1}{3} \end{array} \right]^T .$$

These stochastic vectors can be summarized as simple statements of the optimal randomized strategies for the two players. Indeed, player A's optimal strategy is as follows:

when holding 1, mix lines 1 and 3 in 5:1 proportion;
when holding 2, mix lines 1 and 2 in 1:1 proportion;
when holding 3, mix lines 2 and 3 in 1:1 proportion.

Similarly, player B's optimal strategy can be described as

when holding 1, mix lines 1 and 3 in 2:1 proportion;
when holding 2, mix lines 1 and 2 in 2:1 proportion;
when holding 3, use line 4.

Note that it is optimal for player A to use line 3 when holding a 1 at least some of the time. Since line 3 says to bet, this bet is a bluff. Player B also bluffs sometimes, since betting line 3 is sometimes used when holding a 1. Clearly, the optimal strategies also exhibit some underbidding.

Exercises

11.1 Players A and B each hide a nickel or a dime. If the hidden coins match, player A gets both; if they don't match, then B gets both. Find the optimal strategies. Which player has the advantage? Solve the problem for arbitrary denominations a and b.

11.2 Players A and B each pick a number between 1 and 100. The game is a draw if both players pick the same number. Otherwise, the player who picks the smaller number wins unless that smaller number is one less than the opponent's number, in which case the opponent wins. Find the optimal strategy for this game.

11.3 We say that row r *dominates* row s if $a_{rj} \geq a_{sj}$ for all $j = 1, 2, \ldots, n$. Similarly, column r is said to dominate column s if $a_{ir} \geq a_{is}$ for all $i = 1, 2, \ldots, m$. Show that

(a) If a row (say, r) dominates another row, then the row player has an optimal strategy y^* in which $y^*_r = 0$.

(b) If a column (say, s) is dominated by another column, then the column player has an optimal strategy x^* in which $x^*_s = 0$.

Use these results to reduce the following payoff matrix to a 2×2 matrix:

$$\begin{bmatrix} -6 & 2 & -4 & -7 & -5 \\ 0 & 4 & -2 & -9 & -1 \\ -7 & 3 & -3 & -8 & -2 \\ 2 & -3 & 6 & 0 & 3 \end{bmatrix}.$$

11.4 Solve simplified poker assuming that antes are \$2 and bets are \$1.

11.5 Give necessary and sufficient conditions for the rth pure strategy of the row and the sth pure strategy of the column player to be simultaneously optimal.

11.6 Use the Minimax Theorem to show that

$$\max_x \min_y \; y^T A x = \min_y \max_x \; y^T A x.$$

11.7 *Bimatrix Games*. Consider the following two-person game defined in terms of a pair of $m \times n$ matrices A and B: if the row player selects row index i and the column player selects column index j, then the row player pays a_{ij} dollars and the column player pays b_{ij} dollars. Stochastic vectors x^* and y^* are said to form a *Nash equilibrium* if

$$\begin{aligned} {y^*}^T A x^* &\leq y^T A x^* && \text{for all } y \\ {y^*}^T B x^* &\leq {y^*}^T B x && \text{for all } x. \end{aligned}$$

The purpose of this exercise is to relate Nash equilibria to the problem of finding vectors x and y that satisfy

$$
\begin{aligned}
\begin{bmatrix} 0 & -A \\ -B^T & 0 \end{bmatrix} \begin{bmatrix} y \\ x \end{bmatrix} + \begin{bmatrix} w \\ z \end{bmatrix} &= \begin{bmatrix} -e \\ -e \end{bmatrix}, \\
y_i w_i &= 0, \qquad \text{for all } i, \\
x_j z_j &= 0, \qquad \text{for all } j, \\
x, w, y, z &\geq 0
\end{aligned} \tag{11.5}
$$

(vectors w and z can be thought of as being defined by the matrix equality). Problem (11.5) is called a *linear complementarity problem*.

(a) Show that there is no loss in generality in assuming that A and B have all positive entries.

(b) Assuming that A and B have all positive entries, show that, if (x^*, y^*) is a Nash equilibrium, then

$$x' = \frac{x^*}{{y^*}^T A x^*}, \qquad y' = \frac{y^*}{{y^*}^T B x^*}$$

solves the linear complementarity problem (11.5).

(c) Show that, if (x', y') solves the linear complementarity problem (11.5), then

$$x^* = \frac{x'}{e^T x'}, \qquad y^* = \frac{y'}{e^T y'}$$

is a Nash equilibrium.

(An algorithm for solving the linear complementarity problem is developed in Exercise 18.7.)

11.8 *The Game of Morra*. Two players simultaneously throw out one or two fingers and call out their guess as to what the total sum of the outstretched fingers will be. If a player guesses right, but his opponent does not, he receives payment equal to his guess. In all other cases, it is a draw.

(a) List the pure strategies for this game.

(b) Write down the payoff matrix for this game.

(c) Formulate the row player's problem as a linear programming problem. (*Hint: Recall that the row player's problem is to minimize the maximum expected payout*.)

(d) What is the value of this game?

(e) Find the optimal randomized strategy.

11.9 *Heads I Win—Tails You Lose*. In the classical coin-tossing game, player A tosses a fair coin. If it comes up heads player B pays player A \$2 but if it comes up tails player A pays player B \$2. As a two-person zero-sum game, this game is rather trivial since neither player has anything to *decide* (after agreeing to play the game). In fact, the matrix for this game is a 1×1 matrix with only a zero in it, which represents the expected payoff from player A to B.

Now consider the same game with the following twist. Player A is allowed to peek at the outcome and then decide either to stay in the game or to bow out. If player A bows out, then he automatically loses but only has to pay player B \$1. Of course, player A must inform player B of his decision. If his decision is to stay in the game, then player B has the option either to stay in the game or not. If she decides to get out, then she loses \$1 to player A. If both players stay in the game, then the rules are as in the classical game: heads means player A wins, tails means player B wins.

(a) List the strategies for each player in this game. (Hint: Don't forget that a strategy is something that a player has control over.)
(b) Write down the payoff matrix.
(c) A few of player A's strategies are uniformly inferior to others. These strategies can be ruled out. Which of player A's strategies can be ruled out?
(d) Formulate the row player's problem as a linear programming problem. (*Hints: (1) Recall that the row player's problem is to minimize the maximum expected payout. (2) Don't include rows that you ruled out in the previous part.*)
(e) Find the optimal randomized strategy.
(f) Discuss whether this game is interesting or not.

Notes

The Minimax Theorem was proved by von Neumann (1928). Important references include Gale et al. (1951), von Neumann and Morgenstern (1947), Karlin (1959), and Dresher (1961). Simplified poker was invented and analyzed by Kuhn (1950). Exercises 11.1 and 11.2 are borrowed from Chvátal (1983).