Notes on two person zero sum games

Consider a game where each player chooses one of three options with the outcome represented by the following matrix: $\begin{bmatrix} (+1,-1) & (-1,+1) & (-2,+2) \\ (-1,+1) & (+1,-1) & (+1,-1) \\ (+2,-2) & (-1,+1) & (-1,+1) \end{bmatrix}$. We interpret this as follows: the choice of player A is indicated by the rows and the choice of player B is indicated by columns. The corresponding entry has the first coordinate equal to player A's 'winnings' and the second that of player B. For example, if A chooses row 2 and B chooses column 3 the entry is $(+1,-1)$ so A 'wins' 1 and B 'wins -1. (Think of this as B pays A 1 dollar.) If A chooses row 1 and B chooses column 3 the entry is $(-2,+2)$ so A 'wins' -2 and B 'wins' +2. (Think of this as A pays B 2 dollars.) This is a zero sum game because the winnings for each entry sum to 0. Informally whatever A wins B losses and vice versa.

We consider a situation in which A and B repeatedly play this game. Is it fair in the sense that each player's can in the long term expect to break even? What kind of strategies might be used? It turns out that the best each player can do is play randomly. However there probabilities for each of the choices need not be equal. There is not better strategy. Determining the probabilities and showing that the best expected values for the players follows from work done by John Von Neumann in the 1920's. These ideas are a precurser to linear programming duality.

Since the game is zero sum we can use only the first entry in each pair which gives the 'payoff' to player A. Player A wants to maximize this payoff and player B wants to minimize it (since it is the payoff to A). Play A considers the sequence of plays and breaks them up based on which column B plays. If for example A chooses row 1 with probability .2, row 2 with probability .3 and row 3 with probability .5 then when B chooses column 1, A wins 'on average' (the expected value of A's winnings), +1 with probability .2, -1 with probability .3 and +2 with probability .5 for expected winnings of $(+1).2 + (-1).3 + (+2).5 = .9$. Similarly with these probabilities A expects -.4 when B plays column 2 and -.6 when B plays column 3. So if B 'figures out' what A is doing B will play column 3 and A will lose. On the other hand, if A plays row 1 with probability 0 (i.e., never plays this row), row 2 with probability .6 and row 3 with probability .4 then the expected winnings .2 for each column. No matter what B does, A wins on average at least .2. So this particular game is not fair. It turns out that this is the best that A can do and that B has a random strategy that guarantees that A does not do better than .2.

Player A determines a random strategy that will maximize the minimum expected winnings over all choices that B can make. Then A will on average win this amount no matter what strategy B uses. This is called a maximin strategy. Use variables x_1, x_2, x_3 to represent the probabilities for A for rows 1,2,3. The expected values for columns 1,2,3 are respectively $x_1 - x_2 + 2x_3$ and $-x_1 + x_2 - x_3$ and $-2x_1 + x_2 - x_3$. The coefficients are from the columns of the matrix (the first entries). We introduce a new variable x_0 and make each of these at least x_0. Then maximizing x_0 will maximize the minimum expected winnings. So we want to maximize x_0 subject to the constraints $x_1 - x_2 + 2x_3 \geq x_0$ and $-x_1 + x_2 - x_3 \geq x_0$ and $-2x_1 + x_2 - x_3 \geq x_0$. In addition we need x_1, x_2, x_3 to represent probabilities. So we need them to be nonnegative and to sum to 1. Rearranging these we get the following linear program. The values of x_1, x_2, x_3 are probabilities for playing the rows for A and the value of x_0 is the expected amount A is guaranteed to win. We have

$$\begin{array}{lcccccccccc}
\max & x_0 & + & 0x_1 & + & 0x_2 & + & 0x_3 & & \\
s.t. & & & x_1 & + & x_2 & - & 2x_3 & = & 1 \\
& x_0 & - & x_1 & + & x_2 & - & 2x_3 & \leq & 0 \\
& x_0 & + & x_1 & - & x_2 & + & x_3 & \leq & 0 \\
& x_0 & + & 2x_1 & - & x_2 & + & x_3 & \leq & 0 \\
& & & x_1 & & x_2 & & x_3 & \geq 0 &
\end{array}$$

In a similar manner player B sets up a linear program with variables y_1, y_2, y_3 represented probabilities for the column plays and y_0 representing player A's maximum expected winnings. Player B minimizes this to get a minimax strategy. We have

$$\begin{array}{lcccccccccc}
\min & y_0 & + & 0y_1 & + & 0y_2 & + & 0y_3 & & \\
s.t. & & & y_1 & + & y_2 & + & 2y_3 & = & 1 \\
& y_0 & - & y_1 & + & y_2 & + & 2y_3 & \geq & 0 \\
& y_0 & + & y_1 & - & y_2 & - & y_3 & \geq & 0 \\
& y_0 & - & 2y_1 & + & y_2 & + & y_3 & \geq & 0 \\
& & & y_1 & & y_2 & & y_3 & \geq 0 &
\end{array}$$

Note that the coefficients in the program for A are the negatives of the columns of the payoff matrix and the coefficients in the program for B are the negatives of the rows of the payoff matrix. The optimal solutions are $(x_0, x_1, x_2, x_3) = (.2, 0, .6, .4)$ and $(y_0, y_1, y_2, y_3) = (.2, .4, .6, 0)$. If A plays row 1 with probability 0, row 2 with probability .6 and row 3 with probability .4 they are guaranteed an average payoff of at least .2. If B plays row 1 with probability .4, row 2 with probability .6 and row 3 with probability 0 they guarantee that A has an average payoff of at most .2.

These are linear programming duals of each other. Both a feasible. For player A take arbitrary any probabilities for x_1, x_2, x_3 and take x_0 to have a large negative value and for player B take y_1, y_2, y_3 to be any probabilities and take y_0 large. So both a feasible and by strong duality for linear programming we conclude that players A and B each have optimal strategies which give the same expected value for A. That is, player A has a random strategy that guarantees on average a payoff of at least some value v (no matter what strategy B uses) and player B has a random strategy that guarantees on average that the payoff to A is at most v (no matter what strategy A uses). So these strategies are best possible for both players.

For the general case we consider only the first entries (payoffs to A) in the matrix. The payoffs to B are the negatives of these values. Note that the number of choices for A and B do not need to be the same. Let m be the number of rows in the payoff matrix and n the number of columns. Denote the entries by p_{ij}. That is, if A plays row i and B plays column j then A 'wins' p_{ij}.

The general version of the examples above yields the following linear programs whose solutions give the optimal strategies for A and B. The linear program for B is the dual of that for A. For A we have

$$\begin{array}{lllllll}
\max & x_0 & & & & & \\
s.t. & & & \sum_{i=1}^{m} x_i & = & 1 & \\
& x_0 & + & \sum_{i=1}^{m} -p_{ij} x_i & \leq & 0 & \text{for } j = 1, 2, \ldots, n \\
& & & x_i & \geq & 0 & \text{for } i = 1, 2, \ldots, m
\end{array}$$

For B we have

$$\begin{array}{lllllll}
\min & y_0 & & & & & \\
s.t. & & & \sum_{j=1}^{n} y_i & = & 1 & \\
& y_0 & + & \sum_{j=1}^{n} -p_{ij} x_j & \geq & 0 & \text{for } i = 1, 2, \ldots, m \\
& & & y_j & \geq & 0 & \text{for } j = 1, 2, \ldots, n
\end{array}$$