

Part 4. Minimax and Cake

VIJAY CHANDRU & M.R. RAO

“One of the most striking events in connection with the emergence of modern economic theory was the simultaneous but independent development of linear programming on the one hand and game theory on the other, and the eventual realization of the very close relationship that exists between these two subjects.”

David Gale, 1960

The mother of the twin brothers, Ram and Shyam, has a difficult problem to solve. There is only one cake and two hungry and competitive lads. But she is resourceful (as mothers have to be) and comes up with a brilliant solution. “Ramu”, she says, “since you are older by a few hours, you get to cut the cake into two pieces. And, Shyamu, you get to choose the piece you want”.

“Johnny” von Neumann would have been happy with this arrangement, since it illustrates his famous minimax principle of two-person, zero-sum games. This article will explain this principle in some detail and show its equivalence to the duality theorem of linear programming (proved in part 1 of this series).

You may have guessed what Ram and Shyam would do to share the cake. Ram knows that when the cake is in two pieces, his greedy twin will grab the larger of the two pieces. Hence, Ram’s best strategy has to be to cut the cake into two equal pieces and thus minimise the maximum share of cake that Shyam gets. The game can be represented by the *matrix* below, whose entries are the *payoff* to Shyam (his share of the cake).

		Ram’s Strategies	
		Cut equal Pieces	Cut unequal Pieces
Shyam’s Strategies	Choose Bigger Piece	Half	Big Piece
	Choose Smaller Piece	Half	Small Piece

{\em Figure 1.} Minimax and Cake

If Shyam had to declare his strategy first, the outcome would not be any different. He would simply declare that he would always choose the bigger piece and Ram would respond by cutting

the cake evenly. Thus maximin equals minimax, the minimum row maximum equals the maximum column minimum in the 2×2 matrix above, and the game has an obvious saddle point. *Mother* has shewn herself to be a compleat strategyst.

Mixed Strategies and Johnny's 1928 Theorem

“You know that the best you can expect is to avoid the worst.”
 Italo Calvino, 1979

Unfortunately, not all two-person, zero-sum games have as clean and simple solution as the cake example we saw above. By zero-sum we mean that the payoff to one player is always made by the other player, i.e., there are no externalities. We can always invent a rectangular payoff matrix with entries such that the minimum row maximum exceeds the maximum column minimum. Consider for example,

$$\begin{array}{|c|c|c|}
 \hline
 & 1 & -1 \\
 \hline
 1 & 1 & -1 \\
 \hline
 -1 & -1 & 1 \\
 \hline
 \end{array}$$

{\em Figure 2.} Odd-Even Game

for which the difference between the minimax and the maximin is 2. We can interpret the game corresponding to this matrix as one in which both players simultaneously call out a number between 1 and n . If the total value is even, the row player (R) is awarded a rupee by the column player (C) and vice-versa if the total is odd. So the pure strategies, for both players, are the choices of picking an odd or an even number between 1 and n . The game is played repeatedly and so one of the players can end up very rich (at the expense of the other!).

Assuming that both players are rational, we see that neither could consider playing a pure strategy since the other player would eventually catch on and cash in. Hence, the only possibility is that they play *mixed strategies* which means that they pick the pure strategies at random and according to some probability distribution of their choice. For the payoff matrix above, both R and C should choose a (0.5,0.5) mixed strategy and we see that this ensures that the expected payoff for both players is 0. Notice that the minimax of expected payoffs equals the maximin of expected payoffs.

In 1928, von Neumann proved that this closure of the gap between minimax and maximin of expected payoffs holds for mixed strategies on any payoff matrix. Check that the matrix game

$$\begin{array}{|c|c|c|c|c|}
 \hline
 & 5 & 3 & 4 & 4 \\
 \hline
 3 & 3 & 6 & 1 & 6 \\
 \hline
 \end{array}$$

{\em Figure 3.} Rectangular Game

has a value of 3.5 (expected payoff to the row player R) with mixed strategies of (5/6,1/6) by R and (0,1/2,1/2,0) by C.

We are now ready to state the general form of the theorem. The $(m \times n)$ payoff matrix $A = (a_{ij})$ defines a game for two. The entry a_{ij} represents the payoff to R (“Rose”) when she picks the i^{th} strategy and C (“Colm”) picks his j^{th} , and x_i and y_j represent the probabilities of R and C respectively, picking their i^{th} and j^{th} strategies. The resulting expected payoff to R is given by the expression $\sum_{i=1}^m \sum_{j=1}^n a_{ij}x_iy_j$. An m -vector x is called *stochastic* if $\sum_{i=1}^m x_i = 1$ and $x \geq 0$.

If R picks a stochastic m -vector x as her mixed strategy, she is assured of winning at least $\min_y \{x Ay\}$ per round on average, with the minimum taken over all stochastic n -vectors y . Note that in our notation, x is a row vector and y is a column vector. Similarly, if C’s strategy is y , he is assured that he can expect to pay no more than $\max_x \{x Ay\}$ per round.

The Minimax Theorem: For every $m \times n$ matrix A there are stochastic vectors x^* and y^* such that

$$\max_x \min_y \{x Ay\} = \min_y \max_x \{x Ay\} = x^* Ay^*$$

where the minimum is taken over stochastic n -vectors y and the maximum over stochastic m -vectors x .

This result is known as the *von Neumann Minimax Principle* and is the fundamental result of game theory. We will now see that this theorem is really a simple consequence of the duality theorem of linear programming.

LP Duality Proves the Minimax Theorem

We saw in the last section that having chosen a mixed strategy x , R can expect a payoff of at least $\min_y \{x Ay\}$ on average, where the minimum is over $\{y \geq 0 : \sum_{j=1}^n y_j = 1\}$. This is actually a simple linear programming problem with a single equality constraint on non-negative variables. Hence it follows that

$$\min_y \{x Ay\} = \min_j \left\{ \sum_{i=1}^m a_{ij}x_i \right\}$$

which exhibits an optimal extreme point solution. And similarly,

$$\max_x \{x Ay\} = \max_i \left\{ \sum_{j=1}^n a_{ij}y_j \right\}$$

Thus the problem of R finding her best strategy reduces to

$$\max_x \min_j \left\{ \sum_{i=1}^m a_{ij}x_i \right\}$$

which is equivalent to the linear programme

$$z^* = \max \left\{ z : z - \sum_{i=1}^m a_{ij}x_i \leq 0 \ (j = 1, \dots, n), \sum_{i=1}^m x_i = 1, x \geq 0 \right\} \quad (R)$$

and the problem of C finding his best strategy reduces to

$$w^* = \min\{w : w - \sum_{j=1}^n a_{ij}y_j \geq 0 (i = 1, \dots, m), \sum_{j=1}^n y_j = 1, y \geq 0\} \quad (C)$$

Notice that (R) and (C) are a pair of linear programmes that are dual to one another. Notice also, that both linear programmes are feasible. Consequently, $z^* = w^*$ and the minimax theorem is proved. We say that $z^* = w^*$ is the value of the game. Applying the complementary slackness property of optimal solutions to this dual pair of linear programmes, we observe, for each j and each i , that

$$z^* < \sum_{i=1}^m a_{ij}x_i^* \rightarrow y_j^* = 0$$

$$w^* > \sum_{j=1}^n a_{ij}y_j^* \rightarrow x_i^* = 0$$

These conditions may also be seen as a natural way of interpreting the fact that optimal minimax solutions lead to a stable equilibrium for two-person zero-sum games.

The proof of the minimax theorem, as presented in this section, leads to many useful insights - not the least of which is that the enormous algorithmic machinery of linear programming (parts 1-3 of this series) can be brought to bear on solving matrix games. Two-person zero-sum games differ from other games in that there is no reason for any negotiation between the players. This can be inferred from the observation that if (x^1, y^1) and (x^2, y^2) are both pairs of equilibrium solutions for the game, then the linear programmes (R) and (C) imply that so are (x^1, y^2) and (x^2, y^1) .

The Minimax Theorem Proves LP Duality

The minimax theorem of matrix games and duality in linear programming are of equivalent power. This seems to have been first conjectured by von Neumann (see Box 1). To complete the proof of his conjecture, we now need an argument to show that the minimax theorem implies the duality theorem.

A game is symmetric if, to begin with, the number of pure strategies of the two players are equal (i.e., the payoff matrix is square) and also the payoff to R when she chooses strategy i and C chooses j is equal to the payoff to C when he chooses i and R chooses j , holds for all i and j (i.e., the payoff matrix is *skew-symmetric* or $a_{ij} = -a_{ji}$). This is to be distinguished from the odd-even game (Figure 2.) which has a symmetric payoff matrix but is *not* a symmetric game. It is fairly easy to convince oneself that all symmetric matrix games have the following property,

Lemma: The value of a symmetric matrix game is zero.

Symmetric games capture the duality theorem of linear programming in a very natural way as we shall now see. Consider the pair of dual linear programmes

$$\min\{c'x : Ax \geq b, x \geq 0\}(P)$$

$$\max\{b'y : A'y \leq c, y \geq 0\}(D)$$

where b' , c' and A' are transposes. Now let us construct a skew-symmetric payoff matrix using the data of the linear programmes

	Y	X	t	
	-			-
Y	0	-A	+b	
X	+A'	0	-c	
t	-b'	c'	0	
	-			-

Figure 4. Primal-Dual Payoff

From the lemma we surmise that the value of the game is zero and hence that any optimal mixed strategies (Y^*, X^*, t^*) must satisfy the following linear inequality system.

$$\{-AX + bt \leq 0, A'Y - ct \leq 0, -b'Y + c'X \leq 0, X \geq 0, Y \geq 0\}$$

The justification for this observation is that, for a symmetric game, the linear programmes (R) and (C) are solvable with $z^* = w^* = 0$. In addition, if $t^* > 0$ we can dehomogenise the linear inequality system by substituting $X \leftarrow xt$ and $Y \leftarrow yt$ and we would have the pair (x^*, y^*) satisfying

$$\{-Ax^* + b \leq 0, A'y^* - c \leq 0, -b'y^* + c'x^* \leq 0, x^* \geq 0, y^* \geq 0\}$$

which we recognise as precisely the necessary and sufficient conditions for optimality of the linear programmes (P) and (D).

What if $t^* = 0$ in an optimal strategy for the matrix game? The dehomogenising trick will not work and in general we need additional work to extract optimal solutions to (P) and (D). The details get a bit too technical and we will skip them here. Suffice it to say that this degenerate situation can be avoided altogether by assuming that the primal linear programme (P) has a full-dimensional convex polytope as its feasible region. There is no loss of generality in making this assumption since it can be implemented by introducing an innocuous auxiliary variable into the linear programme. Under this assumption, it can be shown that all solutions to the game must have $t^* > 0$.

Conversely, if we had a pair of optimal solutions (\tilde{x}, \tilde{y}) to the linear programmes (P) and (D), we could homogenise this solution by defining $t^* > 0$ by the identity

$$t(\sum \tilde{x}_i + \sum \tilde{y}_j + 1) = 1$$

and substituting $\tilde{x}t^* \leftarrow X^*$ and $\tilde{y}t^* \leftarrow Y^*$. This would yield a mixed strategy (Y^*, X^*, t^*) that solves the game represented by the matrix of Figure 4. Thus we have established,

The Skew-Symmetric Theorem: Corresponding to any pair of dual linear programmes is a skew-symmetric matrix game such that optimal solutions to the linear programmes can be extracted from any optimal (minimax=maximin) mixed strategy of the game.

G.B. Dantzig developed an algorithm for solving skew-symmetric matrix games using a pivoting procedure akin to the simplex method that works on tableaux (or dictionaries as we called them in

part 2 of this series). This has come to be known as the self-dual parametric method. This approach has been generalised to the setting of *bimatrix games* via the *linear complementarity problem* or LCP (see Box 2).

The Saddle-Point of the Lagrangean

We conclude this article with a two-person zero-sum game interpretation of the duality theorem of linear programming that has been found to be very useful in mathematical economics.

Primal (P) is a producer of goods in a closed economy. Let (x_1, x_2, \dots, x_n) denote the bundle of goods produced by P. Producing goods requires resources and let (b_1, b_2, \dots, b_m) denote the set of resources owned by P to start with. To produce the goods he sells, P requires a_{ij} units of resource i for each unit of good j manufactured. The selling price of good j has been fixed at c_j .

Dual (D) is the adversary of P and represents the “market” of the closed economy. The player D has to pay P for the goods produced. In addition, D can sell (or buy) additional resources to (from) P for his production activities but at a cost. Let y_i denote the “price” charged by D for each additional unit of resource i provided to P.

The game is as follows. The two players P and D are allowed to pick strategies x and y respectively. So if P announces an activity level of $x \geq 0$ in production and D announces a price vector $y \geq 0$, the total payoff to P equals

$$L(x, y) = c'x - (Ax - b)'y = (c' - y'A)x + b'y$$

and is called the Lagrangean. The payoff is the total revenue earned by P, less the cost of the additional resources used by him. Since this is a closed economy, the burden of the payoff rests with D.

Let us now examine the conditions that any stable solution (x^*, y^*) must satisfy.

- A stable solution must correspond to a saddle-point of $L(x, y)$. That is, $\min_y \max_x \{L(x, y)\} = \max_x \min_y \{L(x, y)\} = L(x^*, y^*)$.
- $Ax^* \leq b$ must hold. For if $\sum_{j=1}^n a_{ij}x_j^* > b_i$ for some i , then D would pick an arbitrarily large price y_i^* for the resource i and the payoff would be driven to $-\infty$.
- $A'y^* \geq c$ must hold. For if $\sum_{i=1}^m a_{ij}y_i^* < c_j$ for some j , then P would manufacture arbitrarily large amounts of good j and the payoff would be driven to $+\infty$.
- If $\sum_{j=1}^n a_{ij}x_j^* < b_i$ for some i , then D would pick the price $y_i^* = 0$ for the resource i since otherwise the payoff to P would be unnecessarily high.
- Similarly, if $\sum_{i=1}^m a_{ij}y_i^* > c_j$ for some j , then P would not manufacture good j and hence $x_j^* = 0$.

The alert reader would have recognised that these market equilibrium conditions are precisely the necessary and sufficient conditions for optimality of the primal/dual pair of linear programmes

$$\max\{c'x : Ax \leq b, x \geq 0\}(P)$$

$$\min\{b'y : A'y \geq c, y \geq 0\}(D)$$

This explains why dual solutions are often called “shadow prices” in the literature on linear programming.

Suggested Reading

G.B.Dantzig, *Linear Programming and Extensions*, Princeton University Press, Princeton University Press, 1963.

David Gale, *The theory of linear economic models*, McGraw Hill, 1960.

K.G. Murty, *Linear and Combinatorial Programming*, Wiley, 1976.

Guillermo Owen, *Game Theory*, 3rd Edition, 1995.

William Poundstone, *Prisoner's Dilemma*, Doubleday, 1992.

J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behaviour*, Princeton University Press, 1944.

Williams, *The Compleat Strategyst*

BOX 1. A HISTORIC ENCOUNTER

George Dantzig loves to tell the story of his meeting with John von Neumann on October 3, 1947 at the Institute for Advanced Study at Princeton. Dantzig went to that meeting with the express purpose of describing the linear programming problem to von Neumann and asking him to suggest a computational procedure. He was actually looking for methods to benchmark the simplex method against. Instead, he got a 90 minute lecture on Farkas Lemma and Duality (Dantzig's notes of this session formed the source of the modern perspective on linear programming duality). Not wanting Dantzig to be completely amazed, von Neumann admitted

“I don't want you to think that I am pulling all this out of my sleeve like a magician. I have recently completed a book with Morgenstern on the theory of games. What I am doing is conjecturing that the two problems are equivalent. The theory that I am outlining is an analogue to the one we have developed for games.”

BOX 2. BIMATRIX GAMES AND LINEAR COMPLEMENTARITY

A two person, *non-zero* sum game is one in which the sum of the payoffs to the two players is not zero. In this case we need to consider two payoff matrices, A the payoffs to R and B the payoffs to C. We can assume, with no loss of generality, that the elements of both A and B are strictly positive. Adding the same large positive constant to each of the elements does not change the equilibrium or stable strategies but merely changes the expected payoff by the added constant. If x^* and y^* denote the stable strategies for R and C respectively, we must have $x^*Ay^* \geq xAy^*$ for all stochastic m -vectors x and $x^*By^* \geq x^*By$ for all stochastic n -vectors y . Since x and y are stochastic, we may as well rewrite these conditions as

$$x^*Ay^* \geq \sum_j a_{ij}y_j^* \quad i = 1, 2, \dots, m$$

$$x^*By^* \geq \sum_i b_{ij}x_i^* \quad j = 1, 2, \dots, n$$

which can be further simplified to

$$Av + s = e_m \quad \text{and} \quad uB + r = e_n$$

where e_k denotes a column vector of dimension k with all entries equalling 1 and all variables u, v, r, s are non-negative. The simplification has used the reversible (show this) substitutions $u = (x^*By^*)^{-1}x^*$ and $v = (x^*Ay^*)^{-1}y^*$ and the slack variables r and s . These are now just necessary conditions for stable strategies to the bimatrix game. To make them sufficient we add the complementarity conditions $rv = us = 0$. The reader should verify that complementarity as stated above is the same as what we have encountered before. For example, saying that $x^*Ay^* > \sum_j a_{ij}y_j^*$ would imply that $x_i^* = 0$ is compactly expressed by $u_i s_i = 0$.

The search for optimal/stable strategies for bi-matrix games therefore reduces to solving the system

$$Av + s = e_m, \quad uB + r = e_n, \quad rv = us = 0, \quad u, v, r, s \geq 0$$

which is a special case of the *linear complementarity problem*.