Zero Sum Games

IE 5545: Decision Analysis
Prof. Krishnamurthy Iyer*

May 2, 2026

We now consider a special class of games involving two players, where the sum of their payoffs is always zero for any strategy profile. Thus, the gain of one player is the loss of the other. Such games are called *zero-sum games*.

Definition 1. *A game G involving two players with strategy sets $\mathcal{S}_1$ and $\mathcal{S}_2$ is a zero-sum game if*

$$\pi_1(s_1, s_2) + \pi_2(s_1, s_2) = 0 \qquad \textit{for all } s_1 \in \mathcal{S}_1,\ s_2 \in \mathcal{S}_2.$$

Simple examples of zero-sum games include matching pennies, and rock-paper-scissors. Furthermore, a game of chess is a zero-sum game if we assume each player obtains a payoff of 1 on winning, a payoff of -1 on losing and a payoff of 0 upon drawing the game.

From the definition, it follows immediately that for any mixed strategies σ_1 and σ_2 for the two players, we have

$$\pi_1(\sigma_1, \sigma_2) + \pi_2(\sigma_1, \sigma_2) = 0.$$

i.e., the zero-sum condition also holds for mixed strategies.

Recall that a player's preferences do not change if a constant is added to their payoff and/or their payoffs are multiplied by a positive constant. In particular, these payoff transformations do not affect the set of best responses of a player. Using this observation, we can show that many two-player games for which the sum of the players' payoffs does not equal zero can still be equivalent to a zero-sum game:

1. Constant sum games: Suppose we have a two-player game where the payoffs satisfy

   $$\pi_1(s_1, s_2) + \pi_2(s_1, s_2) = K \qquad \text{for all } s_1 \in \mathcal{S}_1 \text{ and } s_2 \in \mathcal{S}_2,$$

   where $K \neq 0$ is a constant. Such games are called constant-sum games. (If $K = 0$, we get back zero-sum games.)

   Any constant-sum game G is equivalent to a zero-sum game $\widetilde{G}$. (Here, we say two games are equivalent if for each player i and for any strategy profile σ_{-i} of the other players, the set $BR_i(\sigma_{-i})$ of best responses of player i is the same in the two games.) To see this, given a constant-sum game G, define the game $\widetilde{G}$ with the same players and strategy sets, and with payoffs

   $$\begin{aligned}\widetilde{\pi}_1(s_1, s_2) &= \pi_1(s_1, s_2) - K \\ \widetilde{\pi}_2(s_1, s_2) &= \pi_2(s_1, s_2).\end{aligned}$$

   Since subtracting a constant does not affect preferences, we conclude that the two games G and $\widetilde{G}$ are equivalent, and furthermore, now the payoffs in $\widetilde{G}$ add up to zero.

2. Strictly competitive games: Consider a two-player game with payoffs satisfying

$$A\,\pi_1(s_1, s_2) + B\,\pi_2(s_1, s_2) = K \qquad \text{for all } s_1 \in \mathcal{S}_1 \text{ and } s_2 \in \mathcal{S}_2,$$

where $A > 0$, $B > 0$, and K is some constant. Such games are called *strictly competitive games*. Once again, any strictly competitive game G is equivalent to a zero-sum game $\widetilde{G}$, with same strategy sets, and payoff functions $\widetilde{\Pi}_1$ and $\widetilde{\Pi}_2$ such that the

$$\widetilde{\pi}_1(s_1, s_2) = A\pi_1(s_1, s_2),$$
$$\widetilde{\pi}_2(s_1, s_2) = B\pi_2(s_1, s_2) - K.$$

(Exercise: verify that the best response mappings for the two players in G and $\widetilde{G}$ are the same.)

Examples of constant-sum games or strictly competitive games are common, e.g., dividing a resource between two people, apportioning a chore requiring a fixed amount of effort between two people, etc.

Beyond such settings, we can also think of zero-sum games as modeling general strategic situations where players do not have complete information about other players' payoffs. In this case, one can take a "worst-case" view, where a player i considers all other players to be *adversarial*, i.e., that the other players together seek to minimize the payoff of player i. This would correspond to a zero-sum game where the player i is competing against a coalition of all other players, whose joint payoff equals the negative of the player i's payoff. Any guarantee player i can obtain under this worst case scenario, will continue to hold in the original setting, where the players need not be adversarial. Such an adversarial analysis is common in many contexts, e.g., statistics (comparison of estimators), computer science (design and analysis of algorithm), machine learning (model training).

How do we analyze zero-sum games?

Since finite zero-sum games are just a special class of finite games, there exists a mixed Nash equilibrium, and we can use it to predict how players will act in such games. Furthermore, unlike general games, we know that for zero sum games, fictitious play can find the Nash equilibrium (see Julia Robinson, 1951). Thus, there exists a constructive algorithm to find the Nash equilibrium for zero-sum games.

However, the general approach does not give us any insight into the fact that the two players are perfectly competitive in a zero-sum game. Does this fact make the analysis easier? Can we use the fact that the payoffs always add up to zero to find the equilibrium more efficiently? To answer these questions, we analyze the zero-sum games directly. (This is also historically more accurate, since the analysis we present below was developed much earlier (in the 1920s) than the notion of Nash equilibrium.)

A pessimistic world-view

To understand how a player should play in a zero-sum game, we consider a more pessimistic world-view: player 1 assumes no matter which strategy she chooses, player 2 knows the chosen strategy and chooses her own strategy based on this knowledge; since it is a zero-sum game, player 2 will then choose the strategy the minimizes player 1's payoff. The reason for taking this pessimistic view is that it allows player 1 to obtain a guarantee (i.e., a lower bound) on her payoff in the original (simultaneous move) game.

To elaborate, if player 1 chooses a pure strategy s_1, then under this view, player 2 knows player 1 will choose s_1 and chooses a strategy s_2 such that the payoff $\pi_1(s_1, s_2)$ is the minimum over all s_2. (For now, we focus on the case where players only pick pure strategies; we will generalize to mixed strategies later.)

Thus, under this pessimistic view, player 1's payoff on choosing a pure strategy s_1 is given by

$$\min_{s_2 \in S_2} \pi_1(s_1, s_2).$$

If player 1 believes this pessimist view, then she will choose a strategy s_1 that maximizes $\min_{s_2 \in S_2} \pi_1(s_1, s_2)$ and hence player 1's payoff will be

$$\max_{s_1 \in S_1} \left(\min_{s_2 \in S_2} \pi_1(s_1, s_2) \right) = m_P.$$

The value m_P is called the *maximin value with pure strategies* for player 1 (the subscript p is to emphasize the restriction to pure strategies). When player 1 plays the strategy s_1^* that attains the maximum above, player 1's payoff will be at least m_P irrespective of how player 2 plays.

Similarly, we can take the same pessimistic view for player 2. In this case, on playing a strategy s_2, player 2 believes player 1 knows the player 2 will play s_2 and hence will pick strategy that minimizes player 1's payoff. In this scenario, player 2's payoff is

$$\min_{s_1 \in S_1} \pi_2(s_1, s_2).$$

Since player 2 maximizes this payoff, her payoff (with pure strategy) is given by

$$\max_{s_2 \in S_2} \left(\min_{s_1 \in S_1} \pi_2(s_1, s_2) \right).$$

Now, note that in this situation, because we have a zero sum game, player 1's payoff is given by

$$\begin{aligned} -\max_{s_2 \in S_2} \left(\min_{s_1 \in S_1} \pi_2(s_1, s_2) \right) &= \min_{s_2 \in S_2} \left(-\min_{s_1 \in S_1} \pi_2(s_1, s_2) \right) \\ &= \min_{s_2 \in S_2} \max_{s_1 \in S_1} \left(-\pi_2(s_1, s_2) \right) \\ &= \min_{s_2 \in S_2} \max_{s_1 \in S_1} \pi_1(s_1, s_2). \end{aligned}$$

Thus, in the pessimistic view of player 2, player 1's payoff is

$$M_P = \min_{s_2 \in S_2} \max_{s_1 \in S_1} \pi_1(s_1, s_2).$$

We call M_P as the *minimax value with pure strategies* for player 1. Note that $-M_P$ is the maximin value with pure strategies for player 2.

To summarize, we have the following expressions for the maximin value and the minimax value (with pure strategies) for player 1:

$$\begin{aligned} \text{Maximin value with pure strategies } m_p &= \max_{s_1 \in S_1} \left(\min_{s_2 \in S_2} \pi_1(s_1, s_2) \right) \\ \text{Minimax value with pure strategies } M_P &= \min_{s_2 \in S_2} \left(\max_{s_1 \in S_1} \pi_1(s_1, s_2) \right). \end{aligned}$$

The above quantities have the following interpretation: irrespective of how player 2 plays, player 1 can get at least m_P, if she plays the strategy s_1^* that achieves the maximum in the definition of m_p. Similarly, irrespective of how player 1 plays, the payoff of player 1 is at most M_P, if player 2 plays the strategy s_2^* that achieves the minimum in the expression for M_p.

Example. Consider the game

	L	R
U	$(3, -3)$	$(5, -5)$
D	$(4, -4)$	$(1, -1)$

It is easy to see that

$$\min_{s_2 \in S_2} \pi_1(U, s_2) = 3, \qquad \min_{s_2 \in S_2} \pi_1(D, s_2) = 1,$$

and hence

$$m_p = \max_{s_1 \in S_1} \left[\min_{s_2 \in S_2} \pi_1(s_1, s_2) \right] = 3.$$

Since the strategy U attains the maximum, we conclude that if player 1 plays U, her payoff is at least $m_p = 3$ no matter how player 2 plays.

Similarly, we have

$$\max_{s_1 \in S_1} \pi_1(s_1, L) = 4, \qquad \max_{s_1 \in S_1} \pi_1(s_1, R) = 5,$$

and hence

$$M_p = \min_{s_2 \in S_2} \left[\max_{s_1 \in S_1} \pi_1(s_1, s_2) \right] = 4.$$

The strategy that attains the minimum is L, so if player 2 plays L, then player 1's payoff is at most $M_p = 4$, no matter how player 1 plays. $\square$

For this game, we see that $m_P < M_P$. In general, we can show $m_P \leq M_P$ as follows: observe that for any (s_1, s_2) we have

$$\pi_1(s_1, s_2) \geq \min_{s_2' \in S_2} \pi_1(s_1, s_2').$$

This implies that

$$\max_{s_1 \in S_1} \pi_1(s_1, s_2) \geq \max_{s_1 \in S_1} \left(\min_{s_2' \in S_2} \pi_1(s_1, s_2') \right) \qquad \text{for all } s_2.$$

The right hand side is just m_P. Thus,

$$\max_{s_1 \in S_1} \pi_1(s_1, s_2) \geq m_P \qquad \text{for all } s_2.$$

Taking minimum over all s_2, we conclude

$$\min_{s_2 \in S_2} \left(\max_{s_1 \in S_1} \pi_1(s_1, s_2) \right) \geq m_P.$$

The left-hand side is M_P. Hence, $\boxed{M_P \geq m_P.}$

Thus the maximin value (with pure strategies) is less than or equal to the minimax value (with pure strategies), and we already saw an example where it is strictly smaller. This is expected, since the interpretation is the m_p is a lower-bound on player 1's payoff in the original (simultaneous-move) game, whereas M_p is an upper-bound.

Extending to mixed strategies

Until now, we considered only pure strategy, but our discussion carries as is to mixed strategies. If we let σ_1 denote player 1's mixed strategy and σ_2 denote player 2's mixed strategy then we can define

$$m = \max_{\sigma_1}\left(\min_{\sigma_2}\pi_1(\sigma_1,\sigma_2)\right) \qquad = \text{maximin value (for player 1)},$$
$$M = \min_{\sigma_2}\left(\max_{\sigma_1}\pi_1(\sigma_1,\sigma_2)\right) \qquad = \text{minimax value (for player 1)}.$$

In the above definition, each player is free to choose a mixed strategy. Thus, for instance, in this pessimistic view for player 1, player 1 believes that no matter which mixed strategy σ_1 she chooses, player 2 knows this choice and seeks to choose a mixed strategy σ_2 to minimize player 1's payoff $\pi_1(\sigma_1,\sigma_2)$ (i.e., to maximizer her own payoff).

The strategy σ_1^* that achieves the maximum in the definition of the maximin value m is known as the *maximin strategy* for player 1. Note that the strategy σ_2^* that attains the minimum in the definition of the minimax value M is the maximin strategy for player 2.

$$\sigma_1^* \in \arg\max_{\sigma_1}\left(\min_{\sigma_2}\pi_1(\sigma_1,\sigma_2)\right) \qquad = \text{maximin strategy for player 1},$$
$$\sigma_2^* \in \arg\min_{\sigma_2}\left(\max_{\sigma_1}\pi_1(\sigma_1,\sigma_2)\right) \qquad = \text{maximin strategy for player 2}.$$

As before, we have the following interpretations: irrespective of how player 2 plays, if player 1 plays her maximin strategy σ_1^*, then her payoff is at least her maximin value m. Similarly, irrespective of how player 1 plays, if player 2 plays her maximin strategy σ_2^*, then the payoff of player 1 is at most the minimax value M.

Using a similar set of inequalities as before, we can show that $m \leq M$. Furthermore, it is not hard to show that $m_p \leq m$ and $M \leq M_p$. In other words, using mixed strategies, player 1 can get a higher lower-bound on her payoffs, and player 2 can obtain a lower upper-bound on player 1's payoffs.

Example (contd.)**.** Consider the two-player game analyzed earlier.

	L	R
U	$(3,-3)$	$(5,-5)$
D	$(4,-4)$	$(1,-1)$

Denote the mixed strategies for the two players as follows:

$$\sigma_1 = \begin{bmatrix} p \\ 1-p \end{bmatrix} \qquad \text{for } 0 \leq p \leq 1,$$
$$\sigma_2 = \begin{bmatrix} q \\ 1-q \end{bmatrix} \qquad \text{for } 0 \leq q \leq 1.$$

Then, for a given σ_1, we have

$$\begin{aligned}
\min_{\sigma_2}\pi_1(\sigma_1,\sigma_2) &= \min_{\sigma_2}\begin{bmatrix} p & 1-p \end{bmatrix}\begin{bmatrix} 3 & 5 \\ 4 & 1 \end{bmatrix}\begin{bmatrix} q \\ 1-q \end{bmatrix} \\
&= \min_{\sigma_2}\begin{bmatrix} 3p+4(1-p) & 5p+1-p \end{bmatrix}\begin{bmatrix} q \\ 1-q \end{bmatrix} \\
&= \min_{0\leq q\leq 1}\left\{q(4-p)+(1-q)(1+4p)\right\}.
\end{aligned}$$

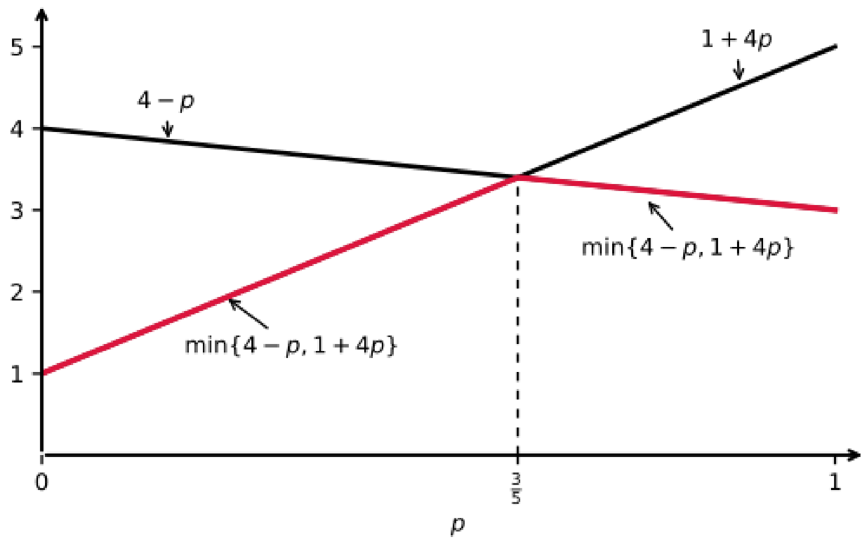

Figure 1: Plot of $\min\{4-p, 1+4p\}$ over $[0,1]$.

Since the expression is linear in q, the minimum will be attained at an endpoint, i.e., either at $q=0$ or at $q=1$. This implies

$$\min_{0\leq q\leq 1}\{q(4-p)+(1-q)(1+4p)\} = \min\{4-p, 1+4p\}.$$

Thus,

$$\min_{\sigma_2} \pi_1(\sigma_1,\sigma_2) = \min\{4-p, 1+4p\}.$$

and hence

$$\max_{\sigma_1}\left(\min_{\sigma_2} \pi_1(\sigma_1,\sigma_2)\right) = \max_{0\leq p\leq 1}\left[\min\{4-p, 1+4p\}\right].$$

The function $\min\{4-p, 1+4p\}$ over $0 \leq p \leq 1$ is plotted in Figure 1. We see that maximum of the function is attained at $p = \frac{3}{5}$, and we have

$$\max_{0\leq p\leq 1}\left(\min\{4-p, 1+4p\}\right) = \frac{17}{5}.$$

Thus,

$$m = \max_{\sigma_1}\left(\min_{\sigma_2} \pi_1(\sigma_1,\sigma_2)\right) = \frac{17}{5}.$$

Thus, the maximin value for player 1 is $m = \frac{17}{5}$. Furthermore, since the maximum value is attained at $p = \frac{3}{5}$, the maximin strategy for player 1 is given by $\sigma_1^* = [\frac{3}{5}\ \ \frac{2}{5}]^\mathsf{T}$. In other words, if player 1 plays strategy σ_1^*, then her payoff will be at least $m = \frac{17}{5}$, no matter what strategy player 2 chooses.

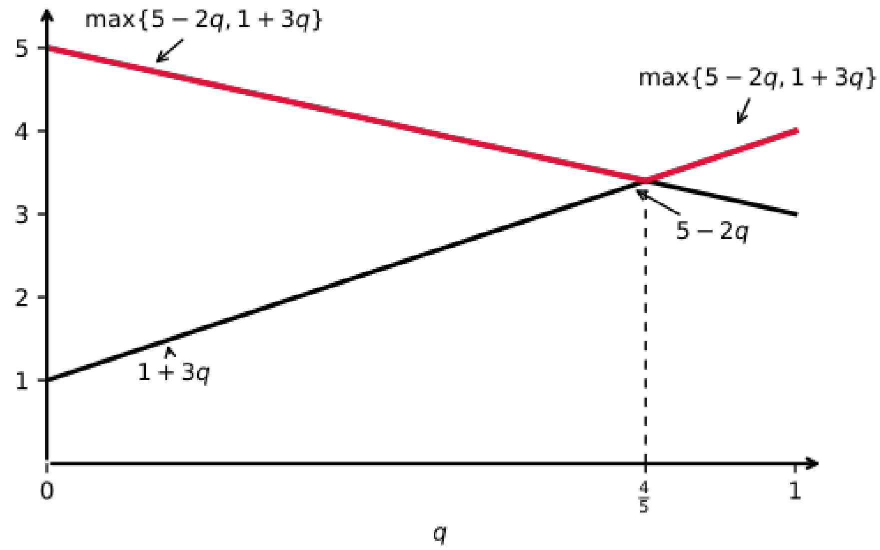

Figure 2: Plot of $\max\{5-2q, 1+3q\}$ over $[0,1]$.

Similarly, for a fixed σ_2, we have

$$
\begin{aligned}
\max_{\sigma_1} \pi_1(\sigma_1, \sigma_2) &= \max_{\sigma_1} \begin{bmatrix} p & 1-p \end{bmatrix} \begin{bmatrix} 3 & 5 \\ 4 & 1 \end{bmatrix} \begin{bmatrix} q \\ 1-q \end{bmatrix} \\
&= \max_{\sigma_1} \begin{bmatrix} p & 1-p \end{bmatrix} \begin{bmatrix} 3q + 5(1-q) \\ 4q + 1 - q \end{bmatrix} \\
&= \max_{0 \le p \le 1} \begin{bmatrix} p & 1-p \end{bmatrix} \begin{bmatrix} 5 - 2q \\ 1 + 3q \end{bmatrix} \\
&= \max_{0 \le p \le 1} p(5-2q) + (1-p)(1+3q) \\
&= \max\{5-2q, 1+3q\}.
\end{aligned}
$$

Here the last equality follows by the same argument as before: since the function being maximized in the penultimate equality is linear in p, the maximum will be attained at endpoint, i.e., either at $p = 0$ or $p = 1$. Thus, we have

$$
\begin{aligned}
M &= \min_{\sigma_2} \left(\max_{\sigma_1} \pi_1(\sigma_1, \sigma_2) \right) \\
&= \min_{\sigma_2} \big(\max\{5-2q, 1+3q\} \big) \\
&= \min_{0 \le q \le 1} \big(\max\{5-2q, 1+3q\} \big).
\end{aligned}
$$

The function $\max\{5-2q, 1+3q\}$ over $0 \le q \le 1$ is plotted in Figure 2. We see the minimum of this function is attained at $q = \frac{4}{5}$, and we have

$$
\begin{aligned}
M &= \min_{\sigma_2} \max_{\sigma_1} \pi_1(\sigma_1, \sigma_2) \\
&= \min_{0 \le q \le 1} \big(\max\{5-2q, 1+3q\} \big) \\
&= \frac{17}{5}.
\end{aligned}
$$

Thus, the minimax value for player 1 is $M = \frac{17}{5}$. Furthermore, since the minimum value is attained at $q = \frac{4}{5}$, the maximin strategy for player 2 is given by $\sigma_2^* = [\frac{4}{5} \;\; \frac{1}{5}]^\mathsf{T}$. In other words, if player 2 plays strategy σ_2^*, then player 1's payoff will be at most $M = \frac{17}{5}$, no matter what strategy player 1 chooses.

We can summarize these findings as below:

1. When we extend our analysis to mixed strategies, we observe that in this game, the maximin value is equal to the minimax value:
$$M = m = \frac{17}{5}.$$
Note that, with pure strategies, we obtained $m_p = 3$ and $M_p = 4$. So we see here an example with $m_p < m$ and $M < m_p$. In other words, using mixed strategies, both players can provide better guarantees on the payoff of player 1.

2. If player 1 plays her maximin strategy σ_1^*, her payoff is at least $m = \frac{17}{5}$, regardless of the strategy of player 2:
$$\pi_1(\sigma_1^*, \sigma_2) \geq m = \frac{17}{5}, \quad \text{for all } \sigma_2. \tag{1}$$

3. If player 2 plays her maximin strategy σ_2^*, the payoff of player 1 is at most $M = \frac{17}{5}$, regardless of the strategy of player 1:
$$\pi_1(\sigma_1, \sigma_2^*) \leq M = \frac{17}{5}, \quad \text{for all } \sigma_1. \tag{2}$$

4. If player 1 plays σ_1^* and player 2 plays σ_2^*, the payoff of player 1 is exactly $m = M = \frac{17}{5}$:
$$\pi_1(\sigma_1^*, \sigma_2^*) = \frac{17}{5}. \tag{3}$$

5. From (2) and (3), we conclude that σ_1^* is the best response of player 1 against σ_2^*:
$$\sigma_1^* \in BR_1(\sigma_2^*).$$

6. From (1), using the fact that $\pi_2 = -\pi_1$, we conclude that
$$\pi_2(\sigma_1^*, \sigma_2) \leq -\frac{17}{5}, \quad \text{for all } \sigma_2.$$
Also, (3) implies that $\pi_2(\sigma_1^*, \sigma_2^*) = -\frac{17}{5}$. Together, we conclude that σ_2^* is a best response against σ_1^*:
$$\sigma_2^* \in BR_2(\sigma_1^*).$$

7. Since $\sigma_1^* \in BR_1(\sigma_2^*)$ and $\sigma_2^* \in BR_2(\sigma_1^*)$, we infer that (σ_1^*, σ_2^*) is a mixed Nash equilibrium. This is surprising, since even with a completely different approach we have arrived at the same equilibrium concept in this game.

We want to see if our observations for this example carry over to general zero sum games. Namely, we have the following two questions:

(1) For general zero-sum games, if we consider mixed strategies for both players, are the maximin value and the minimax value always equal?

(2) Do the maximin strategies of the two players always form a Nash equilibrium?

It turns out that indeed, the answers to the two questions are yes! We will show this in detail in the next section.

Minimax theorem

The positive answer to the first question in the preceding section is an important theorem, which says that for a finite zero sum game, the minimax value and the maximin value are equal.

Theorem 1 (Minimax theorem)**.** *For a finite zero sum game*

$$\min_{\sigma_2}\left(\max_{\sigma_1}\pi_1(\sigma_1,\sigma_2)\right)=\max_{\sigma_1}\left(\min_{\sigma_2}\pi_1(\sigma_1,\sigma_2)\right).$$

This theorem is one of the first in the theory of games, and we will show this using linear programming duality. (In fact, linear programming duality can be proven using the minimax theorem.) Using this theorem, we can prove that the maximin strategies form a Nash equilibrium. Before we proceed, we provide a brief refresher to linear programming duality.

LP duality: Refresher

Recall that every linear program of the form

$$\begin{aligned} \max \quad & c^\top x && \text{(Primal LP)}\\ \text{s.t.} \quad & Ax \le b \\ & x \ge 0, \end{aligned}$$

is related to another linear program, known as its *dual*, given by

$$\begin{aligned} \min \quad & b^\top y && \text{(Dual LP)}\\ \text{s.t.} \quad & A^\top y \ge c \\ & y \ge 0. \end{aligned}$$

The duality relationship is symmetric: the dual of the dual linear program is the original (*primal*) linear program. It is evident that for each constraint $(Ax)_i \le b_i$ in the primal, there is a variable y_i in the dual. Similarly, for each constraint $(A^\top y)_j \ge c_j$ in the dual, there is a variable x_j in the primal.

Weak duality says that if x is feasible for the primal LP, and y is feasible for the dual LP, then we must have $c^\top x \le b^\top y$.

Strong duality says that one of the following must occur: (1) either both primal and the dual are infeasible; (2) primal is infeasible and dual is unbounded; (3) primal is unbounded and dual is infeasible; or (4) both primal and dual are feasible and bounded. In the last case, the optimal values of the two linear programs are equal.

Thus, strong duality implies that if the primal LP is feasible and bounded, then the dual LP is also feasible and bounded, and both have optimal solutions x^* and y^* that satisfy

$$c^\top x^* = b^\top y^*.$$

In this case, the optimal solutions satisfy the *complementary slackness* conditions:

$$\text{if } x_j^* > 0, \text{ then } (A^\top y^*)_j = c_j,$$
$$\text{if } y_i^* > 0, \text{ then } (Ax^*)_i = b_i.$$

Equivalently, (1) if $(A^\top y^*)_j > c_j$, then $x_j^* = 0$ and (2) if $(Ax^*)_i < b_i$ then $y_i^* = 0$.

Proof of the minimax theorem

For a general finite zero-sum game, let player 1's strategy set contain m strategies and player 2's strategy set contain n strategies. Let

$$\mathcal{S}_1 = \{1, 2, \ldots, m\}, \qquad \mathcal{S}_2 = \{1, 2, \ldots, n\}.$$

Let A denote player 1's payoff matrix, with the rows being player 1's strategies, columns being player 2's strategies, and the entry a_{ij} at row i and column j being player 1's payoff when player 1 plays strategy i and player 2 plays strategy j. Note that A is a $m \times n$ matrix.

Denote player 1's mixed strategy σ_1 as

$$p = \begin{bmatrix} p_1 \\ p_2 \\ \vdots \\ p_m \end{bmatrix} \qquad \text{with } p_i \geq 0 \text{ and } \sum_{i=1}^m p_i = 1.$$

Similarly, denote player 2's mixed strategy σ_2 as

$$q = \begin{bmatrix} q_1 \\ \vdots \\ \vdots \\ q_n \end{bmatrix} \qquad \text{with } q_i \geq 0 \text{ and } \sum_{i=1}^n q_i = 1$$

Note that p is an $m \times 1$ probability vector and q is an $n \times 1$ probability vector.

With this notation, we can write player 1's payoff as

$$\pi_1(p, q) = \sum_{i=1}^m \sum_{j=1}^n p_i q_j a_{ij} = p^\top A q.$$

Thus, we have

$$m = \max_{\sigma_1} \left[\min_{\sigma_2} \pi_1(\sigma_1, \sigma_2) \right] = \max_p \left[\min_q p^\top A q \right],$$
$$M = \min_{\sigma_2} \left[\max_{\sigma_1} \pi_1(\sigma_1, \sigma_2) \right] = \min_q \left[\max_p p^\top A q \right].$$

Thus, we want to show that

$$\min_q \left[\max_p p^\top A q \right] = \max_p \left[\min_q p^\top A q \right].$$

We will show using strong duality. In particular, we will reduce the optimization problem in the maximin case to an LP, and show that the dual of that LP is the optimization problem in the minimax case. By strong duality, we will then show that both these LP's have the same optimal value.

Step 1: Converting the maximin problem into a linear program

To begin, let us consider the maximin problem:

$$\max_p \left(\min_q p^\top A q \right).$$

For a fixed p, the inner minimization problem is given by

$$\begin{aligned} &\min_q \, (p^\top A) q \qquad \text{subject to } \sum_{i=1}^n q_i = 1, \; q_i \geq 0. \\ &= \min_q \Big(q_1 (p^\top A)_1 + q_2 (p^\top A)_2 + \cdots + q_n (p^\top A)_n \Big) \qquad \text{subject to } \sum_{i=1}^n q_i = 1, \; q_i \geq 0. \end{aligned}$$

The objective of this minimization problem is a weighted average of m numbers

$$(p^\top A)_1, \; (p^\top A)_2, \; \ldots, \; (p^\top A)_n,$$

where the weights are given by the probability vector q. When we minimize over all weights, the minimum will be attained when q puts all its weight on the minimum of $\{(p^\top A)_1, (p^\top A)_2, \ldots, (p^\top A)_n\}$. Thus, we obtain for every p,

$$\min_q p^\top A q = \min\{(p^\top A)_1, (p^\top A)_2, \ldots, (p^\top A)_n\}.$$

This implies that

$$\max_p \min_q p^\top A q = \max_p \Big[\min\{(p^\top A)_1, (p^\top A)_2, \ldots, (p^\top A)_n\} \Big].$$

Writing out the constraints on p explicitly, we get the following optimization problem:

$$\begin{aligned} &\max_p \Big[\min\{(p^\top A)_1, (p^\top A)_2, \ldots, (p^\top A)_n\} \Big] \\ &\text{subject to } \sum_{i=1}^m p_i = 1, \qquad p_i \geq 0. \end{aligned}$$

The constraints in the optimization problem are linear, but the objective is not a linear function (it is the minimum of m linear functions, which is in general piecewise linear). So we are not yet at an LP. But there is a standard trick we can use to turn this into a linear program. We introduce an auxiliary variable t such that

$$t \leq \min\{(p^\top A)_1, \ldots, (p^\top A)_n\}$$

and then maximize over t. For fixed p, when we maximize over t with the above constraint, its value will be exactly $\min\{(p^\top A)_1, \ldots, (p^\top A)_n\}$. Thus, we get that the above problem is same as

$$\begin{aligned} \max_{p,t} \quad & t \\ \text{subject to} \quad & t \leq \min\{(p^\top A)_1, (p^\top A)_2, \ldots, (p^\top A)_n\} \\ & \sum_{i=1}^m p_i = 1 \\ & p_i \geq 0, \quad t \text{ free}. \end{aligned}$$

This is same as the following linear program:

$$\begin{aligned}
\max_{p,t} \quad & t \\
\text{s.t.} \quad & t \leq (p^\top A)_1 \\
& t \leq (p^\top A)_2 \\
& \quad \vdots \\
& t \leq (p^\top A)_n \\
& \sum_{i=1}^m p_i = 1 \\
& p_i \geq 0, \quad t \text{ free}
\end{aligned}$$

Writing out the LP explicitly, we get

$$\begin{aligned}
\max_{p,t} \quad & 1 \cdot t + 0 \cdot p_1 + \cdots + 0 \cdot p_m \\
\text{s.t.} \quad & t - a_{11}p_1 - a_{21}p_2 - \cdots - a_{m1}p_m \leq 0 \\
& t - a_{12}p_1 - a_{22}p_2 - \cdots - a_{m2}p_m \leq 0 \\
& \vdots \\
& t - a_{1m}p_1 - a_{2m}p_2 - \cdots - a_{mn}p_m \leq 0 \\
& 0t + p_1 + p_2 + \cdots + p_m = 1 \\
& p_i \geq 0, \qquad t \text{ free.}
\end{aligned} \tag{4}$$

We refer to the linear program (4) as the *maximin LP*.

Step 2: Taking the dual

We will now consider the dual to LP (4). Let q_i be the dual variable associated with the constraint $t \leq (p^\top A)_i$, for $i = 1, 2, \ldots, n$. Also, let s be the dual variable associated with the constraint $\sum_{i=1}^m p_i = 1$. Then, the dual is given by

$$\begin{aligned}
\min_{s, q_1, q_2, \ldots, q_m} \quad & 0 \cdot q_1 + 0 \cdot q_2 + \cdots + 0 \cdot q_n + 1 \cdot s \\
& q_1 + q_2 + q_3 + \cdots + q_n + 0 \cdot s = 1 && \text{(coefficients of } t \text{ in (4))} \\
& -a_{11}q_1 - a_{12}q_2 - \cdots - a_{1n}q_n + s \geq 0 && \text{(coefficients of } p_1 \text{ in (4))} \\
& \vdots \\
& -a_{m1}q_1 - a_{m2}q_2 - \cdots - a_{mn}q_n + s \geq 0 && \text{(coefficients of } p_m \text{ in (4))} \\
& q_i \geq 0, \qquad s \text{ free.}
\end{aligned} \tag{5}$$

This LP can be simplified as

$$
\begin{aligned}
\min_{s,q_1,\ldots,q_n} \quad & s \\
\text{s.t.} \quad & q_1 + \cdots + q_n = 1 \\
& -(Aq)_1 + s \geq 0 \\
& -(Aq)_2 + s \geq 0 \\
& \vdots \\
& -(Aq)_m + s \geq 0 \\
& q_i \geq 0, \qquad s \text{ free.}
\end{aligned}
$$

Rewriting, we get the dual LP (5) as

$$
\begin{aligned}
\min_{s,q} \quad & s \\
\text{s.t.} \quad & \sum_{i=1}^{n} q_i = 1 \\
& s \geq (Aq)_1 \\
& s \geq (Aq)_2 \\
& \vdots \\
& s \geq (Aq)_m \\
& q_i \geq 0, \qquad s \text{ free.}
\end{aligned}
$$

Interpreting the dual linear program as the minimax problem

Now, observe the inequalities in the dual LP together imply that

$$s \geq \max\{(Aq)_1, \ldots, (Aq)_m\}.$$

For a fixed q, when we minimize s over the preceding inequalities, its value is exactly $\max\{(Aq)_1, \ldots, (Aq)_m\}$. Thus the dual LP is same as the optimization problem

$$
\begin{aligned}
&\min_{q} \left[\max\{(Aq)_1, \ldots, (Aq)_m\}\right] \\
\text{subject to } &\sum_{i=1}^{n} q_i = 1, \qquad q_i \geq 0.
\end{aligned}
$$

The inner maximization is same as maximizing over all weighted averages

$$p_1(Aq)_1 + p_2(Aq)_2 + \cdots + p_m(Aq)_m = p^\top Aq$$

where the weights satisfy

$$\sum_{i=1}^{m} p_i = 1, \qquad p_i \geq 0.$$

Thus the dual LP is equivalent to

$$\min_{q} \left[\max_{p} p^\top Aq\right],$$

which is the nothing but minimax problem.

Step 4: Applying strong duality

Thus, we have seen that the maximin problem is an LP (4) whose dual, given by (5) is the minimax problem. Because of this reason, we refer to the dual LP (5) as the *minimax LP*.

Now, note that the maximin LP (4) is feasible. This can be seen from the fact that for $p = [1\ \ 0\ \ \ldots 0]^\top$ and any $t < \min_{i,j} a_{ij}$, all the constraints in the maximin LP are satisfied. Similarly, the minimax LP (5) is feasible: set $q = [1\ \ 0\ \ \ldots 0]^\top$ and any $s > \max_{ij} a_{ij}$ to verify that all constraints in the minimax LP are satisfied.

By strong duality, we then conclude that both linear programs are bounded and their optimal values are equal. If (t^*, p^*) is an optimal solution for the maximin LP and (s^*, q^*) is an optimal solution for minimax LP, then $t^* = s^*$.

However, by definition, $t^* = m$, and $s^* = M$. Hence we conclude $M = m$. In other words,

$$\min_q \max_p p^\top A q = \max_p \min_q p^\top A q.$$

Value of the game

Thus, we have shown that for general finite zero-sum games, the maximin value m is same as the minimax value M. We call this common value as the value of the game (for player 1), and denote it by $\mathsf{val}(A)$, where A is the payoff matrix for player 1. Note that the value of the same game for player 2 is just $-\mathsf{val}(A)$.

Maximin strategies and relation to Nash equilibrium

Observe that the maximin LP (4) can be written succinctly in matrix notation as

$$\begin{aligned} \max_{t,p} \quad & t \\ \text{subject to,} \quad & t\mathbf{1} - A^\top p \leq \mathbf{0} \\ & \mathbf{1}^\top p = 1 \\ & p \geq 0, t \text{ free.} \end{aligned}$$

(Here $\mathbf{1}$ is the vector of all ones, and $\mathbf{0}$ is the vector of all zeros.) By finding an optimal solution (t^*, p^*) to this linear program, we obtain the maximin strategy p^* for player 1.

Similarly, the minimax LP (5) can be written in matrix notation as

$$\begin{aligned} \min_{s,q} \quad & s \\ \text{subject to,}\ & s\mathbf{1} - Aq \geq \mathbf{0} \\ & \mathbf{1}^\top q = 1 \\ & q \geq 0, s \text{ free.} \end{aligned}$$

This optimal solution (s^*, q^*) to this linear program yields the maximin strategy for player 2.

Furthermore, by strong duality, we know that the optimal solutions (t^*, p^*) and (s^*, q^*) satisfy the *complementary slackness* conditions:

$$\begin{aligned} &\text{if } s^* > (Aq^*)_i \text{ then } p_i^* = 0 \\ &\text{if } p_i^* > 0 \text{ then } s^* = (Aq^*)_i, \end{aligned}$$

and

$$\text{if } t^* < (A^\top p^*)_j \text{ then } q_j^* = 0$$
$$\text{if } q_j^* > 0 \text{ then } t^* = (A^\top p^*)_j.$$

Now, observe that $(Aq^*)_i = \pi_1(i, q^*) =$ payoff player 1 gets when player 1 plays pure strategy i and player 2 plays q^*. Similarly, $(A^\top p^*)_j = \pi_2(p^*, j) =$ payoff player 2 gets when player 2 plays pure strategy j and player 1 plays p^*. Using these, the first set of complementary slackness conditions become:

$$\text{if } s^* > \pi_1(i, q^*) \text{ then } p_i^* = 0$$
$$\text{if } p_i^* > 0 \text{ then } s^* = \pi_1(i, q^*).$$

Thus, player 1's maximin strategy p^* puts positive probability only on those pure strategy i that maximize $\pi_1(i, q^*)$. This in turn implies that p^* maximizes $\pi_1(p, q^*)$ over all mixed strategies p. Thus, we obtain

$$p^* \text{ is a best response to } q^*, \text{ i.e., } p^* \in BR_1(q^*).$$

By a similar argument using the second set of complementary slackness conditions, we can also show that

$$q^* \in BR_2(p^*).$$

Together, this implies that (p^*, q^*) is a Nash Equilibrium for the zero-sum game.

Thus, if p^* and q^* solve the maximin and the dual minimax LP respectively, then (p^*, q^*) is a Nash equilibrium. Furthermore, the payoff in the Nash Equilibrium is exactly $\mathsf{val}(A)$, the value of the game.

Observe that this above approach gives us a *constructive* method to find a Nash equilibrium in zero sum games: solve the maximin LP (4) and its dual the minimax LP (5) to find the maximin strategies p^* and q^*, and the strategy profile (p^*, q^*) will be an NE. Since there exists efficient algorithms for solving a LP, finding a Nash equilibrium for a zero-sum games is a computationally easy.

Example (contd.)**.** Let us revisit our running example of the zero-sum game:

	L	R
U	$(3, -3)$	$(5, -5)$
D	$(4, -4)$	$(1, -1)$

The payoff matrix for player 1 is given by

$$A = \begin{bmatrix} 3 & 5 \\ 4 & 1 \end{bmatrix}$$

Thus, the maximin LP (4) for this game becomes

$$\begin{aligned} \max_{t, p_1, p_2} \quad & t \\ \text{subject to,} \quad & t - 3p_1 - 4p_2 \leq 0 \\ & t - 5p_1 - p_2 \leq 0 \\ & p_1 + p_2 = 1 \\ & p_1 \geq 0, p_2 \geq 0, t \text{ free} . \end{aligned}$$

Exercise: verify that the optimal solution to the above LP is given by $p_1^* = \frac{3}{5}$, $p_2^* = \frac{2}{5}$ and $t^* = \frac{17}{5}$.

Similarly, the minimax LP (5) for this game is given by

$$\begin{aligned} \min_{a,q_1,q_2} \quad & s \\ \text{subject to,} \quad & s - 3q_1 - 5p_2 \geq 0 \\ & s - 4p_1 - p_2 \geq 0 \\ & q_1 + q_2 = 1 \\ & q_1 \geq 0, q_2 \geq 0, s \text{ free} . \end{aligned}$$

Exercise: verify that the optimal solution to the above LP is given by $q_1^* = \frac{4}{5}$, $q_2^* = \frac{1}{5}$ and $s^* = \frac{17}{5}$.

Thus, we conclude that the maximin strategy for player 1 is $p^* = [\frac{3}{5} \;\; \frac{2}{5}]^\top$, the maximin strategy for player 2 is $q^* = [\frac{4}{5} \;\; \frac{1}{5}]^\top$. Furthermore, the value of this game for player 1 is $\mathsf{val}(A) = \frac{17}{5}$. Finally, (p^*, q^*) is the mixed NE. $\square$