

Minimax theorems

Scribe: Zachary Lau
Presenter: Jeffrey Dawson

October 14, 2026

1 Games

We discuss minimax theorems in the context of games and game theory. First we define a game.

Definition 1.1 (Game). A game consists of

- A set of players $\{1, \dots, k\}$
- For each player, a set of actions S_i
- For each player, a payout function $P_i : \prod_{j=1}^k S_j \rightarrow \mathbb{R}$

The actions reflect choices that the players can make about how they play the game, while the payout function reflects the rewards that each player will receive given their own and others' actions. Decisions may be made sequentially or simultaneously depending on the type of game.

Example 1.2 (Prisoners' Dilemma). In the prisoners' dilemma, two accomplices are captured and each separately given the choice to confess or deny. If they both deny, they will be given a 1-year sentence each; if they both confess they will be given a 5-year sentence each. If one confesses and the other denies, the confessor will get off free, while the denier will get a 10-year sentence. In this scenario, the actions are to confess or deny. The payouts are shown in the following table, where the payout for the row player is shown first (payouts are the negative of the sentence length).

	Confess	Deny
Confess	$(-5, -5)$	$(0, -10)$
Deny	$(-10, 0)$	$(-1, -1)$

In the prisoners' dilemma, each player in isolation benefits by confessing, but if both confess this leads to a sub-optimal outcome. Therefore, prisoners can benefit by cooperating to both deny. However, we will consider a class of games where cooperation is not always possible. In particular, in zero-sum games, one player's gain directly causes another player's loss.

Definition 1.3 (Zero-sum game). A two-player zero-sum game is a game in which

- $k = 2$
- $P_1 + P_2 \equiv 0$

Example 1.4 (Rock paper scissors). We consider a rock paper scissors game where the winner gets a payout of 1, loser gets a payout of -1 and a tie yields 0 payout for both players. The payouts for the column player are shown in the following table. This is a zero-sum game because the sum of payouts is zero in all scenarios.

	Rock	Paper	Scissors
Rock	0	1	-1
Paper	-1	0	1
Scissors	1	-1	0

2 Optimal strategies

We might ask, what is the optimal strategy in the rock paper scissors game? If we consider just deterministic strategies, we may find ourselves bested by an opponent who knows our strategy; if we always pick rock, they will always pick paper. Another strategy is to randomly choose one of rock, paper or scissors. A strategy which is always the same is known as a pure strategy. A strategy which is not pure is called mixed.

Definition 2.1 (Pure/mixed strategy). A pure strategy is one where we choose a specific action 100% of the time. A strategy which is not pure is called “mixed”, and consists of probabilistically choosing an action.

2.1 Payout matrices

We now introduce a useful method of looking at payouts in the presence of mixed strategies. Given a strategy which chooses the i th action with probability x_i , we may represent this strategy by the vector

$$x = [x_1, \dots, x_n]^\top \quad (1)$$

To be a valid strategy we require $\forall i \ x_i \geq 0$ and $1^\top x = 1$. We denote the payouts for the column player as

$$A = \begin{bmatrix} a_{11} & \cdots & a_{1m} \\ \vdots & \ddots & \vdots \\ a_{n1} & \cdots & a_{nm} \end{bmatrix} \quad (2)$$

Here a_{ij} represents the payout when the row player picks action i and the column player picks action j . We use x to denote the strategy of the row player and y the strategy of the column player. Furthermore, because players don't know the final action taken by the other player, we assume independence between the row and column players' sources of randomness. In this scenario, the expected reward for the column player is denoted

$$\sum_{i=1}^n \sum_{j=1}^m a_{ij} x_i y_j = x^\top A y \quad (3)$$

Due to the zero-sum property, the payout for the row player is the negative of this, i.e., $-x^\top A y$. In other words, the row player aims to minimize $x^\top A y$ while the column player aims to maximize $x^\top A y$.

2.2 Von Neumann minimax

We now consider optimal strategies in sequential games. For example, the row player may choose a strategy, then the column player may choose a strategy, with knowledge of the row player's (potentially mixed) strategy. We may expect going second in such a scenario to be an advantage. Perhaps counterintuitively, Von Neumann's minimax theorem tells us that this is not the case.

Theorem 2.2 (Von Neumann's Minimax). *In a two-player zero-sum game payout matrix A for the column player, and row and column strategies x and y ,*

$$\min_x \max_y x^\top A y = \max_y \min_x x^\top A y \quad (4)$$

In the statement of the above theorem, the left-hand side is the column player's expected payout under optimal play when the row player (x) announces their strategy first, while the right-hand side corresponds to the column player's expected payout when the column player (y) announces their strategy first.

Proof. We begin by showing that

$$\min_x \max_y x^\top A y \geq \max_y \min_x x^\top A y \quad (5)$$

Intuitively we expect going second to advantage the y player. We mathematically demonstrate this. Fix y . Then from the definition of $\min$

$$\forall x \quad x^\top Ay \geq \min_{x'} x'^\top Ay \quad (6)$$

This is true for any fixed y so

$$\forall y \forall x \quad x^\top Ay \geq \min_{x'} x'^\top Ay \quad (7)$$

$$\forall x \quad \max_y x^\top Ay \geq \max_y \min_{x'} x'^\top Ay \quad (8)$$

$$\min_x \max_y x^\top Ay \geq \max_y \min_x x^\top Ay \quad (9)$$

We now show that

$$\min_x \max_y x^\top Ay \leq \max_y \min_x x^\top Ay \quad (10)$$

The strategy is to show that for any $t > \max_y \min_x x^\top Ay$, we also have $t > \min_x \max_y x^\top Ay$. In other words, there is no gap in Equation 9 (such a gap would be called a duality gap). Let $t > \max_y \min_x x^\top Ay$. We first note that for a given y , an optimal x strategy is the pure strategy which picks a coordinate i to minimize $(Ay)_i$. I.e.,

$$\min_x x^\top Ay = \min_i (Ay)_i \quad (11)$$

We show this mathematically. For a fixed y we first show that the x player can do no better than $\min_i (Ay)_i$.

$$\forall x \quad x^\top Ay = \sum_i x_i (Ay)_i \quad (12)$$

$$\geq \sum_i x_i \min_j (Ay)_j \quad (13)$$

$$= \min_j (Ay)_j \quad (14)$$

The second line follows because each $x_i \geq 0$, and the last because $\sum_i x_i = 1$. We thus see that

$$t > \max_y \min_i (Ay)_i \quad (15)$$

$$\forall y \quad t > \min_i (Ay)_i \quad (16)$$

$$\nexists y \quad t \leq \min_i (Ay)_i \quad (17)$$

Denote by Q_m the set of probability vectors in m dimensions (and Q_n analogously). Equation 17 is equivalent to the statement on sets

$$\{Ay | y \in Q_m\} \cap \{z | \min_i z_i \geq t\} = \emptyset \quad (18)$$

We notice that both sets are convex sets. Q_m is the convex hull of the finite set of pure strategies, and so its image under an affine transformation is compact and convex. The second set is the intersection of closed half-spaces, and hence also closed and convex. We may apply the separating hyperplane theorem. Denote $C_1 = \{Ay | y \in Q_m\}$ and $C_2 = \{z | \min_i z_i \geq t\}$. By the separating hyperplane theorem, there exists a separating hyperplane, i.e., there is a direction v such that

$$\max_{u \in C_1} \langle u, v \rangle < \min_{z \in C_2} \langle z, v \rangle \quad (19)$$

We first assert that v has to have all nonnegative entries. Heuristically, if v is negative in one direction, then moving in that direction it will eventually cut into C_2 . We formally show this. Suppose that $v_i < 0$. Take any point $p_1 \in C_1$. Denote $w = [t, \dots, t]^\top \in C_2$. We consider the point

$$p_2 = w + e_i \frac{\langle v, w \rangle - \langle v, p_1 \rangle}{-v_i} \quad (20)$$

If $\langle v, w \rangle - \langle v, p_1 \rangle \leq 0$ then v does not define a separating hyperplane, since it doesn't separate $w \in C_2$ and $p_1 \in C_1$. We thus consider only $\langle v, w \rangle - \langle v, p_1 \rangle > 0$. In this case first note that $p_2 \in C_2$ since all coordinates are unchanged except the i coordinate which is increased. Furthermore

$$\langle v, p_2 \rangle = \langle v, w \rangle + v_i \left(\frac{\langle v, w \rangle - \langle v, p_1 \rangle}{-v_i} \right) \quad (21)$$

$$= \langle v, w \rangle - (\langle v, w \rangle - \langle v, p_1 \rangle) \quad (22)$$

$$= \langle v, p_1 \rangle \quad (23)$$

In this case v does not separate $p_1 \in C_1$ and $p_2 \in C_2$. We conclude that v cannot have negative entries if it defines a valid separating hyperplane. The magnitude of v is not important, only its direction. We thus scale it by its L1 norm. We also note that $\|v\|_1 \neq 0$ since the zero-vector does not define a hyperplane. Denote

$$x = \frac{v}{\|v\|_1} \quad (24)$$

Since x has nonnegative entries and $\|x\|_1 = 1$, x is a probability vector. Beyond the fact that x is a probability vector we don't really care much about its properties; its existence is sufficient to continue. We have

$$\exists x \in Q_n \quad \max_y \langle x, Ay \rangle < \langle x, w \rangle \quad (25)$$

Since x is a probability vector and $\forall i \ w_i = t$, then $\langle x, w \rangle = t$, so

$$\exists x \in Q_n \quad \max_y \langle x, Ay \rangle < t \quad (26)$$

$$\min_x \max_y x^\top Ay < t \quad (27)$$

Since this holds for every $t > \max_y \min_x x^\top Ay$, letting t decrease to $\max_y \min_x x^\top Ay$ shows that $\max_y \min_x x^\top Ay \geq \min_x \max_y x^\top Ay$. We have shown both directions of the inequality and so have equality. This completes the proof. $\square$

3 Nash equilibria and saddle points

One of the more famous concepts in game theory is the Nash equilibrium.

Definition 3.1 (Nash equilibrium). In a k player game, a Nash equilibrium consists of a strategy for each player such that given the other players' strategies no player can improve their expected payout by changing only their own strategy.

This concept is commonly applied in areas such as economics. In the specific case of a two-player zero-sum game, the Nash equilibrium is the same thing as a "saddle point".

Definition 3.2 (Saddle point). In a two-player zero-sum game with row strategy x , column strategy y and column player payout $x^\top Ay$, we define a saddle point as a point x_*, y_* at which y_* maximizes the function $g_y(y) = x_*^\top Ay$ and x_* minimizes the function $g_x(x) = x^\top Ay_*$.

We will use the terms minimax and maximin to refer to the points that achieve the minimax and maximin optima, respectively.

Definition 3.3 (Minimax, maximin). In a two-player zero-sum game, a minimax point is a point such that

$$x_* = \arg \min_x \max_y x^\top Ay \quad (28)$$

A maximin point is a point such that

$$y_* = \arg \max_y \min_x x^\top Ay \quad (29)$$

It turns out that saddle points are closely related to minimax and maximin points.

Theorem 3.4. *In a two-player zero-sum game with mixed strategies, if a point (x_*, y_*) is a saddle point, then x_* is a minimax point and y_* is a maximin point. Conversely, if x_* is a minimax point and y_* is a maximin point, then (x_*, y_*) is a saddle point.*

Proof. First we show that a saddle point is necessarily both a minimax and a maximin point. We start with minimax. Suppose that (x_*, y_*) is a saddle point. Then from the saddle point criterion on x_* we know that

$$\forall x \quad \max_y x^\top Ay \geq x^\top Ay_* \geq x_*^\top Ay_* \quad (30)$$

$$\min_x \max_y x^\top Ay \geq x_*^\top Ay_* \quad (31)$$

Note that $y_* = \arg \max_y x_*^\top Ay$, so equality is achieved in the last line at the point $x = x_*$. Thus x_* is a minimax point. Likewise, using the saddle point criterion on y_* we find

$$\forall y \quad \min_x x^\top Ay \leq x_*^\top Ay \leq x_*^\top Ay_* \quad (32)$$

$$\max_y \min_x x^\top Ay \leq x_*^\top Ay_* \quad (33)$$

We get equality at $y = y_*$, since $x_* = \arg \min_x x^\top Ay_*$, and therefore y_* is a maximin point.

Conversely, suppose x_* is a minimax point and y_* is a maximin point. Then we find that

$$\max_y \min_x x^\top Ay = \min_x x^\top Ay_* \quad (34)$$

$$\leq x_*^\top Ay_* \quad (35)$$

$$\leq \max_y x_*^\top Ay = \min_y \max_x x^\top Ay \quad (36)$$

By Von Neumann's minimax theorem, the first and last terms are equal, so the entire chain of inequalities becomes a chain of equalities. In particular,

$$\min_x x^\top Ay_* = x_*^\top Ay_* = \max_y x_*^\top Ay \quad (37)$$

From this we see that the point (x_*, y_*) is necessarily a saddle point; changing x cannot decrease $x^\top Ay_*$ while changing y cannot increase $x_*^\top Ay$. $\square$

Up until now we have only considered mixed strategies. If we only allow pure strategies, the forward direction of Theorem 3.4 still holds. That is, if (x_*, y_*) is a saddle point, then x_* is a minimax point and y_* is a maximin point. This may be useful when trying to find minimax or maximin points without doing nested optimization. However, Von Neumann's minimax theorem doesn't always hold with just pure strategies, and the converse of Theorem 3.4 can fail. For example, consider the rock-paper-scissors game with just pure strategies. Whichever pair of pure strategies is played, one of the players can improve by switching, so there are no pure Nash equilibria or saddle points. In particular, $\min_x \max_y x^\top Ay = 1$ while $\max_y \min_x x^\top Ay = -1$, since a player who moves second against a known pure strategy always wins. Every pure strategy is both a minimax and a maximin point, but no pair of them is a saddle point. In general, if we restrict ourselves to just pure strategies, the following proposition holds.

Proposition 3.5. *In a two-player zero-sum game with only pure strategies, a saddle point exists if and only if*

$$\min_i \max_j a_{ij} = \max_j \min_i a_{ij} \quad (38)$$

Proof. The proof is almost identical to the proof of Theorem 3.4. For the “if” direction, we follow the proof of the converse, but instead of using Von Neumann's minimax theorem to assert equality, we use the assumption $\min_i \max_j a_{ij} = \max_j \min_i a_{ij}$. For the “only if” direction, the proof of the forward direction of Theorem 3.4 shows that if (i_*, j_*) is a saddle point then

$$\min_i \max_j a_{ij} = \max_j \min_i a_{ij} = a_{i_* j_*} \quad (39)$$

$\square$