

18.7 Matrix Games and the Minimax Theorem

Chapter 17 was built on a min–max formula: the Courant–Fischer theorem describes each eigenvalue of a Hermitian matrix as a minimum of maxima over subspaces. This section proves a second min–max theorem that looks similar and works for a completely different reason. It concerns a game in which two players choose simultaneously, and it says that if the players are allowed to **randomize**, it does not matter which of them has to reveal a plan first. The theorem is due to von Neumann, and with linear programming duality in hand (§06) it takes a page to prove.

The two theorems are easy to confuse, and a reader who has just finished Chapter 17 will be tempted to. So the section also says precisely how they differ. Courant–Fischer is a statement about **subspaces**, and its equality comes from a **dimension count**. The minimax theorem is a statement about **simplices**, and its equality comes from **convexity**.

Throughout, the field is $\mathbb{R}$, $\mathbf{1}$ is the all-ones vector of whatever length the context requires, and $\mathbf{J} = \mathbf{1}\mathbf{1}^\top$ is the all-ones matrix. Inequalities between vectors are entrywise, as in §05.

18.7.1 Matrix games

Two players, Row and Column, play a game described by a real matrix $\mathbf{A} \in M_{m \times n}(\mathbb{R})$, the **payoff matrix**. Row secretly chooses a row i , Column secretly chooses a column j , both reveal their choices, and Column pays Row the amount a_{ij} . A negative entry means that Row pays Column. Whatever one player wins, the other loses, so the game is called **zero-sum**. Row wants a_{ij} large and Column wants it small. A choice of a row, or of a column, is a **pure strategy**.

Suppose Row must announce a row in advance. Column will then answer with the smallest entry in that row, so the best Row can guarantee is

$$\max_i \min_j a_{ij}.$$

If instead Column must announce first, the best Column can guarantee is to pay at most $\min_j \max_i a_{ij}$. Announcing first is never an advantage, and that holds for any payoff function whatsoever.

Lemma 18.7.1. Max–Min Never Exceeds Min–Max

Let X and Y be sets and $f: X \times Y \rightarrow \mathbb{R}$ a function for which the extrema below exist. Then

$$\max_{x \in X} \min_{y \in Y} f(x, y) \leq \min_{y \in Y} \max_{x \in X} f(x, y).$$

Proof. Let $x_0 \in X$ and $y_0 \in Y$. Then

$$\min_{y \in Y} f(x_0, y) \leq f(x_0, y_0) \leq \max_{x \in X} f(x, y_0).$$

The left side does not depend on y_0 and the right side does not depend on x_0 . Taking the maximum over x_0 on the left and then the minimum over y_0 on the right proves the lemma. ■

The inequality can be strict. In the game with

$$\mathbf{A} = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix},$$

Row wins when the two choices match and loses when they differ. Every row contains a -1 , so $\max_i \min_j a_{ij} = -1$. Every column contains a 1 , so $\min_j \max_i a_{ij} = 1$. Whoever announces first loses a unit, and there is a gap of 2 between what the players can guarantee.

Some games have no gap. In

$$\mathbf{A} = \begin{pmatrix} 3 & 1 & 4 \\ 2 & 0 & 1 \end{pmatrix}$$

the row minima are 1 and 0, whose maximum is 1, and the column maxima are 3, 1, 4, whose minimum is also 1. The entry $a_{12} = 1$ is the smallest in its row and the largest in its column. Such an entry is a **saddle point**. If Row plays row 1 and Column plays column 2, neither can gain by changing unilaterally, and the game is settled in pure strategies. The first game has no saddle point, and in that game the natural advice is: do not be predictable.

18.7.2 Mixed strategies

The way to be unpredictable is to randomize. Row chooses row i with probability x_i , and Column chooses column j with probability y_j , independently. The probabilities of the m rows form a vector with nonnegative entries summing to 1, that is, a point of the standard simplex of Definition 18.1.2,

$$\Delta_m = \{\mathbf{x} \in \mathbb{R}^m : \mathbf{x} \geq \mathbf{0} \text{ and } x_1 + \cdots + x_m = 1\}.$$

A *mixed strategy* is a probability distribution over the pure strategies, and the *simplex* is the set of all of them.

Definition 18.7.1. Mixed Strategy

In the game with payoff matrix $\mathbf{A} \in M_{m \times n}(\mathbb{R})$, a **mixed strategy** for Row is a vector $\mathbf{x} \in \Delta_m$, a mixed strategy for Column is a vector $\mathbf{y} \in \Delta_n$, and the **expected payoff** is

$$\mathbf{x}^\top \mathbf{A} \mathbf{y} = \sum_{i=1}^m \sum_{j=1}^n x_i a_{ij} y_j.$$

In words: a mixed strategy assigns to each pure strategy a weight that is **nonnegative**, and the weights **sum to exactly** 1. The payoff a_{ij} occurs with probability $x_i y_j$, since the two choices are independent, so $\mathbf{x}^\top \mathbf{A} \mathbf{y}$ is the average payoff over many plays. The pure strategy “row i ” is the mixed strategy $\mathbf{e}_i$, a vertex of the simplex, and $\mathbf{e}_i^\top \mathbf{A} \mathbf{e}_j = a_{ij}$. So mixed strategies extend pure ones, and the expected payoff extends the payoff.

Examples. Δ_2 is the segment from $(1, 0)$ to $(0, 1)$ in the plane, whose points $(p, 1 - p)$ with $0 \leq p \leq 1$ mean “row 1 with probability p ”. Δ_3 is the triangle with vertices $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ in $\mathbb{R}^3$, and its center $\frac{1}{3}\mathbf{1}$ is the uniform strategy. The degenerate case is $\Delta_1 = \{1\}$: a player with only one pure strategy has no choice, and the game is a single row or a single column.

Nonexamples by minimal change. $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ has nonnegative entries, but they sum to $\frac{3}{2}$, so it is not a probability distribution. $(\frac{3}{2}, -\frac{1}{2})$ sums to 1, but it has a negative entry. It lies on the line through Δ_2 , outside the segment.

Why this definition. Δ_n is exactly the set of convex combinations of $\mathbf{e}_1, \dots, \mathbf{e}_n$, since $\mathbf{x} = \sum x_j \mathbf{e}_j$. So it is the convex hull of the pure strategies, by Theorem 18.1.3, and in particular it is a convex set (Definition 18.1.1). That convexity is what makes the theorem below true. The convention of Definition 18.1.2, n weights for n pure strategies, is the one a game needs. The payoff $\mathbf{x}^\top \mathbf{A} \mathbf{y}$ is linear in $\mathbf{y}$ for fixed $\mathbf{x}$, and linear in $\mathbf{x}$ for fixed $\mathbf{y}$, so a mixture of strategies earns the same mixture of payoffs.

Warning. **A mixed strategy is not a compromise move.** Playing $(\frac{1}{2}, \frac{1}{2})$ in the matching game does not mean playing “half of row 1 and half of row 2” in a single round. Each round Row plays one actual row, chosen at random, and $\mathbf{x}^\top \mathbf{A} \mathbf{y}$ is an average over rounds. In the matching game with $\mathbf{x} = \mathbf{y} = (\frac{1}{2}, \frac{1}{2})$, the expected payoff is 0, yet every single round pays +1 or -1, never 0.

Against a fixed mixed strategy, the opponent never needs to randomize. That observation turns every “min over a simplex” into a min over finitely many numbers.

Lemma 18.7.2. A Pure Best Response Exists

Let $\mathbf{A} \in M_{m \times n}(\mathbb{R})$. For every $\mathbf{x} \in \Delta_m$,

$$\min_{\mathbf{y} \in \Delta_n} \mathbf{x}^\top \mathbf{A} \mathbf{y} = \min_{1 \leq j \leq n} (\mathbf{A}^\top \mathbf{x})_j,$$

and the minimum on the left is attained at $\mathbf{y} = \mathbf{e}_j$ for any j attaining the minimum on the right. Likewise, for every $\mathbf{y} \in \Delta_n$, $\max_{\mathbf{x} \in \Delta_m} \mathbf{x}^\top \mathbf{A} \mathbf{y} = \max_{1 \leq i \leq m} (\mathbf{A} \mathbf{y})_i$, attained at a vertex $\mathbf{e}_i$.

Idea. For fixed $\mathbf{x}$, the payoff $\mathbf{x}^\top \mathbf{A} \mathbf{y}$ is an average of the entries of $\mathbf{A}^\top \mathbf{x}$ with weights y_j , and an average is never below the smallest entry, which the vertex $\mathbf{e}_j$ attains.

Proof. Write $\mathbf{w} = \mathbf{A}^\top \mathbf{x} \in \mathbb{R}^n$ and $\mu = \min_j w_j$, a minimum of finitely many numbers. For $\mathbf{y} \in \Delta_n$,

$$\mathbf{x}^\top \mathbf{A} \mathbf{y} = \mathbf{w}^\top \mathbf{y} = \sum_{j=1}^n y_j w_j \geq \sum_{j=1}^n y_j \mu = \mu,$$

because every $y_j \geq 0$ and $\sum_j y_j = 1$. If $w_{j_0} = \mu$, then $\mathbf{e}_{j_0} \in \Delta_n$ gives $\mathbf{x}^\top \mathbf{A} \mathbf{e}_{j_0} = w_{j_0} = \mu$. So μ is a value of $\mathbf{y} \mapsto \mathbf{x}^\top \mathbf{A} \mathbf{y}$ on Δ_n and a lower bound for all its values, hence its minimum. The second statement is the same argument applied to $\mathbf{A} \mathbf{y}$ with the inequality reversed. ■

So once Row fixes $\mathbf{x}$, the amount Row can **guarantee** is

$$g(\mathbf{x}) = \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^\top \mathbf{A} \mathbf{y} = \min_j (\mathbf{A}^\top \mathbf{x})_j,$$

the smallest entry of $\mathbf{A}^\top \mathbf{x}$. Column, in the same way, guarantees by the second half of the lemma to pay at most $h(\mathbf{y}) = \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^\top \mathbf{A} \mathbf{y} = \max_i (\mathbf{A} \mathbf{y})_i$. In the matching game, $\mathbf{x} = (\frac{1}{2}, \frac{1}{2})$ gives $\mathbf{A}^\top \mathbf{x} = (0, 0)$ and $g(\mathbf{x}) = 0$, and the same holds for Column's h . With mixing, both players

can guarantee 0, and the gap of 2 has closed. The theorem says this always happens.

18.7.3 The minimax theorem

Theorem 18.7.3. Von Neumann's Minimax Theorem

Let $\mathbf{A} \in M_{m \times n}(\mathbb{R})$. There are $\mathbf{x}^* \in \Delta_m$, $\mathbf{y}^* \in \Delta_n$, and $v \in \mathbb{R}$ such that

$$\mathbf{x}^\top \mathbf{A} \mathbf{y}^* \leq v \leq \mathbf{x}^{*\top} \mathbf{A} \mathbf{y} \quad \text{for all } \mathbf{x} \in \Delta_m, \mathbf{y} \in \Delta_n. \quad (18.7.1)$$

Consequently

$$\max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^\top \mathbf{A} \mathbf{y} = v = \min_{\mathbf{y} \in \Delta_n} \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^\top \mathbf{A} \mathbf{y},$$

where the outer maximum is attained at $\mathbf{x}^*$ and the outer minimum at $\mathbf{y}^*$.

Idea. Row wants the $\mathbf{x} \in \Delta_m$ that maximizes the smallest entry of $\mathbf{A}^\top \mathbf{x}$. If every entry of $\mathbf{A}$ is positive, this can be rescaled into a linear program. Put $\mathbf{p} = \mathbf{x}/v$, where v is the guaranteed amount. Then “ $\mathbf{A}^\top \mathbf{x} \geq v \mathbf{1}$ with $\mathbf{1}^\top \mathbf{x} = 1$, v as large as possible” becomes “ $\mathbf{A}^\top \mathbf{p} \geq \mathbf{1}$, $\mathbf{p} \geq \mathbf{0}$, $\mathbf{1}^\top \mathbf{p}$ as small as possible”. Column's problem becomes the dual of that program, by the same substitution. Strong duality (§06) says the two optimal values agree, and that common value is $1/v$. The plan:

① shift the entries by a constant to make them positive, which changes every payoff by the same constant; ② write down the two linear programs, check that both are feasible, and apply Theorem 18.6.3; ③ rescale the optimal solutions into strategies and read off Eq. 18.7.1.

Proof. Step 1. We may assume every $a_{ij} > 0$. Choose $c \in \mathbb{R}$ with $a_{ij} + c > 0$ for all i, j , and put $\mathbf{B} = \mathbf{A} + c\mathbf{J}$. For $\mathbf{x} \in \Delta_m$ and $\mathbf{y} \in \Delta_n$,

$$\mathbf{x}^\top \mathbf{J} \mathbf{y} = (\mathbf{x}^\top \mathbf{1})(\mathbf{1}^\top \mathbf{y}) = 1, \quad \text{so} \quad \mathbf{x}^\top \mathbf{B} \mathbf{y} = \mathbf{x}^\top \mathbf{A} \mathbf{y} + c.$$

Hence if $\mathbf{x}^*, \mathbf{y}^*, w$ satisfy Eq. 18.7.1 for $\mathbf{B}$, then $\mathbf{x}^*, \mathbf{y}^*, w - c$ satisfy it for $\mathbf{A}$. So it suffices to prove Eq. 18.7.1 when every entry of $\mathbf{A}$ is positive, and we assume this from now on. Let $\mu = \min_{i,j} a_{ij} > 0$.

Step 2. Two linear programs. Consider

$$(P) \quad \text{maximize } \mathbf{1}^\top \mathbf{q} \quad \text{subject to } \mathbf{A} \mathbf{q} \leq \mathbf{1}, \mathbf{q} \geq \mathbf{0} \quad (\mathbf{q} \in \mathbb{R}^n),$$

$$(D) \quad \text{minimize } \mathbf{1}^\top \mathbf{p} \quad \text{subject to } \mathbf{A}^\top \mathbf{p} \geq \mathbf{1}, \mathbf{p} \geq \mathbf{0} \quad (\mathbf{p} \in \mathbb{R}^m).$$

(P) is a linear program in the inequality form of Definition 18.6.1, with data $(\mathbf{A}, \mathbf{1}, \mathbf{1})$, and (D) is its dual (Definition 18.6.2). (P) is feasible, with $\mathbf{q} = \mathbf{0}$. (D) is feasible, with $\mathbf{p} = \mu^{-1} \mathbf{1}$: the j -th entry of $\mathbf{A}^\top \mathbf{p}$ is $\mu^{-1} \sum_i a_{ij} \geq \mu^{-1} a_{1j} \geq 1$. By Theorem 18.6.3 (a), there are optimal $\mathbf{q}^*$ and $\mathbf{p}^*$ with equal values; call the common value ω , so that

$$\mathbf{1}^\top \mathbf{q}^* = \mathbf{1}^\top \mathbf{p}^* = \omega.$$

Moreover $\omega > 0$. Since $\mathbf{A}^\top \mathbf{p}^* \geq \mathbf{1}$, we have $\mathbf{p}^* \neq \mathbf{0}$, and a nonzero vector with nonnegative entries has a positive entry sum.

Step 3. Strategies from the programs. Put

$$\mathbf{x}^* = \omega^{-1} \mathbf{p}^*, \quad \mathbf{y}^* = \omega^{-1} \mathbf{q}^*, \quad v = \omega^{-1}.$$

Then $\mathbf{x}^* \geq \mathbf{0}$ with $\mathbf{1}^\top \mathbf{x}^* = \omega^{-1} \omega = 1$, so $\mathbf{x}^* \in \Delta_m$, and likewise $\mathbf{y}^* \in \Delta_n$. The constraints of the two programs give

$$\mathbf{A}^\top \mathbf{x}^* = \omega^{-1} \mathbf{A}^\top \mathbf{p}^* \geq v \mathbf{1}, \quad \mathbf{A} \mathbf{y}^* = \omega^{-1} \mathbf{A} \mathbf{q}^* \leq v \mathbf{1}.$$

For every $\mathbf{y} \in \Delta_n$, therefore, $\mathbf{x}^{*\top} \mathbf{A} \mathbf{y} = \sum_j y_j (\mathbf{A}^\top \mathbf{x}^*)_j \geq v \sum_j y_j = v$. For every $\mathbf{x} \in \Delta_m$, $\mathbf{x}^\top \mathbf{A} \mathbf{y}^* = \sum_i x_i (\mathbf{A} \mathbf{y}^*)_i \leq v$. This is [Eq. 18.7.1](#).

Step 4. The two extrema. For $\mathbf{x} \in \Delta_m$ let $g(\mathbf{x}) = \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^\top \mathbf{A} \mathbf{y}$, which exists by [Lemma 18.7.2](#). By [Eq. 18.7.1](#), $g(\mathbf{x}^*) \geq v$, since every value of $\mathbf{y} \mapsto \mathbf{x}^{*\top} \mathbf{A} \mathbf{y}$ is at least v . And $g(\mathbf{x}) \leq \mathbf{x}^\top \mathbf{A} \mathbf{y}^* \leq v$ for every $\mathbf{x}$. So v is a value of g and an upper bound for all its values: $\max_{\mathbf{x}} g(\mathbf{x})$ exists, equals v , and is attained at $\mathbf{x}^*$. In the same way $h(\mathbf{y}) = \max_{\mathbf{x}} \mathbf{x}^\top \mathbf{A} \mathbf{y}$ satisfies $h(\mathbf{y}^*) \leq v \leq \mathbf{x}^{*\top} \mathbf{A} \mathbf{y} \leq h(\mathbf{y})$ for every $\mathbf{y}$, so $\min_{\mathbf{y}} h(\mathbf{y}) = v$, attained at $\mathbf{y}^*$. This proves the theorem. ■

Definition 18.7.2. Value and Optimal Strategies

The number v of [Theorem 18.7.3](#) is the **value** of the game with payoff matrix $\mathbf{A}$, written $v(\mathbf{A})$. A strategy $\mathbf{x} \in \Delta_m$ is **optimal for Row** if it guarantees the value, $\min_{\mathbf{y}} \mathbf{x}^\top \mathbf{A} \mathbf{y} = v(\mathbf{A})$, and $\mathbf{y} \in \Delta_n$ is **optimal for Column** if $\max_{\mathbf{x}} \mathbf{x}^\top \mathbf{A} \mathbf{y} = v(\mathbf{A})$.

The value is well defined: it equals the max–min, and that depends only on $\mathbf{A}$. A strategy $\mathbf{x}$ is optimal for Row exactly when every entry of $\mathbf{A}^\top \mathbf{x}$ is at least $v(\mathbf{A})$, and $\mathbf{y}$ is optimal for Column exactly when every entry of $\mathbf{A} \mathbf{y}$ is at most $v(\mathbf{A})$; both statements are [Lemma 18.7.2](#). These are linear conditions, so each player’s set of optimal strategies is a convex subset of the simplex. Inequality [Eq. 18.7.1](#) says that $(\mathbf{x}^*, \mathbf{y}^*)$ is a **saddle point in mixed strategies**. If Row plays $\mathbf{x}^*$, Column cannot push the expected payment below v . If Column plays $\mathbf{y}^*$, Row cannot push it above v . Neither player gains by deviating alone, and announcing an optimal strategy in advance costs nothing.

Alternatively, the theorem can be proved with no linear programming, by separation. The key fact is an alternative in the style of Farkas: either some $\mathbf{y} \in \Delta_n$ has $\mathbf{A} \mathbf{y} \leq \mathbf{0}$, or some $\mathbf{x} \in \Delta_m$ has $\mathbf{A}^\top \mathbf{x} > \mathbf{0}$. To prove it, separate the origin from the closed convex set $\{\mathbf{A} \mathbf{y} + \mathbf{z} : \mathbf{y} \in \Delta_n, \mathbf{z} \geq \mathbf{0}\}$ using §03, and apply the result to $\mathbf{A} - t\mathbf{J}$ for every real t . That route shows most directly that the theorem is about convexity. It also needs the compactness of Δ_n , (A3) of Chapter 16’s introduction, to know that the set being separated is closed. The route taken above hides all of the analysis inside Farkas’s lemma.

18.7.4 Two min–max theorems

[Theorem 18.7.3](#) and the Courant–Fischer theorem of Chapter 17 (Theorem 17.2.3) both have “min” and “max” in them, and both are called min–max theorems. They are different theorems, with different proofs. Put side by side, for a Hermitian $\mathbf{A} \in M_n(F)$ ($F = \mathbb{R}$ or $\mathbb{C}$) and a real

$\mathbf{B} \in M_{m \times n}(\mathbb{R})$:

$$\begin{aligned}\lambda_k(\mathbf{A}) &= \min_{\dim W = n-k+1} \max_{\mathbf{0} \neq \mathbf{x} \in W} R_{\mathbf{A}}(\mathbf{x}) \\ &= \max_{\dim W = k} \min_{\mathbf{0} \neq \mathbf{x} \in W} R_{\mathbf{A}}(\mathbf{x}), \\ v(\mathbf{B}) &= \min_{\mathbf{y} \in \Delta_n} \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^\top \mathbf{B} \mathbf{y} \\ &= \max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^\top \mathbf{B} \mathbf{y}.\end{aligned}$$

What is being chosen. In Courant–Fischer the outer variable is a **subspace** W , and the inner variable is a vector **inside that subspace**. The inner set depends on the outer choice, so the formula is a nested optimization, not a game between two independent players. In von Neumann’s theorem the two variables range over **fixed** sets, Δ_m and Δ_n , neither of which depends on the other.

What the theorem asserts. Von Neumann’s theorem is an **exchange of order**: the same function, with the max and the min taken in both orders, gives the same number. Courant–Fischer exchanges nothing. Its two forms use different dimensions, $n - k + 1$ and k , and swapping the order while keeping the dimension gives $\lambda_{n+1-k}(\mathbf{A})$, which differs from $\lambda_k(\mathbf{A})$ unless $n = 2k - 1$ or the two eigenvalues happen to be equal (the warning after Theorem 17.2.3). The sharpest contrast is in the inequality between the two forms. In von Neumann’s theorem, “max–min $\leq$ min–max” is [Lemma 18.7.1](#), and it is free for every function. In Courant–Fischer the corresponding inequality

$$\max_{\dim W = k} \min_{\mathbf{0} \neq \mathbf{x} \in W} R_{\mathbf{A}}(\mathbf{x}) \leq \min_{\dim W' = n-k+1} \max_{\mathbf{0} \neq \mathbf{x} \in W'} R_{\mathbf{A}}(\mathbf{x})$$

is **not** free: it is exactly the dimension count. Since $k + (n - k + 1) > n$, any W of dimension k meets any W' of dimension $n - k + 1$ in a nonzero vector $\mathbf{x}$ (Lemma 17.2.1), and then $\min_W R_{\mathbf{A}} \leq R_{\mathbf{A}}(\mathbf{x}) \leq \max_{W'} R_{\mathbf{A}}$.

Where the equality comes from. In Courant–Fischer it is a **dimension count**. Any subspace of dimension $n - k + 1$ must meet the span of the top k eigenvectors, by Lemma 17.2.1, and on the shared vector the Rayleigh quotient is pinned. Convexity plays no role: the subspaces of dimension $n - k + 1$ do not form a convex set in any sense used here, and $R_{\mathbf{A}}$ is not linear. In von Neumann’s theorem it is **convexity**. The strategy sets are simplices, the payoff is linear in each variable separately, and the equality is linear programming duality, which is Farkas’s lemma, which is a separating hyperplane. No dimension is counted anywhere.

What they need. Courant–Fischer needs a Hermitian matrix, since the spectral theorem supplies the eigenvectors that are its witnesses. Von Neumann’s theorem holds for **every** real matrix, square or not.

How each fails when its key ingredient is removed. Replace the simplices by the finite sets of pure strategies, which are not convex, and von Neumann’s equality fails, as in the matching game above. Replace the dimension $n - k + 1$ by $n - k$, which is possible when $k < n$, so that the dimension count no longer forces an intersection, and Courant–Fischer’s formula returns $\lambda_{k+1}(\mathbf{A})$ instead of $\lambda_k(\mathbf{A})$.

18.7.5 Solving a game

Example. A game with an unused column

Find the value and optimal strategies for the game with payoff matrix


$$\mathbf{A} = \begin{pmatrix} 3 & -1 & 2 \\ -2 & 4 & 0 \end{pmatrix}.$$

Solution. No saddle point. The row minima are -1 and -2 , so $\max_i \min_j a_{ij} = -1$. The column maxima are $3, 4, 2$, so $\min_j \max_i a_{ij} = 2$. The two differ, so both players must mix.

Row's side. Write $\mathbf{x} = (p, 1-p)$ with $0 \leq p \leq 1$. Then

$$\mathbf{A}^\top \mathbf{x} = (5p - 2, 4 - 5p, 2p),$$

and by [Lemma 18.7.2](#) Row guarantees $g(p) = \min\{5p - 2, 4 - 5p, 2p\}$. The graph of g is the lower envelope of three lines.



Row's guarantee $g(p)$ is the lower envelope (thick). It is highest at $p = 3/5$, where columns 1 and 2 cross at height 1. The dashed line never touches it

For $p \leq \frac{3}{5}$ the first line is the lowest: $5p - 2 \leq 4 - 5p$ exactly when $p \leq \frac{3}{5}$, and $5p - 2 \leq 2p$ exactly when $p \leq \frac{2}{5}$. For $p \geq \frac{3}{5}$ the second line is the lowest: $4 - 5p \leq 2p$ exactly when $p \geq \frac{4}{7}$, which holds since $\frac{3}{5} > \frac{4}{7}$. So $g(p) = 5p - 2$ on $[0, \frac{3}{5}]$ and $g(p) = 4 - 5p$ on $[\frac{3}{5}, 1]$. With slopes 5 and -5 , g **strictly** increases on the first interval and **strictly** decreases on the second, so its maximum $g(\frac{3}{5}) = 1$ is attained at $p = \frac{3}{5}$ and nowhere else. The only optimal strategy for Row is $\mathbf{x}^* = (\frac{3}{5}, \frac{2}{5})$, and $\mathbf{A}^\top \mathbf{x}^* = (1, 1, \frac{6}{5})$.

Column's side. Against $\mathbf{x}^*$, column 3 costs Column $\frac{6}{5}$, which is more than the value, so we expect Column never to use it. Try $\mathbf{y} = (q, 1-q, 0)$: then $\mathbf{A}\mathbf{y} = (4q - 1, 4 - 6q)$, and the two entries are equal when $q = \frac{1}{2}$. So $\mathbf{y}^* = (\frac{1}{2}, \frac{1}{2}, 0)$, with $\mathbf{A}\mathbf{y}^* = (1, 1)$.

Verification. $\mathbf{A}^\top \mathbf{x}^* = (1, 1, \frac{6}{5}) \geq 1 \cdot \mathbf{1}$ and $\mathbf{A}\mathbf{y}^* = (1, 1) \leq 1 \cdot \mathbf{1}$. As in Step 3 of the proof, these two inequalities give [Eq. 18.7.1](#) with $v = 1$. So the value is $v(\mathbf{A}) = 1$, Row's optimal strategy is $(\frac{3}{5}, \frac{2}{5})$, and $(\frac{1}{2}, \frac{1}{2}, 0)$ is optimal for Column.

Column's optimal strategy is unique too. If $\mathbf{y}$ is optimal for Column, then $\mathbf{A}\mathbf{y} \leq \mathbf{1}$, so $\mathbf{x}^{*\top} \mathbf{A}\mathbf{y} \leq 1$. But $\mathbf{x}^{*\top} \mathbf{A}\mathbf{y} = y_1 + y_2 + \frac{6}{5}y_3 = 1 + \frac{1}{5}y_3$, so $y_3 = 0$. With $\mathbf{y} = (q, 1-q, 0)$, the conditions $4q - 1 \leq 1$ and $4 - 6q \leq 1$ give $q \leq \frac{1}{2}$ and $q \geq \frac{1}{2}$. This is complementary slackness (Theorem 18.6.5) in the language of games: a column whose payoff against an optimal $\mathbf{x}$ exceeds the value is never used.

Quick check. In the game of [Example \(A game with an unused column\)](#), is $\mathbf{x} = (\frac{1}{2}, \frac{1}{2})$ an optimal strategy for Row? What does Column do against it?

Solution. No. $\mathbf{A}^\top \mathbf{x} = (\frac{1}{2}, \frac{3}{2}, 1)$, so $g(\mathbf{x}) = \frac{1}{2} < 1 = v(\mathbf{A})$. Column answers with column 1 and pays only $\frac{1}{2}$ on average. The symmetric-looking strategy is not the safe one.

Warning. “Make the opponent indifferent between all columns” is not a method. In [Example \(A game with an unused column\)](#), no $\mathbf{x} \in \Delta_2$ makes all three entries of $\mathbf{A}^\top \mathbf{x}$ equal: the first two are equal only at $p = \frac{3}{5}$, where the third is $\frac{6}{5}$, not 1. The optimal $\mathbf{x}$ equalizes only the columns that Column actually uses. Deciding **which** columns those are is the real content of solving the game, and it is what the linear program decides. Equalizing the wrong set of columns produces a strategy that is not optimal, or a “solution” with negative probabilities.

Finally, optimal strategies need not be unique, although in the example both were. [Exercise \(B3: Many optimal strategies\)](#) gives a game in which Row has a whole segment of optimal strategies, which is what the convexity of the optimal set allows.

18.7.6 Exercises

18.7.6.1 A. Check your understanding

Exercise. A1

- Define the standard simplex Δ_n , a mixed strategy, and the value of a matrix game.
- State von Neumann’s minimax theorem.
- Determine whether the following is correct, and justify your answer: every matrix game has a saddle point in pure strategies.
- Explain in two sentences why $\max_x \min_y f(x, y) \leq \min_y \max_x f(x, y)$ for every function f for which the extrema exist.
- Name the mechanism that produces the equality in Theorem 17.2.3 and the one that produces it in [Theorem 18.7.3](#).

Solution.

- $\Delta_n = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x} \geq \mathbf{0}, \sum_i x_i = 1\}$. In a game with payoff matrix $\mathbf{A} \in M_{m \times n}(\mathbb{R})$, a mixed strategy for Row is an $\mathbf{x} \in \Delta_m$ and one for Column is a $\mathbf{y} \in \Delta_n$. The value is $v(\mathbf{A}) = \max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^\top \mathbf{A} \mathbf{y}$.
- For every $\mathbf{A} \in M_{m \times n}(\mathbb{R})$, $\max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^\top \mathbf{A} \mathbf{y} = \min_{\mathbf{y} \in \Delta_n} \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^\top \mathbf{A} \mathbf{y}$, both extrema being attained. Equivalently, there are $\mathbf{x}^*, \mathbf{y}^*$ and v with $\mathbf{x}^\top \mathbf{A} \mathbf{y}^* \leq v \leq \mathbf{x}^{*\top} \mathbf{A} \mathbf{y}$ for all $\mathbf{x}, \mathbf{y}$.
- Incorrect. In the matching game with $\mathbf{A} = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$, $\max_i \min_j a_{ij} = -1 < 1 = \min_j \max_i a_{ij}$. An entry that is smallest in its row and largest in its column would make these equal.

- (d) For any x_0, y_0 , $\min_y f(x_0, y) \leq f(x_0, y_0) \leq \max_x f(x, y_0)$. The left side does not involve y_0 and the right side does not involve x_0 , so we may maximize the left over x_0 and minimize the right over y_0 .
- (e) For Courant–Fischer, a dimension count: Lemma 17.2.1 forces every competing subspace to meet a span of eigenvectors. For von Neumann, convexity: the strategy sets are simplices and the payoff is linear in each variable, and the equality is linear programming duality, which rests on separation.

18.7.6.2 B. Practice

Exercise. B1: A two-by-two game

Let $\mathbf{A} = \begin{pmatrix} 2 & -3 \\ -1 & 1 \end{pmatrix}$.

- (a) Show that the game has no saddle point in pure strategies.
- (b) Find the value and optimal strategies for both players, and verify the saddle inequality Eq. 18.7.1.

Solution.

- (a) The row minima are -3 and -1 , so $\max_i \min_j a_{ij} = -1$. The column maxima are 2 and 1 , so $\min_j \max_i a_{ij} = 1$. These differ, so no entry is both smallest in its row and largest in its column.
- (b) For $\mathbf{x} = (p, 1-p)$, $\mathbf{A}^\top \mathbf{x} = (3p-1, 1-4p)$. The first entry increases and the second decreases in p , so their minimum is largest where they are equal: $3p-1 = 1-4p$, that is, $p = \frac{2}{7}$, with common value $-\frac{1}{7}$. For $\mathbf{y} = (q, 1-q)$, $\mathbf{A}\mathbf{y} = (5q-3, 1-2q)$, and the entries are equal when $q = \frac{4}{7}$, with common value $-\frac{1}{7}$. So take $\mathbf{x}^* = (\frac{2}{7}, \frac{5}{7})$, $\mathbf{y}^* = (\frac{4}{7}, \frac{3}{7})$, and $v = -\frac{1}{7}$. Then $\mathbf{A}^\top \mathbf{x}^* = -\frac{1}{7}\mathbf{1}$ and $\mathbf{A}\mathbf{y}^* = -\frac{1}{7}\mathbf{1}$. For every $\mathbf{y} \in \Delta_2$, $\mathbf{x}^{*\top} \mathbf{A}\mathbf{y} = -\frac{1}{7}(y_1 + y_2) = -\frac{1}{7}$, and, since $\mathbf{A}\mathbf{y}^* = -\frac{1}{7}\mathbf{1}$, for every $\mathbf{x} \in \Delta_2$ we have $\mathbf{x}^\top \mathbf{A}\mathbf{y}^* = -\frac{1}{7}(x_1 + x_2) = -\frac{1}{7}$. So Eq. 18.7.1 holds with equality throughout, and $v(\mathbf{A}) = -\frac{1}{7}$. The game slightly favors Column.

Exercise. B2: A saddle point

Let $\mathbf{A} = \begin{pmatrix} 3 & 1 & 4 \\ 2 & 0 & 1 \end{pmatrix}$, the second game of this section. Show that $\mathbf{x}^* = \mathbf{e}_1$, $\mathbf{y}^* = \mathbf{e}_2$ satisfy Eq. 18.7.1 with $v = 1$, and deduce the value of the game.

Solution. For $\mathbf{x} \in \Delta_2$, $\mathbf{x}^\top \mathbf{A}\mathbf{e}_2 = x_1 \cdot 1 + x_2 \cdot 0 = x_1 \leq 1$, since $x_1 \leq x_1 + x_2 = 1$. For $\mathbf{y} \in \Delta_3$, $\mathbf{e}_1^\top \mathbf{A}\mathbf{y} = 3y_1 + y_2 + 4y_3 \geq y_1 + y_2 + y_3 = 1$. So $\mathbf{x}^\top \mathbf{A}\mathbf{e}_2 \leq 1 \leq \mathbf{e}_1^\top \mathbf{A}\mathbf{y}$ for all $\mathbf{x}, \mathbf{y}$, which is Eq. 18.7.1 with $v = 1$. By Step 4 of the proof of Theorem 18.7.3 (whose argument uses only Eq. 18.7.1), the max–min and the min–max both equal 1, so $v(\mathbf{A}) = 1$. A pure saddle point is a special case of a mixed one.

Exercise. B3: Many optimal strategies

Let $\mathbf{A} = \begin{pmatrix} 4 & 0 & 1 \\ -1 & 3 & 1 \end{pmatrix}$.

- (a) Show that $v(\mathbf{A}) = 1$.
- (b) Find **all** optimal strategies for Row.
- (c) Show that Column has exactly one optimal strategy.

Solution.

- (a) Column can play $\mathbf{e}_3$, which gives $\mathbf{x}^\top \mathbf{A} \mathbf{e}_3 = x_1 + x_2 = 1$ for every $\mathbf{x}$. Hence

$$\min_{\mathbf{y}} \max_{\mathbf{x}} \mathbf{x}^\top \mathbf{A} \mathbf{y} \leq \max_{\mathbf{x}} \mathbf{x}^\top \mathbf{A} \mathbf{e}_3 = 1.$$

Row can play $\mathbf{x} = (\frac{1}{2}, \frac{1}{2})$, with $\mathbf{A}^\top \mathbf{x} = (\frac{3}{2}, \frac{3}{2}, 1) \geq \mathbf{1}$, so $\max_{\mathbf{x}} \min_{\mathbf{y}} \mathbf{x}^\top \mathbf{A} \mathbf{y} \geq 1$ by [Lemma 18.7.2](#). By [Theorem 18.7.3](#) both equal $v(\mathbf{A})$, which is therefore 1.

- (b) By the remark after [Definition 18.7.2](#), $\mathbf{x} = (p, 1-p)$ is optimal exactly when $\mathbf{A}^\top \mathbf{x} = (5p-1, 3-3p, 1) \geq \mathbf{1}$, that is, $p \geq \frac{2}{5}$ and $p \leq \frac{2}{3}$. So the optimal strategies for Row are the whole segment $\{(p, 1-p) : \frac{2}{5} \leq p \leq \frac{2}{3}\}$.
- (c) $\mathbf{y} = (a, b, c) \in \Delta_3$ is optimal exactly when $\mathbf{A} \mathbf{y} \leq \mathbf{1}$: $4a + c \leq 1$ and $-a + 3b + c \leq 1$. Substituting $c = 1 - a - b$, these become $3a \leq b$ and $2b \leq 2a$. Hence $3a \leq b \leq a$, so $a \leq 0$, and with $a \geq 0$ we get $a = 0$ and then $b = 0$. So $\mathbf{y} = \mathbf{e}_3$ is the unique optimal strategy for Column.

18.7.6.3 C. Going deeper**Exercise. C1: Symmetric games**

Let $\mathbf{A} \in M_n(\mathbb{R})$ be skew-symmetric, $\mathbf{A}^\top = -\mathbf{A}$. Such a game is **symmetric**: the two players have the same options, and swapping their roles swaps the payoff's sign.

- (a) Prove that $\mathbf{x}^\top \mathbf{A} \mathbf{x} = 0$ for every $\mathbf{x} \in \mathbb{R}^n$, and deduce that $v(\mathbf{A}) = 0$.
- (b) Prove that $\mathbf{x}$ is optimal for Row if and only if it is optimal for Column.
- (c) In rock-paper-scissors, with the pure strategies in that order and payoff 1 for a win, -1 for a loss, and 0 for a tie,

$$\mathbf{A} = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}.$$

Prove that the uniform strategy $\frac{1}{3}\mathbf{1}$ is the **only** optimal strategy.

Solution.

- (a) $\mathbf{x}^\top \mathbf{A} \mathbf{x}$ is a 1×1 matrix, so it equals its transpose: $\mathbf{x}^\top \mathbf{A} \mathbf{x} = \mathbf{x}^\top \mathbf{A}^\top \mathbf{x} = -\mathbf{x}^\top \mathbf{A} \mathbf{x}$. Hence $\mathbf{x}^\top \mathbf{A} \mathbf{x} = 0$. Let $\mathbf{x}^*, \mathbf{y}^*, v$ be as in [Theorem 18.7.3](#). Taking $\mathbf{y} = \mathbf{x}^*$ in the right half of [Eq. 18.7.1](#) gives $v \leq \mathbf{x}^{*\top} \mathbf{A} \mathbf{x}^* = 0$. Taking $\mathbf{x} = \mathbf{y}^*$ in the left half gives $0 = \mathbf{y}^{*\top} \mathbf{A} \mathbf{y}^* \leq v$.

So $v(\mathbf{A}) = 0$.

- (b) By the remark after [Definition 18.7.2](#) and (a), $\mathbf{x}$ is optimal for Row exactly when $\mathbf{A}^\top \mathbf{x} \geq \mathbf{0}$, and optimal for Column exactly when $\mathbf{A}\mathbf{x} \leq \mathbf{0}$. Since $\mathbf{A}^\top \mathbf{x} = -\mathbf{A}\mathbf{x}$, the two conditions are the same.
- (c) With $\mathbf{x} = (x_1, x_2, x_3) \in \Delta_3$, the condition $\mathbf{A}^\top \mathbf{x} \geq \mathbf{0}$ reads

$$x_2 - x_3 \geq 0, \quad x_3 - x_1 \geq 0, \quad x_1 - x_2 \geq 0.$$

So $x_2 \geq x_3 \geq x_1 \geq x_2$, all three are equal, and since they sum to 1, each is $\frac{1}{3}$. By (a) and (b), $\frac{1}{3}\mathbf{1}$ is the unique optimal strategy for either player, and the value is 0.

Exercise. C2: The value depends continuously on the payoffs

For $\mathbf{A}, \mathbf{B} \in M_{m \times n}(\mathbb{R})$, put $\varepsilon = \max_{i,j} |a_{ij} - b_{ij}|$.

- (a) Prove that $|v(\mathbf{A}) - v(\mathbf{B})| \leq \varepsilon$.
- (b) Prove also that if $a_{ij} \geq b_{ij}$ for all i, j , then $v(\mathbf{A}) \geq v(\mathbf{B})$, and that $v(\mathbf{A} + c\mathbf{J}) = v(\mathbf{A}) + c$ for every real c .

Solution.

- (a) For $\mathbf{x} \in \Delta_m$ and $\mathbf{y} \in \Delta_n$,

$$|\mathbf{x}^\top \mathbf{A}\mathbf{y} - \mathbf{x}^\top \mathbf{B}\mathbf{y}| \leq \sum_{i,j} x_i y_j |a_{ij} - b_{ij}| \leq \varepsilon \sum_{i,j} x_i y_j = \varepsilon,$$

since $\sum_{i,j} x_i y_j = (\sum_i x_i)(\sum_j y_j) = 1$. So $\mathbf{x}^\top \mathbf{A}\mathbf{y} \leq \mathbf{x}^\top \mathbf{B}\mathbf{y} + \varepsilon$ for all $\mathbf{x}, \mathbf{y}$. Taking the minimum over $\mathbf{y}$ preserves this inequality between two functions of $\mathbf{y}$, and so does taking the maximum over $\mathbf{x}$. By [Theorem 18.7.3](#) the resulting max-mins are the values, so $v(\mathbf{A}) \leq v(\mathbf{B}) + \varepsilon$. Exchanging $\mathbf{A}$ and $\mathbf{B}$ gives $v(\mathbf{B}) \leq v(\mathbf{A}) + \varepsilon$.

- (b) If $a_{ij} \geq b_{ij}$, then $\mathbf{x}^\top \mathbf{A}\mathbf{y} \geq \mathbf{x}^\top \mathbf{B}\mathbf{y}$ for all strategies, since every $x_i y_j \geq 0$, and the same two steps give $v(\mathbf{A}) \geq v(\mathbf{B})$. For the shift, $\mathbf{x}^\top (\mathbf{A} + c\mathbf{J})\mathbf{y} = \mathbf{x}^\top \mathbf{A}\mathbf{y} + c$ for all strategies, as in Step 1 of the proof of [Theorem 18.7.3](#), and adding a constant to a function adds it to every minimum and maximum. So $v(\mathbf{A} + c\mathbf{J}) = v(\mathbf{A}) + c$.