Von Neumann's Minimax theorem with proof

Von Neumann's Minimax Theorem). *Let A be a $m \times n$ payoff matrix, and let $\Delta_m = \{\mathbf{x} : \mathbf{x} \geq \mathbf{0}; \sum_i x_i = 1\}$, $\Delta_n = \{\mathbf{y} : \mathbf{y} \geq \mathbf{0}; \sum_j y_j = 1\}$, then*

$$\max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^T A \mathbf{y} = \min_{\mathbf{y} \in \Delta_n} \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T A \mathbf{y}.$$

This quantity is called the ***value*** *of the two-person, zero-sum game with payoff matrix A.*

The proof is not examinable.

A result from convex analysis used in the proof:

Definition A set $K \subseteq \mathbb{R}^d$ is **convex** if, for any two points $\mathbf{a}, \mathbf{b} \in K$, the line segment that connects them,

$$\{p\mathbf{a} + (1-p)\mathbf{b} : p \in [0,1]\},$$

also lies in K.

Theorem (**The Separating Hyperplane Theorem**). *Suppose that $K \subseteq \mathbb{R}^d$ is closed and convex. If $\mathbf{0} \notin K$, then there exists $\mathbf{z} \in \mathbb{R}^d$ and $c \in \mathbb{R}$ such that*

$$0 < c < \mathbf{z}^T \mathbf{v},$$

for all $\mathbf{v} \in K$.

Theorem (**The Separating Hyperplane Theorem**). *Suppose that* $K \subseteq \mathbb{R}^d$ *is closed and convex. If* $\mathbf{0} \notin K$, *then there exists* $\mathbf{z} \in \mathbb{R}^d$ *and* $c \in \mathbb{R}$ *such that*

$$0 < c < \mathbf{z}^T \mathbf{v} \, ,$$

for all $\mathbf{v} \in K$.

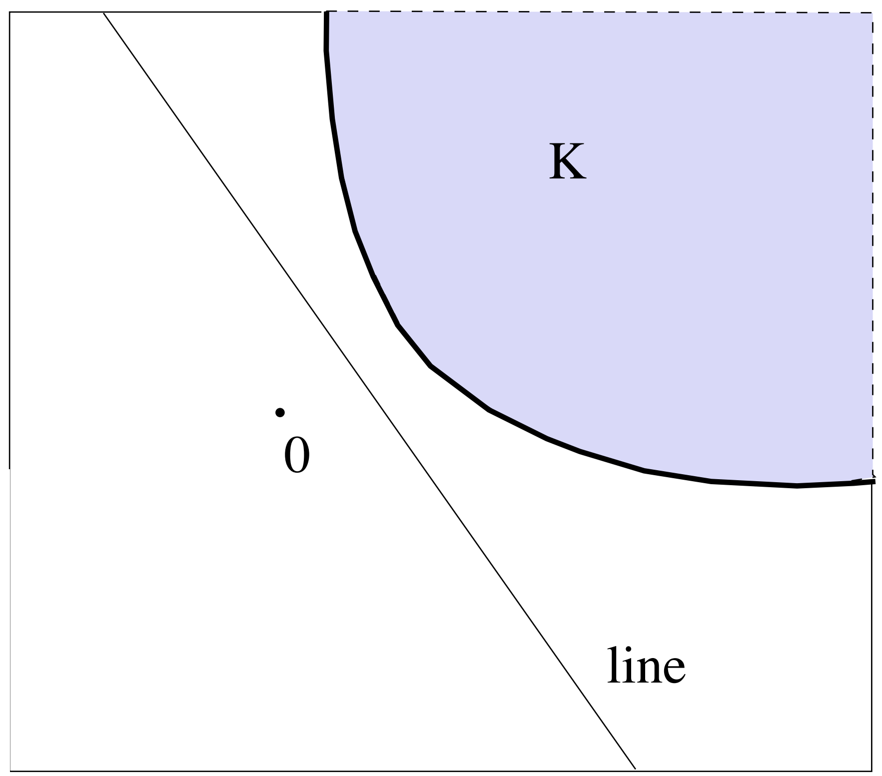

Interpretation of theorem: there is a **hyperplane** that separates $\mathbf{0}$ from K. In particular, on any continuous path from $\mathbf{0}$ to K, there is some point that lies on this hyperplane.

The separating hyperplane is given by $\{\mathbf{x} \in \mathbb{R}^d : \mathbf{z}^T \mathbf{x} = c\}$.

Lemma *Let X and Y be closed and bounded sets in $\mathbb{R}$ and let $(\mathbf{x}^*, \mathbf{y}^*) \in X \times Y$. Let $f : X \times Y \to \mathbb{R}$ be continuous in both coordinates. Then,*

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

Lemma *Let X and Y be closed and bounded sets in $\mathbb{R}$ and let $(\mathbf{x}^*, \mathbf{y}^*) \in X \times Y$. Let $f : X \times Y \to \mathbb{R}$ be continuous in both coordinates. Then,*

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

Proof. Let $(\mathbf{x}^*, \mathbf{y}^*) \in X \times Y$ be given. Clearly we have $f(\mathbf{x}^*, \mathbf{y}^*) \leq \sup_{\mathbf{x}\in X} f(\mathbf{x}, \mathbf{y}^*)$ and $\inf_{\mathbf{x}\in X} f(\mathbf{x}, \mathbf{y}^*) \leq f(\mathbf{x}^*, \mathbf{y}^*)$, which gives us

$$\inf_{\mathbf{y}\in Y} f(\mathbf{x}^*, \mathbf{y}) \leq \sup_{\mathbf{x}\in X} f(\mathbf{x}, \mathbf{y}^*).$$

Because the inequality holds for any $\mathbf{x}^* \in X$, it holds for $\sup_{\mathbf{x}^*\in X}$ of the quantity on the left. Similarly, because the inequality holds for all $\mathbf{y}^* \in Y$, it must hold for the $\inf_{\mathbf{y}^*\in Y}$ of the quantity on the right. We have:

$$\sup_{\mathbf{x}\in X} \inf_{\mathbf{y}\in Y} f(\mathbf{x}, \mathbf{y}) \leq \inf_{\mathbf{y}\in Y} \sup_{\mathbf{x}\in X} f(\mathbf{x}, \mathbf{y}).$$

Von Neumann's Minimax Theorem). *Let A be a $m \times n$ payoff matrix, and let $\Delta_m = \{\mathbf{x} : \mathbf{x} \geq \mathbf{0}; \sum_i x_i = 1\}$, $\Delta_n = \{\mathbf{y} : \mathbf{y} \geq \mathbf{0}; \sum_j y_j = 1\}$, then*

$$\max_{\mathbf{x}\in\Delta_m} \min_{\mathbf{y}\in\Delta_n} \mathbf{x}^T A\mathbf{y} = \min_{\mathbf{y}\in\Delta_n} \max_{\mathbf{x}\in\Delta_m} \mathbf{x}^T A\mathbf{y}.$$

This quantity is called the ***value*** *of the two-person, zero-sum game with payoff matrix A.*

Proof. That

$$\max_{\mathbf{x}\in\Delta_m} \min_{\mathbf{y}\in\Delta_n} \mathbf{x}^T A\mathbf{y} \leq \min_{\mathbf{y}\in\Delta_n} \max_{\mathbf{x}\in\Delta_m} \mathbf{x}^T A\mathbf{y}$$

follows immediately from the lemma because $f(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T A\mathbf{y}$ is a continuous function in both variables and $\Delta_m \subset \mathbb{R}^m$, $\Delta_n \subset \mathbb{R}^n$ are closed and bounded.

Von Neumann's Minimax Theorem *Let A be a $m \times n$ payoff matrix, and let $\Delta_m = \{\mathbf{x} : \mathbf{x} \geq \mathbf{0}; \sum_i x_i = 1\}$, $\Delta_n = \{\mathbf{y} : \mathbf{y} \geq \mathbf{0}; \sum_j y_j = 1\}$, then*

$$\max_{\mathbf{x}\in\Delta_m} \min_{\mathbf{y}\in\Delta_n} \mathbf{x}^T A\mathbf{y} = \min_{\mathbf{y}\in\Delta_n} \max_{\mathbf{x}\in\Delta_m} \mathbf{x}^T A\mathbf{y}.$$

This quantity is called the ***value*** *of the two-person, zero-sum game with payoff matrix A.*

Proof. That

$$\max_{\mathbf{x}\in\Delta_m} \min_{\mathbf{y}\in\Delta_n} \mathbf{x}^T A\mathbf{y} \leq \min_{\mathbf{y}\in\Delta_n} \max_{\mathbf{x}\in\Delta_m} \mathbf{x}^T A\mathbf{y}$$

follows immediately from the lemma because $f(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T A\mathbf{y}$ is a continuous function in both variables and $\Delta_m \subset \mathbb{R}^m$, $\Delta_n \subset \mathbb{R}^n$ are closed and bounded.

Recall: *Lemma* *Let X and Y be closed and bounded sets in $\mathbb{R}$ and let $(\mathbf{x}^*, \mathbf{y}^*) \in X \times Y$. Let $f : X \times Y \to \mathbb{R}$ be continuous in both coordinates. Then,*

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

Theorem's statement: $\max_{\mathbf{x}\in\Delta_m} \min_{\mathbf{y}\in\Delta_n} \mathbf{x}^T A\mathbf{y} = \min_{\mathbf{y}\in\Delta_n} \max_{\mathbf{x}\in\Delta_m} \mathbf{x}^T A\mathbf{y}$

Proof (ct)

For the other inequality, suppose toward a contradiction that

$$\max_{\mathbf{x}\in\Delta_m} \min_{\mathbf{y}\in\Delta_n} \mathbf{x}^T A\mathbf{y} < \lambda < \min_{\mathbf{y}\in\Delta_n} \max_{\mathbf{x}\in\Delta_m} \mathbf{x}^T A\mathbf{y}.$$

We can define a new game with payoff matrix $\hat{A}$ given by $\hat{a}_{ij} = a_{ij} - \lambda$. For this game, we have

$$\max_{\mathbf{x}\in\Delta_m} \min_{\mathbf{y}\in\Delta_n} \mathbf{x}^T \hat{A}\mathbf{y} < 0 < \min_{\mathbf{y}\in\Delta_n} \max_{\mathbf{x}\in\Delta_m} \mathbf{x}^T \hat{A}\mathbf{y}. \qquad (*)$$

Each mixed strategy $\mathbf{y} \in \Delta_n$ for player II yields a payoff vector $\hat{A}\mathbf{y} \in \mathbb{R}^m$. Let K denote the set of all vectors $\mathbf{u}$ for which there exists a payoff vector $\hat{A}\mathbf{y}$ such that u dominates $\hat{A}\mathbf{y}$. That is,

$$K = \Big\{\mathbf{u} = \hat{A}\mathbf{y} + \mathbf{v} : \mathbf{y} \in \Delta_n,\ \mathbf{v} \in \mathbb{R}^m, \mathbf{v} \geq \mathbf{0}\Big\}.$$

$$K = \left\{ \mathbf{u} = \hat{A}\mathbf{y} + \mathbf{v} : \mathbf{y} \in \Delta_n,\ \mathbf{v} \in \mathbb{R}^m, \mathbf{v} \geq \mathbf{0} \right\}.$$

It is easy to see that K is convex and closed: this follows immediately from the fact that Δ_n, the set of probability vectors corresponding to mixed strategies $\mathbf{y}$ for player II, is closed, bounded and convex. Also, K cannot contain the $\mathbf{0}$ vector because if $\mathbf{0}$ were in K, there would be some mixed strategy $\mathbf{y} \in \Delta_n$ such that $\hat{A}\mathbf{y} \leq \mathbf{0}$, hence for any $\mathbf{x} \in \Delta_m$ we have $\mathbf{x}^T \hat{A}\mathbf{y} \leq 0$, which would contradict the right-hand side of (*) .

Recall:

$$\max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^T \hat{A}\mathbf{y} < 0 < \min_{\mathbf{y} \in \Delta_n} \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T \hat{A}\mathbf{y} \qquad (*)$$

Thus, K satisfies the conditions of the separating hyperplane theorem which gives us $\mathbf{z} \in \mathbb{R}^m$ and $c > 0$ such that $0 < c < \mathbf{z}^T\mathbf{w}$ for all $\mathbf{w} \in K$. That is,

$$\mathbf{z}^T(\hat{A}\mathbf{y} + \mathbf{v}) > c > 0 \text{ for all } \mathbf{y} \in \Delta_n, \quad \mathbf{v} \geq \mathbf{0}. \qquad (**)$$

It must be the case that $z_i \geq 0$ for all i because if $z_j < 0$, for some j we could choose $\mathbf{y} \in \Delta_n$, so that $\mathbf{z}^T A\mathbf{y} + \sum_i z_i v_i$ would be negative (let $v_i = 0$ for $i \neq j$ and $v_j \to \infty$), which would contradict (**).

Thus, K satisfies the conditions of the separating hyperplane theorem which gives us $\mathbf{z} \in \mathbb{R}^m$ and $c > 0$ such that $0 < c < \mathbf{z}^T\mathbf{w}$ for all $\mathbf{w} \in K$. That is,

$$\mathbf{z}^T(\hat{A}\mathbf{y} + \mathbf{v}) > c > 0 \text{ for all } \mathbf{y} \in \Delta_n, \quad \mathbf{v} \geq \mathbf{0}. \tag{**}$$

It must be the case that $z_i \geq 0$ for all i because if $z_j < 0$, for some j we could choose $\mathbf{y} \in \Delta_n$, so that $\mathbf{z}^T A\mathbf{y} + \sum_i z_i v_i$ would be negative (let $v_i = 0$ for $i \neq j$ and $v_j \to \infty$), which would contradict (**).

The same condition (**) gives us that not all of the z_i can be zero. This means that $s = \sum_{i=1}^m z_i$ is strictly positive, so that $\mathbf{x} = (1/s)(z_1, \ldots, z_m)^T = (1/s)\mathbf{z} \in \Delta_m$, with $\mathbf{x}^T A\mathbf{y} > c > 0$ for all $\mathbf{y} \in \Delta_n$.

In other words, $\mathbf{x}$ is a mixed strategy for player I that gives a positive expected payoff against any mixed strategy of player II. This contradicts the left hand inequality of (**), which says that player I can assure at best a negative payoff. □