5.8. Mixed Strategies for Zero-Sum Games and The Minimax Theorem.

5.8.1. Definition.

A bimatrix game with matrices A and B is called a zero-sum game, if A = -B.

Hence everything that was said about bimatrix games in general makes sense for zero-sum games. Especially, there are mixed strategies, and a pair of mixed strategies $\left(p_e, q_e\right)$ is a point of equilibrium, if

$$\forall\left(\vec{p} \in \Delta_{m-1}\right)\left(\vec{p}^{\mathrm{T}} A \vec{q}_e \leq \vec{p}_e^{\,T} A \vec{q}_e\right)$$

$$\forall\left(\vec{q} \in \Delta_{n-1}\right)\left(\vec{p}_e^{\,\mathrm{T}} B \vec{q} \leq \vec{p}_e^{\,T} B \vec{q}_e\right)$$

Since $A = -B$, these conditions reduce to

$$\forall\left(\vec{p} \in \Delta_{m-1}\right)\left(\vec{p}^{\mathrm{T}} A \vec{q}_e \leq \vec{p}_e^{\,T} A \vec{q}_e\right)$$

$$\forall\left(\vec{q} \in \Delta_{n-1}\right)\left(\vec{p}_e^{\,\mathrm{T}} A \vec{q} \geq \vec{p}_e^{\,T} A \vec{q}_e\right)$$

These two inequalities are equivalent to

$$\forall\left(\vec{p} \in \Delta_{m-1}, \vec{q} \in \Delta_{n-1}\right)\left(\vec{p}^{\mathrm{T}} A \vec{q}_e \leq \vec{p}_e^{\,T} A \vec{q}_e \leq \vec{p}_e^{\,\mathrm{T}} A \vec{q}\right)$$

5.8.2. Definition.

Let G be a zero-sum game with matrix A. A pair of mixed strategies $\left(p_s, q_s\right)$ is called a saddle point, if

$$\vec{p}^{\mathrm{T}} A \vec{q}_s \leq \vec{p}_s^{\,T} A \vec{q}_s \leq \vec{p}_s^{\,\mathrm{T}} A \vec{q}$$

holds for all $\vec{p} \in \Delta_{m-1}$ and all $\vec{q} \in \Delta_{n-1}$.

The following statement is obvious:

5.8.3. Proposition.

Every point of equilibrium of a zero-sum game is a saddle point, and vice versa, every saddle point is a point of equilibrium.

We know from Nash's theorem that every zero-sum game has at least one point of equilibrium consisting of mixed strategies. Since every point of equilibrium is a saddle point, we obtain.

5.8.4. Theorem (John von Neumann, 1928)

Let G be a zero-sum game with matrix A. Then G has a saddle point.

We can rewrite this condition for a point of equilibrium in the following form:

5.8.5. Theorem.

Let G be a zero-sum game with matrix A, and let (p_s, q_s) *be a saddle point consisting of mixed strategies. Then*

$$\begin{aligned}\vec{p}_s^{\,T} A\vec{q}_s &= \max\left\{\vec{p}^{\,T} A\vec{q}_s : \vec{p} \in \Delta_{m-1}\right\}\\ &= \min\left\{\vec{p}_s^{\,T} A\vec{q} : \vec{q} \in \Delta_{n-1}\right\}\end{aligned}$$

Moreover, in this case

$$\begin{aligned}\vec{p}_s^{\,T} A\vec{q}_s &= \min\left\{\max\left\{\vec{p}^{\,T} A\vec{q} : \vec{p} \in \Delta_{m-1}\right\} : \vec{q} \in \Delta_{n-1}\right\}\\ &= \max\left\{\min\left\{\vec{p}^{\,T} A\vec{q} : \vec{q} \in \Delta_{n-1}\right\} : \vec{p} \in \Delta_{m-1}\right\}\end{aligned}$$

Proof. We only have to verify the second part of the theorem. First, we show that for every choice of sets X and Y and every function $\omega : X \times Y \to \Re \cup \{\pm\infty\}$ we have

$$\max\left\{\min\left\{\omega(x,y) : y \in Y\right\} : x \in X\right\} \le \min\left\{\max\left\{\omega(x,y) : x \in X\right\} : y \in Y\right\}$$

Indeed, if $x_0 \in X$ and $y_0 \in Y$ are given, then

$$\min\left\{\omega(x_0,y) : y \in Y\right\} \le \omega(x_0,y_0) \le \max\left\{\omega(x,y_0) : x \in X\right\}$$

Since this holds for each value of $x_0 \in X$, we conclude that

$$\max\left\{\min\left\{\omega(x_0,y) : y \in Y\right\} : x_0 \in X\right\} \le \max\left\{\omega(x,y_0) : x \in X\right\}$$

Again, this holds for every value of $y_0 \in Y$, hence

$$\max\left\{\min\left\{\omega(x_0,y) : y \in Y\right\} : x_0 \in X\right\} \le \min\left\{\max\left\{\omega(x,y_0) : x \in X\right\} : y_0 \in Y\right\}$$

To complete the proof, we have to verify that for every saddle point (p_s, q_s) the inequality

$$\min\left\{\max\left\{\vec{p}^{\,T} A\vec{q} : \vec{p} \in \Delta_{m-1}\right\} : \vec{q} \in \Delta_{n-1}\right\} \le \vec{p}_s^{\,T} A\vec{q}_s \le \max\left\{\min\left\{\vec{p}^{\,T} A\vec{q} : \vec{q} \in \Delta_{n-1}\right\} : \vec{p} \in \Delta_{m-1}\right\}$$

holds. First, note that

$$\begin{aligned}\vec{p}_s^{\,T} A\vec{q}_s &= \max\left\{\vec{p}^{\,T} A\vec{q}_s : \vec{p} \in \Delta_{m-1}\right\}\\ &= \min\left\{\vec{p}_s^{\,T} A\vec{q} : \vec{q} \in \Delta_{n-1}\right\}\end{aligned}$$

Hence

$$\begin{aligned}\min\left\{\max\left\{\vec{p}^{T}A\vec{q}:\vec{p}\in\Delta_{m-1}\right\}:\vec{q}\in\Delta_{n-1}\right\} &\le \max\left\{\vec{p}^{T}A\vec{q}_s:\vec{p}\in\Delta_{m-s}\right\}\\ &= \vec{p}_s{}^{T}A\vec{q}_s\\ &= \min\left\{\vec{p}_s{}^{T}A\vec{q}:\vec{q}\in\Delta_{n-1}\right\}\\ &\le \max\left\{\min\left\{\vec{p}^{T}A\vec{q}:\vec{q}\in\Delta_{n-1}\right\}:\vec{p}\in\Delta_{m-1}\right\}\end{aligned}$$

This completes the proof.

Another popular version of John von Neumann's theorem is the following corollary:

5.8.6. Corollary.

If G is a zero-sum game with matrix A, then

$$\min\left\{\max\left\{\vec{p}^{T}A\vec{q}:\vec{p}\in\Delta_{m}\right\}:\vec{q}\in\Delta_{n}\right\}=\max\left\{\min\left\{\vec{p}^{T}A\vec{q}:\vec{q}\in\Delta_{n}\right\}:\vec{p}\in\Delta_{m}\right\}$$

5.8.7. Definition.

If G is a zero-sum game with matrix A, then the value of G is defined as

$$\begin{aligned}v(G) &= \min\left\{\max\left\{\vec{p}^{T}A\vec{q}:\vec{p}\in\Delta_{m}\right\}:\vec{q}\in\Delta_{n}\right\}\\ &= \max\left\{\min\left\{\vec{p}^{T}A\vec{q}:\vec{q}\in\Delta_{n}\right\}:\vec{p}\in\Delta_{m}\right\}\end{aligned}$$

5.8.8. Proposition.

If G is a zero-sum game with matrix A, then for every saddle point $(\vec{p}_s,\vec{q}_s)$ of mixed strategies we have

$$v(G)=\vec{p}_e{}^{T}A\vec{q}_e$$

Conversely, if $(\vec{p},\vec{q})$ is a pair of mixed strategies, then the equation $v(G)=\vec{p}^{T}A\vec{q}$ does not imply that $(\vec{p},\vec{q})$ is a saddle point. For example, if

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

then $((1,0),(0,1))$ is a saddle point, hence the value of the game is 1. However, if $(\vec{p},\vec{q})=((0,1),(0,1))$, then we also have

$$(0,1)A\begin{pmatrix}1\\ 0\end{pmatrix}=1$$

even though $((0,1),(0,1))$ is not a saddle point.

Saddle points are easy to find using derivatives. If $\vec{p} = (p_1, \cdots, p_m)$ and $\vec{q} = (q_1, \cdots, q_n)$, and if

$$\omega(p_1, \cdots, p_m, q_1 \cdots, q_n) = (p_1, \cdots p_m) A \begin{pmatrix} q_1 \\ \vdots \\ q_n \end{pmatrix}$$

then, for fixed values of the q_j's, the function ω has a maximum at $\vec{p}_s$, and for fixed values of the p_i's, the function ω has a minimum at $\vec{q}_s$. Hence the derivatives of ω with respect to the p_i's and q_j's has to be *0*. We have to be a bit careful, since the coordinates of $\vec{p}$ and $\vec{q}$ are constrained by the condition that the sum of their coordinates is equal to 1. Therefore the domain of ω is not an open subset of Euclidean space, and boundary points usually deserve extra attention. This is best explained by an example:

Example. Find all saddle points of the game "Rock, Paper, Scissors".

The matrix of this game is given by

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

So the function ω is given by

$$\omega(p_1, p_2, p_3, q_1, q_2, q_3) = p_1 q_2 - p_1 q_3 - p_2 q_1 + p_2 q_3 + p_3 q_1 - p_3 q_2$$

We can replace p_3 by $1 - p_1 - p_2$ and q_3 by $1 - q_1 - q_2$ and obtain

$$\omega^*(p_1, p_2, q_1, q_2) = p_2 - p_1 + q_1 - q_2 + 3p_1 q_2 - 3p_2 q_1$$

The domain of this new function is

$$\{(p_1, p_2, q_1, q_2) : 0 \le p_1, p_2, q_1, q_2, 1 - p_1 - p_2, 1 - p_1 - p_2\}$$

The interior of this set is found by using just strict inequalities. Saddle points in the interior are found by taking derivatives:

$$\frac{\partial \omega^*}{\partial p_1} = -1 + 3q_2 = 0$$

$$\frac{\partial \omega^*}{\partial p_2} = 1 - 3q_1 = 0$$

$$\frac{\partial \omega^*}{\partial q_1} = 1 - 3p_2 = 0$$

$$\frac{\partial \omega^*}{\partial q_2} = -1 + 3p_2 = 0$$

Hence the only saddle point in the interior is given by $\vec{p}_s = \vec{q}_s = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$.

In order to find saddle points on the boundary, we have to use the method of Lagrange multipliers. We will skip this here.