Equilibrium Computation in Games

Bernhard von Stengel

Department of Mathematics
London School of Economics

Overview – Two-Player Games

- Best responses and Nash equilibrium
- Zero-sum games, von Neumann's minimax theorem with proof
- Geometry of Nash equilibria
  - Lemke-Howson, odd number of Nash equilibria
  - Bimatrix games and labeled polytopes
  - Complementary pivoting

- Extensive games
  - Perfect recall and the sequence form
- Correlated equilibria
- PPAD

Zero-sum games: start

A zero-sum game

Max \ min	*l*	*r*
T	**1**	**0**
B	**−2**	**2**

A zero-sum game

Max \ min	l	r
T	1	0
B	-2	2

Best response payoffs / costs:

□ maximizing row player

○ minimizing column player

A zero-sum game

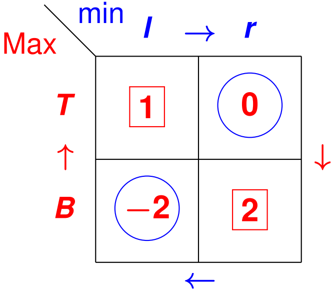

Best response payoffs / costs:

□ maximizing row player

○ minimizing column player

⇒ no "stable" way of playing deterministically

Optimal mixed (= randomized) strategies

Max \ min	$\frac{2}{5}$	$\frac{3}{5}$	
$\frac{4}{5}$	**1**	**0**	$\frac{2}{5}$
$\frac{1}{5}$	**−2**	**2**	$\frac{2}{5}$
	$\frac{2}{5}$	$\frac{2}{5}$	

Optimal mixed (= randomized) strategies

Max \ min	$\frac{2}{5}$	$\frac{3}{5}$	
$\frac{4}{5}$	1	0	$\frac{2}{5}$
$\frac{1}{5}$	−2	2	$\frac{2}{5}$
	$\frac{2}{5}$	$\frac{2}{5}$	

Probabilities found with the **“difference trick”**:

they are inversely proportional to the opponent-payoff differences in the respective rows and columns, and make the opponent **indifferent.**

Payoffs must be expected utilities

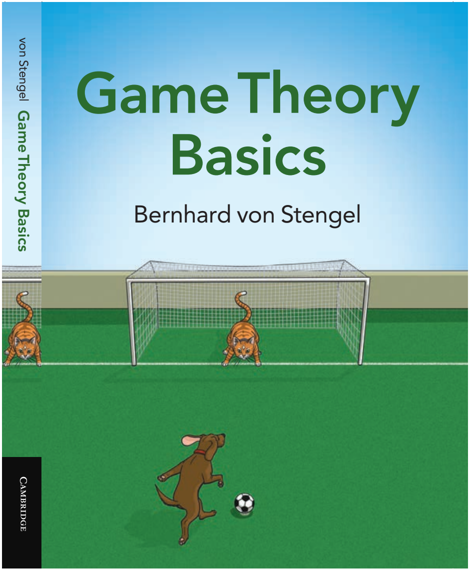

Payoffs must be expected utilities

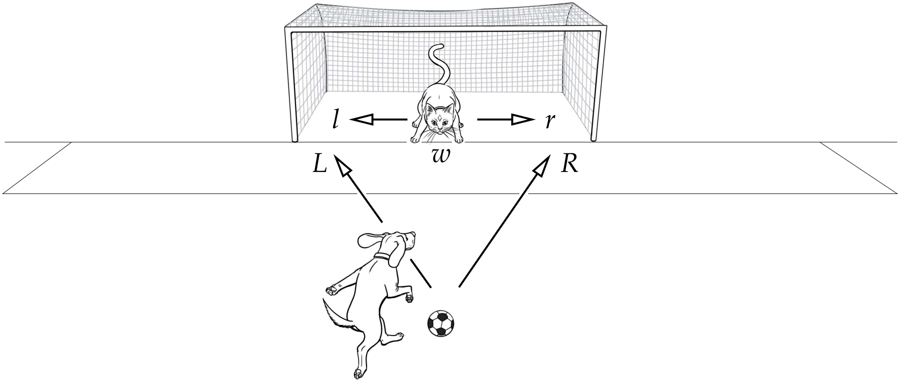

Probability of goal = payoff to striker (Maximizer), cost to goalkeeper (minimizer)

Max \ min	*l*	*w*	*r*
L	0.5	0.7	1.0
R	1.0	0.6	0.4

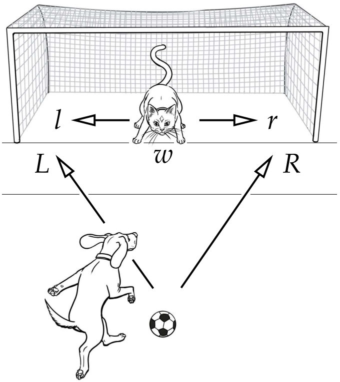

Probability of goal = payoff to striker (Maximizer), cost to goalkeeper (minimizer)

Max \ min	*l*	*w*	*r*
L	0.5	0.7	1.0
R	1.0	0.6	0.4

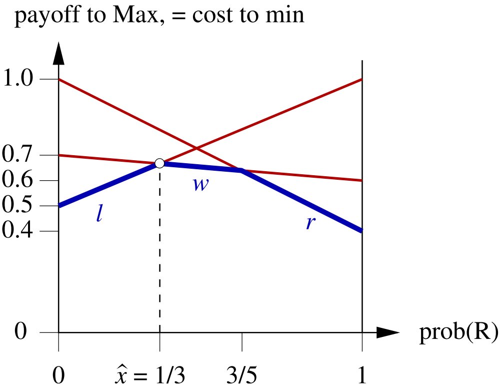

optimal for row player: **max-min strategy $\hat{x}$**

Notation: treat vectors and scalars as matrices

All vectors are column vectors. $\boldsymbol{A}^\top$ = matrix $\boldsymbol{A}$ transposed.

$\mathbf{0} = (\mathbf{0}, \dots, \mathbf{0})^\top, \quad \mathbf{1} = (\mathbf{1}, \dots, \mathbf{1})^\top$.

$\boldsymbol{A}\boldsymbol{y}$ = linear combination of columns of $\boldsymbol{A}$

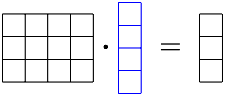

$\boldsymbol{x}^\top\boldsymbol{A}$ = linear combination of rows of $\boldsymbol{A}$

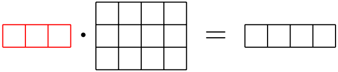

$\boldsymbol{x}^\top\boldsymbol{b}$ = scalar product of $\boldsymbol{x}$ and $\boldsymbol{b}$

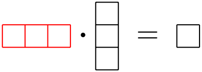

$\boldsymbol{y}\alpha$ = (column) vector $\boldsymbol{x}$ scaled by α

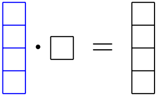

$\alpha\boldsymbol{x}^\top$ = row vector $\boldsymbol{x}^\top$ scaled by α

Example use of notation

Given: $\boldsymbol{A} \in \mathbb{R}^{\boldsymbol{m} \times \boldsymbol{n}}$,

probability vectors $\boldsymbol{x} \in \mathbb{R}^{\boldsymbol{m}}$ for rows, $\boldsymbol{y} \in \mathbb{R}^{\boldsymbol{n}}$ for columns

i.e. $\mathbf{1}^\top \boldsymbol{x} = 1, \quad \mathbf{1}^\top \boldsymbol{y} = 1\,.$

Constant $\alpha \in \mathbb{R}$ added to all entries of $\boldsymbol{A}$ gives $\boldsymbol{A} + \mathbf{1}\alpha\mathbf{1}^\top$.

Then

$$
\begin{aligned}
\boldsymbol{x}^\top(\boldsymbol{A} + \mathbf{1}\alpha\mathbf{1}^\top)\boldsymbol{y} &= \boldsymbol{x}^\top\boldsymbol{A}\boldsymbol{y} + \boldsymbol{x}^\top(\mathbf{1}\alpha\mathbf{1}^\top)\boldsymbol{y} \\
&= \boldsymbol{x}^\top\boldsymbol{A}\boldsymbol{y} + (\boldsymbol{x}^\top\mathbf{1})\alpha(\mathbf{1}^\top\boldsymbol{y}) \\
&= \boldsymbol{x}^\top\boldsymbol{A}\boldsymbol{y} + \alpha\,.
\end{aligned}
$$

The best-response condition

Bimatrix game (A, B)

row player I	column player II
m pure strategies $i = 1, \ldots, m$	n pure strategies $j = 1, \ldots, n$
payoff a_{ij}, payoff matrix A	payoff b_{ij}, payoff matrix B

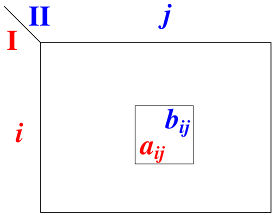

Bimatrix game (A, B)

row player I	column player II
m pure strategies $i = 1, \ldots, m$	n pure strategies $j = 1, \ldots, n$
payoff a_{ij}, payoff matrix A	payoff b_{ij}, payoff matrix B

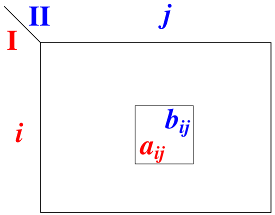

mixed strategy x	mixed strategy y
probabilities $x_1, \ldots, x_m$	probabilities $y_1, \ldots, y_n$
expected payoff $x^\top Ay$	expected payoff $x^\top By$

Expected payoffs

Given: $m \times n$ bimatrix game (A, B) .

mixed strategy vector $x = (x_1, \ldots, x_m)^\top$ for player I

mixed strategy vector $y = (y_1, \ldots, y_n)^\top$ for player II

Expected payoff to player I is

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

Expected payoff to player II is

$$\sum_{i=1}^{m} \sum_{j=1}^{n} x_i \, b_{ij} \, y_j = x^\top B y = \sum_{j=1}^{n} (x^\top B)_j \, y_j$$

Expected payoffs – what the player controls

The expected payoff $\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{y}$ to player I should be read as $\boldsymbol{x}^\top(\boldsymbol{A}\boldsymbol{y})$,

$$\sum_{i=1}^{m} x_i (\boldsymbol{A}\boldsymbol{y})_i$$

because player I chooses $\boldsymbol{x}$, against given $\boldsymbol{y}$ and expected payoff vector $\boldsymbol{A}\boldsymbol{y}$ with entries $(\boldsymbol{A}\boldsymbol{y})_i$ for the rows i.

Expected payoffs – what the player controls

The expected payoff $\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{y}$ to player I should be read as $\boldsymbol{x}^\top(\boldsymbol{A}\boldsymbol{y})$,

$$\sum_{i=1}^{m} x_i (Ay)_i$$

because player I chooses $\boldsymbol{x}$, against given $\boldsymbol{y}$ and expected payoff vector $\boldsymbol{A}\boldsymbol{y}$ with entries $(\boldsymbol{A}\boldsymbol{y})_i$ for the rows i.

Similarly, read the expected payoff $\boldsymbol{x}^\top \boldsymbol{B}\boldsymbol{y}$ to player II as $(\boldsymbol{x}^\top \boldsymbol{B})\boldsymbol{y}$.

Expected payoffs – what the player controls

The expected payoff $\boldsymbol{x}^\top \boldsymbol{Ay}$ to player I should be read as $\boldsymbol{x}^\top(\boldsymbol{Ay})$,

$$\sum_{i=1}^{m} x_i (\boldsymbol{Ay})_i$$

because player I chooses $\boldsymbol{x}$, against given $\boldsymbol{y}$ and expected payoff vector $\boldsymbol{Ay}$ with entries $(\boldsymbol{Ay})_i$ for the rows i.

Similarly, read the expected payoff $\boldsymbol{x}^\top \boldsymbol{By}$ to player II as $(\boldsymbol{x}^\top \boldsymbol{B})\boldsymbol{y}$.

Example

$\boldsymbol{Ay} = (4, 4, 3)^\top$, $\boldsymbol{x}^\top = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, expected payoff $3\frac{2}{3}$.

Expected payoffs – what the player controls

The expected payoff $\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{y}$ to player I should be read as $\boldsymbol{x}^\top(\boldsymbol{A}\boldsymbol{y})$,

$$\sum_{i=1}^{m} x_i (\boldsymbol{A}\boldsymbol{y})_i$$

because player I chooses $\boldsymbol{x}$, against given $\boldsymbol{y}$ and expected payoff vector $\boldsymbol{A}\boldsymbol{y}$ with entries $(\boldsymbol{A}\boldsymbol{y})_i$ for the rows i.

Similarly, read the expected payoff $\boldsymbol{x}^\top \boldsymbol{B}\boldsymbol{y}$ to player II as $(\boldsymbol{x}^\top \boldsymbol{B})\boldsymbol{y}$.

Example

$\boldsymbol{A}\boldsymbol{y} = (4, 4, 3)^\top$, $\boldsymbol{x}^\top = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, expected payoff $3\frac{2}{3}$.

Is this the optimal expected payoff? No, player I could get payoff 4 with $\boldsymbol{x}^\top = (1, 0, 0)$ or $\boldsymbol{x}^\top = (0, 1, 0)$ or $\boldsymbol{x}^\top = (\frac{1}{3}, \frac{2}{3}, 0)$ or . . . .

The best response condition

Theorem

Given: $m \times n$ bimatrix game (A, B) .

Let x be a mixed strategy of player I and
y be a mixed strategy of player II. Then

	x is a best response to y
$\Leftrightarrow$	for all pure strategies i of player I :
	$x_i > 0 \;\Rightarrow\; (Ay)_i = u = \max\{\,(Ay)_k \mid 1 \le k \le m\,\}.$

The best response condition

Theorem

Given: $m \times n$ bimatrix game (A, B) .

Let x be a mixed strategy of player I and
y be a mixed strategy of player II. Then

	x is a best response to y
$\Leftrightarrow$	for all pure strategies i of player I :
	$x_i > 0 \;\Rightarrow\; (Ay)_i = u = \max\{(Ay)_k \mid 1 \le k \le m\}$.

That is, only **pure best responses** may be played with positive probability.

Proof of the best response condition

	$\boldsymbol{x}$ is a best response to $\boldsymbol{y}$
$\Leftrightarrow$	for all pure strategies $\boldsymbol{i}$ of player I :
	$x_i > 0 \;\Rightarrow\; (Ay)_i \;=\; u \;=\max\{\,(Ay)_k \mid 1 \le k \le m\,\}.$

Proof

$$\begin{aligned} x^\top Ay &= \sum_{i=1}^{m} x_i\,(Ay)_i = \sum_{i=1}^{m} x_i\,(u-(u-(Ay)_i)) \\ &= \sum_{i=1}^{m} x_i\,u - \sum_{i=1}^{m} x_i\,(u-(Ay)_i) \end{aligned}$$

Proof of the best response condition

$\boldsymbol{x}$ is a best response to $\boldsymbol{y}$

$\Leftrightarrow$ for all pure strategies i of player I :

$$x_i > 0 \;\Rightarrow\; (Ay)_i \;=\; u \;=\; \max\{\, (Ay)_k \mid 1 \le k \le m \,\}.$$

Proof

$$\begin{aligned} x^\top Ay &= \sum_{i=1}^m x_i\,(Ay)_i = \sum_{i=1}^m x_i\,(u - (u - (Ay)_i)) \\ &= \sum_{i=1}^m x_i\,u - \sum_{i=1}^m x_i\,(u - (Ay)_i) \\ &= u - \sum_{i=1}^m \underbrace{x_i}_{\ge 0}\,\underbrace{(u - (Ay)_i)}_{\ge 0} \;\le\; u\,. \end{aligned}$$

Proof of the best response condition

	x is a best response to y
$\Leftrightarrow$	for all pure strategies i of player I :
	$x_i > 0 \;\Rightarrow\; (Ay)_i \;=\; u \;= \max\{\, (Ay)_k \mid 1 \le k \le m \,\}.$

Proof

$$
\begin{aligned}
x^\top Ay &= \sum_{i=1}^{m} x_i\,(Ay)_i = \sum_{i=1}^{m} x_i\,(u - (u - (Ay)_i)) \\
&= \sum_{i=1}^{m} x_i\,u - \sum_{i=1}^{m} x_i\,(u - (Ay)_i) \\
&= u - \sum_{i=1}^{m} \underbrace{x_i}_{\ge 0}\;\underbrace{(u - (Ay)_i)}_{\ge 0} \;\le\; u\,.
\end{aligned}
$$

So $x^\top Ay = u \;\Leftrightarrow\; x_i > 0$ implies $u - (Ay)_i = 0$, as claimed.

best response condition written as orthogonality = complementarity

$\boldsymbol{x}$ is a best response to $\boldsymbol{y}$

$\Leftrightarrow$

$$\boxed{\boldsymbol{x} \geq \mathbf{0} \quad \perp \quad \boldsymbol{A}\boldsymbol{y} \leq \mathbf{1}\boldsymbol{u}}$$

$:\Leftrightarrow$

$$\boldsymbol{x} \geq \mathbf{0}\,, \quad \boldsymbol{A}\boldsymbol{y} \leq \mathbf{1}\boldsymbol{u}\,, \quad \boldsymbol{x}^\top(\mathbf{1}\boldsymbol{u} - \boldsymbol{A}\boldsymbol{y}) = \mathbf{0}$$

Convex combinations

Lines and line segments

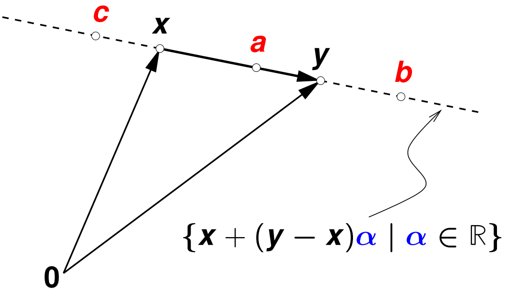

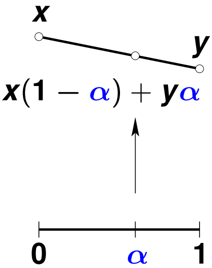

Line through points $\boldsymbol{x}$ and $\boldsymbol{y}$ given by $\boldsymbol{x} + (\boldsymbol{y} - \boldsymbol{x})\alpha$ for $\alpha \in \mathbb{R}$.

Examples: $\boldsymbol{a}$ for $\alpha = 0.6$, $\boldsymbol{b}$ for $\alpha = 1.5$, $\boldsymbol{c}$ for $\alpha = -0.4$.

Line segment that connects $\boldsymbol{x}$ and $\boldsymbol{y}$ $\Leftrightarrow$ $0 \le \alpha \le 1$.

Convexity

Rewrite $\boldsymbol{x} + (\boldsymbol{y} - \boldsymbol{x})\alpha$ as

$$\boldsymbol{x}(\mathbf{1} - \alpha) + \boldsymbol{y}\alpha$$

which for $\alpha \in [\mathbf{0}, \mathbf{1}]$ is called a **convex combination** of $\boldsymbol{x}$ and $\boldsymbol{y}$.

Convexity

Rewrite $\boldsymbol{x} + (\boldsymbol{y} - \boldsymbol{x})\alpha$ as

$$\boldsymbol{x}(1 - \alpha) + \boldsymbol{y}\alpha$$

which for $\alpha \in [0, 1]$ is called a **convex combination** of $\boldsymbol{x}$ and $\boldsymbol{y}$.

Convex sets contain all convex combinations of their points:

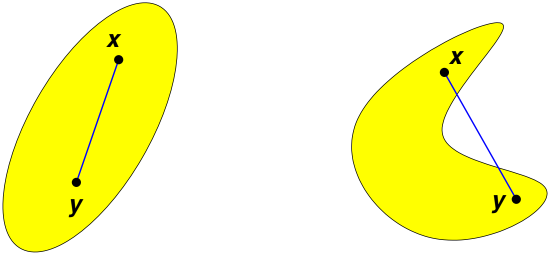

convex   not convex

Mixed strategy sets X and Y

For player I and player II,

$$
\begin{aligned}
X &= \{x \in \mathbb{R}^m \mid x \geq \mathbf{0},\ \mathbf{1}^\top x = 1\}, \\
Y &= \{y \in \mathbb{R}^n \mid y \geq \mathbf{0},\ \mathbf{1}^\top y = 1\},
\end{aligned}
$$

X and Y are **simplices**,

simplex = convex hull of unit vectors.

Mixed strategy line segment X for $m = 2$

$$X = \{x \in \mathbb{R}^m \mid x \geq 0,\ \mathbf{1}^\top x = 1\}$$

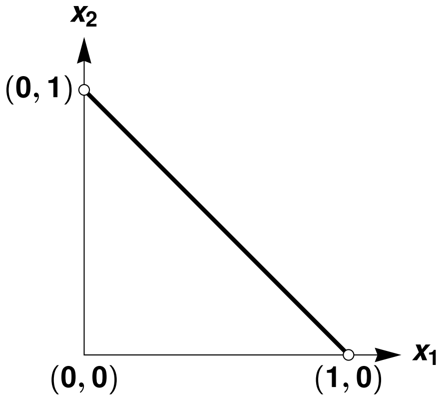

Mixed strategy triangle X for $m = 3$

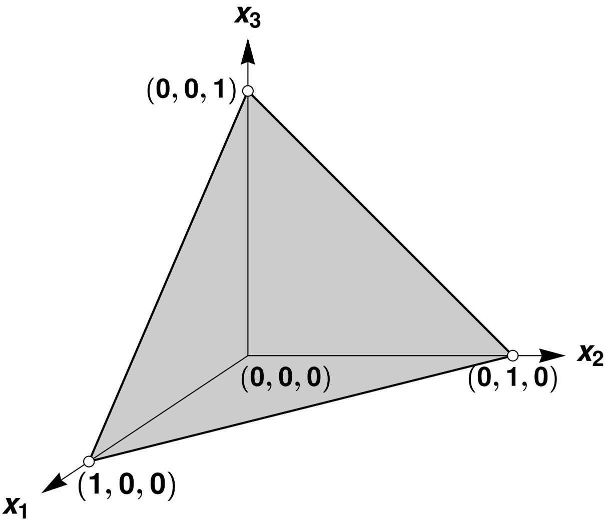

Mixed strategy tetrahedron X for $m = 4$

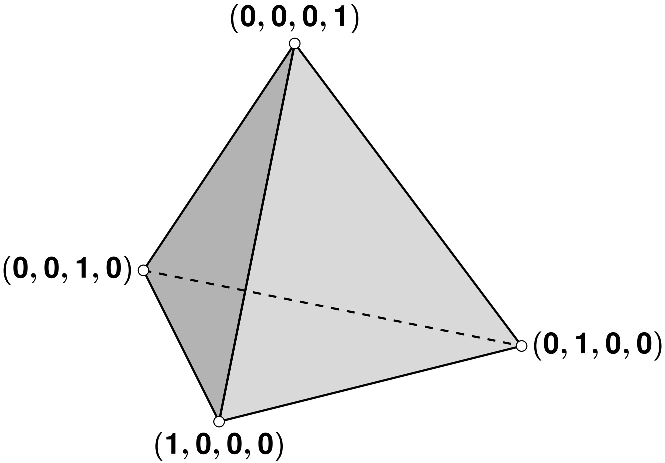

for general m called mixed strategy **simplex** X .

Zero-sum games: continued

Best responses against y

Let $y \in Y$. $(Ay)_i$ = expected payoff to player I in row i.

A **best response** $x \in X$ to y maximizes $x^\top Ay$.

$$
\begin{aligned}
& \max\{x^\top(Ay) \mid x \in X\} \\
= \; & \max\{(Ay)_1, \ldots, (Ay)_m\} \\
= \; & \min\{u \in \mathbb{R} \mid (Ay)_1 \leq u, \ldots, (Ay)_m \leq u\} \\
= \; & \min\{u \in \mathbb{R} \mid Ay \leq \mathbf{1}u\}
\end{aligned}
$$

Best responses against $\boldsymbol{y}$

Let $\boldsymbol{y} \in \boldsymbol{Y}$. $(\boldsymbol{Ay})_i$ = expected payoff to player I in row $\boldsymbol{i}$.

A **best response** $\boldsymbol{x} \in \boldsymbol{X}$ to $\boldsymbol{y}$ maximizes $\boldsymbol{x}^\top \boldsymbol{Ay}$.

$$\begin{aligned} & \max\{\boldsymbol{x}^\top(\boldsymbol{Ay}) \mid \boldsymbol{x} \in \boldsymbol{X}\} \\ = \; & \max\{(\boldsymbol{Ay})_1, \ldots, (\boldsymbol{Ay})_m\} \\ = \; & \min\{\boldsymbol{u} \in \mathbb{R} \mid (\boldsymbol{Ay})_1 \leq \boldsymbol{u}, \ldots, (\boldsymbol{Ay})_m \leq \boldsymbol{u}\} \\ = \; & \min\{\boldsymbol{u} \in \mathbb{R} \mid \boldsymbol{Ay} \leq \mathbf{1}\boldsymbol{u}\} \end{aligned}$$

In a **zero-sum game** $(\boldsymbol{A}, -\boldsymbol{A})$, player II minimizes $\boldsymbol{u}$ with her best choice of $\boldsymbol{y} \in \boldsymbol{Y}$, her **min-max strategy** $\hat{\boldsymbol{y}}$.

max-min and min-max strategies

min-max strategy $\hat{\boldsymbol{y}} \in \boldsymbol{Y}$:

$$\begin{aligned} \max_{\boldsymbol{x} \in \boldsymbol{X}} \boldsymbol{x}^\top \boldsymbol{A} \hat{\boldsymbol{y}} &= \min_{\boldsymbol{y} \in \boldsymbol{Y}} \max_{\boldsymbol{x} \in \boldsymbol{X}} \boldsymbol{x}^\top \boldsymbol{A} \boldsymbol{y} \\ &= \min_{\boldsymbol{y} \in \boldsymbol{Y}} \{ \boldsymbol{u} \in \mathbb{R} \mid \boldsymbol{A}\boldsymbol{y} \leq \mathbf{1}\boldsymbol{u} \} \end{aligned}$$

max-min strategy $\hat{\boldsymbol{x}} \in \boldsymbol{X}$:

$$\begin{aligned} \min_{\boldsymbol{y} \in \boldsymbol{Y}} \hat{\boldsymbol{x}}^\top \boldsymbol{A} \boldsymbol{y} &= \max_{\boldsymbol{x} \in \boldsymbol{X}} \min_{\boldsymbol{y} \in \boldsymbol{Y}} \boldsymbol{x}^\top \boldsymbol{A} \boldsymbol{y} \\ &= \max_{\boldsymbol{x} \in \boldsymbol{X}} \{ \boldsymbol{v} \in \mathbb{R} \mid \boldsymbol{v}\mathbf{1}^\top \leq \boldsymbol{x}^\top \boldsymbol{A} \} \end{aligned}$$

max min $\leq$ min max

The "easy part" of max-min versus min-max payoff:

$$\max_{\boldsymbol{x}\in \boldsymbol{X}} \min_{\boldsymbol{y}\in \boldsymbol{Y}} \boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{y} \leq \min_{\boldsymbol{y}\in \boldsymbol{Y}} \max_{\boldsymbol{x}\in \boldsymbol{X}} \boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{y}$$

Proof

$$\begin{aligned}
\max_{\boldsymbol{x}\in \boldsymbol{X}} \min_{\boldsymbol{y}\in \boldsymbol{Y}} \boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{y} &= \min_{\boldsymbol{y}\in \boldsymbol{Y}} \hat{\boldsymbol{x}}^\top \boldsymbol{A}\boldsymbol{y} \\
&\leq \hat{\boldsymbol{x}}^\top \boldsymbol{A}\hat{\boldsymbol{y}} \\
&\leq \max_{\boldsymbol{x}\in \boldsymbol{X}} \boldsymbol{x}^\top \boldsymbol{A}\hat{\boldsymbol{y}} \\
&= \min_{\boldsymbol{y}\in \boldsymbol{Y}} \max_{\boldsymbol{x}\in \boldsymbol{X}} \boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{y} \qquad \square
\end{aligned}$$

von Neumann's minimax theorem [1928]

Every zero-sum game $\boldsymbol{A}$ has a **value** $\boldsymbol{v}$:

$$\max_{\boldsymbol{x} \in \boldsymbol{X}} \min_{\boldsymbol{y} \in \boldsymbol{Y}} \boldsymbol{x}^\top \boldsymbol{A} \boldsymbol{y} = \boldsymbol{v} = \min_{\boldsymbol{y} \in \boldsymbol{Y}} \max_{\boldsymbol{x} \in \boldsymbol{X}} \boldsymbol{x}^\top \boldsymbol{A} \boldsymbol{y}$$

John von Neumann (1903–1957)

- set theory
- mathematics of quantum mechanics
- minimax theorem [1928], game theory
- stored-program computer

John von Neumann (1903–1957)

- set theory
- mathematics of quantum mechanics
- minimax theorem [1928], game theory
- stored-program computer

from *The Man from the Future (2021):*

"Von Neumann would carry on a conversation with my three-year-old son, and the two of them would talk as equals, and I sometimes wondered if he used the same principle when he talked to the rest of us." Edward Teller, 1966

Minimax theorem: Proof by Loomis [1946]

min-max strategy $\boldsymbol{y} \in \boldsymbol{Y}$: minimize $\boldsymbol{u}$ s.t. $\boldsymbol{Ay} \leq \mathbf{1}\boldsymbol{u}$,

max-min strategy $\boldsymbol{x} \in \boldsymbol{X}$: maximize $\boldsymbol{v}$ s.t. $\boldsymbol{v}\mathbf{1}^\top \leq \boldsymbol{x}^\top \boldsymbol{A}$,

$$\boldsymbol{v} = \boldsymbol{v}\mathbf{1}^\top \boldsymbol{y} \leq \boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{y} \leq \boldsymbol{x}^\top \mathbf{1}\boldsymbol{u} = \boldsymbol{u}.$$

Minimax theorem: Proof by Loomis [1946]

min-max strategy $\boldsymbol{y} \in \boldsymbol{Y}$: minimize $\boldsymbol{u}$ s.t. $\boldsymbol{A}\boldsymbol{y} \leq \mathbf{1}\boldsymbol{u}$,

max-min strategy $\boldsymbol{x} \in \boldsymbol{X}$: maximize $\boldsymbol{v}$ s.t. $\boldsymbol{v}\mathbf{1}^\top \leq \boldsymbol{x}^\top\boldsymbol{A}$,

$$\boldsymbol{v} = \boldsymbol{v}\mathbf{1}^\top\boldsymbol{y} \leq \boldsymbol{x}^\top\boldsymbol{A}\boldsymbol{y} \leq \boldsymbol{x}^\top\mathbf{1}\boldsymbol{u} = \boldsymbol{u}.$$

$\boldsymbol{v}\mathbf{1}^\top = \boldsymbol{x}^\top\boldsymbol{A}$ and $\boldsymbol{A}\boldsymbol{y} = \mathbf{1}\boldsymbol{u} \quad \Rightarrow \quad \boldsymbol{v} = \boldsymbol{u}$, done.

Minimax theorem: Proof by Loomis [1946]

min-max strategy $\boldsymbol{y} \in \boldsymbol{Y}$: minimize $\boldsymbol{u}$ s.t. $\boldsymbol{Ay} \le \mathbf{1}\boldsymbol{u}$,

max-min strategy $\boldsymbol{x} \in \boldsymbol{X}$: maximize $\boldsymbol{v}$ s.t. $\boldsymbol{v}\mathbf{1}^\top \le \boldsymbol{x}^\top \boldsymbol{A}$,

$$\boldsymbol{v} = \boldsymbol{v}\mathbf{1}^\top \boldsymbol{y} \le \boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{y} \le \boldsymbol{x}^\top \mathbf{1}\boldsymbol{u} = \boldsymbol{u}.$$

$\boldsymbol{v}\mathbf{1}^\top = \boldsymbol{x}^\top \boldsymbol{A}$ and $\boldsymbol{Ay} = \mathbf{1}\boldsymbol{u} \quad \Rightarrow \quad \boldsymbol{v} = \boldsymbol{u}$, done.

Assume $(\boldsymbol{Ay})_{\boldsymbol{k}} < \boldsymbol{u}$ for some row $\boldsymbol{k}$, let $\overline{\boldsymbol{A}}$ be $\boldsymbol{A}$ without row $\boldsymbol{k}$.

By **inductive hypothesis**, $\overline{\boldsymbol{A}}$ has game value $\overline{\boldsymbol{u}}$, $\overline{\boldsymbol{A}}\overline{\boldsymbol{y}} \le \mathbf{1}\overline{\boldsymbol{u}}$.
$\overline{\boldsymbol{u}} \le \boldsymbol{v}$, $\overline{\boldsymbol{u}} \le \boldsymbol{u}$, ($\overline{\boldsymbol{A}}$ better than $\boldsymbol{A}$ for minimizer).

Minimax theorem: Proof by Loomis [1946]

min-max strategy $y \in Y$: minimize u s.t. $Ay \leq \mathbf{1}u$,

max-min strategy $x \in X$: maximize v s.t. $v\mathbf{1}^\top \leq x^\top A$,

$$v = v\mathbf{1}^\top y \leq x^\top A y \leq x^\top \mathbf{1} u = u.$$

$v\mathbf{1}^\top = x^\top A$ and $Ay = \mathbf{1}u \Rightarrow v = u$, done.

Assume $(Ay)_k < u$ for some row k, let $\overline{A}$ be A without row k.

By **inductive hypothesis**, $\overline{A}$ has game value $\overline{u}$, $\overline{A}\overline{y} \leq \mathbf{1}\overline{u}$.
$\overline{u} \leq v$, $\overline{u} \leq u$, ($\overline{A}$ better than A for minimizer).

Claim : $\overline{u} = u$. Intuition: maximizer avoids row k of A anyhow.

Proof that $\overline{u} = u$

minimal u s.t. $Ay \leq \mathbf{1}u$, maximal v s.t. $v\mathbf{1}^\top \leq x^\top A$, $v \leq u$.

$(Ay)_k < u$, matrix $\overline{A}$ is A without row k, value $\overline{u} \leq v$, $\overline{u} \leq u$.

Proof that $\overline{u} = u$

minimal u s.t. $Ay \le \mathbf{1}u$, maximal v s.t. $v\mathbf{1}^\top \le x^\top A$, $\quad v \le u$.

$(Ay)_k < u$, matrix $\overline{A}$ is A without row k, value $\overline{u} \le v$, $\overline{u} \le u$.

Suppose $\overline{u} < u$. For $0 < \varepsilon \le 1$,

$$\overline{A}(\underbrace{y(1-\varepsilon) + \overline{y}\varepsilon}_{y(\varepsilon)\in Y\ (\textit{convex})}) \le \mathbf{1}(u(1-\varepsilon)+\overline{u}\varepsilon) = \mathbf{1}(u-\varepsilon(u-\overline{u})) < \mathbf{1}u$$

Proof that $\overline{u} = u$

minimal u s.t. $Ay \le \mathbf{1}u$, maximal v s.t. $v\mathbf{1}^\top \le x^\top A$, $\quad v \le u$.

$(Ay)_k < u$, matrix $\overline{A}$ is A without row k, value $\overline{u} \le v$, $\overline{u} \le u$.

Suppose $\overline{u} < u$. For $\mathbf{0} < \varepsilon \le \mathbf{1}$,

$$\overline{A}(\underbrace{y(1-\varepsilon)+\overline{y}\varepsilon}_{y(\varepsilon)\in Y\ (\textit{convex})}) \le \mathbf{1}(u(1-\varepsilon)+\overline{u}\varepsilon) = \mathbf{1}(u-\varepsilon(u-\overline{u})) < \mathbf{1}u$$

For missing row k of A and sufficiently small $\varepsilon > \mathbf{0}$:

$$(A(y(1-\varepsilon)+\overline{y}\varepsilon))_k = \underbrace{(Ay)_k}_{<u}(1-\varepsilon)+(A\overline{y})_k\varepsilon < u,$$

overall $Ay(\varepsilon) < \mathbf{1}u$, contradicting minimality of u. Hence $\overline{u} = u$.

Proof that $\overline{u} = u$

minimal u s.t. $Ay \leq \mathbf{1}u$, maximal v s.t. $v\mathbf{1}^\top \leq x^\top A$, $\quad v \leq u$.

$(Ay)_k < u$, matrix $\overline{A}$ is A without row k, value $\overline{u} \leq v$, $\overline{u} \leq u$.

Suppose $\overline{u} < u$. For $\mathbf{0} < \varepsilon \leq \mathbf{1}$,

$$\overline{A}(\underbrace{y(1-\varepsilon) + \overline{y}\varepsilon}_{y(\varepsilon)\in Y\ (\textbf{convex})}) \leq \mathbf{1}(u(1-\varepsilon)+\overline{u}\varepsilon) = \mathbf{1}(u-\varepsilon(u-\overline{u})) < \mathbf{1}u$$

For missing row k of A and sufficiently small $\varepsilon > \mathbf{0}$:

$$(A(y(1-\varepsilon) + \overline{y}\varepsilon))_k = \underbrace{(Ay)_k}_{<u}(1-\varepsilon) + (A\overline{y})_k\varepsilon < u,$$

overall $Ay(\varepsilon) < \mathbf{1}u$, contradicting minimality of u. Hence $\overline{u} = u$.

$\Rightarrow \quad \overline{u} \leq v \leq u = \overline{u}$, $\boxed{v = u}$. Induction complete. □

minimax pair = saddle point

Minimax theorem and Nash equilibrium

$$\boxed{\max_{x \in X} \min_{y \in Y} x^\top A y = \min_{y \in Y} \max_{x \in X} x^\top A y}$$

Minimax theorem and Nash equilibrium

$$\boxed{\max_{x \in X} \min_{y \in Y} x^\top A y = \min_{y \in Y} \max_{x \in X} x^\top A y}$$

with max-min strategy $\hat{x}$ and min-max strategy $\hat{y}$:

$$\min_{y \in Y} \hat{x}^\top A y = \hat{x}^\top A \hat{y} = \max_{x \in X} x^\top A \hat{y}$$

$$\Leftrightarrow \quad \forall y \in Y,\ x \in X : \quad \hat{x}^\top A y \geq \hat{x}^\top A \hat{y} \geq x^\top A \hat{y}$$

$\Leftrightarrow$ $(\hat{x}, \hat{y})$ is a **Nash equilibrium**
($\hat{x}$ and $\hat{y}$ are mutual best responses)

Minimax theorem and Nash equilibrium

$$\max_{x \in X} \min_{y \in Y} x^\top A y = \min_{y \in Y} \max_{x \in X} x^\top A y$$

with max-min strategy $\hat{x}$ and min-max strategy $\hat{y}$:

$$\min_{y \in Y} \hat{x}^\top A y = \hat{x}^\top A \hat{y} = \max_{x \in X} x^\top A \hat{y}$$

$$\Leftrightarrow \quad \forall y \in Y,\ x \in X : \quad \hat{x}^\top A y \geq \hat{x}^\top A \hat{y} \geq x^\top A \hat{y}$$

$\Leftrightarrow$ $(\hat{x}, \hat{y})$ is a **Nash equilibrium**
($\hat{x}$ and $\hat{y}$ are mutual best responses)

Exercise: prove that if $(\overline{x}, \overline{y})$ is a Nash equilibrium in the zero-sum game $(A, -A)$, then $\overline{x}$ is a max-min strategy and $\overline{y}$ is a min-max strategy.

Consequences for zero-sum games

Zero-sum game:

equilibrium strategy = max-min / min-max strategy

- equilibrium **exists** $\Leftrightarrow$ $\boxed{\max\min = \min\max}$

- strategies are **optimal**, independent of opponent

- unique equilibrium payoff / cost $\boldsymbol{v}$ $=$ **value** of the game

- $(\boldsymbol{x}, \boldsymbol{y}), (\overline{\boldsymbol{x}}, \overline{\boldsymbol{y}})$ equilibria $\Rightarrow$ $(\overline{\boldsymbol{x}}, \boldsymbol{y}), (\boldsymbol{x}, \overline{\boldsymbol{y}})$ equilibria **(exchangeability)**

- $\boldsymbol{x}, \overline{\boldsymbol{x}}$ equilibrium strategy $\Rightarrow$ so is $\boldsymbol{x}(1-\alpha) + \overline{\boldsymbol{x}}\alpha$ for $\alpha \in [0, 1]$ **(convexity)**