Multiagent decision-making and control

Zero-sum games

Maryam Kamgarpour

Professor of Engineering (IGM, STI), EPFL

ME-429, Spring 2025, EPFL

Course topics

1. Static games
2. Zero-sum games
3. Potential games
4. Dynamic games, dynamic programming principle
5. Dynamic games, dynamic programming for games
6. Dynamic games, linear quadratic games, Markov games
7. Convex games, Nash equilibria characterization
8. Convex games, Nash equilibria computation
9. Auctions
10. Bayesian games
11. Learning in games
12. Extensive form games
13. Feedback games in extensive form
14. Final project presentations

Plan of today's lecture

- Game play: custom game
- Math review: See hand-written notes in class
    - fill in the gap: proof of existence of mixed strategy NE in finite action games
    - review: results on linear program basic solution
    - fill in the gap: proof of completely mixed NE computation from last lecture
- Project information
- Zero-sum games

Review Exercise: Custom Declaration

You are arriving at Geneva international airport and have brought some food from abroad

- Declaring the food costs 10 CHF
- The fine for not declaring it is 20 CHF
- Customs can stop you, but that has a personnel cost of 1 CHF.

3 since, no pure strategy NE & 2 actions only, we can use LP to find mixed strategy

Exercise

1. Is there a pure strategy Nash Equilibrium? No
2. What are the security strategies for each player?
3. How would you compute a mixed strategy Nash Equilibrium?

$\min_i \max_j a_{ij}$: P1

	Declare	Smuggle
Check	$(1, 10)$	$(-19, 20)$
Don't check	$(0, 10)$	$(0, 0)$

$\min_j \max_i b_{ij}$: P2

2 security strategies: (Don't check : P1, Declare : P2)

$$A z^* = p^* 1 \quad , \quad z_1^* + z_2^* = 1$$

$$\Rightarrow \begin{bmatrix} 1 & -19 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} z_1^* \\ 1 - z_1^* \end{bmatrix} = \begin{matrix} p^* \\ p^* \end{matrix} \qquad (1)$$

$$y^{*T} B = q^* 1^T \quad , \quad y_1^* + y_2^* = 1$$

$$\Rightarrow [y_1^* \quad 1 - y_1^*] \begin{bmatrix} 10 & 20 \\ 10 & 0 \end{bmatrix} = [q^* \quad q^*] \qquad (2)$$

$$(1): \left.\begin{matrix} 20 z_1^* - 19 = p^* \\ 0 = p^* \end{matrix}\right\} \Rightarrow z_1^* = \frac{19}{20}$$

$$(2): \left.\begin{matrix} q^* = 10 \\ 20 y_1^* = q^* \end{matrix}\right\} \Rightarrow y^* = \frac{1}{2}$$

Two-Person Zero-Sum Games

Zero-sum games

Two-person games in which the two players have opposite payoffs.

Static game with $B = -A$ (we only indicate one matrix, A)

Column player's actions

Row player's actions

$$\begin{array}{|ccc|} a_{11} & a_{12} & a_{1n} \\ a_{21} & a_{22} & a_{2n} \\ \vdots & & \\ a_{m1} & a_{m2} & a_{mn} \end{array}$$

Payoff matrix

- Row player loses a_{ij}
- Column player gains a_{ij}
- Row player minimizes outcome V
- Column player maximizes outcome V

Why?

- Zero-sum games model a large number of **practical applications**: robust optimization and control, chess, tic-tac-toe
- Nash equilibria in zero-sum games have many **useful properties**
- Nash equilibria in zero-sum games are much **easier to compute**

Examples we have seen

Consider only one round of the game.

$A =$

	Rock	Paper	Scissors
Rock			
Paper			
Scissors			

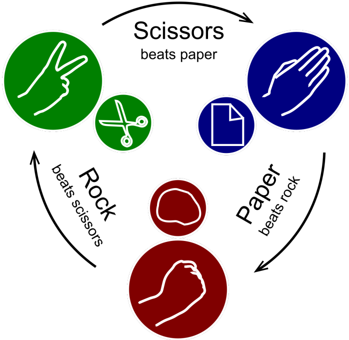

Matching pennies

	Head	Tail
Head	$(1, -1)$	$(-1, 1)$
Tail	$(-1, 1)$	$(1, -1)$

Security levels and strategies

The **security level** of P_1 (minimizer) is defined by

$$\overline{V} := \min_{i \in \{1,\dots,n\}} \max_{j \in \{1,\dots,m\}} a_{ij}$$

$\overline{V}$ is the best possible outcome (**lowest** number) that P_1 can *guarantee* himself, for any **adversarial** choice P_2 can make.

The **security strategy** of P_1 (minimizer) is defined by

$$\bar{i} \in \arg\min_{i \in \{1,\dots,n\}} \max_{j \in \{1,\dots,m\}} a_{ij}$$

Exact same definition as in general non-zero-sum games.
Remember: security levels do not depend on the other players' outcome matrix.

However, this time the other player is truly adversarial (zero-sum)!

Security levels and strategies

The **security level** of P_2 (maximizer) is defined by

$$\underline{V} := \max_{j\in\{1,\dots,n\}} \min_{i\in\{1,\dots,m\}} a_{ij}$$

$\underline{V}$ is the best possible outcome (**highest** number) that P_2 can *guarantee* himself, for any **adversarial** choice P_1 can make.

The **security strategy** of P_2 (maximizer) is defined by

$$\underline{j} \in \arg\max_{j\in\{1,\dots,n\}} \min_{i\in\{1,\dots,m\}} a_{ij}$$

n m

Just a different formulation of the same definition, for $B = -A$

Security levels and strategies

Min-Max Property

For every finite matrix A, the following properties hold:

(*i*) Security levels are well defined and unique

(*ii*) Both players have security strategies (not necessarily unique)

(*iii*) The security levels always satisfy

$$\overbrace{\underline{V} := \max_{j \in \{1,\dots,n\}} \min_{i \in \{1,\dots,m\}} a_{ij}}^{P_2} \leq \overbrace{\overline{V} := \min_{i \in \{1,\dots,m\}} \max_{j \in \{1,\dots,n\}} a_{ij}}^{P_1}$$

Proof: $\underline{V} \leq \overline{V}$ for all matrices A

We will show a more general statement: Let $\mathcal{Y} \subset \mathbb{R}^n$ and $\mathcal{Z} \subset \mathbb{R}^n$, $J : \mathcal{Y} \times \mathcal{Z} \to \mathbb{R}$.

Min-Max Property for General functions

The sup-inf inequality holds: $\sup_{z \in \mathcal{Z}} \inf_{y \in \mathcal{Y}} J(y, z) \leq \inf_{y \in \mathcal{Y}} \sup_{z \in \mathcal{Z}} J(y, z)$.
Furthermore, if $\mathcal{Y}$, $\mathcal{Z}$ are closed and bounded and J is continuous then
$\max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} J(y, z) \leq \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} J(y, z)$.

Proof

$$\inf_{y \in \mathcal{Y}} J(y,z) \leq J(y,z) \quad \forall y, z$$

$$\Rightarrow \inf_{y \in \mathcal{Y}} J(y,z) \leq \sup_{z \in \mathcal{Z}} J(y,z) \quad \forall y, z$$

$$\Rightarrow \sup_{z \in \mathcal{Z}} \inf_{y \in \mathcal{Y}} J(y,z) \leq \sup_{z \in \mathcal{Z}} J(y,z) \; \forall y, \text{ (supremum is the least upper bound)}$$

$$\Rightarrow \sup_{z \in \mathcal{Z}} \inf_{y \in \mathcal{Y}} J(y,z) \leq \inf_{y \in \mathcal{Y}} \sup_{z \in} J(y,z). \text{ (since infimum is the greatest lower bound)}$$

Exercise: Rock-Paper-Scissors

- What are the security levels of the Rock-Paper-Scissors game?
- What is the gap between $\overline{V}$ and $\underline{V}$?
- What are the security strategies?
- What is the outcome if both players play a security strategy?
- Given a matrix A, how do you compute $\overline{V}$ and $\underline{V}$ in Matlab?

```
Vover = min(max(A'))

Vunder = max(min(A))
```

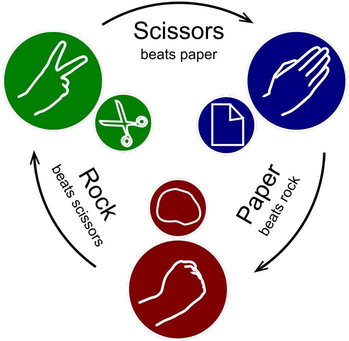

Competence

Evaluate fundamental quantities in a given game and **verify that they are compatible with what the theory says**.

Example: Feedback Control

An engineer must choose a proportional feedback controller gain, K_p, for a system made with components specified to coarse tolerances. For such a resistor R in the system, the resulting feedback performance metric (to be minimized) is as shown:

Engineer \ Nature	$R = 3.5\Omega$	$R = 4.0\Omega$	$R = 4.5\Omega$
$K_p = 1$	6.0	4.0	3.0
$K_p = 5$	5.0	4.0	3.5
$K_p = 10$	$+\infty$	20.0	2.0

- What are $\overline{V}$, $\underline{V}$, $\bar{i}$, $\underline{j}$?
- What is the interpretation in control design terms?

$\overline{V} = 5$, $\bar{i} = 2$ $(K_p = 5)$

$\underline{V} = 5$, $\underline{j} = 1$ $R = 3.5$

Nash equilibrium in zero-sum games

Given a zero-sum game described by the payoff matrix A, we say that the pair of actions $\gamma_{i^*} \in \Gamma$ and $\sigma_{j^*} \in \Sigma$ are a **Nash Equilibrium** if

$$a_{i^*j^*} \leq a_{ij^*} \quad \forall i = 1, \ldots, m \qquad \text{(lowest outcome of column } j^*\text{)}$$

and

$$a_{i^*j^*} \geq a_{i^*j} \quad \forall j = 1, \ldots, n \qquad \text{(highest outcome of row } i^*\text{)}$$

Also known as **saddle-point equilibrium**

$$a_{i^*j} \leq a_{i^*j^*} \leq a_{ij^*}$$

$$\forall i \in \{1, \ldots, m\}, \quad \forall j \in \{1, \ldots, n\}$$

We call $V^* := a_{i^*j^*}$ the **saddle-point value**.

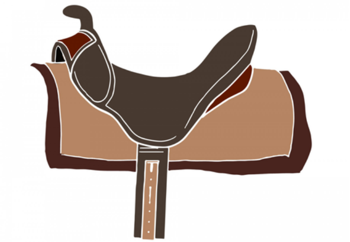

Saddle-point and security levels

Not all zero-sum games have a saddle point (think of Rock-Paper-Scissors).

We can exactly characterize the zero-sum games that have a saddle point.

Theorem (Saddle-point and security levels)

A zero-sum game defined by A has a saddle-point equilibrium <u>*if and only if*</u>

$$\overbrace{\underline{V} := \max_{j\in\{1,\dots,n\}} \min_{i\in\{1,\dots,m\}} a_{ij}}^{P_2} = \overbrace{\overline{V} := \min_{i\in\{1,\dots,m\}} \max_{j\in\{1,\dots,n\}} a_{ij}}^{P_1}$$

and $\underline{V} = \overline{V}$ is the saddle-point value.

Moreover, if $\bar{i}$ and $\underline{j}$ are security strategies, then $(\bar{i}, \underline{j})$ is a saddle-point equilibrium.

Extremely simple Matlab test: `max(min(A)) == min(max(A'))`

Proof: Saddle-point and security levels

We will break it into 2 parts, "⇒" and "⇐"

Start with "⇐": if a (Nash equilibrium) saddle-point exists, then $\underline{V} = \overline{V}$.

- By definition of a saddle-point, $a_{i^*j} \overbrace{\leq}^{(ii)} a_{i^*j^*} \overbrace{\leq}^{(i)} a_{ij^*}, \forall i, j$ holds

$$a_{i^*j^*} \underbrace{=}_{(i)} \min_i \; a_{ij^*} \leq \max_j \; \min_i \; a_{ij} = \underline{V}$$

$$a_{i^*j^*} \underbrace{=}_{(ii)} \max_j \; a_{i^*j} \geq \min_i \; \max_j \; a_{ij} = \overline{V}$$

- We just showed that that $\overline{V} \leq \underline{V}$.
- From the previous proposition we have that $\underline{V} \leq \overline{V}$.
- Therefore it needs to be $\underline{V} = \overline{V}$.

Proof: Saddle-point and security levels

Let us prove "$\Rightarrow$": if $\underline{V} = \overline{V}$, then a Nash equilibrium (saddle point) exists.

- Consider the two security strategies $\bar{i}$ and $\underline{j}$. We have that

$$\overline{V} = \min_i \max_j a_{ij} = \max_j a_{\bar{i}j} \quad \text{and} \quad \underline{V} = \max_j \min_i a_{ij} = \min_i a_{i\underline{j}}$$

- By definition of min and max, we have

$$\underline{V} = \min_i a_{i\underline{j}} \leq a_{\bar{i}\underline{j}} \leq \max_j a_{\bar{i}j} = \overline{V}$$

- Because $\overline{V} = \underline{V}$, then these inequalities must be equalities:

$$\min_i a_{i\underline{j}} = a_{\bar{i}\underline{j}} = \max_j a_{\bar{i}j}.$$

  Hence, $a_{\bar{i}j} \leq a_{\bar{i}\underline{j}} \leq a_{i\underline{j}}$, and $(\bar{i}, \underline{j})$ is a saddle point (Nash equilibrium).

Value of a game and order interchangeability

Remember: security strategies are not unique

In the theorem we showed that if $(\bar{i}, \underline{j})$ is a saddle-point equilibrium then its value is $V^* = \overline{V} = \underline{V}$.

Important consequences follow (only for zero-sum games!)

Value of a game

All saddle-point equilibria (Nash equilibria) of a zero-sum game have the same value V^*, which we denote as the **value of the game**.

Order interchangeability

If (i_1^*, j_1^*) and (i_2^*, j_2^*) are saddle-point equilibria for a zero-sum game, then (i_1^*, j_2^*) and (i_2^*, j_1^*) are also saddle-point equilibria.

proof

$$V^* = a_{i_1^* j_1^*} \leq a_{i_1^* j_2^*} \leq a_{i_2^* j_2^*} = V^*$$

by $a_{i_1^* j_1^*}$ being Nash equilib

by $a_{i_2^* j_2^*}$ being a Nash equilibrium

Mixed Strategies

Let us recall the definition of **mixed strategies**.

A mixed strategy for P_1 is a vector of numbers $(y_1, \dots, y_m)$ chosen from the simplex

$$\mathcal{Y} = \left\{ (y_1, \dots, y_n) \ : \ \sum_{i=1}^{n} y_i = 1, \ y_i \geq 0, \ i = 1, \dots, n \right\}$$

where y_i is the probability with which P_1 selects action $i \in \{1, \dots, m\}$

A mixed strategy for P_2 is a vector $(z_1, \dots, z_n)$ chosen from the simplex

$$\mathcal{Z} = \left\{ (z_1, \dots, z_n) \ : \ \sum_{j=1}^{n} z_j = 1, \ z_j \geq 0, \ j = 1, \dots, n \right\}$$

where z_j is the probability with which P_2 selects action $j \in \{1, \dots, n\}$

Mixed security strategies

- The **mixed security level** for P_1 (the minimizer) is defined as

$$
\begin{aligned}
\overline{V}_{\mathrm{m}} &:= \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} \; y^\top A z \\
&= \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} \; \sum_{i=1}^{m} \sum_{j=1}^{n} \; y_i z_j a_{ij}
\end{aligned}
$$

- A **mixed security strategy** is any

$$
\overline{y} \in \arg\min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} \; y^\top A z
$$

Such a $\overline{y}$ minimizes the expected game value for the worst possible choice of mixed strategy z that P_2 can make.

Mixed security strategies

- The **mixed security level** for P_2 (maximizer) is defined as

$$\underline{V}_{\mathrm{m}} := \max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} \; y^\top A z$$
$$= \max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} \; \sum_{i=1}^{n} \sum_{j=1}^{m} \; y_i z_j a_{ij}$$

- A **mixed security strategy** is any

$$\underline{z} \in \arg\max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} \; y^\top A z$$

Such a $\underline{z}$ maximizes the expected game value for the worst possible choice of mixed strategy y that P_1 can make.

Background for computing mixed security strategies

If $\mathcal{X}$ is the simplex defined as follows:

$$\mathcal{X} := \left\{ x \in \mathbb{R}^n \; : \; \sum_{i=1}^{n} x_i = 1, x_i \geq 0, i = 1, \ldots, n \right\}$$

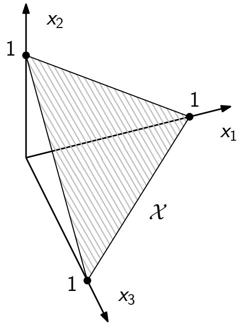

Then these two optimization problems are equivalent

$$\max_{x \in X} a^T x = \max_{x \in \mathcal{X}} \sum_{i=1}^{n} x_i a_i \quad \Longleftrightarrow \quad \max_{i \in \{1, \ldots, n\}} a_i$$

since the optimizer x^* lies at a vertex of the simplex X.

$a \in \mathbb{R}^n$, $x \in \mathbb{R}^n$

$x_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, x_2 \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, x_3 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$

Computing the mixed security level

$$\overline{V}_{\mathrm{m}} = \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} \sum_{i=1}^{m} \sum_{j=1}^{n} y_i z_j a_{ij}$$

- Do you have to search the $\max$ in the entire simplex $\mathcal{Z}$?

for a fixed y

$$\max_{z \in \mathcal{Z}} \sum_{i=1}^{m} \sum_{j=1}^{n} y_i z_j a_{ij} = \max_{j} \sum_{i=1}^{m} y_i a_{ij}$$

$\in \{1, \dots, n\}$

In other words, the worst case is always a pure strategy of Player 2.

How can we use the above to compute $\overline{V}_{\mathrm{m}}$ and $\overline{y}$?

Computing the mixed security level via LP

$$\overline{V}_{\mathrm{m}} := \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} \sum_{i=1}^{m} \sum_{j=1}^{n} y_i z_j a_{ij} = \min_{y \in \mathcal{Y}} \max_{j} \left(\sum_{i=1}^{m} y_i a_{ij} \right)$$

which is equivalent to the following linear program

$$\| \quad \min_{y \in \mathcal{Y}} V_m$$

$$\begin{aligned} \min_{y, V_{\mathrm{m}}} \quad & V_{\mathrm{m}} \\ \text{subject to} \quad & \sum_{i=1}^{m} y_i a_{ij} \leq V_{\mathrm{m}}, \quad j = 1, \ldots, n \\ & y \in \mathcal{Y},\ V_{\mathrm{m}} \in \mathbb{R} \end{aligned}$$

or in compact form

$$\begin{aligned} \min_{y, V_{\mathrm{m}}} \quad & V_{\mathrm{m}} \\ \text{subject to} \quad & A^{\top} y \leq V_{\mathrm{m}} \mathbf{1} \\ & y \in \mathcal{Y},\ V_{\mathrm{m}} \in \mathbb{R} \end{aligned}$$

where $\mathbf{1} = (1, \ldots, 1)^{\top} \in \mathbb{R}^n$ is the vector of ones.

Computing the mixed security level via linear programming (LP)

```
% define the game
m = 3; n = 2;
A = rand(m,n);

% evaluate security levels
Vover = min(max(A'));
Vunder = max(min(A));
Vover==Vunder   % check if a pure NE exists


% set up the LP for finding the mixed NE
f = [zeros(m,1); 1];     % objective function
LB = [zeros(m,1); -inf];    % lower bound
UB = [ones(m,1);inf];       % upper bound
Aeq = [ones(m,1); 0]';      % equality constraints
beq = 1;
Ain = [A', -1*ones(n,1)];   % inequality constraints
bin = zeros(n,1);

[yopt, V_m] = linprog(f,Ain,bin,Aeq,beq,LB,UB);
```

Exercise 1

Write the code for finding the mixed Nash equilibrium strategy of the second player, z^*. Show that alternatively, you can use the dual of the above linear program to find z^*.

Min-Max Property

Min-Max Property

For every (finite) matrix A, the following properties hold:

- Average security levels are well defined and unique
- Both players have mixed security strategies (not necessarily unique)
- The average security levels always satisfy the following inequalities:

$$\underbrace{\max_j \min_i a_{ij}}_{\underline{V}} \leq \underbrace{\max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} y^\top A z}_{\underline{V}_{\mathrm{m}}} \leq \underbrace{\min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} y^\top A z}_{\overline{V}_{\mathrm{m}}} \leq \underbrace{\min_i \max_j a_{ij}}_{\overline{V}}$$

The result is partly intuitive, and can be proven similarly to the previous results. We will prove the first and the second inequalities.

we saw in slide 10

now

Min-Max Property Proof (1)

Let us prove that

$$\underbrace{\underline{V}}_{\max_j \min_i a_{ij}} \leq \underbrace{\underline{V}_{\mathrm{m}}}_{\max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} y^\top A z}$$

We saw that the $\min$ of a linear function over a simplex is always achieved at one of the vertices (corners), and therefore

$$\max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} y^\top A z = \max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} \sum_{i=1}^{m} \sum_{j=1}^{n} a_{ij} y_i z_j = \max_{z \in \mathcal{Z}} \min_{i \in \{1,\ldots,m\}} \Big(\sum_{j=1}^{n} a_{ij} z_j \Big)$$

If we "restrict" the maximization from the simplex $\mathcal{Z}$ to only it's corners ($j \in \{1, \ldots, n\}$ pure strategies), we necessarily obtain a lower value. Therefore

$$\max_{z \in \mathcal{Z}} \min_{i \in \{1,\ldots,m\}} \Big(\sum_{j=1}^{n} a_{ij} z_j \Big) \geq \underbrace{\max_{j \in \{1,\ldots,n\}} \min_{i \in \{1,\ldots,m\}} a_{ij}}$$

which completes this part of the proof.

security level
of P2 in pure strategy

Min-Max Property Proof (2)

How would you verify the next inequality? $\underbrace{\max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} y^\top A z}_{\underline{V}_{\mathrm{m}}} \leq \underbrace{\min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} y^\top A z}_{\overline{V}_{\mathrm{m}}}$?

follows from slide 10

Mixed Strategy Equilibrium

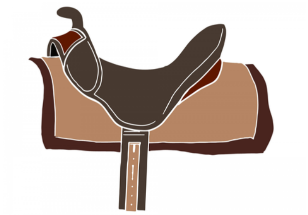

Mixed Nash equilibrium for zero-sum games

A pair of strategies $(y^*, z^*) \in \mathcal{Y} \times \mathcal{Z}$ is called a mixed-strategy saddle-point equilibrium (or Nash equilibrium) if

$$y^{*\top} A z^* \leq y^\top A z^*, \quad \forall y \in \mathcal{Y} \quad \text{(the minimizer)}$$

$$y^{*\top} A z^* \geq y^{*\top} A z, \quad \forall z \in \mathcal{Z} \quad \text{(the maximizer)}$$

$y^{*\top} A z^*$ is called the **saddle point value**.

$\forall z \in \mathcal{Z} \quad y^{*\top} A z \leq y^{*\top} A z^* \leq y^\top A z^* \quad \forall y \in \mathcal{Y}$

Mixed Saddle Point vs. Security Levels

We proved earlier that in the case of **pure strategies** a **pure** saddle-point equilibrium exists if and only if $\underline{V} = \overline{V}$.

An analogous result holds for **mixed strategies**, and can be proved in a very similar fashion.

Theorem (Mixed saddle point and mixed security levels)

*A zero-sum game has a **mixed saddle-point equilibrium** if and only if*

$$\underline{V}_m = \max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} \; y^\top A z = \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} \; y^\top A z = \overline{V}_m$$

If this condition holds, then

- *the saddle-point equilibrium corresponds to the mixed security strategies* $(\overline{y}, \underline{z})$ *satisfying*

$$\overline{y}^\top A z \le \overline{y}^\top A \underline{z} \le y^\top A \underline{z} \quad \forall y \in \mathcal{Y}, z \in \mathcal{Z}$$

- $\underline{V}_m = \overline{V}_m$ *is the saddle point value.*

V ″ value of the game

Computing mixed Nash equilibria

But we know from **Nash theorem** that a mixed Nash equilibrium always exists

$$\Rightarrow \underline{V}_{\mathrm{m}} = \overline{V}_{\mathrm{m}} \quad \text{and} \quad (\bar{i}, \underline{j}) \text{ are mixed Nash equilibria}$$

The correspondence between **mixed security strategies** and **mixed Nash equilibria** is the fundamental reason why zero-sum games are important

- We can compute mixed Nash equilibria from the security strategies
- It's just a Linear Program! (compare to non-zero-sum games)
- All Nash equilibria have the same value ($\rightarrow$ **value of the game**)
- All Nash equilibria are admissible
- All Nash equilibria are interchangeable

verify this

Competence

Compare and contrast zero-sum and non-zero-sum games in terms of theoretical results and computational methods.

Historical note

We proved that $\underline{V}_{\mathrm{m}} = \overline{V}_{\mathrm{m}}$ using Nash Theorem.

Historically, things went the other way around.

John von Neumann first proved the Minimax Theorem in 1928 for zero-sum games.
Von Neumann gave several proofs of this result, some geometric (supporting hyperplane theorem) proofs and some based on Brouwer's fixed point theorem.

He later wrote:

> *"As far as I can see, there could be no theory of games . . . without that theorem . . . I thought there was nothing worth publishing until the Minimax Theorem was proved."*

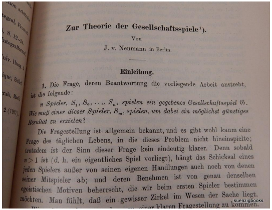

Zur Theorie der Gesellschaftsspiele[1]).

Von

J. v. Neumann in Berlin.

Einleitung.

1. Die Frage, deren Beantwortung die vorliegende Arbeit anstrebt, ist die folgende:

n Spieler, $S_1, S_2, \ldots, S_n$, spielen ein gegebenes Gesellschaftsspiel $\mathfrak{G}$. Wie muß einer dieser Spieler, S_m, spielen, um dabei ein möglichst günstiges Resultat zu erzielen?

Die Fragestellung ist allgemein bekannt, und es gibt wohl kaum eine Frage des täglichen Lebens, in die dieses Problem nicht hineinspielte; trotzdem ist der Sinn dieser Frage kein eindeutig klarer. Denn sobald $n > 1$ ist (d. h. ein eigentliches Spiel vorliegt), hängt das Schicksal eines jeden Spielers außer von seinen eigenen Handlungen auch noch von denen seiner Mitspieler ab; und deren Benehmen ist von genau denselben egoistischen Motiven beherrscht, die wir beim ersten Spieler bestimmen möchten. Man fühlt, daß ein gewisser Zirkel im Wesen der Sache liegt.

John von Neumann

Historical note

In 1949, John Nash generalized von Neumann's result to nonzero-sum games with any number of players.

His proof uses the fixed point theorem, is just one page long, and it won him the Nobel prize.

Apparently when Nash explained his result to von Neumann, the latter said:

> *"That's trivial, you know. That's just a fixed point theorem."*

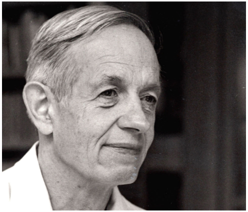

Von Neumann's original minimax theorem is easier (if you consider the complexity of fixed-point theorems in John Nash's proof... Kakutani's fixed point theorem was proved in 1941!). See Lecture 5 of

An Introductory Course in Noncooperative Game Theory
João P. Hespanha

Summary

- Definition of zero-sum games (ZSGs)
- (Pure) security levels and strategies in a ZSG
- Min-Max property of pure security levels in a ZSG
- (Pure) Nash equilibria in a ZSG
- Theorem: Pure Nash equilibria and security levels in a ZSG
- Value of a ZSG
- Order interchangeability of Nash equilibria in a ZSG
- Mixed security levels and strategies in a ZSG
- Computing mixed security levels in a ZSG
- Min-Max property of mixed security levels in a ZSG
- Mixed Nash equilibria in a ZSG
- Theorem: Mixed Nash equilibria and mixed security strategies in a ZSG
- Computing mixed Nash equilibria in a ZSG