نظریه بازیها
Game Theory

ارائه کننده: امیرحسین نیکوفرد
مهندسی برق و کامپیوتر دانشگاه خواجه نصیر

دانشگاه صنعتی خواجه نصیرالدین طوسی

What is Game Theory? (Recap)

1. The **players** are the agents that make decisions (At least 2 Players), with $N := \{1,\ldots,N\}$ denoting the Players set.
2. The **actions** available to each player at each decision point. The decision or action variable of Player i is denoted by $x_i \in X_i$,where X_i is the action set of Player i
3. The information structure specifies what each player knows before making each decision
4. The objective specifies the payoffs of each in the game

Example(Rock-Paper-Scissors):(Recap)

- A game is specified by: players (1…N), actions, and (expected) payoff matrices (functions of joint actions)

B's action

A's action

	R	P	S
R	0	−1	+1
P	+1	0	−1
S	−1	+1	0

A's payoff

	R	P	S
R	0	+1	−1
P	−1	0	+1
S	+1	−1	0

B's payoff

- If payoff matrices are identical, A and B are cooperative, else non-cooperative (zero-sum = purely competitive)

Non-cooperative Finite Games: Two-Person Zero-Sum

Material

- Dynamic Noncooperative Game Theory: Second Edition
  - Chapter 2. 1 and 2. 2
- An Introductory Course in Noncooperative Game Theory
  - Chapter 3, 6. 4, 6. 5

Matrix games

- Two-player simultaneous move games (both zero and non-zero sum types) can be written in matrix form (also called strategic form) as shown below.
- The strategies of one player form the rows of the matrix, while the strategies of the other player form the columns. Each entry in the matrix represents a possible outcome based on a corresponding selection of strategies.

<table>
<tr><td colspan="2" rowspan="2">This is an example of a two-player matrix game where each player has a choice of two possible strategies.</td><td colspan="2">Column Player (player 2)</td></tr>
<tr><td>A</td><td>B</td></tr>
<tr><td rowspan="2">Row Player
(player 1)</td><td>X</td><td>(m1,m2)</td><td>(m3,m4)</td></tr>
<tr><td>Y</td><td>(m5,m6)</td><td>(m7,m8)</td></tr>
</table>

Strategic forms game

Definition :

(**Strategic Form Game**) A strategic forms game is a triplet $< N, (X_i)_{i \in N}, (u_i)_{i \in N} >$ such that

- N is a finite set of players, i.e., $N := \{1, \ldots, N\}$
- X_i is the set of available actions for player i
- $x_i \in X_i$ is an action for player i
- $u_i : X \rightarrow R$ is the payoff function of player i where $X = \prod_i X_i$ the set of all action profiles.
- $x_{-i} = [x_j]_{j \neq i}$: vector of actions for all players except i
- $(x_i, x_{-i}) \in X$ is a **strategy profile**, or outcome.

Nash Equilibrium

- A pure-strategy **Nash equilibrium** of a non-cooperative game $G = < N, (X_i)_{i \in N}, (u_i)_{i \in N} >$ is a strategy profile $x^* \in X$ such that for all $\forall i \in N$

$$(x^*_i, x^*_{-i}) \geq (x_i, x^*_{-i}) \qquad \text{for all } x_i \in X_i$$

In simpler words

- A **set of strategies** is a Nash equilibrium if no player can do better by **unilaterally** changing their strategy
- These strategies are sometimes called **best response strategies**

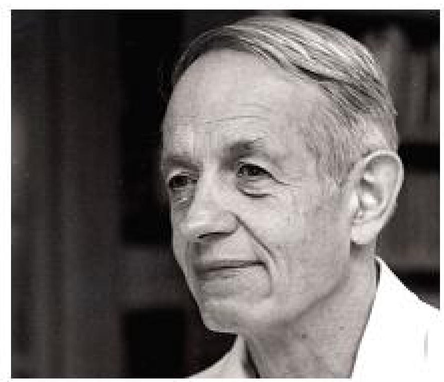

Example (The Prisoner's Dilemma):

- **"The Prisoner's Dilemma":** This is a classic two-player nonzero-sum game, which illustrates many points of game theory.
- In this game, the two players are partners in a crime who have been captured by the police. Each suspect is placed in a separate cell and offered the opportunity to confess to the crime.

Example (The Prisoner's Dilemma):

- **The Prisoner's Dilemma**

		Prisoner 2	
		Confess (Defect)	Hold out (Cooperate)
Prisoner 1	Confess (Defect)	-5 -5	0 -10
	Hold out (Cooperate)	-10 0	-1 -1

Example (The Prisoner's Dilemma):

- Whatever Prisoner 2 does, the best that Prisoner 1 can do is Confess

		Prisoner 2: Confess (Defect)	Prisoner 2: Hold out (Cooperate)
Prisoner 1	Confess (Defect)	-5, -5	0, -10
	Hold out (Cooperate)	-10, 0	-1, -1

Example (The Prisoner's Dilemma):

- Whatever Prisoner 1 does, the best that Prisoner 2 can do is Confess

		Prisoner 2	
		Confess (Defect)	Hold out (Cooperate)
Prisoner 1	Confess (Defect)	-5 -5	0 -10
	Hold out (Cooperate)	-10 0	-1 -1

Example (The Prisoner's Dilemma):

- Each player has a dominant strategy to Confess.
- The dominant strategy equilibrium is (Confess,Confess)

Prisoner 2

		Confess (Defect)	Hold out (Cooperate)
Prisoner 1	Confess (Defect)	-5, -5	0, -10
	Hold out (Cooperate)	-10, 0	-1, -1

Example (The Prisoner's Dilemma):

		Prisoner 2	
		Confess (Defect)	Hold out (Cooperate)
Prisoner 1	Confess (Defect)	-5 / -5	0 / -10
	Hold out (Cooperate)	-10 / 0	-1 / -1

The payoff in the dominant strategy equilibrium (-5,-5) is worse for both players than (-1,-1), the payoff in the case that both players hold out. Thus, the Prisoners' Dilemma Game is a game of social conflict.

Opportunity for multi-agent learning: by learning during repeated play, the Pareto optimal solution (-1,-1) can emerge as a result of learning (also can arise in evolutionary game theory).

Example (Battle of the Sexes):

		Bob	
		Mountain	Cinema
Alice	Mountain	2 1	-1 -1
	Cinema	-5 -5	1 2

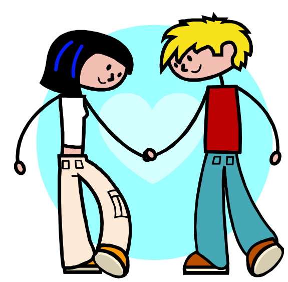

Example (Battle of the Sexes):

		Bob	
		Mountain	Cinema
Alice	Mountain	2 1	-1 -1
	Cinema	-5 -5	1 2

This game has

- no (iterated) dominant strategy equilibrium

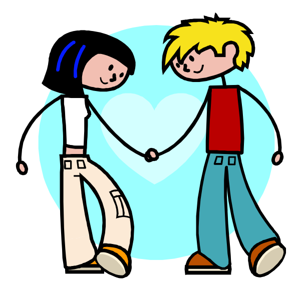

Example (Battle of the Sexes):

Bob

Alice		Mountain	Cinema
	Mountain	2, 1	-1, -1
	Cinema	-5, -5	1, 2

This game has

- no (iterated) dominant strategy equilibrium

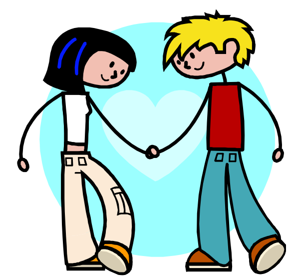

Example (Battle of the Sexes):

		Bob: Mountain	Bob: Cinema
Alice	Mountain	2, 1	-1, -1
	Cinema	-5, -5	1, 2

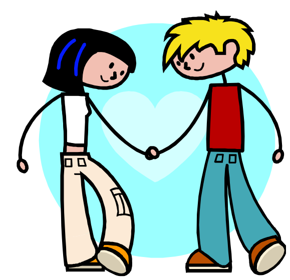

This game has

- no (iterated) dominant strategy equilibrium
- two Nash equilibria (Mountain, Mountain) and (Cinema, Cinema)

Example (Battle of the Sexes):

Bob

Alice		Mountain	Cinema
	Mountain	2, 1	-1, -1
	Cinema	-5, -5	1, 2

This game has two Nash equilibria

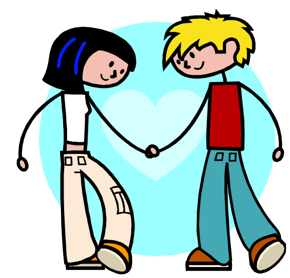

How can these two players coordinate ?

Zero sum games

Normally, if a matrix game is a zero-sum game, we don't need to write the payoffs for both players. Instead we write only one number in each matrix cell.

Player 2

Player 1

	A	B
X	(2,-2)	(2,-2)
Y	(1,-1)	(3,-3)

Player 2

Player 1

	A	B
X	2	2
Y	1	3

- If only one number is written in each cell of the matrix, then the game is understood to be a zero-sum game.

Zero sum games

- For example, in the matrix below, because only one number is written in each cell, it is understood that the payoffs are the given numbers for the row player and the negative of each of these values for the column player.

Player 2

Player 1

	A	B
X	2	2
Y	1	3

- The row player will want to the maximum value in the matrix.
- The column player will want the minimum value in the matrix.

Zero sum games:

Actions:

- each player knows the game setting (available strategies to row player and column player, values of payoff matrix)

Information structure:

- both players simultaneously choose their strategy, that is, without knowing what their opponent chooses

Objective:

- each player chooses a strategy that enables him/her to do best, given that the opponent knows his/her strategy

Assumption:

- both players are rational:
  - they try to maximize their pay off
  - they show no compassion for their opponent

Finding equilibrium point

Player 2

Player1

	A	B	
X	2	2	2
Y	1	3	1

- To look for an equilibrium point in a matrix game, begin by finding the minimum values of each row. (These are called the row minima.)
- Then, choose the maximum value of these minima. (This is called the maximin strategy for the row player.)
- We do this because a rationalization for the row player is to find the best move possible *assuming that the column player chooses his best strategy.* It's like thinking "what's the best I can do when he plays his best."

Finding equilibrium point

Player 2

Player1

	A	B	
X	2	2	2
Y	1	3	1
	2	3	

- Now, look for the maximum values for each column. (These are column maxima.)
- Then, choose the smallest of the column maxima. This is called the minimax strategy for the column player.
- Here, the column player is thinking: "What is the best I can do when the row player is playing his best strategy."

Finding equilibrium point

Player 2

Player1

	A	B	
X	2	2	2
Y	1	3	1
	2	3	

- Because, in this example, the row maximin and the column minimax strategies coincide with the same outcome, these strategies are called a *saddle point* (or an *equilibrium point*).
- The outcome associated with the equilibrium point is called the *value of the game*. In this case, the value is 2.
- We say a game is fair if it's value is 0. Thus, this particular game is not fair.

Saddle points

- In zero-sum games, the terms saddle point and equilibrium point are interchangeable.

- These terms refer to the combination of strategies that are in each player's best interest assuming all other player's use the strategy in their best interest.
- The term saddle point comes from the fact that, in a game with two players, each with two strategies, it represents at one point a highest point for the lowest payoff values and also a lowest point for the highest payoff values.

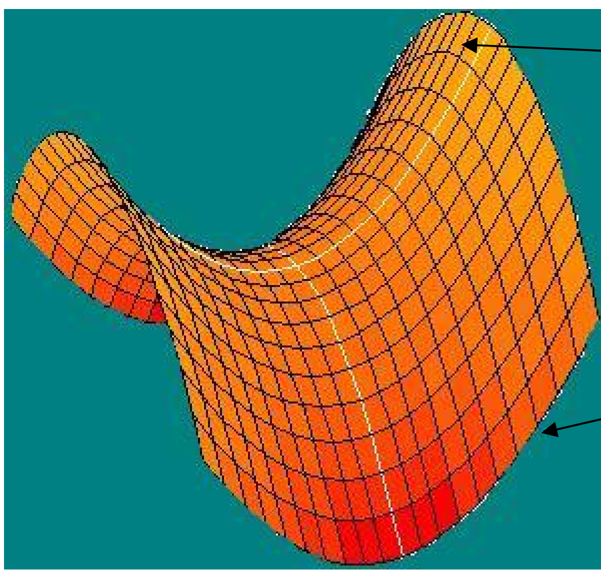

Row player wants highest value in matrix

Column player wants lowest value in matrix

Saddle points

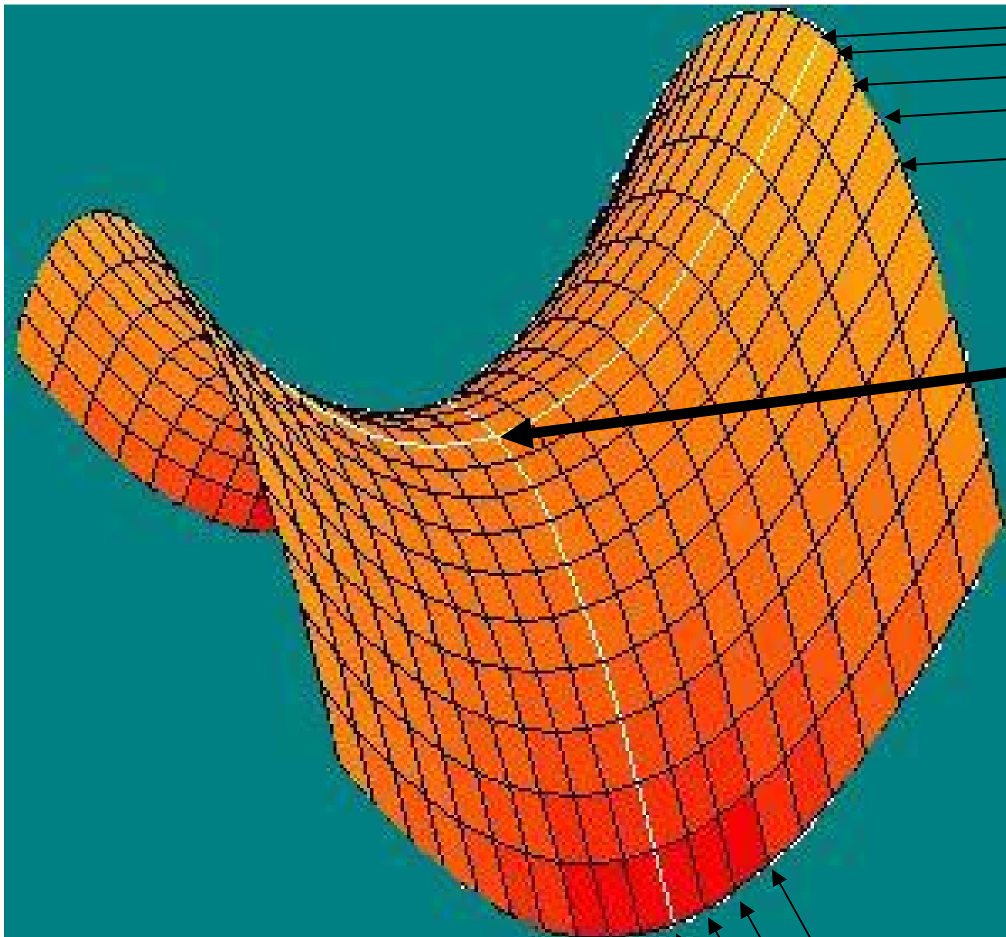

As the row player, imagine these lines are your choices. You want the highest on the curve.

The saddle point is where both players get the best result assuming the other makes the best choice for themselves.

As the column player, imagine these lines as your choices. You want the lowest on the curve.

Security levels and policies

- The security level of P1 (the maximizer) is defined by.

$$\underline{V}(A) := \max_{i \in \{1,\ldots,\mathrm{m}\}} \; \min_{j \in \{1,\ldots,\mathrm{n}\}} a_{i,j}$$

- The security policy of P1 (the maximizer) is defined by

$$i^* \in \arg\max_{i \in \{1,\ldots,\mathrm{m}\}} \; \min_{j \in \{1,\ldots,\mathrm{n}\}} a_{i,j}$$

Security levels and policies

- The security level of P2 (the minimizer) is defined by.

$$\bar{V}(A) = \min_{j\in\{1,\dots,n\}} \max_{i\in\{1,\dots,m\}} a_{i,j}$$

- The security policy of P2 (the minimizer) is defined by

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

Security levels and policies

Proposition (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 policies (not necessarily unique)*

(iii) *The security levels always satisfy*

$$\underline{V}(A) := \max_{i \in \{1,\ldots,m\}} \min_{j \in \{1,\ldots,n\}} a_{i,j} \leq \bar{V}(A) := \min_{j \in \{1,\ldots,n\}} \max_{i \in \{1,\ldots,m\}} a_{i,j}$$

Security levels and policies

Let proof part (iii), i.e., that

$$\underline{V}(A) := \max_{i\in\{1,\dots,m\}} \min_{j\in\{1,\dots,n\}} a_{i,j} \le \bar{V}(A) := \min_{j\in\{1,\dots,n\}} \max_{i\in\{1,\dots,m\}} a_{i,j}$$

Let i^* be a security policy for P1, i.e.,

$$\underline{V}(A) = \min_{j\in\{1,\dots,n\}} a_{i^*,j}$$

Since

$$a_{i^*,j} \le \max_{i\in\{1,\dots,m\}} a_{i,j} \qquad \forall j \in \{1,\dots,n\}$$

Combing this two

$$\underline{V}(A) = \min_{j\in\{1,\dots,n\}} a_{i^*,j} \le \min_{j\in\{1,\dots,n\}} \max_{i\in\{1,\dots,m\}} a_{i,j} = \bar{V}(A)$$

Saddle-point and security levels

Theorem (Saddle-point and security levels)

A matrix game defined by A has a saddle-point equilibrium if and only if

$$\underline{V}(A) := \max_{i\in\{1,\ldots,m\}} \min_{j\in\{1,\ldots,n\}} a_{i,j} = \overline{V}(A) := \min_{j\in\{1,\ldots,n\}} \max_{i\in\{1,\ldots,m\}} a_{i,j}$$

(i^*, j^*) is a saddle-point equilibrium

$\underline{V}(A) = \overline{V}(A)$ *is the saddle-point value;*