

Computing Nash Equilibrium; Maxmin

Lecture 5

Lecture Overview

- 1 Recap
- 2 Computing Mixed Nash Equilibria
- 3 Fun Game
- 4 Maxmin and Minmax

Pareto Optimality

- **Idea:** sometimes, one outcome o is at least as good for every agent as another outcome o' , and there is some agent who strictly prefers o to o'
 - in this case, it seems reasonable to say that o is better than o'
 - we say that o **Pareto-dominates** o' .
- An outcome o^* is **Pareto-optimal** if there is no other outcome that Pareto-dominates it.
 - a game can have more than one Pareto-optimal outcome
 - every game has at least one Pareto-optimal outcome

Best Response

- If you knew what everyone else was going to do, it would be easy to pick your own action
- Let $a_{-i} = \langle a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n \rangle$.
 - now $a = (a_{-i}, a_i)$
- **Best response:** $a_i^* \in BR(a_{-i})$ iff
$$\forall a_i \in A_i, u_i(a_i^*, a_{-i}) \geq u_i(a_i, a_{-i})$$

Nash Equilibrium

- Now we return to the setting where no agent knows anything about what the others will do
- Idea: look for **stable** action profiles.
- $a = \langle a_1, \dots, a_n \rangle$ is a (“pure strategy”) **Nash equilibrium** iff $\forall i, a_i \in BR(a_{-i})$.

Mixed Strategies

- It would be a pretty bad idea to play any deterministic strategy in matching pennies
- Idea: confuse the opponent by playing **randomly**
- Define a **strategy** s_i for agent i as any probability distribution over the actions A_i .
 - **pure strategy**: only one action is played with positive probability
 - **mixed strategy**: more than one action is played with positive probability
 - these actions are called the **support** of the mixed strategy
- Let the set of **all strategies** for i be S_i
- Let the set of **all strategy profiles** be $S = S_1 \times \dots \times S_n$.

Utility under Mixed Strategies

- What is your **payoff** if all the players follow mixed strategy profile $s \in S$?
 - We can't just read this number from the game matrix anymore: we won't always end up in the same cell
- Instead, use the idea of **expected utility** from decision theory:

$$u_i(s) = \sum_{a \in A} u_i(a) Pr(a|s)$$

$$Pr(a|s) = \prod_{j \in N} s_j(a_j)$$

Best Response and Nash Equilibrium

Our definitions of best response and Nash equilibrium generalize from actions to strategies.

- **Best response:**

- $s_i^* \in BR(s_{-i})$ iff $\forall s_i \in S_i, u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$

- **Nash equilibrium:**

- $s = \langle s_1, \dots, s_n \rangle$ is a Nash equilibrium iff $\forall i, s_i \in BR(s_{-i})$

- **Every finite game has a Nash equilibrium!** [Nash, 1950]

- e.g., matching pennies: both players play heads/tails 50%/50%

Lecture Overview

- 1 Recap
- 2 Computing Mixed Nash Equilibria
- 3 Fun Game
- 4 Maxmin and Minmax

Computing Mixed Nash Equilibria: Battle of the Sexes

| | B | F |
|---|------|------|
| B | 2, 1 | 0, 0 |
| F | 0, 0 | 1, 2 |

- It's hard in general to compute Nash equilibria, but it's easy when you can guess the **support**
- For BoS, let's look for an equilibrium where all actions are part of the support

Computing Mixed Nash Equilibria: Battle of the Sexes

| | B | F |
|---|------|------|
| B | 2, 1 | 0, 0 |
| F | 0, 0 | 1, 2 |

- Let player 2 play B with p , F with $1 - p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between F and B (why?)

Computing Mixed Nash Equilibria: Battle of the Sexes

| | | |
|---|------|------|
| | B | F |
| B | 2, 1 | 0, 0 |
| F | 0, 0 | 1, 2 |

- Let player 2 play B with p , F with $1 - p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between F and B (why?)

$$u_1(B) = u_1(F)$$

$$2p + 0(1 - p) = 0p + 1(1 - p)$$

$$p = \frac{1}{3}$$

Computing Mixed Nash Equilibria: Battle of the Sexes

| | B | F |
|---|------|------|
| B | 2, 1 | 0, 0 |
| F | 0, 0 | 1, 2 |

- Likewise, player 1 must randomize to make player 2 indifferent.
 - Why is player 1 willing to randomize?

Computing Mixed Nash Equilibria: Battle of the Sexes

| | B | F |
|---|------|------|
| B | 2, 1 | 0, 0 |
| F | 0, 0 | 1, 2 |

- Likewise, player 1 must randomize to make player 2 indifferent.
 - Why is player 1 willing to randomize?
- Let player 1 play B with q , F with $1 - q$.

$$u_2(B) = u_2(F)$$

$$q + 0(1 - q) = 0q + 2(1 - q)$$

$$q = \frac{2}{3}$$

- Thus the strategies $(\frac{2}{3}, \frac{1}{3})$, $(\frac{1}{3}, \frac{2}{3})$ are a Nash equilibrium.

Interpreting Mixed Strategy Equilibria

What does it mean to play a mixed strategy?

- Randomize to **confuse** your opponent
 - consider the matching pennies example
- Players randomize when they are **uncertain** about the other's action
 - consider battle of the sexes
- Mixed strategies are a concise description of what might happen in **repeated play**: count of pure strategies in the limit
- Mixed strategies describe **population dynamics**: 2 agents chosen from a population, all having deterministic strategies. MS is the probability of getting an agent who will play one PS or another.

Lecture Overview

- 1 Recap
- 2 Computing Mixed Nash Equilibria
- 3 Fun Game**
- 4 Maxmin and Minmax

Fun Game!

| | | |
|----------|----------|----------|
| | <i>L</i> | <i>R</i> |
| <i>T</i> | 80, 40 | 40, 80 |
| <i>B</i> | 40, 80 | 80, 40 |

- Play once as each player, recording the strategy you follow.

Fun Game!

| | L | R |
|-----|---------|--------|
| T | 320, 40 | 40, 80 |
| B | 40, 80 | 80, 40 |

- Play once as each player, recording the strategy you follow.

Fun Game!

| | | |
|----------|----------|----------|
| | <i>L</i> | <i>R</i> |
| <i>T</i> | 44, 40 | 40, 80 |
| <i>B</i> | 40, 80 | 80, 40 |

- Play once as each player, recording the strategy you follow.

Fun Game!

| | <i>L</i> | <i>R</i> |
|----------|--------------------------------|----------|
| <i>T</i> | 80, 40; 320, 40; 44, 40 | 40, 80 |
| <i>B</i> | 40, 80 | 80, 40 |

- Play once as each player, recording the strategy you follow.
- What does row player do in equilibrium of this game?

Fun Game!

| | L | R |
|-----|--------------------------------|--------|
| T | 80, 40; 320, 40; 44, 40 | 40, 80 |
| B | 40, 80 | 80, 40 |

- Play once as each player, recording the strategy you follow.
- What does row player do in equilibrium of this game?
 - row player randomizes 50-50 all the time
 - that's what it takes to make column player indifferent

Fun Game!

| | <i>L</i> | <i>R</i> |
|----------|--------------------------------|----------|
| <i>T</i> | 80, 40; 320, 40; 44, 40 | 40, 80 |
| <i>B</i> | 40, 80 | 80, 40 |

- Play once as each player, recording the strategy you follow.
- What does row player do in equilibrium of this game?
 - row player randomizes 50-50 all the time
 - that's what it takes to make column player indifferent
- What happens when people play this game?

Fun Game!

| | <i>L</i> | <i>R</i> |
|----------|--------------------------------|----------|
| <i>T</i> | 80, 40; 320, 40; 44, 40 | 40, 80 |
| <i>B</i> | 40, 80 | 80, 40 |

- Play once as each player, recording the strategy you follow.
- What does row player do in equilibrium of this game?
 - row player randomizes 50-50 all the time
 - that's what it takes to make column player indifferent
- What happens when people play this game?
 - with payoff of 320, row player goes up essentially all the time
 - with payoff of 44, row player goes down essentially all the time

Lecture Overview

- 1 Recap
- 2 Computing Mixed Nash Equilibria
- 3 Fun Game
- 4 Maxmin and Minmax**

Maxmin Strategies

- Player i 's **maxmin strategy** is a strategy that maximizes i 's worst-case payoff, in the situation where all the other players (whom we denote $-i$) happen to play the strategies which cause the greatest harm to i .
- The **maxmin value** (or **safety level**) of the game for player i is that minimum amount of payoff guaranteed by a maxmin strategy.

Definition (Maxmin)

The **maxmin strategy** for player i is $\arg \max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$, and the **maxmin value** for player i is $\max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$.

- Why would i want to play a maxmin strategy?

Maxmin Strategies

- Player i 's **maxmin strategy** is a strategy that maximizes i 's worst-case payoff, in the situation where all the other players (whom we denote $-i$) happen to play the strategies which cause the greatest harm to i .
- The **maxmin value** (or **safety level**) of the game for player i is that minimum amount of payoff guaranteed by a maxmin strategy.

Definition (Maxmin)

The **maxmin strategy** for player i is $\arg \max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$, and the **maxmin value** for player i is $\max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$.

- Why would i want to play a maxmin strategy?
 - a conservative agent maximizing worst-case payoff
 - a paranoid agent who believes everyone is out to get him

Minmax Strategies

- Player i 's **minmax strategy** against player $-i$ in a 2-player game is a strategy that minimizes $-i$'s best-case payoff, and the **minmax value** for i against $-i$ is payoff.
- Why would i want to play a minmax strategy?

Definition (Minmax, 2-player)

In a two-player game, the **minmax strategy** for player i against player $-i$ is $\arg \min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$, and player $-i$'s **minmax value** is $\min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$.

Minmax Strategies

- Player i 's **minmax strategy** against player $-i$ in a 2-player game is a strategy that minimizes $-i$'s best-case payoff, and the **minmax value** for i against $-i$ is payoff.
- Why would i want to play a minmax strategy?
 - to punish the other agent as much as possible

Definition (Minmax, 2-player)

In a two-player game, the **minmax strategy** for player i against player $-i$ is $\arg \min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$, and player $-i$'s **minmax value** is $\min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$.

Minmax Strategies

- Player i 's **minmax strategy** against player $-i$ in a 2-player game is a strategy that minimizes $-i$'s best-case payoff, and the **minmax value** for i against $-i$ is payoff.
- Why would i want to play a minmax strategy?
 - to punish the other agent as much as possible

We can generalize to n players.

Definition (Minmax, n -player)

In an n -player game, the **minmax strategy** for player i against player $j \neq i$ is i 's component of the mixed strategy profile s_{-j} in the expression $\arg \min_{s_{-j}} \max_{s_j} u_j(s_j, s_{-j})$, where $-j$ denotes the set of players other than j . As before, the **minmax value** for player j is $\min_{s_{-j}} \max_{s_j} u_j(s_j, s_{-j})$.

Minmax Theorem

Theorem (Minimax theorem (von Neumann, 1928))

In any finite, two-player, zero-sum game, in any Nash equilibrium each player receives a payoff that is equal to both his maxmin value and his minmax value.

Minmax Theorem

Theorem (Minimax theorem (von Neumann, 1928))

In any finite, two-player, zero-sum game, in any Nash equilibrium each player receives a payoff that is equal to both his maxmin value and his minmax value.

- 1 Each player's maxmin value is equal to his minmax value. By convention, the maxmin value for player 1 is called the **value of the game**.

Minmax Theorem

Theorem (Minimax theorem (von Neumann, 1928))

In any finite, two-player, zero-sum game, in any Nash equilibrium each player receives a payoff that is equal to both his maxmin value and his minmax value.

- 1 Each player's maxmin value is equal to his minmax value. By convention, the maxmin value for player 1 is called the **value of the game**.
- 2 For both players, the set of maxmin strategies coincides with the set of minmax strategies.

Minmax Theorem

Theorem (Minimax theorem (von Neumann, 1928))

In any finite, two-player, zero-sum game, in any Nash equilibrium each player receives a payoff that is equal to both his maxmin value and his minmax value.

- 1 Each player's maxmin value is equal to his minmax value. By convention, the maxmin value for player 1 is called the **value of the game**.
- 2 For both players, the set of maxmin strategies coincides with the set of minmax strategies.
- 3 Any maxmin strategy profile (or, equivalently, minmax strategy profile) is a Nash equilibrium. Furthermore, these are all the Nash equilibria. Consequently, all Nash equilibria have the same payoff vector (namely, those in which player 1 gets the value of the game).

Saddle Point: Matching Pennies

