

भारतीय प्रौद्योगिकी संस्थान मुंबई

Indian Institute of Technology Bombay

Strategic Multi-agent Artificial Intelligence

Day 1: AI and competitive games

Swaprava Nath

Slide preparation acknowledgments: Drashthi Doshi

ज्ञानम् परमम् ध्येयम् Knowledge is the supreme goal



▶ Part 1

• Sequential move games • Partial information games • Subgame perfection

► Part 2

• Solutions for 2 player zero sum games • Simultaneous move games













Rules of the game are:







Rules of the game are:

Player 1 (agent) picks the bucket







Rules of the game are:

- Player 1 (agent) picks the bucket
- Player 2 (opponent) picks a number from the "Selected" Bucket







Rules of the game are:

- Player 1 (agent) picks the bucket
- Player 2 (opponent) picks a number from the "Selected" Bucket
- Player 1 (agent) utility will be the number picked by player 2





Now, let's construct the **game tree**





Now, let's construct the **game tree**







The opponent is **stochastic**, i.e., the probability of choosing both left and right numbers in a bucket are the same (i.e., $\frac{1}{2}$).



The opponent is **stochastic**, i.e., the probability of choosing both left and right numbers in a bucket are the same (i.e., $\frac{1}{2}$).

• If *A* is chosen, expected utility $=\frac{-50+50}{2}=0$.



The opponent is **stochastic**, i.e., the probability of choosing both left and right numbers in a bucket are the same (i.e., $\frac{1}{2}$).

- If *A* is chosen, expected utility $=\frac{-50+50}{2}=0$.
- If *B* is chosen, expected utility = $\frac{1+3}{2} = 2$.



The opponent is **stochastic**, i.e., the probability of choosing both left and right numbers in a bucket are the same (i.e., $\frac{1}{2}$).

- If *A* is chosen, expected utility $=\frac{-50+50}{2}=0$.
- If *B* is chosen, expected utility $=\frac{1+3}{2}=2$.
- If *C* is chosen, expected utility = $\frac{-10+20}{2} = 5$.



The opponent is **stochastic**, i.e., the probability of choosing both left and right numbers in a bucket are the same (i.e., $\frac{1}{2}$).

- If A is chosen, expected utility $=\frac{-50+50}{2}=0.$
- If *B* is chosen, expected utility $=\frac{1+3}{2}=2$.
- If *C* is chosen, expected utility = $\frac{-10+20}{2} = 5$.

So in this case agent picks **bucket C**, to maximize it's utility.



Consider another strategy that the opponent can adopt.

The opponent is a **min player**, i.e., the opponent always chooses the **minimum number** from the bucket.



Consider another strategy that the opponent can adopt.

The opponent is a **min player**, i.e., the opponent always chooses the **minimum number** from the bucket.





Consider another strategy that the opponent can adopt.

The opponent is a **min player**, i.e., the opponent always chooses the **minimum number** from the bucket.

In this case, the agent picks **Bucket B**, as $\max\{-50, 1, -10\} = 1$.







• **Players:** {Agent, Opponent}



- **Players:** {Agent, Opponent}
- Starting State: S₀

- **Players:** {Agent, Opponent}
- Starting State: S₀
- Actions(s): Possible actions at state 's'

- **Players:** {Agent, Opponent}
- Starting State: S₀
- Actions(s): Possible actions at state 's'
- Player(s): Player who makes the move at state 's'

- **Players:** {Agent, Opponent}
- Starting State: S₀
- Actions(s): Possible actions at state 's'
- Player(s): Player who makes the move at state 's'
- **Successor**(*s*, *a*): Resulting state if action *a* is taken at state '*s*'

- **Players:** {Agent, Opponent}
- Starting State: S₀
- Actions(s): Possible actions at state 's'
- Player(s): Player who makes the move at state 's'
- **Successor**(*s*, *a*): Resulting state if action *a* is taken at state '*s*'
- **isEnd**(*s*): Is state an end state/ Flag to identify terminal state



- **Players:** {Agent, Opponent}
- Starting State: S₀
- Actions(s): Possible actions at state 's'
- Player(s): Player who makes the move at state 's'
- **Successor**(*s*, *a*): Resulting state if action *a* is taken at state '*s*'
- **isEnd**(*s*): Is state an end state / Flag to identify terminal state
- Utility(s): Agent's utility at the "end state"



- **Players:** {Agent, Opponent}
- Starting State: S₀
- Actions(s): Possible actions at state 's'
- Player(s): Player who makes the move at state 's'
- **Successor**(*s*, *a*): Resulting state if action *a* is taken at state '*s*'
- **isEnd**(*s*): Is state an end state / Flag to identify terminal state
- Utility(s): Agent's utility at the "end state"



- **Players:** {Agent, Opponent}
- Starting State: S₀
- Actions(s): Possible actions at state 's'
- Player(s): Player who makes the move at state 's'
- **Successor**(*s*, *a*): Resulting state if action *a* is taken at state '*s*'
- isEnd(s): Is state an end state / Flag to identify terminal state
- Utility(s): Agent's utility at the "end state"

Note:

i) In Actions(*s*), Player(*s*), Successor(*s*, *a*), '*s*' is an intermediate state.



- **Players:** {Agent, Opponent}
- Starting State: S₀
- Actions(s): Possible actions at state 's'
- Player(s): Player who makes the move at state 's'
- **Successor**(*s*, *a*): Resulting state if action *a* is taken at state '*s*'
- isEnd(s): Is state an end state / Flag to identify terminal state
- Utility(s): Agent's utility at the "end state"

Note:

- i) In Actions(*s*), Player(*s*), Successor(*s*, *a*), '*s*' is an intermediate state.
- ii) Utility(*s*) is not defined in other (intermediate) states.



2-player zero sum games



For example in the game of chess

Example: Chess

2-player zero sum games



- Example: Chess
 - **Players:** {White, Black}



- Example: Chess
 - **Players:** {White, Black}
 - State: A board position (positions of the chess pieces at a given time), denoted s



- Example: Chess
 - **Players:** {White, Black}
 - State: A board position (positions of the chess pieces at a given time), denoted s
 - Actions(s): All legal moves possible by Player(s)



- Example: Chess
 - **Players:** {White, Black}
 - State: A board position (positions of the chess pieces at a given time), denoted s
 - Actions(s): All legal moves possible by Player(s)
 - **Player**(*s*): Player who makes the move at state '*s*'



Example: Chess

- **Players:** {White, Black}
- State: A board position (positions of the chess pieces at a given time), denoted s
- Actions(s): All legal moves possible by Player(s)
- **Player(***s***):** Player who makes the move at state '*s*'
- **Successor**(*s*, *a*): Resulting state if action *a* is taken at state '*s*'



Example: Chess

- **Players:** {White, Black}
- State: A board position (positions of the chess pieces at a given time), denoted s
- Actions(s): All legal moves possible by Player(s)
- Player(s): Player who makes the move at state 's'
- **Successor**(*s*, *a*): Resulting state if action *a* is taken at state '*s*'
- **isEnd**(*s*): Whether '*s*' is a checkmate or a draw



Example: Chess

- **Players:** {White, Black}
- State: A board position (positions of the chess pieces at a given time), denoted s
- Actions(s): All legal moves possible by Player(s)
- Player(s): Player who makes the move at state 's'
- **Successor**(*s*, *a*): Resulting state if action *a* is taken at state '*s*'
- **isEnd**(*s*): Whether '*s*' is a checkmate or a draw
- Utility function is defined as

 $\mathbf{Utility}(s) = \begin{cases} +M & \text{if white wins} \\ -M & \text{if black wins/white loses} \\ 0 & \text{if it is a draw} \end{cases}$


► Part 1

• Sequential move games • Partial information games • Subgame perfection

► Part 2

• Solutions for 2 player zero sum games • Simultaneous move games

Games with partial information

Consider the game tree below:



 $a \in \{A, B, C\}$





$$u_{agent}(s) = \begin{cases} \text{utility}(s) & \text{if isEnd}(s) \\ \sum_{a \in actions(s)} \prod_{agent}(s)[a] \ u_{agent}(\text{successor}(s, a)) & \text{if player}(s) = \text{agent} \\ \sum_{a \in actions(s)} \prod_{opponent}(s)[a] \ u_{agent}(\text{successor}(s, a)) & \text{if player}(s) = \text{opponent} \end{cases}$$

Note: $\Pi_{agent}(s)[a]$ and $\Pi_{opponent}(s)[a]$ are probabilities that the agent and the opponent pick action 'a' respectively.



► Part 1

• Sequential move games • Partial information games • Subgame perfection

► Part 2

• Solutions for 2 player zero sum games • Simultaneous move games





A **subgame** at 'S' is a restriction of the game at the subtree rooted at 'S' where isEnd(s) is false





A **subgame** at 'S' is a restriction of the game at the subtree rooted at 'S' where isEnd(s) is false





A **subgame** at 'S' is a restriction of the game at the subtree rooted at 'S' where isEnd(s) is false



• **Player 1** is a utility maximizer (**Max Player**) i.e. plays to maximize his utility. Hence if player(s) = player 1, the utility of Player 1 changes to, $u_1(s) =$ $\max_{a \in actions(s)} u_1(successor(s, a))$



A **subgame** at 'S' is a restriction of the game at the subtree rooted at 'S' where isEnd(s) is false



• **Player 1** is a utility maximizer (**Max Player**) i.e. plays to maximize his utility. Hence if player(s) = player 1, the utility of Player 1 changes to, $u_1(s) =$

 $\max_{a \in actions(s)} u_1(\operatorname{successor}(s, a))$

• **Player 2** is the utility minimiser(**Min Player**) i.e. plays to minimize his utility.

Hence if player(s) = player 2, the utility of Player 1 changes to, $u_1(s) = \min_{a \in actions(s)} u_1(successor(s, a))$







Let us look at the **subgame at node** n_4 :





Let us look at the **subgame at node** n_4 :



In this subgame, it is Player 1's turn and he is a **max** player.



Let us look at the **subgame at node** n_4 :



In this subgame, it is Player 1's turn and he is a **max** player.

Player 1 (max player) should pick **G** here so as to get **2** (> 1, that he would have got by choosing **H**).



Let us look at the **subgame at node** n_4 :



In this subgame, it is Player 1's turn and he is a **max** player.

Player 1 (max player) should pick **G** here so as to get **2** (> 1, that he would have got by choosing **H**).

We are done solving this subgame.





Similarly using the result of the subgame at n_4 , we can solve the subgame at n_3 , and this way we move to the upper levels.

















In this subgame, it's Player 2's turn and he is a **min** player.







In this subgame, it's Player 2's turn and he is a **min** player.

Therefore, he should pick F so that player 1 gets 2 (from n_4 , instead of getting 5 by picking E).

















In this subgame, it's Player 2's turn and he is a **min** player.





In this subgame, it's Player 2's turn and he is a **min** player.

Therefore, he should pick C so that player 1 gets 3 (instead of getting 8 by picking D).













In this subgame, it's Player 1's turn.







In this subgame, it's Player 1's turn.

Player 1 (max player) should pick A here to get 3 (> 2, that he would have got by choosing B).







In this subgame, it's Player 1's turn.

Player 1 (max player) should pick A here to get 3 (> 2, that he would have got by choosing **B**).

Therefore, the **final utility of the Player 1 = 3**







The summary of solving the above game is:

• Player 1 will choose **H** in the subgame at n_4



- Player 1 will choose **H** in the subgame at n_4
- Player 2 will choose **F** in the subgame at *n*₃



- Player 1 will choose **H** in the subgame at n_4
- Player 2 will choose **F** in the subgame at *n*₃
- player 2 will choose **C** in the subgame at *n*₂



- Player 1 will choose **H** in the subgame at n_4
- Player 2 will choose **F** in the subgame at *n*₃
- player 2 will choose **C** in the subgame at *n*₂
- player 1 will choose **A** in the subgame at *n*₁



- Player 1 will choose **H** in the subgame at n_4
- Player 2 will choose **F** in the subgame at *n*₃
- player 2 will choose **C** in the subgame at *n*₂
- player 1 will choose **A** in the subgame at *n*₁



The summary of solving the above game is:

- Player 1 will choose **H** in the subgame at n_4
- Player 2 will choose **F** in the subgame at *n*₃
- player 2 will choose **C** in the subgame at *n*₂
- player 1 will choose **A** in the subgame at *n*₁

The algorithm described in the above example is **Backward Induction**.



We can apply Backward induction on small games like Tic-tac-toe.



We can apply Backward induction on small games like Tic-tac-toe.

But can we apply it to Chess, Go, Checkers, etc.?


But can we apply it to Chess, Go, Checkers, etc.?



But can we apply it to Chess, Go, Checkers, etc.?

We can, but the game tree is huge.

• Checkers game tree $\sim 10^{20}$ nodes



But can we apply it to Chess, Go, Checkers, etc.?

- Checkers game tree $\sim 10^{20}$ nodes
- Chess game tree $\sim 10^{40}$ nodes

But can we apply it to Chess, Go, Checkers, etc.?

- Checkers game tree $\sim 10^{20}$ nodes
- Chess game tree $\sim 10^{40}$ nodes
- Go game tree $\sim 10^{170}$ nodes



But can we apply it to Chess, Go, Checkers, etc.?

- Checkers game tree $\sim 10^{20}$ nodes
- Chess game tree $\sim 10^{40}$ nodes
- Go game tree $\sim 10^{170}$ nodes



But can we apply it to Chess, Go, Checkers, etc.?

We can, but the game tree is huge.

- Checkers game tree $\sim 10^{20}$ nodes
- Chess game tree $\sim 10^{40}$ nodes
- Go game tree $\sim 10^{170}$ nodes

Checkers was solved in 2007 after 18 years of computation and the optimal solution was a draw.



► Part 1

• Sequential move games • Partial information games • Subgame perfection

► Part 2

• Solutions for 2 player zero sum games • Simultaneous move games

Recall: 2-player zero sum game





Depth Limited Search



• Depth-limited search algorithm explores as far as possible along each branch before backtracking.



- Depth-limited search algorithm explores as far as possible along each branch before backtracking.
- A **depth limit** is imposed and when the algorithm reaches the specified depth limit, it stops exploring that branch further and backtracks to explore other branches

Depth Limited Search



- Depth-limited search algorithm explores as far as possible along each branch before backtracking.
- A **depth limit** is imposed and when the algorithm reaches the specified depth limit, it stops exploring that branch further and backtracks to explore other branches
- This prevents the algorithms from getting stuck in deep paths and allows it to handle very deep trees effectively, reducing the computational time.

Depth Limited Search



- Depth-limited search algorithm explores as far as possible along each branch before backtracking.
- A **depth limit** is imposed and when the algorithm reaches the specified depth limit, it stops exploring that branch further and backtracks to explore other branches
- This prevents the algorithms from getting stuck in deep paths and allows it to handle very deep trees effectively, reducing the computational time.



- Depth-limited search algorithm explores as far as possible along each branch before backtracking.
- A **depth limit** is imposed and when the algorithm reaches the specified depth limit, it stops exploring that branch further and backtracks to explore other branches
- This prevents the algorithms from getting stuck in deep paths and allows it to handle very deep trees effectively, reducing the computational time.

The utility of the agent at node s and depth d can be written as

$$u_{agent}(s,d) = \begin{cases} utility(s) & \text{if } isEnd(s) \\ eval(s) & \text{if } d = 0 \\\\ \max_{a \in actions(s)} \{u_{agent}(succ(s,a),d-1)\} & \text{if } player(s) = agent \\\\ \min_{a \in actions(s)} \{u_{agent}(succ(s,a),d-1)\} & \text{if } player(s) = opponent \end{cases}$$



Where eval(s) is a domain specific function denoting the possible utility to the agent.



(1)

Where $\ensuremath{\textit{eval}}(s)$ is a domain specific function denoting the possible utility to the agent.

For example in the game of chess, it can be written as

$$eval(s) = army + mobility + king's \ safety + \dots$$



Where eval(s) is a domain specific function denoting the possible utility to the agent.

For example in the game of chess, it can be written as

$$eval(s) = army + mobility + king's \ safety + \dots$$
 (1)

The contribution of army towards the utility can be represented by the following equation

$$army = 10^{100}(K - K') + q(Q - Q') + r(R - R') + \dots$$
 (2)



Where eval(s) is a domain specific function denoting the possible utility to the agent.

For example in the game of chess, it can be written as

$$eval(s) = army + mobility + king's \ safety + \dots$$
 (1)

The contribution of army towards the utility can be represented by the following equation

$$army = 10^{100}(K - K') + q(Q - Q') + r(R - R') + \dots$$
(2)

Here K, Q, R and K', Q', R' are the number of king, queen and rooks that the agent and the opponent respectively have.





• We initialise α with $-\infty$ and β with $+\infty$ in all action nodes.



- We initialise α with $-\infty$ and β with $+\infty$ in all action nodes.
- we mark subtrees to be pruned based on the values of α and β . If $\alpha \ge \beta$ at a node, we can prune the subtree rooted at that node as it won't affect the final decision.



- We initialise α with $-\infty$ and β with $+\infty$ in all action nodes.
- we mark subtrees to be pruned based on the values of α and β . If $\alpha \ge \beta$ at a node, we can prune the subtree rooted at that node as it won't affect the final decision.



- We initialise α with $-\infty$ and β with $+\infty$ in all action nodes.
- we mark subtrees to be pruned based on the values of α and β . If $\alpha \ge \beta$ at a node, we can prune the subtree rooted at that node as it won't affect the final decision.
- — For 'Max' nodes: $\alpha = \max(\alpha, \text{value})$



- We initialise α with $-\infty$ and β with $+\infty$ in all action nodes.
- we mark subtrees to be pruned based on the values of α and β . If $\alpha \ge \beta$ at a node, we can prune the subtree rooted at that node as it won't affect the final decision.
- — For 'Max' nodes: $\alpha = \max(\alpha, \text{value})$
 - For 'Min' nodes: $\beta = \min(\beta, \text{value})$



Considering the given binary tree, these are the following steps :-

• Start with the subtree rooted at D. After transfering values from its two leaves, $\alpha = 5$, $\beta = \infty$.



An example of α , β pruning





An example of α , β pruning







Taking the value from first child of E, α =6. So, $\alpha > \beta$, hence the rest branch of E is ignored. B(5) С D(5) E(6) F G 5 3 9 2 -1 0 6 1

An example of α , β pruning











• After transfering values from A and F to C, α =5, β =2. So, $\alpha > \beta$. Hence, G subtree is ignored.



Likewise, the whole game can be solved in a computionally less expensive way by pruning.



► Part 1

• Sequential move games • Partial information games • Subgame perfection

► Part 2

• Solutions for 2 player zero sum games • Simultaneous move games



Consider football where the shooter makes a decision on whether to aim towards the left, right, or center of the goal. At the same time the goalkeeper must decide whether to block the left, right or center of the goal



Consider football where the shooter makes a decision on whether to aim towards the left, right, or center of the goal. At the same time the goalkeeper must decide whether to block the left, right or center of the goal





Consider football where the shooter makes a decision on whether to aim towards the left, right, or center of the goal. At the same time the goalkeeper must decide whether to block the left, right or center of the goal



This is an example of a **2-player zero sum simultaneous move game** and is called a **Matrix Game**.





Definition

Equilibrium : A tuple of actions from which no player **gains** (a strict gain) by a **unilateral deviation**. A unilateral deviation from a tuple of actions $(a_1, a_2, ..., a_n)$ is a tuple of actions of the form $(b_1, b_2, ..., b_n)$ where $a_i = b_i$, $\forall i \neq j$ for some j ie, the actions of exactly one player j differs from the previous set.





Definition

Equilibrium : A tuple of actions from which no player **gains** (a strict gain) by a **unilateral deviation**. A unilateral deviation from a tuple of actions $(a_1, a_2, ..., a_n)$ is a tuple of actions of the form $(b_1, b_2, ..., b_n)$ where $a_i = b_i$, $\forall i \neq j$ for some j ie, the actions of exactly one player j differs from the previous set.

Observe that for the game described above, no such equilibrium exists. This can be shown by analysing all possible pairs of actions.




Definition

Equilibrium : A tuple of actions from which no player **gains** (a strict gain) by a **unilateral deviation**. A unilateral deviation from a tuple of actions $(a_1, a_2, ..., a_n)$ is a tuple of actions of the form $(b_1, b_2, ..., b_n)$ where $a_i = b_i$, $\forall i \neq j$ for some j ie, the actions of exactly one player j differs from the previous set.

Observe that for the game described above, no such equilibrium exists. This can be shown by analysing all possible pairs of actions.

• (L,L) : Player 1 gains additional utility by (C, L) or (R, L). Similarly for (C, C) and (R, R), player 1 stands to gain by unilaterally deviating.

Equilibrium



Definition

Equilibrium : A tuple of actions from which no player **gains** (a strict gain) by a **unilateral deviation**. A unilateral deviation from a tuple of actions $(a_1, a_2, ..., a_n)$ is a tuple of actions of the form $(b_1, b_2, ..., b_n)$ where $a_i = b_i$, $\forall i \neq j$ for some j ie, the actions of exactly one player j differs from the previous set.

Observe that for the game described above, no such equilibrium exists. This can be shown by analysing all possible pairs of actions.

- (L,L) : Player 1 gains additional utility by (C, L) or (R, L). Similarly for (C, C) and (R, R), player 1 stands to gain by unilaterally deviating.
- (L,R) : Player 2 gains additional utility by (L, L). Similarly for any tuple (A, B), where $A \neq B$, if player 2 deviates and chooses the action A, there is a utility gain.

Equilibrium



Definition

Equilibrium : A tuple of actions from which no player **gains** (a strict gain) by a **unilateral deviation**. A unilateral deviation from a tuple of actions $(a_1, a_2, ..., a_n)$ is a tuple of actions of the form $(b_1, b_2, ..., b_n)$ where $a_i = b_i$, $\forall i \neq j$ for some j ie, the actions of exactly one player j differs from the previous set.

Observe that for the game described above, no such equilibrium exists. This can be shown by analysing all possible pairs of actions.

- (L,L) : Player 1 gains additional utility by (C, L) or (R, L). Similarly for (C, C) and (R, R), player 1 stands to gain by unilaterally deviating.
- (L,R) : Player 2 gains additional utility by (L, L). Similarly for any tuple (A, B), where $A \neq B$, if player 2 deviates and chooses the action A, there is a utility gain.





Definition

Equilibrium : A tuple of actions from which no player **gains** (a strict gain) by a **unilateral deviation**. A unilateral deviation from a tuple of actions $(a_1, a_2, ..., a_n)$ is a tuple of actions of the form $(b_1, b_2, ..., b_n)$ where $a_i = b_i$, $\forall i \neq j$ for some j ie, the actions of exactly one player j differs from the previous set.

Observe that for the game described above, no such equilibrium exists. This can be shown by analysing all possible pairs of actions.

- (L,L) : Player 1 gains additional utility by (C, L) or (R, L). Similarly for (C, C) and (R, R), player 1 stands to gain by unilaterally deviating.
- (L,R) : Player 2 gains additional utility by (L,L). Similarly for any tuple (A,B), where $A \neq B$, if player 2 deviates and chooses the action A, there is a utility gain.

Therefore there exists no such equilibrium for the game.





Consider a modification of the above game where the shooter can score by aiming to the left regardless of whether the goal keeper defends the left side or not.





Consider a modification of the above game where the shooter can score by aiming to the left regardless of whether the goal keeper defends the left side or not.







Consider a modification of the above game where the shooter can score by aiming to the left regardless of whether the goal keeper defends the left side or not.



In this case , (L, L) is a simultaneous move equilibrium for game 2 because, the agent (player 1) will not gain any extra utility from (C, L) or (R, L) and similarly the opponent (player 2) will not gain any extra utility from (L, C) or (L, R).





The agent tries to maximise his utility at each step and the opponent tries to minimise the agent's utility thereby maximising his own utility. :



The agent tries to maximise his utility at each step and the opponent tries to minimise the agent's utility thereby maximising his own utility. :

• For every row of the matrix calculate the minimum utility possible. The value corresponding to row *i* denotes the least utility for the agent upon performing move *i*. Maximize over these values. This quantity is denoted as $\max_{s1} \min_{s2} u(s1, s2)$.



The agent tries to maximise his utility at each step and the opponent tries to minimise the agent's utility thereby maximising his own utility. :

- For every row of the matrix calculate the minimum utility possible. The value corresponding to row *i* denotes the least utility for the agent upon performing move *i*. Maximize over these values. This quantity is denoted as $\max_{s_1} \min_{s_2} u(s_1, s_2)$.
- Similarly, for every column of the matrix calculate the maximum utility possible. The values corresponding to column *j* denotes the most utility the agent can obtain if the opponent performs move *j*. Minimize over these values. This quantity is denoted as $\min_{s^2} \max_{s_1} u(s_1, s_2)$.



Lemma

 $\max_{s_1} \min_{s_2} u(s_1, s_2) \leqslant \min_{s_2} \max_{s_1} u(s_1, s_2).$



$$\underbrace{\text{Lemma}}_{s_1} \min_{s_2} u(s_1, s_2) \leqslant \min_{s_2} \max_{s_1} u(s_1, s_2).$$

If for some game, these two quantities equal each other, we have an equilibrium and the equilibrium point is called the saddle point.



Proof.

For any (s1, s2),





Proof.

For any (s1, s2),

$$u(s_1, s_2) \leq \max_{s_1} u(s_1, s_2)$$
 and
 $\min_{s_2} u(s_1, s_2) \leq u(s_1, s_2)$
 $\Rightarrow \min_{s_2} u(s_1, s_2) \leq \max_{s_1} u(s_1, s_2)$

Since the previous inequality is true $\forall s_1$, it is also true for s_1^* which maximises $\min_{s_2} u(s_1, s_2)$. Therefore we have,

$$\min_{s_2} u(s_1^*, s_2) \leqslant \max_{s_1} u(s_1, s_2)$$

$$\Rightarrow \max_{s_1} \min_{s_2} u(s_1, s_2) \leqslant \max_{s_1} u(s_1, s_2)$$



Proof.

For any (s1, s2),

$$u(s_1, s_2) \leq \max_{s_1} u(s_1, s_2)$$
 and

$$\Rightarrow \min_{s_2} u(s_1, s_2) \leqslant \max_{s_1} u(s_1, s_2)$$

Since the previous inequality is true $\forall s_1$, it is also true for s_1^* which maximises $\min_{s_2} u(s_1, s_2)$. Therefore we have,

$$\min_{s_2} u(s_1^*, s_2) \leq \max_{s_1} u(s_1, s_2)$$

$$\Rightarrow \max_{s_1} \min_{s_2} u(s_1, s_2) \leq \max_{s_1} u(s_1, s_2)$$



Proof.

For any (s1, s2),

$$u(s_1, s_2) \leq \max_{s_1} u(s_1, s_2)$$
 and
 $\min_{s_2} u(s_1, s_2) \leq u(s_1, s_2)$
 $\Rightarrow \min_{s_2} u(s_1, s_2) \leq \max_{s_1} u(s_1, s_2)$



Proof.

For any (s1, s2),

$$u(s_1, s_2) \leq \max_{s_1} u(s_1, s_2)$$
 and
 $\min_{s_2} u(s_1, s_2) \leq u(s_1, s_2)$
 $\Rightarrow \min_{s_2} u(s_1, s_2) \leq \max_{s_1} u(s_1, s_2)$

Since the previous inequality is true $\forall s_1$, it is also true for s_1^* which maximises $\min_{s_2} u(s_1, s_2)$.



Proof.

For any (s1, s2),

$$u(s_1, s_2) \leq \max_{s_1} u(s_1, s_2)$$
 and
 $\min_{s_2} u(s_1, s_2) \leq u(s_1, s_2)$
 $\Rightarrow \min_{s_2} u(s_1, s_2) \leq \max_{s_1} u(s_1, s_2)$

Since the previous inequality is true $\forall s_1$, it is also true for s_1^* which maximises $\min_{s_2} u(s_1, s_2)$. Therefore we have,

$$\min_{s_2} u(s_1^*, s_2) \leqslant \max_{s_1} u(s_1, s_2)$$

$$\Rightarrow \max_{s_1} \min_{s_2} u(s_1, s_2) \leqslant \max_{s_1} u(s_1, s_2)$$



Proof.

For any (s1, s2),

$$u(s_1, s_2) \leq \max_{s_1} u(s_1, s_2)$$
 and
 $\min_{s_2} u(s_1, s_2) \leq u(s_1, s_2)$
 $\Rightarrow \min_{s_2} u(s_1, s_2) \leq \max_{s_1} u(s_1, s_2)$

Since the previous inequality is true $\forall s_1$, it is also true for s_1^* which maximises $\min_{s_2} u(s_1, s_2)$. Therefore we have,

$$\min_{s_2} u(s_1^*, s_2) \leqslant \max_{s_1} u(s_1, s_2)$$

$$\Rightarrow \max_{s_1} \min_{s_2} u(s_1, s_2) \leqslant \max_{s_1} u(s_1, s_2)$$



Proof.

Similarly, since the previous inequality is true $\forall s_2$, it is also true for s_2^* which minimises $\max_{s_1} u(s_1, s_2)$. Therefore we have,

$$\max_{s_1} \min_{s_2} u(s_1, s_2) \leq \max_{s_1} u(s_1, s_2^*)$$

$$\Rightarrow \max_{s_1} \min_{s_2} u(s_1, s_2) \leq \min_{s_2} \max_{s_1} u(s_1, s_2)$$

And we are done with the proof.



भारतीय प्रौद्योगिकी संस्थान मुंबई Indian Institute of Technology Bombay