



भारतीय प्रौद्योगिकी संस्थान मुंबई
Indian Institute of Technology Bombay

CS 6001: Game Theory and Algorithmic Mechanism Design

Week 3 – Matrix Games and Mixed Strategy Nash Equilibrium

Swaprava Nath

August, 2026

Week 3 Roadmap



Lecture	Chunks (20–25 min each)	Quizzes
Lecture 5	1. Matrix Games & Saddle Points 2. Saddle Point = PSNE 3. When Does a PSNE Exist? 4. Mixed Strategies	Quiz 1 after Chunk 2 Quiz 2 after Chunk 4
Lecture 6	1. MSNE 2. Support & Characterization 3. Proof of Characterization 4. Algorithm, Dominance, Existence	Quiz 1 after Chunk 2 Quiz 2 after Chunk 4



► Lecture 5: Matrix Games, Saddle Points, and Security Values

► Lecture 6: Mixed Strategy Nash Equilibrium



Chunk 1

Matrix Games, Best Replies, and Saddle Points

($\approx$ 20–25 minutes)

Matrix Games: Two-Player Zero-Sum Games



A special class with certain nice **security** and **stability** properties



Matrix Games: Two-Player Zero-Sum Games

A special class with certain nice **security** and **stability** properties

Definition (Two player zero-sum games)

A NFG $\langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$ with $N = \{1, 2\}$ and $u_1 + u_2 \equiv 0$



Matrix Games: Two-Player Zero-Sum Games

A special class with certain nice **security** and **stability** properties

Definition (Two player zero-sum games)

A NFG $\langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$ with $N = \{1, 2\}$ and $u_1 + u_2 \equiv 0$

Where these show up

One side's gain is exactly the other's loss: a red-team attacker racing a blue-team detector, or a penalty shooter against a goalkeeper.



Matrix Games: Two-Player Zero-Sum Games

A special class with certain nice **security** and **stability** properties

Definition (Two player zero-sum games)

A NFG $\langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$ with $N = \{1, 2\}$ and $u_1 + u_2 \equiv 0$

Where these show up

One side's gain is exactly the other's loss: a red-team attacker racing a blue-team detector, or a penalty shooter against a goalkeeper.

Question

Why called **matrix** game?



Matrix Games: Two-Player Zero-Sum Games

A special class with certain nice **security** and **stability** properties

Definition (Two player zero-sum games)

A NFG $\langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$ with $N = \{1, 2\}$ and $u_1 + u_2 \equiv 0$

Where these show up

One side's gain is exactly the other's loss: a red-team attacker racing a blue-team detector, or a penalty shooter against a goalkeeper.

Question

Why called **matrix** game?

Answer

Possible to represent the game with only one matrix considering the utilities of player 1; player 2's utilities are negative of this matrix

The Ghost and the Sentinel



Ghost probes ingress channel L or R; **Sentinel** watches one channel. A match means Sentinel catches Ghost; a mismatch means Ghost slips through.

		Sentinel	
		L	R
Ghost	L	-1, 1	1, -1
	R	1, -1	-1, 1



The Ghost and the Sentinel

Ghost probes ingress channel L or R; **Sentinel** watches one channel. A match means Sentinel catches Ghost; a mismatch means Ghost slips through.

		Sentinel	
		L	R
Ghost	L	-1, 1	1, -1
	R	1, -1	-1, 1

$\Rightarrow$

$$\begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix} =: u$$

		L	R	maxmin
Ghost	L	-1	1	-1
	R	1	-1	-1
	minmax	1	1	



The Ghost and the Sentinel

Ghost probes ingress channel L or R; **Sentinel** watches one channel. A match means Sentinel catches Ghost; a mismatch means Ghost slips through.

		Sentinel		
		L	R	
Ghost	L	-1, 1	1, -1	$\Rightarrow \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix} =: u$
	R	1, -1	-1, 1	

Sentinel's **maxmin** value is the negative of the **minmax** value of this matrix

		L	R	maxmin
Ghost	L	-1	1	-1
	R	1	-1	-1
	minmax	1	1	

Scaling Up: Three Attack Vectors



Now suppose the exercise gets bigger. Ghost has three moves: T (Trojan payload), M (Man-in-the-middle), B (Botnet flood). Sentinel has three defenses: L (Log monitoring), C (Content filtering), R (Rate limiting). Payoffs are Ghost's threat-bounty score.

		Sentinel		
		L	C	R
Ghost	T	3, -3	-5, 5	-2, 2
	M	1, -1	4, -4	1, -1
	B	6, -6	-3, 3	-5, 5

Scaling Up: Three Attack Vectors



Now suppose the exercise gets bigger. Ghost has three moves: T (Trojan payload), M (Man-in-the-middle), B (Botnet flood). Sentinel has three defenses: L (Log monitoring), C (Content filtering), R (Rate limiting). Payoffs are Ghost's threat-bounty score.

		Sentinel					L	C	R	maxmin
		L	C	R			T			
Ghost	T	3, -3	-5, 5	-2, 2	Ghost	T	3	-5	-2	-5
	M	1, -1	4, -4	1, -1		M	1	4	1	1
	B	6, -6	-3, 3	-5, 5		B	6	-3	-5	-5
minmax							6	4	1	

Two Duels Side by Side



Ghost

		L	R	maxmin
L		-1	1	-1
R		1	-1	-1
minmax		1	1	

Ghost

		L	C	R	maxmin
T		3	-5	-2	-5
M		1	4	1	1
B		6	-3	-5	-5
minmax		6	4	1	

Two Duels Side by Side



Ghost

		L	R	maxmin
L		-1	1	-1
R		1	-1	-1
	minmax	1	1	

Ghost

		L	C	R	maxmin
T		3	-5	-2	-5
M		1	4	1	1
B		6	-3	-5	-5
	minmax	6	4	1	

Question

What are the PSNEs for the above games?



Saddle point of a matrix

The value is simultaneously the maximum in its column and minimum in its row i.e., maximum for player 1 and minimum for player 2



Saddle point of a matrix

The value is simultaneously the maximum in its column and minimum in its row i.e., maximum for player 1 and minimum for player 2

Question

What are the saddle points for the previous two duels?

Saddle Point, Continued



		L	R
Ghost	L	-1	1
	R	1	-1

		L	C	R
Ghost	T	3	-5	-2
	M	1	4	1
	B	6	-3	-5

Saddle Point, Continued



		L	R
Ghost	L	-1	1
	R	1	-1

		L	C	R
Ghost	T	3	-5	-2
	M	1	4	1
	B	6	-3	-5

Answer

For the two-channel exercise: no saddle point, for the three-vector exercise: (M,R)

Saddle Point, Continued



		L	R
Ghost	L	-1	1
	R	1	-1

		L	C	R
Ghost	T	3	-5	-2
	M	1	4	1
	B	6	-3	-5

Answer

For the two-channel exercise: no saddle point, for the three-vector exercise: (M,R)

Theorem

In a matrix game with utility matrix u , (s_1^, s_2^*) is a saddle point iff it is a PSNE.*



Chunk 2

Saddle Point = PSNE, and Security Values

($\approx$ 20–25 minutes)

Saddle Point and PSNE



Proof.

Consider (s_1^*, s_2^*) to be a saddle point. By definition of saddle point, this happens iff $u(s_1^*, s_2^*) \geq u(s_1, s_2^*), \forall s_1 \in S_1$ and $u(s_1^*, s_2^*) \leq u(s_1^*, s_2), \forall s_2 \in S_2$. Since, $u \equiv u_1 \equiv -u_2$, the above is equivalent to $u_1(s_1^*, s_2^*) \geq u_1(s_1, s_2^*), \forall s_1 \in S_1$ and $u_2(s_1^*, s_2^*) \geq u_2(s_1^*, s_2), \forall s_2 \in S_2 \Leftrightarrow (s_1^*, s_2^*)$ is a PSNE. □

Saddle Point and PSNE



Proof.

Consider (s_1^*, s_2^*) to be a saddle point. By definition of saddle point, this happens iff $u(s_1^*, s_2^*) \geq u(s_1, s_2^*), \forall s_1 \in S_1$ and $u(s_1^*, s_2^*) \leq u(s_1^*, s_2), \forall s_2 \in S_2$. Since, $u \equiv u_1 \equiv -u_2$, the above is equivalent to $u_1(s_1^*, s_2^*) \geq u_1(s_1, s_2^*), \forall s_1 \in S_1$ and $u_2(s_1^*, s_2^*) \geq u_2(s_1^*, s_2), \forall s_2 \in S_2 \Leftrightarrow (s_1^*, s_2^*)$ is a PSNE. □

Consider **maxmin** and **minmax** values

$$\underline{v} = \max_{s_1 \in S_1} \min_{s_2 \in S_2} u(s_1, s_2)$$

maxmin

$$\bar{v} = \min_{s_2 \in S_2} \max_{s_1 \in S_1} u(s_1, s_2)$$

minmax

Saddle Point and PSNE



Proof.

Consider (s_1^*, s_2^*) to be a saddle point. By definition of saddle point, this happens iff $u(s_1^*, s_2^*) \geq u(s_1, s_2^*), \forall s_1 \in S_1$ and $u(s_1^*, s_2^*) \leq u(s_1^*, s_2), \forall s_2 \in S_2$. Since, $u \equiv u_1 \equiv -u_2$, the above is equivalent to $u_1(s_1^*, s_2^*) \geq u_1(s_1, s_2^*), \forall s_1 \in S_1$ and $u_2(s_1^*, s_2^*) \geq u_2(s_1^*, s_2), \forall s_2 \in S_2 \Leftrightarrow (s_1^*, s_2^*)$ is a PSNE. □

Consider **maxmin** and **minmax** values

$$\underline{v} = \max_{s_1 \in S_1} \min_{s_2 \in S_2} u(s_1, s_2)$$

maxmin

$$\bar{v} = \min_{s_2 \in S_2} \max_{s_1 \in S_1} u(s_1, s_2)$$

minmax

Question

How are the maxmin and minmax values related?

Relationship of the Security Values



Lemma

For matrix games $\bar{v} \geq \underline{v}$.

Relationship of the Security Values



Lemma

For matrix games $\bar{v} \geq \underline{v}$.

Proof.

$$u(s_1, s_2) \geq \min_{t_2 \in S_2} u(s_1, t_2), \quad \forall s_1, s_2,$$

definition of min



Relationship of the Security Values



Lemma

For matrix games $\bar{v} \geq \underline{v}$.

Proof.

$$\begin{aligned} u(s_1, s_2) &\geq \min_{t_2 \in S_2} u(s_1, t_2), \quad \forall s_1, s_2, && \text{definition of min} \\ \Rightarrow \max_{t_1 \in S_1} u(t_1, s_2) &\geq \max_{t_1 \in S_1} \min_{t_2 \in S_2} u(t_1, t_2), \quad \forall s_2 \in S_2 && \text{RHS was dominated for each } s_1 \end{aligned}$$





Relationship of the Security Values

Lemma

For matrix games $\bar{v} \geq \underline{v}$.

Proof.

$$u(s_1, s_2) \geq \min_{t_2 \in S_2} u(s_1, t_2), \quad \forall s_1, s_2,$$

definition of min

$$\Rightarrow \max_{t_1 \in S_1} u(t_1, s_2) \geq \max_{t_1 \in S_1} \min_{t_2 \in S_2} u(t_1, t_2), \quad \forall s_2 \in S_2$$

RHS was dominated for each s_1

$$\Rightarrow \min_{t_2 \in S_2} \max_{t_1 \in S_1} u(t_1, t_2) \geq \max_{t_1 \in S_1} \min_{t_2 \in S_2} u(t_1, t_2)$$

RHS was a constant





Relationship of the Security Values

Lemma

For matrix games $\bar{v} \geq \underline{v}$.

Proof.

$$u(s_1, s_2) \geq \min_{t_2 \in S_2} u(s_1, t_2), \quad \forall s_1, s_2, \quad \text{definition of min}$$

$$\Rightarrow \max_{t_1 \in S_1} u(t_1, s_2) \geq \max_{t_1 \in S_1} \min_{t_2 \in S_2} u(t_1, t_2), \quad \forall s_2 \in S_2 \quad \text{RHS was dominated for each } s_1$$

$$\Rightarrow \min_{t_2 \in S_2} \max_{t_1 \in S_1} u(t_1, t_2) \geq \max_{t_1 \in S_1} \min_{t_2 \in S_2} u(t_1, t_2) \quad \text{RHS was a constant}$$



Reading it in the exercise

Sentinel's minimum highest value (minmax) never falls below Ghost's maximum lowest utility (maxmin).

Revisiting the Two Duels: No PSNE



		L	R	maxmin
Ghost	L	-1	1	-1
	R	1	-1	-1
	minmax	1	1	

Revisiting the Two Duels: No PSNE



		L	R	maxmin
Ghost	L	-1	1	-1
	R	1	-1	-1
	minmax	1	1	

$$\bar{v} = 1 > -1 = \underline{v}$$



Revisiting the Two Duels: No PSNE

		L	R	maxmin
Ghost	L	-1	1	-1
	R	1	-1	-1
	minmax	1	1	

$$\bar{v} = 1 > -1 = \underline{v}$$

PSNE does not exist – Ghost and Sentinel must keep guessing

Revisiting the Two Duels: PSNE Exists



		L	C	R	maxmin
Ghost	T	3	-5	-2	-5
	M	1	4	1	1
	B	6	-3	-5	-5
minmax		6	4	1	

Revisiting the Two Duels: PSNE Exists



		L	C	R	maxmin
Ghost	T	3	-5	-2	-5
	M	1	4	1	1
	B	6	-3	-5	-5
minmax		6	4	1	

$$\bar{v} = 1 = \underline{v}$$



Revisiting the Two Duels: PSNE Exists

		L	C	R	maxmin
Ghost	T	3	-5	-2	-5
	M	1	4	1	1
	B	6	-3	-5	-5
minmax		6	4	1	

$$\bar{v} = 1 = \underline{v}$$

PSNE exists – with three vectors on the table, Sentinel's rate-limiting defense (R) settles the exercise at (M,R)



In-Class Quiz 1

Open the Google Form now.

Questions refer only to what was just discussed –
you cannot answer them without having been in class.



Chunk 3

When Does a Matrix Game Have a PSNE?

($\approx$ 20–25 minutes)



Define the following strategies

$$s_1^* \in \arg \max_{s_1 \in S_1} \min_{s_2 \in S_2} u(s_1, s_2),$$

maxmin strategy of player 1

$$s_2^* \in \arg \min_{s_2 \in S_2} \max_{s_1 \in S_1} u(s_1, s_2),$$

minmax strategy of player 2



Define the following strategies

$$s_1^* \in \arg \max_{s_1 \in S_1} \min_{s_2 \in S_2} u(s_1, s_2),$$

maxmin strategy of player 1

$$s_2^* \in \arg \min_{s_2 \in S_2} \max_{s_1 \in S_1} u(s_1, s_2),$$

minmax strategy of player 2

Theorem

A game has a PSNE (equivalently, a saddle point) if and only if $\bar{v} = \underline{v} = u(s_1^, s_2^*)$, where s_1^* and s_2^* are maxmin and minmax strategies for players 1 and 2 respectively.*

Corollary: (s_1^*, s_2^*) is a PSNE

Proof of the PSNE Theorem



Proof

($\implies$) let (s_1^*, s_2^*) is a PSNE $\implies \bar{v} = \underline{v} = u(s_1^*, s_2^*)$ and s_1^* and s_2^* are maxmin and minmax

Proof of the PSNE Theorem



Proof

($\implies$) let (s_1^*, s_2^*) is a PSNE $\implies \bar{v} = \underline{v} = u(s_1^*, s_2^*)$ and s_1^* and s_2^* are maxmin and minmax
Since (s_1^*, s_2^*) is a PSNE, $u(s_1^*, s_2^*) \geq u(s_1, s_2^*), \forall s_1 \in S_1$.

Proof of the PSNE Theorem



Proof

($\implies$) let (s_1^*, s_2^*) is a PSNE $\implies \bar{v} = \underline{v} = u(s_1^*, s_2^*)$ and s_1^* and s_2^* are maxmin and minmax
Since (s_1^*, s_2^*) is a PSNE, $u(s_1^*, s_2^*) \geq u(s_1, s_2^*), \forall s_1 \in S_1$.

$$\begin{aligned} \implies u(s_1^*, s_2^*) &\geq \max_{t_1 \in S_1} u(t_1, s_2^*) \\ &\geq \min_{t_2 \in S_2} \max_{t_1 \in S_1} u(t_1, t_2), \text{ since } s_2^* \text{ is a specific strategy} \\ &= \bar{v} \end{aligned}$$

Proof of the PSNE Theorem



Proof

($\implies$) let (s_1^*, s_2^*) is a PSNE $\implies \bar{v} = \underline{v} = u(s_1^*, s_2^*)$ and s_1^* and s_2^* are maxmin and minmax
Since (s_1^*, s_2^*) is a PSNE, $u(s_1^*, s_2^*) \geq u(s_1, s_2^*), \forall s_1 \in S_1$.

$$\begin{aligned} \implies u(s_1^*, s_2^*) &\geq \max_{t_1 \in S_1} u(t_1, s_2^*) \\ &\geq \min_{t_2 \in S_2} \max_{t_1 \in S_1} u(t_1, t_2), \text{ since } s_2^* \text{ is a specific strategy} \\ &= \bar{v} \end{aligned}$$

Similarly, using the same argument for player 2, we get $\underline{v} \geq u(s_1^*, s_2^*)$

Proof of the PSNE Theorem



Proof

($\implies$) let (s_1^*, s_2^*) is a PSNE $\implies \bar{v} = \underline{v} = u(s_1^*, s_2^*)$ and s_1^* and s_2^* are maxmin and minmax
Since (s_1^*, s_2^*) is a PSNE, $u(s_1^*, s_2^*) \geq u(s_1, s_2^*), \forall s_1 \in S_1$.

$$\begin{aligned}\implies u(s_1^*, s_2^*) &\geq \max_{t_1 \in S_1} u(t_1, s_2^*) \\ &\geq \min_{t_2 \in S_2} \max_{t_1 \in S_1} u(t_1, t_2), \text{ since } s_2^* \text{ is a specific strategy} \\ &= \bar{v}\end{aligned}$$

Similarly, using the same argument for player 2, we get $\underline{v} \geq u(s_1^*, s_2^*)$

But $\bar{v} \geq \underline{v}$ (from the previous lemma), hence

Proof of the PSNE Theorem



Proof

($\implies$) let (s_1^*, s_2^*) is a PSNE $\implies \bar{v} = \underline{v} = u(s_1^*, s_2^*)$ and s_1^* and s_2^* are maxmin and minmax
Since (s_1^*, s_2^*) is a PSNE, $u(s_1^*, s_2^*) \geq u(s_1, s_2^*), \forall s_1 \in S_1$.

$$\begin{aligned}\implies u(s_1^*, s_2^*) &\geq \max_{t_1 \in S_1} u(t_1, s_2^*) \\ &\geq \min_{t_2 \in S_2} \max_{t_1 \in S_1} u(t_1, t_2), \text{ since } s_2^* \text{ is a specific strategy} \\ &= \bar{v}\end{aligned}$$

Similarly, using the same argument for player 2, we get $\underline{v} \geq u(s_1^*, s_2^*)$

But $\bar{v} \geq \underline{v}$ (from the previous lemma), hence

$$\begin{aligned}u(s_1^*, s_2^*) &\geq \bar{v} \geq \underline{v} \geq u(s_1^*, s_2^*) \\ \implies u(s_1^*, s_2^*) &= \bar{v} = \underline{v}\end{aligned}$$



Proof of the PSNE Theorem

Proof

($\implies$) let (s_1^*, s_2^*) is a PSNE $\implies \bar{v} = \underline{v} = u(s_1^*, s_2^*)$ and s_1^* and s_2^* are maxmin and minmax
Since (s_1^*, s_2^*) is a PSNE, $u(s_1^*, s_2^*) \geq u(s_1, s_2^*), \forall s_1 \in S_1$.

$$\begin{aligned}\implies u(s_1^*, s_2^*) &\geq \max_{t_1 \in S_1} u(t_1, s_2^*) \\ &\geq \min_{t_2 \in S_2} \max_{t_1 \in S_1} u(t_1, t_2), \text{ since } s_2^* \text{ is a specific strategy} \\ &= \bar{v}\end{aligned}$$

Similarly, using the same argument for player 2, we get $\underline{v} \geq u(s_1^*, s_2^*)$

But $\bar{v} \geq \underline{v}$ (from the previous lemma), hence

$$\begin{aligned}u(s_1^*, s_2^*) &\geq \bar{v} \geq \underline{v} \geq u(s_1^*, s_2^*) \\ \implies u(s_1^*, s_2^*) &= \bar{v} = \underline{v}\end{aligned}$$

Also implies that the maxmin for player 1 and minmax for player 2 are s_1^* and s_2^* respectively.

Proof of the PSNE Theorem (contd.)



Proof (contd.)

($\Leftarrow$) i.e. $\bar{v} = \underline{v} = u(s_1^*, s_2^*) = v$ (say) and s_1^* and s_2^* are maxmin and minmax $\implies (s_1^*, s_2^*)$ is a PSNE

Proof of the PSNE Theorem (contd.)



Proof (contd.)

($\Leftarrow$) i.e. $\bar{v} = \underline{v} = u(s_1^*, s_2^*) = v$ (say) and s_1^* and s_2^* are maxmin and minmax $\implies (s_1^*, s_2^*)$ is a PSNE

$$\begin{aligned} u(s_1^*, s_2) &\geq \min_{t_2 \in S_2} u(s_1^*, t_2), \text{ by definition of min} \\ &= \max_{t_1 \in S_1} \min_{t_2 \in S_2} u(t_1, t_2), \text{ since } s_1^* \text{ is the maxmin strategy for player 1} \\ &= v \text{ (given)} \end{aligned}$$

Proof of the PSNE Theorem (contd.)



Proof (contd.)

($\Leftarrow$) i.e. $\bar{v} = \underline{v} = u(s_1^*, s_2^*) = v$ (say) and s_1^* and s_2^* are maxmin and minmax $\implies (s_1^*, s_2^*)$ is a PSNE

$$\begin{aligned} u(s_1^*, s_2) &\geq \min_{t_2 \in S_2} u(s_1^*, t_2), \text{ by definition of min} \\ &= \max_{t_1 \in S_1} \min_{t_2 \in S_2} u(t_1, t_2), \text{ since } s_1^* \text{ is the maxmin strategy for player 1} \\ &= v \text{ (given)} \end{aligned}$$

Similarly, we can show $u(s_1, s_2^*) \leq v, \forall s_1 \in S_1$

Proof of the PSNE Theorem (contd.)



Proof (contd.)

($\Leftarrow$) i.e. $\bar{v} = \underline{v} = u(s_1^*, s_2^*) = v$ (say) and s_1^* and s_2^* are maxmin and minmax $\implies (s_1^*, s_2^*)$ is a PSNE

$$\begin{aligned} u(s_1^*, s_2) &\geq \min_{t_2 \in S_2} u(s_1^*, t_2), \text{ by definition of min} \\ &= \max_{t_1 \in S_1} \min_{t_2 \in S_2} u(t_1, t_2), \text{ since } s_1^* \text{ is the maxmin strategy for player 1} \\ &= v \text{ (given)} \end{aligned}$$

Similarly, we can show $u(s_1, s_2^*) \leq v, \forall s_1 \in S_1$

But $v = u(s_1^*, s_2^*)$. Substitute and get that (s_1^*, s_2^*) is a PSNE.



Chunk 4

Mixed Strategies

($\approx$ 20–25 minutes)

Mixed Strategies



Mixed strategy: *probability distribution* over the set of strategies of that player.

		Sentinel	
		L	R
Ghost	L	$-1, 1$	$1, -1$
	R	$1, -1$	$-1, 1$

Mixed Strategies



Mixed strategy: *probability distribution* over the set of strategies of that player.

		Sentinel	
		L	R
Ghost	$\frac{2}{3}$ L	$-1, 1$	$1, -1$
	$\frac{1}{3}$ R	$1, -1$	$-1, 1$

Mixed Strategies



Mixed strategy: *probability distribution* over the set of strategies of that player.

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	-1, 1	1, -1
	$\frac{1}{3}$ R	1, -1	-1, 1

Mixed Strategies



Mixed strategy: *probability distribution* over the set of strategies of that player.

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	-1, 1	1, -1
	$\frac{1}{3}$ R	1, -1	-1, 1

- Consider a finite set A , define

$$\Delta A = \{p \in [0, 1]^{|A|} : \sum_{a \in A} p(a) = 1\}$$

Mixed Strategies



Mixed strategy: *probability distribution* over the set of strategies of that player.

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	-1, 1	1, -1
	$\frac{1}{3}$ R	1, -1	-1, 1

- Consider a finite set A , define

$$\Delta A = \{p \in [0, 1]^{|A|} : \sum_{a \in A} p(a) = 1\}$$

- Mixed strategy set of Ghost: $\Delta S_1 = \Delta\{L, R\}$, $(\frac{2}{3}, \frac{1}{3}) \in \Delta S_1$

Mixed Strategies, Notation



- **Notation:** σ_i is a mixed strategy of player i

Mixed Strategies, Notation



- **Notation:** σ_i is a mixed strategy of player i
- $\sigma_i \in \Delta S_i$, i.e. , $\sigma_i : S_i \rightarrow [0, 1]$ s.t. $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$

Mixed Strategies, Notation



- **Notation:** σ_i is a mixed strategy of player i
- $\sigma_i \in \Delta S_i$, i.e. , $\sigma_i : S_i \rightarrow [0, 1]$ s.t. $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$
- We are discussing non-cooperative games, the players choose their strategies independently

Mixed Strategies, Notation



- **Notation:** σ_i is a mixed strategy of player i
- $\sigma_i \in \Delta S_i$, i.e. , $\sigma_i : S_i \rightarrow [0, 1]$ s.t. $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$
- We are discussing non-cooperative games, the players choose their strategies independently
- The joint probability of player 1 picking s_1 and player 2 picking $s_2 = \sigma_1(s_1)\sigma_2(s_2)$



Mixed Strategies, Notation

- **Notation:** σ_i is a mixed strategy of player i
- $\sigma_i \in \Delta S_i$, i.e. , $\sigma_i : S_i \rightarrow [0, 1]$ s.t. $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$
- We are discussing non-cooperative games, the players choose their strategies independently
- The joint probability of player 1 picking s_1 and player 2 picking $s_2 = \sigma_1(s_1)\sigma_2(s_2)$
- Utility of player i at a mixed strategy profile (σ_i, σ_{-i}) is

$$u_i(\sigma_i, \sigma_{-i}) = \sum_{s_1 \in S_1} \sum_{s_2 \in S_2} \cdots \sum_{s_n \in S_n} \sigma_1(s_1) \cdot \sigma_2(s_2) \cdots \sigma_n(s_n) \textcolor{red}{u_i(s_1, s_2, \dots, s_n)}$$



Mixed Strategies, Notation

- **Notation:** σ_i is a mixed strategy of player i
- $\sigma_i \in \Delta S_i$, i.e. , $\sigma_i : S_i \rightarrow [0, 1]$ s.t. $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$
- We are discussing non-cooperative games, the players choose their strategies independently
- The joint probability of player 1 picking s_1 and player 2 picking $s_2 = \sigma_1(s_1)\sigma_2(s_2)$
- Utility of player i at a mixed strategy profile (σ_i, σ_{-i}) is

$$u_i(\sigma_i, \sigma_{-i}) = \sum_{s_1 \in S_1} \sum_{s_2 \in S_2} \cdots \sum_{s_n \in S_n} \sigma_1(s_1) \cdot \sigma_2(s_2) \cdots \sigma_n(s_n) \textcolor{red}{u_i(s_1, s_2, \dots, s_n)}$$

- We are *overloading* u_i to denote the utility at *pure* and *mixed* strategies



Mixed Strategies, Notation

- **Notation:** σ_i is a mixed strategy of player i
- $\sigma_i \in \Delta S_i$, i.e. , $\sigma_i : S_i \rightarrow [0, 1]$ s.t. $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$
- We are discussing non-cooperative games, the players choose their strategies independently
- The joint probability of player 1 picking s_1 and player 2 picking $s_2 = \sigma_1(s_1)\sigma_2(s_2)$
- Utility of player i at a mixed strategy profile (σ_i, σ_{-i}) is

$$u_i(\sigma_i, \sigma_{-i}) = \sum_{s_1 \in S_1} \sum_{s_2 \in S_2} \cdots \sum_{s_n \in S_n} \sigma_1(s_1) \cdot \sigma_2(s_2) \cdots \sigma_n(s_n) u_i(s_1, s_2, \dots, s_n)$$

- We are *overloading* u_i to denote the utility at *pure* and *mixed* strategies
- Utility at a mixed strategy is the **expectation** of the utilities at pure strategies; all the rules of expectation hold, e.g., linearity, conditional expectation, etc.

Computing Expected Utility



Ghost probes L with probability $\frac{2}{3}$ and R with probability $\frac{1}{3}$; Sentinel watches L with probability $\frac{4}{5}$ and R with probability $\frac{1}{5}$.

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	-1, 1	1, -1
	$\frac{1}{3}$ R	1, -1	-1, 1

Computing Expected Utility



Ghost probes L with probability $2/3$ and R with probability $1/3$; Sentinel watches L with probability $4/5$ and R with probability $1/5$.

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	$-1, 1$	$1, -1$
	$\frac{1}{3}$ R	$1, -1$	$-1, 1$

$$u_1(\sigma_1, \sigma_2)$$

Computing Expected Utility



Ghost probes L with probability $2/3$ and R with probability $1/3$; Sentinel watches L with probability $4/5$ and R with probability $1/5$.

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	$-1, 1$	$1, -1$
	$\frac{1}{3}$ R	$1, -1$	$-1, 1$

$$u_1(\sigma_1, \sigma_2) = \frac{2}{3} \cdot \frac{4}{5} \cdot (-1)$$

Computing Expected Utility



Ghost probes L with probability $2/3$ and R with probability $1/3$; Sentinel watches L with probability $4/5$ and R with probability $1/5$.

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	$-1, 1$	$1, -1$
	$\frac{1}{3}$ R	$1, -1$	$-1, 1$

$$u_1(\sigma_1, \sigma_2) = \frac{2}{3} \cdot \frac{4}{5} \cdot (-1) + \frac{2}{3} \cdot \frac{1}{5} \cdot (1)$$

Computing Expected Utility



Ghost probes L with probability $2/3$ and R with probability $1/3$; Sentinel watches L with probability $4/5$ and R with probability $1/5$.

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	$-1, 1$	$1, -1$
	$\frac{1}{3}$ R	$1, -1$	$-1, 1$

$$u_1(\sigma_1, \sigma_2) = \frac{2}{3} \cdot \frac{4}{5} \cdot (-1) + \frac{2}{3} \cdot \frac{1}{5} \cdot (1) + \frac{1}{3} \cdot \frac{4}{5} \cdot (1)$$

Computing Expected Utility



Ghost probes L with probability $2/3$ and R with probability $1/3$; Sentinel watches L with probability $4/5$ and R with probability $1/5$.

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	$-1, 1$	$1, -1$
	$\frac{1}{3}$ R	$1, -1$	$-1, 1$

$$u_1(\sigma_1, \sigma_2) = \frac{2}{3} \cdot \frac{4}{5} \cdot (-1) + \frac{2}{3} \cdot \frac{1}{5} \cdot (1) + \frac{1}{3} \cdot \frac{4}{5} \cdot (1) + \frac{1}{3} \cdot \frac{1}{5} \cdot (-1)$$



In-Class Quiz 2

Open the Google Form now.

Questions refer only to what was just discussed –
you cannot answer them without having been in class.

Lecture 5: What We Built



- Matrix (two-player zero-sum) games; saddle point $\Leftrightarrow$ PSNE

Lecture 5: What We Built



- Matrix (two-player zero-sum) games; saddle point $\Leftrightarrow$ PSNE
- Maxmin never exceeds minmax; a PSNE exists exactly when the two coincide, at the maxmin/minmax strategy pair

Lecture 5: What We Built



- Matrix (two-player zero-sum) games; saddle point $\Leftrightarrow$ PSNE
- Maxmin never exceeds minmax; a PSNE exists exactly when the two coincide, at the maxmin/minmax strategy pair
- Mixed strategies: a probability distribution over pure strategies, with utility computed as an expectation

Lecture 5: What We Built



- Matrix (two-player zero-sum) games; saddle point $\Leftrightarrow$ PSNE
- Maxmin never exceeds minmax; a PSNE exists exactly when the two coincide, at the maxmin/minmax strategy pair
- Mixed strategies: a probability distribution over pure strategies, with utility computed as an expectation
- **Next:** Mixed Strategy Nash Equilibrium and how to actually find one



End of Lecture 5

See you next class.



► Lecture 5: Matrix Games, Saddle Points, and Security Values

► Lecture 6: Mixed Strategy Nash Equilibrium



Chunk 1

Mixed Strategy Nash Equilibrium

($\approx$ 20–25 minutes)

Mixed Strategy Nash Equilibrium



Definition (Mixed Strategy Nash Equilibrium)

A **mixed strategy Nash equilibrium (MSNE)** is a mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$, s.t.

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*), \quad \forall \sigma_i \in \Delta S_i \text{ and } \forall i \in N.$$



Mixed Strategy Nash Equilibrium

Definition (Mixed Strategy Nash Equilibrium)

A **mixed strategy Nash equilibrium (MSNE)** is a mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$, s.t.

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*), \quad \forall \sigma_i \in \Delta S_i \text{ and } \forall i \in N.$$

Question

Relation between **PSNE** and **MSNE**?

Mixed Strategy Nash Equilibrium



Definition (Mixed Strategy Nash Equilibrium)

A **mixed strategy Nash equilibrium (MSNE)** is a mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$, s.t.

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*), \quad \forall \sigma_i \in \Delta S_i \text{ and } \forall i \in N.$$

Question

Relation between **PSNE** and **MSNE**?

Answer

PSNE $\implies$ MSNE



An Alternative Definition

Theorem

A mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$, is an **MSNE** if and only if $\forall s_i \in S_i$ and $\forall i \in N$

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*).$$



An Alternative Definition

Theorem

A mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$, is an **MSNE** if and only if $\forall s_i \in S_i$ and $\forall i \in N$

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*).$$

Proof.

($\Rightarrow$): The pure strategy s_i is a special case of the mixed strategy, the mixed strategy with s_i having probability 1. Inequality holds by definition of MSNE





An Alternative Definition

Theorem

A mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$, is an **MSNE** if and only if $\forall s_i \in S_i$ and $\forall i \in N$

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*).$$

Proof.

($\Rightarrow$): The pure strategy s_i is a special case of the mixed strategy, the mixed strategy with s_i having probability 1. Inequality holds by definition of MSNE

($\Leftarrow$) Pick an arbitrary mixed strategy σ_i of player i

$$u_i(\sigma_i, \sigma_{-i}^*) = \sum_{s_i \in S_i} \sigma_i(s_i) \cdot u_i(s_i, \sigma_{-i}^*)$$



An Alternative Definition

Theorem

A mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$, is an **MSNE** if and only if $\forall s_i \in S_i$ and $\forall i \in N$

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*).$$

Proof.

($\Rightarrow$): The pure strategy s_i is a special case of the mixed strategy, the mixed strategy with s_i having probability 1. Inequality holds by definition of MSNE

($\Leftarrow$) Pick an arbitrary mixed strategy σ_i of player i

$$u_i(\sigma_i, \sigma_{-i}^*) = \sum_{s_i \in S_i} \sigma_i(s_i) \cdot \underbrace{u_i(s_i, \sigma_{-i}^*)}_{\leq u_i(\sigma_i^*, \sigma_{-i}^*) \text{ (given)}}$$



An Alternative Definition

Theorem

A mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$, is an **MSNE** if and only if $\forall s_i \in S_i$ and $\forall i \in N$

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*).$$

Proof.

($\Rightarrow$): The pure strategy s_i is a special case of the mixed strategy, the mixed strategy with s_i having probability 1. Inequality holds by definition of MSNE

($\Leftarrow$) Pick an arbitrary mixed strategy σ_i of player i

$$\begin{aligned} u_i(\sigma_i, \sigma_{-i}^*) &= \sum_{s_i \in S_i} \sigma_i(s_i) \cdot u_i(s_i, \sigma_{-i}^*) \\ &\leq \sum_{s_i \in S_i} \sigma_i(s_i) \cdot u_i(\sigma_i^*, \sigma_{-i}^*) \end{aligned}$$



An Alternative Definition

Theorem

A mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$, is an **MSNE** if and only if $\forall s_i \in S_i$ and $\forall i \in N$

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*).$$

Proof.

($\Rightarrow$): The pure strategy s_i is a special case of the mixed strategy, the mixed strategy with s_i having probability 1. Inequality holds by definition of MSNE

($\Leftarrow$) Pick an arbitrary mixed strategy σ_i of player i

$$\begin{aligned} u_i(\sigma_i, \sigma_{-i}^*) &= \sum_{s_i \in S_i} \sigma_i(s_i) \cdot u_i(s_i, \sigma_{-i}^*) \\ &\leq \sum_{s_i \in S_i} \sigma_i(s_i) \cdot u_i(\sigma_i^*, \sigma_{-i}^*) \\ &= u_i(\sigma_i^*, \sigma_{-i}^*) \cdot \sum_{s_i \in S_i} \sigma_i(s_i) = u_i(\sigma_i^*, \sigma_{-i}^*) \end{aligned}$$



Checking a Candidate: Is This an MSNE?

Question

Does mixing $(2/3, 1/3)$ against $(4/5, 1/5)$ give an **MSNE**?

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	-1, 1	1, -1
	$\frac{1}{3}$ R	1, -1	-1, 1

- To answer this, we need to show that there does not exist any better mixed strategy for the player



Checking a Candidate: Is This an MSNE?

Question

Does mixing $(2/3, 1/3)$ against $(4/5, 1/5)$ give an **MSNE**?

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	-1, 1	1, -1
	$\frac{1}{3}$ R	1, -1	-1, 1

- To answer this, we need to show that there does not exist any better mixed strategy for the player
- Sentinel's expected utility from watching $L = 2/3 \cdot 1 + 1/3 \cdot (-1) = 1/3$



Checking a Candidate: Is This an MSNE?

Question

Does mixing $(2/3, 1/3)$ against $(4/5, 1/5)$ give an **MSNE**?

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	-1, 1	1, -1
	$\frac{1}{3}$ R	1, -1	-1, 1

- To answer this, we need to show that there does not exist any better mixed strategy for the player
- Sentinel's expected utility from watching $L = 2/3 \cdot 1 + 1/3 \cdot (-1) = 1/3$
- Sentinel's expected utility from watching $R = 2/3 \cdot (-1) + 1/3 \cdot 1 = -1/3$

Checking a Candidate: What Balances the Duel?



Question

Is the balanced mixed strategy profile an **MSNE**?

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	-1, 1	1, -1
	$\frac{1}{3}$ R	1, -1	-1, 1

- Sentinel's expected utility will increase if some probability is transferred from R to L



Checking a Candidate: What Balances the Duel?

Question

Is the balanced mixed strategy profile an **MSNE**?

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	-1, 1	1, -1
	$\frac{1}{3}$ R	1, -1	-1, 1

- Sentinel's expected utility will increase if some probability is transferred from R to L
- $\Rightarrow$ the $((2/3, 1/3), (4/5, 1/5))$ profile is **not** an MSNE



Checking a Candidate: What Balances the Duel?

Question

Is the balanced mixed strategy profile an **MSNE**?

		Sentinel	
		$\frac{4}{5}$ L	$\frac{1}{5}$ R
Ghost	$\frac{2}{3}$ L	-1, 1	1, -1
	$\frac{1}{3}$ R	1, -1	-1, 1

- Sentinel's expected utility will increase if some probability is transferred from R to L
- $\Rightarrow$ the $((2/3, 1/3), (4/5, 1/5))$ profile is **not** an MSNE
- **Some balance in the utilities is needed**



Checking a Candidate: What Balances the Duel?

Question

Is the balanced mixed strategy profile an **MSNE**?

		Sentinel	
		$\frac{1}{2}$ L	$\frac{1}{2}$ R
Ghost	$\frac{1}{2}$ L	-1, 1	1, -1
	$\frac{1}{2}$ R	1, -1	-1, 1

- Sentinel's expected utility will increase if some probability is transferred from R to L
- $\Rightarrow$ the $((2/3, 1/3), (4/5, 1/5))$ profile is **not** an MSNE
- **Some balance in the utilities is needed**
- **Does there exist any improving mixed strategy for the balanced $(1/2, 1/2)$ profile?**



Chunk 2

Support and the Characterization Theorem

($\approx$ 20–25 minutes)

Support of a Mixed Strategy



Definition (Support of mixed strategy/probability distribution)

For mixed strategy σ_i , the subset of strategy set of i on which σ_i has a positive mass is called the **support** of σ_i and is denoted by $\delta(\sigma_i)$. Formally, $\delta(\sigma_i) = \{s_i \in S_i : \sigma_i(s_i) > 0\}$.

Support of a Mixed Strategy



Definition (Support of mixed strategy/probability distribution)

For mixed strategy σ_i , the subset of strategy set of i on which σ_i has a positive mass is called the **support** of σ_i and is denoted by $\delta(\sigma_i)$. Formally, $\delta(\sigma_i) = \{s_i \in S_i : \sigma_i(s_i) > 0\}$.

Using the definition of support, here is a characterization of MSNE

Theorem

A mixed strategy profile $(\sigma_i^, \sigma_{-i}^*)$ is an MSNE iff^a $\forall i \in N$*

^aThis is a shorthand for 'if and only if'.



Support of a Mixed Strategy

Definition (Support of mixed strategy/probability distribution)

For mixed strategy σ_i , the subset of strategy set of i on which σ_i has a positive mass is called the **support** of σ_i and is denoted by $\delta(\sigma_i)$. Formally, $\delta(\sigma_i) = \{s_i \in S_i : \sigma_i(s_i) > 0\}$.

Using the definition of support, here is a characterization of MSNE

Theorem

A mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE iff^a $\forall i \in N$

• $u_i(s_i, \sigma_{-i}^*)$ is identical $\forall s_i \in \delta(\sigma_i^*)$,

^aThis is a shorthand for 'if and only if'.



Support of a Mixed Strategy

Definition (Support of mixed strategy/probability distribution)

For mixed strategy σ_i , the subset of strategy set of i on which σ_i has a positive mass is called the **support** of σ_i and is denoted by $\delta(\sigma_i)$. Formally, $\delta(\sigma_i) = \{s_i \in S_i : \sigma_i(s_i) > 0\}$.

Using the definition of support, here is a characterization of MSNE

Theorem

A mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE iff^a $\forall i \in N$

- ① $u_i(s_i, \sigma_{-i}^*)$ is identical $\forall s_i \in \delta(\sigma_i^*)$,
- ② $u_i(s_i, \sigma_{-i}^*) \geq u_i(s'_i, \sigma_{-i}^*), \forall s_i \in \delta(\sigma_i^*), s'_i \notin \delta(\sigma_i^*)$.

^aThis is a shorthand for 'if and only if'.

Solving the Ghost–Sentinel Duel: Ruling Out Cases



Back to the two-channel exercise

		Sentinel	
		L	R
Ghost	L	$-1, 1$	$1, -1$
	R	$1, -1$	$-1, 1$

Solving the Ghost–Sentinel Duel: Ruling Out Cases



Back to the two-channel exercise

		Sentinel	
		L	R
Ghost	L	$-1, 1$	$1, -1$
	R	$1, -1$	$-1, 1$

Case 1: support profile $(\{L\}, \{L\})$: for Ghost, $s'_1 = R$ – violates condition 2

Solving the Ghost–Sentinel Duel: Ruling Out Cases



Back to the two-channel exercise

		Sentinel	
		L	R
Ghost	L	-1, 1	1, -1
	R	1, -1	-1, 1

Case 1: support profile $(\{L\}, \{L\})$: for Ghost, $s'_1 = R$ – violates condition 2

Case 2: support profile $(\{L, R\}, \{L\})$ – symmetric for the other case

For Ghost, the expected utility has to be the same for L and R - **not possible** – violates condition 1

Solving the Ghost–Sentinel Duel: The Balanced Case



Case 3: support profile $(\{L, R\}, \{L, R\})$: condition 2 is vacuously satisfied

Solving the Ghost–Sentinel Duel: The Balanced Case



Case 3: support profile $(\{L, R\}, \{L, R\})$: condition 2 is vacuously satisfied

For condition 1, let Ghost chooses L w.p. p and Sentinel chooses L w.p. q

Solving the Ghost–Sentinel Duel: The Balanced Case



Case 3: support profile $(\{L, R\}, \{L, R\})$: condition 2 is vacuously satisfied

For condition 1, let Ghost chooses L w.p. p and Sentinel chooses L w.p. q

For Ghost:

$$u_1(L, (q, 1 - q)) = u_1(R, (q, 1 - q)) \Rightarrow (-1)q + 1 \cdot (1 - q) = 1 \cdot q + (-1)(1 - q) \Rightarrow q = \frac{1}{2}$$

Solving the Ghost–Sentinel Duel: The Balanced Case



Case 3: support profile $(\{L, R\}, \{L, R\})$: condition 2 is vacuously satisfied

For condition 1, let Ghost chooses L w.p. p and Sentinel chooses L w.p. q

For Ghost:

$$u_1(L, (q, 1 - q)) = u_1(R, (q, 1 - q)) \Rightarrow (-1)q + 1 \cdot (1 - q) = 1 \cdot q + (-1)(1 - q) \Rightarrow q = \frac{1}{2}$$

For Sentinel:

$$u_2((p, 1 - p), L) = u_2((p, 1 - p), R) \Rightarrow p = \frac{1}{2}$$

Solving the Ghost–Sentinel Duel: The Balanced Case



Case 3: support profile $(\{L, R\}, \{L, R\})$: condition 2 is vacuously satisfied

For condition 1, let Ghost chooses L w.p. p and Sentinel chooses L w.p. q

For Ghost:

$$u_1(L, (q, 1 - q)) = u_1(R, (q, 1 - q)) \Rightarrow (-1)q + 1 \cdot (1 - q) = 1 \cdot q + (-1)(1 - q) \Rightarrow q = \frac{1}{2}$$

For Sentinel:

$$u_2((p, 1 - p), L) = u_2((p, 1 - p), R) \Rightarrow p = \frac{1}{2}$$

MSNE =

$$\left(\left(\frac{1}{2}, \frac{1}{2} \right), \left(\frac{1}{2}, \frac{1}{2} \right) \right)$$

Exercise 1



Back to Riya and Karan's Chess-or-Football game night from Week 2 – now find the MSNE.

		Friend 2	
		F	C
Friend 1	F	2,1	0,0
	C	0,0	1,2

Exercise 2



Now suppose Friend 2 gets a third option, D (a documentary neither prefers, but at least it is not a total mismatch):

		Friend 2		
		F	C	D
Friend 1	F	2, 1	0, 0	1, 1
	C	0, 0	1, 2	2, 0



In-Class Quiz 1

Open the Google Form now.

Questions refer only to what was just discussed –
you cannot answer them without having been in class.



Chunk 3

Proof of the Characterization Theorem

($\approx$ 25–30 minutes)

MSNE Characterization Theorem



Theorem

A mixed strategy profile is an MSNE iff $\forall i \in N$

- ➊ $u_i(s_i, \sigma_{-i}^*)$ is identical $\forall s_i \in \delta(\sigma_i^*)$,
- ➋ $u_i(s_i, \sigma_{-i}^*) \geq u_i(s'_i, \sigma_{-i}^*), \forall s_i \subseteq \delta(\sigma_i^*), s'_i \notin \delta(\sigma_i^*)$.

MSNE Characterization Theorem



Theorem

A mixed strategy profile is an MSNE iff $\forall i \in N$

- 1 $u_i(s_i, \sigma_{-i}^*)$ is identical $\forall s_i \in \delta(\sigma_i^*)$,
- 2 $u_i(s_i, \sigma_{-i}^*) \geq u_i(s'_i, \sigma_{-i}^*), \forall s_i \subseteq \delta(\sigma_i^*), s'_i \notin \delta(\sigma_i^*)$.

Observations:

- $\max_{\sigma_i \in \Delta S_i} u_i(\sigma_i, \sigma_{-i}) = \max_{s_i \in S_i} u_i(s_i, \sigma_{-i})$
maximizing w.r.t. a distribution $\Leftrightarrow$ whole probability mass at max



MSNE Characterization Theorem

Theorem

A mixed strategy profile is an MSNE iff $\forall i \in N$

- 1 $u_i(s_i, \sigma_{-i}^*)$ is identical $\forall s_i \in \delta(\sigma_i^*)$,
- 2 $u_i(s_i, \sigma_{-i}^*) \geq u_i(s'_i, \sigma_{-i}^*), \forall s_i \in \delta(\sigma_i^*), s'_i \notin \delta(\sigma_i^*)$.

Observations:

- $\max_{\sigma_i \in \Delta S_i} u_i(\sigma_i, \sigma_{-i}) = \max_{s_i \in S_i} u_i(s_i, \sigma_{-i})$
maximizing w.r.t. a distribution $\Leftrightarrow$ whole probability mass at max
- If $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE, then

$$\max_{\sigma_i \in \Delta S_i} u_i(\sigma_i, \sigma_{-i}^*) = \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) = \max_{s_i \in \delta(\sigma_i^*)} u_i(s_i, \sigma_{-i}^*)$$

the maximizer must lie in $\delta(\sigma_i^*)$ – if not, then put all probability mass on that $s'_i \notin \delta(\sigma_i^*)$ that has the maximum value of the utility – $(\sigma_i^*, \sigma_{-i}^*)$ is not a MSNE

Proof of MSNE Characterization Theorem



($\Rightarrow$) Given $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE

$$u_i(\sigma_i^*, \sigma_{-i}^*) = \max_{\sigma_i \in \Delta S_i} u_i(\sigma_i, \sigma_{-i}^*) = \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) = \max_{s_i \in \delta(\sigma_i^*)} u_i(s_i, \sigma_{-i}^*) \quad (1)$$

Proof of MSNE Characterization Theorem



($\Rightarrow$) Given $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE

$$u_i(\sigma_i^*, \sigma_{-i}^*) = \max_{\sigma_i \in \Delta S_i} u_i(\sigma_i, \sigma_{-i}^*) = \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) = \max_{s_i \in \delta(\sigma_i^*)} u_i(s_i, \sigma_{-i}^*) \quad (1)$$

By definition of expected utility

$$u_i(\sigma_i^*, \sigma_{-i}^*) = \sum_{s_i \in S_i} \sigma_i^*(s_i) u_i(s_i, \sigma_{-i}^*) = \sum_{s_i \in \delta(\sigma_i^*)} \sigma_i^*(s_i) u_i(s_i, \sigma_{-i}^*) \quad (2)$$

Proof of MSNE Characterization Theorem



($\Rightarrow$) Given $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE

$$u_i(\sigma_i^*, \sigma_{-i}^*) = \max_{\sigma_i \in \Delta S_i} u_i(\sigma_i, \sigma_{-i}^*) = \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) = \max_{s_i \in \delta(\sigma_i^*)} u_i(s_i, \sigma_{-i}^*) \quad (1)$$

By definition of expected utility

$$u_i(\sigma_i^*, \sigma_{-i}^*) = \sum_{s_i \in S_i} \sigma_i^*(s_i) u_i(s_i, \sigma_{-i}^*) = \sum_{s_i \in \delta(\sigma_i^*)} \sigma_i^*(s_i) u_i(s_i, \sigma_{-i}^*) \quad (2)$$

Equations (1) and (2) are equal, i.e., max is equal to positive weighted average – **can happen only when all values are same: proves condition 1**



For **condition 2**: Suppose for contradiction, there exists $s_i \in \delta(\sigma_i^*)$ and $s'_i \notin \delta(\sigma_i^*)$ s.t.
 $u_i(s_i, \sigma_{-i}^*) < u_i(s'_i, \sigma_{-i}^*)$



For **condition 2**: Suppose for contradiction, there exists $s_i \in \delta(\sigma_i^*)$ and $s'_i \notin \delta(\sigma_i^*)$ s.t.
 $u_i(s_i, \sigma_{-i}^*) < u_i(s'_i, \sigma_{-i}^*)$

We can shift the probability mass $\sigma^*(s_i)$ to s'_i , this new mixed strategy gives a strict higher utility to player i : contradicts MSNE



For **condition 2**: Suppose for contradiction, there exists $s_i \in \delta(\sigma_i^*)$ and $s'_i \notin \delta(\sigma_i^*)$ s.t.
 $u_i(s_i, \sigma_{-i}^*) < u_i(s'_i, \sigma_{-i}^*)$

We can shift the probability mass $\sigma^*(s_i)$ to s'_i , this new mixed strategy gives a strict higher utility to player i : contradicts MSNE

This completes the proof of the necessary direction.



For **condition 2**: Suppose for contradiction, there exists $s_i \in \delta(\sigma_i^*)$ and $s'_i \notin \delta(\sigma_i^*)$ s.t.
 $u_i(s_i, \sigma_{-i}^*) < u_i(s'_i, \sigma_{-i}^*)$

We can shift the probability mass $\sigma^*(s_i)$ to s'_i , this new mixed strategy gives a strict higher utility to player i : contradicts MSNE

This completes the proof of the necessary direction.

($\Leftarrow$) Given the 2 conditions of the theorem, need to show that $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE

Proof (contd.)



For **condition 2**: Suppose for contradiction, there exists $s_i \in \delta(\sigma_i^*)$ and $s'_i \notin \delta(\sigma_i^*)$ s.t.
 $u_i(s_i, \sigma_{-i}^*) < u_i(s'_i, \sigma_{-i}^*)$

We can shift the probability mass $\sigma^*(s_i)$ to s'_i , this new mixed strategy gives a strict higher utility to player i : contradicts MSNE

This completes the proof of the necessary direction.

($\Leftarrow$) Given the 2 conditions of the theorem, need to show that $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE

$$\text{Let } u_i(s_i, \sigma_{-i}^*) = m_i(\sigma_{-i}^*), \forall s_i \in \delta(\sigma_i^*)$$

condition 1

$$\text{Note } m_i(\sigma_{-i}^*) = \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*)$$

condition 2



$$u_i(\sigma_i^*, \sigma_{-i}^*) = \sum_{s_i \in \delta(\sigma_i^*)} \sigma_i^*(s_i) u_i(s_i, \sigma_{-i}^*),$$

by definition of $\delta(\sigma_i^*)$



$$\begin{aligned} u_i(\sigma_i^*, \sigma_{-i}^*) &= \sum_{s_i \in \delta(\sigma_i^*)} \sigma_i^*(s_i) u_i(s_i, \sigma_{-i}^*), \\ &= m_i(\sigma_{-i}^*) \end{aligned}$$

by definition of $\delta(\sigma_i^*)$

previous conclusion



$$\begin{aligned} u_i(\sigma_i^*, \sigma_{-i}^*) &= \sum_{s_i \in \delta(\sigma_i^*)} \sigma_i^*(s_i) u_i(s_i, \sigma_{-i}^*), \\ &= m_i(\sigma_{-i}^*) \\ &= \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) \end{aligned}$$

by definition of $\delta(\sigma_i^*)$

previous conclusion

previous conclusion



$$\begin{aligned} u_i(\sigma_i^*, \sigma_{-i}^*) &= \sum_{s_i \in \delta(\sigma_i^*)} \sigma_i^*(s_i) u_i(s_i, \sigma_{-i}^*), \\ &= m_i(\sigma_{-i}^*) \\ &= \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) \\ &= \max_{\sigma_i \in \Delta S_i} u_i(\sigma_i, \sigma_{-i}^*) \end{aligned}$$

by definition of $\delta(\sigma_i^*)$

previous conclusion

previous conclusion

from the observation



$$\begin{aligned} u_i(\sigma_i^*, \sigma_{-i}^*) &= \sum_{s_i \in \delta(\sigma_i^*)} \sigma_i^*(s_i) u_i(s_i, \sigma_{-i}^*), \\ &= m_i(\sigma_{-i}^*) \\ &= \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) \\ &= \max_{\sigma_i \in \Delta S_i} u_i(\sigma_i, \sigma_{-i}^*) \\ &\geq u_i(\sigma_i, \sigma_{-i}^*), \quad \forall \sigma_i \in \Delta S_i \end{aligned}$$

by definition of $\delta(\sigma_i^*)$

previous conclusion

previous conclusion

from the observation



$$\begin{aligned} u_i(\sigma_i^*, \sigma_{-i}^*) &= \sum_{s_i \in \delta(\sigma_i^*)} \sigma_i^*(s_i) u_i(s_i, \sigma_{-i}^*), \\ &= m_i(\sigma_{-i}^*) \\ &= \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) \\ &= \max_{\sigma_i \in \Delta S_i} u_i(\sigma_i, \sigma_{-i}^*) \\ &\geq u_i(\sigma_i, \sigma_{-i}^*), \quad \forall \sigma_i \in \Delta S_i \end{aligned}$$

by definition of $\delta(\sigma_i^*)$

previous conclusion

previous conclusion

from the observation

This proves the sufficient direction. The result yields an algorithmic way to find MSNE



Chunk 4

Algorithm, Dominance, and Existence of MSNE

($\approx$ 25–30 minutes)

MSNE Characterization Theorem to Algorithm



Consider a NFG $G = \langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$

MSNE Characterization Theorem to Algorithm



Consider a NFG $G = \langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$

The total number of supports of $S_1 \times S_2 \times S_3 \cdots \times S_n$ is

$$K = (2^{|S_1|} - 1) \times (2^{|S_2|} - 1) \times \cdots \times (2^{|S_n|} - 1)$$



MSNE Characterization Theorem to Algorithm

Consider a NFG $G = \langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$

The total number of supports of $S_1 \times S_2 \times S_3 \cdots \times S_n$ is

$$K = (2^{|S_1|} - 1) \times (2^{|S_2|} - 1) \times \cdots \times (2^{|S_n|} - 1)$$

For every support profile $X_1 \times X_2 \times \cdots \times X_n$, where $X_i \subseteq S_i$, solve the following feasibility program

Program

$$\begin{aligned} w_i &= \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} \sigma_j(s_j) \right) \cdot u_i(s_i, s_{-i}), \forall s_i \in X_i, \forall i \in N \\ w_i &\geq \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} \sigma_j(s_j) \right) \cdot u_i(s_i, s_{-i}), \forall s_i \in S_i \setminus X_i, \forall i \in N \\ \sigma_j(s_j) &\geq 0, \forall s_j \in S_j, \forall j \in N, \quad \sum_{s_j \in X_j} \sigma_j(s_j) = 1, \forall j \in N \end{aligned}$$

Remarks on the Algorithm



Program

$$\begin{aligned}w_i &= \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} \sigma_j(s_j) \right) \cdot u_i(s_i, s_{-i}), \forall s_i \in X_i, \forall i \in N \\w_i &\geq \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} \sigma_j(s_j) \right) \cdot u_i(s_i, s_{-i}), \forall s_i \in S_i \setminus X_i, \forall i \in N \\ \sigma_j(s_j) &\geq 0, \forall s_j \in S_j, \forall j \in N, \quad \sum_{s_j \in X_j} \sigma_j(s_j) = 1, \forall j \in N\end{aligned}$$

- This is not a linear program unless $n = 2$

¹Daskalakis, Goldberg, Papadimitriou, "The Complexity of Computing a Nash Equilibrium" [2009]

Remarks on the Algorithm



Program

$$\begin{aligned}w_i &= \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} \sigma_j(s_j) \right) \cdot u_i(s_i, s_{-i}), \forall s_i \in X_i, \forall i \in N \\w_i &\geq \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} \sigma_j(s_j) \right) \cdot u_i(s_i, s_{-i}), \forall s_i \in S_i \setminus X_i, \forall i \in N \\ \sigma_j(s_j) &\geq 0, \forall s_j \in S_j, \forall j \in N, \quad \sum_{s_j \in X_j} \sigma_j(s_j) = 1, \forall j \in N\end{aligned}$$

- This is not a linear program unless $n = 2$
- For general game, there is no poly-time algorithm

¹Daskalakis, Goldberg, Papadimitriou, "The Complexity of Computing a Nash Equilibrium" [2009]



Remarks on the Algorithm

Program

$$\begin{aligned}w_i &= \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} \sigma_j(s_j) \right) \cdot u_i(s_i, s_{-i}), \forall s_i \in X_i, \forall i \in N \\w_i &\geq \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} \sigma_j(s_j) \right) \cdot u_i(s_i, s_{-i}), \forall s_i \in S_i \setminus X_i, \forall i \in N \\ \sigma_j(s_j) &\geq 0, \forall s_j \in S_j, \forall j \in N, \quad \sum_{s_j \in X_j} \sigma_j(s_j) = 1, \forall j \in N\end{aligned}$$

- This is not a linear program unless $n = 2$
- For general game, there is no poly-time algorithm
- Problem of finding an MSNE is PPAD-complete [Polynomial Parity Argument on Directed graphs]¹

¹Daskalakis, Goldberg, Papadimitriou, "The Complexity of Computing a Nash Equilibrium" [2009]

MSNE and Dominance



A gig-platform driver picks a shift type: Standard rides (T), Pooled rides (M), or Premium rides (B). The platform's dispatch algorithm runs in mode L or R. The previous algorithm can be applied to a smaller set of strategies by removing the dominated strategies.

Is there a dominated strategy in this game? Domination can be via mixed strategies too

		Platform	
		L	R
Driver	T	4, 1	2, 5
	M	1, 3	6, 2
	B	2, 2	3, 3



Theorem

If a pure strategy s_i is strictly dominated by a mixed strategy $\sigma_i \in \Delta S_i$, then in every MSNE of the game, s_i is chosen with probability zero.

So, We can remove such strategies without loss of equilibrium

Back to the driver

Neither T nor M strictly dominates B in pure strategies, but a 50-50 mix of T and M beats B against both L and R – B can be dropped before running the algorithm.



Definition (Finite Games)

A game is said to be **finite** when the number of players is finite, and each player has a finite set of strategies.

Existence of MSNE



Definition (Finite Games)

A game is said to be **finite** when the number of players is finite, and each player has a finite set of strategies.

Theorem (Nash 1951)

Every finite game has a (mixed) Nash equilibrium.

Existence of MSNE



Definition (Finite Games)

A game is said to be **finite** when the number of players is finite, and each player has a finite set of strategies.

Theorem (Nash 1951)

Every finite game has a (mixed) Nash equilibrium.

Why this matters

No matter how adversarial the setup – Ghost against Sentinel, a driver against a dispatch algorithm, or two agents locked in an arms race – some equilibrium is guaranteed to exist once randomization is allowed. This is the theoretical bedrock beneath training procedures like GANs and adversarial ML, which search for exactly such an equilibrium.

Proof requires a few tools and a result from real analysis. Proof is given in the lecture notes.

Existence of MSNE: Background



Some background for understanding the proof.

- A set $S \subseteq \mathbb{R}^n$ is **convex** if $\forall x, y \in S$ and $\forall \lambda \in [0, 1]$, $\lambda x + (1 - \lambda)y \in S$.



Existence of MSNE: Background

Some background for understanding the proof.

- A set $S \subseteq \mathbb{R}^n$ is **convex** if $\forall x, y \in S$ and $\forall \lambda \in [0, 1]$, $\lambda x + (1 - \lambda)y \in S$.
- A set $S \subseteq \mathbb{R}^n$ is **closed** if it contains all its **limit points** (points whose **every** neighborhood contains a point in S). Example of a set that is **not closed**: $[0, 1)$ - every ball of radius $\epsilon > 0$ around 1 has a member of $[0, 1)$, but 1 is not in the set $[0, 1)$.



Existence of MSNE: Background

Some background for understanding the proof.

- A set $S \subseteq \mathbb{R}^n$ is **convex** if $\forall x, y \in S$ and $\forall \lambda \in [0, 1]$, $\lambda x + (1 - \lambda)y \in S$.
- A set $S \subseteq \mathbb{R}^n$ is **closed** if it contains all its **limit points** (points whose **every** neighborhood contains a point in S). Example of a set that is **not closed**: $[0, 1)$ - every ball of radius $\epsilon > 0$ around 1 has a member of $[0, 1)$, but 1 is not in the set $[0, 1)$.
- A set $S \subseteq \mathbb{R}^n$ is **bounded** if $\exists x_0 \in \mathbb{R}^n$ and $R \in (0, \infty)$ s.t. $\forall x \in S$, $\|x - x_0\|_2 < R$.



Existence of MSNE: Background

Some background for understanding the proof.

- A set $S \subseteq \mathbb{R}^n$ is **convex** if $\forall x, y \in S$ and $\forall \lambda \in [0, 1]$, $\lambda x + (1 - \lambda)y \in S$.
- A set $S \subseteq \mathbb{R}^n$ is **closed** if it contains all its **limit points** (points whose **every** neighborhood contains a point in S). Example of a set that is **not closed**: $[0, 1)$ - every ball of radius $\epsilon > 0$ around 1 has a member of $[0, 1)$, but 1 is not in the set $[0, 1)$.
- A set $S \subseteq \mathbb{R}^n$ is **bounded** if $\exists x_0 \in \mathbb{R}^n$ and $R \in (0, \infty)$ s.t. $\forall x \in S$, $\|x - x_0\|_2 < R$.
- A set $S \subseteq \mathbb{R}^n$ is **compact** if it is **closed** and **bounded**.



Existence of MSNE: Background

Some background for understanding the proof.

- A set $S \subseteq \mathbb{R}^n$ is **convex** if $\forall x, y \in S$ and $\forall \lambda \in [0, 1]$, $\lambda x + (1 - \lambda)y \in S$.
- A set $S \subseteq \mathbb{R}^n$ is **closed** if it contains all its **limit points** (points whose **every** neighborhood contains a point in S). Example of a set that is **not closed**: $[0, 1)$ - every ball of radius $\epsilon > 0$ around 1 has a member of $[0, 1)$, but 1 is not in the set $[0, 1)$.
- A set $S \subseteq \mathbb{R}^n$ is **bounded** if $\exists x_0 \in \mathbb{R}^n$ and $R \in (0, \infty)$ s.t. $\forall x \in S$, $\|x - x_0\|_2 < R$.
- A set $S \subseteq \mathbb{R}^n$ is **compact** if it is **closed** and **bounded**.



Existence of MSNE: Background

Some background for understanding the proof.

- A set $S \subseteq \mathbb{R}^n$ is **convex** if $\forall x, y \in S$ and $\forall \lambda \in [0, 1]$, $\lambda x + (1 - \lambda)y \in S$.
- A set $S \subseteq \mathbb{R}^n$ is **closed** if it contains all its **limit points** (points whose **every** neighborhood contains a point in S). Example of a set that is **not closed**: $[0, 1)$ - every ball of radius $\epsilon > 0$ around 1 has a member of $[0, 1)$, but 1 is not in the set $[0, 1)$.
- A set $S \subseteq \mathbb{R}^n$ is **bounded** if $\exists x_0 \in \mathbb{R}^n$ and $R \in (0, \infty)$ s.t. $\forall x \in S, \|x - x_0\|_2 < R$.
- A set $S \subseteq \mathbb{R}^n$ is **compact** if it is **closed** and **bounded**.

A result from real analysis (proof omitted):

Brouwer's fixed point theorem

If $S \subseteq \mathbb{R}^n$ is **convex** and **compact** and $T : S \rightarrow S$, is **continuous** then T has a fixed point, i.e., $\exists x^* \in S$ s.t. $T(x^*) = x^*$.



In-Class Quiz 2

Open the Google Form now.

Questions refer only to what was just discussed –
you cannot answer them without having been in class.

Lecture 6 and Week 3: What We Built



- MSNE: no unilateral profitable deviation, even to a pure strategy

Lecture 6 and Week 3: What We Built



- MSNE: no unilateral profitable deviation, even to a pure strategy
- Support-based characterization: equal utility on the support, no better pure deviation off it

Lecture 6 and Week 3: What We Built



- MSNE: no unilateral profitable deviation, even to a pure strategy
- Support-based characterization: equal utility on the support, no better pure deviation off it
- An algorithm: enumerate support profiles, solve a feasibility program on each – but this is PPAD-complete in general

Lecture 6 and Week 3: What We Built



- MSNE: no unilateral profitable deviation, even to a pure strategy
- Support-based characterization: equal utility on the support, no better pure deviation off it
- An algorithm: enumerate support profiles, solve a feasibility program on each – but this is PPAD-complete in general
- Strategies strictly dominated even by a mixed strategy get probability zero in every MSNE, and can be dropped first

Lecture 6 and Week 3: What We Built



- MSNE: no unilateral profitable deviation, even to a pure strategy
- Support-based characterization: equal utility on the support, no better pure deviation off it
- An algorithm: enumerate support profiles, solve a feasibility program on each – but this is PPAD-complete in general
- Strategies strictly dominated even by a mixed strategy get probability zero in every MSNE, and can be dropped first
- Nash's theorem (1951): every finite game has an MSNE, via Brouwer's fixed point theorem

Lecture 6 and Week 3: What We Built



- MSNE: no unilateral profitable deviation, even to a pure strategy
- Support-based characterization: equal utility on the support, no better pure deviation off it
- An algorithm: enumerate support profiles, solve a feasibility program on each – but this is PPAD-complete in general
- Strategies strictly dominated even by a mixed strategy get probability zero in every MSNE, and can be dropped first
- Nash's theorem (1951): every finite game has an MSNE, via Brouwer's fixed point theorem
- **Next week:** onward to mechanism design



End of Lecture 6

See you next class.



भारतीय प्रौद्योगिकी संस्थान मुंबई

Indian Institute of Technology Bombay