



भारतीय प्रौद्योगिकी संस्थान मुंबई
Indian Institute of Technology Bombay

CS 6001: Game Theory and Algorithmic Mechanism Design

Week 6 – Equilibrium in IIEFGs, P2P File Sharing, and Bayesian Games

Swaprava Nath

September, 2026

Week 6 Roadmap



Lecture	Chunks (20–25 min each)	Quizzes
Lecture 11	1. From Subgame Perfection to Beliefs 2. Bayesian Beliefs in Action 3. Sequential Rationality and PBE 4. In Practice: P2P File Sharing	Quiz 1 after Chunk 2 Quiz 2 after Chunk 4
Lecture 12	1. Incomplete Information & Bayesian Games 2. Strategies, Ex-ante and Ex-interim 3. Bayesian Equilibrium = Nash 4. Auctions as Bayesian Games	Quiz 1 after Chunk 2 Quiz 2 after Chunk 4



► Lecture 11: Equilibrium in IIEFGs and P2P File Sharing

► Lecture 12: Bayesian Games and Bayesian Equilibrium



Chunk 1

From Subgame Perfection to Beliefs

($\approx$ 20–25 minutes)



Last week: IIEFGs, behavioral strategies, perfect recall, Kuhn's theorem.

Where We Stand



Last week: IIEFGs, behavioral strategies, perfect recall, Kuhn's theorem.

We can now *describe* play under imperfect information.



Last week: IIEFGs, behavioral strategies, perfect recall, Kuhn's theorem.

We can now *describe* play under imperfect information.

Today's question

What replaces subgame perfection when a player does not know which node of the tree she is standing on?



- Can extend the subgame perfection of PIEFG, but since the nodes/histories are uncertain, we need to extend to mixed strategies



- Can extend the subgame perfection of PIEFG, but since the nodes/histories are uncertain, we need to extend to mixed strategies
- Because of the information sets, best response cannot be defined without the belief of each player



- Can extend the subgame perfection of PIEFG, but since the nodes/histories are uncertain, we need to extend to mixed strategies
- Because of the information sets, best response cannot be defined without the belief of each player

Belief

It is the conditional probability distribution over the histories in an information set - conditioned on reaching the information set.

Example 1: Release or Review



Lab (Player 1) has a model ready. It can **Exit the process** (E) and release under the existing licence, or send it for external review (X).


Example 1: Release or Review



Lab (Player 1) has a model ready. It can **Exit the process** (E) and release under the existing licence, or send it for external review (X).


If it goes to review, the Lab ships either the **Tuned** build (T) or the **Base** build (B). The **Auditor** (Player 2) receives a stripped artefact and cannot tell the two apart, then runs a **Light** (L) or **Rigorous** (R) audit.

Example 1: IIEFG and Subgame Perfection






Example 1: IIEFG and Subgame Perfection



- Two subgames: rooted at x_1 and x_2



Example 1: IIEFG and Subgame Perfection



- Two subgames: rooted at x_1 and x_2
- Backward induction: solve at x_2 : R dominates L for player 2, then T dominates B for player 1

Why Only Two Subgames?



Question

Why is the tree hanging below x_3 not a subgame?



Why Only Two Subgames?

Question

Why is the tree hanging below x_3 not a subgame?

Answer

A subgame must not cut an information set. x_3 and x_4 sit in the same information set I_2^1 of the Auditor. A “subgame” at x_3 alone would hand the Auditor information she does not have.



Why Only Two Subgames?

Question

Why is the tree hanging below x_3 not a subgame?

Answer

A subgame must not cut an information set. x_3 and x_4 sit in the same information set I_2^1 of the Auditor. A “subgame” at x_3 alone would hand the Auditor information she does not have.

- Here we got lucky: I_2^1 lies *entirely* inside the subgame at x_2 , so backward induction still worked.



Why Only Two Subgames?

Question

Why is the tree hanging below x_3 not a subgame?

Answer

A subgame must not cut an information set. x_3 and x_4 sit in the same information set I_2^1 of the Auditor. A “subgame” at x_3 alone would hand the Auditor information she does not have.

- Here we got lucky: I_2^1 lies *entirely* inside the subgame at x_2 , so backward induction still worked.
- In general it will not. If the game does not have any intermediate history with a singleton information set, subgame perfection has nothing to say.



Why Only Two Subgames?

Question

Why is the tree hanging below x_3 not a subgame?

Answer

A subgame must not cut an information set. x_3 and x_4 sit in the same information set I_2^1 of the Auditor. A “subgame” at x_3 alone would hand the Auditor information she does not have.

- Here we got lucky: I_2^1 lies *entirely* inside the subgame at x_2 , so backward induction still worked.
- In general it will not. If the game does not have any intermediate history with a singleton information set, subgame perfection has nothing to say.
- What we need instead is a rule for behaving *inside* an information set.



Chunk 2

Bayesian Beliefs in Action

($\approx$ 20–25 minutes)

Example 2: A Three-Round Agent Duel



An **Agent** (Player 1) plays a repeated interaction against an adaptive **Opponent** (Player 2). The Agent opens with one of three probes, the Opponent responds without seeing which probe was used, and the Agent then acts again from a position it can only partly identify.



Example 2: A Three-Round Agent Duel

An **Agent** (Player 1) plays a repeated interaction against an adaptive **Opponent** (Player 2). The Agent opens with one of three probes, the Opponent responds without seeing which probe was used, and the Agent then acts again from a position it can only partly identify.

Formally this is **Example 7.38 of MSZ**: an IIEFG with perfect recall, so mixed and behavioral strategies are equivalent.




Example 2: A Three-Round Agent Duel

An **Agent** (Player 1) plays a repeated interaction against an adaptive **Opponent** (Player 2). The Agent opens with one of three probes, the Opponent responds without seeing which probe was used, and the Agent then acts again from a position it can only partly identify.

Formally this is **Example 7.38 of MSZ**: an IIEFG with perfect recall, so mixed and behavioral strategies are equivalent.


Player 1 has three information sets I_1^1, I_1^2, I_1^3 ; Player 2 has one, I_2^1 .

Example 2: IIEFG and Sequential Rationality






Example 2: IIEFG and Sequential Rationality



Consider the behavioral strategy profile: σ_1 , at $I_1^1(L\{5/12\}, M\{4/12\}, R\{3/12\})$




Example 2: IIEFG and Sequential Rationality



Consider the behavioral strategy profile: σ_2 , at $I_2^1(l\{1\}, m\{0\}, r\{0\})$ choose l




Example 2: IIEFG and Sequential Rationality



Consider the behavioral strategy profile: σ_1 , at $I_1^2(L_1\{0\}, R_1\{1\})$ choose R_1




Example 2: IIEFG and Sequential Rationality



Consider the behavioral strategy profile: σ_1 , at $I_1^3(L_2\{1\}, R_2\{0\})$ choose L_2



Is This an Equilibrium?




Question

Is this an equilibrium?
which implies

- Are the Bayesian beliefs consistent with P_σ - that visits vertex x with probability $P_\sigma(x)$?
- The actions and beliefs are consistent for every player, i.e., maximizes their expected utility?

Is This an Equilibrium?




Sequential Rationality

Choose an action maximizing expected utility at each information set.

The strategy vector σ induces the following probabilities to the vertices.

$$P_{\sigma}(x_2) = 5/12, P_{\sigma}(x_3) = 4/12, P_{\sigma}(x_4) = 0, P_{\sigma}(x_5) = 0, P_{\sigma}(x_6) = 4/12, P_{\sigma}(x_7) = 0$$


Is This an Equilibrium?



- Player 2 at I_2^1 believes that x_3 is reached w.p. $P_\sigma(x_3|I_2^1) = P_\sigma(x_3)/(P_\sigma(x_2) + P_\sigma(x_3)) = 4/9$
- Similarly $P_\sigma(x_2|I_2^1) = 5/9$



Is This an Equilibrium?



Question


Is the action of player 2 sequentially rational w.r.t. her belief?

Answer

By picking l , expected utility $= \frac{5}{9} \times 1 + \frac{4}{9} \times 2 = \frac{13}{9}$, larger than any other choice of action.



Is This an Equilibrium?



Question

Given all information, what is the sequentially rational strategy for player 1 at I_1^1

Answer

L, M, R all give the same expected utility for player 1 (utility = 2).

A Simple Analogy: The Doctor Behind the Curtain



The picture

A doctor is told: “your next patient has either flu or COVID, but you will only be told the symptoms, not which one it is.” She cannot see the patient’s file (the information set), only the room she has been sent into.

A Simple Analogy: The Doctor Behind the Curtain



The picture

A doctor is told: “your next patient has either flu or COVID, but you will only be told the symptoms, not which one it is.” She cannot see the patient’s file (the information set), only the room she has been sent into.

- Before meeting the patient, she already knows how common each illness is this season – that is her **belief**.



A Simple Analogy: The Doctor Behind the Curtain

The picture

A doctor is told: “your next patient has either flu or COVID, but you will only be told the symptoms, not which one it is.” She cannot see the patient’s file (the information set), only the room she has been sent into.

- Before meeting the patient, she already knows how common each illness is this season – that is her **belief**.
- With no other information, her belief is just last season’s flu-vs-COVID ratio: exactly $\mu_2^1(x_2) = 5/9$ and $\mu_2^1(x_3) = 4/9$ from the previous slide, how often the Agent’s strategy sends the Opponent to x_2 versus x_3 .



A Simple Analogy: The Doctor Behind the Curtain

The picture

A doctor is told: “your next patient has either flu or COVID, but you will only be told the symptoms, not which one it is.” She cannot see the patient’s file (the information set), only the room she has been sent into.

- Before meeting the patient, she already knows how common each illness is this season – that is her **belief**.
- With no other information, her belief is just last season’s flu-vs-COVID ratio: exactly $\mu_2^1(x_2) = 5/9$ and $\mu_2^1(x_3) = 4/9$ from the previous slide, how often the Agent’s strategy sends the Opponent to x_2 versus x_3 .



A Simple Analogy: The Doctor Behind the Curtain

The picture

A doctor is told: “your next patient has either flu or COVID, but you will only be told the symptoms, not which one it is.” She cannot see the patient’s file (the information set), only the room she has been sent into.

- Before meeting the patient, she already knows how common each illness is this season – that is her **belief**.
- With no other information, her belief is just last season’s flu-vs-COVID ratio: exactly $\mu_2^1(x_2) = 5/9$ and $\mu_2^1(x_3) = 4/9$ from the previous slide, how often the Agent’s strategy sends the Opponent to x_2 versus x_3 .

The one rule

Whenever a room can genuinely be visited, the belief about what’s in it must match how often patients like this actually end up there: $\mu_2^1(x) = P_\sigma(x) / \sum_{y \in I_2^1} P_\sigma(y)$. That is all Bayes’ rule is saying – a ratio of how often each room is actually reached.

What If the Room Is Never Visited?



Back to the analogy

Suppose a particular room is one the clinic's booking system *never* actually sends anyone to – a booking error that in practice never happens. What should the doctor believe if she ever finds herself there?

What If the Room Is Never Visited?



Back to the analogy

Suppose a particular room is one the clinic's booking system *never* actually sends anyone to – a booking error that in practice never happens. What should the doctor believe if she ever finds herself there?

- There is no historical ratio to fall back on: nobody has ever been in that room, so “how often does flu vs. COVID end up here” has no answer.

What If the Room Is Never Visited?



Back to the analogy

Suppose a particular room is one the clinic's booking system *never* actually sends anyone to – a booking error that in practice never happens. What should the doctor believe if she ever finds herself there?

- There is no historical ratio to fall back on: nobody has ever been in that room, so “how often does flu vs. COVID end up here” has no answer.
- In our tree, this is exactly $I_1^2 = \{x_4, x_5\}$: the Agent's strategy σ gives $P_\sigma(x_4) = P_\sigma(x_5) = 0$, so there is nothing to divide by.

What If the Room Is Never Visited?



Back to the analogy

Suppose a particular room is one the clinic's booking system *never* actually sends anyone to – a booking error that in practice never happens. What should the doctor believe if she ever finds herself there?

- There is no historical ratio to fall back on: nobody has ever been in that room, so “how often does flu vs. COVID end up here” has no answer.
- In our tree, this is exactly $I_1^2 = \{x_4, x_5\}$: the Agent's strategy σ gives $P_\sigma(x_4) = P_\sigma(x_5) = 0$, so there is nothing to divide by.
- Bayes' rule is silent here. *Any* belief the doctor picks is equally consistent with the data, because there is no data.



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

Sequential Rationality and Perfect Bayesian Equilibrium

($\approx$ 20–25 minutes)



Belief

Let the information sets of player i be $I_i = \{I_i^1, I_i^2, I_i^3, \dots, I_i^{k(i)}\}$.

The belief of player i is a collection of mappings $\mu_i^j : I_i^j \rightarrow [0, 1]$ s.t., $\sum_{x \in I_i^j} \mu_i^j(x) = 1, \forall j = 1, \dots, k(i)$.



Formal definitions

Belief

Let the information sets of player i be $I_i = \{I_i^1, I_i^2, I_i^3, \dots, I_i^{k(i)}\}$.

The belief of player i is a collection of mappings $\mu_i^j : I_i^j \rightarrow [0, 1]$ s.t., $\sum_{x \in I_i^j} \mu_i^j(x) = 1, \forall j = 1, \dots, k(i)$.

Bayesian Belief

A **belief** $\mu_i = \{\mu_i^1, \mu_i^2, \dots, \mu_i^{k(i)}\}$ of player i is **Bayesian** w.r.t. to the behavioral strategy σ , if it is derived from σ using Bayes rule, i.e.,

$$\mu_i^j(x) = P_\sigma(x) / \sum_{y \in I_i^j} P_\sigma(y), \quad \forall x \in I_i^j, \forall j = 1, 2, 3, \dots, k(i)$$



Formal definitions

Belief

Let the information sets of player i be $I_i = \{I_i^1, I_i^2, I_i^3, \dots, I_i^{k(i)}\}$.

The belief of player i is a collection of mappings $\mu_i^j : I_i^j \rightarrow [0, 1]$ s.t., $\sum_{x \in I_i^j} \mu_i^j(x) = 1, \forall j = 1, \dots, k(i)$.

Bayesian Belief

A **belief** $\mu_i = \{\mu_i^1, \mu_i^2, \dots, \mu_i^{k(i)}\}$ of player i is **Bayesian** w.r.t. to the behavioral strategy σ , if it is derived from σ using Bayes rule, i.e.,

$$\mu_i^j(x) = P_\sigma(x) / \sum_{y \in I_i^j} P_\sigma(y), \quad \forall x \in I_i^j, \forall j = 1, 2, 3, \dots, k(i)$$

In the doctor analogy: I_i^j is a room, x a possible patient in it, and $\mu_i^j(x)$ the doctor's belief in that patient – read off from how often patients like this actually appear there.



Sequential Rationality

A strategy σ_i of player i at an information set I_i^j is sequentially rational given σ_{-i} and belief μ_i^j if

$$\sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma_i, \sigma_{-i} | x) \geq \sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma'_i, \sigma_{-i} | x), \quad \forall \sigma'_i.$$



Sequential Rationality

A strategy σ_i of player i at an information set I_i^j is sequentially rational given σ_{-i} and belief μ_i^j if

$$\sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma_i, \sigma_{-i} | x) \geq \sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma'_i, \sigma_{-i} | x), \quad \forall \sigma'_i.$$

In the doctor analogy

Given her belief about which patient she is likely facing, the doctor prescribes the treatment with the best *expected* outcome. She is not asked to guess correctly every time – only to do the best she can given what she believes.



Sequential Rationality

A strategy σ_i of player i at an information set I_i^j is sequentially rational given σ_{-i} and belief μ_i^j if

$$\sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma_i, \sigma_{-i} | x) \geq \sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma'_i, \sigma_{-i} | x), \quad \forall \sigma'_i.$$

In the doctor analogy

Given her belief about which patient she is likely facing, the doctor prescribes the treatment with the best *expected* outcome. She is not asked to guess correctly every time – only to do the best she can given what she believes.

- The tuple (σ, μ) , called an **assessment**, is sequentially rational if it holds for every player at every information set.



Sequential Rationality

A strategy σ_i of player i at an information set I_i^j is sequentially rational given σ_{-i} and belief μ_i^j if

$$\sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma_i, \sigma_{-i} | x) \geq \sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma'_i, \sigma_{-i} | x), \quad \forall \sigma'_i.$$

In the doctor analogy

Given her belief about which patient she is likely facing, the doctor prescribes the treatment with the best *expected* outcome. She is not asked to guess correctly every time – only to do the best she can given what she believes.

- The tuple (σ, μ) , called an **assessment**, is sequentially rational if it holds for every player at every information set.
- A refinement of Nash Equilibrium, coinciding with SPNE on PIEFGs.

Sequential Rationality Is a Threshold Condition



Back to $I_1^3 = \{x_6, x_7\}$. Write p for the Agent's belief that the node is x_7 .

Sequential Rationality Is a Threshold Condition



Back to $I_1^3 = \{x_6, x_7\}$. Write p for the Agent's belief that the node is x_7 .

$$u_1(L_2) = (1 - p) \cdot 2 + p \cdot 0 = 2 - 2p,$$

$$u_1(R_2) = (1 - p) \cdot 0 + p \cdot 5 = 5p.$$

Sequential Rationality Is a Threshold Condition



Back to $I_1^3 = \{x_6, x_7\}$. Write p for the Agent's belief that the node is x_7 .

$$u_1(L_2) = (1 - p) \cdot 2 + p \cdot 0 = 2 - 2p,$$

$$u_1(R_2) = (1 - p) \cdot 0 + p \cdot 5 = 5p.$$

So R_2 beats L_2 exactly when $5p > 2 - 2p$, i.e. when $p > 2/7$.



Sequential Rationality Is a Threshold Condition

Back to $I_1^3 = \{x_6, x_7\}$. Write p for the Agent's belief that the node is x_7 .

$$u_1(L_2) = (1 - p) \cdot 2 + p \cdot 0 = 2 - 2p,$$

$$u_1(R_2) = (1 - p) \cdot 0 + p \cdot 5 = 5p.$$

So R_2 beats L_2 exactly when $5p > 2 - 2p$, i.e. when $p > 2/7$.

Under our σ

$P_\sigma(x_6) = 4/12$ and $P_\sigma(x_7) = 0$, so $p = 0$ and L_2 is right. If the belief had been more than $2/7$ in favour of x_7 , player 1 should have chosen R_2 .



Sequential Rationality Is a Threshold Condition

Back to $I_1^3 = \{x_6, x_7\}$. Write p for the Agent's belief that the node is x_7 .

$$u_1(L_2) = (1 - p) \cdot 2 + p \cdot 0 = 2 - 2p,$$

$$u_1(R_2) = (1 - p) \cdot 0 + p \cdot 5 = 5p.$$

So R_2 beats L_2 exactly when $5p > 2 - 2p$, i.e. when $p > 2/7$.

Under our σ

$P_\sigma(x_6) = 4/12$ and $P_\sigma(x_7) = 0$, so $p = 0$ and L_2 is right. If the belief had been more than $2/7$ in favour of x_7 , player 1 should have chosen R_2 .

The action is not right or wrong on its own. It is right or wrong relative to a belief.



Theorem

In a PIEFG, a behavioral strategy profile σ is an SPNE iff the tuple $(\sigma, \hat{\mu})$ is sequentially rational.



Theorem

In a PIEFG, a behavioral strategy profile σ is an SPNE iff the tuple $(\sigma, \hat{\mu})$ is sequentially rational.

In a PIEFG, every information set is a singleton, $\hat{\mu}$ is the degenerate distribution at that singleton.
The doctor always sees the file: no belief is needed, so this collapses to ordinary backward induction.



Theorem

In a PIEFG, a behavioral strategy profile σ is an SPNE iff the tuple $(\sigma, \hat{\mu})$ is sequentially rational.

In a PIEFG, every information set is a singleton, $\hat{\mu}$ is the degenerate distribution at that singleton.
The doctor always sees the file: no belief is needed, so this collapses to ordinary backward induction.

Equilibrium with Sequential Rationality

Perfect Bayesian Equilibrium: An assessment (σ, μ) is PBE if $\forall i \in N$

- μ_i is Bayesian w.r.t. σ
- σ_i is sequentially rational given σ_{-i} and μ_i



Formal definitions

Theorem

In a PIEFG, a behavioral strategy profile σ is an SPNE iff the tuple $(\sigma, \hat{\mu})$ is sequentially rational.

In a PIEFG, every information set is a singleton, $\hat{\mu}$ is the degenerate distribution at that singleton.
The doctor always sees the file: no belief is needed, so this collapses to ordinary backward induction.

Equilibrium with Sequential Rationality

Perfect Bayesian Equilibrium: An assessment (σ, μ) is PBE if $\forall i \in N$

- μ_i is Bayesian w.r.t. σ
 - σ_i is sequentially rational given σ_{-i} and μ_i
-
- Often represented only with σ , since μ is obtained from σ



Theorem

In a PIEFG, a behavioral strategy profile σ is an SPNE iff the tuple $(\sigma, \hat{\mu})$ is sequentially rational.

In a PIEFG, every information set is a singleton, $\hat{\mu}$ is the degenerate distribution at that singleton.

The doctor always sees the file: no belief is needed, so this collapses to ordinary backward induction.

Equilibrium with Sequential Rationality

Perfect Bayesian Equilibrium: An assessment (σ, μ) is PBE if $\forall i \in N$

- μ_i is Bayesian w.r.t. σ
 - σ_i is sequentially rational given σ_{-i} and μ_i
-
- Often represented only with σ , since μ is obtained from σ
 - Self-enforcing (like the SPNE) in a Bayesian way.



Chunk 4


Game Theory in Practice: P2P File Sharing

($\approx$ 20–25 minutes)

Peer to Peer¹




download rate



traditional

download rate



P2P

¹Slides of this section are adapted from CS186, Harvard

Desired Properties and Terminology



- Scalability

Terminology:

Desired Properties and Terminology



- Scalability
- Failure resilience

Terminology:

Desired Properties and Terminology



- Scalability
- Failure resilience

Terminology:

- **Protocol:** messages that can be sent, actions that can be taken over the network

Desired Properties and Terminology



- Scalability
- Failure resilience

Terminology:

- **Protocol:** messages that can be sent, actions that can be taken over the network
- **Client:** a particular process for sending messages, taking actions

Desired Properties and Terminology



- Scalability
- Failure resilience

Terminology:

- **Protocol:** messages that can be sent, actions that can be taken over the network
- **Client:** a particular process for sending messages, taking actions
- **Reference client:** particular implementation

Desired Properties and Terminology



- Scalability
- Failure resilience

Terminology:

- **Protocol:** messages that can be sent, actions that can be taken over the network
- **Client:** a particular process for sending messages, taking actions
- **Reference client:** particular implementation
- **Peer**



Napster (1999 - 2001)

- Centralized database
- Users download music from each other



Napster (1999 - 2001)

- Centralized database
- Users download music from each other

Gnutella (2000 -)

- Get list of IP addresses of peers from set of known peers (no server)
- To get a file: Query message broadcast by peer A to known peers
- Query response: sent by B if B has the desired file (routed back to requestor)
- A can then download directly from B

The File Sharing Game



		Player 2	
		Share	Free-ride
Player 1	Share	2, 2	-1, 3
	Free-ride	3, -1	0, 0

(Gnutella) File Sharing Game

The File Sharing Game (Contd.)




Image courtesy: Adar and Huberman (2000)

Incentives for Client Developers



- Client developers can ensure file sharing
- But competition among the developers

Incentives for Client Developers



- Client developers can ensure file sharing
- But competition among the developers
- 85% peers free-riding by 2005; Gnutella less than 1% of ww P2P traffic by 2013
- Few other P2P systems met the same fate



BitTorrent (2001 -)

- Approx 85% of P2P traffic in US
- File sharing
- Also used for S/W distribution (e.g., Linux)



BitTorrent (2001 -)

- Approx 85% of P2P traffic in US
- File sharing
- Also used for S/W distribution (e.g., Linux)

Key innovations

- Break file into pieces: A repeated game!
- “If you let me download, I’ll reciprocate.”

BitTorrent Schematic




Figure 5.4.: Starting a download process in the BitTorrent protocol: 1) A user goes to a searchable directory to find a link to a .torrent file corresponding to the desired content; 2) the .torrent file contains metadata about the content, in particular the URL of a tracker; 3) the tracker provides a list of peers participating in the swarm for the content (i.e., their IP address and port); 4) the user's BitTorrent client can now contact all these peers and download content.

BitTorrent Optimistic Unchoking Algorithm



Tracker is a centralized entity that controls the traffic, tracks the connection between peers and their speed of upload, download etc.

BitTorrent Optimistic Unchoking Algorithm



Tracker is a centralized entity that controls the traffic, tracks the connection between peers and their speed of upload, download etc.

Reference Client Protocol:

- Set a threshold r of uploading speed (typically the third maximum speed in the recent past)

BitTorrent Optimistic Unchoking Algorithm



Tracker is a centralized entity that controls the traffic, tracks the connection between peers and their speed of upload, download etc.

Reference Client Protocol:

- Set a threshold r of uploading speed (typically the third maximum speed in the recent past)
- If a peer j uploaded to i at a rate $\geq r$, unchoke j in the next period

BitTorrent Optimistic Unchoking Algorithm



Tracker is a centralized entity that controls the traffic, tracks the connection between peers and their speed of upload, download etc.

Reference Client Protocol:

- Set a threshold r of uploading speed (typically the third maximum speed in the recent past)
- If a peer j uploaded to i at a rate $\geq r$, unchoke j in the next period
- If a peer j uploaded to i at a rate $< r$, choke j in the next period

BitTorrent Optimistic Unchoking Algorithm



Tracker is a centralized entity that controls the traffic, tracks the connection between peers and their speed of upload, download etc.

Reference Client Protocol:

- Set a threshold r of uploading speed (typically the third maximum speed in the recent past)
- If a peer j uploaded to i at a rate $\geq r$, unchoke j in the next period
- If a peer j uploaded to i at a rate $< r$, choke j in the next period
- Every three time periods, optimistically unchoke a random peer from the neighborhood who is currently choked, and leave that peer unchoked for three time periods.

BitTorrent Optimistic Unchoking Algorithm



Tracker is a centralized entity that controls the traffic, tracks the connection between peers and their speed of upload, download etc.

Reference Client Protocol:

- Set a threshold r of uploading speed (typically the third maximum speed in the recent past)
- If a peer j uploaded to i at a rate $\geq r$, unchoke j in the next period
- If a peer j uploaded to i at a rate $< r$, choke j in the next period
- Every three time periods, optimistically unchoke a random peer from the neighborhood who is currently choked, and leave that peer unchoked for three time periods.

Forcing a repeated game by fragmenting the files

BitTorrent Optimistic Unchoking Algorithm



Tracker is a centralized entity that controls the traffic, tracks the connection between peers and their speed of upload, download etc.

Reference Client Protocol:

- Set a threshold r of uploading speed (typically the third maximum speed in the recent past)
- If a peer j uploaded to i at a rate $\geq r$, unchoke j in the next period
- If a peer j uploaded to i at a rate $< r$, choke j in the next period
- Every three time periods, optimistically unchoke a random peer from the neighborhood who is currently choked, and leave that peer unchoked for three time periods.

Forcing a repeated game by fragmenting the files

The leecher-seeder game is a repeated Prisoners' Dilemma

BitTorrent Optimistic Unchoking Algorithm



Tracker is a centralized entity that controls the traffic, tracks the connection between peers and their speed of upload, download etc.

Reference Client Protocol:

- Set a threshold r of uploading speed (typically the third maximum speed in the recent past)
- If a peer j uploaded to i at a rate $\geq r$, unchoke j in the next period
- If a peer j uploaded to i at a rate $< r$, choke j in the next period
- Every three time periods, optimistically unchoke a random peer from the neighborhood who is currently choked, and leave that peer unchoked for three time periods.

Forcing a repeated game by fragmenting the files

The leecher-seeder game is a repeated Prisoners' Dilemma

Strategy of the seeder is tit-for-tat



Illustration



- How often to contact tracker?
- Which pieces to reveal?
- How many upload slots, which peers to unchoke, at what speed?
- What data to allow others to download?
- Possible goals: min upload, max download speed, some balance

Attacks on BitTorrent



- BitThief
- Strategic piece revealer
- BitTyrant



- Goal: download files without uploading
- Keep asking for peers from tracker, grow neighborhood quickly
- Exploit the optimistic unchoking part
- Never upload!



- Goal: download files without uploading
- Keep asking for peers from tracker, grow neighborhood quickly
- Exploit the optimistic unchoking part
- Never upload!
- **Fix: modify the tracker (block same IP address within 30 minutes).**


Ref: Locher et al., “Free Riding in BitTorrent is Cheap”, HotNets 2006



- Reference client: tell neighbors about new pieces, use “rarest-first” to request
- Manipulator strategy: reveal most common piece that reciprocating peer does not have!
- Try to protect a monopoly, keep others interested

Ref: Levin et al., “BitTorrent is an Auction: Analyzing and Improving BitTorrent’s Incentives”, SIGCOMM 2008

Strategic Piece Revealer





- P2P demonstrates importance of game-theory in computer systems



- P2P demonstrates importance of game-theory in computer systems
- Early systems were easily manipulated



- P2P demonstrates importance of game-theory in computer systems
- Early systems were easily manipulated
- BitTorrent's innovation was to break files into pieces, enabling TitForTat.



- P2P demonstrates importance of game-theory in computer systems
- Early systems were easily manipulated
- BitTorrent's innovation was to break files into pieces, enabling TitForTat.
- Still some vulnerabilities, but generally very successful example of incentive-based protocol design.



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.



End of Lecture 11

See you next class.



► Lecture 11: Equilibrium in IIEFGs and P2P File Sharing

► Lecture 12: Bayesian Games and Bayesian Equilibrium



Chunk 1

Incomplete Information and Bayesian Games

($\approx$ 20–25 minutes)



- **Non-cooperative games**



- **Non-cooperative games**
 - *Complete information* – players deterministically know which game they are playing



- **Non-cooperative games**

- *Complete information* – players deterministically know which game they are playing
 - Normal form games: simultaneous move, single stage. SDSE, WDSE, PSNE, MSNE, Correlated



- **Non-cooperative games**

- *Complete information* – players deterministically know which game they are playing
 - Normal form games: simultaneous move, single stage. SDSE, WDSE, PSNE, MSNE, Correlated
 - Extensive form games: multi-stage. SPNE (PIEFG), mixed and behavioral strategies (IIEFG), PBE



- **Non-cooperative games**

- *Complete information* – players deterministically know which game they are playing
 - Normal form games: simultaneous move, single stage. SDSE, WDSE, PSNE, MSNE, Correlated
 - Extensive form games: multi-stage. SPNE (PIEFG), mixed and behavioral strategies (IIEFG), PBE
- *Incomplete information* – players do not deterministically know which game they are playing



- **Non-cooperative games**
 - *Complete information* – players deterministically know which game they are playing
 - Normal form games: simultaneous move, single stage. SDSE, WDSE, PSNE, MSNE, Correlated
 - Extensive form games: multi-stage. SPNE (PIEFG), mixed and behavioral strategies (IIEFG), PBE
 - *Incomplete information* – players do not deterministically know which game they are playing
- **Cooperative games** – players form coalitions and utilities are defined over coalitions



- **Non-cooperative games**
 - *Complete information* – players deterministically know which game they are playing
 - Normal form games: simultaneous move, single stage. SDSE, WDSE, PSNE, MSNE, Correlated
 - Extensive form games: multi-stage. SPNE (PIEFG), mixed and behavioral strategies (IIEFG), PBE
 - *Incomplete information* – players do not deterministically know which game they are playing
- **Cooperative games** – players form coalitions and utilities are defined over coalitions
- **Other types** – repeated, stochastic, etc.

Games with Incomplete Information



Games with Complete Information

- Players deterministically know the game they are playing
- There can be some chance moves but probabilities are known

Games with Incomplete Information



Games with Complete Information

- Players deterministically know the game they are playing
- There can be some chance moves but probabilities are known

Games with Incomplete Information

- Players do not deterministically know which game they are playing
- They receive private signals / **types**

Games with Incomplete Information



Games with Complete Information

- Players deterministically know the game they are playing
- There can be some chance moves but probabilities are known

Games with Incomplete Information

- Players do not deterministically know which game they are playing
 - They receive private signals / **types**
-
- To discuss: a special subclass called games with incomplete information with **common priors** (Harsanyi 1967)

Games with Incomplete Information



Games with Complete Information

- Players deterministically know the game they are playing
- There can be some chance moves but probabilities are known

Games with Incomplete Information

- Players do not deterministically know which game they are playing
 - They receive private signals / **types**
-
- To discuss: a special subclass called games with incomplete information with **common priors** (Harsanyi 1967)
 - Also called **Bayesian games**

Bayesian Games: Example



Football game (two competing teams).

Bayesian Games: Example



Football game (two competing teams).

- Each can choose a gameplan: aim to win (Aggressive) or aim to draw (Passive)

Bayesian Games: Example



Football game (two competing teams).

- Each can choose a gameplan: aim to win (Aggressive) or aim to draw (Passive)
- We call the gameplan their **type**, a private signal caused by external factors, e.g., weather, player injuries, ground conditions

Bayesian Games: Example



Football game (two competing teams).

- Each can choose a gameplan: aim to win (Aggressive) or aim to draw (Passive)
- We call the gameplan their **type**, a private signal caused by external factors, e.g., weather, player injuries, ground conditions
- Four possible type profiles: AA, AP, PA, PP (PA is the mirror of AP)

Bayesian Games: Example



Football game (two competing teams).

- Each can choose a gameplan: aim to win (Aggressive) or aim to draw (Passive)
- We call the gameplan their **type**, a private signal caused by external factors, e.g., weather, player injuries, ground conditions
- Four possible type profiles: AA, AP, PA, PP (PA is the mirror of AP)

		FRA	
		ATT	DEF
ARG	ATT	1,1	2,0
	DEF	0,2	0,0

AA profile

Bayesian Games: Example



Football game (two competing teams).

- Each can choose a gameplan: aim to win (Aggressive) or aim to draw (Passive)
- We call the gameplan their **type**, a private signal caused by external factors, e.g., weather, player injuries, ground conditions
- Four possible type profiles: AA, AP, PA, PP (PA is the mirror of AP)

		FRA	
		ATT	DEF
ARG	ATT	1,1	2,0
	DEF	0,2	0,0

AA profile

		FRA	
		ATT	DEF
ARG	ATT	2,0	2,1
	DEF	0,1	0,0

AP profile

Bayesian Games: Example



Football game (two competing teams).

- Each can choose a gameplan: aim to win (Aggressive) or aim to draw (Passive)
- We call the gameplan their **type**, a private signal caused by external factors, e.g., weather, player injuries, ground conditions
- Four possible type profiles: AA, AP, PA, PP (PA is the mirror of AP)

		FRA	
		ATT	DEF
ARG	ATT	1,1	2,0
	DEF	0,2	0,0

AA profile

		FRA	
		ATT	DEF
ARG	ATT	2,0	2,1
	DEF	0,1	0,0

AP profile

		FRA	
		ATT	DEF
ARG	ATT	0,0	1,0
	DEF	0,1	-1,-1

PP profile

One Story, Four Different Games



- In Γ_{AA} , ATT strictly dominates DEF for both teams. Unique PSNE $(1, 1)$.

One Story, Four Different Games



- In Γ_{AA} , ATT strictly dominates DEF for both teams. Unique PSNE $(1, 1)$.
- In Γ_{PP} , mutual DEF is the worst cell, $(-1, -1)$, ATT weakly dominant.

One Story, Four Different Games



- In Γ_{AA} , ATT strictly dominates DEF for both teams. Unique PSNE $(1, 1)$.
- In Γ_{PP} , mutual DEF is the worst cell, $(-1, -1)$, ATT weakly dominant.
- In Γ_{AP} , for the P team, DEF/ATT when opponent ATT/DEF.

One Story, Four Different Games



- In Γ_{AA} , ATT strictly dominates DEF for both teams. Unique PSNE $(1, 1)$.
- In Γ_{PP} , mutual DEF is the worst cell, $(-1, -1)$, ATT weakly dominant.
- In Γ_{AP} , for the P team, DEF/ATT when opponent ATT/DEF.
- So the “right” action depends on a type nobody else observes.

One Story, Four Different Games



- In Γ_{AA} , ATT strictly dominates DEF for both teams. Unique PSNE $(1, 1)$.
- In Γ_{PP} , mutual DEF is the worst cell, $(-1, -1)$, ATT weakly dominant.
- In Γ_{AP} , for the P team, DEF/ATT when opponent ATT/DEF.
- So the “right” action depends on a type nobody else observes.

Question

A team does not know which of the four matrices it is in. What is it even choosing?



One Story, Four Different Games

- In Γ_{AA} , ATT strictly dominates DEF for both teams. Unique PSNE $(1, 1)$.
- In Γ_{PP} , mutual DEF is the worst cell, $(-1, -1)$, ATT weakly dominant.
- In Γ_{AP} , for the P team, DEF/ATT when opponent ATT/DEF.
- So the “right” action depends on a type nobody else observes.

Question

A team does not know which of the four matrices it is in. What is it even choosing?

Answer

Not an action. A *rule* that maps its own type to an action. That is the whole content of the next chunk.



Assumptions

- The probabilities of choosing different games (or type profiles) come from a **common prior** distribution.



Assumptions

- The probabilities of choosing different games (or type profiles) come from a **common prior** distribution.
- The common prior is common knowledge



Assumptions

- The probabilities of choosing different games (or type profiles) come from a **common prior** distribution.
- The common prior is common knowledge

Definition

A Bayesian game is represented by $\langle N, (\Theta_i)_{i \in N}, P, (\Gamma_\theta)_{\theta \in (\times_{i \in N} \Theta_i)} \rangle$

- N : set of players



Assumptions

- The probabilities of choosing different games (or type profiles) come from a **common prior** distribution.
- The common prior is common knowledge

Definition

A Bayesian game is represented by $\langle N, (\Theta_i)_{i \in N}, P, (\Gamma_\theta)_{\theta \in (\times_{i \in N} \Theta_i)} \rangle$

- N : set of players
- Θ_i : set of types of player i



Assumptions

- The probabilities of choosing different games (or type profiles) come from a **common prior** distribution.
- The common prior is common knowledge

Definition

A Bayesian game is represented by $\langle N, (\Theta_i)_{i \in N}, P, (\Gamma_\theta)_{\theta \in (\times_{i \in N} \Theta_i)} \rangle$

- N : set of players
- Θ_i : set of types of player i
- P : common prior distribution over $\Theta = \times_{i \in N} \Theta_i$
s.t. $\sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_i, \theta_{-i}) > 0, \quad \forall \theta_i \in \Theta_i, \forall i \in N$
i.e., marginals for every type is positive (otherwise we can prune the type set)



Assumptions

- The probabilities of choosing different games (or type profiles) come from a **common prior** distribution.
- The common prior is common knowledge

Definition

A Bayesian game is represented by $\langle N, (\Theta_i)_{i \in N}, P, (\Gamma_\theta)_{\theta \in (\times_{i \in N} \Theta_i)} \rangle$

- N : set of players
- Θ_i : set of types of player i
- P : common prior distribution over $\Theta = \times_{i \in N} \Theta_i$
s.t. $\sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_i, \theta_{-i}) > 0, \quad \forall \theta_i \in \Theta_i, \forall i \in N$
i.e., marginals for every type is positive (otherwise we can prune the type set)
- Γ_θ : NFG for the type profile $\theta \in \Theta$ i.e., $\Gamma_\theta = \langle N, (A_i(\theta))_{i \in N}, (u_i(\theta))_{i \in N} \rangle$
 $u_i : A \times \Theta \rightarrow \mathbb{R}, A = \times_{i \in N} A_i$ [We assume $A_i(\theta) = A_i, \forall \theta$]



Stages of a Bayesian game

- $\theta = (\theta_i, \theta_{-i})$ is chosen randomly according to the common prior P



Stages of a Bayesian game

- $\theta = (\theta_i, \theta_{-i})$ is chosen randomly according to the common prior P
- Each player observes her own type θ_i



Stages of a Bayesian game

- $\theta = (\theta_i, \theta_{-i})$ is chosen randomly according to the common prior P
- Each player observes her own type θ_i
- Player i picks action $a_i \in A_i, \forall i \in N$



Stages of a Bayesian game

- $\theta = (\theta_i, \theta_{-i})$ is chosen randomly according to the common prior P
- Each player observes her own type θ_i
- Player i picks action $a_i \in A_i, \forall i \in N$
- Player i realizes a payoff of $u_i(a_i, a_{-i}; \theta_i, \theta_{-i})$



Stages of a Bayesian game

- $\theta = (\theta_i, \theta_{-i})$ is chosen randomly according to the common prior P
- Each player observes her own type θ_i
- Player i picks action $a_i \in A_i, \forall i \in N$
- Player i realizes a payoff of $u_i(a_i, a_{-i}; \theta_i, \theta_{-i})$

Why the positive marginal is not a technicality

Without it, $P(\theta_{-i}|\theta_i)$ is 0/0 and the player of that type has no belief to best-respond to. It is the same 0/0 we met off the equilibrium path in Lecture 11 – and here we simply legislate it away.



Chunk 2

Strategies, Ex-ante and Ex-interim Utility

($\approx$ 20–25 minutes)



Definition

Strategy is a plan to map type to action.

$$s_i : \Theta_i \rightarrow A_i$$

Pure

$$\sigma_i : \Theta_i \rightarrow \Delta A_i$$

Mixed



Definition

Strategy is a plan to map type to action.

$$s_i : \Theta_i \rightarrow A_i$$

Pure

$$\sigma_i : \Theta_i \rightarrow \Delta A_i$$

Mixed

The player can experience its utility in two stages for Bayesian games (depending on the realization of θ_i).

Ex-ante utility

Before θ_i realizes



Definition

Strategy is a plan to map type to action.

$$s_i : \Theta_i \rightarrow A_i$$

Pure

$$\sigma_i : \Theta_i \rightarrow \Delta A_i$$

Mixed

The player can experience its utility in two stages for Bayesian games (depending on the realization of θ_i).

Ex-ante utility
Before θ_i realizes



Ex-interim utility
 θ_i realized but θ_{-i} did not



Definition

Strategy is a plan to map type to action.

$$s_i : \Theta_i \rightarrow A_i$$

$$\sigma_i : \Theta_i \rightarrow \Delta A_i$$

Pure

Mixed

The player can experience its utility in two stages for Bayesian games (depending on the realization of θ_i).

Ex-ante utility
Before θ_i realizes



Ex-interim utility
 θ_i realized but θ_{-i} did not



Ex-post utility
After θ realizes

Strategy and Utilities



Definition

Strategy is a plan to map type to action.

$$s_i : \Theta_i \rightarrow A_i$$

Pure

$$\sigma_i : \Theta_i \rightarrow \Delta A_i$$

Mixed

The player can experience its utility in two stages for Bayesian games (depending on the realization of θ_i).

Ex-ante utility
Before θ_i realizes



Ex-interim utility
 θ_i realized but θ_{-i} did not



Ex-post utility
After θ realizes

Counting

A player with $|\Theta_i|$ types and $|A_i|$ actions has $|A_i|^{|\Theta_i|}$ pure strategies, not $|A_i|$. The type set is the domain of the plan, not a label on it.



Definition (Ex-ante utility)

Expected utility before observing own type.

$$\begin{aligned} u_i(\sigma) &= \sum_{\theta \in \Theta} P(\theta) u_i(\sigma(\theta); \theta) \\ &= \sum_{\theta \in \Theta} P(\theta) \sum_{(a_1, a_2, \dots, a_n) \in A} \prod_{j \in N} \sigma_j(\theta_j)[a_j] u_i(a_1, \dots, a_n; \theta_1, \dots, \theta_n) \end{aligned}$$



Ex-ante Utility

Definition (Ex-ante utility)

Expected utility before observing own type.

$$\begin{aligned} u_i(\sigma) &= \sum_{\theta \in \Theta} P(\theta) u_i(\sigma(\theta); \theta) \\ &= \sum_{\theta \in \Theta} P(\theta) \sum_{(a_1, a_2, \dots, a_n) \in A} \prod_{j \in N} \sigma_j(\theta_j)[a_j] u_i(a_1, \dots, a_n; \theta_1, \dots, \theta_n) \end{aligned}$$

The belief of player i over others' types changes after observing her own type θ_i :

$$P(\theta_{-i} | \theta_i) = \frac{P(\theta_i, \theta_{-i})}{\sum_{\tilde{\theta}_{-i} \in \Theta_{-i}} P(\theta_i, \tilde{\theta}_{-i})} \quad \text{Bayes rule}$$



Ex-ante Utility

Definition (Ex-ante utility)

Expected utility before observing own type.

$$\begin{aligned} u_i(\sigma) &= \sum_{\theta \in \Theta} P(\theta) u_i(\sigma(\theta); \theta) \\ &= \sum_{\theta \in \Theta} P(\theta) \sum_{(a_1, a_2, \dots, a_n) \in A} \prod_{j \in N} \sigma_j(\theta_j)[a_j] u_i(a_1, \dots, a_n; \theta_1, \dots, \theta_n) \end{aligned}$$

The belief of player i over others' types changes after observing her own type θ_i :

$$P(\theta_{-i} | \theta_i) = \frac{P(\theta_i, \theta_{-i})}{\sum_{\tilde{\theta}_{-i} \in \Theta_{-i}} P(\theta_i, \tilde{\theta}_{-i})} \quad \text{Bayes rule}$$

This is why we needed every marginal to be positive



Definition (Ex-interim utility)

Expected utility after observing one's own type.

$$u_i(\sigma|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_{-i}|\theta_i) u_i(\sigma(\theta); \theta)$$



Definition (Ex-interim utility)

Expected utility after observing one's own type.

$$u_i(\sigma|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_{-i}|\theta_i) u_i(\sigma(\theta); \theta)$$

Special Case: for independent types, observing θ_i does not give any information on θ_{-i}



Definition (Ex-interim utility)

Expected utility after observing one's own type.

$$u_i(\sigma|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_{-i}|\theta_i) u_i(\sigma(\theta); \theta)$$

Special Case: for independent types, observing θ_i does not give any information on θ_{-i}

Relation between the two utilities is given by

$$u_i(\sigma) = \sum_{\theta_i \in \Theta_i} P(\theta_i) u_i(\sigma|\theta_i)$$



Ex-interim utility

Definition (Ex-interim utility)

Expected utility after observing one's own type.

$$u_i(\sigma|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_{-i}|\theta_i) u_i(\sigma(\theta); \theta)$$

Special Case: for independent types, observing θ_i does not give any information on θ_{-i}

Relation between the two utilities is given by

$$u_i(\sigma) = \sum_{\theta_i \in \Theta_i} P(\theta_i) u_i(\sigma|\theta_i)$$

Ex-ante utility is the $P(\theta_i)$ -average of the ex-interim utilities. Keep this identity; the whole of Chunk 3 rests on it.

Example 1: Two Player Bargaining Game



- Player 1 : seller, type θ_1 : price at which he is willing to sell



Example 1: Two Player Bargaining Game

- Player 1 : seller, type θ_1 : price at which he is willing to sell
- Player 2 : buyer, type θ_2 : price at which he is willing to buy



Example 1: Two Player Bargaining Game

- Player 1 : seller, type θ_1 : price at which he is willing to sell
- Player 2 : buyer, type θ_2 : price at which he is willing to buy
- $\Theta_1 = \Theta_2 = \{1, 2, \dots, 100\}$, $A_1 = A_2 = \{1, 2, \dots, 100\}$



Example 1: Two Player Bargaining Game

- Player 1 : seller, type θ_1 : price at which he is willing to sell
- Player 2 : buyer, type θ_2 : price at which he is willing to buy
- $\Theta_1 = \Theta_2 = \{1, 2, \dots, 100\}$, $A_1 = A_2 = \{1, 2, \dots, 100\}$
- If the bid of the seller is smaller or equal to that of the buyer, trade happens at a price average of the two bids. Else, trade does not happen.

Example 1: Two Player Bargaining Game (Contd.)



Suppose type generation is independent and uniform over Θ_1, Θ_2 respectively,

$$P(\theta_2|\theta_1) = P(\theta_2) = \frac{1}{100}, \forall \theta_1, \theta_2$$

$$P(\theta_1|\theta_2) = P(\theta_1) = \frac{1}{100}, \forall \theta_1, \theta_2$$



Example 1: Two Player Bargaining Game (Contd.)

Suppose type generation is independent and uniform over Θ_1, Θ_2 respectively,

$$P(\theta_2|\theta_1) = P(\theta_2) = \frac{1}{100}, \forall \theta_1, \theta_2$$

$$P(\theta_1|\theta_2) = P(\theta_1) = \frac{1}{100}, \forall \theta_1, \theta_2$$

$$u_1(a_1, a_2; \theta_1, \theta_2) = \begin{cases} \frac{a_1 + a_2}{2} - \theta_1 & \text{if } a_2 \geq a_1 \\ 0 & \text{otherwise} \end{cases}$$

$$u_2(a_1, a_2; \theta_1, \theta_2) = \begin{cases} \theta_2 - \frac{a_1 + a_2}{2} & \text{if } a_2 \geq a_1 \\ 0 & \text{otherwise} \end{cases}$$



Example 1: Two Player Bargaining Game (Contd.)

Suppose type generation is independent and uniform over Θ_1, Θ_2 respectively,

$$P(\theta_2|\theta_1) = P(\theta_2) = \frac{1}{100}, \forall \theta_1, \theta_2$$

$$P(\theta_1|\theta_2) = P(\theta_1) = \frac{1}{100}, \forall \theta_1, \theta_2$$

$$u_1(a_1, a_2; \theta_1, \theta_2) = \begin{cases} \frac{a_1 + a_2}{2} - \theta_1 & \text{if } a_2 \geq a_1 \\ 0 & \text{otherwise} \end{cases}$$

$$u_2(a_1, a_2; \theta_1, \theta_2) = \begin{cases} \theta_2 - \frac{a_1 + a_2}{2} & \text{if } a_2 \geq a_1 \\ 0 & \text{otherwise} \end{cases}$$

Common Prior : $P(\theta_1, \theta_2) = \frac{1}{10000}, \forall \theta_1, \theta_2$

Example 2: Sealed Bid Auction



Two players, both willing to buy an object. Their values θ_i and bids b_i lie in $[0, 1]$.



Example 2: Sealed Bid Auction

Two players, both willing to buy an object. Their values θ_i and bids b_i lie in $[0, 1]$.

Allocation Function:

$$O_1(b_1, b_2) = \begin{cases} 1 & \text{if } b_1 \geq b_2 \\ 0 & \text{ow} \end{cases} \quad O_2(b_1, b_2) = \begin{cases} 1 & \text{if } b_2 > b_1 \\ 0 & \text{ow} \end{cases}$$



Example 2: Sealed Bid Auction

Two players, both willing to buy an object. Their values θ_i and bids b_i lie in $[0, 1]$.

Allocation Function:

$$O_1(b_1, b_2) = \begin{cases} 1 & \text{if } b_1 \geq b_2 \\ 0 & \text{ow} \end{cases} \quad O_2(b_1, b_2) = \begin{cases} 1 & \text{if } b_2 > b_1 \\ 0 & \text{ow} \end{cases}$$

Beliefs:

$$f(\theta_2|\theta_1) = 1, \forall \theta_1, \theta_2$$

$$f(\theta_1|\theta_2) = 1, \forall \theta_1, \theta_2$$

$$f(\theta_1, \theta_2) = 1, \forall \theta_1, \theta_2$$



Example 2: Sealed Bid Auction

Two players, both willing to buy an object. Their values θ_i and bids b_i lie in $[0, 1]$.

Allocation Function:

$$O_1(b_1, b_2) = \begin{cases} 1 & \text{if } b_1 \geq b_2 \\ 0 & \text{ow} \end{cases} \quad O_2(b_1, b_2) = \begin{cases} 1 & \text{if } b_2 > b_1 \\ 0 & \text{ow} \end{cases}$$

Beliefs:

$$f(\theta_2|\theta_1) = 1, \forall \theta_1, \theta_2$$

$$f(\theta_1|\theta_2) = 1, \forall \theta_1, \theta_2$$

$$f(\theta_1, \theta_2) = 1, \forall \theta_1, \theta_2$$

$$u_i(b_1, b_2; \theta_1, \theta_2) = O_i(b_1, b_2)(\theta_i - b_i)$$

Winner pays his/her bid.



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

Bayesian Equilibrium = Nash Equilibrium

($\approx$ 20–25 minutes)



Equilibrium concepts in Bayesian games

Ex-ante: before observing her own type

Nash Equilibrium (σ^*, P) : $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i', \sigma_{-i}^*), \forall \sigma_i', \forall i \in N$



Equilibrium concepts in Bayesian games

Ex-ante: before observing her own type

Nash Equilibrium (σ^*, P) : $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i', \sigma_{-i}^*), \forall \sigma_i', \forall i \in N$

Ex-interim: after observing her own type

Bayesian Equilibrium (σ^*, P) : $u_i(\sigma_i^*(\theta_i), \sigma_{-i}^*|\theta_i) \geq u_i(\sigma_i'(\theta_i), \sigma_{-i}^*|\theta_i), \forall \sigma_i', \forall \theta_i \in \Theta_i, \forall i \in N$

$$u_i(\sigma) = \sum_{\theta \in \Theta} P(\theta) u_i(\sigma(\theta); \theta), \quad u_i(\sigma|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_{-i}|\theta_i) u_i(\sigma(\theta); \theta)$$



Equilibrium concepts in Bayesian games

Ex-ante: before observing her own type

Nash Equilibrium (σ^*, P) : $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i', \sigma_{-i}^*), \forall \sigma_i', \forall i \in N$

Ex-interim: after observing her own type

Bayesian Equilibrium (σ^*, P) : $u_i(\sigma_i^*(\theta_i), \sigma_{-i}^*|\theta_i) \geq u_i(\sigma_i'(\theta_i), \sigma_{-i}^*|\theta_i), \forall \sigma_i', \forall \theta_i \in \Theta_i, \forall i \in N$

$$u_i(\sigma) = \sum_{\theta \in \Theta} P(\theta) u_i(\sigma(\theta); \theta), \quad u_i(\sigma|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_{-i}|\theta_i) u_i(\sigma(\theta); \theta)$$

- The RHS of the definitions can be replaced by a pure strategy $a_i, \forall a_i \in A_i$. The reason is exactly the same as that of MSNE (these definitions are equivalent)



Equilibrium concepts in Bayesian games

Ex-ante: before observing her own type

Nash Equilibrium (σ^*, P) : $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i', \sigma_{-i}^*), \forall \sigma_i', \forall i \in N$

Ex-interim: after observing her own type

Bayesian Equilibrium (σ^*, P) : $u_i(\sigma_i^*(\theta_i), \sigma_{-i}^*|\theta_i) \geq u_i(\sigma_i'(\theta_i), \sigma_{-i}^*|\theta_i), \forall \sigma_i', \forall \theta_i \in \Theta_i, \forall i \in N$

$$u_i(\sigma) = \sum_{\theta \in \Theta} P(\theta) u_i(\sigma(\theta); \theta), \quad u_i(\sigma|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_{-i}|\theta_i) u_i(\sigma(\theta); \theta)$$

- The RHS of the definitions can be replaced by a pure strategy $a_i, \forall a_i \in A_i$. The reason is exactly the same as that of MSNE (these definitions are equivalent)
- NE takes expectation over $P(\theta)$



Equilibrium concepts in Bayesian games

Ex-ante: before observing her own type

Nash Equilibrium (σ^*, P) : $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i', \sigma_{-i}^*), \forall \sigma_i', \forall i \in N$

Ex-interim: after observing her own type

Bayesian Equilibrium (σ^*, P) : $u_i(\sigma_i^*(\theta_i), \sigma_{-i}^*|\theta_i) \geq u_i(\sigma_i'(\theta_i), \sigma_{-i}^*|\theta_i), \forall \sigma_i', \forall \theta_i \in \Theta_i, \forall i \in N$

$$u_i(\sigma) = \sum_{\theta \in \Theta} P(\theta) u_i(\sigma(\theta); \theta), \quad u_i(\sigma|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_{-i}|\theta_i) u_i(\sigma(\theta); \theta)$$

- The RHS of the definitions can be replaced by a pure strategy $a_i, \forall a_i \in A_i$. The reason is exactly the same as that of MSNE (these definitions are equivalent)
- NE takes expectation over $P(\theta)$
- BE takes expectation over $P(\theta_{-i}|\theta_i)$

Equivalence of equilibrium concepts



Theorem

In finite Bayesian games, a strategy profile is Bayesian Equilibrium iff it is a Nash equilibrium

Equivalence of equilibrium concepts



Theorem

In finite Bayesian games, a strategy profile is Bayesian Equilibrium iff it is a Nash equilibrium

Proof.

For the forward direction, suppose (σ^*, P) is a Bayesian equilibrium, consider

$$\begin{aligned} u_i(\sigma'_i, \sigma_{-i}^*) &= \sum_{\theta_i \in \Theta_i} P(\theta_i) u_i(\sigma'_i(\theta_i), \sigma_{-i}^* | \theta_i) \\ &\leq \sum_{\theta_i \in \Theta_i} P(\theta_i) u_i(\sigma_i^*(\theta_i), \sigma_{-i}^* | \theta_i), \quad \text{since } (\sigma^*, P) \text{ is a BE} \\ &= u_i(\sigma_i^*, \sigma_{-i}^*) \end{aligned}$$





Equivalence of equilibrium concepts

Theorem

In finite Bayesian games, a strategy profile is Bayesian Equilibrium iff it is a Nash equilibrium

Proof.

For the forward direction, suppose (σ^*, P) is a Bayesian equilibrium, consider

$$\begin{aligned} u_i(\sigma'_i, \sigma_{-i}^*) &= \sum_{\theta_i \in \Theta_i} P(\theta_i) u_i(\sigma'_i(\theta_i), \sigma_{-i}^* | \theta_i) \\ &\leq \sum_{\theta_i \in \Theta_i} P(\theta_i) u_i(\sigma_i^*(\theta_i), \sigma_{-i}^* | \theta_i), \quad \text{since } (\sigma^*, P) \text{ is a BE} \\ &= u_i(\sigma_i^*, \sigma_{-i}^*) \end{aligned}$$



A weighted average of pointwise inequalities is an inequality. That is the whole forward direction.

Equivalence of equilibrium concepts



Theorem

In finite Bayesian games, a strategy profile is a Bayesian Equilibrium iff it is a Nash equilibrium

Equivalence of equilibrium concepts



Theorem

In finite Bayesian games, a strategy profile is a Bayesian Equilibrium iff it is a Nash equilibrium

Proof.

For the reverse direction, proof by contradiction. Suppose (σ^*, P) is not a Bayesian equilibrium, i.e., there exists some $i \in N$, some $\theta_i \in \Theta_i$, some $a_i \in A_i$, s.t.

$$u_i(a_i, \sigma_{-i}^* | \theta_i) > u_i(\sigma_i^*(\theta_i), \sigma_{-i}^* | \theta_i)$$



Equivalence of equilibrium concepts

Theorem

In finite Bayesian games, a strategy profile is a Bayesian Equilibrium iff it is a Nash equilibrium

Proof.

For the reverse direction, proof by contradiction. Suppose (σ^*, P) is not a Bayesian equilibrium, i.e., there exists some $i \in N$, some $\theta_i \in \Theta_i$, some $a_i \in A_i$, s.t.

$$u_i(a_i, \sigma_{-i}^* | \theta_i) > u_i(\sigma_i^*(\theta_i), \sigma_{-i}^* | \theta_i)$$

Construct the strategy $\hat{\sigma}_i$ s.t.,

$$\begin{aligned} \hat{\sigma}_i(\theta'_i) &= \sigma_i^*(\theta'_i), \quad \forall \theta'_i \in \Theta_i \setminus \{\theta_i\} \\ \hat{\sigma}_i(\theta_i)[a_i] &= 1, \quad \hat{\sigma}_i(\theta_i)[b_i] = 0, \quad \forall b_i \in A_i \setminus \{a_i\} \end{aligned}$$



Equivalence of equilibrium concepts



Proof.

Reverse direction proof continued ...

$$\begin{aligned} u_i(\hat{\sigma}_i, \sigma_{-i}^*) &= \sum_{\tilde{\theta}_i \in \Theta_i} P(\tilde{\theta}_i) u_i(\hat{\sigma}_i(\tilde{\theta}_i), \sigma_{-i}^* | \tilde{\theta}_i) \\ &= \sum_{\tilde{\theta}_i \in \Theta_i \setminus \{\theta_i\}} P(\tilde{\theta}_i) u_i(\hat{\sigma}_i(\tilde{\theta}_i), \sigma_{-i}^* | \tilde{\theta}_i) + P(\theta_i) u_i(\hat{\sigma}_i(\theta_i), \sigma_{-i}^* | \theta_i) \\ &> \sum_{\tilde{\theta}_i \in \Theta_i \setminus \{\theta_i\}} P(\tilde{\theta}_i) u_i(\sigma_i^*(\tilde{\theta}_i), \sigma_{-i}^* | \tilde{\theta}_i) + P(\theta_i) u_i(\sigma_i^*(\theta_i), \sigma_{-i}^* | \theta_i) = u_i(\sigma_i^*, \sigma_{-i}^*) \end{aligned}$$

Hence, $(\sigma_i^*, \sigma_{-i}^*)$ is not a Nash equilibrium



Equivalence of equilibrium concepts



Proof.

Reverse direction proof continued ...

$$\begin{aligned} u_i(\hat{\sigma}_i, \sigma_{-i}^*) &= \sum_{\tilde{\theta}_i \in \Theta_i} P(\tilde{\theta}_i) u_i(\hat{\sigma}_i(\tilde{\theta}_i), \sigma_{-i}^* | \tilde{\theta}_i) \\ &= \sum_{\tilde{\theta}_i \in \Theta_i \setminus \{\theta_i\}} P(\tilde{\theta}_i) u_i(\hat{\sigma}_i(\tilde{\theta}_i), \sigma_{-i}^* | \tilde{\theta}_i) + P(\theta_i) u_i(\hat{\sigma}_i(\theta_i), \sigma_{-i}^* | \theta_i) \\ &> \sum_{\tilde{\theta}_i \in \Theta_i \setminus \{\theta_i\}} P(\tilde{\theta}_i) u_i(\sigma_i^*(\tilde{\theta}_i), \sigma_{-i}^* | \tilde{\theta}_i) + P(\theta_i) u_i(\sigma_i^*(\theta_i), \sigma_{-i}^* | \theta_i) = u_i(\sigma_i^*, \sigma_{-i}^*) \end{aligned}$$

Hence, $(\sigma_i^*, \sigma_{-i}^*)$ is not a Nash equilibrium



Q: Where exactly did the strict inequality come from?



Equivalence of equilibrium concepts

Proof.

Reverse direction proof continued ...

$$\begin{aligned} u_i(\hat{\sigma}_i, \sigma_{-i}^*) &= \sum_{\tilde{\theta}_i \in \Theta_i} P(\tilde{\theta}_i) u_i(\hat{\sigma}_i(\tilde{\theta}_i), \sigma_{-i}^* | \tilde{\theta}_i) \\ &= \sum_{\tilde{\theta}_i \in \Theta_i \setminus \{\theta_i\}} P(\tilde{\theta}_i) u_i(\hat{\sigma}_i(\tilde{\theta}_i), \sigma_{-i}^* | \tilde{\theta}_i) + P(\theta_i) u_i(\hat{\sigma}_i(\theta_i), \sigma_{-i}^* | \theta_i) \\ &> \sum_{\tilde{\theta}_i \in \Theta_i \setminus \{\theta_i\}} P(\tilde{\theta}_i) u_i(\sigma_i^*(\tilde{\theta}_i), \sigma_{-i}^* | \tilde{\theta}_i) + P(\theta_i) u_i(\sigma_i^*(\theta_i), \sigma_{-i}^* | \theta_i) = u_i(\sigma_i^*, \sigma_{-i}^*) \end{aligned}$$

Hence, $(\sigma_i^*, \sigma_{-i}^*)$ is not a Nash equilibrium



Q: Where exactly did the strict inequality come from? **A:** From $P(\theta_i) > 0$.

Existence of Bayesian Equilibrium



Theorem

Every finite Bayesian game has a Bayesian equilibrium.

[Finite Bayesian game: set of players, action set and type set are finite]

Existence of Bayesian Equilibrium



Theorem

Every finite Bayesian game has a Bayesian equilibrium.

[Finite Bayesian game: set of players, action set and type set are finite]

Proof.

Proof idea: Transform the Bayesian game into a complete information game treating **each type as a player**, and invoke Nash Theorem for the existence of equilibrium – which is a BE in the original game. [See course notes for details]





Existence of Bayesian Equilibrium

Theorem

Every finite Bayesian game has a Bayesian equilibrium.

[Finite Bayesian game: set of players, action set and type set are finite]

Proof.

Proof idea: Transform the Bayesian game into a complete information game treating **each type as a player**, and invoke Nash Theorem for the existence of equilibrium – which is a BE in the original game. [See course notes for details] □

Same trick, twice

Splitting one player into one player per type turns a hard problem into a game Nash's theorem already covers.



Chunk 4

Auctions as Bayesian Games

($\approx$ 20–25 minutes)



Example 2 : Sealed Bid Auction

Two players, both willing to buy an object. Their values and bids lie in $[0, 1]$.

Allocation Function

$$O_1(b_1, b_2) = \mathbf{1}\{b_1 \geq b_2\}$$

$$O_2(b_1, b_2) = \mathbf{1}\{b_2 > b_1\}$$

Beliefs

$$f(\theta_2|\theta_1) = 1, \forall \theta_1, \theta_2$$

$$f(\theta_1|\theta_2) = 1, \forall \theta_1, \theta_2$$

$$f(\theta_1, \theta_2) = 1, \forall \theta_1, \theta_2$$

First Price Auction



- If $b_1 \geq b_2$, player 1 wins and pays her bid; otherwise, player 2 wins and pays her bid.

$$u_1(b_1, b_2, \theta_1, \theta_2) = (\theta_1 - b_1) \mathbf{1}\{b_1 \geq b_2\}$$

$$u_2(b_1, b_2, \theta_1, \theta_2) = (\theta_2 - b_2) \mathbf{1}\{b_1 < b_2\}$$



- If $b_1 \geq b_2$, player 1 wins and pays her bid; otherwise, player 2 wins and pays her bid.

$$u_1(b_1, b_2, \theta_1, \theta_2) = (\theta_1 - b_1) \mathbf{1}\{b_1 \geq b_2\}$$

$$u_2(b_1, b_2, \theta_1, \theta_2) = (\theta_2 - b_2) \mathbf{1}\{b_1 < b_2\}$$

- $b_1 = s_1(\theta_1), b_2 = s_2(\theta_2)$

Assume $s_i(\theta_i) = \alpha_i \theta_i, \alpha_i > 0, i = 1, 2$

First Price Auction



To find the BE, we need to find the s_i^* (or α_i^*) that maximizes the ex-interim utility of player i . i.e.

$$\max_{\sigma_i} u_i(\sigma_i, \sigma_{-i}^* | \theta_i)$$

$$\text{if } \frac{b_1}{\alpha_2} \leq 1$$

First Price Auction



To find the BE, we need to find the s_i^* (or α_i^*) that maximizes the ex-interim utility of player i . i.e.

$$\max_{\sigma_i} u_i(\sigma_i, \sigma_{-i}^* | \theta_i)$$

For player 1, this reduces to

$$\max_{b_1 \in [0,1]} u_1(b_1, \sigma_2^* | \theta_1)$$

$$\text{if } \frac{b_1}{\alpha_2} \leq 1$$

First Price Auction



To find the BE, we need to find the s_i^* (or α_i^*) that maximizes the ex-interim utility of player i . i.e.

$$\max_{\sigma_i} u_i(\sigma_i, \sigma_{-i}^* | \theta_i)$$

For player 1, this reduces to

$$\max_{b_1 \in [0,1]} u_1(b_1, \sigma_2^* | \theta_1) = \max_{b_1 \in [0,1]} \int_0^1 f(\theta_2 | \theta_1) (\theta_1 - b_1) \mathbf{1}\{b_1 \geq \alpha_2 \theta_2\} d\theta_2$$

if $\frac{b_1}{\alpha_2} \leq 1$

First Price Auction



To find the BE, we need to find the s_i^* (or α_i^*) that maximizes the ex-interim utility of player i . i.e.

$$\max_{\sigma_i} u_i(\sigma_i, \sigma_{-i}^* | \theta_i)$$

For player 1, this reduces to

$$\begin{aligned} \max_{b_1 \in [0,1]} u_1(b_1, \sigma_2^* | \theta_1) &= \max_{b_1 \in [0,1]} \int_0^1 f(\theta_2 | \theta_1) (\theta_1 - b_1) \mathbf{1}\{b_1 \geq \alpha_2 \theta_2\} d\theta_2 \\ &= \max_{b_1 \in [0,1]} (\theta_1 - b_1) \frac{b_1}{\alpha_2} \quad \text{if } \frac{b_1}{\alpha_2} \leq 1 \end{aligned}$$

First Price Auction



To find the BE, we need to find the s_i^* (or α_i^*) that maximizes the ex-interim utility of player i . i.e.

$$\max_{\sigma_i} u_i(\sigma_i, \sigma_{-i}^* | \theta_i)$$

For player 1, this reduces to

$$\begin{aligned} \max_{b_1 \in [0,1]} u_1(b_1, \sigma_2^* | \theta_1) &= \max_{b_1 \in [0,1]} \int_0^1 f(\theta_2 | \theta_1) (\theta_1 - b_1) \mathbf{1}\{b_1 \geq \alpha_2 \theta_2\} d\theta_2 \\ &= \max_{b_1 \in [0,1]} (\theta_1 - b_1) \frac{b_1}{\alpha_2} \quad \text{if } \frac{b_1}{\alpha_2} \leq 1 \\ \implies b_1 &= \frac{\theta_1}{2} \end{aligned}$$



First Price Auction

To find the BE, we need to find the s_i^* (or α_i^*) that maximizes the ex-interim utility of player i . i.e.

$$\max_{\sigma_i} u_i(\sigma_i, \sigma_{-i}^* | \theta_i)$$

For player 1, this reduces to

$$\begin{aligned} \max_{b_1 \in [0,1]} u_1(b_1, \sigma_2^* | \theta_1) &= \max_{b_1 \in [0,1]} \int_0^1 f(\theta_2 | \theta_1) (\theta_1 - b_1) \mathbf{1}\{b_1 \geq \alpha_2 \theta_2\} d\theta_2 \\ &= \max_{b_1 \in [0,1]} (\theta_1 - b_1) \frac{b_1}{\alpha_2} \quad \text{if } \frac{b_1}{\alpha_2} \leq 1 \\ \implies b_1 &= \frac{\theta_1}{2} \end{aligned}$$

Note: α_2 scales the objective but not the maximiser. Player 1's best reply does not depend on how aggressively player 2 shades (underreports).

First Price Auction



From this we get,

$$s_1^*(\theta_1) = \frac{\theta_1}{2}$$
$$s_2^*(\theta_2) = \frac{\theta_2}{2}$$

is a BE.

In the Bayesian Game induced by uniform prior on first price auction, bidding half the true value is a Bayesian equilibrium.

First Price Auction



From this we get,

$$s_1^*(\theta_1) = \frac{\theta_1}{2}$$
$$s_2^*(\theta_2) = \frac{\theta_2}{2}$$

is a BE.

In the Bayesian Game induced by uniform prior on first price auction, bidding half the true value is a Bayesian equilibrium.

Reading it

Shading by a factor of two is the exact trade-off between winning less often and paying less when you do win.

Second Price Auction



Highest bidder wins but pays the second highest bid.

$$u_1(b_1, b_2, \theta_1, \theta_2) = (\theta_1 - b_2) \mathbf{1}\{b_1 \geq b_2\}$$

$$u_2(b_1, b_2, \theta_1, \theta_2) = (\theta_2 - b_1) \mathbf{1}\{b_1 < b_2\}$$

Second Price Auction



Highest bidder wins but pays the second highest bid.

$$u_1(b_1, b_2, \theta_1, \theta_2) = (\theta_1 - b_2) \mathbf{1}\{b_1 \geq b_2\}$$

$$u_2(b_1, b_2, \theta_1, \theta_2) = (\theta_2 - b_1) \mathbf{1}\{b_1 < b_2\}$$

Player 1 has to maximize

$$\begin{aligned} &= \int_0^1 f(\theta_2 | \theta_1) (\theta_1 - s_2(\theta_2)) \mathbf{1}\{b_1 \geq s_2(\theta_2)\} d\theta_2 \\ &= \int_0^1 1 \cdot (\theta_1 - \alpha_2 \theta_2) \mathbf{1}\{\theta_2 \leq \frac{b_1}{\alpha_2}\} d\theta_2 \\ &= \frac{1}{\alpha_2} (b_1 \theta_1 - \frac{b_1^2}{2}) \end{aligned}$$

This is maximized when $b_1 = \theta_1$. Similarly for $b_2 = \theta_2$.

Second Price Auction



If the distribution of θ_1 and θ_2 were arbitrary but independent, the maximization problem would have been

$$\int_0^{\frac{b_1}{\alpha_2}} f(\theta_2)(\theta_1 - \alpha_2 \theta_2) d\theta_2 = \theta_1 F\left(\frac{b_1}{\alpha_2}\right) - \alpha_2 \int_0^{\frac{b_1}{\alpha_2}} \theta_2 f(\theta_2) d\theta_2$$

Second Price Auction



If the distribution of θ_1 and θ_2 were arbitrary but independent, the maximization problem would have been

$$\int_0^{\frac{b_1}{\alpha_2}} f(\theta_2)(\theta_1 - \alpha_2 \theta_2) d\theta_2 = \theta_1 F\left(\frac{b_1}{\alpha_2}\right) - \alpha_2 \int_0^{\frac{b_1}{\alpha_2}} \theta_2 f(\theta_2) d\theta_2$$

Differentiating w.r.t. b_1 , we get

$$\theta_1 \frac{1}{\alpha_2} f\left(\frac{b_1}{\alpha_2}\right) - \alpha_2 \cdot \frac{b_1}{\alpha_2} f\left(\frac{b_1}{\alpha_2}\right) \frac{1}{\alpha_2} = 0 \implies f\left(\frac{b_1}{\alpha_2}\right) (b_1 - \theta_1) = 0 \quad (1)$$

$$\implies b_1 = \theta_1, \text{ if } f\left(\frac{b_1}{\alpha_2}\right) > 0 \quad (2)$$

Similarly for player 2.



Second Price Auction

If the distribution of θ_1 and θ_2 were arbitrary but independent, the maximization problem would have been

$$\int_0^{\frac{b_1}{\alpha_2}} f(\theta_2)(\theta_1 - \alpha_2 \theta_2) d\theta_2 = \theta_1 F\left(\frac{b_1}{\alpha_2}\right) - \alpha_2 \int_0^{\frac{b_1}{\alpha_2}} \theta_2 f(\theta_2) d\theta_2$$

Differentiating w.r.t. b_1 , we get

$$\theta_1 \frac{1}{\alpha_2} f\left(\frac{b_1}{\alpha_2}\right) - \alpha_2 \cdot \frac{b_1}{\alpha_2} f\left(\frac{b_1}{\alpha_2}\right) \frac{1}{\alpha_2} = 0 \implies f\left(\frac{b_1}{\alpha_2}\right) (b_1 - \theta_1) = 0 \quad (1)$$

$$\implies b_1 = \theta_1, \text{ if } f\left(\frac{b_1}{\alpha_2}\right) > 0 \quad (2)$$

Similarly for player 2.

For any independent positive prior, bidding true type is a BE of the induced Bayesian game in Second Price Auction.



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.



- Under imperfect information, an action is optimal only *relative to a belief*. A **belief** is a distribution over an information set; it is **Bayesian** when it comes from σ by Bayes' rule; an assessment that is Bayesian and sequentially rational is a **PBE**.



- Under imperfect information, an action is optimal only *relative to a belief*. A **belief** is a distribution over an information set; it is **Bayesian** when it comes from σ by Bayes' rule; an assessment that is Bayesian and sequentially rational is a **PBE**.
- Bayes' rule is silent off the path. That silence is a real gap, and BitTorrent's attackers lived in it.



- Under imperfect information, an action is optimal only *relative to a belief*. A **belief** is a distribution over an information set; it is **Bayesian** when it comes from σ by Bayes' rule; an assessment that is Bayesian and sequentially rational is a **PBE**.
- Bayes' rule is silent off the path. That silence is a real gap, and BitTorrent's attackers lived in it.
- Under incomplete information, a player's private signal is her **type**, her strategy is a map $\Theta_i \rightarrow A_i$, and **Bayesian equilibrium** is ex-interim optimality for every type.



- Under imperfect information, an action is optimal only *relative to a belief*. A **belief** is a distribution over an information set; it is **Bayesian** when it comes from σ by Bayes' rule; an assessment that is Bayesian and sequentially rational is a **PBE**.
- Bayes' rule is silent off the path. That silence is a real gap, and BitTorrent's attackers lived in it.
- Under incomplete information, a player's private signal is her **type**, her strategy is a map $\Theta_i \rightarrow A_i$, and **Bayesian equilibrium** is ex-interim optimality for every type.
- In finite Bayesian games, BE and NE coincide, and both always exist.



- Under imperfect information, an action is optimal only *relative to a belief*. A **belief** is a distribution over an information set; it is **Bayesian** when it comes from σ by Bayes' rule; an assessment that is Bayesian and sequentially rational is a **PBE**.
- Bayes' rule is silent off the path. That silence is a real gap, and BitTorrent's attackers lived in it.
- Under incomplete information, a player's private signal is her **type**, her strategy is a map $\Theta_i \rightarrow A_i$, and **Bayesian equilibrium** is ex-interim optimality for every type.
- In finite Bayesian games, BE and NE coincide, and both always exist.
- First price: bid $\theta/2$. Second price: bid θ , for any independent positive prior.



- Under imperfect information, an action is optimal only *relative to a belief*. A **belief** is a distribution over an information set; it is **Bayesian** when it comes from σ by Bayes' rule; an assessment that is Bayesian and sequentially rational is a **PBE**.
- Bayes' rule is silent off the path. That silence is a real gap, and BitTorrent's attackers lived in it.
- Under incomplete information, a player's private signal is her **type**, her strategy is a map $\Theta_i \rightarrow A_i$, and **Bayesian equilibrium** is ex-interim optimality for every type.
- In finite Bayesian games, BE and NE coincide, and both always exist.
- First price: bid $\theta/2$. Second price: bid θ , for any independent positive prior.



- Under imperfect information, an action is optimal only *relative to a belief*. A **belief** is a distribution over an information set; it is **Bayesian** when it comes from σ by Bayes' rule; an assessment that is Bayesian and sequentially rational is a **PBE**.
- Bayes' rule is silent off the path. That silence is a real gap, and BitTorrent's attackers lived in it.
- Under incomplete information, a player's private signal is her **type**, her strategy is a map $\Theta_i \rightarrow A_i$, and **Bayesian equilibrium** is ex-interim optimality for every type.
- In finite Bayesian games, BE and NE coincide, and both always exist.
- First price: bid $\theta/2$. Second price: bid θ , for any independent positive prior.

Next week: mechanism design – we stop analysing games and start *designing* them.



End of Lecture 12

See you next class.



भारतीय प्रौद्योगिकी संस्थान मुंबई

Indian Institute of Technology Bombay