भारतीय प्रौद्योगिकी संस्थान मुंबई
Indian Institute of Technology Bombay

CS 6001: Game Theory and Algorithmic Mechanism Design

Week 6

Swaprava Nath

Slide preparation acknowledgments: Ramsundar Anandanarayanan and Harshvardhan Agarwal

ज्ञानम् परमम् ध्येयम्

Knowledge is the supreme goal

Contents

Equilibrium notions in IIEFG

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

Equilibrium notions in IIEFG

- 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

Equilibrium notions in IIEFG

- 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

Equilibrium notions in IIEFG

- 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: An IIEFG with perfect recall

EX 7.38 MSZ: An IIEFG with perfect recall, i.e., mixed and behavioral strategies are equivalent.

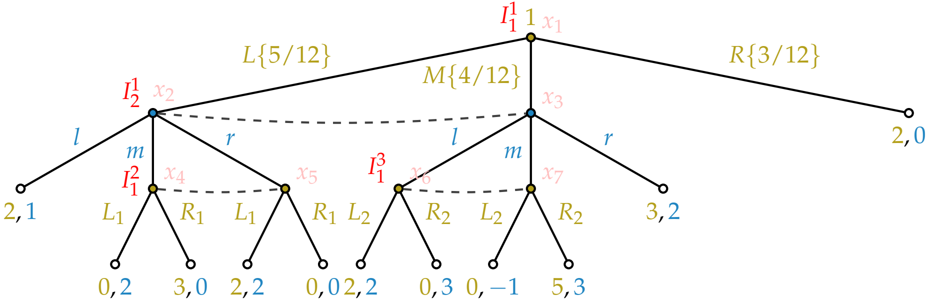

Example: An IIEFG with perfect recall

EX 7.38 MSZ: An IIEFG with perfect recall, i.e., mixed and behavioral strategies are equivalent.

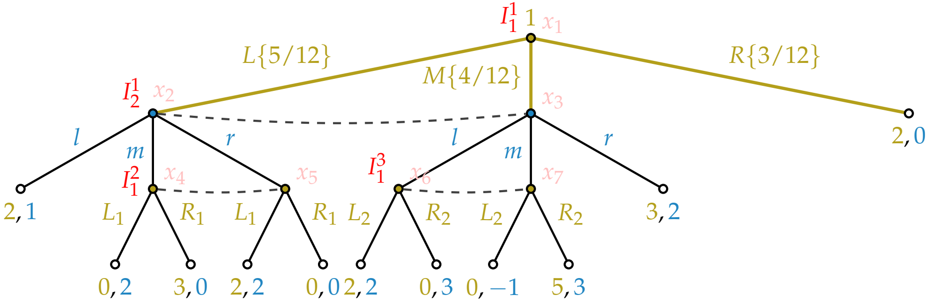

Consider the behavioral strategy profile: σ_1, at $I_1^1(L\{5/12\}, M\{4/12\}, R\{3/12\})$

Example: An IIEFG with perfect recall

EX 7.38 MSZ: An IIEFG with perfect recall, i.e., mixed and behavioral strategies are equivalent.

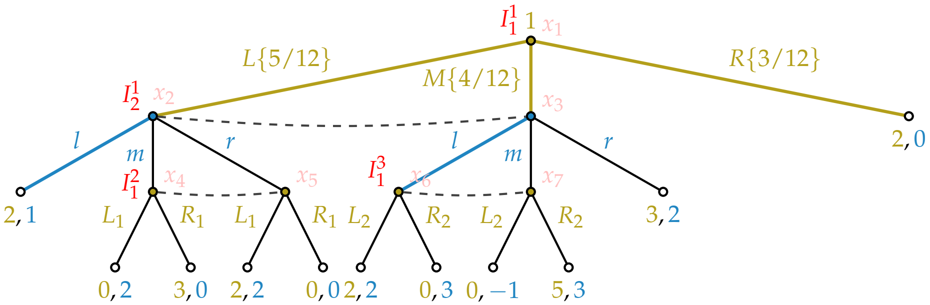

Consider the behavioral strategy profile: σ_2, at $I_2^1(l\{1\}, m\{0\}, r\{0\})$ choose l

Example: An IIEFG with perfect recall

EX 7.38 MSZ: An IIEFG with perfect recall, i.e., mixed and behavioral strategies are equivalent.

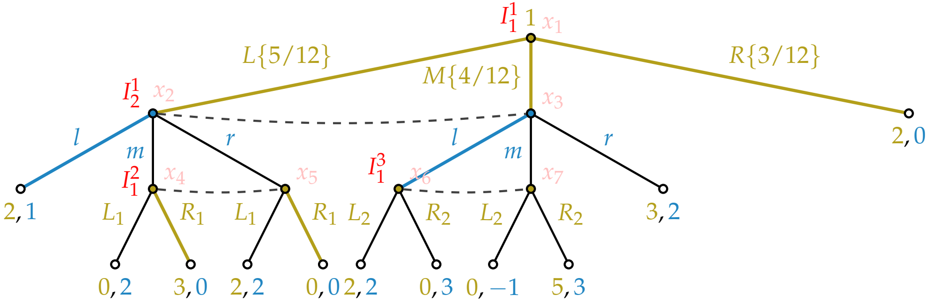

Consider the behavioral strategy profile: σ_1, at $I_1^2(L_1\{0\}, R_1\{1\})$ choose R_1

Example: An IIEFG with perfect recall

EX 7.38 MSZ: An IIEFG with perfect recall, i.e., mixed and behavioral strategies are equivalent.

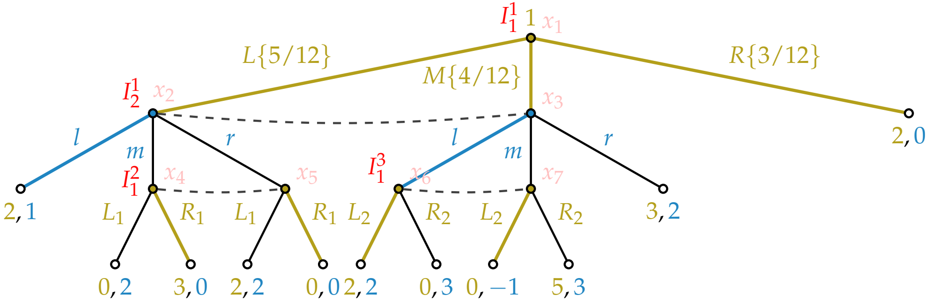

Consider the behavioral strategy profile: σ_1, at $I_1^3(L_2\{1\}, R_2\{0\})$ choose L_2

Example: An IIEFG with perfect recall

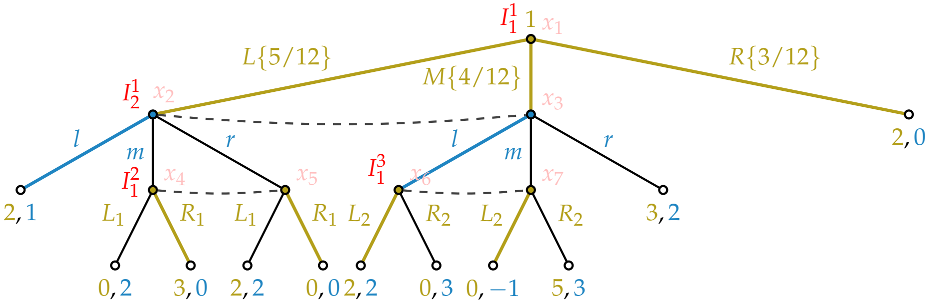

Question

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?

Example: An IIEFG with perfect recall

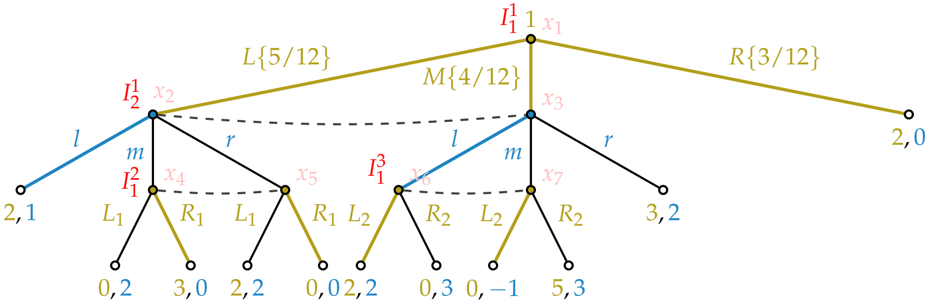

Sequential rationality

Choose an action maximizing expected utility at each information set.

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

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

Example: An IIEFG with perfect recall

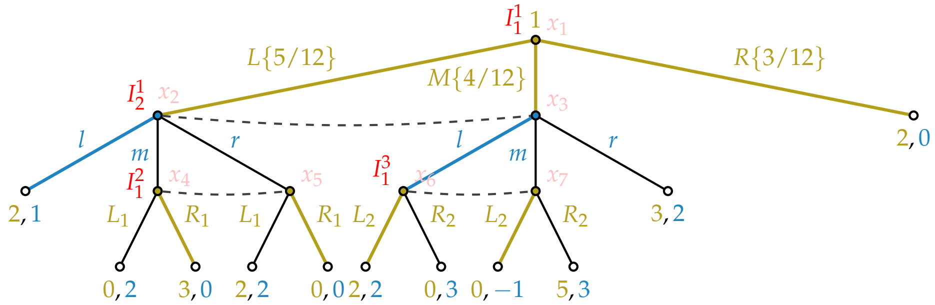

- Player 1 at information set I_1^3, believes that x_6 is reached with probability 1.
- If the belief was $> 2/7$ in favor of x_7, player 1 should have chosen R_2

Example: An IIEFG with perfect recall

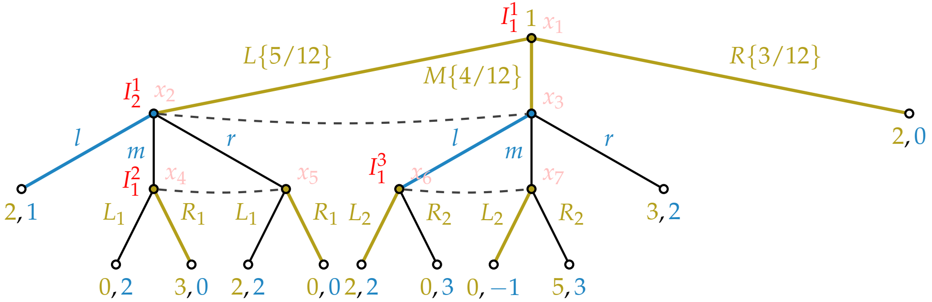

- Player 2 at I_2^1 believes the x_3 is reached w.p. $P_\sigma(x_3|I_2^1) = P_\sigma(x_3)/(P_\sigma(x_2) + P_\sigma(x_3)) = 4/9$
- Similarly $P_\sigma(x_2|I_2^1) = 5/9$

Example: An IIEFG with perfect recall

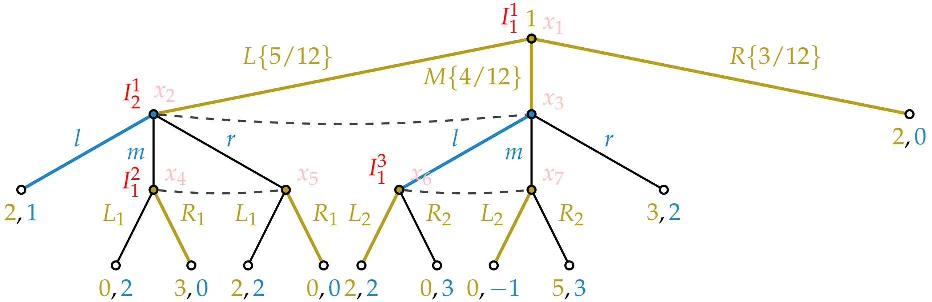

Question

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

Answer

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

Example: An IIEFG with perfect recall

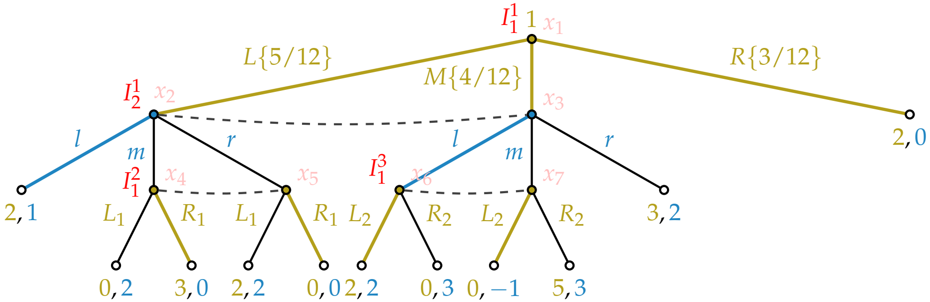

Question

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 $= 2$).

Example: An IIEFG with perfect recall

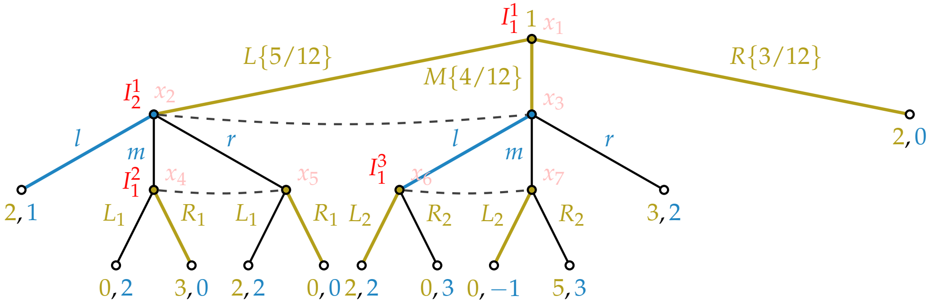

Thus, mixed/behavioral strategy profile σ is sequentially rational for all players.

Formal definitions

Belief

Let the information sets of player i be $I_i = \{I_i^1, I_i^2, I_i^3, ..., I_i^{k(i)}\}$.
The belief of player i is a mapping $\mu_i^j : I_i^j \to [0,1]$ s.t., $\sum_{x \in I_i^j} \mu_i^j(x) = 1$

Formal definitions

Belief

Let the information sets of player i be $I_i = \{I_i^1, I_i^2, I_i^3, ..., I_i^{k(i)}\}$.
The belief of player i is a mapping $\mu_i^j : I_i^j \to [0,1]$ s.t., $\sum_{x \in I_i^j} \mu_i^j(x) = 1$

Bayesian belief

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

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

Formal definitions

Sequential rationality

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

$$\sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma_i, \sigma_{-i} | x) \geqslant \sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma_i', \sigma_{-i} | x)$$

Formal definitions

Sequential rationality

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

$$\sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma_i, \sigma_{-i} | x) \geqslant \sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma_i', \sigma_{-i} | x)$$

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

Formal definitions

Sequential rationality

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

$$\sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma_i, \sigma_{-i} | x) \geqslant \sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma_i', \sigma_{-i} | x)$$

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

Formal definitions

Sequential rationality

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

$$\sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma_i, \sigma_{-i}|x) \geqslant \sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma_i', \sigma_{-i}|x)$$

- The tuple (σ, μ) is sequentially rational if it is sequentially rational for every player at every information set.
- The tuple (σ, μ) is also called an assessment.
- Sequential rationality is a refinement of Nash Equilibrium.

Formal definitions

Sequential rationality

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

$$\sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma_i, \sigma_{-i} | x) \geqslant \sum_{x \in I_i^j} \mu_i^j(x) u_i(\sigma_i', \sigma_{-i} | x)$$

- The tuple (σ, μ) is sequentially rational if it is sequentially rational for every player at every information set.
- The tuple (σ, μ) is also called an assessment.
- Sequential rationality is a refinement of Nash Equilibrium.
- The notion coincides with SPNE when applied to PIEFGs

Formal definitions

Theorem

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

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.

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.

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.

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

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.

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.

Contents

Peer to Peer[1]

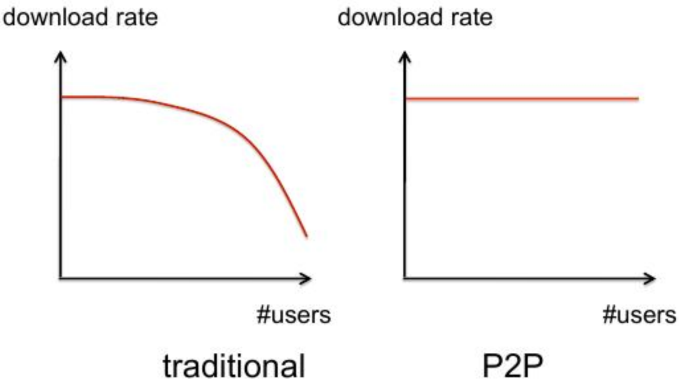

[1]Slides of this section are adapted from CS186, Harvard

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**

Early P2P Technologies

Napster (1999 - 2001)

- Centralized database
- Users download music from each other

Early P2P Technologies

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	$2, 2$	$-1, 3, -1$	$0, 0$

(Gnutella) File Sharing Game

The File Sharing Game (Contd.)

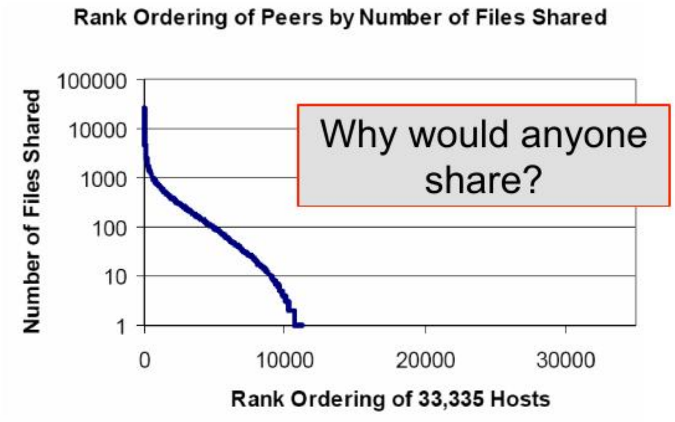

Image courtesy: Adar and Huberman (2000)

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

New Protocol

BitTorrent (2001 -)

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

New Protocol

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

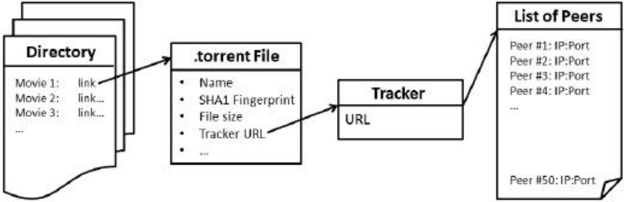

Figure 5.4.: Starting a download process in the BitTorrent protocol: 1) A user goes to a searchable directory to find a link to a .torrent file corresponding to the desired content; 2) the .torrent file contains metadata about the content, in particular the URL of a tracker; 3) the tracker provides a 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.

Image courtesy: Parkes and Seuken (2017)

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 $\geqslant r$, 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 $\geqslant r$, 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 $\geqslant r$, 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 $\geqslant r$, 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 $\geqslant r$, 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 $\geqslant r$, 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

Illustration

Strategic Behaviors

- 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

BitThief

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

BitThief

- 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

Strategic Piece Revealer

- 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

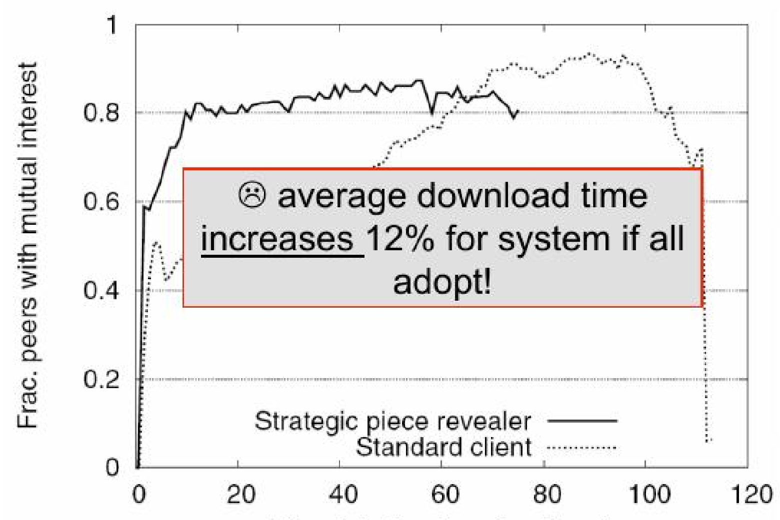

☹ average download time increases 12% for system if all adopt!

Summary

- 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.

Contents

Classification of Games

Games

- **Non-cooperative games**

Classification of Games

Games

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

Classification of Games

Games

- **Non-cooperative games**
  - Complete information - Players **deterministically** know which game they are playing
    - Normal form games
      Appropriate for simultaneous move single-stage games
      Equilibrium notions: SDSE, WDSE, PSNE, MSNE, Correlated

Classification of Games

Games

- **Non-cooperative games**
  - Complete information - Players **deterministically** know which game they are playing
    - Normal form games
      Appropriate for simultaneous move single-stage games
      Equilibrium notions: SDSE, WDSE, PSNE, MSNE, Correlated
    - Extensive form games
      Appropriate for multi-stage games
      Equilibrium notions: SPNE (PIEFG), mixed and behavioral strategies (IIEFG), PBE

Classification of Games

Games

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

Classification of Games

Games

- **Non-cooperative games**
  - Complete information - Players **deterministically** know which game they are playing
    - Normal form games
      Appropriate for simultaneous move single-stage games
      Equilibrium notions: SDSE, WDSE, PSNE, MSNE, Correlated
    - Extensive form games
      Appropriate for multi-stage games
      Equilibrium notions: 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

Classification of Games

Games

- **Non-cooperative games**
  - Complete information - Players **deterministically** know which game they are playing
    - Normal form games
      Appropriate for simultaneous move single-stage games
      Equilibrium notions: SDSE, WDSE, PSNE, MSNE, Correlated
    - Extensive form games
      Appropriate for multi-stage games
      Equilibrium notions: 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 of games - 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

- Players do not know deterministically know which game they are playing

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 know 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 know 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 know 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 (W) or aim to draw (D)

Bayesian Games: Example

Football game (two competing teams)

- Each can choose a gameplan: aim to win (W) or aim to draw (D)
- We will call the gameplan as their **type** which are private signals to them, often caused by external factors, e.g., weather conditions, player injuries, ground conditions etc.

Bayesian Games: Example

Football game (two competing teams)

- Each can choose a gameplan: aim to win (W) or aim to draw (D)
- We will call the gameplan as their **type** which are private signals to them, often caused by external factors, e.g., weather conditions, player injuries, ground conditions etc.
- There are four possible type profiles in this example WW, WD, DW, DD.

Bayesian Games: Example

Football game (two competing teams)

- Each can choose a gameplan: aim to win (W) or aim to draw (D)
- We will call the gameplan as their **type** which are private signals to them, often caused by external factors, e.g., weather conditions, player injuries, ground conditions etc.
- There are four possible type profiles in this example WW, WD, DW, DD.
- The payoff matrices differ as follows (payoff for DW is symmetrically opposite to WD).

Bayesian Games: Example

Football game (two competing teams)

- Each can choose a gameplan: aim to win (W) or aim to draw (D)
- We will call the gameplan as their **type** which are private signals to them, often caused by external factors, e.g., weather conditions, player injuries, ground conditions etc.
- There are four possible type profiles in this example WW, WD, DW, DD.
- The payoff matrices differ as follows (payoff for DW is symmetrically opposite to WD).

Bayesian Games: Example

Football game (two competing teams)

- Each can choose a gameplan: aim to win (W) or aim to draw (D)
- We will call the gameplan as their **type** which are private signals to them, often caused by external factors, e.g., weather conditions, player injuries, ground conditions etc.
- There are four possible type profiles in this example WW, WD, DW, DD.
- The payoff matrices differ as follows (payoff for DW is symmetrically opposite to WD).

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

WW profile

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

WD profile

ARG \ FRA	ATT	DEF
ATT	0, 0	1, 0
DEF	0, 1	−1, −1

DD profile

Bayesian Games

Assumptions

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

Bayesian Games

Assumptions

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

Bayesian Games

Assumptions

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

Bayesian Games

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

Bayesian Games

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

Bayesian Games

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)

Bayesian Games

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., $\Gamma_\theta = \langle N, (A_i(\theta))_{i \in N}, (u_i(\theta))_{i \in N} \rangle$
  $u_i : A \times \Theta \to \mathbb{R}$, $A = \times_{i \in N} A_i$ [We assume $A_i(\theta) = A_i, \forall \theta$]

Bayesian games

Stages of a Bayesian game

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

Bayesian games

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

Bayesian games

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 $a_i \in A_i, \forall i \in N$

Bayesian games

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 $u_i(a_i, a_{-i}; \theta_i, \theta_{-i})$

Contents

Strategy and Utilities

Definition

Strategy is a plan to map type to action.

$$s_i : \Theta_i \to A_i \qquad \textbf{Pure}$$

$$\sigma_i : \Theta_i \to \Delta A_i \qquad \textbf{Mixed}$$

Strategy and Utilities

Definition

Strategy is a plan to map type to action.

$$s_i : \Theta_i \to A_i \qquad \textbf{Pure}$$
$$\sigma_i : \Theta_i \to \Delta A_i \qquad \textbf{Mixed}$$

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

- Ex-ante utility

Strategy and Utilities

Definition

Strategy is a plan to map type to action.

$$s_i : \Theta_i \to A_i \qquad \textbf{Pure}$$
$$\sigma_i : \Theta_i \to \Delta A_i \qquad \textbf{Mixed}$$

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

- Ex-ante utility
- Ex-interim utility

Strategy and Utilities

Definition

Strategy is a plan to map type to action.

$$s_i : \Theta_i \rightarrow A_i \qquad \textbf{Pure}$$
$$\sigma_i : \Theta_i \rightarrow \Delta A_i \qquad \textbf{Mixed}$$

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

- Ex-ante utility
- Ex-interim utility
- Ex-post utility (for complete information game)

Ex-ante Utility

Definition (Ex-ante utility)

Expected utility before observing own type.

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

Ex-ante Utility

Definition (Ex-ante utility)

Expected utility before observing own type.

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

The belief of player i over others' types changes after observing her own type θ_i according to Bayes rule on P.

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

Ex-ante Utility

Definition (Ex-ante utility)

Expected utility before observing own type.

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

The belief of player i over others' types changes after observing her own type θ_i according to Bayes rule on P.

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

This is why we needed every marginal to be positive – otherwise that type can be removed from its type set

Ex-interim utility

Definition (Ex-interim utility)

Expected utility after observing one's own type.

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

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}. Both utilities are the same.

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}. Both utilities are the same.

Relation between the two utilities is given by

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

Example 1: Two Player Bargaining Game

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

Example 1: Two Player Bargaining Game

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

Example 1: Two Player Bargaining Game

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

Example 1: Two Player Bargaining Game

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

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

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

$$u_1(a_1, a_2; \theta_1, \theta_2) = \begin{cases} \frac{a_1+a_2}{2} - \theta_1 & \text{if } a_2 \geqslant a_1 \\ 0 & \text{otherwise} \end{cases}$$
$$u_2(a_1, a_2; \theta_1, \theta_2) = \begin{cases} \theta_2 - \frac{a_1+a_2}{2} & \text{if } a_2 \geqslant a_1 \\ 0 & \text{otherwise} \end{cases}$$

Common Prior : $P(\theta_1, \theta_2) = \frac{1}{1000}, \forall \theta_1, \theta_2$

Example 2: Sealed Bid Auction

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

Allocation Function:

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

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 for his bid.

Contents

Equilibrium concepts in Bayesian games

Ex-ante: before observing her own type

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

$$u_i(\sigma) = \sum_{\theta \in \Theta} P(\theta) u_i(\sigma(\theta); \theta)$$

Equilibrium concepts in Bayesian games

Ex-ante: before observing her own type

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

$$u_i(\sigma) = \sum_{\theta \in \Theta} P(\theta) u_i(\sigma(\theta); \theta)$$

Ex-interim: after observing her own type

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

Equilibrium concepts in Bayesian games

Ex-ante: before observing her own type

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

$$u_i(\sigma) = \sum_{\theta \in \Theta} P(\theta) u_i(\sigma(\theta); \theta)$$

Ex-interim: after observing her own type

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

- The RHS of the definition 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): $u_i(\sigma_i^*, \sigma_{-i}^*) \geqslant u_i(\sigma_i', \sigma_{-i}^*), \forall \sigma_i', \forall i \in N$

$$u_i(\sigma) = \sum_{\theta \in \Theta} P(\theta) u_i(\sigma(\theta); \theta)$$

Ex-interim: after observing her own type

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

- The RHS of the definition 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): $u_i(\sigma_i^*, \sigma_{-i}^*) \geqslant u_i(\sigma_i', \sigma_{-i}^*), \forall \sigma_i', \forall i \in N$

$$u_i(\sigma) = \sum_{\theta \in \Theta} P(\theta) u_i(\sigma(\theta); \theta)$$

Ex-interim: after observing her own type

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

- The RHS of the definition 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

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

□

Equivalence of equilibrium concepts

Theorem

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

Equivalence of equilibrium concepts

Theorem

In finite Bayesian games, a strategy profile is 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 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.,

$$\hat{\sigma}_i(\theta'_i) = \sigma^*_i(\theta'_i), \forall \theta'_i \in \Theta_i \setminus \{\theta_i\}$$

Equivalence of equilibrium concepts

Theorem

In finite Bayesian games, a strategy profile is 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.,

$$\hat{\sigma}_i(\theta'_i) = \sigma^*_i(\theta'_i), \forall \theta'_i \in \Theta_i \setminus \{\theta_i\}$$

$$\hat{\sigma}_i(\theta_i)[a_i] = 1, \hat{\sigma}_i(\theta_i)[b_i] = 0, \forall b_i \in A_i \setminus \{a_i\}$$

□

Equivalence of equilibrium concepts

Theorem

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

Proof.

Reverse direction proof continued ...

□

Equivalence of equilibrium concepts

Theorem

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

Proof.

Reverse direction proof continued ...

$$u_i(\hat{\sigma}_i, \sigma^*_{-i}) = \sum_{\tilde{\theta}_i \in \Theta_i} P(\tilde{\theta}_i) u_i(\hat{\sigma}_i(\tilde{\theta}_i), \sigma^*_{-i} | \tilde{\theta}_i)$$

$\square$

Equivalence of equilibrium concepts

Theorem

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

Proof.

Reverse direction proof continued ...

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

$\square$

Equivalence of equilibrium concepts

Theorem

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

Proof.

Reverse direction proof continued ...

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

$\square$

Equivalence of equilibrium concepts

Theorem

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

Proof.

Reverse direction proof continued ...

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

$\square$

Equivalence of equilibrium concepts

Theorem

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

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 $\square$

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 addendum for details] □

Contents

Example 2 : Sealed Bid Auction

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

$$O_1(b_1, b_2) = I\{b_1 \geqslant b_2\}$$
$$O_2(b_1, b_2) = I\{b_2 > b_1\}$$

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 $b_1 \geqslant b_2$ payer 1 wins and pays her bid otherwise, player 2 wins and pays her bid.

$$u_1(b_1, b_2, \theta_1, \theta_2) = (\theta_1 - b_1)T\{b_1 \geqslant b_2\}$$
$$u_2(b_1, b_2, \theta_1, \theta_2) = (\theta_2 - b_2)T\{b_1 < b_2\}$$

First Price Auction

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

$$u_1(b_1, b_2, \theta_1, \theta_2) = (\theta_1 - b_1)T\{b_1 \geqslant b_2\}$$
$$u_2(b_1, b_2, \theta_1, \theta_2) = (\theta_2 - b_2)T\{b_1 < b_2\}$$

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

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

For player 1, this reduces to

$$\begin{aligned}
max_{\sigma_i} u_i(\sigma_i, \sigma_{-i}^* | \theta_i) &= max_{b_1 \in [0,\alpha_2]} \int_0^1 f(\theta_2|\theta_1)(\theta_1 - b_1) I\{b_1 \geqslant \alpha_2\theta_2\} d\theta_2 \\
&= max_{b_1 \in [0,\alpha_2]} (\theta_1 - b_1) \frac{b_1}{\alpha_2} \\
\implies b_1 &= \begin{cases} \frac{\theta_1}{2} \text{if } \alpha_2 > \frac{\theta_1}{2} \\ \alpha_2 \text{otherwise} \end{cases}
\end{aligned}$$

First Price Auction

From this we get,

$$s_1^*(\theta_1) = min\{\frac{\theta_1}{2}, \alpha_2\}$$
$$s_2^*(\theta_2) = min\{\frac{\theta_2}{2}, \alpha_1\}$$

If $\alpha_1 = \alpha_2 = \frac{1}{2}$, then $(\frac{\theta_1}{2}, \frac{\theta_2}{2})$ is a BE.
In the Bayesian Game induced by uniform prior on first price auction, bidding half the true value is a Bayesian equilibrium.

Second Price Auction

Highest bidder wins but pays the second highest bid.

$$u_1(b_1, b_2, \theta_1, \theta_2) = (\theta_1 - b_2)T\{b_1 \geqslant b_2\}$$
$$u_2(b_1, b_2, \theta_1, \theta_2) = (\theta_2 - b_1)T\{b_1 < b_2\}$$

Player 1 has to maximize

$$= \int_0^1 f(\theta_2|\theta_1)(\theta_1 - s_2(\theta_2))I\{b_1 \geqslant s_2(\theta_2)\}d\theta_2$$
$$= \int_0^1 1 \cdot (\theta_1 - \alpha_2\theta_2)I\{\theta_2 \leq \frac{b_1}{\alpha_2}\}d\theta_2$$
$$= \frac{1}{\alpha_2}(b_1\theta_1 - \frac{\theta_1^2}{2})$$

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

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

Differentiating wrt b_1, we get

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

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

Similarly for 2.
For any independent positive prior, bidding true type is a BE of the induced Bayesian game in Second Price Auction.

भारतीय प्रौद्योगिकी संस्थान मुंबई

Indian Institute of Technology Bombay