Microeconomics IV (Game Theory)

Lecture 6 – Static Games of Incomplete Information

Lasha Chochua

TA: Luka Makharadze

2026

Incomplete Information

Incomplete Information

- A game with *incomplete information* is one where, at the first moment players plan their actions, some players have private information that others do not.
- This private information is not common knowledge at the start of the game.

- Example
  - Suppose Player 1 draws a card and sees it **before** learning how the card will be used.
  - The game now has incomplete information: Player 1 knows something that Player 2 does not.

- Many real-world situations involve players who already have **different private information** at the start.
- It can be unnatural to define the start of the game as a time **before** this information was known.

Deep Private Information

- Some private information may be deeply tied to a player's identity:
  - Gender
  - Language
  - Risk preferences

- It might be **meaningless** to imagine the player acting *before* learning this information.

- Player Types
  - We allow that a player may already hold private information **when** the game begins.
  - This private information is called the **type** of the player.

Representing Incomplete Information

- Games with incomplete information can be modeled in extensive form using a **chance node** that randomly determines players' types.
- This implies that players plan their actions **before** learning their own types.

- But this setup is often awkward, especially when types represent deep traits like preferences or beliefs.
- It can be unrealistic to assume players choose strategies without knowing their type.

- **Harsanyi (1967–68)** proposed a solution: the **Bayesian form** of a game.
- The Bayesian form generalizes the strategic form to allow players to **condition their strategy on their type**.
- It avoids the need to assume players plan before learning private information.
- The Bayesian form is almost as simple as the strategic form, but better suited to games with incomplete information.

Bayesian Games

Bayesian Game: Elements

- A **Bayesian game** consists of:
  - A set of players N
  - For each player $i \in N$:
    - Action set C_i
    - Type set T_i
    - Belief function p_i
    - Utility function u_i

- Define the set of:

$$C = \prod_{i \in N} C_i \quad \text{(action profiles)}$$

$$T = \prod_{i \in N} T_i \quad \text{(type profiles)}$$

Beliefs and Interpretation

- For each player i, define:

$$T_{-i} = \prod_{j \in N \setminus \{i\}} T_j$$

  - Represents all type combinations of players other than i

- The belief function p_i assigns:

$$p_i(\cdot \mid t_i) \in \Delta(T_{-i})$$

  - A probability distribution over T_{-i} given type t_i

- Interpretation:
  - $p_i(t_{-i} \mid t_i)$ is the probability player i assigns to the others having types t_{-i}, if their own type is t_i

Utility and Full Game Structure

- For each player i, the utility function is:

$$u_i : C \times T \to \mathbb{R}$$

  - It assigns a payoff based on the action profile c and type profile t
- The **Bayesian game** is written as:

$$\Gamma^b = (N, (C_i)_{i \in N}, (T_i)_{i \in N}, (p_i)_{i \in N}, (u_i)_{i \in N})$$

- Γ^b is **finite** if N, C_i, and T_i are all finite for every i
- We assume:
  - Each player knows their own type
  - The full structure of the game is **common knowledge**

Strategy vs. Action

- In a Bayesian game, a player chooses an **action** after learning their type
- To avoid confusion:
  - The word **strategy** refers to a full plan for all possible types
  - The word **action** is what a player chooses **once their type is known**
- A **strategy** for player i is a function:

$$s_i : T_i \to C_i$$

- It assigns an action c_i for each type t_i that player i could have

Buyer-Seller Game (Bayesian Example)

- Player 1 is a **seller**, Player 2 is the **only buyer**
- Each player knows the object's value **to himself**, but not to the other
- Each believes the other's valuation is uniform on $\{1, 2, \ldots, 100\}$

- Each player simultaneously submits a **bid** in $\{1, \ldots, 100\}$
- If buyer's bid $\geq$ seller's bid: trade occurs at the **average** of the bids
- Otherwise: **no trade**

- Players are risk-neutral, and utility equals monetary gain from trading

Bayesian Formulation of Buyer-Seller Game

- Set of players: $N = \{1, 2\}$
- Types: $T_i = \{1, 2, \dots, 100\}$ (player i's valuation)
- Actions: $C_i = \{1, \dots, 100\}$ (player i's bid)
- Beliefs: $p_i(t_{-i} \mid t_i) = \frac{1}{100}$ for all t_{-i}

- Utility functions:

$$u_1(c, t) = \begin{cases} (c_1 + c_2)/2 - t_1 & \text{if } c_2 \geq c_1 \\ 0 & \text{otherwise} \end{cases}$$

$$u_2(c, t) = \begin{cases} t_2 - (c_1 + c_2)/2 & \text{if } c_2 \geq c_1 \\ 0 & \text{otherwise} \end{cases}$$

Consistent Beliefs

Consistent Beliefs in Bayesian Games

- Beliefs $p_i(t_{-i} \mid t_i)$ are **consistent** if derived from a common prior P:

$$p_i(t_{-i} \mid t_i) = \frac{P(t)}{\sum_{s_{-i} \in T_{-i}} P(s_{-i}, t_i)}$$

  where $t = (t_i, t_{-i})$

- In the finite buyer-seller game:
    - $T = \{1, \ldots, 100\} \times \{1, \ldots, 100\}$
    - Prior: $P(t) = \frac{1}{10000}$ for all $t \in T$
- In the infinite version:
    - $T_1 = T_2 = [1, 100]$
    - Beliefs are consistent with a **uniform prior** over $[1, 100]^2$
    - For any $1 \leq x \leq y \leq 100$:

$$p_i([x, y] \mid t_i) = \frac{y - x}{100}$$

Bayesian Belief Updating: Example (1/2)

- Players: 1 and 2
- Type sets:
  - Player 1: $\theta_1 \in \{a, b\}$
  - Player 2: $\theta_2 \in \{c, d\}$

- Nature chooses (θ_1, θ_2) according to the joint distribution:

Player 1\Player 2	c	d
a	$\frac{1}{6}$	$\frac{1}{3}$
b	$\frac{1}{3}$	$\frac{1}{6}$

- The table gives $\Pr(\theta_1, \theta_2)$ for each type profile
- This is a **common prior** shared by both players

Bayesian Belief Updating: Example (2/2)

- Suppose player 1 learns that $\theta_1 = a$
- What is player 1's belief about player 2's type?

- Use Bayes' Rule:

$$\phi_1(\theta_2 = c \mid \theta_1 = a) = \frac{\Pr(\theta_1 = a \wedge \theta_2 = c)}{\Pr(\theta_1 = a)} = \frac{\frac{1}{6}}{\frac{1}{6} + \frac{1}{3}} = \frac{1}{3}$$

$$\phi_1(\theta_2 = d \mid \theta_1 = a) = \frac{\frac{1}{3}}{\frac{1}{6} + \frac{1}{3}} = \frac{2}{3}$$

- These are player 1's **posterior beliefs** over player 2's type, conditional on knowing $\theta_1 = a$

Why Are Consistent Beliefs Assumed?

- Most Bayesian games in theory and applications assume **consistent beliefs** based on a common prior over types T
- A common prior simplifies modeling:
  - All belief functions $p_i(\cdot \mid t_i)$ are derived from a single probability distribution P
  - This is easier than specifying multiple subjective probability functions
- Consistency feels natural:
  - Differences in beliefs are explained by differences in **information**
  - In contrast, **inconsistent beliefs** imply differences in opinion not based on data
- Still, inconsistent beliefs are possible:
  - Example: two sports coaches each believe their own team has a $2/3$ chance of winning
  - These beliefs **cannot** come from a common prior
  - See: Aumann (1976) on agreeing to disagree

When Are Two Bayesian Games Equivalent?

- Two Bayesian games are **fully equivalent** if they imply the same decision-making behavior for every type of each player
- Formally, Γ^b and $\Gamma^{b'}$ are equivalent if:

$$q(t_{-i} \mid t_i)\, w(c,t) = A_i(t_i)\, p(t_{-i} \mid t_i)\, u_i(c,t) + B_i(t)$$

  - $A_i(t_i) > 0$ for all t_i
  - $B_i(t)$ can vary with full type profile t
- This implies same incentives and beliefs up to affine transformations

Type-Agent Representation

Type-Agent Representation

- Harsanyi (1967–68), building on Selten's suggestion, showed how to represent a Bayesian game Γ^b as a strategic-form game
- This is called the **type-agent representation** (also known as the **Selten game** or **posterior-lottery model**)
- Each type $t_i \in T_i$ is treated as a separate player (or agent)
- Assume type sets T_i are disjoint: $T_i \cap T_j = \emptyset$ for $i \neq j$
- The new player set is:

$$T^* = \bigcup_{i \in N} T_i$$

Strategies and Payoffs in Type-Agent Game

- Agent $t_i \in T_i$ chooses an action from C_i
- Strategy set for t_i: $D_{t_i} = C_i$
- Agent t_i selects the action that player i would use in Γ^b if their type is t_i

- Utility for agent t_i is the expected utility for player i in Γ^b, given type t_i
- For any strategy profile $d = (d(s))_{s \in T^*}$:

$$v_{t_i}(d) = \sum_{t_{-i} \in T_{-i}} p_i(t_{-i} \mid t_i)\, u_i\left((d(t_j))_{j \in N},\, (t_j)_{j \in N}\right)$$

- The result is a standard strategic-form game equivalent to the original Bayesian game

Bayesian Equilibrium

Bayesian Equilibrium: Concept

- A **Bayesian equilibrium** is a Nash equilibrium of the **type-agent representation** of a Bayesian game
- Each type of each player chooses a strategy to **maximize expected utility**:
  - The type is known to the player
  - The types of others are unknown
- A strategy in Bayesian equilibrium:
  - **Can** depend on a player's own type
  - **Cannot** depend on other players' types
- We must define behavior for **every type**, not just the actual ones
  - Otherwise, we can't evaluate expected utility under uncertainty

Bayesian Equilibrium: Formal Setup

- Let $\Gamma^b = (N, (C_i)_{i\in N}, (T_i)_{i\in N}, (p_i)_{i\in N}, (u_i)_{i\in N})$ be a Bayesian game
- A **randomized strategy profile** is:

$$\sigma = (\sigma_i(c_i \mid t_i))_{c_i\in C_i,\, t_i\in T_i,\, i\in N}$$

- Each $\sigma_i(c_i \mid t_i)$ is the probability that type t_i of player i chooses action c_i
- Conditions:

$$\sigma_i(c_i \mid t_i) \geq 0, \quad \sum_{c_i\in C_i} \sigma_i(c_i \mid t_i) = 1 \quad \text{for all } t_i \in T_i,\ i \in N$$

Bayesian Nash Equilibrium

- A **Bayesian equilibrium** is a randomized strategy profile σ such that each type t_i of each player maximizes its expected utility

- Let $\sigma_i(\cdot \mid t_i)$ denote the strategy of type t_i:

$$\sigma_i(\cdot \mid t_i) = (\sigma_i(c_i \mid t_i))_{c_i \in C_i}$$

- Then σ is a Bayesian equilibrium if, for all $i \in N$, $t_i \in T_i$:

$$\sigma_i(\cdot \mid t_i) \in \arg \max_{\tau_i \in \Delta(C_i)} \sum_{t_{-i} \in T_{-i}} p_i(t_{-i} \mid t_i) \sum_{c \in C} \left(\prod_{j \in N \setminus \{i\}} \sigma_j(c_j \mid t_j) \right) \tau_i(c_i) u_i(c, t)$$

- Each type chooses a **best response** to the strategy profile of others, **given its belief**

Examples of Bayesian Equilibrium

Bayesian Equilibrium: Example I (1/3)

- Two players:
  - $C_1 = \{x_1, y_1\}$, $T_1 = \{1.0\}$ (Player 1 has no private info)
  - $C_2 = \{x_2, y_2\}$, $T_2 = \{2.1, 2.2\}$
- Beliefs of Player 1:
  - $p_1(2.1 \mid 1.0) = 0.6$
  - $p_1(2.2 \mid 1.0) = 0.4$

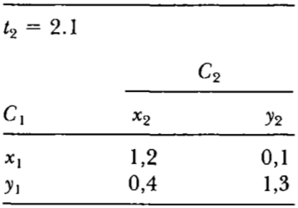

$t_2 = 2.1$

	C_2	
C_1	x_2	y_2
x_1	1,2	0,1
y_1	0,4	1,3

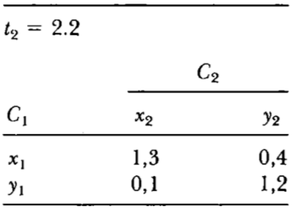

$t_2 = 2.2$

	C_2	
C_1	x_2	y_2
x_1	1,3	0,4
y_1	0,1	1,2

Bayesian Equilibrium: Example I (2/3)

- Strong dominance:
  - y_2 is dominated for type 2.1
  - x_2 is dominated for type 2.2
- Thus, Bayesian equilibrium must be:

$$\sigma_1(\cdot \mid 1.0) = [x_1], \quad \sigma_2(\cdot \mid 2.1) = [x_2], \quad \sigma_2(\cdot \mid 2.2) = [y_2]$$

- Player 1 best responds to beliefs by choosing x_1

Incomplete vs. Complete Information

- If player 1 **knew** player 2's type:
  - If 2.1: he would choose x_1
  - If 2.2: he would choose y_1
- But in Bayesian game, player 1 doesn't know 2's type

- So the prediction:
  "(x_1, x_2) *if* 2.1, (y_1, y_2) *if* 2.2" *is* ***not implementable*** *without communication*

- Player 2 has **no incentive** to reveal her type:

  - She prefers (y_1, y_2) if 2.1
  - She prefers (x_1, x_2) if 2.2

- Hence, **manipulating** communication could occur

- This shows the **danger** of analyzing each matrix separately as if the game had complete information

Example II : Game Setup

- Players: $N = \{\mathsf{I}, \mathsf{II}\}$
- Type sets:
  - Player I: $T_I = \{\mathsf{I}_1, \mathsf{I}_2\}$
  - Player II: $T_{II} = \{\mathsf{II}\}$ (one type)
- Beliefs:
  - $\Pr(\mathsf{I}_1, \mathsf{II}) = \Pr(\mathsf{I}_2, \mathsf{II}) = \frac{1}{2}$
  - Types of Player I are equally likely
- Each player has **two actions**, depending on the state:
  - Player I plays either $\{T_1, B_1\}$ or $\{T_2, B_2\}$
  - Player II always plays from $\{L, R\}$

State Games

Type I$_1$

	L	R
T_1	(1, 0)	(0, 2)
B_1	(0, 3)	(1, 0)

Type I$_2$

	L	R
T_2	(0, 2)	(1, 1)
B_2	(1, 0)	(0, 2)

Strategy Notation and Payoff Matrices

- Player I has types I_1 and I_2
- Player II mixes: q on L, $1-q$ on R
- Strategies:
  - Type I_1: $x(T_1)$, $(1-x)(B_1)$
  - Type I_2: $y(T_2)$, $(1-y)(B_2)$
  - Player II: $q(L)$, $(1-q)(R)$

Matrix for type I_1

$$\begin{array}{c|cc} & q & 1-q \\ \hline x & (1,0) & (0,2) \\ 1-x & (0,3) & (1,0) \end{array}$$

Matrix for type I_2

$$\begin{array}{c|cc} & q & 1-q \\ \hline y & (0,2) & (1,1) \\ 1-y & (1,0) & (0,2) \end{array}$$

Player II Must Mix

- If $q = 1$, then type I$_1$'s best reply is T_1 $(x = 1)$ and type I$_2$'s best reply is B_2 $(y = 0)$ $\Rightarrow$ Player II prefers R, so $q = 1$ cannot be an equilibrium
- If $q = 0$, then type I$_1$'s best reply is B_1 $(x = 0)$ and type I$_2$'s best reply is T_2 $(y = 1)$ $\Rightarrow$ Player II prefers L, so $q = 0$ cannot be an equilibrium

- Therefore, q must be strictly between 0 and 1, and Player II must be indifferent between L and R

Indifference Condition

- Player II's expected payoff from L is:

$$\frac{1}{2} \cdot 3(1-x) + \frac{1}{2} \cdot 2y$$

- Player II's expected payoff from R is:

$$\frac{1}{2} \cdot 2x + \frac{1}{2}\left(y + 2(1-y)\right)$$

- Equating both sides gives:

$$\frac{3(1-x)}{2} + y = x + \frac{2-y}{2}$$

- Simplify to:

$$x = \frac{1+3y}{5} \tag{1}$$

Consistency with Best Responses

- For (x, y) to form part of a Bayesian equilibrium, it must satisfy both Equation (1) and the best response conditions:
- If $q < \frac{1}{2}$, then Player I's best reply is $x = 0$, $y = 1$ which violates (1)
- If $q = \frac{1}{2}$, then both types of Player I are indifferent and any (x, y) satisfying (1) is a best response
- If $q > \frac{1}{2}$, then Player I's best reply is $x = 1$, $y = 0$ which also violates (1)
- Thus, Bayesian equilibrium requires $q = \frac{1}{2}$ and Equation (1)

Continuum of Equilibria and Player II's Payoff

- From $x = \frac{1+3y}{5}$ and $y \in [0, 1]$, we get:

$$\frac{1}{5} \leq x \leq \frac{4}{5}, \quad 0 \leq y \leq 1$$

- This defines a continuum of Bayesian equilibria (x, y, q) with $q = \frac{1}{2}$ and (x, y) satisfying Equation (1)
- Player II's expected payoff is:

$$\frac{1}{2} \cdot 3(1 - x) + \frac{1}{2} \cdot 2y = \frac{12 + y}{10}$$

- Player I's expected payoff (either type) is always $\frac{1}{2}$

Required Reading

- **Myerson (1997).** *Game Theory: Analysis of Conflict.* Harvard University Press.
    - Chapter 2 – 2.8
    - Chapter 3 – 3.9