

# Microeconomics IV (Game Theory)

## Lecture 2 – Basic Models of Game Theory

Lasha Chochua

TA: Luka Makharadze

2026

# Last Week

# The Big Picture

- **Lecture 1 Goal:** show that any rational, intelligent agent behaves as if maximizing expected utility
- The argument has three steps:

Step	Content
1	Define the <b>object of choice</b> : lotteries $f : \Omega \rightarrow \Delta(X)$
2	Impose <b>8 consistency axioms</b> on preferences over lotteries
3	Show preferences satisfying the axioms $\iff$ <b>SEU maximization</b>

- Two outputs of the theorem: a **utility function**  $u(x, t)$  and **subjective beliefs**  $p(t | S)$
- Both are elicited from observable preferences – not assumed

## Lotteries – The Object of Choice

- $X$  = finite set of **prizes**;  $\Omega$  = finite set of **states of the world**
- Two sources of uncertainty:
  - **Risk (objective)**: probabilities known – coin tosses, roulette wheels
  - **Uncertainty (subjective)**: probabilities unknown – competitor's strategy, market demand

### Definition: Lottery

A **lottery** is any function  $f : \Omega \rightarrow \Delta(X)$ , where  $f(x | t)$  is the objective probability of prize  $x$  given state  $t$ .

- Nature picks state  $t$  (subjective uncertainty); then  $f(\cdot | t)$  determines the prize by objective randomization
- The agent chooses  $f$  **before** either is resolved
- This compound structure lets the theorem separately identify  $u(x, t)$  and  $p(t | S)$

# The Eight Axioms

Axiom	Name	What it requires
1A	Completeness	Any two lotteries can be ranked
1B	Transitivity	Rankings are consistent across comparisons
2	Relevance	Only outcomes inside $S$ matter given $S$
3	Monotonicity	More probability on the better option is better
4	Continuity	No lottery is infinitely better than another
5A/B	Objective Substitution	Mixing with the same lottery preserves rankings
6A/B	Subjective Substitution	Prefer $f$ over $g$ in $S$ and in $T \Rightarrow$ prefer $f$ in $S \cup T$
7	Interest	Something is at stake in every state
8	State Neutrality	(Optional) Same prize valued equally across states

- Axioms 5 and 6 are the deepest – violated by the **Allais paradox** (Axiom 5B) and the **Ellsberg paradox** (Axiom 6)

# Subjective Expected Utility

## Subjective Expected Utility

Axioms 1AB, 2, 3, 4, 5AB, 6AB, 7 are jointly satisfied **if and only if**  $\exists$  utility  $u : X \times \Omega \rightarrow \mathbb{R}$  and conditional probabilities  $p : \Xi \rightarrow \Delta(\Omega)$  such that:

$$\max_x u(x, t) = 1, \quad \min_x u(x, t) = 0 \quad \forall t \in \Omega \quad (1)$$

$$p(R | T) = p(R | S) \cdot p(S | T) \quad \forall R \subseteq S \subseteq T \subseteq \Omega \quad (2)$$

$$f \succsim_S g \iff \sum_{t \in S} p(t | S) \sum_{x \in X} u(x, t) f(x | t) \geq \sum_{t \in S} p(t | S) \sum_{x \in X} u(x, t) g(x | t) \quad (3)$$

Adding Axiom 8  $\iff$  (1)–(3) hold with state-independent  $u(x, t) = U(x)$ .

- (1): utility normalized to  $[0, 1]$  in every state
- (2): beliefs satisfy Bayes's formula (chain rule)
- (3): agent always chooses the lottery with higher **expected utility**

# Introduction

# Modeling Games

- Every game begins with a **model** that describes the situation
- A model must not be too simple or too complex:
  - Too simple: misses key features of the game
  - Too complex: hides the essential logic
- The goal is to capture the *minimal* structure needed to identify the strategic problem

## Two Forms of Representation

- Two key ways to model games:
  - **Extensive form:** detailed tree of actions and information
  - **Strategic (normal) form:** simplified table of payoffs
- Extensive form is richer and captures timing and information explicitly
- Strategic form abstracts away timing – it is a planning-stage representation

### **i** Note

Both forms are exact – neither loses information about rational behavior, under the right conditions.

# Extensive-Form Games

# Origins of the Extensive Form

- Standard definition of extensive form is due to Kuhn (1953)
- Built on von Neumann and Morgenstern (1944)
- Strategic and Bayesian forms are simpler but **derived** from the extensive form
  - Expected payoffs are computed by summing over terminal nodes of the tree

## A Simple Card Game

- Two players: *Player 1* and *Player 2*
- Each contributes \$1 to a pot
- Player 1 draws a card (red or black) and decides:
  - **Raise** (add \$1 more), or
  - **Fold** (end the game)
- If Player 1 raises, Player 2 decides:
  - **Meet** (add \$1 more), or
  - **Pass**

# Game Outcomes

- If Player 1 **folds**:
  - Player 1 wins if red, loses if black
- If Player 2 **passes**:
  - Player 1 wins
- If Player 2 **meets**:
  - Player 1 shows the card
  - Red: Player 1 wins, gets \$2
  - Black: Player 2 wins, gets \$2

# Tree Representation

- The game can be shown as a **tree diagram**:
  - **Nodes**: decision points
  - **Branches**: possible actions
  - **Root**: beginning of the game
  - **Terminal nodes**: outcomes with payoffs
- Payoffs shown as pairs  $(x_1, x_2)$  for Players 1 and 2
- Example:  $(2, -2)$  means Player 1 wins 2, Player 2 loses 2

# Tree Diagram Example




Figure 1: Game 1

- Payoffs shown as pairs  $(x_1, x_2)$  for Players 1 and 2
- For example:
  - $(2, -2)$  means Player 1 wins 2, Player 2 loses 2

# Chance and Decision Nodes

- Nodes labeled “0” are **chance nodes** (e.g. drawing a card)
- Others (e.g. “1”, “2”) are **decision nodes** controlled by players
- Probability of red or black is 0.5 each

# Incomplete Information

- Player 1 sees the card color; Player 2 does not
- In the tree, Player 2 must act without knowing the color
- She must choose the same action at both of her nodes

## **i** Note

This is the key departure from perfect-information games: a player may not know exactly where in the tree she is when she moves.

## Information Sets – The Core Idea

- When a player cannot distinguish between two nodes, we say they belong to the **same information set**
- An information set is a group of nodes that are **indistinguishable** to the player at the time of the move
- The player must choose a single action that applies to all nodes in the set
- **Why does this matter?** It is what forces a player to act the same way in situations she cannot tell apart
- The information set is the game-theoretic formalization of *what a player knows*

# Information Sets

- We label nodes with:
  - **Player number**
  - **Information state**
- Example: “2.0” means Player 2, in information state 0
- Information state indicates what a player knows at that node

# Improved Tree with Information




Figure 2: Game 2

- Dashed lines show **information sets**
- Player 2 cannot distinguish between red and black branches

# Move Labels and Player Knowledge

- We keep branch labels for player decisions, but not for chance events
- Move labels matter because players choose actions, not specific branches
- A player must know all her options at any decision point

# Information Sets and Choices

- Player 2 must make the same decision at both of her nodes
- She does not know if Player 1 drew red or black
- So both of her nodes must have the same set of move labels: “Meet” and “Pass”

## Ensuring Meaningful Play

- If we added a third branch only after one of Player 2's nodes, she could exploit hidden information
- To be valid, both nodes must offer the same moves if in the same information set

### **i** Note

The consistency requirement on move labels is not just technical – it encodes the idea that a player cannot act on information she does not have.

## Player 1's Information

- Player 1 knows the card's color
- So he can have different move labels at different nodes
- Player 2 cannot, because she cannot tell nodes apart

## Importance of Move Labels

- Move labels shape the player's choices
- In our game, "Pass" is better for Player 2 if Player 1 has red
- "Meet" is better if Player 1 has black
- But Player 2 does not know the color when choosing

# What is a Path?

- A **path** is a sequence of connected branches in a graph
- Formally, a path is:

$$\{\{x_1, x_2\}, \{x_2, x_3\}, \dots, \{x_{m-1}, x_m\}\} = \{\{x_k, x_{k+1}\}\} \quad \text{for } k = 1, \dots, m - 1$$

- Here:
  - $m \geq 2$
  - Each  $x_k$  is a different node

# Basic Graph Terminology

- A **graph** is a set of:
  - **Nodes** (points)
  - **Branches** (connections between nodes)
- A **branch** connects two nodes

# What is a Tree?

- A **tree** is a graph with a special structure:
  - Between any two nodes, there is **exactly one** path
- This makes it easy to trace the flow of a game or process

# Rooted Tree

- A **rooted tree** has a designated starting node called the **root**
- In diagrams, the root is usually on the **left**
- All play begins at the root

# Paths in Rooted Trees

- The **path to a node** is the **unique sequence of branches** connecting it to the root
- This structure lets us track the history of the game step by step

# Alternatives and Branching

- An **alternative** is any branch **not** on the path to a node
- For example, if the current path is  $\{x_1, x_2, x_3\}$ , then a branch from  $x_2$  to another node is an alternative

# Follows and Immediate Follows

- A branch or node  $x$  **follows**  $y$  if  $y$  is **in the path** to  $x$
- $x$  **immediately follows**  $y$  if:
  - $x$  directly follows  $y$ , and
  - There is a branch  $\{x, y\}$

# Terminal Nodes

- A **terminal node** is a node with **no further branches**
- It represents an **end point** in the game
- At terminal nodes, we specify the **payoffs** for players

## Definition: $n$ -Person Extensive-Form Game (1/3)

An  $n$ -person extensive-form game  $\Gamma^e$  consists of a rooted tree with labeled nodes and branches. The following conditions must be satisfied:

- 1 **Player Labels** Each nonterminal node is labeled with a number in  $\{0, 1, \dots, n\}$ 
  - Label 0: **chance node**, controlled by nature
  - Label  $i \in \{1, \dots, n\}$ : **decision node**, controlled by player  $i$
- 2 **Chance Probabilities** Every branch from a chance node has a nonnegative probability label summing to 1:

$$\sum_{j=1}^k p_j = 1, \quad p_j \geq 0$$

## Definition: $n$ -Person Extensive-Form Game (2/3)

- ③ **Information States** Every node controlled by a player is assigned an **information state**
- A player knows only the information state at the time of the move
  - Nodes are indistinguishable if the player cannot tell which was reached

Notation: node labeled  $i.k$  means player  $i$ , information state  $k$

- ④ **Move Labels** Each alternative (branch) at a node has a **move label**
- If two nodes share the same player and information label, the set of move labels must be **identical**
  - This ensures players know their available actions even without knowing the exact node

## Definition: $n$ -Person Extensive-Form Game (3/3)

- 5 **Terminal Node Payoffs** Each terminal node is labeled with a vector of utilities:

$$(u_1, u_2, \dots, u_n)$$

- $u_i$ : utility (payoff) to player  $i$

6 **Perfect Recall**

Players remember everything they knew earlier:

- All previous moves they made
- All past information states they experienced
- If nodes  $y$  and  $z$  are indistinguishable and a decision was made at  $y$ , a matching indistinguishable move must exist at  $z$
- This avoids contradictions in player memory and ensures coherent strategy

# Perfect Information and Strategies (1/2)

- A game has **perfect information** if:
  - No two nodes share the same information state
  - The player always knows:
    - All previous moves by others
    - Their own past moves
    - What node they are at
- In this case, the player knows exactly what has happened before making any move

## Perfect Information and Strategies (2/2)

### Strategy in Extensive-Form Games

A **strategy** is a complete plan specifying what a player will do at **every** information state – including states that may never be reached in actual play. Formally, let  $S_i$  be the set of information states for player  $i$  and  $D_s$  the set of moves at state  $s$ . The set of strategies for player  $i$  is:

$$\prod_{s \in S_i} D_s$$

# Why Must a Strategy Cover Unreachable States?

- A strategy is a **plan made before the game begins** – a complete contingent decision rule
- Even if an information state is never reached, the player must decide *in advance* what they would do there
- **Why?** Because other players' decisions depend on their *beliefs* about what you would do at every node
  - Player 2 chooses whether to meet partly based on what would happen if she does
  - That belief is only well-defined if Player 1 has a plan at every information state
- **Analogy:** a chess player preparing for a tournament writes down what they would do in every position – including positions they may never reach – because the opponent's strategy depends on it

## Why Must a Strategy Cover Unreachable States? (Cont.)

- A second reason: **consistency in equilibrium analysis**
  - When we check whether a strategy profile is an equilibrium, we ask: does any player want to deviate?
  - To answer this, we must evaluate payoffs from deviating – which requires knowing what happens at information states that are **off the equilibrium path**
  - Without a complete plan, we cannot check whether a player prefers to deviate

### **i** Note

This is why “I would never reach that node anyway” is not a valid reason to leave a strategy undefined. The definition of equilibrium forces us to evaluate the full plan.

## Example: Strategies in the Card Game

- Player 1 has 2 information states:
  - 1.*a*: sees red card
  - 1.*b*: sees black card
- Each strategy is a pair of choices at 1.*a* and 1.*b*:
  - {Rr, Rf, Fr, Ff}

### **i** Note

- First letter: move at 1.*a* (capital for red card)
- Second letter: move at 1.*b* (lowercase for black card)
- Rf = “Raise if red, Fold if black”
- Rr = “Raise regardless of card color”

## Player 2's Strategy

- Player 2 has only 1 information state
- So she has just 2 strategies:
  - M = "Meet"
  - P = "Pass"
- Even though only one path is realized, a strategy covers **every possible contingency**

## Examples




Figure 3: Game 3




Figure 4: Game 4

## Game 3 – Limited Information

- Player 2 must choose between **L** and **R**
- She **does not observe** Player 1's move (T or B)
- Player 1 anticipates this and will choose **T**, since:
  - Against either L or R, T gives a better payoff than B
- Example outcome:
  - If Player 1 chooses T, and Player 2 chooses L, Player 2 gets payoff 2

## Game 4 – Full Information

- Player 2's two decision nodes now have **different information labels**
- She observes whether Player 1 chose T or B
- This changes Player 2's strategy set to:

$$\{Ll, Lr, Rl, Rr\}$$

- Player 2 can now condition her action on Player 1's move

# Strategic Implications

- In Game 3:
  - Player 2's strategies:  $\{L, R\}$
- In Game 4:
  - Player 2's strategies:  $\{Ll, Lr, Rl, Rr\}$  (First letter = move if T, second = move if B)
- If Player 2 is rational:
  - She will follow **Lr**: choose L if T, r if B
  - Player 1 anticipates this and chooses **B**, getting payoff 3

## Strategic Reasoning in Game 4

- Suppose Player 1 plays B
- Player 2 will respond with **r** in both:
  - Strategy Rr (“R if T, r if B”)
  - Strategy Lr (“L if T, r if B”)
- Why should Player 1 choose B?
  - Because Player 2 prefers L if T and r if B  $\Rightarrow$  strategy Lr
  - Then Player 1 expects payoff:
    - 2 from T
    - 3 from B  $\Rightarrow$  **chooses B**

## Information Structure and Equilibrium

- The change from Game 3 to Game 4 increases Player 2's strategy space
- This can:
  - Change Player 1's optimal move
  - Lower Player 2's expected payoff (despite having more info)
- If Player 2 used strategy  $R_r$ :
  - Player 1 would choose T and get payoff 4

### **i** Note

More information does not always help – it can change the equilibrium in ways that harm the better-informed player.

# Strategic Form

# Strategic Form Representation

- The **strategic form** (or normal form) is a simpler way to represent a game
- It specifies:
  - A set of players  $N$
  - Each player's strategy set  $C_i$
  - A utility function  $u_i$  for each player

$$\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$$

- $u_i : \prod_{j \in N} C_j \rightarrow \mathbb{R}$  gives player  $i$ 's payoff from any strategy profile

# Strategy Profiles and Timing

- A **strategy profile**  $c = (c_1, \dots, c_n)$  is a combination of strategies, one for each player
- All players choose their strategies **simultaneously**
- The strategic form is a **static model**:
  - It abstracts from the timing of moves
  - It only captures the strategic planning stage

## From Extensive to Strategic Form

- The strategic form can be derived from an extensive-form game
- Example: the card game in Game 2

- Player 1's strategy set:

$$C_1 = \{Rr, Rf, Fr, Ff\}$$

- Player 2's strategy set:

$$C_2 = \{M, P\}$$

- Players choose these strategies **before the card is revealed**

## Example: Strategy Profile (Rf, M)

- Strategy Rf: Raise if red, Fold if black
- Strategy M: Player 2 meets if Player 1 raises

- Expected payoffs:

$$u_1(\text{Rf}, \text{M}) = 2 \cdot \frac{1}{2} + (-1) \cdot \frac{1}{2} = 0.5$$

$$u_2(\text{Rf}, \text{M}) = -2 \cdot \frac{1}{2} + 1 \cdot \frac{1}{2} = -0.5$$

- These are **expected utilities** given uncertainty about the card draw

# Game 2 in Normal Form (1/2)




Figure 5: Game 2

## Game 2 in Normal Form (2/2)

$C_1$	$C_2$	
	M	P
Rr	0,0	1,-1
Rf	0.5,-0.5	0,0
Fr	-0.5,0.5	1,-1
Ff	0,0	0,0

**Figure 6:** Game 2 in a Normal Form

# Mapping Extensive Form to Strategic Form (1/3)

- Let  $\Gamma^e$  be an extensive-form game
- To construct its strategic form  $\Gamma$ , define:
  - $N$ : set of players
  - $C_i$ : strategy set for player  $i$
  - $c_i \in C_i$ : function assigning a move  $c_i(r)$  at each information state  $r$
  - $c = (c_1, \dots, c_n)$ : a **strategy profile**

## Mapping Extensive Form to Strategic Form (2/3)

- We define  $P(x | c)$  as the probability that node  $x$  is reached given strategy profile  $c$
- Rules:
  - If  $x$  is the root:  $P(x | c) = 1$
  - If  $x$  follows a chance node  $y$ :

$$P(x | c) = q \cdot P(y | c) \quad \text{where } q \text{ is the chance probability}$$

- If  $x$  follows a decision node  $y$  for player  $i$  in state  $r$ :
  - $P(x | c) = P(y | c)$  if  $c_i(r)$  selects move from  $y$  to  $x$
  - $P(x | c) = 0$  otherwise

## Mapping Extensive Form to Strategic Form (3/3)

- Let  $\Omega^*$  be the set of all **terminal nodes** of  $\Gamma^e$
- Let  $w_i(x)$  be the utility for player  $i$  at terminal node  $x$
- Then for any strategy profile  $c$ , the expected payoff for player  $i$  is:

$$u_i(c) = \sum_{x \in \Omega^*} P(x | c) \cdot w_i(x) \quad (4)$$

- This gives the **normal representation**:

$$\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$$

- Each player's payoff depends on the **entire strategy profile**  $c$
- All probabilistic uncertainty is resolved by computing  $P(x | c)$  over the tree

# Convert Both Games to their Normal Forms




Figure 7: Game 3




Figure 8: Game 4

# Equivalence of Strategic-Form Games

- Utility numbers represent **preferences**, not cardinal values
- Two games are equivalent if they reflect **the same preferences**
- Utility functions  $u(\cdot)$  and  $\hat{u}(\cdot)$  are **equivalent** iff:

$$\hat{u}(c) = Au(c) + B, \quad A > 0$$

- Two strategic-form games  $\Gamma$  and  $\hat{\Gamma}$  are **fully equivalent** iff:

$$\hat{u}_i(c) = A_i u_i(c) + B_i, \quad \forall c \in C, \forall i \in N$$

# Implication of Full Equivalence

- Solution concepts must give the **same prediction** for equivalent games
- Each player  $i$  prefers strategy distribution  $\mu$  over  $\lambda$  if:

$$\sum_{c \in C} \mu(c) u_i(c) \geq \sum_{c \in C} \lambda(c) u_i(c)$$

- In the transformed game  $\hat{\Gamma}$ , same preference holds if:

$$\sum_{c \in C} \mu(c) \hat{u}_i(c) \geq \sum_{c \in C} \lambda(c) \hat{u}_i(c)$$

## Best-Response Equivalence (1/3)

- A weaker notion: players may have **same best responses** even if utilities differ
- Let  $C_{-i}$  be strategy combinations of all players **except**  $i$ :

$$C_{-i} = \prod_{j \in N \setminus \{i\}} C_j$$

- Given belief  $\eta \in \Delta(C_{-i})$ , player  $i$  chooses:

$$\arg \max_{d_i \in C_i} \sum_{e_{-i} \in C_{-i}} \eta(e_{-i}) u_i(e_{-i}, d_i)$$

## Probability Distributions over Strategies (2/3)

- The set  $\Delta(C_{-i})$  denotes all probability distributions over  $C_{-i}$ :

$$\Delta(C_{-i}) = \left\{ \eta : C_{-i} \rightarrow [0, 1] \mid \sum_{e_{-i} \in C_{-i}} \eta(e_{-i}) = 1 \right\}$$

- Each  $\eta \in \Delta(C_{-i})$  assigns a probability to every possible combination of strategies chosen by players other than  $i$
- The probabilities must be nonneg and sum to 1

## Best-Response Equivalence (3/3)

- In  $\hat{\Gamma}$ , same rule applies with  $\hat{u}_i$ :

$$\arg \max_{d_i \in C_i} \sum_{e_{-i} \in C_{-i}} \eta(e_{-i}) \hat{u}_i(e_{-i}, d_i)$$

- $\Gamma$  and  $\hat{\Gamma}$  are **best-response equivalent** iff:

$$\arg \max_{d_i \in C_i} \sum_{e_{-i}} \eta(e_{-i}) u_i(e_{-i}, d_i) = \arg \max_{d_i \in C_i} \sum_{e_{-i}} \eta(e_{-i}) \hat{u}_i(e_{-i}, d_i)$$

for all  $i \in N$  and all  $\eta \in \Delta(C_{-i})$

## Reduced Normal Representations

- Some strategic-form games can be simplified by removing **redundant strategies**
- Example: In an extensive-form game, if a node is never reached, actions taken at that node don't affect payoffs
- Such strategies can be **grouped or eliminated**

# Payoff Equivalence

- Given game  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$ :
- Two strategies  $d_i, e_i \in C_i$  are **payoff equivalent** iff:

$$u_j(c_{-i}, d_i) = u_j(c_{-i}, e_i), \quad \forall c_{-i} \in C_{-i}, \forall j \in N$$

- That is, no player is ever affected by whether  $i$  plays  $d_i$  or  $e_i$
- A **purely reduced normal representation** merges all payoff-equivalent strategies into one

## Random Redundancy

- A **randomized strategy**  $\sigma_i \in \Delta(C_i)$  is a probability distribution over  $C_i$
- A **pure strategy**  $c_i \in C_i$  is **randomly redundant** if:
  - There exists  $\sigma_i \in \Delta(C_i)$  such that  $\sigma_i(c_i) = 0$  and:

$$u_j(c_{-i}, c_i) = \sum_{e_i \in C_i} \sigma_i(e_i) u_j(c_{-i}, e_i), \quad \forall c_{-i} \in C_{-i}, \forall j \in N$$

- That is,  $c_i$  can be replaced by a **randomized mixture** without affecting any player's payoff
- If  $\sigma_i(c_i) > 0$  were allowed, the condition would be trivial:
  - The degenerate mixture  $\sigma_i(c_i) = 1$  always replicates  $c_i$  perfectly
- The requirement  $\sigma_i(c_i) = 0$  forces replication by **genuinely other strategies**

## Example (1/3)




Figure 9: Game 5

## Example (2/3)

$C_1$	$C_2$	
	$x_2$	$y_2$
$a_1x_1$	6,0	6,0
$a_1y_1$	6,0	6,0
$a_1z_1$	6,0	6,0
$b_1x_1$	8,0	0,8
$b_1y_1$	0,8	8,0
$b_1z_1$	3,4	7,0

Figure 10: Game 5

## Example (3/3)

$C_1$	$C_2$	
	$x_2$	$y_2$
$a_1$	6,0	6,0
$b_1x_1$	8,0	0,8
$b_1y_1$	0,8	8,0
$b_1z_1$	3,4	7,0

Figure 11: Game 5

# Strong Domination

- In a strategic-form game  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$ :
  - A strategy  $d_i \in C_i$  is **strongly dominated** for player  $i$  if:

$$u_i(c_{-i}, d_i) < \sum_{e_i \in C_i} \sigma_i(e_i) u_i(c_{-i}, e_i), \quad \forall c_{-i} \in C_{-i}$$

for **some** randomized strategy  $\sigma_i \in \Delta(C_i)$

- **A strongly dominated strategy can never be a best response under any belief**

# Iterated Elimination of Strongly Dominated Strategies

- Let  $C_i^{(1)}$  be the subset of strategies in  $C_i$  not strongly dominated
- Let  $\Gamma^{(1)}$  be the game with  $(C_i^{(1)})_{i \in N}$
- Define inductively:

$$\Gamma^{(k)} = \left( N, (C_i^{(k)})_{i \in N}, (u_i)_{i \in N} \right)$$

where  $C_i^{(k)}$  is the set of strategies in  $C_i^{(k-1)}$  that are not strongly dominated

- There exists a smallest  $K$  such that:

$$C_i^{(K)} = C_i^{(K+1)} = \dots \quad \text{for all } i \in N$$

# Iteratively Undominated Strategies

- Let  $C_i^{(\infty)} = C_i^{(K)}$  and  $\Gamma^{(\infty)} = \Gamma^{(K)}$
- Then  $C_i^{(\infty)}$  contains all strategies that are **iteratively undominated** in the strong sense
- Every strategy in  $C_i^{(\infty)}$  is a best response to some belief over  $C_{-i}^{(\infty)}$
- The set  $(C_i^{(\infty)})_{i \in N}$  is maximal under this condition

## Weak Domination

- A strategy  $d_i \in C_i$  is **weakly dominated** for player  $i$  if:

$$u_i(c_{-i}, d_i) \leq \sum_{e_i \in C_i} \sigma_i(e_i) u_i(c_{-i}, e_i) \quad \forall c_{-i} \in C_{-i}$$

and the inequality is strict for at least one  $c_{-i} \in C_{-i}$

- Weak domination allows for a strategy to be a best response for some beliefs
- Iterative elimination of weakly dominated strategies is **order dependent** and may alter outcomes

### **i** Note

Strong domination: the strategy is *never* a best response. Weak domination: it is sometimes a best response but never strictly better than the dominating strategy. The distinction matters for solution concepts.

# Multiagent Representation

## Why Do We Need a Multiagent Representation?

- The strategic form bundles a player's entire plan into a single strategy
- This works well at the planning stage – but it creates a problem for **sequential rationality**
  
- Suppose a player reaches an information state that was supposed to have zero probability under the equilibrium
- Her strategy was formed before the game started, but now she is *actually at this node*
- Should she follow her original plan, or reconsider?
  
- The strategic form cannot answer this – it has no way to separate decisions at different nodes
- The **multiagent representation** separates decisions by information state, making sequential rationality checkable at each node independently

## Multiagent Representation – Key Idea

- Instead of treating each player as a single decision-maker, we treat each **information state** as a separate agent
- These agents are called **temporary agents** – they share the same utility function as the original player, but act independently at their own node
- This matters because:
  - It allows us to ask whether behavior is optimal **at each node separately**
  - It is the natural foundation for **equilibrium refinements** such as sequential equilibrium and trembling-hand perfection
  - It makes off-path behavior explicit – the trembles of a player at one node are **independent** of trembles at another
- Think of each information state as a different “self” of the player, each making an independent decision but all sharing the same payoff function.

# Multiagent Representation – Definition

- The **multiagent representation** maps an extensive-form game into a strategic-form game
- Each **information state** of a player is treated as a separate player (agent)
- These agents, called **temporary agents**, share the same preferences as the original player

## Example (1/3)




Figure 12: Game 6

## Example (2/3)

- Normal Form

$C_1$	$C_2$	
	$y_2$	$z_2$
$a_1 w_1$	5,0	1,1
$a_1 x_1$	4,0	4,0
$b_1 w_1$	8,3	0,1
$b_1 x_1$	7,3	3,0

## Example (3/3)

- Multiagent Representation

	$y_2$		$z_2$	
	$w_1$	$x_1$	$w_1$	$x_1$
$a_1$	5,5,0	4,4,0	1,1,1	4,4,0
$b_1$	8,8,3	7,7,3	0,0,1	3,3,0

## What the Multiagent Representation Reveals

- In the normal form, Player 1's strategies bundle choices at all her information states together
- In the multiagent representation, each information state is a **separate player**
- This separation reveals something the normal form hides:
  - A Nash equilibrium of the multiagent game requires each agent to be **independently optimal** at her own node
  - This is exactly the condition of **sequential rationality**
  - An equilibrium of the normal form need not satisfy this – a player may be best-responding overall while acting suboptimally at a node that is never reached
- The multiagent representation is the bridge between the normal form and the extensive-form refinements we will study later (sequential equilibrium, perfect equilibrium). It makes “off-path rationality” a first-class concept.

# Common Knowledge and Bayesian Games

# Common Knowledge

- A fact is **common knowledge** among players if:
  - Every player knows it,
  - Every player knows that every player knows it,
  - Every player knows that every player knows that every player knows it,
  - and so on...
- Formally, every statement of the form

(every player knows that)<sup>k</sup> every player knows it

is true for all  $k = 0, 1, 2, \dots$

- **Private information:** information held by a player that is *not* common knowledge among all players

## Common Knowledge – An Illustration

- Suppose two people are in a room and both **see a fire alarm go off**
- Each knows there is a fire; each knows the other knows; and so on – this is common knowledge
  
- Now suppose each receives a **private text message** saying “there is a fire”
- Each knows there is a fire, but neither knows whether the other received the message
- This is *not* common knowledge – the public alarm creates it, the private message does not

### **i** Note

This distinction matters in games: players can coordinate on an equilibrium only when the equilibrium itself is common knowledge. Private signals, even about the same event, may not support coordination.

# Intelligence Assumption

- In game theory, we assume players are **intelligent**:
  - Whatever we (the analysts) know about the **structure** of the game, the players also know
  - Players know the model, the strategy sets, and the payoff functions
  - Furthermore, they know that they all know the model, and so on

## **i** Note

The intelligence assumption means the **game structure** is common knowledge – not the realization of private information. Player 2 knows Player 1 drew a card; she does not know which color was drawn.

- Therefore, when constructing a model of a game, the **game structure itself must be common knowledge**

## Bayesian Games – Incomplete Information

- A game has **incomplete information** when players know different things at the time they start planning their moves
- One or more players may have **private information** that others do not know
- Example: A card game becomes one of incomplete information if Player 1 looks at a card before learning how it will be used
- Such games are common in practice

# Modeling Private Information

- Often, players have private information they have known for a long time
- It may feel unnatural to define the game as starting **before** they learned this
- Some private information is part of a player's identity (e.g., gender, language, risk preferences)
- This initial private information is called the player's **type**

# Representing Bayesian Games

- We can model these games using a **historical chance node**:
  - This node determines the players' types before actions begin
- But this can feel artificial if types are basic to identity
- So, we admit that players **start** the game already knowing their types

## Strategic Form vs. Bayesian Form

- In standard extensive-form games, players plan their moves at the **root**, before any private info is learned
- In Bayesian games, this assumption is **invalid**:
  - The root is before type information is known
  - But players begin planning **after** knowing their types
- Harsanyi (1967–68) introduced a solution

# Harsanyi's Bayesian Form

- The **Bayesian form** generalizes the strategic form to handle private information
- Players do not choose strategies before learning their type
- This form is:
  - Simpler than the full extensive form
  - More realistic than assuming full information at the root
- Bayesian games are strategic-form games where each player has a **type**, and strategies are chosen **after** types are realized

# Components of a Bayesian Game

- To define a Bayesian game, we specify:
  - A set of players:  $N$
  - For each player  $i \in N$ :
    - A set of actions:  $C_i$
    - A set of types:  $T_i$
    - A probability function:  $p_i$
    - A utility function:  $u_i$

- We define:

$$C = \prod_{i \in N} C_i, \quad T = \prod_{i \in N} T_i$$

- $C$ : all possible action profiles
- $T$ : all possible type profiles

## Beliefs About Other Players' Types

- Let  $T_{-i} = \prod_{j \in N \setminus \{i\}} T_j$
- The probability function  $p_i$  maps each  $t_i \in T_i$  to a belief over  $T_{-i}$ :

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

- So  $p_i(t_{-i} | t_i)$  is player  $i$ 's belief that the others have types  $t_{-i}$ , given his own type  $t_i$

# Utility Function in a Bayesian Game

- For each player  $i$ , the utility function is:

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

- It assigns a payoff to each combination of:
  - Action profile  $c \in C$
  - Type profile  $t \in T$
- This payoff is on a von Neumann–Morgenstern utility scale

## Formal Definition

### Definition: Bayesian Game

A Bayesian game is:

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

The game is **finite** if all sets  $N$ ,  $C_i$ , and  $T_i$  are finite. It is **common knowledge** that each player knows his own type and the full structure of the game.

# Actions vs. Strategies

- In a Bayesian game, we speak of **actions**, not strategies
- An **action** is a move chosen after the player learns their type
- A **strategy** (in extensive form) covers all contingencies before knowing the type
- So, a strategy for player  $i$  is a function:

$$s_i : T_i \rightarrow C_i$$

assigning an action for each possible type

## Example – A Simple Card Game as a Bayesian Game

- Suppose Player 1 knows the color of a card at the start of the game
- Then the game has **incomplete information** and can be written in Bayesian form

# Formal Representation

- Set of players:  $N = \{1, 2\}$
- Types:
  - $T_1 = \{1.a, 1.b\}$  (red or black card)
  - $T_2 = \{2\}$  (no private info)
- Actions:
  - $C_1 = \{R, F\}$  (raise or fold)
  - $C_2 = \{M, P\}$  (meet or pass)




Figure 13: Game 2

# Beliefs

- Player 2 believes the card is red or black with equal probability:

$$p_2(1.a | 2) = 0.5 = p_2(1.b | 2)$$

- Player 1 knows Player 2's type with certainty:

$$p_1(2 | 1.a) = 1 = p_1(2 | 1.b)$$

# Utility Functions

- Utility functions are of the form:

$$u_i(c, t), \quad \text{where } c = (c_1, c_2), \quad t = (t_1, t_2)$$

- That is, each player's payoff depends on both the **action profile** and the **type profile**
- **Example entries** (card game):
  - $u_1((R, M), (1.a, 2)) = 2$  (Player 1 raises, Player 2 meets, card is red – Player 1 wins)
  - $u_1((R, M), (1.b, 2)) = -2$  (same actions, card is black – Player 2 wins)

# Bayesian Representation

 $t_1 = 1.a$ 

	M	P
R	2,-2	1,-1
F	1,-1	1,-1

 $t_1 = 1.b$ 

	M	P
R	-2,2	1,-1
F	-1,1	-1,1

Figure 14: Bayesian Representation

## Type-Agent Representation (Harsanyi–Selten)

- Harsanyi (1967–68), inspired by Selten, introduced a way to express any Bayesian game  $\Gamma^b$  in **strategic form**
- This method is called the **type-agent representation** (also known as the **Selten game** or **posterior-lottery model**)

## Key Idea

- In this representation:
  - Each **type** of each player is treated as a separate agent
  - So, there is **one agent per type** in the original Bayesian game
- We assume type sets are disjoint:

$$T_i \cap T_j = \emptyset \quad \text{for } i \neq j$$

- Let the set of all agents be:

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

## Strategies and Actions

- Each agent  $t_i \in T_i$  chooses from the same actions as player  $i$ :

$$D_{t_i} = C_i$$

- Agent  $t_i$  is responsible for the action taken **if player  $i$  has type  $t_i$**  in the Bayesian game

## Expected Utility

- Agent  $t_i$ 's payoff is the **expected utility** of player  $i$  given  $t_i$  is their type
- Formally, for any profile of decisions  $d = (d(s))_{s \in T^*}$ :

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

- The outcome depends on:
  - Actions selected by all agents
  - The type profile  $(t_j)_{j \in N}$

## Type-Agent Representation – Normal Form

- Agents:  $T^* = \{1.a, 1.b, 2\}$
- Rows give the action of agent 1.a
- Columns give the action pair  $(d_{1.b}, d_2)$
- Entries are  $(v_{1.a}, v_{1.b}, v_2)$
- Agent 2 averages over Player 1's type using  $p = 0.5$

	$d_2 = M$		$d_2 = P$	
	$R_{1.b}$	$F_{1.b}$	$R_{1.b}$	$F_{1.b}$
$R_{1.a}$	2, -2, 0	2, 1, -0.5	1, 1, -1	1, -1, 0
$F_{1.a}$	1, -2, 0.5	1, 1, 0	1, 1, -1	1, -1, 0

# Modeling Games with Incomplete Information

- Game theory gives us a general framework to model strategic situations
- But modeling **incomplete information** is difficult in practice
- Main issue: players' beliefs are often **subjective probabilities**
  - What one player believes about another's belief becomes complex

## Example – Trivia Quiz Game

- We modify the card game into a **trivia quiz game**:
  - ① Both players put \$1 into the pot
  - ② A question is randomly drawn and announced to both
  - ③ Player 1 can **Raise** or **Fold**
  - ④ If Player 1 raises, Player 2 can **Meet** (add \$1) or **Pass**
  - ⑤ If Player 2 meets, Player 1 answers the question
- Player 1 wins if:
  - They fold and keep \$1, or
  - They raise, Player 2 passes, or
  - They raise, Player 2 meets, and Player 1 answers correctly
- Otherwise, Player 2 wins

## Belief Complexity

- After the question is announced, Player 2 is unsure whether Player 1 knows the answer
- Let  $Q$  be Player 2's **subjective probability** that Player 1 knows the answer:

$$Q \in [0, 1]$$

- But Player 1 may be unsure about Player 2's  $Q$
- Then Player 1 needs a **distribution over  $Q$**
- And Player 2 may be unsure about that – so a **distribution over distributions!**
- This leads to **higher-order beliefs** and infinite regress

# Universal Belief Space

- Can any Bayesian game represent **all** possible beliefs?
- Mertens & Zamir (1985) proved: Yes, using a **universal belief space**
- It is a belief space large enough to encode **every level** of beliefs:
  - About the state
  - About others' beliefs
  - About others' beliefs about your beliefs, etc.

## The Universal Belief Space – A Double-Edged Sword (1/2)

### Result (Mertens–Zamir 1985)

A universal belief space  $T^*$  exists that encodes every possible hierarchy of beliefs – Harsanyi's type-space approach is without loss of generality.

- This is reassuring: **any** belief configuration can be represented
- But the richness of  $T^*$  creates a fragility problem:
  - Rationalizability is the robust baseline
    - It does not rule out much; it only eliminates actions inconsistent with rationality and beliefs in rationality
  - Stronger concepts, such as Bayesian Nash equilibrium, make sharper predictions
  - Weinstein–Yildiz (2007) show that these extra predictions may be fragile
  - Any action still allowed by rationalizability can become **uniquely predicted** after a small perturbation of higher-order beliefs

## The Universal Belief Space – A Double-Edged Sword (2/2)

Concept	Robust in $T^{**}$ ?	Why?
Iterated strict dominance	<b>Yes</b>	Stable under belief perturbations
Iterated weak dominance	<b>No</b>	Sensitive to zero-probability beliefs
Nash equilibrium	<b>Partial</b>	Exists, but selection is fragile
Nash refinements	<b>No</b>	Sensitive to higher-order beliefs

# Limitations and Practical Modeling

- The universal belief space is mathematically valid but **huge** and **intractable**
- In practice, we need to work with **simplified belief structures**
- We rely on **common knowledge** to:
  - Keep type sets manageable
  - Focus on core strategic uncertainty
- So while theory guarantees generality, practice requires simplification

# Required Reading

- **Myerson (1997)** Chapter 2 – Basic Models
- *Optional:*
  - Mertens, J.F., and Zamir, S. (1985). Formulation of Bayesian Analysis for Games with Incomplete Information. *International Journal of Game Theory*, 14(1), 1–29.
  - Weinstein, J., and Yildiz, M. (2007). A Structure Theorem for Rationalizability with Application to Robust Predictions of Refinements. *Econometrica*, 75(2), 365–400.

# Example

