Bayesian Games

CMPUT 654: Modelling Human Strategic Behaviour

S&LB §6.3

Paper Presentation Scheduling

- Starting **October 15**, we will have student presentations of selected papers in behavioural game theory
- The (candidate) papers for each lecture are listed on the schedule page of the course website
- We will assign papers to students **NEXT CLASS**
    - Not every paper will be assigned
    - At least one paper per area (i.e., lecture)
    - We will use a quasilinear mechanism for the assignment :)

Recap: Repeated Games

- A **repeated game** is one in which agents play the same normal form game (the **stage game**) multiple times.
- **Finitely repeated:** Can represent as an **imperfect information extensive form** game.
- **Infinitely repeated:** Life gets more complicated
    - Payoff to the game: either **average** or **discounted** reward
    - Pure strategies map from **entire previous history** to action
- **Folk theorem** characterizes which **payoff profiles** can arise in any equilibrium
    - All profiles that are both **enforceable** and **feasible**

Lecture Outline

1. Logistics & Recap
2. Bayesian Game Definitions
3. Strategies and Expected Utility
4. Bayes-Nash Equilibrium

Fun Game!

- Everyone should have a slip of paper with 2 **dollar values** on it
- Play a sealed-bid first-price auction with three other people
  - **If you win**, utility is your **first dollar value minus your bid**
  - **If you lose**, utility is **0**
- Play again with the same neighbours, same valuation
- Then play again with same neighbours, valuation #2
- **Question:** How can we model this interaction as a game?

Payoff Uncertainty

- Up until now, we have assumed that the following are always **common knowledge**:
    - **Number** of players
    - **Actions** available to each player
    - **Payoffs** associated with each pure strategy profile
- Bayesian games are games in which there is uncertainty about the very **game being played**

Bayesian Games

We will assume the following:

1. In every possible game, number of **actions** available to each player is the same; they differ only in their **payoffs**

2. Every agent's **beliefs** are posterior beliefs obtained by conditioning a **common prior** distribution on private signals.

There are at least three ways to define a Bayesian game.

Bayesian Games via Information Sets

Definition:

A **Bayesian game** is a tuple (N, G, P, I), where

- N is a set of n ***agents***
- G is a set of ***games*** with N agents such that if $g, g' \in G$ then for each agent $i \in N$ the actions available to i in g are identical to the actions available to i in g'
- $P \in \Delta(G)$ is a **common prior** over games in G
- $I = (I_1, I_2, \ldots, I_n)$ is a tuple of **partitions** over G, one for each agent

Information Sets Example

$I_{2,1}$ $I_{2,2}$

$I_{1,1}$

MP

$2, 0$	$0, 2$
$0, 2$	$2, 0$

$p = 0.3$

PD

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

$p = 0.1$

$I_{1,2}$

Coord

$2, 2$	$0, 0$
$0, 0$	$1, 1$

$p = 0.2$

BoS

$2, 1$	$0, 0$
$0, 0$	$1, 2$

$p = 0.4$

Bayesian Games via Imperfect Information with Nature

- Could instead have a special agent **Nature** who plays according to a commonly-known mixed strategy
- **Nature** chooses the game at the outset
- Cumbersome for simultaneous-move Bayesian games
- Makes more sense for **sequential-move** Bayesian games, especially when players learn from other players' moves

Imperfect Information with Nature Example

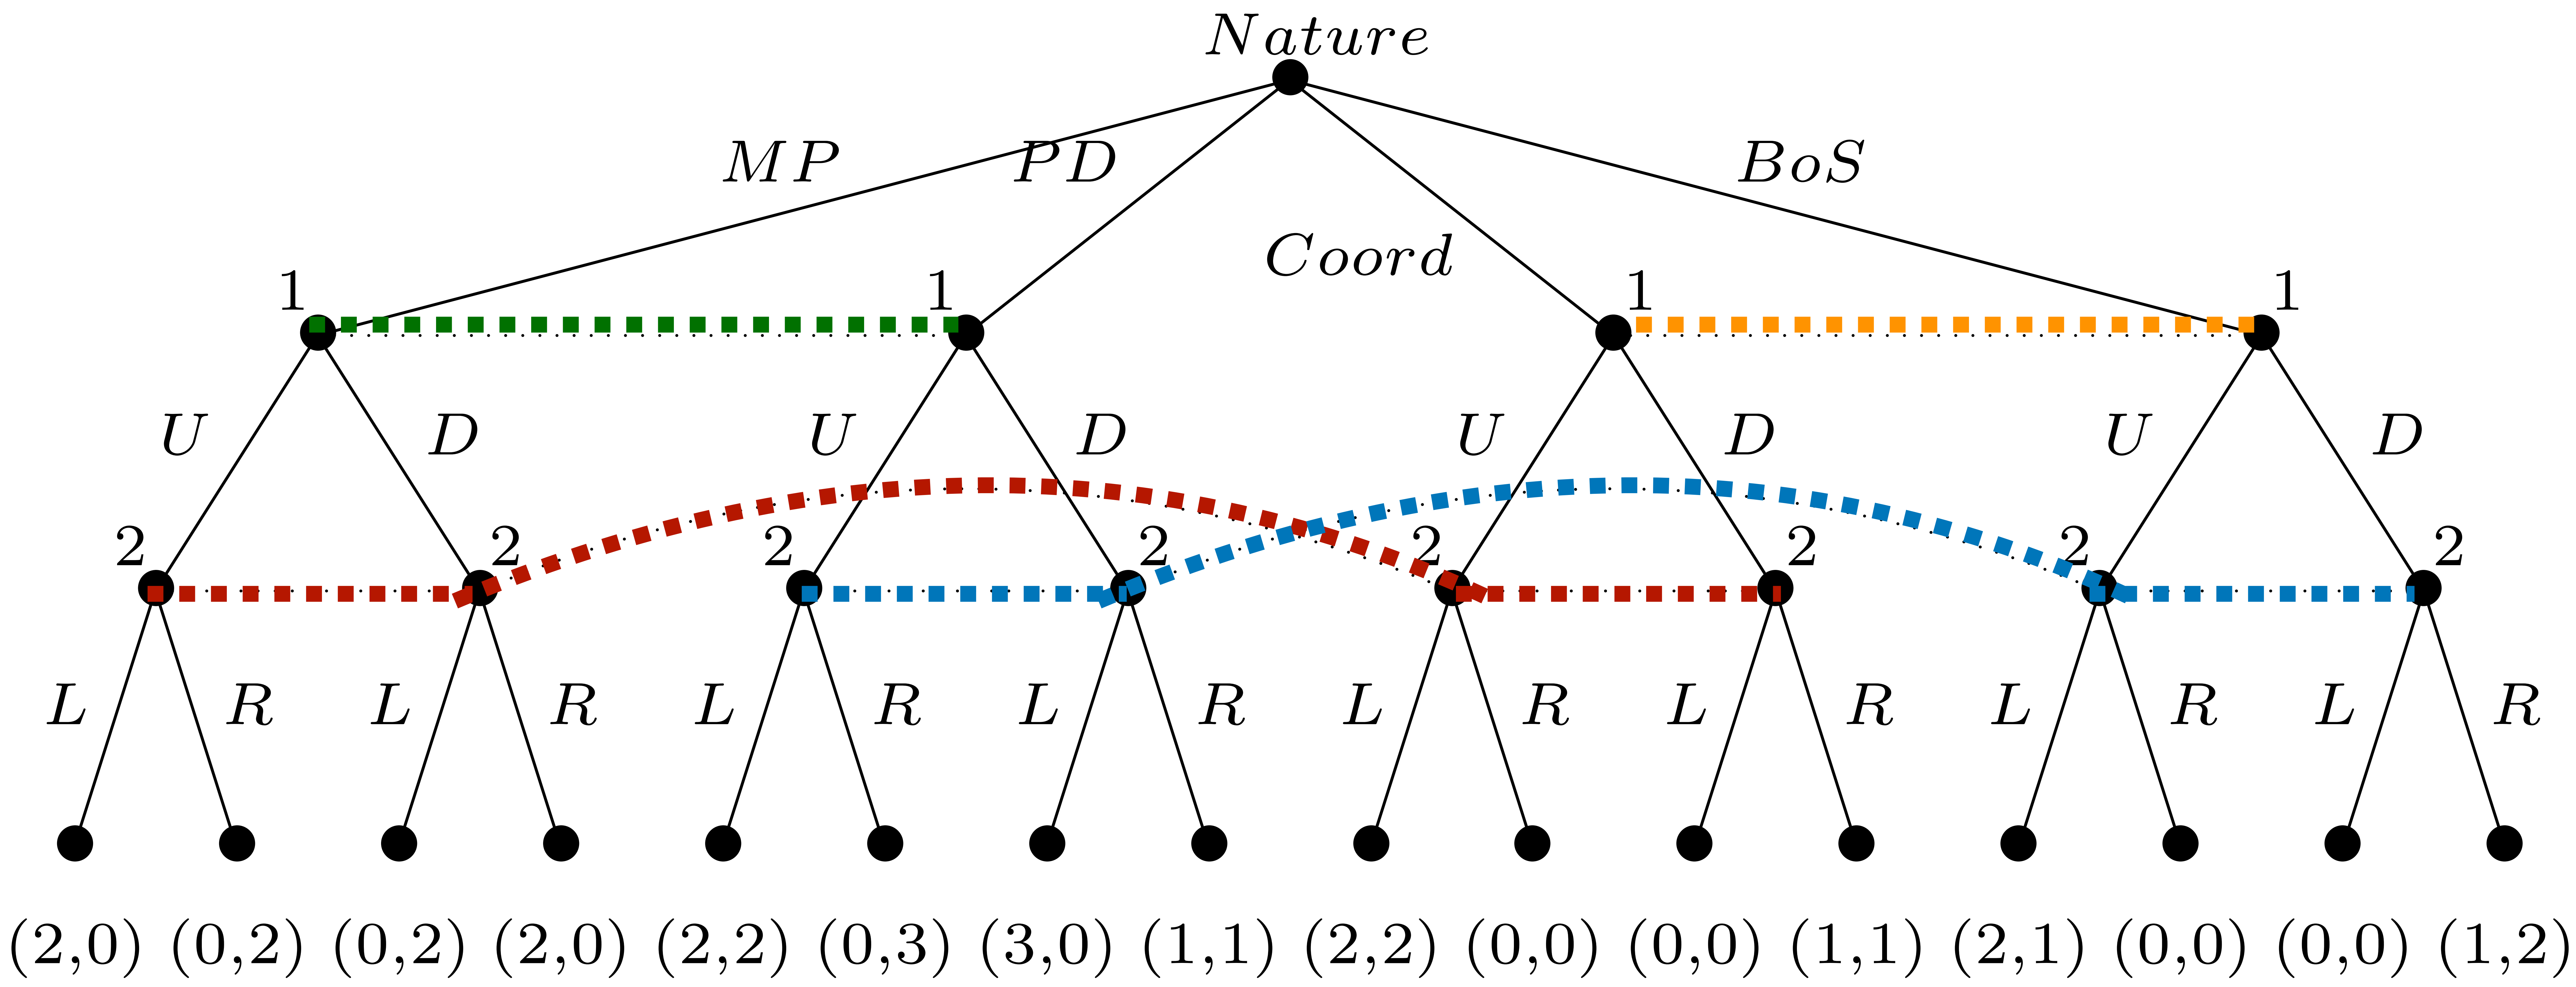

Bayesian Games via Epistemic Types

Definition:
A **Bayesian game** is a tuple (N, A, Θ, p, u) where

- N is a set of n ***players***
- $A = A_1 \times A_2 \times \cdots \times A_n$ is the set of **action profiles**
  - A_i is the **action set** for player i
- $\Theta = \Theta_1 \times \Theta_2 \times \cdots \times \Theta_n$ is the set of **type profiles**
  - Θ_i is the **type space** of player i
- $p \in \Delta(\Theta)$ is a **prior distribution** over type profiles
- $u = (u_1, u_2, \ldots, u_n)$ is a tuple of **utility functions**, one for each player
  - $u_i : A \times \Theta \to \mathbb{R}$

What is a Type?

- All of the elements in the previous definition are **common knowledge**
  - Parameterizes utility functions in a **known way**
- Every player knows their **own type**
- Type encapsulates all of the knowledge that a player has that is **not common knowledge**:
  - Beliefs about **own payoffs**
  - But also beliefs about **other player's payoffs**
  - But *also* beliefs about **other player's beliefs** about own payoffs

Epistemic Types
Example

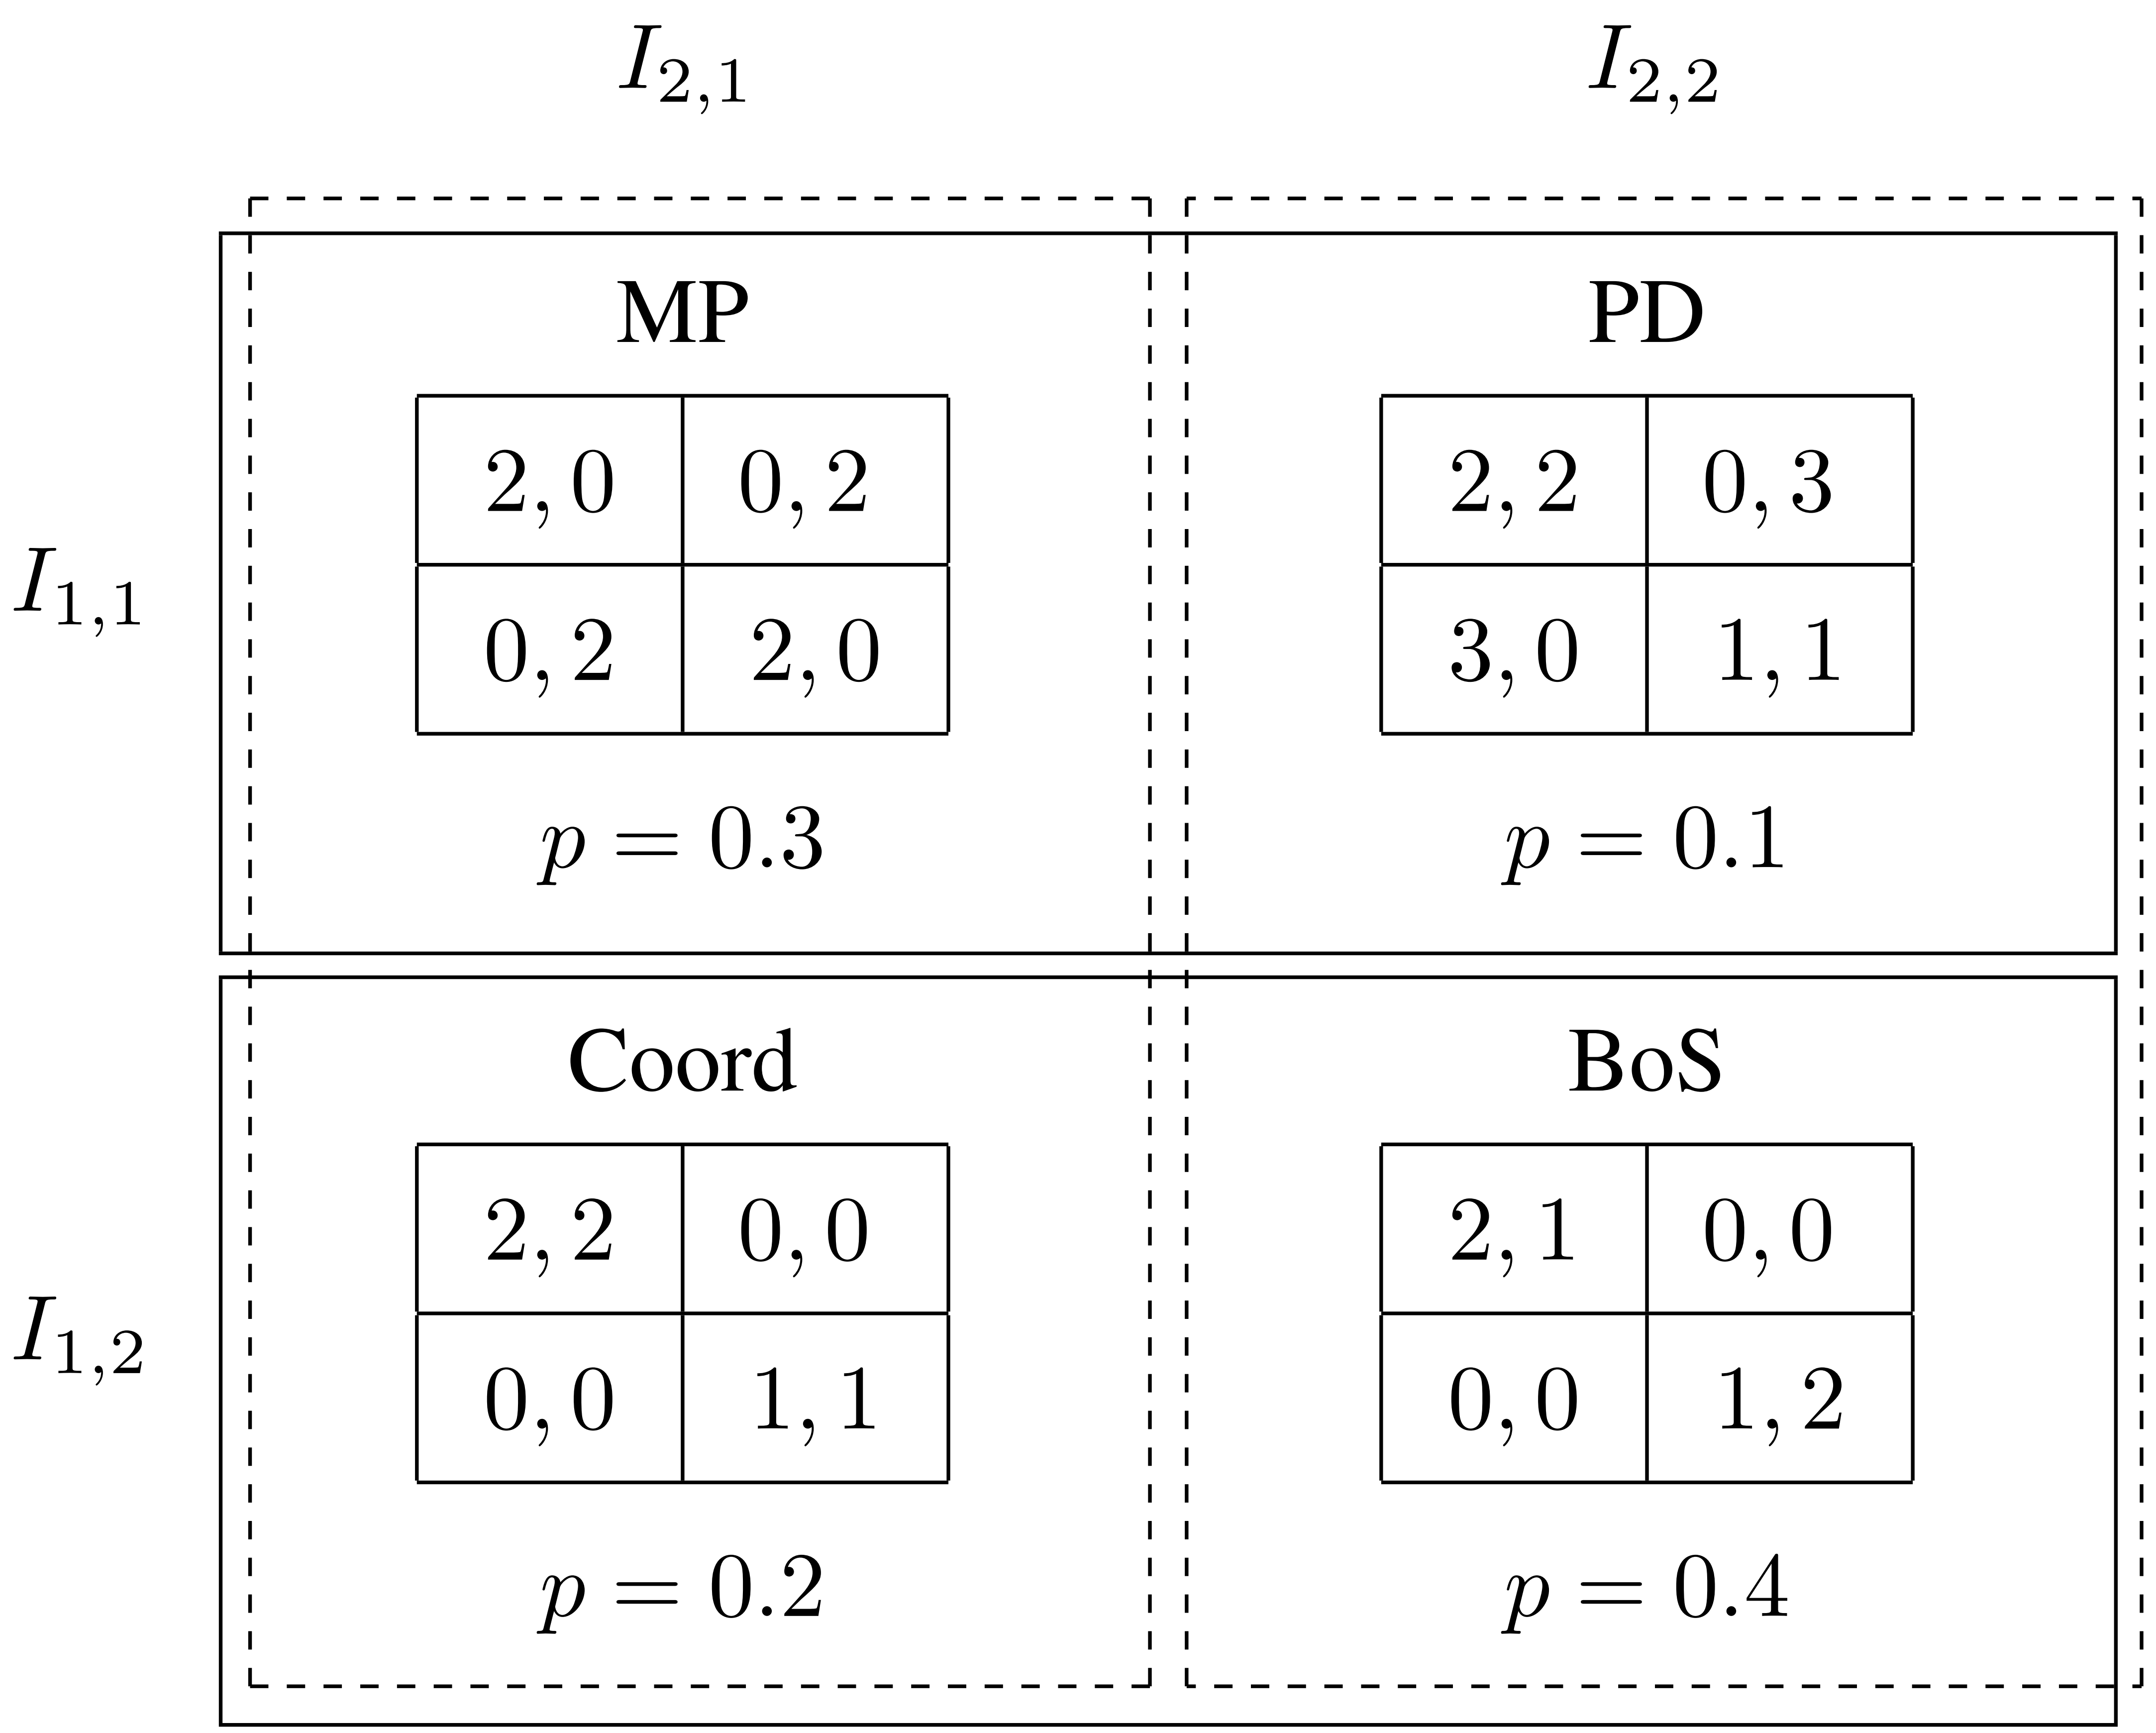

a_1	a_2	θ_1	θ_2	u_1	u_2
U	L	$\theta_{1,1}$	$\theta_{2,1}$	2	0
U	L	$\theta_{1,1}$	$\theta_{2,2}$	2	2
U	L	$\theta_{1,2}$	$\theta_{2,1}$	2	2
U	L	$\theta_{1,2}$	$\theta_{2,2}$	2	1
U	R	$\theta_{1,1}$	$\theta_{2,1}$	0	2
U	R	$\theta_{1,1}$	$\theta_{2,2}$	0	3
U	R	$\theta_{1,2}$	$\theta_{2,1}$	0	0
U	R	$\theta_{1,2}$	$\theta_{2,2}$	0	0

a_1	a_2	θ_1	θ_2	u_1	u_2
D	L	$\theta_{1,1}$	$\theta_{2,1}$	0	2
D	L	$\theta_{1,1}$	$\theta_{2,2}$	3	0
D	L	$\theta_{1,2}$	$\theta_{2,1}$	0	0
D	L	$\theta_{1,2}$	$\theta_{2,2}$	0	0
D	R	$\theta_{1,1}$	$\theta_{2,1}$	2	0
D	R	$\theta_{1,1}$	$\theta_{2,2}$	1	1
D	R	$\theta_{1,2}$	$\theta_{2,1}$	1	1
D	R	$\theta_{1,2}$	$\theta_{2,2}$	1	2

Strategies

- **Pure strategy:** mapping from agent's **type** to an **action**

$$s_i : \Theta_i \to A_i$$

- **Mixed strategy:** **distribution** over an agent's **pure strategies**

$$s_i \in \Delta(A^{\Theta_i})$$

  - *or:* mapping from type to **distribution over actions**

$$s_i : \Theta_i \to \Delta(A)$$

  - **Question:** is this equivalent? Why or why not?
  - We can use conditioning notation for the probability that i plays a_i given that their type is θ_i

$$s_i(a_i \mid \theta_i)$$

Expected Utility

The agent's expected utility is different depending on **when** they compute it, because it is taken with respect to different **distributions**.

Three relevant timeframes:

1. ***Ex-ante***: **nobody's** type is known
2. ***Ex-interim***: **own** type is known but **not others'**
3. ***Ex-post***: **everybody's** type is known

Ex-post Expected Utility

Definition:

Agent i's ***ex-post* expected utility** in a Bayesian game (N, A, Θ, p, u), where the agents' strategy profile is s and the agents' type profile is θ, is defined as

$$EU_i(s, \theta) = \sum_{a \in A} \left(\prod_{j \in N} s_j(a_j \mid \theta_j) \right) u_i(a, \theta).$$

The only source of uncertainty is in which **actions** will be realized from the mixed strategies.

Ex-interim Expected Utility

Definition:
Agent i's ***ex-interim* expected utility** in a Bayesian game(N, A, Θ, p, u), where the agents' strategy profile is s and i's type is θ_i, is defined as

$$EU_i(s, \theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i} \mid \theta_i) \sum_{a \in A} \left(\prod_{j \in N} s_j(a_j \mid \theta_j) \right) u_i(a, \theta),$$

or equivalently as

$$EU_i(s, \theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i} \mid \theta_i) EU_i(s, (\theta_i, \theta_{-i})).$$

Uncertainty over both the **actions** realized from the mixed strategy profile, and the **types** of the other agents.

Ex-ante Expected Utility

Definition:
Agent i's ***ex-ante* expected utility** in a Bayesian game (N, A, Θ, p, u), where the agents' strategy profile is s, is defined as

$$EU_i(s) = \sum_{\theta \in \Theta} p(\theta) \sum_{a \in A} \left(\prod_{j \in N} s_j(a_j \mid \theta_j) \right) u_i(a, \theta),$$

or equivalently as

$$EU_i(s) = \sum_{\theta_i \in \Theta_i} p(\theta_i) EU_i(s, \theta_i),$$

or again equivalently as

$$EU_i(s) = \sum_{\theta \in \Theta} p(\theta) EU_i(s, \theta).$$

Question:

Why are these three expressions equivalent?

Best Response

Question: What is a **best response** in a Bayesian game?

Definition:
The set of agent i's **best responses** to mixed strategy profile s_{-i} are given by

$$BR_i(s_{-i}) = \arg\max_{s_i' \in S_i} EU_i(s_i', s_{-i}).$$

Question: Why is this defined using ***ex-ante*** expected utility?

Bayes-Nash Equilibrium

Question: What is the **induced normal form** for a Bayesian game?

Question: What is a **Nash equilibrium** in a Bayesian game?

Definition:
A **Bayes-Nash equilibrium** is a mixed strategy profile s that satisfies

$$\forall i \in N : s_i \in BR_i(s_{-i}).$$

Ex-post Equilibrium

Definition:
An ***ex-post* equilibrium** is a mixed strategy profile s that satisfies

$$\forall \theta \in \Theta \;\; \forall i \in N : s_i \in \arg\max_{s_i' \in S_i} EU_i((s_i', s_{-i}), \theta).$$

- *Ex-post* equilibrium is similar to dominant-strategy equilibrium, but **neither implies the other:**
    - **Dominant strategy equilibrium**: agents need not have accurate beliefs about others' **strategies**
    - ***Ex-post* equilibrium:** agents need not have accurate beliefs about others' **types**

Question:

Why isn't *ex-post* equilibrium implied by dominant strategy equilibrium?

Dominant Strategy Equilibrium vs Ex-post Equilibrium

Question: What is a **dominant strategy** in a Bayesian game?

Example:
A game in which a dominant strategy equilibrium is not an ex-post equilibrium:

$$N = \{1,2\}$$

$$A_i = \Theta_i = \{H, L\} \qquad \forall i \in N$$

$$p(\theta) = 0.25 \qquad \forall \theta \in \Theta$$

$$u_i(a,\theta) = \begin{cases} 10 \text{ if } a_i = \theta_{-i} = \theta_i, \\ 2 \quad \text{if } a_i = \theta_{-i} \neq \theta_i, \\ 0 \quad \text{otherwise.} \end{cases} \qquad \forall i \in N$$

Summary

- Bayesian games represent settings in which there is uncertainty about the **very game being played**
- Can be defined as **game of imperfect information** with a **Nature** player, *or* as a **partition and prior** over games
- Can be defined using **epistemic types**
- **Expected utility** evaluates against three different distributions:
    - *ex-ante*, *ex-interim*, and *ex-post*
- **Bayes-Nash equilibrium** is the usual solution concept
    - ***Ex-post* equilibrium** is a stronger solution concept