Bayesian Games

Michal Jakob
Artificial Intelligence Center,
Dept. of Computer Science and Engineering,
FEE, Czech Technical University

CGT Autumn 2022 - Lecture 7

(based on Jackson, Leyton-Brown and Shoham)

Assumptions so far

All players know **what game is being played**, i.e., everyone knows *fully*:

- the number of players
- the actions available to each player
- the payoff associated with each action vector

In real-word strategic situations, this is often not the case

- salary negotiation,

Games with incomplete information

Various models of incomplete information games proposed in the literature.

We will focus on the following, practically highly useful case:

1. All games have the **same** number of **players** and the same **strategy space**. The difference is only in **payoffs** (this is without the loss of generality)
2. Agents have **beliefs** about the values of the payoffs. These believes are obtained by conditioning a **common prior** on individual private signals.

This setting is called the **Bayesian game**.

Bayesian Games

Bayesian Game: Definition 1

Definition: **Bayesian Game** (explicit partitions)

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

- N is a set of **players**
- G is a set of **games** with N players each such that: if $g, g' \in G$ then for each player $i \in \mathrm{N}$ the strategy space in g is *identical* to the strategy space in g'
- $P \in \prod(G)$ is a **common prior** over games, where $\prod(G)$ is the set of all probability distributions over G
- $I \ = \ (I_1, \dots, I_N)$ is a set of **partitions** of G, one for each agent.

Example

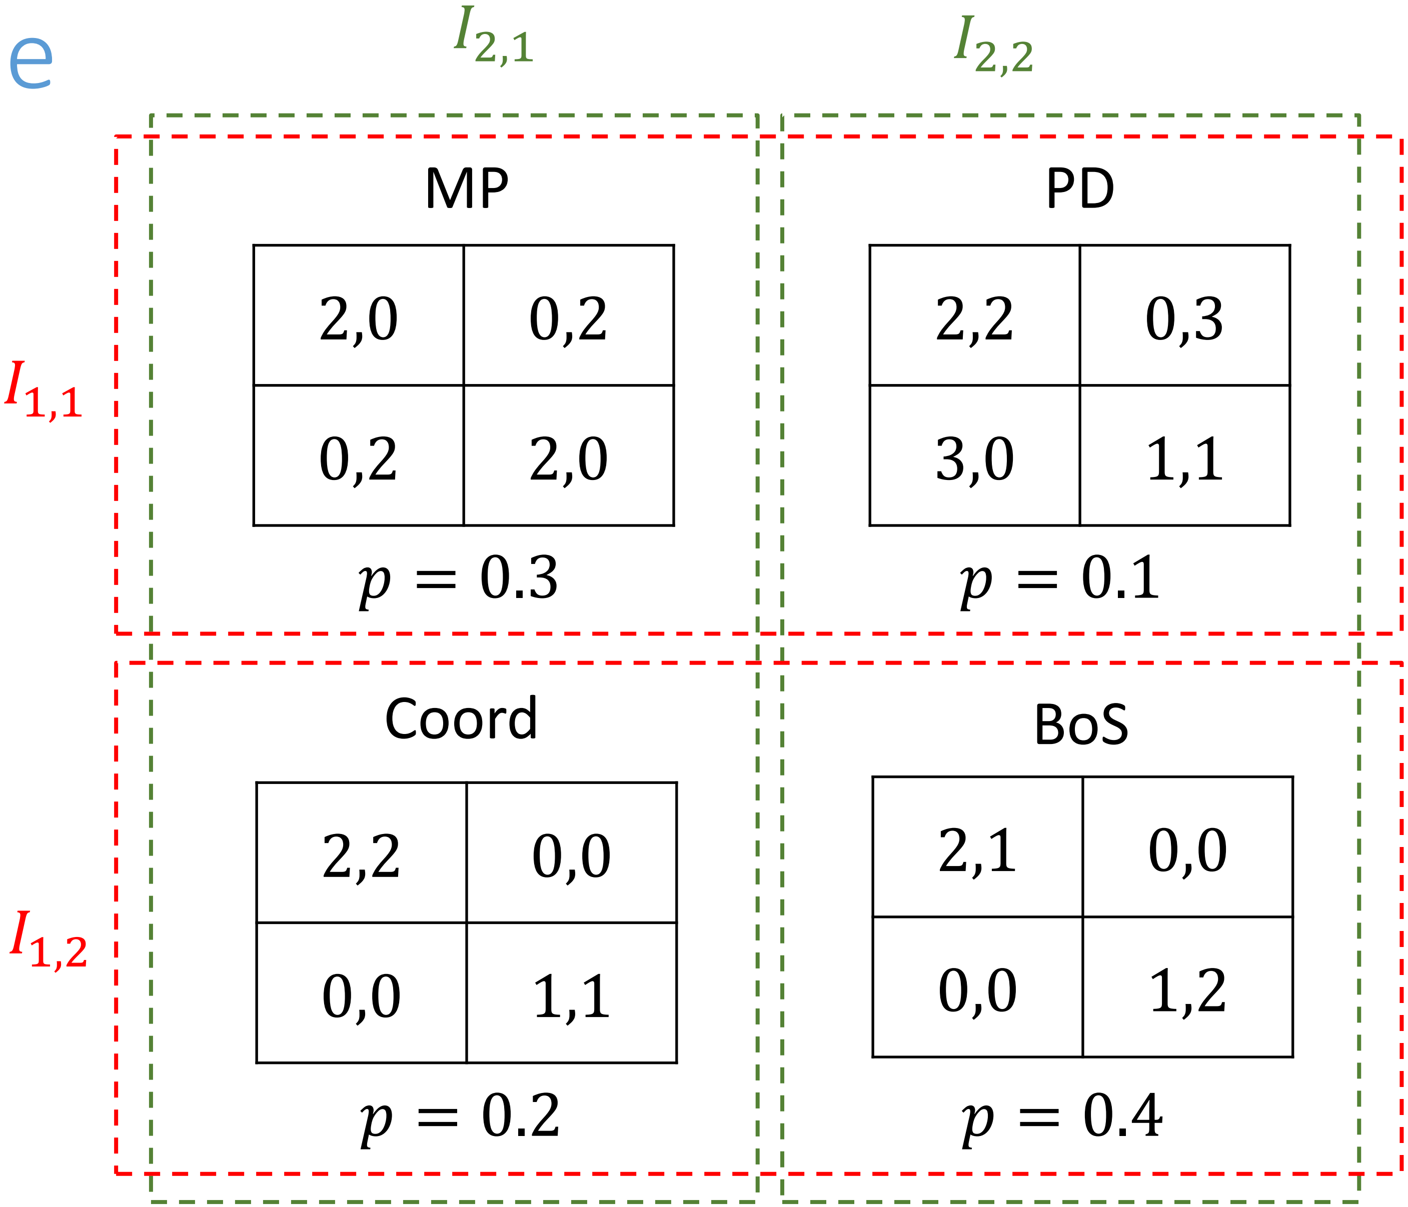

$I_{2,1}$ $I_{2,2}$

$I_{1,1}$

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

$p = 0.4$

Two players: Row player’s actions = {**T**op, **B**ottom}; Column player’s actions = {**L**eft, **R**ight}

Example

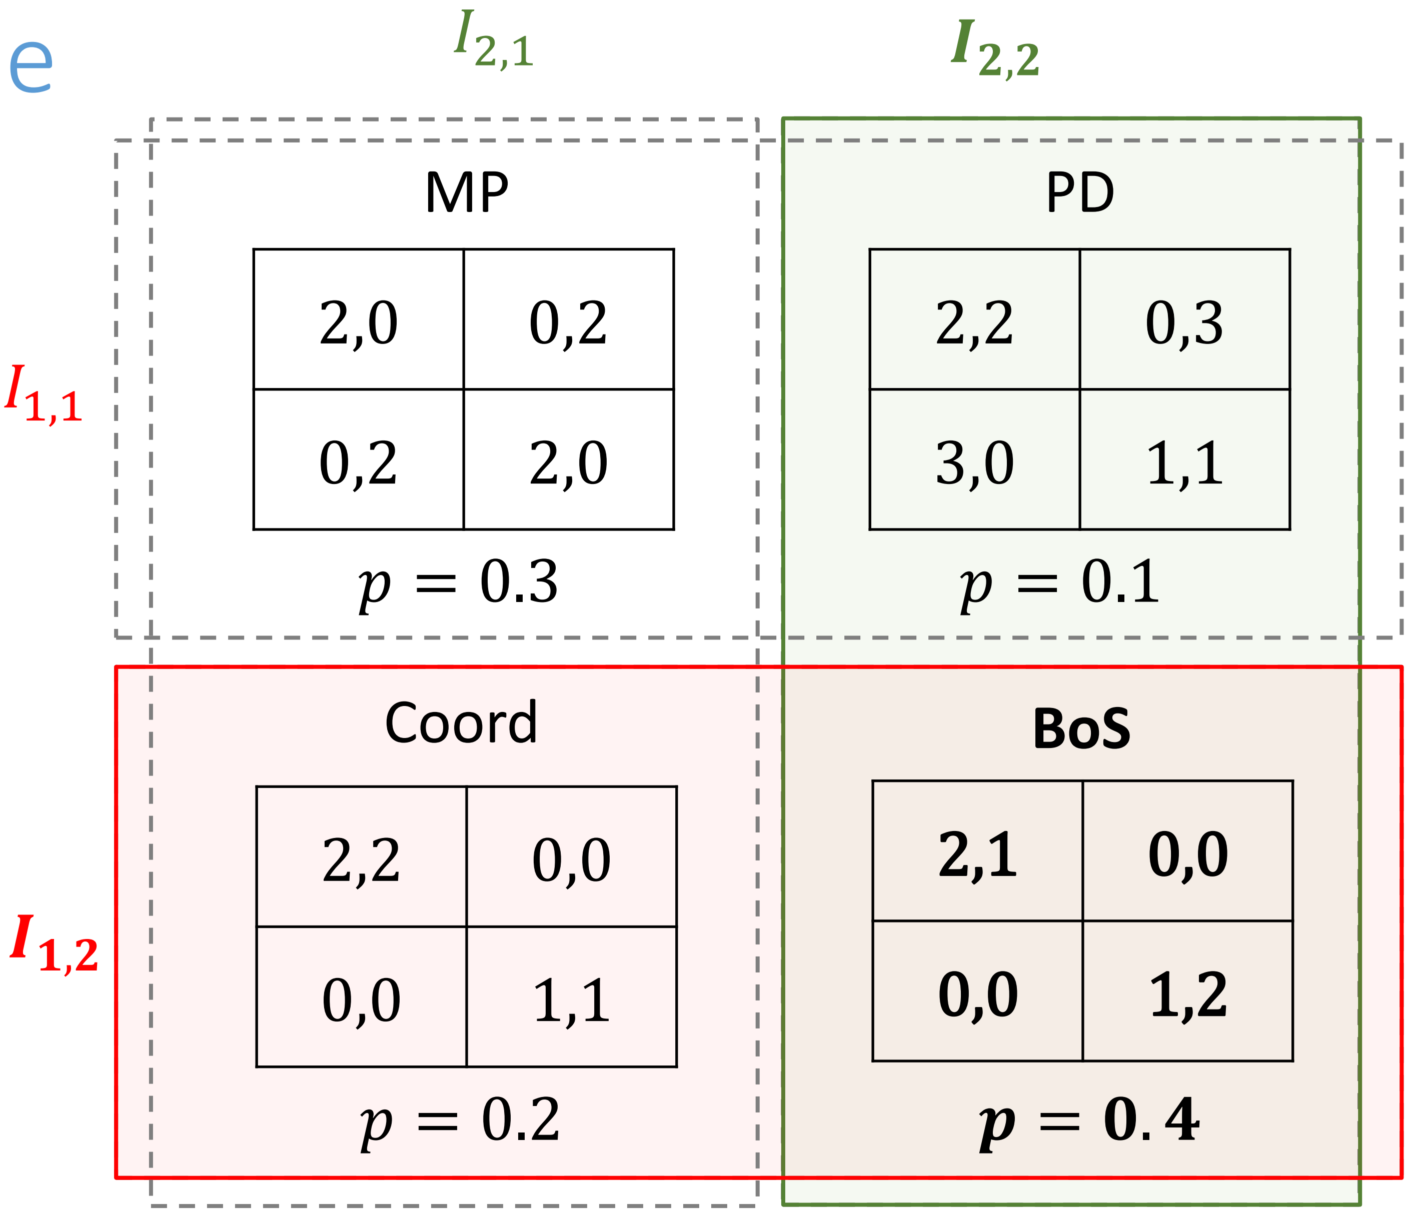

$I_{2,1}$

$\boldsymbol{I_{2,2}}$

$I_{1,1}$

MP

2,0	0,2
0,2	2,0

$p = 0.3$

PD

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

$p = 0.1$

$\boldsymbol{I_{1,2}}$

Coord

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

$p = 0.2$

BoS

2,1	**0,0**
0,0	**1,2**

$\boldsymbol{p = 0.4}$

Assume **BoS** is being played

Example

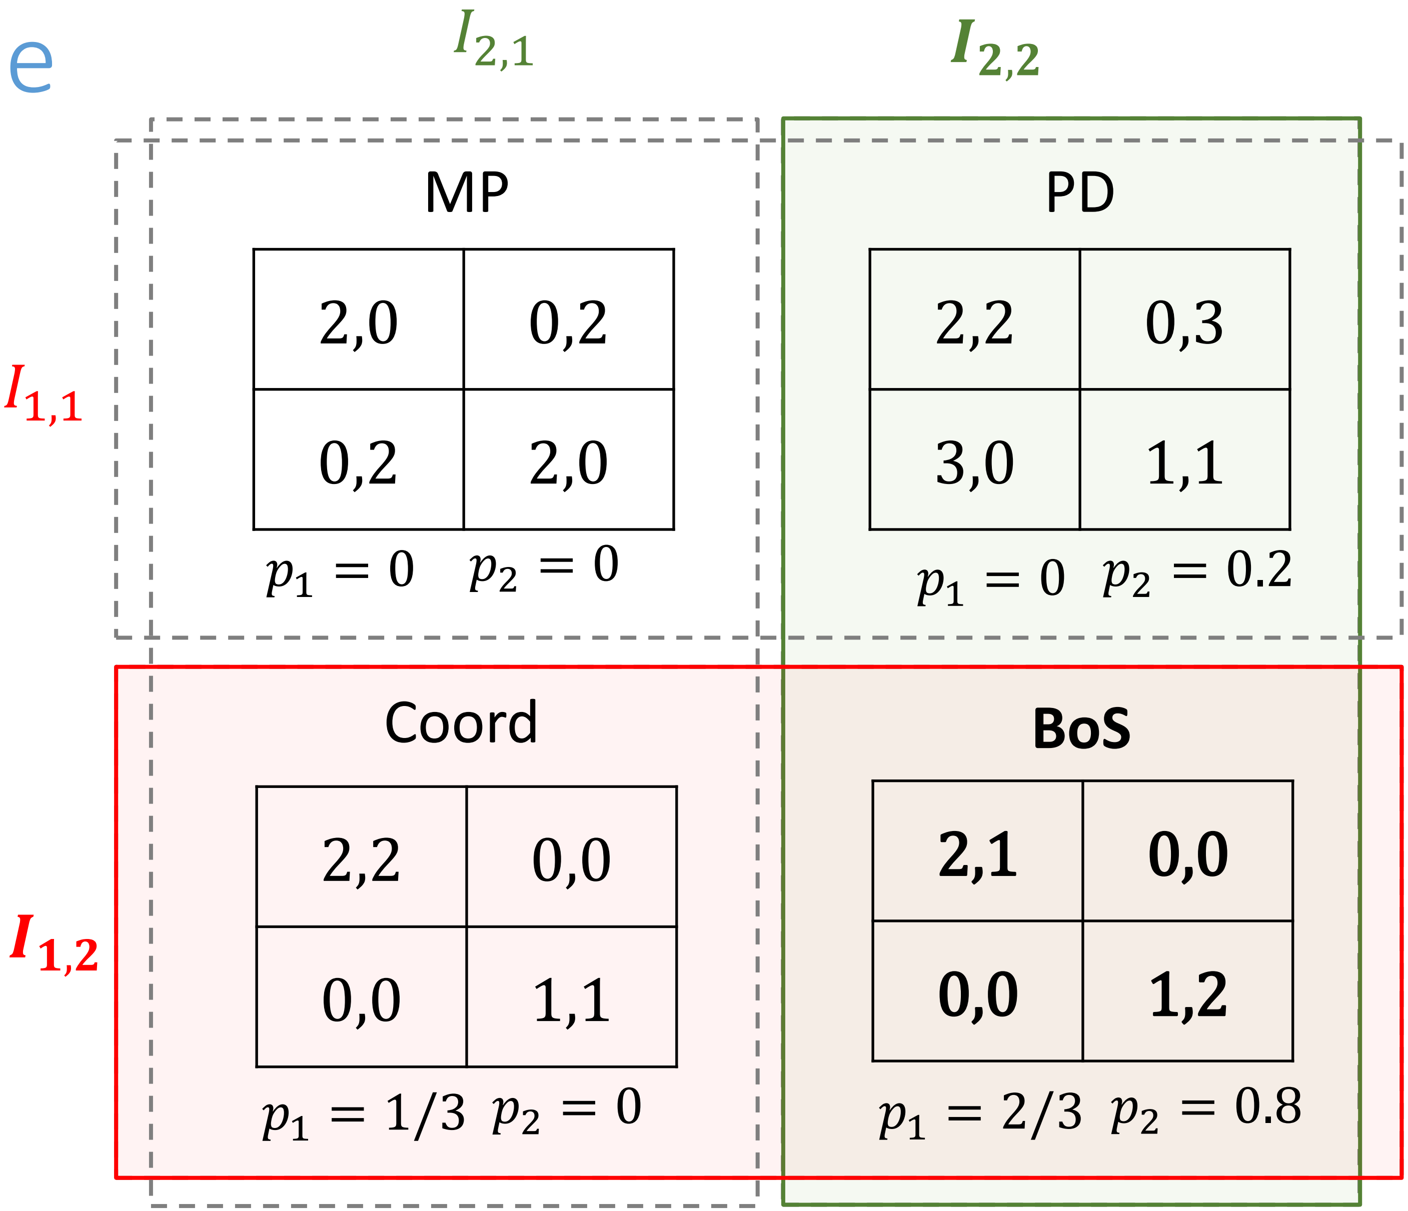

MP

2,0	0,2
0,2	2,0

$p_1 = 0$ $p_2 = 0$

PD

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

$p_1 = 0$ $p_2 = 0.2$

Coord

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

$p_1 = 1/3$ $p_2 = 0$

BoS

2,1	**0,0**
0,0	**1,2**

$p_1 = 2/3$ $p_2 = 0.8$

p_1 and p_2 … Player 1's / 2's posterior beliefs (after the private signal has been received) about which game is being played.

Example

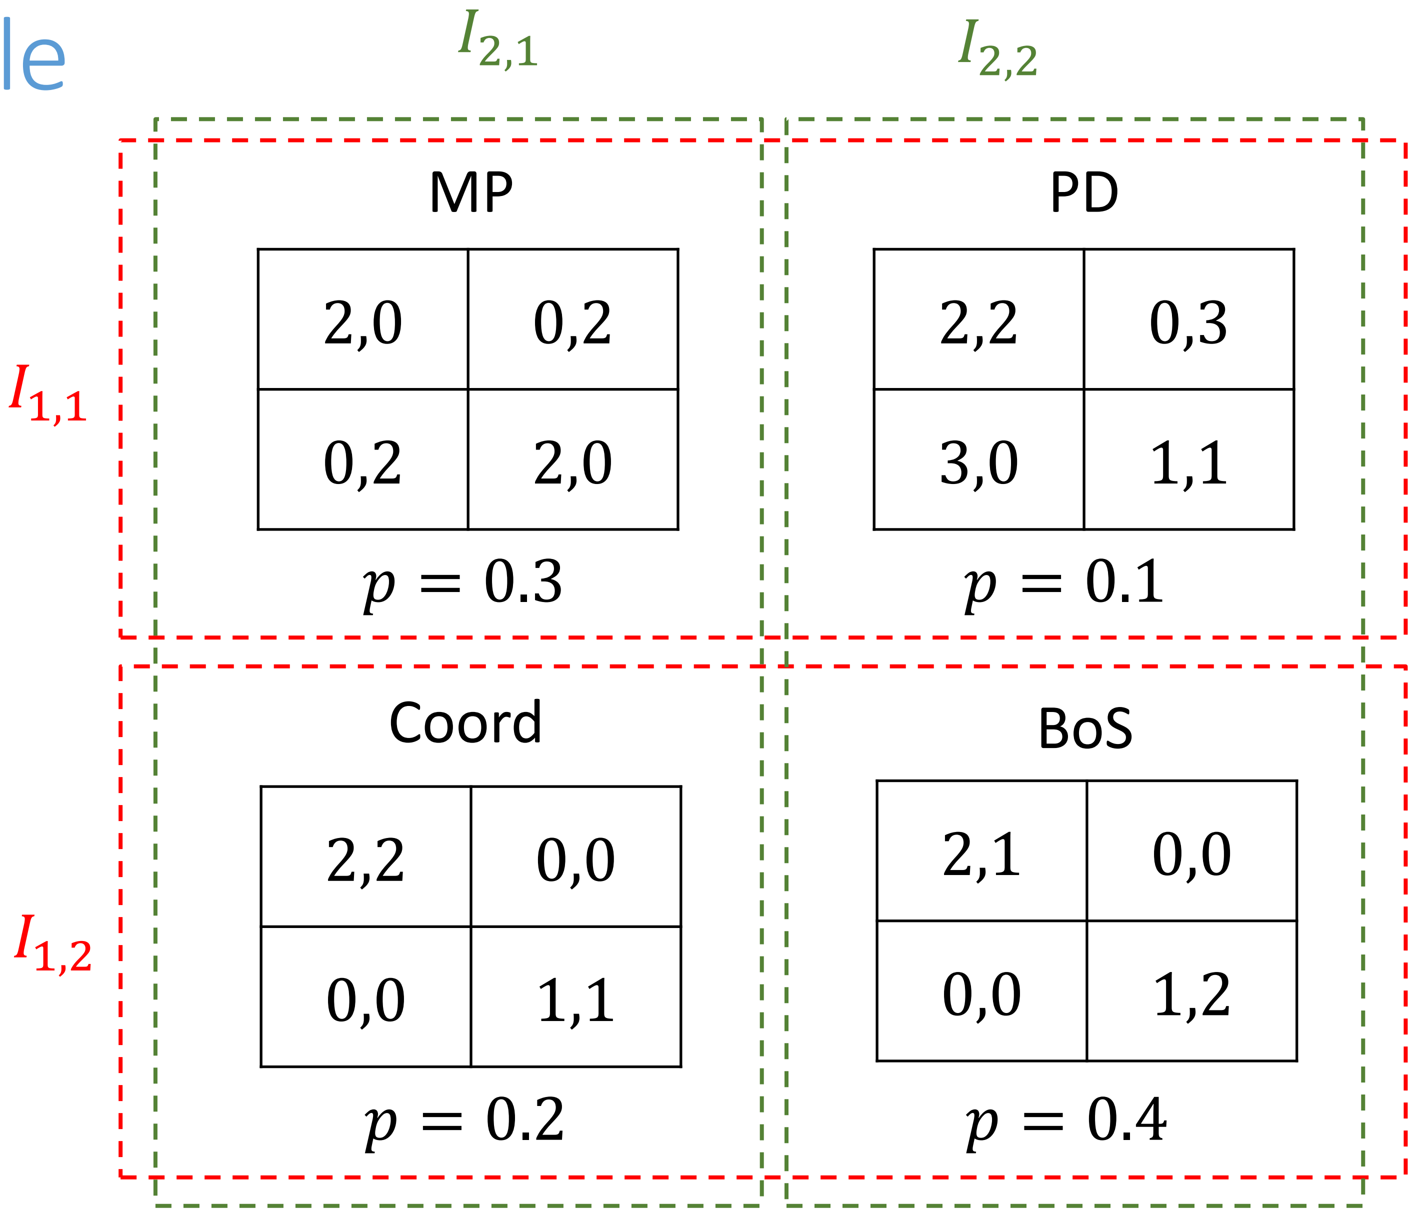

MP

2,0	0,2
0,2	2,0

$p = 0.3$

PD

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

$p = 0.1$

Coord

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

$p = 0.2$

BoS

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

$p = 0.4$

The whole infinite hierarchy of **nested beliefs** is **common knowledge.**

Another definition

This was a definition based on an explicit partitioning of the games into information sets.

There is an equivalent, mathematically more compact definition.

Bayesian Game: Definition 2

Directly represent uncertainty over utility function using the notion of **epistemic type**.

Definition: **Bayesian Game** (type-based)

Bayesian game is a tuple $\langle N, A, \Theta, p, \boldsymbol{u} \rangle$ where

- N is the set of **players**
- $\Theta = \Theta_1 \times \Theta_2 \times \cdots \times \Theta_n$, Θ_i is the **type space** of player i
- $A = A_1 \times A_2 \times \cdots \times A_n$ where A_i is the **set of actions** for player i
- $p: \Theta \mapsto [0,1]$ is a **common prior over types**
- $\boldsymbol{u} = (u_1, \ldots, u_n)$, where $u_i: \mathrm{A} \times \Theta \mapsto \mathbb{R}$ is the **utility function** of player i

The type captures all the information private to a player.

Example

(using Definition 2)

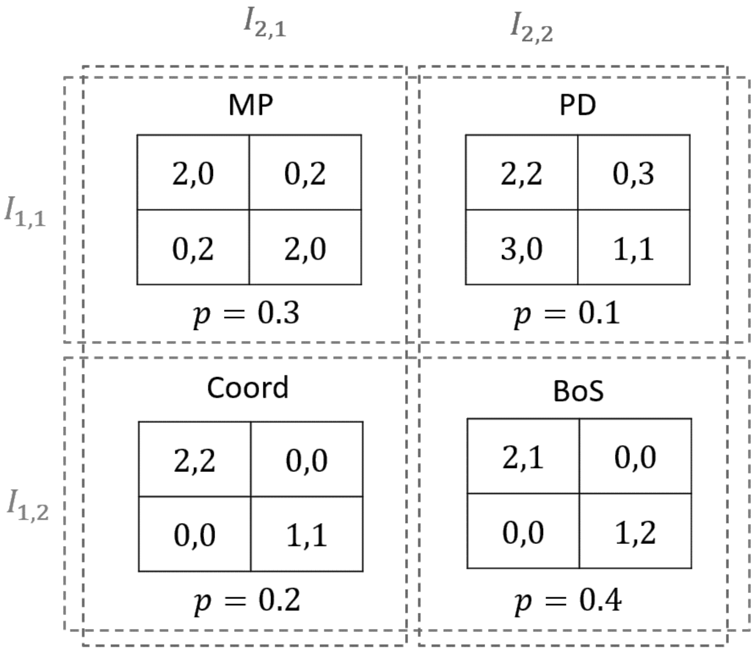

$I_{2,1}$ $I_{2,2}$

$I_{1,1}$

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

$p = 0.4$

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

Analysing Bayesian Games

Bayesian (Nash) Equilibrium

A plan of action for each player as a function of types that **maximize each type's expected utility**:

1. expecting over the **actions** of other players,
2. expecting over the **types** of (other) players.

Strategies

Given a Bayesian game (N, A, θ, p, u) with *finite* sets of players, actions, and types, strategies are defined as functions of player types as follows:

- **Pure strategy**: $s_i: \Theta_i \rightarrow A_i$
- **Mixed strategy**: $s_i: \Theta_i \rightarrow \prod A_i$

We denote $\boldsymbol{s_i(a_i|\theta_i)}$ the probability under a mixed strategy s_i that player i plays action a_i, given that i's type is θ_i.

Can be generalized to *infinite sets* (both countable and uncountable) but need to be careful about details (in particular measurability).

Expected Utility in Bayesian Games

Three standard notions of **expected utility**:

- **ex-ante**: the player knows nothing about *anyone's* actual type (including her)
- **interim:** the player knows her own type but not the types of the other players;
- **ex-post**: the player knows all players' types
(→ corresponds to a complete information game)

Interim expected utility

Given a Bayesian game (N, A, Θ, p, u) with *finite* sets of players, actions, and types, player i's ***interim* expected utility** with respect to *type* θ_i and a *mixed strategy profile* s is

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

Ex-ante expected utility

Given a Bayesian game (N, A, θ, p, u) with finite sets of players, actions, and types, player i’s ***ex-ante* expected utility** with respect a *mixed strategy* profile s is

$$EU_i(s) = \sum_{\theta_i \in \theta_i} p(\theta_i) \underbrace{EU_i(s|\theta_i)}_{\text{interim expected utility}}$$

Note: Ex-ante expected utility is not conditioned on the player’s type.

Bayesian Equilibrium (or Bayes-Nash equilibrium)

Definition (Bayes Nash Equilibrium)

Bayesian equilibrium is a mixed strategy profile s that satisfies

$$s_i \in \arg\max_{s_i'} \sum_{\theta_i} p(\theta_i) EU_i(s_i', s_{-i} | \theta_i)$$

for each i and $\theta_i \in \Theta_i$.

This definition is based on interim maximization of utility.

Bayesian Equilibrium (ex-ante)

Assuming all types occur with *positive probability*, i.e., every $p(\theta_i) > 0$ for all $\theta_i \in \Theta_i$, then for each i:

$$s_i \in \arg\max_{s_i'} EU_i(s_i', s_{-i}) = \arg\max_{s_i'} \sum_{\theta_i} p(\theta_i) EU_i(s_i', s_{-i} | \theta_i)$$

I.e. the Bayes-Nash equilibrium strategy should maximize **ex-ante** expected utility.

Bayes-Nash Equilibrium

Explicitly models behavior in uncertain environment.

Playes choose strategies to maximize their payoffs in response to others accounting for:

- **strategic uncertainty** about how others will play
- **payoff uncertainty** about the values of their actions

Example: Sheriff’s Dilemma

A sheriff is faces an armed suspect and they each must (simultaneously) decide whether to shoot the other or not, and:

the suspect is either a **criminal** (with probability p) or **innocent** with probability $1 - p$.

vs.

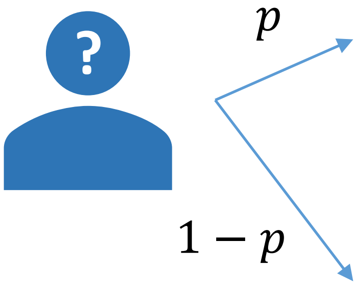

the **criminal:** would rather shoot even if the sheriff does not, as the criminal would be caught if he does not shoot.

the **innocent suspect**: would rather not shoot even if the sheriff shoots.

the sheriff would rather shoot if the suspect shoots, but not if the suspect does not.

Sheriff’s Dilemma: Baysesian Game Formulation

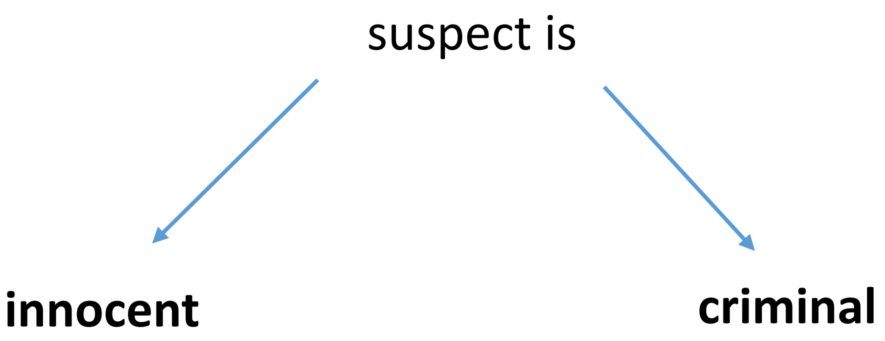

suspect is

innocent

criminal

sheriff

		Shoot	Not
suspect	Shoot	$-3,-1$	$-1,-2$
	Not	$-2,-1$	$0,0$

sheriff

		Shoot	Not
suspect	Shoot	$-0,-0$	$2,-2$
	Not	$-2,-1$	$-1,1$

Sheriff’s Dilemma: Suspect’s strategy

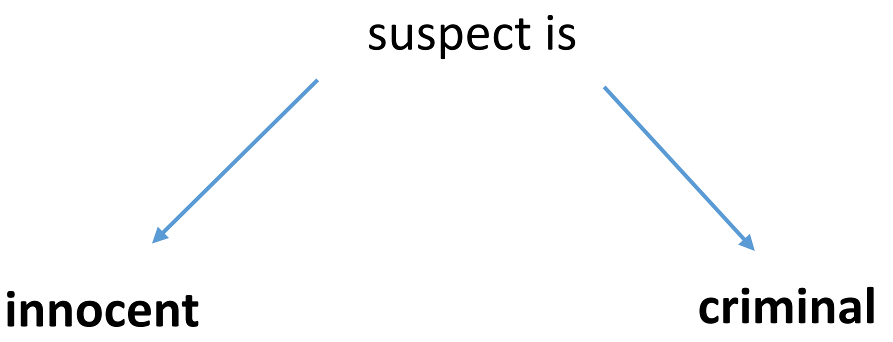

suspect is

innocent

criminal

		sheriff	
		Shoot	Not
suspect	~~Shoot~~	~~$3, -1$~~	~~$1, -2$~~
	Not	$-2, -1$	$0, 0$

dominant strategy for innocent suspect

		sheriff	
		Shoot	Not
suspect	Shoot	$-0, -0$	$2, -2$
	~~Not~~	~~$-2, -1$~~	~~$-1, 1$~~

dominant strategy for criminal

Sheriff’s Dilemma: Sheriff’s strategy

innocent

sheriff

suspect	Shoot	Not	
~~Shoot~~	~~$-3, -1$~~	~~$-1, -2$~~	
Not	$-2, -1$	$0, 0$	$1 - p$

criminal

sheriff

suspect	Shoot	Not	
Shoot	$-0, -0$	$2, -2$	p
~~Not~~	~~$-2, -1$~~	~~$-1, 1$~~	

Sheriff’s expected payoff :

Shoot	Not
$-1(1-p)$	$0(1-p)$
$+\ 0p$	$-\ 2p$
$= \boldsymbol{p - 1}$	$= \boldsymbol{-2p}$

$\Rightarrow$

Sheriff’s best response:

$p > \frac{1}{3}$: shoot

$p < \frac{1}{3}$: do NOT shoot

$p = \frac{1}{3}$: any mixture

Sheriff's Dilemma: Bayes-Nash Equilibrium

Bayes-Nash equilibrium for the Sheriff's game depends on p:

- $p > \frac{1}{3}$: sheriff should shoot; suspect should shoot if criminal and not shoot if innocent (unique equilibrium)
- $p < \frac{1}{3}$: sheriff should NOT shoot; suspect same as above (unique equilibrium)
- $p = \frac{1}{3}$: sheriff any mixture; suspect same as above

Bayesian Equlibrium Summary

Explicitly models behavior in an uncertain environment

Players choose strategies to maximize their payoffs in response to others accounting for:

- *strategic uncertainty* about how others will play and
- *payoff uncertainty* about the value to their actions

Payoff uncertainty common in real-world strategic situations.