Lecture 10:
Bayesian Game

Yi, Yung (이융)
KAIST, Electrical Engineering
http://lanada.kaist.ac.kr
yiyung@kaist.edu

Overview of Bayesian Game

- What happens I don't know about you?
- Incomplete information game
- Static Bayesian Game
- Example:
  - Battle of Sexes with incomplete information
  - Cournot Duopoly with incomplete information
- Dynamic case (extensive form): Not covered in this lecture

What is Bayesian Game?

What is Bayesian Game?

Game in strategic form

- Complete information (each player has perfect information regarding the element of the game)
- Iterated deletion of dominated strategy, Nash equilibrium: solutions of the game in strategic form

Bayesian Game

- A game with incomplete information
- Each player has initial private information, **type**.
- Bayesian equilibrium: solution of the Bayesian game

Battle of Sexes Game

	B	S
B	2,1	0,0
S	0,0	1,2

Battle of the sexes with incomplete information

Player 1 would like to match player 2's action

Player 1 is unsure about player 2's preferences:

a) may like to match player 1

b) may like to avoid player 1

Player 2 knows that player 1 is unsure

Player 1's View of Player 2

- Player 1's unsure belief
  - Probabilistic distribution for each "type"
  - "A-ha, P2 has two types, but I cannot differentiate between them"

matching

	B	S
B	2,1	0,0
S	0,0	1,2

state $\omega_1, \Pr(\omega_1) = \pi$

avoiding

	B	S
B	2,0	0,2
S	0,1	1,0

state $\omega_2, \Pr(\omega_2) = 1 - \pi$

Bayesian Game Setup

1. Nature chooses state $\omega \in \{\omega_1, \omega_2\}$, with prob $(\pi, 1-\pi)$

2. Player 2 observes the realized state ω; player 1 does not.

3. The two player's simultaneously choose actions

4. Payoffs given by the actions chosen and state, as in tables.

Structure of this game (1-4) is commonly known by both players

e.g. 2 knows that 1 knows that $\Pr(\omega_1) = \pi$

What is a strategy?

- Player 2 knows the state, but Player 1 does not

- Thus, a strategy is a triple of actions
  - One for player 1
  - Two for player 2

- Example: [B;(B,S)]
  - Player 1 chooses B
  - Player 2 chooses B for the state 1 and S for the state 2

- Well, that's "pure" strategy

Bayesian Game Setup: Interpretation

- Player 2
  - Initially, P2 does not know the state; she is informed of the state by a s ignal that depends on the state
  - Before receiving the signal she carries out her planned action for that signal

- Player 1
  - P1 also receives the signal, but it is uninformative: it must be the sam e in each state.
  - Given her signal, she is unsure of the state; when choosing an action s he takes into account her belief about the likelihood of each state, giv en her signal.

What is a strategy? Pure? Mixed?

- Mixed strategy

- Example: $[p;(q_1, q_2)]$. It means …

- Nash Equilibrium
  - Pure strategy NE
  - Mixed strategy NE

  - Study this for our particular example
  - Then, we will formalize this later.

NE of BoS (1)

	B	S
B	2,1	0,0
S	0,0	1,2

state $\omega_1, \Pr(\omega_1) = \pi$

	B	S
B	2,0	0,2
S	0,1	1,0

state $\omega_2, \Pr(\omega_2) = 1 - \pi$

Pure strategy equilibrium: action for player 1, a_1

& pair of actions for player 2, $a_2(\omega_1), a_2(\omega_2)$:

Player 1's action is optimal:

$$\pi u_1[a_1, a_2(\omega_1)] + (1-\pi) u_1[a_1, a_2(\omega_2)] \geq \pi u_1[a_1', a_2(\omega_1)] + (1-\pi) u_1[a_1', a_2(\omega_2)], \ \forall a_1'$$

(player 1's payoff does not depend upon state)

NE of BoS (2)

	B	S
B	2,1	0,0
S	0,0	1,2

state ω_1, $\Pr(\omega_1) = \pi$

	B	S
B	2,0	0,2
S	0,1	1,0

state ω_2, $\Pr(\omega_2) = 1 - \pi$

Player 2's action is optimal at ω_1 :

$$u_2(a_1, a_2(\omega_1), \omega_1) \geq u_2(a_1, a_2', \omega_1), \forall a_2'$$

Player 2's action is optimal at ω_2 :

$$u_2(a_1, a_2(\omega_2), \omega_2) \geq u_2(a_1, a_2', \omega_2), \forall a_2'$$

Player 2 knows the state, so the probability π is not relevant for his calculation.

NE of BoS (3)

	B	S
B	2,1	0,0
S	0,0	1,2

state $\omega_1, \Pr(\omega_1) = \pi$

	B	S
B	2,0	0,2
S	0,1	1,0

state $\omega_2, \Pr(\omega_2) = 1 - \pi$

Strategy for player 1: $a_1 \in A_1$

Strategy for 2: $(a_2(\omega_1), a_2(\omega_2))$

Is $[B; (B, S)]$ an equilibrium?

If 1 plays B, optimal for 2 to play B at ω_1 and S at ω_2.

For 1,

$$\mathbf{E}u_1[B; (B, S)] = \pi \times 2 + (1 - \pi) \times 0 = 2\pi.$$

NE of BoS (3)

For 1,

$$\mathbf{E}u_1[B;(B,S)] = \pi \times 2 + (1-\pi) \times 0 = 2\pi.$$

$$\mathbf{E}u_1[S;(B,S)] = \pi \times 0 + (1-\pi) \times 1 = 1-\pi.$$

Optimal to play B as long as

$$2\pi \geq 1-\pi \Rightarrow \pi \geq \frac{1}{3}.$$

If $\pi \geq \frac{1}{3}$, $[B;(B,S)]$ a Nash equilibrium.

If $\pi < \frac{1}{3}$, $[B;(B,S)]$ is not a Nash equilibrium.

NE of BoS (3)

Suppose $\pi \geq \frac{1}{3}$. Is there another pure strategy Nash equilibrium?

$[S;(S,B)]$. What conditions on π must be satisfied for this to be a NE?

Similar caclulation: if $\pi \geq \frac{2}{3}$, $[S;(S,B)]$ is an equilibrium, not if $\pi < \frac{2}{3}$.

So if $\pi \geq \frac{2}{3}$, two pure strategy equilibria

What happens if $\pi < \frac{1}{3}$?

Neither is an equilibrium, so equilibrium must be in mixed strategies

Fix $\pi = 0.25$.

NE of BoS (3)

Solve for a mixed strategy Nash equilibrium.

Player 1 chooses a prob. of B, and each type of player 2 chooses a prob. of B.

(one type of player 1 might choose a pure action)

$[p; (q_1, q_2)]$

Another Example

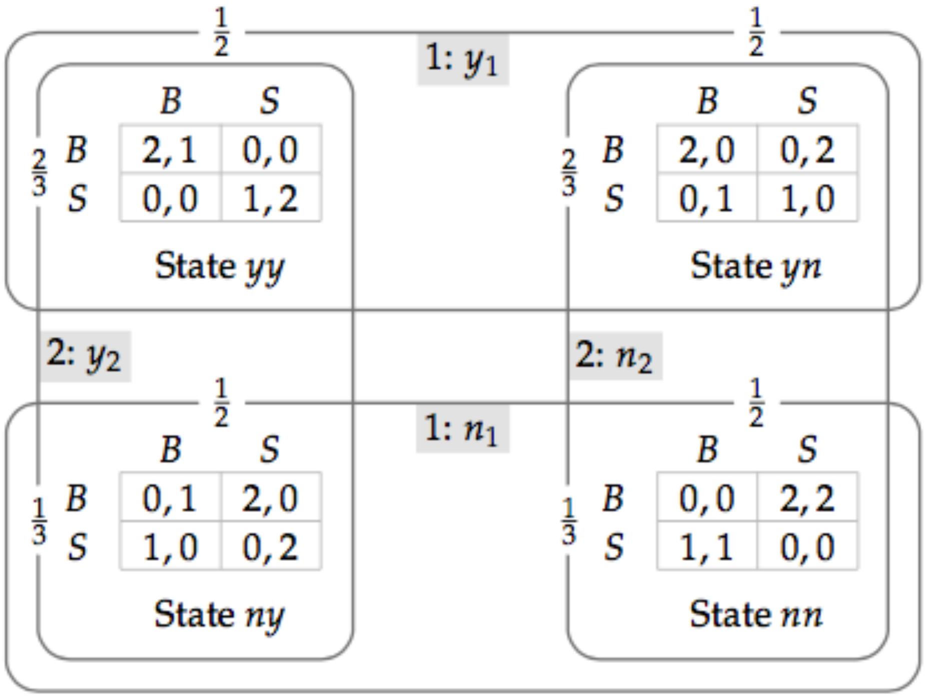

P1: cannot differentiate between yy and yn
or between ny, nn
P2: cannot differentiate between yy and ny
or between yn and nn

Understanding:
P1 receives the same "signal" y1 for yy and yn
a different "signal" n1 for ny and nn
Similarly, for P2 for the signals y2 and n2

Expected Payoff

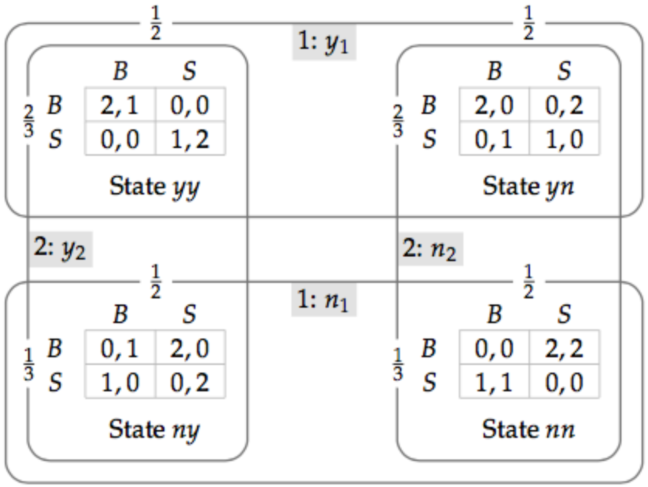

Player 1's expected payoff

P2

P1

	(B,B)	(B,S)	(S,B)	(S,S)
(B,B)				
(B,S)				
(S,B)				
(S,S)				

Player 2's expected payoff

P2

P1

	(B,B)	(B,S)	(S,B)	(S,S)
(B,B)				
(B,S)				
(S,B)				
(S,S)				

Bayesian Game

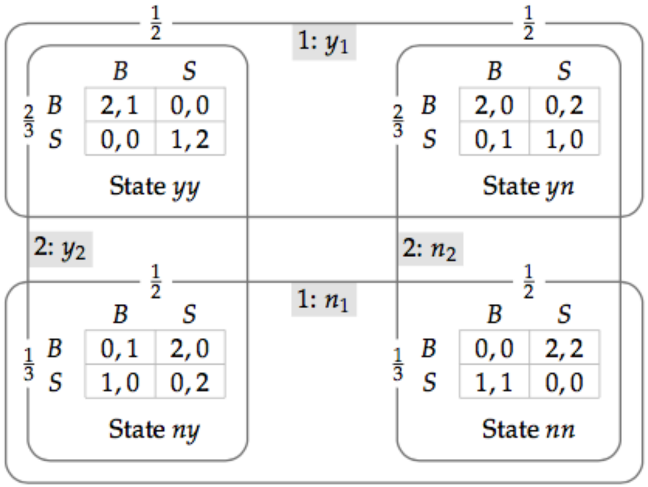

Definition

A Bayesian game consists of

- *A set of players* $\mathcal{I}$;
- *A set of actions (pure strategies) for each player* i: S_i;
- *A set of types for each player* i: $\theta_i \in \Theta_i$;
- *A payoff function for each player* i: $u_i(s_1, \ldots, s_I, \theta_1, \ldots, \theta_I)$;
- *A (joint) probability distribution* $p(\theta_1, \ldots, \theta_I)$ *over types (or* $P(\theta_1, \ldots, \theta_I)$ *when types are not finite).*

Note that the players' types may not independent

Bayesian Game

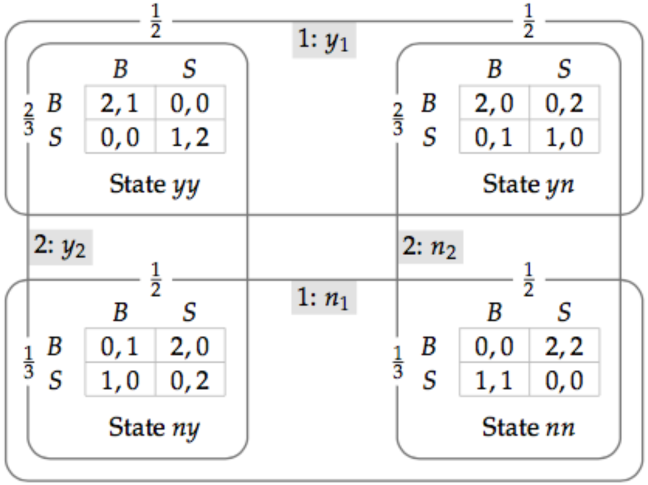

- Importantly, throughout in Bayesian games, the strategy spaces, the payoff functions, possible types, and the prior probability distribution are assumed to be **common knowledge**.
  - Very strong assumption.
  - But very convenient, because any private information is included in the description of the type and others can form beliefs about this type and each player understands others' beliefs about his or her own type, and so on, and so on.

Definition

A (pure) strategy for player i is a map $s_i : \Theta_i \to S_i$ prescribing an action for each possible type of player i.

Bayesian Game

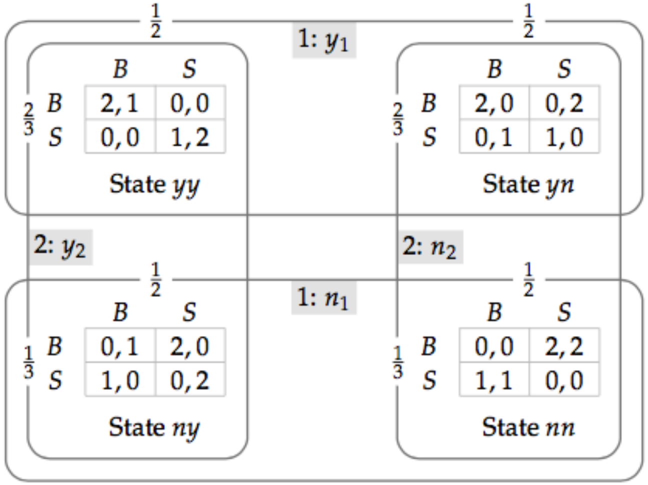

- Recall that player types are drawn from some prior probability distribution $p(\theta_1, \ldots, \theta_I)$.
- Given $p(\theta_1, \ldots, \theta_I)$ we can compute the conditional distribution $p(\theta_{-i} \mid \theta_i)$ using Bayes rule.
  - Hence the label "Bayesian games".
  - Equivalently, when types are not finite, we can compute the conditional distribution $P(\theta_{-i} \mid \theta_i)$ given $P(\theta_1, \ldots, \theta_I)$.
- Player i knows her own type and evaluates her expected payoffs according to the **conditional distribution** $p(\theta_{-i} \mid \theta_i)$, where $\theta_{-i} = (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_I)$.

Bayesian Game

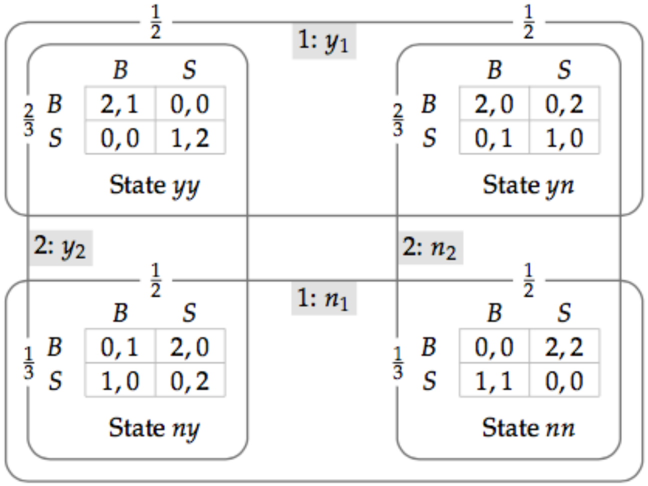

- Since the payoff functions, possible types, and the prior probability distribution are common knowledge, we can compute expected payoffs of player i of type θ_i as

$$U\left(s_i', s_{-i}, \theta_i\right) = \sum_{\theta_{-i}} p(\theta_{-i} \mid \theta_i) u_i(s_i', s_{-i}(\theta_{-i}), \theta_i, \theta_{-i})$$

when types are finite

$$= \int u_i(s_i', s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) P(d\theta_{-i} \mid \theta_i)$$

when types are not finite.

NEP Example

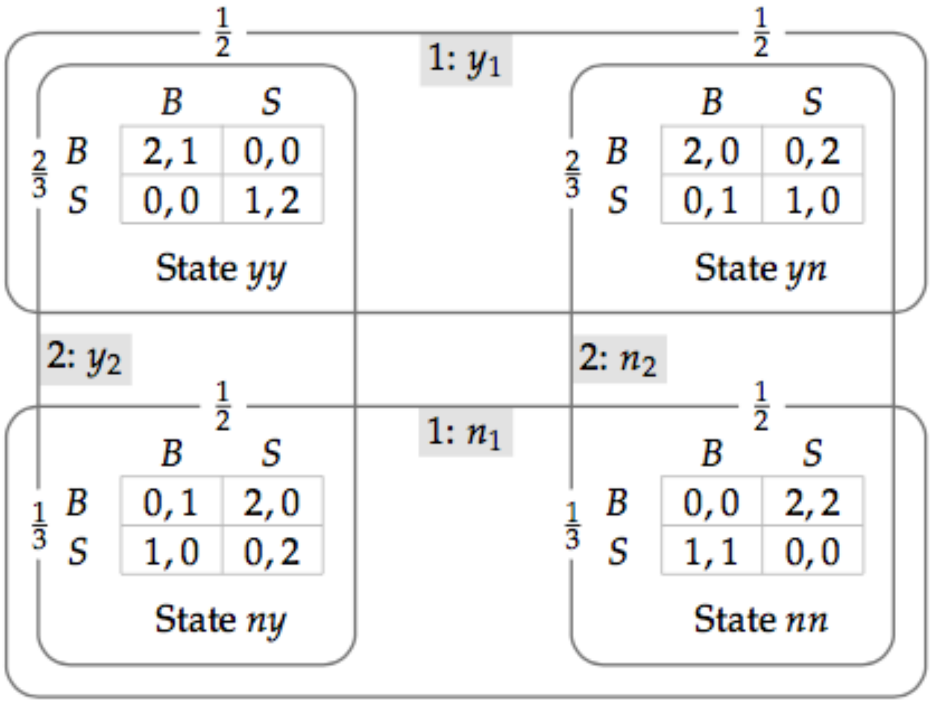

- [(B,B), (B,S)] is a NEP?
- [(S,B), (S,S)] is a NEP?

Bayesian Equilibrium

Definition

(Bayesian Nash Equilibrium) The strategy profile $s(\cdot)$ is a (pure strategy) Bayesian Nash equilibrium if for all $i \in \mathcal{I}$ and for all $\theta_i \in \Theta_i$, we have that

$$s_i(\theta_i) \in \arg\max_{s_i' \in S_i} \sum_{\theta_{-i}} p(\theta_{-i} \mid \theta_i) u_i(s_i', s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}),$$

or in the non-finite case,

$$s_i(\theta_i) \in \arg\max_{s_i' \in S_i} \int u_i(s_i', s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) P(d\theta_{-i} \mid \theta_i)\ .$$

- Hence a Bayesian Nash equilibrium is a Nash equilibrium of the "expanded game" in which each player i's space of pure strategies is the set of maps from Θ_i to S_i.

Bayesian Equilibrium: Existence

Theorem

Consider a finite incomplete information (Bayesian) game. Then a mixed strategy Bayesian Nash equilibrium exists.

Theorem

Consider a Bayesian game with continuous strategy spaces and continuous types. If strategy sets and type sets are compact, payoff functions are continuous and concave in own strategies, then a pure strategy Bayesian Nash equilibrium exists.

Example: Cournot Duopoly

- Suppose that two firms both produce at constant marginal cost.
- Demand is given by $P(Q)$ as in the usual Cournot game.
- Firm 1 has marginal cost equal to C (and this is common knowledge).
- Firm 2's marginal cost is private information. It is equal to C_L with probability θ and to C_H with probability $(1-\theta)$, where $C_L < C_H$.

Summary