Bayesian games

1: Definition and equilibrium

Universidad Carlos III de Madrid

The Bayesian game

- Until now, in all the games we have seen, all the elements of the game are known by all players.
  - Strategies (actions).
  - Order of play.
  - Payoff function.
- This means that, when we define a game all players know which game they are playing.
  - Also: all players know that all players know which game they are playing.
  - Also: every one knows that every one know that every one knows … which game they are playing.
- This is called common knowledge. This means that, once a game is defined, it is common knowledge.

The Bayesian game

- Very often, some of the characteristics of the game is not known by some player.
- This means that some players have private information that is not known by other players.
- In these cases, the models seen until now are not, in principle, adequate.
- A new model: the Bayesian game, where payoffs are not common knowledge.
- The private information over some of the other elements may also be modeled as a case of private information over payoffs.
- E.g., if Player 1 does not know if the strategy set for Player 2 is {A, B} or {A, B, C} we can transform the game into another one where Player 2 has strategy set {A, B, C}, but Player 1 does not know if the payoffs for Player 2 in case she chooses C are the actual ones or some very low payoffs (that would imply that C is never chosen).
- Depending on whether the game about which there is asymmetric information is static or dynamic, the Bayesian game will also be static or dynamic. We will only see the static games.

The Bayesian game

Examples of Bayesian games.

- Cournot duopoly where the rival's costs are unknown.
- Bertrand oligopoly where demand is uncertain, but one of the firms has privileged information about it.
- Auction where other participants' valuations are not known.
- Private contributions to the provision of a public good without knowing other people's costs or valuations.
- Negotiation not knowing the rival's discount factor.
- Battle of the sexes where one player does not know if her companion prefers to be left alone or not.
- Prisoners' dilemma not knowing if the other player has altruistic preferences.

The Bayesian game

In the study of Bayesian games we will see:

- The definition of a Bayesian game to include asymmetric information.
- The use of best replies to define the Bayesian equilibrium.
- The redefinition of a Bayesian game as a dynamic game of imperfect information.
- The equivalence of the Bayesian equilibrium to the Nash equilibrium of the normal form corresponding to that dynamic game.

Example 1

- Player 1 chooses between two actions A and B.
- Player 2 chooses between two actions I and D.
- Payoffs depend on players' types.
- Player 1 has only one type, known to Player 2.
- Player 2 may be of type *x* or of type *y.*
- Player 2 knows her own type, but Player 1 does not know for certain the type of Player 2.
- Player 1 knows that Player 2 is of type *x* with probability 2/3, and of type *y* with probability 1/3.
- Payoffs, depending of the chosen actions and of the type of Player 2, are shown in the next tables.

Example 1

Model "not knowing payoffs" as "not knowing types"

$p = 2/3$ 2 type x

		I	D
1	A	4 , 3	3 , 1
	B	3 , 6	2 , 3

$p = 1/3$ 2 type y

		I	D
1	A	3 , 3	1 , 6
	B	1 , 1	5 , 3

- This is a Bayesian game.
- These are not two normal form game: they are not two independent matrices and cannot be analyzed as such.
- Player 1 does not know the matrix in which they are.
- Player 2 does know the matrix in which they are.

Example 1

- Show the Bayesian equilibrium using the best reply correspondences.
- Start with Player 2, who knows her type (and the type of Player 1):
  - If 2 is of type *x*:
    - The strategy D is dominated by strategy I. Her best reply against any strategy by Player 1 is I.
  - If 2 is of type *y*:
    - The strategy I is dominated by strategy D. Her best reply against any strategy by Player 1 is D.
  - This means that the best reply by Player 2 against any strategy by Player 1 is (I, D).
- Now we can calculate the best replay of Player 1 against (I, D):
  - $U_1(\text{A, ID}) = (2/3)4 + (1/3)1 = 9/3$
  - $U_1(\text{B, ID}) = (2/3)3 + (1/3)5 = 11/3$
- The Bayesian equilibrium is: (B, (I, D)).

Example 1

- The example was quite easy because Player 2 has a dominant strategy. In general, this will not be the case and we will have to calculate all best replies for Player 1 against all possible strategies of Player 2, $S_2 = \{\text{II, ID, DI, DD}\}$.
- From there, we will look for strategies such that one is a best reply against the other and vice versa (we will see this more carefully in Example 2).

Expected payoff from playing A:

$$U_1(\text{A, II}) = (2/3)4 + (1/3)3 = 11/3$$
$$U_1(\text{A, ID}) = (2/3)4 + (1/3)1 = 9/3$$
$$U_1(\text{A, DI}) = (2/3)3 + (1/3)3 = 9/3$$
$$U_1(\text{A, DD}) = (2/3)3 + (1/3)1 = 7/3$$

Expected payoff from playing B:

$$U_1(\text{B, II}) = (2/3)3 + (1/3)1 = 7/3$$
$$U_1(\text{B, ID}) = (2/3)3 + (1/3)5 = 11/3$$
$$U_1(\text{B, DI}) = (2/3)2 + (1/3)1 = 5/3$$
$$U_1(\text{B, DI}) = (2/3)2 + (1/3)1 = 5/3$$

	II	ID	DI	DD
A	$\underline{11/3}$	3	$\underline{3}$	7/3
B	7/3	$\underline{11/3}$	5/3	$\underline{3}$

Representation of the Bayesian game

A Bayesian game must have the following elements:

- The players set: $N = \{1, 2, \dots, n\}$.
- The players' types (one type for each piece of private information a player may have): $T_i = \{t_i^1, t_i^2, \dots, t_i^{m_i}\}$ for all $i \in N$.
- Each type of each player must have defined a probability distribution over combinations types (representing beliefs over other players' types):

  $$p\big((t_1, t_2, \dots, t_{i-1}, t_{i+1}, \dots, t_n)|t_i\big) = p(t_{-i}|t_i)$$

  for all $t_{-i} \in T_{-i}$, for all $t_i \in T_i$ and for all $i \in N$.
- Possible actions for each type. For convenience, we will assume that all types of a player have the same set of actions: $A_{t_i} = A_i = \left(A_i^1, A_i^2, \dots, A_i^{k_i}\right)$ for all $t_i \in T_i$ and for all $i \in N$.
- A strategy for Player i is a vector of actions, one for each type: $\left(a_{t_i^1}, a_{t_i^2}, \dots, a_{t_i^{m_i}}\right)$.
- Payoff functions that depend not only on actions, but also on types:
  - Payoff for type t_i: $u_{t_i}(a_1, a_2, \dots, a_n; t_1, t_2, \dots t_{i-1}, t_{i+1}, \dots t_n)$.
  - Payoff for Player i:

  $$u_{t_i}(a_1, a_2, \dots, a_n; t_1, t_2, \dots, t_n) = \sum_{t_i \in T_i} p(t_i) u_{t_i}(a_1, a_2, \dots, a_n;\ t_1, t_2, \dots, t_{i-1}, t_{i+1}, \dots, t_n).$$

Representation of the Bayesian game

A technical note:

- Strictly speaking, as a primitive element of the Bayesian game we should have a probability distribution over all combinations of types.
- We would calculate all other probabilities after this distribution.
- For example, let us have a game with 3 players with types $T_1 = \{a, b\}$, $T_2 = \{x, y\}$, $T_2 = \{r, t\}$
- Let the probability distribution over types be: $\big(p(a,x,r), p(a,x,t), p(a,y,r), p(a,y,t), p(b,x,r), p(b,x,t), p(b,y,r), p(b,y,t)\big)$. These probabilities are numbers between 0 and 1, with sum 1.
- Now we can calculate, for instance:
  - $p(a) = p(a,x,r) + p(a,x,t) + p(a,y,r) + p(a,y,t)$.
  - $p(x,t|a) = \frac{p(a,x,t)}{p(a)}$.

Representation of the Bayesian game

Let us see the Bayesian representation of Example 1:

- Players, $N = \{1, 2\}$.
- Types: $T_1 = \{1\}$, $T_2 = \{x, y\}$.
- Probabilities over types: each one of the three types has beliefs over others.
  - $p(x|1) = \frac{2}{3}$, $p(y|1) = \frac{1}{3}$,
  - $p(1|x) = 1$,
  - $p(1|y) = 1$.
- Actions: $A_1 = \{\text{A, B}\}$, $A_x = A_y = \{\text{I, D}\}$.
- Strategies:
  - $S_1 = A_1 = \{\text{A, B}\}$,
  - $S_2 = A_x \times A_y = \{\text{II, ID, DI, DD}\}$.
- Payoffs: the 2 matrices in slide 7.

Representation of the Bayesian game

The Bayesian game can be represented through its extensive form

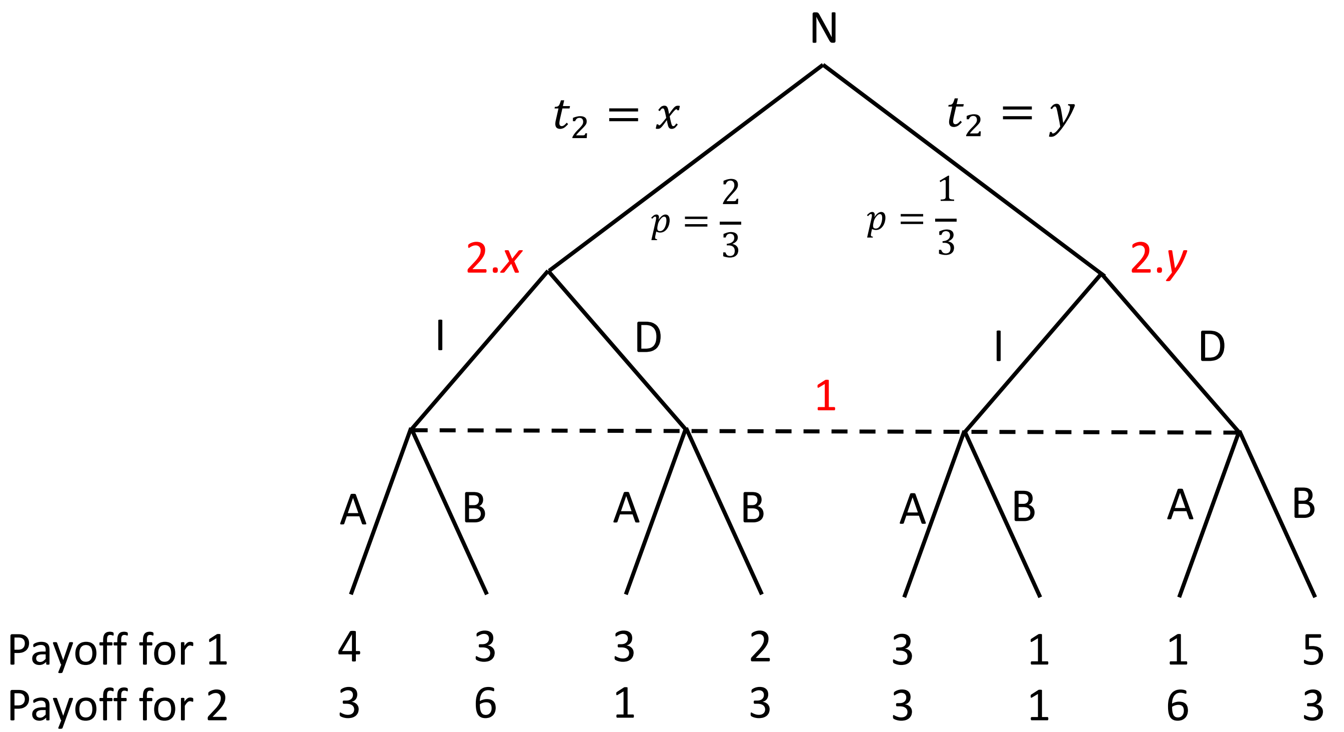

Information sets. Player 1: {1}, Player 2: {2.*x*, 2.*y*}.

Strategies: $S_1 = \{\text{A, B}\}$, $S_2 = \{\text{II, ID, DI, DD}\}$.

Example 2

- The battle of the sexes with asymmetric information.
- She prefers that they go together rather than separated, and she prefers football to opera.
- Preferences for He depend on whether He is stressed out.
  - If He is stressed out He prefers to spend the evening alone.
  - If He is not stressed out (normal) He prefers to spend the evening with She.
  - In both cases He prefers opera to football.
- Payoffs are shown in the next tables.
- She assigns equal probabilities to He being stressed out and being normal.

Example 2

Payoffs if He fells normal

Prob. = 1/2

He

She		F	O
	F	2 , 1	0 , 0
	O	0 , 0	1 , 2

Payoffs if He fells stressed out

Prob. = 1/2

He

She		F	O
	F	2 , 0	0 , 2
	O	0 , 1	1 , 0

Example 2

Let us analyze the best reply for He.

He normal

		F	O
She	F	2, $\underline{1}$	0, 0
	O	0, 0	1, $\underline{2}$

He stressed out

		F	O
She	F	2, 0	0, $\underline{2}$
	O	0, $\underline{1}$	1, 0

If She chooses football, the best reply for He is: football if normal and opera if stressed out.

If She chooses opera, the best reply for He is: opera if normal and football if stressed out.

$$BR_{\mathrm{He}}(\mathrm{F}) = \mathrm{FO},\ BR_{\mathrm{He}}(\mathrm{O}) = \mathrm{OF}.$$

Example 2

Best reply for She.

	FF	FO	OF	OO
F	$\underline{2}$	$\underline{1}$	$\underline{1}$	0
O	0	0,5	0,5	$\underline{1}$

$BR_{\text{She}}(\text{FF}) = \text{F}$
$BR_{\text{She}}(\text{FO}) = \text{F}$
$BR_{\text{She}}(\text{OF}) = \text{F}$
$BR_{\text{She}}(\text{OO}) = \text{O}$

$BR_{\text{He}}(\text{F}) = \text{FO}$
$BR_{\text{He}}(\text{O}) = \text{OF}$

BNE: (football, (football if normal, opera if stressed out)) = (F, FO)

Example 2

Extensive form of the battle of the sexes with asymmetric information.

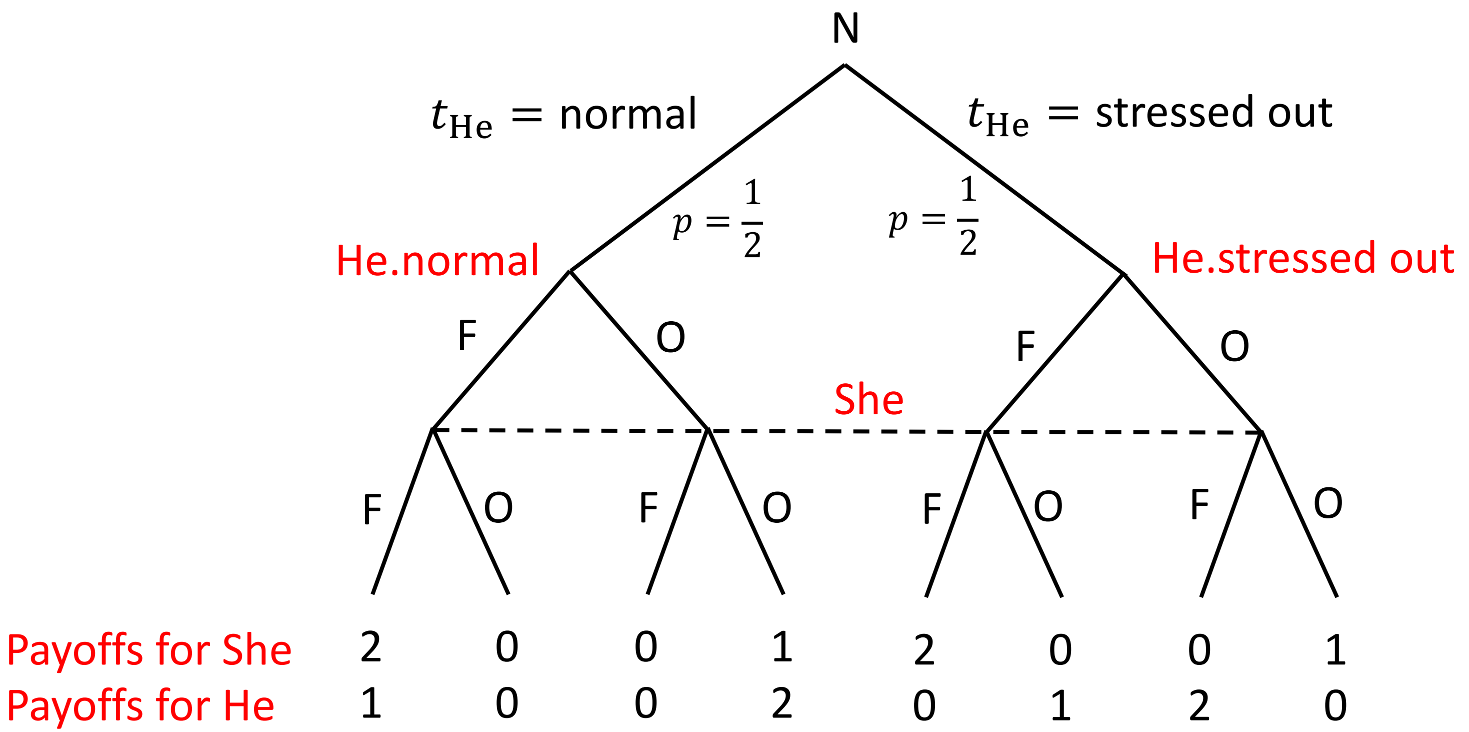

Let us find the SPNE of this extensive form game.

Since there is only one subgame, it will coincide with the NE.

Example 2

Let us produce the normal form:

	FF	FO	OF	OO
F	$\underline{2}$, 0,5	$\underline{1}$, $\underline{1{,}5}$	$\underline{1}$, 0	0, 1
O	0, 0,5	0,5, 0	0,5, $\underline{1{,}5}$	$\underline{1}$, 1

- NE: (F, FO). It coincides with the Bayesian equilibrium.
- In fact, the NE of the normal form that corresponds to the extensive form of a Bayesian game is always the Bayesian equilibrium of the Bayesian game.

Extensions

- She has one type, He has three types. The matrix of the normal form is 2x8.
- She and He have two types each.
- She and He have an arbitrary number of types each.
- There are meta beliefs: He has 2 types, She has beliefs about these two types. He does not know She's believes about He's types, but He has believes about She's believes.
- This may get as much complicated as one wants. According to Harsanyi's theorem, it can always be modeled as we have done, just multiplying the number of types.
- There are more than two players.
- Players choose sequentially. The player playing second can observe the action, but not the type of the player playing first.
- These extensions will not be seen in this course.