Bayesian Games

Roberto Lucchetti

Politecnico di Milano

Summary of the slides

1. Examples of Bayesian games
2. Setting the model
3. Types
4. Beliefs
5. Bayes-Nash equilibria
6. Examples
7. Auctions

Example 1

Player One does not know if player Two wants to spend the evening together at C or at S

Probability $\frac{1}{2}$

$$\begin{pmatrix} (2, 1) & (0, 0) \\ (0, 0) & (1, 2) \end{pmatrix}$$

Probability $\frac{1}{2}$

$$\begin{pmatrix} (2, 0) & (0, 2) \\ (0, 1) & (1, 0) \end{pmatrix}$$

Expected payoffs player One

	(C,C)	(C,S)	(S,C)	(S,S)
C	$\underline{2}$	$\underline{1}$	$\underline{1}$	0
S	0	$\frac{1}{2}$	$\frac{1}{2}$	$\underline{1}$

Nash equilibrium (C, CS), since the best reaction to C of the player 2 willing to meet is C unwilling to meet is S. For player 1, the best reaction to (C, S) is C as seen in the above table

Example 2

Both Players do not know if the other wants to spend the evening together at C or at S. Player One assigns probability $\frac{1}{2}$, $\frac{1}{2}$ to Player Two being "(meeting, alone)", Player Two assigns probability $\frac{2}{3}$, $\frac{1}{3}$ to Player one being "(meeting, alone)"

Here Player One is "meeting"

$$\begin{pmatrix} (2,1) & (0,0) \\ (0,0) & (1,2) \end{pmatrix}, \begin{pmatrix} (2,0) & (0,2) \\ (0,1) & (1,0) \end{pmatrix}.$$

Player One is "alone"

$$\begin{pmatrix} (0,1) & (2,0) \\ (1,0) & (0,2) \end{pmatrix}, \begin{pmatrix} (0,0) & (2,2) \\ (1,1) & (0,0) \end{pmatrix}$$

Expected payoffs of the players
Player One (on the left when meeting)

	(C,C)	(C,S)	(S,C)	(S,S)
C	2	1	1	0
S	0	$\frac{1}{2}$	$\frac{1}{2}$	1

	(C,C)	(C,S)	(S,C)	(S,S)
C	0	1	1	2
S	1	$\frac{1}{2}$	$\frac{1}{2}$	0

Player 2 (on the left when meeting)

	(C,C)	(C,S)	(S,C)	(S,S)
C	1	$\frac{2}{3}$	$\frac{1}{3}$	0
S	0	$\frac{2}{3}$	$\frac{4}{3}$	2

	(C,C)	(C,S)	(S,C)	(S,S)
C	0	$\frac{1}{3}$	$\frac{2}{3}$	1
S	2	$\frac{4}{3}$	$\frac{2}{3}$	0

The way to model the situation

- The variants of the game are possible states of the nature
- We can assume Nature selects the game to play
- The nature sends a signal to the players, telling them which type of players they are (in the first example Player 1 is of one possible type, willing to meet, Player 2 is of two types)
- The players must have probability distribution any time they have uncertainty on the state of the world
- The difference among types of players is (only) in the payoffs

Bayesian games

Definition

A ***Bayesian game*** *is*

1. *a set of players*
2. *a set Ω of states*

To each player are associated:

1. *a set of strategies*
2. *a set of signals assigning a signal to each state*
3. *a probability distribution over the set of states associated with each signal*
4. *a payoff function defined on the pairs (a, ω), where a is a strategy profile and ω is a state.*

Notation

Denote by $\tau_i(\cdot)$ the player i's signal function

Let $\{t_i^1, \ldots, t_i^{k(i)}\} = \tau_i(\Omega)$ be the set of the types of Player i

For each player strategies are the same regardless the type determined by the state

The probability distribution over the set of states associated with each signal received by Player i is called belief of player i

Back to Example 1

1. Players: the pair of people having to decide to meet or not
2. States: $\Omega = \{\omega_1, \omega_2\}$, where we can assume that ω_1 represents "meeting" , ω_2 "alone".
3. Strategies: C and S for both players
4. Player One receives only one signal: $\tau_1(\omega_1) = \tau_1(\omega_2)$, Player Two receives two signals: $\tau_2(\omega_1) \neq \tau_2(\omega_2)$: thus there is only one type of player One, and two types of player Two.
5. Beliefs: player One assigns probability $\frac{1}{2}$ to each of the two states associated with the unique signal received, player Two, types one and two, assigns probability 1 to each (unique) state associated with the signal received
6. Payoffs: bimatrices (1) and (2)

Back to Example 2

1. Players: the pair of people having to decide to meet or not
2. States: $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, where we can assume that ω_1 represents "meeting,meeting" , $\dots \omega_4$ "(alone, alone)".
3. Strategies: C and S for both players
4. Player One receives two signals: $\tau_1(\omega_1) = \tau_1(\omega_2) = t_1^1$, $\tau_1(\omega_3) = \tau_1(\omega_4) = t_1^2$. Player Two receives two signals: $\tau_2(\omega_1) = \tau_2(\omega_3) = t_2^1$, $\tau_2(\omega_2) = \tau_2(\omega_4) = t_2^2$: there are two types (t_1^1, t_1^2) of player One, and two types (t_2^1, t_2^2) of player Two.
5. Beliefs: player One, no matter her type is, assigns probability $\frac{1}{2}$ to each of the two types of player Two,, each type of player Two assigns probability $\frac{2}{3}$ to "player One is meeting", $\frac{1}{3}$ to "player One is "alone "
6. Payoffs: bimatrices (1) (2) (3) (4): the state decides the matrix, a strategy profile the entry of the matrix

Visualizing signals and types in Example 2

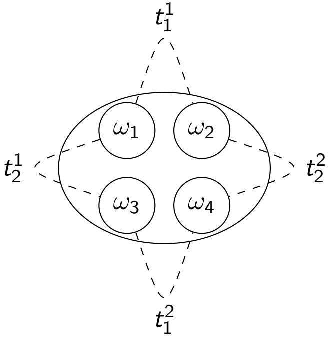

- $\omega_1 = \mathsf{mm}$
- $\omega_2 = \mathsf{ma}$
- $\omega_3 = \mathsf{am}$
- $\omega_4 = \mathsf{aa}$

- $\tau_1(\omega_1) = \tau_1(\omega_2) = t_1^1$
- $\tau_1(\omega_3) = \tau_1(\omega_4) = t_1^2$
- $\tau_2(\omega_1) = \tau_2(\omega_3) = t_2^1$
- $\tau_2(\omega_2) = \tau_2(\omega_4) = t_2^2$

Observe Player One has same probability distributions on types of Player Two, and conversely. But this is specific here, in general this is not the case

Bayes Nash equilibrium: notation

Let $p(\omega, t_i)$ be the probability that type t_i of player i assigns to the state ω

Let $a(j, \tau_j(\omega))$ be the strategy used by player j when she observes signal $\tau_j(\omega)$, let $\hat{a}_j(\omega) = a(j, \tau_j(\omega))$.

The expected payoff of type t_i if she selects strategy a_i and a strategy profile $\hat{a}(\omega)$ is fixed:

$$\sum_{\omega \in \Omega} p(\omega, t_i) u_i((a_i, \hat{a}_{-i}(\omega)), \omega), \tag{1}$$

where $u_i(a, \omega)$ is the utility of player i when the strategy profile a is given, and under the state ω

Bayes Nash equilibrium: definition

Definition

A Nash equilibrium in a Bayesian game is the Nash equilibrium of the following strategic game:

1. *Players: each pair (i, t_i)*
2. *Strategies: each pair (i, t_i) has the set A_i of strategies of Player i in the Bayesian game*
3. *Payoffs: each pair (i, t_i) has payoff defined as in (1)*

$$\sum_{\omega \in \Omega} p(\omega, t_i) u_i((a_i, \hat{a}_{-i}(\omega))\omega)$$

An interesting example: first case

First case: both Players have the same beliefs:

Probability $\frac{1}{2}$ (T, B strategies of the first, a, b, c of the second)

$$\begin{pmatrix} (4,2) & (4,0) & (4,3) \\ (8,8) & (0,0) & (0,12) \end{pmatrix}$$

Probability $\frac{1}{2}$

$$\begin{pmatrix} (4,2) & (4,3) & (4,0) \\ (8,8) & (0,12) & (0,0) \end{pmatrix}$$

Since $\mathrm{BR}_2(T) = \mathrm{BR}_2(B) = a$ and $\mathrm{BR}_1(a) = B$ then the unique equilibrium provides $(8, 8)$

An interesting example: second case

Now Player 1 has the same beliefs as before, Player 2 is informed of the state:

Probability $\frac{1}{2}$ (T, B strategies of the first, a, b, c of the second)

$$\begin{pmatrix} (4,2) & (4,0) & (4,3) \\ (8,8) & (0,0) & (0,12) \end{pmatrix}$$

Probability $\frac{1}{2}$

$$\begin{pmatrix} (4,2) & (4,3) & (4,0) \\ (8,8) & (0,12) & (0,0) \end{pmatrix}$$

Payoffs Player One:

	(a,a)	(a,b)	(a,c)	(b,a)	(b,b)	(b,c)	(c,a)	(c,b)	(c,c)
T	4	4	4	4	4	4	4	4	4
B	8	4	4	4	0	0	4	0	0

An interesting example: conclusion

Payoffs Player One:

	(a,a)	(a,b)	(a,c)	(b,a)	(b,b)	(b,c)	(c,a)	(c,b)	(c,c)
T	4	$\underline{4}$	$\underline{4}$	$\underline{4}$	$\underline{4}$	$\underline{4}$	$\underline{4}$	$\underline{4}$	$\underline{4}$
B	$\underline{8}$	$\underline{4}$	$\underline{4}$	$\underline{4}$	0	0	$\underline{4}$	0	0

Since $\mathrm{BR}_2(T) = (c, b)$ and $\mathrm{BR}_2(B) = (c, b)$, then

$\{(T, (c, b)\}$ unique BN equilibrium with payoffs $(4, 3, 3)$

Auctions

Auctions: remember, there are several types, since ancient times. . .

1. Sequential offers
2. Sealed
3. First price
4. Second price
5. Different termination rules

Auctions: complete information

Suppose there are n bidders, each one has a valuation v of the object, and suppose $v_1 > v_2 > \cdots > v_n$

Each bidder proposes a (non negative) bid. Some rule must handle ties

An assignment rule for the payment must be done

We consider only auctions where the winner is the highest bid

First price

For the first price auction the payment rule is: the winner j offering the bid b_j pays her bid. Other players pay nothing

Game:

1. Players: the n bidders ($n \geq 2$)
2. Strategies: $[0, +\infty)$ for each player
3. Payoffs: let b_i be the bid of player i and let $\hat{b} = \max b_{-i}$. If either $b_i > \hat{b}$ or $b_i = \hat{b}$ and the breaking rule assigns the object to i, the payoff for i is $v_i - b_i$. Otherwise it is 0

Nash equilibria for the first price auction

1. One Nash equilibrium is $(v_2, v_2, v_3, \dots, v_n)$
2. In all equilibria the winner is Player One
3. The two highest bids are the same and one is made by Player One. The highest bid b_1 satisfies $v_2 \leq b_1 \leq v_1$. All such bid profiles are Nash equilibria
4. for i bidding more than v_i is weakly dominated

Second price

For the second price auction the payment rule is: the winner j offering the bid b_j pays the second best bid. Other players pay nothing

Game:

1. Players: the n bidders ($n \geq 2$)
2. Strategies: $[0, +\infty)$ for each player
3. Payoffs: Let b_i be the bid of player i and let $\hat{b} = \max b_{-i}$. If either $b_i > \hat{b}$ or $b_i = \hat{b}$ and the breaking rule assigns the object to i, the payoff for i is $v_i - \hat{b}$. Otherwise it is 0

Nash equilibria for the second price auction

1. One Nash equilibrium is $(v_1, v_2, v_3, \dots, v_n)$
2. Other equilibria: $(v_1, 0, 0, \dots, 0)$, $(v_2, v_1, v_3, \dots, v_n)$
3. A player's bid equalizing her evaluation is a weakly dominant strategy

Auctions with incomplete information

Assumptions

1. There are $\underline{v}, \bar{v}$ such that the evaluation v_i of each Player i fulfills $\underline{v} \leq v_i \leq \bar{v}$
2. each Player knows that all other evaluations are in $[\underline{v}, \bar{v}]$
3. there is a (common) function $f; [0, \infty) \to [0, 1]$, increasing, with $f(\underline{v}) = 0, f(\bar{v}) = 1$ such that the probability that any evaluation is less than v is $f(v)$
4. breaking rule: in case of multiple winners they share the earning (evaluation -bid)

The game

1. Players: the bidders
2. States: all possible profiles $(v_1, \dots, v_n)$ with $v_i \in [\underline{v}, \bar{v}]$
3. $\tau_i(v_1, \dots, v_n) = v_i$
4. Beliefs: Every type of Player i assigns probability $f(v_1) \times f(v_{i-1}) \times f(v_{i+1}) \times f(v_n)$ to the event that the evaluation of Player j is at most v_j
5. Payoffs: For Player i in state $(v_1, \dots, v_n)$ it is 0 if $b_i < \hat{b}$, otherwise it is $\frac{v_i - p(b)}{m}$, where $p(b)$ is the paid price (depending on the type of auction), and m is the number of people bidding $\hat{b}$

Nash equilibria for the first price

1. bid v_i of type v_i weakly dominates greater bids,
2. bid v_i of type v_i is weakly dominated by a lower bid
3. a symmetric BN equilibrium is

$$E(v) = v - \frac{\int_{\underline{v}}^{v} [f(x)]^{n-1}\, dx}{[f(v)]^{n-1}}$$

for $v \in [\underline{v}, \bar{v}]$, showing that $E(v) < v$

The two bidders case and uniform distribution

Suppose the bidders are two and f is uniform distribution in $[0,1]$. In this case the symmetric equilibrium provides $\frac{1}{2}v$

Suppose each type of Player two plays in this way. Then bids of Player two are uniformly distributed in $[0,1/2]$. Thus Player One wins for sure offering more than $1/2$. If she offers $b_1 < 1/2$, she wins if the evaluation of the second player is less than $2b_1$. Thus the payoff function of type v_1 of Player one is

$$\begin{cases} 2b_1(v_1 - b_1) & \textit{if} \quad 0 \leq b_1 \leq 1/2 \\ v_1 - b_1 & \textit{otherwise} \end{cases}$$

The maximum of the payoff function is when $b_1 = \frac{1}{2}v_1$

Nash equilibria for the second price

For every type of every bidder bidding her real evaluation is a weakly dominant strategy

Second price is non manipulable