

Notes on Bayesian Games

Course Econ 201B

Rodolfo G. Campos
UCLA

February 4, 2005

1 Bayesian Games

1.1 Types

The type of a player consists of all the private information that is relevant to decision making. Types are denoted as $\theta_i \in \Theta_i$, drawn from a finite set. A profile of types is $\theta = (\theta_1, \theta_2, \dots, \theta_n)$. In a Bayesian game we will always assume that players know the probability function from which types are drawn, $p(\theta)$. This implies that they agree on the probability of any given type. The type θ_i is observed by player i only, who then updates his beliefs about the likelihood of the opponents' types by calculating the conditional probability $p(\theta_i|\theta_{-i})$.

1.2 Strategies

Since every player knows his own type, a strategy for him assigns an action to each type (since it may be optimal to use different actions for different types). Formally, a strategy is a function from the type space to the action space, $s_i : \Theta_i \rightarrow A_i$, with element $s_i(\theta_i)$.

1.3 Utilities

Players calculate their utilities by taking expectations over types. The notation we will use is $E_\theta[\cdot]$ if the expectation is taken over the complete profile of types, and $E_{-\theta_i}[\cdot|\theta_i]$ if the expectation is over the opponents' types.

2 Bayesian Nash Equilibrium

A Bayesian Nash Equilibrium is the equivalent in a Bayesian game to a Nash Equilibrium in a regular game. The only difference in the definition is that players take expectations over types. In what follows we will consider only pure strategies.

Definition 1 (BNE) *A Bayesian Nash Equilibrium (BNE) is a Nash Equilibrium of a Bayesian Game, i.e.*

$$\sum_{\theta} u_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) p(\theta) \geq \sum_{\theta} u_i(s'_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) p(\theta) \quad (1)$$

for all $s'_i(\theta_i) \in S_i$

Sometimes it is useful to use a different condition for finding a Bayesian Nash Equilibrium. The following proposition shows that there is an equivalent condition that we can check.

Proposition 2 *Condition (1) holds iff the following expression holds*

$$\sum_{\theta_{-i}} u_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) p(\theta_{-i}|\theta_i) \geq \sum_{\theta_{-i}} u_i(s'_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) p(\theta_{-i}|\theta_i) \quad (2)$$

for all $s'_i(\theta_i) \in S_i$ and for all types θ_i occurring with positive probability.

Proof. (\Leftarrow) Realize that equation (1) can be written with the expectation operator

$$E_{\theta} [u_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i})] \geq E_{\theta} [u_i(s'_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i})] \quad (3)$$

Use the Law of Iterated Expectation on both sides to get

$$E_{\theta_i} [E_{\theta_{-i}} [u_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i})] | \theta_i] \geq E_{\theta_i} [E_{\theta_{-i}} [u_i(s'_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i})] | \theta_i] \quad (4)$$

Converting this equation into summations again,

$$\sum_{\theta_i} \left[\sum_{\theta_{-i}} u_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) p(\theta_{-i}|\theta_i) \right] p(\theta_i) \geq \sum_{\theta_i} \left[\sum_{\theta_{-i}} u_i(s'_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) p(\theta_{-i}|\theta_i) \right] p(\theta_i) \quad (5)$$

Now, it is easy to see that equation (2) implies equation (1) because $p(\theta_i) \geq 0$.

(\Rightarrow) (by contradiction) Suppose that equation (1) holds while equation (2) does not. Then $\exists j, \hat{\theta}_j, \hat{s}_j(\hat{\theta}_j) \in S_j$ and $p(\hat{\theta}_j) > 0$ such that

$$\sum_{\theta_{-j}} u_j(s_j(\hat{\theta}_j), s_{-j}(\theta_{-j}), \hat{\theta}_j, \theta_{-j}) p(\theta_{-j}|\hat{\theta}_j) < \sum_{\theta_{-j}} u_j(\hat{s}_j(\hat{\theta}_j), s_{-j}(\theta_{-j}), \hat{\theta}_j, \theta_{-j}) p(\theta_{-j}|\hat{\theta}_j) \quad (6)$$

Construct a new strategy for player j

$$s'_j(\theta_j) = \begin{cases} s_j(\theta_j) & \text{if } \theta_j \neq \hat{\theta}_j \\ \hat{s}_j(\theta_j) & \text{if } \theta_j = \hat{\theta}_j \end{cases} \quad (7)$$

This new strategy must give higher expected utility to player j than $s_j(\theta_j)$, because $s'_j(\theta_j)$ is strictly better than the latter for type $\hat{\theta}_j$ who has positive probability and equal for all other types. This contradicts statement (1). ■

Thanks to this equivalent condition, it will sometimes be possible to find a BNE in a simpler way.¹ The following examples illustrate both methods of finding a BNE.

2.1 Example 1: 2 ways of solving a bayesian game

2.1.1 Setup

Consider the game in which player 2 can be of two types: $\theta_2 \in \{a, b\}$. His type is observed only by him at the beginning of the game and both types are equally likely.

If player 2 is of type a then the payoff matrix is

	L	R
U	10,8	0,9
D	9,0	5,5

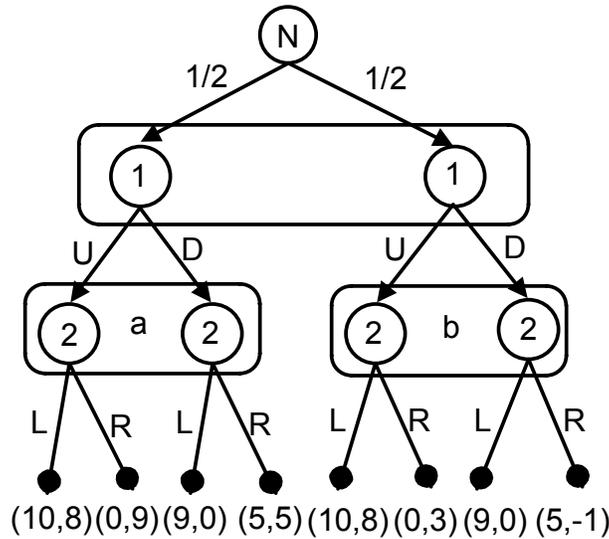
If player 2 is of type b then the payoff matrix is

	L	R
U	10,8	0,3
D	9,0	5,-1

¹Especially in those cases where we are dealing with a large set of types.

2.1.2 First way of finding a BNE: Construction of the normal form

To construct the normal form, it is useful to take a look at the extensive form of the game first (since we know how to construct normal forms from extensive forms)



Strategies are $S_1 = \{U, D\}$ and $S_2 = \{LL, LR, RL, RR\}$. For player 2 the first action is for when he is of type a and the second if he is of type b .

	LL	LR	RL	RR
U	10,8	5,5.5	5,8.5	0,6
D	9,0	7,-0.5	7,2.5	5,2

Now we have to find the best responses

	LL	LR	RL	RR
U	10,8	5,5.5	5, 8.5	0,6
D	9,0	7,-0.5	7,2.5	5,2

Therefore, (D, RL) is a BNE for this game. It is also the only profile that survives iterated elimination of strictly dominated strategies. Therefore, it is the unique BNE of this game.

2.1.3 Second way of finding a BNE: Conditional expectations

Start with player $i = 2$. Since player 1 has only one type there is no need for expectations. In this case,

$$E_{\theta_{-i}} [u_i (s_i (\theta_i), s_{-i} (\theta_{-i}), \theta_i, \theta_{-i})] = u_i (s_i (\theta_i), s_{-i}, \theta_i) \quad (8)$$

Consider $\theta_2 = a$. Then the game is

	L	R
U	10,8	0,9
D	9,0	5,5

Player 2 of type a has a dominant strategy,

$$u_2 (R, s_1, a) > u_2 (L, s_1, a) \quad (9)$$

for $s_1 = U, D$.

Now consider $\theta_2 = b$. Then the game is

	L	R
U	10,8	0,3
D	9,0	5,-1

Player 2 of type b has a dominant strategy again,

$$u_2 (L, s_1, b) > u_2 (R, s_1, b) \quad (10)$$

for $s_1 = U, D$.

Therefore, we have concluded that player 2 will choose R if he is type a and L if he is type b .

Now, let's deal with player $i = 1$. We'll have to deal with expectations

$$E_{\theta_{-i}} [u_i (s_i (\theta_i), s_{-i} (\theta_{-i}), \theta_i, \theta_{-i})] = E_{\theta_{-i}} [u_i (s_i, s_{-i} (\theta_{-i}), \theta_{-i})] \quad (11)$$

Writing it out,

$$E_{\theta_{-i}} [u_i (s_i, s_{-i} (\theta_{-i}), \theta_{-i})] = \frac{1}{2} u_1 (s_1, s_{-i} (a), a) + \frac{1}{2} u_1 (s_1, s_{-i} (b), b)$$

But, we already know how player 2 plays in each of his types.

$$E_{\theta_{-i}} [u_i (s_i, s_{-i} (\theta_{-i}), \theta_{-i})] = \frac{1}{2} u_1 (s_1, R, a) + \frac{1}{2} u_1 (s_1, L, b) \quad (12)$$

The only thing left to do is calculating this expectation for $s_1 = U, D$ and choosing the higher number.

$$\begin{aligned}
E_{\theta_{-i}} [u_i (U, s_{-i} (\theta_{-i}), \theta_{-i})] &= \frac{1}{2} u_1 (U, R, a) + \frac{1}{2} u_1 (U, L, b) \\
&= \frac{1}{2} \times 0 + \frac{1}{2} \times 10 \\
&= 5
\end{aligned}$$

And,

$$\begin{aligned}
E_{\theta_{-i}} [u_i (D, s_{-i} (\theta_{-i}), \theta_{-i})] &= \frac{1}{2} u_1 (D, R, a) + \frac{1}{2} u_1 (D, L, b) \\
&= \frac{1}{2} \times 5 + \frac{1}{2} \times 9 \\
&= 7
\end{aligned}$$

Thus, player 1 will play D.
The BNE is (D, RL) .

2.2 Example 2: Complicating the game

2.2.1 Setup

Modify the game to have the following payoffs.

If player 2 is of type a then the payoff matrix is

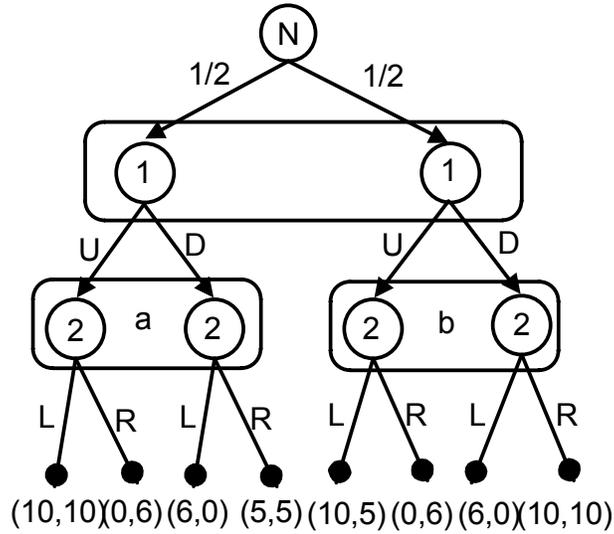
	L	R
U	10,10	0,6
D	6,0	5,5

If player 2 is of type b then the payoff matrix is

	L	R
U	10,10	0,6
D	6,0	10,10

2.2.2 First way of finding a BNE: Construction of the normal form

The extensive form is similar to the case before. We only have to change some of the payoffs.



As before, strategies are $S_1 = \{U, D\}$ and $S_2 = \{LL, LR, RL, RR\}$. For player 2 the first action is for when he is of type a and the second if he is of type b .

	LL	LR	RL	RR
U	10,7.5	5,8	5,5.5	0,6
D	6,0	8,5	5.5,2.5	7.5,7.5

Now we have to find the best responses

	LL	LR	RL	RR
U	10,7.5	5,8	5,5.5	0,6
D	6,0	8,5	5.5,2.5	7.5,7.5

Therefore, (D, RR) is the unique pure BNE for this game.

2.2.3 Second way of finding a BNE: Conditional expectations

Start with player $i = 2$. Since player 1 has only one type there is no need for expectations. In this case,

$$E_{\theta_{-i}} [u_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i})] = u_i(s_i(\theta_i), s_{-i}, \theta_i) \quad (13)$$

Consider $\theta_2 = a$. Then the game is

	L	R
U	10,10	0,6
D	6,0	5,5

The best responses for player 2 are

	L	R
U	10,10	0,6
D	6,0	5,5

Therefore, $BR_2(U) = L$ and $BR_2(D) = R$
 Now consider $\theta_2 = b$. Then the game is

	L	R
U	10,5	0,6
D	6,0	10,10

Now there is a dominant strategy for Player 2 of type b , $s_2 = R$.

Therefore, we have concluded that player 2 will choose $BR_{2a}(U) = L$ and $BR_{2a}(D) = R$ if he is type a and $BR_{2b}(U) = BR_{2b}(D) = R$ if he is type b .

Now, let's deal with player $i = 1$. We'll have to deal with expectations

$$E_{\theta_{-i}} [u_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i})] = E_{\theta_{-i}} [u_i(s_i, s_{-i}(\theta_{-i}), \theta_{-i})] \quad (14)$$

Writing it out,

$$E_{\theta_{-i}} [u_i(s_i, s_{-i}(\theta_{-i}), \theta_{-i})] = \frac{1}{2}u_1(s_1, s_{-i}(a), a) + \frac{1}{2}u_1(s_1, s_{-i}(b), b)$$

In a BNE (just like in a NE) player 2's strategy must be a best response to player 1's strategy. Using that information,

$$E_{\theta_{-i}} [u_i(s_i, s_{-i}(\theta_{-i}), \theta_{-i})] = \frac{1}{2}u_1(s_1, BR_{2a}(s_1), a) + \frac{1}{2}u_1(s_1, BR_{2b}(s_1), b) \quad (15)$$

The only thing left to do is calculating this expectation for $s_1 = U, D$ and choosing the higher number.

$$\begin{aligned} E_{\theta_{-i}} [u_i(U, s_{-i}(\theta_{-i}), \theta_{-i})] &= \frac{1}{2}u_1(U, BR_{2a}(U), a) + \frac{1}{2}u_1(U, BR_{2b}(U), b) \\ &= \frac{1}{2}u_1(U, L, a) + \frac{1}{2}u_1(U, R, b) \\ &= \frac{1}{2} \times 10 + \frac{1}{2} \times 0 \\ &= 5 \end{aligned}$$

And,

$$\begin{aligned} E_{\theta_{-i}} [u_i(D, s_{-i}(\theta_{-i}), \theta_{-i})] &= \frac{1}{2} u_1(D, BR_{2a}(s_1), a) + \frac{1}{2} u_1(D, R, b) \\ &= \frac{1}{2} u_1(D, R, a) + \frac{1}{2} u_1(D, R, b) \\ &= \frac{1}{2} \times 5 + \frac{1}{2} \times 10 \\ &= 7.5 \end{aligned}$$

Thus, player 1 will play D.
The BNE is (D, RR) .