Game Theory

Lecture Notes By

Y. Narahari

Department of Computer Science and Automation

Indian Institute of Science

Bangalore, India

March 2007

Chapter 10: Bayesian Games

Note: *This is a only a draft version, so there could be flaws. If you find any errors, please do send email to* `hari@csa.iisc.ernet.in`. *A more thorough version would be available soon in this space.*

- A game with *incomplete information* is one in which, at the first point in time when the players can begin to plan their moves in the game, some players have *private information* about the game that other players do not know.
- In contrast, in *complete information* games, there is no such private information and all information is publicly known to everybody.
- Incomplete information games are more realistic, more practical.
- The initial private information that a player has *at the first point in time when he begins to plan his moves in the game* is called the *type* of the player. For example, in an auction involving a single indivisible item, each player would have a valuation for the item and typically the player himself would know this valuation deterministically while the other players may only have a guess about how much this player values the item.
- John Harsanyi (Joint Nobel Prize winner in Economic Sciences in 1994 with John Nash and Richard Selten) proposed in 1968, *Bayesian form* games to represent games with incomplete information.

1 Bayesian Games: Definition, Notation, and Examples

A Bayesian game is a tuple $\Gamma = \langle N, (\Theta_i), (S_i), (p_i), (u_i)\rangle$, where the components are described in Table 2.

- $N = \{1, 2, \ldots, n\}$ is a set of players
- Θ_i is the set of types of player i where $i = 1, 2, \ldots, n$.
- S_i is the set of actions or pure strategies of player i where $i = 1, 2, \ldots, n$.

N	A set of players, $\{1, 2, \ldots, n\}$
Θ_i	Set of types of player i
S_i	Set of actions or pure strategies of player i
p_i	A probability (belief) function of player i A function from Θ_i into $\Delta(\Theta_{-i})$
Θ	Set of all type profiles $\times_{i \in N} \Theta_i$
θ	$\theta = (\theta_1, \ldots, \theta_n) \in \Theta$; a type profile
Θ_{-i}	Set of type profiles of agents except i; $\times_{j \neq i \Theta_j}$
θ_{-i}	$\theta_{-i} \in \Theta_{-i}$; a profile of types of agents except i
S	Set of all pure strategy profiles $\times_{i \in N} S_i$
u_i	Utility function of player i; $u_i : S \times \Theta \to \Re$

Table 1: Notation for a Bayesian game

- The probability function p_i is a function from Θ_i into $\Delta(\Theta_{-i})$ the set of probability distributions over Θ_i. That is, for any possible type $\theta_i \in \Theta_i$, p_i specifies a probability distribution $p_i(\bullet|\theta_i)$ over the set Θ_{-i} representing what player i would believe about the types of the other players if his own type were θ_i.
- The payoff function $u_i : \Theta \times S \to \Re$ is such that, for any profile of actions and any profile of types $(\theta, s) \in \Theta \times S$, $u_i(\theta, s)$ specifies the payoff that player i would get, in some *von Neumann - Morgenstern utility scale*, if the players' actual types were all as in Θ and the players all chose their actions according to s.
- Γ is a finite game iff N, $(\Theta_i)_{i \in N}$, and $(S_i)_{i \in N}$ are all finite.
- When we study a Bayesian game, we assume that
  1. Each player i knows the entire structure of the game as defined above
  2. Each player knows his own type in Θ_i
  3. The above facts are common knowledge among all the players in N.
  4. The exact type of a player is not known deterministically to the other players who however have a probabilistic guess of what this type is; note that the belief functions p_i (which are conditional probability distributions) are also common knowledge among the players.
- The phrases *actions* and *strategies* are used differently in the Bayesian game context. A strategy for a player i in Bayesian games is defined as a mapping from Θ_i to S_i. A strategy s_i of a player i, therefore, specifies a pure action for each type of player i. Thus $s_i(\theta_i)$ for a given $\theta_i \in \Theta_i$ would specify the pure action that player i would play if his type were θ_i. The notation $s_i(.)$ is used to refer to the pure action of player i corresponding to an arbitrary type from his type set.

Example: A Bargaining Game

There are two players, player 1 and player 2. Player 1 is the seller of some object and player 2 is the only potential buyer. Each player knows what the object is worth to himself but thinks that its value to the other player may be any integer from 1 to 100 with probability $\frac{1}{100}$. In this game, each player will simultaneously name a bid between 0 and 100 for trading the object. If the buyer's bid is greater than the seller's bid they will trade the object at a price equal to the average of their bids; otherwise no trade occurs. For this game:

$$
\begin{aligned}
&N = \{1,2\} \\
&\Theta_1 = \Theta_2 = \{1,2, \ \ldots \ ,100\} \\
&S_1 = S_2 = \{0,1,2, \ \ldots \ ,100\} \\
&p_i(\theta_{-i}|\theta_i) = \frac{1}{100} \forall i \in N \quad \forall(\theta_i, \theta_{-i}) \in \Theta \\
&u_1(\theta_1,\theta_2,s_1,s_2) = \frac{s_1+s_2}{2} - \theta_1 \quad \text{if } s_2 \geq s_1 \\
&\qquad\qquad\qquad\quad = 0 \qquad\qquad\quad \text{if } s_2 < s_1 \\
&u_2(\theta_1,\theta_2,s_1,s_2) = \theta_2 - \frac{s_1+s_2}{2} \quad \text{if } s_2 \geq s_1 \\
&\qquad\qquad\qquad\quad = 0 \qquad\qquad\quad \text{if } s_2 < s_1
\end{aligned}
$$

Bayesian Games with Infinite Type Sets

- It is often easier to analyze examples with infinite type sets than those with large finite type sets.
- The only notational complication is that, in the infinite case, the probability distributions $p_i(\bullet|\theta_i)$ must be defined on all measurable subsets of Θ_{-i} instead of just individual elements of Θ_{-i}.
- For example, if R_{-i} is a subset of Θ_{-i}, we define $p_i(R_{-i}|\theta_i)$ as the subjective probability that player i would assign to the event that the profile of others' types is in R_{-i}, if his own type were θ_i

Example: Bargaining Game with Continuous Types

Consider the bargaining game as above but with real intervals as type sets. For example, $\Theta_1 = \Theta_2 = S_1 = S_2 = [0,100]$. For each player i and each $\theta_i \in \Theta_i$, let $p_i(\bullet|\theta_i)$ be the uniform distribution over [0,100]. Then for any two numbers x and y such that $0 \leq x \leq y \leq 100$, the probability that any type θ_i of player i would assign to the event that the other player's type is between x and y is:

$$p_i([x,y]|\theta_i) = \frac{y-x}{100}$$

Consistency of Beliefs

We say beliefs $(p_i)_{i\in N}$ in a Bayesian game are *consistent* iff there is some common prior distribution over the set of type profiles Θ such that each player's beliefs given his type are just the conditional probability distributions that can be computed from the prior distribution by the Bayes' formula. In the finite case, beliefs are consistent if $\exists$ some probability distribution $P \in \Delta(\Theta)$ such that

$$p_i(\theta_{-i}|\theta_i) = \frac{P(\theta_i,\theta_{-i})}{\sum\limits_{\theta_{-i}\in\Theta_{-i}} P(\theta_i,\theta_{-i})} \quad \forall \theta \in \Theta; \forall i \in N$$

Consistency simplifies the definition of the model. The common prior on Θ determines all the probability functions. In a consistent model, differences in beliefs among players can be explained by differences in information whereas inconsistent beliefs involve differences of opinion that cannot be derived from any differences in observation and must be simply assumed a priori.

Example: Consistent Beliefs in the Bargaining Game

In the bargaining problem discussed above, the beliefs are consistent with the prior

$$P(\theta) = \frac{1}{10000} \ \ \forall \theta \in \Theta$$

where

$$\Theta = \{1, \ \ldots \ , 100\} \times \{1, \ \ldots \ , 100\}$$

In the infinite version, the beliefs are consistent with a uniform prior on $[0, 100] \times [0, 100]$.

Example: A Game with Inconsistent Beliefs

If it is *common knowledge* that the coaches of two teams in a cricket match believe that his own team has a 2/3 probability of winning the match, then the beliefs of the coaches cannot be consistent with any common prior. It is important to note that, in a consistent model, it can happen that each coach believes that his team has a 2/3 probability of winning but this difference of beliefs cannot be common knowledge.

2 Type Agent Representation and the Selten Game

Richard Selten proposed a representation of Bayesian games that enables a Bayesian game to be transformed to a strategic form game (with complete information). The idea is to represent every possible type of every player as an *agent* or player in the new game. Given a Bayesian game

$$\Gamma = \langle N, (\Theta_i), (S_i), (p_i), (u_i) \rangle$$

the Selten game is an equivalent strategic form game

$$\Gamma^s = \langle N^s, (S^s_j), (U_j) \rangle$$

The idea used by Selten is to have *type agents*. Each player in the original Bayesian game is now replaced with a number of type agents; in fact, a player is replaced by exactly as many type agents as the number of types in the type set of that player. We can safely assume that the type sets of the players are mutually disjoint. The set of players in the Selten game is given by:

$$N^s = \bigcup_{i \in N} \Theta_i$$

Note that each type agent of a particular player can play precisely the same actions as the player himself. This means that for every $\theta_i \in \Theta_i$,

$$S^s_{\theta_i} = S_i$$

From now on, we will use S^s and S interchangeably whenever there is no confusion.

The pay off function U_{θ_i} for each $\theta_i \in \Theta_i$ is the conditionally expected utility to player i in the Bayesian game given that θ_i is his actual type. It is a mapping with the following domain and co-domain:

$$U_{\theta_i} : \underset{i \in N}{\times} \underset{\theta_i \in \Theta_i}{\times} S_i \rightarrow \Re$$

We will explain the way U_{θ_i} is derived using the following example.

Example: Selten Game

Consider the following Bayesian game.

$$\begin{aligned}
& N = \{1, 2\} \\
& \Theta_1 = \{x_1\} \\
& \Theta_2 = \{x_2, y_2\} \\
& S_1 = \{a_1, b_1\} \\
& S_2 = \{a_1, b_2\} \\
& p_1(x_2|x_1) = 0.6 \\
& p_1(y_2|x_1) = 0.4 \\
& p_2(x_1|x_2) = 1 \\
& p_2(x_1|y_2) = 1
\end{aligned}$$

Note that, since Θ_1 is a singleton set, player 1 has only one type, which implies that his type information is common knowledge. To complete the definition of the Bayesian game, we now have to specify the utility functions. Let the utility functions for the two possible type profiles $\theta_1 = x_1,\ \theta_2 = x_2$ and $\theta_1 = x_1,\ \theta_2 = y_2$ be defined as follows.

	2	
1	a_2	b_2
a_1	1,2	0, 1
b_1	0,4	1, 3

u_1 and u_2 for $\theta_1 = x_1; \theta_2 = x_2$

	2	
1	a_2	b_2
a_1	1,3	0, 4
b_1	0,1	1, 2

u_1 and u_2 for $\theta_1 = x_1; \theta_2 = y_2$

This completes the description of the Bayesian game. We now derive the equivalent Selten game:

$$\langle N^s, (S_{\theta_i})_{\substack{\theta_i \in \Theta_i \\ i \in N}}, (U_{\theta_i})_{\substack{\theta_i \in \Theta_i \\ i \in N}} \rangle$$

We have

$$\begin{aligned}
N^s &= \Theta_1 \cup \Theta_2 = \{x_1, x_2, y_2\} \\
S_{x_1} &= S_1 = \{a_1, b_1\} \\
S_{x_2} &= S_{y_2} = S_2 = \{a_2, b_2\}
\end{aligned}$$

Note that

$$U_{\theta_i} : S_1 \times S_2 \times S_2 \to \Re \;\; \forall \theta_i \in \Theta_i, \forall i \in N$$

$$\begin{aligned}
S_1 \times S_2 \times S_2 \;\; = \;\; & \{(a_1, a_2, a_2), (a_1, a_2, b_2), (a_1, b_2, a_2), (a_1, b_2, b_2), (b_1, a_2, a_2), (b_1, a_2, b_2), \\
& (b_1, b_2, a_2), (b_1, b_2, b_2)\}
\end{aligned}$$

The above set gives the set of all strategy profiles of all the type agents. A typical strategy profile can be represented as $(s_{x_1}, s_{x_2}, s_{y_2})$. This could also be represented as $(s_1(.), s_2(.))$ where the strategy s_1 is a mapping from Θ_1 to S_1 and the strategy s_2 is a mapping from Θ_2 to S_2. In general, for an n player Bayesian game, a pure strategy profile is of the form

$$((s_{\theta_1})_{\theta_1 \in \Theta_1}, (s_{\theta_2})_{\theta_2 \in \Theta_2}, \ldots, (s_{\theta_n})_{\theta_n \in \Theta_n})$$

Another way to write this would be $(s_1(.), s_2(.), \ldots, s_n(.))$, where s_i is a mapping from Θ_i to S_i for $i = 1, 2, \ldots, n$.

The payoffs for type agents (in the Selten game) are obtained as conditional expectations over the type profiles of the rest of the agents. For example, let us compute the payoff $U_{x_1}(a_1, a_2, a_2)$, which is the expected payoff obtained by type agent x_1 (belonging to player 1) when this type agent plays action a_1 and the type agents x_2 and y_2 of player 2 play the actions a_2 and a_2 respectively. In this case, the type of player 1 is known but the type of player could be x_2 or y_2 with probabilities given by the belief function $p_1(.|x_1)$. The following conditional expectation gives the required payoff.

$$\begin{aligned}
U_{x_1}(a_1, a_2, a_2) &= p_1(x_2|x_1) u_1(x_1, x_2 a_1, a_2) \\
&\quad + p_1(y_2|x_1) u_1(x_1, y_2, a_1, a_2) \\
&= (1)(0.6) + (1)(0.4) \\
&= 0.6 + 0.4 \\
&= 1
\end{aligned}$$

Similarly, the payoff $U_{x_1}(a_1, a_2, b_2)$ can be computed as follows.

$$\begin{aligned}
U_{x_1}(a_1, a_2, b_2) &= p_1(x_2|x_1) u_1(x_1, x_2, a_1, a_2) \\
&\quad + p_1(y_2|x_1) u_1(x_1, y_2, a_1, b_2) \\
&= (1)(0.6) + (0)(0.4) \\
&= 0.6
\end{aligned}$$

Payoff Computation in Selten Game

In general, given: **(1)** a Bayesian game $\Gamma = \langle N, (\Theta_i), (S_i), (p_i), (u_i) \rangle$, **(2)** its equivalent Selten game $\Gamma^s = \langle N^s, (S_{\theta_i}), (U_{\theta_i}) \rangle$, and **(3)** a pure strategy profile in the Selten game is of the form

$$((s_{\theta_1})_{\theta_1 \in \Theta_1}, (s_{\theta_2})_{\theta_2 \in \Theta_2}, \ldots, (s_{\theta_n})_{\theta_n \in \Theta_n}),$$

the payoffs U_{θ_i} for $\theta_i \in \Theta_i$ $(i \in N)$ are computed as follows.

$$U_{\theta_i}(s_{\theta_i}, s_{\theta_{-i}}) = \sum_{t_{-i} \in \Theta_{-i}} p_i(t_{-i}|\theta_i) u_i(\theta_i, t_{-i}, s_{\theta_i}, s_{t_{-i}})$$

where $s_{t_{-i}}$ is the strategy profile corresponding to the type agents t_i. A concise way of writing the above would be:

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

Another notation which is also often used for $U_{\theta_i}(s_{\theta_i}, s_{\theta_{-i}})$ is $U_i(s_{\theta_i}, s_{\theta_{-i}}|\theta_i)$.

3 Equilibria in Bayesian Games

Pure Strategy Bayesian Nash Equilibrium

A pure strategy Bayesian Nash equilibrium in a Bayesian game

$$\Gamma = \langle N, (\Theta_i), (S_i), (p_i), (u_i) \rangle$$

can be defined in a natural way as a pure strategy Nash equilibrium of the equivalent Selten game. That is, a profile of type agent strategies

$$s^* = ((s^*_{\theta_1})_{\theta_1 \in \Theta_1}, (s^*_{\theta_2})_{\theta_2 \in \Theta_2}, \ldots, (s^*_{\theta_n})_{\theta_n \in \Theta_n})$$

is said to be a pure strategy Bayesian Nash equilibrium of Γ if $\forall i \in N$, $\forall \theta_i \in \Theta_i$,

$$U_{\theta_i}(s^*_{\theta_i}, s^*_{-\theta_i}) \geq U_{\theta_i}(s_i, s^*_{-\theta_i}) \;\; \forall s_i \in S_i$$

Alternatively, a strategy profile $(s_1^*(.), s_2^*(.), \ldots, s_n^*(.))$ is said to be a Bayesian Nash equilibrium if

$$U_{\theta_i}(s_i^*(\theta_i), s^*_{-i}(\theta_{-i})) \geq U_{\theta_i}(s_i, s^*_{-i}(\theta_{-i})) \;\; \forall s_i \in S_i \;\; \forall \theta_i \in \Theta_i \;\; \forall i \in N$$

Example: Pure Strategy Bayesian Nash Equilibrium

Consider the example being discussed. We make the following observations.

- When $\theta_2 = x_2$, the strategy b_2 is strongly dominated by a_2. Thus player 2 chooses a_2 when $\theta_2 = x_2$.
- When $\theta_2 = y_2$, the strategy a_2 is strongly dominated by b_2 and therefore player 2 chooses b_2 for player 2 when $\theta_2 = y_2$.
- When the profiles are (a_1, a_2) or (b_1, b_2), player 1 has payoff 1 regardless of the type of player 2. In all other profiles, pay off of player 1 are zero.
- Since $p_1(x_2|x_1) = 0.6$ and $p_1(y_2|x_1) = 0.4$, player 1 thinks that the type x_2 of player 2 is more likely than type y_2.

The above arguments show that the unique pure strategy Bayesian Nash equilibrium in the above example is given by:

$$(s^*_{x_1} = a_1, s^*_{x_2} = a_2, s^*_{y_2} = b_2)$$

The above example illustrates the danger of analyzing each matrix separately as shown by the arguments below.

- If it is common knowledge that player 2's type is x_2, then the unique Nash equilibrium is (a_1, a_2). If it is common knowledge that player 2 has type y_2, then we get (b_1, b_2) as the unique Nash equilibrium.
- In a Bayesian game, the type of player 2 is not common knowledge and hence the above prediction based on analyzing the matrices separately is wrong.
- Note that if player 2's type is x_2, then the preferred strategy for player 1 is a_1 while it is b_1 if player 2's type is y_2.
- This implies that player 1 could not behave as predicted unless he received some information from player 2.
- Thus the prediction (a_1, a_2) if player 2's type is x_2 and (b_1, b_2) if player 2's type is y_2 can be fulfilled only if appropriate communication is added to the structure of the game.
- Another important point to note is that player 2 prefers (b_1, b_2) over (a_1, a_2) if $\theta_2 = x_2$ while he prefers (a_1, a_2) over (b_1, b_2) otherwise. This implies that even if communication is allowed, player 2 would not be willing to communicate the information that is necessary to fulfill this prediction because it would always give him the outcome that he prefers less. Player 2 would prefer to manipulate his communications to get the outcome (b_1, b_2) if his type is x_2 and the outcome (a_1, a_2) otherwise.

Dominant Strategy Equilibria

The dominant strategy equilibria of Bayesian games can again be defined using the Selten game representation. Given a Bayesian game

$$\Gamma = \langle N, (\Theta_i), (S_i), (p_i), (u_i) \rangle$$

a profile of type agent strategies $(s_1^*(.), s_2^*(.), \ldots, s_n^*(.))$ is said to be a strongly dominant strategy equilibrium if

$$U_{\theta_i}(s_i^*(\theta_i), s_{-i}(\theta_{-i})) > U_{\theta_i}(s_i, s_{-i}(\theta_{-i})) \;\; \forall s_i \in S_i \setminus \{s_i^*(\theta_i)\} \;\; \forall \theta_i \in \Theta_i \;\; \forall \theta_{-i} \in \Theta_{-i} \;\; \forall i \in N$$

Similarly, a profile of type agent strategies $(s_1^*(.), s_2^*(.), \ldots, s_n^*(.))$ is said to be a weakly dominant strategy equilibrium if

$$U_{\theta_i}(s_i^*(\theta_i), s_{-i}(\theta_{-i})) \geq U_{\theta_i}(s_i, s_{-i}(\theta_{-i})) \;\; \forall s_i \in S_i \;\; \forall \theta_i \in \Theta_i \;\; \forall \theta_{-i} \in \Theta_{-i} \;\; \forall i \in N$$

with strict inequality satisfied for at least one $s_i in S_i$. The notion of dominant strategy equilibrium is independent of the belief functions and this is what makes it a very powerful notion and a very strong property. The notion of a weakly dominant strategy equilibrium is used extensively in mechanism design theory to define *dominant strategy implementation* of mechanisms.

Mixed Strategy Nash Equilibrium

Consider a Bayesian game

$$\Gamma = \langle N, (\Theta_i), (S_i), (p_i), (u_i) \rangle$$

. A mixed strategy profile (or randomized strategy profile) associates a mixed strategy to every type agent:

$$\sigma \in \underset{i \in N}{\times} \underset{\theta_i \in \Theta_i}{\times} \Delta(S_i)$$

This implies any σ such that

$$\sigma = ((\sigma_i(s_i|\theta_i))_{s_i \in S_i})_{\substack{\theta_i \in \Theta_i \\ i \in N}}$$

such that

$$\sigma_i(s_i|\theta_i) \geq 0 \qquad \forall s_i \in S_i \quad \forall \theta_i \in \Theta_i \quad \forall i \in N$$
$$\sum_{s_i \in S_i} \sigma_i(s_i|\theta_i) = 1 \qquad \forall \theta_i \in \Theta_i \quad \forall i \in N$$

$\sigma_i(s_i|\theta_i)$ is the conditional probability that a player i would use action s_i if his type were θ_i. The mixed strategy for type θ_i of player i is:

$$\sigma_i(\bullet|\theta_i) = (\sigma_i(s_i|\theta_i))_{s_i \in S_i}$$

A randomized strategy profile

$$\sigma^* = (\sigma^*_{\theta_i})_{\substack{\theta_i \in \Theta_i \\ i \in N}} \in \underset{i \in N}{\times} \underset{\theta_i \in \Theta_i}{\times} \Delta(S_i)$$

is said to be a mixed strategy Bayesian Nash equilibrium if $\forall i \in N$, $\forall \theta_i \in \Theta_i$,

$$U_{\theta_i}(\sigma^*_{\theta_i}, \sigma^*_{-\theta_i}) \geq U_{\theta_i}(\sigma_{\theta_i}, \sigma^*_{-\theta_i}) \ \ \forall \sigma_{\theta_i} \in \Delta(S_i)$$

A Bayesian Nash equilibrium specifies for each type of each player a randomized strategy such that each type of each player would be maximizing his own expected utility when he knows his own type but does not know the types of other players.

4 Purification of Randomized Strategies in Equilibria

Consider the following game:

	2	
1	L	R
T	0,0	0, −1
B	1,0	−1,3

The unique (mixed strategy) equilibrium of this game is

$$\sigma_1(T) = 0.75\sigma_1(B) = 0.25$$
$$\sigma_2(L) = 0.5\sigma_2(R) = 0.5$$

- The necessity for player 1 to randomly choose among T and B with probabilities 0.75 and 0.25 might not seem to coincide with any compulsion that people experience in real life.
- Of course, if player 1 thinks that player 2 is equally likely to choose either L or R, then player 1 is willing to randomize. But what could make player 1 actually want to use exact probabilities 0.75 and 0.25?

- Harsanyi (1973) showed that Nash equilibria that involve randomized strategies can be interpreted as limits of Bayesian equilibria in which each player is (almost) always choosing his uniquely optimal action.
- Harsanyi's idea is to modify the game slightly so that each player has some private information about his own payoffs.

Suppose the example game is modified slightly as follows.

	2	
1	L	R
T	$\varepsilon\bar{\alpha}$, $\varepsilon\bar{\beta}$	$\varepsilon\bar{\alpha}$, -1
B	1, $\varepsilon\bar{\beta}$	-1, 3

This is now a game with incomplete information.

- ε is some given number such that $0 < \varepsilon < 1$ and $\bar{\alpha}, \bar{\beta}$ are i.i.d. random variables, uniform over [0,1].
- when the game is played, player 1 knows the value of $\bar{\alpha}$ but not $\bar{\beta}$ and player 2 knows the value of $\bar{\beta}$ but not $\bar{\alpha}$.
- ε is some very small positive number and note that the table becomes the original previous table when $\varepsilon = 0$. Then $\bar{\alpha}$ and $\bar{\beta}$ represent minor factors that have a small influence on the players' payoffs when T or L is chosen.
- Notice that T becomes better for player 1 as $\bar{\alpha}$ increases and L becomes better for player 2 as $\bar{\beta}$ increases.
- Thus there should exist numbers p and q such that player 1 chooses T if $\bar{\alpha}$ is greater than p and chooses B if $\bar{\alpha} < p$. Similarly player 2 chooses L if $\bar{\beta} > q$ and chooses R if $\bar{\beta} < q$.
- Then from player 1's perspective, the probability that 2 will choose L is $1 - q$.
- To make player 1 indifferent between T and B at the critical value of $\bar{\alpha} = p$, we need

$$\varepsilon p = 1(1 - q) + (-1)q$$
$$\varepsilon p = 1 - 2q$$

- Similarly, to make player 2 indifferent between L and R at the critical value $\bar{\beta} = q$, we need

$$\varepsilon q = (-1)(1 - p) + (3)p$$
$$\varepsilon q = 4p - 1$$

- The solution to the above two equations is

$$p = \frac{2 + \varepsilon}{8 + \varepsilon 2}$$
$$q = \frac{4 + \varepsilon}{8 + \varepsilon 2}$$

- There is a unique Bayesian equilibrium for this game:
  - player 1 chooses T if he observes $\bar{\alpha} > (2+\varepsilon)/(8+\varepsilon^2)$ and he chooses B otherwise
  - Player 2 chooses L if she observes $\bar{\beta} > (4-\varepsilon)/(8+\varepsilon^2)$ and she chooses R otherwise.
- Thus the Bayesian equilibrium satisfies

$$\begin{aligned} \sigma_1(\bullet|\bar{\alpha}) &= [T] \text{ if } \bar{\alpha} > \frac{2+\varepsilon}{8+\varepsilon^2} \\ &= [B] \text{ if } \bar{\alpha} < \frac{2+\varepsilon}{8+\varepsilon^2} \end{aligned}$$

$$\begin{aligned} \sigma_2(\bullet|\bar{\beta}) &= [L] \text{ if } \bar{\beta} > \frac{4-\varepsilon}{8+\varepsilon^2} \\ &= [R] \text{ if } \bar{\beta} < \frac{4+\varepsilon}{8+\varepsilon^2} \end{aligned}$$

$\sigma_1(\bullet|\bar{\alpha})$ and $\sigma_2(\bullet|\bar{\beta})$ can be chosen arbitrarily in the zero-probability events that

$$\bar{\alpha} = \frac{2+\varepsilon}{8+\varepsilon^2} \quad \bar{\beta} = \frac{4-\varepsilon}{8+\varepsilon^2}$$

- when player 2 uses the equilibrium strategy σ_2 in this game, player 1 would be indifferent between T and B only if

$$\bar{\alpha} = \frac{2+\varepsilon}{8+\varepsilon^2}$$

  otherwise his expected utility is uniquely maximized by the action designated for him by $\sigma_1(\bullet|\bar{\alpha})$.
- When player 1 uses the strategy σ_1 in this game, player 2 would be indifferent between L and R only if

$$\bar{\beta} = \frac{4-\varepsilon}{8+\varepsilon^2}$$

- Otherwise her expected utility is uniquely maximized by the action designated for her by $\sigma_2(\bullet|\beta)$.
- Thus, each player's expected behavior makes the other player almost indifferent between his two actions; therefore the minor factor that he observes independently can determine the unique optimal action for him.
- Notice that as $\varepsilon \to 0$, this Bayesian Nash equilibrium converges to the unique mixed strategy Nash equilibrium of the original game.
- In general, when we study an equilibrium that involves randomized strategies, we can interpret each player's randomization as depending on minor factors that have been omitted from the description of the game.
- When a game has no equilibrium in pure strategies, we should expect that a player's optimal strategy may be determined by some minor factors that he observes independently of the other players. Thus minor private information may be decisive.