Game Theory

Lecture Notes By

Y. Narahari

Department of Computer Science and Automation

Indian Institute of Science

Bangalore, India

August 2012

Chapter 13: 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.*

We have so far studied strategic form games with complete information. We will now study games with incomplete information, which are crucial to the theory of mechanism design. In particular, we study Bayesian games and introduce different equilibrium notions in Bayesian games such as Bayesian Nash equilibrium.

1 Games with Incomplete Information

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. Clearly, incomplete information games are more realistic, more practical.

The initial private information that a player has, just before making a move 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 Reinhard Selten) proposed in 1968, *Bayesian form* games to represent games with incomplete information.

Definition 1 (Bayesian Game) . *A Bayesian game* Γ *is defined as a tuple* $\Gamma = \langle N, (\Theta_i), (S_i), (p_i), (u_i) \rangle$ *where*

- $N = \{1, 2, \ldots, n\}$ *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* $i = 1, 2, \ldots, n$.

- *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 $p_i(.|\theta_i)$ 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 $u_i : \Theta \times S \rightarrow \mathbb{R}$ 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.*

The notation for Bayesian games is described in Table 1.

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
Θ	Set of all type profiles $= \Theta_1 \times \Theta_2 \times \ldots \times \Theta_n$
θ	$\theta = (\theta_1, \ldots, \theta_n) \in \Theta$; a type profile
Θ_{-i}	Set of type profiles of agents except $i = \Theta_1 \times \ldots \Theta_{i-1} \times \Theta_{i+1} \times \ldots \times \Theta_n$
θ_{-i}	$\theta_{-i} \in \Theta_{-i}$; a profile of types of agents except i
S	Set of all action profiles $= S_1 \times S_2 \times \ldots \times S_n$
p_i	A probability (belief) function of player i A function from Θ_i into $\Delta(\Theta_{-i})$
u_i	Utility function of player i; $u_i : \Theta \times S \rightarrow \mathbb{R}$

Table 1: Notation for a Bayesian game

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 $\theta_i \in \Theta_i$.
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. The belief functions p_i describe these conditional probabilities. Note that the belief functions p_i 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; $s_i(\theta_i)$ 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.

Definition 2 (Consistency of Beliefs) . *We say beliefs $(p_i)_{i \in N}$ in a Bayesian game are* consistent *if 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.*

If the game is finite, beliefs are consistent if there exists some probability distribution $P \in \Delta(\Theta)$ such that

$$p_i(\theta_{-i}|\theta_i) = \frac{P(\theta_i, \theta_{-i})}{\sum_{t_{-i} \in \Theta_{-i}} P(\theta_i, t_{-i})} \quad \forall \theta_i \in \Theta_i; \ \forall \theta_{-i} \in \Theta_{-i}; \ \forall i \in N.$$

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.

2 Examples of Bayesian Games

2.1 A Two Player Bargaining Game

This example is taken from the book by Myerson [1]. There are two players, player 1 and player 2. Player 1 is the seller of some object, and player 2 is a 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 $\frac{1}{100}$. Assume that each player will simultaneously announce a bid between 0 and 100 for trading the object. If the buyer's bid is greater than or equal to 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:

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

Note that the type of the seller indicates the willingness to sell (minimum price at which the seller is prepared to sell the item), and the type of the buyer indicates the willingness to pay (maximum price the buyer is prepared to pay for the item). Also, note that the beliefs are consistent with the prior:

$$P(\theta_1, \theta_2) = \frac{1}{10000} \quad \forall \theta_1 \in \Theta_1 \ \ \forall \theta_2 \in \Theta_2$$

where

$$\Theta_1 \times \Theta_2 = \{1, \ \ldots \ , 100\} \times \{1, \ \ldots \ , 100\}.$$

2.2 A Sealed Bid Auction

Consider a seller who wishes to sell an indivisible item through an auction. Let there be two prospective buyers who bid for this item. The buyers have their individual valuations for this item. These valuations could be considered as the types of the buyers. Here the game consists of the two bidders, namely the buyers, so $N = \{1, 2\}$. The two bidders submit bids, say s_1 and s_2 for the item. Let us say that the one who bids higher is awarded the item with a tie resolved in favor of bidder 1. The winner determination function therefore is:

$$\begin{aligned} f_1(s_1, s_2) &= 1 \text{ if } s_1 \geq s_2 \\ &= 0 \text{ if } s_1 < s_2 \\ \\ f_2(s_1, s_2) &= 1 \text{ if } s_1 < s_2 \\ &= 0 \text{ if } s_1 \geq s_2. \end{aligned}$$

Assume that the valuation set for each buyer is the real interval $[0, 1]$ and also that the strategy set for each buyer is again $[0, 1]$. This means $\Theta_1 = \Theta_2 = [0, 1]$ and $S_1 = S_2 = [0, 1]$. If we assume that each player believes that the other player's valuation is chosen according to an independent uniform distribution, then note that

$$p_i([x, y]|\theta_i) = y - x \quad \forall \ 0 \leq x \leq y \leq 1; \quad i = 1, 2.$$

In a first price auction, the winner will pay what is bid by her, and therefore the utility function of the players is given by

$$u_i(\theta_1, \theta_2, s_1, s_2) = f_i(s_1, s_2)(\theta_i - s_i); \quad i = 1, 2.$$

This completes the definition of the Bayesian game underlying a first price auction involving two bidders. One can similarly develop the Bayesian game for the second price sealed bid auction.

2.3 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 $p_i(.|\theta_i)$ 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 $p_i(.|\theta_i)$ 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:

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

3 Type Agent Representation and the Selten Game

This is a representation of Bayesian games that enables a Bayesian game to be transformed to a strategic form game (with complete information). 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

$$\Gamma^s = \langle N^s, (S_j^s), (U_j) \rangle.$$

The idea used in formulating a Selten game 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 $\theta_i \in \Theta_i$,

$$S_{\theta_i}^s = S_i.$$

The payoff 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:

$$U_{\theta_i} : \underset{i \in N}{\times} \; \underset{\theta_i \in \Theta_i}{\times} S_i \rightarrow \mathbb{R}.$$

We will explain the way U_{θ_i} is derived using an example. This example is developed, based on the illustration in the book by Myerson [1].

3.1 Selten Game for a Bayesian Pricing Game

Consider two firms, company 1 and company 2. Company 1 produces a product x_1 whereas company 2 produces either product x_2 or product y_2. The product x_2 is somewhat similar to product x_1 while the product y_2 is a different line of product. The product to be produced by company 2 is a closely guarded secret, so it can be taken as private information of company 2. We thus have $N = \{1, 2\}$, $\Theta_1 = \{x_1\}$, and $\Theta_2 = \{x_2, y_2\}$. Each firm has to choose a price for the product it produces, and this is the strategic decision to be taken by the company. Company 1 has the choice of choosing a low price a_1 or a high price b_1 whereas company 2 has the choice of choosing a low price a_2 or a high price b_2. We therefore have $S_1 = \{a_1, b_1\}$ and $S_2 = \{a_2, b_2\}$. The type of company 1 is common knowledge since Θ_1 is a singleton. Therefore, the belief probabilities of company 2 about company 1 are given by $p_2(x_1|x_2) = 1$ and $p_2(x_1|y_2) = 1$. Let us assume the belief probabilities of company 1 about company 2 to be $p_1(x_2|x_1) = 0.6$ and $p_1(y_2|x_1) = 0.4$. 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 given as in Tables 2 and 3.

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

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

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

Table 2: 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

Table 3: u_1 and u_2 for $\theta_1 = x_1; \theta_2 = y_2$

We have

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

Note that

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

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

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.

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

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

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

It can be similarly shown that

$$\begin{aligned} U_{x_1}(b_1, a_2, a_2) &= 0 \\ U_{x_1}(b_1, a_2, b_2) &= 0.4 \\ U_{x_2}(a_1, a_2, b_2) &= 2 \\ U_{x_2}(a_1, b_2, b_2) &= 1 \\ U_{y_2}(a_1, a_2, b_2) &= 4 \\ U_{y_2}(a_1, a_2, a_2) &= 3. \end{aligned} \tag{1}$$

From the above, we see that

$$\begin{aligned} U_{x_1}(a_1, a_2, b_2) &> U_{x_1}(b_1, a_2, b_2) \\ U_{x_2}(a_1, a_2, b_2) &> U_{x_2}(a_1, b_2, b_2) \\ U_{y_2}(a_1, a_2, b_2) &> U_{y_2}(a_1, a_2, a_2). \end{aligned} \tag{2}$$

From this, we can conclude that the action profile (a_1, a_2, b_2) is a Nash equilibrium of the type agent representation.

3.2 Payoff Computation in Selten Game

From now on, when there is no confusion, we will use u instead of U. 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)** an action profile in the type agent representation 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 in t_{-i}. A concise way of writing the above would be:

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

The notation u_{θ_i} refers to the utility of player i conditioned on the type being equal to θ_i. We will be using this notation frequently in this section. With this setup, we now look into the notion of an equilibrium in Bayesian games.

3.3 Equilibria in Bayesian Games

Definition 3 . *[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*

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

3.4 Example 1: Bayesian Pricing Game

Consider the Bayesian pricing game 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 when $\theta_2 = y_2$.
- When the action 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, payoff of player 1 is 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:

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

thus validating what we have already shown. Note that the equilibrium strategy for company 1 is always to price the product low whereas for company 2, the equilibrium strategy is to price it low if it produces x_2 and to price it high if it produces y_2.

The above example also illustrates the danger of analyzing each matrix separately. 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. However, in a Bayesian game, the type of player 2 is not common knowledge, and hence the above prediction based on analyzing the matrices separately would be wrong.

3.5 Example 2: First Price Sealed Bid Auction

Consider an auctioneer or a seller and two potential buyers as in Example 2.23. Here each buyer submits a sealed bid, $s_i \geq 0$ $(i = 1, 2)$. The sealed bids are looked at, and the buyer with the higher bid is declared the winner. If there is a tie, buyer 1 is declared the winner. The winning buyer pays to the seller an amount equal to his bid. The losing bidder does not pay anything.

Let us make the following assumptions:

1. θ_1, θ_2 are independently drawn from the uniform distribution on $[0, 1]$.
2. The sealed bid of buyer i takes the form $s_i(\theta_i) = \alpha_i \theta_i$, where $\alpha_i \in [0, 1]$. This assumption implies that player i bids a fraction α_i of his value; this is a reasonable assumption that implies a linear relationship between the bid and the value.

Buyer 1's problem is now to bid in a way to maximize his expected payoff:

$$\max_{s_1 \geq 0} (\theta_1 - s_1) P\{s_2(\theta_2) \leq s_1\}.$$

Since the bid of player 2 is $s_2(\theta_2) = \alpha_2 \theta_2$ and $\theta_2 \in [0, 1]$, the maximum bid of buyer 2 is α_2. Buyer 1 knows this and therefore $s_1 \in [0, \alpha_2]$. Also,

$$\begin{aligned} P\{s_2(\theta_2) \leq s_1\} &= P\{\alpha_2 \theta_2 \leq s_1\} \\ &= P\{\theta_2 \leq \frac{s_1}{\alpha_2}\} \\ &= \frac{s_1}{\alpha_2} \text{ (since } \theta_2 \text{ is uniform over } [0, 1]). \end{aligned}$$

Thus buyer 1's problem is:

$$\max_{s_1 \in [0, \alpha_2]} (\theta_1 - s_1) \frac{s_1}{\alpha_2}.$$

The solution to this problem is

$$s_1(\theta_1) = \begin{cases} \frac{1}{2}\theta_1 & \text{if } \frac{1}{2}\theta_1 \leq \alpha_2 \\ \alpha_2 & \text{if } \frac{1}{2}\theta_1 > \alpha_2. \end{cases}$$

e can show on similar lines that

$$s_2(\theta_2) = \begin{cases} \frac{1}{2}\theta_2 & \text{if } \frac{1}{2}\theta_2 \leq \alpha_1 \\ \alpha_1 & \text{if } \frac{1}{2}\theta_2 > \alpha_1. \end{cases}$$

Let $\alpha_1 = \alpha_2 = \frac{1}{2}$. Then we get

$$\begin{aligned} s_1(\theta_1) &= \frac{\theta_1}{2} \qquad \forall\ \theta_1 \in \Theta_1 = [0, 1] \\ s_2(\theta_2) &= \frac{\theta_2}{2} \qquad \forall\ \theta_2 \in \Theta_2 = [0, 1]. \end{aligned}$$

Note that if $s_2(\theta_2) = \frac{\theta_2}{2}$, the best response of buyer 1 is $s_1(\theta_1) = \frac{\theta_1}{2}$ and vice-versa. Hence the profile $\left(\frac{\theta_1}{2}, \frac{\theta_2}{2}\right)$ is a Bayesian Nash equilibrium.

3.6 Dominant Strategy Equilibria

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

Definition 4 (Strongly Dominant Strategy Equilibrium) . *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*

$$u_{\theta_i}(s_i^*(\theta_i), s_{-i}(\theta_{-i})) > u_{\theta_i}(s_i, s_{-i}(\theta_{-i}))$$

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

Definition 5 (Weakly Dominant Strategy Equilibrium) . *A profile of type agent strategies* $(s_1^*(.), s_2^*(.), \ldots, s_n^*(.))$ *is said to be a weakly dominant strategy equilibrium if*

$$u_{\theta_i}(s_i^*(\theta_i), s_{-i}(\theta_{-i})) \geq u_{\theta_i}(s_i, s_{-i}(\theta_{-i}))$$

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

and strict inequality satisfied for at least one $s_i \in S_i$.

Definition 6 (Very Weakly Dominant Strategy Equilibrium) . *A profile of type agent strategies* $(s_1^*(.), s_2^*(.), \ldots, s_n^*(.))$ *is said to be a very weakly dominant strategy equilibrium if*

$$u_{\theta_i}(s_i^*(\theta_i), s_{-i}(\theta_{-i})) \geq u_{\theta_i}(s_i, s_{-i}(\theta_{-i}))$$

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

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*. Often very weakly dominant strategy equilibrium is used in these settings.

3.7 Example: Weakly Dominant Strategy Equilibrium of Second Price Auction

We have shown above that the first price sealed bid auction has a Bayesian Nash equilibrium. Now we consider the second price sealed bid auction with two bidders and show that it has a weakly dominant strategy equilibrium. Let us say buyer 2 announces his bid as $\hat{\theta}_2$. There are two cases.

1. $\theta_1 \geq \hat{\theta}_2$.
2. $\theta_1 < \hat{\theta}_2$.

Case 1: $\theta_1 \geq \hat{\theta}_2$

Let $\hat{\theta}_1$ be the announcement of buyer 1. Here there are two cases.

- If $\hat{\theta}_1 \geq \hat{\theta}_2$, then the payoff for buyer 1 is $\theta_1 - \hat{\theta}_2 \geq 0$.
- If $\hat{\theta}_1 < \hat{\theta}_2$, then the payoff for buyer 1 is 0.
- Thus in this case, the maximum payoff possible is $\theta_1 - \hat{\theta}_2 \geq 0$.

If $\hat{\theta}_1 = \theta_1$ (that is, buyer 1 announces his true valuation), then payoff for buyer 1 is $\theta_1 - \hat{\theta}_2$, which happens to be the maximum possible payoff as shown above. Thus announcing θ_1 is a best response to buyer 1 whatever the announcement of buyer 2.

Case 2: $\theta_1 < \hat{\theta}_2$

Here again there are two cases: $\hat{\theta}_1 \geq \hat{\theta}_2$ and $\hat{\theta}_1 < \hat{\theta}_2$.

- If $\hat{\theta}_1 > \hat{\theta}_2$, then the payoff for buyer 1 is $\theta_1 - \hat{\theta}_2$, which is negative.
- If $\hat{\theta}_1 < \hat{\theta}_2$, then buyer 1 does not win and payoff for him is zero.
- Thus in this case, the maximum payoff possible is 0.

If $\hat{\theta}_1 = \theta_1$, payoff for buyer 1 is 0. By announcing $\hat{\theta}_1 = \theta_1$, his true valuation, buyer 1 gets zero payoff, which in this case is a best response.

We can now make the following observations about this example.

- Bidding his true valuation is optimal for buyer 1 regardless of what buyer 2 announces.
- Similarly bidding his true valuation is optimal for buyer 2 whatever the announcement of buyer 1.
- This means truth revelation is a weakly dominant strategy for each player, and (θ_1, θ_2) is a weakly dominant strategy equilibrium.

4 To Probe Further

The material discussed in this chapter is mainly drawn from the the book by Myerson [1]. John Harsanyi wrote a series of three classic papers introducing, formalizing, and elaborating upon Bayesian games. These papers [2, 3, 4] appeared in 1967 and 1968.

5 Problems

1. *(An exchange game)* Each of two players receives a ticket on which there is a number in some finite subset S of the interval $[0, 1]$. The number on a player's ticket is the size of a prize that he may receive. The two prizes are identically and independently distributed, with distribution function$F$. Each player is asked independently and simultaneously whether he wants to exchange his prize for the other player's prize. If both players agree then the prizes are exchanged; otherwise each player receives his own prize. Each player's objective is to maximize his expected payoff. Model this situation as a Bayesian game and show that in any Nash equilibrium the highest prize that either player is willing to exchange is the smallest possible prize.

2. Consider the following strategic situation. Two opposed armies are poised to seize an island. Each army's general can choose either "attack" or "not attack". In addition, each army is either "strong" or "weak" with equal probability (the draws for each army are independent), and an army's type is known only to its general. Payoffs are as follows: The island is worth M if captured. An army can capture the island either by attacking when its opponent does not or by attacking when its rival does if it is strong and its rival is weak. If two armies of equal strength both attack, neither captures the island. An army also has a "cost" of fighting, which is s if it is strong and w if it is weak, where $s < w$. There is no cost of attacking if its rival does not. Identify all pure strategy Bayesian Nash equilibria of this game.

References

[1] Roger B. Myerson. *Game Theory: Analysis of Conflict.* Harvard University Press, Cambridge, Massachusetts, USA, 1997.

[2] John C. Harsanyi. Games with incomplete information played by Bayesian players. Part I: The basic model. *Management Science*, 14:159–182, 1967.

[3] John C. Harsanyi. Games with incomplete information played by Bayesian players. Part II: Bayesian equilibrium points. *Management Science*, 14:320–334, 1968.

[4] John C. Harsanyi. Games with incomplete information played by Bayesian players. Part III: The basic probability distribution of the game. *Management Science*, 14:486–502, 1968.