ECE-GY 6263 Game Theory	Fall 2019

Lecture 9 — November 8, 2019

Prof. Quanyan Zhu | *Scribe: Yuhan Zhao*

1 Overview

In the last lecture we talked about the continuous kernel game and robust estimation.

In this lecture we will discuss the formulation of the games with incomplete information, the concept of Bayesian equilibrium and auction models.

2 Bayesian Games and Two Related Models

A game with incomplete information can is called the **Bayesian game**, which can written as the following strategic form:

$$\langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$$

where N is the set of the players, A_i is the action set of player i, and u_i is the utility function of player i. The incomplete information means that the player i only knows his private information such as utility function and preference, but don't know other players' information. In other words, there is no common information set like the previous games we discussed in the class.

There are many knowledge models to model this type of the game. Two of them are **Aumann model of incomplete information** and **Harsanyi game with incomplete information**. We use the definitions in chapter 9 of [1] directly. In addition, [1] also proved that these two models are **equivalent**.

2.1 Aumann Model of Incomplete Information

Definition 1. *Let S be a finite set of states of nature. An Aumann model of incomplete information (over the set S of states of nature) consists of four components $(N, Y, (\mathcal{F}_i)_{i \in N}, s)$, where:*

- *N is a finite set of players;*
- *Y is a finite set of elements called states of the world;*
- *$\mathcal{F}_i$ is a partition of Y , for each $i \in N$ (i.e., a collection of disjoint nonempty subsets of Y whose union is Y);*
- *$s : Y \to S$ is a function associating each state of the world with a state of nature.*

The Aumann model relates to a new filed called the epistemic game theory. The description of this model is tough because we have to model what a player knows other players and what other players know about the player. The description is complete, but it is computationally hard.

2.2 Harsanyi Game With Incomplete Information

Compared with the Aumann model, Harsanyi's framework captures all the uncertainties and the hierarchy knowledge into a one single variable, which is called the **type** or **attribute** parameter. In other words, each player is characterized by its own type, which is private information. The player i knows his own type, but he doesn't know other players' type; other players don't know the player i's exact type. But we assume that the distribution of type of each player is commonly know. We give the definition of Harsanyi's framework of the game with incomplete information.

Definition 2. *A Harasanyi game with incomplete information is a vector* $(N, (T_i)_{i\in N}, p, S, (s_t)_{t\in\prod_{i\in N} Ti})$ *where:*

- *N is a finite set of players.*
- *T_i is a finite set of types for player i, for each $i \in N$. The set of type vectors is denoted by $T = \prod_{i\in N} T_i$.*
- *$p \in \Delta(T)$ is a probability distribution over the set of type vectors that satisfies $p(t_i) := \sum_{t_{-i}\in T_{-i}} p(t_i, t_{-i}) > 0$ for every player $i \in N$ and every type $t_i \in T_i$.*
- *S is a set of states of nature, which will be called state games. Every state of nature $s \in S$ is a vector $s = (N, (A_i)_{i\in N}, (u_i)_{i\in N})$, where A_i is a nonempty action set of player i and $u_i : \prod_{i\in N} A_i \to \mathbb{R}$ is the payoff function of player i.*
- *$s_t = (N, (A_i(t_i))_{i\in N}, (u_i(t))_{i\in N}) \in S$ is the state game for the type vector t, for every $t \in T$. Thus, player is action set in the state game st depends on his type t_i only, and is independent of the types of the other players.*

3 Two-bidder Auction As An Example

3.1 Description of The Auction

We first put a two-bidder auction into a Bayesian model to illustrate the Harasanyi's framework. The game proceeds as follows:

(1) We have one seller and two bidder. Each buyer has a valuation v_i, $i = 1, 2$, which is privately known information only to the players themselves.

(2) v_i is a realization of the random variable V_i with support $[0, v_{i,max}]$, and associated probability density function[1] is p_i, $i = 1, 2$.

(3) Each bidder chooses a bid b_i, which should be related to the type. Here the type is the valuation v_i.

(4) All bids are submitted simultaneously. The bidders don't know other bidders bid.

(5) Define the allocation rules: the bidder with the highest bid wins. Define the payment rules: the winner pays the highest bid[2].

[1]Note that (v_1, v_2) can be correlated, so there might be a joint distribution function, but here for simplicity, we assume that they are independent.

[2]We may also use other payments rules such as the second price rule or the third price rules.

(6) Define the payoff function for the bidder i, $i = 1, 2$.

$$u_i(b_1, b_2, v_1, v_2) = \begin{cases} v_i - b_i & b_i > b_{-i} \\ 0 & b_i < b_{-1} \\ \frac{1}{2}(v_i - b_i) & b_i = b_{-i} \end{cases}.$$

The last row of u_i means each bidder has half chance to win if they bid the same.

(7) **Ex ante solution**: the strategy or the plan that considers all the contingencies before the nature starts to play, i.e., the bidder doesn't know his true type. Mathematically, we write this as

$$b_i = \mu_i(v_i), \quad i = 1, 2$$

where $\mu_i : [0, v_{i,max}] \mapsto \mathbb{R}_+$ is called the **bidding function**. Since each bidder doesn't know what type he will be at the ex ante stage, he needs to prepare for all the contingencies. So find an ex ante solution is equivalent to finding a bidding function μ_i.

Ex post solution: the strategy or the plan after the nature plays, i.e., the bidder knows his true type. Mathematically, the ex post solution is just a number $\mu_i(v_i)$, $i = 1, 2$, because at the ex post stage, each bidder already knows his type.

(8) Looking for the bidding equilibrium.

- For bidder 1, the ex ante solution is

$$\max_{\mu_1(\cdot)} \mathbb{E}_{p_2} \left[u_1(\mu_1(v_1), \mu_2(v_2), v_1, v_2) \right].$$

This optimization is hard because we are looking for a function. But it is the same as solving the following problem:

$$\max_{b_1 \in \mathbb{R}} \mathbb{E}_{p_2} \left[u_1(b_1, \mu_2(v_2), v_1, v_2) | v_1 \right].$$

which is the ex post solution. We can solve for b_1 by fixing v_1 as b_1 is a real number. Now b_1 is a function of v_1. By solving all possible b_1, we can find a function $\mu_1^*(\cdot)$ such that $b_1^* = \mu_1^*(v_1)$.

- For bidder 2, we can find $b_2^* = \mu_2^*(v_2)$ that solves

$$\max_{b_2 \in \mathbb{R}} \mathbb{E}_{p_1} \left[u_2(\mu_1(v_1), b_2, v_1, v_2) | v_2 \right].$$

3.2 Cournot Game With Incomplete Information

The Cournot game can be found in Chapter 6 of [2]. There are two players P1 and P2 with the following utility functions

$$u_1 = q_1(\theta_1 - q_1 - q_2), \quad u_2 = q_2(\theta_2 - q_1 - q_2),$$

where θ_i, $i = 1, 2$, is the type of each player. The distribution of θ_i, $i = 1, 2$, is commonly known:

$$\theta_1 = 1 \text{ w.p.1}, \quad \theta_2 = \begin{cases} \frac{3}{4} & \text{w.p.} \frac{1}{2} \\ \frac{5}{4} & \text{w.p.} \frac{1}{2} \end{cases}.$$

We want to find the equilibrium of this game.

For P1, he will have two possibilities when predicting q_2 at the ex ante stage (although eventually he meets one of them), denoting as q_2^a and q_2^b. Then P1 wants to maximize the expected value:

$$\max_{q_1} \frac{1}{2} q_1 \left(\theta_1 - q_1 - q_2^a\right) + \frac{1}{2}\left(\theta_1 - q_1 - q_2^b\right). \tag{1}$$

Note that (1) is concave in q_1, we have

$$q_1 = \frac{1}{2} - \frac{1}{4}\left(q_2^a + q_2^b\right). \tag{2}$$

For P2, he knows his type. When P2 is in type a, he wants to solve

$$\max_{q_2^a} q_2^a \left(\frac{3}{4} - q_1 - q_2^a\right) \quad \Rightarrow \quad q_2^a = \frac{3}{8} - \frac{1}{2} q_1. \tag{3}$$

When P2 is in type b, he wants to solve

$$\max_{q_2^b} q_2^b \left(\frac{5}{4} - q_1 - q_2^b\right) \quad \Rightarrow \quad q_2^b = \frac{5}{8} - \frac{1}{2} q_1. \tag{4}$$

The equilibrium can be computed via (2),(3) and (4):

$$\begin{cases} q_1^* = \frac{1}{2} \\ q_2^{a*} = \frac{5}{24} \\ q_2^{b*} = \frac{11}{24} \end{cases} \quad \Rightarrow \quad \begin{array}{l} q_1 = \mu_1^*(\theta_1) = \dfrac{1}{3} \quad \forall \theta_1 \\ q_2 = \mu_2^*(\theta_2) = \begin{cases} \frac{5}{24} & \theta_2 = \frac{3}{4} \\ \frac{11}{24} & \theta_2 = \frac{5}{4} \end{cases} \end{array}$$

4 N-person Bayesian Nash Equilibrium Problem

This section generalizes the notion of the Bayesian game and Bayesian Nash equilibrium (BNE) in Section 3. The definition of BNE can be found in Chapter 6.4 of [2].

4.1 Description of The Bayesian Game

(1) The players set is $N = \{1, 2, \ldots, N\}$.

(2) Each player is associated a type/attribute (private information) $\theta_i \Theta_i$. Note that the realization θ_i is private but the distribution of the type and the associated support is commonly known.

(3) The type θ_i is drawn from the distribution p_i, $i = 1, \ldots, k$. Note that p_i can be either depends only on Θ_i or depends on other Θ_{-i}. We don't explicitly make assumptions of what p_i depends on.

(4) Each player has an action space A_i, the set of all possible actions for player i, $i = 1, \ldots, N$.

(5) The **strategy** for player i is a mapping $\mu_i L \Theta_i \mapsto A_i$, which is

$$a_i = \mu_i(\theta_i) \quad i = 1, \ldots, N.$$

(6) The payoff for player i is

$$u_i(a_1, \ldots, a_n; \theta_i, \theta_{-i}) = \tilde{u}_i\left(\mu_1(\theta_1), \ldots, \mu_N(\theta_N); \theta_i, \theta_{-i}\right) \quad i = 1, \ldots, N.$$

Note that $u_i(\cdot)$ is defined on the action space while $\tilde{u}_i(\cdot)$ is defined on the strategy space.

(7) $\{\mu_i^*\}_{i=1}^N$ is a (pure strategy) Bayesian Nash Equilibrium (BNE) if $a_i^* = \mu_i^*(\theta_i)$ and

$$\max_{a_i \in A_i} \mathbb{E}_{p_{-1}|\theta_i}\Big[u_i\big(\mu_1(\theta_1), \ldots, \mu_{i-1}(\theta_{i-1}), a_i, \mu_{i+1}(\theta_{i+1}), \ldots, \mu_N(\theta_N); \theta_i, \theta_{-i}\big)|\theta_i\Big] \quad \forall \theta_i \in \Theta_i, i \in N$$

Note that here we capture the uncertainty using the type variables in the utility function, but uncertainty can appear anywhere such action spaces. There are many other ways to capture and model the uncertainty in different situations.

4.2 Existence of BNE

For finite games, the existence of BNE is an immediate consequence of the Nash existence theorem [2] because the entire game can be thought as a matrix game. The finite means two things: first, the number of players is finite; second, each type variable has a discrete distribution and its associated support is also finite.

For infinite games, we have to use the continuous kernel arguments to find the existence. In this case proving the existence is complicated because we are looking for a function and it depends on so many things. We need more work to prove this. There is a paper [3] talking about the existence of behavioral strategy equilibrium in the Bayesian game.

5 Incomplete Information As a Possible Interpretation of Mixed Strategy

The mixed strategy was not very clear to people for some time because in real world decision makers don't "flip coins". Harsanyi shows that, according to Chapter 6.7 of [2], "the mixed strategy equilibria of complete information games can be can usually be interpreted as the limits of pure strategy equilibria of slightly perturbed games of incomplete information". We will show this using the following example. More detailed information can be found in Chapter 9.5 of [1].

Consider the matrix game with two players. Both of the players are maximizers.

	L	R	
T	$0,0$	$0,-1$	y
B	$1,0$	$-1,3$	$1-y$
	z	$1-z$	

The optimal mixed strategy of the game is $y^* = \frac{3}{4}$, $z^* = \frac{1}{2}$. Now we perturb the game as follows.

$\epsilon\alpha, \epsilon\beta$	$\epsilon\alpha, -1$
$1, \epsilon\beta$	$-1, 3$

(5)

where α, β are i.i.d. random variables over the interval $[-1, 1]$ with uniform distribution. ϵ is a small number which can be interpreted as the magnitude of the perturbation. So we can think of this as a Bayesian game:

(1) The players set is $\{1, 2\}$.

(2) Player 1 and player 2 have continuous type spaces $T_1 = [-1, 1]$ and $T_2 = [-1, 1]$ respectively.

(3) The type random variables (T_1, T_2) has a uniform distribution over the unite square.

(4) For every realization of α, β, we have the matrix game (5).

(5) The strategy of the player 1 is $s_1(\alpha)$; the strategy of the player 2 is $s_2(\beta)$.

Now player 1 wants to solve

$$\max_{a_1} \mathbb{E}_\beta \big[u_1(a_1, s_1(\beta); \alpha, \beta)|\alpha\big].$$

Player 2 wants to solve

$$\max_{a_2} \mathbb{E}_\alpha \big[u_2(s_1(\alpha), a_2; \alpha, \beta)|\beta\big].$$

Solve the preceding optimization problems is hard because α and β are all continuous. Thus we investigate the threshold strategy:

$$s_1(\alpha) = \begin{cases} T & \alpha > \alpha_0 \\ B & \alpha \le \alpha_0 \end{cases}, \quad s_2(\beta) = \begin{cases} L & \beta > \beta_0 \\ R & \beta \le \beta_0 \end{cases}.$$

Player 1 and player 2 should determine α_0 and β_0 respectively. Then for player 1, if $\alpha > \alpha_0$, we have

$$\mathbb{E}_\beta \big[u_1(a_1, s_1(\beta); \alpha, \beta)|\alpha\big] = \epsilon\alpha\, \mathbb{P}(\beta > \beta_0) + \epsilon\alpha\, \mathbb{P}(\beta \le \beta_0) = \epsilon\alpha.$$

Otherwise, we have

$$\mathbb{E}_\beta \big[u_1(a_1, s_1(\beta); \alpha, \beta)|\alpha\big] = 1\, \mathbb{P}(\beta > \beta_0) + (-1)\, \mathbb{P}(\beta \le \beta_0) = -\beta_0.$$

If player 1 chooses T, we know $\alpha > \alpha_0 \Rightarrow \epsilon\alpha \ge -\beta_0$. If player 2 chooses B, we have $\epsilon\alpha \le -\beta_0$. To show player 1 has no inventive to deviate from the choice of α_0, we have

$$\epsilon\alpha_0 = -\beta_0. \tag{6}$$

Similarly, for player 2, if he has no incentive to deviate from the choice of β_0, we have

$$\epsilon\beta_0 = 1 + 2\alpha_0. \tag{7}$$

From (6) and (7), we can obtain the equilibrium:

$$\alpha_0 = -\frac{1}{2+\epsilon^2}, \quad \beta_0 = \frac{\epsilon}{2+\epsilon^2}.$$

Therefore, we can obtain the following threshold strategy:

$$s_{1,\epsilon}(B) = \mathbb{P}\left(\alpha \le -\frac{1}{2+\epsilon^2}\right) = \frac{1+\epsilon^2}{4+2\epsilon^2} = \frac{1}{4} \text{ as } \epsilon \to 0$$
$$s_{2,\epsilon}(R) = \mathbb{P}\left(\beta \le \frac{2}{2+\epsilon^2}\right) = \frac{2+\epsilon+\epsilon^2}{4+2\epsilon^2} = \frac{1}{2} \text{ as } \epsilon \to 0$$

which are exactly the mixed strategy of the original game as $\epsilon \to 0$.

The limit shows there is a continuity in mixed strategies. It also shows that we can interpret mixed strategies with pure strategies with incomplete information.

6 Auction Revisited

In this section we study the N-bidder auction under some mild symmetric assumptions, and derive the closed form solution of the BNE of the game (Myerson). The Chapter 12 of [1] gives a systematic introduction and analysis of suctions.

Description of the auction:

(1) The players set is $\{1, 2, \ldots, N\}$.

(2) Each bidder's valuation of the good is a random variable Θ_i. Its realization has the support $\theta_i \in [0, \theta_{max}]$ for all i. All Θ_i are i.i.d. according to the distribution function p.

(3) Each bidder has his own bid b_i, and the associated bidding function is $b_i = \mu_i(\theta_i)$, $i = 1, \ldots, N$.

Assumptions:

(1) μ_i is strictly increasing[3], continuous and differentiable.

(2) Every bidder has the same bidding function[4], i.e., $\mu_i(\cdot) = \mu(\cdot)$ for all $i = 1, \ldots, N$. Thus we are looking for a **symmetric equilibrium**.

The utility function of the first bidder is

$$\begin{aligned} u_1(b_1, \mu(\theta_2), \ldots, \mu\theta_N; \theta) &= (\theta_1 - b_1)\mathbf{1}_{\{b_1 > \max(\mu(\theta_2), \ldots, \mu(\theta_N))\}} \\ &= \begin{cases} \theta_1 - b_1 & \text{if } \{b_1 > \max(\mu(\theta_2), \ldots, \mu(\theta_N))\} \text{ is true} \\ 0 & \text{otherwise} \end{cases} . \end{aligned}$$

Similarly, we can define other bidders utility functions. For the condition in u_1, using the properties of $\mu(\cdot)$ in the assumption, we rewrite it as follows.

$$\begin{aligned} & b_1 > \max(\mu(\theta_2), \ldots, \mu(\theta_N)) \\ \Rightarrow \quad & b_1 > \mu\left(\max(\theta_2, \ldots, \theta_N)\right) \\ \Rightarrow \quad & \mu^{-1}(b_1) > \max(\theta_2, \ldots, \theta_N). \end{aligned}$$

Thus, u_1 becomes

$$u_1(b_1, \mu(\theta_2), \ldots, \mu\theta_N; \theta) = (\theta_1 - b_1)\mathbf{1}_{\{\mu^{-1}(b_1) > \max(\theta_2, \ldots, \theta_N)\}}.$$

Taking expectation over all other players, we can obtain the probability of bidder 1 winning.

$$\mathbb{P}\left(\max(\theta_2, \ldots, \theta_N) < \mu^{-1}(b_1)\right) = \prod_{i=1}^{N-1} \mathbb{P}\left(\theta_i < \mu^{-1}(b_1)\right) = \left[p(\mu^{-1}(b_1))\right]^{N-1},$$

where $p(\cdot)$ is a CDF.

[3]The monotone property can be interpreted in this way: if a bidder values less, he has no incentive to bid more.

[4]Note that the bidding function $\mu_i(\cdot)$ of each bidder is the same, but the input of $\mu(\cdot)$ may be different, and that's how we distinguish different bidders, i.e., $b_1 = \mu(\theta_1)$, $b_2 = \mu(\theta_2)$.

Define $Y = \max(\theta_2, \dots, \theta_N)$. Let $F_Y(y)$ and $f_Y(y)$ be the CDF and PDF of Y respectively. Then we have

$$F_Y(y) = \left[p(\mu^{-1}(b_1))\right]^{N-1}.$$

Then, the expectation of u_1 given θ_1 can be written as

$$\mathbb{E}\left[u_1(b_1, \mu(\theta_2), \dots, \mu(\theta_N); \theta)|\theta_1\right] = (\theta_1 - b_1)F_Y\left(\mu^{-1}(b_1)\right). \tag{8}$$

The first term in (8) can be interpreted as the benefit of winning while the second term is the probability of winning.

Next, the first bidder wants to maximize the expectation w.r.t. b_1. Using FOC, we can obtain

$$\begin{aligned}
&\frac{d}{db_1}(\theta_1 - b_1)F_Y\left(\mu^{-1}(b_1)\right) \\
\Rightarrow\quad & F_Y'\left(\mu^{-1}(b_1)\right)\left[\mu^{-1}(b_1)\right]'(\theta_1 - b_1) - F_Y\left(\mu^{-1}(b_1)\right) = 0.
\end{aligned}$$

Note that $F_Y'\left(\mu^{-1}(b_1)\right) = f_Y\left(\mu^{-1}(b_1)\right)$. To find $\left[\mu^{-1}(b_1)\right]'$, we use the identity

$$\mu\left(\mu^{-1}(b_1)\right) = b_1.$$

Taking derivative w.r.t. b_1 on both sides, we have

$$\mu'\left(\mu^{-1}(b_1)\right)\left[\mu^{-1}(b_1)\right]' = 1 \quad \Rightarrow \quad \left[\mu^{-1}(b_1)\right]' = \frac{1}{\mu'\left(\mu^{-1}(b_1)\right)}.$$

Therefore,

$$\frac{f_Y\left(\mu^{-1}(b_1)\right)}{\mu'\left(\mu^{-1}(b_1)\right)}(\theta_1 - b_1) - F_Y\left(\mu^{-1}(b_1)\right) = 0. \tag{9}$$

Note that $\theta_1 = \mu^{-1}(b_1)$, so (9) can be written as

$$\begin{aligned}
&\frac{f_Y(\theta_1)}{\mu'(\theta_1)}(\theta_1 - b_1) - F_Y(\theta_1) = 0 \\
\Rightarrow\quad & \theta_1 f_Y(\theta_1) - \mu(\theta_1)f_Y(\theta_1) - F_Y(\theta_1)\mu'(\theta_1) = 0 \\
\Rightarrow\quad & \theta_1 f_Y(\theta_1) = \frac{d}{d\theta_1}\left[\mu(\theta_1)F_Y(\theta_1)\right] \\
\Rightarrow\quad & \mu(\theta_1)F_Y(\theta_1) = \int_0^{\theta_1} \theta f_Y(\theta)d\theta + const.
\end{aligned}$$

We also assume that $\mu(0) = 0$, then the constant term goes to 0. So we have

$$\mu(\theta_1) = \int_0^{\theta_1} \frac{\theta f_Y(\theta)}{F_Y(\theta_1)}d\theta \quad \text{or} \quad \mu(\theta_1)\mathbb{E}(Y|Y \leq \theta_1),$$

which is the conditional expectation.

We can use SOC to verify that $\mu(\theta_1)$ is indeed optimal for the first bidder.

References

[1] E. Solan, M. Maschler, and S. Zamir, *Game Theory*, New York: Cambridge University Press, 2013.

[2] D. Fudenberg and J. Tirole, *Game Theory*, Cambridge, MA: MIT Press, 1991.

[3] O. Carbonell-Nicolau and R. P. McLeana, "On the existence of Nash equilibrium in Bayesian games." *Mathematics of Operations Research*, vol. 43, no. 1, Feb., pp. 100-129, 2018.