به نام خدا

An introduction to Game Theory

Lecture 5

By: M. Khakestari

University of Economics Sciences

Fall 2013

Bayesian Games

Outline

- Incomplete information.
- Bayes rule and Bayesian inference.
- Bayesian Nash Equilibria.
- Auctions.
- Extensive form games of incomplete information.
- Perfect Bayesian (Nash) Equilibria.
- Introduction to social learning and herding.

Reading:
Osborne, Chapter 9.
Eric Rasmusen, Chapter 6.

Incomplete Information

- In many game theoretic situations, one agent is unsure about the payoffs or preferences of others.
- Incomplete information introduces additional strategic interactions and also raises questions related to "learning".
- Examples:
  - Bargaining (how much the other party is willing to pay is generally unknown to you)
  - Auctions (how much should you bid for an object that you want, knowing that others will also compete against you?)
  - Market competition (firms generally do not know the exact cost of their competitors)
  - Signaling games (how should you infer the information of others from the signals they send)
  - Social learning (how can you leverage the decisions of others in order to make better decisions)

Bayesian Games

So far we have been assuming that everything in the game was common knowledge for everybody playing. But in fact players may have private information about their own payoffs, about their type or preferences, etc. The way to modelling this situation of asymmetric or incomplete information is by recurring to an idea generated by Harsanyi (1967) [1].

[1]. J. Harsanyi (1967, 68) games of incomplete information, Bayesian equilibrium

Example: Incomplete Information Battle of the Sexes

- Recall the battle of the sexes game, which was a complete information "coordination" game.

	B	F
B	$(2,1)$	$(0,0)$
F	$(0,0)$	$(1,2)$

- In this game there are two pure strategy equilibria (one of them better for player 1 and the other one better for player 2), and a mixed strategy equilibrium.
- Now imagine that player 1 does not know whether player 2 wishes to meet or wishes to avoid player 1. Therefore, this is a situation of **incomplete information**—also sometimes called *asymmetric information*.

Example: Incomplete Information Battle of the Sexes

- We represent this by thinking of player 2 having two different **types**, one type that wishes to meet player 1 and the other wishes to avoid him.
- More explicitly, suppose that these two types have probability $1/2$ each. Then the game takes the form one of the following two with probability $1/2$.

	B	F
B	$(2,1)$	$(0,0)$
F	$(0,0)$	$(1,2)$

	B	F
B	$(2,0)$	$(0,2)$
F	$(0,1)$	$(1,0)$

- Crucially, player 2 knows which game it is (she knows the **state of the world**), but player 1 does not.
- What are strategies in this game?

Example: Incomplete Information Battle of the Sexes

- Let us consider the following strategy profile $(B, (B, F))$, which means that player 1 will play B, and while in state 1, player 2 will also play B (when she wants to meet player 1) and in state 2, player 2 will play F (when she wants to avoid player 1).
- Clearly, given the play of B by player 1, the strategy of player 2 is a best response.
- Let us now check that player 2 is also playing a best response.
- Since both states are equally likely, the expected payoff of player 2 is

$$\mathbb{E}[B, (B, F)] = \frac{1}{2} \times 2 + \frac{1}{2} \times 0 = 1.$$

- If, instead, he deviates and plays F, his expected payoff is

$$\mathbb{E}[F, (B, F)] = \frac{1}{2} \times 0 + \frac{1}{2} \times 1 = \frac{1}{2}.$$

- Therefore, the strategy profile $(B, (B, F))$ is a (Bayesian) Nash equilibrium.

Example: Incomplete Information Battle of the Sexes

- Interestingly, meeting at Football, which is the preferable outcome for player 2 is no longer a Nash equilibrium. Why not?
- Suppose that the two players will meet at Football when they want to meet. Then the relevant strategy profile is $(F, (F, B))$ and

$$\mathbb{E}[F, (F, B)] = \frac{1}{2} \times 1 + \frac{1}{2} \times 0 = \frac{1}{2}.$$

- If, instead, player 1 deviates and plays B, his expected payoff is

$$\mathbb{E}[B, (F, B)] = \frac{1}{2} \times 0 + \frac{1}{2} \times 2 = 1.$$

- Therefore, the strategy profile $(F, (F, B))$ is not a (Bayesian) Nash equilibrium.

Bayesian Games

Definition 1 *A Bayesian Game is a game in normal form with incomplete information that consists of:*

1) Players $i \in \{1, 2, ..., I\}$

2) Finite action set for each player $a_i \in A_i$

3) Finite type set for each player $\theta_i \in \Theta_i$

4) A probability distribution over types $p(\theta)$ *(common prior beliefs about the players' types)*

5) Utilities $u_i : A_1 \times A_2 \times ... \times A_I \times \Theta_1 \times \Theta_2 \times ...\Theta_I \to \mathbb{R}$

Bayesian Games

Definition

A (pure) strategy for player i is a map $s_i : \Theta_i \to S_i$ prescribing an action for each possible type of player i.

Bayesian Games

- Recall that player types are drawn from some prior probability distribution $p(\theta_1, \ldots, \theta_I)$.
- Given $p(\theta_1, \ldots, \theta_I)$ we can compute the conditional distribution $p(\theta_{-i} \mid \theta_i)$ using Bayes rule.

- Player i knows her own type and evaluates her expected payoffs according to the the **conditional distribution** $p(\theta_{-i} \mid \theta_i)$, where $\theta_{-i} = (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_I)$.

Bayesian Games

- Since the payoff functions, possible types, and the prior probability distribution are common knowledge, we can compute expected payoffs of player i of type θ_i as

$$\begin{aligned} U\left(s_i', s_{-i}(\cdot), \theta_i\right) &= \sum_{\theta_{-i}} p(\theta_{-i} \mid \theta_i) u_i(s_i', s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) \\ &\qquad \text{when types are finite} \\ &= \int u_i(s_i', s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) P(d\theta_{-i} \mid \theta_i) \\ &\qquad \text{when types are not finite.} \end{aligned}$$

Bayes Rule

- Let $\Pr(A)$ and $\Pr(B)$ denote, respectively, the probabilities of events A and B; $\Pr(B \mid A)$ and $\Pr(A \cap B)$, conditional probabilities (one event conditional on the other one), and $\Pr(A \cap B)$ be the probability that both events happen (are true) simultaneously.
- Then Bayes rule states that

$$\Pr(A \mid B) = \frac{\Pr(A \cap B)}{\Pr(B)}. \qquad \text{(Bayes I)}$$

- Intuitively, this is the probability that A is true given that B is true.
- When the two events are **independent**, then $\Pr(B \cap A) = \Pr(A) \times \Pr(B)$, and in this case, $\Pr(A \mid B) = \Pr(A)$.

Bayes Rule

- Bayes rule also enables us to express conditional probabilities in terms of each other. Recalling that the probability that A is not true is $1 - \Pr(A)$, and denoting the event that A is not true by A^c (for A "complement"), so that $\Pr(A^c) = 1 - \Pr(A)$, we also have

$$\Pr(A \mid B) = \frac{\Pr(A) \times \Pr(B \mid A)}{\Pr(A) \times \Pr(B \mid A) + \Pr(A^c) \times \Pr(B \mid A^c)}. \quad \text{(Bayes II)}$$

- This equation directly follows from (Bayes I) by noting that

$$\Pr(B) = \Pr(A) \times \Pr(B \mid A) + \Pr(A^c) \times \Pr(B \mid A^c),$$

and again from (Bayes I)

$$\Pr(A \cap B) = \Pr(A) \times \Pr(B \mid A).$$

Bayes Rule

- More generally, for a finite or countable partition $\{A_j\}_{j=1}^{n}$ of the event space, for each j

$$\Pr(A_j \mid B) = \frac{\Pr(A_j) \times \Pr(B \mid A_j)}{\sum_{i=1}^{n} \Pr(A_i) \times \Pr(B \mid A_i)}.$$

- For continuous probability distributions, the same equation is true with densities

$$f(A' \mid B) = \frac{f(A') \times f(B \mid A')}{\int f(B \mid A) \times f(A)\, dA}.$$

Existence of Bayesian Nash Equilibria

A Bayesian Nash Equilibrium is simply a Nash Equilibrium of the game where Nature moves first, chooses $\theta \in \Theta$ from a distribution with probability p (θ) and reveals θi to player i.

- ***Theorem***

Consider a finite incomplete information (Bayesian) game. Then a mixed strategy Bayesian Nash equilibrium exists.

Every finite Bayesian Game has a Bayesian Nash Equilibrium

Auctions

- A major application of Bayesian games is to auctions.
- This corresponds to a situation of incomplete information because the valuations of different potential buyers are unknown.
- We made the distinction between:
- Private value auctions: valuation of each agent is independent of others' valuations;
- Common value auctions: the object has a potentially common value, and each individual's signal is imperfectly correlated with this common value.
- We have analyzed private value first-price and second-price sealed bid auctions.
- Each of these two auction formats defines a static game of incomplete information (Bayesian game) among the bidders.
- We determined Bayesian Nash equilibria in these games and compared the equilibrium bidding behavior.

Auctions (continued)

- Games of incomplete information
- First Price Sealed Bid Auction
  - Buyers simultaneously submit their bids
  - Buyers' valuations of the good unknown to each other
  - Highest Bidder wins and gets the good at the amount he bid
  - Nash Equilibrium: Each person would bid less than what the good is worth to you
- Second Price Sealed Bid Auction
  - Same rules
  - Exception – Winner pays the second highest bid and gets the good
  - Nash equilibrium: Each person exactly bids the good's valuation

Modeling Auctions

- Model of auction:
  - a valuation structure for the bidders (i.e., private values for the case of private-value auctions),
  - a probability distribution over the valuations available to the bidders.
- Let us focus on first and second price sealed bid auctions, where bids are submitted simultaneously.
- Each of these two auction formats defines a static game of incomplete information (Bayesian game) among the bidders.
- We determine Bayesian Nash equilibria in these games and compare the equilibrium bidding behavior.

Modeling Auctions(continued)

Two bidders are trying to purchase the same item in a sealed bid auction. The bidders simultaneously submit bids b1 and b2 and the auction is sold to the highest bidder at his bid price (this is called a "**first price**" auction). If there is a tie, there is a coin flip to determine the winner. Suppose the players utilities are:

$$u_i(b_i, b_{-i}) = \begin{cases} v_i - b_i & \text{if } b_i > b_{-i} \\ \frac{1}{2}(v_i - b_i) & \text{if } b_i = b_{-i} \\ 0 & \text{if } b_i < b_{-i} \end{cases}$$

Exercise

Consider the following situation - called a "first-price, sealed bid auction":

There are 2 bidders, with a valuation v_i of bidder $i = 1, 2$ of a good. The good is indivisible and the supply is a single unit. The two bidders' valuation are independently, uniformly distributed on the interval $[0, 1]$. Each bidder only knows his own value of the good. If the bidder i obtains the good and pays a price of p, the value to the bidder is $v_i - p$. If he does not obtain the good the value is 0. The rules of the game is as follows: Each bidder

simultanously submit a bid. The highest bidder is granted the good and pays the bid. If they bid the same the good is randomly allocated between the two, i.e., there is a probability of $\frac{1}{2}$ of each bidder getting the good.

a) Formulate this situation as a Bayesian game

b) Solve for a Bayesian Nash equilibrium (Hint: Assume that the strategy is of the affin form $b_i(v_i) = a_i + b_i v_i$. Solve then for a_i and b_i)

c) Do the bidders bid their value of the good?

Second-Price Auction

- Second price auctions will have the structure very similar to a complete information auction discussed earlier in the lectures.
- There we saw that each player had a weakly dominant strategy. This will be true in the incomplete information version of the game and will greatly simplify the analysis.
- In the auction, each bidder submits a sealed bid of b_i, and given the vector of bids $b = (b_i, b_{-i})$ and evaluation v_i of player i, its payoff is

$$U_i\left(\left(b_i, b_{-i}\right), v_i\right) = \begin{cases} v_i - \max_{j \neq i} b_j & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{if } b_i > \max_{j \neq i} b_j . \end{cases}$$

- Let us also assume that if there is a tie, i.e., $b_i = \max_{j \neq i} b_j$, the object goes to each winning bidder with equal probability.

Second-Price Auction (continued)

- This can be established with the same graphical argument as the one we had for the complete information case.
- The first graph shows the payoff for bidding one's valuation, the second graph the payoff from bidding a lower amount, and the third the payoff from bidding higher amount.
- In all cases B^* denotes the highest bid excluding this player.