نظریه بازیها
Game Theory

ارائه کننده: امیرحسین نیکوفرد
مهندسی برق و کامپیوتر دانشگاه خواجه نصیر

Bayesian games

Material

- Fudenberge D., Birole J., Game Theory, MIT Press, Cambridge, Massachusetts, 1991,Sections 6.1-6.5.

Bayesian games

- Zero sum games
- Non-zero sum games
- Infinite Games
- Infinite Dynamic Games
- Stochastic games
- **Bayesian games**
  - **Incomplete information.**
  - **Bayes rule and Bayesian inference.**
  - **Bayesian Nash Equilibria.**
  - **Auctions.**

Incomplete Information

- All the kinds of games we've looked at so far have assumed that
  - everything relevant about the game being played is common knowledge to all the players:
  - the number of players,
  - the actions available to each , and
  - the payoff vector associated with each action vector
- True even for imperfect-information games
  - The actual moves aren't common knowledge, but the game is
- We'll now consider games of **incomplete** (not imperfect) information
  - Players are uncertain about the game being played

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)

Incomplete Information

- We will consider cases in which a party may have some information that is not known by some other party. Such games are called games of **incomplete information** or **asymmetric information**. The informational asymmetries are modeled by **Nature's moves.** Some players can distinguish certain moves of nature while some others cannot. Consider the following simple example, where a firm is contemplating the hiring of a worker, without knowing how able the worker is.

Example

- Consider the game in the Figure. There are a **Firm** and a **Worker**. Worker can be of **High** ability, in which case he would like to Work when he is hired, or of **Low** ability, in which case he would rather Shirk. Firm would want to Hire the worker that will work but not the worker that will shirk. Worker knows his ability level. Firm does not know whether the worker is of high ability or low ability. Firm believes that the worker is of high ability with probability **p** and low ability with probability **1-p**. **Most importantly, the firm knows that the worker knows his own ability level**.

Example

- To model this situation, we let Nature choose between High and Low, with probabilities p and 1-p, respectively. We then let the worker observe the choice of Nature, but we do not let the firm observe Nature's choice.

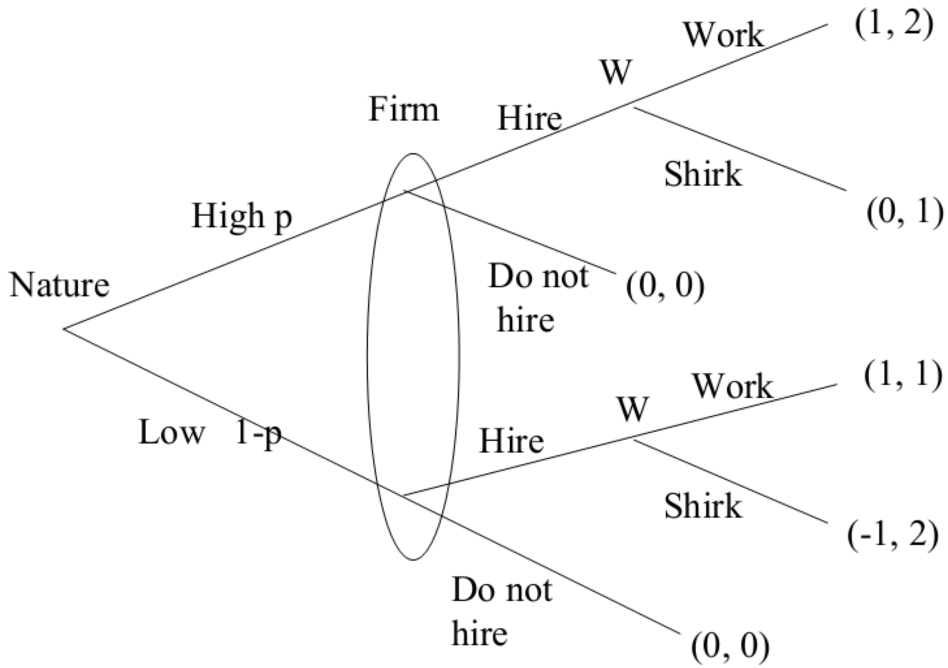

Example

- A player's **private information** is called his "**type**". For instance, in the above example Worker has two types**: High and Low.**
- Since Firm does not have any private information, Firm has only **one type**. As in the above example, incomplete information is modeled via imperfect-information games where Nature chooses each player's type and privately informs him. These games are called **incomplete-information game** or **Bayesian game**.

Bayesian Games

- A **Bayesian** game or "incomplete information games" includes the following:
  - set of actions (pure strategies) for each player i : A_i
  - A set of players or agents: $N = \{1, 2, \ldots, n\}$
  - A set of types for each player i (t_i in T_i)
  - A payoff function foreach player i : $u_i(a_1, \ldots, a_n, t_1, \ldots, t_n)$
  - A (joint) probability distribution $P(t_1, \ldots, t_n)$ over types or belief about the other players
- One can write the game in the example above as a Bayesian game by setting

Bayesian Games

- $N = \{F, W\}.$
- $T_F = \{t_f\}, \mathrm{T}_W = \{\mathrm{high}, \mathrm{low}\}.$
- $P(t_f, high) = \mathrm{p}, \quad P(t_f, low) = 1 - \mathrm{p}.$
- $A_F = \{hire, dont\}, A_W = \{work, shirk\}.$
- and the utility functions u_F and u_W are defined by the following tables, where the first entry is the payoff of the firm and the table on the left corresponds to $(t_f, high)$

t_W=high	Work	Shirk
hire	1,2	0,1
dont	0,0	0,0

t_W=low	Work	Shirk
hire	1,1	-1,2
dont	0,0	0,0

Bayesian games

- Importantly, throughout in Bayesian games, the strategy spaces, the payoff functions, possible types, and the prior probability distribution are assumed to be **common knowledge**.
- A (pure) strategy for player i is a map $s_i : T_i \rightarrow A_i$ prescribing an action for each possible type of player I
- Recall that player types are drawn from some prior probability distribution $P(t_1, \ldots, t_n)$
- Given $P(t_1, \ldots, t_n)$ we can compute the conditional distribution $P(t_{-i} \mid t_i)$ using **Bayes rule**.
- Player i knows her own type and evaluates her expected payoffs according to the conditional distribution $P(t_{-i} \mid t_i)$

Bayesian games

- **Bayes' Rule:** Let A and B be two events, then probability that A occurs conditional on B occurring is

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)}$$

- where $P(A \cap B)$ is the probability that A and B occur simultaneously and $P(B)$:the (unconditional) probability that B occurs.
- In static games of incomplete information, the application of Bayes' Rule will often be trivial, but a very good understanding of the Bayes' Rule is necessary to follow the treatment of the dynamic games of incomplete information later.

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 t_i as

$$U(s_i', s_{-i}, t_i) = \sum_{t_{-i}} P(t_{-i} \mid t_i) u_i(s_i', s_{-i}(\mathrm{t}_{-i}), t_i, t_{-i})$$

when types are finite

$$U(s_i', s_{-i}, t_i) = \int u_i(s_i', s_{-i}(\mathrm{t}_{-i}), t_i, t_{-i}) \quad P(\,\mathrm{d}t_{-i} \mid t_i)$$

when types are not finite

Bayesian Nash equilibrium

Definition :

(**Bayesian Nash Equilibrium**) The strategy profile s ( ) is a (pure strategy) Bayesian Nash equilibrium if for all $i \in N$ and for all $t_i \in T_i$, we have that

$$s_i(\mathrm{t}_i) \in \arg\max_{s_i' \in S_i} \sum_{t_{-i}} P(t_{-i} | t_i) u_i(s_i', s_{-i}(\mathrm{t}_{-i}), t_i, t_{-i})$$

or in the non-finite case,

$$s_i(\mathrm{t}_i) \in \arg\max_{s_i' \in S_i} \int u_i(s_i', s_{-i}(\mathrm{t}_{-i}), t_i, t_{-i}) P(\mathrm{d}t_{-i} | t_i)$$

Example: Incomplete Information Cournot

- Suppose that two firms both produce at constant marginal cost.
- Demand is given by P (Q) as in the usual Cournot game.
- Firm1 has marginal cost equal to C (and this is common knowledge).
- Firm 2's marginal cost is private information. It is equal to C_L with probability p and to C_H with probability (1-p), where $C_L<C_H$
- This game has 2 players, 2 states (L and H) and the possible actions of each player are $q_i \in [0,\infty)$, but firm 2 has two possible types.

Example: Incomplete Information Cournot

- The payoff functions of the players, after quantity choices are made, are given by

$$u_1((q_1, q_2), t) = q_1(P(q_1 + q_2) - C)$$
$$u_2((q_1, q_2), t) = q_2(P(q_1 + q_2) - C_t)$$

- where $t \in \{L, H\}$ is the type of player 2.
- A strategy profile can be represented as (q_1^*, q_L^*, q_H^*) [or equivalently as $(q_1^*, q_2^*(t_2))$], where q_L^* *and* q_H^* denote the actions of player 2 as a function of its possible types.
- We now characterize the Bayesian Nash equilibria of this game by computing the best response functions (correspondences) and finding their intersection

Example: Incomplete Information Cournot

- There are now three best response functions and they are are given by

$$B_1(q_L, q_H) = \arg\max_{q_1 \geq 0}\{p(P(q_1 + q_L) - C)q_1 + (1-p)(P(q_1 + q_H) - C)q_1\}$$

$$B_L(q_1) = \arg\max_{q_L \geq 0}\{(P(q_1 + q_L) - C_L)q_L\}$$

$$B_H(q_1) = \arg\max_{q_H \geq 0}\{(P(q_1 + q_H) - C_H)q_H\}$$

- The Bayesian Nash equilibria of this game are vectors (q_1^*, q_L^*, q_H^*) such that

$$B_1(q_L^*, q_H^*) = q_1^*, \qquad B_L(q_1^*) = q_L^*, \qquad B_H(q_1^*) = q_H^*$$

Example: Incomplete Information Cournot

- To simplify the algebra, let us assume that $P(Q)=\alpha-Q, \quad \alpha \geq Q$

  Then we can compute:

$$q_1^* = \frac{1}{3}(\alpha - 2\mathrm{C} + \mathrm{pC}_L + (1-\mathrm{p})\mathrm{C}_H)$$

$$q_L^* = \frac{1}{3}(\alpha - 2\mathrm{C}_L + C) - \frac{1}{6}(1-\mathrm{p})(\mathrm{C}_H - \mathrm{C}_L)$$

$$q_H^* = \frac{1}{3}(\alpha - 2\mathrm{C}_H + C) + \frac{1}{6}\mathrm{p}(\mathrm{C}_H - \mathrm{C}_L)$$

- Note that $q_L^* > q_H^*$.This reflects the fact that with lower marginal cost, the firm will produce more.

Example: Incomplete Information Cournot

- However, incomplete information also affects firm 2's output choice.
- Recall that, given this demand function, if both firms knew each other's marginal cost, then the unique Nash equilibrium involves output of firm i given by $\frac{1}{3}(\alpha - 2C_i + C_J)$
- With incomplete information, firm 2's output is more if its cost is C_L and less if its cost is C_H. If firm 1 knew firm 2's cost is high, then it would produce more. However, its lack of information about the cost of firm 2 leads firm 1 to produce a relatively moderate level of output, which then allows firm 2 to be more "aggressive".
- Hence, in this case, firm 2 benefits from the lack of information of firm 1 and it produces more than if 1 knew his actual cost.

Example: Incomplete Information Cournot

- An auction is a way to sell a fixed supply of a **commodity** (an item to be sold) for which there is no well-established ongoing market
- **Bidders make bids**
  - proposals to pay various amounts of money for the commodity
- The commodity is sold to the bidder who makes the largest bid
- Example applications
  - Real estate, art, oil leases, electricity, eBay, google ads
- Several kinds of auctions are incomplete-information, and can be modeled as **Bayesian games**