Econ 421
Fall, 2023
Li, Hao
UBC

Lecture 7. Static Games of Incomplete Information

1. Bayesian game

Static games of incomplete information, where players

- simultaneously choose actions and the game ends;
- have preferences over actions chosen by all of the players;
- have incomplete, and often asymmetric, information about game.

- **Example** (Cournot duopoly with private cost). Consider two firms, 1 and 2, choosing quantities q_1 and q_2 simultaneously. The inverse demand function is given by $p = a - (q_1 + q_2)$. Firm 1's marginal cost is known to be c, with no fixed cost, while Firm 2 has private knowledge about its marginal cost. In particular, Firm 1 only knows that Firm 2's cost is either c_H or $c_L < c_H$, with probabilities p_H and p_L respectively ($p_H + p_L = 1$).

- **Example** (First price auction with independent private values). Consider a sealed-bid first price auction. Two potential buyers bid for an indivisible good. with his own valuation being 0; Each has private knowledge about his own valuation for the good, but knows only that the other bidder's valuation is either v_H or $v_L < v_H$, with probabilities p_H and p_L respectively ($p_H + p_L = 1$). The bidder with the higher bid wins the object and pays his bid; if two bids are same, each gets object with probability $\frac{1}{2}$ and pays the bid. The payoff for the winning bidder is his valuation less his bid, and the losing bidder pays nothing and has payoff 0.

- Model of static games of incomplete information: *Bayesian* game

- As in static games of complete information, we need to specify

  - A set of players.

  - A set of available actions for each player.

  - A payoff function for each player.

- New element is how to model incomplete information of players have about relevant aspects of the game.

Formally, a Bayesian game consists of

- A set of players, $i = 1, \ldots, n$.
- Each player i can be any *type* from a *type space* T_i.
  - *Prior probability* of each *type profile* $t = (t_1, \ldots, t_n)$ is $p(t)$.
  - Each i knows his type own t_i and *conditional probability distribution* $p(t_{-i}|t_i)$ of opponents' type profile.
- Each player i chooses a_i from action space A_i.
- Player i's vNM payoff from action profile $a = (a_1, \ldots, a_n)$ is $u_i(a; t)$.

Remarks on Bayesian game

- Incomplete information is modeled by uncertainty about types of opponents.
  - Type is payoff relevant: $u_i(a;t)$ depends on type profile as well as on action profile.
  - Type is general concept used to model incomplete information about: which game you are playing; whether your opponents are "rational" or not.

- Conditional probability distribution $p(t_{-i}|t_i)$ is derived from $p(t)$ according to Bayes rule

$$p(t_{-i}|t_i) = \frac{p(t_i, t_{-i})}{\sum_{\tilde{t}_{-i} \in T_{-i}} p(t_i, \tilde{t}_{-i})}.$$

  - A special case is independent types: no Bayesian updating, with $p(t_{-i}|t_i)$ does not depend on t_i, because $p(t) = \prod_i p_i(t_i)$ for some given profile of probability distributions $(p_1, \ldots, p_n)$.

- General model of Bayesian games: state space approach.
  - In most economic applications, payoff type is sufficient to capture all relevant incomplete information.
  - This is true both in Cournot duopoly with private cost, and in First price auction with independent private values.
  - When payoff type is insufficient, a more general approach based on state space is needed.

A general model of Bayesian games.

- A set of players, $i = 1, \ldots, n$.
- A state space Ω, with a prior probability $\pi(\omega)$ of each state $\omega \in \Omega$.
  - A type space T_i for each player i, and a signal function $\tau_i(\cdot)$ that maps each state $\omega \in \Omega$ to a type $\tau_i(\omega) \in T_i$.
- Each player i chooses a_i from action space A_i.
- Player i's vNM payoff from action profile $a = (a_1, \ldots, a_n)$ is $u_i(a; \omega)$.

Remarks on the state-space model.

- State space Ω is a complete description of all states of the world that are relevant to the game, not just in terms of payoffs.

- Incomplete information of each player is modeled by the signal function, which partitions the state space into type space.

- Type-space model is a special case of state-space model: take Ω to be set of all type profiles, with $\pi(\omega) = p(t)$, and for each i, take the signal function $\tau_i(\cdot)$ be such that $\tau_i(\omega) = \tau_i(t_1, \ldots, t_n) = t_i$.

An example of state-space model of Bayesian games.

- Two players, R and C, and one of two games Γ_1 and Γ_2 is to be played, where each player has the same two possible actions, r_1 and r_2 for R and c_1 and c_2 for C, with different payoffs.

- Consider following different scenarios: neither R nor C knows which game is being played; only R knows which game is being played; R knows which game is being played but does not know if C knows which game is being played; R knows which game is being played and C does not know if R knows that C knows which game is being played.

2. Bayesian Nash equilibrium

- Each player i knows his type t_i, has belief $p(t_{-i}|t_i)$ about the opponents' types, and must choose an action from A_i.

- A *pure strategy* of player i is a mapping $s_i : T_i \to A_i$; denote the set of all strategies as S_i.

  - A strategy s_i of player i is a complete plan: it specifies an action $s_i(t_i)$ for each of his possible type t_i.

- Why is it necessary to define a strategy as a complete plan, even though each player playing the Bayesian game knows his own type?

- *Bayesian Nash equilibrium* is simply NE of Bayesian game.

- Formally, a Bayesian Nash equilibrium is a strategy profile $(s_1^*, \ldots, s_n^*)$, such that for each i, the strategy s_i^* solves

$$\max_{s_i \in S_i} \sum_t u_i(s_i(t_i), s_{-i}^*(t_{-i}); t)p(t),$$

where $s_{-i}^*(t_{-i})$ stands for $(s_1^*(t_1), \ldots, s_{i-1}^*(t_{i-1}), s_{i+1}^*(t_{i+1}), \ldots, s_n^*(t_n))$.

- Bayesian Nash equilibrium is a profile of plans, made before each i is informed of his type t_i, such that no player has a profitable unilateral deviation from his plan.

Proposition (Alternative definition of BNE). A strategy profile $(s_1^*, \ldots, s_n^*)$ is a Bayesian Nash equilibrium if and only if for each i and each $t_i \in T_i$, the action $s_i^*(t_i)$ solves $\max_{a_i \in A_i} \sum_{t_{-i}} u_i(a_i, s_{-i}^*(t_{-i}); t_i, t_{-i}) p(t_{-i}|t_i)$.

- *Proof.* Writing $p(t_i) = \sum_{t_{-i}} p(t_i, t_{-i})$, for each i and each $s_i \in S_i$,

$$\sum_t u_i(s_i(t_i), s_{-i}^*(t_{-i}); t_i, t_{-i}) p(t) = \sum_{t_i} p(t_i) \left[\sum_{t_{-i}} u_i(s_i(t_i), s_{-i}^*(t_{-i}); t) p(t_{-i}|t_i) \right].$$

- Interpretation: each type of a player is a separate player.

- The alternative definition is simpler to apply in practice.

3. Cournot duopoly game with incomplete information

- The Bayesian game is given by:
  - Players: Firm 1 and Firm 2.
  - Firm 1's type space is a singleton set; and Firm 2's type space is $\{c_H, c_L\}$ with probabilities p_H and p_L.
  - Action space for both firms is $A_1 = A_2 = [0, \infty)$.
  - Payoff function for Firm 1 is $\pi_1(q_1, q_2) = (a - c - q_1 - q_2)q_1$, and for Firm 2 is $\pi_2(q_1, q_2; c_2) = (a - c_2 - q_1 - q_2)q_2$ for each type $c_2 = c_H, c_L$.

- Firm 1's strategy is just a quantity q_1; Firm 2's strategy is a plan (q_{2H}, q_{2L}).

- Best responses.

  - Type c_H Firm 2's best response to q_1 is

$$b_H(q_1) = \begin{cases} \frac{1}{2}(a - c_H - q_1) & \text{if } q_1 < a - c_H, \\ 0 & \text{if } q_1 \geq a - c_H. \end{cases}$$

  - Type c_L Firm 2's best response function $b_L(q_1)$ is similar, with c_L replacing c_H in above.

– Firm 1's best response to Firm 2's strategy (q_{2H}, q_{2L}) solves

$$\max_{q_1} p_H(a - c - q_{2H} - q_1)q_1 + p_L(a - c - q_{2L} - q_1)q_1.$$

– Solution is

$$b_1(q_{2H}, q_{2L}) = \begin{cases} \frac{1}{2}(a - c - p_H q_{2H} - p_L q_{2L}) & \text{if postive,} \\ 0 & \text{otherwise.} \end{cases}$$

- BNE is the intersection of best response functions:

$$\begin{cases} q_{2H}^* = \frac{1}{2}(a - c_H - q_1^*), \\ q_{2L}^* = \frac{1}{2}(a - c_L - q_1^*), \\ q_1^* = \frac{1}{2}(a - c - p_H q_{2H}^* - p_L q_{2L}^*). \end{cases}$$

- Bayesian Nash equilibrium if solutions are positive:

$$\begin{cases} q_{2H}^* = \frac{1}{3}(a - 2c_H + c) + \frac{1}{6}(c_H - c_L)p_L, \\ q_{2L}^* = \frac{1}{3}(a - 2c_L + c) - \frac{1}{6}(c_H - c_L)p_H, \\ q_1^* = \frac{1}{3}(a - 2c + p_H c_H + p_L c_L). \end{cases}$$

- Comparison with Nash equilibrium when $c_2 = c_H$ and known: $q_1 = \frac{1}{3}(a - 2c + c_H)$, $q_2 = \frac{1}{3}(a - 2c_H + c)$.
  - Firm 1 produces less under incomplete information, because it is unsure about c_2 and c_2 could be low; knowing this, Firm 2 produces more than it would than when $c_2 = c_H$ is known.
  - Incomplete information hurts Firm 1 and benefits Firm 2.
- Opposite conclusions hold in comparison with NE when $c_2 = c_L$ and known: $q_1 = \frac{1}{3}(a - 2c + c_L)$, $q_2 = \frac{1}{3}(a - 2c_L + c)$.

4. Bayesian Nash equilibrium in state-space model

- A strategy s_i of i is still a complete plan: it specifies an action $s_i(t_i) \in A_i$ for each type $t_i \in T_i$.

- Expected payoff for i from strategy s_i against opponents strategy profile s_{-i} is

$$\sum_{\omega \in \Omega} u_i(s_i(\tau_i(\omega)), s_{-i}(\tau_{-i}(\omega)); \omega)\pi(\omega),$$

where $s_{-i}(\tau_{-i}(\omega))$ is action profile where each player $j \neq i$ plays action $s_j(\tau_j(\omega))$ in state ω.

- Strategy profile s^* is Bayesian Nash equilibrium if each strategy s_i^* maximizes i's expected payoff given s_{-i}^*.

- Alternative definition: s^* is Bayesian Nash equilibrium if for each player i and each type t_i, action $s_i^*(t_i)$ maximizes t_i's expected payoff $\sum_{\omega \in \Omega} u_i(a_i, s_{-i}(\tau_{-i}(\omega)); \omega)\pi(\omega|t_i)$ over action set A_i, where $\pi(\omega|t_i)$ is t_i's updated belief of state ω, given by Bayes' rule

$$\pi(\omega|t_i) = \frac{\pi(\omega)}{\sum_{\tilde{\omega} \in \Omega : \tau_i(\tilde{\omega}) = t_i} \pi(\tilde{\omega})}.$$