INSE6441	Jia Yuan Yu
6: Bayesian Games	
Concordia	March 14, 2017

Last time, we introduced the notion of subgame, and subgame perfection. Subgame perfect equilibria can be found by backward induction (iterated best response starting from the last player): we have illustrated this approach with the bargaining game, and Stackelberg duopoly. We also illustrated how information sets model the notion of private information or signals from Nature in a version of one-card poker (Borel's game of relance). Let's see how backward induction works in games with private information.

Example 0.1 (Voting game)**.** Vote for preferred midterm date.

1 Borel's game of relance

Let a, b denote the strategies for players 1 and 2. This is a zero-sum game, so $u_1(a, b) = -u_2(a, b)$. The expected payoff is computed as follows:

$$\begin{aligned} u_1(a,b) &= \mathbb{P}(\text{P1 folds})(-1) + \mathbb{P}(\text{P1 bets, P2 folds})(1) \\ &\quad + \mathbb{P}(\text{P1 bets, P2 calls, } X > Y)(1+\beta) + \mathbb{P}(\text{P1 bets, P2 calls, } X < Y)(-1-\beta) \\ &= -a + (1-a)b \\ &\quad + (1-b)(1-b)\frac{1}{2}(1+\beta) - \left((b-a)(1-b) + (1-b)(1-b)\frac{1}{2}\right)(1+\beta) \\ &= a\beta - b\beta - ab(2+\beta) + b^2(1+\beta), \\ u_2(a,b) &= -u_1(a,b). \end{aligned}$$

We start with P2's best response[1]. Given a fixed strategy a for P1, the best-response for P2 is[2]

$$\begin{aligned} b^*(a) &\in \arg\max_b u_2(a,b) && (1) \\ (\frac{\mathrm{d}}{\mathrm{d}b} u_2(a,b) = 0) \quad &= \frac{\beta + a(2+\beta)}{2(1+\beta)}. && (2) \end{aligned}$$

By plugging in b^* in u_1, we obtain:

$$\begin{aligned} u_1(a, b^*) &= a\beta - b^*\big(\beta + a(2+\beta)\big) + (b^*)^2(1+\beta) \\ &= a\beta - \frac{\beta + a(2+\beta)}{2(1+\beta)}\big(\beta + a(2+\beta)\big) + \left(\frac{\beta + a(2+\beta)}{2(1+\beta)}\right)^2 (1+\beta). \end{aligned}$$

[1] What would happen if we start with P1?

[2] Check that second derivative is positive for a minimum.

Solve $\frac{\mathrm{d}}{\mathrm{d}a}u_1(a, b^*) = 0$[3]:

$$\begin{aligned}
\beta - \frac{(2+\beta)\beta}{2(1+\beta)} - \frac{\beta(2+\beta)}{2(1+\beta)} - \frac{2a(2+\beta)^2}{2(1+\beta)} + \frac{2\beta(2+\beta) + 2a(2+\beta)^2}{4(1+\beta)} &= 0 \\
\implies \quad \beta - \frac{2(2+\beta)\beta}{2(1+\beta)} + \frac{-a(2+\beta)^2 + (2+\beta)\beta}{2(1+\beta)} &= 0 \\
\implies \quad \beta - \frac{(2+\beta)\beta}{2(1+\beta)} &= \frac{a(2+\beta)^2}{2(1+\beta)} \\
\implies \quad 2(1+\beta)\beta - (2+\beta)\beta &= a(2+\beta)^2 \\
\implies \quad a^*(b^*) &= \frac{\beta^2}{(2+\beta)^2}.
\end{aligned}$$

In turn, by substituting into (2), we obtain

$$\begin{aligned}
b^*(a^*) &= \frac{\beta + a^*(2+\beta)}{2(1+\beta)} \\
&= \frac{\beta + \frac{\beta^2}{(2+\beta)}}{2(1+\beta)} \\
&= \frac{\beta}{2+\beta}.
\end{aligned}$$

Remark 1. As we can see by plugging in (a^*, b^*) into u_1, P1 has an advantage.

Remark 2. If P2's card comes from the same deck as P1, then X and Y are not independent random variables!

Remark 3 (Strategic form of relance). The game of relance also has an equivalent strategic form, where both players fix their strategics simultaneously, and then receive random cards.

Dynamic games where some moves (by Nature or an opponent) are hidden from a player are said to have imperfect information[4]. We next introduce the notion of player types, where players do not possess full information about their opponents' payoff or utility (i.e., players have private information or signals). For example, in auctions, the valuations are private. These are called games with incomplete information. Some games with incomplete information (Bayesian games) can be modeled as extensive-form games with imperfect information with Nature moves.

2 Static games with incomplete information

In this section, we consider settings where players possess private information. First, we consider the setting of Bayesian games, where player i's private information or player type is revealed before they play simultaneously (as in static—also called one-shot or strategic-form—games).

Formally, a Bayesian game is described by

[3] Check that second derivative is negative.

[4] Nonsingleton information sets.

- I players
- Set of actions S_i for each player i
- Set of possible types[5] T_i for each player i
- A belief function p_i for each player i
- A utility function u_i for each player i

The belief function p_i models the player i's belief about the probability distribution of other players' types[6]:

$$p_i : T_i \to \Delta(T_{-i}).$$

For instance, $p_i(t_{-i} \mid t_i)$ denotes the estimated probability that the other players have drawn cards $t_{-i} = (t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_I)$ from a deck, given that player i drew the card t_i from the deck.

The utility function

$$u_i : S \times T \to \mathbb{R},$$

maps the strategy profile and type profile to real values, so that $u_i(s, t)$ denotes the payoff to player i corresponding to the action profile $s \in S$ when the type profile is $t \in T$.

A pure strategy for player i in a Bayesian game is a mapping[7]:

$$\pi_i : T_i \to S_i.$$

A mixed strategy is:

$$\sigma_i : T_i \to \Delta(S_i).$$

For simplicity, we only consider the case where all beliefs are consistent.

Assumption 2.1 (Consistent beliefs)**.** Suppose that T_i is finite and that there exists a probability distribution p such that, for all i,

$$p_i(t_{-i} \mid t_i) = \frac{p(t_i, t_{-i})}{p(t_i)}$$

and the marginal and conditional probabilities are well-defined. In this case, we write $p_i(t_{-i} \mid t_i) \triangleq p(t_{-i} \mid t_i)$.

[5] cf. cards in a deck

[6] cf. probability distribution of opponent's poker hand

[7] cf. the equilibrium strategy in Borel's Relance.

Definition 2.1 (Bayesian equilibrium)**.** A pure Bayesian equilibrium is a strategy profile $(\pi_1, \ldots, \pi_I)$ such that for every player i and every type $t_i \in T_i$, we have

$$\pi_i(t_i) \in \arg\max_{z \in S_i} \underbrace{\sum_{t_{-i} \in T_{-i}} p(t_{-i} \mid t_i) u_i(z, \pi_{-i}(t_{-i}), t)}_{\mathbb{E}_{\mathbf{t}_{-i} \sim p(\cdot \mid t_i)} u_i(z, \pi_{-i}(\mathbf{t}_{-i}), t_i, \mathbf{t}_{-i}))}. \tag{3}$$

A mixed Bayesian equilibrium is a profile $\sigma = (\sigma_1, \ldots, \sigma_I)$ such that for every player i and every type $t_i \in T_i$, we have

$$\sigma_i(t_i) \in \arg\max_{z \in \Delta(S_i)} \sum_{t_{-i} \in T_{-i}} p(t_{-i} \mid t_i) \tilde{u}_i(z, \sigma_{-i}(t_{-i}), t).$$

Example 2.1 (Cournot duopoly with incomplete information (Example 6.2 of FT))**.** Let t_1 and t_2 denote the type signals of Firm 1 and 2. The price function is $\pi(q) = 1 - q$. It is common knowledge that Firm 1's production cost is $t_1 = c_1$ per unit (complete information) and whereas Firm 2's per-unit production cost is a random variable with the distribution[8]:

$$t_2 = \begin{cases} w & \text{w.p. } x, \\ v & \text{w.p. } 1 - x. \end{cases}$$

The two firms decide on q_1 and q_2 simultaneously. The utility functions of the two firms are

$$\begin{aligned} u_1(q_1, q_2, t_1, t_2) &= q_1 \pi(q_1 + q_2) - t_1 q_1 = q_1(1 - t_1 - q_1 - q_2) \\ u_2(q_1, q_2, t_1, t_2) &= q_2(1 - t_2 - q_1 - q_2). \end{aligned}$$

From the Bayesian equilibrium definition, the best-response functions in (3) are:

$$\begin{aligned} q_2(w) &= \arg\max_z \quad z(1 - w - q_1(c_1) - z), \\ q_2(v) &= \arg\max_z \quad z(1 - v - q_1(c_1) - z), \\ q_1(c_1) &= \arg\max_z \quad xz(1 - c_1 - z - q_2(w)) + (1 - x) z(1 - c_1 - z - q_2(v)). \end{aligned}$$

By taking derivatives and setting to zero, we obtain:

$$\begin{aligned} q_2(v) &= \frac{1 - v - q_1(c_1)}{2}, \\ q_2(w) &= \frac{1 - w - q_1(c_1)}{2}, \\ q_1(c_1) &= \frac{x(1 - c_1 - q_2(w)) + (1 - x)(1 - c_1 - q_2(v))}{2}. \end{aligned}$$

[8]The information available to players are asymmetric.

For instance, if $x = 1/2$, $c_1 = 1/2$, $w = 3/4$, and $v = 1/4$, we get

$$q_2(1/4) = \frac{3/4 - q_1(1/2)}{2}$$
$$q_2(3/4) = \frac{1/4 - q_1(1/2)}{2}$$
$$q_1(1/2) = \frac{1/2 - q_2(1/4)/2 - q_2(3/4)/2}{2}.$$

Solving these three equations in three unknowns gives the following Bayesian equilibrium:

$$q_1(1/2) = 1/6, \quad q_2(1/4) = 7/24, \quad q_2(3/4) = 1/24.$$

3 Reading material

- Chapter 4.4, 6 of Fudenberg and Tirole.