Advanced Microeconomics

ECON5200 - Fall 2014

Static Games with Incomplete Information

There are many circumstances in which agents have private information. Some examples are:

- A bidder does not know the other bidders' value in an auction;
- Parties do not know the voters' preferences;
- An employer does not know the skills of the employee;
- The incumbent firm does not know whether the entrant is aggressive or not;
- ....

Bayesian Games

- N players with $i \in I \equiv \{1, ..., N\}$;
- $\omega \in \Omega$ finite set of "states of nature";
- $\tau_i : \Omega \rightarrow T_i$ types (signal) profile with $t_i \in T_i$;
- $p_i : \Omega \rightarrow [0, 1]$ prior belief with $p_i (\omega | t_i) \geq 0$
- $\sigma \in \Delta (S) \equiv \prod_{i=1,..,N} \Delta (S_i)$ strategy profile with $\sigma_i : T_i \rightarrow \Delta (S_i)$;
- $\upsilon_{t_i} \equiv \sum_{\omega \in \Omega} p_i (\omega | t_i) u_i (\sigma, \omega)$ the expected payoff of type t_i;
- $G \equiv \left\langle I, \Omega, \{S_i\}_i, \{T_i\}_i, \{\tau_i\}_i, \{p_i\}_i, \{\upsilon_{t_i}\}_{t_i} \right\rangle$

Bayesian Games: Interpretation

- Ω is a set of possible states of nature that determine the physical setup of the game (payoffs);
- T_i is the set of i 's private types that encode player i 's information/knowledge;
- p_i is player i 's interim belief about the state and the other players' types.

Battle of the Sexes Revisited

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

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

- $\omega \in \Omega \equiv \{\omega_1, \omega_2\}$ with $\omega_1 = meet$ and $\omega_2 = avoid$;
- $\tau_1(\omega_1) = \tau_1(\omega_2) = z$;
- $m = \tau_2(\omega_1) \neq \tau_2(\omega_2) = x$;
- $p_1(\omega_1|z) = p_1(\omega_2|z) = 1/2$, $p_2(\omega_1|m) = p_2(\omega_2|x) = 1$;
- $(1/2)\,Eu_1((B, \sigma_2), \omega_1) + (1/2)\,Eu_1((B, \sigma_2), \omega_2)$ player 1's ex-ante utility if she plays B.

Bayesian Nash Equilibrium

Definition (Harsanyi (1967/1968))

A Nash equilibrium of a Bayesian Game is a Nash equilibrium of a strategic game characterized by:

- Set of players (i, t_i) with $i \in I$ and $t_i \in T_i$;
- Set of strategies for each (i, t_i);
- Payoff function for each (i, t_i) is given by υ_{t_i}.

Following Harsanyi (1967/1968) we transform a game of *incomplete information* in a game with *imperfect information* where Nature moves first.

Bayesian Nash Equilibrium

Definition

$\sigma^* \in \Delta(S)$ is a Bayesian Nash Equilibrium if:

$$E\left[v_{t_i}\left(\sigma_i^*\left(t_i\right), \sigma_{-i}^*\left(\tilde{t}_{-i}\right), \tilde{\omega}\right)\right] \geq E\left[v_{t_i}\left(\sigma_i, \sigma_{-i}^*\left(\tilde{t}_{-i}\right), \tilde{\omega}\right)\right]$$

for each $\sigma_i \in \Delta(S_i)$, $t_i \in T_i$ and $i \in I$.

Example (5 - Building New Capacity - See notes!)

	Enter	*Don't*
Build	0, −1	2, 0
Don't	2, 1	3, 0

high-cost

	Enter	*Don't*
Build	1.5, −1	3.5, 0
Don't	2, 1	3, 0

low-cost

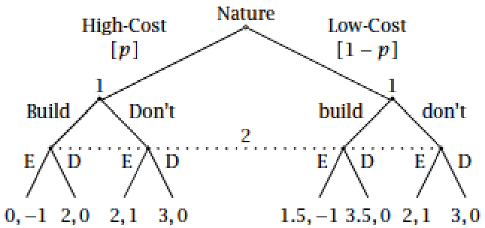

Example (6 - Public Good Provision - Proposed as exercise!)

- There are two players, $i = 1, 2$, who may either cooperate or defeat in the provision of a public good;
- $s_i \in S_i \equiv \{0, 1\}$ is the players' strategy space, where 0 stands for "defeat" and 1 for "cooperate";
- If agents decide to cooperate, then they sustain a cost c_i, which is private information;
- Common-Knowledge: $c_i \sim P(\cdot)$ over $[\underline{c}, \overline{c}]$ with $\underline{c} < 1 < \overline{c}$;
- The individual payoff is $u_i(s_i, s_j, c_i) = \max(s_1, s_2) - c_i s_i$;
- Find the BNE of the public good game.

Example (7 - Second-Price vs First-Price Auction - *Proposed as exercise*)

- n bidders whose private evaluation is $\underline{v} \leq v_i \leq \overline{v}$ make a bid $b_i \geq 0$;
- Each bidder observes only his own evaluation but believes that the others' evaluations are iid and distributed according to $F \sim [\underline{v}, \overline{v}]$;
- The player with the highest bid wins the auction by paying the second highest bid;
- Find:
  1. that $b_i = v_i$ is a weakly dominant strategy;
  2. the BNE of a first-price auction (i.e. the player with the highest bid wins the auction by paying his own bid).

Dynamic Games with Perfect Information

- We study dynamic games where players make a choice sequentially;
- We assume perfect information: Each player can perfectly observe the past actions;
- Best representation by using extensive form games.

Dynamic Games with Perfect Information

Chain-Store Game

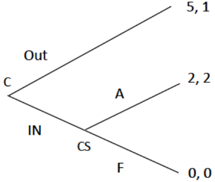

Stackelberg-Cournot Competition

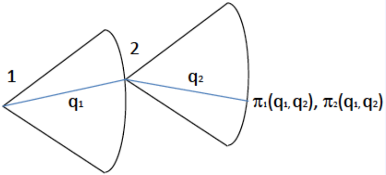

Dynamic Games in Extensive Form

- N players with $i \in I \equiv \{1, ..., N\}$;
- H set of histories with a^k equal to an action taken by a player:
  - $\emptyset \in H$;
  - if $\left(a^1, ...a^k\right) \in H$ then $\left(a^1, ...a^l\right) \in H$ for each $l < k$;
  - if $\left(a^1, ...a^k, ...\right)$ is an infinite sequence such that $\left(a^1, ...a^k\right) \in H$ for each $k \in \mathbb{N}$ then $\left(a^1, ...a^k, ...\right) \in H$.
- Z set of terminal histories:
  - $\left(a^1, ...a^k\right) \in Z$ if it is an infinite sequence or $\nexists\ a^{k+1}$ such that $\left(a^1, ...a^{k+1}\right) \in H$.

Dynamic Games in Extensive Form

- $P : H \backslash Z \rightarrow I$ assignment function;
- $A(h) = \{a | (h, a) \in H\}$ set of actions available to $P(h)$;
- $v_i : Z \rightarrow \mathbb{R}$;
- $\Gamma \equiv \langle I, H, P, \{v_i\}_i \rangle$.

Strategies

Definition

A strategy of player $i \in I$ in Γ, σ_i, is a mapping from H to a distribution on the set of available action, $\sigma_i(h) \in \Delta(A_i(h))$ for each non terminal history $h \in H \backslash Z$ for which $P(h) = i$ (complete contingent plan).

For each strategy profile in Γ, let $\mathcal{O}(\sigma)$ the outcome of σ.

Nash Equilibrium

Definition

A Nash equilibrium of a dynamic game with perfect information Γ is a strategy profile σ^* such that for each $i \in I$ and for each σ_i, $\mathcal{O}\left(\sigma^*\right) \geq_i \mathcal{O}\left(\sigma_i, \sigma^*_{-i}\right)$.

Theorem (Zermelo 1913, Kuhn 1953)

A finite dynamic game of perfect information has a pure-strategy Nash equilibrium.

Backward Induction

Backward induction is the following procedure:

- Let $L < \infty$ be the maximum length of all histories;
- Find all nonterminal histories of $L - 1$ length and assign an optimal action there. Eliminate unreached L-length terminal histories and regard other $L-$length terminal histories as $L - 1$-length terminal histories;
- Find all nonterminal histories of $L - 2$ length and assign an optimal action there. Eliminate unreached $L - 1$-length terminal histories and regard other $L - 1$-length terminal histories as $L - 2$-length terminal histories;
- ....

Example (9 - Stackelberg-Cournot Game - See notes!)

....

Example (10 - Hotelling Game and Product Differentiation - Proposed as exercise!)

- Consumers are distributed uniformly along the interval $[0, 1]$;
- Two firms are located at the extremes and compete on prices;
- c is the cost of 1 unit of good and t is the transportation cost by unit of distance squared;
- Consumers' payoff is $U = s - p - td^2$ where s is the max willingness to pay, p is the market price and d is the distance;
- Find;
    1. The NE of the game when firms' location is exogenously given;
    2. The SPE of the game when firms decide first their location and then compete on prices.