Advanced Microeconomics

ECON5200 - Fall 2014

Subgame

Definition

The subgame of Γ following $h \in H$ is the extensive-form game $\Gamma(h) \equiv \left\langle I, H|_h, P_{|h}, \left\{v_{i|h}\right\}_i \right\rangle$ where:

- $h' \in H|_h \Leftrightarrow (h, h') \in H$;
- $P_{|h}(h') = P(h, h')$ for each $h' \in H|_h$;
- $v_{i|h}(h') = v_i(h, h')$ for each $h' \in Z|_h \subset H|_h$.

Let $\sigma_{i|h}$ a strategy for player i of $\Gamma(h)$ and $\mathcal{O}_{|h}\left(\sigma_{|h}\right)$ the outcome of $\sigma_{|h}$.

Subgame Perfect Equilibrium

Definition

A subgame perfect equilibrium of an extensive form game with perfect information Γ is a strategy σ^* such that for any $i \in I$ and non terminal history $h \in H \backslash Z$ for which $P(h) = i$, one has:

$$\mathcal{O}_h\left(\sigma^*_{|h}\right) \geq_{i|h} \mathcal{O}_h\left(\sigma_i, \sigma^*_{-i|h}\right)$$

for all strategy σ_i in the subgame $\Gamma(h)$.

One-Shot-Deviation Principle

- To find a SPE we need to check a very large number of incentive constraints;
- We can apply a principle of dynamic programming: OSDP;

Definition

σ_i' in $\Gamma(h)$ at $h \in H \backslash Z$ for $i \in P(h)$ is called one-shot deviation from σ_i if $\sigma_{i|h}$ and σ_i' prescribe a different action *only* at the initial history (i.e. $\sigma_i'(\emptyset) \neq \sigma_i(h)$ and $\sigma_i'(h') = \sigma_i(h, h')$ for any $h' \neq \emptyset$ with $(h, h') \in H \backslash Z$).

One-Shot-Deviation Principle

Theorem

In an extensive form game with perfect information Γ *a strategy* σ^* *is a SPE iff:*

$$\mathcal{O}_h\left(\sigma^*_{|h}\right) \succsim_{i|h} \mathcal{O}_h\left(\sigma_i, \sigma^*_{-i|h}\right)$$

for any one-shot deviation σ_i *from* $\sigma^*_{i|h}$ *at any* $h \in H \backslash Z$ *for* $i \in P(h)$.

Example (11 - Bargaining Game - Proposed as exercise!)

- Two players use the following procedure to split $1kr$:
  - Players 1 offers player 2 an amount $x \in [0, 1]$;
  - If player 2 accepts, then 1 gets $1 - x$, if 2 refuses neither receives any money;
- Find:
  1. The SPE of the bargaining game;
  2. Introduce the possibility of player 2 to make a counter-offer. Let δ_i be the individual discount factor. Find the SPE;
  3. Find the SPE of the infinitely repeated version.

Infinite Repeated Game

Through the infinite repeated version of the dynamic game with perfect information we can answer to the following questions:

- When can people cooperate in a long-term relationship?
- What is the most efficient outcome that arises as an equilibrium?
- What is the set of all outcomes that can be supported in equilibrium?

Infinite Repeated Game

- N players with $i \in I \equiv \{1, ..., N\}$;
- $a^t \in A \equiv \prod_{i=1,..,N} A_i$, $a_i^t \in A_i$ finite set;
- $u_i\left(a^t\right)$ payoff or utility;
- $G^t \equiv \langle I, \{A_i\}_i, \{u_i\left(a^t\right)\}_i\rangle$ stage-game;
- $\mathcal{F}^t \equiv co\left\{u\left(a^t\right) | \forall a^t \in A\right\}$ set of feasible payoffs;
- An infinite repeated game, $G^\infty\left(\delta\right)$, is equal to the infinite repetition of G^t, where $\delta \in (0, 1)$ is the individual discount factor.

Strategy

- See def. of strategy and SPE for dynamic games with complete information;
- An equilibrium strategy profile σ generates an infinite sequence of action profiles $\left(a^1, a^2, ...\right) \in A^\infty$;
- The discounted average payoff is given by:

$$V_i\left(\sigma\right) = (1-\delta)\sum_{t=0}^{\infty}\delta^t u_i\left(a^t\right)$$

Min-Max Payoff

Definition

The min-max payoff is equal to:

$$\underline{v}_i = \min_{a_{-i}} \max_{a_i} u_i\left(a\right)$$

- In the prisoner dilemma $\underline{v}_i = 0$ for $i = 1, 2$;
- The min-max payoff serves as a lower bound on equilibrium payoffs in a repeated game.

Lemma

Player i's payoff in any NE for $G^{\infty}\left(\delta\right)$ is at least as large as $\underline{v}_i$.

Example (12 - Infinite Repeated Prisoner Dilemma - See notes!)

	C	D
C	$1,1$	$-\ell, 1+g$
D	$1+g, -\ell$	$0,0$

$g>0$, $l>0$

When can (C, C) be played in every period in equilibrium?

Folk Theorem

We know that player i 's (pure strategy) SPE payoff is never strictly below $\underline{v}_i$. The Folk Theorem shows that every feasible v_i strictly above $\underline{v}_i$ can be supported by SPE.

Definition

$v \in \mathcal{F}$ is strictly individually rational if v_i is strictly larger than $\underline{v}_i$ for all $i \in I$. Let $\mathcal{F}^* \subset \mathcal{F}$ be the set of feasible and strictly individually rational payoff profiles.

Folk Theorem

Theorem (Fudenberg and Maskin (1986))

Suppose that $\mathcal{F}^*$ is full-dimensional. For any $v \in \mathcal{F}^*$, there exists a strategy profile σ and $\underline{\delta} \in (0,1)$ such that σ is a SPE and achieves v for any $\delta \in (\underline{\delta}, 1)$.

Example (13 - Optimal Collusion - Infinite Cournot Competition - Proposed as exercise!)

- Dynamic Cournot duopoly model;
- Stage game equal to the static Cournot game and $\delta \in (0, 1)$;
- Using a "stick and carrot" strategy find the strongly symmetric SPE.

Dynamic Games with Incomplete Information

- We consider dynamic games where past actions (by players or nature) are imperfectly observed;
- We treat them as an extension of dynamic game with complete information.

Dynamic Games with Incomplete Information

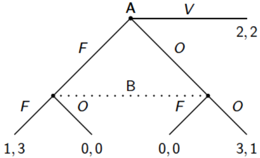

- There are no subgame out of the game itself;
- The pure NE are (O, O) and (V, F), but is the latter credible?

Dynamic Games with Incomplete Information

- N players with $i \in I \equiv \{1, ..., N\}$ and c denotes Nature;
- $h_t = \left(a^1, a^2, ..., a^k\right) \in H$ set of histories with a^k equal to an action taken by a player;
- $P : H \backslash Z \rightarrow I \cup \{c\}$, $A(h) = \{a | (h, a) \in H\}$ set of actions available to $P(h)$;
- $f_c(a|h)$ is the probability that a occurs after h for which $P(h) = c$;
- $v_i : Z \rightarrow \mathbb{R}$.

Dynamic Games with Incomplete Information

- $\mathcal{I}_i$ a partition of $\{h \in H | P(h) = i\}$ with the property $A(h) = A(h')$ if $h, h' \in \mathcal{I}_i$;
- Each $I_i \in \mathcal{I}_i$ is player i's information set: the set of histories that player i cannot distinguish;
- $A(I_i)$ the set of action available at I_i;
- $\Gamma \equiv \langle I, H, P, f_c, \{\mathcal{I}_i\}_i, \{v_i\}_i \rangle$.

Extensive Games with Imperfect Information

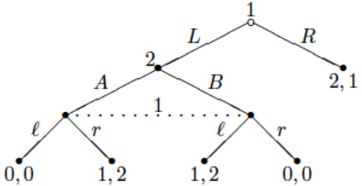

- $P(\varnothing) = P(L, A) = P(L, B) = 1$ and $P(L) = 2$;
- $\mathcal{I}_1 = \{\{\varnothing\}, \{(L, A), (L, B)\}\}$ and $\mathcal{I}_1 = \{\{L\}\}$.

Mixed and Behavioral Strategies

Definition

A *mixed strategy* of player i in an extensive game $\langle I, H, P, f_c, \{\mathcal{I}_i\}_i, \{v_i\}_i\rangle$ is a probability measure over the set of player i's pure strategies. A *behavioral strategy* of player i is a collection $\{\beta_i(I_i)\}_{I_i \in \mathcal{I}_1}$ of independent probability measures, where $\beta_i(I_i)$ is a probability measure over $A(I_i)$.

Theorem

For any mixed strategy of a player in a finite extensive game with perfect recall there is an outcome-equivalent behavioral strategy.

The Nash equilibrium of the game can be found in the usual way. We need a reasonable refinement of NE.

Perfect Bayesian Nash Equilibrium

- Recall that in games with complete information some NE may be based on the assumption that some players will act sequentially irrationally at certain information sets off the path of equilibrium;
- In those games we ignored these equilibria by focusing on SPE;
- We extend this notion to the games with incomplete information by requiring sequential rationality at each information set: PBNE as equilibrium refinement of BNE;
- For each information set, we must specify the beliefs of the agent who moves at that information set.

Sequential Rationality

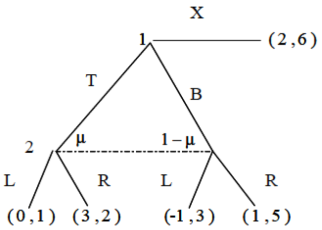

Definition

A player is said to be sequentially rational iff, at each information set he is to move, he maximizes his expected utility given his beliefs at the information set (and given that he is at the information set) - even if this information set is precluded by his own strategy.

Consistency

Definition

Given any strategy profile σ, an information set is said to be on the path of play iff the information set is reached with positive probability according to σ.

Definition

Given any strategy profile σ and any information set I_i on the path of play of σ, a player's beliefs at I_i is said to be consistent with σ iff the beliefs are derived using the Bayes' rule and σ.

This definition does not apply off the equilibrium path because otherwise we cannot apply the Bayes' rule.

Consistency

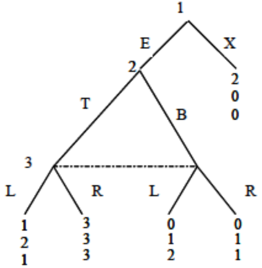

- Can the strategy (X, T, L) be considered not consistent by using our definition of consistency?
- We need to check consistency also off the equilibrium path by "trembling handing".

Perfect Bayesian Nash Equilibrium

Definition

A strategy profile is said to be sequentially rational iff, at each information set, the player who is to move maximizes his expected utility given:

1. his beliefs at the information set;
2. given that the other players play according to the strategy profile in the continuation game.

Definition

A Perfect Bayesian Nash Equilibrium is a pair (σ, μ) of strategy profile and a set of beliefs such that:

1. σ is sequentially rational given beliefs μ;
2. μ is consistent (also off the equilibrium path) with σ.

Perfect Bayesian Nash Equilibrium

Example (13 - PBNE - See notes!)