Epistemic Game Theory

Handout Lecture 1

ESSLLI'12, Opole

Eric Pacuit Olivier Roy

TiLPS, Tilburg University MCMP, LMU Munich

ai.stanford.edu/~epacuit

http://olivier.amonbofis.net

August 6, 2012

Strategic Games

Definition

A game in strategic form $\mathbb{G}$ is a tuple $\langle \mathcal{A}, S_i, u_i \rangle$ such that :

- $\mathcal{A}$ is a finite set of agents.
- S_i is a finite set of *actions* or *strategies* for i. A *strategy profile* $\sigma \in \Pi_{i \in \mathcal{A}} S_i$ is a vector of strategies, one for each agent in I. The strategy s_i which i plays in the profile σ is noted σ_i.
- $u_i : \Pi_{i \in \mathcal{A}} S_i \longrightarrow \mathbb{R}$ is an *utility function* that assigns to every strategy profile $\sigma \in \Pi_{i \in \mathcal{A}} S_i$ the utility valuation of that profile for agent i.

Extensive form games

Definition

A *game in extensive form* $\mathcal{T}$ is a tuple $\langle I, T, \tau, \{u_i\}_{i \in I} \rangle$ such that:

- T is finite set of finite sequences of *actions*, called *histories*, such that:
  - The empty sequence $\emptyset$, the *root* of the tree, is in T.
  - T is prefix-closed: if $(a_1, \ldots, a_n, a_{n+1}) \in T$ then $(a_1, \ldots, a_n) \in T$.
- A history h is *terminal in* T whenever it is the sub-sequence of no other history $h' \in T$. Z denotes the set of terminal histories in T.
- $\tau : (T - Z) \longrightarrow I$ is a *turn function* which assigns to every non-terminal history h the player whose turn it is to play at h.
- $u_i : Z \longrightarrow \mathbb{R}$ is a *payoff function* for player i which assigns i's payoff at each terminal history.

Strategies

Definition

- A *strategy* s_i for agent i is a function that gives, for every history h such that $i = \tau(h)$, an action $a \in A(h)$. S_i is the set of strategies for agent i.
- A *strategy profile* $\sigma \in \Pi_{i \in I} S_i$ is a combination of strategies, one for each agent, and $\sigma(h)$ is a shorthand for the action a such that $a = \sigma_i(h)$ for the agent i whose turn it is at h.
- A history h' is *reachable* or *not excluded* by the profile σ from h if $h' = (h, \sigma(h), \sigma(h, \sigma(h)), ...)$ for some finite number of application of σ.
- We denote $u_i^h(\sigma)$ the value of $util_i$ at the unique terminal history reachable from h by the profile σ.

Nash Equilibrium - General Definition

Definition

A profile of mixed strategy σ is a *Nash equilibrium* iff for all i and all mixed strategy $\sigma_i' \neq \sigma_i$:

$$EU_i(\sigma_i, \sigma_{-i}) \geq EU_i(\sigma_i', \sigma_{-i})$$

Where EU_i, the *expected utility of the strategy* σ_i against σ_{-i} is calculated as follows ($\sigma = (\sigma_i, \sigma_{-i})$):

$$EU_i(\sigma) = \Sigma_{s \in \Pi_j S_j} \left((\Pi_{j \in \mathrm{Ag}} \sigma_j(s_j)) u_i(s) \right)$$

Some Facts about Nash Equilibrium

- Nash equilibria in Pure Strategies do not always exist.
- Every game in strategic form has a Nash equilibrium in mixed strategies.
  - The proof of this make use of Kakutani's Fixed point thm.
- Some games have multiple Nash equilibria.

von Neumann's minimax theorem

For every two-player zero- sum game with finite strategy sets S_1 and S_2, there is a number v, called the **value** of the game such that:

$$\begin{aligned} v \quad &= \max_{p\in\Delta(S_1)} \min_{q\in\Delta(S_2)} u_1(s_1, s_2) \\ &= \min_{q\in\Delta(S_2)} \max_{p\in\Delta(S_1)} u_1(s_1, s_2) \end{aligned}$$

Furthermore, a mixed strategy profile (s_1, s_2) is a Nash equilibrium if and only if

$$s_1 \in \mathrm{argmax}_{p\in\Delta(S_1)} \min_{q\in\Delta(S_2)} u_1(p, q)$$

$$s_2 \in \mathrm{argmax}_{q\in\Delta(S_2)} \min_{p\in\Delta(S_1)} u_1(p, q)$$

Finally, for all mixed Nash equilibria (p, q), $u_1(p, q) = v$

Maximization of Expected Utility

Let $DP = \langle S, O, u, p \rangle$ be a decision problem. S is a finite set of states and O a set of outcomes. An action $a : S \longrightarrow O$ is a function from states to outcomes, u_i a real-valued utility function on O, and p_i a probability measure over S. The **expected utility** of $a \in A$ with respect to p_i is defined as follows:

$$EU_p(a) \; := \; \Sigma_{s \in S} p(s) u(a(s))$$

An action $a \in A$ **maximizes expected utility** with respect to p_i provided for all $a' \in A$, $EU_p(a) \geq EU_p(a')$. In such a case, we also say a is a **best response** to p in game DP.

Some facts about strict dominance

- **Strict dominance is downward monotonic**: If a_i is strictly dominated with respect to $X \subseteq S$ and $X' \subseteq X$, then a_i is strictly dominated with respect to X'.

  - Intuition: the condition of being strictly dominated can be written down in a first-order formula of the form $\forall x\varphi(x)$, where $\varphi(x)$ is quantifier-free. Such formulas are downward monotonic: If $\mathcal{M}, s \models \forall x\varphi(x)$ and $\mathcal{M}' \subseteq \mathcal{M}$ then $\mathcal{M}', s \models \forall x\varphi(x)$

Some facts about strict dominance

- **Relation with MEU**:
  Suppose that $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ is a strategic game. A strategy $s_i \in S_i$ is strictly dominated (possibly by a mixed strategy) with respect to $X \subseteq S_{-i}$ iff there is no probability measure $p \in \Delta(X)$ such that s_i is a best response with respect to p.

Suppose that $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ is a finite strategic game.

Suppose that $s_i \in S_i$ is strictly dominated with respect to X:

$$\exists s_i' \in S_i, \forall s_{-i} \in X, \quad u_i(s_i', s_{-i}) > u_i(s_i, s_{-i})$$

Let $p \in \Delta(X)$ be any probability measure. Then,

$$\forall s_{-i} \in X, \quad p(s_{-i}) \cdot u_i(s_i', s_{-i}) \geq p(s_{-i}) \cdot u_i(s_i, s_{-i})$$

$$\exists s_{-i} \in X, \quad p(s_{-i}) \cdot u_i(s_i', s_{-i}) > p(s_{-i}) \cdot u_i(s_i, s_{-i})$$

Hence,

$$\sum_{s_{-i} \in S_{-i}} p(s_{-i}) \cdot u_i(s_i', s_{-i}) > \sum_{s_{-i} \in S_{-i}} p(s_{-i}) \cdot u_i(s_i, s_{-i})$$

So, $EU(s_i', p) > EU(s_i, p)$: s_i is not a best response to p.

For the converse direction, we sketch the proof for two player games and where $X = S_{-i}$. [1]

Let $G = \langle S_1, S_2, u_1, u_2 \rangle$ be a two-player game.
(Let $U_i : \Delta(S_1) \times \Delta(S_2) \to \mathbb{R}$ be the expected utility for i)

Suppose that $\alpha \in \Delta(S_1)$ is not a best response to any $p \in \Delta(S_2)$.

$$\forall p \in \Delta(S_2) \;\; \exists q \in \Delta(S_1), \quad U_1(q, p) > U_1(\alpha, p)$$

We can define a function $b : \Delta(S_2) \to \Delta(S_1)$ where, for each $p \in \Delta(S_2)$, $U_1(b(p), p) > U_1(\alpha, p)$.

[1] The proof of the more general statement uses the *supporting hyperplane theorem* from convex analysis.

Consider the game $G' = \langle S_1, S_2, \overline{u}_1, \overline{u}_2 \rangle$ where

$$\overline{u}_1(s_1, s_2) = u_1(s_1, s_2) - U_1(\alpha, s_2) \text{ and } \overline{u}_2(s_1, s_2) = -\overline{u}_1(s_1, s_2)$$

By the minimax theorem, there is a Nash equilibrium (p_1^*, p_2^*) such that for all $m \in \Delta(S_2)$,

$$\overline{U}(p_1^*, m) \geq \overline{U}_1(p_1^*, p_2^*) \geq \overline{U}_1(b(p_2^*), p_2^*)$$

We now prove that $\overline{U}_1(b(p_2^*), p_2^*) > \overline{U}_1(\alpha, p_2^*)$:

$$
\begin{aligned}
\overline{U}_1(b(p_2^*), p_2^*) &= \textstyle\sum_{x \in S_1} \sum_{y \in S_2} b(p_2^*)(x) p_2^*(y) \overline{u}_1(x, y) \\
&= \textstyle\sum_{x \in S_1} \sum_{y \in S_2} b(p_2^*)(x) p_2^*(y) [u_1(x, y) - U_1(\alpha, y)] \\
&= \textstyle\sum_{x \in S_1} \sum_{y \in S_2} b(p_2^*)(x) p_2^*(y) u_1(x, y) \\
&- \textstyle\sum_{x \in S_1} \sum_{y \in S_2} b(p_2^*)(x) p_2^*(y) U_1(\alpha, y) \\
&= U_1(b(p_2^*), p_2^*) - \textstyle\sum_{x \in S_1} \sum_{y \in S_2} b(p_2^*)(x) p_2^*(y) U_1(\alpha, \\
&> U_1(\alpha, p_2^*) - \textstyle\sum_{x \in S_1} \sum_{y \in S_2} b(p_2^*)(x) p_2^*(y) U_1(\alpha, y) \\
&> U_1(\alpha, p_2^*)
\end{aligned}
$$

Relation with MEU: Proof

Since $p_2^*(y)U_1(\alpha, y)$ does not depend on x, we have

$$\sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)U_1(\alpha, y) = \sum_{x\in S_1} b(p_2^*)(x)\sum_{y\in S_2} p_2^*(y)U_1(\alpha, y)$$

$$= \sum_{x\in S_1} b(p_2^*)(x)U_1(\alpha, p_2^*) = U_1(\alpha, p_2^*)\sum_{x\in S_1} b(p_2^*)(x) = U_1(\alpha, p_2^*)$$

Hence, for all $m \in \Delta(S_2)$ we have $\overline{U}_1(s_1^*, m) > 0$ which implies for all $m \in \Delta(S_2)$, $U_1(s_1^*, m) > U_1(\alpha, m)$, and so α is strictly dominated by p_1^*.