Game Theory

Georg Moser

Institute of Computer Science @ UIBK

Winter 2009

Summary of Last Lecture

Definition

a two-person zero-sum game Γ in strategic form is a game $\Gamma = (\{1,2\}, C_1, C_2, u_1, u_2)$: $u_1(c_1,c_2) = -u_2(c_1,c_2)$ $\forall c_1 \in C_1$, $\forall c_2 \in C_2$

Example

	C_2	
C_1	M	F
Rr	$0, 0$	$1, -1$
Rp	$0.5, -0.5$	$0, 0$
Pr	$-0.5, 0.5$	$1, -1$
Pp	$0, 0$	$0, 0$

Observation

$$u_1(c_1,c_2) = -u_2(c_1,c_2) \qquad \forall c_1 \in \{Rr, Rp, Pr, Pp\} \quad \forall c_2 \in \{M, F\}$$

Min-Max Theorem

Theorem

(σ_1, σ_2) is an equilibrium of a finite two-person zero-sum game $\Gamma = (\{1,2\}, C_1, C_2, u_1, -u_1)$ if and only if

$$\sigma_1 \in \text{argmax}_{\tau_1 \in \Delta(C_1)} \min_{\tau_2 \in \Delta(C_2)} u_1(\tau_1, \tau_2)$$

$$\sigma_2 \in \text{argmin}_{\tau_2 \in \Delta(C_2)} \max_{\tau_1 \in \Delta(C_1)} u_1(\tau_1, \tau_2)$$

furthermore if (σ_1, σ_2) an equilibrium of Γ, then

$$u_1(\sigma_1, \sigma_2) = \max_{\tau_1 \in \Delta(C_1)} \min_{\tau_2 \in \Delta(C_2)} u_1(\tau_1, \tau_2) = \min_{\tau_2 \in \Delta(C_2)} \max_{\tau_1 \in \Delta(C_1)} u_1(\tau_1, \tau_2)$$

Observation

without randomised strategies, the existence of an equilibrium cannot be guaranteed and the min-max theorem fail

Bayesian Equilibria

consider

$$\Gamma^b = (N, (C_i)_{i \in N}, (T_i)_{i \in N}, (p_i)_{i \in N}, (u_i)_{i \in N})$$

such that

- T_i is the set of types of player i; $T = \prod_{i \in N} T_i$
- $p_i(\cdot|t_i) \in \Delta(T_{-i})$ is the probability distribution over the types of the other players T_{-i}
- for each i: $u_i \colon C \times T \to \mathbb{R}$ is the expected utility payoff

Definition

- strategy for player i is a function $f \colon T \to C$
- randomised strategy profile $\sigma \in \prod_{i \in N} \prod_{t_i \in T_i} \Delta(C_i)$

Definition

Bayesian equilibrium

$$\sigma_i(\cdot|t_i) \in \text{argmax}_{\tau_i \in \Delta(C_i)} \sum_{t_{-i} \in T_{-i}} p_i(t_{-i}|t_i) \sum_{c \in C} \Big(\prod_{j \in N \setminus \{i\}} \sigma_j(c_j|t_j) \Big) \tau_i(c_i) u_i(c, t)$$

Content

motivation, introduction to decision theory, decision theory

basic model of game theory, dominated strategies, common knowledge, Bayesian games, incomplete information, Nash equilibrium

two-person zero-sum games, Bayesian equilibria, sequential equilibria of extensive-form games, subgame-perfect equilibria

(efficient) computation of Nash equilibria, complexity class PPAD, complexity of Nash equilibria, refinements of equilibrium in strategic form, persistent equilibria, games with communication, sender-receiver games

Purification of Randomised Strategies

Example

complete information

consider the following game

	C_2	
C_1	L	R
T	$0, 0$	$0, -1$
B	$1, 0$	$-1, 3$

Observation

- the unique equilibrium is
$$(\frac{3}{4}[T] + \frac{1}{4}[B], \frac{1}{2}[L] + \frac{1}{2}[R])$$
- $([T], [L])$ are pay-off equivalent to equilibrium
- $([T], [L])$ is not an equilibrium

Example

incomplete information

let $\alpha, \beta, \epsilon \in [0,1]$, α is known to player 1 (not player 2), β is known to player 2 (not player 1)

C_1 \ C_2	L	R
T	$\epsilon \cdot \alpha, \epsilon \cdot \beta$	$\epsilon \cdot \alpha, -1$
B	$1, \epsilon \cdot \beta$	$-1, 3$

given ϵ $\exists$ unique Bayesian equilibrium (σ_1, σ_2)

$$\sigma_1(\cdot|\alpha) = \begin{cases} [T] & \alpha > \frac{2+\epsilon}{8+\epsilon^2} \\ [B] & \alpha < \frac{2+\epsilon}{8+\epsilon^2} \end{cases} \qquad \sigma_2(\cdot|\beta) = \begin{cases} [L] & \beta > \frac{4-\epsilon}{8+\epsilon^2} \\ [B] & \beta < \frac{4-\epsilon}{8+\epsilon^2} \end{cases}$$

Observation

if $\epsilon \longrightarrow 0$, the Bayesian equilibrium (σ_1, σ_2) becomes the unique equilibrium in the game with complete information

Auctions

Example

consider the following Bayesian game

- there are n bidders in an auction for a single object
- each player submits a sealed bid b_i
- each player know the value of the object to him
- the highest bid wins
- let $b = (b_1, \dots, b_n)$ the profile of bids; $v = (v_1, \dots, v_n)$ the profile of values
- expected payoff for player i

$$u_i(b, v) = \begin{cases} v_i - b_i & \{i\} = \operatorname{argmax}_{j \in [1,n]} b_j \\ 0 & \text{otherwise} \end{cases}$$

Definition

- let F be an increasing and differentiable function
- $\forall$ players, $F(w)$ denotes the probability that player values the object with less than w

Equilibrium Analysis

- let M be the maximal value and set $\beta\colon [1, n] \to [0, M]$; β is the bidding function, assumed to be increasing and differentiable
- in the equilibrium, player i expects the other players to bid in the interval $[0, \beta(M)]$
- suppose player i's value is v_i, but submits bid $\beta(w_i)$
- player j will submit bid $< \beta(w_i)$ if $\beta(v_j) < \beta(w_i)$
- hence $v_j < w_i$ and probability that $\beta(w_i)$ wins is $F(w_i)^{n-1}$

Lemma

expected payoff to player i bidding $\beta(w_i)$ is

$$(v_i - \beta(w_i)) \cdot F(w_i)^{n-1}$$

for value v_i, bid ought to be $\beta(v_i)$, hence

$$0 = (v_i - \beta(v_i))[F(v_i)]'(n-1)F(v_i)^{n-2} - \beta'(v_i)F(v_i)^{n-1}$$

Lemma

let β, F as above, then

$$\beta(x)F(x)^{n-1} = \int_0^x y(n-1)F(y)^{n-2}F'(y)dy$$

Lemma

assume types/bids are uniformly distributed ($F(y) = \frac{y}{M}$):

$$\beta(v_i) = (1 - \frac{1}{n})v_i \qquad \forall v_i \in [0, M]$$

Definition

- an auction where the private values are independent is called independent private values

  the auction studied is of this form
- if the value of the object is the same for all bidders, but the bidders have different private information, the auction is an common value auction

Example

consider a two-bidder auction with a single object with unknown common value

- x_0, x_1, x_2 independent random variables
- value of object for highest bidder

$$A_0 x_0 + A_1 x_1 + A_2 x_2 \qquad A_i \text{ are nonnegative constants}$$

- A_i is publicly known; player 1 knows x_0, x_1, player 2 know x_0, x_2
- bids c_1, c_2, if the bids tie a coin toss decides
- utility payoff function for player i (the other player is denoted as j)

$$u_i(c_1, c_2, (x_0, x_1), (x_0, x_2)) = \begin{cases} A_0 x_0 + A_1 x_1 + A_2 x_2 - c_i & c_i > c_j \\ \frac{1}{2}(A_0 x_0 + A_1 x_1 + A_2 x_2 - c_i) & c_i = c_j \\ 0 & \text{otherwise} \end{cases}$$

the unique Bayesian equilibrium is (for player 1, 2 respectively)

$$A_0 x_0 + 0.5(A_1 + A_2) x_1 \qquad A_0 x_0 + 0.5(A_1 + A_2) x_2$$

Infinite Strategy Sets

we extend the set of strategies to comprise the real interval $[0, 1]$

Definition

- a metric space is a set M together with the metric $\delta \colon M \times M \to \mathbb{R}$ such that

$$\delta(x, y) = \delta(y, x) \geqslant 0$$
$$\delta(x, y) = 0 \qquad \text{if } x = y$$
$$\delta(x, y) + \delta(y, z) \geqslant \delta(x, z)$$

- $B(x, \epsilon)$ denotes the open sphere with distance ϵ
- a set is an open subset of M if for every $x \in S$ $\exists$ $B(x, \epsilon) \subseteq S$
- a set $N \subseteq M$ is closed if $M \setminus N$ is open
- a metric space is compact if every collection of open sets that covers M has a finite sub-collection that covers M

Definition

let $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$ be a game such that

- C_i is a compact metric space

Borel Subsets

probability distributions need to be defined for subsets of C_i; for that we consider measurable sets

> *unfortunately, for technical reasons, it may be mathematically impossible to consistently assign probabilities to all subsets [...]*

Definition

- the class of measurable subsets of C_i is the smallest class of subsets that include open subsets, closed subsets and all finite or countable infinite unions and intersections of sets in the class
- $\sigma_i \in \Delta(C_i)$ if σ_i is a function that assigns a non-negative number $\sigma_i(Q)$ to each measurable subsets $Q \subseteq C_i$
- $\sigma_i(C_i) = 1$ and

$$\sigma_i(\bigcup_{k \geqslant 1} Q_k = \sum_{k \geqslant 1} \sigma_i(Q_k)$$

- $\exists$ metric (the Prohorov metric) such that $\Delta(C_i)$ is compact metric space

Existence of Equilibrium

Definition

- a function $g\colon C \to R$ is measurable if $\forall\ x \in \mathbb{R}$ the following is measurable

$$\{c \in C \mid g(c) \geqslant x\}$$

- a function g is bounded if $\exists\ K$ such that $|g(c)| \leqslant K$
- utility functions are bounded and measurable

Definition

utility function

let $\sigma \in \prod_{i \in N} \Delta(C_i)$ be a randomised strategy profile in, then

$$u_i(\sigma) = \int_{c_n \in C_n} \cdots \int_{c_1 \in C_1} u_i(c) d\sigma_1(c_1) \dots d\sigma_n(c_n)$$

Theorem

Nash's theorem of the existence of an equilibrium is extensible to games over infinite strategy sets