Game Theory

Georg Moser

Institute of Computer Science @ UIBK

Winter 2009

Summary of Last Last Lecture

Definition (Bayesian games)

a Bayesian game is a tuple $\Gamma^b = (N, (C_i)_{i \in N}, (T_i)_{i \in N}, (p_i)_{i \in N}, (u_i)_{i \in N})$ such that

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

Definition

a strategy for player i in Γ^b is a function $f \colon T \to C$

Definition

let $\sigma \in \prod_{i \in N} \Delta(C_i)$, let $u_i(\sigma)$ denote the expected utility payoff for player i, when players choose strategies according to σ:

$$u_i(\sigma) = \sum_{c \in C} \Big(\prod_{j \in N} \sigma_j(c_j) \Big) u_i(c) \quad \text{for all } i \in N$$

for $\tau_i \in \Delta(C_i)$, let (σ_{-i}, τ_i) denote the randomised strategy profile, where τ_i is substituted for σ_i, thus

$$u_i(\sigma_{-i}, \tau_i) = \sum_{c \in C} \Big(\prod_{j \in N \setminus \{i\}} \sigma_j(c_j) \Big) \tau_i(c_i) u_i(c)$$

Definition (Nash equilibrium)

a randomised strategy profile σ is a Nash equilibrium of Γ if the following holds for all $i \in N$, and every $\tau_i \in \Delta(C_i)$

$$u_i(\sigma) \geqslant u_i(\sigma_{-i}, \tau_i)$$

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 equilibra of extensive-form games, subgame-perfect equilibra

(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

Existence of Nash Equilibrium

Theorem

Nash 1951

given a finite game Γ in strategic form, there exists at least one (Nash) equilibrium in $\prod_{i \in N} \Delta(C_i)$

Example

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

then no pure equilibrium exists, and we can only eliminated Pp

Fact

randomised strategies are needed for this theorem

Definition

the outcome of a game in Pareto efficient if there is no other outcome that would make all players better of

a game may have equilibria that are inefficient, and a game may have multiple equilibria

Example

prisoner dilemma

C_1	C_2: g_2	f_2
g_1	5,5	0,6
f_1	6,0	1,1

- the only equilibrium is $([f_1], [f_2])$ which is inefficient

Example

battle of the sexes

C_1	C_2: f_2	s_2
f_1	3,1	0,0
s_1	0,0	1,3

- the game as two pure equilibria

$$([f_1], [f_2]) \qquad ([s_1], [s_2])$$

- and one (inefficient) mixed equilibria

$$(0.75[f_1] + 0.25[s_1], 0.25[f_2] + 0.75[s_2])$$

The Focal-Point Effect

Definition

focal-point effect

anything that tends to focus the players' attention on one equilibrium may make them all expect it and hence fulfil it; this is called focal-point effect

Example

battle of the sexes with communication

C_1	C_2: $f_2 f_2$	$f_2 s_2$	$s_2 f_2$	$s_2 s_2$
Ff_1	3,1	3,1	0,0	0,0
Fs_1	0,0	0,0	1,3	1,3
Sf_1	3,1	0,0	3,1	0,0
Ss_1	0,0	1,3	0,0	1,3

Definition

if a game can be influence by preplay communication, the player whose words are headed is called focal arbitrator

Example

battle of the sexes (2)

C_1 \ C_2	f_2	s_2
f_1	3,1	0,0
s_1	0,0	**1,3**

Example

battle of the sexes (3)

C_1 \ C_2	f_2	s_2
f_1	3,1	0,0
s_1	0,0	1,3

- assumption: the man is Dr. Taub and he has recently confessed his adultery

Example

divide the dollar

- there are two players
- both can make demands for sum $[1, 100]$ in €i.e.,

$$C_1 = C_2 = \{x \in \mathbb{R} \mid 0 \leqslant x \leqslant 100\}$$

- the payoff function is defined as follows:

$$u_i(c_1, c_2) = \begin{cases} 0 & \text{if } c_1 + c_2 > 100 \\ c_i & \text{if } c_1 + c_2 \leqslant 100 \end{cases}$$

Analysis

- any pair $(x, 100 - x)$ is an equilibrium, on the other hand also the pair $(100, 100)$ is an equilibrium
- an impartial moderator may suggest $(50, 50)$ as it is efficient
- moreover $(50, 50)$ has strong incentive, it is a focal equilibrium even if the moderator is absent

Evolution

Idea

Axelrod 1984

identify good strategies by a biological evolutionary criterion

Definition

- $L_i \subseteq \Delta(C_i)$ of promising randomised strategies
- $\forall$ player i
- $\exists$ i-animals that implement a strategy $\sigma_i \in L_i$
- each i-animal plays the game repeatedly using σ_i
- $\forall$ player $j \neq i$
- let the j-animals randomly choose among the strategies in L_j
- define

$$q_j^k(\sigma_j) = \frac{j\text{-animals that implement } \sigma_j}{\text{all } j\text{-animals}} \qquad (\text{in generation } k)$$

Definition

- define

$$\overline{\sigma}_{-i}^k(c_j) = \sum_{\sigma_j \in L_j} q_j^k(\sigma_j)\sigma_j(c_j) \qquad \forall j \in N,\ \forall c_j \in C_j$$

- set $\overline{\sigma}^k = (\overline{\sigma}_j^k)_{j \in N}$
- and $\overline{u}_i^k(\sigma_i) = u_i(\overline{\sigma}_{-i}^k, \sigma_i)$

Definition

the number of children in the next generation $k+1$ depends on the expected payoff:

$$q_i^{k+1}(\sigma_i) = \frac{q_i^k(\sigma_i)\overline{u}_i^k(\sigma_i)}{\sum_{\tau_i \in L_i} q_i^k(\sigma_i)\overline{u}_i^k(\tau_i)}$$

"Definition"

strategies that survive in the end, are good

strategies that behave poorly can be crucial to determine which strategy reproduces best

Risk Dominance

Idea ② Harsanyi, Selten 1988

overcome this dependency on poor strategies, using risk dominance of strategies

Definition

- $\forall$ games Γ is strategic form
- $\forall$ σ, τ equilibria in $\prod_{i \in N} \Delta(C_i)$ the resistance of σ against τ is the largest $\lambda \in [0,1]$ such that $\forall j \in N$:

$$u_i((\lambda\tau_j + (1-\lambda)\sigma_j)_{j \in N-\{i\}}, \sigma_i) \geqslant u_i((\lambda\tau_j + (1-\lambda)\sigma_j)_{j \in N-\{i\}}, \tau_i)$$

- an equilibrium σ risk dominates another equilibrium τ if the resistance of σ against τ is greater than the resistance of τ against σ

Note

the resistance measure the "evolutionary" strength of an equilibrium

Two-Person Zero-Sum Games

Example

C_1	C_2: 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

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

Definition

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

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

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

Proof

easy ■

Observation

withouth randomised strategies, the existence of an equilibrium cannot be guranteed and the min-max theorem fail

Example

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

- allow only the pure strategies
- we obtain

$$\max_{c_1 \in \{Rr,Rp,Pr,Pp\}} \min_{c_2 \in \{M,F\}} u_1(c_1,c_2) = \max\{0,0,-0.5,0\} = 0$$

$$\min_{c_2 \in \{M,F\}} \max_{c_1 \in \{Rr,Rp,Pr,Pp\}} u_1(c_1,c_2) = \min\{0.5,1\} = 0.5 \neq 0$$

- Γ doesn't admit a pure equilibrium

Example (cont'd)

- proof of the theorem uses the existence of a Nash equilibrium, this is essential
- we need this for

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

Definition

an optimisation problem is defined as

$$\text{minimise}_{x \in \mathbb{R}^n} f(x) \qquad \text{subject to } g_i(x) \geqslant 0 \quad \forall i \in \{1, \dots, m\}$$

where f, g_1, ..., g_m are functions from $\mathbb{R}^n \to \mathbb{R}$

Observation

two-person zero-sum games and optimisation problems are closely linked

Lemma

the optimisation problem

$$\text{minimise}_{x \in \mathbb{R}^n} f(x) \qquad \text{subject to } g_i(x) \geqslant 0 \quad \forall i \in \{1, \dots, m\}$$

is equivalent to

$$\text{minimise}_{x \in \mathbb{R}^n} (\max_{y \in \mathbb{R}^m_+} f(x) - \sum_{i=1}^{m} y_i g_i(x)) \tag{1}$$

here $R^m_+ = \{(y_1, \dots, y_m) \mid y_i \geqslant 0\}$

Proof

observe that $\max_{y \in \mathbb{R}^m_+} (f(x) - \sum_{i=1}^m y_i g_i(x)) = f(x)$ if the constraints are met, otherwise it is $+\infty$ ■

Definition

the dual of (1) is defined as

$$\text{maximise}_{y \in \mathbb{R}^m_+} (\min_{x \in \mathbb{R}^n} f(x) - \sum_{i=1}^{m} y_i g_i(x))$$

Bayesian Equlibria

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

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