Game Theory and Planning (PhD seminar)

Georg Moser & Radu Prodan

Institute of Computer Science @ UIBK

Winter 2010

Summary First Lecture

Definition

- games in extensive form
- games in strategic form
- fully reduced normal form

Example

consider the following game

	C_2		
C_1	x_2	y_2	z_2
x_1	3,0	0,2	0,3
y_1	2,0	1,1	2,0
z_1	0,3	0,2	3,0

the unique Nash equilibrium is (y_1, y_2) as it is the best-response to all other strategies

Summary Last Lecture

Topics

- Learning, regret minimisation and equilibria
- Computation of market equilibria by convex programming
- Graphical games
- Mechanism design
- Combinatorial auctions
- Routing games
- Load balancing or job allocation schemes
- Price of anarchy and the design of scalable resource allocation mechanisms
- Cascading behaviour in networks: algorithmic and economic issues
- Sponsored search auctions

Randomised (or Mixed) Strategies

let Z be a finite set, the probability distributions $\Delta(Z)$ over Z are defined as follows:

$$\Delta(Z) = \{q \colon Z \to \mathbb{R} \mid \sum_{y \in Z} q(y) = 1 \text{ and } \forall z \in Z \; q(z) \geqslant 0\}$$

Definition

let $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$

- a randomised strategy for player i, is a probability distribution $\Delta(C_i)$ over C_i
- $c_i \in C_i$ is a pure strategy
- a randomised strategy profile $\sigma \in \prod_{i \in N} \Delta(C_i)$ specifies a randomised strategy for every player

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:

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

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

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

define $[c_i] \in \Delta(C_i)$ such that

$$[c_i](x) = \begin{cases} 1 & x = c_i \\ 0 & \text{otherwise} \end{cases}$$

Nash Equilibrium

$\forall$ Z and $\forall$ $f : Z \to \mathbb{R}$, define

$$\text{argmax}_{y \in Z} f(y) = \{y \in Z \mid f(y) = \max_{z \in Z} f(z)\}$$

Definition (informal)

- a best response of player i to a randomised strategy profile σ is a randomised strategy τ_i that maximises the expected utility $u(\sigma_{-i}, \tau_i)$ of player i
- a (mixed) Nash equilibrium is a strategy profile σ such that all mixed strategies are best responses to each other

Definition

a randomised strategy profile σ is a (mixed) Nash equilibrium of Γ if $\forall$ $i \in N$, and $\forall$ $c_i \in C_i$

$$\text{if } \sigma_i(c_i) > 0, \text{ then } c_i \in \text{argmax}_{d_i \in C_i} u_i(\sigma_{-i}, [d_i])$$

Lemma

- *$\forall\ \sigma \in \prod_{i\in N} \Delta(C_i)$ and $\forall$ player i*

$$\max_{c_i\in C_i} u_i(\sigma_{-i}, [c_i]) = \max_{\tau_i\in\Delta(C_i)} u_i(\sigma_{-i}, \tau_i)$$

- *furthermore, $p_i \in \mathrm{argmax}_{\tau_i\in\Delta(C_i)}\, u_i(\sigma_{-i}, \tau_i)$ if and only if $p_i(c_i) = 0$ for every $c_i \notin \mathrm{argmax}_{c_i\in C_i}\, u_i(\sigma_{-i}, c_i)$*

the highest expected utility player i can get is independent of the fact whether player i used randomised strategies for herself

Definition

a pure strategy profile $c \in C$ is a pure Nash equilibrium if for all $i \in N$, and every $d_i \in C_i$

$$u_i(c) \geqslant u_i(c_{-i}, d_i)$$

Existence of Nash Equilibrium

Theorem

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_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	$0, 0$	$0, 0$

no pure equilibrium exists

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 Dillema/Routing problem

C_1	C_2: f_2	g_2
f_1	3,3	0,4
g_1	4,0	1,1

the only equilibrium is $([g_1],[g_2])$ 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 battle of sexes is an example of a coordination game
- similar phenomena occur in routing games, which can be conceived as anti-coordination game

Two-Person Zero-Sum Games

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

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

Definition

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

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

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

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

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

Example

C_1	C_2: M	C_2: 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

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

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

Observation

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

Example

C_1	C_2: x_2	C_2: y_2
x_1	(3,3)	(3,2)
y_1	(2,2)	(5,6)
z_1	(0,3)	(6,1)

Γ is representable as two matrices A, B

$$A = \begin{pmatrix} 3 & 3 \\ 2 & 5 \\ 0 & 6 \end{pmatrix} \qquad B = \begin{pmatrix} 3 & 2 \\ 2 & 6 \\ 3 & 1 \end{pmatrix}$$

Notation

- M denotes the set of m pure strategies of player 1
- N denotes the set of n pure strategies of player 2

$$M = \{1, \ldots, m\} \qquad N = \{m + 1, \ldots, m + n\}$$

- let c be a pure strategy profile and $\sigma \in \prod_{i \in N} \Delta(C_i)$ a randomised strategy profile
- using linear algebra notation, we write:

$$\sigma_i = \sum_{c_i \in C_i} \sigma(c_i)[c_i]$$

- only the vector $x := (\sigma(c_{i1}), \ldots, \sigma(c_{i|C_i|}))$ is important
- we call x a mixed strategy

Lemma

let x, y be be mixed strategies, then x is best response to y iff

$$x_i > 0 \quad \textit{implies} \quad (Ay)_i = u = \max\{(Ay)_k \mid k \in M\} \qquad \forall i \in M$$

Definition

the support of a mixed strategy x is the set

$$\prod_{i \in N} \{c_i \in C_i \mid x_i > 0\}$$

Example

in the battle of sexes

1. the support of $([f_1],[f_2])$ is $\{f_1\} \times \{f_2\}$ and the support of $([s_1],[s_2])$ is $\{s_1\} \times \{s_2\}$
2. the support of $(0.75[f_1] + 0.25[s_1], 0.25[f_2] + 0.75[s_2])$ is $\{f_1, s_1\} \times \{f_2, s_2\}$

Definition

a (two-player) game is non-degenerate if no mixed strategy of support size k has more than k pure best responses.

Theorem

for a Nash equilibrium (x, y) of a non-degenerated bimatrix game, x and y have support of equal size

Equilibria by Support Enumeration

Algorithm

- INPUT: a non-degenerate bimatrix game
- OUTPUT: all Nash equilibria

Method

1. $\forall\ k \in \{1, \ldots, \min\{m, n\}\}$
2. $\forall$ k-sized subsets (I, J) of M, N
3. solve the following equation

$$\sum_{i \in I} x_i b_{ij} = v \quad \text{for } j \in J \qquad \sum_{j \in J} a_{ij} y_j = u \quad \text{for } i \in I$$

$$\sum_{i \in I} x_i = 1 \qquad \sum_{j \in J} y_j = 1$$

such that $x \geqslant \mathbf{0}$, $y \geqslant \mathbf{0}$ and the best response condition is fulfilled for x and y

Games with Incomplete Information

- a game has incomplete information if some players have private information before the game starts
- the initial private information is called the type of the player

Definition

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_{-1}
6. for each i: $u_i : C \times T \to \mathbb{R}$ is the expected utility payoff

Strategies in Bayesian Games

Definition

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

Example

consider bargaining game: player 1 is the seller, player two is the buyer

- each player knows the value of the object to himself; assumes the value to the other is $\in [1, 100]$ with uniform probability
- each player bids a number $\in [0, 100]$
- assume utility $=$ monetary profit

any Bayesian game is representable as strategic game by conceiving each type as a player

Applications of Bayesian games: Auctions

auctions are not really a new idea

- used by the Babylonians (500 BC)
- first Roman fire brigade offered to buy the burning house and only extinguished the fire if the offer was accepted
- after having killed Emperor Pertinax, the Prätorian Guard auctioned off the Roman Empire (193)
- Johann Wolfgang von Goethe sold a manuscript through a second-price auction (1797)
- biggest revenue yet was generated by the US FCC spectrum auctions (1994–2008)

Observation

- game theoretic analysis of first price auctions shows non-efficiency of this auction
- mechanism design aims at the design of auctions where Bayesian-Nash eq. is Pareto efficient

Assignment of Topics

Topics

1. Learning, regret minimisation and equilibria
2. Computation of market equilibria by convex programming
3. Graphical games
4. Mechanism design
5. Combinatorial auctions
6. Routing games
7. Load balancing or job allocation schemes
8. Price of anarchy and the design of scalable resource allocation mechanisms
9. Cascading behaviour in networks: algorithmic and economic issues
10. Sponsored search auctions