CSCI7000-016: Optimization and Control of Networks

Static Games and Classical Mechanism Design

Lijun Chen

01/12/2016

Outline

- Strategic games and their solution concepts
  - Strategic form games and dominated strategies
  - Nash equilibrium and correlated equilibrium
- Classical mechanism design
  - Incomplete information games
  - Incentive-compatible mechanism
  - VCG mechanism

Strategic game

- Def: a game in strategic form is a triple

$$G = \{N, S_{i\in N}, u_{i\in N}\}$$

  - N is the set of players (agents)
  - S_i is the player i strategy space
  - $u_i : S \to R$ is the player i payoff function
- Notations
  - $S = S_1 \times S_2 \times \cdots \times S_N$: the set of all profiles of player strategies
  - $s = (s_1, s_2, \cdots, s_N)$: profile of strategies
  - $s_{-i} = (s_1, s_2, \cdots, s_{i-1}, s_{i+1}, \cdots, s_N)$: the profile of strategies other than player i

- Implicitly assume that players have preferences over different outcomes, which can be captured by assigning payoffs to the outcomes
- The basic model of **rationality** is that of a payoff maximizer
- First consider pure strategy, will consider mixed strategy later

Example: finite game

		column L	column M	column R
row	U	4,3	5,1	6,2
row	M	2,1	8,4	3,6
row	D	3,0	9,6	2,8

Example: Continuous strategy game

- Cournot competition
  - Two players: firm 1 and firm2
  - Strategy $s_i \in [0, \infty]$: the amount of widget that firm i produces
  - The payoff for each firm is the net revenue

$$u_i(s_1, s_2) = s_i p(s_1 + s_2) - c_i s_i$$

where p is the price, c_i is the unit cost for firm i

Dominated strategies

- How to predict the outcome of a game?
- Prisoner's Dilemma

	D	C
D	-2,-2	-5,-1
C	-1,-5	-4,-4

- Two prisoners will play (C,C)
- Def: a strategy s_i is (weakly) dominated for player i if there exists $s_i' \in S_i$ such that

$$u_i(s_i', s_{-i}) \geq u_i(s_i, s_{-i}) \text{ for all } s_{-i} \in S_{-i}$$

Iterated elimination of dominated strategies

- Iterated elimination of dominated strategies

	L	~~M~~	~~R~~
U	4,3	~~5,1~~	~~6,2~~
~~M~~	~~2,1~~	~~8,4~~	~~3,6~~
~~D~~	~~3,0~~	~~9,6~~	~~2,8~~

- However, most of games are not solvable by iterated elimination of dominated strategies

Nash equilibrium

- Def: a strategy profile s^* is a Nash equilibrium, if for all i,

$$u_i(s^*{}_i, s^*{}_{-i}) \geq u_i(s_i, s^*{}_{-i}) \text{ for all } s_i \in S_i$$

- For any $s_{-i} \in S_{-i}$, define best response function

$$B_i(s_{-i}) = \{s_i \in S_i \mid u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \quad \forall s_i' \in S_i\}.$$

Then a strategy profile s^* is a Nash equilibrium iff $s^*{}_i \in B_i(s^*{}_{-i})$.

Examples

- Battle of the Sexes

	Ballet	Soccer
Ballet	2,1	0,0
Soccer	0,0	1,2

Two Nash equilibria (Ballet, Ballet) and (Soccer, Soccer)

Cournot Competition

- Suppose a price function $p(s_1 + s_2) = \max\{0, 1 - (s_1 + s_2)\}$
- Suppose cost $0 \le c_1 = c_2 = c \le 1$
- Then, the best response function

$$B_1(s_2) = (1 - s_2 - c)/2$$
$$B_2(s_1) = (1 - s_1 - c)/2$$

- Nash equilibrium satisfies $\begin{array}{l} s_1 = B_1(s_2) \\ s_2 = B_2(s_1) \end{array}$, i.e., $\begin{array}{l} s_1 = (1 - c)/3 \\ s_2 = (1 - c)/3 \end{array}$

Second price auction

- An object to be sold to a player in N
- Each player i has a valuation v_i of the object. We further assume $v_1 > v_2 > \cdots > v_N > 0$
- The players simultaneously submit bids $b_1, \cdots, b_N$
- The object is given to the player with highest bid. The winner pays the second highest bid.
- The payoff of the winner is his valuation of the object minus the price he pays. All other players' payoff is zero.

- $(b_1,\cdots,b_N)=(v_1,\cdots,v_N)$ is Nash equilibrium
  - Player 1 receives the object and pay v_2 , and has payoff $v_1 - v_2 > 0$. Player 1 has no incentive to deviate, since his payoff can only decrease
  - For other players, the payoff is zero. In order to change his payoff, he needs to bid more than v_1, but that will result in negative payoff. So, no player has incentive to change
- Question: are they more Nash equilibria?

- Not all games have (pure) Nash equilibrium
- Matching Pennies

	Heads	Tails
Heads	1,-1	-1,1
Tails	-1,1	1,-1

Mixed strategies

- Let Σ_i denote the set of probability distribution over player i strategy space S_i
- A mixed strategy $\sigma_i \in \Sigma_i$ is a probability mass function over pure strategies $s_i \in S_i$
- The payoff of a mixed strategy is the expected value of the pure strategy profiles

$$u_i = \sum_{s \in S} \left(\prod_{j \in N} \sigma_j(s_j) \right) u_i(s)$$

Mixed strategy Nash equilibrium

- Def: a mixed strategy profile σ^* is a (mixed strategy) Nash equilibrium if for all i

$$u_i(\sigma^*{}_i, \sigma^*{}_{-i}) \geq u_i(\sigma_i, \sigma^*{}_{-i}) \text{ for all } \sigma_i \in \Sigma_i$$

- A mixed strategy profile σ^* is a (mixed strategy) Nash equilibrium if for all i

$$u_i(\sigma^*{}_i, \sigma^*{}_{-i}) \geq u_i(s_i, \sigma^*{}_{-i}) \text{ for all } s_i \in S_i$$

  - The payoff $u_i(s_i, \sigma^*{}_{-i})$ is the same for all $s_i \in \text{supp}(\sigma^*{}_i)$
  - The payoff $u_i(s_i, \sigma^*{}_{-i})$ for each $s_i \notin \text{supp}(\sigma^*{}_i)$ is not larger

Example

	Ballet	Soccer
Ballet	2,1	0,0
Soccer	0,0	1,2

- Assume row (column) player choose “ballet” with probability p (q) and “soccer” with probability $1-p$ ($1-q$)
- Mixed strategy Nash equilibrium is $\begin{cases} p = 2/3 \\ q = 1/3 \end{cases}$

$$2 \times q + 0 \times (1-q) = 0 \times q + 1 \times (1-q)$$
$$1 \times p + 0 \times (1-p) = 0 \times p + 2 \times (1-p)$$

Existence of Nash equilibrium

- Theorem (Nash '50): Every finite strategic game has a mixed strategy Nash equilibrium.
- Example: Matching Pennies game has a mixed strategy Nash equilibrium (½, ½; ½, ½)

	Heads	Tails
Heads	1,-1	-1,1
Tails	-1,1	1,-1

- Proof: using Kakutani's fixed point theorem. See section 1.3.1 of the book by Fudenburg & Tirole

Continuous strategy game

- Theorem (Debreu '52; Glicksberg '52; Fan '52): Consider a strategic game $\{N, S_{i\in N}, u_{i\in N}\}$ with continuous strategy space. A pure strategy Nash equilibrium exists if
  - S_i is nonempty compact convex set
  - u_i is continuous in S and quasi-concave in S_i
- Theorem (Glicksberg '52): Consider a strategic game $\{N, S_{i\in N}, u_{i\in N}\}$ with nonempty compact strategy space. A mixed strategy Nash equilibrium exists if u_i is continuous.

Correlated equilibrium

- In Nash equilibrium, players choose strategies **independently**. How about players observing some common signals?
- Traffic intersection game

	Stop	Go
Stop	2,2	1,3
Go	3,1	0,0

  - Two pure Nash equilibria: (stop, go) and (go, stop)
  - One mixed strategy equilibrium: (½, ½; ½, ½)
  - If there is a traffic signal such that with probability ½ (red light) players play (stop, go) and with probability ½ (green light) players play (go, stop). This is a correlated equilibrium.

- Def: correlated equilibrium is a probability distribution $p(\cdot)$ over the pure strategy space such that for all i

$$\sum_{s_{-i}} p(s_i, s_{-i})[u_i(s_i, s_{-i}) - u_i(t_i, s_{-i})] \ge 0 \text{ for all } s_i, t_i \in S_i$$

- A mixed strategy Nash equilibrium is a correlated equilibrium
- The set of correlated equilibria is convex and contains the convex hull of mixed strategy Nash equilibria

Dynamics in games

- Nash equilibrium is a very strong concept. It assumes player strategies, payoffs and rationality are "common knowledge"
- "Game theory lacks a general and convincing argument that a Nash outcome will occur"
- One justification is that equilibria arise as a result of adaptation (learning)
  - Consider repeated play of the strategic game
  - Players are myopic, and adjust their strategies based on the strategies of other players in previous rounds.

- Best response

$$s_i(t+1) = B_i(s_{-i}(t))$$

- Fictitious play, regret-based heuristics, etc
- Many if not most network algorithms are repeated and adaptive, and achieving some equilibria. Will discuss these and networking games later in this course

Classical mechanism design (MD)

- Mechanism: Protocols to implement an outcome (equilibrium) with desired system-wide properties despite the self-interest and private information of agents
- Mechanism design: the design of such mechanisms
- Provide an introduction to game theoretic approach to mechanism design

Game theoretic approach to MD

- Start with a strategic model of agent behavior
- Design rules of a game, so that when agents play as assumed the outcome with desired properties will happen

induce the desired outcome among self-interested agents

Incomplete information games

- Players have private type $(\theta_1, \theta_2, \cdots, \theta_N) \in \Theta$
- Strategy $s_i(\theta_i) \in S_i$ is a function of a player's type
  - Players of different types may react differently in face of the same situation
- Payoff $u_i(s(\theta), \theta_i) \in R$ is a function of player's type
  - Players of different types may have different preferences over the same strategy profile
- All information except actual types of players is common knowledge
  - If a player's type is known, its payoff is known

Incomplete information games

- Assume types are drawn from some objective distribution $p(\theta_1, \theta_2, \cdots, \theta_N)$
- Definition: a strategy profile s^* is a Bayesian-Nash equilibrium if every player i plays a best response to maximize expected payoff given its belief about distribution $p(\theta_{-i} \mid \theta_i)$, i.e.,

$$s_i^*(\theta_i) \in \arg\max_{s_i} \sum_{\theta_{-i}} p(\theta_{-i} \mid \theta_i) u_i(s_i, s_{-i}^*(\theta_{-i}), \theta_i)$$

Example: Variant of Battle of the Sexes

- Two types: either wants to meet the other or does not
- Assume row player wants to meet column player, but not sure if column player want to meet her or not (assign ½ probability to each case); and column player knows row player's type
- If column player want to meet row player, the payoffs are

	Ballet	Soccer
Ballet	2,1	0,0
Soccer	0,0	1,2

- If column player does not want to meet row player, the payoffs are

	Ballet	Soccer
Ballet	2,0	0,2
Soccer	0,1	1,0

- The Bayesian-Nash equilibrium?

  (Ballet, (Ballet, Soccer))

  - E[Ballet, (Ballet, Soccer)]= $\frac{1}{2}$x2+ $\frac{1}{2}$x0=1
  - E[Soccer, (Ballet, Soccer)]= $\frac{1}{2}$x0+ $\frac{1}{2}$x1= $\frac{1}{2}$

Stronger solution concepts

- Definition: a strategy profile s^* is ex post Nash equilibrium if every player i's strategy is best response whatever the type of others

$$s^*_i(\theta_i) \in \arg\max_{s_i} u_i(s_i, s^*_{-i}(\theta_{-i}), \theta_i) \text{ for all } \theta_{-i}$$

- Definition: a strategy profile s^* is dominant strategy equilibrium if every player i's strategy is best response whatever the type and whatever the strategy of others

$$s^*_i(\theta_i) \in \arg\max_{s_i} u_i(s_i, s_{-i}(\theta_{-i}), \theta_i) \text{ for all } \mathrm{s}_{-\mathrm{i}}, \theta_{-i}$$

Dominant strategy equilibrium

- Very robust solution concept
- Make no assumption about information availability
- Do not require an agent to believe others will behave rationally
- A dominant strategy implementation is much more desirable than Nash equilibrium

Example: second price auction

- The type is player valuation v_i
- Each player submit bid $b_i(v_i)$
- A dominant strategy is to bid $b^*{}_i(v_i) = v_i$
- Players don't need to know valuations (types), or strategies of others

Model of Mechanism Design

- Set of alternative outcomes O
- Player i has private information (type) θ_i
- Type defines a value function $v_i(o;\theta_i) \in R$ for outcome $o \in O$ for each player i
- Player payoff $u_i(o;\theta_i) = v_i(o;\theta_i) - p_i$ for outcome o and payment p_i
- The desired properties are encapsulated in the social choice function $f : \Theta \rightarrow O$
  - e.g., choose o to maximize social welfare, i.e.,

$$f(\theta) = \arg\max_{o \in O} \sum_i u_i(o;\theta)$$

- The goal is to implement social choice function $f(\theta)$

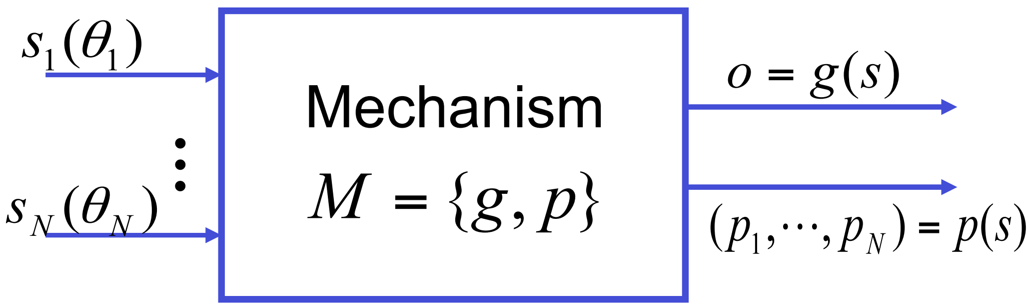

  - A mechanism is defined by an outcome rule $g : S \rightarrow O$ and a payment rule $p : S \rightarrow R^n$
  - A mechanism M implements social choice function $f(\theta)$ if $g(s^*{}_1(\theta_1), \cdots, s^*{}_N(\theta_N)) = f(\theta)$, where the strategy profile $(s^*{}_1, \cdots, s^*{}_N)$ is an equilibrium solution of the game induced by M

Properties of social choice functions and mechanisms

- Pareto optimal:

  if for every $a \neq f(\theta),\ u_i(a,\theta) > u_i(scf(\theta),\theta) \Rightarrow \exists \mathrm{j}\ u_j(a,\theta) < u_j(scf(\theta),\theta)$

- Efficient: if $f(\theta) \in \arg\max_a \sum_i v_i(a,\theta_i)$

- Budget-balance: if $\sum_i p_i(\theta) = 0$

- A mechanism that implements the corresponding social choice functions is called Pareto optimal, efficient, or budget-balanced mechanisms, respectively

Incentive-compatible mechanism

- Revelation principle: any mechanism can be transformed into an incentive compatible, direct-revelation mechanism that implements the same social choice function
- Direct-revelation mechanism is a mechanism in which player strategy space is restricted to their types

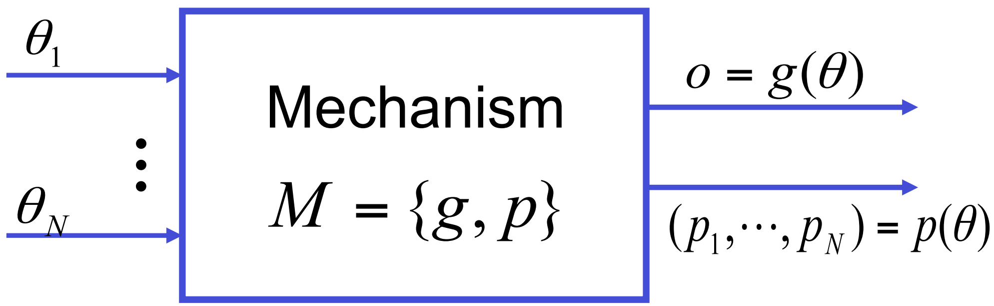

- Incentive-compatible means the equilibrium strategy is to report truthful information about their types (truth-revelation)
  - First price auction is not incentive-compatible. In first price auction, the buyer with highest bid gets the object and pays his bid
  - The second price auction is incentive compatible, direct-revelation mechanism
- Captures the essence of designing a mechanism to overcome the self-interest of agents
  - Report its private information truthfully, out of its own self-interest

Truthful mechanism

- Truthful (aka “strategy-proof”) mechanism: truth-revelation is a dominant strategy equilibrium.
  - Dominant strategy implementation removes game theoretic complexity from mechanism design
  - Very robust to assumption about agent rationality and information about each other
  - An agent can compute its optimal strategy without modeling the types and strategies of others

Vickrey-Clarke-Groves mechanisms

- VCG mechanism:
  - Collect $\theta = (\theta_1, \theta_2, \cdots, \theta_N)$ from agents
  - $g(\theta)$: select an outcome $o^* \in \arg\max_{o \in O} \sum_i v_i(o; \theta_i)$
  - $p(\theta)$: agent i pays $\sum_{j \neq i} v_j(o^{-i}; \theta_j) - \sum_{j \neq i} v_j(o^*; \theta_j)$, where

$$o^{-i} \in \arg\max_{o \in O} \sum_{j \neq i} v_j(o; \theta_j)$$

- Theorem: VCG mechanism is efficient and truthful

- Proof: $u_i(\theta_i, \theta_{-i}) = v_i(o^*; \theta_i) + \sum_{j \neq i} v_j(o^*; \theta_j) - \sum_{j \neq i} v_j(o^{-i}; \theta_j)$

- VCG mechanism is the only mechanism that is efficient and strategy-proof among direct-revelation mechanisms

Combinatorial auction

- Goods P
- Outcomes: allocations $A = (A_1, \cdots, A_N)$, where $A_i \subseteq P$ and are not overlapped
- Agent valuation $v_i(A_i;\theta_i)$ for $A_i \subseteq P$
- Goal: allocate goods to maximize $\sum_i v_i(A_i;\theta_i)$
- Applications: wireless spectrum auction, course scheduling, …

- Two items A and B; 3 agents
- Valuation

	A	B	AB
1	5	0	5
2	0	5	5
3	0	0	12

- Outcome?

  agent 3 wins AB and pays 10-0=10

- Another valuation

	A	B	AB
1	5	0	5
2	0	5	5
3	0	0	7

- Outcome?

  agents 1 and 2 win and each pays 7-5=2

Remarks

- Only consider the incentive issue: to overcome the self-interest of agents
- Not discuss computational and informational issues
  - Tractability (algorithmic MD)
  - Distributed computation (distributed MD)
  - Minimal information revelation
  - Bounded-rational agents
  - ……

Problem features