CSC304 Lecture 7

End of Game Theory

Begin Mechanism Design w/ Money: Intro, Basic Framework

Game Theory Recap

- Normal form games
  - Strictly/weakly dominant strategies
  - Iterated elimination of strictly/weakly dominated strategies
  - Pure/mixed Nash equilibrium
  - Lots of examples
  - Nash's theorem
  - Finding pure NE using best response diagrams
  - Finding mixed NE using the indifference principle

Game Theory Recap

Price of Anarchy (PoA)

- Worst NE vs social optimum

$$\frac{\text{max social welfare}}{\text{min social welfare in NE}}$$

$$\frac{\text{max social cost in NE}}{\text{min social cost}}$$

Price of Stability (PoS)

- Best NE vs social optimum

$$\frac{\text{max social welfare}}{\text{max social welfare in NE}}$$

$$\frac{\text{min social cost in NE}}{\text{min social cost}}$$

$$\text{PoA} \geq \text{PoS} \geq 1$$

Game Theory Recap

- Cost Sharing Games
  - Potential function ⇒ existence of a pure NE
  - PoS = $O(\log n)$, PoA = $\Theta(n)$
- Congestion games
  - Braess' paradox
- Zero-sum games
  - The minimax theorem
- Stackelberg games
  - Security games

Mechanism Design

- A principal who wants the agents to choose certain actions

- Designs the payoff matrix such that rational agents will choose the desired actions

- E.g., in the prisoner’s dilemma, the police setting the payoffs for the four outcomes “betray” and “silent” to get both agents to betray

Mechanism Design

- Formally, a set of outcomes/alternatives A
- Each agent has preferences over A
  - Cardinal values: $v_i : A \to \mathbb{R}$
  - Could also be ranked preferences (later!)
- The principal wants to implement an outcome a^*
  - Social choice theory: "Which outcome is socially good given agent preferences?"
  - Various metrics: efficiency, fairness, stability, revenue, …

Mechanism Design

- For now, we focus on social welfare maximization
  - $a^* \in \text{argmax}_{a \in A} \sum_i v_i(a)$
- But agents want to maximize their own value
  - Might try to feed bad information to the principal
- Key advantage: the principal can charge payments

Mechanism Design

- Focus on direct revelation mechanisms

- Principle declares a pair (f, p)
  - Once all agents report their valuations $v = (v_i)_{i=1}^n$
  - The outcome is $f(v)$
  - The payment vector is $p(v)$: agent i pays $p_i(v)$

- Utility to agent i is quasi-linear
  - $u_i(v) = v_i\big(f(v)\big) - p_i(v)$

Mechanism Design

- Not only that, we want…
  - to choose $f(v) = \text{argmax}_{a \in A} \sum_i v_i(a)$
  - the agents to correctly report their v_i
- You all tell me the truth. I'll compute the social best outcome.
  - Yeah, right.
- How do we get the agents to tell the truth?
  - Use the $p(v)$ correctly!

Mechanism Design

- Dominant strategy incentive compatibility (DSIC)
  - It should be a dominant strategy for the agent to report truthfully
- Bayes-Nash incentive compatibility (BNIC)
  - The agents share a common prior : each v_i is drawn from a distribution ($v_i \sim D_i$)
  - Agent i knows v_i, but takes expectation over other v_j
- The revelation principle (in short)
  - Any outcome that can be achieved as dominant strategy / Bayes-Nash equilibrium can be achieved by a direct revelation mechanism.

Mechanism Design

- Recap
  - We want to maximize social welfare
  - We want to do so using a direct revelation mechanism
  - We want it to be truthful
  - The last two are w.l.o.g. given the revelation principle
- Wait. Why do we want to maximize $\sum_i v_i(a)$?
  - What about payments? We don't really care about them.
  - Alternatively, you can cancel them out if you add the principal/auctioneer as an agent in the system
  - $\left(\sum_i v_i(a) - p_i\right) + \left(\sum_i p_i\right) = \sum_i v_i(a)$