Mechanism Design as an Optimization Problem

Game Theory Course:
Jackson, Leyton-Brown & Shoham

Mechanism Design as an Optimization Problem

- We can understand mechanism design as the problem of finding the best possible mechanism, given constraints about how it operates
- We'll now consider some typical choices for
  - these constraints
  - this notion of "best"

Truthfulness

Definition (Truthfulness)

A transferrable utility mechanism is truthful if it is direct and $\forall i \forall v_i$, agent i's equilibrium strategy is to adopt the strategy $\hat{v}_i = v_i$.

Efficiency

Definition (Efficiency)

A transferrable utility mechanism is strictly Pareto efficient, or just efficient, if in equilibrium it selects a choice x such that

$$\forall v \forall x', \sum_i v_i(x) \geq \sum_i v_i(x').$$

- An efficient mechanism selects the choice that maximizes the sum of agents' utilities, disregarding monetary payments.
- How is this related to the standard game-theoretic definition of Pareto efficiency?

Efficiency

Definition (Efficiency)

A transferrable utility mechanism is strictly Pareto efficient, or just efficient, if in equilibrium it selects a choice x such that

$$\forall v \forall x', \sum_i v_i(x) \geq \sum_i v_i(x').$$

- An efficient mechanism selects the choice that maximizes the sum of agents' utilities, disregarding monetary payments.
- How is this related to the standard game-theoretic definition of Pareto efficiency?
  - if we include the mechanism as an agent, all Pareto-efficient outcomes involve the same choice (and different payments)
  - any outcome involving another choice is Pareto dominated: some agents could pay others such that all would prefer the swap

Efficiency

Definition (Efficiency)

A transferrable utility mechanism is strictly Pareto efficient, or just efficient, if in equilibrium it selects a choice x such that

$$\forall v \forall x', \sum_i v_i(x) \geq \sum_i v_i(x').$$

- Called economic efficiency to distinguish from other (e.g., computational) notions
- Also called social-welfare maximization
- Note: defined in terms of true (not declared) valuations.

Budget Balance

Definition (Budget balance)

A transferrable utility mechanism is budget balanced when

$$\forall v, \sum_i p_i(s(v)) = 0,$$

where s is the equilibrium strategy profile.

- regardless of the agents' types, the mechanism collects and disburses the same amount of money from and to the agents

Budget Balance

Definition (Budget balance)

A transferrable utility mechanism is budget balanced when

$$\forall v, \sum_i p_i(s(v)) = 0,$$

where s is the equilibrium strategy profile.

- regardless of the agents' types, the mechanism collects and disburses the same amount of money from and to the agents
- relaxed version: weak budget balance:

$$\forall v, \sum_i p_i(s(v)) \geq 0$$

  - the mechanism never takes a loss, but it may make a profit

Budget Balance

Definition (Budget balance)

A transferrable utility mechanism is budget balanced when

$$\forall v, \sum_i p_i(s(v)) = 0,$$

where s is the equilibrium strategy profile.

- regardless of the agents' types, the mechanism collects and disburses the same amount of money from and to the agents
- Budget balance can be required to hold *ex ante*:

$$\mathbb{E}_v \sum_i p_i(s(v)) = 0$$

  - the mechanism must break even only on expectation

Individual-Rationality

Definition (*Ex interim* individual rationality)

A mechanism is ex interim individual rational when
$\forall i \forall v_i,\ \mathbb{E}_{v_{-i}|v_i} v_i(\chi(s_i(v_i), s_{-i}(v_{-i}))) - p_i(s_i(v_i), s_{-i}(v_{-i})) \geq 0,$
where s is the equilibrium strategy profile.

- no agent loses by participating in the mechanism.
- *ex interim* because it holds for *every* possible valuation for agent i, but averages over the possible valuations of the other agents.

Individual-Rationality

Definition (*Ex interim* individual rationality)

A mechanism is ex interim individual rational when $\forall i \forall v_i,\ \mathbb{E}_{v_{-i}|v_i} v_i(\chi(s_i(v_i), s_{-i}(v_{-i}))) - p_i(s_i(v_i), s_{-i}(v_{-i})) \geq 0$, where s is the equilibrium strategy profile.

- no agent loses by participating in the mechanism.
- *ex interim* because it holds for *every* possible valuation for agent i, but averages over the possible valuations of the other agents.

Definition (*Ex post* individual rationality)

A mechanism is ex post individual rational when $\forall i \forall v,\ v_i(\chi(s(v))) - p_i(s(v)) \geq 0$, where s is the equilibrium strategy profile.

Tractability

Definition (Tractability)

A mechanism is tractable when $\forall \hat{v}$, $\chi(\hat{v})$ and $p(\hat{v})$ can be computed in polynomial time.

- The mechanism is (guaranteed to be) computationally feasible.

Revenue Maximization

We can also add an objective function to our mechanism. One example: revenue maximization.

Definition (Revenue maximization)

A mechanism is revenue maximizing when, among the set of functions χ and p that satisfy the other constraints, the mechanism selects the χ and p that maximize $\mathbb{E}_\theta \sum_i p_i(s(\theta))$, where $s(\theta)$ denotes the agents' equilibrium strategy profile.

- The mechanism designer can choose among mechanisms that satisfy the desired constraints by adding an objective function such as revenue maximization.

Revenue Minimization

- The mechanism may not be intended to make money.
- Budget balance may be impossible to satisfy.
- Set weak budget balance as a constraint and add the following objective.

Definition (Revenue minimization)

A transferrable utility mechanism is revenue minimizing when, among the set of functions χ and p that satisfy the other constraints, the mechanism selects the χ and p that minimize $\max_v \sum_i p_i(s(v))$ in equilibrium, where $s(v)$ denotes the agents' equilibrium strategy profile.

- Note: this considers the worst case over valuations; we could consider average case instead.

Fairness

- Fairness is hard to define. What is fairer:
  - charging all agents \$100 and making a choice they all hate equally?
  - charging all agents \$0 and making a choice that some hate and some like?

Fairness

- Fairness is hard to define. What is fairer:
  - charging all agents \$100 and making a choice they all hate equally?
  - charging all agents \$0 and making a choice that some hate and some like?
- Maxmin fairness: make the least-happy agent the happiest.

Definition (Maxmin fairness)

A transferrable utility mechanism is maxmin fair when, among the set of functions χ and p that satisfy the other constraints, the mechanism selects the χ and p that maximize

$$\mathbb{E}_v \left[\min_{i \in N} v_i(\chi(s(v))) - p_i(s(v)) \right],$$

where $s(v)$ denotes the agents' equilibrium strategy profile.

Price of Anarchy Minimization

- When efficiency is impossible, we can try to get close
- Minimize the worst-case ratio between optimal social welfare and the social welfare achieved by the given mechanism.

Definition (Price-of-anarchy minimization)

A transferrable utility mechanism minimizes the price of anarchy when, among the set of functions χ and p that satisfy the other constraints, the mechanism selects the χ and p that minimize

$$\max_{v \in V} \frac{\max_{x \in X} \sum_{i \in N} v_i(x)}{\sum_{i \in N} v_i\left(\chi(s(v))\right)},$$

where $s(v)$ denotes the agents' equilibrium strategy profile in the *worst* equilibrium of the mechanism—i.e., the equilibrium in which $\sum_{i \in N} v_i(\chi(s(v)))$ is the smallest.