Mechanism Design

Roberto Lucchetti

Politecnico di Milano

Contents of the week

- What is mechanism design?
- Direct mechanisms and the induced game
- Truthfulness and efficiency
- Example: Second-price auctions
- Vickrey-Clarke-Groves mechanisms
- Examples
  - Single-item auctions
  - Multiple-item auctions
  - Provision of a public good

What is mechanism design?

- In a game the rules are given and the goal is to find an equilibrium.
- Inverse problem: *design the rules of the game* so that equilibria have good properties – *social efficiency*, dominant strategies, low price-of-anarchy. . .
- The main issue is to create incentives to make convenient for the player to tell the truth, in order to achieve a socially efficient outcome.

Example: Auctions

Players in an auction normally bid amounts that differ from their valuation of the object. How can we make sure that the object is assigned efficiently to the player that values it the most?

Which auction rules provide incentives for players to reveal their true valuations?

Example: Provision of a public good

A government is willing to invest in a project only if the sum of social benefits exceeds the investment costs.

How can we find out the true valuations of the users, knowing that they may over-declare to make the project look socially profitable?

Direct mechanisms and the induced game

- A central planner must choose an action $a \in A$ among a finite set of options. Her goal is to maximize the social welfare, that is, the sum $\sum_{i=1}^{n} v_i(a)$ of the utilities $v_i(a)$ of all the individuals $i = 1, \ldots, n$ in the society.
- The valuation function $v_i : A \to \mathbb{R}$ of each agent is private information, and we only know that it belongs to a certain family of functions V_i. A direct mechanism asks all the players to reveal their valuations v_i.

Example: Ann, Bob, Carol and David are discussing which dessert to have for dinner. Their valuations for the different options differ as shown below

	A	B	C	D	Total
Gelato	3	4	2	10	19
Tiramisù	2	5	7	4	18
Torta	5	5	2	2	14

Optimal social choice = Gelato!David is happy... but Carol may want to misreport her valuation for Gelato to be 0 so that the choice changes to Tiramisù. But then Ann and David want to lie...

Direct mechanisms and the induced game

- Individuals may try to manipulate the planner's decision by misreporting their true valuations v_i and declare whichever $s_i \in V_i$ leads to an action a that gives them a higher utility $v_i(a)$. This generates a game!

Goal: neutralize the ability of players to manipulate the decision.

How: introduce side payments that eliminate the benefits of declaring $s_i \neq v_i$ and generate incentives to declare truthfully $s_i = v_i$.

- Let $V = V_1 \times \cdots \times V_n$ be the set of player's valuations. A direct mechanism is defined by a pair of maps (f, p) where $f : V \to A$ is an action rule, and $p : V \to \mathbb{R}^n$ determines the side payments for the players.
- For a profile of declared valuations $(s_1, \ldots, s_n) \in V$, the utility for player i is

$$u_i(s_1, \ldots, s_n) = v_i(f(s_1, \ldots, s_n)) - p_i(s_1, \ldots, s_n).$$

This induces a game where the strategy set of player i is V_i.

Truthfulness and efficiency

- For a profile of declared valuations $(s_1, \ldots, s_n) \in V$, the utility for player i is

$$u_i(s_1, \ldots, s_n) = v_i(f(s_1, \ldots, s_n)) - p_i(s_1, \ldots, s_n).$$

- A direct mechanism (f, p) is said to be truthful or incentive compatible if declaring $s_i = v_i$ is a weakly dominant strategy, that is

$$v_i(f(v_i, s_{-i})) - p_i(v_i, s_{-i}) \geq v_i(f(s)) - p_i(s) \qquad \forall\, s \in S.$$

We look for truthful mechanisms (f, p) whose equilibria lead to an efficient action $a = f(s)$ that maximizes the true social welfare $\sum_{i=1}^{n} v_i(a)$.

Example: Second-price auctions

In an auction the set of actions $A = \{1, \dots, n\}$ corresponds to the possible assignments of the object to the different players. The valuation of player i is, given her valuation w_i of the object,

$$v_i(a) = \begin{cases} w_i & \text{if } a = i \\ 0 & \text{if not.} \end{cases}$$

Second-price auction:

- each player bids an offer $\tilde{w}_i$
- the item is assigned to the maximum bidder
- the winner pays the second highest bid
- losers pay nothing and their utility is zero

Example: 4 bidders with valuations $w_1 = 100, w_2 = 80, w_3 = 75, w_4 = 50$. The object is assigned to bidder 1, and she will pay 80.

Example: Second-price auctions

In the second price auction each player has as a dominant strategy $\tilde{w}_i = w_i$, that is, bid exactly her true valuation.

Proof fix $s = (s_1, \dots, s_n)$. We need to prove that for player i declaring w_i is a weakly dominant strategy. Let $s_j = \max\{s_k : k \neq i\}$.

- Case $w_i > s_j$. In this case declaring any value $s_i > s_j$ (including $s_i = w_i$) provides the player the same positive utility $w_i - s_j$, while declaring any $s_i < s_j$ provides utility zero: thus the claim is verified in this case
- Case $w_i < s_j$. In this case declaring any value $s_i > s_j$ provides the player negative utility $w_i - s_j$, while declaring any $s_i < s_j$ (including $s_i = w_i$ provides utility zero: thus the claim is verified also in this case. ■

Example: Second-price auctions

Truthfulness allows to assign the object to the player who values it the most.
This is equivalent to maximizing

$$\max_{a \in A} w_a = \max_{a \in A} \sum_{i=1}^{n} v_i(a).$$

Observe that second-price auctions are susceptible to collusion. In the example player 1 gets the object and pays 80. Player 2, knowing she is lost, can negotiate with player 1 to declare a smaller $\tilde{w}_2 = 75$ and to split half and half the 5 extra units of utility that player 1 will save.

VCG mechanisms (Vickrey-Clarke-Groves)

Definition

A VCG mechanism (with Clarke's pivot) is a pair (f, p) such that

- *it implements a social optimum* $f(s) \in \mathrm{argmax}_{a\in A} \sum_{i=1}^{n} s_i(a)$,
- *each player i has to pay*

$$p_i(s) = h_i(s_{-i}) - \sum_{j\neq i} s_j(f(s))$$

where $h_i(s_{-i})$ is the so-called Clarke's pivot which is given by

$$h_i(s_{-i}) = \max_{a\in A} \sum_{j\neq i} s_j(a).$$

Note that Clarke's pivot yields $p_i(s) \geq 0$, charging to each player i the *externality* or loss of utility that her presence imposes on the rest of the society.

Example

Ann, Bob, Carol and David are choosing a dessert for dinner. Their true valuations are

	A	B	C	D	Total
Gelato	3	4	2	10	19
Tiramisù	2	5	7	4	18
Torta	5	5	2	2	14

and the optimal social choice is Gelato

Payments and utilities: If David were not present the optimal choice would be Tiramisù and Ann, Bob and Carol would get a total utility of 14. Because of David they only get 9 so that David must pay $p_D = 14 - 9 = 5$ and his utility is $u_D = 10 - 5 = 5$.

If Ann or Bob or Carol were not there the optimal decision would still be Gelato. Their presence does not change the utility of the rest of the society, so they pay nothing and their final utilities are 3, 4 and 2.

What if Carol misreports her Gelato valuation as 0 to change the decision to Tiramisù? The rest of society would get 11 instead of 17. Hence Ann would have to pay 6 and her utility would be $1 = 7 - 6$, less than the utility she gets by declaring truthfully!

VCG mechanisms are thruthful

Theorem

Every VCG mechanism with Clarke's pivot is truthful.

Proof. The utility of the player i is

$$u_i(s) = v_i(f(s)) + \sum_{j \neq i} s_j(f(s)) - h_i(s_{-i}).$$

Her strategy s_i only affects the utility through $f(s)$ so the player's interest is that the chosen action $a = f(s)$ maximizes $v_i(a) + \sum_{j \neq i} s_j(a)$. Since VCG implements a social optimum $f(s) \in \operatorname{argmax}_{a \in A} \sum_{i=1}^{n} s_i(a)$, then the player interest is achieved by declaring $s_i = v_i$. □

Remark: Instead of side payments $p_i(s) = h_i(s_{-i}) - \sum_{j \neq i} s_j(f(s))$ with Clarke's pivot, we may take any function $h_i : V_{-i} \to \mathbb{R}$. This more general form of VCG provides much flexibility to define side payments.

VCG with Clarke's pivot – Voluntary participation

The payments based on Clarke's pivot have another important property: the utilities are non-negative which encourages players to participate in the game.

Proposition

In a VCG mechanism with Clarke's pivot, if the valuations are non-negative the dominant strategy v_i satisfies $u_i(v_i, s_{-i}) \geq 0$.

Proof. Consider a profile $s = (s_1, \ldots, s_n)$ with $s_i = v_i$. The first property of VCG implies

$$\begin{aligned} u_i(s) &= \sum_{j=1}^n s_j(f(s)) - h_i(s_{-i}) \\ &= \max_{a \in A} \sum_{j=1}^n s_j(a) - \max_{a \in A} \sum_{j \neq i} s_j(a) \end{aligned}$$

whose non-negativity results from taking maximum over $a \in A$ in the inequality $\sum_{j \neq i} s_j(a) \leq \sum_{j=1}^n s_j(a)$. □

Example: Single-item auction

In an auction of an object with valuations

$$v_i(a) = \begin{cases} w_i & \text{if } a = i \\ 0 & \text{if not} \end{cases}$$

the sum of utilities $\sum_{i=1}^{n} v_i(a) = w_a$ is equal to the valuation of player a.

In a direct mechanism the players are asked to declare their valuations $w_i \geq 0$. Eventually, player might misreport their valuations declaring $\tilde{w}_i \neq w_i$.

The VCG mechanism assigns the object to the highest bidder, while Clarke's pivot generates payments $p_i = \max_{j \neq i} \tilde{w}_j$ for the winner and $p_j = 0$ for the others.

This is the second-price auction.

Example: Multiple-item auctions

Suppose now that we auction k identical objects.

The set of possible assignments in this case is

$$A = \{(a_1, \ldots, a_n) : a_i \in \mathbb{N}, a_1 + \cdots + a_n = k\}.$$

Each player i has a valuation $v_i(a) = u_i(a_i) \geq 0$ with $u_i(\cdot)$ increasing in the number a_i of items assigned to her. We assume that the marginal valuation $m_i(a) = u_i(a) - u_i(a-1)$ of the a-th unit decreases with a.

VCG assigns the items to the highest k marginal valuations, while Clarke's pivot determines a payment p_i equal to the sum of the a_i higher marginal valuations of the other players who were excluded because of i.

Example: Multiple-item auctions

Consider a 3-item auction between 3 players with marginal valuations

m_1	m_2	m_3
9	4	2
10	5	1
3	2	0

The 3 highest values 10, 9 and 5 are chosen, assigning one item to the first player and two items to the second.

The payments are $p_1 = 3$ for the first player (if this player was not there the item would have been assigned to the third player with value 3).

Similarly $p_2 = 4 + 3 = 7$ and $p_3 = 0$.

Example: Provision of a public good

The Government must decide to execute or not a project of public interest, the options are $A = \{0, 1\}$. The total cost of the project is c. There are $i = 1, \ldots, n$ users whose true valuations for the project are $w_i \geq 0$ so that

$$v_i(a) = \begin{cases} w_i & \text{if } a = 1 \\ 0 & \text{if } a = 0 \end{cases}$$

The project is socially efficient if the sum of valuations $\sum_{i=1}^{n} w_i$ exceeds the cost of the project c. Users will tend to declare a much larger valuation $\tilde{w}_i \gg w_i$ in order to make the project look socially efficient.

To extract the true valuations we implement a VCG mechanism. To this end we consider the government as an additional player 0 with valuation

$$v_0(a) = \begin{cases} -c & \text{if } a = 1 \\ 0 & \text{if } a = 0 \end{cases}$$

Example: Provision of a public good

The VCG mechanism decides to execute the project $(a = 1)$ iff $\sum_{i=1}^{n} w_i \geq c$, whereas Clarke's pivot determines the payments

$$p_i(w) = \begin{cases} c - \sum_{j \neq i} w_j & \text{if } \sum_{j \neq i} w_j < c \leq \sum_{j=1}^{n} w_j \\ 0 & \text{otherwise} \end{cases}$$

Thus: a player only pays if her valuation happens to be pivotal for the project. Note that these payments do not ensure to cover the costs: it could even be the case that no player is pivotal and the payments are all zero, in which case the Government will have to pay for the full project.

Example: Buying a shortest path

- Consider an undirected graph $G = (V, E)$ in which each edge $e \in E$ has an owner who values it $c_e > 0$.
- We wish to buy a path r connecting o to d at the lowest possible cost, so that the set A of actions consists of all such possible paths.
- Define the valuation of player e as

$$v_e(r) = \begin{cases} -c_e & \text{if } e \in r \\ 0 & \text{if } e \notin r \end{cases}$$

- The VCG mechanism asks each arc its price c_e and implements a route r^* of minimal cost $C^* = \sum_{e \in r^*} c_e$.
- The payments generated by Clarke's pivot are $p_e = 0$ for $e \notin r^*$ while for $e \in r^*$ we have $p_e = C^* - C^*_{-e}$ with C^*_{-e} the cost of a minimum path in the graph $(V, E \setminus \{e\})$.
- Note that $p_e \leq 0$. In other words, each arc $e \in r^*$ receives a payment equivalent to its contribution to reducing the cost of the minimum route.