Single-Parameter Mechanisms

Instructor: Xiaohui Bei

In the previous lecture, we learned basic concepts about mechanism design. The goal in this area is to design incentive compatible mechanisms that can implement certain desired social choice functions. In the most general setting such task may not always be possible. Hence we turn our focus on special subclass of utility functions.

In particular, we introduced the quasilinear environment, in which the mechanism needs to determine an outcome as well as the payment for each agent, and agents all have quasilinear utilities, which are equal to a agent's valuation for the outcome minus the money that he needs to pay.

A good mechanism in such quasilinear environment should have the following properties:

- **Incentive Compatible**. Bidding his utility function truthfully should be a dominant strategy for every agent.
- **Performance Guarantee**. Given reported utility functions $\mathbf{u} = (u_1, \ldots, u_n)$, the mechanism should output an outcome $o \in O$ that maximizes social welfare, defined as $\sum_i u_i(o)$.
- **Computational Efficiency.** The mechanism can be implemented in polynomial time.

We have already seen that in a simple but representative example – single-item auctions, there exists such a good mechanism – the second-price auction, where the bidder with the highest bid wins the item and pays the price equal to the second highest bid. In this lecture, we will extend these guarantees beyond single-item auctions to a more general setting, and provide some good characterizations of such desired mechanisms.

1 Single-Parameter Domain

Recall that in the quasilinear environment, the the set of outcomes has form $O = A \times \mathbb{R}^n$, where A is the set of possible allocations, and $\mathbb{R}^n$ represents the payments that agents have to pay (or receive) for a given allocation. And the utility function u_i of each agent i has form $u_i(a, (p_1, \ldots, p_n)) = v_i(a) - p_i$ for any $(a, (p_1, \ldots, p_n)) \in O$. Here $v_i(\cdot)$ is also called the valuation function of agent i.

A natural way to generalize the single-item auctions is to consider the case where every agent's valuation function can be described by one single parameter. This is called the *single-parameter domains.*

Definition 9.1 (Single-Parameter Domains)**.** *In single-parameter domains, the valuation function $v_i(\cdot)$ of each agent i can be presented by a "winning set" $W_i \subseteq A$, and a single value $t_i \in \mathbb{R}$. Such that*

$$v_i(a) = \begin{cases} t_i & \text{if } a \in W_i \\ 0 & \text{otherwise} \end{cases}$$

An intuitive interpretation of such domains is as following: assume that for each agent i, any allocation can represent either "winning" or "losing" to him. This agent has a private scalar value t_i for "winning", and value 0 for "losing". The set of "winning" allocations is modeled by a commonly known set W_i. The main point is that all "winning" allocations are equivalent to each other for this agent, and the same is true for all "losing" allocations.

A mechanism in single-parameter domains can then be described by two parts, assuming that $\mathbf{b} = (b_1, \ldots, b_n)$ is the set of bids reported by the agents to the mechanism.

- **Allocation.** Choose an allocation function $f(\mathbf{b}) \in A$ to decide an allocation.
- **Payments.** Choose payment functions $p_i(\mathbf{b}) \in R$ to decide the payment of each agent i.

With slight abuse of notations, in the following we will use v_i as the scalar t_i.

1.1 Examples

Single-Minded Combinatorial Auctions. Assume that we have n agents and m non-identical items. For each agent i, there is a publicly known subset A_i of items the agent desires, and he has a valuation $t_i \in \mathbb{R}$ indicating his value for receiving a bundle that contains A_i (0 otherwise). The set of allocations contains all possible partitions of items among agents. Our goal is to find an allocation that maximizes the social welfare, which is the sum of values of agents who receive their desired bundles.

Shortest Path Procurement. Assume that we have a network with designated source s and sink t. Each edge in this network is an agent, and agent i has a private cost c_i. Every possible outcome consists of a path from s to t to buy, as well as the payment to each agent along the path. The utility of an agent is equal to the payment to him minus his cost c_i if his edge belongs to the path. Our goal is to buy a path with lowest total cost.

2 Characterization of Incentive Compatibility

In this section we provide a complete characterization of the incentive compatibility condition of a mechanism in single-parameter domains. First, we introduce two important definitions for allocation functions.

Definition 9.2 (Monotonicity)**.** *An allocation function f is called* **monotone** *for agent i if*

$$f(b_i, b_{-i}) \in W_i \implies f(b'_i, b_{-i}) \in W_i$$

for every b_{-i} and every $b_i \leq b'_i \in \mathbb{R}$.

Intuitively, monotonicity states that for every bids b_{-i} such that agent i can both win and lose, there always exists a value, which we called the *critical value*, such that agent i always loses with bid b_i below this value, and wins with b_i above it.

Definition 9.3 (Critical Value)**.** *The* **critical value** *for agent i of a monotone social choice function f is defined as*

$$c_i(b_{-i}) = \sup_{b_i : f(b_i, b_{-i}) \notin W_i} b_i.$$

In this lecture, we will focus on payment functions with the following properties:

(1) $p_i(\mathbf{b})$ is always non-negative.

(2) If $f(\mathbf{b}) \notin W_i$, then $p_i(\mathbf{b}) = 0$.

These are both natural assumptions. Condition (1) requires that the mechanism always collects money from the agents. Condition (2) guarantees that losing agents should not pay any money. It is not difficult to see that every incentive compatible mechanism can be easily turned in to a mechanism that satisfies these two conditions.

Theorem 9.4. *A mechanism $(f, p_1, \ldots, p_n)$ in a single-parameter domain is incentive compatible if and only if the following conditions hold:*

(1) f is monotone for every agent i,.

(2) Every winner agent pays the critical value. That is, for every i such that $f(\mathbf{b}) \in W_i$, we have $p_i(\mathbf{b}) = c_i(b_{-i})$.

Proof. **[If part]** Fixing an agent i and the bids from other agents b_{-i}. According to the mechanism, agent i will receive a utility of $v_i - c_i(b_{-i})$ if he wins, or 0 if he loses. Hence he should prefer winning if $v_i > c_i(b_{-i})$ and losing if $v_i < c_i(b_{-i})$. By bidding the truth $b_i = v_i$ he can achieve exactly this.

[Only if part] This part contains two cases.

First, assume that f is not monotone for some agent i, that is, there exists some b_{-i} and values $b'_i > b_i$ such that $f(b'_i, b_{-i}) \notin W_i$ but $f(b_i, b_{-i}) \in W$. If the mechanism is still incentive compatible, then an agent i with valuation b_i should not benefit by bidding b'_i. That is, $b_i - p_i(b_i, b_{-i}) \geq 0$. Also an agent i with valuation b'_i should not benefit by bidding b_i, this gives us $b'_i - p(b_i, b_{-i}) \leq 0$. Together we have $b_i \geq b'_i$, which is a contradiction.

Second, assume that the price of some winning agent i is not $c_i(b_{-i})$. It is easy to see that since all winning outcomes have the same value to an agent, he should make the same payment for all winning bids. Let this price be $p_i(b_{-i})$. If $p_i(b_{-i}) > c_i(b_{-i})$, then an agent with valuation $p_i(b_{-i}) > v_i > c_i(b_{-i})$ would win by bidding his true value, but paying a value higher than his valuation, hence resulting a negative utility. Thus he is better off bidding some value $v' < c_i(b_{-i})$. In the other direction, if $p_i(b_{-i}) < c_i(b_{-i})$. Then an agent with losing valuation $p_i(b_{-i}) < v_i < c_i(b_{-i})$ would be better off bidding a value $v' > c_i(b_{-i})$ that can make him win. Thus in both case the mechanism is not incentive compatible. □

3 Randomized Mechanisms

So far all mechanisms that we talked about are deterministic mechanisms. It is a natural thought to introduce randomness into mechanism design. For example, the mechanism could return a distribution over outcomes and payments. Another way to think about this, is to consider a randomized mechanism as a distribution over deterministic mechanisms.

How should we refine the definition of incentive compatibility for a randomized mechanism? It turns out that we have two possible definitions.

Definition 9.5. *A randomized mechanism is* **universally incentive compatible** *if every deterministic mechanism in its support is incentive compatible.*

Definition 9.6. *A randomized mechanism is* **incentive compatible in expectation** *if truth-telling is a dominant strategy in the game induced by expectation.*

That is, for every agent i and every v_i, v'_i, v_{-i}, let (a, p_i) and (a', p'_i) be the random variables denoting the outcome and payment of agent i when he bids v_i and v'_i respectively. Then

$$E[v_i(a) - p_i] \geq E[v_i(a') - p'_i],$$

where $E[\cdot]$ denotes expectation over the randomization of the mechanism.

Clearly universal incentive compatibility is a stronger notion than incentive compatibility in expectation. Any mechanism that is universally incentive compatible must also be incentive compatible in expectation. In most scenarios, incentive compatibility in expectation would be enough to serve the purpose. In general randomized mechanisms can often be more useful and achieve more than deterministic ones.

4 Myerson's Lemma

Next we try to characterize randomized incentive compatible mechanism over single-parameter domains. For every $\mathbf{b} = (b_1, \ldots, b_n)$, we use $x_i(\mathbf{b}) = \Pr[f(\mathbf{b}) \in W_i]$ to denote the probability that agent i wins when bidding b_i. We also use $p_i(\mathbf{b})$ to directly denote the expected payment for agent i. With these notations, the expected utility for agent i given bids $\mathbf{b}$ is $v_i \cdot x_i(\mathbf{b}) - p(\mathbf{b})$. For ease of notation we write $x(b_i)$ and $p(b_i)$ when b_{-i} is fixed.

An alternative but equivalent way to look at randomized mechanisms over single-parameter domains is the following: instead of considering only "win" or "lose" for one agent, assume that there is a single homogenous resources, and there are constraints on how this resource can be divided among agents. Each agent's private information v_i now represents his "value per unit resource". Formally speaking, each allocation $a \in A$ is a vector $(x_1, \ldots, x_n) \in [0,1]^n$, where x_i denotes the amount of resources given to agent i, and agent i's valuation function is $v_i(a) = v_i x_i$. Thus upon collecting all the bids $\mathbf{b} = (b_1, \ldots, b_n)$ from the agents, the mechanism will decide: (1) a feasible allocation $\mathbf{x}(\mathbf{b}) \in A \subseteq \mathbb{R}^n$, and (2) payments $p_i(\mathbf{b})$ for each agent i.

The incentive compatibility characterization of randomized mechanisms is given by the following celebrated Myerson's Lemma.

Theorem 9.7 (Myerson's Lemma, 1981)**.** *A randomized mechanism $(f, p_1, \ldots, p_n)$ in a single-parameter domain is incentive compatible in expectation if and only if the following conditions hold:*

- $x(\mathbf{b})$ *is monotone for every agent i.*
- *The payment for agent i is* $p(b_i) = b_i \cdot x(b_i) - \int_0^{b_i} x_i(t)\,\mathrm{d}t$

Proof.

- **[Only if part]** Since the mechanism is incentive compatible in expectation, for any values $b_i' > b_i'$, agent i with true private value b_i would not receive a better utility by reporting b_i' instead of b_i. Thus we have

$$b_i x(b_i) - p(b_i) \geq b_i x(b_i') - p(b_i').$$

The same is true if we switch b_i' and b_i. Hence we also have

$$b_i' x(b_i') - p(b_i') \geq b_i' x(b_i) - p(b_i).$$

Combining the above two equations gives us

$$b_i'[x(b_i) - x(b_i')] \leq p(b_i) - p(b_i') \leq b_i[x(b_i) - x(b_i')].$$

The first and third term in above inequality directly implies the monotonicity of the allocation function.

For the payment function, we divide the above inequality by $b_i - b'_i$, and take the limit as b'_i goes to b_i, we get

$$b_i \frac{\mathrm{d}x}{\mathrm{d}b_i}(b_i) \leq \frac{\mathrm{d}p}{\mathrm{d}b_i}(b_i) \leq b_i \frac{\mathrm{d}x}{\mathrm{d}b_i}(b_i)$$

which implies $b_i \frac{\mathrm{d}x}{\mathrm{d}b_i}(b_i) = \frac{\mathrm{d}p}{\mathrm{d}b_i}(b_i)$. Hence we have

$$\begin{aligned} p(b_i) &= \int_0^{b_i} \frac{\mathrm{d}p}{\mathrm{d}t}(t)\,\mathrm{d}t = \int_0^{b_i} t\frac{\mathrm{d}x}{\mathrm{d}t}(t)\,\mathrm{d}t = \int_0^{b_i} t\,\mathrm{d}x(t) \\ &= t \cdot x(t)\big|_0^{b_i} - \int_0^{b_i} x(t)\,\mathrm{d}t \\ &= b_i x(b_i) - \int_0^{b_i} x(t)\,\mathrm{d}t \end{aligned}$$

- **[If part]** Here we give a proof by picture that with monotone allocation function and payment functions as described, the mechanism is incentive compatible. The following figures demonstrate the scenarios where an agent bids truthfully/underbids. We can easily observe that the agent's utility achieves maximal when he bids truthfully. The case where the agent overbids is left as an exercise to the readers.

□

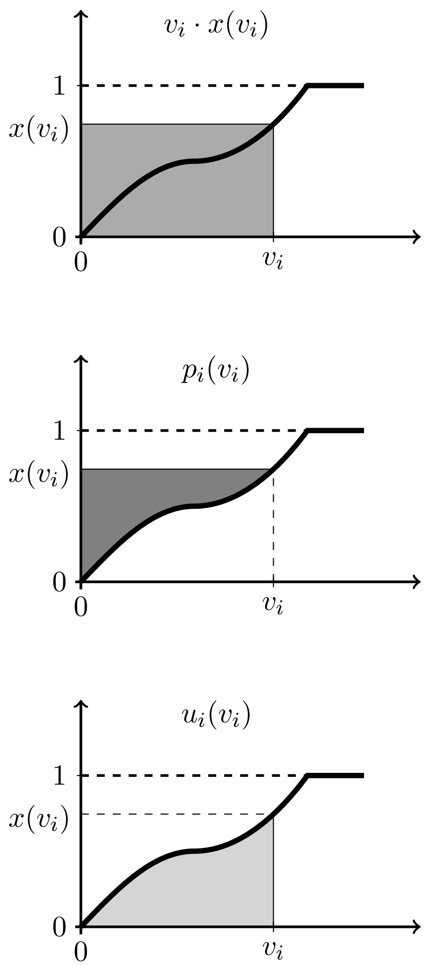

Bidding $b_i = v_i$.

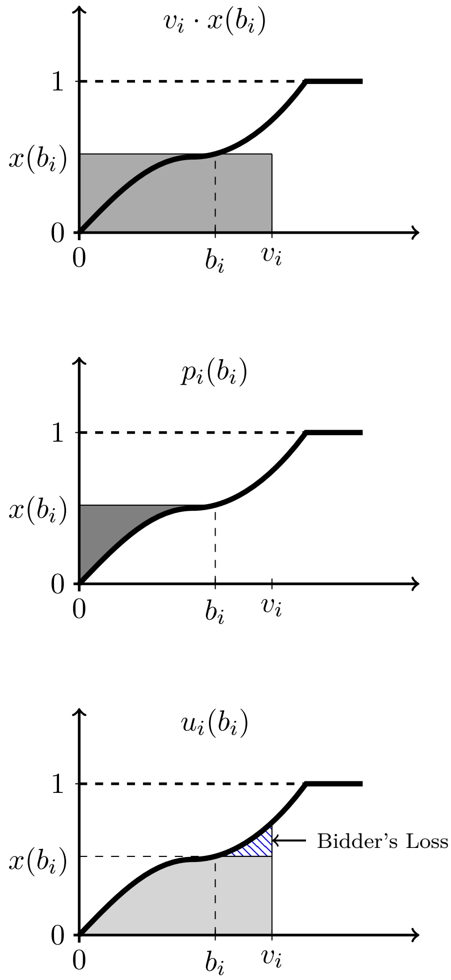

Bidding $b_i < v_i$.

Recommended Literature

- Chapter 9.5 in the AGT book.
- Tim Roughgarden's lecture notes `http://theory.stanford.edu/~tim/f13/l/l3.pdf` and lecture video `https://youtu.be/9qZwchMuslk`
- R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58-73, 1981.