Lecture 11: June 13

Lecturer: Amos Fiat Scribe: Lucas Majerowicz, Sharon Kotler Pikaz and Yaron Velner

11.1 Introduction

In this lecture we'll review optimal mechanism designs for maximizing social surplus and profit. The optimal mechanism for social surplus is VCG, where the item is awarded to the highest bidder. For profit maximization there isn't a single optimal mechanism. The optimality of the mechanism depends on the distribution of the players.

11.1.1 Motivation

Let us consider using second-price auction on the case where there are agents with valuations v_1, v_2 drawn independently and at random from $U[0, 1]$. Recall that under this setting, the item will be sold to the highest bidder for the second highest bid value. Here is a proof that the expected revenue using this mechanism is $\frac{1}{3}$. By the definition of uniform random variables we have that:

$$\Pr[v_1 < x] = \Pr[v_2 < x] = x$$

which implies that

$$\Pr[v_1 > x] = \Pr[v_2 > x] = 1 - x$$

Since v_1 and v_2 are independent we get

$$\Pr[\min(v_1, v_2) > x] = \Pr[v_1 > x \wedge v_2 > x] = \Pr[v_1 > x] \cdot \Pr[v_2 > x] = (1-x)(1-x) = (1-x)^2$$

Recall that for a non-negative random variable X

$$E[X] = \int_0^\infty \Pr[X \geq z]\, dz$$

Hence,

$$E[\min(v_1,v_2)] = \int_0^\infty \Pr[\min(v_1,v_2) > x]\,dx = \int_0^1 (1-x)^2\,dx = \int_0^1 (1-2x+x^2)\,dx =$$

$$= x - x^2 + \frac{x^3}{3}\Big|_0^3 = \frac{1}{3}$$

It's interesting to note that by making a small modification to the second-auction mechanism we considered above we can get higher revenue. Consider a second auction mechanism with reserve price of $\frac{1}{2}$. That is, if none of the two agents bids more than $\frac{1}{2}$, the seller gets to keep the item. Otherwise, the price that the winning agent pays is the maximum between the second biggest bid and $\frac{1}{2}$. Let's see that under this new settings the expected revenue is $\frac{5}{12}$, which is more than $\frac{1}{3}$. Let's consider the following three distinct cases:

- Case 1: both agents bid less than $\frac{1}{2}$. This happens with probability $\frac{1}{4}$ and in this case the revenue is 0, as the seller doesn't sell the item.
- Both bids are above $\frac{1}{2}$. This also happens with probability $\frac{1}{4}$ and in this case the item is sold and the expected revenue is $\frac{2}{3}$.
- Case 3: one of the bids is less than $\frac{1}{2}$ and the other bids is above $\frac{1}{2}$. This happens with probability $\frac{1}{2}$ and the expected revenue is $\frac{1}{2}$ (the reservation price).

Therefore, the expected revenue is

$$\frac{1}{4}\cdot 0 + \frac{1}{4}\cdot\frac{2}{3} + \frac{1}{2}\cdot\frac{1}{2} = \frac{5}{12}$$

which, as pointed out above, is higher than $\frac{1}{3}$.

11.2 The Single-Dimensional Environment

In this setting we have n agents, each of them has a single valuation v_i for receiving certain item or service. The outcome of a single game is described by two vectors:

- The allocation vector $x = (x_1, \ldots, x_n)$, where x_i is an indicator variable that receives value 1 if agent i has been served, and value 0 otherwise.

- The payment vector $p = (p_1, \ldots, p_n)$ where each p_i denotes the payment made by agent i.

Using the vectors defined above, we compute each agent's utility as $u_i = v_i x_i - p_i$. In other words, each agent would like to maximize the difference between their valuation and how much they actually pay for the item or service.

11.2.1 Environment classification

Depending on the constraints and goals of the environment, they can be classified as:

- A *general cost environment* is one where the designer must pay a service cost $c(x)$ for the allocation x produced.
- A *General feasibility environment* is one where there is a feasibility constraint over the set of agents that can be simultaneously served.
- A *downward-closed feasibility constraint* is one where subsets of feasible sets are feasible.

11.3 Design Objectives

The two main objectives to maximize are *social surplus* and *profit*, which are defined as follows:

$$\mathit{Surplus}(v, x) = \sum_i v_i x_i - c(x)$$

$$\mathit{Profit}(v, x) = \sum_i p_i - c(x)$$

11.4 Social Surplus Maximization

The optimization problem of surplus maximization is that of finding x to maximize:

$$\mathit{Surplus}(v, x) = \sum_i v_i x_i - c(x)$$

Let OPT be the optimal value of the maximization problem, namely

$$OPT(v) = \max_x Surplus(v, x)$$

Lemma 11.1 *For each agent i and all values of other agents v_{-i}, the allocation rule of OPT for agent i is a step function.*

Proof Denote (v_{-i}, z) as the vector v with the i-th coordinate replaced with z. For any agent i, we will consider the cases where: (1) i is served by $OPT(v)$; and (2) i is not served by $OPT(v)$.

- Case 1: $i \in OPT(v)$

  $$OPT(v) = \max_x Surplus(v, x) = v_i + \max_{x_{-i}} Surplus((v_{-i}, 0), (x_{-i}, 1))$$

  Let us define $OPT_{-i}(v) = \max_{x_{-i}} Surplus((v_{-i}, 0), (x_{-i}, 1))$. Then we have that

  $$OPT(v) = v_i + OPT_{-i}(v)$$

  Notice that $OPT_{-i}(v)$ is not a function of i.

- Case 2: $i \notin OPT(v)$

  $$OPT(v) = \max_x Surplus(v, x) = \max_{x_{-i}} Surplus((v_{-i}, 0), (x_{-i}, 0))$$

  Let us define $OPT(v_{-i}) = \max_{x_{-i}} Surplus((v_{-i}, 0), (x_{-i}, 0))$. Then we have that

  $$OPT(v) = OPT(v_{-i})$$

  Notice that $OPT(v_{-i})$ is not a function of i.

OPT chooses whether to allocate agent i so as to maximize the surplus. That means that i will be allocated by OPT whenever the surplus from case 1 is greater than the surplus from case 2, i.e., when

$$v_i + OPT_{-i}(v) \geq OPT(v_{-i})$$

Solving for v_i, we get that OPT allocates to agent i when

$$v_i \geq OPT(v_{-i}) - OPT_{-i}(v)$$

None of the terms in the right side of the above inequality contains v_i. Therefore the allocation rule for i is a step function with critical value $\tau_i = OPT(v_{-i}) - OPT_{-i}(v)$, as illustrated in Figure 11.1. ∎

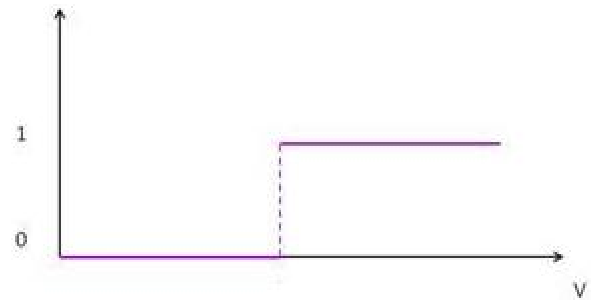

Figure 11.1: **Critical value:** $\tau_i = OPT(v_{-i}) - OPT_{-i}(v)$

11.4.1 Single dimensional surplus maximization mechanism

The mechanism for Surplus maximization is:

1. Solicit and accept scaled bids b
2. $x \leftarrow OPT(b)$
3. For each i, $p_i \leftarrow OPT(b_{-i}) - OPT_{-i}(b)$

The single dimensional surplus maximization mechanism is a special case of the Vickrey-Clarke-Groves (VCG) mechanism (with Clarke Pivot Payments). With Clarke Pivot Payments (i.e., the "critical values") truth-telling is a dominant strategy equilibrium.

Lemma 11.2 *A single dimensional deterministic mechanism is Dominant Strategy Incentive Compatible (DSIC) if and only if for all i:*

1. *(step-function) x_i^M steps from 0 to 1 at*

$$\tau_i = \inf\{z \mid x_i^M(v_{-i}, z) = 1\}$$

and

2. *(critical value)*

$$p_i^M(v_{-i}, v_i) = p_i^M(v_{-i}, 0) + \begin{cases} \tau_i & \text{if } x_i^M(v_{-i}, v_i) = 1 \\ 0 & \text{otherwise} \end{cases}$$

By Lemmas 11.1 and 11.2 we have the following theorem:

Theorem 11.3 *The surplus maximization mechanism is dominant strategy incentive compatible (DSIC)*

By the optimality of OPT and the assumption that agents follow the dominant truth-telling strategy, we have the following corollary:

Corollary 11.4 *The surplus maximization mechanism optimizes social surplus in dominant strategy equilibrium (DSE).*

11.5 Profit Maximization

11.5.1 The model

For the maximization of surplus objective there is unique mechanism that is optimal regardless of distributional assumptions. However, for the maximization of other objectives we should not expect to be so lucky.

Profit maximization **depends on the distribution**. When the distribution of agent values is specified, and the designer has knowledge of this distribution, the profit can be optimized. The mechanism that results from such an optimization is said to be **Bayesian optimal**.

Formally, the value of each agent is denote by v_i, but this value is not visible to the other agents and to the designer; instead, only a distribution on the value of the agent is known. The actions of the players are a vector $b = (b_1, \ldots, b_n)$, and a mechanism M maps actions (i.e., *bids*) to an allocation vector and a payment vector. That is, for every agent i and a vector of bids $b = (b_1, \ldots, b_n)$, the following functions are defined by the mechanism:

- $x_i^M(b)$ is the expected outcome for agent i.
- $p_i^M(b)$ is the payment rule for agent i.

We are interested in two basic properties of mechanisms. The first property is that, given the mechanism, the agents dominate strategy is to bid according to their true value. The second property is that the mechanism maximizes the revenue. We will study these properties in the next two subsections.

11.5.2 Bayesian insensitive compatible mechanisms

A mechanism induces an expected utility function $u_i(b)$, namely

$$u_i^M(b) = v_i x_i^M(b) - p_i^M(b)$$

Intuitively, a mechanism is *Bayesian insensitive compatible* (BIC), if for every agent i, bidding $b_i = v_i$ is a dominating strategy. Formally, M is BIC if for every agent i and every action b_i we have

$$u_i^M(v_i, b_{-i}) \geq u_i^M(b_i, b_{-i})$$

The next theorem characterizes the BIC mechanisms

Theorem 11.5 *A mechanism M is BIC if and only if, for any agent i and bids of other agents b_{-i} fixed,*

1. *$x_i^M(b_i)$ is monotone non-decreasing; and*
2. *$p_i^M(b_i) = b_i x_i^M(b_i) - \int_0^{b_i} x_i^M(z)\, dz$*

Proof ($\Leftarrow$) We prove only the direction from left to right (the proof for the converse direction was not given in class). We consider two possible bid values b_i and v_i with $b_i < v_i$, and we argue that agent i does not benefit by shading his bid down to b_i. Indeed, recall that

$$u_i(z) = v_i x_i(z) - p_i(z)$$

Let us observe the expression $u_i(v_i) - u_i(b_i)$, and we shall prove that its value is non-negative. Indeed, by the definition of x_i and p_i we get

$$u_i(v_i) - u_i(b_i) =$$

$$v_i x_i(v_i) - v_i x_i(v_i) + \int_0^{v_i} x_i(y)\, dy$$

$$-$$

$$v_i x_i(b_i) - b_i x_i(b_i) + \int_0^{b_i} x_i(y)\, dy$$

$$=$$

$$\int_0^{v_i} x_i(y)\, dy - \int_0^{b_i} x_i(y)\, dy - x_i(b_i)(v_i - b_i)$$

$$= \int_{b_i}^{v_i} x_i(y)\, dy - x_i(b_i)(v_i - b_i)$$

We prove that the last expression is non-negative by Figure 11.2.

The proof for the case where $v_i < b_i$ is by similar arguments. ∎

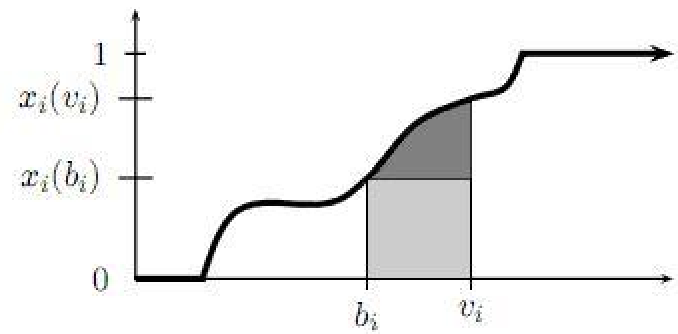

Figure 11.2: The bright area is exactly $x_i(b_i)(v_i - b_i)$, the bright + dark area is exactly $\int_{b_i}^{v_i} x_i(y)\,dy$. Thus, for every monotone non-decreasing function $x_i()$, we have $\int_{b_i}^{v_i} x_i(y)\,dy - x_i(b_i)(v_i - b_i) \geq 0$.

11.5.3 Revenue maximization

In the context of revenue maximization, it is easier to work with *quantile spaces*. Intuitively, given a distribution function over the values of all agents, the *quantile* of a value of an agent is the percentage of agents with higher values than the agent. For example, an agent with value $v = 0.9$ drawn from $U[0,1]$ is stronger than 90% of agents and weaker than 10% with values drawn from the same distribution. Formally, let F be a a one-to-one mapping between the agent's value and his strength relative to the distribution of agent's values, then the quantile q of a value v in the support of F is the probability that $v \leq r$, for a random variable $r \sim F$. By definition, we get that $q = 1 - F(v)$, and thus we can express v, x and p as functions of q, namely:

- $v(q) = F^{-1}(1-q)$
- $x(q) = x(v(q))$
- $p(q) = p(v(q))$

By changing variables and integrating by parts, we can rederive Theorem 11.5 in the setting of quantile spaces, as follows.

Theorem 11.6 *A mechanism M is BIC if and only if, for any agent i:*

1. *$x_i^M(q_i)$ is monotone non-increasing; and*
2. $p_i^M(q_i) = -\int_{q_i}^{1} v_i(r) \frac{d}{dq} x_i^M(r)\, dr + p_i(1)$

To simplify the model, we assume that we have only a single agent, and his value v is drawn from distribution F. In this case, the revenue curve $R(q)$ specifies the revenue as a function of the ex ante probability of sale, where $R(q)$ is

$$R(q) = v(q) \cdot q$$

Hence, to find the mechanism that optimizes the revenue (recall that the mechanism corresponds to q), we need to find q that brings $R(q)$ to a maximum.

Example Let F be the uniform distribution $U[0, 1]$, then

- $F(z) = z$
- $v(q) = 1 - q$
- $R(q) = q - q^2$

To find the maximum of $R(q)$ we compute $\frac{dR}{dq}$, and take q such that $\frac{dR}{dq} = 0$. In this example, $\frac{dR}{dq} = 1 - 2q$ and therefore $q = \frac{1}{2}$ maximizes the revenue to $R(\frac{1}{2}) = \frac{1}{4}$.

We note that in some cases, the equation $\frac{dR}{dq} = 0$ might have multiple solutions, and we will show how to deal with such cases in the next lecture.