Computational Game Theory	Spring Semester, 2009/10

Lecture 13:Revenue Maximization

Lecturer:Yishay Mansour *Scribe:Tovi Almozlino, Alon Baraam*

13.1 Introduction

In previous lectures, we have studied the design of truthful mechanisms that implement social choice functions, such as social welfare maximization. Another fundamental objective, and the focus of this chapter, is the design of mechanisms in which the goal of the mechanism designer is profit maximization. In economics this topic is referred to as *optimal mechanism design*.

Our focus will be on the design of profit-maximizing auctions in settings in which an auctioneer is selling (respectively, buying) a set of goods/services. Formally, there are n agents, each of whom desires some particular service. We assume that agents are *single-parameter*; i.e., agent i's *valuation* for receiving service is v_i and their valuation for no service is normalized to zero. A mechanism takes as input sealed bids from the agents, where agent i's bid b_i represents his valuation v_i, and computes an outcome consisting of an allocation $\mathbf{x} = (x_1, \ldots, x_n)$ and prices $\mathbf{p} = (p_1, \ldots, p_n)$. Setting $x_i = 1$ represents agent i being allocated service whereas $x_i = 0$ is for no service, and p_i is the amount agent i is required to pay the auctioneer. We assume that agents have quasi-linear *utility* expressed by $u_i = v_i x_i - p_i$. Thus, an agent's goal in choosing his bid is to maximize the difference between his valuation and his payment.

Our goal is to design the mechanism, i.e., the mapping from bid vectors to price/allocation vectors so that the auctioneer's profit, defined as

$$Profit = \sum_i p_i,$$

is maximized, and the mechanism is *truthful*.

13.1.1 Preliminaries

In this section we review basic properties of truthful mechanisms.

We will place two standard assumptions on our mechanisms. The first, that they are *individually rational*, means that no agent has negative expected utility for taking part in the mechanism. The second condition we require is that of no *positive transfers* which restricts the mechanism to not pay the agents when they do not win, i.e., $x_i = 0 \rightarrow p_i = 0$.

In general, we will allow our mechanisms to be randomized. In a randomized mechanism, x_i is the probability that agent i is allocated the good, and p_i is agent i's expected payment. Since x_i and p_i are outputs of the mechanism, it will be useful to view them as functions of the input bids as follows. We let $x_i(\mathbf{b})$, $p_i(\mathbf{b})$, and $u_i(\mathbf{b})$ represent agent i's probability of allocation, expected price, and expected utility, respectively. Let $\mathbf{b}_{-i} = \left(b_1 \ldots, b_{i-1}, ?, b_{i+1}, \ldots, b_n\right)$ represent the vector of bids excluding bid i. Then with $\mathbf{b}_{-i}$ fixed, we let $x_i(b_i)$, $p_i(b_i)$, and $u_i(b_i)$ represent agent i's probability of allocation, expected price, and expected utility, respectively, as a function of their own bid. We further define the convenient notation $x_i(b_i, \mathbf{b}_{-i}) = x_i(\mathbf{b})$, $p_i(b_i, \mathbf{b}_{-i}) = p_i(\mathbf{b})$, and $u_i(b_i, \mathbf{b}_{-i}) = u_i(\mathbf{b})$.

Definition 13.1. A mechanism is *strategyproof in expectation* if and only if for all i, v_i, b_i, and $\mathbf{b}_{-i}$, agent i's expected utility for bidding their valuation, v_i, is at least their expected utility for bidding any other value. In other words,

$$u_i(v_i, \mathbf{b}_{-i}) \geq u_i(b_i, \mathbf{b}_{-i}).$$

Theorem 13.2. *A mechanism is strategyproof in expectation if and only if, for any agent i and any fixed choice bids by the other agents $\mathbf{b}_{-i}$,*

(i) *$x_i(b_i)$ is monotone nondecreasing.*

(ii) *$p_i(b_i) = b_i x_i(b_i) - \int_0^{b_i} x_i(z)\, dz$.*

Proof. In the proof we will simplify notation by removing the index i everywhere. In this notation, to show strategyproofness we need to establish that

$$vx(v) - p(v) \geq vx(v') - p(v')$$

for every v'. Plugging in the formula for p we get

$$\cancel{vx(v)} - \cancel{vx(v)} + \int_0^v x(z)\, dz \geq vx(v') - v'x(v') + \int_0^{v'} x(z)\, dz = (v - v')\, x(v') + \int_0^{v'} x(z)\, dz.$$

For $v' > v$ this is equivalent to

$$(v' - v)\, x(v') \geq \int_v^{v'} x(z)\, dz,$$

which is true due to the monotonicity of x. For $v > v'$ we get

$$(v - v')\, x(v') \leq \int_{v'}^{v} x(z)\, dz,$$

which is again true due to the monotonicity of x.

In the other direction, combining the strategyproofness constraint at v,

$$vx(v) - p(v) \geq vx(v') - p(v'),$$

with the strategyproofness constraint at v',

$$v'x(v) - p(v) \leq v'x(v') - p(v'),$$

and subtracting the inequalities, we get

$$(v' - v)\, x(v) \leq (v' - v)\, x(v')$$

which implies monotonicity of x, since $v > v' \rightarrow x(v) \geq x(v')$ and $v' > v \rightarrow x(v') \geq x(v)$.

To derive the formula for p, we can rearrange the two strategyproofness constraints as

$$v\,(x(v) - x(v')) \leq p(v') - p(v) \leq v'\,(x(v) - x(v')).$$

Now by letting $v' = v + \epsilon$, dividing throughout by ϵ, and taking the limit $\epsilon \rightarrow 0$, we get

$$v\frac{d}{dv}x(v) \leq \frac{d}{dv}p(v) \leq v\frac{d}{dv}x(v).$$

Both sides approach the same value, $v\frac{d}{dv}x(v)$, and we get $\frac{d}{dv}p(v) = v\frac{d}{dv}x(v)$. Thus, taking into account the normalization condition $p(0) = 0$, we have that

$$p(v) = \int_0^v z x'(z)\, dz,$$

and integrating by parts completes the proof that

$$p(v) = vx(v) - \int_0^v x(z)\, dz.$$

(This seems to require the differentiability of x, but as x is monotone this holds almost everywhere, which suffices since we immediately integrate). □

Given this theorem, we see that once an allocation rule $\mathbf{x}(\cdot)$ is fixed, the payment rule $\mathbf{p}(\cdot)$ is also fixed. Thus, in specifying a mechanism we need specify only a monotone allocation rule and from it the truth-inducing payment rule can be derived.

It is useful to specialize thm:A-mechanism-is to the case where the mechanism is deterministic. In this case, the monotonicity of $x_i(b_i)$ implies that, for $\mathbf{b}_{-i}$ fixed, there is some threshold bid t_i such that $x_i(b_i) = 1$ for all $b_i > t_i$ and 0 for all $t_i < b_i$. Moreover the second part of the theorem then implies that for any $b_i > t_i$,

$$p_i(b_i) = b_i - \int_{t_i}^{b_i} dz = b_i - (b_i - t_i) = t_i.$$

We conclude the following.

Corollary 13.3. *Any deterministic truthful auction is specified by a set of functions $t_i(\mathbf{b}_{-i})$ which determine, for each bidder i and each set of bids $\mathbf{b}_{-i}$, an offer price to bidder i such that bidder i wins and pays price t_i if $b_i > t_i$, or loses and pays nothing if $b_i < t_i$. (Ties can be broken arbitrarily.)*

13.2 Bayesian Optimal Mechanism Design

In this section we describe the conventional economics approach of *Bayesian optimal mechanism design* where is assumed that the valuations of the agents are drawn from a known distribution. The mechanism we describe is known as the MVS mechanism: it is the truthful mechanism that maximizes the auctioneer's expected profit, where the expectation is taken over the randomness in the agents' valuations.

Consider, for example, a single-item auction with two bidders whose valuations are known to be drawn independently at random from the uniform distribution on $[0, 1]$. In Lecture 10, it was shown that the expected revenue of both the Vickrey (second-price) auction and of the first price auction is $\frac{1}{3}$. In fact, it was observed that any auction that always allocates the item to the bidder with the higher valuations achieves the same expected revenue.

It was also shown that the expected profit of $VA_{\frac{1}{2}}$, the Vickrey auction with reservation price $r = \frac{1}{2}$, is $\frac{5}{12}$. Thus, it is possible to get higher expected profit than the Vickrey auction by sometimes not allocating the item! This raises the problem of identifying, among the class of all truthful auctions, the auction that gives the optimal profit in expectation. The derivation in the next section answers this question and shows that in fact for this scenario $VA_{\frac{1}{2}}$ is the optimal auction.

13.2.1 Virtual Valuations, Virtual Surplus, and Expected Profit

We assume that the valuations of the agents, $v_{1,\ldots,}v_n$, are drawn independently at random from known (but not necessarily identical) continuous probability distributions $D_1, \ldots, D_n$. For simplicity, we assume that $v_i \in [0, h]$ for all i. We denote by F_i the distribution function from which bidder i's valuation, v_i, is drawn (i.e., $F_i(z) = \Pr[v_i \le z]$) and by f_i its density function (i.e., $f_i(z) = \frac{d}{dz} F_i(z)$). Since the agents' valuations are independent, the joint distribution from which $\mathbf{v}$ is drawn is just the product distribution $\mathbf{F} = F_1 \times \cdots \times F_n$.

We now define two key notions: *virtual valuations* and *virtual surplus*.

Definition 13.4. The *virtual valuation* of agent i with valuation v_i is

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}.$$

In particular, for $D_i = U[0, 1]$,

$$\phi_i(v_i) = v_i - \frac{1 - v_i}{1} = 2v_i - 1.$$

Definition 13.5. Given valuations, v_i, and corresponding virtual valuations, $\phi_i(v_i)$, the *virtual surplus* of allocation $\mathbf{x}$ is $\sum_i \phi_i(v_i) x_i(\mathbf{v})$.

As the *surplus* of an allocation is $\sum_i v_i x_i$, the virtual surplus of an allocation is the surplus of the allocation with respect to agents whose valuations are replaced by their virtual valuations, $\phi_i(v_i)$.

We now show that any truthful mechanism has expected profit equal to its expected virtual surplus. Thus, to maximize expected profit, the mechanism should choose an allocation which maximizes virtual surplus. In so far as this allocation rule is monotone, this gives the optimal truthful mechanism!

Theorem 13.6. *The expected profit of any truthful mechanism, $\mathcal{M}$, is equal to its expected virtual surplus, i.e.,*

$$\mathbf{E}_{\mathbf{v}}[\mathcal{M}(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}\left[\sum_i \phi_i(v_i) x_i(\mathbf{v})\right].$$

Thus, if the mechanism on each bid vector $\mathbf{b}$, chooses an allocation, $\mathbf{x}$, which maximizes $\sum_i \phi_i(v_i) x_i(\mathbf{v})$, the auctioneer's profit will be maximized. Notice that if we employ a deterministic tie-breaking rule then the resulting mechanism will be deterministic. thm:The-expected-profit follows from lem:Consider-any-truthful below, and the independence of the agents' valuations.

Lemma 13.7. *Consider any truthful mechanism and fix the bids $\mathbf{b}_{-i}$ of all bidders except for bidder i. The expected payment of a bidder i satisfies:*

$$\mathbf{E}_{b_i}[p_i(b_i)] = \mathbf{E}_{b_i}[\phi_i(v_i) x_i(b_i)].$$

Proof. To simplify notation, we drop the subscript i and refer simply to the bid b being randomly chosen from distribution F with density function f.

By thm:A-mechanism-is, we have

$$\mathbf{E}_b\left[p\left(b\right)\right] = \int_{b=0}^{h} p\left(b\right) f\left(b\right) db = \int_{b=0}^{h} \left(bx\left(b\right) - \int_{z=0}^{b} x\left(z\right) dz\right) f\left(b\right) db$$

$$= \int_{b=0}^{h} bx\left(b\right) f\left(b\right) db - \int_{b=0}^{h} \int_{z=0}^{b} x(z) f\left(b\right) dz db.$$

Focusing on the second term and switching the order of integration, we have

$$\mathbf{E}_b\left[p\left(b\right)\right] = \int_{b=0}^{h} bx\left(b\right) f\left(b\right) db - \int_{z=0}^{h} x\left(z\right) \int_{b=z}^{h} f\left(b\right) db dz$$

$$= \int_{b=0}^{h} bx\left(b\right) f\left(b\right) db - \int_{z=0}^{h} x\left(z\right) \left[1 - F\left(z\right)\right] dz.$$

Now we rename z to b and factor out $x\left(b\right) f\left(b\right)$ to get

$$\mathbf{E}_b\left[p\left(b\right)\right] = \int_{b=0}^{h} bx\left(b\right) f\left(b\right) db - \int_{b=0}^{h} x\left(b\right) \left[1 - F\left(b\right)\right] db$$

$$= \int_{b=0}^{h} \left[b - \frac{1 - F\left(b\right)}{f\left(b\right)}\right] x\left(b\right) f\left(b\right) db$$

$$= \mathbf{E}_b\left[\phi\left(b\right) x\left(b\right)\right].$$

□

13.2.2 Truthfulness of Virtual Surplus Maximization

Of course, it is not immediately clear that maximizing virtual surplus results in a truthful mechanism. By thm:A-mechanism-is this depends on whether or not virtual surplus maximization results in a monotone allocation rule. Recall that the VCG mechanism, which maximizes the actual surplus, i.e., $\sum_i v_i x_i\left(\mathbf{v}\right)$, is truthful precisely because surplus maximization results in a monotone allocation rule. Clearly then, virtual surplus maximization gives an allocation that is monotone in agent valuations precisely when virtual valuation functions are monotone in agent valuations. Indeed, it is easy to find examples of the converse which show that nonmonotone virtual valuations result in a nonmonotone allocation rule. Thus, we conclude the following lemma.

Lemma 13.8. *Virtual surplus maximization is truthful if and only if, for all i, $\phi_i(v_i)$ is monotone nondecreasing in v_i.*

A sufficient condition for monotone virtual valuations is implied by the *monotone hazard rate* assumption. The hazard rate of a distribution is defined as $hr(z) = \frac{f(z)}{(1-F(z))}$. Clearly, if the hazard rate is monotone nondecreasing, then the virtual valuations are monotone nondecreasing as well, since $\phi_i(v_i) = v_i - \frac{1}{hr(v_i)}$. There is a technical construction that extrends these results to the nonmonotone case, but we do not cover it here.

Definition 13.9. Let $\mathbf{F}$ be the prior distribution of agents' valuations satisfying the monotone hazard rate assumption. We denote by $MVS_{\mathbf{F}}(\mathbf{b})$ the *MVS mechanism*: on input $\mathbf{b}$, output $\mathbf{x}$ to maximize the virtual surplus (defined with respect to the distribution $\mathbf{F}$).

Thus, for single parameter problems, profit maximization in a Bayesian setting reduces to virtual surplus maximization. This allows us to describe the MVS mechanism, $MVS_{\mathbf{F}}(\mathbf{b})$, as follows:

(i) Given the bids $\mathbf{b}$ and $\mathbf{F}$, compute "virtual bids": $b_i' = \phi_i(b_i)$.

(ii) Run VCG on the virtual bids $\mathbf{b}'$ to get $\mathbf{x}'$ and $\mathbf{p}'$

(iii) Output $\mathbf{x} = \mathbf{x}'$ and $\mathbf{p}$ with $p_i = \phi_i^{-1}\left(p_i'\right)$.

13.2.3 Applications of the MVS Mechanism

The formulation of virtual valuations and the statement that the optimal mechanism is the one that maximizes virtual surplus is not the end of the story. In many relevant cases this formulation allows one to derive very simple descriptions of the optimal mechanism. We now consider a couple of examples to obtain a more precise understanding of $MVS_{\mathbf{F}}(\mathbf{b})$ and illustrate this point.

Example 13.10. (single-item auction) In a single-item auction, the surplus maximizing allocation gives the item to the bidder with the highest valuation, unless the highest valuation is less than 0 in which case the auctioneer keeps the item. Usually, we assume that all bidders' valuations are at least zero, or they would not want to participate in the auction, so the auctioneer never keeps the item.

However, when we maximize virtual surplus, it may be the case that a bidder has positive valuation but negative virtual valuation. Thus, for allocating a single item, the optimal