

Lecture 14

Lecturer: Aaron Roth

Scribe: Aaron Roth

Auction Design

Last lecture we studied *pricing equilibria*. In this lecture, we continue our study of money as a means of exchange, from the perspective of mechanism design. Specifically, we begin our study of how to design *auctions*, which will be mechanisms for choosing outcomes, while managing the incentives of individuals to report to the mechanism their true preferences.

We will consider a very general setting:

1. We have a set of possible *alternatives* A that we want to choose from.
2. We have a set of n agents i each of whom have a valuation function $v_i \in V$. Each valuation function $v_i : A \rightarrow \mathbb{R}_{\geq 0}$.
3. An *outcome* $o = (a, p)$ denotes an alternative $a \in A$ together with a payment vector $p = (p_1, \dots, p_n) \in \mathbb{R}^n$ specifying a payment p_i for each agent.
4. Agents have quasilinear utility functions. The utility that agent i experiences for outcome $o = (a, p)$ is:

$$u_i(o) = v_i(a) - p_i$$

For example, this could model an allocation problem – we could have some set of goods, and the alternative a could represent a feasible allocation of the goods. Alternatively, it could model a public goods problem – a city could be choosing whether or not to build a library (which everyone gets to enjoy if it is built), together with how to fund it.

A *mechanism* is a method of mapping agent's reported valuations to an outcome:

Definition 1 A *mechanism* is a pair of functions:

1. A choice rule $X : V^n \rightarrow A$
2. A payment rule $P : V^n \rightarrow \mathbb{R}^n$

Any choice of these two functions yields some mechanism or auction. Lets lay out a “wish list” of desiderata that our dream auction would satisfy:

First, at a minimum, we would like the auction to be safe to participate in – nobody should ever end up with negative utility. Otherwise we will find that we have no takers:

Definition 2 (Individual Rationality) A *mechanism* is *individually rational (IR)* if for every agent i and for every $v \in V^n$:

$$v_i(X(v)) \geq P(v)_i$$

i.e. nobody is ever asked to pay more than their (reported) value for the outcome.

Second, if we want to have any idea what our auction rule is doing over the *real* valuation functions as opposed to the *reported* valuation functions, we would like that the agents are incentivized to report their true valuations:

Definition 3 (Dominant Strategy Truthfulness) A *mechanism* is *dominant strategy truthful* if for every agent i , for every $v \in V^n$, and for every alternative report $v'_i \in V$, we have:

$$u_i(X(v), P(v)) \geq u_i(X(v'_i, v_{-i}), P(v'_i, v_{-i}))$$

or equivalently:

$$v_i(X(v)) - P(v)_i \geq v_i(X(v'_i, v_{-i})) - P(v'_i, v_{-i})_i$$

Third, we would actually like the mechanism to compute a high quality outcome!

Definition 4 (Allocative Efficiency) *A mechanism is allocatively efficient, or “Social Welfare Maximizing”, if for all $v \in V^n$, if $a = X(v)$, then for all $a' \in A$ we have:*

$$\sum_i v_i(a) \geq \sum_i v_i(a')$$

Finally – we’re not saints – we want to achieve all of this without the mechanism itself having to lose money.

Definition 5 (No Deficit) *A mechanism is no deficit if for all $v \in V^n$:*

$$\sum_i P(v)_i \geq 0$$

i.e. in total, the mechanism does not have to pay to run the auction.

We will start by illustrating some of these issues with a simple example – a single item auction.

Here we have $A = [n]$ (representing which of the n agents get the single item for sale). Valuations are single dimensional. We will abuse notation by writing $V = \mathbb{R}_{\geq 0}$, but these valuations will really be functions of the form:

$$v_i(a) = \begin{cases} v_i, & a = i; \\ 0, & \text{otherwise.} \end{cases}$$

where $v_i \in \mathbb{R}_{\geq 0}$ is agent i ’s value for the item for sale.

So – can we satisfy all of our desiderata?

First, we must design our allocation rule. If it is to be allocatively efficient, our hands are tied! We must choose $X(v) = \arg \max_i v_i$. What about the payment rule? Lets think how we are constrained.

1. By individual rationality, we must have $p(v)_j \leq 0$ for all $j \neq X(v)$. Lets try $p(v)_j = 0$, so it only remains to fix $p(v)_i$ for $i = X(v)$. Similarly, we know $p(v)_i \leq v_i$.
2. We could try $p(v)_i = v_i$. Does this lead to an incentive compatible auction? Why not?
3. What about $p(v)_i = \arg \max_{j \neq X(v)} v_j$. Is this incentive compatible? Yes? Why? (Informally: Consider $i = X(v)$. Raising his bid does not change his payment or his allocation, so he has no incentive to do it. Lowering his bid does not change his payment or his allocation until he lowers $v'_i < v_j$, at which point he goes from winning to losing, at a price he would have been willing to pay – so he also has no incentive to do this. For $j \neq X(v)$, lowering his bid does not change his payment or allocation. Raising his bid doesn’t either until he raises it to $v'_j > v_i$, at which point he goes to winning, but at a price $v_i > v_j$ which he would not want to pay...)

Observe that this “second price” is also no deficit, since it only asks for non-negative payments, so it satisfies all of our desiderata, at least in this simple setting. This is called the “Vickrey auction”. Note that it results in the same allocation and payment as the “Ascending price” or “English” auction you may have seen on TV.

What about other pricing rules? What if the winner pays the 3rd highest price?

Lets see if we can generalize this beyond single item auctions...

Definition 6 *The Groves Mechanism has choice rule:*

$$X(v) = \arg \max_{a \in A} \sum_i v_i(a)$$

and payment rule:

$$P(v)_i = h_i(v_{-i}) - \sum_{j \neq i} v_j(a^*)$$

where h_i is an arbitrary function (crucially, independent of v_i), and $a^* = X(v)$ is the socially optimal outcome.

We note that the Groves mechanism is really a family of mechanisms, instantiated by a choice of h_i . This can be anything – even $h_i \equiv 0$ is a valid choice.

We start by observing that the Groves mechanism satisfies at least two of our desiderata:

Theorem 7 *The Groves mechanism is dominant strategy incentive compatible and Allocatively efficient.*

Proof It is allocatively efficient by definition, so it remains to verify that it is dominant strategy incentive compatible.

Fix any agent i , and reports v_{-i} of the other players. We have:

$$u_i(X(v), P(v)) = v_i(a^*) + \sum_{j \neq i} v_j(a^*) - h_i(v_{-i})$$

where $a^* = \arg \max_{a \in A} \left(\sum_{j \neq i} v_j(a) + v'_i(a) \right)$. Agent i wishes to report v'_i to maximize his utility. Note that $h_i(v_{-i})$ has no dependence on his report, so equivalently, agent i wishes to report v'_i to maximize:

$$v_i(a^*) + \sum_{j \neq i} v_j(a^*) = \sum_i v_i(a^*)$$

But note that if agent i truthfully reports $v'_i = v_i$, then a^* maximizes this quantity by definition. Hence, it is a dominant strategy for all agents to report truthfully. ■

The intuition here is that the payment scheme of the Groves mechanism aligns the incentives of the agents and the mechanism designer: both prefer higher social welfare outcomes.

Lets consider an example, instantiating the Groves mechanism in a single item auction setting (will, recall $A = [n]$). Lets take $h_i(v_{-i}) = 0$ for all i . Suppose we have two bidders, with values for the item $v_1 = 5$ and $v_2 = 8$. Truthful bidding results in $X(v) = 2$, resulting in social welfare 8. The payment rule mandates:

$$P(v)_1 = -8 \quad P(v)_2 = 0$$

Both bidders get utility 8 (exactly equal to the social welfare), and have no beneficial deviations. Note however that the auction is *not* no-deficit, because it pays the losing bidder \$8! Note however that the mechanism is trivially individually rational – nobody can ever be required to make a positive payment...

We get truthfulness no matter how we pick the functions h_i . The question is whether we can make a clever choice of h_i to achieve the no-deficit property, without breaking individual rationality! (Note it would be easy to break individual rationality with a bad choice of h_i ...) This is what the VCG mechanism does:

Definition 8 (The Vickrey-Clarke-Groves (VCG) Mechanism) *The VCG mechanism is an instantiation of the Groves mechanism with*

$$h_i(v_{-i}) = \sum_{j \neq i} v_j(a_{-i}^*)$$

where $a_{-i}^* = \arg \max_{a \in A} \sum_{j \neq i} v_j(a)$ is the alternative that maximizes social welfare among all agents other than agent i . In other words, the VCG mechanism has payment rule:

$$P(v)_i = \sum_{j \neq i} v_j(a_{-i}^*) - \sum_{j \neq i} v_j(a^*)$$

The idea behind VCG payments is that every agent i is charged the “negative externality” that he imposes on the market – i.e. the difference between the welfare of everyone else when i is in the market, compared to if he were not. We will show that the VCG mechanism satisfies all of our desiderata.

Theorem 9 *The VCG mechanism is allocatively efficient and dominant strategy incentive compatible.*

Proof It is an instantiation of the Groves mechanism. ■

Theorem 10 *The VCG mechanism is individually rational.*

Proof We need to show that Agent i ’s utility satisfies:

$$u_i(o) = v_i(a^*) + \sum_{j \neq i} v_j(a^*) - \sum_{j \neq i} v_j(a_{-i}^*) \geq 0$$

Or equivalently:

$$\sum_i v_i(a^*) \geq \sum_{j \neq i} v_j(a_{-i}^*)$$

But note that if this is not the case, since v_i is non-negative, we would have:

$$\sum_i v_i(a_{-i}^*) \geq \sum_{j \neq i} v_j(a_{-i}^*) > \sum_i v_i(a^*)$$

But this would contradict the allocative efficiency of a^* ! ■

Finally, to complete the picture:

Theorem 11 *The VCG mechanism is no-deficit.*

Proof We will in fact show the stronger claim that for all i , $P(v)_i \geq 0$. Recall that:

$$P(v)_i = \sum_{j \neq i} v_j(a_{-i}^*) - \sum_{j \neq i} v_j(a^*)$$

This is non-negative whenever:

$$\sum_{j \neq i} v_j(a_{-i}^*) \geq \sum_{j \neq i} v_j(a^*)$$

But note that this is always the case, since a_{-i}^* is explicitly defined to be the maximizer of $\sum_{j \neq i} v_j(a)$ over all $a \in A$. ■

So the VCG mechanism satisfies all of our wildest dreams, in an extremely general setting! Perhaps we can end the class here?

Not quite – we will see that the VCG mechanism still leaves a bit to be desired. It doesn’t maximize other objectives (like e.g. revenue), and it isn’t always computationally efficient.