**UPenn NETS 412: Algorithmic Game Theory** April 10 2018

Lecture 17

*Lecturer: Bo Waggoner* *Scribe: Aaron Roth, Bo Waggoner*

## Auction Design

Now that we have defined price of stability and price of anarchy, we can use them to analyze and design mechanisms involving agents collectively allocating items or making decisions. 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.

This will in some sense follow up on our previous discussions of truthful exchange of items (Lecture 10) and stable matching (Lecture 11). A key difference here is that we will be considering mechanisms involving *money* and will be aiming for high social welfare, price of stability, and/or price of anarchy.

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 each of whom has a *valuation function* $v_i \in V$. Each valuation function $v_i : A \to \mathbb{R}_{\geq 0}$.
3. An *outcome* $o = (a, p)$ denotes an alternative $a \in A$ together with a payment vector $p = (p_1, \ldots, p_n) \in \mathbb{R}^n$ 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 $X : V^n \to A$
2. *A* payment rule $P : V^n \to \mathbb{R}^n$

Any choice of these two functions yields some mechanism or auction.

**Goal: total utility.** For now, our main goal is designing the auction will be to obtain high social welfare. Recall this was defined in the previous lecture as the sum of all player utilities. However, in auction design, we also consider the utility of the auctioneer, which we define to be equal to the total payment. So the total welfare for outcome $o = (a, p)$ is defined to be

$$\left( \sum_{i=1}^{n} \underbrace{v_i(a) - p_i}_{i\text{'s utility}} \right) + \left( \underbrace{\sum_{i=1}^{n} p_i}_{\text{auctioneer's utility}} \right)$$
$$= \sum_{i=1}^{n} v_i(a).$$

In other words, the payments don't matter to overall utility because they're just money changing hands, and all agents are assumed to have quasilinear utility, so one agent losing $p_i$ is exactly balanced out by the auctioneer gaining $p_i$. (One can debate if this model is always reasonable.) But the choice of alternative $a$ does matter.

Therefore, the optimal welfare is

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

We would like to design mechanisms with a high price of stability and/or price of anarchy. We are using utilities, so e.g. the price of stability of a mechanism will be the ratio of its total welfare to OPT; this ratio will be between zero and one with higher being better.

**Truthfulness.** Our main approach will be to design auctions that incentivize the bidders to report their valuation functions truthfully. Then, we can just compute a high-welfare alternative. As for what we mean by truthfulness, we will adapt two definitions we have seen before to this setting.

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

$$v_i(X(v)) - P(v)_i \geq 0$$

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

Note here that $P(v)$ is the payment vector when agents report $v$, and $P(v)_i$ is the $i$th component of this vector: the payment that agent $i$ makes.

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** *A mechanism is* ***dominant strategy incentive compatible (DSIC)*** *if for every agent $i$, for every $v \in V^n$, and for every alternative report $\hat{v}_i \in V$, we have:*

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

*or equivalently:*

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

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

**Definition 4** *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 can view this as stating that individual rationality must also hold for the auctioneer.

**Example: single-item auction.** Here we have $A = 1, \ldots, n$, representing which of the $n$ agents gets the single item. Valuations are simple: each agent has a number in mind, and if their valuation is that number if they get the item and zero otherwise. So we will abuse notation by saying that $v_i \in \mathbb{R}$ is their value for the object. But to match the general formalism, one would define:

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

To maximize social welfare truthfully, let's consider the **second price auction**. It works like this:

1. Each bidder $i$ submits a bid $\hat{v}_i \geq 0$. If they are acting truthfully, then $\hat{v}_i = v_i$, but of course they might lie as far as we know.
2. The item is awarded to the highest bidder, call her $i^* = \arg\max_i v_i$.
3. The bidder $i^*$ pays the second-highest bid, i.e. $\max_{i \neq i^*} v_i$.
4. All others get nothing and pay nothing.

**Theorem 5** *The second-price auction is DSIC, IR, and no deficit.*

**Proof** First, the mechanism is no deficit because the winner pays the second-highest bid, and we only allow nonnegative bids, so the mechanism makes at least 0 dollars.

Next, we show dominant strategy incentive compatibility. Consider any bidder $i$ and any reports $v_{-i}$ of the other bidders. (It does not matter if these other reports are truthful or not; we need to show that it is a best response for $i$ to be truthful no matter what.) Let $w = \max_{j \neq i} v_j$ be the highest bid among the other bidders. Then if $i$ bids $\hat{v}_i \geq w$, then $i$ wins and pays $w$, so her utility is

$$v_i - w.$$

If $i$ bids $\hat{v}_i < w$, then $i$ loses and gets nothing, pays nothing, so has utility 0.

So if $v_i \geq w$, then any bid greater than $w$ is a best response because it gets utility $v_i - w \geq 0$, whereas any bid smaller than $w$ gets 0. But if $v_i \leq w$, then any bid smaller than $w$ is a best response because it gets utility 0, whereas any bid greater than $w$ gets $v_i - w \leq 0$.

In both of these cases, bidding $\hat{v}_i = v_i$ is a best response. So it follows that bidding $v_i$ truthfully is a dominant strategy. Furthermore, it is actually individually rational, as $i$ always gets nonnegative utility in the above argument. ■

**Theorem 6** *The price of stability of the second-price auction is always* 1*, but the price of anarchy cannot be bounded above* 0*.*

**Proof** To show the PoS is 1, we must show that there always exists an equilibrium where the welfare is equal to OPT. Here,

$$\begin{aligned} \text{OPT} &= \max_{a \in A} \sum_{i=1}^{n} v_i(a) \\ &= \max_i v_i, \end{aligned}$$

that is, the max achievable social welfare is just the highest valuation of any player.

Now, if every player plays her dominant strategy of reporting truthfully, this must be an equilibrium (in fact, a pure dominant-strategy equilibrium!). In this case, the player with the highest valuation is the one who wins the auction, so the welfare is equal to OPT.

To show the PoA cannot be bounded above 0, we give examples of low-welfare equilibria. Suppose Alex has value 1 for the item while Boris has value 100. But Alex bids 101 while Boris bids 0. This is actually an equilibrium: check that neither player can improve utility by deviating to any other bid! But in this equilibrium, social welfare is 1 because Alex wins the item, while OPT = 100 by giving it to Boris. So the PoA of this example is $\frac{1}{100}$. By changing the 100 to 1000 and 101 to 1001, we get the same result but with a PoA of $\frac{1}{1000}$. Continuing in this way can give a PoA arbitrarily close to zero. ■

**VCG.** Lets see if we can generalize these ideas beyond single item auctions.

**Definition 7** *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 some of our desiderata.

**Theorem 8** *The Groves mechanism is DSIC and has price of stability* 1.

**Proof** We first show DSIC . Fix any agent $i$, and reports $v_{-i}$ of the other players. Let $\hat{v}_i$ be the report of player $i$, and write $v = (v_i, v_{-i})$ for the reported profile when $i$ is truthful and $\hat{v} = (\hat{v}_i, v_{-i})$ for when $i$ reports $\hat{v}_i$. We have

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

where $a^*$ is the choice made by the mechanism, i.e.

$$a^* = X(\hat{v}) = \arg\max_{a \in A}\ \hat{v}_i(a) + \sum_{j \neq i} v_j(a).$$

Agent $i$ wishes to report $\hat{v}_i$ to maximize his utility. Note that $h_i(v_{-i})$ has no dependence on his report, so equivalently, agent $i$ wishes to report $\hat{v}_i$ 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 $\hat{v}_i = v_i$, then $a^*$ maximizes this quantity by definition! Hence, it is a dominant strategy for all agents to report truthfully.

Now, we can easily show the price of stability is 1. There is an equilibrium where all players report truthfully (as we just showed is a dominant strategy). In this equilibrium, the choice rule $X$ by definition picks the highest social welfare outcome, so the welfare is equal to OPT. ■

The intuition here is that the payment scheme of the Groves mechanism aligns the incentives of the agents and the mechanism designer. Every agent's utility ends up being the social welfare, minus some function $h_i$ they can't control. If they report truthfully, this will enable to mechanism designer to choose the alternative $a^*$ with highest social welfare, which is best for them.

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 $X(v) = 2$ (bidder 2 gets the item), 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! On the other hand, it is trivially individually rational – nobody can ever be required to make a positive payment...

We showed above that we get truthfulness (DSIC) 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, as we will see next time.