Incentives and Machine Learning (CSC 482A/581A)
Lectures 6–8

Nishant Mehta

1 Auctions

In an auction, the party running the auction — the seller — has some goods to be divided among the agents. These goods could consist of a single item or multiple items. In general, some of the goods might be retained by the auction, but we avoid such alternatives, which could be viewed as a self-allocation

Let us refer to each way of allocating the goods among the agents as an *alternative*. Let A be the set of alternatives.

For example, in a single-item auction, there is a single, non-divisible item to be given to a single agent (imagine auctioning an iPhone). The set of alternatives is $A = [n] = \{1, 2, \ldots, n\}$, where if $a = i$, then agent i gets the item.

Another example is selling a collection of fruit. Suppose that there are two agents and the goods consist of a mango, a passion fruit, and a crabapple; however, to ensure no one is stuck with just the crabapple, the crabapple can only be sold together with either the mango or the passion fruit. The constraint related to the crabapple induces some combinatorial structure. The number of alternatives is the same as the number of ways to gives goods to the first agent (the second agent receives the complement), of which there are $8 - 2 = 6$ since the first agent cannot receive only the crabapple and the first agent cannot receive only the mango and passion fruit.

Each agent i has a *valuation function* $v_i \colon A \to \mathbb{R}$ which specifies, for each alternative $a \in A$, the agent's valuation $v_i(a)$ of the item. It will be convenient to specify a set of valuation functions V, so that $v_i \in V$.[1]

The mechanism is composed of:

- an *allocation rule*: $X \colon V^n \to A$;
- a *payment rule* $p \colon V^n \to \mathbb{R}^n$.

We make no assumption that agents truthfully report their valuation functions. Therefore, each agent i gives a *reported valuation*, or *bid*, $b_i \in V$, resulting in a *bid profile* $b \in V^n$. A given bid profile induces an allocation $X(b) \in A$ and payment vector $p(b)$, where $p_i(b)$ is the amount that agent i pays.

We adopt the assumption that each agent has a *quasilinear utility function*:

$$u_i(b) = v_i(X(b)) - p_i(b)$$

It is worth reflecting on what this type of assumption rules out. First, the amount that agent pays for an alternative cannot affect their valuation of that alternative. Assuming that utility is quasilinear is a natural starting point. Even so, can you think of realistic scenarios where an agent's

[1] For our purposes, it is fine that we use the same set V for all agents.

utility function might not be quasilinear? A not-so-realistic example would be if an agent has 2 million dollars and is bidding on the right to gain access to a market to buy small islands. If the cheapest small island costs a million dollars, the agent might not want to spend (say) 1.5 million dollars to buy the right to enter the market!

Let's return to the example of single-item auctions. There is a natural allocation rule: give the item to highest bidder, so $X(b) = \arg\max_{i \in [n]} b_i$, with ties broken arbitrarily.

Having decided the allocation rule, it remains to specify the payment rule. It also seems natural to charge nothing to anyone who doesn't receive the item. What should we charge the highest bidder? Let's consider a few options:

1. Charge the highest bidder b_i (the highest bid). This is a first-price auction and poses many difficulties. If the agent is truthful, then when they win they get 0 utility. So, the agent has no incentive to be truthful and is in fact incentivized to underbid!

2. Charge the winning bidder the bid of the second highest bidder. This is a second-price auction and turns out to be truthful, meaning that each agent is incentivized to truthfully report their valuation.

2 Second-price auction

The first mechanism that we analyze is a second-price auction. Does agent i have a dominant strategy? Yes! We will see that agent i's dominant strategy is to truthfully bid its valuation.

Lemma 1. *In a second-price auction, for any player i, truthful reporting ($b_i = v_i$) is a dominant strategy.*

Proof. We will show that $b_i = v_i$ is a dominant strategy for player i. Let $b^* = \max_{j \in [n] \setminus \{i\}} b_j$ be the highest bid among the other players. We consider two cases: either $v_i \geq b^*$ or $v_i < b^*$.

Case 1: $v_i \geq b^*$.

Bidding $b_i \geq b^*$ results in utility $v_i - b^* \geq 0$, whereas bidding $b_i < b^*$ results in utility 0. Hence, for all b_{-i} satisfying $v_i \geq b^*$, any bid $b_i \geq b^*$ is a best response. In particular, the truthful bid $b_i = v_i$ is a best response.

Case 2: $v_i < b^*$.

Bidding $b_i < b^*$ results in utility 0, whereas bidding $v_i \geq b^*$ results in utility $v_i - b^* < 0$. So, for all b_{-i} satisfying $b_i < b^*$, any bid $b_i < b^*$ is a best response. Once again, the truthful bid $b_i = v_i$ is one particular choice for a best response.

Since the cases together cover all b_{-i}, it follows that $b_i = v_i$ is a dominant strategy. □

The previous lemma shows that truthful reporting is a dominant strategy for every player i. It therefore follows that there is a dominant strategy equilibrium, and, moreover, this equilibrium is a *truthful dominant strategy equilibrium.* The next definition is common way of referring to mechanisms that have such equilibria.

Definition 1. If each agent has a dominant strategy that is truthful, then we say there is a truthful dominant strategy equilibrium. A mechanism that has a truthful dominant strategy equilibrium is said to be *dominant strategy incentive compatible (DSIC).*

The next corollary is immediate.

Corollary 1. *Second-price auctions are DSIC.*

3 What makes a good auction?

In a certain sense, auctions which are DSIC are nice from the perspective of the party running the auction: the bidders have no reason to lie, and so it is plausible that bidders reveal their true valuations. Also, if an auction is DSIC, it is straightforward to maximize social welfare.

Definition 2. The *social welfare* of an allocation $a \in A$ is

$$W(a) = \sum_{i=1}^{n} v_i(a).$$

It will sometimes be convenient to refer to social welfare in the counterfactual world where the bidders' bids b are equal to their true valuations. Let us refer to this social welfare as *social welfare given valuations* b, denoted mathematically as $W(a \mid b) = \sum_{i=1}^{n} b_i(a)$. In the very special case of truthful bids, we have $W(a \mid v) = W(a)$.

Definition 3. An auction is *welfare-maximizing* if it selects an allocation a^* that maximizes social welfare:

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

If an auction is DSIC, then bidders' bids should be equal to their valuation functions. Consequently, the allocation rule can directly select a^* as above to maximize social welfare.

In mechanism design, including auctions, it is important to consider whether the various actors involved with the mechanism are incentivized to participate in the mechanism. In auctions, these actors consist of the agents (the bidders) as well as the mechanism itself (the seller).

We first consider the bidders. From the perspective of the bidders, a very important property is that a mechanism is individually rational.

Definition 4. An auction is *(ex-post) individually rational (IR)* if, for every bidder i, every $v_i \in V$, and every $b_{-i} \in V^{n-1}$, it holds that

$$u_i(v_i, b_{-i}) \geq 0.$$

The term "ex-post" more or less means "no matter what": as long as agent i reports truthfully, the agent can be assured to never have negative utility (no matter what the other bidders bid). There are weaker forms of IR as well but these are beyond our scope for now.[2] If an fails to be IR, some agents might not want to participate (which is bad).

Next, let's consider the seller. From the seller's perspective, it is desirable that the auction is *no deficit*.

Definition 5. An auction is *no deficit* if the sum of the payments is nonnegative. That is, for all $b \in V^n$,

$$\sum_{i=1}^{n} p_i(b) \geq 0.$$

[2]For example, if all bidders' valuation functions are drawn from a common prior distribution that is known to all bidders, then a bidder who observes her own valuation function maintains a posterior distribution over the other bidders' valuation functions. In this situation, we might ask for "ex-interim individual rationality: the bidder's expected utility (by taking an expectation with respect to the posterior distribution over the other bidders' valuation functions) is nonnegative, which is does not mean that the bidder's utility is *always* nonnegative.

If an auction fails to be no deficit, then the seller could lose both goods and money! The seller might be better off not participating at all. Of course, there are situations where one wishes to pay for someone to take something from them (imagine if you had nuclear waste), but these situations are rather exceptional.

A very special case of being no deficit is being budget balanced.

Definition 6. An auction is *budget balanced* if the sum of the payments is equal to zero. That is, for all $b \in V^n$,

$$\sum_{i=1}^{n} p_i(b) = 0.$$

We generally will not ask for the auction to be budget balanced. Indeed, budget balance implies that if at least one agent pays a positive amount, then another agent pays a negative amount (meaning, they actually receive money). We likely will see budget balance later in the course when we cover some other topics.

So far, our wish list for auctions is that the mechanism be DSIC, welfare-maximizing, IR, and no deficit. These aren't the only criteria that matter. We live in reality, and it is important to be able to run an auction in a reasonable amount of time. Therefore, we should seek auctions that are computationally efficient. In real-world applications, auctions might need to run in digital marketplaces, and they might need to run very quickly. The number of agents and items (or bundles, in the case of combinatorial structure) could be very large. We might want an auction to run in near-linear time in the size of the input. Even that might not be enough, however. Consider what happens when the set of alternatives is exponentially large; it becomes infeasible for each agent to report their valuation function, and even if they could do so, the size of the valuation function would be exponentially large!

Computational efficiency ultimately is very important. Yet, in this introductory course, we will not give much attention to computational efficiency in the interest of exploring other topics. There is a large literature on trying to approximate computationally expensive mechanisms.

So far, we have seen second-price auctions for the single-item auction setting. It is an easy exercise to verify that, in addition to being DSIC (as we proved), second-price auctions are welfare-maximizing, IR, no deficit, and computationally efficient. It's a good idea to do this exercise now if you haven't already!

4 General auctions

4.1 Groves mechanism

We now shift to thinking about how to design good auctions in general. We would like for our auction to be DSIC, welfare-maximizing, and IR (let's consider computational efficiency as a bonus).[3]

Our strategy for designing auctions will be to first assume that the auction is DSIC, so that bidders are truthful. Operating under this assumption, we select an allocation rule that is welfare-maximizing. Finally, we hope (!) that it's possible to choose a payment rule that ensures that the auction really is DSIC and also that it is IR.

Amazingly, this strategy does work. The key to making it work is the payment rule. We will slowly develop the payment rule, first seeing a generic rule that ensures the auction is DSIC and then refining the rule to make it IR as well.

[3] No deficit is also important; we leave it to the reader to verify no deficit rather than explicitly discussing it here.

Let's first develop some intuition. Suppose that the auction selects an alternative that maximizes social welfare given valuations b, i.e., that maximizes $W(a \mid b)$. An agent i, in deciding what to report, will select a bid that maximizes her utility. Now, if somehow the payment rule p_i could be chosen in such a way that agent i's utility becomes equal to $W(X(b) \mid b)$ except that agent i's true valuation always appears, then agent i's utility and the social welfare (based on the bids) are guaranteed to be equal if agent i's bid is truthful. Put another way, if agent i is truthful, then the mechanism maximizes agent i's utility! There can't be anything better than that for agent i. The question then becomes how we can choose p_i to make agent i's utility as desired.

Let's try to derive a suitable payment rule. We now reason more formally. The question is:

> Does there exist a payment rule p_i such that if bidder i is truthful (so $b_i = v_i$), it holds that $u_i(b) = \sum_{i=1}^{n} b_i(X(b))$?

Such a payment rule would be useful because $X(b)$ maximizes social welfare given valuations b, i.e.,

$$\sum_{i=1}^{n} b_i(X(b)) = \max_{a \in A} \sum_{i=1}^{n} b_i(a).$$

This alignment of the agent's objective and the mechanism's objective is enough to blindly go forward and set the payment rule. We will reason in the proof of the next result why lying can only hurt the agent.

To set the payment rule, observe that under truthfulness of bidder i, we have

$$u_i(b) = v_i(X(b)) - p_i(b) = b_i(X(b)) - p_i(b).$$

So, we simply set $p_i(b) = -\sum_{j \neq i} b_j(X(b))$.

Lemma 2. *Let the mechanism (X, p) be defined by*

$$X(b) \in \arg\max_{a \in A} \sum_{i=1}^{n} b_i(a)$$

and, for all $i \in [n]$,

$$p_i(b) = -\sum_{j \neq i} b_j(X(b)).$$

Then the mechanism is DSIC.

Proof. Consider an arbitrary bidder $i \in [n]$.

By definition,

$$\begin{aligned} u_i(b) = v_i(X(b)) - p_i(b) &= v_i(X(b)) - \left(-\sum_{j \neq i} b_j(X(b)) \right) \\ &= v_i(X(b)) + \sum_{j \neq i} b_j(X(b)). \end{aligned}$$

If $b_i = v_i$, then

$$u_i(b) = v_i(X(b)) + \sum_{j \neq i} b_j(X(b)) = \sum_{j=1}^{n} b_j(X(b)).$$

Note that $\sum_{j=1}^{n} b_j(a)$ is exactly the objective being maximized by the mechanism. Now, consider what happens if $b_i \neq v_i$. Then agent i's utility function is still $u_i(b) = v_i(X(b)) + \sum_{j \neq i} b_j(X(b))$, but the mechanism selects an alternative a^* to maximize the different objective $\sum_{j=1}^{n} b_j(a)$, and this alternative a^* need not maximize the agent's utility. Therefore, truthful reporting is a dominant strategy for bidder i and likewise for all other agents. Hence, the mechanism is DSIC. $\square$

What we have presented so far is a special case of the *Groves mechanism*. An obvious disadvantage of this mechanism is that, assuming bids are nonnegative, all agents payments are nonpositive. Assuming that valuations are nonnegative, the mechanism is certainly IR, but it certainly is *not* no deficit!

To fix this, let's explore an additional degree of freedom present in the payment rule without compromising the DSIC property. Suppose that we adjust the payment rule to be

$$p_i(b) = f_i(b_{-i}) - \sum_{j \neq i} b_j(X(b))$$

for some function $f_i \colon V^{n-1} \to \mathbb{R}$ that only depends on the other bidders' bids. Since bidder i cannot affect $f_i(b_{-i})$ in any way, the resulting mechanism is still DSIC.

Definition 7. The Groves mechanism is the family of mechanisms (X, p) defined by taking, for any choice of functions $f_1, \ldots, f_n$ all mapping V^{n-1} to $\mathbb{R}$,

$$X(b) \in \arg\max_{a \in A} \sum_{i=1}^{n} b_i(a)$$

and, for all $i \in [n]$,

$$p_i(b) = f_i(b_{-i}) - \sum_{j \neq i} b_j(X(b)).$$

The following corollary is immediate.

Corollary 2. *The Groves mechanism is DSIC.*

The question now is whether we can choose $f_i(b_{-i})$ in such a way so that the resulting mechanism is no deficit while still being IR. The answer is yes!

4.2 Vickrey-Clarke-Groves mechanism

We now introduce the Vickrey-Clarke-Groves (VCG) mechanism, an instance of the Groves mechanism.

Definition 8. The Vickrey-Clarke-Groves mechanism (X, p) is defined by taking

$$X(b) \in \arg\max_{a \in A} \sum_{i=1}^{n} b_i(a)$$

and, for all $i \in [n]$,

$$p_i(b) = f_i(b_{-i}) - \sum_{j \neq i} b_j(X(b)),$$

with

$$f_i(b_{-i}) = \max_{a \in A} \sum_{j \neq i} b_j(a).$$

In words, the VCG mechanism instantiates the Groves mechanism by setting $f_i(b_{-i})$ to be equal to the maximum social welfare in a world where agent i does not exist. Let's write the explicit form of agent i's payment rule to gain a better understanding of what it is doing. We have

$$p_i(b) = \max_{a \in A} \sum_{j \neq i} b_j(a) - \sum_{j \neq i} b_j(X(b)),$$

Now, let's denote an arbitrary (in case there are multiple) social welfare-maximizing alternative as $a^* = X(b) \in \arg\max_{a \in A} \sum_{j=1}^{n} b_i(a)$. Similar, considering the world where bidder i does not exist, we denote an arbitrary social welfare-maximizing alternative as $a^*_{j \neq i} \in \arg\max_{a \in A} \sum_{j \neq i} b_j(a)$. We can now rewrite the payment rule as

$$p_i(b) = \sum_{j \neq i} b_j(a^*_{j \neq i}) - \sum_{j \neq i} b_j(a^*).$$

Thus, we charge agent i the *externality* the agent imposes on the other agents, i.e, we charge agent i the amount by which the social welfare of the other agents drops due to agent i's participation.

Theorem 1. *The VCG mechanism is DSIC, welfare-maximizing, IR, and no deficit.*

Proof. The VCG mechanism is both DSIC and welfare-maximizing since it is an instance of the Groves mechanism.

Let's check that the mechanism is IR. Observe that for any agent i, assuming that agent i is truthful (so $b_i = v_i$), we have

$$\begin{aligned} u_i(b) &= v_i(X(b)) + \sum_{j \neq i} b_j(X(b)) - \max_{a \in A} \sum_{j \neq i} b_j(a) \\ &= \sum_{j=1}^{n} b_j(X(b)) - \max_{a \in A} \sum_{j \neq i} b_j(a) \\ &= \max_{a \in A} \sum_{j=1}^{n} b_j(a) - \max_{a \in A} \sum_{j \neq i} b_j(a) \\ &\geq 0. \end{aligned}$$

Finally, the VCG mechanism is no deficit because each agent i's payment

$$p_i(b) = \max_{a \in A} \sum_{j \neq i} b_j(a) - \sum_{j \neq i} b_j(X(b)).$$

is clearly nonnegative. □

The VCG mechanism seems great. It satisfies most of our criteria for good auctions. Unfortunately, however, in many situations, the VCG mechanism is not computationally tractable. For example, in some combinatorial settings, even selecting a^* (a social welfare-maximizing alternative) can be computationally intractable. Moreover, in many settings, we don't realy want to ask bidders to report their entire valuation functions as these could be incredibly large objects. Another issue is that social welfare maximization is not always the goal. In the real world, companies might want to maximize revenue instead. There is yet another issue with VCG in the real world. On the one hand, VCG is DSIC. Another common term for DSIC is *strategyproof*: no individual agent can strategize (act strategically, by lying) to gain utility. Unfortunately, however, the VCG mechanism is not *group strategyproof*, meaning that a group of agents that collude to decide on false bids might obtain higher utility than if they were both truthful. You will explore the concept of group strategyproofness in the first problem set.