Incentives and Machine Learning (CSC 482A/581A)
Lectures 9–11

Nishant Mehta

For *general* mechanism design problems, the VCG mechanism told us how we can design a DSIC mechanism that is welfare-maximizing (as well as individually rational and no-deficit). What more could one ask for? Well, among various issues with VCG, a major one is that the VCG mechanism restricts the allocation rule to a single choice, namely, the rule that maximizes social welfare given bids b. What if someone wants to use a different allocation rule? For example, a seller might not want to sell an item unless a reserve price is met. How can we determine whether a given allocation rule X can be extended — by appropriate choice of the payment rule P — to a DSIC mechanism?[1]

In this lecture, we will give a definitive answer to the previous question for a class of mechanism design problems known as *single-parameter environments*. These environments can instantiate many interesting problems. Let's first introduce this family, see a few examples to convince ourselves that they're worthy of study, and then we'll develop a far-reaching result. This result is Myerson's Lemma. In short, for direct-revelation mechanisms[2], Myerson's Lemma characterizes which allocation rules X can be extended to a mechanism (X, P) that is DSIC.

1 Single-Parameter Environments

In a single-parameter environment:

- Each alternative a is represented as $(a_1, \ldots, a_n) \in \mathbb{R}^n$, where a_i indicates how much stuff agent i gets.
- Each agent i has a single parameter valuation $v_i \in \mathbb{R}$, where v_i indicates agent i's valuation per unit of stuff. Therefore, each agent's bid b_i is in $\mathbb{R}$.

Since alternatives are vectors, any allocation rule X can be expressed in terms of "partial allocation rules" X_i, where $X(b) = (X_1(b), \ldots, X_n(b))$. Any agent i's utility under bid profile $b \in \mathbb{R}^n$ can now be expressed as

$$u_i(b) = v_i \cdot X_i(b) - P_i(b).$$

Example 1. In single-item auctions, each alternative is a binary vector $a \in \{0, 1\}^n$ with at most one 1. Note that for welfare maximization, we have been assuming that there is precisely one 1. More generally, if there is a reserve price for example, it could make sense to not give the item to any bidder.

Example 2. Consider an auction with k identical items where each agent can get at most one item. Then the set of alternatives are binary vectors $a \in \{0, 1\}^n$ with at most k ones.

[1] Note the switch to P instead of p; this will make some later parts of the lecture clearer.

[2] In a direct-revelation mechanism, agents are asked to report their valuation functions. All the auction mechanisms we considered thus far were direct-revelation mechanisms. We will talk about other types of mechanisms a bit later.

Example 3. The previous two examples considered binary alternative vectors, but in general the alternatives need not be binary. Suppose a power company runs an auction for a fixed amount W (for Watts) of electricity to n factories. The electricity is infinitely. The set of alternatives are all vectors $a \in \mathbb{R}_+$ such that $\sum_{i=1}^n a_i \leq W$.

Example 4. Extending the previous example, in a knapsack auction, k identical items are being sold. Each bidder i has a demand d_i, meaning the bidder needs precisely d_i items. The bidder also has a valuation v_i for getting d_i items (meaning, they get a total of v_i units of happiness for receiving d_i items). The set of alternatives are binary vectors $a \in \{0,1\}^n$ such that $\sum_{i=1}^n d_i \cdot a_i \leq k$. A typical goal is to maximize social welfare, which is $\sum_{i=1}^n v_i \cdot a_i$.

Example 5. Imagine there is a disease outbreak and a health organization is trying to decide how to allocate vaccines to n different cities in Canada. Each city will either vaccines or not get vaccines, any city's allocation will be binary. For political reasons, the health organization takes an all-or-nothing approach for each province: either all cities in a province get vaccines, or none of them do. So, any alternative a is a binary vector $a \in \{0,1\}^n$ with grouped structure: all cities in a province must get the same allocation. Finally, the allocation is decided based on how each city (somehow...) reports their valuation for getting vaccines.

2 From agent utilities to proper scoring rules

Let's return to the main question for today's lecture. How can we characterize the set of allocation rules X that can be extended to a DSIC mechanism (X, P)? In answering this question, we turn to proper scoring rules. As we now establish, for fixed bids b_{-i} of the other agents, agent i's utility function can be viewed as a scoring rule. This connection will allow us to give structure to those utility functions for which truthful reporting of agent i's valuation v_i is a dominant strategy.

2.1 Re-expression of single parameter proper scoring rules

Consider a binary outcome space $\mathcal{Y} = \{0,1\}$, so that $\Delta(\mathcal{Y}) = [0,1]$. Let $S\colon [0,1] \times \{0,1\} \to \mathbb{R}$ be a scoring rule. Observe that S is fully determined by defining the "partial scoring rules" $S_1(r) = S(r,1)$ and $S_0(r) = S(r,0)$. As usual, we extend the second argument of S (the outcome) to $[0,1]$ by writing

$$S(r,q) = \mathsf{E}_{Y \sim q}\left[S(r,Y)\right] = qS(r,1) + (1-q)S(r,0) = qS_1(r) + (1-q)S_0(r).$$

To make a connection to an agent's quasi-linear utility function in a single-parameter environment, we re-express the scoring rule as

$$S(r,q) = q(S_1(r) - S_0(r)) + S_0(r).$$

When viewing a scoring rule as a payment rule (in the information elicitation setting, wherein the agent *receives* the payment), we may interpret the scoring rule as always paying the agent $S_0(r)$ and then correcting the payment by the difference of partial scores $D(r) := S_1(r) - S_0(r)$ if outcome 1 happens. Using D, we finally re-express S as

$$S(r,q) = qD(r) + S_0(r).$$

Let us compare the above form to agent i's quasi-linear utility function when all other agent's bids are fixed to b_{-i}. With some abuse of notation, we write the agent's allocation as $x_i(b_i) =$

$X_i(b_i, b_{-i})$, the agent's payment as $p_i(b_i) = P_i(b_i, b_{-i})$, and we explicitly indicate the agent's (true) valuation by writing the agent's utility function as $\bar{u}_i(b_i, v_i) = u_i(b_i, b_{-i})$. Note that $\bar{u}_i$ leaves b_{-i} implicit, whereas u_i leaves v_i implicit. With this new notation, agent i's utility (with b_{-i} implicit) is

$$\bar{u}_i(b_i, v_i) = v_i \cdot x_i(b_i) - p_i(b_i).$$

Now, consider the correspondence:

v_i	$\longleftrightarrow$	q	(belief)
x_i	$\longleftrightarrow$	D	(generalized allocation)
b_i	$\longleftrightarrow$	r	(reported belief)
$-p_i$	$\longleftrightarrow$	S_0	(baseline payment)

With this correspondence, $\bar{u}_i(b_i, v_i)$ is nothing more than a scoring rule $S(r, q)$. Hence, we can apply our characterization of proper scoring rules! Our characterization tells us which scoring rules are proper (which is the analogue of the DSIC property in auctions). For auctions, the characterization tells us which mechanisms are DSIC (the analogue of properness in information elicitation). We have the following lemma.

Lemma 1. *Consider a mechanism (X, P) in a single-parameter environment. The mechanism (X, P) is DSIC if and only if, for all agents $i \in [n]$ and for all bids b_{-i} of the other agents, there exists a convex function $G\colon \mathbb{R}_? \to \mathbb{R}$ such that agent i's utility can be expressed as*

$$\bar{u}_i(b_i, v_i) = G(b_i) + G'(b_i)(v_i - b_i),$$

where G' is the derivative of G.

Recall that when we showed that any proper scoring rule can be represented in terms of a convex function, we actually only proved the nice case where the proper scoring rule is differentiable and the convex function is differentiable. Yet, the characterization does hold in general. We similarly gloss over the non-differentiable case in the above characterization of DSIC mechanisms. Everything can be made to work in general, but with added technical difficulty. Yet, it is worth mentioning the other end of the spectrum, which is when agent i's utility is discontinuous in her bid. Such situations are common: discontinuity happens even in second-price auctions (think about why). To see a proof that directly accommodates things like second-price auctions (and many more complicated auction mechanisms), I direct the reader to Section 3.4 of Tim Roughgarden's Twenty Lectures on Algorithmic Game Theory. If there is a time, I may try to extend the above lemma to indicate how to use its form (with a subgradient in place of G') to accommodate mechanisms like the second-price auction.

It is worth reflecting on the meaning of G. In analogy with the proper scoring rules, wherein $G(q) = S(q, q)$, for auctions it will hold that $G(v_i) = v_i X_i(v_i, b_{-i}) - P_i(v_i, b_{-i})$. In the above, nowhere did we restrict ourselves to mechanisms that are individually rational (IR). If a mechanism is not only DSIC but also IR, does the mechanism being IR further limit the possibilities for the function G? If so, then in what way?

We are almost ready to present Myerson's Lemma, a foundational result in mechanism design. We just need one definition.

Definition 1 (Monotone)**.** An allocation rule X is *monotone* if, for all agents i and bids $b_{-i} \in V^{n-1}$ of the other agents, the partial allocation rule $X_i(b_i, b_{-i})$ is non-decreasing in agent i's bid b_i.

Intuitively, constraining an allocation rule to be monotone is a very mild restriction. If a an allocation rule is not monotone, an agent who increases her bid might actually receive less stuff! Such a situation seems particularly bad if the agent reports truthfully.

We now present Myerson's Lemma. Fortunately, most of the heavy lifting in the proof was already done by the previous lemma.[3]

Theorem 1 (Myerson's Lemma)**.** *Consider an allocation rule X for a single-parameter environment. Then*

(i) *There exists a payment rule P such that the mechanism (X, P) is DSIC if and only if X is monotone.*

(ii) *If X is monotone and we assume that for all i that $P_i(0, b_{-i}) = 0$, then there is a unique payment rule P such that (X, P) is DSIC, with P taking the explicit form*

$$P_i(b_i, b_{-i}) = b_i X_i(b_i, b_{-i}) - \int_0^{b_i} X_i(z, b_{-i}) dz = \int_0^{b_i} z \frac{\partial X_i(z, b_{-i})}{\partial z} dz;$$

the second equality holds if any partial allocation rule X_i is differentiable in agent i's bid.

Proof. By definition, agent i's utility is

$$\bar{u}_i(b_i, v_i) = v_i \cdot x_i(b_i) - p_i(b_i).$$

From Lemma 1, there exists a convex function G such that this same utility can be expressed as

$$\bar{u}_i(b_i, v_i) = G(b_i) + G'(b_i)(v_i - b_i) = G'(b_i)v_i + G(b_i) - b_i G'(b_i)$$

The only possibility for x_i is $x_i(b_i) = G'(b_i)$. Since G is convex, G' is non-decreasing[4] and hence monotone, which proves part *(i)* of the theorem.

We now prove part *(ii)*. Let us rewrite the utility as

$$\bar{u}_i(b_i, v_i) = x_i(b_i)v_i + G(b_i) - b_i x_i(b_i).$$

We will absorb $b_i x_i(b_i)$ into the payment rule. To get an explicit form of $G(b_i)$ in terms of x_i, observe that[5]

$$\int_0^{b_i} x_i(z) dz = \int_0^{b_i} G'(z) dz = G(b_i) - G(0),$$

and so

$$G(b_i) = \int_0^{b_i} x_i(z) dz + G(0).$$

We set the payment rule as

$$p_i(b_i) = b_i x_i(b_i) - \int_0^{b_i} x_i(z) dz - G(0).$$

[3] The proof might look a bit dense because all of the symbols, but I assure you that (aside from one part of the proof, which I point out), there is nothing fancy happening inside the proof.

[4] This is a property of convex functions.

[5] This is the fancy step.

Now, what value should $G(0)$ take? Well, since we want $p_i(0) = 0$, the only choice for $G(0)$ is $G(0) = 0$, which yields the payment rule in the theorem. Moreover, if x is differentiable (so G is twice differentiable), we may write

$$b_i x_i(b_i) - \int_0^{b_i} x_i(z)dz = \int_0^{b_i} z x_i'(z).$$

To last equality can be confirmed by applying integration by parts to the RHS. □

3 Revelation Principle

All the mechanisms we have considered so far were *direct-revelation mechanisms*. In the context of auctions, a direct-revelation mechanism is a mechanism that asks each agent to reveal her true valuation function. More generally, when each agent has a *type* — some private information, such as the agent's belief in information elicitation or the agent's valuation function in auctions — a direct revelation mechanism is one which asks each agent to reveal her type.

In our coverage of auctions, we have always sought mechanisms that are DSIC. If a mechanism is DSIC, then it has a dominant strategy equilibrium. Moreover, this dominant strategy equilibrium is a truthful equilibrium, meaning that every agent is incentivized to report her valuation. Note that a mechanism can be DSIC only if it is a direct revelation mechanism (since otherwise the agent's set of strategies is not equal to her set of valuation functions).

Yet, there are many examples of non-direct-revelation mechanisms. For example, consider any of a number of ascending auctions, in which a price is iteratively increased and the last bidder wins the item, or descending auctions, in which a price is iteratively decreased and the first bidder wins the item. In general, in a non-direct-revelation mechanism, each agent i plays strategies s_i in some set of strategies S_i, with the set of strategy profiles being $S = S_1 \times \dots S_n$. The allocation rule X and each payment rule P_i now take elements from S as inputs.

Ignoring computational issues, does the mechanism designer wield more power by going beyond direct-revelation mechanisms? If we restrict attention to mechanisms that have a dominant strategy equilibrium, then in a certain sense, the answer is no.

Theorem 2 (Revelation Principle). *For any non-direct-revelation mechanism (X, P) that has a dominant strategy equilibrium, there is a DSIC direct-revelation mechanism $(\bar{X}, \bar{P})$ that* implements *(X, P).*

By "implements", we mean that the mechanism $(\bar{X}, \bar{P})$ works by using (X, P) "under the hood". This will be made precise in the proof. Before getting to the proof, let's reflect on the theorem statement. The theorem applies to non-direct-revelation mechanisms that have a dominant strategy equilibrium. Such mechanisms are nice in the sense that the designer can reasonably infer what agents will do, i.e., play their dominated strategies. The theorem is interesting in that it takes the "dominant strategy equlibrium" part of the DSIC definition and says that a mechanism satisfying this first part of the definition can be "upgaded" to a direct-revelation mechanism that satisfies the full definition.

Proof. For any agent i, set $s_i(v_i)$ be the agent's dominant strategy given valuation $v_i \in V_i$. We design the mechanism $(\bar{X}, \bar{P})$ as follows:

1. Ask each agent i to report her valuation, giving bid $b_i \in V_i$.
2. For each agent i, feed $s_i(b_i)$ into (X, P).

The above mechanism simply asks each agent to report her valuation and then plays the agent's dominant strategy (given the reported valuation) on the agent's behalf.

We see that "implements" means that the mechanism $(\bar{X}, \bar{P})$ is really just using mechanism (X, P), after converting bids to dominant strategies (since (X, P) takes inputs in S instead of V).

It is not hard to see why $(\bar{X}, \bar{P})$ is DSIC. Suppose that an agent's utility were higher by reporting some $b_i \neq v_i$. Since the agent's utility is higher, it must be the case of that $s_i(b_i) \neq s_i(v_i)$; however, if $s_i(b_i)$ gives more utility (under mechanism (X, P) than $s_i(v_i)$ when the agent's (true) valuation is v_i, then $s_i(v_i)$ could not have been a dominant strategy, a contradiction. $\square$