Lecture 19

Lecturer: Bo Waggoner *Scribe: Bo Waggoner*

Truthfulness and Single-Parameter Domains

In this lecture, we discuss truthfulness in auction design, including an important "revelation principle" and Myerson's Lemma for single-parameter environments.

Revelation principle. We call the mechanisms defined so far *direct-revelation* mechanisms because they ask agents to report their valuation functions. But what about other kinds of mechanisms? In general, we can define a **(non-revelation) mechanism** where each agent i has some action space B_i, the choice rule is $Y : B_1 \times \cdots \times B_n \to A$, and the payment rule is $Q : B_1 \times \cdots \times B_n \to \mathbb{R}^n$.

For example, the *descending-price* auction for a single item starts with a high price, then slowly lowers it while the bidders observe. At any time, a bidder can shout "Stop", claim the item, and pay the current price. Here B_i can be modeled as determining when to claim the item, assuming it is unclaimed.

In this class, we won't worry too much about non-revelation mechanisms, due to the following theorem.

Theorem 1 *Let (Y, Q) be a non-revelation mechanism with a dominant strategy equilibrium $b(v)$, i.e. each player i plays $b_i(v_i)$. Then there exists a direct-revelation mechanism (X, P) with the same choice and payment rule that is dominant-strategy incentive compatible.*

Proof The mechanism (X, P) will work as follows.

1. Ask all agents to report $v_1, \dots, v_n$.
2. Compute the actions $b_i(v_i)$ they would have taken under the mechanism (Y, Q).
3. Compute the choice $a = Y(b_1, \dots, b_n)$ and payments $p = Q(b_1, \dots, b_n)$.

So in the end, we have $X(v) = Y(b(v))$ and $P(v) = Q(b(v))$.

To see that this is DSIC, consider any deviation $\hat{v}_i$. The mechanism (X, P) will compute some simulated action $\hat{b}_i = b_i(\hat{v}_i)$. This may be different than $b_i(v_i)$. Meanwhile, all the other agents play some b_{-i}. So being truthful gives the outcome $(Y(b_i, b_{-i}), Q(b_i, b_{-i})$ while misreporting gives $(Y(\hat{b}_i, b_{-i}), Q(\hat{b}_i, b_{-i}))$. By assumption, the mechanism Y, P had a dominant strategy of b_i, so deviating to $\hat{b}_i$ cannot improve i's utility. ■

Corollary 2 *Any welfare achievable in dominant strategy equilibrium by a non-revelation mechanism is also achievable in a direct-revelation DSIC mechanism.*

In fact, one can extend this principle further to equilibria that are not dominant strategy, but you get the idea.

Truthfulness in single-parameter environments. So we can only worry about truthful mechanisms, but how do we design those in general? Here, we'll look at a pretty general environment called *single-parameter* settings.

Definition 3 (Single Parameter Domain) *A* ***single parameter domain*** *is a mechanism-design setting where:*

1. *We can write each alternative $a \in A$ as $a = (a_1, \dots, a_n)$ where a_i is interpreted as the "amount" that i gets when a is selected.*

2. *Each valuation function can be captured by a real number w_i, interpreted as i's "value per unit" amount that i gets.*

Thus, i's utility for outcome $o = (a, p)$ is $w_i a_i - p_i$.

Many things are single parameter domains. For example:

1. Single item auctions. We already know that agent i prefers to win the item than to lose it – all that needs to be specified is how much agent i values the item. Here the set of alternatives A looks like all vectors of the form $(0, 0, 1, 0, \ldots, 0)$. That is:

$$a_i = \begin{cases} 1, & i \text{ wins the item} \\ 0, & \text{otherwise.} \end{cases}$$

2. Selling a divisible item. The seller has, say, a liter of ice cream and can split it among the agents in any way. Agent i has utility w_i per liter of ice cream (suppose it is linear for the sake of the example). Here a consists of all vectors $(a_1, \ldots, a_n)$ that sum to one liter, where a_i is the amount of ice cream allocated to i.

3. Randomized single-item auctions. Suppose that, given a set of bids, the auctioneer may decide to allocate the item randomly. The set of alternatives $A = \Delta_n$, the set of probability distributions over n agents. The agent's value for $a = (a_1, \ldots, a_n)$ is her expected value, or w_i times the probability of getting the item: $w_i a_i$.

4. Buying a path in a network: In this problem, agents correspond to edges in a network, and will experience some cost if they are used. The mechanism would like to buy service from a set of agents that form a path in the network, to optimize some objective (minimize social cost, maximize throughput, etc.) Here an alternative a is a set of edges and:

$$a_i = \begin{cases} 1, & i\text{'s edge is in } A; \\ 0, & \text{otherwise.} \end{cases}$$

5. Job Scheduling: In this problem, the agents correspond to machines i, each of whom has a different cost c_i for running one unit of computation. Jobs j have different sizes ℓ_j (i.e. a job that would cost machine i $\ell_j \cdot c_i$ to run), and the task is to allocate jobs to machines to optimize some objective. We write $x_{ij} = 1$ if job j is allocated to machine i. Then:

$$a_i = \sum_j x_{ij} \ell_j$$

Now, we want to understand all DSIC mechanisms in single-parameter domains. To do this, we will use a similar argument to that we used to characterize truthful proper scoring rules.

Fix all reports w_{-i} of agents except i. Let us use the notation $a_i(w_i), p_i(w_i)$ as shorthand for the amount and payment for i when she reports w_i. More formally, we have

$$a_i(w_i) = X(w_i, w_{-i})_i$$
$$p_i(w_i) = P(w_i, w_{-i})_i.$$

Then we can define i's utility for reporting $\hat{w}_i$ when her true value is w_i:

$$u_i(\hat{w}_i; w_i) = w_i a_i(\hat{w}_i) - p_i(\hat{w}_i).$$

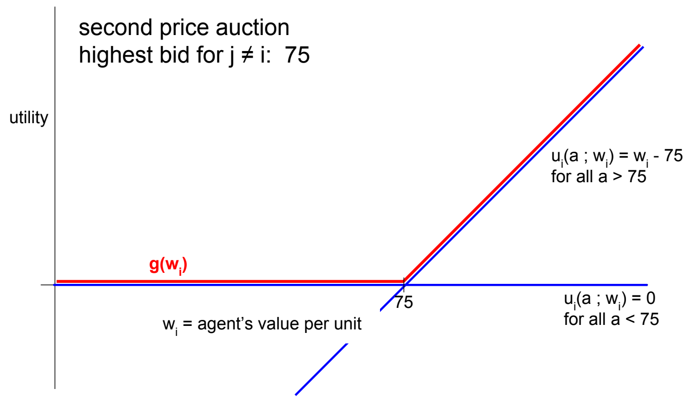

Figure 1: Illustrating Lemma 4 with a second-price auction where the highest bid was 75 from all bidders except i. All of i's bids map to two possibilities: if i bids anything $a \leq 75$, then she doesn't get the item. So she gets utility $u_i(a; w_i) = 0$ no matter what her value w_i is (horizontal blue line). If she bids anything $a \geq 75$, she gets the item and pays 75, so she gets utility $u_i(a; w_i) = w_i - 75$ (other blue line).

Lemma 4 *The mechanism (X, P) is DSIC if and only if, for all fixed w_{-i}, there exists a convex function $g : \mathbb{R}_{\geq 0} \to \mathbb{R}$ such that*

$$u_i(\hat{w}_i; w_i) = g(\hat{w}_i) + \frac{dg(\hat{v}_i)}{d\hat{w}_i} (w_i - \hat{w}_i) .$$

(Note: to be technically formal, $\frac{dg_i}{d\hat{w}_i}$ should be a subgradient rather than a derivative, but we won't be picky about this.)

Proof "Only if" direction: We must show that any DSIC mechanism satisfies the above. For each fixed report $\hat{w}_i$, the function $u_i(\hat{w}_i; \cdot)$ is an affine function of its second argument, i.e. linear plus a constant. An agent with valuation w_i best-responds by picking the report with maximum utility, which is the maximum over these affine functions, so

$$g(w_i) := u_i(w_i; w_i) = \max_{\hat{w}_i} u_i(\hat{w}_i; w_i).$$

Because g is a maximum over affine functions, it is a convex function. Furthermore, the line $u_i(\hat{w}_i; \cdot)$ is tangent to g at $\hat{w}_i$, so we can write $u_i(\hat{w}_i; w_i) = g(\hat{w}_i) + \frac{dg(\hat{w}_i)}{d\hat{w}_i} (w_i - \hat{w}_i)$.

"If" direction: Suppose that, for each fixed w_{-i}, there exists a convex g such that $u_i(\hat{w}_i; w_i)$ of the form given. Then the agent's best response by definition is $\arg\max_{\hat{w}_i} u_i(\hat{w}_i; w_i)$. As noted above, each $u_i(\hat{w}_i; \cdot)$ is an affine function and the maximum of them at w_i is $u_i(w_i; \cdot)$. So this mechanism is DSIC, that is, for every action of the other players, truthfulness is a best response. ■

This implies the more famous characterization:

Theorem 5 ("Myerson's Lemma") *The allocation rule X can be implemented as part of a DSIC mechanism if and only if, for each fixed bids w_{-i} of agents other than i, $a_i(\hat{w}_i)$ is monotone increasing. If so, it can be truthfully implemented only by the payment rule of the form*

$$p_i(\hat{w}_i) = C + a_i(\hat{w}_i)\hat{w}_i - \int_0^{\hat{w}_i} a_i(z)dz$$

for some constant C depending only on the other bids.

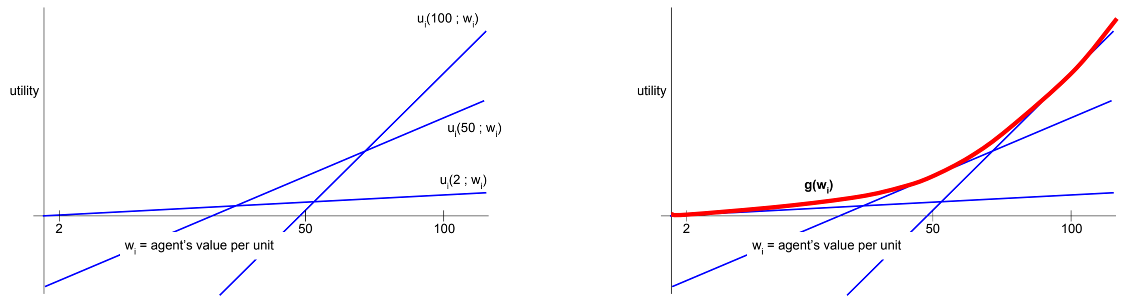

Figure 2: Illustrating Lemma 4 in some more general setting.

Proof Recall that

$$u_i(\hat{w}_i; w_i) := w_i a_i(\hat{w}_i) - p_i(\hat{w}_i)$$

We have from the above lemma that a mechanism with this pair a_i, p_i is truthful mechanism if and only if

$$u_i(\hat{w}_i; w_i) = g(\hat{w}_i) + \frac{dg(\hat{w}_i)}{d\hat{w}_i}(w_i - \hat{w}_i)$$

for some convex function g. Putting these together, we get that

$$a_i(\hat{w}_i) = \frac{dg(\hat{w}_i)}{d\hat{w}_i}$$

and

$$\begin{aligned} p_i(\hat{w}_i) &= \frac{dg(\hat{w}_i)}{d\hat{w}_i}\hat{w}_i - g(\hat{w}_i) \\ &= a_i(\hat{w}_i)\hat{w}_i - \int_0^{\hat{w}_i} a_i(z)dz. \end{aligned}$$

(We are omitting some technicalities, but one can assume some niceness or generalize the proof to get rid of them.)

Now, there exists a convex function g with $a_i = \frac{dg}{d\hat{w}_i}$ if and only if a_i is monotone increasing (this is one definition of convex function). This proves the result. ■