Algorithmic Game Theory, Summer 2017 Lecture 8 (4 pages)

Single-Parameter Mechanisms

Instructor: Thomas Kesselheim

1 Myerson's Lemma

Our main question today will be to identify those outcome rules f, for which we can find payment rules p such that $M = (f, p)$ is a truthful mechanism. We will call these outcome rules *implementable.*

It turns out that there is a very satisfying answer to this question, if we confine ourselves to single-parameter environments.

Definition 8.1. *An allocation rule f for a single-parameter mechanism design problem is monotone if for each player $i \in \mathcal{N}$ and for all bids b_{-i} of the players other than i, the allocation $f_i(z, b_{-i})$ to player i is non-decreasing in bid z.*

Theorem 8.2 (Myerson 1981)**.** *For single parameter environments, the following three claims hold: (1) An allocation rule is implementable if and only if it is monotone. (2) If allocation rule f is monotone, then the exists a unique payment rule p such that the mechanism $M = (f, p)$ is truthful, assuming that a zero bid implies a zero payment. It is given by*

$$p_i(b_i, b_{-i}) = b_i f_i(b_i, b_{-i}) - \int_0^{b_i} f_i(t, b_{-i}) \mathrm{d}t \enspace .$$

This result is remarkable for several reasons: (1) It reduces the rather abstract problem of deciding whether a certain allocation rule can be implemented, to the far more operational question of whether a given allocation rule is monotone. (2) It leaves essentially no ambiguity in regard to the payments. If we require that an agent with value zero pays nothing, then there is a unique payment rule that turns a given allocation rule into a truthful mechanism. (3) It gives an explicit formula for the payments that achieve this.

Proof. Let us consider any allocation rule f, whether monotone or not, and let us study how truthful payments could look like. Truthfulness requires that the utility of each bidder is maximized by bidding truthfully, no matter who bids and no matter what the other players' bids are, where the utility of player i for bid z is $u_i((z, b_{-i}), v_i) = v_i \cdot f_i(z, b_{-i}) - p_i(z, b_{-i})$ for b_{-i} denoting the bids of the other players.

Observe that for two possible valuations y and z the respective truthfulness inequalities imply

$$\begin{aligned} y f_i(y, b_{-i}) - p_i(y, b_{-i}) &\geq y f_i(z, b_{-i}) - p_i(z, b_{-i}) \\ z f_i(z, b_{-i}) - p_i(z, b_{-i}) &\geq z f_i(y, b_{-i}) - p_i(y, b_{-i}) \end{aligned}$$

The first inequality states that if the true value is y then the bidder does not want to instead bid z. The second inequality states that deviation to y is not beneficial if the true value is z. Rearranging terms and writing both inequalities together, we get lower and upper bounds on the payment difference for both bids

$$y \left(f_i(z, b_{-i}) - f_i(y, b_{-i}) \right) \leq p_i(z, b_{-i}) - p_i(y, b_{-i}) \leq z \left(f_i(z, b_{-i}) - f_i(y, b_{-i}) \right) \enspace . \tag{1}$$

This inequality is often called payment difference sandwich.

Ignoring the middle part, we already get that if $y \leq z$ then $f_i(y, b_{-i}) \leq f_i(z, b_{-i})$. This is the forward direction of part (1) of the theorem.

For the sake of simplicity, let us limit ourselves to allocation rules that are piecewise constant, as in the Vickrey auction of a single item, or in sponsored search. See Figure 1 for an illustration.

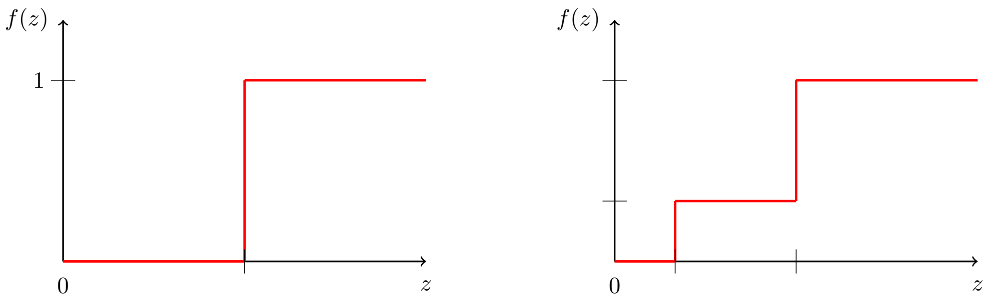

Figure 1: Piece-wise constant allocation curves

Proposition 8.3. *Given a truthful single-parameter mechanism $M = (f, p)$. Suppose $f(\cdot, b_{-i})$ is a monotone function that is piecewise constant on intervals $[z_j, z_{j+1})$ for $0 = z_0 < z_1 < \dots$. If $p_i(0, b_{-i}) = 0$, then*

$$p_i(b_i, b_{-i}) = \sum_{j: z_j \le b_i} z_j \left(f_i(z_j, b_{-i}) - f_i(z_{j-1}, b_{-i}) \right).$$

Proof. We use the payment differences sandwich (1). First, let us consider any b_i that is not a breakpoint z_j. Setting $v_i = b_i$ and $v_i' = b_i - \epsilon$ for small enough ϵ, we get from (1) that $p_i(b_i, b_{-i}) = p_i(b_i - \epsilon, b_{-i})$. This implies that $p_i(\cdot, b_{-i})$ is constant on $[z_j, z_{j+1})$.

Next, consider any breakpoint z_j. Now, by definition for any $\epsilon > 0$ that is small enough, we have $f_i(z_j - \epsilon, b_{-i}) = f_i(z_{j-1}, b_{-i})$. By the above consideration $p_i(z_j - \epsilon, b_{-i}) = p_i(z_{j-1}, b_{-i})$. That is for all $\epsilon > 0$ that are small enough

$$(z_j - \epsilon) \left(f_i(z_j, b_{-i}) - f_i(z_{j-1}, b_{-i}) \right) \le p_i(z_j, b_{-i}) - p_i(z_{j-1}, b_{-i}) \le z_j \left(f_i(z_j, b_{-i}) - f_i(z_{j-1}, b_{-i}) \right) \quad .$$

This means that

$$p_i(z_j, b_{-i}) - p_i(z_{j-1}, b_{-i}) = z_j \left(f_i(z_j, b_{-i}) - f_i(z_{j-1}, b_{-i}) \right)$$

because for $\epsilon \to 0$ the limits of the left and right part are identical. □

We can also rearrange the explicit formula. If $z_k \le b_i < z_{k+1}$, then

$$p_i(b_i, b_{-i}) = z_k f(z_k, b_{-i}) - \sum_{j=1}^{k} (z_j - z_{j-1}) f_i(z_j, b_{-i}).$$

Note that this matches exactly the integral expression in the theorem statement. By the same technique but playing around with more ϵ terms it can be shown to also hold for general functions f.

It remains to show that any monotone allocation rule combined with the payments $p_i(b_i, b_{-i}) = b_i f_i(b_i, b_{-i}) - \int_0^{b_i} f_i(t, b_{-i}) \mathrm{d}t$ is truthful. To this end, observe that in the mechanism $M = (f, p)$, we have

$$u_i(b, v_i) = (v_i - b_i) f_i(b) + \int_0^{b_i} f_i(t, b_{-i}) \mathrm{d}t$$

If $b_i \le v_i$ then

$$\begin{aligned} u_i(b, v_i) - u_i((v_i, b_{-i}), v_i) &= (v_i - b_i) f_i(b) + \int_0^{b_i} f_i(t, b_{-i}) \mathrm{d}t - \int_0^{v_i} f_i(t, b_{-i}) \mathrm{d}t \\ &= (v_i - b_i) f_i(b) - \int_{b_i}^{v_i} f_i(t, b_{-i}) \mathrm{d}t \le (v_i - b_i) f_i(b) - \int_{b_i}^{v_i} f_i(b) \mathrm{d}t = 0 \quad , \end{aligned}$$

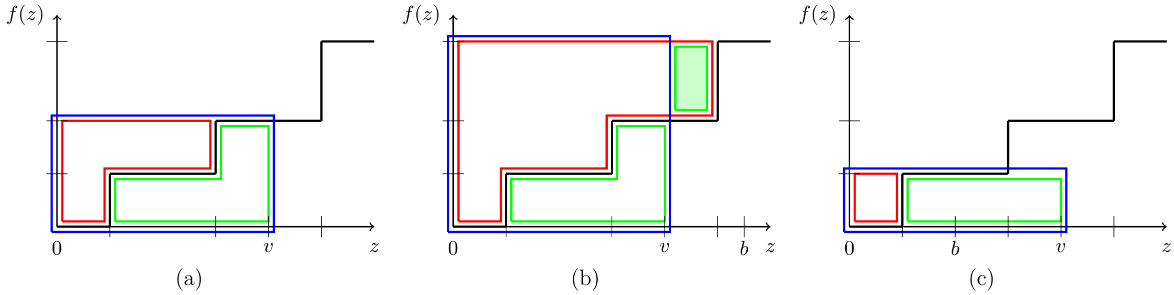

Figure 2: Visualization of the value (blue), the payment (red), and the utility (green) when bidding truthfully (on the left) and for over- and underbidding (in the middle and on the right). Shaded areas contribute negatively.

where the inequality uses monotonicity of f. If $b_i \geq v_i$ then by the same argument

$$\begin{aligned} u_i(b, v_i) - u_i((v_i, b_{-i}), v_i) &= (v_i - b_i) f_i(b) + \int_0^{b_i} f_i(t, b_{-i}) \mathrm{d}t - \int_0^{v_i} f_i(t, b_{-i}) \mathrm{d}t \\ &= (v_i - b_i) f_i(b) + \int_{v_i}^{b_i} f_i(t, b_{-i}) \mathrm{d}t \leq (v_i - b_i) f_i(b) - \int_{v_i}^{b_i} f_i(b) \mathrm{d}t = 0 \enspace . \end{aligned}$$

So, in any case $u_i((v_i, b_{-i}), v_i) \geq u_i(b, v_i)$.

We could also convince ourselves pictorially that this payment scheme is truthful, see Figure 2. In all three parts of Figure 2, the allocation curve is the same, as well as the true value of our player. Figure 2 (a) shows what happens in a truthful bid: Our bidder gets the surplus indicated by the area of the blue rectangle, with the red area showing her payment and the green area her utility. Figure 2 (b) shows what happens when she overbids: For bid b with $v < b$, her allocation goes up and therefore her surplus goes up (blue), but her pay (red) goes up by more than her surplus, resulting in a utility that is lower (the lower green L-shape minus the small green rectangle). On the other hand, underbidding (Figure 2 (c)) leads to a smaller allocation, smaller surplus (blue), smaller pay (red), but also smaller utility (green). That is, the player's utility is indeed maximized by her true bid, which proves the theorem.

□

2 Examples

We are now ready to apply the tools that we developed in this lecture to the three examples mentioned last time.

Example 8.4 (Single-Item Auction)**.** *We have already seen that the Vickrey (second-price) auction is truthful. We can recover this result from Myerson's lemma. We know that the payment for winning is the critical value at which a player becomes a winner. This is the second highest bid.*

Example 8.5 (Sponsored Search Auction)**.** *In sponsored search social welfare is maximized by greedily assigning position* 1 *through* k *to the bidders with the* 1*-st to* k*-th highest bid. Denoting the* j*-th highest bid by* $b_{(j)}$*, Myerson's lemma yields the following graphical representation of a player's payment whose bid is highest:*

More generally, the externality of a player i *that is assigned position* j *is the loss in welfare incurred on the players assigned slots below. If player* i *was not present they could all move one position up. In other words, setting* $\alpha_{k+1} = 0$*, player* i*'s payment is given by*

$$p_i(b_i, b_{-i}) = \sum_{\ell=j}^{k} (\alpha_j - \alpha_{j+1}) \cdot b_{(j+1)} \enspace .$$

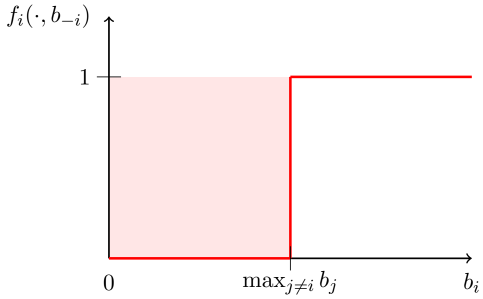

Figure 3: Allocation curve in a single-item auction

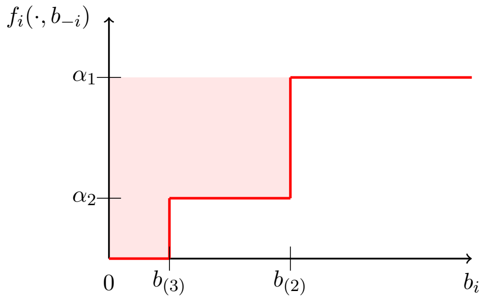

Figure 4: Allocation curve in a sponsored search auction

Example 8.6 (Scheduling on Related Machines)**.** *Minimizing the makespan on related machines is intractable (unless $P = NP$). So let's consider the following simple algorithmic strategy: Arrange the jobs in decreasing work order, and then greedily assign the next job to the machine that finishes it earliest. For example, three jobs with* 2, 1.1, *and* 1.05 *work units will be placed on two machines with* 0.4 *and* 0.45 *processing times per unit of work as follows. The heaviest job will go to the faster machine, and the other two lighter jobs will go to the slower machine. Now consider what happens if the slower machine claims to have a processing time of* 0.5*. Then it will receive the heaviest job, while the two smaller jobs are assigned to the other machine. So by claiming a higher speed (shorter processing time), a machine can reduce its workload. A contradiction to monotonicity.*

Let us conclude with two important orthogonal observations: (1) In many practical applications to which Myerson's Lemma applies, other (non-truthful) mechanisms are used in practice. For example, the mechanism used by Google to sell sponsored search results is not truthful. So there must be other reasons, in addition to truthfulness, that play a role. We will return to this point and non-truthful mechanisms later. (2) Myerson's lemma tells us that we can find the best truthful polynomial-time mechanism for a problem by searching for the best polynomial-time algorithm that is *monotone*. An important question thus is, does this additional requirement make the problem any harder?

Recommended Literature

- Tim Roughgarden's lecture notes `http://theory.stanford.edu/~tim/f13/l/l3.pdf`
- R. Myerson, Optimal Mechnism Design, Mathematics of Operations Research, 6:58–73, 1981. (Original characterization of truthful mechanisms)
- A. Archer and É. Tardos, Truthful Mechanisms for One-Parameter Agents. FOCS 2001. (Characterization of truthful mechanisms, which is deemed more accessible to computer scientists)