CS 6840 Algorithmic Game Theory January 22, 2020

Lecture 24: Smoothness in first price in auctions

Instructor: Eva Tardos *Scribe: Sander Aarts, Sophia Liu*

Recap

Last lecture we analyzed all-pay auction and defined smoothness in auctions. We looked at the case where there is one item and n bidders. Each bidder i has a valuation v_i and submits a bit b_i. We saw that in *all-pay* auctions the Nash equilibria could look relatively complicated, even in a two-player setting. We then defined smoothness in auctions; an auction is (λ, μ)-smooth if each bidder has a bid b_i^* s.t.

$$\sum_i u_i(b_i^*, b_{-i}) \geq \lambda \mathrm{OPT} - \mu \mathrm{Rev}(b). \tag{1}$$

We also proved that this implies that the social welfare at a Nash equilibrium b satisfies

$$\mathrm{SW}(b) \geq \frac{\lambda}{\max \mu, 1}.$$

In this lecture we are interested whether such a (λ, μ)-pair can be found. We show that in *first-price* auctions this is indeed the case; these auctions are $(\frac{1}{2}, 1)$-smooth.

Smoothness in first-price auctions

Recap of the model: We assume there are n bidders competing for one item. Each bidder i has values the item at some v_i. Bidders simultaneously place bids b_i. The player with the highest bid wins the item, and pays her bit; losing players pay nothing. Ties are resolved arbitrarily.

To prove smoothness, it suffices to find some b^* s.t. the smoothness condition (1) holds for some (λ, μ).

To this end, player i bids $b_i^* = v_i/2$. If the others bid b_{-i}, what do we know about $u_i(b_i^*, b_{-i})$?

There are three cases to consider:

1. If $b_i^* < \max_{j \neq i} b_j$ then $u_i(b_i^*, b_{-i}) = 0$.
2. If If $b_i^* = \max_{j \neq i} b_j$ then the outcome depends on how ties are broken, but certainly $u_i(b_i^*, b_{-i}) \geq 0$.
3. If $b_i^* > \max_{j \neq i} b_j$ then $u_i(b_i^*, b_{-i}) = v_i - b_i^* = v_i/2$.

We can add the utilities $u_i(b_i^*, b_{-i})$ over all players. Then if k is the player with the highest valuation, $k = \arg\max_j \{v_j\}$, we have

$$\sum_i u_i(b_i^*, b_{-i}) \geq u_k(b_k^*, b_{-k}).$$

In this setting the optimal social welfare OPT is exactly v_k, since this is the value to a bidder that the item can generate. The revenue to the auctioneer under b is $\mathrm{Rev}(b) = \max_j \{b_j\}$.

Claim: For player k with the highest valuation $u_k(b_k^*, b_{-k}) \geq \frac{1}{2}v_k - \max_j\{b_j\}$.

Proof. If the players bid (b_k^*, b_{-k}) and k wins then $u_k(b_k^*, b_{-k}) = v_k/2 \geq v_k/2 - \max_j\{b_j\}$.

If players bid (b_k^*, b_{-k}) and k loses, there is some other player $j \neq k$ that bids higher than b_k^* so

$$u_k(b_k^*, b_{-k}) = 0 \geq b_k^* - \max_{j \neq k}\{b_j\} \geq v_k/2 - \max_j\{b_j\}.$$

■

If we put the above arguments together, we've shown that there is a strategy b^* such that for any b

$$\sum_i u_i(b_i^*, b_{-i}) \geq u_k(b_k^*, b_{-k}) \geq \tfrac{1}{2}v_k - \max_j\{b_j\} = \tfrac{1}{2}\mathrm{OPT} - \mathrm{Rev}(b).$$

I.e. First-Price auctions are $(\frac{1}{2}, 1)$-smooth. This implies that the social welfare under any Nash equilibrium is at least half of the social optimum.

Question: In what settings does this PoA bound apply?

1. Pure strategy Nash equilibrium?

   Clearly the result holds here. If we let b be a PSNE then

   $$\sum_i u_i(b) \geq \sum_i u_i(b_i^*, b_{-i}) \geq \tfrac{1}{2}\mathrm{OPT} - \mathrm{Rev}(b)$$

2. Mixed strategy Nash equilibrium?

   The same arguments work when b is random. We just have to add expectations where appropriate.

   $$\sum_i \mathop{\mathbb{E}}_b[u_i(b)] \geq \sum_i \mathop{\mathbb{E}}_b[u_i(b_i^*, b_{-i})] \geq \tfrac{1}{2}\mathrm{OPT} - \mathop{\mathbb{E}}_b[\mathrm{Rev}(b)]$$

   Where the last inequality applies via smoothness and the linearity of expectation.

3. Learning outcomes? (We'll likely see whether this holds later in the course).

4. Bayes' Nash equilibrium?

5. Bayes' Nash equilibrium with values from a joint distribution?

   We can focus on part 5. as it includes 4. as a special case. First, note that in a Bayes' Nash equilibrium a strategy can depend on a player's valuation, so $b_i(v_i)$ is a function that maps i's valuations to her bids. The map should not depend on others' valuations; these are unknown to i, though player i knows the (joint) distribution of the valuations. A strategy $b(\cdot) = (b_1(\cdot), \ldots, b_n(\cdot))$ is Nash if

   $$\mathop{\mathbb{E}}_{b, v_{-i}}[u_i(b(v)) \mid v_i] \geq \mathop{\mathbb{E}}_{b, v_{-i}}[u_i(b_i', b_{-i}(v)) \mid v_i] \quad \text{for all } b_i' \text{ and } v_i$$

   So if $b(\cdot)$ is Nash we have:

   $$\begin{aligned} \mathop{\mathbb{E}}_{v_i}\mathop{\mathbb{E}}_{b, v_{-i}}[u_i(b(v)) \mid v_i] &\geq \mathop{\mathbb{E}}_{v_i}\mathop{\mathbb{E}}_{b, v_{-i}}[u_i(b_i^*, b_{-i}(v)) \mid v_i] \\ \implies \mathop{\mathbb{E}}_{b,v}[u_i(b(v))] &\geq \mathop{\mathbb{E}}_{b,v}[u_i(b_i^*(v), b_{-i}(v))] \\ \implies \sum_i \mathop{\mathbb{E}}_{b,v}[u_i(b(v))] &\geq \sum_i \mathop{\mathbb{E}}_{b,v}[u_i(b_i^*(v), b_{-i}(v))] \end{aligned}$$

As before, smoothness and linearity of expectation gives

$$\sum_i \mathop{\mathbb{E}}_{b,v}[u_i(b_i^*(v), b_{-i}(v))] \geq \tfrac{1}{2} \mathop{\mathbb{E}}_{v}[\mathrm{OPT}] - \mathop{\mathbb{E}}_{v,b}[\mathrm{Rev}(b)].$$

Note that throughout we only required that b_i^* depend on v_i; this is quite general, and works even if players know more about their opponents' valuations.