CS 6840: Algorithmic Game Theory **Spring 2017**

Lecture 23: March 22

Lecturer: Éva Tardos *Scribe: Jacob Bigenwald*

23.1 Limits of Price of Anarchy Analysis for Smooth Auctions

Given a notion of auction smoothness, what are the limits of analyzing price of anarchy through this property? Most notably, how do arguments by smoothness hold up under conditions of correlation? This can imply two things: correlated actions (as in a Coarse Correlated Equilibrium), or correlated participant valuations (as in a incomplete-information Bayes-Nash equilibrium).

As we'll see, the nice POA bounds achieved by smoothness for other auction setups holds for some instances of correlation, but not all of them.

23.1.1 Smoothness

For some auction, let $v_1 ... v_n$ be bidder valuations for bidder i (potentially a vector of valuations, for multi-item auctions). Bidders bid some value b_i, such that $b = (b_1, ... b_n)$. Bidders seek to maximize some utility $u_i(b, v_i)$, dependent on both the auction outcomes caused by other bids and their own valuation profile. For some distribution over bids N (where b can be sampled from N), this implies bidders seek to maximize $\mathbb{E}_{b \sim N} u_i(b, v_i)$.

With quasi-linear utility:

$$\mathbb{E}_{b \sim N}(u_i(b, v_i)) = \mathbb{E}_{b \sim N}(v_i(S_i(b)) - p_i(b))$$

where $S_i(b)$ is the set of items received by bidder i for a certain outcome, while $p_i(b)$ is the price paid by i under the same outcome. v_i in this situation is the total value gained from the items received, which may vary under assumptions of unit-demand.

Focusing only on non-distributed bids b, an auction (first-price, or all-pay) is smooth if there exists some bidding deviation $b_i^*(v)$ such that, for all b:

$$\sum_i u_i(b_i^*(v), b_{-i}, v_i) \geq \lambda Opt(v) - \mu \cdot rev(b)$$

where $rev(b)$ is the auction revenue, and $Opt(v)$ is the optimal allocation given bidder valuations. For a second-price auction, the alternative form exists:

$$\sum_i u_i(b_i^*(v), b_{-i}, v_i) \geq \lambda Opt(v) - \mu(\sum_i b_i(S_i(b)))$$

When b is sampled from a distribution N, the same holds. In this case, the sums are performed over expected utility, and expected revenue is used.

23.1.2 POA bounds in Smooth Auctions

So, with the above definition of smoothness, lets prove some preliminary POA bounds under certain conditions. First, lets consider price of anarchy for a full-info mixed-nash equilibrium, σ (where σ is a distribution over b).

Lemma 23.1 *The price of anarchy for a (λ, μ)-smooth auction in a first-price (or all-pay) full-info mixed-Nash equilibrium is* $\frac{\lambda}{max(\mu,1)}$

Proof: To start, consider the definition of a full-info mixed-Nash equilibrium:

$$\sum_i \mathbb{E}_{b\sim\sigma}(u_i(b,v)) \geq \sum_i \mathbb{E}_{b\sim\sigma}(u_i(b_i^*(v), b_{-i}, v))$$

Note our bidding deviation b_i^* does rely on information of the other bidder's valuations, but not on their bids. Also note that the right-hand side is equivalent (with expectations added) to the LHS of our above definition of smoothness. Hence, swapping that in:

$$\sum_i \mathbb{E}_{b\sim\sigma}(u_i(b,v)) \geq \lambda Opt(v) - \mu\mathbb{E}_{b\sim\sigma}(rev(b))$$

On the RHS, we can swap in $max(1, \mu)$ for μ, as this will only serve to decrease the value of the RHS, still preserving the inequality. On the LHS, we can multiply by $max(1, \mu)$ as this can only increase the already-positive utility, again preserving the inequality. Hence:

$$max(1,\mu)\cdot\sum_i \mathbb{E}_{b\sim\sigma}u_i(b,v) \geq \lambda Opt(v) - max(1,\mu)\cdot\mathbb{E}_{b\sim\sigma}rev(b)$$

$$max(1,\mu)\cdot\sum_i \mathbb{E}_{b\sim\sigma}u_i(b,v) + max(1,\mu)\cdot\mathbb{E}_{b\sim\sigma}rev(b) \geq \lambda Opt(v)$$

$$max(1,\mu)\cdot\mathbb{E}_{b\sim\sigma}(\sum_i u_i(b,v) + rev(b)) \geq \lambda Opt(v)$$

Because social welfare ($SW(b, v)$) is just the sum of bidder utilities plus revenue:

$$max(1,\mu)\cdot\mathbb{E}_{b\sim\sigma}SW(b,v) \geq \lambda Opt(v)$$

$$\mathbb{E}_{b\sim\sigma}SW(b,v) \geq \frac{\lambda}{max(1,\mu)}Opt(v)$$

Hence, the expected social welfare at a full-info Nash equilibrium σ is at least $\frac{\lambda}{max(1,\mu)}$ times the optimal social welfare ■

Lemma 23.2 *The price of anarchy for a (λ, μ)-smooth auction in a second-price full-info mixed-Nash equilibrium is $\frac{\lambda}{\mu+1}$, assuming bidders big below their value, that is for all sets S, and all bidders i, the total bid for set S, denoted by $b_i(S)$ is at most $v_i(S)$, the value for bidder i for the set.*

Proof: Like in the above, we can use the equilibrium definition, now combined with the second smoothness definition, to get us the following:

$$\sum_i \mathbb{E}_{b\sim\sigma}(u_i(b,v)) \geq \lambda Opt(v) - \mu(\sum_i \mathbb{E}_{b\sim\sigma}(b_i(S_i(b))))$$

Social welfare (again $SW(b,v)$) is at least the total utility of the bidders (as revenue cannot be negative), so the following still holds:

$$\mathbb{E}_{b\sim\sigma}(SW(b,v)) \geq \lambda Opt(v) - \mu(\sum_i \mathbb{E}_{b\sim\sigma}(b_i(S_i(b))))$$

By the assumption that no bidder is bidding above their value, so for $b \sim \sigma$ $\sum_i b_i(S_i(b)) \leq \sum_i v_i(S_i(b)) = SW(b,v)$, where $S_i(b)$ is the set that bidder i won with bids b. Hence:

$$\mathbb{E}_{b\sim\sigma}(SW(b,v)) \geq \lambda Opt(v) - \mu(\mathbb{E}_{b\sim\sigma}(SW(b,v)))$$
$$(\mu+1)\cdot\mathbb{E}_{b\sim\sigma}(SW(b,v)) \geq \lambda Opt(v)$$
$$\mathbb{E}_{b\sim\sigma}(SW(b,v)) \geq \frac{\lambda}{\mu+1} Opt(v)$$

Hence, for some Nash equilibrium σ in a second-price auction, social welfare at least $\frac{\lambda}{\mu+1}$ that of the optimal welfare is achieved.

■

23.1.3 Introducing Correlation

But what if we want to know price of anarchy in a CCE? For our purposes, let's only consider a first-price pure Nash, just because it makes the proof simpler (it does, in fact, work for the same auctions as above). For this, the proof is similar, except it begins with our definition of no-regret learning (as this corresponds to a CCE):

$$\sum_t u_i(b^t, v_i) \geq (1-\epsilon)(\sum_t u_i(b_i^*(v), b_{-i}^t)) - \frac{\log d}{\epsilon}$$

Then, summing over all bidders:

$$\sum_i \sum_t u_i(b^t, v_i) \geq (1-\epsilon)(\sum_i \sum_t u_i(b_i^*(v), b_{-i}^t)) - \sum_i \frac{\log d}{\epsilon}$$

$$\sum_i \sum_t u_i(b^t, v_i) \geq (1-\epsilon)(\sum_t [\sum_i u_i(b_i^*(v), b_{-i}^t)]) - \sum_i \frac{\log d}{\epsilon}$$

Note the bracketed expression on the RHS. This matches nicely with our first definition of smoothness, for some b^t. So we swap in that inequality (still preserving the current one), to get:

$$\sum_i \sum_t u_i(b^t, v_i) \geq (1-\epsilon)(\sum_t (\lambda Opt(v) - \mu \cdot rev(b^t))) - \sum_i \frac{\log d}{\epsilon}$$

It was here we stopped our proof, under the reason that the proof now can continue just the same way we did in he previous case. With that in mind, we showed that correlated bids from a CCE don't impact much our POA calculations when using smoothness. Concretely, we get the following if we ran leaning for T time steps, and there are n players

$$\frac{\max(1,\mu)}{T} \sum_t SW(b^t, W) \geq (1-\epsilon)\lambda Opt(v) - n\frac{\log d}{T\epsilon}$$

23.1.4 Correlated valuations

23.1.4.1 Uncorrelated Bayes-Nash

The other case to consider is when bidder valuations are unknown, but correlated. To start, let's just consider when bidder valuations are unknown, but also uncorrelated. A value v_i is distributed by some F_i, a known distribution on valuations. Consider the Bayes-Nash condition:

$$\sum_i \mathbb{E}_{v\sim F, b_j = b_j(v_j)}(u_i(b, v_i)) \geq \sum_i \mathbb{E}_{v\sim F, b_i = b_i(v_i), v'\sim F}(u_i(b_i^*(v_i, v'_{-i}), b_{-i}, v_i))$$

Let's dissect this. The first part is the $b_j = b_j(v)$. This is just a statement that the bid for a player is determined by only their value, as it's the only one they know. Next is the $v' \sim F$ in the RHS. Since bidder i doesn't know v_{-i}, he cannot consider the bid $b_i^*(v_i, v_{-i})$. Instead, he can re-sample v'_{-i} from the same distribution, and use bid $b_i^*(v_i, v'_{-i})$.

With that, the above states that the expected utility at some equilibrium (determined by b_i) is greater than the expected utility for any deviation, basing that deviation on v'_{-i} as explained above. With a change of variables we get:

$$\sum_i \mathbb{E}_{v\sim F, b_j = b_j(v_j)}(u_i(b, v_i)) \geq \sum_i \mathbb{E}_{b\sim G, v\sim F}(u_i(b_i^*(v), b_{-i}, v_i)) = \mathbb{E}_{b\sim G, v\sim F}(\sum_i u_i(b_i^*(v), b_{-i}, v_i))$$

where b is drawn from some distribution G composed of the original bidding functions b, and v is drawn from F, but the two drawn independently. Notice that this is just a change of variables (ra-naming v'_{-i} as v_{-i}: for the expectation of the utility of bidder i ($u_i(b_i^*(v_i, v'_{-i}), b_{-i}, v_i)$, this depends only on b_{-i}, v'_{-i}, and v_i, doesn't depend on b_i and v_{-i}, so for no index j does this depend on v_j and b_j at the same time, hence it is the same as our independently drawn v and b variables on the RHS.

With the RHS value, we can again break this down into a smoothness representation, and complete the proof from there. We have omitted the rest of the proof, but you do indeed get a same POA bound out of this as we got for the full information Nash case.

23.1.4.2 Correlated Bayes-Nash

However, this bound doesn't hold when bidder values are correlated. Consider the following example, for a second-price auction where each bidder is only allowed to bid on a single item:

There are n bidders, and $\sqrt{n} + n$ items to bid on. $\sqrt{n}$ of those items are "shared", and the other n are "unique". "Shared" items are valued by every bidder, with a value of 1. "Unique" items are only valued by a single bidder, with one per bidder, again with a value of 1. Bidders have unit-demand, and hence only want one item.

The optimal outcome, in this case, is assigning every bidder to their unique item, for a total welfare of n.

Now consider a Bayesian distribution of values, all as stated above, with the $\sqrt{n}$ shared items selected at random, and the assignment of unique items to bidders also done uniformly at random. For any bidder, they know which items they like (for a total number of $\sqrt{n} + 1$), but not which of those is unique to them. A Nash equilibrium in this setting is to pick a random item, and bid 1 on it. Since the auction is second price, once a bidder selected an item to bid on, bidding his true value is dominant strategy, and with no information about which item may be unique to him, this random selection is Nash.

Across all the bidders, you would expect $n/(\sqrt{(}n) + 1)$ unique items to be won (n times the probability they chose their unique item), and no more than $\sqrt{(}n)$ shared items to be won, because there are no more shared items. Hence the expected social welfare at this equilibrium is at most $n/(\sqrt{(}n)+1)+\sqrt{(}n) \leq 2\sqrt{(}n)$.

Therefore, plugging in for POA gives you $2/\sqrt{(}n)$, which can become arbitrarily bad as n increases. So, clearly, a Bayes-Nash equilibrium with correlated values doesn't have the same nice POA guarantees. But where in our proof is this issue introduced?

Well, in the above we considered the deviation $u_i(b_i^*(v_i, v'_{-i}), b_{-i}, v_i))$ for a sampled $v' \sim F$. With the correlation between items, we could instead sample $v'_{-i} \sim (F_{-i}|v_i)$, the rest of the value distribution conditioned on the value v_i known to the player. If we do this, we get the following, using $(F_{-i}|v_i)$ and $(G_{-i}|v_i)$ the distribution of values and bids conditioned on value v_i.

$$\sum_i \mathbb{E}_{v \sim F, b_j = b_j(v_j)}(u_i(b, v_i)) \geq \sum_i \mathbb{E}_{v \sim F, b_{-i} \sim (G_i|v_i)}(u_i(b_i^*(v), b_{-i}, v_i))$$

With correlated values we cannot do the next step to move the expectation outside of the sum, as each term of the expectation takes the bids b with a different conditioning.