CS364B: Frontiers in Mechanism Design
Lecture #16: The Price of Anarchy in First-Price Auctions *

Tim Roughgarden†

February 26, 2014

1 First-Price Auctions

1.1 Comparison to Second-Price Rules

In this lecture we pass from second-price to first-price auctions — auctions where you pay your bid. We haven't said much about pay-as-bid auctions in this or the previous course. The reason is that we've been focusing on auctions that are easy to participate in and that have strong incentive guarantees, like dominant-strategy incentive-compatibility.

In this part of the course, where we focus on simple auction formats in which it's not necessarily easy to figure out how to bid, the second-price auction loses its luster. The simultaneous second-price auctions studied in the last two lectures offers bidder with non-additive valuations no dominant strategies, and we didn't even try to figure out what their equilibria look like. Indeed, given the choice between a first-price and a second-price pricing rule for a simultaneous single-item auction format, the former might well strike many bidders as simpler to reason about. In any case, first-price payment rules are now at least as well motivated as second-price rules, and we study them in this lecture.

1.2 Bayes-Nash Equilibria of First-Price Single-Item Auctions

Even with only $m = 1$ item, the Bayes-Nash equilibrium of a first-price auction depends in an intricate way on the details of the environment. Recall from a previous exercise that with n bidders with valuations drawn i.i.d. from the uniform distribution on $[0, 1]$, then one Bayes-Nash equilibrium is for every bidder i to bid $\sigma_i(v_i) = \frac{n-1}{n} v_i$. That is, as the competition

†Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu.

grows stronger, all bidders shade their bids less. This is the unique Bayes-Nash equilibrium, though this fact is not easy to prove (see [1] and the references therein).

More generally, with all valuations drawn i.i.d. from a distribution that satisfies some mild technical conditions, the unique Bayes-Nash equilibrium is symmetric, with bidder i bidding the expected highest valuation by another bidder, conditioned on i's valuation being the highest (see [4]). For example, with two bidders and uniformly distribution valuations, if a bidder with valuation v_i has the highest valuation, than the conditional expected valuation of the other bidder is $v_i/2$. A consequence of this fact is that the Bayes-Nash POA in single-item first-price auctions with i.i.d. bidders equals 1: in the unique Bayes-Nash equilibrium, all bidders using the same bidding function is strictly increasing in valuation, so the highest bidder in equilibrium is always the bidder with the highest valuation.

The story is much different when bidders have different valuation distributions, however, even when the valuations are independent and there is only one item. First, the Bayes-Nash POA need not be one — some of the time, the highest bidder in equilibrium will not be the bidder with the highest valuation. This is true even when there are two bidders with valuations uniform in $[0, 1]$ and in $[0, 2]$. Here is a rough intuition for this fact. If the bidders were i.i.d., then each would bid half its value. From the first bidder's perspective, the other bidder represents stiffer competition than an i.i.d. bidder, so it is incentivized to bid more aggressively than in the i.i.d. case. The opposite reasoning applies to the second bidder, so it is incentivized to bid less aggressively than in the i.i.d. case. This opens up the possibilities of bid inversions (w.r.t. the valuations) and hence welfare loss. See the Exercises for details.

Another complication of the i.i.d. case is that it's almost impossible to solve explicitly for Bayes-Nash equilibria. For example, Vickrey's original paper [6], in addition to proposing the Vickrey auction, solves for the Bayes-Nash equilibrium of a first-price auction with i.i.d. bidders with uniform distributions. He proposed as an open question solve for the Bayes-Nash equilibrium of a first-price auction with two bidders with valuation distributions that are uniform with different supports $[a_1, b_1]$ and $[a_2, b_2]$. This problem was only solved *a half-century later* [3]. Clearly, any approach that relies on explicitly solving for the Bayes-Nash equilibria won't get very far. Fortunately, as we've seen, good price-of-anarchy bounds do not require a characterization of the equilibria.

2 The Bayes-Nash POA in Simultaneous First-Price Auctions

The rest of this lecture proves bounds on the Bayes-Nash POA of first-price auctions, both in single-item auctions and in simultaneous first-price auctions (S1A's). We can calibrate our expectations using our experience with S2A's in the last lecture. Like with S2A's, we'd rather not have to manipulate directly any priors or Bayes-Nash equilibria. We'd prefer to establish a "smoothness-type" inequality for auction games of complete information, and then conclude bounds on the Bayes-Nash POA through an extension theorem. The key inequality that drove our Bayes-Nash POA analysis for S2A's was the following: for every valuation profile $\mathbf{v}$, there exist hypothetical deviations $\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*$ such that, for every bid vector $\mathbf{b}$,

$$\sum_{i=1}^{n} u_i(\mathbf{b}_i^*, \mathbf{b}_{-i}) \geq \text{OPT welfare}(\mathbf{v}) - \sum_{i=1}^{n} \sum_{j \in S_i(\mathbf{b})} b_{ij}. \tag{1}$$

Importantly, the hypothetical deviations $\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*$ depend on $\mathbf{v}$ but not on $\mathbf{b}$.

2.1 Single-Item First-Price Auctions

Let's warm up with the case of a single item ($m = 1$); as we'll see, there an elegant way to extend the following analysis to an arbitrary number of items. Motivated by (1), we're looking for an analog of the following: for every valuation profile $\mathbf{v}$, there exist bids $b_1^*, \ldots, b_n^*$ such that

$$\sum_{i=1}^{n} u_i(\mathbf{b}_i^*, \mathbf{b}_{-i}) \geq \max_{i=1}^{n} v_i - \max_{i=1}^{n} b_i. \tag{2}$$

We emphasize that we want (2) to hold for every bid vector $\mathbf{b}$, no matter how weird.

In the second-price auctions studied in the past two lectures, the deviating bids $\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*$ were always constructed in the same way: compute the welfare-maximizing allocation for $\mathbf{v}$, and define $\mathbf{b}_i^*$ as going "all-in" for the bundle i receives in this allocation. In the present single-item context, this corresponds to setting $b_i^* = v_i$ if i has the highest valuation and zero otherwise. This choice of $b_1^*, \ldots, b_n^*$ fails utterly to satisfy (2): the left-hand side is guaranteed to be 0 (since if i wins in $(b_i^*, \mathbf{b}_{-i})$ it pays is valuation) while the right-hand side can be arbitrarily large.

Happily, transferring all our work for second-price auctions over to first-price auctions requires only the smallest change: to compensate for the higher prices, we just cut our previous deviating bids in half. Thus, for a single item, we set $b_i^* = v_i/2$ if i has the highest valuation (breaking ties arbitrarily) and 0 otherwise. The following analog of (2) holds.

Lemma 2.1 *For every valuation profile $\mathbf{v}$, if $b_i^* = v_i/2$ when i has the highest valuation and $b_i^* = 0$ otherwise, then*

$$\sum_{i=1}^{n} u_i(\mathbf{b}_i^*, \mathbf{b}_{-i}) \geq \frac{1}{2} \max_{i=1}^{n} v_i - \max_{i=1}^{n} b_i. \tag{3}$$

Proof: Since $b_i^* \leq v_i$ for every i, the left-hand side of (3) is nonnegative. Thus, we can assume that $\frac{1}{2} \max_{i=1}^{n} v_i > \max_{i=1}^{n} b_i$. In this case, if bidder i has the highest valuation, then it wins the item in the outcome $(b_i^*, \mathbf{b}_{-i})$ at a price of $\frac{1}{2}v_i$, yielding utility $\frac{1}{2}v_i \geq \frac{1}{2}v_i - \max_{k=1}^{n} b_k$. This verifies the inequality (3). ■

If Lemma 2.1 and its proof strike you as a little loose and ripe for optimization, then you're right — see Section 2.3. We also note that the deviations $b_i^* = v_i/2$ for *all* i would work equally well in Lemma 2.1; we use this fact in Corollary 2.3 below.

Simple as it is, Lemma 2.1 is enough to deduce good POA bounds for Bayes-Nash equilibria in S1A's. Let's start simply, with the pure Nash equilibria of single-item auctions.

You might expect that the coefficient of $\frac{1}{2}$ in (3) would lead to worse POA bounds that for second-price auctions, but this is not the case.

Corollary 2.2 *Every pure Nash equilibrium of a first-price single item auction has welfare at least 50% of the maximum possible.*

The actual statement of Corollary 2.2 is uninteresting, since it's not hard to prove that every pure Nash equilibrium of a first-price single-item auction is optimal (see Exercises). Corollary 2.2 also ignores the issue of the existence of pure Nash equilibrium, which depends on the details of the tie-breaking rule (see Exercises). What is interesting is the proof of Corollary 2.2, which we'll soon extend to settings where the POA is not 1 (like Bayes-Nash equilibria). Also note that there is no "no overbidding" hypothesis in Corollary 2.2 — in first-price auctions, bidder automatically avoid overbidding in equilibrium (since it results in negative utility).

Proof of Corollary 2.2: Let $v_i(\mathbf{b})$ denote the welfare contributed by player i in the outcome $\mathbf{b}$ — v_i if i is the highest bidder, 0 otherwise. Similarly define $p_i(\mathbf{b})$ for i's payment and $u_i(\mathbf{b}) = v_i(\mathbf{b}) - p_i(\mathbf{b})$ for i's utility. If $\mathbf{b}$ is a pure Nash equilibrium, then

$$\sum_{i=1}^{n} v_i(\mathbf{b}) = \sum_{i=1}^{n} u_i(\mathbf{b}) + \underbrace{\sum_{i=1}^{n} p_i(\mathbf{b})}_{=\max_{i=1}^{n} b_i} \tag{4}$$

$$\geq \sum_{i=1}^{n} u_i(b_i^*, \mathbf{b}_{-i}) + \max_{i=1}^{n} b_i \tag{5}$$

$$\geq \frac{1}{2} \max_{i=1}^{n} v_i, \tag{6}$$

where (5) follows from the pure Nash equilibrium condition and (6) follows from the smoothness condition (1). ■

Since $b_1^*, \ldots, b_n^*$ are independent of $\mathbf{b}$, the proof of Corollary 2.2 is a smoothness proof in the sense of CS364A. Hence, this POA bound of $\frac{1}{2}$ extends automatically to the usual full-information suspects, like the set of coarse correlated equilibria.

It may seem surprising that we proved a weaker-looking smoothness condition (1) for first-price auctions than for second-price auctions (where there was no "$\frac{1}{2}$") and yet obtained the same POA bound of $\frac{1}{2}$. To explain this, recall that our analysis of S2A's began with the inequality $\sum_{i=1}^{n} v_i(S_i(\mathbf{b})) \geq \sum_{i=1}^{n} u_i(S_i(\mathbf{b}))$, which effectively throws out the revenue term $\sum_{i=1}^{n} p_i(\mathbf{b})$ in (4). In the derivation (4)–(6) we keep this term around, and are eventually able to cancel it out with the $\max_{i=1}^{n} b_i$ term on the right-hand side. This cancellation works out in first-price auctions because the sum of the prices paid (on the left-hand side) equals the sum of the winning bids (on the right-hand side). With second-price auctions, the former quantity can be arbitrarily smaller than the latter quantity. This is why the former quantity is not directly useful and can be dropped without harm in the S2A analysis (the guarantee of $\frac{1}{2}$ is tight in the worst case). This is also why a no overbidding condition is needed to control the otherwise arbitrarily large second quantity in the S2A analysis.

We now turn to the Bayes-Nash POA. Remarkably, with a single item, the welfare guarantee of $\frac{1}{2}$ extends to Bayes-Nash equilibria with respect to an arbitrary *correlated* prior distribution.

Corollary 2.3 *For every (correlated) prior distribution* $\mathbf{F}$, *every (mixed) Bayes-Nash equilibrium of a first-price single item auction has expected welfare at least 50% of the maximum possible.*

Corollary 2.3 is a bit of an aside — it can't be extended to S1A's with two or more items — but it nicely illustrates conditions under which there is a stronger extension theorem for Bayes-Nash equilibria, for correlated and not only product prior distributions. This stronger extension theorem is enabled by the fact that each hypothetical deviation b_i^* defined above depends only on i's valuation v_i and does *not* depend on $\mathbf{v}_{-i}$.[1] Because such deviations can be executed without knowledge of $\mathbf{v}_{-i}$, the "doppelganger trick" from last time is unnecessary, and the simpler linearity arguments from last quarter's extension theorems can be used instead.

Proof of Corollary 2.3: For each i, define the $b_i^*(v_i) = v_i/2$. If σ is a Bayes-Nash equilibrium, then

$$\begin{aligned}
\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^n v_i(\sigma(\mathbf{v}))\right] &= \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^n u_i(\sigma(\mathbf{v}))\right] + \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\underbrace{\sum_{i=1}^n p_i(\sigma(\mathbf{v}))}_{\max_{i=1}^n \sigma_i(v_i)}\right] &(7)\\
&\geq \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^n u_i(b_i^*(v_i), \sigma_{-i}(\mathbf{v}_{-i}))\right] + \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\max_{i=1}^n \sigma_i(v_i)\right] &(8)\\
&\geq \frac{1}{2}\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\max_{i=1}^n v_i\right], &(9)
\end{aligned}$$

where (8) follows from linearity of expectation and the Bayes-Nash equilibrium condition, with the hypothetical (and well defined) deviations $b^1(v_1), \ldots, b^*(v_n)$; and (9) follows from Lemma 2.1 and linearity of expectation. ■

2.2 Simultaneous First-Price Auctions

We now consider the general case of S1A for m items. As with S2A's, we'll be able to handle XOS valuations — valuations of the form $v_i(S) = \max_{\ell=1}^r\{\sum_{j\in S} a_{ij}^\ell\}$, where $a_i^1, \ldots, a_i^r$ are additive valuations on U.

[1] Actually, the definition of $b_1^*, \ldots, b_m^*$ in Lemma 2.1 does not satisfy this property. But the hypothetical deviations $\mathbf{b}_i^* = v_i/2$ do, and work equally well for Lemma 2.1 and Corollary 2.2.

As usual, the key point is to identify hypothetical deviations $\mathbf{b}_1^*(\mathbf{v}), \ldots, \mathbf{b}_n^*(\mathbf{v})$ for the bidders, which can depend on the bidders' valuations (but on the choice of another bid vector $\mathbf{b}$). We combine the ideas from S2A's and from a first-price auction is the most straightforward way possible: $\mathbf{b}_i^*$ corresponds to bidder i going "all in" for the bundle S_i^* that it gets in the optimal allocation for $\mathbf{v}$, where "all in" corresponds to bids that are half of what they were for S2A's (to compensate for the first-price payment rule).

Formally, for a profile $\mathbf{v}$ of XOS valuations and optimal allocation $S_1^*, \ldots, S_n^*$, for $i = 1, 2, \ldots, n$ let a_i^* be an additive valuation that satisfies

$$\sum_{j \in S_i^*} a_{ij}^* = v_i(S_i^*) \tag{10}$$

and

$$\sum_{j \in S} a_{ij}^* \leq v_i(S) \tag{11}$$

for all $S \subseteq U$. Define

$$b_{ij}^* = \begin{cases} \frac{a_{ij}^*}{2} & \text{if } j \in S_i^* \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

In contrast to the single-item case (Corollary 2.3), these deviating bids depend on v_i *and* $\mathbf{v}_{-i}$.

We now show that our smoothness inequality (1) for first-price single-item auctions extends effortless to S1A's, essentially by the additivity of prices and effective additivity of XOS valuations with respect to the fixed allocation $S_1^*, \ldots, S_n^*$. Let $u_i^*(\mathbf{b}) = \sum_{j \in S_i(\mathbf{b})} a_{ij}^* - p_i(\mathbf{b})$ denote the hypothetical utility of player i if its valuation was a_i^* instead of v_i. By (11), $u_i^*(\mathbf{b}) \leq u_i(\mathbf{b})$ for every bid vector $\mathbf{b}$. Let u_{ij}^* denote the contribution to i's utility from the j single-item auction; this is well defined because both a_i^* and p_i are additive over items.

For every bid vector $\mathbf{b}$, and $\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*$ defined as in (12), we can derive the following smoothness-type inequality:

$$\begin{aligned}
\sum_{i=1}^n u_i(\mathbf{b}_i^*, \mathbf{b}_{-i}) &\geq \sum_{i=1}^n u_i^*(\mathbf{b}_i^*, \mathbf{b}_{-i}) \\
&= \sum_{j=1}^m \sum_{i=1}^n u_{ij}^*(b_{ij}^*, (\mathbf{b}_{-i})_j) \\
&\geq \sum_{j=1}^m \left(\frac{1}{2} \max_{i=1}^n a_{ij}^* - \max_{i=1}^n b_{ij} \right) && (13) \\
&\geq \frac{1}{2} \sum_{i=1}^n \sum_{j \in S_i^*} a_{ij}^* - \max_{i=1}^n \sum_{j \in S_i(\mathbf{b})} b_{ij} \\
&= \frac{1}{2} \cdot \text{OPT welfare}(\mathbf{v}) - \max_{i=1}^n \sum_{j \in S_i(\mathbf{b})} b_{ij}. && (14)
\end{aligned}$$

Inequality (13) follows from Lemma 2.1, applied to the single-item auction for each item j with bid profile $(b_{1j},\dots,b_{nj})$, valuations a^*_{ij} for the player i with $j \in S^*_i$ and 0 otherwise, and deviating bids $b^*_{ij} = a^*_{ij}/2$ for this player and 0 otherwise. Inequality (14) follows from (10). Summarizing, the $(\frac{1}{2},1)$-smoothness inequality (1) for single-item auctions extends to the $(\frac{1}{2},1)$-smoothness inequality

$$\sum_{i=1}^{n} u_i(\mathbf{b}^*_i,\mathbf{b}_{-i}) \ge \frac{1}{2}\cdot \text{OPT welfare}(\mathbf{v}) - \max_{i=1}^{n} \sum_{j\in S_i(\mathbf{b})} b_{ij}, \tag{15}$$

S1A's with XOS valuations, with $(b^*_{1j},\dots,b^*_{nj})$ defined as in (12). Syrgkanis and Tardos [5] refer to this extension from single-item to simultaneous single-item auctions with XOS valuations as *simultaneous composition.*[2]

Just as (1) yields Corollary 2.2 for single-item auctions, inequality (15) implies the following (using the same proof).

Corollary 2.4 *For every profile of XOS valuations, every pure Nash equilibrium of a S1A has welfare at least 50% of the maximum possible.*

Like Corollary 2.4, the POA bound in Corollary 2.4 extends to mixed Nash, correlated, and coarse correlated equilibria.

Last lecture we showed that, even with unit-demand valuations, the Bayes-Nash POA of S2A's with a correlated valuation distribution can be inverse polynomial in n. A variant of this example proves the same point for S1A's (see Exercises). Thus, Corollary 2.3 cannot extend to S1A's. We do, however, have the following analog of our Bayes-Nash POA bound for S2A's.

Theorem 2.5 ([2]) *For every product prior distribution* $\mathbf{F}$ *over XOS valuations, and every (mixed) Bayes-Nash equilibrium* σ*,*

$$\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\mathit{welfare}(\sigma(\mathbf{v}))] \ge \frac{1}{2}\cdot \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\mathit{OPT\ welfare}(\mathbf{v})]\,. \tag{16}$$

We've stated the guarantee (16) for pure Bayes-Nash equilibria; the statement for mixed Nash equilibria just has an extra expectation (over the random actions of $\sigma(\mathbf{v})$) on the left-hand side.

The proof of Theorem 2.5 stiches together arguments that you've already seen. Given a Bayes-Nash equilibrium σ, the first step (7) is the same as in Corollary 2.3. The rest of the proof follows the Bayes-Nash POA bound for S2A's given last lecture. In more detail, the second step aims to invoke the Bayes-Nash equilibrium hypothesis with the hypothetical deviations $\mathbf{b}^*_1(\mathbf{v}),\dots,\mathbf{b}^*_n(\mathbf{v})$ defined in (12). This can't be done directly, since these deviations depend on the full valuation profile $\mathbf{v}$ and a bidder acts knowing only its own valuation

[2] The $(1,1)$-smoothness inequality $\sum_{i=1}^n u_i(\mathbf{b}^*_i,\mathbf{b}_{-i}) \ge \text{OPT welfare}(\mathbf{v}) - \max_{i=1}^n \sum_{j\in S_i(\mathbf{b})} b_{ij}$ we proved for S2A's can also be derived via simultaneous composition, by first proving the inequality only for second-price auctions and then using the derivation above to extend it to S2A's (see Exercises).

(and the prior $\mathbf{F}$ and the equilibrium strategies σ). For this reason, the proof reuses the doppelganger trick from last lecture — i's hypothetical deviating strategy $\sigma_i^*(v_i)$ is to sample a valuation profile $\mathbf{w} \sim \mathbf{F}$ and bid according to $\mathbf{b}_i^*(v_i, \mathbf{w}_{-i})$. Because σ is a Bayes-Nash equilibrium, adopting strategy σ_i^* only lowers bidder i's expected utility. Expanding and rearranging like last lecture, and using that $\mathbf{F}$ is a product distribution, gives

$$\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^{n} u_i(\sigma(\mathbf{v}))\right] \geq \frac{1}{2}\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\text{OPT welfare}(\mathbf{v})] - \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^{n}\sum_{j\in S_i(\sigma(\mathbf{v}))} \sigma_i(v_i)_j\right]; \tag{17}$$

the details are left as an exercise. With first-price auctions, we can cancel the prices paid in (7) with the winning bids in (17), and Theorem 2.5 follows.

2.3 S1A's vs. S2A's: From $\frac{1}{2}$ to $\frac{e-1}{e}$

We've studied in detail two different simple auction formats, S2A's and S1A's. Which one is better? We proved exactly the same worst-case Bayes-Nash POA bound of $\frac{1}{2}$ for both, although for S2A's the bound holds only for equilibria that satisfy a no overbidding condition. We saw two lectures ago that the bound of $\frac{1}{2}$ is tight in the worst case for S2A's, even for pure Nash equilibria with unit-demand valuations. We've given no lower bounds for the POA in S1A's, opening the possibility of superior POA guarantees. Indeed, pure Nash equilibria of full-information S1A's, when they exist are fully efficient.

For first-price single-item auctions, lemma 2.1 can be improved to the following.

Lemma 2.6 ([5]) *For every valuation profile $\mathbf{v}$, there are distributions $D_1^*, \ldots, D_n^*$ over deviations $b_1^*, \ldots, b_n^*$ such that*

$$\mathbf{E}_{\mathbf{b}^*\sim\mathbf{D}^*}\left[\sum_{i=1}^{n} u_i(\mathbf{b}_i^*, \mathbf{b}_{-i})\right] \geq \frac{e-1}{e}\max_{i=1}^{n} v_i - \max_{i=1}^{n} b_i. \tag{18}$$

We leave the proof of Lemma 2.6 as an exercise, using the distribution XXX

The reader should verify that the improvement from $\frac{1}{2}$ to $\frac{e-1}{e} \approx .63$ in Lemma 2.6 carries over to Corollaries 2.2, 2.3, and 2.4, and to Theorem 2.5. This approximation guarantee is essentially as good as those provided by the best-known polynomial-time approximation algorithms for welfare maximization with bidders with XOS valuations. Thus, worst-case POA analysis advocates the first-price rule over the second-price rule in simultaneous single-item auctions.

3 Toolbox Recap: Extension and Composition Theorems

In this lecture and the previous one, we learned two types of powerful tools for proving POA bounds in auction formats. The first is *extension theorems*, which extend smoothness-type inequalities to bounds on the Bayes-Nash POA. We saw our first extension theorem last quarter, for full-information equilibrium concepts like coarse correlated equilibria. Here, we've seen a number of closely related extension theorems for games of incomplete information. We emphasize that these extensions theorems are in no way particular to simultaneous single-item auctions; they apply to any setting in which a smoothness condition (see (19) below) holds. We'll apply extension theorems to a different auction format next lecture.

Each extension theorem takes as input a smoothness inequality of the form: for every valuation profile $\mathbf{v}$, there exist hypothetical deviations $\mathbf{b}_1^*(\mathbf{v}), \ldots, \mathbf{b}_n^*(\mathbf{v})$ such that, for every bid profile $\mathbf{b}$,

$$\sum_{i=1}^{n} u_i(\mathbf{b}_i^*, \mathbf{b}_{-i}) \geq \lambda \cdot \text{OPT welfare}(\mathbf{v}) - \mu \cdot \sum_{i=1}^{n} \text{sum of } i\text{'s winning bids in } \mathbf{b}. \qquad (19)$$

As in Lemma 2.6, distributions over deviating bids such that (19) holds in expectation are also sufficient.[3]

Each extension theorem produces as output some type of Bayes-Nash POA bound. They vary along two axes: the set of prior distributions covered (product vs. arbitrary) and the exact formula of the POA bound as a function of λ and μ ($\lambda / \max\{\mu, 1\}$ vs. $\lambda/(1+\mu)$).

1. When each deviating bid $\mathbf{b}_i^*(\mathbf{v}) = \mathbf{b}_i^*(v_i)$ depends only on v_i and is independent of $\mathbf{v}_{-i}$, then the corresponding POA bound holds for (mixed) Bayes-Nash equilibria with respect to an *arbitrary*, not necessarily product, prior distribution. Corollary 2.3 was an example of this extension theorem. The conclusion is strong and the proof is relatively simple because a bidder i has sufficient information (namely, v_i) to execute the deviation $\mathbf{b}^*(v_i)$.

2. In the more general and common case where the deviating bid $\mathbf{b}_i^*(\mathbf{v})$ depends on v_i *and* $\mathbf{v}_{-i}$, the corresponding POA bound holds for (mixed) Bayes-Nash equilibria with respect to an arbitrary *product* prior distribution. The proof uses the doppelganger trick, whereby a bidder i executes the deviation $\mathbf{b}_i^*(v_i, \mathbf{w}_{-i})$ for doppelgangers $\mathbf{w}_{-i}$ — bidder i can do this because it knows v_i and the valuation distribution $\mathbf{F}$.

3. In "pay-as-bid" auctions, the term $\sum_{i=1}^{n}$ sum of i's winning bids in $\mathbf{b}$ equals the revenue of the auction. As we first saw in the proof of Corollary 2.2, this term can be cancelled with the revenue term that arises naturally when switching from welfare (our objective function) to the sum of player utilities (where the equilibrium hypothesis applies). When $\mu = 1$, as in our examples, this results in a POA bound of λ. For general $\mu > 0$, as long as bidders can guarantee nonnegative utility (e.g., by bidding 0 everywhere), the POA bound is $\frac{\lambda}{\max\{\mu,1\}}$; see the Exercises.

4. In auctions where payments can be less than bids, the term $\sum_{i=1}^{n}$ sum of i's winning bids in $\mathbf{b}$ can be much larger than the auction's revenue, and even much larger than the auction's welfare. This difference has two consequences. First, since we can't cancel this term

[3] In our examples, $\mu = 1$ and $\lambda \in \{\frac{1}{2}, \frac{e-1}{e}, 1\}$.

with the auction revenue, we simply ignore the auction revenue completely. Second, we prove POA bounds only for equilibria that satisfy a no overbidding condition, under which we can relate the term $\sum_{i=1}^{n}$ sum of i's winning bids in $\mathbf{b}$ to the welfare in the bid profile $\mathbf{b}$. The resulting POA bound is $\lambda/(1+\mu)$.

The second tool we acquire is a *composition theorem*, which extends a smoothness inequality of the form (19) of a single auction (like a single-item auction) to an analogous smoothness inequality with the same parameters for an arbitrary number of copies of the auction run in parallel. The derivation (13)–(14) establishes a simultaneous composition theorem for bidders with XOS valuations over the outcomes of the constituent auctions. There is analogous sequential composition theorem for auctions run in series, for bidders with unit-demand valuations over the outcomes of the constituent auctions [5].

Extension and composition theorems work well together, and can reduce proving Bayes-Nash POA bounds for non-trivial auction formats (like S2A's and S1A's) to the relatively trivial task of proving a smoothness condition (19) of an extremely simple auction (like a single-item auction). In hindsight, we can phrase our analysis of S2A's and S1A's as follows:

1. Prove the appropriate smoothness condition for a single-item auction. The details of the parameters of the deviating bids will differ with the auction format, but working them out is a relatively trivial task (cf., Lemmas 2.1 and 2.6).

2. Use simultaneous composition to extend the smoothness condition to simultaneous single-item auctions with bidders with XOS valuations.

3. Use the appropriate extension theorem to bound the Bayes-Nash POA with respect to an arbitrary product distribution over XOS valuations (in S2A's, under a suitable no overbidding condition).

References

[1] S. Chawla and J. D. Hartline. Auctions with unique equilibria. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, pages 181–196, 2013.

[2] A. Hassidim, H. Kaplan, M. Mansour, and N. Nisan. Non-price equilibria in markets of discrete goods. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*, pages 295–296, 2011.

[3] T. Kaplan and S. Zamir. Asymmetric first-price auctions with uniform distributions: Analytic solutions to the general case. *Economic Theory*, 50(2):269–302, 2012.

[4] V. Krishna. *Auction Theory*. Academic Press, second edition, 2010.

[5] V. Syrgkanis and É. Tardos. Composable and efficient mechanisms. In *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*, pages 211–220, 2013.

[6] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.