CS364B: Frontiers in Mechanism Design
Lecture #17: Part I: Demand Reduction in Multi-Unit Auctions Revisited *

Tim Roughgarden†

March 5, 2014

1 Recall: Multi-Unit Auctions

The last several lectures focused on simultaneous single-item auctions, with both second-price and first-price payment rules. The toolbox we developed for bounding the POA in auctions can also be applied to many other auction formats. In this lecture we revisit an old setting (from Lecture #4) and derive some new insights via this toolbox.

Recall the setting of *multi-unit auctions with downward-sloping valuations.*

Scenario #4:

- A set m *identical* items.
- Each bidder i has a private *marginal valuation* $\mu_i(j)$ for a jth item. Thus, bidder i's total valuation for ℓ units is $v_i(\ell) := \sum_{j=1}^{\ell} \mu_i(j)$.
- Valuations are *downward-sloping*, meaning that $\mu_i(1) \geq \mu_i(2) \geq \cdots \geq \mu_i(m)$ for every i. Thus, additional units provide diminishing returns.

Recall that this is the special case of gross substitutes valuations (Lecture #5) where all items are identical. Recall also that, in this scenario, the welfare-maximizing allocation can be computed using a simple greedy algorithm. First, identify the set of the top m $\mu_i(j)$'s. Since the bids are downward-sloping, each bidder i will have a prefix of its first x_i reported marginal valuations in this set. A straightforward exchange argument shows that giving x_i items to each bidder i maximizes the welfare with respect to the reported marginal valuations.

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

As always, the welfare-maximizing allocation rule can be extended to a dominant-strategy incentive compatible (DSIC) mechanism by charging VCG payments. Our concern back in Lecture #4 was the clinching auction, an ascending ex post incentive-compatible (EPIC) implementation of the VCG mechanism (both allocation and payments). Recall that these mechanisms charge different prices for different items. For example, in the clinching auction, there is a current price that rises over time, and each bidder clinches items at successively higher prices.

We've already previously discussed the *uniform price mechanism* for multi-unit auctions, which naively equates supply (i.e., m) with demand. Formally, the allocation rule of this mechanism is the same as that of the VCG mechanism — award the m items to the highest m $b_i(j)$'s — and charge a common price for every item, equal to the $(m+1)$th highest $b_i(j)$. Another way to think about the uniform price mechanism is that it implements the lowest Walrasian equilibrium (i.e., market-clearing price) with respect to the reported valuations. The Walrasian equilibria in this scenario correspond to the prices between the mth highest and the $(m+1)$th highest $b_i(j)$'s. Note that, in contrast to simultaneous single-item auctions, the bid space of the uniform price mechanism is the same as the valuation space (all downward-sloping valuations) — its "simplicity" is of a different type.

The uniform price mechanism and the VCG mechanism coincide when all bidders are unit-demand or when all bidders are additive. Even with a mixture of unit-demand and additive bidders, however, the uniform price mechanism uses a simpler payment rule than the VCG mechanism, and as a result the DSIC guarantee is lost. Let's recall a concrete illustration of this.

Example 1.1 (Demand Reduction) Suppose there are $m = 2$ identical items and $n = 2$ bidders. Suppose the true valuations are $\mu_1(1) = \mu_1(2) = 3$, $\mu_2(1) = 2$, and $\mu_2(2) = 0$. Thus, the first bidder is additive and the second bidder is unit demand. Suppose the second bidder bids truthfully. If the first bidder also bids truthfully, then it wins both items at a price of 2 each, and its utility is 2. (By contrast, the VCG mechanism would only charge 2 for both items combined.) The first bidder is better off bidding $b_1(1) = 3$ and $b_1(2) = 0$ — it gets one item for free, for a utility of 3.

Example 1.1 illustrates *demand reduction* — the first bidder misreports to receive fewer items but at a much cheaper price. It shows that the uniform pricing rule is not DSIC — in effect, it has the wrong payment rule for the welfare-maximization allocation.

We have now have the tools to understand the consequences of adopting the simpler uniform-price payments rule — to analyze the POA of the uniform-price mechanism.

2 The POA of Demand Reduction

We prove the following guarantee for the uniform price mechanism.

Theorem 2.1 ([1, 2]) *For every (possibly correlated) prior distribution* $\mathbf{F}$ *over downward-sloping valuations, the expected welfare of every Bayes-Nash equilibrium that satisfies a no overbidding condition is at least 25% that of the expected optimal welfare.*

As in our previous Bayes-Nash POA examples, Theorem 2.1 boils down to proving the following smoothness condition: for every valuation profile μ, there exist hypothetical deviations $\mathbf{b}_1^*(\mu_1), \ldots, \mathbf{b}_n^*(\mu_n)$ such that, for every bid profile $\mathbf{b}$,

$$\sum_{i=1}^{n} u_i(\mathbf{b}_i^*(\mu_i), \mathbf{b}_{-i}) \geq \frac{1}{2} \cdot \text{OPT welfare}(\mu) - \sum_{i=1}^{n} \sum_{j=1}^{x_i(\mathbf{b})} b_i(j), \tag{1}$$

where $x_i(\mathbf{b})$ denotes the number of items that bidder i receives in an outcome $\mathbf{b}$. The proof that (1) implies Theorem 2.1 is exactly the same as the one we used to prove a POA bound of $\frac{1}{2}$ for S2A's in Lecture #15, and we leave those details as an exercise. The no overbidding condition is the same as the one for S2A's, stating that the expected welfare of every bidder i is at most the expected value of its winning bids:

$$\mathbf{E}_{\mu \sim \mathbf{F}}\left[\sum_{j=1}^{x_i(\sigma(\mathbf{v}))} \mu_i(j)\right] \leq \mathbf{E}_{\mu \sim \mathbf{F}}\left[\sum_{j=1}^{x_i(\sigma(\mathbf{v}))} b_i(j)\right].$$

We obtain a bound of $\frac{1}{4}$ here, as opposed to the $\frac{1}{2}$ bound for S2A's, because of the extra coefficient of $\frac{1}{2}$ on the first term of the right-hand side of (1). We obtain this bound for all correlated valuation distributions, as opposed to merely all product valuation distributions for S2A's, because in (1) each hypothetical deviation $b_i^*(\mu_i)$ depends only on μ_i's valuation, and not on the full valuation profile μ.[1]

Proving (1) boils down to exhibiting good hypothetical deviations $\mathbf{b}_1^*(\mu_1), \ldots, \mathbf{b}_n^*(\mu_n)$. Let's try adapt the ideas that have worked in the past. For S2A's (Lecture #14), we defined b_i^* as follows: compute the optimal allocation for the full valuation profile, and go "all in" for the bundle of items that i gets in this allocation. In the present context, this initial idea translates to the deviation

$$b_{ij}^* = \begin{cases} \mu_i(j) & \text{if } j \leq x_i^*(\mu) \\ 0 & \text{otherwise,} \end{cases} \tag{2}$$

where $x_i^*(\mu)$ denotes the number of items that i gets in a welfare-maximizing allocation for the valuation profile μ. Leaving aside the criticism that this deviation demands on the full profile μ, and not just the valuation μ_i, there is a more basic problem: there are bid profiles $\mathbf{b}$

[1] Recall the discussion of the different extension theorems at the end of last lecture. When a hypothetical deviation depends only on the deviator's valuation, it can execute it directly. The corresponding extension theorem relies only linearity of expectation and applies to all correlated distributions. When a hypothetical deviation depends also on others' valuations, which are unknown to the deviator, we resort to the "doppelganger trick" and consider executing the deviation with randomly sampled valuations. The resulting extension theorem applies only to product valuation distributions.

We've only seen good POA bounds for correlated valuation distributions once before, for first-price single-item auctions (Lecture #16).

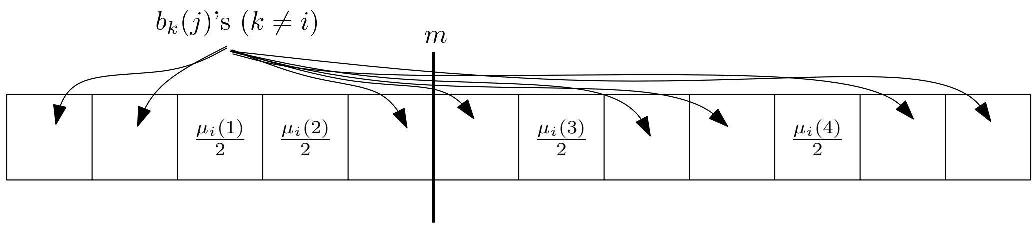

Figure 1: Winning bids for the strategy $b_i^* = \mu_i/2$.

such that (1) does not hold. Intuitively, this choice of b_i^* is too aggressive, and the left-hand side of (1) might be zero while the right-hand side is arbitrarily large (see Exercises). We encountered a similar setback when we studied S1A's last lecture. The obvious second idea is to reuse the solution for S1A's from last lecture: just cut the definition in (2) in half. This choice of $b_1^*, \ldots, b_n^*$ satisfies (1), as the following derivation shows, but we're still left with the drawback that these deviations depend on the full valuation profile μ, and so are useful only for deriving Bayes-Nash POA bounds with respect to product distributions. Our final definition of $b_1^*, \ldots, b_n^*$ takes its cue from our single-item first-price auction analysis, and implements the "bid half your value" idea in the obvious way: set $b_i^* = \mu_i/2$ for every i.[2] We now proceed to verify (1) for the deviations, which then implies Theorem 2.1 via the appropriate extension theorem.

Fix an arbitrary bid vector $\mathbf{b}$. Consider a term $u_i(\mathbf{b}_i^*, \mathbf{b}_{-i}) = u_i(\frac{\mu_i}{2}, \mathbf{b}_{-i})$ on the left-hand side of (1). Some prefix of i's bids are in the top m (out of the m $\frac{\mu_i(j)}{2}$'s and the $(n-1)m$ $b_k(j)$'s for $k \neq i$), and the rest are not. The price of each item that i gets equals the $(m+1)$th highest bid. See also Figure 1.

For technical convenience, we imagine that there is an $(n+1)$th bidder who submits the bids $\mathbf{b}_i$ that bidder i originally submitted in $\mathbf{b}$. Bidder i's utility in the new outcome $(\mathbf{b}_i^*, \mathbf{b})$ is only lower than in $(\mathbf{b}_i^*, \mathbf{b}_{-i})$ — the two effects of the new m bids by the extra bidder are to cause bidder i to lose some items that it previously won (for nonnegative utility), and to increase the price of the items that i still wins.

For each bidder i, let x_i^* and x_i denote the number of items that i gets in an optimal allocation for the valuation profile μ and in the outcome $(\mathbf{b}_i^*, \mathbf{b})$, respectively. Each item j that bidder i wins in $(\mathbf{b}_i^*, \mathbf{b})$ contributes at least $\mu_i(j) - \frac{\mu_i(j)}{2} = \frac{\mu_i(j)}{2}$ to its utility, where we are using that the price for an item is always at most the winning bid for it. Thus, in the happy event that $x_i \geq x_i^*$, we have

$$u_i(\tfrac{\mu_i}{2}, \mathbf{b}_{-i}) \geq \frac{1}{2} \sum_{j=1}^{x_i^*} \mu_i(j). \qquad (3)$$

[2] Given that this deviation is so simple, why didn't we just do it for simultaneous single-item auctions? The reason is that for S2A's and S1A's, $v_i/2$ is not generally an allowable bid. Each bidder is required to bid separately on each item — equivalently, each is forced to declare an additive bid despite having a non-additive valuation.

The right-hand side of the inequality is i's contribution to OPT welfare(μ), so this lower bound is even stronger than we we're shooting for in (1). When $x_i < x_i^*$, we can still write

$$u_i(\tfrac{\mu_i}{2}, \mathbf{b}_{-i}) \geq \frac{1}{2}\sum_{j=1}^{x_i} \mu_i(j) + \sum_{j=x_i+1}^{x_i^*} \left(\frac{1}{2}\mu_i(j) - b^{(m-j+1)}\right), \tag{4}$$

where $b^{(\ell)}$ denote the ℓ biggest of the $b_k(j)$'s. The first term on the right-hand side of (4) is, by (3), a lower bound on i's utility in the outcome $(\frac{\mu_i}{2}, \mathbf{b})$. We claim that each summand in the second term is nonpositive; this verifies (4). To see this, observe that none of bidder i's bids $\frac{\mu_i(x_i+1)}{2}, \ldots, \frac{\mu_i(x_i^*)}{2}$ are amongst the top m bids. Thus, even the smallest $(x_i^* - x_i)$ of the top $m - x_i$ $b_k(j)$'s are bigger.

For convenience, we simplify the inequality in (4) and make it even more true by subtracting additional terms from the right-hand side:

$$u_i(\tfrac{\mu_i}{2}, \mathbf{b}_{-i}) \geq \frac{1}{2}\sum_{j=1}^{x_i^*} \mu_i(j) - \sum_{j=1}^{x_i^*} b^{(m-j+1)}. \tag{5}$$

Write X_i^* for $\sum_{k \leq i} x_i^*$. We complete the verification of (1) by summing over all bidders i and manipulating the final term; we derive

$$\begin{aligned}
\sum_{i=1}^{n} u_i(\tfrac{\mu_i}{2}, \mathbf{b}_{-i}) &\geq \frac{1}{2}\text{OPT welfare}(\mu) - \sum_{i=1}^{n}\sum_{j=1}^{x_i^*} b^{(m-j+1)} \\
&\geq \frac{1}{2}\text{OPT welfare}(\mu) - \sum_{i=1}^{n}\sum_{j=X_{i-1}^*+1}^{X_i^*} b^{(m-j+1)} \\
&= \frac{1}{2}\text{OPT welfare}(\mu) - \sum_{j=1}^{m} b^{(m-j+1)}. \qquad (6)
\end{aligned}$$

Since the final terms of (1) and (5) are both equal to the sum of the m highest $b_k(j)$'s, this verifies (1) (and implies Theorem 2.1).

The POA bound in Theorem 2.1 can be extended in several ways [1, 2]. For example, the bound of $\frac{1}{4}$ extends to the case where both the valuation space and the bid space are the set of gross substitutes valuations (Lecture #5), where the mechanism outputs a Walrasian equilibrium with respect to the reported valuations.[3] This POA bound can be interpreted as a welfare guarantee for the venerable Kelso-Crawford auction. As noted in Lecture #5, straightforward bidding in that auction leads to a (fully efficient) Walrasian equilibrium, but is not generally optimal for players. This POA bound of $\frac{1}{4}$ limits the damage of strategic behavior in the Kelso-Crawford auction when bidders have GS valuations and do not overbid.

[3] This requires generalizing the arguments in (3)–(6), but this is not overly difficult given all the structure of GS valuations that we identified in several previous lectures.

References

[1] M. Babaioff, B. Lucier, N. Nisan, and R. Paes Leme. On the efficiency of the Walrasian mechanism. In *Proceedings of the 15th ACM Conference on Electronic Commerce (EC)*, 2014.

[2] B. de Keijzer, E. Markakis, G. Schäfer, and O. Telelis. Inefficiency of standard multi-unit auctions. In *Proceedings of the 18th Annual European Symposium on Algorithms (ESA)*, pages 385–396, 2013.