CS364B: Frontiers in Mechanism Design
Lecture #17: Part II: Beyond Smoothness and XOS Valuations *

Tim Roughgarden†

March 5, 2014

1 Subadditive Valuations

1.1 The Setup

In this lecture we study a scenario that generalizes almost all of the ones that we've studied in the course.

Scenario #9:

- A set U of m non-identical items.
- Each bidder i has a private valuation $v_i : 2^U \rightarrow \mathcal{R}^+$ that is *subadditive*, meaning that for every pair of sets $S, T \subseteq U$,

$$v_i(S \cup T) \leq v_i(S) + v_i(T). \tag{1}$$

  As always, we also assume that every valuation satisfies $v_i(\emptyset) = 0$ and is monotone (i.e., $S \subseteq T$ implies $v_i(S) \leq v(T)$).

Subadditivity is yet another way to formalize the idea that items are not complements — that getting some items don't suddenly make other items more valuation. Of all the articulations of this idea that we've seen, subadditivity is the most general; see Figure 1.

Proposition 1.1 *The set of subadditive valuations strictly contains the set of XOS valuations.*

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

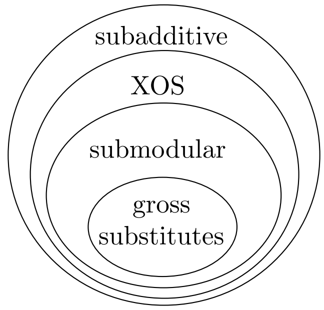

Figure 1: Subadditive valuations are the most general valuation class we have seen.

Proof: For containment, fix an item set U. First, let v be an XOS valuation, meaning there are additive valuations $a^1, \ldots, a^r$ on U such that

$$v(S) = \max_{i=1}^{r} \sum_{j \in S} a_j^i \tag{2}$$

for every S. We show that v is subadditive. Fix subsets $S, T \subseteq U$. Since v is monotone, we can assume that S and T are disjoint. Let a^ℓ determine the maximum in (2) for $S \cup T$. Since a^ℓ is additive, $a^\ell(S \cup T) = a^\ell(S) + a^\ell(T)$. By (2), $v(S) + v(T) \geq a^\ell(S) + a^\ell(T)$, as desired.

We leave as an exercise a proof that the containment is strict. ■

Welfare maximization for bidders with subadditive valuations appears to be strictly harder than for bidders with submodular or XOS valuations. No constant-factor DSIC mechanism for bidders with subadditive valuations is known. With a priori known valuations, the best polynomial-time approximation algorithm uses demand oracles and has guarantee of 2 [3]. No simple constant-factor approximation algorithm is known. Subadditive valuations are close to the most general valuation class for which computationally tractable constant-factor approximations are known. What happens if just sell items using simultaneous single-item auctions? What is the POA of S1A's and S2A's when bidder have subadditive valuations?

1.2 Smoothness Gets Stuck

At this point in the course, you've been trained to immediately try to prove a suitable smoothness condition. For subadditive valuations, however, smoothness arguments seem to get stuck at a guarantee of $\Theta(1/\log m)$. To see why, recall that a smoothness condition has the following 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 $\mathbf{b}$, a certain inequality holds. In effect, the deviations $\mathbf{b}_1^*(\mathbf{v}), \ldots, \mathbf{b}_n^*(\mathbf{v})$ are required to achieve some type of guarantee for worst-case bidding behavior of the other bidders. With S2A's with XOS valuations, for example, we defined $\mathbf{b}_i^*(\mathbf{v})$ by targeting the bundle $S_i^*(\mathbf{v})$ that i gets in an optimal allocation for $\mathbf{v}$. By "targeting" a bundle S we mean that the sum of i's bids on items $j \in S$ is comparable to its value for S. The XOS assumption allows us to target a fixed bundle (like $S_i^*(\mathbf{v})$ without overbidding on any subsets of that bundle. Avoiding overbidding is crucial to obtaining a utility guarantee (despite worst-case bidding behavior by others) in cases where, after deviating to $(\mathbf{b}_i^*, \mathbf{b}_{-i})$, the bidder i only receives a strict subset of the bundle $S_i^*(\mathbf{v})$ it was targeting. When a valuation v is merely subadditive, targeting a specific bundle can require overbidding on some subset of the bundle by a $\log m$ factor (see Exercises). For this reason, extending the smoothness-based POA bounds for S2A's and S1A's from XOS to subadditive valuations results in a loss of roughly $\log m$ in the bounds [1, 5].

1.3 A Direct Argument

We next show how to bypass the smoothness-based approach to prove the following remarkable guarantees for simultaneous single-item auctions.

Theorem 1.2 ([4]) *For every product valuation distribution over subadditive valuations:*

(a) every Bayes-Nash equilibrium of a S1A has expected welfare at least 50% of the maximum possible;

(b) every Bayes-Nash equilibrium of a S2A that satisfies a no overbidding condition has expected welfare at least 25% of the maximum possible.

We prove part (a); the proof of (b) is along the same lines, with some additional details. Remarkably, the bound in (a) is as good as the best-known approximation algorithm for welfare maximization with subadditive bidder valuations.

If we don't prove a Bayes-Nash POA bound using a smoothness condition, then how would we do it? Recall that a smoothness condition requires that the bid deviations $\mathbf{b}_1^*(\mathbf{v}), \ldots, \mathbf{b}_n^*(\mathbf{v})$ be chosen independently of $\mathbf{b}$ — in effect, the same deviations are used no matter which equilibrium we're arguing about. If all we care about is a Bayes-Nash POA bound and not the smoothness condition per se, then we're free to choose a different collection of bid deviations to bound the expected welfare of each equilibrium. This is how the following analysis proceeds. A similar idea can be used to bound the POA of correlated equilibria in the full-information model [2].

The following key lemma will substitute for a smoothness condition in the proof of Theorem 1.2(a).

Lemma 1.3 *In a S1A with item set U, fix a bidder i with subadditive valuation v_i, a distribution D over the bids $\mathbf{b}_{-i}$ of the other bidders, a subset $S \subseteq U$. There exists a bid vector $\mathbf{b}_i^*$ such that*

$$\mathbf{E}_{\mathbf{b}_{-i}\sim D}[u_i(\mathbf{b}_i^*, \mathbf{b}_{-i})] \geq \frac{1}{2} \cdot v_i(S) - \mathbf{E}_{\mathbf{b}_{-i}\sim D}\left[\sum_{j\in S} \max_{k\neq i} b_{kj}\right]. \quad (3)$$

When we apply Lemma 1.3, the distribution D will be $\sigma_{-i}(\mathbf{v}_{-i})$ in a Bayes-Nash equilibrium $\sigma(\mathbf{v})$, and S will be i's bundle in a hypothetical welfare-maximizing allocation. Note that the hypothetical deviation $\mathbf{b}_i^*$ in Lemma 1.3 is a function of D (in addition to S), and in this sense is not a smoothness condition. We proceed to the very neat proof.

Proof of Lemma 1.3: It is enough to show that a randomly chosen bid vector $\mathbf{b}_i^*$ satisfies (3) in expectation — this implies that there exists a choice of $\mathbf{b}_i^*$ for which (3) holds.

We generate $\mathbf{b}_i^*$ using the following randomized algorithm. First, choose $\mathbf{a}_{-i} \sim D$. Second, set

$$b_{ij}^* = \begin{cases} \max_{k \neq i} a_{kj} & \text{if } j \in S \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

In effect, bidder i simulates the behavior of the other bidders under the bidding distribution D, and bids on each $j \in S$ as if it was the highest other bidder.

To lower bound the expected value (over $\mathbf{b}_i^*$) of the left-hand side of (3), we consider the expected payment and expected welfare of bidder i separately. Its expected payment (over $\mathbf{b}_i^*$ and $\mathbf{b}_{-i}$) is at most the expected sum of its bids (over $\mathbf{b}_i^*$), which by definition is

$$\mathbf{E}_{\mathbf{a}_{-i} \sim D}\left[\sum_{j \in S} \max_{k \neq i} b_{kj}\right],$$

which equals the final term of (3).

Next, by the symmetry of $\mathbf{a}_{-i}$ and $\mathbf{b}_{-i}$, we claim that

$$\mathbf{Pr}_{\mathbf{b}_i^*, \mathbf{b}_{-i}}[i \text{ wins set } A \text{ in } (\mathbf{b}_i^*, \mathbf{b}_{-i})] = \mathbf{Pr}_{\mathbf{b}_i^*, \mathbf{b}_{-i}}[i \text{ wins set } S \setminus A \text{ in } (\mathbf{b}_i^*, \mathbf{b}_{-i})] \tag{5}$$

for every $A \subseteq S$.[1] This follows from the definition of $\mathbf{b}_i^*$: the items of S that i wins are precisely those on which its sample $\mathbf{a}_{-i}$ is bigger than $\mathbf{b}_{-i}$, and the realizations ($\mathbf{a}_{-i} = \mathbf{b}_{-i}^{(1)}, \mathbf{b}_{-i} = \mathbf{b}_{-i}^{(2)}$) and ($\mathbf{a}_{-i} = \mathbf{b}_{-i}^{(2)}, \mathbf{b}_{-i} = \mathbf{b}_{-i}^{(1)}$) are equally likely for every pair $\mathbf{b}_{-i}^{(1)}, \mathbf{b}_{-i}^{(2)}$ of bid vectors.[2]

Equation (5) suggests pairing up the contributions of complementary item sets when computing i's expected welfare. Formally, letting $j^* \in S$ be an arbitrary item of S, i's expected welfare (over $\mathbf{b}_i^*$ and $\mathbf{b}_{-i}$) is

$$\begin{aligned} \sum_{A \subseteq S} \mathbf{Pr}_{\mathbf{b}_i^*, \mathbf{b}_{-i}}[i \text{ wins } S \setminus A\ ]\, v_i(A) &= \sum_{A\,:\, j^* \in A \subseteq S} \mathbf{Pr}_{\mathbf{b}_i^*, \mathbf{b}_{-i}}[i \text{ wins set } S \setminus A \text{ in } (\mathbf{b}_i^*, \mathbf{b}_{-i})]\,(v_i(A) + v_i(S \setminus A)) \\ &\geq v_i(S) \underbrace{\sum_{A\,:\, j^* \in A \subseteq S} \mathbf{Pr}_{\mathbf{b}_i^*, \mathbf{b}_{-i}}[i \text{ wins set } S \setminus A \text{ in } (\mathbf{b}_i^*, \mathbf{b}_{-i})]}_{\mathbf{Pr}[i \text{ wins } j^*]} \\ &= \frac{1}{2} \cdot v_i(S), \end{aligned}$$

where the inequality follows from the subadditivity of v_i the last equation follows the fact that $\mathbf{b}_j^*$ and $\max_{k \neq i} b_{kj}$ are identically distributed.

[1] We can ignore items outside S that i wins (at price 0), with can only contribute additional expected utility.

[2] For simplicity, we are ignoring the possibility of ties.

Summarizing, we've exhibited a distribution over bids $\mathbf{b}_i^*$ such that

$$\mathbf{E}_{\mathbf{b}_i^*,\mathbf{b}_{-i}\sim D}[u_i(\mathbf{b}_i^*,\mathbf{b}_{-i})] \geq \frac{1}{2}\cdot v_i(S) - \mathbf{E}_{\mathbf{b}_{-i}\sim D}\left[\sum_{j\in S}\max_{k\neq i} b_{kj}\right].$$

Hence, there is a choice of $\mathbf{b}_i^*$ satisfying (3), which proves the lemma. ■

We now prove part (a) of Theorem 1.2. It proceeds similarly to our extension theorems for Bayes-Nash equilibria for product distributions (based on the doppelganger trick), with Lemma 1.3 substituting for a smoothness condition, although some of the details differ.

Proof of Theorem 1.2: Let σ denote an arbitrary Bayes-Nash equilibrium. As usual, to minimize notation we write the following derivation for pure Bayes-Nash equilibria. Adding an extra expectation over players' random actions extends the derivation to mixed Bayes-Nash equilibria.

First, we write

$$\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\text{welfare}(\sigma(\mathbf{v}))] = \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^n u_i(\sigma(\mathbf{v}))\right] + \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\underbrace{\sum_{i=1}^n p_i(\sigma(\mathbf{v}))}_{=\sum_{j\in U}\max_{i=1}^n(\sigma_{ij}(v_i))}\right], \tag{6}$$

where $\sigma_{ij}(v_i)$ denotes bidders i's bid on item j when its valuation is v_i.

As always, the next step to derive a lower bound on each bidder's equilibrium through a judicious choice of a hypothetical deviation. Lemma 1.3 is an obvious tool for choosing a deviation. The distribution D over $\mathbf{b}_{-i}$ in Lemma 1.3 naturally corresponds to the equilibrium bids $\sigma_{-i}(\mathbf{v}_{-i})$ of bidders other than i. But what about the set S? To relate the equilibrium welfare to the optimal welfare, the natural choice of S is bundle i gets in a welfare-maximizing allocation. But this does make sense: when bidder i contemplates deviations, it knows only its own valuation v_i and not the others $\mathbf{v}_{-i}$, so it is does not have enough information to compute a welfare-maximizing allocation. As in our previous extension theorems, we salvage this idea using the doppelganger trick.

Formally, for each bidder i and valuation v_i, we define the (mixed) deviation $\mathbf{b}_i^*$ according to the following randomized algorithm:

1. Sample doppelganger valuations $\mathbf{w}\sim\mathbf{F}$.[3]

2. Let $S_i^*(v_i,\mathbf{w}_{-i})$ denote the bundle i receives in a welfare-maximizing allocation for the valuation profile $(v_i,\mathbf{w}_{-i})$.

3. Bid $b_i^*(v_i,\mathbf{w}_{-i},\sigma_{-i}(\mathbf{v}_{-i}))$ as in Lemma 1.3, with target bundle $S = S_i^*(v_i,\mathbf{w}_{-i})$ and opposing bid distribution $D = \{\sigma_{-i}(\mathbf{v}_{-i})\}_{\mathbf{v}_{-i}\sim\mathbf{F}_{-i}}$.

[3]Since $\mathbf{F}$ is a product distribution, there is no need to condition on v_i, and $(v_i,\mathbf{w}_{-i})$ is distributed according to $\mathbf{F}$.

Since σ is a Bayes-Nash equilibrium, for every i and v_i, the unilateral deviation $\mathbf{b}_i^*(v_i, \mathbf{w}_{-i}, \sigma_{-i}(\mathbf{v}_{-i}))$ can only decrease i's expected utility, where the expectation is over the randomness in others' valuations and in i's action:

$$\begin{aligned} \mathbf{E}_{\mathbf{v}_{-i}\sim\mathbf{F}_{-i}}[u_i(\sigma(\mathbf{v}))] &\geq \mathbf{E}_{\mathbf{v}_{-i}\sim\mathbf{F}_{-i},\mathbf{w}\sim\mathbf{F}}[u_i(b_i^*(v_i, \mathbf{w}_{-i}, \sigma_{-i}(\mathbf{v}_{-i})), \sigma_{-i}(\mathbf{v}_{-i}))] \qquad (7) \\ &\geq \mathbf{E}_{\mathbf{w}\sim\mathbf{F}}\left[\frac{1}{2}v_i(S_i^*(v_i, \mathbf{w}_{-i}))\right] - \mathbf{E}_{\mathbf{v}_{-i}\sim\mathbf{F}_{-i}}\left[\sum_{j\in S_i^*(v_i,\mathbf{w}_{-i})} \max_{k\neq i} \sigma_{kj}(v_k)\right], \qquad (8) \end{aligned}$$

where the second inequality follows from the guarantee provided by Lemma 1.3.

Next, we integrate the inequality (7)–(8) over $v_i \sim F_i$ and sum over all the bidders i. Consider the term $\mathbf{E}_{\mathbf{w}\sim\mathbf{F}}[v_i(S_i^*(v_i, \mathbf{w}_{-i}))]$ in (8). Since $S_i^*(v_i, \mathbf{w}_{-i})$ is i's contribution to the optimal welfare when the valuation profile is $(v_i, \mathbf{w}_{-i})$, this term is i's expected contribution to the optimal welfare when its valuation in v_i. After integrating over $\mathbf{v}_{-i} \sim F_i$, the term becomes i's contribution to the expected optimal welfare. Thus, summing over all bidders i yields the following:

$$\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^n u_i(\sigma(\mathbf{v}))\right] \geq \frac{1}{2} \cdot \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\text{OPT welfare}(\mathbf{v})] - \sum_{i=1}^n \mathbf{E}_{\mathbf{v},\mathbf{w}\sim\mathbf{F}}\left[\sum_{j\in S_i^*(v_i,\mathbf{w}_{-i})} \max_{k\neq i} \sigma_{kj}(v_k)\right].$$

If you think about it, we are free to replace the sum over $j \in S_i^*(v_i, \mathbf{w}_{-i})$ by a sum over $j \in S_i^*(\mathbf{w})$. After all, the summands $\max_{k\neq i} \sigma_{kj}(v_k)$ are just numbers (for fixed $\mathbf{v}_{-i}$), independent of i's valuation, and $(v_i, \mathbf{w}_{-i})$ and $(w_i, \mathbf{w}_{-i})$ are identically distributed. This, with linearity of expectation, gives

$$\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^n u_i(\sigma(\mathbf{v}))\right] \geq \frac{1}{2} \cdot \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\text{OPT welfare}(\mathbf{v})] - \mathbf{E}_{\mathbf{v},\mathbf{w}\sim\mathbf{F}}\left[\sum_{i=1}^n \sum_{j\in S_i^*(\mathbf{w})} \max_{k\neq i} \sigma_{kj}(v_k)\right]. \qquad (9)$$

The bundles $\{S_i^*(\mathbf{w})\}_{i=1}^n$ are by definition disjoint for every $\mathbf{w}$, so

$$\begin{aligned} \sum_{i=1}^n \sum_{j\in S_i^*(\mathbf{w})} \max_{k\neq i} \sigma_{kj}(v_k) &\leq \sum_{j\in U} \max_{k=1}^n \sigma_{kj}(v_k) \\ &= \sum_{i=1}^n p_i(\sigma(\mathbf{v})) \qquad (10) \end{aligned}$$

for every $\mathbf{w}$ and $\mathbf{v}$, where the equality follows from the first-price payment rule. Substituting (10) into (9) and integrating out over $\mathbf{w}$ yields

$$\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^n u_i(\sigma(\mathbf{v}))\right] \geq \frac{1}{2} \cdot \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\text{OPT welfare}(\mathbf{v})] - \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^n p_i(\sigma(\mathbf{v}))\right], \qquad (11)$$

and combining (11) with (6) proves the theorem. ■

Lemma 1.3 can be modified to hold for S2A's as well [4]. Combining with the proof above and imposing the same no overbidding condition as in Lecture #15 yields part (b) of Theorem 1.2. The POA bound is only $\frac{1}{4}$ because, with second-price payment rules, the revenue in (6) cannot be canceled with sum of winning bids in (11) (which might be much larger).

References

[1] K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 700–709, 2011.

[2] Paul Dütting, Monika Henzinger, and Martin Starnberger. Valuation compressions in vcg-based combinatorial auctions. In *Proceedings of the 7th International Workshop on Internet and Network Economies (WINE)*, pages 146–159, 2013.

[3] U. Feige. On maximizing welfare when utility functions are subadditive. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 41–50, 2006.

[4] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient. In *45th ACM Symposium on Theory of Computing (STOC)*, pages 201–210, 2013.

[5] 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.