

CS364B: Frontiers in Mechanism Design

Lecture #7: Submodular Valuations*

Tim Roughgarden[†]

January 29, 2014

1 Introduction: DSIC Approximation Mechanisms

With this lecture we commence Part II of the course. We continue to pursue very strong incentive guarantees (i.e., DSIC mechanisms) in scenarios where the underlying welfare-maximization problem is NP-hard. This turns out to be a very difficult quest; the goal of this part of the course is to highlight the challenges involved and survey what few design techniques and positive results are known. Subsequent parts of the course will relax the DSIC requirement to obtain stronger results and simpler mechanisms.

2 Submodular Valuations

We next take a seemingly small step beyond gross substitutes valuations, but this step is enough to introduce a host of fundamental complications.

Scenario #6:

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

$$v_i(T \cup \{j\}) - v_i(T) \leq v_i(S \cup \{j\}) - v_i(S). \quad (1)$$

See also Figure 1. As always, we also assume that every valuation satisfies $v_i(\emptyset) = 0$ and is monotone (i.e., $S \subseteq T$ implies $v_i(S) \leq v_i(T)$).

*©2014, Tim Roughgarden.

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




Figure 1: Submodularity: adding j to T cannot increase the valuation more than adding j to S .

The submodularity condition in (1) is a form of diminishing returns: the more stuff you’ve already got, the less additional value you get from new stuff. A non-example is a “single-minded” bidder with valuation $v(S) = 1$ if $S = U$ and 0 otherwise (for $|U| \geq 2$) — the marginal value of an item is 0 unless it “completes the set,” in which case its marginal value is 1.

It is not immediately clear for the gross substitutes condition relates to the submodularity condition (1). Both conditions seems to be saying the same thing, in different ways: as items are taken away from you, your value for the items you still have only goes up. We next result shows that the gross substitutes conditions is at least as strong as submodularity.

Proposition 2.1 ([3]) *If a valuation v satisfies the gross substitutes condition, then v is submodular.*

Proof: Fix $S \subseteq T$ and an item j . The only interesting case is when $j \notin T$. To invoke the gross substitutes conditions, we need to propose some price vectors over items and argue about the corresponding demand sets $D(\cdot)$. With an eye toward (1), we begin with a price vector designed to get the set $T \cup \{j\}$ into the picture. Define $p(j) = 0$ on $T \cup \{j\}$ and $+\infty$ (or sufficiently large) everywhere else (Figure 2(a)). Since v is monotone, $T + \{j\} \in D(\mathbf{p})$.

Next we need to determine the relationship between $v(T \cup \{j\})$ and $v(T)$. To do this, consider raising the price on item j . The Utility of all bundles that include j decrease, while the utility of bundles that exclude j stay the same. Thus, there will be a price $\epsilon \geq 0$ at which the utility of a subset of T equals that of $T \cup \{j\}$. Let \mathbf{p}^ϵ denote the price vector equal to \mathbf{p} except $p(j) = \epsilon$ (Figure 2(b)). By construction, $T \cup \{j\}$ belongs to $D(\mathbf{p}^\epsilon)$ as does at least one subset of T . If we raise the price of j any further, then all bundles that include j (including $T \cup \{j\}$) will drop out of the demand set. The gross substitutes condition now implies that $T \in D(\mathbf{p}^\epsilon)$: if not, after raising the price of j further the demand set will contain only strict subsets of T , violating GS (since $T \cup \{p\} \in D(\mathbf{p})$). Since $T, T \cup \{j\} \in D(\mathbf{p}^\epsilon)$, we have

$$v(T \cup \{j\}) - v(T) = \mathbf{p}^\epsilon(j) = \epsilon.$$

To zoom in on the sets S and $S \cup \{j\}$, we raise the prices on the items $T \setminus S$ to $+\infty$; call the resulting price vector \mathbf{q}^ϵ (Figure 2(c)). Clearly, only subsets of $S \cup \{j\}$ belong to $D(\mathbf{q}^\epsilon)$. Since $T \cup \{j\} \in D(\mathbf{p}^\epsilon)$ and only prices on the items of $T \setminus S$ have increased, the gross




Figure 2

substitutes condition implies that $S \cup \{j\} \in D(\mathbf{q}^\epsilon)$. In particular, the utility of $S \cup \{j\}$ is at least that of S under the price vector \mathbf{q}^ϵ , so

$$v(S \cup \{j\}) - v(S) \geq \mathbf{q}^\epsilon(j) = \epsilon = v(T \cup \{j\}) - v(T),$$

which completes the proof. ■

Does the converse of Proposition 2.1 hold? When there are only two items A and B , the answer is yes: both the gross substitutes condition and submodularity are equivalent to the condition that v is subadditive, meaning $v(\{A, B\}) \leq v(\{A\}) + v(\{B\})$ (see Exercises). With three or more items, however, the gross substitutes condition is a strictly stronger requirement.

Example 2.2 (Submodularity Does Not Imply Gross Substitutes) Consider a set $U = \{a, b, c\}$ of three items. Define $v_a = 1$, $v_b = 2$, $v_c = 3$, and

$$v(S) = \min \left\{ \sum_{j \in S} v_j, 3 \right\}. \quad (2)$$

This type of valuation is sometimes called “budgeted additive”. It is submodular — the marginal value of an item j is the minimum of j and the “residual capacity,” which is decreasing the the number of other items. To see that it fails the gross substitutes condition, observe that at the price vector $\mathbf{p} = (\frac{1}{2}, \frac{1}{2}, 2)$, the demand set $D(\mathbf{p})$ contains only $\{a, b\}$. After raising the price on b to 2, however, the bidder wants to relinquish the good a : the unique demanded set is $\{c\}$.

Conceptually, submodularity is a relatively modest generalization of the gross substitutes condition. But this small step is enough to cross the tractability frontier for computing a welfare-maximizing allocation.

Proposition 2.3 *Computing a welfare-maximizing allocation for bidders with budgeted additive valuations is NP-hard, even when there are only two bidders.*

As in (2), a budgeted additive valuation has the form $v(S) = \min\{B, \sum_{j \in S} v_j\}$, where v_1, \dots, v_m and B are nonnegative numbers. Every such valuation is submodular and can be described using a linear (in m) number of parameters. Proposition 2.3 follows from an easy reduction from Partition (see Exercises). Stronger intractability results, which rule out constant-factor approximation for constants sufficiently close to 1, are also known [1].

3 Part II Goal: DSIC Approximation Mechanisms

Generalizing beyond gross substitutes valuations is a well motivated goal: preferences of bidders in real combinatorial auctions need not be well approximated by such valuations. Proposition 2.3 shows that even a small step past gross substitutes requires redefining our goals — we can no longer get everything we want. To recap, here are the properties that guided us in Part I of the course:

1. **(Pt I Incentive guarantee.)** DSIC for direct-revelation mechanisms; EPIC for iterative auctions.
2. **(Pt I Performance guarantee.)** Direct revelation/sincere bidding yields a welfare-maximizing outcome.
3. **(Pt I Tractability guarantee.)** Simple and/or polynomial running time.

These are essentially the strongest guarantees of these types that we could hope for. Recall also that the design problem is at least as hard for EPIC iterative auctions as for DSIC direct-revelation auctions. Indeed, we will not see another EPIC iterative auction in the rest of the course — essentially none are known for the scenarios and desiderata that we consider.

Proposition 2.3 teaches us that, beyond gross substitutes, the second and third guarantees are fundamentally in conflict (assuming $P \neq NP$), even if we drop the first guarantee entirely. In fact, beyond gross substitutes, few fully satisfying positive results are known (and in some cases we can prove that none are possible). Parts II–IV of the course explore the Pareto frontier of the state-of-the-art solutions.

In Part II of the course we’ll make only the obviously necessary concession to the second property.

1. **(Pt II Incentive guarantee.)** DSIC. We’ll consider only direct-revelation mechanisms in this part of the course.
2. **(Pt II Performance guarantee.)** Direct revelation yields an allocation that *approximately* maximizes welfare over all feasible allocations.
3. **(Pt II Tractability guarantee.)** Simple and/or polynomial running time.

There is a deep and mature literature that studies approximation algorithms for NP-hard problems, including for welfare maximization for bidders with different types of valuations. We'll see that there are polynomial-time algorithms with good approximation guarantees for several cases of interest, including bidders with submodular valuations. Unfortunately, there are far fewer positive results for DSIC approximation mechanisms.¹ Moreover, the positive results known are generally achieved by very complex and impractical (though still polynomial time) mechanisms. These drawbacks motivate relaxing the DSIC guarantee to a weaker incentive guarantee; we do this in Parts III and IV of the course to obtain stronger approximation guarantees and simpler mechanisms.²

4 $\frac{1}{2}$ -Approximation Algorithm for Welfare-Maximization

Before trying to design any DSIC approximation mechanisms, let's begin with the simpler but still non-trivial problem of designing a good approximation algorithm for welfare maximization. For bidders with submodular valuations, a simple greedy algorithm gives a $\frac{1}{2}$ -approximation [4]; see the Exercises. Here, we show that the Kelso-Crawford auction (Lecture #5) also gives a $\frac{1}{2}$ -approximation for the problem.

Theorem 4.1 ([2]) *If all bidders have submodular valuations and bid sincerely, then the KC auction terminates with an allocation with welfare at least 50% of that of an optimal allocation (up to $m\epsilon$).*³

With gross substitutes valuations, the KC auction terminates at Walrasian equilibrium and obtains the maximum-possible welfare (up to $n\epsilon$). Theorem 4.1 asserts that this welfare guarantee approximately carries over to bidders with submodular valuations, even though there be no Walrasian equilibria.

In Theorem 4.1, we are using the KC auction purely as an algorithm, and make no claims about the incentives of bidding sincerely. As an algorithm, the KC auction uses demand queries and runs in pseudo-polynomial time.⁴

We next review the salient properties of the Kelso-Crawford auction; Lecture #5 has the full description. We introduced the KC auction as the natural extension of the Crawford-Knoer (CK) auction from unit-demand bidders to bidders with general valuations. The first important property of the auction is that there are no bid withdrawals; the only way a bidder relinquishes an item is when it is outbid by a different bidder. A consequence is that an

¹Of particular interest are mechanisms that guarantee welfare within a constant factor of the maximum possible.

²Here by “complex” and “simple” we're referring to the description of the mechanism, and not to the cognitive burden faced by participants. Complex mechanisms can be simple for bidders (if DSIC, for example) while simple mechanism can be difficult to reason about (e.g., a first-price auction).

³If the guarantee of 50% seems unimpressive, keep in mind that for many NP-hard valuations, including welfare maximization with sufficiently general valuations, no polynomial-time constant-factor approximation algorithms exist (assume $P \neq NP$). We advocate interpreting a constant-factor guarantee as evidence that an optimization problem is “tractable to approximate well.”

⁴The greedy algorithm described in the exercises uses only value queries and runs in polynomial time.

item goes unsold only if no bidder ever bid for it, in which case its price at termination is 0. The definition of the gross substitutes condition ensures that no bidder wants to relinquish a good in any case; when this condition is violated, as with submodular valuations, a bidder might wish for the opportunity to withdraw a previous bid.

The second key property of the KC auction is that, in an iteration where bidder i bids sincerely, it bids for the utility-maximizing set T_i of new items given the current prices and the items S_i already assigned to it. Precisely, if the current prices are \mathbf{p} , T_i maximizes $v_i(S_i \cup T) - \sum_{j \in S} p(j) - \sum_{j \in T} (p(j) + \epsilon)$ over $T \subseteq U \setminus S_i$. The items of T_i are taken from their current owners, reassigned to i , and their prices are incremented by ϵ .

Our proof of Theorem 4.1 begins with the following key lemma.

Lemma 4.2 *If all bidders have submodular valuations and bid sincerely, then the KC auction terminates with an allocation in which every bidder has non-negative utility.*

The guarantee in Lemma 4.2 might seem weak, but let's not forget that we've already seen an example where it fails: in Lecture #5, when we motivated the gross substitutes condition, we observed that a (non-submodular) single-minded bidder that bids sincerely in the KC auction might end up with negative utility.

Proof of Lemma 4.2: We prove that following invariant of the KC auction with submodular valuations and sincere bidding:

Invariant: for every bidder i with assigned items S_i and for every subset $T \subseteq S_i$,

$$v_i(T) \geq \sum_{j \in T} p(j), \quad (3)$$

where \mathbf{p} is the current set of item prices.

This invariant implies the lemma (take $T = S_i$ for each i at termination). The base case, when $S_i = \emptyset$ for each i , is trivial.

For the inductive step, consider an iteration of the KC auction in which the prices are p and bidder i has the items S_i and then acquires the set T_i of additional items. Since i bids sincerely, the set T_i maximizes

$$v_i(S_i \cup T) - \sum_{j \in S_i} p(j) + \sum_{j \in T} (p(j) + \epsilon) \quad (4)$$

over all $T \subseteq U \setminus S_i$.

Bidders other than i may lose some of their items, but the prices on the items they retain stays the same, and so the invariant (3) continues to hold for all subsets of the items they still possess. To prove that the invariant still applies to bidder i , suppose for contradiction that there is a set $A \subseteq S_i \cup T_i$ with

$$v_i(A) < \sum_{j \in A \cap S_i} p(j) + \sum_{j \in A \cap T_i} (p(j) + \epsilon). \quad (5)$$




Figure 3

Write X and Y for $A \cap S_i$ and $A \cap T_i$, respectively (Figure 3). By the inductive hypothesis,

$$v_i(X) \geq \sum_{j \in X} p(j). \quad (6)$$

Label the items of $Y = \{1, 2, \dots, \ell\}$ and imagine adding them to X one by one. Comparing (5) and (6), we eventually get to an item j whose addition increases the left-hand side less than the right-hand side:

$$v_i(X \cup \{1, 2, \dots, j\}) - v_i(X \cup \{1, 2, \dots, j-1\}) < p(j) + \epsilon.$$

Since v_i is submodular, the same is only more true with respect to the bigger set $S_i \cup T_i - \{j\}$:

$$v_i(S_i \cup T_i) - v_i(S_i \cup T_i \setminus \{j\}) < p(j) + \epsilon.$$

But this implies that adding T_i to S_i results in strictly less utility than adding $T_i \setminus \{j\}$, which contradicts the choice of T_i as the maximizer of (4). ■

Lemma 4.2 is used in the proof of Theorem 4.1 through the following corollary.

Corollary 4.3 *If all bidders have submodular valuations and bid sincerely, then the KC auction terminates with an allocation (S_1, \dots, S_n) and prices \mathbf{p} such that*

$$\sum_{j \in U} p(j) \leq \sum_{i=1}^n v_i(S_i).$$

Proof: Recall that every unsold item j satisfies $p(j) = 0$. Since the S_i 's are disjoint, we have

$$\sum_{j \in U} p(j) = \sum_{i=1}^n \sum_{j \in S_i} p(j) \leq \sum_{i=1}^n v_i(S_i),$$

with the inequality following from Lemma 4.2. ■

Proof of Theorem 4.1: Let (S_1, \dots, S_n) and \mathbf{p} be the allocation and prices of the KC auction at termination. Let (S_1^*, \dots, S_n^*) be a welfare-maximizing allocation. At termination, every bidder i is uninterested in augmenting its items S_i by the items $S_i^* \setminus S_i$, and so

$$v_i(S_i^* \cup S_i) - v_i(S_i) \leq \sum_{j \in S_i^* \setminus S_i} (p(j) + \epsilon).$$

Since v_i is monotone and prices are nonnegative, this implies that

$$v_i(S_i^*) - v_i(S_i) \leq \sum_{j \in S_i^*} p(j) + |S_i^*|\epsilon.$$

Summing over all bidders i and rearranging yields

$$\begin{aligned} \sum_{i=1}^n v_i(S_i) &\geq \sum_{i=1}^n v_i(S_i^*) - \sum_{i=1}^n \sum_{j \in S_i^*} p(j) - m\epsilon \\ &\geq \sum_{i=1}^n v_i(S_i^*) - \sum_{j \in U} p(j) - m\epsilon \\ &\geq \sum_{i=1}^n v_i(S_i^*) - \sum_{i=1}^n v_i(S_i) - m\epsilon, \end{aligned}$$

where the first inequality follows from the fact that S_i^* 's are disjoint and the second inequality follows from Corollary 4.3. Rearranging terms completes the proof. ■

References

- [1] S. Dobzinski and J. Vondrak. Communication complexity of combinatorial auctions with submodular valuations. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1205–1215, 2013.
- [2] H. Fu, R. D. Kleinberg, and R. Lavi. Conditional equilibrium outcomes via ascending price processes with applications to combinatorial auctions with item bidding. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, page 586, 2011.
- [3] F. Gul and E. Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87(1):95–124, 1999.
- [4] B. Lehmann, D. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006.