

Lecture 4: Sealed-Bid Combinatorial Auctions

*Lecturer: Sébastien Lahaie**Scribe: Sébastien Lahaie*

In this lecture we examine two different generalizations of the Vickrey auction to a setting where there are multiple distinct items to be allocated, and the agents' values are defined over bundles rather than just individual items. Auctions that allocate multiple distinct items are commonly called “combinatorial” or “package” auctions. The properties of the Vickrey auction we would like to generalize are:

1. The items are allocated to those who value them the most (in some formal sense).
2. The agents are incentivized to truthfully report their values.
3. The agents have no reason not to participate in the auction (they will derive non-negative utility).
4. The revenue to the seller is non-negative.

We will examine a classic sealed-bid auction from the economics literature that successfully generalizes all these properties, the VCG mechanism. However, the auction implicitly requires the solution of computationally difficult (NP-hard) problems, so we will then look at an auction, the LOS mechanism, that tries to achieve all the properties above using polynomial-time computation. First we provide the standard model used to study combinatorial auctions.

1 The Combinatorial Allocation Problem

There is a set of agents N and a set of items M , held by a single seller. Let $n = |N|$ and $m = |M|$ be the numbers of agents and items. The items are distinct and indivisible. A *bundle* is simply a subset of items. An *allocation* is a vector of bundles $R = (R_i)_{i \in N}$, where R_i is the bundle allocated to agent i , and similarly R_j is the bundle allocated to agent j , etc. An allocation is *feasible* if $R_i \cap R_j = \emptyset$ for $i \neq j$. We denote the set of feasible allocations among agents N by $\Gamma(N)$, and similarly for subsets of N . When we speak of an allocation R among agents $N - i$ (the set of agents without i), we mean an allocation such that $R_i = \emptyset$.

Each agent i has a *valuation* function over bundles $v_i : 2^M \rightarrow \mathbf{R}_+$. (The notation 2^M denotes the set of all subsets of M —the set of all bundles including the empty bundle.) Valuations are defined over bundles rather than just items because we want to allow for the possibility of complements and substitutes. Informally, items are complements when their value together is more than the sum of their individual values, and they are substitutes when the reverse holds.

An *outcome* is a pair (R, q) consisting of a feasible allocation R and a vector of payments $q = (q_i)_{i \in N}$ from the agents to the seller. The *utility* to agent j when the outcome is (R, q) is defined as

$$u_j(R, q; v) = v_j(R_j) - q_j.$$

A utility function is parametrized by the valuation profile $v = (v_i)_{i \in N}$, but this is usually suppressed when clear from context. There are several things to note about the form of the utility function:

Private values. Each agent i knows its own valuation v_i , and its utility does not change if it learns of any other agent's values.

No externalities. An agent only cares about the bundle it acquires and payment it makes, not the bundles and payments of other agents.

Quasi-linearity. Utilities are denoted in a common currency, and utility can be transferred to and from agents in the form of payments.

In particular, agents are not constrained by any budget. For this to hold, it is sufficient to assume that the budget of each agent is at least $\max_{i \in N} v_i(M)$, so quasi-linearity is a reasonable assumption when agents have ample liquidity or credit. Meanwhile, the seller, indexed by 0, only derives utility from payments received:

$$u_0(R, q) = \sum_{i \in N} q_i.$$

The problem at hand is to find a feasible allocation R together with payments q that maximizes the total utility to the agents and seller:

$$\begin{aligned} & \sum_{i \in N} u_i(R, q) + u_0(R, q) \\ &= \sum_{i \in N} [v_i(R_i) - q_i] + \sum_{i \in N} q_i \\ &= \sum_{i \in N} v_i(R_i). \end{aligned} \tag{1}$$

We see that payments cancel out, and the problem of finding an allocation and payments that maximize total utility is equivalent to the problem of finding an allocation that maximizes the total value to the agents—namely, an *efficient* allocation. We will refer to the problem of efficiently allocating the items as the *combinatorial allocation problem* (CAP).

2 The VCG Mechanism

Ignoring for the moment issues of communication and computation, we first ask how to motivate the agents to truthfully reveal their valuation functions given that we want to implement an efficient allocation. If we can make this happen, we can be sure the resulting allocation is efficient with respect to the agents' true valuations. The Vickrey-Clarke-Groves (VCG) mechanism, a natural generalization of the Vickrey auction, is a solution to this problem [1, 2, 5].

Throughout we will consider the class of *general* valuations and its subclasses. It is denoted $\mathcal{V}$.

Definition 1 A valuation $v_i : 2^M \rightarrow \mathbf{R}_+$ is general if it is

1. Each agent $i \in N$ reports a valuation $\tilde{v}_i \in \mathcal{V}$.
2. Select an allocation R that is efficient with respect to reported values,

$$R \in \arg \max_{R' \in \Gamma} \sum_{i \in N} \tilde{v}_i(R'_i) \quad (2)$$

3. Charge each agent i a payment of

$$\hat{q}_i = \sum_{j \neq i} \tilde{v}_j(R_j^{-i}) - \sum_{j \neq i} \tilde{v}_j(R_j) \quad (3)$$

where R^{-i} is an efficient allocation among agent $N - i$ (i.e., we remove agent i from the set of agents).

Figure 1: The VCG Mechanism

Monotone: For bundles S, T such that $S \subseteq T$, we have $v_i(S) \leq v_i(T)$.

Normalized: The value for the empty set is $v_i(\emptyset) = 0$.

The VCG mechanism is given in Figure 1. The seller is just a passive participant that receives payments, because it has only one possible utility function, which it therefore cannot misrepresent. A “bid” is another term for a reported value. The VCG mechanism can thus be interpreted as a sealed-bid auction that allocates the items efficiently according to bids.

Before turning to the formal properties of the VCG mechanism, let’s consider how (3) comes about. The VCG payments are defined so that the report of an agent i does not affect the total utility of the others (including the seller). Without agent i , the payoff to the others is

$$\sum_{j \neq i} \tilde{v}_j(R_j^{-i}). \quad (4)$$

With i present, the payoff to the others then becomes

$$\hat{q}_i + \sum_{j \neq i} \tilde{v}_j(R_j). \quad (5)$$

Equating (4) with (5) and solving for the payment yields (3). Thus whatever impact agent i ’s report might have on the others is reflected in the payment; economists call this “internalizing the externality”. This feature makes truthful reporting—setting $\tilde{v}_i = v_i$ —an optimal strategy for each bidder,

Proposition 1 *In the VCG mechanism, truthful reporting is always an optimal strategy for each bidder, no matter what the reports of the other agents.*

Proof. Consider an arbitrary bidder i , and fix the reports of the other bidders to v_{-i} . Let R^{-i} be an efficient allocation among bidders $N - i$ with respect to v_{-i} . Let R be the

allocation dictated by the VCG mechanism if i reports v_i , his true valuation. Its utility from this report is

$$\begin{aligned} & v_i(R_i) - \left[\sum_{j \neq i} v_j(R_j^{-i}) - \sum_{j \neq i} v_j(R_j) \right] \\ &= \sum_{j \in N} v_j(R_j) - \sum_{j \neq i} v_j(R_j^{-i}). \end{aligned} \quad (6)$$

Let R' be the allocation that results if i reports $\tilde{v}_i \neq v_i$. The utility to i is then

$$\sum_{j \in N} v_j(R'_j) - \sum_{j \neq i} v_j(R_j^{-i}). \quad (7)$$

Subtracting (7) from (6) we obtain

$$\sum_{j \in N} v_j(R_j) - \sum_{j \in i} v_j(R'_j) \geq 0. \quad (8)$$

The inequality here follows because R is an efficient allocation with respect to the true profile of valuations $(v_i)_{i \in N}$. This completes the proof, because the choice of the other agents' reports was arbitrary. $\square$

In words, here is the intuition behind why the VCG mechanism induces truthful reporting. If an agent i reports truthfully, the VCG mechanism maximizes the total value to the agents. If i makes a false report, the total value will be weakly less. But the VCG payments ensure that i 's report cannot affect the payoff to the remaining agents and seller. Thus the overall decrease in total value due to a false report must be borne entirely by agent i . It follows that truthful reporting is an optimal strategy.

Although truthful reporting is an optimal strategy, we have not ruled out the possibility that other reports might also be optimal. This is a concern, because we would like to be able to exactly predict the agent's behavior to confirm other properties of the VCG mechanism. If truthful reporting is the only strategy that is always optimal (i.e., for any vector of the other agents' valuations), we say that it is a *dominant strategy*.

It is standard in game theory to use the notation v_{-i} to denote the vector v with its i th component removed. Thus $v = (v_i, v_{-i})$ and $(\tilde{v}_i, v_{-i})$ is the vector v with its i th component replaced with $\tilde{v}_i$. We say that a false report $\tilde{v}_i$ is *potentially pivotal* if there is a vector v_{-i} of the other agents' reports such that there is an efficient allocation with respect to $(\tilde{v}_i, v_{-i})$ that is not efficient with respect to (v, v_{-i}) . Recall that v_i is i 's true valuation.

Proposition 2 *If every report of agent i is potentially pivotal, then truthful reporting is a dominant strategy for agent i .*

Proof. If every report is potentially pivotal, then for every misreport $\tilde{v}_i$, there is a vector of the other agents' valuations v_{-i} such that inequality (8) is strict. Thus $\tilde{v}_i$ cannot be an always optimal strategy. Hence v_i is the only always optimal strategy: it is a dominant strategy. $\square$

It is not hard to check that if we consider general valuations, every report is potentially pivotal. Thus we expect the agents to truthfully report their values in the VCG mechanism.

Proposition 3 *The VCG mechanism is efficient.*

Proof. Immediate from the fact that agents report their values truthfully in the VCG mechanism, and that it implements an efficient outcome with respect to these reports. $\square$

If an agent chooses not to participate in the mechanism, it derives utility 0: the empty bundle $\emptyset$ and no payment. The following thus shows that the VCG mechanism is *individually-rational*: agents have no incentive not to participate in the mechanism.

Proposition 4 *In the VCG mechanism, the utility derived by a truthful agent is non-negative.*

Proof. Given that the others' reports are v_{-i} , the utility to i of reporting its true valuation v_i is

$$\sum_{j \in N} v_j(R_j) - \sum_{j \neq i} v_j(R_j^{-i}).$$

We claim this is non-negative. Recall that R^{-i} is a feasible allocation that gives $\emptyset$ to agent i . By step 2 of the VCG mechanism, R is efficient for the profile of valuations v submitted. By definition, the total value of R is at least that of any other feasible allocation, in particular R^{-i} . Thus we have

$$\sum_{j \in N} v_j(R_j) \geq \sum_{j \in N} v_j(R_j^{-i}) = \sum_{j \neq i} v_j(R_j^{-i}),$$

proving the result. $\square$

Finally, the VCG mechanism is also individually-rational for the seller.

Proposition 5 *In the VCG mechanism, the revenue to the seller is non-negative.*

Proof. Observe that the VCG payment (3) of each agent i is non-negative from the fact that R^{-i} is a efficient allocation with respect to reports $\tilde{v}_{-i}$. $\square$

You should check for yourself that if there is a single item, then the VCG mechanism reduces to the Vickrey auction.

3 The LOS Mechanism

The computational problem of finding an efficient allocation in the second step of the VCG mechanism is commonly called the “winner determination” (WD) problem. Here we establish that even for a restricted subset of the general valuations that is concise to represent (the description length is polynomial in m), we cannot hope for a polynomial-time algorithm to solve WD.

Definition 2 *A valuation v_i is single-minded if there is a bundle S_i such that*

$$v_i(T_i) = \begin{cases} v_i(S_i) & \text{if } T_i \supseteq S_i \\ 0 & \text{otherwise} \end{cases}$$

for each bundle T_i .

1. Each agent $i \in N$ reports a single-minded valuation (T_i, b_i) .
2. Reindex the agents so that

$$\frac{b_1}{\sqrt{|T_1|}} \geq \frac{b_2}{\sqrt{|T_2|}} \geq \dots \geq \frac{b_n}{\sqrt{|T_n|}}. \quad (9)$$

3. For $i = 1, \dots, n$ in order, check whether no items of T_i have been allocated to a previous agent. If so, set $R_i = T_i$; if not, set $R_i = \emptyset$. Allocate R_i to agent i .
4. For each bidder i , if i does not uniquely block any other bid, charge payment $q_i = 0$. Otherwise, let $\sigma(i) > i$ be the first bid in the ordering uniquely blocked by i , and charge

$$q_i = \frac{b_{\sigma(i)}}{\sqrt{|T_{\sigma(i)}|}} \sqrt{|T_i|}. \quad (10)$$

Figure 2: The LOS Mechanism

A bidder with a single-minded valuation is interested in acquiring all the items in S , and no more. This is perhaps the simplest kind of valuation that exhibits complementarity between items. We represent a typical single-minded valuation by a pair (S_i, w_i) , where w_i is the value of S_i , the set of relevant items.

Here are some facts about the complexity of WD with single-minded agents [4]:

- The WD problem is NP-hard (by reduction from weighted independent set).
- For every $\epsilon > 0$, there is no $O(m^{\frac{1}{2}-\epsilon})$ -approximation algorithm for WD, unless NP=ZPP.

ZPP is the class of decision problems which can be solved by a randomized algorithm with polynomial expected running time. It is highly unlikely that NP=ZPP; this would imply a host of other improbable complexity results such as NP = co-NP. Thus, even ignoring incentive issues, the best we can reasonably hope for is a $O(\sqrt{m})$ -approximation algorithm for the CAP with single-minded bidders. The Lehmann-O'Callaghan-Shoham (LOS) mechanism, given in Figure 2, achieves this ratio and also motivates truthful bidding [3]. Let's first prove the ratio, assuming truthful bids (so we write them as (S_i, w_i) instead of (T_i, b_i)).

Theorem 1 *The LOS algorithm is a $\sqrt{m}$ -approximation algorithm for the CAP.*

Proof. Let $I \subseteq \{1, 2, \dots, n\}$ be the set of indices of the bids granted by the LOS algorithm, and let I^* be those of an optimal set of bids. We will show that

$$\sum_{i \in I^*} w_i \leq \sqrt{m} \sum_{i \in I} w_i. \quad (11)$$

We say that a bid $i \in I$ *blocks* a bid $i' \in I^*$ if $S_i \cap S_{i'} \neq \emptyset$. For a bid $i \in I$, let F_i be the subset of bids of I^* such that i is the first bid in the greedy ordering to block them:

$$F_i = \{i' \in I^* \mid S_i \cap S_{i'} \neq \emptyset, \text{ and there is no } j < i \text{ with this property}\}.$$

We may have $i \in F_i$. Now note that if $i' \in F_i$, then when the greedy algorithm chose bid i , it could have chosen i' , because no previous bid blocked i' . Thus by (9) we have

$$\frac{w_{i'}}{\sqrt{|S_{i'}|}} \leq \frac{w_i}{\sqrt{|S_i|}}. \quad (12)$$

Note also that each $i' \in I^*$ must lie in exactly one of the F_i 's, because a bid cannot be “first” blocked by two different bids. Thus we can write

$$\sum_{i \in I^*} w_i = \sum_{i \in I} \sum_{i' \in F_i} w_{i'}. \quad (13)$$

We therefore consider each $i \in I$ separately to reach the bound (11). Summing (12) over all $i' \in F_i$ we obtain

$$\sum_{i' \in F_i} w_{i'} \leq \frac{w_i}{\sqrt{|S_i|}} \left(\sum_{i' \in F_i} \sqrt{|S_{i'}|} \right). \quad (14)$$

We want to bound the expression in parentheses. Since the bids in F_i are part of a feasible solution, they are mutually disjoint, and thus $\sum_{i' \in F_i} |S_{i'}| \leq m$. In the worst-case this holds with equality, and to maximize the term in parentheses we would distribute the items evenly among the elements of F_i . Thus we can weaken (14) to

$$\sum_{i' \in F_i} w_{i'} \leq \frac{w_i}{\sqrt{|S_i|}} \left(\sum_{i' \in F_i} \sqrt{\frac{m}{|F_i|}} \right) = \frac{w_i}{\sqrt{|S_i|}} \sqrt{m} \sqrt{|F_i|} \quad (15)$$

We are almost there: all we need now is to show $\sqrt{|F_i|} \leq \sqrt{|S_i|}$. But we must have $|F_i| \leq |S_i|$ since S_i must contain an item of each bid in F_i (because it blocks all of them), and the bids in F_i are disjoint. So (15) leads to

$$\sum_{i' \in F_i} w_{i'} \leq \sqrt{m} w_i,$$

and summing over all $i \in I$ and applying (13), we have (11). $\square$

To understand the payment rule of the LOS mechanism, we need to define what it means for a bid to “uniquely block” another.

Definition 3 *A bid i uniquely blocks a bid j if j is not granted when i is included in the input, but is granted when i is removed from the input.*

As an exercise, you should prove to yourself that if i uniquely blocks j , we must have $j > i$ in the greedy ordering.

To prove that payments (10) make truthful bidding an optimal strategy, we proceed in two steps. First we show there is no benefit to any bid (T_i, b_i) over the bid (S_i, b_i) —there is no reason to misreport the bundle S_i . Then we show that setting $b_i = w_i$ is optimal. First, however, we establish that the LOS mechanism is individually rational.

Lemma 1 *In the LOS mechanism, the utility derived by a truthful agent is non-negative.*

Proof. The lemma holds if a bidder is always charged at most its bid. Because the index $\sigma(i)$ defined in step 4 of the LOS mechanism occurs after i in the greedy ordering, we have

$$\frac{b_i}{\sqrt{|T_i|}} \geq \frac{b_{\sigma(i)}}{\sqrt{|T_{\sigma(i)}|}}$$

and rearranging gives $b_i \geq q_i$. □

The following is a key insight for the following proofs.

Lemma 2 *Suppose bid i uniquely blocks at least one other bid, and let j be the highest indexed bid in the ordering that i uniquely blocks. Let I be the set of bids ranked before j , other than i . Then, when i is removed from the input, the LOS mechanism grants and rejects bids in I exactly as it would were i included.*

Proof. By induction on the indices in I , from highest ranked to lowest. For the base case, if the highest index bid is granted with i in the input, then it is granted with i not part of the input, because no bid blocks it. If it is not granted with i in the input, it cannot be granted when i is not part of the input, because this would contradict the fact that j is the highest bid in the ordering uniquely blocked by i .

Assume bids are rejected and granted identically with and without i in the input for the first $k - 1$ bids in I , where $k < j$. If k is granted when i is present, this means none of the first $k - 1$ bids selected besides i block it. These bids are selected identically when i is removed (and no other bid is selected) by the induction hypothesis, thus bid k remains unblocked and is granted. If k is not granted with i present, it cannot be granted when i is removed, because this would contradict the fact that j is the highest bid in the ordering uniquely blocked by i . □

The following simple fact is also useful.

Lemma 3 *In the LOS mechanism, if i uniquely blocks at least one bid, then $T_i \cap T_{\sigma(i)} \neq \emptyset$.*

Proof. Bid $\sigma(i)$ is rejected when i is included in the input, but granted when i is removed. Because bids in i are granted and rejected identically whether i is included in the input or not, by Lemma 2, we must therefore have $T_i \cap T_{\sigma(i)} \neq \emptyset$. □

The next lemma shows we can assume $T_i = S_i$ for all $i \in N$.

Lemma 4 *In the LOS mechanism, if a player can benefit over the truthful strategy by bidding (T_i, b_i) , it can benefit just as much by bidding (S_i, b_i) .*

Proof. Consider an arbitrary bidder i and fix the others' bids. Assume for the sake of contradiction that the false bid (T_i, b_i) leads to strictly higher utility than truthful bidding. By Lemma 1, this means the bid must yield positive utility. In particular, the bid must be granted, and we must also have $T_i \supseteq S_i$ (recall the definition of single-minded bidders).

To prove the result it suffices to show that: (a) the bid (S_i, b_i) would be granted, and (b) the payment charged with this bid would be no more than the payment given bid (T_i, b_i) .

To see (a), note that $S_i \subseteq T_i$ implies that bid (S_i, b_i) would lie higher in the ordering than (T_i, b_i) according to (9), and thus would also be granted.

To see (b), let (T_j, b_j) be the highest new bid granted if we remove bid (T_i, b_i) , and let (T_k, b_k) be the highest new bid granted if we remove bid (S_i, b_i) . We argue that (T_j, b_j) lies higher than (T_k, b_k) . Assume that the contrary holds, and assume first that (T_i, b_i) lies before (T_k, b_k) . By Lemma 3, we have $S_i \cap T_k \neq \emptyset$, and since $T_i \supseteq S_i$, we then have $T_i \cap T_k \neq \emptyset$. Now if we remove (S_i, b_i) , by Lemma 2 the algorithm proceeds as if this bid were there until (T_k, b_k) is reached, and then the latter is granted. Thus the same occurs if we remove (T_i, b_i) . This means that (T_k, b_k) is uniquely blocked by (T_i, b_i) . But if (T_i, b_i) is present, bid (T_k, b_k) cannot be granted because $T_i \cap T_k \neq \emptyset$. But this contradicts the fact that (T_j, b_j) is the first bid blocked by (T_i, b_i) .

On the other hand, suppose (T_i, b_i) lies after (T_k, b_k) . Because (T_k, b_k) is granted when we remove (S_i, b_i) , this means it is granted when (T_i, b_i) is included. But this is a contradiction because (T_i, b_i) is granted, and $T_i \cap T_k \neq \emptyset$.

Thus (T_k, b_k) , the first bid uniquely blocked by (S_i, b_i) , lies after (T_j, b_j) , the first bid uniquely blocked by (T_i, b_i) . We therefore have

$$\frac{b_j}{\sqrt{|T_j|}} \sqrt{|T_i|} \geq \frac{b_k}{\sqrt{|T_k|}} \sqrt{|T_i|} \geq \frac{b_k}{\sqrt{|T_k|}} \sqrt{|S_i|}.$$

Thus (S_i, b_i) leads to lower payment than (T_i, b_i) , completing the proof. $\square$

Finally, the main result.

Theorem 2 *In the LOS mechanism, truthful reporting is an always optimal strategy for each bidder.*

Proof. Consider a bidder i , and hold the others' bids fixed (to some arbitrary bids). We consider all possible bids (T_i, b_i) by case analysis and show that none can yield higher utility than (S_i, w_i) . By Lemma 4 it suffices to consider bids of the form (S_i, b_i) . First suppose (S_i, w_i) would be granted. If that bid does not uniquely block any other, the payment is 0 and the utility is w_i , the highest possible. So assume the truthful bid uniquely blocks at least one other bid, and let j be the highest-ranked such bid. Let q_i be the payment under truthful bidding. There are three cases.

- (a) $b_i > q_i$. We claim that in this range, the bid is granted and the payment is always q_i . We show that the selection of bids before j does not depend on the presence of bid (S_i, b_i) . Note that the selection of bids up to j is identical when we remove (S_i, b_i) and when we remove (S_i, w_i) . By Lemma 2 this selection is also identical (other than bid i of course) when (S_i, w_i) is present. Consider the first bid k before j that is not identical to this selection when we introduce (S_i, b_i) . Necessarily k occurs after the latter. If k was selected before introducing (S_i, b_i) , but not after, then we must have $S_i \cap T_k \neq \emptyset$. But this contradicts the fact that bid k is selected when we introduce (S_i, w_i) . If k is not selected before (S_i, b_i) is added, some higher ranked bid blocks it. This bid remains granted when (S_i, b_i) is introduced, as k is the highest ranked bid that changes status, and thus k is still not granted.

- (b) $b_i < q_i$. We claim that in this range, the bid is not granted and so the utility is 0, no better than truthful bidding by Lemma 1. Let (T_j, b_j) be the first bid uniquely blocked by (S_i, b_i) . Since (S_i, b_i) lies below (T_j, b_j) for this range of b_i , the latter is granted by the fact that it is granted when (S_i, w_i) is removed. But we have $S_i \cap T_j \neq \emptyset$ by Lemma 3, and so (S_i, b_i) cannot be granted.
- (c) $b_i = q_i$. This depends on how the tie between bids (S_i, b_i) and (T_j, b_j) is broken. If the former is granted, we have case (a), if the latter is granted, we have case (b).

Now suppose that (S_i, w_i) would not be granted. The only way to gain positive utility is to bid $b_i > w_i$ such that the bid is granted, so consider such a bid. Since (S_i, w_i) is not granted, there is a bid (T_j, b_j) ranked higher that is granted and such that $T_j \cap S_i \neq \emptyset$. Since (S_i, b_i) is granted, (T_j, b_j) cannot be granted when (S_i, b_i) is placed; thus the latter is not granted. But when (S_i, b_i) is removed, (T_j, b_j) is granted. By definition, (S_i, b_i) uniquely blocks (T_j, b_j) , and thus the payment under this bid is at least

$$\frac{b_j}{\sqrt{|T_j|}} \sqrt{|S_i|} \geq w_i$$

where the inequality follows from the fact that (S_i, w_i) is ranked lower than (T_j, b_j) . This implies that the utility from bidding (S_i, b_i) is at most 0. $\square$

Again, you should check that when there is a single item, the LOS mechanism reduces to the Vickrey auction.

References

- [1] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [2] Theodore Groves. Efficient collective choice when compensation is possible. *Review of Economic Studies*, 46:227–241, 1979.
- [3] Daniel Lehmann, Liadan Ita O’Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, 2002.
- [4] Tuomas Sandholm. An algorithm for optimal winner determination in combinatorial auctions. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence*, Stockholm, Sweden, 1999.
- [5] William Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.