

Posted Prices with Incomplete Information

Instructor: Thomas Kesselheim

Our question is still “How well can prices coordinate markets?” Last time, we got a glimpse at the classic economic theory: If there is a Walrasian equilibrium, it defines prices, which make everybody happy and yield maximum social welfare.

Today, we will turn to much more recent results, in which the perspective of computer science comes into play. We turn to a setting of incomplete information. Our goal is to post prices for items without knowing which buyers will be present eventually. Buyers then show up one after the other and buy their preferred item(s).

1 Model

Recall the definition of a combinatorial auction. There are n buyers $\mathcal{N} = \{1, \dots, n\}$ and m items M . Each buyer has a private valuation function $v_i: 2^M \rightarrow \mathbb{R}_{\geq 0}$. Each item can be assigned to at most one buyer.

We assume that the valuation functions v_i are *unit demand*, that is, they are of the form $v_i(S) = \max_{j \in S} v_{i,j}$.

Buyer i 's valuation v_i is drawn from a publicly known distribution \mathcal{D}_i . The outcome v_i , however, is private. We use the knowledge of the distributions $(\mathcal{D}_i)_{i \in \mathcal{N}}$ to compute item prices $(p_j)_{j \in M}$. The mechanism then looks as follows:


- Approach the buyers in order $i = 1, \dots, n$
- Buyer i buys whatever set S_i of unsold items maximizes $v_i(S_i) - \sum_{j \in S_i} p_j$, pays $\sum_{j \in S_i} p_j$

Note that this mechanism still consists of an allocation function f and a payment function p . Of course, buyers could decide to lie about their valuation v_i and buy another set. But this can only reduce the utility because the choice of the set S_i is exactly so that it maximizes utility.

Observation 20.1. *The posted-prices mechanism is truthful for any choice of prices.*

We are interested to what extent such a mechanism can optimize social welfare. That is, how does $\sum_{i \in \mathcal{N}} v_i(S_i)$ compare to $\text{OPT}(v) = \max_{\text{allocation}} \sum_{i \in \mathcal{N}} v_i(S_i^*)$. In the last Lecture, we have seen that if valuations are unit-demand, a Walrasian equilibrium always exists. However, the following example illustrates that, even in the full information setting, using prices given by a Walrasian equilibrium does not necessarily lead to a social welfare maximizing allocation in the posted prices mechanism.

Example 20.2. *Consider items A, B, C and bidders 1, 2 that have unit-demand valuations. Assigning the items is just the same as finding a matching in a complete bipartite graphs whose vertices are $N \cup M$. The edge between $i \in N$ and $j \in M$ has weight $v_{i,j}$. We only draw edges of positive value.*



First of all, one can verify that there exist two Walrasian equilibria using the depicted price vector, namely, allocating either A, B or B, C to the bidders which are therefore welfare maximizing allocations. Nevertheless, using the posted prices mechanism with the given item prices it can occur that bidder 1 buys item A and bidder 3 buys item C which only gives a $\frac{2}{k+1}$ fraction of the optimal social welfare.

In the remainder of this lecture, we will look how to set prices of the items in order to get a good fraction of the optimal social welfare.

2 Step 1: Full Information

We will first assume that we actually know the valuation functions $(v_i)_{i \in \mathcal{N}}$. How can we set prices in this case that still optimize social welfare?

Let $\text{OPT}_i(v)$ denote the item that buyer i gets in optimal solution on v . We define the price for item j depending on who gets it in the optimal allocation by setting

$$p_j^v = \begin{cases} \frac{1}{2}v_{i,j} & \text{if buyer } i \text{ gets item } j \text{ in optimal solution on } v \\ 0 & \text{if item } j \text{ is unassigned in optimal solution on } v \end{cases}$$

Note that equivalently we could write

$$p_j^v = \frac{1}{2} \sum_{i \in \mathcal{N}} \mathbf{1}_{\text{OPT}_i(v)=j} v_{i,j} . \quad (1)$$

Define $T_i(v)$ as the set of items that are sold to buyers $1, \dots, i$ on v . The revenue is given by

$$\text{revenue}(v) = \sum_{j \in M} p_j^v \mathbf{1}_{j \in T_n(v)} \geq \sum_{i \in \mathcal{N}} p_{\text{OPT}_i(v)}^v \mathbf{1}_{\text{OPT}_i(v) \in T_n(v)} .$$

One option for buyer i is to buy nothing. Therefore, $u_i(v) \geq 0$. If $\text{OPT}_i(v)$ has not been sold yet, that is, $\text{OPT}_i(v) \notin T_{i-1}(v)$, then buyer i could also buy $\text{OPT}_i(v)$. This gives us

$$u_i(v) \geq \left(v_{i, \text{OPT}_i(v)} - p_{\text{OPT}_i(v)}^v \right) \mathbf{1}_{\text{OPT}_i(v) \notin T_{i-1}(v)} \geq \left(v_{i, \text{OPT}_i(v)} - p_{\text{OPT}_i(v)}^v \right) \mathbf{1}_{\text{OPT}_i(v) \notin T_n(v)} ,$$

where in the second step we use that $T_{i-1}(v) \subseteq T_n(v)$.

Taking the sum of revenue and buyers' utilities

$$\begin{aligned} \sum_{i \in \mathcal{N}} v_i(S_i) &= \text{revenue}(v) + \sum_{i \in \mathcal{N}} u_i(v) \\ &\geq \sum_{i \in \mathcal{N}} p_{\text{OPT}_i(v)}^v \left(\mathbf{1}_{\text{OPT}_i(v) \in T_n(v)} + \mathbf{1}_{\text{OPT}_i(v) \notin T_n(v)} \right) = \sum_{i \in \mathcal{N}} p_{\text{OPT}_i(v)}^v = \frac{1}{2} \text{OPT}(v) . \end{aligned}$$

3 Step 2: Incomplete Information

It is very easy to turn the above posted-price mechanism into one for the setting of incomplete information. Let \tilde{v} be another sample from the known distributions. Then set the price of item j to $p_j = \mathbf{E} \left[p_j^{\tilde{v}} \right]$. That is, we set it to the expected price, using an independent fresh sample.

Theorem 20.3 (Feldman/Gravin/Lucier, 2015). *The expected social welfare of the posted-prices mechanism is a $\frac{1}{2}$ fraction of the expected optimal social welfare.*

Proof. For the revenue, we have again

$$\text{revenue}(v) = \sum_{j \in M} p_j \mathbf{1}_{j \in T_n(v)} .$$

So, by linearity of expectation

$$\mathbf{E} [\text{revenue}(v)] = \mathbf{E} \left[\sum_{j \in M} p_j \mathbf{1}_{j \in T_n(v)} \right] = \sum_{j \in M} p_j \mathbf{E} \left[\mathbf{1}_{j \in T_n(v)} \right] .$$

Note that we could also replace $\mathbf{E} \left[\mathbf{1}_{j \in T_n(v)} \right] = \Pr [j \in T_n(v)]$ but we will keep the indicator because it nicely cancels out eventually.

Lower bounding the utilities is more complicated because we have to avoid dependencies. To this end, draw another valuation profile $v_{-i}^{(i)}$ for every $i \in \mathcal{N}$. Buyer i could buy the item she gets in the optimal solution on $(v_i, v_{-i}^{(i)})$. So, this is the optimal solution on the valuation consisting of the actual valuation v_i but the “hallucinated” other valuations $v_{-i}^{(i)}$. The utility is at least

$$u_i(v) \geq \sum_{j \in M} \mathbf{1}_{j = \text{OPT}_i(v_i, v_{-i}^{(i)})} (v_{i,j} - p_j) \mathbf{1}_{j \notin T_{i-1}(v)} .$$

By linearity of expectation, this implies

$$\begin{aligned} \mathbf{E} [u_i(v)] &\geq \mathbf{E} \left[\sum_{j \in M} \mathbf{1}_{j = \text{OPT}_i(v_i, v_{-i}^{(i)})} (v_{i,j} - p_j) \mathbf{1}_{j \notin T_{i-1}(v)} \right] \\ &= \sum_{j \in M} \mathbf{E} \left[\mathbf{1}_{j = \text{OPT}_i(v_i, v_{-i}^{(i)})} (v_{i,j} - p_j) \mathbf{1}_{j \notin T_{i-1}(v)} \right] . \end{aligned}$$

Observe that the first part of the expectation only depends on $v_{-i}^{(i)}$ and v_i whereas the second part only depends on v_1, \dots, v_{i-1} .

$$\underbrace{\mathbf{1}_{j = \text{OPT}_i(v_i, v_{-i}^{(i)})} (v_{i,j} - p_j)}_{\text{only depends on } v_i \text{ and } v_{-i}^{(i)}} \underbrace{\mathbf{1}_{j \notin T_{i-1}(v)}}_{\text{only depends on } v_1, \dots, v_{i-1}}$$

Therefore, we can write

$$\begin{aligned} &\mathbf{E} \left[\mathbf{1}_{j = \text{OPT}_i(v_i, v_{-i}^{(i)})} (v_{i,j} - p_j) \mathbf{1}_{j \notin T_{i-1}(v)} \right] \\ &= \mathbf{E} \left[\mathbf{1}_{j = \text{OPT}_i(v_i, v_{-i}^{(i)})} (v_{i,j} - p_j) \right] \mathbf{E} \left[\mathbf{1}_{j \notin T_{i-1}(v)} \right] . \end{aligned}$$

Finally, we use that $v_{-i}^{(i)}$ and v_{-i} are identically distributed to get

$$\mathbf{E} \left[\mathbf{1}_{j = \text{OPT}_i(v_i, v_{-i}^{(i)})} (v_{i,j} - p_j) \right] = \mathbf{E} \left[\mathbf{1}_{j = \text{OPT}_i(v)} (v_{i,j} - p_j) \right] ,$$

and that $T_{i-1}(v) \subseteq T_n(v)$ to get

$$\mathbf{E} \left[\mathbf{1}_{j \notin T_{i-1}(v)} \right] \geq \mathbf{E} \left[\mathbf{1}_{j \notin T_n(v)} \right] .$$

So overall

$$\mathbf{E} [u_i(v)] \geq \sum_{j \in M} \mathbf{E} \left[\mathbf{1}_{j = \text{OPT}_i(v)} (v_{i,j} - p_j) \right] \mathbf{E} \left[\mathbf{1}_{j \notin T_n(v)} \right] .$$

Now, we take the sum over all $i \in \mathcal{N}$

$$\begin{aligned}
\mathbf{E} \left[\sum_{i \in \mathcal{N}} u_i(v) \right] &= \sum_{i \in \mathcal{N}} \mathbf{E} [u_i(v)] \\
&\geq \sum_{i \in \mathcal{N}} \sum_{j \in M} \mathbf{E} \left[\mathbf{1}_{j=\text{OPT}_i(v)} (v_{i,j} - p_j) \right] \mathbf{E} \left[\mathbf{1}_{j \notin T_n(v)} \right] \\
&= \sum_{j \in M} \mathbf{E} \left[\mathbf{1}_{j \notin T_n(v)} \right] \sum_{i \in \mathcal{N}} \mathbf{E} \left[\mathbf{1}_{j=\text{OPT}_i(v)} (v_{i,j} - p_j) \right] \\
&= \sum_{j \in M} \mathbf{E} \left[\mathbf{1}_{j \notin T_n(v)} \right] \left(\mathbf{E} \left[\sum_{i \in \mathcal{N}} \mathbf{1}_{j=\text{OPT}_i(v)} v_{i,j} \right] - \mathbf{E} \left[\sum_{i \in \mathcal{N}} \mathbf{1}_{j=\text{OPT}_i(v)} p_j \right] \right)
\end{aligned}$$

Observe that by (1)

$$\mathbf{E} \left[\sum_{i \in \mathcal{N}} \mathbf{1}_{j=\text{OPT}_i(v)} v_{i,j} \right] = \mathbf{E} \left[2p_j^v \right] = 2p_j .$$

Furthermore, note that $\sum_{i \in \mathcal{N}} \mathbf{1}_{j=\text{OPT}_i(v)} \leq 1$ because the optimum may allocate item j at most once. Therefore

$$\mathbf{E} \left[\sum_{i \in \mathcal{N}} \mathbf{1}_{j=\text{OPT}_i(v)} p_j \right] \leq p_j .$$

So, in combination

$$\sum_{i \in \mathcal{N}} \mathbf{E} \left[\mathbf{1}_{j=\text{OPT}_i(v)} (v_{i,j} - p_j) \right] \geq p_j .$$

For the sum of buyer utilities this means

$$\mathbf{E} \left[\sum_{i \in \mathcal{N}} u_i(v) \right] \geq \sum_{j \in M} \mathbf{E} \left[\mathbf{1}_{j \notin T_n(v)} \right] p_j .$$

Summarizing

$$\mathbf{E} \left[\sum_{i \in \mathcal{N}} v_i(S_i) \right] = \mathbf{E} \left[\text{revenue}(v) + \sum_{i \in \mathcal{N}} u_i(v) \right] \geq \sum_{j \in M} p_j = \frac{1}{2} \mathbf{E} [\text{OPT}(v)] . \quad \square$$

4 Optimality

Note that any posted-prices mechanism inherently works in a sequential way. Therefore, if we show optimality for any sequential algorithm, then this also implies our choice of prices is optimal.

Theorem 20.4. *There are distributions such that the expected social welfare of any sequential/online algorithm is no better than $\frac{1}{2}$ fraction of the expected optimal social welfare.*

Proof. Consider a single item. Buyer 1 has value 1, buyer 2 has value $\frac{1}{\epsilon}$ with probability ϵ , 0 otherwise. The optimal social welfare is achieved by giving the item to buyer 2 if he has high value, to buyer 1 otherwise. The expected value is

$$\epsilon \cdot \frac{1}{\epsilon} + (1 - \epsilon) \cdot 1 = 2 - \epsilon .$$

In contrast, an algorithm that sequentially makes the decisions, has to decide whether to give the item to buyer 1 without knowing buyer 2's value. No matter if it decides to give the item to buyer 1 or not (in which case it goes to buyer 2), the expected value is always 1. \square

References

- M. Feldman, N. Gravin, B. Lucier, Combinatorial Auctions via Posted Prices, SODA 2015. (Original proof in a more general form.)
- P. Dütting, M. Feldman, T. Kesselheim, B. Lucier, Prophet Inequalities Made Easy: Stochastic Optimization by Pricing Non-Stochastic Inputs, FOCS 2017. (Improved and generalized proof; formalization that it is enough to consider full-information setting.)