Lecture 19: Bayesian Combinatorial Auctions

Lecturer: *Sahil Singla* Last updated: March 21, 2022

In this lecture, we'll take a Bayesian view on Combinatorial Auctions, i.e., we'll assume that the bidder valuations are drawn from some known probability distributions. This stochastic assumption will allow us to design truthful approximation mechanisms that are much simpler (hence more practical) and with much better approximation ratios.

1 Model

There is a set of n bidders and a set $[m]$ of m items. The valuation $v_i(\cdot) : 2^{[m]} \to \mathbb{R}_{\geq 0}$ of the i-th bidder is drawn from some known probability distribution $\mathcal{D}_i$[1]. That is, player i receives value $v_i(S)$ for receiving the set S of items. We make the standard assumption that valuation v_i of each bidder i is normalized, i.e., $v_i(\emptyset) = 0$, and monotone, i.e., $v_i(S) \leq v_i(T)$ for every $S \subseteq T \subseteq [m]$. Your optimization problem is to find an allocation, that is a partition of $[m]$ into $S_1, \ldots, S_n$ so as to maximize the *social welfare*, $\sum_i v_i(S_i)$.

Similar to the last lecture, we don't know the valuations $v_i(\cdot)$ (the distributions $\mathcal{D}_i$ are known) and have to design a "truthful" mechanism/algorithm that incentivizes the bidders to report truthfully. We'll assume the bidders are quasi-linear, i.e., they want to maximize their utility which is the difference of their value minus payment. A mechanism is called *truthful* if the dominant strategy of each bidder is to reveal their true valuation in response to given queries (same as Definition 1 from last lecture).

Query access to valuations. Since valuations have size exponential in m, a common assumption is that valuations are specified via certain queries instead, in particular, value queries and demand queries. A *value query* to valuation v on bundle S reveals the value of $v(S)$. A *demand query* specifies a price vector $\mathbf{p} \in \mathbb{R}^m_{\geq 0}$ on items and the answer is the "most demanded" bundle under this pricing, i.e., a bundle $S \in \arg\max_{S'}\{v(S') - \mathbf{p}(S')\}$.

2 Fixed-Price Auctions

For a price vector $\mathbf{p}$ and a set of items $M' \subseteq M$, we define $\mathbf{p}(M') := \sum_{j \in M'} p_j$. For an allocation $A = (A_1, \ldots, A_n)$, we sometimes abuse the notation and use A to denote the set of allocated items.

[1] We will not worry about how this distribution is specified. For concreteness, you may think that $\mathcal{D}_i$ is distributed over a polynomial number of valuation functions.

2.1 XOS Valuations and Supporting Prices

We are interested in the case when bidders valuations are *submodular* and hence capture the notion of "diminishing marginal utility" of items for bidders. A valuation v is submodular iff $v(S \cup T) + v(S \cap T) \leq v(S) + v(T)$ for any $S, T \subseteq M$.

Submodular functions are a strict subset of *XOS* valuations also known as *fractionally additive* valuations (you'll prove this in HW-4) that are defined as follows. A valuation a is additive/linear iff $a(S) = \sum_{j \in S} a(\{j\})$ for every bundle S, i.e., valuation of a set is the sum of item valuations inside the set. A valuation function v is XOS iff there exist t additive valuations $\{a_1, \ldots, a_t\}$ such that $v(S) = \max_{r \in [t]} a_r(S)$ for every $S \subseteq M$. (Note that t could be arbitrarily large, like exponential in the number of items.) Each a_r is referred to as a *clause* of v. If $a \in \arg\max_{r \in [t]} a_r(S)$, then a is called a *maximizing clause* for S and $a(\{j\})$ is a *supporting price* of item j in this maximizing clause.

Definition 1 (Supporting prices)**.** *We say that an allocation $A = (A_1, \ldots, A_n)$ of items to n bidders with XOS valuation is* supported *by prices $\boldsymbol{q} = (q_1, \ldots, q_m)$ iff each q_j is a supporting price for item j in the maximizing clause of the bidder i to whom j is allocated, i.e., $j \in A_i$.*

2.2 A Fixed-Price-Auction for XOS Bidders

We use a standard fixed-price auction as a subroutine in our mechanism. For an *ordered* set N of bidders, M of items, and a price vector $\mathbf{p}$, $\mathsf{FPA}(N, M, \mathbf{p})$ is defined as follows.

$\mathsf{FPA}(N, M, \mathbf{p})$

1. Iterate over the bidders i of the ordered set N in the given order:
   (a) Allocate $A_i \in \arg\max_{S \subseteq M}\{v_i(S) - \mathbf{p}(S)\}$ to bidder i and update $M \leftarrow M \setminus A_i$.
2. Return the allocation $A = (A_1, \ldots, A_n)$.

It is easy to see that FPA can be implemented using one demand query per bidder. Its truthfulness is also easy to check as bidders have no influence on the pricing mechanism.

The following theorem gives a key property of this auction. It shows that for XOS valuations there always exist good prices such that running FPA w.r.t. those prices gives a good-approximation to the welfare.

Theorem 1. *Suppose $O = (O_1, \ldots, O_n)$ is an allocation with supporting prices $\boldsymbol{q}$ (think of O being the optimal allocation). Then, setting prices $\boldsymbol{p}$ in the FPA with item j priced at $p_j = q_j/2$ gives an allocation A of total welfare $\mathsf{val}(A) \geq \sum_{j \in M} p_j = \sum_{j \in M} q_j/2$.*

Proof. Define $\overline{A}_i = O_i \setminus A$ for every $i \in N$ as the set of items in O_i that are never allocated by FPA. Bidder i could have chosen $\overline{A}_i$ in FPA but decided to pick another bundle A_i

instead. This implies that the utility of the i-th bidder can be lower bounded as

$$\text{Utility(i)} \;:=\; v_i(A_i) - \mathbf{p}(A_i) \;\geq\; v_i(\overline{A}_i) - \mathbf{p}(\overline{A}_i). \tag{1}$$

Since welfare can be viewed as the sum of revenue and utility, we have

$$\begin{aligned}\mathsf{val}(A) \;=\; \sum_{i=1}^{n} v_i(A_i) \;&=\; \underbrace{\mathbf{p}(A)}_{\text{Revenue}} + \sum_{i=1}^{n} \underbrace{(v_i(A_i) - \mathbf{p}(A_i))}_{\text{Utility(i)}} \\ &\overset{(1)}{\geq}\; \mathbf{p}(A) + \sum_{i=1}^{n}(v_i(\overline{A}_i) - \mathbf{p}(\overline{A}_i)) \;\geq\; \mathbf{p}(A) + \sum_{i=1}^{n}(\mathbf{q}(\overline{A}_i) - \mathbf{p}(\overline{A}_i))\end{aligned}$$

since $\overline{A}_i \subseteq O_i$. Now using $q_j = 2p_j$ for all $j \in \overline{A} \subseteq O$, we get

$$\mathsf{val}(A) \;\geq\; \mathbf{p}(A) + \sum_{i=1}^{n} \mathbf{p}(\overline{A}_i) \;=\; \mathbf{p}(O) \;=\; 1/2 \cdot \mathbf{q}(O),$$

where the second-last equality uses $O_i = A_i \cup \overline{A}$ and $\overline{A} \cap A = \emptyset$. □

3 Prophet Inequalities for Combinatorial Auctions

Now we return to our model of Bayesian Combinatorial Auction from Section 1 where the XOS valuation of the i-th bidder is drawn from a known distribution $\mathcal{D}_i$. We will design a simple FPA that gives 2-approximation for this setting.

Formally, for any XOS bidder valuations $\mathbf{v} = (v_1, \ldots, v_n)$, let $O_{\mathbf{v}}$ denote the optimal allocation with supporting prices $\mathbf{q}_{\mathbf{v}}$. The following is the main result.

Theorem 2 ([1])**.** *If we run* FPA *with prices* $\boldsymbol{p} = 1/2 \cdot \mathbb{E}_{\boldsymbol{v}}[\boldsymbol{q}_{\boldsymbol{v}}]$ *then the obtained allocation* A *has total total expected welfare* $\mathbb{E}[\mathsf{val}(A)] \geq \sum_{j \in M} p_j = 1/2 \cdot \mathbb{E}_{\boldsymbol{v}}[\boldsymbol{q}(O_{\boldsymbol{v}})]$.

This result can also be interpreted as a Prophet Inequality result (as studied in Lecture 13). Consider the following online problem: n bidders arrive one-by-one where the i-th bidder valuation $v_i \sim \mathcal{D}_i$. On arrival of a bidder the goal is to immediately and irrevocably allocate them a subset S_i of the remaining items while trying to maximize the total expected welfare $\mathbb{E}[\sum_i v_i(S_i)]$. Now, Theorem 2 says that we can obtain a 2-competitive-ratio for this problem by letting the bidders choose their favorite set of remaining items at prices $\mathbf{p}$. So, Theorem 2 is a massive generalization of Theorem 2 from Lecture 13.

Proof of Theorem 2. The proof will again go by viewing the welfare as the sum of revenue and utility. That is, total welfare

$$\mathsf{val}(A) \;=\; \sum_{i=1}^{n} v_i(A_i) \;=\; \underbrace{\mathbf{p}(A)}_{\text{Revenue}} + \sum_{i=1}^{n} \underbrace{(v_i(A_i) - \mathbf{p}(A_i))}_{\text{Utility(i)}}. \tag{2}$$

We will prove lower bounds on both the Revenue and the Utility.

Let f_j denote the probability that item j is free/unsold by $\mathsf{FPA}(N, M, \mathbf{p})$ till the end. This means that item j gets sold with probability $(1 - f_j)$, so we get expected Revenue equals

$$\mathbb{E}[\mathbf{p}(A)] \;=\; \sum_{j \in M} (1 - f_j) \cdot p_j. \tag{3}$$

The following lemma proves a lower bound on total utility in terms of f_j.

Lemma 3. *The total expected utility $\sum_{i=1}^{n}$ Utility$(i) \geq \sum_{j \in M} f_j \cdot p_j$.*

Before proving the lemma, we complete the proof of Theorem 2: combining (2) with (3) and Lemma 3 gives $\mathbb{E}[\mathsf{val}(A)] \quad \geq \quad \sum_j (1 - f_j) \cdot p_j + \sum_j f_j \cdot p_j \quad = \quad \sum_{j \in M} p_j$. ☐

Next, we prove the missing Lemma 3.

Proof of Lemma 3. The main idea in the proof to lower bound utility is to define a set $\widetilde{O}_i$ that bidder i could have bought. To define $\widetilde{O}_i$, we consider the optimal allocation where every bidder except bidder i draws independent fresh valuations, i.e., consider the optimal allocation w.r.t. valuations $\mathbf{v}_{-i} := (\widetilde{v}_1, \ldots, \widetilde{v}_{i-1}, v_i, \widetilde{v}_{i+1}, \widetilde{v}_n)$ and let $\widetilde{O}_i$ denote the set of items that bidder i receives in this allocation. Let $q_j^{(i)}$ denote the supporting prices for bidder i w.r.t. set $\widetilde{O}_i$. Now we can lower bound

$$\begin{aligned}
\text{Utility}(i) &\geq \mathbb{E}\Big[\sum_j \mathbf{1}_{j \text{ is unsold till } i} \cdot \mathbf{1}_{j \in \widetilde{O}_i} \cdot (q_j^{(i)} - p_j)^+\Big] \\
&\geq \sum_j \Pr[j \text{ is unsold till } i] \cdot \mathbb{E}_{\mathbf{v}_{-i}}\Big[\mathbf{1}_{j \in \widetilde{O}_i} \cdot (q_j^{(i)} - p_j)^+\Big]
\end{aligned}$$

where we use that $q_j^{(i)}$ and $\widetilde{O}_i$ are independent of whether item j is unsold by the time bidder i arrives. Now using $\Pr[j \text{ is unsold till } i] \geq \Pr[j \text{ is free}] = f_j$ (i.e., unsold till the end) and summing over all bidders, we get

$$\begin{aligned}
\sum_i \text{Utility}(i) \;\geq\; \sum_{i,j} f_j \cdot \mathbb{E}_{\mathbf{v}_{-i}}\Big[\mathbf{1}_{j \in \widetilde{O}_i} \cdot (q_j^{(i)} - p_j)^+\Big] \;&\geq\; \sum_{i,j} f_j \cdot \mathbb{E}_{\mathbf{v}_{-i}}\Big[\mathbf{1}_{j \in \widetilde{O}_i} \cdot (q_j^{(i)} - p_j)\Big] \\
&=\; \sum_j f_j \cdot \sum_i \mathbb{E}_{\mathbf{v}}\Big[\mathbf{1}_{j \in O_i} \cdot (q_j^{(i)} - p_j)\Big],
\end{aligned}$$

where the last equality uses the fact that $\mathbf{v}_{-i}$ and $\mathbf{v}$ are i.i.d.. Finally, using $\mathbb{E}_{\mathbf{v}}[\mathbf{1}_{j \in O_i} \cdot q_j^{(i)}] = 2p_j$ by definition of p_j, we get

$$\sum_i \text{Utility}(i) \geq \sum_j f_j \cdot (2p_j - p_j),$$

which completes the proof of the lemma. ☐

References

[1] M. Feldman, N. Gravin, and B. Lucier. "Combinatorial auctions via posted prices". Proceedings of ACM-SIAM Symposium on Discrete Algorithms (SODA), 2014.