DS 574 Algorithmic Mechanism Design	**Lecture #23**
Fall 2022	Prof. Kira Goldner

Selling Separately or Bundling

Consider the setting with a revenue-maximizing monopolist who has two heterogenous items to sell (an apple and an orange) and a single additive buyer with values v_1 and v_2 for item 1 and 2 respectively. What mechanism should the monopolist use?

Two of the simplest possible mechanisms are as follows:

- (SREV) Sell the items separately. Post a price of p_i on each item i, optimizing these prices to maximize expected revenue from your distribution. The revenue from optimally selling items separately SREV $\geq \sum_i p_i Pr_{v_i}[v_i \geq p_i]$ for valid prices p_i.
- (BREV) Sell the items together in one grand bundle as if they were a single item with some price p. Optimize this price to maximize expected revenue from your distribution. The revenue from optimally selling the grand bundle BREV $\geq p \cdot \Pr_v[\sum_i v_i \geq p]$.

Why Getting the Optimal Mechanism is Tricky

Example 1: Bundling Can Be Better. Suppose $v_1, v_2 \sim U\{1,2\}$.
Selling separately revenue: 2, by selling at 1 or 2 per item.
Grand bundle revenue: 9/4, by selling at 3 w.p. 3/4.

Example 2: Better than Bundling. Suppose $v_1, v_2 \sim U\{0,1,2\}$.
Selling separately revenue: 4/3, by selling at 0 or 1 per item.
Bundle revenue: 4/3, by selling at 2 w.p. 2/3.
Different option: One item at 2, both at 3. Revenue of 13/9.

Example 3: Randomization is Necessary. Suppose $v_1, v_2 \sim F$ where

$$F = \begin{cases} 1 & \text{w.p. } \frac{1}{6} \\ 2 & \text{w.p. } \frac{1}{2} \\ 4 & \text{w.p. } \frac{1}{3}. \end{cases}$$

This gives

$v_1 \backslash v_2$	1	2	4
1	2	3	5
2	3	4	6
4	5	6	8

Deterministic outcomes: Nothing (must be priced 0), the first item, the second item, or both items. We then calculate the probability of sale and expected revenue for each of the following:

1 item:

Price	Pr	Rev
1	1	1
2	5/6	5/3
4	1/3	4/3

2 items:

Price	Pr	Rev
2	1	1
3	35/36	35/12
4	29/36	29/9
5	5/9	25/9
6	4/9	8/3
8	1/9	8/9

Selling separately is best at a price of 2 for a revenue of 5/3; grand bundling is best at 4 for a revenue of 29/9. The optimal deterministic revenue comes from either selling the bundle for 4, or selling a single item for 4 or the bundle for 5, both options earning 29/5.
Randomized option:

- Pay 1 for a lottery ticket that gets the first item w.p. .5.
- Pay 1 for a lottery ticket that gets the first item w.p. .5.
- Pay 4 to get both items with certainty.

This auction yields revenue $3\frac{17}{36}$.

Example 4: But Sometimes Selling Separately is Best. Suppose we have m items where

$$v_i = \begin{cases} 2^i & \text{w.p. } 2^{-i} \\ 0 & \text{otherwise.} \end{cases}$$

Selling separately revenue: Post 2^i for item i for expected revenue m.
Grand bundle revenue: For any price $b \in [2^k, 2^{k+1})$, the buyer will only be willing to pay b for the bundle if they have a high value for some item $i \geq k$. In this case, the selling-separately revenue already captures for revenue for the high items. There is no concentration of values, which is why bundling doesn't help here.

Approximately Optimal Revenue for an Additive Buyer [1]

Note: Refer to Lecture 14 from October 25.

Bounding OPT. Previously, we used the Lagrangian duality framework of CDW '16 [2], weak duality, and the Myersonian-like flow of dual variables to achieve the following upper bound on optimal revenue:

$$\begin{aligned} \textsc{Rev}(F) \leq & \sum_{v \in V} \sum_j f(v) \cdot \hat{x}_j(v) \cdot v_j \cdot \mathbb{1}[v \notin R_j] \quad (\textsc{Non-Favorite}) \\ & + \sum_{v \in V} \sum_j f(v) \cdot \hat{x}_j(v) \cdot \bar{\varphi}_j(v_j) \cdot \mathbb{1}[v \in R_j] \quad (\textsc{Single}) \end{aligned}$$

where R_j is the set of valuations v under which item j is the bidder's favorite item, breaking ties lexicographically (by smallest item index), and $\bar{\varphi}_j(v_j)$ is the standard ironed Myersonian virtual value for item j.

Bounding Single. We also saw that $\textsc{Single} \leq \textsc{opt}^{\textsc{Copies}}$ where the $\textsc{Copies}$ setting is a *single-dimensional* setting with nm single-dimensional bidders, where copy (i, j)'s value for winning is v_{ij} (just one parameter—which is still drawn from F_{ij}). $\textsc{opt}^{\textsc{Copies}}(F)$ is the revenue of Myerson's optimal auction in the copies setting induced by F.

The Core-Tail Split of Non-Favorite. When the bidder is additive, we further decompose non-favorite into two terms we call core and tail by partitioning the valuations by those with $v_j \leq r$ and those with $v_j > r$ for some threshold r, where we will choose r to be the optimal revenue earned by posting a separate price on each item, $r = \textsc{srev}$.

$$\begin{aligned}
&\sum_{v \in V} \sum_j f(v) \cdot \hat{x}_j(v) \cdot v_j \cdot \mathbb{1}[v \notin R_j] \quad (\textsc{Non-Favorite}) \\
&\quad \leq \sum_{v \in V} \sum_j f(v) \cdot v_j \cdot \mathbb{1}[v \notin R_j] \\
&\quad = \sum_j \sum_{v_j > r} f_j(v_j) \cdot v_j \cdot \sum_{v_{-j}} f_{-j}(v_{-j}) \cdot \mathbb{1}[v \notin R_j] + \sum_j \sum_{v_j \leq r} f_j(v_j) \cdot v_j \cdot \sum_{v_{-j}} f_{-j}(v_{-j}) \cdot \mathbb{1}[v \notin R_j] \\
&\quad \leq \sum_j \sum_{v_j > r} f_j(v_j) \cdot v_j \cdot \Pr_{v_{-j}}[v \notin R_j] \quad (\textsc{Tail}) \quad + \sum_j \sum_{v_j \leq r} f_j(v_j) \cdot v_j \quad (\textsc{Core})
\end{aligned}$$

Bounding the Tail.

Lemma 1. $\textsc{Tail} \leq \textsc{srev}$.

Proof. By the definition of R_j, for any given v_j,

$$\Pr_{v_{-j}}[v \notin R_j] \quad \leq^1 \quad \Pr_{v_{-j}}[\exists k \neq j, v_k \geq v_j].$$

Posting a price of v_j on each item separately earns revenue at least $v_j \cdot \Pr_{v_{-j}}[\exists k \neq j, v_k \geq v_j]$, as in this case, the buyer will purchase at price v_j with certainty. Then

$$v_j \cdot \Pr_{v_{-j}}[v \notin R_j] \leq v_j \cdot \Pr_{v_{-j}}[\exists k \neq j, v_k \geq v_j] \leq \textsc{srev}$$

for all v_j. Thus

$$\textsc{Tail} \leq \textsc{srev} \cdot \sum_j \sum_{v_j > r} f_j(v_j) = \sum_j r \cdot \Pr_{v_j}[v_j > r] \leq \textsc{srev}.$$

because r is a valid price to post on each item, and $\textsc{srev}$ is the revenue from the optimal item prices. □

[1] It would be $=$ but for lexicographical tie-breaking to determine R_j.

Lemma 2. *If we sell the grand bundle at price* $\textsc{Core} - 2r$*, the bidder will purchase it with probability at least* $1/2$*. Then* $\textsc{brev} \geq \frac{\textsc{Core}}{2} - r$*, or* $\textsc{Core} \leq 2\textsc{brev} + 2\textsc{srev}$*.*

To prove this, we need the following fact from a technical lemma in CDW.

Technical Fact (CDW '16). Let x be a positive signle dimensional random variable drawn from F of finite support such that for any number a, $a \cdot \Pr_{x\sim F}[x \geq a] \leq \mathcal{B}$ where $\mathcal{B}$ is an absolute constant. Then for any positive number s, the second moment of the random variable $x_s = x \cdot \mathbb{1}[x \leq s]$ is upper-bounded by $2\mathcal{B} \cdot s$.

Proof. For each item j define a new random variable c_j that 0's out the part of the distribution not in the core as follows: Draw a sample v_j. If $v_j \in [0, r]$, then $c_j = v_j$. Otherwise, $c_j = 0$.

Let $c = \sum_j c_j$ be the sum of these new core random variables. Notice that $\mathbb{E}[c] = \sum_j \sum_{v_j \leq r} f_j(v_j) \cdot v_j$. We now show that c, the sum of item values drawn only from the core, concentrates, because it has small variance.

We wish to show that pricing the grand bundle at $\mathbb{E}[c] - 2r$ sells with probability at least $1/2$, that is, that

$$\Pr_{v\sim F}\left[\sum_j v_j \geq \mathbb{E}[c] - 2r\right] \geq \frac{1}{2}.$$

You may wish to use the following:

- $\mathrm{Var}[X] = \mathbb{E}[X^2] - \mathbb{E}[X]^2$.
- If $Z = \sum_i X_i$ and the X_i's are drawn independently, then $\mathrm{Var}[Z] = \sum_i \mathrm{Var}[X_i]$.
- Chebyshev's inequality: $\Pr[|X - \mathbb{E}[X]| \geq t] \leq \mathrm{Var}[X]/t^2$ where $t > 0$.
- The above technical fact.
- Let $r_j = \max_x\{x \cdot \Pr_{v_j}[v_j \geq x]\}$ be the optimal selling-separately revenue from item j, and $r = \sum_j r_j$.

Because the c_j's are independently drawn, then

$$\mathrm{Var}[c] = \sum_j \mathrm{Var}[c_j] \leq \sum_j \mathbb{E}[c_j^2].$$

We will bound each $\mathbb{E}[c_j^2]$ separately. Let $r_j = \max_x\{x \cdot \Pr_{v_j}[v_j \geq x]\}$. By the above fact, we can upper bound $\mathbb{E}[c_j^2]$ by $2r_j \cdot r$. On the other hand, $r = \sum_j r_j$ (as this is the definition of $\textsc{srev}$), so $\mathrm{Var}[c] \leq 2r^2$. By the Chebyshev inequality,

$$\Pr[c \leq \mathbb{E}[c] - 2r] \leq \frac{\mathrm{Var}[c]}{4r^2} \leq \frac{1}{2}.$$

Therefore,

$$\Pr_{v\sim F}\left[\sum_j v_j \geq \mathbb{E}[c]-2r\right] \geq \Pr[c \geq \mathbb{E}[c]-2r] \geq \frac{1}{2}.$$

Then $\textsc{brev} \geq \frac{\mathbb{E}[c]-2r}{2}$, as we can sell the grand bundle at price $\mathbb{E}[c]-2r$, and it will be purchased with probability at least 1/2. □

Theorem 1. *For a single additive bidder, the optimal revenue is* $\leq 2\textsc{brev} + 4\textsc{srev}$.

Proof. As shown previously, $\textsc{opt} \leq \textsc{Single} + \textsc{Core} + \textsc{Tail}$ and $\textsc{Single} \leq \textsc{opt}^{\textsc{Copies}}$. Note that in the additive setting, the optimal revenue from the copies setting is equal to $\textsc{srev}$, as each item is sold separately with no feasibility constraint.

Lemma 1 shows that $\textsc{Tail} \leq \textsc{srev}$ and Lemma **??** shows that $\textsc{Core} \leq 2\textsc{srev} + 2\textsc{brev}$, hence

$$\textsc{opt} \leq 2\textsc{brev} + 4\textsc{srev} \leq 6\max\{\textsc{srev}, \textsc{brev}\},$$

so the better of selling separately or selling the grand bundle gives a 6-approximation to the optimal revenue for a single additive buyer. □

References

[1] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A Simple and Approximately Optimal Mechanism for an Additive Buyer. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 21–30. IEEE, 2014.

[2] Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. A Duality Based Unified Approach to Bayesian Mechanism Design. In *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing*, STOC '16, pages 926–939, New York, NY, USA, 2016. ACM.