CS880: Algorithmic Mechanism Design

Lecture 6: Simplicity vs Optimality in Multi-parameter Settings **Scribe:** Jeremy McMahan

6.1 Unit-demand Setting

In a *Multi-paramter* setting, each agent has a value for many different outcomes. We usually focus on the context of combinatorial auctions, where each buyer may want to obtain several items instead of just one. The allocation of a mechanism assigns to each agent a set describing which items this agent obtains. The total value of the set of items obtained is defined by a value function $v(S)$, which could be arbitrarily complicated. In contrast, in the single parameter setting, each agent only cares about a single outcome (getting the desired item or not) and can describe his value for this outcome with a single number instead of needing an entire function. We will focus on two simple multi-parameter settings. In the first, called the *unit-demand buyer setting*, each agent, also called a buyer, will only receive the value of the highest valued item in the allocated set. Formally, $v(S) = \max_{i \in S} v_i$ where v_i is the buyer's value for item i. In the *additive setting*, we have the value of an allocated set is exactly the total value of each time in the set, $v(S) = \sum_{i \in S} v_i$. Note, we will assume that agent's values for different items are drawn from independent distributions.

In the multi-paramter setting, the optimal mechanism is described by menus that list allocation, price pairs. Hence, each agent may purchase different sets of items at different prices. However, the optimal mechanism may require lotteries which sell a distribution of items at some fixed price. In this context it was shown that certain value distributions require unbounded sized menus, so we focus on approximation mechanisms. In particular, we will try to find simple mechanisms, notably per-item posted-pricings, continuing with the theme of simplicity vs optimality in mechanism design.

Let's first consider a single unit-demand buyer. We want to show that fixing a price for each item leads to an approximately optimal mechanism. Most of the proof was given last time, so we just need to complete the final claim. We refer to *Setting 1* as the setting where we have a single unit-damand buyer and the values for all the items are drawn from a product distribution, $v \sim F_1 \times \ldots \times F_n$. We let OPT denote the optimal mechanism's allocation in Setting 1. On the other hand, we refer to *Setting 2* as the setting where we have m agents, a single item, and the value of each agents value is drawn from an independent distribution, $v_i \sim F_i$. Note, the optimal mechanism in Setting 2 is exactly Myerson's mechanism as discussed before, so we refer to this mechanism as MYE. We argued previously that item pricing form Setting 1 is has revenue at least 1/2 that of Myerson's mechanism applied to Setting 2. We need to show that:

Claim 6.1.1 $MYE \geq \frac{1}{2} OPT$

Note that this claim is not easily generalizable to other settings like the additive one. After this claim has been established, we get overall:

Corollary 6.1.2 *Item pricing achieves revenue at least* $1/4$ *that of* OPT

This is the best approximation known at this point in time. Let's prove the claim.

Proof:

We proceed by using OPT to construct an incentive compatible mechanism for Setting 2, $MECH$. Being that the best mechanism for Setting 2, MYE, must achieve at least as much revenue as any other mechanism for the setting, we will have the claim.

Recall, in Setting 1, the optimal menu may have lotteries and so the utility of the buyer may look like $\sum_{i=1}^{n} x_i v_i - p$ where p is the price of the lottery and x_i is the probability of getting item i having value v_i. We want to use the same allocation so that social surplus of the n agents in Setting 2 is exactly $\sum_{i=1}^{n} x_i v_i$. However, to make this incentive compatible and achieve good revenue we need to break the single payment p into payments to each individual agent.

For a fixed value vector v, we define the favorite item of the unit-demand agent as $i^* = \underset{i \in [n]}{\text{argmax}}\ v_i$. Our constructed mechanism for Setting 2 will then take the reported vector of values of each agent v, and assign to i^* a fraction $x_{i^*}(v)$ to this agent and charge price $p - \sum_{i \neq i^*} x_i(v) v_i$. The utility of friend i in Setting 2, will be

$$
\begin{aligned}
& x_{i^*}(v) v_{i^*} - (p - \sum_{i \neq i^*} x_i(v) v_i) \\
& = \sum_i x_i(v) - p \\
& = \text{Utility of unit-demand agent in Setting 1}
\end{aligned}
$$

Thus, if friend i wants to increase his utility by deviating from the colluding strategy, then the unit-demand agent in Setting 1 could have also increased his utility. Thus, incentive compatibility of OPT guarantees that of our constructed mechanism.

The revenue of our mechanism is

$$
\begin{aligned}
& E_v[p - \sum_{i \neq i^*} x_i(v) v_i] \\
& = E_v[p] - E_v[\sum_{i \neq i^*} x_i(v) v_i] \\
& = OPT - E_v[\sum_{i \neq i^*} x_i(v) v_i] \\
& \geq OPT - E_v[\max_{i \neq i^*} v_i]
\end{aligned}
$$

Note since i^* was the favorite item in the original setting, the second term above is exactly the second favorite item's value. Thus, the Vickrey auction in Setting 2, would give exactly revenue $E_v[\max_{i \neq i^*} v_i]$. Hence, we get the revenue of the constructed mechanism is at least $OPT - Rev(Vickrey)$. Consequently, the revenue of $MECH$ plus the revenue of Vickrey gives at least OPT. Now, both of $MECH$ and Vickrey are IC mechanisms for Setting 2, and so MYE has larger revenue than both of them. Thus, $2MYE \geq OPT$.

■

6.2 Additive Setting Examples

Recall, we have m items, one agent, and the values for each item is drawn independently, $v \sim F_1 \times \ldots \times F_m$. The value of a set of items, S, is $v(S) = \sum_{i \in S} v_i$. We consider 4 different examples in this setting.

1. Suppose we have two items with values $v_1, v_2 \sim U\{1, 2\}$. What would the optimal price be for item 1? In this case, we could price item 1 at 1 or 2 and get equally good revenue. Overall, we have a revenue of 1 dollar per item, for a total revenue of 2 dollars. However, if we sell a bundle of items we can do better. Note the distribution of a bundle of items is $U\{2, 3, 3, 4\}$. Then, the optimal bundling revenue results from posting a price of 3 for the bundle. The revenue achieved is $\$3 \cdot \frac{3}{4} > \$2$.

2. Now, if we had 100 items all drawn from the same $U\{1, 2\}$ distribution, we could again sell them separately and get a revenue $SRev = 100$. If we bundle them together, we get a bundling revenue $BRev$. Note that we have a sum of random variables, and so they are very concentrated around the mean, 150. In particular, the deviation around the mean is at most 10 for 100 items. Thus, we can post a price slightly under the mean, say at 140, and still make a sale the majority of the time yet get very close to 150 in revenue. Hence, $BRev \approx 150 -$ something small. Though, bundling is not always optimal unless the items are distributed identically as we see next.

3. Suppose we have n items and the distribution for item i is:

$$v_i = \begin{cases} 0 & w.p.\ 1 - 1/(n \cdot 2^i) \\ 2^i & w.p.\ 1/(n \cdot 2^i) \end{cases}$$

   The probability of having one item of non-zero value in a set is very small, so we expect bundling to give a very small revenue. Specifically, it behaves like the equal revenue distribution which is defined by $F(x) = 1 - 1/x$ giving a flat revenue curve $R(p) - 1$. In this case, we have the selling separately strategy's revenue will be much better than bundling. So, we have that sometimes selling separately is good and sometimes bundling is good. The question becomes which items to bundle and which to sell separately. Though we can also sell lotteries in this context too giving even stranger phenomenon. We consider this next.

4. Again we have two items but now having different value distributions $v_1 \sim U\{1, 2\}, v_2 \sim U\{2, 3\}$. Note, $v_1 + v_2 \sim U\{2, 3, 4, 5\}$. The best mechanism that sells separately achieves revenue $SRev = \$1 \cdot 1 + \$3 \cdot 1/2 = \$2.5$. The best mechanism that bundles everything

achieves revenue $BRev = \$3 \cdot 3/4 = \$2.25$. Lastly, the optimal mechanism for this example uses lotteries.

$$OPT \text{ menu} = \begin{cases} (1,1) & \$4 \\ (1,\frac{1}{2}) & \$2.5 \\ (0,0) & \$0 \end{cases}$$

The optimal revenue is then $\$4 \cdot \frac{1}{2} + \$2.5 \cdot \frac{1}{4} = \$2.625$.

6.3 Additive Setting Approximation

The main result is that the better of the two mechanisms that sell everything separately and bundle everything will be approximately optimal [1].

Theorem 6.3.1 $\max(SRev, BRev) \geq \frac{1}{6}OPT$.

Lemma 6.3.2 (Lemma 1) *Let S and T be disjoint sets of items. Let F be a distribution over the values of S and F' a distribution over the values of T. Then,*

$$Rev(F \times F') \leq Rev(F) + Val(F')$$

where $Val(F') = E_{F'}[\sum_{i\in T} v_i]$

Proof:

The proof follows similarly to Claim 7.1.1 from before. Let MP be the optimal mechanism for $F \times F'$. We will define a mechanism MF for F. Given $v_{|S} \in F$ we draw $v_{|T} \in F'$, to construct a total vector v for $F \times F'$. Suppose now MP outputs (x, p) when given values from v. We define MF so that MF will assign $x_{|S}$ on $v_{|S}$ and charge a price of $p - \sum_{i\in T} x_i v_i$. A similar argument as before implies that MF will be IC. Also, we have

$$\begin{aligned} Rev(MF) &= Rev(MP) - E_{v\sim F'}[\sum_{i\in T} x_i v_i] \\ &\leq Rev(MP) - E_{v\sim F'}[\sum_{i\in T} v_i] \\ &= Rev(MP) - Val(F') \end{aligned}$$

■

We will use a *Core-Tail decomposition* defined as follows. Given a random $v \sim F_1 \times \ldots \times F_m = F$ and given a fixed threshold t_i for each i, we say i is in the *core* of F if $v_i \leq t$ and we say i is in the *tail* if $v_i > t_i$. Note that $Pr[i \in \text{tail}] = Pr[v_i > t] = q_i$ and hence the name tail. Let r_i be the revenue from an item pricing for i which is equal to $\max_p p(1 - F_i(p))$. We also define $r = SRev = \sum_i r_i$. We have $t_i q_i = t_i Pr[v_i > t_i] \leq r_i$. So, $q_i \leq \frac{r_i}{t_i}$.

Lemma 6.3.3 (Lemma 2) *Suppose we partition the space of all value vectors into components $C_1, C_2, \ldots$. Let $F_{|C_j}$ be the distribution over values conditioned on being in C_j. Then,*

$$Rev(F) \leq \sum_j Pr[C_j] Rev(F_{|C_j})$$

Lemma 1 discusses a partition over a set of items for a fixed value vector whereas Lemma 2 partitions the set of possible value vectors itself.

For every set $A \subseteq [m]$, define $C_A = \{v | \forall i \in A, v_i > t_i \wedge \forall i \notin A, v_i \leq t_i\}$. In words, every item in A must be in the tail and all items not in A must be in the core, so C_A are all the value vectors having tail A. Let $F_A = F_{|C_A}$. From Lemma 2, we have that

$$Rev(F) \leq \sum_A Pr[tail = A] Rev(F_A)$$

We can apply Lemma 1 to the last revenue term to get

$$\begin{aligned} Rev(F) &\leq \sum_A Pr[A]\{Rev(F_A^{tail}) + Val(F_A^{core})\} \\ &\leq \sum_A Pr[A] Rev(F_A^{tail}) + Val(F_A^{core}) \end{aligned}$$

We first want to bound the revenue of the tail. Recall, in the unit-demand setting, $Rev(D) \leq 4SRev(D)$. For some set A and additive values over this set,

$$Rev(D_{|A}) \leq |A| Rev(\text{unit-demand buyer over } A) \leq 4SRev(D_{|A})$$

Hence,

$$\sum_A Pr[A] Rev(F_A^{tail}) \leq 4SRev(F) E[|tail|]$$

We can bound the expectation term by

$$E[|tail|] = \sum_i q_i \leq \sum_i \frac{r_i}{t_i} = \frac{\sum_i r_i}{SRev} = \frac{SRev}{SRev} = 1$$

We see that we have a small tail in expectation, so the first part of the $Rev(F)$ expression is bounded by $4SRev$. Now, we can focus on the value of the core term. Note, if i is in the core, then $v_i \in [0, r]$. We want to show that $S = \sum_{i \in core} v_i$ is concentrated around its mean to use the bundling results from before. Our strategy will be to relate the expectation of S to its median.

Lemma 6.3.4 (Lemma 3) *Let $S = \sum_{i \in core} v_i$. There exists constants c_1, c_2 so that*

$$E[S] \leq c_1 Median(S) + c_2 r$$

Note in this context, the median of the sum will be $BRev$ and r is $SRev$.

Recall, the best revenue from pricing i is r_i. Consequently, for any p, we know that $p(1-F_i(p)) \leq r_i$. Rearranging this inequality gives $F_i(p) \geq 1 - \frac{r_i}{p}$, so $f_i = \frac{r_i}{p^2}$. We will eventually use Chebyshev's inequality, so we need a bound on the variance of S. We compute such a bound by bounding the variance of each term in the sum S. For a fixed i in the core, we have

$$Var(v_i) \leq E[v_i^2] \leq \int_0^r t^2 \frac{r_i}{t^2} dt + r^2 \frac{r_i}{r} = 2r_i r$$

Thus,

$$Var(S) \leq \sum_i Var(v_i) \leq \sum_i 2r_i r = 2r^2$$

Let $\mu = E[S]$. Applying Chebyshev's inequality gives

$$Pr[S \leq \frac{1}{4}\mu] \leq Pr[|S-\mu| > \frac{3}{4}\mu] \leq \frac{Var(S)}{(\frac{3}{4}\mu)^2} \leq \frac{2r^2 16}{9\mu^2} = \frac{32}{9}\frac{r^2}{\mu^2}$$

Now, this bound gives different implications depending on whether $SRev$ is close to the mean or not.

1. If $SRev \geq \frac{1}{4}\mu = \frac{1}{4}Val(core)$, we are done.
2. If $SRev < \frac{1}{4}\mu$, then $Pr[S \leq \frac{\mu}{4}] \leq \frac{2}{9}$. Hence, $BRev \geq \frac{\mu}{4}\frac{7}{9} \geq \frac{\mu}{4} = \frac{1}{4}Val(core)$.

In either case, $\max(SRev, BRev) \geq \frac{1}{4}Val(core)$. Now, we use this to bound the overall optimal revenue, $Rev(F)$.

$$\begin{aligned} OPT &\leq Rev(tail) + Val(core) \\ &\leq Const \cdot SRev + Const \cdot \max(SRev, BRev) \\ &\leq Const \cdot \max(SRev, BRev) \end{aligned}$$

Using the straightforward bounds presented here would give a constant of 8, though working out the details more strictly yields the best known constant of 6.

6.4 Extensions

1. The single buyer setting with independent item values can be extended subadditively to sets. In particular, it is shown that the better of item pricing and bundle pricing is still a constant approximation.

2. The multiple buyers setting with independent item values can similarly be extended subadditiviely to sets [2]. Approximate optimality is achieved through a sequential mechanism that offers to each buyer a *two-part tariff*. Two-part tariff mechanisms use an entry fee with per

item prices. For example, to buy from Amazon you first need an Amazon membership which costs a fee. Then, you buy items at a fixed price per item. The per item pricing gives us the behavior of a selling individually mechanism while the entry fee gives behaviour of a grand bundling mechanism since the entry fee acts as a payment for access to all items like buying a bundle.

Using an Ex-ante Relaxation yields the above result. For a single buyer, define the revenue of this buyer, $Rev_{q_i}(F_i)$, where $q_i = (q_{i1}, \ldots, q_{im})$ is a vector to be the maximum revenue we can achieve from buyer i under the constraint that item j is allocated to i with probability at most q_{ij}. Note the q_i represent the ex-ante supply constraints. In particular, this constraint requires that $\sum_i q_i \leq (1, \ldots, 1)$. We then have the total revenue is

$$Rev(F) \leq \max_{\sum_i q_i \leq (1,\ldots,1)} \sum_i Rev_{q_i}(F_i)$$

Note independent item values and subadditivity is crucial in all of these results.

References

[1] Moshe Babaioff et al. *A Simple and Approximately Optimal Mechanism for an Additive Buyer.* 2020. arXiv: `1405.6146 [cs.GT]`.

[2] Shuchi Chawla and J. Benjamin Miller. *Mechanism Design for Subadditive Agents via an Ex-Ante Relaxation.* 2016. arXiv: `1603.03806 [cs.GT]`.