TTIC 31260 - Algorithmic Game Theory (Winter 2026)

Lecturer: Avrim Blum — Lecture 11: February 16, 2026

Mechanism Design III: Revenue - Pricing Digital Goods

In the last two lectures, we considered the problem of maximizing social welfare. Today, we will be looking at the problem of maximizing revenue. We will be focusing in particular on the case where we have what are called *digital goods* which are things that we can make an unlimited number of copies of (like music or software). Our goal will be to assign prices to these items in order to maximize our revenue. Note that since we are in the unlimited-supply setting, maximizing social welfare is trivial by giving everything away at a price of 0. But our goal instead is to maximize *revenue*. Digital goods are in some sense the cleanest case for studying this because it's clear that if you want good revenue, you need to do something very different than charging people their externality (which is 0).

1 Selling a single digital item

We first consider the case where we're selling a single item with unlimited supply. We assume each bidder/buyer i has some value v_i on getting the item, and nobody wants more than one copy. We will also assume that all valuations are between 1 and H. The question we will examine is what fraction of the optimal social welfare can we hope to extract as revenue?

We will consider the following very simple kind of mechanism: assigning a price to the item, and charging that price to everyone.

First, one useful thing to notice is that if buyers were arriving from a fixed probability distribution, then we would optimize revenue by charging:

$$\arg\max_p [p \cdot \Pr(v_i \geq p)]. \tag{1}$$

What can we say about how this revenue compares to social welfare? First, we have the following lower bound.

Theorem 1 *There exist distributions over $[1, H]$ such that the expected revenue for any price is at most a $\frac{1}{\log H}$ fraction of the expected optimal social welfare.*

Proof: A distribution meeting this lower bound is what's called the "equal revenue distribution". Specifically, suppose that each buyer's valuation is drawn uniformly from $H, H/2, H/3, ..., H/H$.

Notice that if we set price $p = H$ then there is a $1/H$ chance that a buyer has $v \geq p$, giving us expected revenue of 1 per buyer. If we set price $p = H/2$ then there is a $2/H$ chance that a buyer has $v \geq p$, also giving us expected revenue 1 per buyer. More generally, if we set $p = H/i$ then there is an i/H chance that a buyer has valuation $v \geq p$, again giving us expected revenue 1 per buyer. Any other price $\frac{H}{i} > p > \frac{H}{i+1}$ is strictly worse because it has the same probability for $v \geq p$ as $\frac{H}{i}$, and so will only produce less revenue.

But the expected *social welfare* is $1 + \frac{1}{2} + \frac{1}{3} + ... + \frac{1}{H} \approx \log H$, giving us get the result. ■

We now give an upper bound: in fact, a randomized algorithm such that no matter what the buyer's valuation—so long as it is in the range $[1, H]$—the expected revenue is within an $\Omega(1/\log H)$ factor of the social welfare (i.e., of the buyer's valuation).

Theorem 2 *There exists a randomized algorithm that for any buyer (or any sequence of buyers) gets revenue within an* $\Omega(\frac{1}{\log H})$ *factor of the optimal social welfare.*

Proof: In order to do this we need a randomized strategy, since for any fixed price, we can have an adversary that gives us a bidder who has valuation just slightly too low to buy.

Let v denote the (unknown to us) valuation of the bidder. Notice that if we can choose a price between $v/2$ and v, then we will get revenue at least half of the social welfare. So, what we'll do is pick the price to be a random power of 2 between 1 and H (i.e., equally likely $1, 2, 4, 8, ..., 2^{\lfloor \log_2 H \rfloor}$). This gives us at least a $\frac{1}{1+\log_2 H}$ chance of picking a price between $v/2$ and v, giving us expected revenue $\Omega(\frac{v}{\log_2 H})$ as desired. ■

2 Multiple items

Now, we consider the case where we have n items. Buyers have an arbitrary valuation function over subsets $v_i : S \subseteq \{1, 2, ..., n\} \to \mathbb{R}^+$, where the maximum valuation for any subset is $1 \leq \max_S v_i(S) \leq H$.

First, suppose we were allowed to assign prices to "bundles" or subsets of items, and we assume free disposal so that $\arg\max_S v_i(S) = \{1, ..., n\}$. Then, this problem can be converted to the original setting, by bundling everything together and selling it as a single item.

However, suppose we can only assign prices to items, not bundles. It turns out that we can do almost as well. We can get $\Omega(\frac{1}{\log(nH)})$ of the optimal social welfare. This algorithm and analysis comes from [BBM08].

Before describing the algorithm, we first present a lower bound showing why the $1/\log n$ term is necessary: in particular, this is needed even if you know the buyers' valuations completely.

Theorem 3 *There exist buyer valuation functions such that any item pricing will produce revenue that is at most a* $\frac{1}{\log n}$ *fraction of the maximum social welfare.*

Proof: Suppose the buyers all have the following valuation function. Any one item is worth 1 to the buyers. Any two items is worth $1 + \frac{1}{2}$ to the buyers. Any three items is worth $1 + \frac{1}{2} + \frac{1}{3}$ to the buyers, and so on.

If we price everything at \$1 dollar, we make \$1 per buyer, as everybody only wants to buy one. If we price everything at \$0.5, then we make \$1 per buyer, since everybody wants two items.

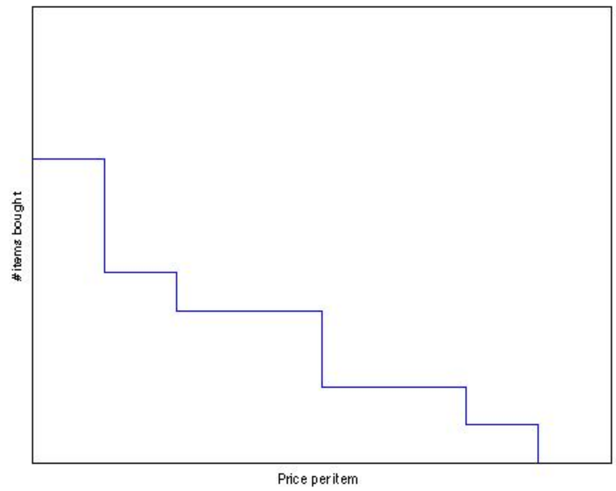

Figure 1: A function showing the number of items a buyer would purchase as a function of the price.

In fact, we can never make more than \$1 per buyer even if we assign different prices to different items. Any buyer always want to sort items by price, and then buy them in that order starting with the cheapest, buying a new item if and only if their utility increases (because the marginal value for each new item is going down and the marginal cost is going up). This means the buyer will only buy the i'th item if its price is less than or equal to $\frac{1}{i}$, and this means that the total amount spent is at most 1. On the other hand, the maximum social welfare for each buyer is $1 + 1/2 + 1/3 + \ldots + 1/n \approx \log(n)$. ■

We now present the upper bound. First, a few helpful lemmas.

Consider one buyer and the function $f(p) : \mathbb{R} \to \mathbb{Z}^+$, giving the number of items purchased for any given price p (all items have the same price p). This function is a step function, and an example is shown in Figure 1, where the x axis is the price of the item and the y axis is the number of items purchased.

Lemma 1 *$f(p)$ is non-increasing.*

Proof: Consider some price p, and let S be the bundle purchased at that price. For any bundle S' such that $|S'| > |S|$, S' only becomes less attractive compared to S as prices are increased, since an increase in price Δp increases the price of S' by $|S'| \cdot \Delta p$, whereas the price of S only increases by $|S| \cdot \Delta p$. Hence, if the buyer was not buying S' before, they certainly won't want to switch to it as the price increases. Another way to think of this is that the utility for any given bundle is a decreasing function of price, with slope equal to the size of the bundle. So, the size of the maximum utility bundle can only drop as price is increased. ■

Lemma 2 *$\max_S v(S)$ is equal to the area under the curve of $f(p)$.*

Proof: The buyer's utility at price p is $\max_S[v(S) - p|S|]$. So, consider price $p = 0$ where the buyer's utility is $\max_S v(S) = v(S_0)$ where S_0 is the bundle purchsed at price 0. Now, let p_1 be the price at which the first step occurs. Notice that in the interval $[0, p_1]$, set S_0 is still a utility-maximizing bundle, so the buyer's utility at price p_1 is $v(S_0) - p_1|S_0|$; that is, by price p_1, the buyer's utility has decreased by exactly the area under the curve for the interval $[0, p_1]$. Now, let S_1 be the new utility-maximizing bundle in the interval $[p_1, p_2]$ where p_2 is the price at which the next step occurs. In that interval, the utility of the buyer decreases from $v(S_1) - p_1|S_1|$ to $v(S_1) - p_2|S_1|$, so again it decreases by exactly $f(p)$ times the width of the interval, which is the area under the curve for that interval. This continues until finally we reach a price where the buyer purchases nothing and the utility of the buyer is 0. Since for each interval, the buyer's utility decreases by the area under the curve for that interval, the area under the curve equals the initial utility, $\max_S v(S)$. ■

With these lemmas in hand, we now give our algorithm: we will pick a price p uniformly at random from $H, H/2, H/4, ..., \frac{1}{4n}$, and price every item at that price.

Theorem 4 *Choosing a random price $p \in \{H, H/2, H/4, ..., \frac{1}{4n}\}$ yields expected revenue within a factor $\Omega(\frac{1}{\log(nH)})$ of the optimal social welfare.*

Proof: For any given price p, we receive revenue equal to p times the number of items bought, i.e., the area of the associated rectangle. So, consider the rectangles corresponding to the prices $H, H/2, H/4, ..., \frac{1}{4n}$ and let's consider the sum of their areas. This sum of areas is at least the area of the union of the rectangles. The area of the union of the rectangles is at least $\frac{1}{2}$(area under curve)$-\frac{1}{4}$, because each region that is uncovered by the union of the rectangles has area at most that of the rectangular region to its left, except for the thin region at the very top of the leftmost rectangle, whose area is at most $n \cdot \frac{1}{4n} = \frac{1}{4}$. Since we are considering $\log(4nH)$ different prices, it follows that picking uniformly between these prices gives us average revenue at least $\frac{\frac{1}{2}(optimal\ social\ welfare) - \frac{1}{4}}{\log(4nH)}$. ■

References

[BBM08] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour, "Item Pricing for Revenue Maximization." *Proc. 9th ACM Conference on Electronic Commerce*, pp. 50-59. 2008.