Profit Maximization, Digital Goods, and the Random Sampling Auction

In previous lectures, we have extensively studied auction design for *welfare maximization.* In this lecture, we consider a different objective: *profit maximization* for the auctioneer. As a case study, we will consider a particular setting (digital goods auctions) in which the two objectives are in dramatic conflict.

Informally, digital goods auctions (also known as *unlimited supply auctions*) model the setting in which a single seller is selling goods that have zero marginal cost of production (think software, which can be copied for free once produced). Hence, there is no natural constraint on how many individuals can be selected to "win" the auction.

Definition 1 *A digital goods auction is a single parameter domain with a set of alternatives $A = \{S \subseteq [n]\}$ – i.e.* any *set of bidders is a feasible outcome, because there is no supply constraint. For $a \in A$ we write $a_i = \begin{cases} 1, & \text{if } i \in S; \\ 0, & \text{otherwise.} \end{cases}$. Since it is a single parameter domain, each bidder's valuation function is parameterized by $v_i \in \mathbb{R}_{\geq 0}$, and $v_i(a) := v_i \cdot a_i$.*

First, lets observe that truthful welfare maximization and profit maximization are in conflict here. The VCG mechanism maximizes social welfare while making truth-telling a dominant strategy. What would it do?

Note that since supply is unlimited, the social welfare maximizing outcome (i.e. the one always chosen by the VCG mechanism) sets $S = \{1, \ldots, n\}$ – i.e. everyone is a winner! Since everybody wins no matter what, no bidder has any negative externality on any other, and so the VCG payments are $p_i = 0$ for all bidders i. The welfare maximizing auction obtains zero revenue!

So if we want to maximize revenue, we need to do something else – in particular, we'll need to artificially limit supply. Before we think about *how* to do that though, lets think about what our benchmark should be. What is the "optimal" revenue?

Certainly no individually rational mechanism can obtain revenue higher than the welfare it obtains, so $\sum_{i=1}^{n} v_i$ is an upper bound on the revenue we can obtain. But its not a reasonable upper bound – no truthful mechanism can obtain full welfare extraction. Instead, we focus on a simple benchmark (that turns out can be theoretically justified as the "optimal" benchmark in some sense) – the revenue we could obtain with the best *fixed* price.

If se set a price of p, everyone with value $v_i \geq p$ will buy. Hence, the revenue of price p is $p \cdot |\{i : v_i \geq p\}|$. Its not hard to see that the best fixed price in hindsight is always $p \in \{v_1, \ldots, v_n\}$. (why?) Hence, we can define the revenue of the best fixed price as:

$$\text{OPT}(v) = \max_i v_i \cdot |\{j : v_j \geq v_i\}| = \max_i (i \cdot v_{(i)})$$

where $v_{(i)}$ is the i'th highest valuation in sorted order.

However, even this benchmark is not achievable by a truthful mechanism if the optimal profit is obtained with $i = 1$ – i.e. by setting the price equal to the value of the highest bidder. To see this, consider the case of just a single bidder – to obtain this benchmark, we would have to charge him his value exactly. But observe that a truthful mechanism can't charge any bidder a price that depends on his own bid, so this is impossible.

Instead, we focus on a relaxed benchmark – the revenue of the best fixed price *that sells to at least 2 people.* Note that we shouldn't think of this as a very serious restriction in large markets, since typically the optimal revenue is not obtained by having only 1 buyer... We write:

$$\text{OPT}^{\geq 2}(v) = \max_{i \geq 2} \left(i \cdot v_{(i)}\right)$$

Our first thought might simply be to take in the bids v, compute the best fixed price v_j, and charge that – but this won't yield a truthful mechanism, because bidder j could change the price he has to pay by manipulating his bid.

Our *next* thought might be to offer each i he price p_i corresponding to $\text{OPT}^{\geq 2}(v_{-i})$ – i.e. the best fixed price excluding agent i. This yields a truthful mechanism, but does terribly compared to our benchmark. Consider the following example:

Example 1 *Suppose we have 90 "low value" agents with $v_i = 1$, and 10 "high value" agents with $v_i = 10$. $OPT^{\geq 2}(v) = 100$, achieved by charging either $p = 10$ or $p = 1$. But for $v_i = 1$, $OPT^{\geq 2}(v_{-i}) \leftrightarrow p_i = 10$, and for $v_i = 10$, $OPT^{\geq 2}(v_{-i}) \leftrightarrow p_i = 1$. So this auction gets profit only 10... (And the ratio to $OPT^{\geq 2}(v)$ can be made arbitrarily bad.)*

Profit Extractors We will start with an intermediate goal. Given a target profit R, we want a mechanism that will obtain profit R *if* $\text{OPT}^{\geq 2}(v) \geq R$. Otherwise we won't require any revenue guarantee for the mechanism. We call such a mechanism a profit extractor.

Definition 2 *The digital goods profit extractor with target profit R, which we write as Extract(R, v), does the following: it finds the largest value k such that $v_{(k)} \geq R/k$, and then sells to the top k bidders at price R/k. If there is no such k, it sells to nobody.*

Lemma 3 *Extract(R, v) is dominant strategy truthful.*

Proof We can view the profit extractor as running the following process: It starts with $i = n$, and offers a price of $p = R/i$ to the top i bidders. If the i'th highest bidder *rejects* the offer (i.e. $v_{(i)} < R_i$), then it sets $i \leftarrow i - 1$ and repeats the offer (now a higher offer, to 1 fewer bidders). If the i'th highest bidder *accepts* the offer, then the top i bidders receive the good and pay the last offer price. Note that if any bidder rejects the offer, she can never win in any future round, so rejecting any offer of $p < v_i$ is a dominated strategy. Similarly, accepting an offer of $p > v_i$ is a dominated strategy since it ends the auction and forces everyone to pay p_i. Hence the dominant strategy for every bidder i is to report their true value. ■

Lemma 4 *Extract(R, v) obtains revenue R if $OPT^{\geq 2}(v) \geq R$, and otherwise obtains revenue 0.*

Proof Recall that $\text{OPT}^{\geq 2}(v) = k \cdot v_{(k)}$ for some $k \in \{2, \ldots, n\}$. If $\text{OPT}^{\geq 2}(v) \geq R$ then $v_{(k)} \geq \frac{R}{k}$. Hence, the profit extractor finds some $k' \geq k$ such that $v_{(k')} \geq R/k'$, and obtains profit $k' \cdot R/k' = R$.

On the other hand, if $R > \text{OPT}^{(2)}(v) = \max_k k \cdot v_{(k)}$, then by definition there is no k such that $v_{(k)} \geq R/k$, and so the mechanism halts without selling to anybody (and hence obtains 0 revenue). ■

Ok – so we now have a useful tool. With a profit extractor, we can always obtain revenue R *if* we know that it is possible to obtain revenue R with a fixed price. But we're not done, since we don't know R. But what we have done is reduced our problem to finding a good *estimate* of the true optimal revenue R^*. Recall however that it is important that R be defined independently of the bidders we run the profit extractor on – it was only truthful assuming R was independent of their bids.

This is the idea behind the random sampling auction – we will try and estimate R^* from a random sample of the bidders, and then run the profit extractor parameterized with this estimate on the remaining bidders. First we observe that the random sampling auction is truthful:

Theorem 5 *The random sampling auction is dominant strategy truthful.*

Proof This follows because Extract(R, v) is truthful whenever it is run with a value R computed independently of the bidders it is run on. ■

Now on to the revenue:

Algorithm 1 The Random Sampling Auction

RS(v):

Randomly partition the agents by assigning each agent uniformly at random to one of two sets: S' or S''.
Calculate $R' = \text{OPT}^{\geq 2}(v_{S'})$ and $R'' = \text{OPT}^{\geq 2}(v_{S''})$.
Run Extract$(R', v_{S''})$ on S'' and Extract$(R'', v_{S'})$ on S'.

Lemma 6 *The revenue of the random sampling auction is at least* $\min(R', R'')$.

Proof It must be that either $R' \geq R''$ or $R'' \geq R'$ (or possibly both). So at least one copy of Extract succeeds. ■

So to bound the revenue of the auction, we only need to understand the quantity $\min(R', R'')$ as a function of $R := \text{OPT}^{\geq 2}(v)$.

Lets start with some simple math:

Theorem 7 *If we flip $k \geq 2$ coins, then* $\mathbb{E}[\min(\#heads, \#tails)] \geq k/4$.

Proof Let M_i be $\mathbb{E}[\min(\#\text{heads}, \#\text{tails})]$ after i coin flips. Some direct calculations show: $M_1 = 0$ and $M_2 = 1/2$.

Now define $X_i := M_i - M_{i-1}$ to be the expected change to $\min(\#\text{heads}, \#\text{tails})$ after we flip the i'th coin. Note that by linearity of expectation:

$$M_k = \sum_{i=1}^{k} X_i$$

so we are done if we can compute X_i for all i. We have two cases:

Case 1: i is even In this case, $i-1$ is odd, and we have $\#\text{heads} \neq \#\text{tails}$ after $i-1$ coin flips. Hence $X_i = 1/2$, since with probability $1/2$, the coin flip contributes to the smaller of the two quantities.

Case 2: i is odd In this case, $i-1$ is even, so it could be that $\#\text{heads} = \#\text{tails}$. Still, at the very least we can say $X_i \geq 0$. So:

$$M_k = \sum_{i=1}^{k} X_k \geq \frac{k}{2} \cdot \frac{1}{2} = \frac{k}{4}$$

which completes the proof[1]. ■

We're now ready to prove our main theorem:

Theorem 8 *Let Rev be the expected revenue of the random sampling auction. Then:*

$$Rev \geq \frac{OPT^{\geq 2}(v)}{4}.$$

Proof Recall that we have shown:

$$Rev \geq \mathbb{E}[\min(R', R'')]$$

[1] Actually, we were a little sloppy... we only showed that $M_k \geq \lfloor \frac{k}{2} \rfloor \cdot \frac{1}{2}$, which might be a little less than $k/4$. To be fully rigorous, we have to directly verify that $X_3 = 1/4$ (i.e. not 0), which makes up the difference

and we know that $\text{OPT}^{\geq 2}(v) = k \cdot p$ for some $k \geq 2$ and some price p. Of the k winners when using price p, let k' be the number in S' and k'' be the number in S''. Observe that $R' \geq k' \cdot p$ and $R'' \geq k'' \cdot p$ (they could be higher if better prices are available). Hence:

$$\begin{aligned}
\frac{Rev}{\text{OPT}^{\geq 2}(v)} &\geq \frac{\mathbb{E}[\min(R', R'')]}{k \cdot p} \\
&\geq \frac{\mathbb{E}[\min(k' \cdot p, k'' \cdot p)]}{k \cdot p} \\
&\geq \frac{\mathbb{E}[\min(k', k'')]}{k} \\
&\geq \tfrac{1}{4}
\end{aligned}$$

which completes the proof. ■