Game Theory and Algorithms*
Lecture 14: Maximizing Auctioneer Revenue, Part 2

April 19, 2011

Summary: Last class we looked at results of Myerson on finding truthful mechanisms which maximize profit when the bidders' distributions are known and independent. Today we look at the case where the bidders' distributions are unknown and we have *digital goods* where we can sell as many copies as we like. Among truthful auctions, we find one to "approximately maximize" the income no matter how the players bid.

1 Worst-Case Analysis

Consider the following scenario, which is the one-bidder case of what we will study today. A single bidder walks in to your store, and you want to offer them a truthful mechanism for buying an item. Recall that for a single player, truthful mechanisms are basically the same as selling an item at a fixed price. Last class, we assumed their valuation was drawn from some known distribution and in this case we were able to compute the optimal truthful auction (sell at price P to maximize $P \cdot \Pr[b \geq P]$). But if we have no information whatsoever about the bidder's valuation, to what extent can we still maximize our revenue? It turns out that we cannot even *approximately* maximize our revenue within a constant factor, compared to the optimal price (which would be b if we knew it in advance).

Remark: today we consider "randomized truthful auctions" which means that the algorithm can flip some coins, and then choose a truthful auction depending on the outcome of the coin flips. This is a smaller class than the the class of truthful-in-expectation auctions. It's also sometimes called "a probability distribution over truthful mechanisms" or a "universally truthful randomized mechanism." A good way to think about it is that we still have the same qualitative restrictions as a deterministic auction, but we allow the algorithm to become randomized which makes probabilistic methods become available to us.

Claim 1. *There does not exist a randomized truthful one-person auction with the following property: for some constant C, for every bid b, the expected profit is at least b/C.*

* Lecture Notes for a course given by David Pritchard at EPFL, Lausanne. Main source of today's lecture: "Lectures on Optimal Mechanism Design," J. Hartline.

Proof. Suppose for the sake of contradiction that this randomized truthful one-person auction did exist. Then in particular, if we randomize the bidder then the expected profit should still be at least $\mathsf{E}[b]/C$.

Let H be a positive integer to be fixed later. Pick the following distribution for the bid b: the possible bids are $b \in \{1, 2, \ldots, H\}$ and

$$\text{for all } i = 1, \ldots, H : \Pr[b \geq i] = \frac{1}{i}.$$

In other words, explicitly, we have $\Pr[b = i] = \frac{1}{i} - \frac{1}{i+1}$ for $i < H$ and $\Pr[b = H] = 1/H$. On the one hand, it is straightforward to calculate that the expected bid is $\mathsf{E}[b] = \sum_{i=1}^{H} \frac{1}{i}$. On the other hand, observe that every truthful (deterministic) auction does not have a very high profit: selling at price $x \in \{1, \ldots, H\}$ gives expected profit

$$x \cdot \Pr[b \geq x] = x \cdot \frac{1}{x} = 1$$

and it is not hard to see no other price can do better than this. Likewise, every probability distribution over truthful mechanisms has expected profit at most 1.

Since $1 < \mathsf{E}[b]/C = \sum_{i=1}^{H} \frac{1}{i}/C$ provided H is sufficiently large relative to C, we are done. $\square$

So, no truthful mechanism can always (for all b) extract a constant fraction of the optimal profit (b).

Exercise. Show that the above result also holds for truthful-in-expectation mechanisms.

1.1 Digital Goods/Lowering our Standards

For the rest of lecture we will consider selling *digital goods* to single-parameter agents: we can declare any arbitrary subset of players as the winners (compare this to k-item auctions, where we could only have k winners).

The proof at the end of the previous section rules out the type of mechanism which we really desire for profit maximization, if we require truthfulness. Therefore, it makes sense to "lower our standards:" rather than try to obtain the absolute maximum profit (sum of the bids), we try to achieve nearly as much profit as a good "benchmark."

Definition 2.

- For a vector x, let $x^{(i)}$ be the i^{th}-largest number in x.
- The *optimal single-price profit with hindsight* $\mathcal{S}$ of a bid vector b is

$$\mathcal{S}(b) := \max_{i=1}^{n} i \cdot b^{(i)}.$$

- The same as the previous, but requiring two or more sales, gives $\mathcal{S}^{(2)}$ defined by

$$\mathcal{S}^{(2)}(b) := \max_{i=2}^{n} i \cdot b^{(i)}.$$

Claim 1 shows that it is impossible for a truthful mechanism to extract a constant fraction of $\mathcal{S}$. Intuitively, $\mathcal{S}^{(2)}$ is a somewhat more reasonable benchmark since we can use one bidder as a reference point for another bidder, if we are comparing to sets of two or more bidders. In fact, the main result of this lecture is:

Theorem 3. *There is a truthful randomized mechanism for digital goods which obtains expected profit at least $\mathcal{S}^{(2)}(b)/4$ on all bid vectors b.*

The algorithm is "Random-Sample, Profit-Extraction" or RSPE. We thus say that RSPE is *4-competitive*, and likewise we call a mechanism *k-competitive* if it obtains profit at least $\mathcal{S}^{(2)}(b)/k$ on all bid vectors b; equivalently we call (the minimal) k its *competitive ratio*.

Example 4. Let us calculate the competitive ratio of the *second-price auction*, when there are just two bidders. We claim that the competitive ratio is 2. First, the definition of $\mathcal{S}^{(2)}(b)$ becomes very simple for $n = 2$: $\mathcal{S}^{(2)}$ is just 2 times the lowest bid. Second, the profit extracted by a second-price auction equals the lowest bid. So in fact, the competitive ratio is exactly 2 (even for any particular vector b).

We will work towards proving Theorem 3.

2 Profit Extraction

One basic tool we will use in RSPE is a truthful *profit extractor*. It takes a target profit R as input and tries to extract it from the bidders.

Procedure PROFIT-EXTRACT(R)
Initialize I to the set of all bidders
loop
    If any $i \in I$ has $b_i < R/|I|$, remove i from I
Declare the set of winners to be I
Charge each winner $R/|I|$

We use the following facts:

- PROFIT-EXTRACT is truthful. To prove this, we apply the (monotony + payment rule) characterization from last lecture.

- PROFIT-EXTRACT gives profit R when $\mathcal{S} \geq R$, and profit 0 otherwise.

Remark 5. PROFIT-EXTRACT is an important special case of the *Moulin(-Shenker) mechanism.* It is useful in more general *cost-sharing* games where we charge different players different prices, and where these prices can depend on the set of winners in more general ways. If the cost-sharing functions satisfy certain properties then it is even group-strategyproof, but we don't need this in today's lecture.

Exercise. Show that the specific mechanism PROFIT-EXTRACT is group-strategyproof: if a subset I of players deviate from truthfulness so that one of their utilities increase, show that another of the players' utilities decrease.

3 The RSPE Algorithm

If we need to offer a truthful mechanism against unknown bidders, the general idea which gets us the furthest is to partition the bidders, and use one group as a reference point for the others. This idea is used in many papers and it is for this reason that we are highlighting the specific example of RSPE, which is defined as follows.

Procedure RANDOM-SAMPLE, PROFIT-EXTRACTION
1: For each bidder, flip a coin: if it lands heads put that bidder in set I', otherwise put that bidder in set I''
2: Let b' be the bids of I' and b'' be the bids of I''
3: Compute the optimal single-price profits with hindsight: $S' := \mathcal{S}(b'), S'' := \mathcal{S}(b'')$
4: Run PROFIT-EXTRACT(S') on I'' and PROFIT-EXTRACT(S'') on I'

For any choice of coin flips, the algorithm is truthful to each player (although possibly non-group-strategyproof if players from I' collude with players from I''). In order to prove that this algorithm is 4-competitive, we need the following lemma.

Lemma 6. *If we flip $j \geq 2$ coins, then* $\mathsf{E}[\min\{\#heads, \#tails\}] \geq j/4$.

We will prove it later, but here is how it is used.

Theorem 7. *RSPE is 4-competitive (against $\mathcal{S}^{(2)}$).*

Proof. Due to the properties of the profit extractor, the profit is the minimum of S' and S''. It remains to show that $\mathsf{E}[\min\{S', S''\}] \geq \mathcal{S}^{(2)}/4$.

From the definition of $\mathcal{S}^{(2)}$, let j be the index achieving the optimal revenue, i.e. choose $j \geq 2$ such that $\mathcal{S}^{(2)} = j \cdot b^{(j)}$; and let J be the set of winners, $J = \{i \mid b_i \geq b^{(j)}\}$. In order to use the lemma, let k be $|J \cap I'|$, so $|J \cap I''| = j - k$. Note that the distribution of $(k, j-k)$ is the same as that obtained by flipping j coins and counting (#heads, #tails).

For any fixed k, we have that $S' \geq k \cdot b^{(j)}$ since one possible fixed sale price for I' is to offer the price $b^{(j)}$, which everyone in I' would accept. Similarly $S'' \geq (j-k) \cdot b^{(j)}$ and $\min\{S', S''\} \geq \min\{k, j-k\} b^{(j)}$. So taking the expectation and factoring out the price, the expected profit is

$$\mathsf{E}[\min\{S', S''\}] \geq \mathsf{E}[\min\{k, j-k\} b^{(j)}] = b^{(j)} \mathsf{E}[\min\{k, j-k\}] \geq b^{(j)} \cdot j/4$$

where the last inequality uses the Lemma. $\square$

Finally, we give the proof of Lemma 6.

Proof. We will use double induction on j. To this end, we must manually verify the lemma for $j = 2$ and $j = 3$. For $j = 2$, the expected value is $\frac{1}{4}0 + \frac{1}{2}1 + \frac{1}{4}0 = \frac{1}{2} \geq j/4$ as needed. For $j = 3$, the expected value is $\frac{1}{8}0 + \frac{3}{8}1 + \frac{3}{8}1 + \frac{1}{8}0 = \frac{3}{4} \geq j/4$ as needed.

Next, we show that conditional on any specific sequence of the first $j - 2$ flips, the expected increase in $\min\{\#\text{heads}, \#\text{tails}\}$ from the next two flips is at least $2/4$. Removing the conditioning and using linearity of expectation, this will complete the proof.

Case 1 : In the first $j - 2$ flips, there are an equal number of heads and tails. Then the next two flips increase $\min\{\#\text{heads}, \#\text{tails}\}$ by 0 with probability $1/2$, and by 1 with probability $1/2$, so the total expected increase is $1/2$ as needed.

Case 2 : In the first $j - 2$ flips, the number of heads and tails is unequal. Then even just one more flip causes $\min\{\#\text{heads}, \#\text{tails}\}$ to increase by $1/2$ in expectation; the second flip does not decrease the value. □

3.1 Related Facts

- The competitive ratio of RSPE is exactly 4, even for just two players. This can be verified by running it on the bid vector $(1, 2)$.

- The best known truthful auction has competitive ratio 3.25. It is based on randomly splitting the bidders into three groups, and then running the optimal three-player auction on these "group bids."

- No *symmetric deterministic* truthful mechanism for n players has competitive ratio less than n; asymmetric deterministic ones have been obtained recently via derandomization techniques. See Theorem 3.15 in Hartline.

- No randomized truthful mechanism for 2 players has competitive ratio less than 2. (So even $\mathcal{S}^{(2)}$ is somewhat too optimistic, as a worst-case benchmark; and the second-price auction for 2 players has the optimal competitive ratio. The proof is similar to that of Claim 1, see Lemma 3.22 in Hartline.)

- More generally, there are constants $c_2 = 2, c_3 = 13/6, \ldots$ tending to 2.42 so that no randomized truthful mechanism for n players has competitive ratio less than c_n.