CS280 Project Report: Competitive Auctions

based on *A.V. Goldberg et al./Games and Economic Behavior 55 (2006) 242–269*

SungHeon Jeong, Kenneth Zhao, Gnana Heemmanshuu Dasari, Jun Li, Feilong Hou

1 Motivation

The authors discuss the design and analysis of profit-maximising auctions that perform well in *worst case scenarios* in unknown markets. There are a number of advantages to this approach:

- Addresses issues arising from incomplete knowledge of distribution of bidder valuations
- Provides provable performance guarantees in total uncertainty, thus providing a baseline for performance which could be improved with additional information
- Assists in representing the quantitative value of "information" in such a market, etc.

The authors focus on the analysis of goods in *unlimited supply*, where the auctioneer has negligible difficulty in duplication/distribution of the goods. This focus helps visualise the case of digital markets and goods such as pay-per-view TV, streaming services, digital media such as ebooks and music, software licenses, mobile apps and much more.

2 Definitions, notations, prerequisites

2.1 Auctions

The authors consider *single round, sealed bid, truthful* auctions for items available in unlimited supply. Such auctions consist of

1. n number of bidders
2. A bid vector $\mathbf{b}$ with b_i being the bidding of bidder i
3. Allocation and price vectors $\mathbf{x}, \mathbf{p}$ where x_i, p_i represent bidder i's receipt of the item and the price he pays for it, respectively
4. Profit of the auctioneer, given by $\mathcal{A}(\mathbf{b}) = \sum_i p_i$

Further, natural assumptions such as each bidder bidding to maximise their utility $u_i = v_i x_i - pi$, where v_i is the private valuation of bidder i, non-collusion and full knowledge of auction mechanism among bidders, indistinguishability of bidders from auctioneer's perspective, no positive transfers ($0 \leq p_i \leq b_i$) and voluntary participation ($p_i = 0$ if $x_i = 0\ \forall i$) are made.

From **Myerson's revelation principle**, analysis of *truthful* auctions will be enough (utility is maximised for a bidder by bidding their valuation). From **Myerson's Lemma**, it is known that a truthful auction is equivalent to a *bid-independent auction* $\mathcal{A}_f(\mathbf{b})$, where f is a function from masked bid vector $\mathbf{b}_{-i}$ to price t_i such that bidder i wins if $t_i \leq b_i$.

2.2 Metrics for evaluation

The performance of these truthful auctions to be designed are compared with two types of *omniscient* auctions: single priced and multi priced. *Omniscient* auctions are purely hypothetical auctions used for benchmarking purposes, where the auctioneer knows everything about bidders.

Single priced omniscient auctions sell at a common price that returns maximum profit to the auctioneer (making its profit $\mathcal{F}(\mathbf{b}) = \max_{1 \leq i \leq n} iz_i$), whereas multiple priced ones sell to each bidder at their bid value with profit $\mathcal{T}(\mathbf{b}) = \sum_{1 \leq i \leq n} b_i$

The Competitive Analysis Framework

Worst case analysis is the only way to give provable performance guarantees of auction mechanisms when there is total uncertainty. Approaches like Bayesian analysis use expected profit which hides the possibility of the auction performing poorly on specific inputs. As there is total uncertainty, assumptions about bidder valuations used in Bayesian analysis can lead to errors.

The profit of the designed auction is compared with that of the optimal omniscient auctions for each bid set (including the *worst case*). The auction is **competitive** if it achieves a profit that is a fraction of optimal on every input. The goal is to minimize the *competitive ratio*, given by

$$\frac{\text{Profit of best possible optimal auction that can be competed with}}{\text{Profit of designed auction}}$$

The auction that minimises this value is a *competitive auction.*

It can be shown that any truthful auction cannot really be competitive with the optimal single price auction $\mathcal{F}$. Thus, competitiveness is defined against a weaker form of the single price auction $\mathcal{F}^{(m)}$, where the auctioneer wants to sell at least m units of the good, i.e., $\mathcal{F}^{(m)}(\mathbf{b}) = \max_{m \leq i \leq n} iz_i$. Then, an auction $\mathcal{A}$ is $\beta-$competitive against $\mathcal{F}^{(m)}$ if

$$\mathbf{E}[\mathcal{A}(\mathbf{b})] \geq \frac{\mathcal{F}^{(m)}(\mathbf{b})}{\beta}$$

β is the *competitive ratio* of auction $\mathcal{A}$. Just stating $\beta-$competitive implies $\beta-$competitiveness against $\mathcal{F}^{(2)}$.

3 Deterministic auctions are not competitive

This section defines a symmetric auction, and utilizes this concept to show that no symmetric deterministic auction can be competitive. An auction $\mathcal{A}$ is symmetric if given a permutation of the bid vector, the price vector and allocation vector is permuted the same way. In other words, the outcome of the auction does not depend on the order of the bids.

Theorem 4.1 For any symmetric deterministic auction $\mathcal{A}_f$ defined by bid-independent function f, it is not competitive: for any $1 \leq m \leq n$ there exists a bid vector $\mathbf{b}$ of length n such that the profit of $\mathcal{A}_f$ on $\mathbf{b}$ is at most $\mathcal{F}^{(m)}(\mathbf{b})\frac{m}{n}$.

This means a deterministic mechanism's performance at worst is $\frac{m}{n}$ of a single price omniscient auction for the same bid vector. For a randomized mechanism, these scenarios can still occur, as it only randomizes over deterministic auctions. However, through careful control of the internal randomness, they could be designed in ways to minimize the probability of these worst case outcomes.

4 A lower bound on the competitive ratio

This section focuses on proving that the best possible performance of any truthful auction, even a randomized one, compared to $\mathcal{F}^{(2)}$ has a competitive ratio of at least 2.42. Specifically, for any truthful auction $\mathcal{A}$, there exists an input bid vector $\mathbf{b}$ such that

$$\mathbf{E}[\mathcal{A}(\mathbf{b})] \leq \frac{\mathcal{F}^{(2)}(\mathbf{b})}{2.42}.$$

The proof of the lower bound comes from analysis of the behavior of $\mathcal{A}$ over randomly chosen bid vectors. The outcome of the auction is a random variable depending on $\mathcal{A}$ and $\mathbf{b}$. The key result comes from showing that $\mathbf{E_b}[\mathbf{E}_{\mathcal{A}}[\mathcal{A}(\mathbf{b})]] \leq \mathbf{E_b}[\mathcal{F}^{(2)}(\mathbf{b})]/2.42$.

Lemma 5.3. Under the specified bid parameters defined in the section, we have

$$\frac{\mathbf{E}[\mathcal{F}^{(2)}(\mathbf{b}^{(n)})]}{\mathbf{E}[\mathcal{A}(\mathbf{b}^{(n)})]} = 1 - \sum_{i=2}^{n}(\frac{-1}{n})^{i-1}\frac{i}{i-1}\binom{n-1}{i-1}.$$

For $n = 2$, this gives a lower bound of 2, which matches the 1-unit Vickrey auction; for $n = 3$, the lower bound is $\frac{13}{6}$. Taking the limit as n approaches infinity, we get the lower bound of the competitive ratio for a general n, giving us the result:

Theorem 5.5. The competitive ratio of any truthful randomized auction is at least 2.42.

5 Competitive auctions via random sampling

This section introduces two auction design techniques based on random sampling. Both techniques involve partitioning the set of input bids into two groups by randomly assigning each bid to a partition with equal probability. One partition is used for market analysis, and the resulting insights are applied to a sub-auction executed on the other partition. The process is symmetric, as each partition alternates between analysis and execution roles.

1. **Random Sampling Optimal Price (*RSOP*):**
   The *RSOP* sub-auction employs a single "take-it-or-leave-it" offer made uniformly to all bidders.
2. **Random Sampling Profit Extraction (*RSPE*):**
   The *RSPE* sub-auction leverages a specialized version of the cost-sharing mechanism developed by [1].

5.1 *RSOP*, Random sampling optimal price auction

Definition: The *RSOP* auction follows these steps:

- **Random Sampling**: Partition the input bids randomly into two groups, b' and b''.
  - Assign each bidder to a group by flipping a fair coin.
- **Optimal Price Calculation**:
  - For b', compute $opt(b')$, the optimal sale price, denoted as p'.
  - For b'', compute $opt(b'')$, the optimal sale price, denoted as p''.
- **Auction Execution**:
  - Apply p' to group $b'' \rightarrow$ accept only bids in b'' that are at least p'.
  - Apply p'' to group $b' \rightarrow$ accept only bids in b' that are at least p''.

Features and Performance Guarantees: *RSOP* is a "dual-priced auction," applying different prices to the two groups. If a single price is required, one group can be ignored, reducing the expected profit by half. *RSOP* is truthful, meaning bidders are incentivized to bid their true valuation. Regarding performance, **Theorem 6.3** states that *RSOP* achieves a constant fraction of the profit of the optimal single-price auction ($F(2)$). Additionally, **Theorem 6.4** demonstrates that when bids are within a bounded range ($[1, h]$), *RSOP* achieves near-optimal performance. As the number of bidders increases, *RSOP*'s performance converges to that of the optimal single-price auction (F).

5.2 $RSPE$, Random sampling profit extraction auction

The $RSPE$ auction is a truthful and simple mechanism that achieves a competitive ratio of 4, extending a cost-sharing technique by [1] to focus on profit extraction.

Definition of $ProfitExtract_R$:

- Finds the largest k such that the top k bids are at least R/k.
- Charges these k bidders R/k and rejects the others.
- Ensures truthfulness and achieves revenue R if $R \leq F(b)$, otherwise earns none.

$RSPE$ randomly partitions bids into two groups, b' and b'', computes the optimal single-price profits for each group (F' and F''), and runs $\text{ProfitExtract}_{F''}$ on b' and $\text{ProfitExtract}_{F'}$ on b''. $RSPE$ is truthful, as it relies on truthful mechanisms for each partition, and it is 4-competitive, with this bound being tight. When $F' \neq F''$, $RSPE$ achieves $\min(F', F'')$, ensuring at least $F^{(2)}/4$. However, $RSPE$ may lose at least half the profit when one partition yields no revenue. To address this, a parameter $\gamma < 1$ can be introduced to adjust the extracted profit ($\text{ProfitExtract}(\gamma F')$ and $\text{ProfitExtract}(\gamma F'')$), improving performance with appropriate knowledge of bidder distributions.

6 Limited Supply

In limited supply auctions, the auctioneer has a fixed number of items, k, fewer than the number of bidders. The goal is to maximize profit without requiring all k items to be sold. The problem can be reduced to the unlimited supply model by sorting bids, retaining the top k, and applying an unlimited supply auction algorithm.

A common approach is the A_k mechanism: Firstly, simulate a k-Vickrey auction to retain the top k bids and set the $k+1$-th highest bid as the price. Secondly, treat non-winning bids as zero and apply an unlimited supply algorithm A to the top k bids. Finally, combine the results of the k-Vickrey auction and A, using the higher price and intersecting their allocation rules.

This method ensures truthfulness and competitiveness. For example, the Random Sampling Optimal Price (RSOP) algorithm remains constant-competitive, and the Random Sampling Profit Extraction (RSPE) algorithm achieves a 4-competitive ratio. These mechanisms often outperform k-Vickrey in terms of profit while being simple extensions of the unlimited supply framework.

7 Beyond $\mathcal{F}$

The optimal single-price auction $\mathcal{F}$ sets a uniform price to maximize profit, assuming full knowledge of bidders' valuations. While multi-price auctions theoretically allow for higher revenue by setting different prices for different bidders, their performance is constrained by strategy-proofness and practical benchmarks.

Existence of Stronger Benchmarks: The optimal multi-price omniscient auction T charges each bidder their exact valuation ($T(b) = \sum b_i$), making it exponentially stronger than $\mathcal{F}$ in some cases. For instance, with bid vector $b = \{n, \frac{n}{2}, \ldots, 1\}$:

$$F(b) = n, \quad T(b) = n \cdot H_n \approx n \ln n.$$

However, T is impractical for real-world use.

Performance of Monotone Auctions: Monotone auctions ensure that higher bidders are more likely to win and pay no less than lower bidders. Yet, no monotone auction can systematically exceed $\mathcal{F}$'s revenue:

$$E[R_A(b)] \leq F(b),$$

where A is any monotone randomized auction. This is because the expected revenue of A is bounded by the revenue of $\mathcal{F}$ under monotonicity constraints.

Hard-Coded Auctions: Specific bid vectors b^* allow for tailored auctions achieving $T(b^*)$. However, these mechanisms generalize poorly, supporting $\mathcal{F}$ as the most practical and robust benchmark.

References

[1] H. Moulin and S. Shenker, “Strategyproof sharing of submodular costs: budget balance versus efficiency,” *Economic Theory*, vol. 18, pp. 511–533, 2001.