CS 6840 Algorithmic Game Theory September 23, 2024

Lecture 12: Single Item Auction

Instructor: Eva Tardos *Scribe: Oluwasola Ogundare*

1 Recap

We recall the structure of a single-item auction, in which n players bid with values $v_1, \ldots, v_n$, where v_i represents the bid of player i for the item. The Social Welfare (SW) of awarding the item to bidder i at price p is defined as:

$$SW = (v_i - p) + p + 0 + \ldots + 0 = v_i$$

Here, $v_i - p$ represents the utility for the winning player, p denotes the payment made to the auctioneer, and the remaining $n - 1$ zeros represent the zero utility for the other participants.

This formulation for Social Welfare is considered reasonable, as it aligns with the principle of maximizing the total economic value exchanged.

Price of Anarchy (PoA)

Our objective is to analyze the Price of Anarchy (PoA) for a single-item auction without explicitly computing the equilibrium of the game. Let $u_i(s_i^*, s_{-i})$ represent the utility of player i, where s^* denotes a special strategy, and s refers to any other strategy vector.

The term **special strategy** can be interpreted as a recommended action. Following such a recommendation ensures that a player receives a guaranteed payoff, regardless of the strategies employed by others. This concept will prove useful in subsequent analysis.

2 First-Price Auction

A *First-Price Auction* is an auction in which the winning player pays their bid, $p = b_i$, while all other participants pay nothing.

A natural question arises: *What should a player bid in a First-Price Auction?* A potential strategy is to bid $v_i/2$, where v_i represents the player's valuation of the item. Bidding the full value, v_i, yields an expected utility of zero, whereas bidding zero guarantees a loss. Consequently, players tend to *shade their bids*, reducing them by some factor.

We now present a formal claim regarding the utility of this strategy.

Theorem 1. *For any strategy or bid vector s, the following holds:*

$$u_i\left(\frac{v_i}{2}, s_{-i}\right) \geq \frac{v_i}{2} - \text{PRICE}(s)$$

where $\text{PRICE}(s)$ represents the highest bid in strategy vector s.

Proof. Let $\text{PRICE}(s) = \max_j s_j$. We now consider two cases:

Case 1: The bid $\frac{v_i}{2}$ wins. Since the winning bid was half of player i's valuation, we have that $u_i\left(\frac{v_i}{2}, s_{-i}\right) = v_i - \frac{v_i}{2} = \frac{v_i}{2}$. As $\textsc{price}(s)$ is non-negative, the theorem holds.

Case 2: The bid $\frac{v_i}{2}$ loses. In this case, player i pays nothing and gains no utility, so $u_i\left(\frac{v_i}{2}, s_{-i}\right) = 0$. Additionally, $\textsc{price}(s) > \frac{v_i}{2}$, since a player $j \neq i$ outbid i. Thus, $\frac{v_i}{2} - \textsc{price}(s) < 0$, and the theorem holds. ■

Furthermore, we observe that the utility is always non-negative, as either the player wins and gains $\frac{v_i}{2}$, or they lose and pay nothing.

Next, we use these results to establish that the Price of Anarchy (PoA) in this setting is bounded above by a factor of 2.

Theorem 2. *The following inequality holds:*

$$\sum_i u_i\left(\frac{v_i}{2}, s_{-i}\right) \geq \frac{1}{2}\max_j v_j - \textsc{price}(s)$$

Proof. By applying the results of the previous claim to all players, and using the highest bid for the player with the maximum valuation, we can derive:

$$\begin{aligned}
&\sum_i u_i(s) \geq \sum_i u_i\left(\frac{v_i}{2}, s_{-i}\right) \geq \frac{1}{2}\max_j v_j - \textsc{price}(s), \\
&SW(s) \geq \frac{1}{2}\max_j v_j, \\
&SW(s) \geq \frac{1}{2}\text{OPT}, \\
&\frac{SW}{SW \text{ at OPT}} \leq 2.
\end{aligned} \tag{1}$$

■

This framework can be generalized to analyze the PoA across auctions more broadly. Let OPT denote the optimal Social Welfare, and let $\textsc{Rev}(s)$ represent the revenue obtained by the auctioneer under strategy vector s. Furthermore, let s_i^* represent the special strategy. The generalized inequality can be expressed as:

$$\sum_i u_i(s_i^*, s_{-i}) \geq \lambda\text{OPT} - \textsc{Rev}(s), \quad SW(s) \geq \lambda\text{OPT} \quad \forall \text{ Nash equilibria } s, \text{ PoA} \leq \lambda. \tag{2}$$

If we establish that the first inequality holds for any auction type, it follows that this applies across all strategy vectors s and for the special strategy s^* chosen for each player. This demonstrates that, regardless of the specific auction format or strategy profile, the aggregate utility obtained by all players, when following their respective special strategies, is always at least as large as $\lambda \cdot \text{OPT}$ minus the auctioneer's revenue. Given this result, the second and third lines naturally follow: since this holds for any Nash equilibrium, the social welfare $SW(s)$ of any equilibrium strategy s is always at least $\lambda \cdot \text{OPT}$. Consequently, the PoA is bounded above by λ.

In particular, this result holds for pure Nash equilibria, and we will now extend it to the case of mixed Nash equilibria by introducing expected utilities over strategies $s \sim \sigma$.

$$\begin{aligned}&\sum_i \mathbb{E}_{s\sim\sigma}[u_i(s)] \geq \sum_i \mathbb{E}_{s\sim\sigma}[u_i\left(\frac{v_i}{2}, s_{-i}\right)] \geq \mathbb{E}_{s\sim\sigma}[\lambda \text{OPT}] - \mathbb{E}_{s\sim\sigma}[\text{PRICE}(s)],\\&\mathbb{E}_{s\sim\sigma}[u_i(s)] \geq \lambda\text{OPT} - \mathbb{E}_{s\sim\sigma}[\text{PRICE}(s)].\end{aligned} \tag{3}$$

This extension to mixed Nash equilibria is established by first considering σ as a Nash equilibrium strategy. The first inequality holds because, by the definition of Nash equilibrium, no player can unilaterally improve their utility in expectation by deviating from σ. The second inequality follows as it applies to all strategy vectors s, and hence holds in expectation as well.

3 No-Regret Learning in First-Price Auctions

We now analyze first-price auctions under a no-regret learning framework. In this setting, each of the n players iteratively submits bids over time, and these bids can be updated across rounds.

The central idea in this analysis is to sum the utilities obtained by all players over time and take advantage of the fact that this sum is at least as large as the utility obtained by following a fixed strategy in hindsight, i.e.,

$$\sum_t u_i(s^t) \geq \sum_t u_i\left(\frac{v_i}{2}, s^t_{-i}\right) - \text{REG}.$$

Since Nash equilibrium requires a finite strategy space, we address the issue of infinite bid spaces by discretizing the bids into increments of ϵv_i, resulting in $\frac{1}{\epsilon}$ possible bids.

Given this, we obtain a regret bound of $\text{REG} = \sqrt{2T \ln \frac{1}{\epsilon}}$, as the regret bound in learning follows $\sqrt{2T \ln K}$, where K represents the number of strategies or experts (as we saw in Lecture 4 and 5). Formally, we can state: Formally, we can state:

$$\begin{aligned}&\sum_t \sum_i u_i(s^t) \geq \sum_t \sum_i u_i\left(\frac{v_i}{2}, s^t_{-i}\right) - n\cdot\text{REG} \geq \sum_t \frac{1}{2}\max_j v_j - \sum_t \text{PRICE}(s^t) - n\cdot\text{REG},\\&\sum_t SW(s^t) \geq \frac{1}{2} T \max_j v_j - n\cdot\text{REG}.\end{aligned} \tag{4}$$

Here, the term $n \cdot \text{REG}$ arises because each player i has an associated regret.

Note: The factor of 2 in the PoA for pure Nash equilibria may not be optimal, and mixed strategies can yield superior results. This is left as an exercise.

4 All-Pay Auctions

We now shift our focus to the All-Pay Auction, a type of auction in which every player pays their bid, regardless of whether they win the item. Recall from equation (2) that the utility bounds can be generalized to hold in various auction formats, and we aim to show that this same inequality applies in the context of an All-Pay Auction.

$$\sum_i u_i(s_i^*, s_{-i}) \geq \lambda\text{OPT} - \text{REV}(s).$$

We aim to demonstrate that this inequality also holds in the context of an All-Pay Auction.

First, let j represent any player whose bid is not the maximum. We can define the special strategy for player j as $s_j^* = 0$. Consider a trivial scenario in which player i bids 0. In this case, $u_i(0, s_{-i}) \geq 0$, as the player pays nothing, and in the best case, they might win the item if all others also bid 0.

Next, let j represent the player with the highest valuation. This player can adopt a mixed strategy, bidding according to a uniform distribution $s_j^* \sim \text{Uni}[0, v_j]$. The expected utility of this strategy can be computed as follows:

$$\begin{aligned} &\mathbb{E}_{s_j^*}[u_j(s_j^*, s_{-i})] = \frac{1}{v_j}\int_0^{v_j} u_j(x, s_{-j})\, dx, \\ &u_j(x, s_{-i}) = v_j \cdot \frac{v_j - p}{v_j}, \text{ where } p = \max_{i \neq j} s_i, \\ &\mathbb{E}[u_j(s_j^*, s_{-i})] = v_j - p. \end{aligned} \tag{5}$$

Note that this expression represents the expected utility before accounting for payments. After subtracting the expected payment, we obtain an expected payoff of $\frac{v_j}{2} - p$. Therefore, we conclude that:

$$\mathbb{E}\left[\sum_i u_i(s_i^*, s_{-i})\right] \geq \frac{1}{2}\max_j v_j - \textsc{Rev}(s) \implies SW(s) \geq \frac{1}{2}\text{OPT}.$$

Thus, the generalized Price of Anarchy bounds extend to pure Nash, mixed Nash, and no-regret learning settings.