The Almighty Adversary in Combinatorial Posted Price Auctions

Cui, Ziyong and D'Amico-Wong, Luca
ziyongcui@college.harvard.edu ldamicowong@college.harvard.edu

November 18, 2023

Abstract

The performance of sequential posted price mechanisms in combinatorial settings, where finding welfare-optimal allocations is NP-hard and optimal mechanisms can easily become extremely complex, has been well-studied, typically operating within a Bayesian framework. However, existing results require that the buyers arrive in a fixed ordering, not allowing for the ordering of the buyers to change according to their realized valuations. As incentives exist for buyers to change their positions given their valuations, settings in which such an adversarial order might arise are natural to consider.

To this end, we consider the performance of sequential posted price mechanisms in a Bayesian setting, with the addition of an almighty adversary. The almighty adversary possesses full knowledge of the buyers' realized values and the posted prices, and for each possible realization of the buyers' values, is allowed to order the buyers such as to minimize the welfare or revenue attained by the mechanism. Under this absolute worst-case scenario, we show positive results when buyers have additive valuations, recovering the best possible $\frac{1}{2}$*-guarantees of revenue and welfare and provide constructive prices to do so. Additionally, we bound the adversarial price of anarchy, the ratio of the worst-case performance given an almighty adversary when compared with an adversary which much select a fixed ordering.*

Finally, we consider possible extensions when faced with more general subadditive valuations and propose additional directions for future research. We conjecture that the worst-case guarantees under the almighty adversary are identical to those in the fixed order case; namely, we believe that there exist posted prices that deliver $\frac{1}{2}$*-guarantees of welfare even under XOS valuations with an almighty adversary.*

I. Introduction

Single-item auctions are typically poorly equipped for settings in which one wants to sell multiple items, in particular when buyers have complex preferences over different bundles of these items due to their relative substitutability/complementarity. As a result, combinatorial auctions have attracted significant attention from both researchers and practitioners, perhaps most notably in the FCC's spectrum auctions, which generated billions in revenue for the American government [17].

While combinatorial auctions allow for much increased expressivity, they face several additional challenges that do not arise in the single-item setting. We consider the natural question of designing an incentive compatible (IC) combinatorial auction that maximizes social welfare. It is well-known that even finding the welfare-maximizing allocation (equivalent to solving the winner determination problem) is NP-hard when bidders have submodular valuation functions [15]. While poly-time greedy algorithms that provide constant approximations to the optimal welfare exist, mechanisms based on these greedy allocations do not lend themselves nicely to incentive compatibility.

With this in mind, it is natural to look for a simple and transparent IC mechanisms that respects computational constraints while also providing good welfare guarantees. One natural class of mechanisms that has been well studied is *posted price mechanisms*. In a posted price mechanism, prices are set beforehand by the auctioneer, with buyers arriving one-by-one and selecting their most preferred bundle. Posted price mechanisms are appealing both for their simplicity and their economic properties, being trivially DSIC and weakly group strategyproof.

Operating in the Bayesian setting, Feldman et al. [9] showed positive results for XOS valuations, providing constructive posted prices that guarantee $\frac{1}{2}$ of the expected optimal welfare under any fixed ordering of the buyers, possibly adversarially selected [1]. Despite these strong guarantees, they did not, however, provide any guarantees when faced with variable orderings that may depend on the valuations of the buyers.

To better understand when and how these variable orderings might arise, note that the items that are available to a given bidder depend on the valuations of the bidders that came before her. Thus, while these posted-price mechanisms are DSIC once an ordering is set, there still exist incentives for bidders to try and modify their position in the order if possible. In a different vein, one could imagine a malicious third party that is interested in sabotaging an auction, either wanting to minimize the welfare provided by the auction or wanting to allocate items to specific bidders who may not receive these items in the optimal allocation.

Motivated by these concerns over potential tampering with the orderings, we seek to extend upon the work of Feldman et al. and establish robustness guarantees for posted-price mechanisms in the worst possible case. Specifically, we allow for an *almighty adversary* that can set the ordering of buyers after observing both the posted prices and the crystallized valuation functions for each bidder. We prove positive results for the additive case (both in the single-unit and multi-unit setting), conjecture a price of anarchy bound on the separation between the fixed adversarial ordering and the almighty adversary case, and propose further directions of research.

i. The Model

Formally, we consider the following setting:

- We have a set of items $M = \{1, 2, \cdots, m\}$ and a set $N = \{1, 2, \cdots, n\}$ of buyers.
- Each buyer has a valuation function $v_i(\cdot) : 2^M \to \mathbb{R}^{\geq 0}$, mapping bundles to their value for that bundle.
- We focus on the Bayesian setting, where each bidder's valuation function is drawn independently from some known distribution $\mathcal{F}_i$.
- We restrict our attention primarily to subadditive valuation functions which do not allow for bidders to have complementary valuations for certain goods. Positive results in settings with high degrees of complementarity are rare [14].

In particular, we consider the following classes of valuation functions, following the standard hierarchy of subadditive functions.

- **Additive:** A buyer has a value for each item, and their value for a bundle is given by $v_i(S) = \sum_{j \in S} v_i(j)$.

[1] In fact, these results hold even for a stronger online adversary that can modify the ordering of the bidders based on the previous results of the auction and observed valuations in an online fashion.

- **Submodular:** The marginal value of an item is decreasing in bundle "size." For any $S \subseteq T \subseteq 2^M$ and item $j \notin S \cup T$,
$$v_i(S \cup \{j\}) - v_i(S) \geq v_i(T \cup \{j\}) - v_i(j).$$
- **XOS:** There exist additive functions $A_1, \cdots, A_k$ such that for all $S \subseteq 2^M$, $v(S) = \max A_i(S)$.
- **Subadditive:** For any $S, T \subseteq 2^M$, $v_i(S) + v_i(T) \geq v_i(S \cup T)$.

It can be verified that additive $\subset$ submodular $\subset$ XOS $\subset$ subadditive.

ii. Sequential Posted Price Mechanisms

Sequential posted price mechanisms proceed in the following steps:

1. The designer determines prices $p_{i,j}$ for each item $j \in \{1, 2, \cdots, m\}$ and each agent i. The price to agent i for a bundle of items $S \subseteq 2^M$ is defined to be $\sum_{j \in S} p_{i,j}$. If $p_{i,j} = p_j$ for all i – that is, there is no price discrimination – then the prices are said to be anonymous.

2. The N buyers arrive sequentially. Upon arriving, each buyer purchases their most preferred bundle (in terms of standard quasilinear utility), with ties broken arbitrarily.

3. Let $\mathbf{X} = (X_1, \ldots, X_n)$ denote the final allocation of goods after the conclusion of the auction. The final social welfare is defined as $\text{SW}(\mathbf{X}) = \sum_{i=1}^{n} v_i(X_i)$. The final revenue is given by $\text{Rev}(\mathbf{X}) = \sum_{i=1}^{n} \sum_{j \in X_i} p_{i,j}$.

iii. The Almighty Adversary

We focus on posted price mechanisms under an almighty adversary who can set the ordering of bidders after witnessing both the set prices and the crystallized valuation functions of all the bidders. Specifically, the posted price mechanism under an almighty adversary works as follows:

1. The posted prices are set by the mechanism designer, and each bidder's valuation function is drawn from their corresponding distribution $\mathcal{F}_i$.

2. The almighty adversary observes the set prices and crystallized valuation functions v_i for each buyer $i \in \{1, 2, \cdots, n\}$. The almighty adversary determines an ordering $\pi = \{\pi_1, \cdots, \pi_n\}$ for the N buyers with the goal of minimizing social welfare or revenue.

3. The designer runs the auction and determines the resulting social welfare and revenue.

Crucially, the ordering π set by the almighty adversary can change depending on the realized valuation functions. Thus, for given distributions $\mathcal{F}_i$ over potential valuation functions for the bidders, the designer no longer faces a fixed ordering over bidders but rather an ordering that may depend on the realizations of these random distributions.

iv. Related Work

Posted price mechanisms are widely used in large markets where the demand of buyers are well-understood. Following their introduction to department stores in the 1840s, posted prices have become the standard in retail markets [7]. In the case where the buyers' valuations are public knowledge and satisfy gross substitutes conditions, it has been shown that there exists anonymous prices that efficiently clear the market [12]. While it is less clear how to best price items in the combinatorial case, a significant amount of recent literature has attempted to tackle this problem.

As previously mentioned, Feldman et al. (2014) give constant approximation guarantees on the welfare under a posted price mechanism, assuming bidders have XOS valuations. As their result provides much of the foundation for our work, we state it in full below.

Theorem 1. *(Feldman et al., 2015): Given a black-box access to an algorithm $\mathcal{A}$ for XOS valuations, sample access to XOS distributions $\mathcal{F}$, an XOS query for the valuations in the support of $\mathcal{F}$, and a demand oracle for the valuations, there exists a posted price mechanism that, for every XOS valuation profile v, and every $\epsilon > 0$, returns an outcome that gives expected social welfare of at least $\frac{1}{2}\mathbb{E}_{v\sim\mathcal{F}}[\mathcal{A}(v)] - \epsilon$ and runs in time $POLY(n, m, 1/\epsilon)$ [9].*

Assuming only that bidder valuations are monotonic, Correa et al. [5] give posted prices that yield $\frac{1}{m+1}$ of the optimal welfare and show that this approximation is tight. Cai and Zhao [2] consider a similar setting, but they do so in the context of revenue maximization rather than welfare maximization. Specifically, they show the existence of an anonymous posted price mechanism (with entry fees) that guarantees a constant fraction of the optimal BIC revenue when bidders have XOS valuations and at least $\Omega(\frac{1}{\log m})$ of the optimal BIC revenue under subadditive valuations, leveraging duality and flow arguments similar to those presented by Babaioff et al. [1]. Unlike Cai and Zhao, we do not allow for the possibility of entry fees.

All three of the above results hold for any fixed ordering, possibly adversarial. However, they do not hold in the stronger setting of an almighty adversary. A number of papers have also considered the ordering in which the bidders arrive, either under weaker assumptions or stronger adversaries.

Ehsani et al. [6] offer stronger guarantees when buyers arrive in a uniformly random order and are assumed to be unit-demand. Closely related to the original prophet secretary problem [8], they show the existence of dynamic prices that guarantee at least $(1 - 1/e)$ of the optimal social welfare. In a recent[2] paper, Correa and Cristi resolve a longstanding open question, showing the existence of an online algorithm that provides a constant approximation to the optimal welfare under subadditive valuations [4]. While they rely on dynamic prices, they show that their results can be turned into an IC mechanism that holds *even in the case of an almighty adversary.*

As far as we know, our work is the first to consider the performance of sequential posted price mechanisms under an almighty adversary.

II. The Additive Setting

We begin by focusing on the simplest setting, in which bidders have additive valuations. We prove positive results for optimal welfare and revenue approximations, leveraging existing results from prophet inequalities. Additionally, we demonstrate separation between the fixed adversarial ordering setting and the almighty adversary, bounding the Bayesian price of anarchy when faced with an almighty adversary.

[2] In fact, the paper was made available as a preprint just this week.

i. Extending Welfare and Revenue Guarantees

We first look at the question of whether Feldman et al. [9] and Cai and Zhao's [2] results can be extended to the case of the almighty adversary. In both cases, their proofs require assuming a fixed ordering, and it is not immediately clear how to alter the proofs to hold under our stronger adversary.

However, in the additive setting, we can borrow tools from the prophet inequality literature. We rely on the key observation that the additive setting effectively allows us to translate single-item results into results about multi-item combinatorial auctions. We use the following classic result from the single-item prophet inequality problem, modified to match our setting.

Theorem 2 (Samuel-Cahn, 1984 [20])**.** *In the single-item setting, there exist anonymous posted prices that guarantee $\frac{1}{2}$ of the optimal welfare, even under an almighty adversary.*

Proof. While Samuel-Cahn never specifically considered adding an almighty adversary, it can be shown that the guarantees hold even if the algorithm happens to select the smallest realization above the threshold each time ([18], Lecture #6). This precisely matches the case of the almighty adversary. □

noindent Next, we make use of the additive structure of the bidders' valuations to reduce our combinatorial auction to multiple copies of the single-item setting, thus proving the desired welfare (and revenue) guarantees. The basic reduction used in the proof of the following theorem forms the basis for all of the later results in this section.

Theorem 3. *For additive buyers, there exist anonymous posted prices that guarantee $\frac{1}{2}$ of the optimal welfare/revenue, even with an almighty adversary dictating the order in which the buyers arrive.*

Proof. As the buyers have additive valuations, buyer i will purchase an item j if and only if $v_i(j) \geq p_j$. Thus, we can think of our m-item auction as running m separate single-item auctions, with the same adversarial order used in all m auctions.

Let $\text{SW}_{almighty}(\mathcal{F}, p)$ denote the welfare attained under an almighty adversary, given joint valuation distribution $\mathcal{F}$ and using posted prices p. Let $X_p(v, \pi)$ denote the allocation when faced with valuations v and ordering π, using posted prices p Formally, we have that

$$
\begin{aligned}
\text{SW}_{almighty}(\mathcal{F}, p) &= \mathbb{E}_{v\sim\mathcal{F}} \min_{\pi} \text{SW}(X_p(v,\pi)) \\
&= \mathbb{E}_{v\sim\mathcal{F}} \min_{\pi} \sum_{j=1}^{M} \sum_{i=1}^{n} v_i(j) \cdot \mathbb{I}\left(j \in X_p(v,\pi)\right) && \text{(Additive Welfare)} \\
&\geq \mathbb{E}_{v\sim\mathcal{F}} \sum_{j=1}^{M} \min_{\pi} \sum_{i=1}^{n} v_i(j) \cdot \mathbb{I}\left(j \in X_p(v,\pi)\right) && \text{(Almightier Adversary)} \\
&= \sum_{j=1}^{m} \mathbb{E}_{v\sim\mathcal{F}} \min_{\pi} \sum_{i=1}^{n} v_i(j) \cdot \mathbb{I}\left(j \in X_p(v,\pi)\right) && \text{(Linearity)} \\
&= \sum_{j=1}^{m} \mathbb{E}_{v^j\sim\mathcal{F}} \min_{\pi} \sum_{i=1}^{n} v_i(j) \cdot \mathbb{I}\left(j \in X_p(v,\pi)\right)
\end{aligned}
$$

where v^j denotes the valuations for item j. The third inequality holds as we are actually giving the adversary more power in this setting, allowing them to set a different ordering of buyers for each item in each of the

m separate single-item auctions.

However, note that the final expression corresponds to the maximum achievable welfare when selling m separate items, assuming the almighty adversary can set any ordering in each of the m auctions. By setting $p_j = \text{Med}\left(\max_i v_i(j)\right)$, Theorem 2 guarantees that we get half of the optimal welfare for each item, and thus half of the optimal welfare for all m items combined. Let p^* denote these prices from the prophet inequality. Then we have that

$$\begin{aligned} \text{SW}_{almighty}(\mathcal{F}, p^*) &\geq \sum_{j=1}^{m} \mathbb{E}_{v^j \sim \mathcal{F}} \min_{\pi} \sum_{i=1}^{n} v_i(j) \cdot \mathbb{I}\left(j \in X_{p^*}(v, \pi)\right) \\ &\geq \sum_{j=1}^{m} \frac{1}{2} \cdot \mathbb{E}_{v^j \sim \mathcal{F}} \max_i v_i(j) = \frac{1}{2}\text{SW}_{opt}(\mathcal{F}) \qquad \text{(Theorem 2)} \end{aligned}$$

Replacing values with virtual values concludes the proof for revenue as well. □

Not only does this resolve the question for the additive setting, but it also shows that there is no worst case difference between having an almighty adversary and facing a fixed ordering.

Remark 1. It is well known that this $\frac{1}{2}$ bound is tight even for the non-adversarial case, meaning that there is no separation between the worst case under a fixed adversarial ordering or under an almighty adversary.

In fact, we can make further use of results from the prophet inequality literature to yield welfare-approximating posted prices given only nm queries of the buyers' valuations. That is, even if we don't have full knowledge of the joint distribution over valuations $\mathcal{F}$, it requires relatively few samples to construct "good" posted prices. We will need the following theorem from Rubinstein et al. [19].

Theorem 4 (Rubinstein et al., 2019 [19])**.** *Consider the standard prophet inequality problem with n random variables. Given a single sample $\tilde{X}_i$ from each random variable, setting $T = \max_i \tilde{X}_i$ yields $\frac{1}{2}$ of the optimal value in expectation, even against an almighty adversary.*

With the single-sample prophet inequality in hand, we can now transfer it over to our combinatorial setting.

Corollary 1. *Given mn queries to the buyers' valuation functions, we can construct posted prices that guarantee $\frac{1}{2}$ of the optimal welfare in expectation (including the randomness from the sampling), even against an almighty adversary.*

Proof. Following Theorem 4, we can take n samples for each of the m items, setting posted prices $p_j = \max_i \tilde{X}_{i,j}$ where $X_{i,j}$ is the sample for buyer i for item j. Note that we cannot reuse the same sample for multiple items as this may induce unwanted dependencies across the thresholds. Note that since we assume that we do not know $\mathcal{F}_i$, we cannot calculate a buyer's virtual value, meaning we cannot apply the same argument for revenue approximation. □

Remark 2. In fact, Feldman et al. [9] only guarantee the construction of prices within a factor of $\frac{1}{2}$ up to some additive constant ϵ in time POLY$(m, n, 1/\epsilon)$. Specifically, they require $O(4m^3 n/\epsilon^2 \cdot (\log m + \log n - \log \epsilon)$ samples for their Chernoff + union bound approach to go through. Our results in the additive setting give prices that deliver a $\frac{1}{2}$ ratio in expectation using far fewer samples.

Having established welfare and revenue guarantees for the single-unit setting, one might also consider what would happen in the multi-unit setting, when we have k copies of each item to sell. We can again leverage standard prophet inequality results to provide welfare and revenue guarantees in this case.

Theorem 5. *(Chawla et al., 2020 [3]): There exist anonymous posted prices that guarantee* $1 - O\left(\sqrt{\frac{\log k}{k}}\right)$ *of the optimal welfare when selling k units of a single item, even against an almighty adversary.*

The following result for the multi-unit, additive case then follows quite naturally.

Corollary 2. *In the multi-unit setting with k copies of each item and additive buyers, there exist posted prices that guarantee* $1 - O\left(\sqrt{\frac{\log k}{k}}\right)$ *of optimal welfare and optimal revenue, even with an almighty adversary dictating the order in which the buyers arrive.*

Proof Sketch. The proof follows exactly along the same lines as the proof of Theorem 3, reducing the larger combinatorial multi-unit auction to m separate multi-unit auctions. □

ii. Bounding the Adversarial Price of Anarchy

Given the positive results of the previous section, one might ask how much worse things can get under an almighty adversary when compared to a weaker adversary that must commit to a fixed ordering. Is there any gap at all between the almighty adversary and this weaker adversary? To quantify the power of the stronger almighty adversary, we consider the adversarial price of anarchy, defined below.

Definition 1. We define the **adversarial price of anarchy** as the worst-case ratio between the welfare achieved by the optimal mechanism under a fixed adversarial order and the welfare achieved by the optimal mechanism under an almighty adversary. Formally, we have that

$$\text{APoA} = \max_{\mathcal{F}} \frac{\max_p \min_\pi \mathbb{E}_{v\sim\mathcal{F}} \text{SW}(X_p(v,\pi))}{\max_p \mathbb{E}_{v\sim\mathcal{F}} \min_\pi \text{SW}(X_p(v,\pi))}$$

where $X_p^*(v, \pi)$ denotes the allocation of goods under prices p for valuations v and ordering π.

We have the following simple upper bound from the results of the previous section.

Proposition 1. *In the additive setting, the* APoA *is bounded above by* 2*.*

Proof. This is a simple corollary of Theorem 3 as for any valuation distributions $\mathcal{F}$ for the bidders, we have that

$$\max_p \mathbb{E}_{v\sim\mathcal{F}} \min_\pi \text{SW}(X_p(v,\pi)) \geq \frac{1}{2}\text{SW}_{opt}(v) \geq \frac{1}{2} \max_p \min_\pi \mathbb{E}_{v\sim\mathcal{F}} \text{SW}(X_p(v,\pi))$$

□

The more interesting case comes when trying to give a lower bound. We provide the following example, which we conjecture is tight.

Proposition 2. *In the additive setting, the* APoA *is bounded below by* $\frac{3}{2}$*.*

Proof. Consider the following simple example with $n = 2$ bidders and $m = 1$ goods. Suppose the bidders have identical distributions over valuation functions given by

$$v_i(1) = \begin{cases} 1 & \text{w.p. } \frac{2x-2}{2x-1} \\ x & \text{w.p. } \frac{1}{2x-1} \end{cases}$$

where $x > 1$. It can be shown that as $x \to \infty$, the ratio between the performance under a fixed adversarial ordering and an almighty adversary goes to $\frac{3}{2}$. The details of the proof are relegated to the appendix. □

In fact, we conjecture that this bound is tight.

Conjecture 1. *In the additive setting, the APoA is precisely equal to $\frac{3}{2}$. That is, the example from Proposition 2 is the worst possible case with an almighty adversary.*

Intuition. To give some intuition for why we believe this to be the case, note first that the i.i.d setting effectively renders the weaker adversary powerless as each fixed ordering is identical. Additionally, any auction with more than one item can be effectively translated into multiple single-item auctions, as demonstrated in the previous section [3]. The example from Proposition 2 is the worst possible case when we have 2 bidders, so an example demonstrating a worse ratio that satisfies i.i.d value distributions must depend on there being a larger number of bidders, which we were unable to produce. If we attempt to reproduce our example with $n = 3$ bidders and $m = 1$ goods, we find a ratio of $\frac{4}{3} < \frac{3}{2}$, suggesting that similar examples with a larger number of bidders will fail to find a larger ratio. □

III. Considering Submodular Valuations

Progressing up the hierarchy of subadditive functions, the natural next class to consider consists of submodular valuations. Unfortunately, once we enter the setting of submodular valuations, we can no longer reduce our larger auction into multiple single-item auctions, meaning that the same prophet inequality analysis used in Section II fails.

The most natural first step is to consider modifying the original proof from Feldman et al. [9] to account for the possibility of an almighty adversary. Unfortunately, their proof hinges on the key fact that the set of items available to a bidder upon arrival does not depend on their own valuations for the items. This is true for any fixed ordering, and in fact, it is even true even for a weaker online adversary that observes the same outcomes as the designer and can change the ordering in an online fashion. However, this is not the case for an almighty adversary, where the bidders' valuations very well may change their own position in the ordering.

i. Finding a Lower Bound

Given the increased power of the almighty adversary, one might conjecture that there exists a counterexample showing that $\frac{1}{2}$-guarantees of the optimal welfare are no longer possible. However, the worst case example under an adversarial fixed ordering does not get any worse under an almighty adversary.

Specifically, the canonical lower bound example proceeds as follows. Suppose we have one item and two buyers who value the item as follows:

$$v_1(1) = 1 \qquad\qquad v_2(1) = \begin{cases} 0 & \text{w.p. } 1-\epsilon \\ \epsilon & \text{w.p. } 1/\epsilon \end{cases}$$

[3] In fact, as seen in the proof of Theorem 3, the adversary has less power in these separate single-item auctions, having to commit to a single order for all auctions.

The optimal allocation yields welfare $2 - \epsilon$ while any posted prices under this ordering yield welfare 1.

In this case, the almighty adversary lacks any true power – while they can rearrange the bidders, this does not affect the final welfare from posted prices. Changing the distribution over valuations to let buyer 2 have a small nonzero valuation in the $1 - \epsilon$ case still does not prevent the designer from simply setting the prices precisely equal to 1 to ignore this small valuation. This same sort of logic seems to hold for most of the potential counterexamples one could try.

For this reason, we pose the following conjecture.

Conjecture 2. *Even under an almighty adversary, there still exist posted prices that yield $\frac{1}{2}$ of the optimal welfare for submodular valuations.*

We propose one potential approach to resolve this question in the following subsection.

ii. Online Content Resolution Schemes

One of the most promising directions to consider involves techniques from the literature on *online contention resolution schemes* (OCRSs) [10, 11, 13, 16]. First introduced by Feldman et al. [10], OCRSs yield simple schemes that provide guarantees even against an almighty adversary. At a high level, OCRSs work by optimizing a convex realization of the corresponding combinatorial optimization problem and using this optimal solution to derive a "good" algorithm for the original problem.

For a special class of OCRSs, termed *c-selectable greedy* OCRSs, their guarantees hold even against the almighty adversary from our setting. In fact, Feldman et al. [11] use OCRSs to guarantee constant revenue approximations for submodular bidders under various different types of matroid constraints, guaranteeing an $O(1)$ gap between the mechanism derived from their OCRS and the optimal mechanism.

Given that they are able to generate robust mechanisms in the revenue approximation case, one natural next step is to consider whether these OCRSs can also be used to make guarantees about welfare.

IV. Further Exploration

In addition to extending our welfare and revenue guarantees to submodular valuations, there are a number of other directions worthy of consideration. We present a few of what we believe to be the most interesting and tractable questions below.

- While guarantees for general submodular valuations may be difficult to find, can we guarantee anything under weaker assumptions, perhaps assuming unit demand buyers?
- Can we prove that our bound on the adversarial price of anarchy in the additive setting is tight? Namely, can we formalize the intuition we presented to do so?
- How does the adversarial price of anarchy behave for different subadditive functions? Given our conjecture that we can retain $\frac{1}{2}$-welfare guarantees, we conjecture that the adversarial price of anarchy is also the same as in the additive setting.

References

[1] Moshe Babaioff et al. "A simple and approximately optimal mechanism for an additive buyer". In: *Journal of the ACM (JACM)* 67.4 (2020), pp. 1–40.

[2] Yang Cai and Mingfei Zhao. "Simple mechanisms for subadditive buyers via duality". In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 2017, pp. 170–183.

[3] Shuchi Chawla, Nikhil Devanur, and Thodoris Lykouris. "Static pricing for multi-unit prophet inequalities". In: *arXiv preprint arXiv:2007.07990* (2020).

[4] José Correa and Andrés Cristi. "A Constant Factor Prophet Inequality for Online Combinatorial Auctions". In: (2023).

[5] José Correa et al. "Optimal item pricing in online combinatorial auctions". In: *Integer Programming and Combinatorial Optimization: 23rd International Conference, IPCO 2022, Eindhoven, The Netherlands, June 27–29, 2022, Proceedings*. Springer. 2022, pp. 126–139.

[6] Soheil Ehsani et al. "Prophet Secretary for Combinatorial Auctions and Matroids". In: *CoRR* abs/1710.11213 (2017). arXiv: `1710.11213`. URL: `http://arxiv.org/abs/1710.11213`.

[7] Liran Einav et al. "Auctions versus Posted Prices in Online Markets". In: *Journal of Political Economy* 126.1 (2018).

[8] Hossein Esfandiari et al. "Prophet secretary". In: *SIAM Journal on Discrete Mathematics* 31.3 (2017), pp. 1685–1701.

[9] Michal Feldman, Nick Gravin, and Brendan Lucier. "Combinatorial auctions via posted prices". In: *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*. SIAM. 2014, pp. 123–135.

[10] Moran Feldman, Ola Svensson, and Rico Zenklusen. "Online contention resolution schemes". In: *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*. SIAM. 2016, pp. 1014–1033.

[11] Moran Feldman, Ola Svensson, and Rico Zenklusen. "Online contention resolution schemes with applications to bayesian selection problems". In: *SIAM Journal on Computing* 50.2 (2021), pp. 255–300.

[12] Faruk Gul and Ennio Stacchetti. "Walrasian equilibrium with gross substitutes". In: *Journal of Economic Theory* 87.1 (1999).

[13] Euiwoong Lee and Sahil Singla. "Optimal online contention resolution schemes via ex-ante prophet inequalities". In: *arXiv preprint arXiv:1806.09251* (2018).

[14] Benny Lehmann, Daniel Lehmann, and Noam Nisan. "Combinatorial auctions with decreasing marginal utilities". In: *Proceedings of the 3rd ACM conference on Electronic Commerce*. 2001, pp. 18–28.

[15] Daniel Lehmann, Rudolf Müller, and Tuomas Sandholm. "The winner determination problem". In: *Combinatorial auctions* (2006), pp. 297–318.

[16] Vasilis Livanos. "Simple and Optimal Greedy Online Contention Resolution Schemes". In: *Advances in Neural Information Processing Systems*. 2022.

[17] David Porter et al. "Combinatorial auction design". In: *Proceedings of the National Academy of Sciences*. 2003.

[18] Tim Roughgarden. *Twenty lectures on algorithmic game theory*. Cambridge University Press, 2016.

[19] Aviad Rubinstein, Jack Z Wang, and S Matthew Weinberg. "Optimal single-choice prophet inequalities from samples". In: *arXiv preprint arXiv:1911.07945* (2019).

[20] Ester Samuel-Cahn. "Comparison of threshold stop rules and maximum for independent nonnegative random variables". In: *the Annals of Probability* (1984), pp. 1213–1216.

A. Appendix

i. Proof of Theorem 2

Recall that we are selling $m = 1$ item with $n = 2$ bidders, each with identical distributions over valuation functions given by

$$v_i(1) = \begin{cases} 1 & \text{w.p. } \frac{2x-2}{2x-1} \\ x & \text{w.p. } \frac{1}{2x-1} \end{cases}$$

Note that we only need to consider two possible choices of posted prices – the price for the item can either be set to be below 1 or in the interval $(1, x]$. Any higher price would result in zero welfare. It can be easily verified that in both the case of an adversarially fixed ordering and in the case of the almighty adversary, setting a price of 1 is optimal [4].

As our buyers are i.i.d., an adversarially fixed ordering is indifferent between either of the two orderings. Meanwhile, the almighty adversary will always choose to place a bidder with realized value 1 before a bidder with realized value x. Thus, comparing the welfare achieved under each adversary, we have that

$$\max_p \mathbb{E}_{v\sim\mathcal{F}} \min_\pi \text{SW}(X_p(v,\pi)) = 1 - \frac{1}{(2x-1)^2} + x \cdot \frac{1}{(2x-1)^2}$$
$$\max_p \min_\pi \mathbb{E}_{v\sim\mathcal{F}} \text{SW}(X_p(v,\pi)) = \frac{2x-2}{2x-1} + \frac{x}{2x-1}$$

Sending $x \to \infty$, we have that

$$\lim_{x\to\infty} \max_p \mathbb{E}_{v\sim\mathcal{F}} \min_\pi \text{SW}(X_p(v,\pi)) = 1$$
$$\lim_{x\to\infty} \max_p \min_\pi \mathbb{E}_{v\sim\mathcal{F}} \text{SW}(X_p(v,\pi)) = \frac{3}{2}$$

concluding the proof of our lower bound.

[4] In the case of an almighty adversary, the designer is actually indifferent between setting either of the two prices.