Online Combinatorial Auctions

Yuan Deng* Debmalya Panigrahi† Hanrui Zhang†

Abstract

We study combinatorial auctions in online environments with the goal of maximizing social welfare. In this problem, new items become available on each day and must be sold before their respective expiration dates. We design online auctions for the widely studied classes of submodular and XOS valuations, and show the following results:

- For submodular valuations, we give an $O(\log m)$-competitive mechanism for adversarial valuations and an $O(1)$-competitive mechanism for Bayesian valuations, where m is the total number of items. Both these mechanisms are computationally efficient and universally truthful for myopic agents, i.e., agents with no knowledge of the future.

- For XOS valuations, we show that there is no online mechanism that can achieve a competitive ratio of $o\left((m/\log m)^{1/3}\right)$ even in a Bayesian setting. Our lower bound holds even if we do not require truthfulness and/or computational efficiency of the mechanism.

This establishes a sharp separation between XOS valuations and its subclass of submodular valuations for online combinatorial auctions. In contrast, no such separation exists for offline auctions, where the best bounds for both submodular and XOS valuations are $O((\log\log m)^3)$ for adversarial settings (Assadi and Singla, FOCS 2019) and $O(1)$ for Bayesian settings (Dütting *et al.*, FOCS 2017).

In contrast to the above, if items do not expire and only need to be sold before the market closes, then we give a reduction from offline to online mechanisms that preserves the competitive ratio for all subadditive valuations (that includes XOS and submodular valuations), thereby achieving the same bounds as the respective best offline mechanisms.

1 Introduction

Inspired by the popularity of selling online advertising opportunities via repeated auctions, there has been a growing body of literature on dynamic mechanism design for additive valuations in the past decade (e.g., [1, 25, 27]). Dynamic mechanisms create the possibility of boosting revenue and/or welfare by evolving the auctions across time. In spite of the success of dynamic mechanisms for additive buyers, little is known for dynamic mechanism design with buyers whose valuations are combinatorial. This is frequently the case for buyers participating in marketplaces such as Amazon and eBay, where large volumes of heterogeneous items are sold. In many cases, these items may be perishable, or have to be sold within a stipulated time frame before they lose value, and therefore, cannot be sold via offline auctions. In light of these, we initiate the study of combinatorial auctions in *online* environments in this paper. In this problem, items arrive every day and have expiry dates before which they need to be sold. The goal is to design an auction that maximizes social welfare, where the auction dynamically evolves over time but does not have knowledge of the future.

Combinatorial auctions (see, e.g., [6]) are a central object of study in the field of algorithmic game theory. In a combinatorial auction, we are given a set U of m items and a set of n buyers with respective valuation functions $(v_1, \cdots, v_n)$ defined on all subsets of items. The goal is to design an auction that allocates the items to the buyers $S = (S_1, \ldots, S_n)$ (i.e., buyer i receives the subset of items S_i where $S_i \cap S_j = \emptyset$) such that the social welfare, defined as $\sum_i v_i(S_i)$, is maximized. In a seminal work, Dobzinski et al. [14] provided the first *truthful*[1] and computationally efficient mechanism that approximates the social welfare to a factor of $O(\sqrt{m})$ for general monotone combinatorial valuations and $O(\log^2 m)$ when restricted to XOS valuations. Since then, welfare-optimal combinatorial auctions have been extensively studied and the current best approximations are $O(\log m \log\log m)$ for subadditive valuations [9] and $O\left((\log\log m)^3\right)$ [2] for XOS valuations. This line of work establishes a relatively clear picture of combinatorial auctions in *static* settings.

In the online environment, we assume (wlog[2]) that

*Google Research. Email: dengyuan@google.com.

†Department of Computer Science, Duke University. Email: {debmalya,hrzhang}@cs.duke.edu.

[1] We use *truthful* and *incentive compatible* interchangeably to mean that a buyer cannot profit by misreporting their valuation.

[2] This is without loss of generality: for a buyer who may join and leave the market multiple times, her marginal valuation can

there is a fixed set of buyers throughout the entire time horizon. The items arrive online in batches over time (e.g., a new batch of items arrives every day). Neither the platform (or sellers) nor the buyers have knowledge of items that will appear in the future. Each item has an expiration date before which the item must to sold. Moreover, once an item is sold to a buyer, the seller cannot retrieve the item and reallocate it later. The central question that we address in this paper is: *can we design welfare-optimal combinatorial auctions that are truthful and computationally efficient in an online environment?*

1.1 Our Results In this paper, we focus on complement-free buyers, i.e., buyers with subadditive valuations, its widely studied subclass of XOS valuations, and a further well-studied subclass of submodular valuations. We assume that in addition to *value queries*, where a buyer is asked to report her value for a particular set of items, the seller is allowed to use *demand queries* to ask for a buyer's favorite bundle given specific item prices. We consider both the prior-free and Bayesian settings. In both settings, the buyers' valuations for future goods are chosen by an adversary in an adaptive manner. In the Bayesian setting, the seller additionally has access to the distributions of the buyers' valuations over the existing goods.

Online Allocation for XOS Buyers. We begin with online allocations for complement-free buyers even without requiring truthfulness and/or efficiency. For XOS buyers (and by generalization also for subadditive buyers), we show a lower bound of $\Omega\left((m/\log m)^{1/3}\right)$ on the competitive ratio[3] of randomized algorithms even in the Bayesian setting (Section 3). (Recall that m denotes the total number of items.) This polynomial lower bound sharply contrasts with upper bounds of 2 [15] and $O((\log\log m)^3)$ [2] known for *offline* combinatorial auctions with XOS valuations in the Bayesian and prior-free settings respectively.

Online Mechanism for Submodular Buyers. Given the lower bound above, we restrict our attention to a widely studied subclass of XOS valuations, namely submodular valuations. It is well-known that an algorithm that allocates each item immediately when it becomes available to the buyer with the highest marginal valuation is 2-competitive for submodular valuations [21, 23]. Hence, the above lower bound does not hold for submodular valuations. But, what can we achieve in terms of *welfare-optimal mechanisms* for submodular buyers, i.e., where we also desire the allocation to be incentive compatible and computationally efficient? Since neither the seller nor the buyers have any prior knowledge about the future, we model the buyers' strategic behavior by assuming that they maximize their utility for the current stage because they do not know about future opportunities. (We call this *myopic* behavior.) Even assuming myopic behavior, the design of a truthful, competitive mechanism for submodular buyers turns out to be challenging. For instance, the greedy algorithm has a competitive ratio of 2, but is not incentive compatible. On the other hand, selling each batch in separate second-price auctions is truthful, but is not competitive in terms of social welfare.

Our main contributions are to design online mechanisms that achieve competitive ratios of $O(\log m)$ in the prior-free setting and 8 in the Bayesian setting (Section 4 and Section 5). Note that these are (at least) exponentially better than the lower bound for XOS buyers, thereby establishing a clear separation between these two classes. This is in sharp contrast to offline combinatorial auctions, where the best bounds for submodular valuations are identical to those for XOS valuations. Interestingly, our mechanisms sell every item as soon as they become available and do not depend on the expiry dates. Moreover, they are universally truthful for myopic buyers, and are computationally efficient.

No Expiration. We also consider a setting where the items do not expire, but must be sold before the market closes (the closing date being unknown in advance to the platform and the sellers). Our negative results for XOS valuations rely on the fact that all the items must be sold immediately on becoming available, and therefore, do not apply in this case. Indeed, we give a reduction to convert offline mechanisms to this setting while preserving the competitive ratio up to a constant factor. As a result, the best bounds for subadditive [9, 16] and XOS/submodular [2, 15] buyers for both the prior-free and Bayesian offline settings generalize to the online case when there are no expiry dates (Section 6).

Our results are summarized in Table 1. (We have also provided existing offline results for comparison.)

1.2 Technical Overview We demonstrate a sharp separation between XOS and submodular valuations in terms of the approximation ratio. XOS and submodular valuations are often considered to be at the same level of complexity in truthful welfare-maximizing combinatorial auction design, since most techniques developed for submodular valuations can be directly generalized to XOS valuations (e.g., [2, 11, 14]). This is mostly due to the fact that prior work relies heavily on a (by now, standard) revenue-utility decomposition argument (see, e.g., [15, 17, 20]). However, in online environments, this ar-

be set to be 0 for stages when she is absent from the market.

[3]The *competitive ratio* of an online allocation is the worst case ratio of its (expected) welfare and that of the optimal allocation.

		Online	No Expiration	Offline
Prior-free	Subadditive	$\Omega\left((m/\log m)^{1/3}\right)$	$O(\log m \log\log m)$	$O(\log m \log\log m)$ [9]
	XOS	$\Omega\left((m/\log m)^{1/3}\right)$	$O\left((\log\log m)^3\right)$	$O\left((\log\log m)^3\right)$ [2]
	Submodular	$O(\log m)$	$O\left((\log\log m)^3\right)$	$O\left((\log\log m)^3\right)$ [2]
Bayesian	Subadditive	$\Omega\left((m/\log m)^{1/3}\right)$	$O(\log\log m)$	$O(\log\log m)$ [16]
	XOS	$\Omega\left((m/\log m)^{1/3}\right)$	≤ 8	2 [15]
	Submodular	≤ 8	≤ 8	2 [15]

Table 1: Summary of our results. We include the results of offline versions for reference.

gument fails in a fundamental way that is demonstrated and exploited by our lower bound for XOS valuations.

Instead, our online mechanisms for submodular valuations rely on a carefully designed online pricing scheme that is truthful for myopic buyers. Our analysis relates the welfare accumulated online to the offline optimum via a proxy benchmark that is a constant approximation of the offline optimum. This proxy benchmark is useful because we ensure that, unlike the actual offline optimum, this benchmark evolves in a relatively "stable" manner over time as new items are added. Also key to our analysis is an associated novel revenue-utility decomposition argument that is tailored to submodular (as opposed to XOS) valuations (see Lemma 4.6 for the prior-free setting and Lemma 5.3 for the Bayesian setting), and might be useful in arguing about submodular valuations in other problem settings.

The mechanisms above are truthful only for myopic buyers, i.e., who cannot foresee the future. While this is a reasonable (and perhaps the most natural) assumption in an online setting, let us briefly also consider the case of *omniscient* buyers, i.e., who can plan with knowledge of the future. Interestingly, designing an "interesting" truthful mechanism (i.e., one that depends on the valuations in a non-trivial way) in this case *even without welfare or efficiency guarantees* seems non-trivial. In fact, note that running separate second-price auctions for each batch of items is *not* truthful for omniscient submodular buyers. (This is in contrast to additive buyers, for whom this auction is indeed truthful even if they are omniscient.) We leave the case of omniscient buyers as an interesting direction for future work.

2 Preliminaries

We consider a setting with n buyers and one seller. The goods arrive in batches in T stages where T is unknown to the seller. We let the set of newly arrived goods at stage t be B^t. The entire set of goods is $U = \bigcup_{t=1}^{T} B^t$. For convenience, we will use the notation $U^{(t,t')} = \bigcup_{\tau=t}^{t'} B^\tau$ to represent the items arriving between stage t and stage t'. Let $m_t = |U^{(1,t)}|$ be the total number of items in the first t stages. As usual, we use $-i$ to indicate the buyers other than buyer i.

The buyers' valuations are combinatorial and we assume the valuations are normalized, i.e., $v_i(\emptyset) = 0$ for all buyers i, and monotone, i.e., $v_i(S) \geq v_i(S')$ for all $S \supseteq S'$. For convenience, we give each item an index j in a chronological order of its arrival. More precisely, the items in B^t are indexed between $(m_{t-1}+1)$ and m_t. Moreover, for item j, we will write $v_i(j) = v_i(\{j\})$ for short. Each item j has an expiration date $e(j)$ and for an item j without expiration date, $e(j) = \infty$. We focus on subadditive valuations in this paper.

Definition 2.1. *A valuation v is subadditive if for every bundle S and S' such that $S \cap S' = \emptyset$, we have $v(S) + v(S') \geq v(S \cup S')$.*

Definition 2.2. *A valuation v is additive if for every bundle $S \subseteq U$, we have $v(S) = \sum_{j \in S} v(\{j\})$. A valuation v is XOS if there exist additive valuations $a_1, \cdots, a_q$ such that for every bundle $S \subseteq U$, we have $v(S) = \max_r a_r(S)$. A valuation v is submodular if for every bundle S and S', we have $v(S) + v(S') \geq v(S \cup S') + v(S \cap S')$.*

For an XOS valuation v with associated additive valuations $a_1, \cdots, a_q$, each a_r is called a *clause* of v. If $a^* \in \arg\max_{a_r} a_r(S)$, then we say a^* is a *maximizing clause* of S and $a^*(\{j\})$ is the *supporting price* of good j in this maximizing clause. It is well-known that submodularity implies XOS, and XOS implies subadditivity. The marginal valuation of a buyer on an additional bundle S' given that she already has a bundle S is represented by $v(S' \mid S) = v(S' \cup S) - v(S)$.

Online Environments. We describe the online environments for submodular valuations; the corresponding environments for XOS and subadditive valuations can be defined in a similar way. Given a submodular valuation v over $U^{(1,t)}$ and a submodular valuation v' over $U^{(1,t+1)}$, we say v' is *extendable* from v if for all $S \subseteq U^{(1,t)}$, $v'(S) = v(S)$. In a prior-free environment, we consider a setting where the valuations are selected by an (oblivious) adversary. In other words, the buyers' valuations for future goods can be arbitrary but must be extendable from the valuation over the existing goods.

In the Bayesian setting, given a distribution F_i^t of submodular valuations over $U^{(1,t)}$ with support $\mathcal{V}^t$ and a distribution F_i^{t+1} of submodular valuations over $U^{(1,t+1)}$ with support $\mathcal{V}^{t+1}$, we say F_i^{t+1} is *extendable* from F_i^t if there exists a partition of $\mathcal{V}^{t+1}$ as $\{Q_v\}_{v\in\mathcal{V}^t}$, such that for each $v \in \mathcal{V}^t$, we have: (1) Q_v is non-empty; (2) for all $v' \in Q_v$, v' is extendable from v; and (3) $\sum_{v'\in Q_v} \Pr[v'|F_i^{t+1}] = \Pr[v|F_i^t]$. For F_i^{t+1} that is extendable from F_i^t and the buyer's valuation v over $U^{(1,t)}$, the buyer's valuation v' over $U^{(1,t+1)}$ is randomly drawn from Q_v such that the probability of choosing $v' \in Q_v$ is $\frac{\Pr[v'|F_i^{t+1}]}{\Pr[v|F_i^t]}$. We assume that F_i^t is publicly known and independent across buyers for all t and the buyers' distributions for the future goods can be arbitrarily chosen but must be extendable from the distributions over the existing goods.

Let V be some set of valuations. We use $\langle f, p\rangle$ to denote a deterministic online mechanism. $f_i^t : V^n \to 2^{U^{(1,t)}}$ is the allocation function that maps the valuation profile $\vec{v} = (v_1, \cdots, v_n)$ to a subset of goods, indicating the set of goods allocated to buyer i in the first t stages. The payment function $p_i^t : V^n \to \mathbb{R}$ maps the valuation profile to buyer i's cumulative payment for the first t stages. An allocation rule is valid if for all t, and two different buyers i, i', $f_i^t(\vec{v}) \cap f_{i'}^t(\vec{v}) = \emptyset$. Moreover, once a good is sold, the seller cannot retrieve the good and reallocate it in the future, i.e., $\forall t$, $f_i^t(\vec{v}) \subseteq f_i^{t+1}(\vec{v})$ for each buyer i; and the item must be sold before its expiration date, i.e., for any j, we have $j \in f_i^{t'}(\vec{v})$ for all $t' \ge e(j)$ if and only if $j \in f_i^{e(j)}(\vec{v})$. Furthermore, for a stage t and two different valuation profiles $\vec{v}$ and $\vec{v}'$ satisfying $v_i(S) = v_i'(S)$ for all buyer i and $S \subseteq U^{(1,t)}$, we must have $f_i^t(\vec{v}) = f_i^t(\vec{v}')$ and $p_i^t(\vec{v}) = p_i^t(\vec{v}')$ for all buyer i.

Universally Truthful Mechanisms. We consider myopic buyers, and therefore, incentive compatibility only concerns the current stage without taking the future into account. In both the prior-free and Bayesian settings, we are interested in designing universally truthful mechanisms.

DEFINITION 2.3. *A deterministic mechanism $\langle f, p\rangle$ is truthful if for every stage t, every buyer i, and any valuations $v_i, v_i' \in V$ with $v_i'(S) = v_i(S)$ for all $S \subseteq U^{(1,t-1)}$, and any $\vec{v}_{-i} \in V^{n-1}$, we have*

$$v_i\big(f_i^t(v_i, \vec{v}_{-i})\big) - p_i^t(v_i, \vec{v}_{-i}) \ge v_i\big(f_i^t(v_i', \vec{v}_{-i})\big) - p_i^t(v_i', \vec{v}_{-i}).$$

A randomized mechanism $\langle f, p\rangle$ is universally truthful if it is a probability distribution over truthful deterministic mechanisms.

Competitive Ratio. Let S_i^t be the set of items allocated to buyer i at the end of stage t. We will use the vectorized symbol without subscript $\vec{S}^t = (S_1^t, \cdots, S_n^t)$ to represent the overall allocation at the end of stage t. For convenience, we will use $A^t = \bigcup_{i=1}^n S_i^t$ to represent the set of items sold in the first t stages. The welfare with respect to an allocation S is denoted by $v(\vec{S}) = \sum_i v_i(S_i)$. For a set of items U', the welfare-optimal allocation with respect to a valuation profile $\vec{v}$ is represented by $\mathrm{OPT}(U', \vec{v}) = \big(\mathrm{OPT}_1(U', \vec{v}), \cdots, \mathrm{OPT}_n(U', \vec{v})\big)$. We will drop $\vec{v}$ from the notation when it is clear from the context. The performance of our mechanism is measured by its competitive ratio:

DEFINITION 2.4. (COMPETITIVE RATIO) *For a set V of valuations, in a prior-free setting, an online mechanism $\langle f, p\rangle$ is κ-competitive if for any $(v_1, \cdots, v_n) \in V^n$ and $1 \le t \le T$:*

$$\kappa \cdot \mathbb{E}\left[v\left(f_1^t(U^{(1,t)}), \cdots, f_n^t(U^{(1,t)})\right)\right] \ge v\left(\mathrm{OPT}(U^{(1,t)})\right)$$

where the expectation is taken over the randomness of the mechanism. In the Bayesian setting, an online mechanism $\langle f, p\rangle$ is κ-competitive if for any independent distributions $F_1, \cdots, F_n \in \Delta(V)$ and $1 \le t \le T$:

$$\kappa \cdot \mathop{\mathbb{E}}_{\vec{v}}\left[v\left(f_1^t(U^{(1,t)}), \cdots, f_n^t(U^{(1,t)})\right)\right] \ge \mathop{\mathbb{E}}_{\vec{v}}\left[v\left(\mathrm{OPT}(U^{(1,t)})\right)\right]$$

where the expectation is additionally taken over $\vec{v}$ randomly drawn from the prior $\prod_i F_i$.

3 Lower bound for XOS valuations

We show that for XOS valuations, no (randomized) truthful mechanism is $o\left((m_T/\log m_T)^{1/3}\right)$-competitive even in the Bayesian setting (the same lower bound naturally holds in the prior free case as well). Our lower bound is information-theoretic, which means that it holds even if we do not require the mechanism to be truthful or computationally efficient. Moreover, our construction works even when the buyers are symmetric, i.e., all the buyers have the same valuation distributions.

THEOREM 3.1. *When the buyers' valuations are XOS and all items expire immediately, no randomized mechanism is $o(n)$-competitive for $m_T = \Omega(n^3 \log n)$, even if all buyers have i.i.d. valuations.*

Before proving the theorem, we provide the high level idea about our construction.We consider an online environment with T stages with a single new item at each stage, i.e., item j arrives at stage j. So the total number of items is T. For each buyer, we will construct an XOS valuation with n (same as the number of buyers) clauses. For ease of presentation, we represent buyer k's XOS valuation by a matrix Z_k with n rows and (eventually) T columns, such that row i corresponds to the i-th clause and column j corresponds to the j-th item, i.e., $Z_k(i,j)$ is the value of item j in the i-th clause of buyer k. On the arrival of item j, for each buyer k, we add a new column in Z_k as follows: pick a row i uniformly at random (and independent of any other choice) and assign $Z_k(i,j) = 1$ and $Z_k(i',j) = 0$ for all $i' \neq i$.

We first argue the performance of the optimal offline allocation on these valuations. Let c_k be the k-th clause of buyer k, i.e., the k-th row in Z_k. We will allocate item j to any buyer k with $c_k(j) = 1$; and if such a buyer does not exist, then we allocate item j arbitrarily. Notice that for each pair (k,j) of buyer k and item j, we have $\Pr[c_k(j) = 1] = 1/n$, and therefore, $\Pr[\exists k,\ c_k(j) = 1] \geq 1 - e^{-1} = \Omega(1)$. By linearity of the expectation, the expected welfare of the optimal offline allocation is $\Omega(T)$.

We are left to bound the welfare generated by an online algorithm. In order to build intuition, let us first make the simplifying, but false, assumption that the online algorithm cannot observe the realization of $Z_k(i,j)$ for all buyers k and clauses i when a new item j arrives. Now, suppose that the online algorithm assigns s_k items in total to buyer k. Note that for each of these s_k items, exactly one clause chosen uniformly at random has a valuation of 1, and all other clauses have valuation of 0, for buyer k. If we think of the n clauses as bins, and a valuation of 1 for each of the s_k items as balls being thrown uniformly at random into the bins, then the clause with the maximum valuation for these s_k items corresponds to the bin with the most balls. Using this correspondence, a simple calculation then shows that the welfare of the online algorithm summed over all the buyers concentrates around T/n, thereby giving us the lower bound we are after.

But, our simplifying assumption is false because the adversary must reveal the realization of $Z_k(i,j)$ for all clauses i and buyers k when item j arrives. Recall that our goal, for any buyer k, is to extend one clause with 1 and the other $n-1$ clauses with 0s without revealing which clause got a 1. To this end, on the arrival of item j, we create a temporary matrix Z'_k with $2n$ rows such that both the $(2i-1)$-th and $(2i)$-th rows of Z'_k are copies of the i-th row of the current Z_k. For each i, we will assign $Z'_k(2i-1,j) = 1$ and $Z'_k(2i,j) = 0$, and present Z'_k to the online algorithm at stage j. After the end of stage j, we pick an index $i \in \{1, 2, \ldots, n\}$ uniformly at random, and reconstruct Z_k as follows: $Z_k(i,\cdot) = Z'_k(2i-1,\cdot)$ while for all $i' \neq i$, we have $Z_k(i',\cdot) = Z'_k(2i',\cdot)$. Such a procedure successfully hides the random choice of i at stage j since Z'_k does not contain any information about the random choice. We "discard" the remaining n clauses in Z'_k by giving them valuations of 0 for all items henceforth.

Now we give the full proof of the Theorem 3.1.

Proof. Assume that there is one item arriving per stage. We construct the prior distribution by designing a scheme to generate the collection of additive valuations randomly, proceeding from the first item to the last item. At stage t, a clause is either *outstanding* or *done*. If a clause c is done at stage t, then for all $j > t$, $c(j) = 0$. Our scheme maintains a set of n outstanding clauses at every stage.

Suppose at the end of stage t, the n outstanding clauses are $\{c_1^t, \ldots, c_n^t\}$. Upon the arrival of the $(t+1)$-th item at stage $(t+1)$, we create $2n$ new clauses $c_1'^{t+1}, \ldots, c_n'^{t+1}$ and $c_1''^{t+1}, \ldots, c_n''^{t+1}$. Clause $c_a'^{t+1} = [c_a^t, 1]$ is obtained by simply appending 1 to the end of c_a^t, and $c_a''^{t+1} = [c_a^t, 0]$ is obtained by appending 0 to c_a^t. After the mechanism allocates the $(t+1)$-th item, the adversary flips a coin to choose uniformly at random some $i^{t+1} \in \{1, \cdots, n\}$. The new outstanding clauses are then

$$c_a^{t+1} = \begin{cases} c_a'^{t+1}, & a = i^{t+1} \\ c_a''^{t+1}, & \text{otherwise} \end{cases}.$$

We now analyze this construction after T stages. Let $Z_{k,i}$ be the i-th outstanding clause at time T in buyer k's valuation. First observe that the offline optimal allocation can achieve welfare $\Omega(T)$. This is because for each item j and any buyer k, with probability $\frac{1}{n}$, $Z_{k,k}(j) = 1$, and with probability

$$1 - \left(1 - \frac{1}{n}\right)^n \approx 1 - \frac{1}{e},$$

there is some buyer, denoted by k_j,[4] whose valuation satisfies $Z_{k_j,k_j}(j) = 1$. If such a buyer does not exist, let $k_j = 0$.

Knowing the future realization of all valuations, the offline optimal allocation could assign item j to buyer k_j, and discard j (possibly by assigning j to an arbitrary buyer) when $k_j = 0$. Whenever $k_j \neq 0$, item j contributes 1 to the total welfare. By the definition of

[4] If there are multiple such buyers, let k_j be the one with the smallest index, or simply any of them.

XOS valuations, the expected welfare would be

$$\mathbb{E}\left[\sum_{k\in[n]} v(\{j \mid k_j = k\})\right]$$
$$\geq \mathbb{E}\left[\sum_{k\in[n]} Z_{k,k}(\{j \mid k_j = k\})\right] = \mathbb{E}\left[\sum_{k\in[n]}\sum_{j:k_j=k} Z_{k,k}(j)\right]$$
$$= \mathbb{E}\left[\sum_j \mathbb{I}[k_j \neq 0]\right] = \sum_j \left(1 - \left(1 - \frac{1}{n}\right)^n\right) = \Omega(T).$$

Now consider the welfare obtained by any online mechanism. We upper bound the welfare by upper bounding the value of each buyer separately.

Fix some buyer k_0, random indices i^t for all t, and outstanding clauses $\{c_a^t\}$ drawn for k_0 for all t, and a realized allocation $S_{k_0}^T$ generated by the online mechanism. Let t_ℓ be the ℓ-th item assigned to k_0 for $1 \leq \ell \leq |S_{k_0}^T|$. When item t_ℓ is assigned to k_0, the value of clause $c_a^{t_\ell}(S_{k_0}^{t_\ell})$ increases by 1 if and only if $i^{t_\ell} = a$, which happens with probability exactly $\frac{1}{n}$. Therefore,

$$\mathbb{E}[c_a^T(S_{k_0}^T)] = \frac{|S_{k_0}^T|}{n}. \tag{3.1}$$

Moreover, by the Chernoff bound,

$$\Pr\left[c_a^T(S_{k_0}^T) - \frac{|S_{k_0}^T|}{n} \geq \sqrt{|S_{k_0}^T| \log n}\right] \leq \exp(-2\log n) = \frac{1}{n^2}.$$

Taking union bound over the outstanding clauses after stage T, we have

$$\max_a c_a^T(S_{k_0}^T) \leq \frac{|S_{k_0}^T|}{n} + \sqrt{|S_{k_0}^T| \log n} \tag{3.2}$$

with probability at least $1 - \frac{1}{n}$, and with probability at most $\frac{1}{n}$, $\max_a c_a^T(S_{k_0}^T)$ is at most $|S_{k_0}^T|$. Moreover, notice that we have

$$\mathbb{E}[v_{k_0}(S_{k_0}^T)] \leq \mathbb{E}[\max_a c_a^T(S_{k_0}^T) + 1].$$

The extra 1 in the inequality is due to the fact that the maximizing clauses at stage T might be one of the clauses created at stage T rather than one of the outstanding clause after stage T. Nevertheless, note that the difference between their valuations is at most 1. Using (3.1) and (3.2), we have

$$\mathbb{E}[v_{k_0}(S_{k_0}^T)] \leq \frac{n-1}{n} \cdot \left(\frac{|S_{k_0}^T|}{n} + \sqrt{|S_{k_0}^T| \log n}\right) + \frac{1}{n} \cdot |S_{k_0}^T| + 1$$
$$\leq \frac{2|S_{k_0}^T|}{n} + \sqrt{|S_{k_0}^T| \log n} + 1.$$

Finally, we sum over k_0:

$$\mathbb{E}[v(\vec{S}^T)] \leq \sum_{k_0}\left(\frac{2|S_{k_0}^T|}{n} + \sqrt{|S_{k_0}^T| \log n} + 1\right)$$
$$\leq \frac{2T}{n} + \sqrt{Tn\log n} + n.$$

As long as $m_T = T = \Omega(n^3 \log n)$, the above inequality implies that $\mathbb{E}[v(\vec{S}^s)] = O\left(\frac{T}{n}\right)$. ☐

4 Prior-free Online Mechanism for Submodular Valuations

In this section, we proceed to design our online mechanisms for submodular valuations. Our online mechanism for the prior-free setting PriorFreeOnline consists of a second price auction and a random fixed-price auction. Before the first stage, the seller flips a fair coin and runs the *second price auction* if it is heads, and the *random fixed-price auction* otherwise.

Second Price Auction. At stage t, we will run a second price auction without a reserve price for the bundle B^t that arrives at stage t. Notice that the second price auction is deterministic and dominant-strategy incentive-compatible, and thus, at stage t, each buyer i will truthfully report her marginal valuation over the bundle: $v_i(B^t \mid S_i^{t-1})$. (Recall that S_i^{t-1} is the set of items allocated to buyer i till stage $t-1$.)

Random Fixed-Price Auction. We first divide the buyers into two groups: STAT and MECH where the group for each buyer is chosen independently and uniformly at random. Note that such a partitioning is done only once before the arrival of the first item. For convenience, let $\mathcal{O}(U', \mathcal{C})$ be the set of all possible allocations for items U' and buyers $\mathcal{C} \subseteq [n]$ and $\mathrm{OPT}_{\mathcal{C}}^{(t_1,t_2)} = \arg\max_{\vec{S}\in\mathcal{O}(U^{(t_1,t_2)},\mathcal{C})} \sum_i v_i(S_i)$ represent the welfare maximizing allocation of items between stage t_1 and t_2 to buyers $\mathcal{C}$. As a shorthand, let $\mathrm{OPT}^{(t_1,t_2)} = \mathrm{OPT}_{[n]}^{(t_1,t_2)}$ be the welfare maximizing allocation to all buyers. In addition, let $\mathrm{OPT}^{(t_1,t_2)}|_{\mathcal{C}}$ be the optimal allocation for all buyers restricted to buyers in $\mathcal{C}$, such that $\mathrm{OPT}_i^{(t_1,t_2)}|_{\mathcal{C}} = \mathrm{OPT}_i^{(t_1,t_2)}$ if $i \in \mathcal{C}$ and otherwise, $\mathrm{OPT}_i^{(t_1,t_2)}|_{\mathcal{C}} = \emptyset$. Note that in general, $\mathrm{OPT}_{\mathcal{C}}^{(t_1,t_2)} \neq \mathrm{OPT}^{(t_1,t_2)}|_{\mathcal{C}}$. For each stage t, we maintain an estimate est_t of the optimal welfare of allocating $U^{(1,t)}$ to buyers in STAT.

THEOREM 4.1. ([18, 21, 23]) *For complement-free valuations, there exists a computationally efficient 2-approximation estimation algorithm for the optimal welfare using demand queries and value queries. The estimate* est_t *obtained from the algorithm satisfies* $\frac{1}{2}v(\text{OPT}^{(1,t)}|_{\text{STAT}}) \le \text{est}_t \le v(\text{OPT}^{(1,t)})$.

Given the estimate est_t, let $P_t = \{(\frac{\text{est}_t}{c \cdot m_t^2}), (\frac{\text{est}_t}{\frac{c}{2} \cdot m_t^2}), \cdots, (\frac{c}{2} \cdot m_t^2 \cdot \text{est}_t), (c \cdot m_t^2 \cdot \text{est}_t)\}$ be a set of prices, where c is a sufficiently large constant and $m_t = |U^{(1,t)}|$. Intuitively, the size of P_t is $O(\log m_t)$ and the price grows geometrically such that the j-th price is $\frac{\text{est}_t}{c \cdot m_t^2} \cdot 2^{j-1}$. We are now ready to describe our random fixed-price auction for the buyers in MECH (see Algorithm 4.1). Notice that given a fixed price p^t, the fixed-price auction is truthful at stage t.

ALGORITHM 4.1. (RANDOM FIXED-PRICE AUCTIONS)

for each stage t **do**
  Set p^t to be a price drawn from P_t uniformly at random
  Let $M = B^t$
  for each buyer $i \in \text{MECH}$ in some arbitrary order **do**
    Let $D_i = \arg\max_{S \in M} v_i(S \mid S_i^{t-1}) - p^t|S|$
    Allocate D_i to buyer i, i.e., $S_i^t = S_i^{t-1} \cup D_i$, and charge her $p^t|D_i|$
    $M = M \setminus D_i$
  end for
end for

Compared to the mechanism used in [14, 22] for the static setting, one main difference is that our mechanism samples a new price per stage instead of using only one price throughout all stages. Moreover, the set of prices we sample from per stage is updated dynamically. Sampling new prices per stage also introduces new challenges into the analysis. Nonetheless, we manage to show that:

THEOREM 4.2. PRIORFREEONLINE *is universally truthful and* $O(\log m_T)$*-competitive.*

Our analysis uses submodularity in both the second price auction and the random fixed-price auction. We consider two situations depending on whether a *dominant* buyer exists. Buyer i is a dominant buyer if $v_i\left(U^{(1,T)}\right) \ge \frac{v\left(\text{OPT}^{(1,T)}\right)}{10^4 \log m_T}$. When there exists a dominant buyer, it is easy to show that the welfare of running the second price auction is at least the valuation of the dominant buyer i^* over the entire bundle, i.e., $v_{i^*}\left(U^{(1,T)}\right)$, which immediately yields a $O(\log m_T)$ approximation.

LEMMA 4.1. *For a set V of submodular valuations, when there exists a dominant buyer, the second price auction yields welfare at least* $\frac{v\left(\text{OPT}^{(1,T)}\right)}{10^4 \log m_T}$.

Proof. Let buyer i^* be one of the dominant buyers. For convenience, let i_t be the buyer winning the second price auction at time t, and we have $v(\vec{S}^T) = \sum_i v_i(S_i^T) = \sum_t v_{i_t}(B^t \mid S_{i_t}^{t-1})$. In second price auctions, the marginal value of the winner is no less than that of any other buyer, so we have

$$(4.3)\quad v(\vec{S}^T) = \sum_t v_{i_t}(B^t \mid S_{i_t}^{t-1}) \ge \sum_t v_{i^*}(B^t \mid S_{i^*}^{t-1}) \ge \sum_t v_{i^*}(B^t \mid U^{t-1}) = v_{i^*}(U^{(1,T)}),$$

where the second inequality follows from submodularity of v_{i^*} and the fact that $S_{i^*}^{t-1} \subseteq U^{t-1}$. ∎

From now on, we assume there is no dominant buyer. To analyze the performance of our algorithm, we will use $v\left(\text{OPT}^{(t^*,T)}\right)$ as a benchmark where t^* is chosen from Lemma 4.2 such that $v\left(\text{OPT}^{(t^*,T)}\right) \ge \frac{v\left(\text{OPT}^{(1,T)}\right)}{2}$ and $v\left(\text{OPT}^{(1,t)}\right) \ge \frac{v\left(\text{OPT}^{(1,T)}\right)}{2}$ for all $t \ge t^*$. Such a choice is necessary, because intuitively, the initial stages are too sensitive for our analysis to work effectively.

LEMMA 4.2. $\exists t^*$ *s.t.* $v(\text{OPT}^{(t^*,T)}) \ge \frac{v(\text{OPT}^{(1,T)})}{2}$ *and* $v(\text{OPT}^{(1,t)}) \ge \frac{v(\text{OPT}^{(1,T)})}{2}$ *for all* $t \ge t^*$.

Proof. Let t^* be the earliest stage such that $v(\text{OPT}^{(1,t^*)}) \ge \frac{v(\text{OPT}^{(1,T)})}{2}$. Notice that $t^* \le T$. The first property follows immediately from the definition of t^*. For the second property, because of the optimality of $\text{OPT}^{(t^*,T)}$ and subadditivity of v_i,

$$v(\text{OPT}^{(t^*,T)}) \ge \sum_i v_i(\text{OPT}_i \cap U^{(t^*,T)}) \ge v(\text{OPT}^{(1,T)}) - \sum_i v_i(\text{OPT}_i \cap U^{(1,t^*-1)}).$$

Again, because of the optimality of $\text{OPT}^{(1,t^*-1)}$,

$$v(\text{OPT}^{(t^*,T)}) \ge v(\text{OPT}^{(1,T)}) - v(\text{OPT}^{(1,t^*-1)}) \ge \frac{1}{2}v(\text{OPT}^{(1,T)}).$$

∎

We consider the set $\mathcal{T} = \{(t_1, t_2) \mid t_1 \le t_2 \le T, v(\text{OPT}^{(t_1,t_2)}) \ge v(\text{OPT}^{(1,T)})/256\}$. Intuitively,

$(t_1, t_2) \in \mathcal{T}$ if the optimal welfare restricted to the items appearing between stages t_1 and t_2 is a constant fraction of the optimal welfare over all items. The next lemma shows that for any $(t_1, t_2) \in \mathcal{T}$, both $v(\mathrm{OPT}^{(t_1,t_2)}|_{\mathrm{STAT}})$ and $v(\mathrm{OPT}^{(t_1,t_2)}|_{\mathrm{MECH}})$ are a constant fraction of $v(\mathrm{OPT}^{(t_1,t_2)})$.

LEMMA 4.3. *For any $(t_1, t_2) \in \mathcal{T}$, with probability at least $1 - \frac{1}{m_T}$, we have*

$$\min\{v(\mathrm{OPT}^{(t_1,t_2)}|_{\mathrm{STAT}}), v(\mathrm{OPT}^{(t_1,t_2)}|_{\mathrm{MECH}})\} \geq \frac{1}{4} v(\mathrm{OPT}^{(t_1,t_2)}). \tag{4.4}$$

Proof. Fix $t_1 \leq t_2 \leq T$ where $(t_1, t_2) \in \mathcal{T}$. First note that for any i,

$$v_i(U^{(t_1,t_2)}) \leq v_i(U^{(1,T)}) \leq \frac{v(\mathrm{OPT}^{(1,T)})}{10^4 \log m_T} \leq \frac{v(\mathrm{OPT}^{(t_1,t_2)})}{24 \log m_T}.$$

where the second inequality follows the fact that there is no dominant buyer and the last inequality is due to the definition of $\mathcal{T}$. Let $X_i = \mathbf{1}[i \in \mathrm{STAT}]$. Observe that $\{X_i v_i(\mathrm{OPT}_i^{(t_1,t_2)})\}_i$ are independent random variables, where $X_i v_i(\mathrm{OPT}_i^{(t_1,t_2)})$ is in range $[0, v_i(\mathrm{OPT}_i^{(t_1,t_2)})] \subseteq [0, \frac{v(\mathrm{OPT}^{(t_1,t_2)})}{24 \log m_T}]$, and

$$\mathbb{E}\left[\sum_i X_i v_i(\mathrm{OPT}_i^{(t_1,t_2)})\right] = \frac{1}{2} v(\mathrm{OPT}^{(t_1,t_2)}).$$

Applying Hoeffding's inequality, we have

$$\begin{aligned}
&\Pr\left[\left|v(\mathrm{OPT}^{(t_1,t_2)}|_{\mathrm{MECH}}) - \mathbb{E}[v(\mathrm{OPT}^{(t_1,t_2)}|_{\mathrm{MECH}})]\right| \geq \frac{1}{4} v(\mathrm{OPT}^{(t_1,t_2)})\right] \\
= \ &\Pr\left[\left|\sum_i X_i v_i(\mathrm{OPT}_i^{t_1,t_2}) - \frac{1}{2} v(\mathrm{OPT}^{(t_1,t_2)})\right| \geq \frac{1}{4} v(\mathrm{OPT}^{(t_1,t_2)})\right] \\
\leq \ &\exp\left(-\frac{2 \cdot \left(\frac{1}{4} v(\mathrm{OPT}^{(t_1,t_2)})\right)^2}{\sum_i v_i(\mathrm{OPT}_i^{(t_1,t_2)})^2}\right) \\
\leq \ &\exp\left(-\frac{1}{8} \cdot \frac{v(\mathrm{OPT}^{(t_1,t_2)})^2}{24 \log m_T (v(\mathrm{OPT}^{(t_1,t_2)})^2 / (24 \log m_T))^2}\right) \\
= \ &\exp(-3 \log m_T) = \frac{1}{m_T^3}.
\end{aligned}$$

Exactly the same argument implies the same concentration for $v(\mathrm{OPT}^{(t_1,t_2)}|_{\mathrm{STAT}})$.

Observe that there are at most $\binom{T}{2} \leq \binom{m_T}{2} \leq \frac{1}{2} m_T^2$ pairs of t_1 and t_2 satisfying (4.4). Taking union bound over all such pairs and STAT and MECH, we have that (4.4) holds with probability at least $1 - \frac{1}{2} m_T^2 \cdot 2 \cdot \frac{1}{m_T^3} = 1 - \frac{1}{m_T}$. □

Note that Lemma 4.2 implies that (a) $(t^*, T) \in \mathcal{T}$, and (b) for any $t \geq t^*$, $(1, t) \in \mathcal{T}$. Therefore, Lemma 4.3 can be applied to all these intervals. The key lemma we are going to establish next is that with $\Omega(\frac{1}{\log m_T})$ probability, the item goes to the market with a desirable price constructed from additive valuation functions that represent the submodular valuations.

DEFINITION 4.1. ([8]) *A set V of valuations can be point-wise β-approximated by additive valuations if for any $v \in V$ and $S \subseteq U$, v can be point-wise β-approximated at S by an additive valuation v' such that $\beta \cdot v'(S) \geq v(S)$ and $\forall S' \subseteq U$, $v'(S') \leq v(S')$.*

It is well-known that submodular valuations are point-wise 1-approximated by additive valuations (see, e.g., [18]). However, an 1-approximated additive valuation v' is not enough for our analysis since the smallest non-zero entry $v_{\min} > 0$ could be arbitrarily small such that we can no longer guarantee the random price is within $[c_1 v_{\min}, c_2 v_{\min}]$ for some constants $0 < c_1 < c_2 < 1$ with $\Omega(\frac{1}{m_T})$ probability. To overcome this difficulty, we trim the additive valuations *in an online manner*; roughly speaking, our criteria for trimming each item become looser and looser as more items arrive. This is key for the trimming procedure to be compatible with the online environment.

LEMMA 4.4. *There exist additive valuations $(v'_1, \cdots, v'_n)$ such that:*

- *$v'_i(S) \leq v_i(S)$ for any buyer i and $S \subseteq U^{(1,T)}$;*
- *$v'(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}}) \geq v(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})/10 \geq v(\mathrm{OPT}^{(t^*,T)})/40$;*
- *if $j \notin \mathrm{OPT}_i^{(t^*,T)}|_{\mathrm{MECH}}$, $v'_i(j) = 0$;*
- *for $j \in \mathrm{OPT}_i^{(t^*,T)}|_{\mathrm{MECH}}$, if $v'_i(j) > 0$, then $v'_i(j) \geq v'(\mathrm{OPT}_i^{(t^*,T)}|_{\mathrm{MECH}})/(2j^2)$.*

Proof. Let v''_i be any additive valuation that point-wise 1-approximates v_i at $\mathrm{OPT}_i^{(t^*,T)}|_{\mathrm{MECH}}$. We construct v'_i such that $v'_i(j) = v''_i(j)$ if $j \in \mathrm{OPT}_i^{(t^*,T)}|_{\mathrm{MECH}}$ and $v''_i(j) \geq \frac{v''(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})}{2j^2}$, while all other entries are simply 0. Clearly v'_i satisfies the first, third and fourth properties in the lemma. For the second property, notice that we have

$$\begin{aligned}&v'(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})\\ &= \sum_{i\in \mathrm{MECH}} \sum_{j\in \mathrm{OPT}_i^{(t^*,T)}} v_i'(j)\\ &\geq \sum_{i\in \mathrm{MECH}} \sum_{j\in \mathrm{OPT}_i^{(t^*,T)}} \left(v_i''(j) - \frac{v''(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})}{2j^2}\right)\\ &\geq v''(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}}) - \sum_{1\leq j\leq m} \frac{v''(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})}{2j^2}.\end{aligned}$$

Recall that $\sum_{1\leq j\leq\infty} \frac{1}{j^2} = \frac{\pi^2}{6}$, and therefore, we have

$$\begin{aligned}&v'(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})\\ &\geq v''(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}}) - \frac{\pi^2 v''(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})}{12}\\ &\geq \frac{v(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})}{10}.\end{aligned}$$

□

We construct $(v_1', \cdots, v_n')$ satisfying the properties defined in Lemma 4.4. Let the supporting price of item j be $p_j = v_i'(j)$ for $j \in \mathrm{OPT}_i^{(t^*,T)}$. We say an item $j \in B^t$ is a *hit-item* if the random price p^t satisfies $\frac{1}{4}p_j \leq p^t \leq \frac{1}{2}p_j$.

Lemma 4.5. *For $j \in B^t$ with $p_j > 0$ and $t \geq t^*$, with probability $\Omega(\frac{1}{\log m_t})$, $\frac{1}{4}p_j \leq p^t \leq \frac{1}{2}p_j$.*

Proof. Observe that $|P^t| = O(\log m_t)$, so each price is chosen with probability $\Omega\left(\frac{1}{\log m_t}\right)$. It suffices to show that there exists $p \in P^t$ satisfying $\frac{1}{4}p_j \leq p \leq \frac{1}{2}p_j$, which is equivalent to showing that

$$\frac{\mathrm{est}^t}{c\cdot m_t^2} \leq \frac{v'(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})}{4m_t^2} \leq \frac{v'(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})}{4j^2} \leq \frac{1}{2}p_j, \tag{4.5}$$

where the second inequality follows the fact that $j \leq m_t$ and

$$c\cdot m_t^2 \cdot \mathrm{est}_t \geq \frac{v'(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})}{4} \geq \frac{1}{4}p_j. \tag{4.6}$$

Let $c = 2048$. For the first inequality in (4.5), since $\mathrm{est}_t \leq v(\mathrm{OPT}^{(1,T)})$ and $v(\mathrm{OPT}^{(1,T)}) \leq 2v(\mathrm{OPT}^{(t^*,T)})$ (Lemma 4.2),

$$\frac{\mathrm{est}_t}{2048m_t^2} \leq \frac{v(\mathrm{OPT}^{(1,T)})}{2048m_t^2} \leq \frac{v(\mathrm{OPT}^{(t^*,T)})}{1024m_t^2}.$$

Now by Lemma 4.3 and Lemma 4.4, we have $v(\mathrm{OPT}^{(t^*,T)}) \leq 4v(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})$ and $v(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}}) \leq 10v'(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})$, so

$$\begin{aligned}\frac{\mathrm{est}_t}{2048m_t^2} \leq \frac{v(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})}{256m_t^2} &\leq \frac{v'(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})}{25.6m_t^2}\\ &\leq \frac{v'(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})}{4m_t^2}.\end{aligned}$$

For the first inequality in (4.6), because $\mathrm{est}_t \geq v(\mathrm{OPT}^{(1,t)}|_{\mathrm{STAT}})$ and $v(\mathrm{OPT}^{(1,t)}|_{\mathrm{STAT}}) \geq \frac{1}{4}v(\mathrm{OPT}^{(1,t)})$ (Lemma 4.3),

$$\begin{aligned}2048m_t^2 \cdot \mathrm{est}_t &\geq 1024m_t^2 \cdot v(\mathrm{OPT}^{(1,t)}|_{\mathrm{STAT}})\\ &\geq 256m_t^2 \cdot v(\mathrm{OPT}^{(1,t)}).\end{aligned}$$

Again by Lemma 4.4, $v(\mathrm{OPT}^{(1,t)}) \geq \frac{1}{2}v(\mathrm{OPT}^{(1,T)})$, so

$$\begin{aligned}2048m_t^2 \cdot \mathrm{est}_t &\geq 128m_t^2 \cdot v(\mathrm{OPT}^{(1,T)})\\ &\geq \frac{v'(\mathrm{OPT}^{(t^*,T)}|_{\mathrm{MECH}})}{4}.\end{aligned}$$

□

Let G be the set of hit-items in $U^{(t^*,T)}$ and SOLD be the set of items that are sold. Notice that if $j \in G \cap \mathrm{SOLD}$, it contributes revenue p_j to the welfare. All that remains to show is that the buyers' utilities can capture the welfare generated by the unsold items. While this is quite well-understood in static environments, in the online environment that we consider, one additional difficulty is to summarize the contribution of unsold items *over stages*. Moreover, in light of our impossibility results, for any such summarization argument to be useful, it must apply only to submodular valuations. Below we present such an argument. Recall that S_i^T is a set of items allocated to buyer i at the end of stage T.

Lemma 4.6. $\sum_i v_i(S_i^T) \geq \frac{1}{2}\sum_{G\setminus\mathrm{SOLD}} p_j$.

Proof. Let $\Gamma_i = (\mathrm{OPT}_i^{(t^*,T)} \cap G) \setminus \mathrm{SOLD}$. Consider the telescoping sum of $v_i(\Gamma_i \mid S_i^T)$:

$$v_i\left(\Gamma_i \mid S_i^T\right) = \sum_{t\geq t^*} v_i\left(\Gamma_i \cap B^t \mid \Gamma_i \cap U^{(1,t-1)} \cup S_i^T\right).$$

Because of the submodularity of v_i and the fact that $S_i^t \subseteq S_i^T \subseteq \Gamma_i \cap U^{(1,t-1)} \cup S_i^T$,

$$v_i\left(\Gamma_i \mid S_i^T\right) \leq \sum_{t\geq t^*} v_i\left(\Gamma_i \cap B^t \mid S_i^t\right). \tag{4.7}$$

Now consider the behavior of buyer i at time t. After buying $S_i^t \setminus S_i^{t-1}$ at stage t, buyer i could still choose to buy $\Gamma_i \cap B^t$, which would give her $v_i\left(\Gamma_i \cap B^t \mid S_i^t\right)$ value with payment $\sum_{j\in\Gamma_i\cap B^t} p^t$. The only reason that buyer i does not do so is that her marginal gain is at most her payment, i.e., $v_i\left(\Gamma_i \cap B^t \mid S_i^t\right) \le \sum_{j\in\Gamma_i\cap B^t} p^t$. So given that $p^t \le p_j/2$ for $j \in G$, we have

$$v_i\left(\Gamma_i \mid S_i^T\right) \le \sum_{t\ge t^*} v_i\left(\Gamma_i \cap B^t \mid S_i^t\right)$$
$$\le \sum_{t\ge t^*} \sum_{j\in\Gamma_i\cap B^t} p^t \le \frac{1}{2}\sum_{j\in\Gamma_i} p_j = \frac{1}{2}v_i'(\Gamma_i).$$

Therefore, buyer i's value is $v_i\left(S_i^T\right) = v_i\left(S_i^T \cup \Gamma_i\right) - v_i\left(\Gamma_i \mid S_i^T\right) \ge v_i(\Gamma_i) - \frac{1}{2}v_i'(\Gamma_i) \ge \frac{1}{2}\sum_{j\in\Gamma_i} p_j$. □

We can now proceed to prove Theorem 4.2 by combining the contributions from items that are sold and items that are not sold. Note that our proof breaks for XOS valuations since we use the property of submodularity in (4.3) for the welfare guarantee of the second price auction (Lemma 4.1), and in (4.7) for our revenue-utility decomposition argument for submodular valuations (Lemma 4.6).

5 Bayesian Online Mechanism for Submodular Valuations

In this section, we extend our results to a Bayesian setting, where the buyers' valuations are drawn independently from prior distributions that are common knowledge.

We design a computationally efficient and universally truthful mechanism which guarantees $\frac{1}{8}$ of the the optimal welfare in expectation. Compared to the static setting [20], the first challenge is to establish a benchmark and its corresponding supporting prices in an online environment. Moreover, such a benchmark must be *stable*, in the sense that, roughly speaking, as soon as an item arrives, the benchmark restricted to this item can be calculated immediately, and is no longer affected by any future items. To tackle this difficulty, we first establish an *offline benchmark* that guarantees a constant fraction of the welfare produced by the optimal offline allocation algorithm. Our benchmark is inspired by the greedy algorithm to online optimization for submodular valuations [21, 23]. We then show that our online truthful mechanism can approximate the offline benchmark with a constant ratio.

Offline benchmark and supporting prices. For an offline allocation algorithm $\mathcal{A}$, let the allocation that $\mathcal{A}$ generates with items U and valuations $\vec{v}$ as input be $\mathcal{A}(U,\vec{v}) = (\mathcal{A}_1(U,\vec{v}),\dots,\mathcal{A}_n(U,\vec{v}))$. Consider the greedy allocation rule $\mathcal{A}$ defined inductively, as follows.

$$\mathcal{A}_i\big([j],\vec{v}\big) = \begin{cases} \mathcal{A}_i\big([j-1],\vec{v}\big) \cup \{j\}, \\ \qquad \text{if } i = \operatorname{argmax}_{i'} v_{i'}\Big(\{j\} \mid \mathcal{A}_{i'}\big([j-1],\vec{v}\big)\Big) \\ \mathcal{A}_i\big([j-1],\vec{v}\big), \\ \qquad \text{otherwise} \end{cases}$$

where $[j] = \{1,\cdots,j\}$ and ties are broken in any consistent manner. In other words, $\mathcal{A}$ allocates items in a greedy manner such that item j is allocated to the buyer with the largest marginal value for item j. It is known that $\mathcal{A}$ always produces a 2-approximation of the optimal offline allocation.

Lemma 5.1. ([21, 23]) *For any U and submodular valuations $\vec{v} \in V^n$, $v\big(\mathcal{A}(U,\vec{v})\big) \ge \frac{1}{2}v\big(\mathrm{OPT}(U,\vec{v})\big)$.*

We define the supporting prices with respect to $\vec{v}$ for item j from the greedy allocation $\mathcal{A}$:

$$\mathrm{SP}_j\left(\mathcal{A}(U^{(1,t)},\vec{v}),\vec{v}\right) = \sum_i \mathbb{I}\left[j \in \mathcal{A}_i\big(U^{(1,t)},\vec{v}\big)\right] \cdot v_i\Big(\{j\} \mid \mathcal{A}_i\big([j-1],\vec{v}\big)\Big).$$

That is, SP_j is the marginal value of j for the buyer who receives j according to the greedy allocation $\mathcal{A}$. These prices *support* the welfare generated by $\mathcal{A}$ in the following sense:

Lemma 5.2. *The supporting prices $\mathrm{SP}_j\left(\mathcal{A}(U^{(1,T)},\vec{v}),\vec{v}\right)$ satisfy for any buyer i,*

$$\sum_{j\in\mathcal{A}_i\left(U^{(1,T)},\vec{v}\right)} \mathrm{SP}_j\left(\mathcal{A}(U^{(1,T)},\vec{v}),\vec{v}\right) = v_i\Big(\mathcal{A}_i\big(U^{(1,T)},\vec{v}\big)\Big).$$

Moreover, for any buyer i and $S \subseteq \mathcal{A}_i\big(U^{(1,T)},\vec{v}\big)$, we have $\sum_{j\in S}\mathrm{SP}_j\left(\mathcal{A}(U^{(1,T)},\vec{v}),\vec{v}\right) \le v_i(S)$.

In the greedy allocation algorithm $\mathcal{A}$, for $j \in B^t$, $\mathbb{I}\left[j \in \mathcal{A}_i\big(U^{(1,T)},\vec{v}\big)\right] = \mathbb{I}\left[j \in \mathcal{A}_i\big(U^{(1,t)},\vec{v}\big)\right]$, which implies that $\mathrm{SP}_j\left(\mathcal{A}\big(U^{(1,t)},\vec{v}\big),\vec{v}\right) = \mathrm{SP}_j\left(\mathcal{A}\big(U^{(1,T)},\vec{v}\big),\vec{v}\right)$. For each item $j \in B^t$, we will set p_j to be half of the expectation of $\mathrm{SP}_j\left(\mathcal{A}\big(U^{(1,t)},\vec{v}\big),\vec{v}\right)$, i.e., $p_j = \frac{1}{2}\cdot\mathbb{E}_{\vec{v}}\left[\mathrm{SP}_j\left(\mathcal{A}\big(U^{(1,t)},\vec{v}\big),\vec{v}\right)\right]$. We are now ready to present our online mechanism BayesianOnline for the Bayesian environment (Algorithm 5.1). We emphasize that it is crucial to approach the buyers in the same ordering for all stages and we choose the natural order $\{1,\cdots,n\}$ for ease of presentation.

ALGORITHM 5.1. (BAYESIANONLINE)

for each stage t **do**
  Let $M = B^t$
  for each buyer i in the ordering of $\{1, \cdots, n\}$ **do**
    Let $D_i = \arg\max_{S \in M} v_i(S \mid S_i^{t-1}) - \sum_{j \in S} p_j$
    Allocate D_i to buyer i ($S_i^t = S_i^{t-1} \cup D_i$) and charge her $\sum_{j \in D_i} p_j$
    $M = M \setminus D_i$
  end for
end for

Our mechanism is truthful since for each stage, the posted-price auction is truthful. To implement the mechanism, notice that for $j \in B^t$, $p_j = \frac{1}{2} \cdot \mathbb{E}_{\vec{v}}\left[\mathrm{SP}_j\left(\mathcal{A}(U^{(1,t)}, \vec{v}), \vec{v}\right)\right]$ only depends on $F_1^t, \cdots, F_n^t$, which are already known to the mechanism upon the arrival of item j.

THEOREM 5.1. BAYESIANONLINE *is universally truthful and* 8*-competitive.*

For convenience, let $\vec{S}^T(\vec{v})$ represent the allocation after stage T by our online mechanism BAYESIANONLINE when the realized valuation profile is $\vec{v}$. To prove Theorem 5.1, we first generalize our revenue-utility decomposition argument (Lemma 4.6) to the Bayesian setting, allowing for a summarization of the contributions from items over stages. Fix a buyer i's valuation v_i and consider two arbitrary valuation profiles $\vec{v} = (v_i, \vec{v}_{-i})$ and $\vec{v}' = (v_i, \vec{v}'_{-i})$. Let $W_i(\vec{v}, \vec{v}') = \mathcal{A}_i(U^{(1,T)}, \vec{v}') \cap S_i^T(\vec{v})$ and $Y_i(\vec{v}, \vec{v}') = \mathcal{A}_i(U^{(1,T)}, \vec{v}') \setminus \left(\bigcup_{i' \leq i} S_{i'}^T(\vec{v})\right)$.

LEMMA 5.3. *For any buyer i and two valuation profiles $\vec{v} = (v_i, \vec{v}_{-i})$ and $\vec{v}' = (v_i, \vec{v}'_{-i})$,*

$$v_i\left(S_i^T(\vec{v})\right) \geq v_i\Big(W_i(\vec{v}, \vec{v}') \cup Y_i(\vec{v}, \vec{v}')\Big) - \sum_{j \in Y_i(\vec{v}, \vec{v}')} p_j.$$

Proof. For ease of presentation, let $S_i^T = S_i^T(\vec{v})$, $W_i = W_i(\vec{v}, \vec{v}')$, and $Y_i = Y_i(\vec{v}, \vec{v}')$. Then, we have

$$\begin{aligned} v_i\left(S_i^T\right) &= v_i\left(S_i^T \cup Y_i\right) - v_i\left(Y_i \mid S_i^T\right) \\ &= v_i\left(S_i^T \cup Y_i\right) \\ &\quad - \sum_t v_i\left(Y_i \cap B^t \mid S_i^T \cup \left(Y_i \cap U^{(1,t-1)}\right)\right) \end{aligned}$$

where the last inequality is due to the telescoping sum representation of $v_i\left(Y_i \mid S_i^T\right)$. Notice that $S_i^t \subseteq S_i^T \subseteq S_i^T \cup \left(Y_i \cap U^{(1,t-1)}\right)$ and we have

$$v_i\left(S_i^T\right) \geq v_i\left(S_i^T \cup Y_i\right) - \sum_t v_i\left(Y_i \cap B^t \mid S_i^t\right), \tag{5.8}$$

where the inequality follows submodularity. Since buyer i did not purchase bundle $Y_i \cap B^t$ when she has already purchased S_i^t, the price for acquiring $Y_i \cap B^t$ must be larger than her marginal value. Therefore, we have $v_i\left(S_i^T\right) \geq v_i\left(S_i^T \cup Y_i\right) - \sum_t \sum_{j \in Y_i \cap B^t} p_j$. We finish the proof by noticing that $\sum_t \sum_{j \in Y_i \cap B^t} p_j = \sum_{j \in Y_i} p_j$ and $v_i\left(S_i^T \cup Y_i\right) \geq v_i\left(W_i \cup Y_i\right)$ since $W_i \subseteq S_i^T$. ◻

We are ready to prove Theorem 5.1 by noticing that Lemma 5.3 implies $v_i\left(S_i^T(\vec{v})\right)$ can be lower bounded as $v_i\left(S_i^T(\vec{v})\right) \geq \mathbb{E}_{\vec{v}'_{-i}}\left[v_i\Big(W_i(\vec{v}, \vec{v}') \cup Y_i(\vec{v}, \vec{v}')\Big) - \sum_{j \in Y_i(\vec{v}, \vec{v}')} p_j\right]$, where $\vec{v}'_{-i} \sim \prod_{i' \neq i} F_i^T$.

Proof. [Proof of Theorem 5.1] Let $\mathrm{SOLD}(\vec{v})$ be the set of items that are sold in BAYESIANONLINE when the realized valuation profile is $\vec{v}$. Moreover, let $\mathrm{SOLD}_i(\vec{v}_{-i})$ be the set of items that are sold to some buyer in $[i-1]$ when the realized valuation profile for other buyers is $\vec{v}_{-i}$. Notice that if item j is sold, then item j contributes revenue p_j to the welfare. Therefore, we have

$$\begin{aligned} \mathop{\mathbb{E}}_{\vec{v}}\left[v\left(\vec{S}^T(\vec{v})\right)\right] &\geq \sum_j \Pr_{\vec{v}}\left[j \in \mathrm{SOLD}(\vec{v})\right] \cdot p_j \\ &= \frac{1}{2} \sum_j \Pr_{\vec{v}}\left[j \in \mathrm{SOLD}(\vec{v})\right] \cdot \mathop{\mathbb{E}}_{\vec{v}}\left[\mathrm{SP}_j\left(\mathcal{A}(U^{(1,T)}, \vec{v}), \vec{v}\right)\right]. \end{aligned}$$

We now proceed to show that,

$$\begin{aligned} \mathop{\mathbb{E}}_{\vec{v}}\left[v\left(\vec{S}^T(\vec{v})\right)\right] \geq \frac{1}{2} \sum_j \Pr_{\vec{v}}\left[j \notin \mathrm{SOLD}(\vec{v})\right] \\ \cdot \mathop{\mathbb{E}}_{\vec{v}}\left[\mathrm{SP}_j\left(\mathcal{A}(U^{(1,T)}, \vec{v}), \vec{v}\right)\right]. \end{aligned}$$

From Lemma 5.3, we have that $\mathbb{E}_{\vec{v}}\left[v_i\left(S_i^T(\vec{v})\right)\right]$ is at least:

$$\mathop{\mathbb{E}}_{v_i,\vec{v}_{-i},\vec{v}'_{-i}}\left[v_i\Big(W_i(\vec{v},\vec{v}')\cup Y_i(\vec{v},\vec{v}')\Big) - \sum_j \mathbb{I}\big[j\in Y_i(\vec{v},\vec{v}')\big]\cdot p_j\right]$$

$$\geq \mathop{\mathbb{E}}_{v_i,\vec{v}_{-i},\vec{v}'_{-i}}\left[\sum_j \mathrm{SP}_j\Big(\mathcal{A}\big(U^{(1,T)},v_i,\vec{v}'_{-i}\big),v_i,\vec{v}'_{-i}\big)\Big) \cdot \mathbb{I}\Big[j\notin \mathrm{SOLD}_i(\vec{v}_{-i})\Big]\cdot \mathbb{I}\Big[j\in\mathcal{A}_i\big(U^{(1,T)},v_i,\vec{v}'_{-i}\big)\Big] - \sum_j \mathbb{I}\Big[j\in\mathcal{A}_i\big(U^{(1,T)},v_i,\vec{v}'_{-i}\big)\Big] \cdot \mathbb{I}\Big[j\notin \mathrm{SOLD}_i(\vec{v}_{-i})\Big]\cdot p_j\right]$$

$$= \sum_j \Pr_{\vec{v}_{-i}}\Big[j\notin \mathrm{SOLD}_i(\vec{v}_{-i})\Big] \cdot \mathop{\mathbb{E}}_{v_i,\vec{v}'_{-i}}\left[\mathbb{I}\Big[j\in\mathcal{A}_i\big(U^{(1,T)},v_i,\vec{v}'_{-i}\big)\Big] \cdot \Big(\mathrm{SP}_j\big(\mathcal{A}(U^{(1,T)},v_i,\vec{v}'_{-i}),v_i,\vec{v}'_{-i}\big)-p_j\Big)\right]$$

$$\geq \sum_j \Pr_{\vec{v}}\Big[j\notin \mathrm{SOLD}(\vec{v})\Big] \cdot \mathop{\mathbb{E}}_{\vec{v}}\left[\mathbb{I}\Big[j\in\mathcal{A}_i\big(U^{(1,T)},\vec{v}\big)\Big] \cdot \Big(\mathrm{SP}_j\big(\mathcal{A}(U^{(1,T)},\vec{v}),\vec{v}\big)-p_j\Big)\right].$$

where the first inequality follows that $W_i(\vec{v},\vec{v}')\cup Y_i(\vec{v},\vec{v}') = \mathcal{A}_i(U^{(1,T)},\vec{v}')\setminus\big(\bigcup_{i'<i}S_{i'}(\vec{v})\big)$, Lemma 5.2, and $W_i(\vec{v},\vec{v}')\cup Y_i(\vec{v},\vec{v}')\supseteq Y_i(\vec{v},\vec{v}')$, while the equality follows the independence between $\mathrm{SOLD}_i(\vec{v}_{-i})$ and $(v_i,\vec{v}'_{-i})$. Particularly, the independence between $\mathrm{SOLD}_i(\vec{v}_{-i})$ and $(v_i,\vec{v}'_{-i})$ is established from the fact that the mechanism approaches the buyers in the same ordering for all stages. This is the reason why approaching the buyers in the same ordering for all stages is crucial.

Summing over i, we have $\sum_i \mathbb{E}_{\vec{v}}\left[v_i\left(S_i^T(\vec{v})\right)\right]$ is at least

$$\sum_j \Pr_{\vec{v}}\Big[j\notin \mathrm{SOLD}(\vec{v})\Big] \cdot \mathop{\mathbb{E}}_{\vec{v}}\left[\sum_i \mathbb{I}\Big[j\in\mathcal{A}_i\big(U^{(1,T)},\vec{v}\big)\Big] \cdot \Big(\mathrm{SP}_j\big(\mathcal{A}(U^{(1,T)},\vec{v}),\vec{v}\big)-p_j\Big)\right]$$

$$= \sum_j \Pr_{\vec{v}}\Big[j\notin \mathrm{SOLD}(\vec{v})\Big] \cdot \mathop{\mathbb{E}}_{\vec{v}}\Big[\mathrm{SP}_j\Big(\mathcal{A}\big(U^{(1,T)},\vec{v}\big),\vec{v}\Big)-p_j\Big]$$

$$= \frac{1}{2}\sum_j \Pr_{\vec{v}}\Big[j\notin \mathrm{SOLD}(\vec{v})\Big] \cdot \mathop{\mathbb{E}}_{\vec{v}}\Big[\mathrm{SP}_j\Big(\mathcal{A}\big(U^{(1,T)},\vec{v}\big),\vec{v}\Big)\Big].$$

Putting the two parts for sold items and unsold items together, we have

$$\mathop{\mathbb{E}}_{\vec{v}}\left[v\left(\vec{S}^T(\vec{v})\right)\right] \geq \frac{1}{4}\sum_j \mathop{\mathbb{E}}_{\vec{v}}\Big[\mathrm{SP}_j\Big(\mathcal{A}\big(U^{(1,T)},\vec{v}\big)\Big)\Big] = \frac{1}{4}\mathop{\mathbb{E}}_{\vec{v}}\Big[v\Big(\mathcal{A}\big(U^{(1,T)},\vec{v}\big)\Big)\Big] \geq \frac{1}{8}\mathop{\mathbb{E}}_{\vec{v}}\Big[v\Big(\mathrm{OPT}\big(U^{(1,T)}\big)\Big)\Big]$$

where the equality follows Lemma 5.2 and the last inequality follows Lemma 5.1. ◻

Note that our proof breaks for XOS valuations since our offline benchmark highly relies on submodular valuations (Lemma 5.1 and Lemma 5.2) and we use the property of submodularity in (5.8) for our revenue-utility decomposition for submodular valuations (Lemma 5.3).

6 Online Mechanisms with No Expiration Date

In this section, we describe our reduction from the setting with no expiration date to the classical offline environment. The only condition required by the reduction is that the offline mechanism needs to be *approximately monotone*, which roughly says that if we give buyers some items before the mechanism starts, then the (expected) welfare after running the mechanism is not much smaller than the welfare from running the mechanism without the initial items. This condition holds for most, if not all, existing mechanisms for subadditive (including XOS) buyers. As long as the condition holds, our reduction preserves the approximation ratio of the offline mechanism up to a constant factor in the no expiration environment.

6.1 The Reduction We first state the requirement of our reduction.

Definition 6.1. (Approximate Monotonicity) *A truthful mechanism $\mathcal{M}$, which maps a set of items U*

and valuations $\vec{v}$ to a (randomized) allocation $\mathcal{M}(U, \vec{v})$ is approximately monotone for a class $\mathcal{V}$ of valuations, if there exists a constant $C > 0$, such that for any $\vec{U}^0 = \{U_1^0, \cdots, U_n^0\}$ and U with $U \cap U_i^0 = \emptyset$ for all i, and $\vec{v} \in \mathcal{V}^n$ where the domain of v_i is over $\bigcup_{i=1}^n U_i^0 \cup U$,

$$\mathbb{E}\left[\sum_i v_i\Big(\mathcal{M}_i(U, \vec{v})\Big)\right] \leq C \cdot \mathbb{E}\left[\sum_i v_i\Big(U_i^0 \cup \mathcal{M}_i(U, v_{i|U_i^0})\Big)\right].$$

where $v_{i|U_i^0}(S) = v_i(S \cup U_i^0) - v_i(U_i^0)$ for all i and $S \subseteq U$.

As discussed above, the condition can be interpreted as giving items U_i^0 to buyer i for free before running $\mathcal{M}$ does not hurt the expected welfare more than a multiplicative constant factor. While this interpretation makes the condition appear trivially true, we emphasize that it is technically non-trivial to prove an offline mechanism is approximately monotone since $v_{i|U_i^0}$ is not necessarily a member of the class $\mathcal{V}$, e.g., when $\mathcal{V}$ is the class of XOS valuations. We show that fortunately, most existing mechanisms, including those inducing the desired competitive ratios in the literature, are in fact approximately monotone. The proofs are deferred to Section 6.3.

6.1.1 Prior-free setting

We are now ready to give our reduction in the prior-free setting. Before the first stage, the seller flips a fair coin. If the result is heads, she implements *the second price auction with reserve*; otherwise if the result is tails, she runs an estimation scheme and makes repeated calls upon the offline mechanism.

Second Price Auction with Reserve At stage t, we will run a second price auction for the entire bundle of unsold items $U^{(1,t)} \setminus A^{t-1}$ with a reserve price set to the total welfare of the allocated items, i.e., $v(\vec{S}^{t-1})$. Notice that the second price auction is deterministic and dominant-strategy incentive-compatible, and thus, at stage t, each buyer i will truthfully report her marginal valuation over the bundle: $v_i(U^{(1,t)} \setminus A^{t-1} \mid S_i^{t-1})$. Once the bundle is sold to buyer i at stage t, the seller can update the total welfare of the allocated items to $v(\vec{S}^t) = v(\vec{S}^{t-1}) + v_i(U^{(1,t)} \setminus A^{t-1} \mid S_i^{t-1})$.

The Estimation Scheme We first divide the buyers into two groups, STAT and MECH, where the group for each buyer is chosen independently and uniformly at random. For each stage t, we maintain an estimate est_t of the optimal welfare of allocating $U^{(1,t)}$ to agents in STAT, using the 2-approximation algorithm for subadditive valuations by Feige [18] (see Theorem 4.1). The estimation scheme works in the following way:

- Initialize $k = 0$, $\text{est}_0 = 0$, and $t_0 = 0$.
- At each time t, compute est_t. If $\text{est}_t \geq 8\text{est}_{t_k}$, set $k = k + 1$ and $t_k = t$; call the offline mechanism with items $U^{(t_{k-1}+1,t)}$ and buyers MECH.

In words, we implement the offline mechanism on a new batch of items when the current estimate est_t is at least 8 times the estimate when the previous allocation happened. Intuitively, this guarantees that there is high enough welfare to be allocated in the current batch of items and at the same time, the welfare loss is low, if the market terminates with these items unallocated.

6.1.2 Bayesian setting

For the Bayesian setting, there is no need to flip a coin to implement two different mechanisms. In fact, it suffices to implement the estimation scheme only:

The Estimation Scheme For each stage t, we compute an estimate est_t of the expected optimal welfare of allocating $U^{(1,t)}$ from the prior. The estimation scheme works in the following way:

- Initialize $k = 0$, $\text{est}_0 = 0$, and $t_0 = 0$.
- At each time t, compute est_t. If $\text{est}_t \geq 2\text{est}_{t_k}$, set $k = k + 1$ and $t_k = t$; call the offline mechanism with items $U^{(t_{k-1}+1,t)}$, and all the buyers.

In words, we implement the offline mechanism on a new batch of items when the current estimate est_t is at least twice the estimate when the previous allocation happened.

6.2 Analysis

In the prior-free setting, the second price auction with reserve is clearly incentive-compatible. As for the estimation scheme in the prior-free setting, observe that the scheme queries only buyers in STAT, who get no items whatsoever, and therefore will answer all queries truthfully. On the other hand, the offline mechanism interacts only with buyers in MECH, whom the estimation scheme does not query at all. These buyers, being myopic, will act truthfully as long as the offline mechanism itself is truthful. Therefore, the estimation scheme is also incentive-compatible, and thus, the combination of these two subroutines is universally truthful. A similar argument can demonstrate that the estimation scheme in the Bayesian setting is also universally truthful.

The following theorem, which is the main result of this section, translates approximation guarantees in the classical offline setting to the no expiration environment. The only requirement, as stated above, is that the offline mechanism must be approximately monotone, which is indeed satisfied by almost all existing results, and in particular, by the state-of-the-art mechanisms for subadditive and XOS buyers, respectively,

THEOREM 6.1. *For a set $\mathcal{V}$ of complement-free valuations, suppose there exists a truthful offline $\beta(m_T)$-approximate mechanism with m_T items and the offline mechanism is approximately monotone. Then, there exists a truthful online $O(\beta(m_T))$-competitive mechanism for the no expiration environment.*

The rest of this section is devoted to providing high-level proof ideas for Theorem 6.1. The proof for the Bayesian setting follows the fact that the state-of-art mechanisms are approximately monotone and an argument presented below for the estimation scheme similar to the prior-free setting. We will focus on the prior-free setting from now on. Note that it suffices to show that for a fixed end of horizon T, the expected welfare generated by our reduction is at least $\Omega\left(\frac{1}{\beta(m_T)}\right)$ fraction of the welfare of the optimal allocation. Our analysis is divided into two parts depending on whether a *dominant* buyer exists: buyer i is dominant if $v_i(U^{(1,T)}) \geq \frac{v\left(\mathrm{OPT}(U^{(1,T)})\right)}{10^4}$.

LEMMA 6.1. *When there exists a dominant buyer, the second price auction with reserve guarantees $\Omega(1)$ fraction of the optimal welfare.*

Proof. Let i^* be a dominant buyer. For convenience, let t_k be the k-th stage in which the bundle is sold in the auction, i.e., there exists a buyer i such that $v_i(U^{(1,t)} \setminus A^{t-1}|S_i^{t-1}) \geq v(\vec{S}^{t-1})$ for $t = t_k$. We show inductively that for every t_k, the welfare $v(\vec{S}^{t_k})$ satisfies $v(\vec{S}^{t_k}) \geq \frac{v_{i^*}(U^{(1,t_k)})}{2}$.

By our definition of t_k, we have $A^{t_k} = U^{(1,t_k)}$ and $A^{t_k-1} = U^{(1,t_{k-1})}$. Assume that at t_k, the bundle $U^{(t_{k-1}+1,t_k)}$ is allocated to agent i_k. Since the bundle is sold, we have

(6.9)
$$v_{i_k}(U^{(t_{k-1}+1,t_k)} \mid S_{i_k}^{t_{k-1}}) \geq v(\vec{S}^{t_{k-1}}) \geq v_{i^*}(S_{i^*}^{t_{k-1}}).$$

On the other hand, since buyer i_k wins the second price auction, her bid must be at least the bid submitted by buyer i^*:

(6.10)
$$v_{i_k}(U^{(t_{k-1}+1,t_k)} \mid S_{i_k}^{t_{k-1}}) \geq v_{i^*}(U^{(t_{k-1}+1,t_k)} \mid S_{i^*}^{t_{k-1}}).$$

Combining (6.9) and (6.10), we have

(6.11)
$$\begin{aligned}2v_{i_k}(U^{(t_{k-1}+1,t_k)} \mid S_{i_k}^{t_{k-1}}) &\geq v_{i^*}(U^{(t_{k-1}+1,t_k)}) \\ &\geq v_{i^*}(U^{(t_{k-1}+1,t_k)} \mid U^{(1,t_{k-1})}),\end{aligned}$$

where the last inequality is by subadditivity. Therefore, we have:

$$\begin{aligned}v(\vec{S}^{t_k}) &= v(\vec{S}^{t_{k-1}}) + v_{i_k}(U^{(t_{k-1}+1,t_k)} \mid S_{i_k}^{t_{k-1}}) \\ &\text{(induction hypothesis and (6.11))} \\ &\geq \frac{1}{2}v_{i^*}(U^{(1,t_{k-1})}) + \frac{1}{2}v_{i^*}(U^{(t_{k-1}+1,t_k)} \mid U^{(1,t_{k-1})}) \\ &= \frac{1}{2}v_{i^*}(U^{(1,t_k)}).\end{aligned}$$

We will finish our proof by showing that at stage T,

$$v(\vec{S}^T) \geq \frac{1}{5}v_{i^*}(U^{(1,T)}) \geq \frac{1}{5} \cdot \frac{v\left(\mathrm{OPT}(U^{(1,T)})\right)}{10^4}.$$

Let t_K is the last stage in which the bundle is sold in the auction. If $t_K = T$, then the above inequality from induction implies

$$v(\vec{S}^T) = v(\vec{S}^{t_K}) \geq \frac{1}{2}v_{i^*}(U^{(1,t_K)}) = \frac{1}{2}v_{i^*}(U^{(1,T)}).$$

Otherwise, at time T no item is allocated. Using the induction hypothesis and the property of a second-price auction with reserve, we have that

$$v(\vec{S}^T) \geq \max\left\{\frac{1}{2}v_{i^*}(U^{(1,t_K)}), v_{i^*}(U^{(t_K+1,T)} \mid S_{i^*}^{t_K})\right\}.$$

Therefore,

$$\begin{aligned}v(\vec{S}^T) &\geq \frac{4}{5} \cdot \frac{1}{2}v_{i^*}(U^{(1,t_K)}) + \frac{1}{5}v_{i^*}(U^{(t_K+1,T)} \mid S_{i^*}^{t_K}) \\ &= \frac{1}{5}v_{i^*}(U^{(1,t_K)}) + \frac{1}{5}v_{i^*}(S_{i^*}^{t_K}) + \frac{1}{5}v_{i^*}(U^{(1,t_K)} \mid S_{i^*}^{t_K}) \\ &\quad + \frac{1}{5}v_{i^*}(U^{(t_K+1,T)} \mid S_{i^*}^{t_K}).\end{aligned}$$

Using monotonicity of the valuation functions, we get

$$\begin{aligned}v(\vec{S}^T) &\geq \frac{1}{5}v_{i^*}(U^{(1,t_K)}) + \frac{1}{5}v_{i^*}(S_{i^*}^{t_K}) \\ &\quad + \frac{1}{5}v_{i^*}(U^{(t_K+1,T)} \mid S_{i^*}^{t_K}) \\ &= \frac{1}{5}v_{i^*}(U^{(1,t_K)}) + \frac{1}{5}v_{i^*}(U^{(t_K+1,T)}) \\ &\geq \frac{1}{5}v_{i^*}(U^{(1,T)})\end{aligned}$$

where the last inequality follows the subadditivity of the valuation functions. □

From now on, we will focus on the case in which there is no dominant buyer. Moreover, let K be the final value of k at the end of the estimation scheme, which is the number of calls made to the offline mechanism. The estimation scheme divides items into batches, and runs one auction for each batch. The approximation guarantee of the offline mechanism then applies with respect to the welfare supported by these individual batches. We first need one of these batches to be large enough to support a *constant* fraction of the welfare given by the offline optimal allocation. To this end, we consider batches which overlap the time interval $[T_1, T_2]$, on which the optimal welfare from the prefix of items $U^{(1,t)}$ for $t \in [T_1, T_2]$ grows from $\frac{v(\mathrm{OPT}^{(1,T)})}{1000}$ to $\frac{v(\mathrm{OPT}^{(1,T)})}{100}$. Subadditivity guarantees that the optimal welfare from $U^{(T_1,T_2)}$ is a constant fraction of $v(\mathrm{OPT}^{(1,T)})$. By standard concentration bounds, this welfare is distributed almost equally into STAT and MECH. As a result, est_{T_1} and est_{T_2} are within a constant factor of each other, and there are only a constant number of batches overlapping $[T_1, T_2]$, since est can only increase so much. Thus, the largest batch among these provides a constant fraction of $v(\mathrm{OPT}^{(1,T)})$ to buyers in MECH:

LEMMA 6.2. *Suppose there is no dominant agent, i.e., for any agent i, $v_i(U^{(1,T)}) < \frac{v(\mathrm{OPT}^{(1,T)})}{10^4}$, and then with constant probability, there is some k such that $v(\mathrm{OPT}_{\mathrm{MECH}}^{(t_{k-1}+1,t_k)}) = \Theta(v(\mathrm{OPT}^{(1,T)}))$, so that the batch supports enough welfare.*

Proof. Let

$$T_1 = \min\{t \mid v(\mathrm{OPT}^{(1,t)}) \geq v(\mathrm{OPT}^{(1,T)})/1000\},$$

and

$$T_2 = \min\{t \mid v(\mathrm{OPT}^{(1,t)}) \geq v(\mathrm{OPT}^{(1,T)})/100\}.$$

Observe that $T_1 \leq T_2 \leq T$, and it is possible that $T_1 = T_2$ or $T_2 = T$ or $T_1 = T_2 = T$. Also, by the Hoeffding bound, with constant probability, simultaneously for all $t \in \{T_1, T_2, T\}$,

$$\begin{aligned} &0.2v(\mathrm{OPT}^{(1,t)}) \\ \leq\ &\min(v(\mathrm{OPT}^{(1,t)}|_{\mathrm{STAT}}), v(\mathrm{OPT}^{(1,t)}|_{\mathrm{MECH}})) \\ \leq\ &\max(v(\mathrm{OPT}^{(1,t)}|_{\mathrm{STAT}}), v(\mathrm{OPT}^{(1,t)}|_{\mathrm{MECH}})) \\ \leq\ &0.8v(\mathrm{OPT}^{(1,t)}). \end{aligned}$$

We condition on this from now on.

Note that there are only a constant number of batches ending between T_1 and T, inclusively. This is simply because $\mathrm{est}^{T_1} = \Omega(v(\mathrm{OPT}^{(1,T)}))$, and $\mathrm{est}^{t_K} \leq v(\mathrm{OPT}^{(1,T)})$. We argue that one of these batches satisfies the conditions of the proposition.

We first show there is enough welfare between T_1 and T_2 (inclusively) for agents in MECH.

(subadditivity of OPT)

$$v(\mathrm{OPT}_{\mathrm{MECH}}^{(T_1,T_2)}) \geq v(\mathrm{OPT}_{\mathrm{MECH}}^{(1,T_2)}) - v(\mathrm{OPT}_{\mathrm{MECH}}^{(1,T_1-1)})$$

(optimality and monotonicity w.r.t. agents of OPT)

$$\geq v(\mathrm{OPT}^{(1,T_2)}|_{\mathrm{MECH}}) - v(\mathrm{OPT}^{(1,T_1-1)})$$

(concentration at T_2)

$$\geq 0.2v(\mathrm{OPT}^{(1,T_2)}) - v(\mathrm{OPT}^{(1,T_1-1)})$$

(choice of T_1 and T_2)

$$\begin{aligned} &\geq \frac{0.2}{100}v(\mathrm{OPT}^{(1,T)}) - \frac{1}{1000}v(\mathrm{OPT}^{(1,T)}) \\ &= \Omega(v(\mathrm{OPT}^{(1,T)})). \end{aligned}$$

Now intuitively, the remaining issue is that maybe the final unclosed batch starts before T_2 (inclusively), and contains most of the above welfare. We show that this is impossible. In particular, there must be a batch ending after T_2 (inclusively). Suppose otherwise, i.e., $t_K < T_2$. We show that $\mathrm{est}^T \geq 8\mathrm{est}^{t_K}$, leading to a contradiction. In fact,

$$\begin{aligned} &\text{(2-approximation)} && \mathrm{est}^T \geq \frac{1}{2}v(\mathrm{OPT}^{(1,T)}|_{\mathrm{STAT}}) \\ &\text{(concentration at } T) && \geq 0.1v(\mathrm{OPT}^{(1,T)}) \\ & && \geq 8 \times \frac{v(\mathrm{OPT}^{(1,T)})}{100} \\ &\text{(choice of } T_2) && \geq 8 \times v(\mathrm{OPT}^{(1,T_2-1)}) \\ &\text{(definition of } \mathrm{est}^t) && \geq 8 \times \mathrm{est}^{T_2-1} \\ &(t_K < T_2 \text{ and monotonicity of } \mathrm{est}^t) && \\ & && \geq 8 \times \mathrm{est}^{t_K}. \end{aligned}$$

Now we know:

- $v(\mathrm{OPT}_{\mathrm{MECH}}^{(T_1,T_2)}) = \Theta(v(\mathrm{OPT}^{(1,T)}))$, and
- there are only $O(1)$ batches overlapping $[T_1, T_2]$, whose indices are $k_1, \ldots, k_2$ where $t_{k_1-1} < T_1 \leq t_{k_1}$, $t_{k_2} \geq T_2$, and $k_2 - k_1 = O(1)$.

We only need to show that for some $k \in \{k_1, \ldots, k_2\}$,

$$v(\mathrm{OPT}_{\mathrm{MECH}}^{(t_{k-1}+1,t_k)}) = \Theta(v(\mathrm{OPT}^{(1,T)})).$$

By subadditivity and monotonicity w.r.t. items of OPT,

$$\begin{aligned} \sum_{k \in \{k_1,\ldots,k_2\}} v(\mathrm{OPT}_{\mathrm{MECH}}^{(t_{k-1}+1,t_k)}) &\geq v(\mathrm{OPT}_{\mathrm{MECH}}^{(t_{k_1-1}+1,t_{k_2})}) \\ &\geq v(\mathrm{OPT}_{\mathrm{MECH}}^{(T_1,T_2)}). \end{aligned}$$

Since there are only $O(1)$ summands, for some k,

$$\begin{aligned} v(\mathrm{OPT}_{\mathrm{MECH}}^{(t_{k-1}+1,t_k)}) &= \Theta(v(\mathrm{OPT}_{\mathrm{MECH}}^{(T_1,T_2)})) \\ &= \Theta(v(\mathrm{OPT}^{(1,T)})). \end{aligned}$$

This is our desired batch. ∎

We then focus on this constant-approximate batch guaranteed by Lemma 6.2. We argue that the approximation guarantee of the offline mechanism still holds for this batch, so the welfare from this batch alone is a good approximation of the offline optimal welfare. While this may appear trivially true, we note that by the time the offline mechanism is called, the buyers may already possess some items, which may lower their interest in purchasing new items. Such a change of their behavior has a potential to ruin the welfare guarantee. This, however, will not happen if the offline mechanism is approximately monotone, which concludes the proof of Theorem 6.1.

6.3 Approximate Monotonicity of Mechanisms Based on Posted-Price Auctions

In this section, we argue that if the offline mechanism is "essentially based on posted-price auctions and standard revenue-surplus arguments," then the mechanism is approximately monotone. The argument presented here applies in particular for the $O(\log m \log\log m)$-approximate mechanism for subadditive buyers [9] and the $O((\log\log m)^3)$-approximate mechanism for XOS buyers [2] in the prior-free environment, and the $O(\log\log m)$-approximate mechanism for subadditive buyers [16] and 2-approximate mechanism for XOS buyers [15] in the Bayesian environment.

For brevity we refrain from unnecessarily repeating the entire arguments of the offline mechanisms. The key property we need to prove is that in a posted-price auction, if enough "under-priced" items remain unsold, then the allocation supports reasonably large welfare, *no matter what items buyers already possess before the auction.* This can be formalized as the following lemma.

LEMMA 6.3. *Given a set of buyers $\mathcal{C}$ with valuations $\vec{v}$, suppose buyer $i \in \mathcal{C}$ already has items S_i^0. Consider a posted-price auction that is run with items U and prices p_j for $j \in U$ as input and after the auction, buyer i has items $S_i^0 \cup S_i^1$.*

Let OPT be an allocation maximizing the welfare $\sum_{i \in \mathcal{C}} v_i(\mathrm{OPT}_i \cup S_i^0)$. Suppose $\{q_j\}_{j\in U}$ satisfy: for any i and $T' \subseteq \mathrm{OPT}_i$, $\sum_{j \in T'} q_j \le v_i(T')$.

Let $T \subseteq U$ be a set of items satisfying: T is not sold in the auction, and for any $j \in T$, $p_j \le \frac{1}{2} q_j$, then

$$\sum_{i \in \mathcal{C}} v_i(S_i^0 \cup S_i^1) \ge \frac{1}{4} \sum_{j \in T} q_j.$$

Before proving the lemma, we briefly discuss the offline counterpart of Lemma 6.3 and the connection between them. $\{q_j\}$ in the lemma can be viewed as supporting prices for OPT, and T is the unsold set of items whose prices are sufficiently smaller compared to the supporting prices. In the offline environment, when the posted-price auction happens, no buyer has any item, i.e., $S_i^0 = \emptyset$. In such cases, it is easy to show that the outcome of the auction satisfies

$$\sum_{i \in \mathcal{C}} v_i(S_i^0 \cup S_i^1) \ge \frac{1}{2} \sum_{j \in T} q_j.$$

The intuition is that the unsold items provided an option for all buyers, which would guarantee each buyer some surplus (i.e., value minus payment). The buyers, however, did not choose this option, so it must be the case that the buyers chose something more desirable, which gave them only larger surplus. The above lemma essentially says, even if the buyers already have some items before the auction, this bound can only be worse by a factor of 2.

Proof. [Proof of Lemma 6.3] For each $i \in \mathcal{C}$, we show

$$v_i(S_i^0 \cup S_i^1) \ge \frac{1}{4} \sum_{j \in T \cap \mathrm{OPT}_i} q_j.$$

The lemma then follows by summing over i. By purchasing $T \cap \mathrm{OPT}_i$ instead of S_i^1, the marginal utility of i is at least

$$\sum_{j \in T \cap \mathrm{OPT}_i} (q_j - p_j) - v_i(S_i^0) \ge \sum_{j \in T \cap \mathrm{OPT}_i} \frac{1}{2} q_j - v_i(S_i^0),$$

which lower bounds i's value $v_i(S_i^0 \cup S_i^1)$. On the other hand, by monotonicity, i's value is at least $v_i(S_i^0)$. Putting the two bounds together,

$$\begin{aligned} v_i(S_i^0 \cup S_i^1) &\ge \max\left\{ \sum_{j \in T \cap \mathrm{OPT}_i} \frac{1}{2} q_j - v_i(S_i^0), v_i(S_i^0) \right\} \\ &\ge \frac{1}{2} \sum_{j \in T \cap \mathrm{OPT}_i} \frac{1}{2} q_j - v_i(S_i^0) + \frac{1}{2} v_i(S_i^0) \\ &= \frac{1}{4} \sum_{j \in T \cap \mathrm{OPT}_i} q_j. \end{aligned}$$

□

On may check that given the above lemma, the entire arguments in [9] and [2] remain valid even with free items dispensed beforehand.

7 Other Related Work

Initiated by the seminal work of Dobzinski et al. [14], offline truthful combinatorial auctions have been extensively studied in the last decade. For general monotone valuations with demand queries, Dobzinski et al. [14] gave an $O(\sqrt{m})$-approximation, which matches the communication complexity lower bound by Nisan [26]. Restricted to complement-free buyers, the first nontrivial $O(\log^2 m)$ upper bound for XOS valuations was also given by Dobzinski et al. [14]. Dobzinski [9] later improved the upper bound to $O(\log m \log\log m)$ for subadditive buyers. For XOS buyers, Krysta and Vöcking [22] obtained an upper bound of $O(\log m)$ that betters the more general bound for subadditive buyers. Later, Dobzinski [11] further improved this bound to $O(\sqrt{\log m})$ for XOS buyers. In a very recent paper, Assadi and Singla [2] gave an $O((\log\log m)^3)$-approximate mechanism by combining existing techniques with a novel learning procedure, which iteratively estimates the supporting prices of individual items. No super-constant lower bound is known in this setting. Instead of both demand and value queries, if one were restricted only to value queries, Dobzinski et al. [13] gave an $O(\sqrt{m})$ upper bound for submodular buyers, which is matched by information-theoretic [10] and complexity-theoretic [12] lower bounds.

From a pure algorithmic point of view, the problem of computing a welfare maximizing combinatorial allocation has also been extensively studied. For submodular valuations, Vondrák [28] gave an $(e/(e-1))$-approximation using value queries only, with a matching lower bound by Mirrokni et al. [24]. For the more general classes of XOS and subadditive valuations, it is impossible to achieve $O(\sqrt{m})$-approximation using polynomially many value queries [24], which matches an upper bound by Dobzinski et al. [13]. With demand queries, for submodular buyers, a slightly better upper bound was given by Feige and Vondrák [19], while the best known lower bound is $((2e)/(2e-1))$ [12]. For XOS and subadditive buyers, Feige [18] gave an $(e/(e-1))$-approximation and a 2-approximation respectively, using both value and demand queries. Another line of related research considers an online setting with sequentially arriving buyers and b identical copies of each item, which was initiated by Bartal et al. [4] and Awerbuch et al. [3]. In particular, Krysta and Vöcking [22] gave truthful mechanisms that are $O(m^{1/(b+1)}\log(bm))$-competitive for general buyers for any $b \geq 1$, and $O(\log m)$-competitive for XOS buyers when $b = 1$. Cole et al. [5] consider a related setting, where each buyer is present during some time interval, and design prompt mechanisms in this setting. In Bayesian settings where the distributions of buyers' valuations is known, the model with buyers arriving online can be viewed as a combinatorial variant of prophet inequalities. In this setting, Feldman et al. [20] gave a truthful $((2e)/(e-1))$-competitive mechanism for XOS buyers, which was later improved to 2-competitive by Dütting et al. [15]. Ehsani et al. [17] further showed that the ratio improves to $e/(e-1)$ when buyers arrive in a uniformly random order. For subadditive valuations, the $O(\log m)$ competitive ratio by Feldman et al. [20] has been improved to $O(\log m/\log\log m)$ by Zhang [29], and to $O(\log\log m)$ by Dütting et al. [16].

There has been a large body of research on dynamic mechanism design concerning forward-looking additive buyers in the past decade [1, 7, 25]. For a Bayesian environment, Mirrokni et al. [25] propose non-clairvoyant mechanisms in which it is always an optimal strategy for the buyers to report truthfully no matter what the future would be, when the valuations are additive and distributions are independent across the stages. It is later generalized to a setting with public valuation correlations in which the distributions can vary with any publicly observable information from the past of the mechanism [7]. Nonetheless, their model cannot capture submodular valuations.

8 Future Directions

Our works open up several interesting future research directions. We have already remarked on the open problem of designing a truthful online mechanism for omniscient buyers, i.e., who can plan with future knowledge, with submodular valuations. Future research can also consider to improve the upper bound on the welfare approximation in the prior-free online environment with submodular valuations, given the recent breakthrough in the offline setting by Assadi and Singla [2]. On the other hand, it would also be interesting if one could establish a super-constant lower bound for the prior-free online environment with submodular valuations, particularly since no such lower bound is known offline. One can also consider the middle ground between immediate expiration and no expiration to investigate the effect of expiration dates. Note that our negative result continues to hold when the shelf life of each item is a constant in relation to the entire time horizon T. But, it would be interesting to investigate the case when the shelf life is longer, say a polynomial function such as $\sqrt{T}$.

Acknowledgments.

Part of YD's work was done as a graduate student at the Department of Computer Science, Duke University where he was supported by a Google PhD Fellowship in Algorithms, Optimizations, and Markets. DP was supported in part by NSF grants CCF-1955703, CCF-1535972, and an NSF CAREER Award CCF-1750140. HZ was supported by NSF grant IIS-1814056.

References

[1] I. Ashlagi, C. Daskalakis, and N. Haghpanah. Sequential mechanisms with ex-post participation guarantees. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 213–214. ACM, 2016. 1, 17

[2] S. Assadi and S. Singla. Improved truthful mechanisms for combinatorial auctions with submodular bidders. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 233–248. IEEE, 2019. 1, 2, 3, 16, 17

[3] B. Awerbuch, Y. Azar, and A. Meyerson. Reducing truth-telling online mechanisms to online optimization. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 503–510. ACM, 2003. 17

[4] Y. Bartal, R. Gonen, and N. Nisan. Incentive compatible multi unit combinatorial auctions. In *Proceedings of the 9th conference on Theoretical aspects of rationality and knowledge*, pages 72–87. ACM, 2003. 17

[5] R. Cole, S. Dobzinski, and L. Fleischer. Prompt mechanisms for online auctions. In *International Symposium on Algorithmic Game Theory*, pages 170–181. Springer, 2008. 17

[6] P. Cramton, Y. Shoham, and R. Steinberg. *Combinatorial Auctions*. The MIT Press, 2006. ISBN 0262033429. 1

[7] Y. Deng, V. Mirrokni, and S. Zuo. Non-clairvoyant dynamic mechanism design with budget constraints and beyond. *Available at SSRN 3383231*, 2019. 17

[8] N. Devanur, J. Morgenstern, V. Syrgkanis, and S. M. Weinberg. Simple auctions with simple strategies. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 305–322. ACM, 2015. 8

[9] S. Dobzinski. Two randomized mechanisms for combinatorial auctions. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 89–103. Springer, 2007. 1, 2, 3, 16, 17

[10] S. Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 139–148. ACM, 2011. 17

[11] S. Dobzinski. Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 940–948. ACM, 2016. 2, 17

[12] S. Dobzinski and J. Vondrák. The computational complexity of truthfulness in combinatorial auctions. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 405–422. ACM, 2012. 17

[13] S. Dobzinski, N. Nisan, and M. Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 610–618. ACM, 2005. 17

[14] S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 644–652. ACM, 2006. 1, 2, 7, 16, 17

[15] P. Dütting, M. Feldman, T. Kesselheim, and B. Lucier. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 540–551. IEEE, 2017. 2, 3, 16, 17

[16] P. Dütting, T. Kesselheim, and B. Lucier. An O(log log m) prophet inequality for subadditive combinatorial auctions. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2020. 2, 3, 16, 17

[17] S. Ehsani, M. Hajiaghayi, T. Kesselheim, and S. Singla. Prophet secretary for combinatorial auctions and matroids. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 700–714. SIAM, 2018. 2, 17

[18] U. Feige. On maximizing welfare when utility functions are subadditive. *SIAM Journal on Computing*, 39(1):122–142, 2009. 7, 8, 13, 17

[19] U. Feige and J. Vondrák. Approximation algorithms for allocation problems: Improving the factor of 1-1/e. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, 2006. 17

[20] M. Feldman, N. Gravin, and B. Lucier. Combinatorial auctions via posted prices. In *Proceedings*

of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms, pages 123–135. Society for Industrial and Applied Mathematics, 2015. 2, 10, 17

[21] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey. An analysis of approximations for maximizing submodular set functions—ii. In *Polyhedral combinatorics*, pages 73–87. Springer, 1978. 2, 7, 10

[22] P. Krysta and B. Vöcking. Online mechanism design (randomized rounding on the fly). In *International Colloquium on Automata, Languages, and Programming*, pages 636–647. Springer, 2012. 7, 17

[23] B. Lehmann, D. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006. 2, 7, 10

[24] V. Mirrokni, M. Schapira, and J. Vondrák. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *Proceedings of the 9th ACM conference on Electronic commerce*, pages 70–77. ACM, 2008. 17

[25] V. Mirrokni, R. Paes Leme, P. Tang, and S. Zuo. Non-clairvoyant dynamic mechanism design. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, pages 169–169. ACM, 2018. 1, 17

[26] N. Nisan. The communication complexity of approximate set packing and covering. In *International Colloquium on Automata, Languages, and Programming*, pages 868–875. Springer, 2002. 17

[27] C. Papadimitriou, G. Pierrakos, C.-A. Psomas, and A. Rubinstein. On the complexity of dynamic mechanism design. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1458–1475. Society for Industrial and Applied Mathematics, 2016. 1

[28] J. Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 67–74. ACM, 2008. 17

[29] H. Zhang. Improved prophet inequalities for combinatorial welfare maximization with (approximately) subadditive agents. In *Proceedings of the 28th European Symposium on Algorithms*, 2020. 17