

An $O(\log \log m)$ Prophet Inequality for Subadditive Combinatorial Auctions

Paul Dütting ^{*†}, Thomas Kesselheim [‡], Brendan Lucier [§]

^{*}Google Research, Zürich, Switzerland, duetting@google.com

[†]Department of Mathematics, London School of Economics, London, UK, p.d.duetting@lse.ac.uk

[‡]Institute of Computer Science, University of Bonn, Bonn, Germany, thomas@kesselheim@uni-bonn.de

[§]Microsoft Research, Cambridge, MA, USA, brlucier@microsoft.com

Abstract—Prophet inequalities compare the expected performance of an online algorithm for a stochastic optimization problem to the expected optimal solution in hindsight. They are a major alternative to classic worst-case competitive analysis, of particular importance in the design and analysis of simple (posted-price) incentive compatible mechanisms with provable approximation guarantees.

A central open problem in this area concerns subadditive combinatorial auctions. Here n agents with subadditive valuation functions compete for the assignment of m items. The goal is to find an allocation of the items that maximizes the total value of the assignment. The question is whether there exists a prophet inequality for this problem that significantly beats the best known approximation factor of $O(\log m)$.

We make major progress on this question by providing an $O(\log \log m)$ prophet inequality. Our proof goes through a novel primal-dual approach. It is also constructive, resulting in an online policy that takes the form of static and anonymous item prices that can be computed in polynomial time given appropriate query access to the valuations. As an application of our approach, we construct a simple and incentive compatible mechanism based on posted prices that achieves an $O(\log \log m)$ approximation to the optimal revenue for subadditive valuations under an item-independence assumption.

Keywords—prophet inequality; posted prices; combinatorial auction; mechanism design

I. INTRODUCTION

We study the following online stochastic allocation problem. There is a set M of m objects to be divided among n agents. Each agent has a valuation function that assigns a value to every subset of objects. These valuation functions are random, drawn independently from known (but not necessarily identical) distributions. Agents arrive one by one in an arbitrary order and when an agent arrives her valuation is revealed. The decision-maker must choose which subset of objects to give each agent when she arrives. The goal is to maximize the total value of the assignment.

The special case of a single object is precisely the setup of the famous prophet inequality due to Krenkel, Sucheston, and Samuel-Cahn [1], [2], [3]. They show that there exists an online policy whose expected value is at least half of the expected optimal solution in hindsight, which in this case is simply the expected maximum value held by any agent for the object. The policy also has a surprisingly simple form: allocate to the first agent whose value exceeds

a fixed threshold calculated in advance from the known distributions. Interest in prophet inequalities has surged recently, in part due to applications in pricing and auction design driven by the observation that the fixed threshold can be viewed as a posted price. This has led to a line of literature studying prophet inequalities for more general instances of the allocation problem, yielding approximately-optimal online policies and incentive compatible auctions for increasingly general problem instances with many objects and rich classes of valuation functions (e.g., [4], [5], [6], [7], [8], [9], [10], [11], [12]).

One of the more vexing open problems in this space concerns subadditive valuations. A valuation function v is subadditive if $v(S) + v(T) \geq v(S \cup T)$ for all sets of objects S and T . This captures the property that items are not complementary, in the sense that objects are not more valuable together than they are apart. This is a natural and important property in many contexts, and the subadditive allocation problem has thus received considerable attention from both the algorithmic and economic perspectives. For the former, there is a known $O(1)$ -approximate polynomial-time algorithm for the offline problem [13]. For the latter, it is known that running a sealed-bid auction for each object separately yields an $O(1)$ approximation to the optimal welfare at any Bayes-Nash equilibrium [14]. But despite these results, the best-known prophet inequality bound is $O(\log m)$ [6]. Correspondingly, the best known incentive compatible mechanisms for Bayesian subadditive combinatorial auctions achieve $O(\log m)$ approximations to the expected welfare or revenue [6], [15]. This leaves us with a significant asymptotic gap in the known power of truthful versus non-truthful mechanisms, and closing it is a central open challenge.

We make substantial progress on this problem by obtaining an $O(\log \log m)$ -approximate prophet inequality for subadditive valuations. The policy we construct is threshold-based: the designer constructs static and anonymous prices for the objects, and each agent is assigned the set of objects (from among those that remain) that maximizes her value minus the posted prices. The analysis is constructive, and the appropriate prices can be computed in polynomial

time given access to demand queries for the valuations.¹ This prophet inequality directly implies a polynomial-time incentive compatible posted-price mechanism for subadditive combinatorial auctions that achieves an $O(\log \log m)$ approximation to the expected optimal welfare.

Our construction can also be applied to the problem of constructing simple and approximately revenue-optimal mechanisms. A recent line of work has studied the power of simple posted-price-based mechanisms to approximate optimal revenue in auctions with multiple items. Recently, Cai and Zhao [15] showed that under an item-independence assumption, it is possible to obtain an $O(\log m)$ approximation for subadditive valuations using a posted-price mechanism with either (a) an up-front entry fee for each agent, or (b) a restriction on the number of objects each agent can buy. We show that a modification of our prophet inequality can be used to obtain an improved $O(\log \log m)$ approximation, using the same two classes of mechanisms.

A. Our Results and Techniques

An $O(\log \log m)$ Prophet Inequality: Our first result is an existential $O(\log \log m)$ -approximate price-based prophet inequality for subadditive combinatorial auctions. The pricing interpretation is that we are given access to distributions \mathcal{D}_i over subadditive valuation functions v_i , we will precompute item prices p_j for each item $j \in M$, and then agents will arrive one-by-one, each with valuation v_i drawn from \mathcal{D}_i , and buy a subset of items S that maximizes their utility $v_i(S) - \sum_{j \in S} p_j$. Our result is that we can find prices so that the expected welfare of the resulting allocation will be an $O(\log \log m)$ approximation to the expected welfare that could be achieved by an optimal offline algorithm with advance knowledge of all valuations.

Theorem I.1 (Welfare, Existential). *For subadditive valuations drawn independently from known distributions, there exist static anonymous item prices that yield a $O(\log \log m)$ approximation to the optimal expected welfare.*

We prove the theorem for complete information (known valuations), and then extend it to the Bayesian case with incomplete information. At the heart of our approach is the following lemma, which for a given and fixed subadditive valuation v_i asserts the existence of item prices p_j for a given set U that satisfy a certain inequality.

Lemma I.1 (Key Lemma). *For every $i \in N$, subadditive function v_i , and set $U \subseteq M$ there exist prices p_j for $j \in U$ and a probability distribution λ over $S \subseteq U$ such that for all $T \subseteq U$*

$$\sum_{S \subseteq U} \lambda_S \left(v_i(S \setminus T) - \sum_{j \in S} p_j \right) \geq \frac{v_i(U)}{\alpha} - \sum_{j \in T} p_j,$$

¹Given a valuation function v , a set of objects M , and a price $p_j \geq 0$ for all $j \in M$, a demand query returns the set S that maximizes $v(S) - \sum_{j \in S} p_j$.

where $\alpha \in O(\log \log m)$.

Given this lemma it is relatively straightforward to show Theorem I.1. The idea is to let (U_1, \dots, U_n) be the welfare-maximizing allocation, and for each U_i and $j \in U_i$ use the prices p_j from Lemma I.1 with $U = U_i$. The welfare argument then proceeds by rewriting the welfare as the sum of buyer utilities and revenue, with Lemma I.1 providing a tool to lower bound the buyer utilities.

In this lower bound argument the set T from Lemma I.1 can be interpreted as the set of items which are already gone when we consider agent i , and λ is a distribution over sets of items S that agent i considers to buy. Of course, agent i can only buy items that are still available, so she only derives value from $S \setminus T$. The lemma therefore establishes that the utility that can be obtained by agent i is at least a factor α of her contribution to the optimal welfare, less the revenue obtained from selling the items from T .

To prove Lemma I.1 we write down an LP and use strong LP duality to show the following equivalent condition: There exists a probability distribution λ over set of items S so that for every probability distribution μ with $\sum_{T: j \in T} \mu_T \leq \sum_{S: j \in S} \lambda_S$, i.e., that puts at most the same probability mass on each item j as distribution λ , it holds that

$$\sum_{S, T} \lambda_S \cdot \mu_T \cdot v_i(S \setminus T) \geq \frac{1}{\alpha} \cdot v_i(U). \quad (1)$$

We interpret the left-hand side of (1) as a zero-sum game, in which the protagonist chooses λ and the antagonist chooses μ , and the protagonist's goal is to maximize $\sum_{S, T} \lambda_S \cdot \mu_T \cdot v_i(S \setminus T)$. This has a natural interpretation: the designer's goal is to find a purchasing strategy for the buyer that maximizes the value of the set they obtain, and the adversary's goal is to arrange the purchasing outcomes so that removing all previously-sold items (i.e., T) steals most of the value from the buyer, leaving their realized value $v_i(S \setminus T)$ as small as possible.

We prove a lower bound on the value of this game by restricting attention to distributions λ that put the same probability mass q on each item. The crux of our argument is that for each such "equal-marginals distribution" λ with corresponding probability q , the value of the zero-sum game is at least $f(q) - f(q^2)$, where $f(q)$ is the optimal expected social welfare that can be achieved by a distribution over sets of items $S \subseteq U$ that puts probability mass at most q on each item. Intuitively, if it's possible for the adversary to choose some distribution μ over T that is guaranteed to "steal the value" from the buyer's distribution λ over S , then it must be that the set $S \cap T$ has high expected value. But if λ and μ each place probability at most q on each item, then the distribution over $S \cap T$ places probability at most q^2 on each item. Thus, if the adversary can perform well in the zero-sum game for some q , this directly implies that we should consider the game with the significantly smaller

marginal probability q^2 .

To turn this intuition into an $O(\log \log m)$ bound, we let the protagonist consider such “equal marginal distributions” for $q = 2^{-2^k}$ for $k = 0$ to $k = O(\log \log m)$, and obtain a lower bound on the value of the zero-sum game by taking the average of the sum of the corresponding lower bounds $f(q) - f(q^2)$. Now by the choice of the q this sum has $O(\log \log m)$ terms, and the sum is a telescoping sum which evaluates to $f(1/2) - f(1/m^2)$. The proof is completed by observing that the latter is at least $(1/2 - o(1)) \cdot v_i(U)$.

Polynomial-Time Computation of Prices: Our second result shows how to turn this existential proof into a polytime result, assuming appropriate demand query access to the valuation functions.

Theorem I.2 (Welfare, Computational). *For subadditive combinatorial valuations drawn independently from known distributions and any $\epsilon > 0$, there is a polytime (in n , m , and $1/\epsilon$) algorithm to compute static and anonymous item prices for which the resulting posted-price mechanism achieves an $O(\log \log m)$ approximation to the optimal expected welfare up to an additive error of ϵ .*

We prove this theorem by reformulating our optimization problem in a way that avoids having to compute the equilibrium distributions in our zero-sum game, and instead draws a connection to the classic configuration LP for combinatorial assignment. We then use the fact that the configuration LP can be solved in polynomial time using the Ellipsoid method, since a separation oracle can be implemented with demand queries.

Specifically, we draw q uniformly from the set of q ’s introduced above. For the resulting q we consider the dual LP to the configuration LP for $f(q^2)$ — the optimal welfare that can be achieved with a probability distribution that puts probability mass at most q^2 on each item—and we use the prices from this dual LP scaled by q .

The intuition behind this construction is as follows. In our argument above we bounded the value of the zero-sum game by $f(q) - f(q^2)$. Here $f(q)$ is the highest expected value that the protagonist could obtain from a choice of λ if the adversary abstained, and $f(q^2)$ is an upper bound on how much value the antagonist can take away by choosing μ optimally. By taking the dual prices for the configuration LP for $f(q^2)$ and scaling them by q , we are effectively setting prices that approximate the welfare loss due to the antagonist’s strategy, which is to say the worst-case loss from excluding items that have already been sold.

Revenue Maximization: We also show how to leverage our new insights to make progress on another important frontier in algorithmic mechanism design. Namely, the question of how well the optimal revenue that can be obtained by a Bayesian incentive compatible (BIC) mechanism can be approximated by a simple and dominant strategy incentive compatible (DSIC) mechanism. Our revenue approximation

makes use of a framework for constructing simple mechanisms due to Cai and Zhao [15], which builds upon a recent literature applying a duality approach to revenue maximization [16]. Cai and Zhao established an $O(\log m)$ approximation under a natural item independence assumption. Under the same assumption we show:

Theorem I.3 (Revenue). *When buyers have subadditive valuations over independent items, there is a simple DSIC mechanism that yields an $O(\log \log m)$ approximation to the optimal BIC revenue.*

A key step in the proof of Cai and Zhao [15] invokes a posted-price-based prophet inequality for welfare maximization, and indeed their approximation factor of $O(\log m)$ is driven by the $O(\log m)$ -approximate prophet inequality that they apply. However, one cannot apply a prophet inequality to their framework as a black box. The reason is that the prophet inequality is invoked to argue that the value of a certain interim allocation rule—the core of a revenue-optimal mechanism—can be approximated by posted prices. We obtain the improved bound by extending our prophet inequality so that it can handle arbitrary (not necessarily equal) constraints on the marginal probability of allocating each item, and using this to obtain a better price-based approximation to the core.

Going Beyond $O(\log \log m)$: Finally, we demonstrate that the $O(\log \log m)$ -factor that shows up in all our bounds is best possible using our approach. In particular, our analysis restricts to distributions that set the same marginal probability q of allocating each item. We show by way of example that such distributions (and their associated dual prices) can suffer loss as high as $\Omega(\log \log m)$.

Theorem I.4 (Lower Bound). *There exists a subadditive valuation function v_i over m items such that for any $q \in [0, 1]$, any $\alpha \in o(\log \log m)$, and any distribution λ that puts probability at most q on each item there exists a distribution μ that puts probability at most q on each item that violates inequality (1).*

Our restriction to equal-marginal distributions was crucial for our approach to optimizing over distributions. Of course, it is natural to wonder whether our bound could be improved by relaxing the equal-marginals assumption and permitting an arbitrary profile of marginal distributions. Indeed, we conjecture that an $O(1)$ -approximate prophet inequality can be achieved using item prices that are dual to a distribution with unequal marginals. But we leave resolving this conjecture as an open problem.

Discussion: Connection to Balanced Prices: The main difference between our $O(\log \log m)$ approximation and the earlier state-of-the-art $O(\log m)$ prophet inequality from [6] is that this earlier approach constructed prices by approximating subadditive valuations through fractionally subadditive (a.k.a. XOS) functions. This leads to “balanced prices”

in the sense of [11], where the sum of all prices matches the optimal allocation precisely.

More generally, the balanced prices framework of [11] entails constructing prices for any fixed valuation profile, such that (a) the prices of any subset of items partially offset the value lost due to allocating these items, and (b) the sum of all prices is upper bounded by the total value of the optimal allocation of all items. With parameters $1/\alpha$ and β for (a) and (b) this leads to $O(1/(\alpha\beta))$ price-based prophet inequalities. However, such balanced prices cannot lead to a better than $O(\log m)$ approximation for subadditive combinatorial auctions [6].

Our prices are different, and will generally be much higher. The basic intuition is that, under balanced prices, the sum of all prices approximates the optimal welfare, so in a sense the set of all items is “affordable” and the prices facilitate an outcome where most of the items are purchased. However, depending on the curvature of the subadditive valuations, it may be better to target much smaller sets for purchase if they already capture most of the value. By looking at different marginals q we are basically considering different sizes $q \cdot m$ of sets of items to go after, and our key lemma establishes that there is always a good choice of q . As q becomes smaller, the prices we construct are tailored to facilitate purchases of smaller sets of items, and hence the item prices tend to increase.

B. Further related work

From a purely algorithmic perspective social welfare with fractionally subadditive (or XOS) and subadditive valuations can be approximated to within a constant factor assuming demand queries. The state-of-the-art for both XOS and subadditive valuations is a $1 - 1/e$ approximation due to [13]. These approximation guarantees are best possible in the sense that they match the integrality gap of the LP formulations they are based on.

An important question in algorithmic mechanism design concerns the gap between the best (worst-case) approximation guarantee that can be obtained without incentives (i.e., purely algorithmically) and with a truthful mechanism. Recent breakthroughs for XOS valuations were obtained by Dobzinski [17], who gave an $O(\sqrt{\log m})$ truthful approximation mechanism, and then subsequently by Assadi and Singla [18], who gave an improved $O((\log \log m)^3)$ truthful approximation mechanism. Finding DSIC approximation mechanisms with constant worst-case approximation guarantees for either XOS or subadditive valuations or disproving their existence is a major open problem.

A related question concerns the relative power of truthful direct-revelation mechanisms and general mechanisms at equilibrium. The latter can be analyzed using the price of anarchy framework, in which the expected optimal solution is compared with the worst-case expected outcome at any Bayes-Nash equilibrium of the mechanism. Christodoulou,

Kovacs, and Schapira [19] established an $O(1)$ price of anarchy bound for simultaneous item auctions for XOS valuations, which was subsequently extended to a variety of auction formats and related solution concepts [20]. In particular, Feldman, Fu, Gravin, and Lucier [14] established an $O(1)$ price of anarchy for simultaneous item auctions under subadditive valuations.

Prophet inequalities for XOS and subadditive combinatorial auctions in which agents arrive one by one were previously given in [6], [11] and [21]. For XOS valuations an optimal factor 2 is shown in [6], [11], and this can be improved to $1 - 1/e$ by additionally assuming agents arrive in random order [21]. For subadditive valuations, Feldman et al. [6] give an $O(\log m)$ approximation. Rubinfeld and Singla [10] consider a related but different problem, where there is one subadditive function across all entities that arrive over time. They give an $O(\log m \log^2 r)$ prophet inequality, where r is the rank of an arbitrary downward closed feasibility constraint.

In concurrent and independent work, [22] was able to improve the $O(\log m)$ prophet inequality for subadditive combinatorial auctions to $O(\log m / \log \log m)$. This marks an important breakthrough as it shows that it is possible to improve upon the $O(\log m)$ bound.

Our application to revenue maximization builds upon a recent literature on approximately revenue-optimal mechanisms for buyers with multi-dimensional types. For unit-demand buyers, one can obtain a constant approximation to the optimal mechanism with multiple buyers [23], [4], [24]. Simple constant approximations are known for additive buyers with independent valuations, using a technique known as a tail-core decomposition which bounds separately the revenue contribution from rare outlier values and from “expected” valuation profiles [25], [26], [27], [28]. Chawla and Miller showed how to combine both approaches to develop a general class of approximately optimal mechanisms based on posted prices with per-buyer entry fees [29]. Cai, Devanur, and Weinberg further unify these approaches using a flexible duality framework to effectively “linearize” valuations with respect to revenue [16]. The ideas behind these mechanisms have since been extended to more general valuation classes, including XOS and subadditive valuations [15], [30]. Most related to the current paper is the work of Cai and Zhao [15], which (among other things) uses this framework to design an $O(\log m)$ -approximate mechanism for subadditive valuations, based on posted item prices with per-buyer entry fees.

II. MODEL AND DEFINITIONS

Subadditive Combinatorial Auctions: We are given a set N of n buyers and a set M of m goods. Each buyer $i \in N$ has a valuation function $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$, which is assumed to be normalized and monotone, i.e., $v_i(\emptyset) = 0$ and $v_i(S) \leq v_i(T)$ for $S \subseteq T \subseteq M$. A valuation function v_i is *subadditive* if $v_i(S) + v_i(T) \geq v_i(S \cup T)$ for $S, T \subseteq M$.

We use $\mathbf{v} = (v_1, \dots, v_n)$ to denote a vector of valuation functions. We will occasionally write $\mathbf{v} = (v_i, \mathbf{v}_{-i})$, where we use \mathbf{v}_{-i} to denote the valuations of all buyers except buyer i .

We assume a *Bayesian setting*, in which the valuation function of each buyer i is drawn independently from distribution \mathcal{D}_i . We write $\mathcal{D} = \prod_i \mathcal{D}_i$ for the joint distribution. We emphasize that the independence here is *across bidders*. Valuations of a fixed agent can be arbitrarily correlated across items (though we revisit this when discussing applications to revenue maximization in Section V). We assume that the designer knows the distributions from which the valuation functions are drawn, but not the realizations of the random draws.

An allocation $\mathbf{x} = (x_1, \dots, x_n)$ defines for each buyer $i \in N$ a set of goods $x_i \subseteq M$ that he receives. We require that no good is assigned more than once, i.e., that $x_i \cap x_j = \emptyset$ whenever $i \neq j$. We write $\mathbf{x}_{<i}$ for partial allocations to buyers $s < i$, i.e., for $t \geq i$ we have $x_t = \emptyset$.

We evaluate allocations by the welfare they achieve. The welfare of an allocation \mathbf{x} is $\sum_i v_i(x_i)$. We write $\text{OPT}(\mathbf{v})$ for the welfare-maximizing allocation, and $\mathbf{v}(\text{OPT}(\mathbf{v})) = \sum_{i=1}^n v_i(\text{OPT}_i(\mathbf{v}))$ for the welfare it achieves.

Posted-Price Mechanisms: A posted-price mechanism uses a set of functions $p_i(\cdot \mid \mathbf{x}_{<i}) : 2^M \rightarrow \mathbb{R}_{\geq 0}$ which assign a non-negative price to each set of items $S \subseteq M$. Note that these functions can be personalized, they can be “per set” rather than “per item” (i.e., the price of a set of items need not be a sum of prices of individual items), and they may depend on which items were already allocated.

Of particular interest will be posted-price mechanisms that use static anonymous item prices. A posted-price mechanism has *anonymous* prices if there exists a single set of functions $p(\cdot \mid \mathbf{x}_{<i})$ such that $p_i(S \mid \mathbf{x}_{<i}) = p(S \mid \mathbf{x}_{<i})$ for all i, S , and $\mathbf{x}_{<i}$. It uses *item prices* if $p_i(S \mid \mathbf{x}_{<i}) = \sum_{j \in S} p_i(\{j\} \mid \mathbf{x}_{<i})$ for all i, S , and $\mathbf{x}_{<i}$. Finally, prices are *static* if for each i there is a single function $p_i(\cdot)$ such that $p_i(S \mid \mathbf{x}_{<i}) = p_i(S)$ for all S and $\mathbf{x}_{<i}$.

An important advantage of posted-price mechanism with static anonymous item prices is that they can be succinctly described by a single vector $\mathbf{p} = (p_1, \dots, p_m) \in \mathbb{R}_{\geq 0}^m$.

A posted-price mechanism proceeds as follows. The buyers arrive sequentially, and for notational convenience we assume they are indexed according to their arrival order.² Upon arrival of buyer i the mechanism posts a price $p_i(S, \mathbf{x}_{<i})$ for each set of items S . Buyer i buys any set of items x_i that maximizes her utility $u_i(x_i, \mathbf{p}) = v_i(x_i) - \sum_{j \in x_i} p_j$ among all such sets.

Given a fixed choice of item prices, we will tend to write ALG for the corresponding posted-price mechanism,

²All of our results continue to hold in a more general setting where the arrival order is arbitrary and unknown to the designer (but still fixed in advance) and is revealed online as the buyers arrive.

$\text{ALG}(\mathbf{v})$ for the resulting allocation of items when valuations are \mathbf{v} , and use $\mathbf{v}(\text{ALG}(\mathbf{v})) = \sum_{i=1}^n v_i(\text{ALG}_i(\mathbf{v}))$ to denote the welfare it achieves.

Prophet Inequalities: We will follow the “prophet-inequality paradigm” to evaluate the performance of posted-price mechanisms. That is, we will evaluate the performance of a posted-price mechanism ALG by comparing its expected welfare $\mathbf{E}_{\mathbf{v} \sim \mathcal{D}}[\mathbf{v}(\text{ALG}(\mathbf{v}))]$ to the expected optimal welfare $\mathbf{E}_{\mathbf{v} \sim \mathcal{D}}[\mathbf{v}(\text{OPT}(\mathbf{v}))]$. Extending the notion of competitive ratio from the worst-case analysis of online algorithms, we define the (*stochastic*) *competitive ratio* of a posted-price mechanism as

$$\sup_{\mathcal{D}} \frac{\mathbf{E}_{\mathbf{v} \sim \mathcal{D}}[\mathbf{v}(\text{OPT}(\mathbf{v}))]}{\mathbf{E}_{\mathbf{v} \sim \mathcal{D}}[\mathbf{v}(\text{ALG}(\mathbf{v}))]}.$$

III. AN $O(\log \log m)$ PRICE-BASED PROPHET INEQUALITY FOR WELFARE

We start by establishing the existence of an $O(\log \log m)$ -competitive price-based prophet inequality for subadditive combinatorial auctions and the goal of maximizing welfare.

Theorem III.1. *For subadditive combinatorial auctions there is a $O(\log \log m)$ -competitive posted-price mechanism that uses static anonymous item prices.*

It suffices to show Theorem III.1 for $m > 2$ as for $m = O(1)$ the competitive ratio is constant. We will prove the theorem in two steps. In Section III-A, we show the claim for complete information. That is, we assume valuations are fixed and known. In Section III-B, we sketch our proof for the Bayesian case with incomplete information.

A. Proof for Complete Information

Our key lemma and driver of the improved competitive ratio is the following lemma. We prove this lemma using LP-duality, and derive the existence of appropriate prices and the corresponding probability distribution over sets of items through a zero-sum game formulation. Lemma III.1 is a restatement of Lemma I.1 from the introduction. Recalling the discussion after Lemma I.1, the intuition is that the revenue raised by selling items that would typically be allocated to buyer i (set T), plus the utility that buyer i can obtain from the remaining items (by buying set $S \setminus T$), approximates buyer i ’s contribution to the expected optimal welfare (i.e., $v_i(U)$).

Lemma III.1. *For every $i \in N$, subadditive function v_i , and set $U \subseteq M$ there exist prices p_j for $j \in U$ and a probability distribution λ over $S \subseteq U$ such that for all $T \subseteq U$*

$$\sum_{j \in T} p_j + \sum_{S \subseteq U} \lambda_S \left(v_i(S \setminus T) - \sum_{j \in S} p_j \right) \geq \frac{v_i(U)}{\alpha},$$

where $\alpha \in O(\log \log m)$.

Before we prove Lemma III.1, let’s see how it implies the desired result.

Proof of Theorem III.1 (complete information): Let $\text{OPT}(\mathbf{v}) = (U_1, \dots, U_n)$ be the welfare-maximizing allocation for valuations \mathbf{v} . Define a vector \mathbf{p} of item prices as follows: For $i \in N$ and $j \in U_i$ use price p_j from Lemma III.1. Let ALG be the posted-price mechanism that uses prices \mathbf{p} . Denote the allocation of ALG on valuation profile \mathbf{v} by $\text{ALG}_1(\mathbf{v}), \dots, \text{ALG}_n(\mathbf{v})$ and let $\text{SOLD}(\mathbf{v}) = \cup_{i=1}^n \text{ALG}_i(\mathbf{v})$ denote the set of items sold by the mechanism.

To derive a lower bound on the welfare achieved by the posted-price mechanism we will use that the welfare can be decomposed into utility and revenue. Namely, if we write $u_i((v_i, \mathbf{v}_{-i}), \mathbf{p})$ for the utility of buyer i and $r(\mathbf{v}, \mathbf{p})$ for the revenue, then

$$\mathbf{v}(\text{ALG}(\mathbf{v})) = \sum_{i=1}^n u_i((v_i, \mathbf{v}_{-i}), \mathbf{p}) + r(\mathbf{v}, \mathbf{p}).$$

We begin by deriving a lower bound on the sum of the utilities. To this end consider an arbitrary buyer i . Let λ be the probability distribution over sets of items $S \subseteq U_i$ from Lemma III.1 and let $T = \cup_{\ell < i} \text{ALG}(\mathbf{v})_\ell \cap U_i \subseteq \text{SOLD}(\mathbf{v}) \cap U_i$. Now because buyer i could draw a set of items S from λ and buy set $S \setminus T$ (or no set at all if this gives her negative utility),

$$\begin{aligned} u_i((v_i, \mathbf{v}_{-i}), \mathbf{p}) &\geq \sum_{S \subseteq U_i} \lambda_S \left(v_i(S \setminus T) - \sum_{j \in S} p_j \right) \\ &\geq \frac{v_i(U_i)}{\alpha} - \sum_{j \in T} p_j, \end{aligned}$$

where the inequality holds by Lemma III.1.

Summing over all buyers i , we obtain

$$\sum_{i=1}^n u_i((v_i, \mathbf{v}_{-i}), \mathbf{p}) \geq \frac{\mathbf{v}(\text{OPT}(\mathbf{v}))}{\alpha} - \sum_{j \in \text{SOLD}(\mathbf{v})} p_j. \quad (2)$$

On the other hand, the revenue obtained by the posted-price mechanism is

$$r(\mathbf{v}, \mathbf{p}) = \sum_{j \in \text{SOLD}(\mathbf{v})} p_j. \quad (3)$$

Adding (2) and (3) shows the claim. \blacksquare

To prove Lemma III.1 we will write down an LP that captures the claim. The lemma statement will be satisfied whenever the optimal solution to the LP has non-negative value, and we will show that this is indeed the case using strong duality.

Consider an arbitrary buyer i and an arbitrary set U . Let $\gamma = 1/\alpha$. In order to establish Lemma III.1, we have to show that there are prices p_j for $j \in U$ and a distribution λ over sets of items $S \subseteq U$ such that for all $T \subseteq U$

$$\sum_{j \in T} p_j + \sum_{S \subseteq U} \lambda_S \left(v_i(S \setminus T) - \sum_{j \in S} p_j \right) \geq \gamma v_i(U).$$

or equivalently

$$\sum_{S \subseteq U} \lambda_S \sum_{j \in S} p_j - \sum_{j \in T} p_j \leq \sum_{S \subseteq U} \lambda_S v_i(S \setminus T) - \gamma v_i(U). \quad (4)$$

To show inequality (4), we will consider the following LP. The LP is for a fixed λ and has variables $p_j \geq 0$ for $j \in U$ and two more variables $\ell_+ \geq 0$ and $\ell_- \geq 0$. The extra variables model a slack term $\ell_+ - \ell_-$ of arbitrary sign which we consider adding to the the left-hand side of the inequality.³ The LP maximizes the slack term.

$$\begin{aligned} \max \quad & \ell_+ - \ell_- \\ \text{s.t.} \quad & \sum_{S \subseteq U} \lambda_S \sum_{j \in S} p_j - \sum_{j \in T} p_j + (\ell_+ - \ell_-) \\ & \leq \sum_{S \subseteq U} \lambda_S v_i(S \setminus T) - \gamma v_i(U) \quad \text{for all } T \subseteq U \\ & p_j \geq 0 \quad \text{for all } j \in U \\ & \ell_+ \geq 0 \\ & \ell_- \geq 0. \end{aligned}$$

As a non-negative slack means that inequality (4) is fulfilled, we know that there are prices fulfilling inequality (4) if and only if this LP has an optimal solution with non-negative objective value.

Our strategy for showing this will be to go through the dual. Indeed by strong duality it is equivalent to show that every feasible solution to the following dual LP with variables $\mu_T \geq 0$ for $T \subseteq U$ has non-negative value.

$$\begin{aligned} \min \quad & \sum_T \mu_T \left(\sum_{S \subseteq U} \lambda_S v_i(S \setminus T) - \gamma v_i(U) \right) \\ \text{s.t.} \quad & - \sum_{T: j \in T} \mu_T + \sum_T \sum_{S: j \in S} \lambda_S \mu_T \geq 0 \quad \text{for all } j \in U \\ & \sum_T \mu_T = 1 \\ & \mu_T \geq 0 \quad \text{for all } T \subseteq U. \end{aligned}$$

We note that the dual constraints are equivalent to $\sum_T \mu_T = 1$ and $\sum_{T: j \in T} \mu_T \leq \sum_{S: j \in S} \lambda_S$ for all $j \in U$. So they naturally define probability distributions over sets of items $T \subseteq U$, which can put at most the same probability mass on each item $j \in U$ as the probability distribution λ .

This means that the LP has non-negative value if and only if for every probability distribution μ with $\sum_{T: j \in T} \mu_T \leq \sum_{S: j \in S} \lambda_S$ it holds that

$$\sum_{S, T} \lambda_S \mu_T v_i(S \setminus T) \geq \gamma v_i(U). \quad (5)$$

³The use of two slack variables is simply to have the LP in canonical form; they could equivalently be replaced by a single slack variable of arbitrary sign.

We can now formulate the search for an appropriate λ as a zero-sum game in which pure strategies correspond to subsets of items. The maximizing player chooses S , the minimizing player chooses T , and the payoff associated with two sets S and T is $v_i(S \setminus T)$. We want to show that there is mixed strategy λ for the maximizing player such that when the minimizing player is constrained to use a mixed strategy μ which puts at most the same probability mass on each item j as λ , then the value of the game is at least $\gamma v_i(U)$.

To this end let $\mathbf{q} = (q_1, \dots, q_{|U|})$ where $q_j \in [0, 1]$ and let $\Delta(\mathbf{q})$ denote all probability distributions ν over sets $S \subseteq U$ such that $\sum_{T \ni j} \nu_T \leq q_j$ for all j . Define

$$g(\mathbf{q}) = \max_{\lambda \in \Delta(\mathbf{q})} \min_{\mu \in \Delta(\mathbf{q})} \sum_{S, T \subseteq U} \lambda_S \mu_T v_i(S \setminus T). \quad (6)$$

Inequality (5) and hence inequality (4) and Lemma III.1 are therefore equivalent to there being a \mathbf{q} such that $g(\mathbf{q}) \geq v_i(U)/O(\log \log m)$.

The following lemma shows that it is in fact possible to achieve this with a uniform vector \mathbf{q} in which $q_i = q_j$ for all i and j .

Lemma III.2. *There exists a $q \in [0, 1]$ such that for $\mathbf{q} = (q, \dots, q) \in [0, 1]^{|U|}$ we have $g(\mathbf{q}) \geq \frac{1}{O(\log \log m)} v_i(U)$.*

Proof: For $q \in [0, 1]$ define

$$g(q) = \max_{\lambda \in \Delta(q, \dots, q)} \min_{\mu \in \Delta(q, \dots, q)} \sum_{S, T \subseteq U} \lambda_S \mu_T v_i(S \setminus T), \quad \text{and}$$

$$f(q) = \max_{\lambda \in \Delta(q, \dots, q)} \sum_{S \subseteq U} \lambda_S v_i(S).$$

We now use subadditivity, that $\sum_T \mu_T = 1$, and finally that $\mu_T \lambda_S$ defines a probability distribution on $S \cap T$ that puts at most probability mass q^2 on each item to obtain

$$\begin{aligned} g(q) &= \max_{\lambda \in \Delta(q, \dots, q)} \min_{\mu \in \Delta(q, \dots, q)} \sum_{S, T \subseteq U} \lambda_S \mu_T v_i(S \setminus T) \\ &\geq \max_{\lambda \in \Delta(q, \dots, q)} \min_{\mu \in \Delta(q, \dots, q)} \sum_{S, T \subseteq U} \lambda_S \mu_T (v_i(S) - v_i(S \cap T)) \\ &= \max_{\lambda \in \Delta(q, \dots, q)} \left(\sum_{S \subseteq U} \lambda_S v_i(S) - \max_{\mu \in \Delta(q, \dots, q)} \sum_{S, T \subseteq U} \lambda_S \mu_T v_i(S \cap T) \right) \\ &\geq \max_{\lambda \in \Delta(q, \dots, q)} \left(\sum_{S \subseteq U} \lambda_S v_i(S) - \max_{\nu \in \Delta(q^2, \dots, q^2)} \sum_{S \subseteq U} \nu_S v_i(S) \right) \\ &= f(q) - f(q^2), \end{aligned}$$

For any ℓ we thus have,

$$\begin{aligned} \sum_{i=0}^{\ell} g(2^{-2^i}) &\geq \sum_{i=0}^{\ell} \left(f(2^{-2^i}) - f(2^{-2^{i+1}}) \right) \\ &= f(2^{-1}) - f(2^{-2^{\ell+1}}), \end{aligned}$$

by a telescoping sum argument.

With $\ell = \log \log m$,

$$2^{-2^{\ell+1}} = 2^{-2 \log m} = \frac{1}{m^2}.$$

Hence for $\ell = \log \log m$,

$$\sum_{i=0}^{\ell} g(2^{-2^i}) \geq f\left(\frac{1}{2}\right) - f\left(\frac{1}{m^2}\right).$$

We conclude the proof by showing a lower bound on $f(1/2)$ and an upper bound on $f(1/m^2)$. A lower bound on $f(1/2)$ follows from the fact that λ could take the set U_i with probability $1/2$. So

$$f\left(\frac{1}{2}\right) \geq \frac{1}{2} \cdot v_i(U).$$

For the upper bound on $f(1/m^2)$ we exploit the trivial upper bound on $v_i(S)$ namely $v_i(U)$. Using this we obtain,

$$\begin{aligned} f\left(\frac{1}{m^2}\right) &= \max_{\lambda \in \Delta(1/m^2, \dots, 1/m^2)} \sum_{S \subseteq U} \lambda_S v_i(S) \\ &\leq \max_{\lambda \in \Delta(1/m^2, \dots, 1/m^2)} \left(\sum_{S \subseteq U, S \neq \emptyset} \lambda_S \right) v_i(U) \\ &\leq \sum_{j \in U} \left(\sum_{S: S \ni j} \lambda_S \right) v_i(U) \\ &\leq m \cdot \frac{1}{m^2} \cdot v_i(U) \\ &= \frac{1}{m} \cdot v_i(U). \end{aligned}$$

We conclude that

$$\max_q g(q) \geq \frac{1}{\ell+1} \sum_{i=0}^{\ell} g(2^{-2^i}) \geq \frac{1}{\ell+1} \left(\frac{1}{2} - \frac{1}{m} \right) v_i(U),$$

as claimed. \blacksquare

B. Proof for Incomplete Information

Our proof for the Bayesian case is based on the following variant of Lemma III.1. The proof of this lemma follows the same basic steps as the proof of Lemma III.1, but requires some additional care when deriving a lower bound on the value of the zero-sum game. In particular, since the zero-sum game now has payoffs that are defined with respect to the distribution over valuations, our argument requires that we relate the value of the game to the expected value of the distribution over full-information games.

Lemma III.3. *For every probability distribution $\mathcal{D} = \prod_i \mathcal{D}_i$ over subadditive valuation functions, there exist prices p_j for $j \in M$ and probability distributions $\lambda^{i, \mathbf{v}}$ over $S \subseteq M$ for all i and \mathbf{v} such that for all $T \subseteq M$*

$$\begin{aligned} \sum_{j \in T} p_j + \mathbf{E}_{\mathbf{v}} \left[\sum_{i=1}^n \sum_S \lambda_S^{i, \mathbf{v}} \left(v_i(S \setminus T) - \sum_{j \in S} p_j \right) \right] \\ \geq \frac{1}{\alpha} \mathbf{E}_{\mathbf{v}} [v(\text{OPT}(\mathbf{v}))], \end{aligned}$$

where $\alpha \in O(\log \log m)$.

Using Lemma III.3 it is straightforward to prove Theorem III.1 using a hallucination trick similar to that in the price of anarchy literature [20], [31], the literature on algorithmic stability [32], and in the balanced prices framework for prophet inequalities [11].

IV. COMPUTING PRICES IN POLYNOMIAL TIME

Theorem III.1 shows the existence of static anonymous item prices that yield an $O(\log \log m)$ approximation to the optimal welfare. In this section, we establish the following computational version of this result. It shows how to achieve the same approximation guarantee in polynomial time. The proof also reveals an alternative for choosing prices that yield an $O(\log \log m)$ approximation. We assume to have access to demand oracles for the valuation functions in the support of \mathcal{D} . Recall that a demand oracle for valuation function v_i takes as input item prices p_1, \dots, p_m and returns the set $S \subseteq M$ that maximizes $v_i(S) - \sum_{j \in S} p_j$.

Theorem IV.1. *For subadditive combinatorial auctions and any $\epsilon > 0$, there is a polynomial-time (in n , m , and $1/\epsilon$) algorithm to compute static and anonymous item prices for which the resulting posted-price mechanism achieves expected welfare at least $\frac{1}{\alpha} \cdot \mathbf{E}_{\mathbf{v}}[\mathbf{v}(\text{OPT}(\mathbf{v}))] - \epsilon$ where $\alpha = O(\log \log m)$.*

Our proof is based on the following version of Lemma III.3 that includes computations in polynomial time using demand oracles.

Lemma IV.1. *For every probability distribution $\mathcal{D} = \prod_i \mathcal{D}_i$ over subadditive valuation functions and every $\epsilon > 0$, there exist prices p_j for $j \in M$ and probability distributions $\lambda^{i,\mathbf{v}}$ over $S \subseteq M$ for all i and \mathbf{v} such that for all $T \subseteq M$*

$$\begin{aligned} \sum_{j \in T} p_j + \mathbf{E}_{\mathbf{v}} \left[\sum_{i=1}^n \sum_S \lambda_S^{i,\mathbf{v}} \left(v_i(S \setminus T) - \sum_{j \in S} p_j \right) \right] \\ \geq \frac{1}{\alpha} \mathbf{E}_{\mathbf{v}}[\mathbf{v}(\text{OPT}(\mathbf{v}))] - \epsilon, \end{aligned}$$

where $\alpha \in O(\log \log m)$. Moreover, assuming that $\mathbf{v}(\text{OPT}(\mathbf{v})) \leq 1$ for all \mathbf{v} in the support, for any $\zeta > 0$, there is an algorithm that uses $\text{poly}(n, m, 1/\epsilon, \log(1/\zeta))$ demand oracle queries that computes such prices with probability at least $1 - \zeta$.

We will prove the lemma under the assumption that we can compute expectations $\mathbf{E}_{\mathbf{v}}[\cdot]$. In the full version we show how to simulate this ability by sampling from the distributions.

Computation of Prices: Fix some valuation profile \mathbf{v} . As in the case of complete information, we will consider an f -function that maximizes the expected value $v_i(S)$, where S is drawn from a constrained distribution over sets of items. To this end, for each $q \in [0, 1]$ let $\Gamma(q) = \{\{\nu^i\}_{i=1}^n \mid$

$\sum_{i=1}^n \sum_{S: S \ni j} \nu_S^i \leq q$ for all j in $M\}$ be the collection of distribution profiles for which the marginal probability that each item is allocated is at most q . Then we will define

$$f^{\mathbf{v}}(q) = \sup_{\lambda \in \Gamma(q)} \sum_{i=1}^n \sum_S \lambda_S^i v_i(S).$$

Fix $q \in [0, 1]$ and let $\lambda^{i,\mathbf{v},q}$ be some choice of $\lambda \in \Gamma(q)$ that achieves the maximum in the definition of $f^{\mathbf{v}}(q)$.

We now consider the following linear program:

$$\begin{aligned} \max \quad & \sum_{i=1}^n \sum_S x_S^i v_i(S) \\ \text{s.t.} \quad & \sum_i \sum_{S: j \in S} x_S^i \leq q^2 \quad \text{for all } j \in M \\ & \sum_S x_S^i \leq 1 \quad \text{for all } i \\ & x_S^i \geq 0 \quad \text{for all } i, S \end{aligned}$$

The value of this program is precisely equal to $f^{\mathbf{v}}(q^2)$, from the definition of $f^{\mathbf{v}}$. The dual of this linear program is given by

$$\begin{aligned} \min \quad & \sum_{j \in M} q^2 y_j + \sum_{i=1}^n u_i \\ \text{s.t.} \quad & u_i + \sum_{j \in S} y_j \geq v_i(S) \quad \text{for all } i, S \\ & u_i \geq 0 \quad \text{for all } i \\ & y_j \geq 0 \quad \text{for all } j \end{aligned}$$

Let $(y_j^{\mathbf{v},q})_{j \in M}$, $(u_i^{\mathbf{v},q})_{i \in N}$ denote an optimal solution to this dual. We note that this optimal solution can be computed in polynomial time. Indeed, it suffices to find a separation oracle for the dual program. A separation oracle query is equivalent to finding a set S that maximizes $v_i(S) - \sum_{j \in S} y_j$ for a given choice of $(y_j)_{j \in M}$. But this is precisely a demand query, interpreting the dual variables y_j as item prices. Thus, given access to a demand oracle for v_i , one can solve the dual program and compute $(y_j^{\mathbf{v},q})_{j \in M}$.

Our algorithm now chooses q so as to maximize $\mathbf{E}_{\mathbf{v}}[f^{\mathbf{v}}(q)] - \mathbf{E}_{\mathbf{v}}[f^{\mathbf{v}}(q^2)]$ among all $q = 2^{-2^X}$ for some $X \in \{0, \dots, \ell\}$, where $\ell = \log \log m$. We then define our prices according to $p_j^q = \mathbf{E}_{\mathbf{v}}[q y_j^{\mathbf{v},q}]$. That is, p_j^q is equal to the expected value of $y_j^{\mathbf{v},q}$ scaled by q .

Approximation Guarantee: To prove the approximation guarantee, we first show a property of the prices that holds regardless of the choice of q .

Claim IV.1. *For any choice of $q \in [0, 1]$, the prices defined by $p_j^q = \mathbf{E}_{\mathbf{v}}[q y_j^{\mathbf{v},q}]$ fulfill*

$$\begin{aligned} \mathbf{E}_{\mathbf{v}} \left[\sum_{i,S} \lambda_S^{i,\mathbf{v},q} \left(v_i(S \setminus T) - \sum_{j \in S} p_j^q \right) + \sum_{j \in T} p_j^q \right] \\ \geq \mathbf{E}_{\mathbf{v}}[f^{\mathbf{v}}(q)] - \mathbf{E}_{\mathbf{v}}[f^{\mathbf{v}}(q^2)]. \end{aligned}$$

Claim IV.1 holds for a fixed but arbitrary choice of $q \in [0, 1]$. Our algorithm chooses among all $q = 2^{-2^X}$ for $X \in \{0, \dots, \ell\}$, where $\ell = \log \log m$, the one that maximizes $\mathbf{E}_{\mathbf{v}} [f^{\mathbf{v}}(q)] - \mathbf{E}_{\mathbf{v}} [f^{\mathbf{v}}(q^2)]$. Note that by this choice of q , we have

$$\begin{aligned} & \mathbf{E}_{\mathbf{v}} [f^{\mathbf{v}}(q)] - \mathbf{E}_{\mathbf{v}} [f^{\mathbf{v}}(q^2)] \\ & \geq \frac{1}{\ell+1} \sum_{X=0}^{\ell} \left(\mathbf{E}_{\mathbf{v}} [f^{\mathbf{v}}(2^{-2^X})] - \mathbf{E}_{\mathbf{v}} [f^{\mathbf{v}}((2^{-2^X})^2)] \right) \\ & = \frac{1}{\ell+1} \mathbf{E}_{\mathbf{v}} \left[\left(f^{\mathbf{v}}\left(\frac{1}{2}\right) - f^{\mathbf{v}}\left(\frac{1}{m^2}\right) \right) \right]. \end{aligned}$$

We have already established $f^{\mathbf{v}}\left(\frac{1}{2}\right) \geq \frac{1}{2} \mathbf{v}(\text{OPT}(\mathbf{v}))$ because one particular choice for λ would be $\text{OPT}(\mathbf{v})$ with probability $\frac{1}{2}$ and the empty allocation otherwise. Furthermore, $f^{\mathbf{v}}\left(\frac{1}{m^2}\right) \leq \frac{1}{m} \mathbf{v}(\text{OPT}(\mathbf{v}))$ because for any $\lambda \in \Gamma\left(\frac{1}{m^2}\right)$, we have $\sum_{i=1}^n \sum_S \lambda_S^i v_i(S) \leq \sum_{i=1}^n \sum_S \lambda_S^i \sum_{j \in S} v_i(\{j\}) \leq \frac{1}{m^2} \sum_{j \in M} v_i(\{j\}) \leq \frac{1}{m} \mathbf{v}(\text{OPT}(\mathbf{v}))$.

So, these calculations in combination with Claim IV.1 yield that for the value of q chosen by the algorithm

$$\begin{aligned} & \mathbf{E}_{\mathbf{v}} \left[\sum_{i,S} \lambda_S^{i,\mathbf{v},q} \left(v_i(S \setminus T) - \sum_{j \in S} p_j^q \right) + \sum_{j \in T} p_j^q \right] \\ & \geq \frac{1}{\ell+1} \left(\frac{1}{2} - \frac{1}{m} \right) \mathbf{E}_{\mathbf{v}} [\mathbf{v}(\text{OPT}(\mathbf{v}))]. \end{aligned}$$

The proof of Theorem IV.1 now follows from Lemma IV.1 in a similar way as Theorem III.1 followed from Lemma III.3.

V. AN $O(\log \log m)$ APPROXIMATION TO OPTIMAL REVENUE

We next show how to extend our arguments to obtain a posted-price mechanism that achieves near-optimal revenue rather than welfare. We will follow the approach of Cai and Zhao [15]. Fix the valuation distribution $\mathcal{D} = \prod_i \mathcal{D}_i$. We will make an independence assumption on each distribution \mathcal{D}_i , which is that the valuations are *subadditive over independent items*. Roughly speaking, this means that for any S and T with $S \cap T = \emptyset$, the random variables $v_i(S)$ and $v_i(T)$ are distributed independently. In particular, this implies $v_i(\{j\})$ is distributed independently for each item j . We'll write $v_i(j) = v_i(\{j\})$ for convenience.

Theorem V.1. *When buyers have subadditive valuations over independent items, there exists a simple, deterministic, and DSIC mechanism that achieves an $\Omega(1/\alpha)$ approximation to the optimal BIC revenue where $\alpha = O(\log \log m)$.*

Just like the mechanism of Cai and Zhao [15], our mechanism will be from one of the following two classes:

- 1) Rationed sequential posted-price mechanism (RSPM): The buyers are approached in a fixed order. For each buyer, each item is assigned a static and potentially

personalized posted price. Each buyer can purchase at most a single item at its listed price.

- 2) Anonymous sequential posted-price with entry fee mechanism (ASPE): Each item is assigned a static anonymous posted price. The buyers are approached in a fixed order, and each buyer faces an entry fee that can depend on the set of items that have not yet been sold when they arrive. If the buyer pays the entry fee, they can purchase any set of items at their posted prices.

We will write POSTREV for the optimal revenue attainable using a RSPM, APOSTENREV for the optimal revenue attainable using an ASPE, and REV for the optimal BIC revenue.

In Section V-A and Section V-B we describe the high-level approach of Cai and Zhao [15] and the key facts of their construction that we will reuse. In Section V-B we also state our key lemma, Lemma V.6, and show how it implies the improved bound in Theorem V.1. We prove Lemma V.6 in Section V-C.

A. A First Core-Tail Decomposition and a Benchmark

In this section we summarize a particular core-tail decomposition due to Cai and Zhao [15]. First some notation. We'll use σ to describe an interim allocation rule of mechanism, where $\sigma_S^i(v_i)$ is the probability that agent i is allocated set S when agent i has valuation v_i . We'll write $\pi_j^i(v_i) = \sum_{S \ni j} \sigma_S^i(v_i)$ for the probability that agent i is allocated item j under valuation v_i .

Fix personalized item thresholds $\beta_{ij} \geq 0$ for all i and j , whose values will be chosen later. We now describe what is meant by the core. For each buyer i , define a threshold c_i as follows:

$$c_i = \inf \{ x \geq 0 : \sum_j \Pr_{v_i} [v_i(j) \geq \beta_{ij} + x] \leq 1/2 \}.$$

Let $C_i(v_i) = \{j : v_i(j) < \beta_{ij} + c_i\}$. That is, $C_i(v_i)$ contains all items j for which agent i does not have too large a value for item j as a singleton, relative to threshold β_{ij} .

Fix a mechanism M with interim allocation rule σ . We can now define the core of a mechanism as follows:

$$\text{CORE}(M, \beta) = \mathbf{E}_{\mathbf{v}} \left[\sum_{i,S} \sigma_S^i(v_i) v_i(S \cap C_i(v_i)) \right].$$

That is, the core with respect to M and β is the total welfare generated by M , excluding any item j assigned to a buyer i that has too large a value for it, relative to β_{ij} .

We can interpret this more explicitly as welfare under a valuation transformation. Define $v'_i(S) = v_i(S \cap C_i(v_i))$. Note that with this definition, omitting the dependence of σ and \mathbf{v}' on \mathbf{v} from the notation, $\text{CORE}(M, \beta) = \mathbf{E} \left[\sum_{i,S} \sigma_S^i v'_i(S) \right]$. An important observation due to Cai and Zhao [15] is the following:

Lemma V.1 (Cai and Zhao [15]). *Valuations \mathbf{v}' are subadditive over independent items.*

The following lemma summarizes the implications of a construction due to Cai and Zhao [15]

Lemma V.2 (Cai and Zhao [15]). *Fix an arbitrary constant $b \in (0, 1)$. Then there exists a mechanism M with interim allocation rule σ and corresponding item allocation rules π and personalized item thresholds β_{ij} for all i and j such that the following are true:*

- 1) $\text{REV} \leq \left(\frac{8}{1-b} + 12\right) \cdot \text{POSTREV} + 4 \cdot \text{CORE}(M, \beta)$,
- 2) $\sum_{i=1}^n \frac{c_i}{2} \leq \frac{2}{1-b} \cdot \text{POSTREV}$,
- 3) $\sum_{k \neq i} \Pr_{v_k} [v_k(j) \geq \beta_{kj}] \leq b$,
- 4) $\mathbf{E}_{v_i} [\pi_j^i(v_i)] \leq \Pr_{v_i} [v_i(j) \geq \beta_{ij}] / b$.

From now on we'll fix b and consider the corresponding mechanism M and thresholds β from Lemma V.2, and we'll simply write CORE to mean $\text{CORE}(M, \beta)$. Given Lemma V.2, what remains in order to prove Theorem V.1 is to argue that CORE can be approximated (up to a factor of $O(\log \log m)$) by either POSTREV or APOSTENREV.

B. The Core Within the Core, Key Lemma, and How it Implies the Bound

The CORE as defined above is just the welfare for a specific allocation under transformed but still subadditive valuation functions (namely the v'_i 's). A key idea in the literature on simple, near-optimal posted price mechanisms for revenue is to turn posted-price mechanisms that achieve some approximation guarantee for welfare into mechanisms that achieve the same (up to constant factors) approximation guarantee for revenue by augmenting the mechanism with entry fees. The idea is that if each buyer's surplus is sufficiently concentrated, then that surplus can be extracted as revenue using entry fees. One way to show concentration is to argue that no single item's contribution to the surplus is too large (i.e., a Lipschitz condition), but as it turns out individual items can contribute significantly to surplus under valuations v'_i .

Cai and Zhao [15] therefore invoke a second core restriction (called $\widehat{\text{CORE}}$ in their paper), resulting in a further restricted valuation \hat{v}_i for each agent. This further restricted valuation has a sufficiently small Lipschitz constant.

To define this second restriction, consider a set of item prices p_j (we will fix these later). Then, for each agent i , define

$$\tau_i = \inf \left\{ x : \sum_j \Pr [v_i(j) \geq \max\{\beta_{ij}, p_j + x\}] \leq 1/2 \right\}.$$

Write $Y_i(v_i) = \{j : v_i(j) < p_j + \tau_i\}$.

With these definitions in place, we can formalize what we meant by a further restriction of each agent's valuation v_i . Define $\hat{v}_i(S) = v_i(S \cap Y_i(v_i))$. So \hat{v} is like v , but with any "very high-valued individual items" removed from the

valuation. Unlike the initial core decomposition, now the high-valuedness is with respect to prices p_j rather than the thresholds β_{ij} .

For each agent i , set S , and valuation v_i , set $\mu_i(v_i, S) = \max_{S' \subseteq S} (\hat{v}_i(S') - \sum_{j \in S'} p_j)$. This is the surplus enjoyed by an agent with valuation \hat{v}_i when S is the set of available items, priced according to p_j . (We use the notation $\mu_i(v_i, S)$ for consistency with [15]. It's not related to μ_T as it appears in the zero-sum games.)

The following lemmata summarize key properties of the core-within-the-core that are shown in [15] and that we will re-use in our analysis.

Lemma V.3 (Cai and Zhao [15]). *For any choice of item prices p_j and subadditive functions v_i , the surplus μ_i is monotone, subadditive, and τ_i -Lipschitz.*

Lemma V.4 (Cai and Zhao [15]). *For any choice of item prices p_j and subadditive functions v_i ,*

$$\sum_{i=1}^n \sum_{j \in M} \max\{\beta_{ij}, p_j + \tau_i\} \Pr_{v_i} [v_i(j) \geq \max\{\beta_{ij}, p_j + \tau_i\}] \leq \frac{2}{b(1-b)} \cdot \text{POSTREV}.$$

Lemma V.5 (Cai and Zhao [15]). *For any choice of item prices p_j and subadditive functions v_i , $\sum_i \tau_i \leq \frac{4}{1-b} \cdot \text{POSTREV}$.*

The crux and key innovation of our analysis is now the following lemma, which establishes the existence of an ASPE with an appropriate approximation guarantee.

Our mechanism is actually from the same class of ASPE mechanisms as the mechanism of Cai and Zhao [15], but uses a different set of item prices. This class of ASPE mechanisms, less call them *median ASPE*, is parametrized by a set of item prices p_j and the mechanisms within that class proceed as follows. The agents are approached sequentially in a fixed order. We write $S_i(\mathbf{v})$ for the set of items still available when agent i is approached. (Note that $S_i(\mathbf{v})$ only depends on the entries of \mathbf{v} corresponding to agents that arrived before i .) Agent i faces an entry fee equal to the median (over randomness in v_i) of $\mu_i(v_i, S_i(\mathbf{v}))$. If agent i pays the entry fee, she can then purchase any desired subset of items at prices p_j .

Lemma V.6. *There is a set of item prices p_j such that the median ASPE with these prices achieves expected revenue at least*

$$\text{APOSTENREV} \geq \frac{1}{4\alpha} \cdot \text{CORE} - \frac{1+6b}{2b(1-b)} \cdot \text{POSTREV}.$$

with $\alpha = O(\log \log m)$.

With Lemma V.6 at hand it's rather straightforward to show Theorem V.1. The idea is to distinguish between the case where $\text{POSTREV} \geq c/\alpha \cdot \text{CORE}$ where c is some constant, and the case where it doesn't. In the former case

Lemma V.2 implies a lower bound on POSTREV in terms of REV, and in the latter case it implies a lower bound on APOSTENREV in terms of REV. The respective bounds depend on α , b , and c . An appropriate choice of b and c allows us to balance the two cases, and yields the desired bound.

C. Proof of Key Lemma, Pricing Beyond Additive Supporting Functions

It remains to show Lemma V.6, which boils down to finding appropriate item prices. We would like to find item prices that guide the allocation toward one that approximates the allocation that defines $\text{CORE} = \mathbf{E}[\sum_{i,S} \sigma_S^i v'_i(S)]$. To this end we will extend Lemma III.3 (our key lemma in the incomplete information case) to allow arbitrary constraints on the marginal probability of allocating each item to each agent. Given $\mathbf{z} = \{z_{ij}\}$, where each z_{ij} is a function with $z_{ij}(v_i) \in [0, 1]$ for each i, j , and v_i , we'll write

$$\Lambda(\mathbf{z}) = \left\{ \{\lambda^{\mathbf{v}}\}_{\mathbf{v}} : \mathbf{E}_{\mathbf{v}-i} \left[\sum_{S \ni j} \lambda_S^{i,\mathbf{v}} \right] \leq z_{ij}(v_i) \ \forall i, j, v_i \right\}.$$

That is, $\Lambda(\mathbf{z})$ is the set of all collection of $\lambda^{\mathbf{v}}$'s satisfying the upper bounds described by \mathbf{z} .

Analogous to Lemma III.3 we can now show the following lemma.

Lemma V.7. *For every independent probability distribution $\mathcal{D} = \prod_i \mathcal{D}_i$ over subadditive valuation functions, and $z_{ij}(v_i) \in [0, 1]$ for each i and j and v_i , there exist prices p_j for $j \in M$ and probability distributions $\lambda^{i,\mathbf{v}}$ over $S \subseteq M$ for all i and \mathbf{v} such that $\lambda \in \Lambda(\mathbf{z})$ and, for all $T \subseteq M$,*

$$\begin{aligned} \sum_{j \in T} p_j + \mathbf{E}_{\mathbf{v}} \left[\sum_{i=1}^n \sum_S \lambda_S^{i,\mathbf{v}} (v'_i(S \setminus T) - \sum_{j \in S} p_j) \right] \\ \geq \frac{1}{\alpha} \cdot \mathbf{E}_{\mathbf{v}} \left[\max_{\lambda \in \Lambda(\mathbf{z})} \sum_{i,S} \bar{\lambda}_S^{i,\mathbf{v}} v'_i(S) \right], \end{aligned}$$

where $\alpha \in O(\log \log m)$.

Invoke this lemma with $z_{ij}(v_i) = \pi_j^i(v_i)$, resulting in some p_j and λ . Note then that the RHS of Lemma V.7, $\mathbf{E}_{\mathbf{v}}[\max_{\lambda \in \Lambda(\mathbf{z})} \sum_{i,S} \bar{\lambda}_S^{i,\mathbf{v}} v'_i(S)]$, is at least CORE.

We'd now like to claim that if we replace v'_i by \hat{v}_i in Lemma V.7, this does not change the welfare bound by too much.

Lemma V.8. *For any $\{\lambda^{\mathbf{v}}\}_{\mathbf{v}} \in \Lambda(\mathbf{z})$ with $z_{ij}(v_i) = \pi_j^i(v_i)$ and for all $T \subseteq M$,*

$$\begin{aligned} \mathbf{E}_{\mathbf{v}} \left[\sum_{i=1}^n \sum_S \lambda_S^{i,\mathbf{v}} v'_i(S \setminus T) \right] - \mathbf{E}_{\mathbf{v}} \left[\sum_{i=1}^n \sum_S \lambda_S^{i,\mathbf{v}} \hat{v}_i(S \setminus T) \right] \\ \leq \frac{2(1+b)}{b(1-b)} \cdot \text{POSTREV}. \end{aligned}$$

Similarly to the way we prove Theorem III.1 from Lemma III.3 we can now prove:

Lemma V.9. *Using the prices p_j that result from invoking Lemma V.7 with $z_{ij}(v_i) = \pi_j^i(v_i)$ we have*

$$\begin{aligned} \sum_{i=1}^n \mathbf{E}_{\mathbf{v}} [\mu_i(v_i, S_i(\mathbf{v}))] \geq \frac{1}{\alpha} \text{CORE} - \mathbf{E}_{\mathbf{v}} \left[\sum_{j \in \text{SOLD}(\mathbf{v})} p_j \right] \\ - \frac{2(1+b)}{b(1-b)} \cdot \text{POSTREV}. \end{aligned}$$

We also need the following result from [15] concerning the revenue collected by a median ASPE through the entry fees, which exploits that the μ_i are τ_i -Lipschitz. Recall that we used $S_i(\mathbf{v})$ to denote the set of items that are still available when agent i is approached.

Lemma V.10 (Cai and Zhao [15]). *For any choice of item prices p_j and subadditive valuations v_i , the expected revenue that the median ASPE that uses these prices collects through the entry fees is at least*

$$\frac{1}{4} \sum_{i=1}^n \mathbf{E}_{\mathbf{v}} [\mu_i(v_i, S_i(\mathbf{v}))] - \frac{5}{8} \sum_{i=1}^n \tau_i.$$

Lemma V.6 now follows by observing that APOSTENREV decomposes into the revenue from the posted prices, which is $\mathbf{E}_{\mathbf{v}} [\sum_{j \in \text{SOLD}(\mathbf{v})} p_j]$ by definition, and the revenue from the entry fees, which is lower bounded by $\frac{1}{4} \sum_{i=1}^n \mathbf{E}_{\mathbf{v}} [\mu_i(v_i, S_i(\mathbf{v}))] - \frac{5}{8} \sum_{i=1}^n \tau_i$ through Lemma V.10. Using Lemma V.9 to lower bound $\sum_{i=1}^n \mathbf{E}_{\mathbf{v}} [\mu_i(v_i, S_i(\mathbf{v}))]$ and Lemma V.5 to upper bound $\sum_{i=1}^n \tau_i$ the claim follows.

VI. GOING BEYOND $O(\log \log m)$

We leave it as an open problem whether the $O(\log \log m)$ factor could be further reduced, possibly even to a constant. As a matter of fact, many techniques presented in this paper seem to be very useful to reach such an improved guarantee. In this section we discuss to what extent they can be applied and where there are barriers.

In all of our proofs, the $O(\log \log m)$ -factor originates from variants of the same technical lemma. In particular, the complete-information proof of Theorem III.1 as provided in Section III-A is mainly based on Lemma III.1. Similar lemmas are used for all other proofs as well. This is why we will now revisit the proof of Lemma III.1.

The proof of Lemma III.1 shows that its statement for any value of α is indeed equivalent to there being a vector of probabilities $\mathbf{q} = (q_j)_{j \in U}$, $q_j \in [0, 1]$, one for each item such that the λ player has value at least $v_i(U)/\alpha$ in the zero-sum game induced by the \mathbf{q} vector. More formally, $g(\mathbf{q}) \geq v_i(U)/\alpha$ for the function $g(\mathbf{q}) = \max_{\lambda \in \Delta(\mathbf{q})} \min_{\mu \in \Delta(\mathbf{q})} \sum_{S, T \subseteq U} \lambda_S \mu_T v_i(S \setminus T)$, where $\Delta(\mathbf{q})$ denotes all probability distributions ν over sets $S \subseteq U$ such that $\sum_{T \ni j} \nu_T \leq q_j$ for all j . This also means that to show the existence of prices for an $o(\log \log m)$ -approximation it suffices to show that there always is a vector \mathbf{q} such that $g(\mathbf{q}) \geq v_i(U)/\alpha$ for $\alpha \in o(\log \log m)$.

Our proof continues by Lemma III.2, showing that there exists a $q \in [0, 1]$ such that for $\mathbf{q} = (q, \dots, q) \in [0, 1]^{|U|}$ we have $g(\mathbf{q}) \geq \frac{1}{\alpha} v_i(U)$ for $\alpha \in O(\log \log m)$. That is, we put the same probability mass q on every item. As we will show now, the $O(\log \log m)$ bound is indeed tight for strategies that put the same probability mass q on all items. In other words, for a $o(\log \log m)$ bound, one would have to devise a more sophisticated way to choose the \mathbf{q} vector.

Theorem VI.1. *There exists a subadditive function v_i such that for $U = M$ we have $\max_{q \in [0,1]} g(q, \dots, q) \leq \frac{1}{\Omega(\log \log m)} v_i(M)$.*

REFERENCES

- [1] U. Krengel and L. Sucheston, "Semiamarts and finite values," *Bull. Am. Math. Soc.*, vol. 83, pp. 745–747, 1977.
- [2] —, "On semiamarts, amarts, and processes with finite value," *Adv. in Prob. and Relat. Top.*, vol. 4, pp. 197–266, 1978.
- [3] E. Samuel-Cahn, "Comparison of threshold stop rules and maximum for independent nonnegative random variables," *Ann. Probab.*, vol. 12, pp. 1213–1216, 1984.
- [4] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan, "Multi-parameter mechanism design and sequential posted pricing," in *STOC*, 2010, pp. 311–320.
- [5] R. Kleinberg and S. M. Weinberg, "Matroid prophet inequalities," in *STOC*, 2012, pp. 123–136.
- [6] M. Feldman, N. Gravin, and B. Lucier, "Combinatorial auctions via posted prices," in *SODA*, 2015, pp. 123–135.
- [7] P. Dütting and R. Kleinberg, "Polymatroid prophet inequalities," in *ESA*, 2015, pp. 437–449.
- [8] M. Feldman, O. Svensson, and R. Zenklusen, "Online contention resolution schemes," in *SODA*, 2016, pp. 1014–1033.
- [9] A. Rubinstein, "Beyond matroids: Secretary problem and prophet inequality with general constraints," in *STOC*, 2016, 324–332.
- [10] A. Rubinstein and S. Singla, "Combinatorial prophet inequalities," in *SODA*, 2017, pp. 1671–1687.
- [11] P. Dütting, M. Feldman, T. Kesselheim, and B. Lucier, "Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs," *SIAM J. Comput.*, vol. 49, no. 3, pp. 540–582, 2020.
- [12] S. Chawla, J. B. Miller, and Y. Teng, "Pricing for online resource allocation: Intervals and paths," in *SODA*, 2019, pp. 1962–1981.
- [13] U. Feige, "On maximizing welfare when utility functions are subadditive," *SIAM J. Comput.*, vol. 39, no. 1, pp. 122–142, 2009.
- [14] M. Feldman, H. Fu, N. Gravin, and B. Lucier, "Simultaneous auctions are (almost) efficient," in *STOC*, 2013, pp. 201–210.
- [15] Y. Cai and M. Zhao, "Simple mechanisms for subadditive buyers via duality," in *STOC*, 2017, pp. 170–183.
- [16] Y. Cai, N. R. Devanur, and S. M. Weinberg, "A duality based unified approach to bayesian mechanism design," in *STOC*, 2016, pp. 926–939.
- [17] S. Dobzinski, "Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders," in *STOC*, 2016, pp. 940–948.
- [18] S. Assadi and S. Singla, "Improved truthful mechanisms for combinatorial auctions with submodular bidders," in *FOCS*, 2019, pp. 233–248.
- [19] G. Christodoulou, A. Kovács, and M. Schapira, "Bayesian combinatorial auctions," *J. ACM*, vol. 63, no. 2, pp. 1–19, 2016.
- [20] V. Syrgkanis and É. Tardos, "Composable and efficient mechanisms," in *STOC*, 2013, pp. 211–220.
- [21] S. Ehsani, M. Hajiaghayi, T. Kesselheim, and S. Singla, "Prophet secretary for combinatorial auctions and matroids," in *SODA*, 2018, pp. 700–714.
- [22] H. Zhang, "Improved prophet inequalities for combinatorial welfare maximization with (approximately) subadditive agents," in *ESA*, 2020, pp. 82:1–82:17.
- [23] S. Chawla, J. D. Hartline, and R. Kleinberg, "Algorithmic pricing via virtual valuations," in *EC*, 2007, pp. 243–251.
- [24] S. Chawla, D. L. Malec, and B. Sivan, "The power of randomness in bayesian optimal mechanism design," in *EC*, 2010, pp. 149–158.
- [25] S. Hart and N. Nisan, "Approximate revenue maximization with multiple items," *J. Econ. Theory*, vol. 172, pp. 313 – 347, 2017.
- [26] X. Li and A. C.-C. Yao, "On revenue maximization for selling multiple independently distributed items," *Proc. Natl. Acad. Sci. U.S.A.*, vol. 110, no. 28, pp. 11 232–11 237, 2013.
- [27] M. Babaioff, N. Immorlica, B. Lucier, and S. M. Weinberg, "A simple and approximately optimal mechanism for an additive buyer," in *FOCS*, 2014, pp. 21–30.
- [28] A. C.-C. Yao, "An n-to-1 bidder reduction for multi-item auctions and its applications," in *SODA*, 2015, pp. 92–109.
- [29] S. Chawla and J. B. Miller, "Mechanism design for subadditive agents via an ex ante relaxation," in *EC*, 2016, pp. 579–596.
- [30] A. Rubinstein and S. M. Weinberg, "Simple mechanisms for a subadditive buyer and applications to revenue monotonicity," *ACM Trans. Econ. Comput.*, vol. 6, no. 3-4, pp. 1–25, 2018.
- [31] T. Roughgarden, "Intrinsic robustness of the price of anarchy," *J. ACM*, vol. 62, no. 5, p. 32, 2015.
- [32] M. Hardt, B. Recht, and Y. Singer, "Train faster, generalize better: Stability of stochastic gradient descent," in *ICML*, 2016, pp. 1225–1234.