

# Online Combinatorial Allocations and Auctions with Few Samples

Paul Dütting\*, Thomas Kesselheim†, Brendan Lucier‡, Rebecca Reiffenhäuser§, Sahil Singla¶

\*Google Research, Zurich, Switzerland. Email: duetting@google.com

†University of Bonn, Germany. Email: thomas.kesselheim@uni-bonn.de

‡Microsoft Research New England, Cambridge, MA, USA. Email: brlucier@microsoft.com

§University of Amsterdam, The Netherlands. Email: r.e.m.reiffenhauser@uva.nl

¶Georgia Tech, Atlanta, GA, USA. Email: ssingla@gatech.edu

**Abstract**—In online combinatorial allocations/auctions,  $n$  bidders sequentially arrive, each with a combinatorial valuation (such as submodular/XOS) over subsets of  $m$  indivisible items. The aim is to immediately allocate a subset of the remaining items to maximize the total *welfare*, defined as the sum of bidder valuations. A long line of work has studied this problem when the bidder valuations come from known independent distributions. In particular, for submodular/XOS valuations, we know 2-competitive algorithms/mechanisms that set a fixed price for each item and the arriving bidders take their favorite subset of the remaining items given these prices. However, these algorithms traditionally presume the availability of the underlying distributions as part of the input to the algorithm. Contrary to this assumption, practical scenarios often require the learning of distributions, a task complicated by limited sample availability. This paper investigates the feasibility of achieving  $O(1)$ -competitive algorithms under the realistic constraint of having access to only a limited number of samples from the underlying bidder distributions.

Our first main contribution shows that a mere single sample from each bidder distribution is sufficient to yield an  $O(1)$ -competitive algorithm for submodular/XOS valuations. This result leverages a novel extension of the secretary-style analysis, employing the sample to have the algorithm compete against itself. Although online, this first approach does not provide an online truthful mechanism. Our second main contribution shows that a polynomial number of samples suffices to yield a  $(2+\epsilon)$ -competitive online truthful mechanism for submodular/XOS valuations and any constant  $\epsilon > 0$ . This result is based on a generalization of the median-based algorithm for the single-item prophet inequality problem to combinatorial settings with multiple items.

**Index Terms**—Online algorithms, prophet inequality, posted-price mechanisms

## I. INTRODUCTION

Online combinatorial allocation is a fundamental problem in stochastic optimization, capturing a wide variety of resource allocation problems. In such problems, requests arrive online for a common pool of indivisible items/resources. Each request  $i$  can be assigned some combination of items/resources. The satisfaction of the request is modeled by a valuation function  $v_i$  (e.g. unit-demand, submodular, or XOS), drawn independently

Full version available at: <https://arxiv.org/abs/2409.11091>. T. Kesselheim has been funded by the Federal Ministry of Education and Research of Germany and the state of North Rhine-Westphalia as part of the Lamarr Institute for Machine Learning and Artificial Intelligence, Bonn, Germany. The work of S. Singla was supported in part by NSF award CCF-2327010.

from some probability distribution  $D_i$ , quantifying the value of different combinations. A decision-maker must immediately and irrevocably allocate to each request a subset of the available items/resources, while maximizing the total sum of valuations. The challenge is that allocation decisions must be made online without perfect foreknowledge of future requests. Such online allocation problems capture wide-ranging applications, from cloud computation scheduling to e-commerce, and have attracted substantial attention from the theory community [e.g., 1, 2, 3, 4].

Importantly, the performance achievable by the decision-maker depends crucially on how much information they have about the request distributions  $D_i$ . If nothing is known (i.e., the distributions are adversarial) then no bounded approximation to the offline optimal solution is possible even for allocating a single item. On the other end of the spectrum, when valuations are drawn independently from *known* distributions (the so-called prophet setting), sweeping positive results are possible. For example, in a celebrated result, Feldman, Gravin, and Lucier [1] have shown that when the valuations  $v_i$  are submodular (or even XOS), it is possible to achieve a (tight) factor 2-approximation against the expected offline optimum.

In the prophet setting, many known solutions for online combinatorial allocation have the additional desirable property of being *truthful*, meaning that a self-interested requester (or bidder) can obtain no benefit by strategically manipulating their reported valuation  $v_i$ . This is a non-trivial property even when the distributions are fully known in advance, since classical solutions from economics such as the VCG auctions [5, 6, 7] do *not* work for online bidders. The result of Feldman, Gravin, and Lucier [1], for example, achieves truthfulness by taking the form of a “posted-price mechanism.” This type of mechanism uses the advance knowledge of the valuation distributions to compute fixed item prices, and achieves truthfulness by letting each arriving bidder buy a utility-maximizing bundle of items.

The two extremes discussed so far, no information about the distributions or full information, are respectively too pessimistic and overly optimistic. Indeed, in real-world scenarios, one would expect to have some (but not perfect) distributional knowledge through observational data. Motivated by this, Azar, Kleinberg, and Weinberg [8] suggest taking a sample complexity perspective and seek guarantees that apply when the

underlying distributions are known only through a limited number of samples. This has led to a sequence of efforts [9, 10, 11] expanding sample-based prophet inequalities for allocation problems, but fairly little was known about general submodular (and, more generally, XOS) valuations. To address this gap, we ask:

*For submodular (or even XOS) valuations, how many samples from the distributions suffice to obtain constant-competitive online allocations? How many samples are needed if we additionally require the algorithm to be truthful?*

Given existing work that solves the online allocation problem for known distributions, a natural approach involves first learning the distribution, followed by algorithm creation tailored to the learned distribution. Yet, given the inherent imperfection in distribution learning, this strategy often culminates in suboptimal performance. Indeed, our investigation indicates that prevailing posted-price mechanisms are intrinsically reliant on precise knowledge of distributions. This dependency, in essence, stems from the need to learn the “mean” of certain random variables, a feat unattainable with a limited sample set unless we adopt stringent assumptions, such as bounded support for valuations.

An alternative approach, which is rooted in the realm of Data-Driven Algorithm Design [12], proposes bypassing distribution learning and instead directly learning algorithms from samples. This approach, especially when applied to online settings with combinatorial assignments, introduces significant complexities even when we set aside truthfulness and concentrate solely on constructing an online allocation algorithm to optimize welfare.

In response to the above questions, our contributions are twofold:

- 1) *Online Combinatorial Allocation:* In the non-strategic setting with request valuations that are submodular (or even XOS), we establish that even a single sample from each distribution suffices to obtain an  $O(1)$ -competitive online allocation algorithm. This presents as significant advancement, given that prior single-sample results were limited to specific submodular valuations like unit-demand or budget-additive [10, 11].
- 2) *Posted-Price Mechanisms:* We prove that, given bidder valuations that are submodular (or more generally XOS), a polynomial number of samples can effectively inform the choice of fixed item prices that yield  $(2 + \epsilon)$ -competitive posted-price mechanisms for any constant  $\epsilon > 0$ . Given the intrinsic unlearnability of a random variable’s “mean”, this achievement is anchored in pioneering a “median” pricing strategy for combinatorial settings—potentially a topic of independent interest. Prior to this work, it was only known how to obtain truthful, constant-factor approximations with poly-many samples under stringent assumptions on the distributions, such as boundedness [1].

#### A. Median Pricing and Secretary-Style Analysis for Single Item

To build intuition for our results, we examine the scenario of a single-item allocation/auction, analogous to the single-choice prophet inequality problem [13, 14]. Here, we encounter  $n$  bidders arriving sequentially, each with a valuation  $v_i$  for the item drawn from an independent distribution. We need to design an online algorithm to maximize the value of the bidder who receives the item. A well-known strategy achieves a 2-competitive ratio by assigning the item to the first bidder whose valuation exceeds half the expected maximum value [15]. The posted-price mechanisms introduced by Feldman, Gravin, and Lucier [1] extend this mean-based philosophy to the multi-item, combinatorial auction framework.

However, the landscape changes when we possess only samples of the bidders’ valuation distributions. Learning the mean of such distributions is infeasible with a sample-based approach, particularly for distributions with small, yet pivotal probabilities. For instance, consider an instance with  $n = 2$  bidders where  $v_1 = 1$  w.p. 1 and  $v_2 = 1/\epsilon$  w.p.  $\delta$  and is 0 otherwise, where  $\epsilon, \delta$  are small unknown constants. We need at least  $\approx 1/\delta$  samples to see a non-zero  $v_2$ ; however,  $\epsilon, \delta$  could be arbitrarily small, so we cannot (approximately) learn  $\mathbb{E}[\max\{v_1, v_2\}]$  for this distribution.

Nonetheless, even with limited samples, achieving  $O(1)$ -competitive algorithms is plausible through two principal methods, both circumventing the need to ascertain the distribution means and instead focusing on constant probability selection of bids.

**Secretary-Style Pricing.** One methodology is to select a threshold based on the maximum value among a single sample from each bidder’s distribution [e.g., 8, 9]. The reason this algorithm is  $O(1)$ -competitive is due to a Secretary-Style argument: among all samples and real values, with constant probability, the maximum is among the real values and the second-highest is among the samples. Although it has a simple analysis for single-item, the difficulty of extending this to combinatorial valuations is unclear since there is no natural notion of “maximum” and “second-highest” due to item dependencies. Our single-sample  $O(1)$ -competitive online combinatorial allocation algorithms are based on generalizing this secretary-style approach.

**Median Pricing.** Alternatively, we could pivot from mean to median valuation estimations. A classic result of Samuel-Cahn [16] shows that selecting the first bidder with value above the median  $\tau$  is 2-competitive, i.e.,  $\tau$  is the solution to  $\mathbb{P}[\max v_i \geq \tau] = 1/2$ . With a finite sample set, a threshold  $\hat{\tau}$  can be inferred, which aligns the maximum valuation probability within an acceptable  $\epsilon$  range of the ideal half, thus yielding a  $2 + O(\epsilon)$  competitive algorithm. However, the challenge in extending median-based approaches to combinatorial auctions is that it’s not obvious how to define median for combinatorial valuations due to item dependencies. Our polynomial-sample  $O(1)$ -competitive posted-price mechanism for online combinatorial auctions is based on generalizing this median pricing approach.

### B. Online Combinatorial Allocation

Our first result extends existing single-sample methodologies for the online allocation problem (such as the one for a single item discussed above) all the way to submodular (and even XOS) valuations. Note that a single sample from the product distribution over bidder valuations corresponds to drawing one sample from each of the bidder's (independent) valuation distributions.

**Theorem I.1.** *Given a single sample from any product distribution over submodular (or even XOS) bidder valuations, there exists an  $O(1)$ -competitive online algorithm for the combinatorial allocation problem.*

We actually prove an even stronger result, showing that the same guarantee can be obtained in the Game of Googol model, introduced in [17], with  $k = 1$  sample per buyer (Theorem III.1). In this model, an adversary inscribes  $k + 1$  valuation functions on  $k + 1$  facets of a die for each buyer. The die is then tossed, revealing  $k$  valuation functions to the online algorithm, while the  $(k + 1)$ -st—the true valuation—remains hidden. The algorithm must then make online allocations that are competitive with the best possible allocation in expectation over the die rolls. Any guarantee made in this model naturally extends to the independent samples model. Moreover, via a reduction from [11], this framework implies the first  $O(1)$ -competitive “order-oblivious” secretary algorithm for XOS combinatorial auctions: here, after receiving a constant fraction of the online elements, chosen uniformly at random, as a sample, the leftover elements arrive in adversarial order.

**Corollary I.2.** *There exists an  $O(1)$ -competitive order-oblivious secretary algorithm for the combinatorial allocation problem with submodular (and even XOS) buyers.*

Before describing our techniques, we discuss why simple approaches don't work. We outline the hardness examples here and discuss them in detail in Section C.

**Limitations of Prior Approaches.** Probably the most natural idea, inspired from single-item analysis, is to set the largest observed value in sample  $S$  as a threshold, i.e., item  $j$ 's price equals  $\max_{i \in S} v_i(j)$ . This fails when a single unit-demand buyer has a large value for every item but has only a small contribution to the optimum welfare. A second idea would be to compute the offline optimum allocation for the sample, then price each item according to its “supporting price” under that allocation.<sup>1</sup> This doesn't work (even for deterministic valuations) since a unit-demand buyer that contributes most of the welfare might now purchase a suboptimal item at such a high price.

Motivated by balanced pricing [1, 2], a third idea is to instead compute an offline optimal allocation on the single sample, and then set each item's price to be half of its supporting price. This also does not work. Suppose  $n = m$  and consider a distribution  $D$  over profiles of unit-demand valuations, such that player  $n$

<sup>1</sup>For unit-demand buyers, the supporting price of an item  $j$  allocated to agent  $i$  is  $v_i(j)$ . See Section II for the general definition.

has value 1 for one of the items (uniformly at random), and value  $\epsilon \ll 1/n$  for the remaining items. Each other player has value  $\epsilon$  for every item. Then this pricing approach will result in price  $1/2$  for one of the items (uniformly at random), say item  $j$ , and price  $\epsilon/2$  for all other items. On the real valuation, the optimal allocation has value  $1 + O(n\epsilon)$  but unless  $j = i$ , player  $n$  will receive no items, so the total value obtained will be  $O(n\epsilon)$ .

Given the limitations of prior approaches, we take a totally different approach by first considering a simplified setting where items can be reallocated to the current bidder.

**Greedy with Reallocations.** Consider a simplified setting in which we are permitted to reassign items previously allocated to the  $i$ -th bidder. This relaxation substantially simplifies the problem, allowing a straightforward constant approximation *bidder-wise greedy* algorithm, known since at least [18]. Here, each  $i$ -th bidder may claim any item but must compensate for the value lost by the item's prior owner. In essence, the  $i$ -th bidder encounters item prices  $\mathbf{p}^{(i)}$ , with each item's price reflecting its current contribution to the welfare as determined by its supporting XOS valuation, and may acquire any item (including previously allocated) by remunerating this price. This approach is effective because any highly-valued item can always be repurchased, ensuring that either the current or the previous bidder holds significant valuation for it.

**Two-Samples Algorithm.** Implementing the buyer-wise greedy strategy online without the luxury of reallocating items poses the principal challenge. Let's approach this by presuming the availability of two samples from each distribution. We start by using the first sample to establish item prices, employing the buyer-wise greedy algorithm on this sample to set the prices  $\mathbf{p}^{(i)}$  that the  $i$ -th real bidder faces. The underlying thought is that the ‘marginal’ value ascribed by both the  $i$ -th sample and the real  $i$ -th bidder should be analogous, since they face the same prices and are drawn from the same distribution. However, this alone does not preclude the possibility of item reallocation. Ideally, we wish to allocate the  $i$ -th bidder a subset of these items that are unlikely to be reassigned subsequently.

**Competing with Yourself: Secretary-Style Argument.** Our foremost insight is the conceptualization of the algorithm as *competing against itself*. While still utilizing the first sample to set item prices  $\mathbf{p}^{(i)}$  for the  $i$ -th bidder, we observe the price dynamics on the *second sample* with reallocations. Specifically, we determine which bidder secures each item  $j$  post-reallocations in the second sample, and use the supporting value of this bidder to set a ‘base price’  $b_j$  for every item  $j$ . As the actual  $i$ -th bidder arrives, they aim to select their optimal bundle  $B_i$  based on the initial sample prices  $\mathbf{p}^{(i)}$  (assuming full item availability), yet we only assign them a subset of these items that surpass the base prices  $b_j$  and remain unclaimed. The intuition is that since the algorithm on the second-sample and the real-sample are facing the same prices, a secretary-style argument should imply that with constant probability only the max-bidder for an item  $j$  will be allocated the item.

**Max vs. Sum.** Although utilizing the second-sample to set base prices averts reallocation, it poses a risk of diminishing welfare relative to the buyer-wise greedy model. More precisely, our Lemma III.6 only implies that the expected contribution of any item  $j$  to the welfare of the algorithm is at least a constant fraction of the *maximum* of the max  $j$ -bid in the real valuations and in the second sample.. Conversely, we only know that the *sum* of the  $j$ -bids in the real valuations and the second sample is at least a constant fraction of the buyer-wise greedy algorithm. Our key technical contribution, encapsulated in Lemma III.7, demonstrates that rounding prices to powers-of-two allows the max and the sum of bids to be comparable. The lemma's proof employs a carefully constructed Martingale argument to circumvent dependencies that might emerge between the samples and actual valuations, such as those arising from base pricing.

Finally, we extend this proof to single sample by observing that any  $k$ -sample result can be turned into a single sample result by losing only an  $O(k)$  term (Theorem III.10).

### C. Posted-Price Mechanisms

Our second main result addresses the design of truthful posted-price mechanisms using only sample access to bidder distributions. As mentioned earlier, posted-price mechanisms are highly desirable due to their practical appeal. Existing posted-price mechanism designs depend on mean-based pricing, which is untenable to implement with unknown and potentially unbounded distributions. Our single-sample approach from the previous section is inherently non-truthful since it provisionally assigns a utility-maximizing set of items to a bidder assuming all items are available, but ultimately only assigns them a subset of the items that are still available.

We bypass these issues by establishing a median-based pricing system for combinatorial settings. In the single-item case, the median price  $\tau$  satisfies the condition  $\mathbb{P}[\max v_i \geq \tau] = 1/2$ . In combinatorial settings, we will likewise seek prices such that each item sells with probability  $1/2$ . We then use median-based pricing to design posted-price mechanisms from samples.

**Theorem I.3.** *Given polynomially many samples from any product distribution over submodular (or even XOS) bidder valuations and a constant  $\epsilon > 0$ , there exists a  $(2 + \epsilon)$ -competitive truthful posted-price mechanism for online combinatorial auctions.*

We note that the approximation factor in this theorem is nearly tight. This is because even when considering the sale of a single item with known product distributions, online combinatorial auctions capture single item prophet inequality, where the tight competitive ratio is 2.

An important subtlety in posted-price mechanisms is the handling of agent indifferences. Even with a single item, if the value distributions have atoms, the exact probability of sale may vary based on an agent's decision-making when the price matches her valuation. To navigate the complexities tied to these tie-breaking scenarios, we initially consider *generic* valuation distributions. This term refers to a generalization of

the atomless condition, ensuring that for any specified price vector  $\mathbf{p}$ , occurrences of indifference are extremely unlikely. This assumption is practical and realistic, as all distributions become generic with the introduction of even minimal noise. Nonetheless, we later extend our results to include general distributions. For the sake of clarity, we will temporarily set aside the issue of tie-breaking.

**Combinatorial Median Prices.** The median price in single-item sales ensures that the item is sold exactly half the time, and unsold for the remainder. We extend this concept to define a median price vector  $\mathbf{p}$  for combinatorial settings such that using  $\mathbf{p}$  in a posted-price mechanism ensures every item  $j$  is sold with a probability of exactly one-half. The existence of such a median price vector is not immediately apparent, but in Section IV-A we apply the Kakutani fixed-point theorem to affirm that median price vectors always exist. The reason median prices are interesting is because they possess several amazing properties that we describe next.

**Posted-Pricing with Median Prices.** The first nice property of median prices that we discover is that they yield a posted-price mechanism that is  $O(1)$ -competitive for auctions with submodular (and even XOS) buyers. This result greatly generalizes the classical Optimal Stopping Theory result of Samuel-Cahn [16] for the case of a single-item. The high-level intuition why median prices help is that they ensure any item  $j$  will be available (not sold before) for its optimal bidder  $i$  with at least  $1/2$  probability. Thus, if item  $j$ 's median price is low then bidder  $i$  can easily purchase it to imply good welfare. Else, when  $j$ 's median price is very high, we can again use the median-price property that  $j$  gets sold with at least  $1/2$  a probability, so the revenue (and hence welfare) from  $j$  should be large. The adaptation of this intuitive understanding to a formal proof requires careful consideration; e.g., since the bidders are stochastic, the optimal bidder  $i$  for any fixed item  $j$  is a random bidder and we cannot directly apply the above proof argument.

**Learning Median Prices.** Another beneficial characteristic of median prices is their learnability in combinatorial auction settings with submodular/XOS valuations. We establish that a polynomial number of samples are sufficient to simultaneously approximate the sale probability for each item associated with every price vector  $\mathbf{p}$  within a small margin of  $\pm\epsilon$ . This proof employs uniform convergence principles and involves bounding the VC dimension for learning the performance of all price vectors. Our argument for bounding the VC dimension builds on the techniques of [19, 20]. Consequently, we demonstrate that the median prices based on empirical data closely estimate those for the underlying unknown distributions.

Finally, we shift our focus to the computational aspects of implementing posted-price mechanisms with combinatorial median prices. To illustrate the usefulness of our approach, we present efficient fully polytime algorithms for unit-demand bidders.

#### D. Further Related Work

**Prophet Inequality and Combinatorial Extensions.** Our work and questions are closely related to the literature on prophet inequalities. The classic prophet inequality problem corresponds to online combinatorial auction with a single item (i.e.,  $m = 1$ ). Prophet inequalities were greatly studied in the Optimal Stopping community in the late 70s early 80s, and the basic prophet inequality problem was solved by Krengel and Sucheston (who also credit Garling) [13, 14] by giving a tight 2-approximation. Later, Samuel-Cahn [16] showed that this optimal guarantee can be obtained with simple price-based algorithms.

Pioneering work in Computer Science by [21] and [22] demonstrated the usefulness of the prophet inequality paradigm for algorithmic mechanism design. The connection motivated extensions of the prophet inequality paradigm to combinatorial settings such as matroids [15, 23] or combinatorial auctions [1, 2]. These extensions also established the two main frameworks for proving prophet inequalities, namely *online contention resolution* [23] and *balanced pricing* [15, 1, 2].

Most relevant to our work is the aforementioned work of [1], who give a tight 2-approximation for combinatorial auctions with XOS buyers. This algorithm is based on static item prices, and is therefore truthful. The same approach yields a  $O(\log m)$  approximation for subadditive buyers. Two recent breakthroughs for combinatorial auctions with subadditive buyers, first improved this bound to  $O(\log \log m)$  [3], and then to  $O(1)$  [4].

Very recently, [24] introduced a general approach for turning any prophet inequality algorithm into a pricing-based truthful mechanism. For combinatorial auctions, given a non-truthful algorithm as input, this implies a bundle-pricing algorithm (instead of item pricing). This reduction requires a polynomial number of samples in the size of outcome space, which is exponential in  $m$  for combinatorial auctions.

**Prophet Inequalities with Few Samples.** The study of prophet inequalities from samples was pioneered in [8]. Their main result is an approximation-factor preserving reduction from single-sample prophet inequalities to *order-oblivious* secretary problems. See the book chapter [25] for further discussion on order-oblivious secretary algorithms. Using this, [8] obtain constant-sample  $O(1)$ -competitive prophet inequalities for several (restricted) matroid settings and for constant-degree bipartite matching with edge arrivals. They also obtain a 4-competitive single-sample algorithm for the single choice problem. An  $e$ -competitive secretary algorithm for the online combinatorial allocation problem with XOS buyers was given in [26]. However, prior to Corollary I.2, no order-oblivious secretary algorithm was known for this problem.

Starting with [27] and [9] a significant amount of work has examined the single-choice/single-item prophet inequality problem with samples. Key findings of this line of work include that the factor-2 single-choice prophet inequality can be attained with only a single sample per buyer (even in the stronger Game of Googol model) [9], while a constant number of samples

suffice to get arbitrarily close to the worst-case approximation guarantees for the iid version, the random order, and free order variant of the problem [9, 28, 29].

In independent work, [30] introduced the *competitive analysis with a sample* framework. In this framework, an adversary first writes down  $m = h + n$  numbers. Afterwards, a randomly chosen subset of size  $h$  of these numbers is revealed to the algorithm, who then either observes the remaining  $n$  numbers in adversarial order or in random order. The goal is to be competitive with the expected maximum value in this second set of numbers. For the case where  $h/n \geq 1$ , they give a tight 2-approximation for the adversarial order version of the problem. For the case where the actual values arrive in random order they give a  $(e + 1)/e$  approximation.

Closely related to our work, [11] study the single-sample problem for combinatorial settings. They present a general framework for deriving single sample results from greedy algorithms and obtain  $O(1)$ -competitive prophet inequalities for matchings with vertex or edge arrivals. Beyond this, they obtain improved single-sample constant-factor guarantees for a variety of (restricted) matroids and for combinatorial auctions with budget additive bidders. Budget-additive valuations are a strict subclass of submodular valuations, which are themselves a strict subclass of XOS valuations. In another closely related independent paper, [10] (also see the arXiv version for improved bounds) show  $O(1)$ -competitive algorithms across a range of the aforementioned settings.

For the general case of online combinatorial auctions with XOS buyers, an  $O(1)$ -approximation up to an additive error term of  $\epsilon$  can be obtained with  $\text{poly}(n, m, 1/\epsilon)$  samples [1, 2]. This multiplicative/additive approximation guarantee turns into a purely multiplicative guarantee only under stringent boundedness assumptions. In addition, a *non-truthful* purely multiplicative constant-factor approximation for XOS bidders with  $O(n)$  samples per bidder was known (e.g., [4]). The bounds of [31, 32] for XOS bidders arriving in a random order can be interpreted as giving a truthful  $\text{poly}(\log \log m)$ -approximation with a single sample. However, it's unclear how to extend them to an  $O(1)$ -approximation, even with more samples. Finally, [33] study prophet inequalities and posted-pricing when we are given distributions that are close to the actual distributions (in some metric).

## II. PRELIMINARIES

**Online Combinatorial Auctions.** We have a set of buyers  $[n]$ , a set of items  $[m]$ , and valuation functions  $v_i : 2^{[m]} \rightarrow \mathbb{R}_+$  that assign each buyer  $i \in [n]$  a value for each subset of items. We assume that the valuation functions are normalized, so that  $v_i(\emptyset) = 0$  for all  $i \in [n]$ , and monotone increasing, i.e., for every  $i \in [n]$ ,  $S, T \subseteq [m]$  with  $S \subseteq T$  it holds that  $v_i(S) \leq v_i(T)$ . We write  $\mathbf{v} = (v_1, v_2, \dots, v_n)$  for a valuation profile. We further assume that each valuation function  $v_i$  is independently drawn from an according, unknown distribution  $D_i$ , and call the tuple of all  $D_i$  distribution  $D$ .

An assignment of the items to the buyers is a collection of sets  $A_i \subseteq [m]$  such that  $A_i \cap A_j = \emptyset$  for all  $i, j \in [n]$  with

$i \neq j$ . We write  $\mathbf{A}$  for the tuple of  $A_i$  with  $i \in [n]$ , and  $\mathcal{A}$  for all possible allocations. The (social) welfare of an assignment is  $\sum_{i \in [n]} v_i(A_i)$ .

Our goal is to find a welfare-maximizing assignment of items to buyers, when items are available offline but buyers arrive online in some (previously known, but adversarially chosen) order. We assume for convenience that buyer  $i$  will arrive at time  $i$ , for  $i \in [n]$ .

**Competitive Ratio.** Our focus in this work is on online algorithms ALG for this problem that have sample access to  $D$ . For this we will assume that ALG receives  $k$  sample valuation functions for each buyer  $i \in [n]$ , sampled independently from  $D_i$ . We will typically refer to these as  $s_i^1, s_i^2, \dots, s_i^k$ , and use  $v_i$  for the actual valuation function. (We will also sometimes write  $S_i^1, S_i^2, \dots, S_i^k$  and  $V_i$  to emphasize that these are random variables.)

Let's denote by  $\text{ALG}_i(\mathbf{v}; \mathbf{s}^1, \dots, \mathbf{s}^k)$  the allocation to buyer  $i \in [n]$  on input  $\mathbf{v} = (v_1, v_2, \dots, v_n)$  and samples  $\mathbf{s}^\ell = (v_1^\ell, \dots, v_n^\ell)$  for  $\ell \in [k]$ . We say that ALG is  $\alpha$ -competitive if, for all distributions  $D$ , it holds that

$$\begin{aligned} & \mathbb{E}_{\mathbf{v} \sim D, (\mathbf{s}^1, \dots, \mathbf{s}^k) \sim D^k} \left[ \sum_i v_i(\text{ALG}_i(\mathbf{v}; \mathbf{s}^1, \dots, \mathbf{s}^k)) \right] \\ & \geq \alpha \cdot \mathbb{E}_{\mathbf{v} \sim D} \left[ \max_{\mathbf{A} \in \mathcal{A}} \sum_{i \in [n]} v_i(A_i) \right]. \end{aligned}$$

The existence of such an online algorithm ALG establishes an  $\alpha$ -competitive  $k$ -sample prophet inequality (or  $\alpha$ -competitive  $k$ SPI).

**Game of Googol.** We will also employ a somewhat stronger model, the Game of Googol with a  $(k+1)$ -faceted dice. In this model an adversary, for each buyer  $i \in [n]$ , writes down  $k+1$  valuation functions  $v_i^1, v_i^2, \dots, v_i^{k+1}$ , and then these valuation functions are assigned to  $S_i^1, S_i^2, \dots, S_i^k$  and  $V_i$  uniformly at random. The  $S_i^j$  are the samples, and  $V_i$  is the actual valuation. The goal is then to show that for any possible choice of the adversary,

$$\mathbb{E} \left[ \sum_i v_i(\text{ALG}_i(\mathbf{V}; \mathbf{S}^1, \dots, \mathbf{S}^k)) \right] \geq \alpha \cdot \mathbb{E} \left[ \max_{\mathbf{A} \in \mathcal{A}} \sum_{i \in [n]} v_i(A_i) \right].$$

All other aspects of the problem, including the fixed (worst-case) arrival order, remain untouched. We refer to such a result as an  $\alpha$ -competitive pointwise  $k$ -sample prophet inequality (or an  $\alpha$ -competitive P- $k$ SPI). Observe that the existence of an  $\alpha$ -competitive P- $k$ SPI implies an  $\alpha$ -competitive  $k$ SPI (simply because instead of choosing the valuations adversarially, we could also have drawn them from distributions  $D$ ).

**Classes of Valuations.** All our valuation functions  $v : 2^{[m]} \rightarrow \mathbb{R}_+$  are monotone, i.e., for any item bundles  $A \subseteq B \subseteq [m]$ , we have  $v(A) \leq v(B)$ . We mostly focus on XOS valuation functions [34], but also briefly touch upon the subclasses of submodular and unit-demand valuations:

- **XOS valuations:** Valuation function  $v : 2^{[m]} \rightarrow \mathbb{R}_+$  is XOS if there exist a collection of additive functions

$\{a_\ell : 2^{[m]} \rightarrow \mathbb{R}_+\}_{\ell=1, \dots, k}$  such that for each bundle of items  $B \subseteq [m]$  it holds that  $v(B) = \max_{\ell=1, \dots, k} a_\ell(B) = \max_{\ell=1, \dots, k} \sum_{j \in B} a_\ell(j)$ .

For an XOS valuation function  $v$  we refer to the additive function that defines  $v$  on the set of items  $B \subseteq [m]$  as the additive supporting function of  $v$  on  $B$ . Let's denote this function by  $a_B$ . Note that this function satisfies  $\sum_{j \in B'} a_B(j) \leq v(B')$  for all bundles  $B' \subseteq [m]$ , and  $\sum_{j \in B'} a_B(j) = v(B')$  for  $B' = B$ . We will also require that  $a_B(j) = 0$  for  $j \notin B$ .

- **Submodular valuations:** A valuation function  $v : 2^{[m]} \rightarrow \mathbb{R}_+$  is submodular if for any two item bundles  $A \subseteq B \subseteq [m]$  and any item  $j \in [m] \setminus B$ , we have  $v(A \cup j) - v(A) \geq v(B \cup j) - v(B)$ .
- **Unit-demand valuations:** A unit-demand valuation function  $v : 2^{[m]} \rightarrow \mathbb{R}_+$  is such that for each bundle of items  $B \subseteq [m]$ , we have  $v(B) = \max_{j \in B} v(j)$ .

We will denote the contribution of item  $j$  to the value of a bundle of items  $B \subseteq [m]$  given valuation function  $v$  as  $a_j(v, B)$ , and define this to be the value assigned to  $j$  by the additive supporting function of  $v$  on  $B$ .

**Prices, Utilities, and Demand.** For a vector of item prices  $\mathbf{p} = (p_1, \dots, p_m) \in \mathbb{R}_+^m$  and a set of items  $M \subseteq [m]$ , we write  $\mathcal{D}(v, \mathbf{p}, M)$  for the set of bundles  $B \subseteq M$  that maximizes the utility  $u^v(B, \mathbf{p}) = v(B) - \sum_{j \in B} p_j$ . We refer to  $\mathcal{D}(v, \mathbf{p}, M)$  as the demand at valuation  $v$  given prices  $\mathbf{p}$  when the set of available items is  $M$ . We use the shorthand  $\mathcal{D}(v, \mathbf{p})$  when  $M = [m]$ , or when  $M$  is clear from the context. (Note that we can always simulate that some item  $j \in [m]$  is not available by setting  $p_j$  high enough.)

**Oracle Access.** Since any explicit description of a XOS valuation function  $v : 2^{[m]} \rightarrow \mathbb{R}_+$  would have size that is exponential in  $m$ , we consider three types of oracle access:

- **Value oracle:** A value oracle takes as input a valuation function  $v$  and a bundle of items  $B \subseteq [m]$ , and returns  $v(B)$ .
- **Demand oracle:** A demand oracle takes as input a valuation function  $v$  and a vector of item prices  $\mathbf{p} = (p_1, \dots, p_m) \in \mathbb{R}_+^m$ , and returns a bundle of items  $B \in \mathcal{D}(v, \mathbf{p})$ .
- **XOS oracle (only defined for XOS valuation functions):** A XOS oracle takes as input a valuation function  $v$  and a bundle of items  $B \subseteq [m]$ , and returns the additive supporting function of  $v$  on  $B$ .

**Posted-Price Mechanisms and Truthfulness.** We say that an online algorithm is a posted-price algorithm if for each buyer  $i \in [n]$ , it proceeds as follows: Let  $M_i \subseteq [m]$  be the set of items that are still available when buyer  $i$  arrives. The algorithm defines a vector of prices  $\mathbf{p}^i = (p_1^i, \dots, p_m^i) \in \mathbb{R}_+^m$ , and lets buyer  $i$  choose a bundle of items  $B \subseteq M_i$  that maximizes their utility, i.e., a set  $B \in \mathcal{D}(v_i, \mathbf{p}^i, M_i)$ .

Posted-price algorithms where the vector of prices  $\mathbf{p}^i$  that bidder  $i$  sees does not depend on buyer  $i$ 's valuation  $v_i$  (but may depend on the valuations  $v_{i'}$  of buyers  $i' < i$ , as well

as on the distributions  $D = (D_1, \dots, D_n)$  of all buyers) are known to be dominant strategy incentive compatible (DSIC) — or truthful [e.g., 1]. Hence, they are also called sequential posted-price mechanisms.

### III. ONLINE COMBINATORIAL ALLOCATION WITH A SINGLE SAMPLE

In this section we design  $O(1)$ -competitive algorithms for online combinatorial auctions with XOS buyers, given a single sample from the distributions of the bidders. We will in fact show a slightly stronger result, namely that this can be achieved in the Game of Googol model. I.e., our algorithm is a *pointwise*  $O(1)$ -SPI.

**Theorem III.1** (Restatement of Theorem I.1). *The problem of online combinatorial auctions with XOS buyers admits a pointwise 576-competitive single-sample prophet inequality.*

We remark that we did not try to optimize the constant factor in Theorem I.1, and instead focused on the simplicity of exposition. Next, we first give an  $O(1)$ -competitive result assuming we have access to *two samples*, and later we modify this algorithm to a single sample algorithm.

#### A. Buyer-Wise Greedy

We start with defining a key ingredient of our prophet inequality, the buyer-wise greedy algorithm (due to Dobzinski, Nisan, and Schapira [18]). This algorithm simply assigns every new buyer their demand set when the prices are given by the items' current contributions to the overall solution.

**Buyer-Wise Greedy Algorithm.** We define  $G(\mathbf{v})$ , for valuation profile  $\mathbf{v}$ , as the assignment (and interchangeably, its value) determined as follows:

- For all  $i \in [n]$ , define  $\mathbf{p}^{(i)} = (p_1^{(i)}, \dots, p_m^{(i)})$  to be the prices faced by the  $i$ -th buyer.  
Set  $p_j^{(1)} = 0$  for all  $j \in [m]$ .
- For each buyer  $i = 1, \dots, n$ : (i.e., arriving in the fixed adversarial order)
  - 1) Assign  $i$  the bundle  $A_i = d_i(\mathbf{p}^{(i)}) \in \mathcal{D}(v_i, \mathbf{p}^{(i)}, [m])$ , i.e., some demand set at prices  $\mathbf{p}^{(i)}$ , with ties broken in an arbitrary but consistent manner. (Re-assigning items in  $A_{i'} \cap A_i$  from buyers  $i' < i$  to buyer  $i$ .)
  - 2) Raise prices by setting  $p_j^{(i+1)} = a_j(v_i, A_i)$  for all  $j \in A_i$ .  
(Keeping  $p_j^{(i+1)} = p_j^{(i)}$  for all  $j \notin A_i$ .)

The following lemma of [18] shows that the value attained by the buyer-wise greedy algorithm  $G(\mathbf{v})$  is a 2-approximation to  $\text{OPT}(\mathbf{v})$ . For completeness, we provide a proof in Appendix A.

**Lemma III.2** (Dobzinski, Nisan, and Schapira [18]). *For any XOS  $\mathbf{v}$ , the output of the buyer-wise greedy algorithm  $G(\mathbf{v})$  is a 2-approximation to the offline optimum  $\text{OPT}(\mathbf{v})$ .*

**Powers-of-Two Prices.** We modify the above procedure by making prices increase in an exponential fashion. Concretely, let

$$\mathcal{P} = \{2^k \mid k \in \mathbb{Z}\}.$$

Now, in the above greedy procedure  $G(\mathbf{v})$ , we will raise the prices of items in the demand set  $A_i$  to the closest power of two higher than  $a_j(v_i, A_i)$ . That is, in Step 2 of the algorithm's main loop, rather than setting  $p_j^{(i+1)} = a_j(v_i, A_i)$ , we instead set  $p_j^{(i+1)} = \min\{p \in \mathcal{P} : p \geq a_j(v_i, A_i)\}$  for all  $j \in A_i$ . Call the outcome of this modified procedure  $\hat{G}(\mathbf{v})$ .

**Lemma III.3.** *For any XOS  $\mathbf{v}$ , the output of the modified buyer-wise greedy algorithm  $\hat{G}(\mathbf{v})$  where prices increase as powers of two is a 3-approximation to the offline optimum  $\text{OPT}(\mathbf{v})$ .*

The proof of Lemma III.3 is a simple variation on the argument of [18] and appears in Section A.

**Decoupling Allocation and Pricing.** We conclude by observing that we can decouple allocation and pricing as follows. Suppose that we have two valuations  $f_i^1, f_i^2$  for each buyer  $i \in [n]$  and that these are assigned to samples  $\mathbf{s}$  and valuations  $\mathbf{r}$  by tossing an independent fair coin for each buyer. Let  $\mathbf{p}^{(i)}$  be the prices in the modified buyer-wise greedy run on  $\mathbf{s}$ , and let buyer  $i$  buy the set of items  $A_i$  that maximizes  $r_i(A_i) - \sum_{j \in B_i} p_j^{(i)}$ . Note how this decouples allocation from pricing: while  $\mathbf{s}$  is used to set and update prices,  $\mathbf{r}$  is used to make allocation decisions.

Analogously to Lemma III.3 we can now show the following lemma. Note that unlike Lemma III.3, however, the following lemma does *not* yet imply the existence of an  $O(1)$ -approximate allocation.

**Lemma III.4.** *Consider any pair of XOS valuations  $f_i^1, f_i^2$  for each buyer  $i \in [n]$  that are assigned independently and uniformly at random to  $\mathbf{s}, \mathbf{r}$ . Then, for the modified buyer-wise greedy algorithm, it holds  $\mathbb{E} \left[ \sum_{i \in [n]} r_i(A_i) \right] \geq \frac{1}{3} \mathbb{E} [\text{OPT}(\mathbf{s})]$ .*

The proof (which appears in Section A) exploits the way that  $\mathbf{s}$  and  $\mathbf{r}$  are generated, to argue that in expectation allocations and prices align. So that the approximation guarantee can be established in a similar way as that in Lemma III.3.

#### B. Two Samples Algorithm and Proof Overview

We are now ready to prove the following theorem, which claims existence of an online procedure that attains the guarantee in Theorem III.1 with two samples per buyer.

**Theorem III.5.** *The problem of online combinatorial auctions with XOS buyers admits a pointwise 192-competitive 2-sample prophet inequality.*

Recall what proving a statement like this entails: We need to consider a setting where an adversary writes down three valuation profiles  $\mathbf{v}^1, \mathbf{v}^2, \mathbf{v}^3$ . Then, for each buyer  $i \in [n]$  independently, the valuation functions  $v_i^1, v_i^2, v_i^3$  are assigned uniformly at random to samples  $\mathbf{s}, \mathbf{s}'$  and actual valuations

r. The goal is to show that there is an online algorithm ALG whose expected value when receiving samples  $\mathbf{s}, \mathbf{s}'$  and actual valuations  $\mathbf{r}$  is a  $O(1)$ -approximation to the expected offline optimum OPT on the actual valuations  $\mathbf{r}$ .

Consider the following online procedure:

---

**Two Samples Algorithm.** Given two samples  $\mathbf{s}$  and  $\mathbf{s}'$ , and arriving bidders with valuations  $\mathbf{r}$ :

- 1) Run the modified buyer-wise greedy algorithm  $\hat{\mathbf{G}}$  on the sample  $\mathbf{s}$ . For each buyer  $i \in [n]$ , let  $\mathbf{p}^{(i)}$  denote the prices set by this algorithm on input  $\mathbf{s}_{<i}$  (i.e., on input  $\mathbf{s}$ , but only considering the first  $i-1$  buyers).  
(We emphasize that these prices increase by powers of 2; this will be crucial for the max vs. sum argument in our analysis.)
- 2) For each buyer  $i \in [n]$  with valuation  $s'_i \in \mathbf{s}'$ , determine their demand bundle  $A'_i$  given prices  $\mathbf{p}^{(i)}$ . Define *base price*  $b_j$  for any item  $j \in [m]$  to be the largest contribution of  $j$  to any bundle  $A'_i$  with respect to  $s'_i$ . That is,  $b_j := \max_{i:j \in A'_i} a_j(s'_i, A'_i)$ .  
(Note that these base prices are not powers of 2. In fact, we need all supporting prices to be distinct for a secretary-style argument.)
- 3) For each arriving buyer  $i \in [n]$  with valuation  $r_i \in \mathbf{r}$ :
  - Compute their demand set  $B_i$  given the prices  $\mathbf{p}^{(i)}$ .
  - Assign them the still-available part of  $B_i \cap \{j \in [m] \mid a_j(r_i, B_i) > b_j\}$ .

Let us call the resulting allocation (and interchangeably, its value) ALG.

---

Our goal is to now prove that ALG for any fixed valuation profiles  $\mathbf{v}^1, \mathbf{v}^2, \mathbf{v}^3$  in expectation over the coin tosses that lead to  $\mathbf{s}, \mathbf{s}', \mathbf{r}$  is a constant-factor approximation to the offline optimum. Below, for an item  $j \in [m]$  assigned to buyer  $i \in [n]$  in ALG, we use the shorthand  $\text{ALG}_j = a_j(r_i, B_i)$  to denote item  $j$ 's contribution to the algorithm's welfare. (Note that we define  $\text{ALG}_j$  with respect to the additive supporting function for bundle  $B_i$ , so technically this is a lower bound on item  $j$ 's contribution to the welfare.)

**Proof Overview:** Our proof is broken into two lemmas.

The first lemma shows that the expected contribution of any item  $j$  to the algorithm's welfare is at least a constant fraction of  $j$ 's contribution to the max-bundle in  $\mathbf{r}$  and  $\mathbf{s}'$ . It is similar in spirit to arguments that appeared in [8] and [9] for the single-choice problem.

**Lemma III.6.** For any item  $j \in [m]$ ,

$$\mathbb{E}[\text{ALG}_j] \geq \frac{1}{4} \cdot \mathbb{E} \left[ \max \left( \max_{i:j \in B_i} a_j(r_i, B_i), \max_{i:j \in A'_i} a_j(s'_i, A'_i) \right) \right].$$

Our second lemma shows that the contributions of any item  $j$  to any demand bundle in the actual valuations  $\mathbf{r}$  or the second sample  $\mathbf{s}'$  are dominated by just the largest of them (up to a constant factor). This lemma and its proof is our key technical innovation.

**Lemma III.7.** For any item  $j \in [m]$ ,

$$\begin{aligned} \mathbb{E} \left[ \max \left( \max_{i:j \in B_i} a_j(r_i, B_i), \max_{i:j \in A'_i} a_j(s'_i, A'_i) \right) \right] \\ \geq \frac{1}{32} \cdot \mathbb{E} \left[ \left( \sum_{i:j \in B_i} a_j(r_i, B_i) + \sum_{i:j \in A'_i} a_j(s'_i, A'_i) \right) \right]. \end{aligned}$$

Together, the lemmas yield the constant competitive ratio, by summing up over all items  $j$  and using the fact that the demand bundles are determined in exactly the same way as those in the modified buyer-wise greedy algorithm (which is a const. approx.).

*Proof of Theorem III.5.* Summing over all items  $j \in [m]$  in Lemma III.6 and Lemma III.7 implies

$$\begin{aligned} \sum_{j \in [m]} \mathbb{E}[\text{ALG}_j] \\ \geq \frac{1}{128} \cdot \sum_{j \in [m]} \mathbb{E} \left[ \sum_{i:j \in B_i} a_j(r_i, B_i) + \sum_{i:j \in A'_i} a_j(s'_i, A'_i) \right] \\ = \frac{1}{128} \cdot \mathbb{E} \left[ \sum_{i \in [n]} (r_i(B_i) + s'_i(A'_i)) \right]. \end{aligned}$$

Note that in expectation the RHS is at least the sum of the expected welfare of  $\hat{\mathbf{G}}(\mathbf{r}; \mathbf{s})$  and  $\hat{\mathbf{G}}(\mathbf{s}'; \mathbf{s})$ . Hence, Lemma III.4 implies that the RHS is at least  $\frac{1}{128} \cdot \frac{2}{3} \cdot \mathbb{E}[\text{OPT}(\mathbf{r})]$ , which proves the theorem.  $\square$

### C. Proof of Lemma III.6: Secretary-Style Argument

In this proof we will assume no ties in the additive supporting valuations, which can be ensured by adding noise. We prove the lemma for any fixed sample  $\mathbf{s}$ , and in expectation over the assignment of the remaining two valuation functions  $v_i^1, v_i^2$  for each buyer  $i \in [n]$  to  $\mathbf{s}'$  and  $\mathbf{r}$ .

First fix sample  $\mathbf{s}$ , which fixes all prices  $\mathbf{p}^{(i)}$ . Note that this also fixes the demand bundles  $B_i(\mathbf{s}_{<i}, v_i^k)$  for all  $i \in [n]$  with respect to prices  $\mathbf{p}^{(i)}$  (set by  $\mathbf{s}_{<i}$ ) and valuations  $v_i^k$  for  $k \in \{1, 2\}$ . It also defines

$$a_j^{\max} := \max_k \left( \max_{i:j \in B_i} a_j(v_i^k, B_i(\mathbf{s}_{<i}, v_i^k)) \right).$$

Below we will show that with probability at least  $\frac{1}{4}$  (where the probability is over the assignment of  $v_i^1, v_i^2$  for  $i \in [n]$  to  $\mathbf{s}', \mathbf{r}$ ), the contribution of  $j$  to the algorithm,  $\text{ALG}_j$ , is at least  $a_j^{\max}$ .

Call the valuation that realizes  $a_j^{\max}$ , given it occurs at step  $i^*$ , simply  $v_{i^*}^{\max}$ . Now we determine whether  $v_{i^*}^{\max}$  belongs to  $\mathbf{s}'$  or to  $\mathbf{r}$  with a fair independent coin flip, therefore it will be in  $\mathbf{r}$  with probability  $1/2$ . Consider now the second-highest  $j$ -bid  $a_j^{2nd}$ . Either, it also belongs to buyer  $i^*$ , in which case if  $a_j^{\max} \in \mathbf{r}$ , it will belong to  $\mathbf{s}'$  with probability 1. Or, it belongs to some other buyer  $i \neq i^*$ , in which case another independent

fair coin decides its belonging, and it will end up in  $s'$  with probability  $1/2$ . All in all,

$$\Pr[a_j^{\max} \in \mathbf{r} \text{ and } a_j^{2\text{nd}} \in \mathbf{s}'] \geq \frac{1}{4}.$$

Assume indeed both is the case. Then,  $a_j^{2\text{nd}}$  will determine the cutoff/final price  $b_j$  of  $j$ , by definition. Also, by definition, among all  $\mathbf{r}$ -buyers attempting to obtain  $j$ , only  $i^*$  will be successful. Therefore, we actually obtain value  $a_j^{\max}$  for  $j$ . This yields  $\mathbb{E}[\text{ALG}_j] \geq \frac{1}{4}a_j^{\max}$ , where the expectation is over the coin tosses that decide for each buyer  $i \in [n]$  the assignment of  $v_i^1, v_i^2$  to  $\mathbf{s}', \mathbf{r}$ , respectively.

#### D. Proof of Lemma III.7: Handling Max vs. Sum

We first compare the largest contribution of item  $j$  in  $\mathbf{r}$  with the final price of item  $j$  as determined by  $\mathbf{s}$ . Afterwards, we compare this final price to the sum of the  $j$ -bids in  $\mathbf{r}$ .

#### Claim III.8.

$$\mathbb{E}\left[\max_{i:j \in B_i} a_j(r_i, B_i)\right] \geq \frac{1}{4}\mathbb{E}\left[p_j^{(n+1)}\right]$$

i.e., we expect the largest  $j$ -bid in  $\mathbf{r}$  to be not (too much) smaller than the greedy price of  $j$  in  $\mathbf{s}$ , after all  $n$   $\mathbf{s}$ -buyers have been considered.

*Proof.* To prove the claim, it suffices to show:

$$\begin{aligned} \forall \tau \geq 0 : \quad 2 \cdot \mathbb{P}\left[2 \max_{i:j \in B_i} a_j(r_i, B_i) \geq \tau\right] \\ \geq \mathbb{P}\left[\max_{i \in \{1, \dots, n+1\}} p_j^{(i)} \geq \tau\right]. \quad (1) \end{aligned}$$

This is because

$$\begin{aligned} \mathbb{E}\left[\max_{i:j \in B_i} a_j(r_i, B_i)\right] &= \int_{\tau=0}^{\infty} \mathbb{P}\left[\max_{i:j \in B_i} a_j(r_i, B_i) \geq \tau\right] d\tau \\ &= \frac{1}{2} \int_{\tau=0}^{\infty} \mathbb{P}\left[2 \max_{i:j \in B_i} a_j(r_i, B_i) \geq \tau\right] d\tau \\ &\geq \frac{1}{4} \int_{\tau=0}^{\infty} \mathbb{P}\left[\max_{i \in \{1, \dots, n+1\}} p_j^{(i)} \geq \tau\right] d\tau \\ &= \frac{1}{4}\mathbb{E}\left[\max_{i \in \{1, \dots, n+1\}} p_j^{(i)}\right], \end{aligned}$$

where the inequality uses Inequality (1).

To prove Inequality (1), suppose we have already decided the samples  $\mathbf{s}'$ . We still need to determine the random assignment of the remaining two valuations  $v_i^1, v_i^2$  for each buyer  $i \in [n]$  to  $\mathbf{r}$  and  $\mathbf{s}$ , respectively. The decision tree for this stochastic process is a perfect binary tree with  $n+1$  levels. In level  $i$  we toss a fair coin and the two children at level  $i+1$  correspond to the two possible assignments  $r_i = v_i^1, s'_i = v_i^2$  resp.  $r_i = v_i^2, s'_i = v_i^1$ . Note that there are  $2^n$  root-leave paths, and that these are in one-to-one correspondence with the realizations of the coin tosses.

We will now color (a subset of the vertices) in a top-down manner. Consider a node at level  $i$ . At this node we have already fixed  $\mathbf{r}_{<i}$  and  $\mathbf{s}_{<i}$ . We now decide the membership of  $v_i^1, v_i^2$ . Whenever one of  $v_i^1, v_i^2$  has item  $j$  in their demand

bundle, we color the two children of the current node and all of its descendants

- yellow, if, after the coin flip,  $2 \max_{i' \leq i: j \in B_{i'}} a_j(r_{i'}, B_{i'}) \geq \tau$ .
- red, if, after the coin flip,  $\max_{i'=1, \dots, i+1} p_j^{(i')} \geq \tau$ .

Note that a node can remain uncolored, may have only one of the two colors, or both.

The rationale is: We color a node yellow (resp. red) when it becomes clear that on any root-leave path that passes through this node condition  $2 \max_{i: j \in B_i} a_j(r_i, B_i) \geq \tau$  (or condition  $\max_{i \in \{1, \dots, n+1\}} p_j^{(i)} \geq \tau$ ) is met.

To establish Inequality (1) it thus suffices to argue that the number of yellow leaves is at least  $1/2$  of the number of red leaves. To this end, first consider a node at level  $i$  that is not yet colored. The only way one of its children (and with them the respective subtrees) can become red is when at least one of  $v_i^1, v_i^2$  has a  $j$ -bid of at least  $\tau/2$ . If both valuations have a  $j$ -bid of at least  $\tau/2$ , then both children (and their descendants) will have both colors. If only one of them has such a bid, then one subtree will be yellow and the other will be red.

So far this argument leads to the same number of yellow and red leaves. However, it is also possible that a node that is already colored gains a second color. Specifically, a node that is already yellow but not yet red can have one or both of its subtrees be colored red. If so, in the worst-case, (almost) all yellow leaves are also red.

Together this shows that the total number of red leaves is at most twice the number of yellow leaves, as claimed. This completes the proof.  $\square$

#### Claim III.9.

$$\mathbb{E}\left[p_j^{(n+1)}\right] \geq \frac{1}{4}\mathbb{E}\left[\sum_{i \in [n]} a_j(r_i, B_i)\right],$$

i.e., we expect the the final greedy price of item  $j$  with prices set by  $\mathbf{s}$  to be not (too much) smaller then the sum of the  $j$ -bids in  $\mathbf{r}$ .

*Proof.* The proof of this claim is based on a Martingale approach. Let

$$X_i = \sum_{i' < i: j \in B_{i'}} a_j(r_{i'}, B_{i'}) - 4 \cdot \max_{i' \in \{1, \dots, i\}} p_j^{(i')}.$$

Then, it suffices to show that  $\mathbb{E}[X_{n+1}] \leq 0$ . Since  $X_0 = 0$ , to prove the claim it suffices to show that the sequence  $X_0, X_1, \dots, X_{n+1}$  forms a supermartingale, i.e.,

$$\mathbb{E}\left[X_i - X_{i-1} | X_1, \dots, X_{i-1}\right] \leq 0.$$

To prove this, consider the random step where we decide between  $r_i$  and  $s_i$ . Either, none of the two  $j$ -bids exceed the current price, in which case  $X_i = X_{i-1}$ . Otherwise, at least one  $j$ -bid exists in the current step, the highest such bid  $b$  is associated with  $\mathbf{r}$  or  $\mathbf{s}$  with the same probability. If it ends up in  $\mathbf{s}$ , price  $p_j^{(i+1)}$  will at least be twice of  $p_j^{(i)}$ , and be at least the bid value  $b$  afterwards. Hence,  $p_j$  increases by at least  $b/2$ .

On the other hand, the sum always increases by at most  $b$ . This gives a supermartingale with a factor of 4.  $\square$

Combining Claim III.8 and Claim III.9, we obtain

$$\mathbb{E} \left[ \max_{i:j \in B_i} a_j(r_i, B_i) \right] \geq \frac{1}{16} \mathbb{E} \left[ \sum_{i \in [n]} a_j(r_i, B_i) \right].$$

Since  $a_j(s'_i, A'_i)$  are identically distributed as  $a_j(r_i, B_i)$ , we similarly get

$$\mathbb{E} \left[ \max_{i:j \in B_i} a_j(s'_i, A'_i) \right] \geq \frac{1}{16} \mathbb{E} \left[ \sum_{i \in [n]} a_j(s'_i, A'_i) \right].$$

Summing these two equations completes the proof of Lemma III.7.

#### E. Extending to a Single Sample Algorithm

We round up the proof of Theorem III.1 by showing that we can turn any pointwise  $k$ -sample prophet inequality into a pointwise single-sample prophet inequality, only using a  $O(k)$  term in the approximation guarantee.

**Theorem III.10.** *For any  $\alpha$ -competitive pointwise  $k$ -SPI with the objective of maximizing the sum over valuations of chosen elements in the single online steps, there exists an  $\alpha(k+1)$ -competitive pointwise 1-SPI.*

*Proof.* As input to the Game of Googol with  $k = 1$  we are given two valuations  $f_i^1, f_i^2$  for each buyer  $i \in [n]$ . For each buyer  $i \in [n]$ , there is an independent fair coin toss to decide sample  $s_i$  and actual valuation  $r_i$ . The algorithm sees all the  $s_i$  at the outset, while it can observe the  $r_i$  only one-by-one in an online fashion.

We will show how to construct from this an input to the Game of Googol with  $k > 1$ . Recall that we need to construct, for each buyer,  $k$  samples  $\bar{s}_i^1, \dots, \bar{s}_i^k$  and one actual valuation  $\bar{r}_i$ .

Here's a first attempt at this: We take  $s_i$  and  $k$  zero functions, and assign them uniformly at random to  $\bar{s}_i^1, \dots, \bar{s}_i^k$  and  $\bar{r}_i$ . We feed  $\bar{s}_i^1, \dots, \bar{s}_i^k$  and  $\bar{r}_i$  into the algorithm ALG for the Game of Googol with  $k > 1$ . Then, the expected optimum in the resulting instance with  $k$  samples is a  $(k+1)$ -approximation to that of the original instance of our 1-sample game: consider any buyer  $i$ 's allocated bundle in the optimal solution on  $s$ , and their valuation  $OPT_i(s)$  for it. With probability  $1/(k+1)$  (independently),  $s_i$  will be allocated to  $\bar{r}_i$ , implying  $\mathbb{E}[OPT(\bar{r})] \geq \frac{1}{k+1} \mathbb{E}[OPT(s)]$  (and the same holds for the  $OPT(\bar{s}_i)$ ,  $l \in [k]$ ).

There is one issue with this construction though: This is completely ignoring the actual valuations  $r_i$ . To make the construction useful in the online setting, whenever the coin flips were such that  $\bar{r}_i = s_i$ , we will substitute  $s_i$  with  $r_i$ . While in general, the input functions to the Game of Googol cannot depend on the coin flips made later to assign them, we will now justify that in our case, this is actually equivalent to our above first attempt.

Since all coin flips, those for the 1-sample game as well as those for the  $k$ -sample game, are independent, we can imagine

them to be drawn in the following order: first, the assignments of functions to all  $\bar{s}_i$  and  $\bar{r}_i$  are decided. Only then, we imagine the coins for the assignment of functions in the 1-sample game to be drawn, i.e., each two input functions  $f_i^1, f_i^2$  are assigned to  $s_i$  or  $r_i$ , respectively. With this, for each  $i$ , exchanging  $r_i$  for  $s_i$  makes no difference. Both is simply the result (to be drawn later) of whatever the fair coin flip between  $f_i^1$  and  $f_i^2$  results in.

Since ALG is an  $\alpha$ -approximation to the Game of Googol with  $k > 1$  samples, and the expected optimum of the instance that we feed into this algorithm is a  $(k+1)$ -approximation to the expected optimum of the original instance, we get an  $\alpha(k+1)$  approximation as claimed.  $\square$

We now have all the ingredients to prove Theorem III.1.

*Proof of Theorem III.1.* By Theorem III.5 there is a pointwise  $O(1)$ -competitive two-sample prophet inequality, applying Theorem III.10 we can turn this into a pointwise single-sample prophet inequality while losing only a factor 2 in the approximation guarantee.  $\square$

We conclude by noting that Theorem III.1 together with Theorem 6.2 in [11] implies Corollary I.2: the existence of a  $O(1)$ -competitive algorithm for online combinatorial auctions with XOS buyers in the order-oblivious secretary model.

## IV. POSTED-PRICE MECHANISMS WITH POLYNOMIAL SAMPLES

In this section we design  $(2+\epsilon)$ -competitive truthful combinatorial auctions for bidders with independent XOS valuations via sequential posted-pricing. The prices we construct will have the additional property that each item sells with half probability, over randomness in the agent valuations. We will call such prices *Median Prices*, defined formally below. We show that, even for non-XOS valuations, such prices always exist and are learnable with a polynomial number of samples that is independent of the valuation distributions, and in particular does not require bounded valuations. Finally, we show that the median property directly implies a constant approximation to welfare for XOS valuations.

#### A. Generic Distributions and Median Prices

Informally, median prices have the property that when agents make sequential purchase decisions, each item is sold with probability 1/2. An important subtlety is how agent indifferences are handled. We will first bypass these issues of tie-breaking by focusing on *generic* valuation distributions, a generalization of the atomless property, where for any fixed price vector  $\mathbf{p}$  indifferences occur with probability zero. Then, in Section IV-F we relax this generic assumption and extend our results to general distributions.

**Definition IV.1** (Generic Distribution). Distribution  $D_i$  over valuations is *generic* if for any price vector  $\mathbf{p}$  and any subset  $M_i \subseteq [m]$ , the demand correspondence of agent  $i$  (given prices

$\mathbf{p}$  and set  $M_i$  of available items) contains only a single set almost surely. That is,

$$\Pr_{v_i} [\exists B_1, B_2 \subseteq M_i: B_1 \neq B_2, \{B_1, B_2\} \subseteq \mathcal{D}_i(v_i, \mathbf{p}, M_i)] = 0.$$

Distribution  $D$  over valuation profiles is generic if its marginal distribution over each agent's valuation is generic.

We interpret the genericness condition as saying that a buyer has zero probability of being indifferent between two bundles, and hence between whether or not to include a given item  $j$  in their purchase. This generalizes the atomless property to valuations. We further note that one can make any distribution over valuations generic by adding an arbitrarily small noise to the agent valuations; see Section B for a formal construction. For technical convenience we will first restrict attention to generic distributions, then in Section IV-F we will show how to relax this assumption.

Given a fixed profile of generic valuation distributions, the probability that item  $j$  is sold under price vector  $\mathbf{p}$  is well-defined. We will write  $\pi_j(\mathbf{p})$  for this probability, omitting the dependence on the distributions. We are now ready to define what is meant by median prices.

**Definition IV.2** (Median Prices). Given  $\alpha \in [0, 1/2]$  and generic valuation profile distribution  $D$ , prices  $\mathbf{p}$  are  $\alpha$ -median prices if  $\pi_j(\mathbf{p}) \leq 1 - \alpha$  for each  $j$ , and  $\pi_j(\mathbf{p}) \geq \alpha$  for each  $j$  with  $p_j > 0$ .

Note the interpretation of parameter  $\alpha$ :  $(1/2)$ -median prices have the property that each item is sold with probability exactly  $1/2$ , whereas  $(1/4)$ -median prices have the property that each item is sold with probability between  $1/4$  and  $3/4$ . Lower values of  $\alpha$  are less restrictive, and every price vector  $\mathbf{p}$  is trivially  $0$ -median for any distribution  $D$ . Also note that we permit an item with price  $0$  to be sold with probability less than  $\alpha$ ; this edge case ensures that median prices exist even if there is no demand for a given item even if it is free.

### B. Truthful Combinatorial Auctions via Median Prices

We now show that median prices can be used to construct approximately welfare-efficient posted pricing mechanisms.

**Theorem IV.3.** Let  $D$  be a generic product distribution over XOS valuation profiles, and let  $\mathbf{p}$  be  $\alpha$ -median prices. Then the expected social welfare of a sequential posted-price mechanism using prices  $\mathbf{p}$  is at least  $\alpha$  times the expected optimal social welfare.

*Proof.* Let  $OPT$  denote the optimal social welfare. We will also write  $M_i(\mathbf{p}, \mathbf{v})$  for the set of items still available for purchase for agent  $i$  given the choice of prices and the agent valuations. We will write  $S(\mathbf{p}, \mathbf{v})$  for the set of items that are allocated to any agent.

We now consider separately the expected revenue and the expected buyer surplus generated by the mechanism.

The expected revenue generated by the mechanism is

$$\begin{aligned} \mathbb{E}_{\mathbf{v}} \left[ \sum_j p_j \times \mathbb{1}[j \in S(\mathbf{p}, \mathbf{v})] \right] &= \sum_j p_j \times \Pr_{\mathbf{v}} [j \in S(\mathbf{p}, \mathbf{v})] \\ &\geq \alpha \sum_j p_j, \end{aligned}$$

where we first use linearity of expectation and then that each item sells with probability at least  $\alpha$ .

For buyer surplus, write  $u_i$  for the random variable denoting the utility obtained by agent  $i$ . Note that since each buyer chooses a utility-maximizing set,  $u_i$  must be at least the utility obtained by agent  $i$  from any set of available goods. In particular, if we write  $A^*(\mathbf{v})$  for the welfare-maximizing allocation for valuation profile  $\mathbf{v}$ , we will consider a strategy of agent  $i$  in which they consider drawing a virtual choice of valuations  $\mathbf{v}'_{-i}$  for the other agents, then purchase those items from  $A^*(v_i, \mathbf{v}'_{-i})$  that are still available and for which the prices  $\mathbf{p}$  are not too high.

To bound the utility from this strategy we will make use of the supporting prices property of XOS valuations. Fix any valuation profile  $\mathbf{v}$ . Since the valuations are XOS, there exist supporting prices  $w_{ij}(\mathbf{v}) \geq 0$  such that, for each  $i$ , we have  $v_i(A_i^*(\mathbf{v})) = \sum_{j \in A_i^*(\mathbf{v})} w_{ij}(\mathbf{v})$ , and moreover for any  $T \subseteq A_i^*(\mathbf{v})$  we have  $v_i(T) \geq \sum_{j \in T} w_{ij}(\mathbf{v})$ .

Given an agent  $i$ , prices  $\mathbf{p}$ , valuation profile  $\mathbf{v}$ , and a hallucinated valuation profile  $\mathbf{v}'$  of the other agents, we will define a subset  $T_i(\mathbf{p}, \mathbf{v}, \mathbf{v}'_{-i})$  of  $A_i^*(v_i, \mathbf{v}'_{-i})$  to be  $A_i^*(v_i, \mathbf{v}'_{-i}) \cap M_i(\mathbf{p}, \mathbf{v}) \cap \{j: w_{ij}(\mathbf{v}) \geq p_j\}$ . That is,  $T_i(\mathbf{p}, \mathbf{v}, \mathbf{v}'_{-i})$  is the subset of agent  $i$ 's optimal allocation (under the hallucinated profile of the other agents) that are available (under the true valuations) and for which the prices no larger than  $w_{ij}(\mathbf{v})$ . We then have

$$\begin{aligned} \mathbb{E}_{\mathbf{v}}[u_i] &\geq \mathbb{E}_{\mathbf{v}, \mathbf{v}'_{-i}} \left[ v_i(T_i(\mathbf{p}, \mathbf{v}, \mathbf{v}'_{-i})) - \sum_{j \in T_i(\mathbf{p}, \mathbf{v}, \mathbf{v}'_{-i})} p_j \right] \\ &\geq \mathbb{E}_{\mathbf{v}, \mathbf{v}'_{-i}} \left[ \sum_{j \in T_i(\mathbf{p}, \mathbf{v}, \mathbf{v}'_{-i})} (w_{ij}(v_i, \mathbf{v}'_{-i}) - p_j) \right] \\ &= \mathbb{E}_{v_i, \mathbf{v}'_{-i}} \left[ \sum_{j \in A_i^*(v_i, \mathbf{v}'_{-i})} ((w_{ij}(v_i, \mathbf{v}'_{-i}) - p_j)^+ \times \Pr_{\mathbf{v}'_{-i}} [j \in M_i(\mathbf{p}, \mathbf{v})]) \right] \\ &= \mathbb{E}_{\mathbf{v}} \left[ \sum_{j \in A_i^*(\mathbf{v})} ((w_{ij}(\mathbf{v}) - p_j)^+ \times \Pr_{\mathbf{v}'_{-i}} [j \in M_i(\mathbf{p}, (v_i, \mathbf{v}'_{-i}))]) \right] \\ &\geq \mathbb{E}_{\mathbf{v}} \left[ \sum_{j \in A_i^*(\mathbf{v})} (w_{ij}(\mathbf{v}) - p_j)^+ \times \alpha \right]. \end{aligned}$$

Taking a sum over all items, we conclude (writing  $OPT$  for the expected optimal social welfare)

$$\mathbb{E} \left[ \sum_i u_i \right] \geq \alpha \left( OPT - \sum_j p_j \right).$$

Since the expected social welfare is the expected revenue plus the expected buyer surplus, we conclude that the expected welfare is at least

$$\alpha \sum_j p_j + \alpha(OPT - \sum_j p_j) = \alpha OPT,$$

as required.  $\square$

### C. Existence of Median Prices

Theorem IV.3 shows that median prices guarantee high welfare when used as posted prices. But it is not immediately obvious that such prices are even guaranteed to exist. We now prove the existence of  $(1/2)$ -median prices for generic distributions over (possibly non-XOS) valuations.

**Theorem IV.4.** *For any generic distribution  $D$  over valuations there exist  $(1/2)$ -median prices.*

Fix a generic distribution  $D$ . In order to show the existence we will use the Kakutani fixed-point theorem on an appropriately defined function. Recall that  $\pi_j(\mathbf{p})$  is the (well-defined) probability that item  $j$  is allocated under prices  $\mathbf{p}$  and write  $\pi(\mathbf{p})$  for the profile of allocation probabilities for the items. We will first show that  $\mathbf{p} \mapsto \pi(\mathbf{p})$  is a continuous non-decreasing function.

**Lemma IV.5.** *The function  $\pi(p_j, \mathbf{p}_{-j})$  is continuous and, for each  $j$ ,  $\pi_j(p_j, \mathbf{p}_{-j})$  is weakly decreasing in  $p_j$ .*

*Proof.* Write  $A_i(\mathbf{v}, \mathbf{p})$  for the allocation to agent  $i$  under prices  $\mathbf{p}$  when the valuations are  $\mathbf{v}$ .

To see that it is weakly decreasing, take any  $p'_j > p_j$  and any valuation  $\mathbf{v}$ . If the allocations  $\mathbf{A}(\mathbf{v}, \mathbf{p})$  and  $\mathbf{A}(\mathbf{v}, (p'_j, \mathbf{p}_{-j}))$  are identical, then either both allocate  $j$  or neither does. Otherwise, take the smallest  $i$  for which  $A_i(\mathbf{v}, \mathbf{p}) \neq A_i(\mathbf{v}, (p'_j, \mathbf{p}_{-j}))$ . Then agent  $i$  faces the same set of available items under  $\mathbf{p}$  and  $(p'_j, \mathbf{p}_{-j})$ . Since agent  $i$  always receives a set from her demand correspondence, and agent  $i$  faces the same set of available items and the same prices except for a higher price on item  $j$ , we must have  $j \in A_i(\mathbf{v}, \mathbf{p})$  and  $j \notin A_i(\mathbf{v}, (p'_j, \mathbf{p}_{-j}))$ . In particular, this means that if  $\mathbf{A}(\mathbf{v}, \mathbf{p})$  does not allocate item  $j$ , then neither does  $\mathbf{A}(\mathbf{v}, (p'_j, \mathbf{p}_{-j}))$ . Integrating over all choices of  $\mathbf{v}$  then implies that the probability of allocating item  $j$  is weakly lower under price  $p'_j$  than under price  $p_j$ .

To show continuity of  $\pi_j$ , we proceed as follows. For any agent  $i$  and sets of items  $S$  and  $M'$ , let  $f_{i,S,M'}(\mathbf{p})$  be the probability that under price vector  $\mathbf{p}$  buyer  $i$  buys the set  $S$  conditional on the set of remaining items being  $M'$ .<sup>2</sup> It is now enough to show that for each  $i$ ,  $S$ , and  $M'$  such that  $M'$  is the set of remaining items for  $i$  with positive probability, the function  $f_{i,S,M'}$  is continuous. This is because we can write  $\pi_j(\mathbf{p})$  as

$$\pi_j(\mathbf{p}) = \sum_{\substack{S_1, \dots, S_n \subseteq [m] \\ j \in S_1 \cup \dots \cup S_n}} \prod_{i=1}^n f_{i,S_i,[m] \setminus (S_1 \cup \dots \cup S_{i-1})}(\mathbf{p}).$$

<sup>2</sup>Note that if  $D$  is not a product distribution, the conditioning on  $M'$  impacts the conditional distribution over  $v_i$ .

To show continuity of  $f_{i,S,M'}$  for a fixed choice of  $i$ ,  $S$ , and  $M'$ , let

$$F(\mathbf{p}, t) = \Pr \left[ v_i(S) - \sum_{j \in S} p_j - \max_{\substack{S' \subseteq M', \\ S \neq S'}} (v_i(S') - \sum_{j \in S'} p_j) \leq t \right].$$

We observe that for all  $\mathbf{p}$ , we have

$$F(\mathbf{p}, t) - \lim_{\substack{t \rightarrow 0 \\ t < 0}} F(\mathbf{p}, t) =$$

$$\Pr \left[ v_i(S) - \sum_{j \in S} p_j - \max_{\substack{S' \subseteq M', \\ S \neq S'}} (v_i(S') - \sum_{j \in S'} p_j) = 0 \right] = 0$$

because the distribution is generic.<sup>3</sup> That is, the function  $F(\mathbf{p}, \cdot)$  is continuous at 0. Furthermore, if  $\|\mathbf{p} - \mathbf{p}'\|_\infty \leq \delta$ , then  $F(\mathbf{p}, t - m\delta) \leq F(\mathbf{p}', t) \leq F(\mathbf{p}, t + m\delta)$ . So, in order to show that  $f_{i,S,M'}$  is continuous at  $\mathbf{p}$ , let  $\epsilon > 0$  and choose  $\delta > 0$  such that  $|F(\mathbf{p}, mx) - F(\mathbf{p}, 0)| < \epsilon$  for all  $x \in [-\delta, \delta]$ . Now, for any  $\mathbf{p}'$  with  $\|\mathbf{p} - \mathbf{p}'\|_\infty \leq \delta$ , we have  $F(\mathbf{p}', 0) \leq F(\mathbf{p}, m\delta) < F(\mathbf{p}, 0) + \epsilon$  and  $F(\mathbf{p}', 0) \geq F(\mathbf{p}, -m\delta) > F(\mathbf{p}, 0) - \epsilon$ . The statement now follows because  $f_{i,S,M'}(\mathbf{p}) = F(\mathbf{p}, 0)$  for all  $\mathbf{p}$ .  $\square$

Using monotonicity and continuity of the allocation property, we can complete the proof of existence of  $(1/2)$ -median prices.

*Proof of Theorem IV.4.* We will construct a mapping  $\Psi$ , which takes as input a vector of item prices and returns a vector of closed intervals of item prices. Informally, we wish to define  $\Psi(\mathbf{p})$  to be vector of intervals  $\mathbf{Q}$  where each  $Q_j$  is the closed interval of prices  $[q_j^1, q_j^2]$  such that item  $j$  sells with probability  $1/2$  on any price vector  $(\mathbf{p}_{-j}, q_j)$  with  $q_j \in [q_j^1, q_j^2]$ . We'll then argue that  $\Psi$  has a fixed point, and this will be our  $1/2$ -median prices.

There are a few edge cases to consider when defining the mapping formally. First, write  $\bar{p}$  to be the median of  $\max_i v_i(M)$ . Then it must be that  $\pi_j(\bar{p}, \mathbf{p}_{-j}) \leq 1/2$  for any  $\mathbf{p}$  because for item  $j$  to sell at price  $\bar{p}$  it is necessary that  $v_i(M) \geq \bar{p}$  for an  $i$ . We will focus attention on price vectors lying in  $[0, \bar{p}]^m$ . Fix input  $\mathbf{p}$ , and define  $\phi_j(q_j) := \pi_j(q_j, \mathbf{p}_{-j})$ . We have two cases. Either  $\phi_j(0) < 1/2$ , or  $\phi_j(0) \geq 1/2$ . In the first case, we will define  $\Psi_j(\mathbf{p}) = \{0\}$ . In the latter case, must exist some  $q > 0$  at which  $\phi_j(q) = 1/2$ , by continuity and the fact that  $\phi_j(\bar{p}) \leq 1/2$ . Note there may be more than one such  $q$ , so write  $Q_j$  for the set of  $q$  that satisfy this property, which must be a closed interval since  $\phi_j$  is continuous and weakly decreasing. Then we will define  $\Psi_j(\mathbf{p}) = Q_j$ .

This concludes the definition of  $\Psi$ . We now note that  $\Psi$  is a set-valued mapping from a compact space (i.e.,  $[0, \bar{p}]^m$ ) to itself with closed values. Moreover,  $\Psi$  is upper hemi-continuous. To see why, consider some sequence of prices vectors  $\mathbf{p}^k$  that converge to  $\mathbf{p}$  as  $k \rightarrow \infty$ , with a corresponding sequence  $\mathbf{q}^k$  converging to  $\mathbf{q}$ . We then have that item  $j$  sells with

<sup>3</sup>Here we used the assumption that  $M'$  has non-zero probability of being the set of remaining items for agent  $i$ . If agent  $i$  has non-zero probability of indifference between  $S$  and  $S'$  conditional on the set of remaining items being  $M'$ , then they have non-zero probability of this event unconditionally, contradicting genericness.

probability  $1/2$  on price vector  $(q_j^k, \mathbf{p}_{-j}^k)$  for all  $k$ . Since  $(q_j^k, \mathbf{p}_{-j}^k)$  converges to  $(q_j, \mathbf{p}_{-j})$ , we conclude by continuity of  $\pi$  that  $j$  sells with probability  $1/2$  on price vector  $(q_j, \mathbf{p}_{-j})$  for each  $j$ , and hence  $\mathbf{q} \in \Psi(\mathbf{p})$ .

Since  $\Psi$  is upper hemi-continuous, by the Kakutani fixed-point theorem it has a fixed point  $\mathbf{p}$ . That is, there exists  $\mathbf{p}$  such that  $p_j \in \Psi_j(\mathbf{p})$  for all  $j$ . This fixed point is a choice of  $1/2$ -median prices, by definition.  $\square$

#### D. Learning Median Prices

Next, we can also find median prices with polynomially many samples. Note that in contrast to similar learnability results for prices defined on the mean (e.g., in [1]), our error is only multiplicative. That is, we do not need the distribution of values to be bounded.

**Theorem IV.6.** *For any generic distribution  $D$ , with probability  $1 - \delta$ ,  $O\left(\frac{1}{\epsilon^2}(m^2 + m \log n + \log(1/\delta))\right)$  samples suffice to find  $(\frac{1}{2} - \epsilon)$ -median prices.*

Suppose we are given  $k$  vectors of  $n$  valuation functions that are drawn from the unknown distribution. Construct the empirical distribution  $\hat{D}$ , in which one of these  $k$  vectors is drawn uniformly at random. (Note that there are correlations between the agents in this empirical distribution.)

For draws from the generic distribution  $D$ , regardless of the price vector  $\mathbf{p}$ , almost surely no tie occurs. This is not true for the empirical distribution  $\hat{D}$ . In the following, we will fix any deterministic tie-breaking rule that does not depend on prices or the set of remaining items. That is, every agent has a permutation of the power set of the set of items and whenever there is a tie between  $S$  and  $S'$  agent  $i$  will prefer one over the other.

Based on this tie-breaking, we let  $\hat{\pi}_j(\mathbf{p})$  be analogously the probability that item  $j$  gets sold under prices  $\mathbf{p}$  if a valuation profile  $\mathbf{v}$  is drawn from  $\hat{D}$ . Note that it is safe to assume that the same tie-breaking rule is applied to define the probability  $\pi_j(\mathbf{p})$  that item  $j$  gets sold under price vector  $\mathbf{p}$  if the valuation profile  $\mathbf{v}$  is drawn from  $D$ .

The following lemma states a uniform convergence property that with high probability, the sale probability of each item given *any* price vector is approximately equal between the empirical distribution and the true distribution.

**Lemma IV.7** (Uniform Convergence). *Consider the empirical distribution over  $k$  sampled valuation profiles, with  $k \geq C \frac{1}{\epsilon^2} (m^2 + m \log n + \log(1/\delta))$  for a constant  $C$ . Then, with probability at least  $1 - \delta$ , we have  $\|\hat{\pi}(\mathbf{p}) - \pi(\mathbf{p})\|_\infty < \epsilon$ , that is, there is no price vector on which the probability that an item sells differs by more than  $\epsilon$  between the actual and the empirical distribution.*

Before proving the lemma, we note that this implies Theorem IV.6. We know that there exists a  $(1/2)$ -median price vector  $\mathbf{p}^*$  for  $D$ . So with probability  $1 - \delta$ ,  $\mathbf{p}^*$  is  $(1/2 - \epsilon)$ -median for  $\hat{D}$ , and hence in particular a  $(1/2 - \epsilon)$ -median price profile exists. Let  $\mathbf{p}$  be *any*  $(1/2 - \epsilon)$ -median price profile for  $\hat{D}$ . Then, again conditioning on the event of probability  $1 - \delta$ ,

$\mathbf{p}$  will be  $(1/2 - 2\epsilon)$ -median for  $D$ . Taking  $k$  samples therefore suffices to obtain  $(1/2 - 2\epsilon)$ -median prices with probability  $1 - \delta$ .

*Proof of Lemma IV.7.* Consider the space of all possible valuation vectors  $\mathbf{v}$ . Every price vector  $\mathbf{p}$  defines a subset  $S_j^\mathbf{p}$  of this space, equal to the collection of valuation vectors for which item  $j$  is sold under price vector  $\mathbf{p}$ . By this definition  $\pi_j(\mathbf{p}) = \Pr_{\mathbf{v} \sim D} [\mathbf{v} \in S_j^\mathbf{p}]$  and  $\hat{\pi}_j(\mathbf{p}) = \Pr_{\mathbf{v} \sim \hat{D}} [\mathbf{v} \in S_j^\mathbf{p}]$ .

Standard bounds on uniform convergence (e.g., in [35] combine Theorem 4.10 and Equation (5.50)) imply that with  $k$  samples, we have for any  $\gamma > 0$ ,

$$\Pr_{\mathbf{p}} \left[ \sup_{\mathbf{p}} |\hat{\pi}_j(\mathbf{p}) - \pi_j(\mathbf{p})| \geq C \sqrt{\frac{\nu}{k}} + \gamma \right] \leq \exp(-k\gamma^2/2),$$

where  $C$  is a constant and  $\nu$  is the VC dimension of the set system  $\{S_j^\mathbf{p}\}$ . Below, in Lemma IV.8, we show that the VC dimension of the set system defined by all of these subsets is at most  $O(m^2 + m \log n)$ .

So for any item  $j$  and any  $\epsilon, \delta' > 0$ , this implies that  $O\left(\frac{\nu + \ln(1/\delta')}{\epsilon^2}\right)$  samples are sufficient for  $\Pr[\sup_{\mathbf{p}} |\hat{\pi}_j(\mathbf{p}) - \pi_j(\mathbf{p})| \geq \epsilon] \leq \delta'$ .

By furthermore taking a union bound with  $\delta' = \frac{\delta}{m}$  over all items  $j$ ,  $O\left(\frac{\nu + \ln(1/\delta)}{\epsilon^2}\right)$  samples are sufficient for  $\Pr[\sup_{\mathbf{p}} \|\hat{\pi}(\mathbf{p}) - \pi(\mathbf{p})\|_\infty \geq \epsilon] \leq \delta$ , meaning that with probability at least  $1 - \delta$ , we have that  $|\hat{\pi}_j(\mathbf{p}) - \pi_j(\mathbf{p})| < \epsilon$  for all items  $j$  and all price vectors  $\mathbf{p}$ .  $\square$

So, it remains to bound the VC dimension of the set system  $\{S_j^\mathbf{p}\}$  defined above. We use the technique by Balcan et al. [19] (see also [20]).

**Lemma IV.8.** *The VC dimension of the set system  $\{S_j^\mathbf{p}\}$  is at most  $O(m^2 + m \log n)$ .*

*Proof.* Let  $\mathbf{v}^{(1)}, \dots, \mathbf{v}^{(N)}$  be  $N$  valuation profiles that are shattered. That is, for every set  $T \subseteq [N]$ , there is a price vector  $\mathbf{p}^T$  such that item  $j$  gets sold on valuation profile  $\mathbf{v}^{(t)}$  if and only if  $t \in T$ . We will give an upper bound on the number of price vectors that lead to different outcomes on  $\mathbf{v}\mathbf{v}^{(1)}, \dots, \mathbf{v}\mathbf{v}^{(N)}$ .

Note that if on a valuation profile  $\mathbf{v}$ , item  $j$  gets sold under price vector  $\mathbf{p}$  but doesn't get sold under price vector  $\mathbf{p}'$ , there have to be  $i$ ,  $S$ , and  $S'$  such that agent  $i$  prefers  $S$  under prices  $\mathbf{p}$  but  $S'$  under prices  $\mathbf{p}'$  because otherwise the outcomes would be identical. That is,  $v_i(S) - \sum_{j' \in S} p_{j'} \geq v_i(S') - \sum_{j' \in S'} p_{j'}$  but  $v_i(S) - \sum_{j' \in S} p_{j'} \leq v_i(S') - \sum_{j' \in S'} p_{j'}$ , where one of the two inequalities is strict because the tie-breaking is independent of the prices. So,  $v_i(S) - v_i(S') - \sum_{j' \in S \setminus S'} p_{j'} + \sum_{j' \in S' \setminus S} p_{j'} \geq 0$  but  $v_i(S) - v_i(S') - \sum_{j' \in S \setminus S'} p_{j'} + \sum_{j' \in S' \setminus S} p_{j'} \leq 0$ , where again one of the two inequalities is strict.

In other words, if we consider the space  $\mathbb{R}_+^m$  of all price vectors, the price vectors  $\mathbf{p}$  and  $\mathbf{p}'$  lie on different sides of the hyperplane defined by  $v_i(S) - v_i(S') - \sum_{j' \in S \setminus S'} p_{j'}^* + \sum_{j' \in S' \setminus S} p_{j'}^* = 0$ .

The valuation profiles  $\mathbf{v}^{(1)}, \dots, \mathbf{v}^{(N)}$  define  $N \cdot n \cdot (2^m)^2$  such hyperplanes. Due to a result by Buck [36], this means there are no more than  $m(N \cdot n \cdot (2^m)^2)^m$  different regions. If  $\mathbf{v}^{(1)}, \dots, \mathbf{v}^{(N)}$  is shattered, each of the  $2^N$  price vectors  $\mathbf{p}^T$  for  $T \subseteq [N]$  has to lie in a different region.

We conclude that  $2^N \leq m(N \cdot n \cdot (2^m)^2)^m$ . Taking logs on both sides yields  $N \leq \log m + m \log N + m \log n + 2m^2$ , which implies  $N = O(m^2 + m \log n)$ .  $\square$

Generally, the polynomial dependence in Lemma IV.8 is unavoidable. Indeed, for the special case of unit-demand valuations, even with  $n = 1$  the VC dimension of the proposed set system is  $\Omega(m)$ . Given  $m$  items, and any  $t$  with  $1 \leq t < m$ , let  $v_1^{(t)}$  be unit-demand function that has value 1 for item 1, value 2 for item  $t+1$  and value 0 for all other items. Consider price vectors with  $p_1 = 0$ . Now, depending on the price of item  $t+1$ ,  $v_1^{(t)}$  will prefer item 1 (if  $p_{t+1} < 1$ ) or item  $t+1$  (if  $p_{t+1} > 1$ ). Thus, for every subset of the valuations  $v_1^{(1)}, \dots, v_1^{(m-1)}$ , there is a price vector such that exactly this subset buys item 1.

#### E. Computational Challenges

The uniform convergence result from the previous section also provides us a very powerful tool to compute median prices, namely by considering the empirical distribution of polynomially many samples—which means that one has to argue only about polynomially many scenarios.

To demonstrate this strength, we now describe a Tatonnement style polynomial-time algorithm for computing median prices for unit demand valuations. We leave it as a challenging but interesting open problem to extend this result to XOS valuations.

**Theorem IV.9.** *For any generic distribution  $D$  over unit demand valuations and any  $\epsilon > 0$ , there is a fully polynomial-time algorithm that, given polynomially many samples from  $D$ , generates a  $(1/2 - \epsilon)$ -median distribution over item prices with probability  $1 - \delta$ . The runtime is polynomial in  $n, m, 1/\epsilon$ , and  $\log(1/\delta)$ .*

*Proof.* Fix generic distribution profile  $D$ . Our algorithm will begin by drawing  $k$  sampled valuation profiles  $(\mathbf{v}^1, \dots, \mathbf{v}^k)$  from  $D$ ; let  $\hat{D}$  denote the empirical distribution. By Lemma IV.7, the probability of sale between  $D$  and  $\hat{D}$  agree within  $\epsilon$  on all price vectors, with probability  $1 - \delta$ . We will condition on this event for the remainder of this proof.

Write  $\hat{\pi}(\mathbf{p})$  for the profile of probabilities that each item sells under prices  $\mathbf{p}$  given the empirical distribution  $\hat{D}$ , assuming an arbitrary tie-breaking rule. Note that  $\hat{\pi}_j(\mathbf{p})$  is always a multiple of  $1/k$ , and can equivalently be thought of as counting how many of the  $k$  samples result in item  $j$  being purchased.

Given any price vector  $\mathbf{p}$ , we will define a potential value  $f(\mathbf{p})$  that depends on the probabilities of allocation. This will be the sum of potential values  $f_j(\mathbf{p})$  for each item  $j$ . Assume for convenience that  $k$  is even. If  $\hat{\pi}_j(\mathbf{p}) \leq 1/2$  then  $f_j(\mathbf{p}) = 0$ . If  $1/2 < \hat{\pi}_j(\mathbf{p}) \leq 1/2 + \epsilon$ , say  $\hat{\pi}_j(\mathbf{p}) = 1/2 + t/k$  where  $t \leq \epsilon k$ , then  $f_j(\mathbf{p}) = t^2$ . If  $\hat{\pi}_j(\mathbf{p}) > 1/2 + \epsilon$ , say  $\hat{\pi}_j(\mathbf{p}) =$

$1/2 + t/k$  where  $t > \epsilon k$ , then  $f_j(\mathbf{p}) = (\epsilon k)t$ . Note then that  $0 \leq f(\mathbf{p}) \leq \epsilon k^2 m$  for all prices  $\mathbf{p}$ .

We are now ready to describe our algorithm. We will initialize all item prices  $p_j$  to 0.

If  $\hat{\pi}_j(\mathbf{p}) \leq 1/2 + \epsilon$  for all  $j$  then we are done. Otherwise, since  $k/\epsilon > m$ , there must exist some multiple of  $1/k$ , say  $t/k$ , such that  $t/k \in (1/2, 1/2 + \epsilon)$  and  $\hat{\pi}_j(\mathbf{p}) \neq t/k$  for all  $j$ . We will take  $S$  to be the set of all  $j$  such that  $\hat{\pi}_j(\mathbf{p}) > t/k$ . Note that  $S$  is non-empty, since in particular any item  $j$  with  $\hat{\pi}_j(\mathbf{p}) > 1/2 + \epsilon$  must be in  $S$ .

We will then raise the price of each item in  $S$ , uniformly, until some agent  $i$  in some sample  $k$  switches from demanding an item in  $S$  to demanding an item in  $[m] - S$  (or demanding no item). Note that since the valuations are unit-demand, these are the only changes to the demand of an agent that can occur.

The amount we need to raise the price for this to occur is the minimum, over all  $i, \ell$ , and  $j \in S$  and  $j' \notin S$ , of  $(v_i^\ell(j) - p_j) - (v_i^\ell(j') - p_{j'})$ . (Where here we can allow  $j' = 0$  to represent taking the empty set, with value and price both equal to 0.) We can therefore find this uniform price increment in polynomial time by enumerate each of these polynomially-many options.

Under this update, the allocation probability of some item  $j$  with  $\hat{\pi}_j(\mathbf{p}) > t/k$  reduces by  $1/k$ , and the allocation probability of some item  $j'$  with  $\hat{\pi}_j(\mathbf{p}) < t/k$  increases by  $1/k$ . Thus, if we write  $\mathbf{p}'$  for the new price vector, we conclude that  $f(\mathbf{p}) - f(\mathbf{p}') \geq (t+1)^2 + (t-1)^2 - 2t^2 = 2$ .

We then repeat, choosing a new set  $S$  and iterating, terminating if  $\hat{\pi}_j(\mathbf{p}) \geq 1/2 + \epsilon$  for all  $j$ . Note that if we do not terminate, then in particular we have  $\hat{\pi}_j(\mathbf{p}) > 1/2 + \epsilon$  for some  $j$ , so  $f(\mathbf{p}) \geq f_j(\mathbf{p}) > \epsilon k^2$ . On the other hand, since  $f(\mathbf{p})$  is at most  $\epsilon k^2 m$  initially and reduces by at least 2 on each iteration, we conclude that the procedure must terminate after at most  $\epsilon k^2(m-1)/2$  iterations. This algorithm therefore finds the desired price vector in polynomially many iterations, as claimed.  $\square$

#### F. Beyond Generic Distributions

To this point in our discussion of median prices we have focused on the case of generic distributions. We now show how to remove this assumption. We will show how to define median prices to non-generic distributions, then show that our welfare bound in Theorem IV.3, existence result in Theorem IV.4, and learning result in Theorem IV.6 all extend to general non-generic settings.

For non-generic distributions, an agent  $i$  may be indifferent between multiple sets with positive probability. The probability with which an item sells therefore depends on how agents choose between sets in their demand correspondence. It will thus be convenient to define, for each agent, a choice rule that describes how indifferences are to be broken.

A *choice rule*  $A_i$  for agent  $i$  is a mapping from a demand correspondence  $\mathcal{D} \subseteq 2^{[m]}$  (which recall is a set of sets of items) to a distribution over elements of  $\mathcal{D}$ . We think of  $A_i(\mathcal{D})$  as the (possibly randomized) set chosen by agent  $i$  in the face of indifference. Given a profile of choice rules  $\mathbf{A} = (A_1, \dots, A_n)$ ,

a valuation profile  $\mathbf{v}$ , and prices  $\mathbf{p}$ , we will write  $\mathbf{A}(\mathbf{v}, \mathbf{p})$  for the (random) allocation that occurs when each agent  $i$  sequentially selects from their demand correspondence (given  $v_i$ ,  $\mathbf{p}$ , and the set of remaining items  $M_i$ ) according to  $A_i$ . We will also write  $A_i(\mathbf{v}, \mathbf{p})$  for the (random) allocation to agent  $i$  under allocation profile  $\mathbf{A}(\mathbf{v}, \mathbf{p})$ .

**Definition IV.10** (Median Prices for General Distributions). Given  $\alpha \in [0, 1/2]$  and any profile  $D$  of distributions of bidder valuations, we say that prices  $\mathbf{p}$  are  $\alpha$ -median prices if there exists a choice rule profile  $\mathbf{A}$  such that, when agents acquire bundles according to  $\mathbf{A}$ , the probability that item  $j$  is allocated lies in  $[\alpha, 1 - \alpha]$  for every  $j$ .

Note that prices  $\mathbf{p}$  are median prices if there exists an appropriate corresponding choice rule  $\mathbf{A}$ . We say that such a choice rule is a *witness* for prices  $\mathbf{p}$ .

**Example IV.11.** Suppose there is a single item  $\{1\}$  and two agents  $\{1, 2\}$ . Agent values are deterministic (i.e.,  $D$  is a point mass): agent 1 has value 3 and agent 2 has value 4. Then the price  $p_1 = 4$  is a median price, witnessed by any choice rule such that  $A_1$  is arbitrary and  $A_2(\{\{1\}, \emptyset\})$  randomizes uniformly between  $\{1\}$  and  $\emptyset$ .

**Welfare of Median Prices for General Distributions.** We are now ready to extend Theorem IV.3, our approximation result for median prices, to non-generic settings. Indeed, under our definition of median prices for general distributions, the proof of Theorem IV.3 proceeds unchanged. We require only that in our sequential posted-price mechanism, indifferences are broken according to a choice rule  $\mathbf{A}$  that witnesses the median prices. The reason is that genericity is never used explicitly in the proof; we use only that each item is purchased with probability lying in  $[\alpha, 1 - \alpha]$  under the posted-price mechanism, which is guaranteed if the witnessing choice rule is followed. We obtain the following result.

**Theorem IV.12.** *Let  $D$  be a product distribution over XOS valuation profiles, and let  $\mathbf{p}$  be  $\alpha$ -median prices witnessed by choice rule  $\mathbf{A}$ . Then the expected social welfare of a sequential posted-price mechanism using prices  $\mathbf{p}$  and choice rule  $\mathbf{A}$  to resolve indifferences is at least  $\alpha$  times the expected optimal social welfare.*

**Existence of Median Prices for General Distributions.** Next, we will extend Theorem IV.4 and show that  $(1/2)$ -median prices exist even for non-generic distributions over (possibly non-XOS) valuations. The idea behind the proof is to consider perturbations of the non-generic distribution, yielding generic distributions for which median prices exist. We then take a limit of vanishing perturbations to find median prices for the original distribution. To find the witnessing choice rule, we take an appropriate limit of allocations over the sequence of perturbations. We obtain the following result.

**Theorem IV.13.** *For any (not necessarily generic) distribution  $D$  over valuation profiles there exist  $(1/2)$ -median prices.*

*Proof.* Following the outline above, we will construct a perturbed distribution and consider its median prices. Given distribution  $D$  over valuation profiles and  $\epsilon > 0$ , write  $D^{(\epsilon)}$  for the profile of generic distributions from the proof of Proposition B.1. Recall that, to draw a valuations  $\mathbf{v}'$  from distribution  $D^{(\epsilon)}$ , we first draw  $\mathbf{v}$  from  $D$ , then draw an additive value  $w_{ij}$  uniformly from  $[0, \epsilon]$  for each agent  $i$  and item  $j$ , and then finally set  $v'_i(S) = v_i + \sum_{j \in S} w_{ij}$  for each  $S \subseteq [m]$ . In Proposition B.1 we prove that  $D^{(\epsilon)}$  is generic for any  $\epsilon > 0$ , so by Theorem IV.4 it admits  $(1/2)$ -median prices. Write  $\mathbf{p}^{(\epsilon)}$  for these prices.

First, we argue that the price vectors  $\mathbf{p}^{(\epsilon)}$  are uniformly bounded. Consider the optimal welfare obtainable from a valuation profile drawn from  $D$ . This is a random variable; let  $\mu$  denote its highest median value, which must be finite. If  $\mu^{(\epsilon)}$  is the corresponding median for distribution  $D^{(\epsilon)}$ , then note we must have  $\mu^{(\epsilon)} \leq \mu + m\epsilon$  (as the optimal welfare can increase by at most  $m\epsilon$  pointwise under the addition of  $\epsilon$ -bounded random noise for each item). So  $\mu^{(\epsilon)}$  is at most a finite constant depending on  $D$ , for any  $0 < \epsilon < 1$ . Finally, note that we must have  $p_j^{(\epsilon)} \leq \mu^{(\epsilon)}$  for all  $j$ , as the probability that any set sells at a price higher than  $\mu^{(\epsilon)}$  is less than  $1/2$  by definition. We therefore conclude that  $p^{(\epsilon)}$  lies in  $[0, \mu + m]^m$  for all  $0 < \epsilon < 1$ .

Since the prices  $\mathbf{p}^{(\epsilon)}$  lie in a compact set, the sequence  $\mathbf{p}^{(\epsilon)}$  as  $\epsilon \rightarrow 0$  must have a convergent subsequence. Write  $\mathbf{p}^*$  for the limit of this convergent subsequence.

We claim that  $\mathbf{p}^*$  is a profile of  $(1/2)$ -median prices for valuation profile distribution  $D$ . To prove this claim we must construct a witnessing choice rule  $\mathbf{A}^*$ . We begin by constructing a choice rule  $\mathbf{A}^{(\epsilon)}$  for each  $\epsilon > 0$ .<sup>4</sup> For any  $\epsilon > 0$ , any agent  $i$ , and any collection of sets  $\mathcal{D}$ , we will define

$$A_i^{(\epsilon)}(\mathcal{D}) \in \arg \max_{S \in \mathcal{D}} \left\{ \sum_{j \in S} (w_j + p_j^* - p_j^{(\epsilon)}) \right\}$$

where each  $w_j$  is drawn uniformly from  $[0, \epsilon]$  independently for each  $j$ . Note that if the argmax is not unique in the definition of  $A_i^{(\epsilon)}(\mathcal{D})$  then we can choose an element in the argmax arbitrarily, as this happens with probability 0 for generic distribution  $D^{(\epsilon)}$ .

To give some intuition into the definition of  $A_i^{(\epsilon)}$ , note that for any  $v_i$  and any set  $S$ , if we write  $v'_i$  for the perturbed version of  $v_i$  in the definition of  $A_i^{(\epsilon)}$ , then  $\sum_{j \in S} (w_j + p_j^* - p_j^{(\epsilon)}) = (v'_i(S) - \sum_{j \in S} p_j^{(\epsilon)}) - (v_i(S) - \sum_{j \in S} p_j^*)$ . In other words,  $A_i^{(\epsilon)}$  breaks indifferences according to preference under the perturbed valuations and the corresponding median prices.

As the space of choice rules is compact (mapping a discrete set to probability distributions over a discrete set), there must be a subsequence of  $\mathbf{A}^{(\epsilon)}$  as  $\epsilon \rightarrow 0$  that converges with respect to total variation distance. Let  $\mathbf{A}^*$  denote the limit, which must

<sup>4</sup>These will not be witnesses for prices  $\mathbf{p}^{(\epsilon)}$ ; they will simply be helpful for defining  $\mathbf{A}^*$ .

be a choice rule. We claim that  $\mathbf{A}^*$  is a witnessing choice rule for  $\mathbf{p}^*$ .

To prove the claim, fix some  $\epsilon$  and a valuation profile  $\mathbf{v}$  drawn from  $D$ . For each  $i$  we will couple  $D_i^{(\epsilon)}$  and  $A_i^{(\epsilon)}$  by using the same uniform weights  $w_j \sim U[0, \epsilon]$  in the definition of each. Write  $v'_i$  for the resulting perturbed valuation from  $D_i^{(\epsilon)}$ . Let  $(S_1, \dots, S_n)$  be the sets allocated to the agents under valuations  $\mathbf{v}$ , prices  $\mathbf{p}$ , and choice rule  $A_i^{(\epsilon)}$ , with each  $S_i$  being chosen from a corresponding demand set  $D_i$ . Let  $(T_1, \dots, T_n)$  be the sets allocated under perturbed valuations  $\mathbf{v}'$  and prices  $\mathbf{p}^{(\epsilon)}$ . Since  $D^{(\epsilon)}$  is generic, each  $T_i$  uniquely maximizes  $u_i^{v'_i}(T_i, \mathbf{p}^{(\epsilon)})$  almost surely, so we condition on that event. Our goal now is to show that for all  $i$  and all sufficiently small  $\epsilon$ ,  $S_i = T_i$  with high probability.

Suppose there is some smallest  $i$  such that  $S_i \neq T_i$ . Set  $T_i$  uniquely maximizes  $u_i^{v'_i}(T_i, \mathbf{p}^{(\epsilon)})$  over all items available to agent  $i$ . On the other hand,  $S_i \in \arg \max_{S \in D_i} \{ \sum_{j \in S} (w_j + p_j^* - p_j^{(\epsilon)}) \} = \arg \max_{S \in D_i} \{ u_i^{v'_i}(S, \mathbf{p}^{(\epsilon)}) \}$ . Thus  $T_i \neq S_i$  implies  $T_i \notin D_i$ . So in particular we must have  $u_i^{v'_i}(T_i, \mathbf{p}^*) < u_i^{v_i}(S_i, \mathbf{p}^*)$  but  $u_i^{v'_i}(T_i, \mathbf{p}^{(\epsilon)}) > u_i^{v'_i}(T_i, \mathbf{p}^{(\epsilon)})$ . Since  $v'_i$  and  $v_i$  differ by at most  $m\epsilon$  on the value of any set, we conclude that  $T_i \neq S_i$  implies

$$\begin{aligned} u_i^{v_i}(T_i, \mathbf{p}^*) &< u_i^{v_i}(S_i, \mathbf{p}^*) \\ &< u_i^{v_i}(T_i, \mathbf{p}^*) + m\epsilon + \sum_j |p_j^* - p_j^{(\epsilon)}|. \end{aligned} \quad (2)$$

Our strategy will be to argue that (2) can occur only with small probability.

For any  $\gamma > 0$ ,  $i \in [m]$ , and  $S, T \subseteq [m]$ , let  $E(\gamma, i, S, T)$  denote the event that, for prices  $\mathbf{p}^*$  and valuation  $v_i$  drawn from  $D_i$ ,  $0 < u_i^{v_i}(S, \mathbf{p}^*) - u_i^{v_i}(T, \mathbf{p}^*) \leq \gamma$ . We claim that  $\lim_{\gamma \rightarrow 0} \Pr[E(\gamma, i, S, T)] = 0$  for any distribution over the values  $v_i(S)$  and  $v_i(T)$ . To see this, let  $F(x) = \Pr[u_i^{v_i}(S, \mathbf{p}^*) - u_i^{v_i}(T, \mathbf{p}^*) \leq x]$ . We have  $\Pr[E(\gamma, i, S, T)] = F(\gamma) - F(0)$  and, as  $F$  is the CDF of a random variable, it is right-side continuous, so  $\lim_{\gamma \rightarrow 0, \gamma > 0} F(\gamma) = F(0)$ .

Thus, if we define event  $E(\gamma)$  to be  $\cup_{i, S, T} E(\gamma, i, S, T)$ , a union bound implies that  $\lim_{\gamma \rightarrow 0} E(\gamma) = 0$  as well. So for any  $\delta > 0$  there exists  $\gamma_\delta > 0$  such that  $\Pr[E(\gamma_\delta)] < \delta$ . In particular, with probability at least  $1 - \delta$ , the loss in utility from any non-demanded set taken by any agent is at least  $\gamma_\delta$ .

Since  $\lim_{\epsilon \rightarrow 0} \mathbf{p}^{(\epsilon)} = \mathbf{p}^*$ , There exists some  $\epsilon_1 > 0$  such that for all  $\epsilon < \epsilon_1$  and all  $j$ ,  $\sum_j |p_j^* - p_j^{(\epsilon)}| < \gamma_\delta/2$ . We conclude that for all  $\epsilon < \min\{\epsilon_1, \gamma_\delta/(2m)\}$ , if event  $E(\gamma')$  does not occur then inequality (2) cannot hold, since if  $u_i^{v_i}(T_i, \mathbf{p}^*) < u_i^{v_i}(S_i, \mathbf{p}^*)$  then, from the definition of  $E(\gamma')$ ,

$$\begin{aligned} u_i^{v_i}(S_i, \mathbf{p}^*) &> u_i^{v_i}(T_i, \mathbf{p}^*) + \gamma_\delta \\ &= u_i^{v_i}(T_i, \mathbf{p}^*) + \gamma_\delta/2 + \gamma_\delta/2 \\ &> u_i^{v_i}(T_i, \mathbf{p}^*) + m\epsilon + \sum_j |p_j^* - p_j^{(\epsilon)}|. \end{aligned}$$

We conclude that if event  $E(\gamma_\delta)$  does not occur, then for all sufficiently small  $\epsilon$  it must be that  $S_i = T_i$  for all  $i$ , and

hence the distribution over allocations (and in particular the probability that each item  $j$  sells) is identical. Some notation: write  $\pi_j(D, \mathbf{p}, \mathbf{A})$  for the probability of sale of item  $j$  under valuation distributions  $D$ , prices  $\mathbf{p}$ , and choice rule  $\mathbf{A}$ , where we can omit  $\mathbf{A}$  from the notation for generic distributions. Then since  $\Pr[E(\gamma_\delta)] < \delta$ , we conclude that  $|\pi_j(D, \mathbf{p}^*, \mathbf{A}^{(\epsilon)}) - \pi_j(D^{(\epsilon)}, \mathbf{p}^{(\epsilon)})| < \delta$  for all  $j$  and all sufficiently small  $\epsilon$ . As  $\mathbf{A}^{(\epsilon)}$  converges to  $\mathbf{A}^*$  in total variation distance, we conclude that  $|\pi_j(D, \mathbf{p}^*, \mathbf{A}^*) - \pi_j(D^{(\epsilon)}, \mathbf{p}^{(\epsilon)})| < 2\delta$  for all sufficiently small  $\epsilon$ . Since  $\mathbf{p}^{(\epsilon)}$  are median prices for  $D^{(\epsilon)}$ , taking  $\delta \rightarrow 0$  completes the proof.  $\square$

**Learning Median Prices for General Distributions.** Finally, we will show that median prices not only exist for general distributions, we can also find them with polynomially many samples.

**Theorem IV.14.** *Let  $D$  be any (not necessarily generic) distribution over valuation profiles. With probability  $1 - \delta$ ,  $O(\frac{1}{\epsilon^2} (m^2 + m \log n + \log(1/\delta)))$  samples suffice to find  $(\frac{1}{2} - \epsilon)$ -median prices and their witnessing choice rule.*

We will consider choice rules that are determined by a tie-breaking vector  $\mathbf{q} \in \mathbb{R}^m$ . The choice rule  $\mathbf{A}^{\mathbf{q}}$  is defined as follows. Each agent  $i$  draws  $w_{i,1}, \dots, w_{i,m} \stackrel{\text{i.i.d.}}{\sim} U[0, 1]$  and then chooses whatever  $S$  in their demand correspondence maximizes  $\sum_{j \in S} (w_j - q_j)$  (any remaining ties broken arbitrarily in a fixed way, say lexicographically). A consequence of the proof of Theorem IV.13 is that for every  $\gamma > 0$  there are  $(\frac{1}{2} - \gamma)$ -median prices  $\mathbf{p}^*$  that are witnessed by  $\mathbf{A}^{\mathbf{q}^*}$  for some  $\mathbf{q}^* \in \mathbb{R}^m$ , namely setting  $q_j = \frac{p_j^* - p_j^{(\epsilon)}}{\epsilon'}$  the choice rules  $\mathbf{A}^{\mathbf{q}}$  and  $\mathbf{A}^{(\epsilon')}$  are equivalent. Now consider  $\epsilon'$  that is small enough so that the probability of sale for every item  $j$  under prices  $\mathbf{p}^*$  and  $\mathbf{A}^{(\epsilon')}$  are within  $[\frac{1}{2} - \gamma, \frac{1}{2} + \gamma]$ .

For any choice of a price vector  $\mathbf{p}$  and tie-breaking vector  $\mathbf{q}$ , we let  $\pi_j(\mathbf{p}, \mathbf{q})$  be the probability that item  $j$  is sold under  $\mathbf{p}$  with choice rule  $\mathbf{A}^{\mathbf{q}}$  being applied. Our goal is to find  $\mathbf{p}$  and  $\mathbf{q}$  based on samples from  $D$  such that  $\pi_j(\mathbf{p}, \mathbf{q})$  is close to  $\frac{1}{2}$ .

Let us call  $(v_i, w_i)$  the *extended valuation* of agent  $i$ . For a fixed price vector  $\mathbf{p} \in \mathbb{R}_+^m$  and a fixed tie-breaking vector  $\mathbf{q} \in \mathbb{R}^m$ , the extended valuation profile  $(\mathbf{v}, \mathbf{w})$  completely defines the outcome on  $\mathbf{p}$  is choice rule  $\mathbf{A}^{\mathbf{q}}$  is applied.

Consider again the empirical distribution defined by  $k$  drawn for extended valuation profiles  $(\mathbf{v}, \mathbf{w})$ . Note that  $k$  draws from  $D$  can easily be extended to draws of extended valuation profiles by appending  $w_{i,j}$  drawn i.i.d. from  $U[0, 1]$ . For the empirical distribution of extended valuation profiles, we let  $\hat{\pi}_j(\mathbf{p}, \mathbf{q})$  be the fraction of extended valuation profiles such that item  $j$  gets sold under prices  $\mathbf{p}$  with tie-breaking  $\mathbf{q}$ . We again have a uniform-convergence guarantee.

**Lemma IV.15 (Uniform Convergence).** *Consider the empirical distribution over  $k$  sampled extended valuation profiles, with  $k \geq C \frac{1}{\epsilon^2} (m^2 + m \log n + \log(1/\delta))$  for a constant  $C$ . Then, with probability at least  $1 - \delta$ , we have  $\|\hat{\pi}(\mathbf{p}, \mathbf{q}) - \pi(\mathbf{p}, \mathbf{q})\|_{\infty} < \epsilon$ .*

This bound implies Theorem IV.14. We know that there exists a  $(1/2 - \epsilon)$ -median price vector  $\mathbf{p}^*$  for  $D$  witnessed by  $\mathbf{A}^{\mathbf{q}^*}$  for some  $\mathbf{q}^*$ . So with probability  $1 - \delta$ ,  $\pi_j(\mathbf{p}^*, \mathbf{q}^*) \in [\frac{1}{2} - 2\epsilon, \frac{1}{2} + 2\epsilon]$  for all  $j$ . Let  $\mathbf{p}$  and  $\mathbf{q}$  be any vectors such that  $\hat{\pi}_j(\mathbf{p}, \mathbf{q}) \in [\frac{1}{2} - 2\epsilon, \frac{1}{2} + 2\epsilon]$  for all  $j$ . Then, again conditioning on the event of probability  $1 - \delta$ ,  $\hat{\pi}_j(\mathbf{p}, \mathbf{q}) \in [\frac{1}{2} - 3\epsilon, \frac{1}{2} + 3\epsilon]$  for all  $j$ .

*Proof of Lemma IV.15.* Consider the space of all possible extended valuation vectors  $(\mathbf{v}, \mathbf{w})$ . Every pair of a price vector  $\mathbf{p}$  and tie-breaking vector  $\mathbf{q}$  defines a subset  $S_j^{\mathbf{p}, \mathbf{q}}$  of this space, equal to the collection of valuation vectors for which item  $j$  is sold under price vector  $\mathbf{p}$ . By this definition  $\pi_j(\mathbf{p}, \mathbf{q}) = \Pr_{\mathbf{v} \sim D, \mathbf{w} \sim U^m[0,1]}[(\mathbf{v}, \mathbf{w}) \in S_j^{\mathbf{p}, \mathbf{q}}]$  and  $\hat{\pi}_j(\mathbf{p}, \mathbf{q}) = \Pr_{\mathbf{v} \sim \hat{D}, \mathbf{w} \sim U^m[0,1]}[(\mathbf{v}, \mathbf{w}) \in S_j^{\mathbf{p}}]$ .

Let  $\nu$  be the VC dimension of the set system  $\{S_j^{\mathbf{p}, \mathbf{q}}\}$ . We will show below in Lemma IV.16 that also the VC dimension of the set system defined by all of these subsets is upper bounded by  $O(m^2 + m \log n)$ .

Therefore, also here, for any item  $j$  and any  $\epsilon, \delta' > 0$ , this implies that  $O\left(\frac{\nu + \ln(1/\delta')}{\epsilon^2}\right)$  samples are sufficient for  $\Pr[\sup_{\mathbf{p}, \mathbf{q}} |\hat{\pi}_j(\mathbf{p}, \mathbf{q}) - \pi_j(\mathbf{p}, \mathbf{q})| \geq \epsilon] \leq \delta'$ .

Again, taking a union bound with  $\delta' = \frac{\delta}{m}$  over all items  $j$ ,  $O\left(\frac{\nu + \ln(1/\delta)}{\epsilon^2}\right)$  samples are sufficient for  $\Pr[\sup_{\mathbf{p}, \mathbf{q}} \|\hat{\pi}(\mathbf{p}, \mathbf{q}) - \pi(\mathbf{p}, \mathbf{q})\|_\infty \geq \epsilon] \leq \delta$ , meaning that with probability at least  $1 - \delta$ , we have that  $|\hat{\pi}_j(\mathbf{p}, \mathbf{q}) - \pi_j(\mathbf{p}, \mathbf{q})| < \epsilon$  for all items  $j$  and all price vectors  $\mathbf{p}$  and all tie-breaking vectors  $\mathbf{q}$ .  $\square$

The bound on the VC dimension works in a similar way as before.

**Lemma IV.16.** *The VC dimension of the set system  $\{S_j^{\mathbf{p}, \mathbf{q}}\}$  is at most  $O(m^2 + m \log n)$ .*

*Proof.* Let  $(\mathbf{v}^{(1)}, \mathbf{w}^{(1)}), \dots, (\mathbf{v}^{(N)}, \mathbf{w}^{(1)})$  be  $N$  extended valuation profiles that are shattered. That is, for every set  $T \subseteq [N]$ , there is a price vector  $\mathbf{p}^T$  and a tie-breaking vector  $\mathbf{q}^T$  such that item  $j$  gets sold on  $(v^{(t)}, w^{(t)})$  if and only if  $t \in T$ .

Note that if on an extended valuation profile  $(\mathbf{v}, \mathbf{w})$ , item  $j$  gets sold under  $(\mathbf{p}, \mathbf{q})$  but doesn't get sold under price vector  $(\mathbf{p}', \mathbf{q}')$ , there have to be  $i$ ,  $S$ , and  $S'$  such that agent  $i$  prefers  $S$  under  $(\mathbf{p}, \mathbf{q})$  but  $S'$  under  $(\mathbf{p}', \mathbf{q}')$  because otherwise the outcomes would be identical. That is,  $v_i(S) - \sum_{j' \in S} p_{j'} \geq v_i(S') - \sum_{j' \in S'} p_{j'}$  but  $v_i(S) - \sum_{j' \in S} p'_{j'} \leq v_i(S') - \sum_{j' \in S'} p'_{j'}$  and one of the two inequalities is strict or  $\sum_{j' \in S} (w_{i,j'} - p_{j'}) \geq \sum_{j' \in S'} (w_{i,j'} - p'_{j'})$  but  $\sum_{j' \in S} w_{i,j'} - p'_{j'} \leq \sum_{j' \in S'} (w_{i,j'} - p'_{j'})$  and one of the two inequalities is strict.

That is, if we consider the  $2m$ -dimensional space  $\mathbb{R}_+^m \times \mathbb{R}^m$  of all price vectors and tie-breaking vectors, the vectors  $(\mathbf{p}, \mathbf{q})$  and  $(\mathbf{p}', \mathbf{q}')$  either lie on different sides of the hyperplane defined by  $v_i(S) - v_i(S') - \sum_{j' \in S \setminus S'} p_{j'}^* + \sum_{j' \in S' \setminus S} p_{j'}^* = 0$  or they lie both on the hyperplane but on different sides of the one defined by  $\sum_{j' \in S} w_{i,j} - \sum_{j' \in S'} w_{i,j} - \sum_{j' \in S \setminus S'} q_{j'}^* + \sum_{j' \in S' \setminus S} q_{j'}^* = 0$ .

The extended valuation profiles  $(\mathbf{v}^{(1)}, \mathbf{w}^{(1)}), \dots, (\mathbf{v}^{(N)}, \mathbf{w}^{(N)})$  define  $N \cdot n \cdot (2^m)^2$  such hyperplanes of each kind, so  $2 \cdot N \cdot n \cdot (2^m)^2$  hyperplane in total. This means there are no more than  $2m(2 \cdot N \cdot n \cdot (2^m)^2)^{2m}$  different regions in  $\mathbb{R}^{2m}$ . If  $(\mathbf{v}^{(1)}, \mathbf{w}^{(1)}), \dots, (\mathbf{v}^{(N)}, \mathbf{w}^{(N)})$  is shattered, each of the  $2^N$  vectors  $(\mathbf{p}^T, \mathbf{q}^T)$  for  $T \subseteq [N]$  has to lie in a different region.

We conclude that  $2^N \leq 2m(N \cdot n \cdot (2^m)^2)^{2m}$ . Taking logs on both sides yields  $N \leq \log m + m \log N + m \log n + 2m^2$ , which implies  $N = O(m^2 + m \log n)$ .  $\square$

## V. OPEN PROBLEMS AND FUTURE DIRECTIONS

In this work we showed that a single sample from each bidder's value distribution is sufficient to obtain an  $O(1)$ -approximation for the online combinatorial allocation problem with XOS bidders, while a polynomial number of samples per bidder suffices to get a  $(2 + \epsilon)$ -approximate truthful mechanism. An exciting frontier opened by our work is whether the problem with XOS bidders admits a *truthful*  $O(1)$ -approximate mechanism with a poly-logarithmic number of samples, or even only a single sample from each bidder's value distribution. Another natural next step, would be to explore whether the online combinatorial allocation problem with *subadditive* bidders can be solved with poly-many samples, or even with a single sample.

## APPENDIX A OMITTED PROOFS

### A. Proof of Lemma III.2

*Proof of Lemma III.2.* First observe that the prices in the buyer-wise greedy algorithm are non-decreasing. This can be shown by contradiction. If we had  $a_j(v_i, A_i) < p_j^{(i)}$  for some  $i \in [n]$  and  $j \in A_i$ , then buyer  $i$  would strictly prefer bundle  $A_i \setminus \{j\}$  over  $A_i$ , in contradiction to  $A_i = d_i(\mathbf{p}^{(i)}) \in \mathcal{D}(v_i, \mathbf{p}^{(i)}, [m])$ . Indeed, if we had  $a_j(v_i, A_i) < p_j^{(i)}$ , then

$$\begin{aligned} v_i(A_i \setminus \{j\}) - \sum_{j \in A_i \setminus \{j\}} p_j^{(i)} &\geq \sum_{j \in A_i \setminus \{j\}} (a_j(v_i, A_i) - p_j^{(i)}) \\ &> \sum_{j \in A_i} (a_j(v_i, A_i) - p_j^{(i)}) \\ &= v_i(A_i) - \sum_{j \in A_i} p_j^{(i)}, \end{aligned}$$

where we used that  $a(v_i, A_i)$  is an additive supporting function for the first and last step, and the strict inequality comes from the assumption.

Using this, we now show the approximation guarantee. We will first lower bound  $\mathbf{G}(\mathbf{v}) \geq \sum_{j \in [m]} p_j^{(n+1)}$ . Afterwards, we will show that  $\text{OPT}(\mathbf{v}) \leq \sum_{j \in [m]} p_j^{(n+1)} + \sum_{i \in [n]} u_i(A_i, \mathbf{p}^{(i)})$  and that  $\sum_{i \in [n]} u_i(A_i, \mathbf{p}^{(i)}) \leq \mathbf{G}(\mathbf{v})$ . Together this shows that  $\text{OPT}(\mathbf{v}) \leq 2\mathbf{G}(\mathbf{v})$ , as claimed.

Let's begin with the lower bound on  $\mathbf{G}(\mathbf{v})$ . Since  $\mathbf{G}_i(\mathbf{v}) \subseteq A_i$  for all  $i \in [n]$ ,  $a_j(v_i, A_i) = p_j^{(n+1)}$  for all  $i \in [n]$  and  $j \in$

$G_i(\mathbf{v})$ , and  $p_j^{(n+1)} = 0$  for all  $j \in [m]$  that remain unassigned, we obtain

$$G(\mathbf{v}) = \sum_{i \in [n]} v_i(G_i(\mathbf{v})) \geq \sum_{i \in [n]} \sum_{j \in G_i(\mathbf{v})} p_j^{(n+1)} = \sum_{j \in [m]} p_j^{(n+1)}. \quad (3)$$

For the upper bound on  $\text{OPT}(\mathbf{v})$ , we proceed in two steps. First, observe that for each  $i \in [n]$ , since  $A_i = d(\mathbf{p}^{(i)}) \in \mathcal{D}(v_i, \mathbf{p}^{(i)}, [m])$ , and prices are monotonically increasing, it holds that

$$\begin{aligned} v_i(\text{OPT}_i(\mathbf{v})) - \sum_{j \in \text{OPT}_i(\mathbf{v})} p_j^{(n+1)} &\leq v_i(\text{OPT}_i(\mathbf{v})) - \sum_{j \in \text{OPT}_i(\mathbf{v})} p_j^{(i)} \\ &\leq v_i(A_i) - \sum_{j \in A_i} p_j^{(i)}. \end{aligned} \quad (4)$$

Note that the right-hand side  $v_i(A_i) - \sum_{j \in A_i} p_j^{(i)} = u_i(A_i, \mathbf{p}^{(i)})$ . We now use that in each step of the algorithm, the sum of the additive supporting valuations increases exactly by the updating buyer's utility. Using this, we obtain

$$\begin{aligned} G(\mathbf{v}) &\geq \sum_{i \in [n]} \sum_{j \in G_i(\mathbf{v})} a_j(v_i, A_i) \\ &= \sum_{i=1}^n \left( v_i(A_i) - \sum_{j \in A_i} p_j^{(i)} \right) \\ &= \sum_{i \in [n]} \left( v_i(A_i) - \sum_{j \in A_i} p_j^{(i)} \right). \end{aligned} \quad (5)$$

By summing Inequality (4) over all  $i \in [n]$  and using Equality (5), we obtain

$$\text{OPT}(\mathbf{v}) - \sum_{j \in [m]} p_j^{(n+1)} \leq \sum_{i \in [n]} \left( v_i(A_i) - \sum_{j \in A_i} p_j^{(i)} \right) \leq G(\mathbf{v}).$$

Rearranging this, shows  $\text{OPT}(\mathbf{v}) \leq \sum_{j \in [m]} p_j^{(n+1)} + G(\mathbf{v})$ , as claimed.  $\square$

### B. Proof of Lemma III.3

*Proof of Lemma III.3.* The proof of the approximation guarantee follows from the following observations. First note that as in the unmodified algorithm, prices are monotonically increasing. We can then use that  $a_j(v_i, A_i) \geq p_j^{(n+1)}/2$  to conclude, similar to Inequality (3), that  $\hat{G}(\mathbf{v}) \geq \frac{1}{2} \sum_{j \in [m]} p_j^{(n+1)}$ . Moreover, the arguments that lead to Inequality (4) also apply to the modified algorithm, so that

$$\text{OPT}(\mathbf{v}) \leq \sum_{j \in [m]} p_j^{(n+1)} + \sum_{i \in [n]} u_i(A_i, \mathbf{p}^{(i)}) \quad (6)$$

The argument is completed by observing that in the modified algorithm, although the sum of the additive supporting valuations no longer increase by exactly the updating buyer's utility, it still grows by *at least* this amount. With this we can proceed

as in Inequality (5), except that the first equality becomes an inequality, and conclude that

$$\hat{G}(\mathbf{v}) = \sum_{i \in [n]} \sum_{j \in \hat{G}(\mathbf{v})} a_j(v_i, A_i) \geq \sum_{i \in [n]} u_i(A_i, \mathbf{p}^{(i)}). \quad (7)$$

Putting everything together, we obtain  $\text{OPT}(\mathbf{v}) \leq \sum_{j \in [m]} p_j^{(n+1)} + \sum_{i \in [n]} u_i(A_i, \mathbf{p}^{(i)}) \leq 3\hat{G}(\mathbf{v})$  as claimed.  $\square$

### C. Proof of Lemma III.4

*Proof of Lemma III.4.* Denote the demand set of agent  $i$  with respect to valuation  $r_i$  when prices  $\mathbf{p}^{(i)}$  are set through  $\mathbf{s}_{<i}$  by  $B_i$ ; and let's write  $A_i$  for the demand set of agent  $i$  with respect to valuation  $s_i$  when prices  $\mathbf{p}^{(i)}$  are set through  $\mathbf{s}_{<i}$ . Since  $\mathbf{r}$  and  $\mathbf{s}$  are generated one-by-one by flipping an independent fair coin, for all  $i \in [n]$ , it holds that  $\mathbb{E}[r_i(B_i)] = \mathbb{E}[s_i(A_i)]$  and  $\mathbb{E}\left[\sum_{j \in B_i} p_j^{(i)}\right] = \mathbb{E}\left[\sum_{j \in A_i} p_j^{(i)}\right]$ .

Summing this over all  $i \in [n]$  yields

$$\begin{aligned} \mathbb{E}\left[\sum_{i \in [n]} r_i(B_i)\right] &\geq \mathbb{E}\left[\sum_{i \in [n]} \left( r_i(B_i) - \sum_{j \in B_i} p_j^{(i)} \right) \right] \\ &= \mathbb{E}\left[\sum_{i \in [n]} \left( s_i(A_i) - \sum_{j \in A_i} p_j^{(i)} \right) \right] \\ &\geq \mathbb{E}\left[\text{OPT}(\mathbf{s}) - \sum_{j \in [m]} p_j^{(n+1)}\right], \end{aligned}$$

where the inequality holds by the same arguments as those that lead to Inequality (6).

Moreover, we have that

$$\begin{aligned} \mathbb{E}\left[\sum_{j \in [m]} p_j^{(n+1)}\right] &= \mathbb{E}\left[\sum_{i \in [n]} \sum_{j \in A_i} (p_j^{(i+1)} - p_j^{(i)})\right] \\ &\leq \mathbb{E}\left[\sum_{i \in [n]} \sum_{j \in A_i} p_j^{(i+1)}\right] \leq \mathbb{E}\left[\sum_{i \in [n]} 2s_i(A_i)\right]. \end{aligned}$$

We conclude that

$$\mathbb{E}\left[\sum_{i \in [n]} r_i(B_i)\right] + \mathbb{E}\left[\sum_{i \in [n]} 2s_i(A_i)\right] \geq \mathbb{E}[\text{OPT}(\mathbf{s})],$$

which shows that  $3\mathbb{E}\left[\sum_{i \in [n]} r_i(B_i)\right] \geq \mathbb{E}[\text{OPT}(\mathbf{s})]$  as claimed.  $\square$

## APPENDIX B MAKING DISTRIBUTIONS GENERIC

**Proposition B.1.** *Given any distribution  $D$  over valuations and any  $\epsilon > 0$ , there exists a generic distribution  $D'$  and a coupling  $F$  between  $D$  and  $D'$  such that for all  $(v, v') \sim F$ ,  $|v(S) - v'(S)| < \epsilon$  for all  $S \subseteq [m]$ .*

*Proof.* Given any distribution  $D$  over a class of valuations, consider the following distribution  $D'$ . To draw  $v' \sim D'$ , first draw  $v \sim D$ , then draw  $w_j \sim U[0, \epsilon]$  for each item  $j$ . Then we define  $v'(S) = v(S) + \sum_{j \in S} w_j$  for each  $S \subseteq [m]$ .

We need to prove that  $D'$  is generic. Choose any price vector  $\mathbf{p}$  and any  $M \subseteq [m]$ . Choose any  $S, T \subseteq M$  with  $S \neq T$ . We will show that the probability that  $S$  and  $T$  both lie in  $\mathcal{D}(v', \mathbf{p}, M)$  is 0. To prove the claim, note that in order for both  $S$  and  $T$  to be in  $\mathcal{D}(v', \mathbf{p}, M)$ , a necessary condition is that  $v'(S) - \sum_{j \in S} p_j = v'(T) - \sum_{j \in T} p_j$ . From the definition of  $v'$ , this occurs only if

$$\sum_{j \in S - T} w_j - \sum_{j \in T - S} w_j = (v(T) - v(S)) + \left( \sum_{j \in S} p_j - \sum_{j \in T} p_j \right).$$

Once we draw  $v \sim F$  but before we draw the  $w_j$  terms, the right-hand side is a fixed constant and the left-hand side is a random variable with no atoms (since  $S \neq T$ ). So the probability of equality is 0. Taking a union bound over all  $S$  and  $T$  yields the desired result.  $\square$

## APPENDIX C EXAMPLES FOR SIMPLE ALGORITHMS

We give examples demonstrating the difficulties arising for ‘simple’ sample-based algorithms (for both the allocation and the auction problem), outlining the need for our more elaborate approach.

**The Natural Single-Choice Extension.** Probably the most natural idea is to directly extend the single-choice SSPI that simply sets the largest observed sample value as a threshold. As a first idea, one might attempt to accordingly set a price for each item that equals  $\max_{v_s \in S} v_s(j)$ . However, imagine an instance where every buyer has an additive value of 1 for each item, except for buyer 1, who has a value of 2 for each nonempty subset of the items, i.e. also for each singleton: then, the algorithm never sells more than one item, resulting in a ratio of  $\Omega(m)$ .

**Item Prices as Contribution to Offline Optimum.** Consider the following sequential posted pricing algorithm based on a single sample. Given a single sample, compute the offline optimum allocation for that sample, then price each item according to its supporting price under that allocation (capturing its ‘contribution’ to said optimum). This approach doesn’t work. Suppose  $n = m$  and consider a distribution  $D$  over profiles of unit-demand valuations, such that player  $n$  always has value 1 for item 1 and value  $\epsilon$  for every other item. Each other player and item has value  $\epsilon/2$ . Then assuming  $\epsilon \ll 1/n$ , the optimum for any sample assigns item 1 to player 1 for a total value of  $1 + O(n\epsilon)$ . So this approach sets  $p_1 = 1$  and  $p_j = \epsilon/2$  for all  $j \neq 1$ . But under these prices, player 1 never chooses item 1, so the total value generated is  $O(n\epsilon)$  (since a truthful mechanism must assign buyers their favourite bundle under the given prices). Without truthfulness, the algorithm could instead e.g. assign each player the largest-valued item that is still price-feasible. While this approach seems promising, one needs to find a way of showing that such an algorithm will not sell out high-valued items before the according buyers arrive.

**Balanced Prices from a Single Sample.** Motivated by balanced pricing, suppose we instead take a single sample, compute an offline optimal allocation, then set each item’s price to be

half of its supporting price. This also does not work. Suppose  $n = m$  and consider a distribution  $D$  over profiles of unit-demand valuations, such that player  $n$  has value 1 for one of the items (uniformly at random), and value  $\epsilon$  for the remaining items. Each other player has value  $\epsilon$  for every item. Then this pricing approach will result in price  $1/2$  for one of the items (uniformly at random), say item  $j$ , and price  $\epsilon/2$  for all other items. In the real valuation, the optimal allocation has value  $1 + O(n\epsilon)$  by assigning to player  $n$  whichever item is valued at 1 (say item  $i$ ). But unless  $j = i$ , player  $n$  will receive no items: all items other than  $j$  will be taken by the first  $n - 1$  players, leaving only item  $j$  which has price  $1/2 > \epsilon$ . So the total value obtained will be  $O(n\epsilon)$ .

The construction above assumed a deterministic order in which agents choose their items. Note that if we modify the example by taking  $n \gg m$ , then the analysis would extend to (a) an algorithm where the agent order is randomized for sequential posted pricing, and even (b) scenarios where there isn’t deterministically a single agent with high value for an item, but all agent valuations are drawn iid from a distribution with probability  $O(1/n)$  of having the ‘high’ type (value 1 for some item uniformly at random).

In conclusion, it is not possible to correctly estimate the value our algorithms ‘should’ obtain from an item to compute prices. Instead, our algorithms’ key property is using the fact that both samples and actual valuation stem from the same distribution to bound the probability that items are prematurely assigned.

## ACKNOWLEDGEMENTS

We are thankful to Shuchi Chawla, Trung Dang, and anonymous reviewers for comments on earlier versions of the paper. The last author is thankful to Ellen Vitercik for explaining [20].

## REFERENCES

- [1] M. Feldman, N. Gravin, and B. Lucier, ‘‘Combinatorial auctions via posted prices,’’ in *Proceedings of SODA*, 2015. 1, 2, 3, 5, 7, 13
- [2] P. Dütting, M. Feldman, T. Kesselheim, and B. Lucier, ‘‘Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs,’’ in *Proceedings of FOCS*, 2017. 1, 3, 5
- [3] P. Dütting, T. Kesselheim, and B. Lucier, ‘‘An  $o(\log \log m)$  prophet inequality for subadditive combinatorial auctions,’’ in *Proceedings of FOCS*, 2020, pp. 306–317. 1, 5
- [4] J. R. Correa and A. Cristi, ‘‘A constant factor prophet inequality for online combinatorial auctions,’’ in *Proceedings of STOC*, 2023, pp. 686–697. 1, 5
- [5] W. Vickrey, ‘‘Counterspeculation, auctions, and competitive sealed tenders,’’ *The Journal of finance*, vol. 16, no. 1, pp. 8–37, 1961. 1
- [6] E. H. Clarke, ‘‘Multipart pricing of public goods,’’ *Public choice*, vol. 11, no. 1, pp. 17–33, 1971. 1
- [7] T. Groves *et al.*, ‘‘Incentives in teams,’’ *Econometrica*, vol. 41, no. 4, pp. 617–631, 1973. 1

[8] P. D. Azar, R. Kleinberg, and S. M. Weinberg, “Prophet inequalities with limited information,” in *Proceedings of SODA*, 2014, pp. 1358–1377. 1, 2, 5, 8

[9] A. Rubinstein, J. Z. Wang, and S. M. Weinberg, “Optimal single-choice prophet inequalities from samples,” in *Proceedings of ITCS*, ser. LIPIcs, vol. 151, 2020, pp. 60:1–60:10. 2, 5, 8

[10] H. Kaplan, D. Naori, and D. Raz, “Online weighted matching with a sample,” in *Proceedings of SODA* 2022, 2022, pp. 1247–1272. 2, 5

[11] C. Caramanis, P. Dütting, M. Faw, F. Fusco, P. Lazos, S. Leonardi, O. Papadigenopoulos, E. Pountourakis, and R. Reiffenhäuser, “Single-sample prophet inequalities via greedy-ordered selection,” in *Proceedings of SODA*, 2022. 2, 3, 5, 10

[12] M. Balcan, “Data-driven algorithm design,” in *Beyond the Worst-Case Analysis of Algorithms*, T. Roughgarden, Ed. Cambridge University Press, 2020, pp. 626–645. 2

[13] U. Krengel and L. Sucheston, “Semiamarts and finite values,” *Bulletin of the American Mathematical Society*, vol. 83, p. 745–747, 1977. 2, 5

[14] ———, “On semiamarts, amarts, and processes with finite value,” *Advances in Probability and Related Topics*, vol. 4, p. 197–266, 1978. 2, 5

[15] R. Kleinberg and S. M. Weinberg, “Matroid prophet inequalities,” in *Proceedings of STOC*, 2012, p. 123–136. 2, 5

[16] E. Samuel-Cahn, “Comparison of threshold stop rules and maximum for independent nonnegative random variables,” *The Annals of Probability*, vol. 12, p. 1213–1216, 1984. 2, 4, 5

[17] J. R. Correa, A. Cristi, B. Epstein, and J. A. Soto, “The two-sided game of googol and sample-based prophet inequalities,” in *Proceedings of SODA*, 2020, pp. 2066–2081. 3

[18] S. Dobzinski, N. Nisan, and M. Schapira, “Approximation algorithms for combinatorial auctions with complement-free bidders,” in *Proceedings of STOC*, 2005, pp. 610–618. 3, 7

[19] M. Balcan, T. Sandholm, and E. Vitercik, “A general theory of sample complexity for multi-item profit maximization,” in *Proceedings of EC*, 2018, pp. 173–174. 4, 13

[20] M. Balcan, D. F. DeBlasio, T. Dick, C. Kingsford, T. Sandholm, and E. Vitercik, “How much data is sufficient to learn high-performing algorithms? generalization guarantees for data-driven algorithm design,” in *Proceedings of STOC*, 2021, pp. 919–932. 4, 13, 19

[21] M. T. Hajiaghayi, R. D. Kleinberg, and T. Sandholm, “Automated online mechanism design and prophet inequalities,” in *Proceedings of AAAI*, 2007, pp. 58–65. 5

[22] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan, “Multi-parameter mechanism design and sequential posted pricing,” in *Proceedings of STOC*, 2010, pp. 311–320. 5

[23] M. Feldman, O. Svensson, and R. Zenklusen, “Online contention resolution schemes,” in *Proceedings of SODA*, 2016, pp. 1014–1033. 5

[24] K. Banihashem, M. Hajiaghayi, D. R. Kowalski, P. Krysta, and J. Olkowski, “Power of posted-price mechanisms for prophet inequalities,” in *Proceedings of SODA*, 2024, pp. 4580–4604. 5

[25] A. Gupta and S. Singla, “Random-order models,” in *Beyond the Worst-Case Analysis of Algorithms*, T. Roughgarden, Ed. Cambridge University Press, 2020, pp. 234–258. 5

[26] T. Kesselheim, K. Radke, A. Tönnis, and B. Vöcking, “An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions,” in *Proceedings of ESA*, vol. 8125, 2013, pp. 589–600. 5

[27] J. R. Correa, P. Dütting, F. A. Fischer, and K. Schewior, “Prophet inequalities for I.I.D. random variables from an unknown distribution,” in *Proceedings of EC*. ACM, 2019, pp. 3–17. 5

[28] C. Guo, Z. Huang, Z. G. Tang, , and X. Zhang, “Generalizing complex hypotheses on product distributions: Auctions, prophet inequalities, and pandora’s problem,” in *Proceedings of COLT*, 2021, pp. 248–2288. 5

[29] A. Cristi and B. Ziliotto, “Prophet inequalities require only a constant number of samples,” in *Proceedings of STOC*, 2024, pp. 491–502. 5

[30] H. Kaplan, D. Naori, and D. Raz, “Competitive analysis with a sample and the secretary problem,” in *Proceedings of SODA*, 2020, pp. 2082–2095. 5

[31] S. Assadi and S. Singla, “Improved truthful mechanisms for combinatorial auctions with submodular bidders,” in *Proceedings of FOCS*, 2019. 5

[32] S. Assadi, T. Kesselheim, and S. Singla, “Improved truthful mechanisms for subadditive combinatorial auctions: Breaking the logarithmic barrier,” in *Proceedings of SODA*, 2021, pp. 653–661. 5

[33] P. Dütting and T. Kesselheim, “Posted pricing and prophet inequalities with inaccurate priors,” in *Proceedings of EC*, 2019, pp. 111–129. 5

[34] U. Feige, “On maximizing welfare when utility functions are subadditive,” *SIAM J. Comput.*, vol. 39, no. 1, pp. 122–142, 2009. 6

[35] M. J. Wainwright, *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, ser. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019. 13

[36] R. C. Buck, “Partition of space,” *The American Mathematical Monthly*, vol. 50, no. 9, pp. 541–544, 1943. [Online]. Available: <https://doi.org/10.1080/00029890.1943.11991447> 14