

Prophet Inequalities Made Easy: Stochastic Optimization by Pricing Non-Stochastic Inputs

Paul Dütting*, Michal Feldman[†], Thomas Kesselheim[‡], Brendan Lucier[§]

^{*}*Department of Mathematics, London School of Economics, London, UK, Email: p.d.uetting@lse.ac.uk*

[†]*Blavatnic School of Computer Science, Tel Aviv University, Tel Aviv, Israel, Email: mfeldman@tau.ac.il*

[‡]*Department of Computer Science, TU Dortmund, Dortmund, Germany, Email: thomas.kesselheim@tu-dortmund.de*

[§]*Microsoft Research, New England, Cambridge, MA, USA, Email: brlucier@microsoft.com*

Abstract—We present a general framework for stochastic online maximization problems with combinatorial feasibility constraints. The framework establishes prophet inequalities by constructing price-based online approximation algorithms, a natural extension of threshold algorithms for settings beyond binary selection. Our analysis takes the form of an extension theorem: we derive sufficient conditions on prices when all weights are known in advance, then prove that the resulting approximation guarantees extend directly to stochastic settings. Our framework unifies and simplifies much of the existing literature on prophet inequalities and posted price mechanisms, and is used to derive new and improved results for combinatorial markets (with and without complements), multi-dimensional matroids, and sparse packing problems. Finally, we highlight a surprising connection between the smoothness framework for bounding the price of anarchy of mechanisms and our framework, and show that many smooth mechanisms can be recast as posted price mechanisms with comparable performance guarantees.

Keywords—mechanism design; posted prices; price of anarchy; prophet inequalities; smoothness

I. INTRODUCTION

A concert is being held in a local theatre, and potential audience members begin calling to reserve seats. The organizer doesn't know individuals' values for seats in advance, but has distributional knowledge about their preferences. Some need only a single seat, others require a block of seats. Some think seats are very valuable, others are only willing to attend if tickets are very cheap. Some prefer front-row seats, some prefer to sit a few rows back, and some prefer the balcony. The organizer needs to decide which seats, if any, to allocate to each individual as they call. The goal is to maximize the total value (i.e., social welfare) of the seating arrangement.

Such stochastic online optimization problems have been studied for decades. A common goal is to attain “prophet

The work of M. Feldman was partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122.

This work was done while T. Kesselheim was at Max Planck Institute for Informatics and Saarland University, supported in part by the DFG through Cluster of Excellence MMCI, and while he was visiting the Simons Institute for the Theory of Computing.

inequalities” that compare the performance of an online algorithm to that of an omniscient offline planner. A classic result is that if the goal is to choose exactly one element (i.e., there is only a single seat to allocate), then a simple threshold strategy—choosing the first value higher than a certain pre-computed threshold—yields at least half of the expected maximum value [1], [2], [3]. This solution has the appealing property that it corresponds to posting a take-it-or-leave-it price and allocating to the first interested buyer. A natural question is whether more complex allocation problems (like the concert example above) can be approximated by posting prices and allowing buyers to select their preferred outcomes in sequence.

Driven in part by this connection to posted prices, prophet inequalities have seen a resurgence in theoretical computer science. Recent work has established new prophet inequalities for a variety of allocation problems, including matroids [4], [5], unit-demand bidders [4], [6], and combinatorial auctions [7]. In this paper we develop a framework for proving prophet inequalities and constructing posted-price mechanisms. Our framework, which is based on insights from economic theory, unifies and simplifies many existing results and gives rise to new and improved prophet inequalities in a host of online settings.

A. Example: Combinatorial Auctions

To introduce our framework we will consider a combinatorial auction problem. There is a set M of m items for sale and n buyers. Each buyer i has a valuation function $v_i: 2^M \rightarrow \mathbb{R}_{\geq 0}$ that assigns non-negative value to every subset of at most d items.¹ Valuations are non-decreasing and normalized so that $v_i(\emptyset) = 0$, but otherwise arbitrary. The goal is to assign items to buyers to maximize total value. Write $\mathbf{v}(\mathbf{x}) = \sum_{i=1}^n v_i(x_i)$ for the total value of allocation $\mathbf{x} = (x_1, \dots, x_n)$, where $x_i \subseteq M$ for all i . There is a simple $O(d)$ -approximate greedy algorithm for this problem and a lower bound of $\Omega(d/\log d)$ assuming $P \neq NP$ [8]. Our goal is to match this $O(d)$ approximation as a prophet inequality

¹Alternatively, we can suppose that there is a cardinality constraint that no buyer can receive more than d items.

with posted item prices. That is, given distributions over the valuations, compute prices for the items so that, when buyers arrive in an arbitrary order and each chooses his most-desired bundle from among the unsold items, the expected total value is an $O(d)$ approximation to the expected optimum.²

Let's first consider the simpler full information case where all valuations are known in advance. This problem is still non-trivial, and in fact there may not exist prices that lead to the optimal allocation.³ Intuitively, what we need for an approximation result are prices that balance between two forces. They should be small enough that high-valued buyers are willing to purchase their optimal bundles if available, but also large enough that those items will not first be scooped up by bidders with much lower values. Such "balanced" prices can be obtained as follows: Given valuation profile \mathbf{v} , consider the welfare-maximizing allocation \mathbf{x}^* (which we can assume allocates all items). Then for each item j , say $j \in x_i^*$, set the price of j to $p_j = v_i(x_i^*)/(2|x_i^*|)$. These prices are low enough that the total price of all items is at most $1/2 \cdot \mathbf{v}(\mathbf{x}^*)$, which is significantly less than the total value of \mathbf{x}^* . At the same time, prices are high enough that, for any set of goods S , the total price of S is at least $1/2d$ of the value of allocations in the optimal allocation \mathbf{x}^* that intersect S . So, in particular, a bidder that purchases S must have value at least that high.

To see why these prices yield an $O(d)$ approximation, let \mathbf{x} denote the purchase decisions of the players and let $I \subseteq N$ be the set of players i such that x_i^* intersects with \mathbf{x} . The welfare achieved by \mathbf{x} is equal to the revenue generated plus the sum of buyer utilities. The revenue is the sum of prices of the items sold, and since prices are "balanced" this is at least $(1/2d) \cdot \sum_{i \in I} v_i(x_i^*)$. Also, each buyer $i \notin I$ could have chosen to purchase x_i^* , and therefore must get at least as much utility as they would by purchasing x_i^* , which is $v_i(x_i^*)$ minus the price of x_i^* . Again, since prices are balanced, this means the sum of buyer utilities is at least $\sum_{i \notin I} v_i(x_i^*) - 1/2 \cdot \mathbf{v}(\mathbf{x}^*)$. Multiplying this by $1/2d$ and

²There is a straightforward lower bound of $\Omega(d)$ on the approximation of any posted item prices. Suppose there are d items and two agents. The first agent is unit-demand and has value 1 for any single item. The second agent values the set of all d items for value d , and has value 0 for any subset. If all items have price greater than 1, then neither agent purchases anything. If any item has price less than 1, then the unit-demand agent (who chooses first) will purchase the cheapest single item and the other agent will purchase nothing, generating a total value of 1 whereas the optimum is d . One can avoid issues of tie-breaking by perturbing the values by an arbitrarily small amount.

³For example, suppose there are three items and four single-minded bidders. The first three bidders each have value 2 for a different pair of items, and the last bidder has value 3 for the set of all three items, so at most one bidder can get positive value, and it is optimal to allocate all items to the last bidder. However, at any item prices where the last bidder is willing to purchase, one of the other bidders will purchase first if arriving before the last bidder. This leads to a $3/2$ approximation in the worst arrival order.

adding the revenue gives an $O(d)$ approximation.⁴

The argument above was for the full information case. Perhaps surprisingly, the existence of sufficiently "balanced" prices for full information instances *also* establishes an $O(d)$ -approximate prophet inequality for the general stochastic problem, where one has only distributional knowledge about valuations. Our main result is this reduction from the stochastic setting to the full information setting, which holds for a broad class of allocation problems.

B. A Framework for Prophet Inequalities

Consider a more general combinatorial allocation problem, where the cardinality constraint d is replaced with an arbitrary downward-closed feasibility constraint \mathcal{F} and each v_i is drawn independently from an arbitrary distribution \mathcal{D}_i . While our framework applies for more general outcome spaces (see Sections II and III), combinatorial allocation problems provide a sweet spot between expressiveness and clarity. Our key definition is the following notion of balanced prices for full-information instances. For each $\mathbf{x} \in \mathcal{F}$ we write $\text{OPT}(\mathbf{v} \mid \mathbf{x})$ for the *optimal residual allocation*: the allocation that maximizes $\sum_i v_i(x'_i)$ over $\mathbf{x}' \in \mathcal{F}$ with \mathbf{x}, \mathbf{x}' disjoint and $\mathbf{x} \cup \mathbf{x}' \in \mathcal{F}$. Given a fixed valuation profile \mathbf{v} , a *pricing rule* defines a price $p_i^{\mathbf{y}}(x_i)$ for every bundle that we can assign to buyer i . For example, the item prices described in Section I-A define a pricing rule $p_i^{\mathbf{y}}(x_i) = \sum_{j \in x_i} p_j$. Below we also extend the definition to dynamic prices, i.e., prices that depend on which allocations have already been made.

Key Definition (special case) ((α, β) -balanced prices). *Let $\alpha, \beta > 0$. A pricing rule $\mathbf{p}^{\mathbf{v}} = (p_1^{\mathbf{y}}, \dots, p_n^{\mathbf{y}})$ defined by functions $p_i^{\mathbf{y}} : 2^M \rightarrow \mathbb{R}_{\geq 0}$ is (α, β) -balanced with respect to valuation profile \mathbf{v} if for all $\mathbf{x} \in \mathcal{F}$ and all $\mathbf{x}' \in \mathcal{F}$ with \mathbf{x}, \mathbf{x}' disjoint and $\mathbf{x} \cup \mathbf{x}' \in \mathcal{F}$,*

- (a) $\sum_i p_i^{\mathbf{y}}(x_i) \geq \frac{1}{\alpha} (\mathbf{v}(\text{OPT}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v} \mid \mathbf{x})))$,
- (b) $\sum_i p_i^{\mathbf{y}}(x'_i) \leq \beta \mathbf{v}(\text{OPT}(\mathbf{v} \mid \mathbf{x}))$.

The first condition formalizes what it means that prices are high enough: the sum of prices for \mathbf{x} should partially cover the welfare lost due to allocating \mathbf{x} . The second condition formalizes "low enough": the sum of prices for any \mathbf{x}' that is still feasible "after" allocating \mathbf{x} should not be much higher than the optimal residual welfare.

Our main result is that the existence of balanced prices for full information instances directly implies a price-based prophet inequality for the stochastic setting. The idea to choose balanced prices is a natural one and has appeared in the prophet inequality literature before, most explicitly in the notion of balanced thresholds of Kleinberg and

⁴For simplicity we assumed here that $\sum_{i \notin I} v_i(x_i^*) - 1/2 \cdot \mathbf{v}(\mathbf{x}^*) \geq 0$. More generally, since utilities are non-negative, the sum of buyer utilities is at least $\max\{\sum_{i \notin I} v_i(x_i^*) - 1/2 \cdot \mathbf{v}(\mathbf{x}^*), 0\}$. If the maximum is attained at 0, then $\sum_{i \in I} v_i(x_i^*) > 1/2 \cdot \mathbf{v}(\mathbf{x}^*)$ and the revenue alone exceeds $(1/4d) \cdot \mathbf{v}(\mathbf{x}^*)$, as desired.

Weinberg [5]. Previous definitions, however, applied to the stochastic setting directly, which made the construction and analysis of balanced thresholds inherently probabilistic. A main advantage of our framework is that it suffices to reason about the simpler full-information setting.

Main Theorem (informal). *Consider the setting where valuations are drawn from product distribution \mathcal{D} . Suppose that the pricing rule $\mathbf{p}^{\mathbf{v}}$ is (α, β) -balanced with respect to valuation profile \mathbf{v} . Then posting prices $p_i(x_i) = \frac{\alpha}{1+\alpha\beta} \mathbf{E}_{\tilde{\mathbf{v}} \sim \mathcal{D}} [p_i^{\tilde{\mathbf{v}}}(x_i)]$ achieves welfare at least $\frac{1}{1+\alpha\beta} \mathbf{E}[\mathbf{v}(\text{OPT}(\mathbf{v}))]$.*

In other words, to construct appropriate prices for a stochastic problem instance, it suffices to construct balanced prices for the full-information instances in its support and then post the expected values of those prices, scaled by an appropriate factor. The proof of our main theorem is similar in spirit to proofs in the price of anarchy literature [9], [10] or for establishing algorithmic stability [11], in that it uses “ghost samples.” It is, however, considerably more involved because of the sequential, online aspect of our problem.

Remark 1 (Weakly Balanced Prices). We also define a notion of weakly balanced prices, in which it suffices to upper bound the prices by $\beta \mathbf{v}(\text{OPT}(\mathbf{v}))$. In this case, we can show that posting an appropriately scaled version of the expected prices yields a $4\alpha\beta$ -approximate prophet inequality.

Remark 2 (Computation). It is sometimes easier to compute prices that are balanced with respect to an approximation algorithm ALG rather than OPT. Our result still applies in this case, with OPT replaced by ALG in the welfare guarantee. We also note that if the price rule \mathbf{p} in the main theorem is perturbed to some $\hat{\mathbf{p}}$ with $\|\mathbf{p} - \hat{\mathbf{p}}\|_{\infty} < \epsilon$, then the welfare guarantee degrades by at most an additive $O(n\epsilon)$ term. This robustness is desirable in itself, and also implies that appropriate prices can be computed for bounded values with $\text{POLY}(n, m, 1/\epsilon)$ samples using standard concentration bounds, as has been observed for various posted price settings [12], [7].

Remark 3 (Static vs. Dynamic, Anonymous vs. Discriminatory, Bundle vs. Item Pricing). We have described our framework for static, discriminatory, bundle prices. In general, our construction has the property that if the full-information balanced prices $\mathbf{p}^{\mathbf{v}}$ are dynamic, anonymous, and/or take the form of item prices, then the derived prices for the stochastic setting will have these properties as well. For example, our result holds also for dynamic prices, replacing $p_i(x_i)$ and $p_i(x'_i)$ with $p_i(x_i | \mathbf{x}_{[i-1]})$ and $p_i(x'_i | \mathbf{x}_{[i-1]})$ where the conditioning on $\mathbf{x}_{[i-1]}$ indicates that the price to player i may depend on the purchase decisions of players that precede him. See Sections II and III for details.

Remark 4 (Arrival Order). Balancedness can depend on

player arrival order. In the applications we consider, our results hold even if the arrival order is chosen by an adaptive adversary that observes previous realized values and purchase decisions before selecting the next player to arrive.

Let’s return to our example from Section I-A. We established the existence of weakly $(d, 1)$ -balanced prices (simply undo the scaling by $1/2$), so our main result implies a $O(d)$ -approximate prophet inequality. What about computation? We can compute prices in polynomial time by basing them on the $O(d)$ -approximate greedy algorithm rather than the optimal allocation, but then we only get a $O(d^2)$ -approximate solution. It turns out that we can further improve this to $O(d)$ in polynomial time, as we hoped for in Section I-A, by applying our main theorem to a fractional relaxation of the auction problem. See Section IV for details.

Composition: In the full version of this paper [13], we also show that balanced prices “compose,” as was shown for mechanism smoothness in [10]. This means that to derive a prophet inequality for a complex setting it often suffices to show balancedness for a simpler problem.

C. Unification of Existing Prophet Inequality Proofs

Our framework unifies and simplifies many of the existing prophet inequality proofs. We list some representative examples below. We discuss these examples in more detail in the full version.

Example 1 (Classic Prophet Inequality, [1], [2]). The goal is to pick the single highest-value element v_i . The pricing rule $\mathbf{p}^{\mathbf{v}}$ defined by $p_i^{\mathbf{v}}(x_i) = \max_i v_i$ for all i is $(1, 1)$ -balanced.

Example 2 (Matroids, [5]). The goal is to pick a maximum weight independent set in a matroid. Encode sets S by n -dimensional vectors \mathbf{x} over $\{0, 1\}$ such that $x_i = 1$ if $i \in S$. Then one can define a dynamic pricing rule $\mathbf{p}^{\mathbf{v}}$ by $p_i(x_i | \mathbf{y}) = \mathbf{v}(\text{OPT}(\mathbf{v} | \mathbf{y})) - \mathbf{v}(\text{OPT}(\mathbf{v} | \mathbf{y} \cup x_i))$ for all i , where \mathbf{y} is the set of previously-selected elements. This pricing rule is $(1, 1)$ -balanced.

Example 3 (XOS Combinatorial Auctions, [7]). The goal is to assign m goods to n buyers with XOS valuations.⁵ Let $\mathbf{x}^* = \text{OPT}(\mathbf{v})$ and let a_1, \dots, a_n be the corresponding additive supporting functions. Set item prices $p_j = a_i(j)$ for $j \in x_i^*$. This pricing rule is $(1, 1)$ -balanced.⁶

The final example also illustrates the power of our composition results. Namely, as we show in the full version, the existence of $(1, 1)$ -balanced prices for XOS combinatorial auctions follows directly from the existence of $(1, 1)$ -balanced prices for a single item.

⁵A valuation v is XOS if there is a collection of additive functions $a_1(\cdot), \dots, a_k(\cdot)$, such that for every set S , $v(S) = \max_{1 \leq i \leq k} a_i(S)$. This is a generalization of submodular valuations [14].

⁶Note that this approach also yields a $O(\log m)$ -approximate prophet inequality for subadditive valuations by approximating subadditive valuations with XOS valuations [15], [16].

D. New and Improved Prophet Inequalities

We also establish new prophet inequalities using our framework; see Table I. Our first result is a poly-time $(4k - 2)$ -approximate prophet inequality for MPH- k combinatorial auctions.⁷

Theorem 1 (Combinatorial auctions with MPH- k valuations). *For combinatorial auctions with MPH- k valuations, a $(4k - 2 + \epsilon)$ -approximate posted-price mechanism, with static item prices, can be computed in $\text{POLY}(n, m, 1/\epsilon)$ demand and MPH- k queries.*

Theorem 1 improves the poly-time result of [7] from $O(k^2)$ to $O(k)$. We note two interesting special cases. First, combinatorial auctions with bundle size d (from Section I-A) belong to MPH- d , so Theorem 1 captures the poly-time $O(d)$ approximation discussed above. Second, XOS valuations coincide with MPH-1, so Theorem 1 improves the previously best known poly-time result of [7] from $2e/(e - 1)$ to 2, matching the existential lower bound.

The second set of new results includes Knapsack feasibility constraints and d -sparse Packing Integer Programs (PIPs), for which we obtain a constant- and a $O(d)$ -approximation, respectively.

Theorem 2 (Knapsack). *For Knapsack constraints, a factor $(5 + \epsilon)$ -approximate posted-price mechanism, with static prices, can be computed in time $\text{POLY}(n, 1/\epsilon)$. This improves to a $(3 + \epsilon)$ approximation if no individual demands more than half of the total capacity.*

Theorem 3 (Sparse PIPs). *For d -sparse Packing Integer Programs (PIPs) with constraint matrix $\mathbf{A} \in \mathbb{R}_{\geq 0}^{m \times n}$ where $a_{j,i} \leq 1/2$ for all i, j and unit capacities, a factor $(8d + \epsilon)$ -approximate posted-price mechanism, with static prices, can be computed in time $\text{POLY}(n, m, 1/\epsilon)$.*

To the best of our knowledge, Theorems 2 and 3 are the first prophet inequalities for these settings. We note that [18] derived a prophet inequality for closely-related fractional knapsack constraints, with approximation factor ≈ 11.657 . We obtain an improved prophet inequality for this fractional setting: a corollary of Theorem 1 (with $k = 1$) is that one can obtain a 2-approximation for a fractional knapsack constraint using a static per-unit price, even when knapsack weights are private and arbitrarily correlated with buyer values. See Section IV for more details.

Finally, we generalize the matroid prophet inequalities of Kleinberg and Weinberg [5] to settings where players make choices regarding multiple elements of a matroid, and have submodular preferences over subsets of elements.

Theorem 4 (Multi-Dimensional Matroids). *For matroid feasibility constraints and submodular valuations, there is a*

⁷The maximum over positive hypergraphs- k (MPH- k) hierarchy of valuations [17] is an inclusive hierarchy, where k measures the degree of complementarity.

$(4 + \epsilon)$ -approximate posted-price mechanism, with dynamic prices, that can be computed in $\text{POLY}(n, 1/\epsilon)$ value queries.

We discuss Theorem 1 and Theorem 2 in more detail in Section IV. Additional details regarding the other results can be found in the full version.

E. From Price of Anarchy to Prophet Inequalities

In the proof sketch in Section I-A, we derived a lower bound on buyer utility by considering a deviation to a certain purchasing decision. This deviation argument, which appears in the proof of our main result, is also useful for establishing price of anarchy bounds [9], [10]. There is a subtle but important difference, however. In smoothness proofs one considers deviations against a fixed strategy profile, while the prophet inequality problem is inherently temporal and agents deviate at different points in time. As it turns out, many smoothness proofs have a built-in charging scheme (which we refer to as outcome smoothness) that, under the assumption that critical payments are monotonically increasing, implies prophet inequalities with the same (asymptotic) approximation guarantee. Both outcome smoothness and monotonicity are necessary for this result to hold.

Theorem 5 (informal). *For general multi-parameter problems, if the first-price (i.e., pay-your-bid) mechanism based on declared welfare maximization has a price of anarchy of $O(\gamma)$ provable via outcome smoothness, and critical payments are monotonically increasing, then posting a scaled version of the critical payments yields a $O(\gamma)$ -approximate price-based prophet inequality.*

We also provide two “black-box reductions” for binary single-parameter settings, where price of anarchy guarantees of $O(\gamma)$ established by (normal) smoothness imply $O(\gamma^2)$ -approximate prophet inequalities. See Section V.

Using these results we can, for example, rederive the classic prophet inequality [1], [2] from the smoothness of the first-price single-item auction [10] or the matroid prophet inequality [5] from the smoothness of the pay-your-bid, declared welfare maximizing mechanism for selecting a maximum-weight basis [19].

F. Further Related Work

Prophet inequalities and their applicability as posted-price mechanisms were (re-)discovered in theoretical computer science by [20]. Subsequently, threshold-based prophet inequalities and posted-price mechanisms were developed for matroids and matroid intersection [4], [5], [21], polymatroids [22], unit-demand bidders [4], [6], and combinatorial auctions [6], [7].

Not all prophet inequalities in the literature are based on explicit thresholds. Examples include prophet inequalities for the generalized assignment problem [23], [24], matroids and matroid intersection [25], and for general binary feasibility constraints [26]. On the other hand, many posted-price

Table I

OVERVIEW OF APPLICATIONS. RESULTS ARE COMPUTATIONAL UNLESS OTHERWISE STATED. THE QUERY MODEL REFERS TO THE VALUATION ACCESS NEEDED FOR THE COMPUTATIONAL UPPER BOUNDS, WHERE “EXPLICIT” INDICATES THAT VALUATIONS CAN BE DESCRIBED EXPLICITLY. ALL RESULTS ARE ORDER OBLIVIOUS (SEE SECTION II).

Feasibility Constraint	Valuation Class	Pricing Model	Upper Bound	Query Model
Combinatorial Auction	XOS	Static, anonymous item prices	$\frac{2e}{e-1}$ [7] 2 [this work]	XOS, Demand
Combinatorial Auction	MPH- k	Static, anonymous item prices	$O(k^2)$ [7] $4k - 2$ [this work]	MPH, Demand
Matroid	Submodular	Dynamic prices	2 (existential) 4 (computational)	Value
Knapsack	Additive	Static, anonymous prices	3	Explicit
d -Sparse PIPs	Additive	Static, anonymous prices	$8d$	Explicit

mechanisms from the literature are constructed either without explicit reference to prophet inequalities or via different techniques. Chawla et al. [12] developed approximately-optimal (revenue-wise) posted-price mechanisms for unit-demand buyers. Posted-price mechanisms have subsequently been developed for a variety of other auction settings [27], [28], [29], [30]. Dynamic posted prices that give optimal welfare for unit-demand buyers were established in [31]. Recently, dynamic posted prices for various online settings have been considered, including k -server on the line and metrical task systems [32], and makespan minimization for scheduling problems [33].

Most recently, and in parallel to this work combinatorial prophet inequalities were developed in [34] and [35]. The former, amongst others, proves prophet inequalities for subadditive combinatorial auctions, but considers a different allocation model and is therefore incomparable. The latter, in turn, focuses on revenue and not welfare as we do here. Finally, [36] and [37] re-consider the classic prophet inequality setting, but focus on identical distributions or a large market setting with random or best arrival order.

The notion of smooth games was introduced by Roughgarden [9] as a tool for bounding the price of anarchy, which measures the inefficiency that can be incurred in equilibrium. This notion has been extended to mechanisms by Syrgkanis and Tardos [10]. Notions of outcome smoothness were considered in [38], [39].

II. GENERAL MODEL AND NOTATION

Problem Formulation: There is a set N of n agents. For each agent $i \in N$ there is an outcome space X_i containing a null outcome \emptyset . We write $X = X_1 \times \dots \times X_n$ for the joint outcome space. Given outcome profile $\mathbf{x} \in X$ and a subset of agents $S \subseteq N$, we will write \mathbf{x}_S for the outcome in which each $i \in S$ receives x_i and each $i \notin S$ receives \emptyset . Specifically, we will write $\mathbf{x}_{[i-1]}$ for allocation \mathbf{x} with the outcomes of agents i, \dots, n set to \emptyset . There is a subset $\mathcal{F} \subseteq X$ of *feasible* outcomes. We assume that \mathcal{F} is downward-closed, so that if $\mathbf{x} \in \mathcal{F}$ then also $\mathbf{x}_S \in \mathcal{F}$ for all $S \subseteq N$.

A *valuation function* for agent i is a function $v_i: X_i \rightarrow \mathbb{R}_{\geq 0}$. We will assume values are bounded, and without loss of generality scaled to lie in $[0, 1]$. Each agent i 's valuation v_i is drawn independently from a publicly known distribution \mathcal{D}_i . We write $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_n$ for the product distribution over the set $V = V_1 \times \dots \times V_n$ of valuation profiles. We often suppress dependence on \mathcal{D} from our notation when clear from context. Agent utilities are quasilinear: if agent i receives outcome x_i and makes a payment π_i , his utility is $u_i = v_i(x_i) - \pi_i$.

The *welfare* of outcome \mathbf{x} is $\mathbf{v}(\mathbf{x}) = \sum_i v_i(x_i)$. An outcome rule ALG maps each valuation profile to a feasible outcome. $\text{ALG}_i(\mathbf{v})$ denotes the outcome of agent i on input \mathbf{v} . We will write $\text{OPT}(\mathbf{v}, \mathcal{F}) = \arg \max_{\mathbf{x} \in \mathcal{F}} \{\mathbf{v}(\mathbf{x})\}$ for the welfare-maximizing outcome rule for \mathcal{F} , omitting the dependence on \mathcal{F} when it is clear from context.

Pricing Rules and Mechanisms: A *pricing rule* is a profile of functions $\mathbf{p} = (p_1, \dots, p_n)$ that assign prices to outcomes. We write $p_i(x_i | \mathbf{y})$ for the (non-negative) price assigned to outcome $x_i \in X_i$, offered to agent i , given partial allocation $\mathbf{y} \in \mathcal{F}$. Define $p_i(x_i) = p_i(x_i | \emptyset)$ for convenience. We require that $p_i(x_i | \mathbf{y}) = \infty$ for any x_i such that $(x_i, \mathbf{y}_{-i}) \notin \mathcal{F}$. A pricing rule is said to be *monotone non-decreasing* if $p_i(x_i | \mathbf{y}) \geq p_i(x_i | \mathbf{y}_S)$ for all i , $x_i \in X_i$, $\mathbf{y} \in X$, $(x_i, \mathbf{y}_{-i}) \in \mathcal{F}$, and $S \subseteq N$. In general, we allow prices to be dynamic and discriminatory. We refer to prices that do not depend on the partial allocation (apart from feasibility) as *static* and to prices that do not depend on the identity of the agent as *anonymous*.

A *posted-price mechanism* is defined by a pricing rule \mathbf{p} and an ordering over the agents. This pricing rule can, in general, depend on the distributions \mathcal{D} . The agents are approached sequentially. Each agent i is presented the menu of prices determined by p_i , given all previous allocations, and selects a utility-maximizing outcome. A posted-price mechanism is *order-oblivious* if it does not require the agents to be processed in a specific order. In all of the applications we consider, the mechanisms we construct are order-oblivious. It is well-known that every posted-price

mechanism is *truthful* [4].

Online Allocations and Prophet Inequalities: We consider stochastic allocation algorithms that can depend on the value distributions \mathcal{D} . That is, an allocation algorithm \mathcal{A} maps a value profile and distribution to a feasible outcome. We say \mathcal{A} is an online allocation algorithm if $\mathcal{A}_i(\mathbf{v}, \mathcal{D})$ does not depend on the entries of \mathbf{v} that occur after i in some ordering over the indices. Extending the notion of competitive ratio from the worst-case analysis of online algorithms, we'll say the (*stochastic*) *competitive ratio* of online allocation algorithm \mathcal{A} is

$$\max_{\mathcal{D}} \frac{\mathbb{E}_{\mathbf{v} \sim \mathcal{D}}[\mathbf{v}(\text{OPT}(\mathbf{v}))]}{\mathbb{E}_{\mathbf{v} \sim \mathcal{D}}[\mathbf{v}(\mathcal{A}(\mathbf{v}, \mathcal{D}))]}.$$

We sometimes refer to a competitive ratio using its inverse, when convenient. A *prophet inequality* for constraint \mathcal{F} is an upper bound on the stochastic competitive ratio of an online allocation algorithm for \mathcal{F} . We note that a posted-price mechanism describes a particular form of an online allocation algorithm.

III. A FRAMEWORK FOR PROPHET INEQUALITIES

In this section we state and prove our main result, which reduces prophet inequalities to finding balanced prices for the simpler full information setting. We say that a set of outcome profiles $\mathcal{H} \subseteq X$ is *exchange compatible* with $\mathbf{x} \in \mathcal{F}$ if for all $\mathbf{y} \in \mathcal{H}$ and all $i \in N$, $(y_i, \mathbf{x}_{-i}) \in \mathcal{F}$. We call a family of sets $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$ *exchange compatible* if $\mathcal{F}_{\mathbf{x}}$ is exchange compatible with \mathbf{x} for all $\mathbf{x} \in X$.

Definition 1. Let $\alpha > 0$, $\beta \geq 0$. Given a set of feasible outcomes \mathcal{F} and a valuation profile \mathbf{v} , a pricing rule \mathbf{p} is (α, β) -*balanced* with respect to an allocation rule ALG, an exchange-compatible family of sets $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$, and an indexing of the players $i = 1, \dots, n$ if for all $\mathbf{x} \in \mathcal{F}$

- (a) $\sum_i p_i(x_i | \mathbf{x}_{[i-1]}) \geq \frac{1}{\alpha} \cdot (\mathbf{v}(\text{ALG}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})))$,
- (b) for all $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$:
 $\sum_i p_i(x'_i | \mathbf{x}_{[i-1]}) \leq \beta \cdot \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})).$

The definition provides flexibility in the precise choice of $\mathcal{F}_{\mathbf{x}}$. As $\mathcal{F}_{\mathbf{x}}$ becomes larger (more permissive), both inequalities become easier to satisfy since $\mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))$ increases. On the other hand, a larger set $\mathcal{F}_{\mathbf{x}}$ means that the second condition must be satisfied for more outcomes $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$. We say that a collection of pricing rules $(p^{\mathbf{v}})_{\mathbf{v} \in V}$ is (α, β) -balanced if there exists a choice of $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$ such that, for each \mathbf{v} , the pricing rule $p^{\mathbf{v}}$ is balanced with respect to $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$.

The definition of (α, β) -balancedness captures sufficient conditions for a posted-price mechanism to guarantee high welfare when agents have a known valuation profile \mathbf{v} . Our interest in (α, β) -balanced pricing rules comes from the fact that this result extends to Bayesian settings.

Theorem 6. *Suppose that the collection of pricing rules $(p^{\mathbf{v}})_{\mathbf{v} \in V}$ for feasible outcomes \mathcal{F} and valuation profiles*

$\mathbf{v} \in V$ is (α, β) -balanced with respect to allocation rule ALG and indexing of the players $i = 1, \dots, n$. Then for $\delta = \frac{\alpha}{1+\alpha\beta}$ the posted-price mechanism with pricing rule $\delta \mathbf{p}$, where $p_i(x_i | \mathbf{y}) = \mathbb{E}_{\tilde{\mathbf{v}}} [p_i^{\tilde{\mathbf{v}}}(x_i | \mathbf{y})]$, generates welfare at least $\frac{1}{1+\alpha\beta} \cdot \mathbb{E}_{\mathbf{v}}[\mathbf{v}(\text{ALG}(\mathbf{v}))]$ when approaching players in the order they are indexed.

Proof: We denote the exchange-compatible family of sets with respect to which the collection of pricing rules $(p^{\mathbf{v}})_{\mathbf{v} \in V}$ is balanced by $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$. We will first use Property (b) to show a lower bound on the utilities of the players, and Property (a) to show a lower bound on the revenue of the posted-price mechanism. We will then add these together to obtain a bound on the social welfare.

We will write $\mathbf{x}(\mathbf{v})$ for the allocation returned by the posted-price mechanism on input valuation profile \mathbf{v} and $\mathbf{x}'(\mathbf{v}, \mathbf{v}') = \text{OPT}(\mathbf{v}', \mathcal{F}_{\mathbf{x}(\mathbf{v})})$ for the welfare-maximizing allocation with respect to valuation profile \mathbf{v}' under feasibility constraint $\mathcal{F}_{\mathbf{x}(\mathbf{v})}$.

Utility bound: We obtain a lower bound on the expected utility of a player as follows. We sample valuations $\mathbf{v}' \sim \mathcal{D}$. Player i now considers buying $\text{OPT}_i((v_i, \mathbf{v}'_{-i}), \mathcal{F}_{\mathbf{x}(v'_i, \mathbf{v}'_{-i})})$ at price $\delta \cdot p_i(\text{OPT}_i((v_i, \mathbf{v}'_{-i}), \mathcal{F}_{\mathbf{x}(v'_i, \mathbf{v}'_{-i})}) | \mathbf{x}_{[i-1]}(\mathbf{v}))$. Taking expectations and exploiting that $\mathbf{x}_{[i-1]}(\mathbf{v})$ does not depend on v_i we obtain

$$\begin{aligned} \mathbb{E}_{\tilde{\mathbf{v}}} [u_i(\mathbf{v})] &\geq \mathbb{E}_{\mathbf{v}, \mathbf{v}'} \left[v_i (\text{OPT}_i((v_i, \mathbf{v}'_{-i}), \mathcal{F}_{\mathbf{x}(v'_i, \mathbf{v}'_{-i})})) - \right. \\ &\quad \left. \delta p_i \left(\text{OPT}_i((v_i, \mathbf{v}'_{-i}), \mathcal{F}_{\mathbf{x}(v'_i, \mathbf{v}'_{-i})}) \mid \mathbf{x}_{[i-1]}(\mathbf{v}) \right) \right] \\ &= \mathbb{E}_{\mathbf{v}, \mathbf{v}'} \left[v'_i (x'_i(\mathbf{v}, \mathbf{v}')) - \delta p_i (x'_i(\mathbf{v}, \mathbf{v}') \mid \mathbf{x}_{[i-1]}(\mathbf{v})) \right]. \end{aligned}$$

Summing the previous inequality over all agents we get

$$\begin{aligned} \mathbb{E}_{\tilde{\mathbf{v}}} \left[\sum_{i \in N} u_i(\mathbf{v}) \right] &\geq \mathbb{E}_{\mathbf{v}, \mathbf{v}'} \left[\sum_{i \in N} v'_i (x'_i(\mathbf{v}, \mathbf{v}')) \right] - \\ &\quad \mathbb{E}_{\mathbf{v}, \mathbf{v}'} \left[\sum_{i \in N} \delta \cdot p_i (x'_i(\mathbf{v}, \mathbf{v}') \mid \mathbf{x}_{[i-1]}(\mathbf{v})) \right] \\ &= \mathbb{E}_{\mathbf{v}, \mathbf{v}'} \left[\mathbf{v}' (\text{OPT}(\mathbf{v}', \mathcal{F}_{\mathbf{x}(\mathbf{v})})) \right] - \\ &\quad \mathbb{E}_{\mathbf{v}, \mathbf{v}'} \left[\sum_{i \in N} \delta \cdot p_i (x'_i(\mathbf{v}, \mathbf{v}') \mid \mathbf{x}_{[i-1]}(\mathbf{v})) \right]. \end{aligned} \tag{1}$$

We can upper bound the last term in the previous inequality by using Property (b). This gives

$$\sum_{i \in N} \delta p_i (x'_i(\mathbf{v}, \mathbf{v}') \mid \mathbf{x}_{[i-1]}(\mathbf{v})) \leq \delta \beta \cdot \mathbb{E}_{\tilde{\mathbf{v}}} \left[\tilde{\mathbf{v}} (\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})})) \right]$$

pointwise for any \mathbf{v} and \mathbf{v}' , and therefore also

$$\begin{aligned} \mathbb{E}_{\mathbf{v}, \mathbf{v}'} \left[\sum_{i \in N} \delta \cdot p_i (x'_i(\mathbf{v}, \mathbf{v}') \mid \mathbf{x}_{[i-1]}(\mathbf{v})) \right] \\ \leq \delta \beta \cdot \mathbb{E}_{\tilde{\mathbf{v}}} \left[\tilde{\mathbf{v}} (\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})})) \right]. \end{aligned} \tag{2}$$

Replacing \mathbf{v}' with $\tilde{\mathbf{v}}$ in Inequality (1) and combining it with Inequality (2) we obtain

$$\mathbb{E}_{\tilde{\mathbf{v}}} \left[\sum_{i \in N} u_i(\mathbf{v}) \right] \geq (1 - \delta\beta) \cdot \mathbb{E}_{\tilde{\mathbf{v}}, \mathbf{v}} \left[\tilde{\mathbf{v}} \left(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})}) \right) \right]. \quad (3)$$

Revenue bound: The second step is a lower bound on the revenue achieved by the posted-price mechanism. Applying Property (a) we obtain

$$\begin{aligned} \sum_{i \in N} \delta \cdot p_i(x_i(\mathbf{v}) \mid \mathbf{x}_{[i-1]}(\mathbf{v})) \\ &= \delta \cdot \sum_{i \in N} \mathbb{E}_{\tilde{\mathbf{v}}} [p_i^{\tilde{\mathbf{v}}}(x_i(\mathbf{v}) \mid \mathbf{x}_{[i-1]}(\mathbf{v}))] \\ &\geq \frac{\delta}{\alpha} \cdot \mathbb{E}_{\tilde{\mathbf{v}}} \left[\tilde{\mathbf{v}} (\text{ALG}(\tilde{\mathbf{v}})) - \tilde{\mathbf{v}} (\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})})) \right]. \end{aligned}$$

Taking expectation over \mathbf{v} this shows

$$\begin{aligned} \mathbb{E}_{\tilde{\mathbf{v}}} \left[\sum_{i \in N} \delta \cdot p_i(x_i(\mathbf{v}) \mid \mathbf{x}_{[i-1]}(\mathbf{v})) \right] \\ \geq \frac{\delta}{\alpha} \cdot \mathbb{E}_{\tilde{\mathbf{v}}} [\tilde{\mathbf{v}} (\text{ALG}(\tilde{\mathbf{v}}))] - \frac{\delta}{\alpha} \cdot \mathbb{E}_{\tilde{\mathbf{v}}, \mathbf{v}} [\tilde{\mathbf{v}} (\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})}))]. \end{aligned} \quad (4)$$

Combination: It remains to show how the two bounds can be combined so that they imply the approximation guarantee. By quasi-linearity we can rewrite the expected social welfare that is achieved by the posted-price mechanism as the sum of the expected utilities plus the expected revenue. Using $\delta = \alpha/(1 + \alpha\beta)$ and Inequalities (3) and (4), this gives

$$\begin{aligned} \mathbb{E}_{\tilde{\mathbf{v}}} \left[\sum_{i \in N} v_i(x_i(\mathbf{v})) \right] \\ &\geq \mathbb{E}_{\tilde{\mathbf{v}}} \left[\sum_{i \in N} u_i(\mathbf{v}) \right] + \mathbb{E}_{\tilde{\mathbf{v}}} \left[\sum_{i \in N} \delta \cdot p_i(x_i(\mathbf{v}) \mid \mathbf{x}_{[i-1]}(\mathbf{v})) \right] \\ &\geq (1 - \delta\beta) \cdot \mathbb{E}_{\tilde{\mathbf{v}}, \mathbf{v}} [\tilde{\mathbf{v}} (\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})}))] + \\ &\quad \frac{\delta}{\alpha} \cdot \mathbb{E}_{\tilde{\mathbf{v}}} [\tilde{\mathbf{v}} (\text{ALG}(\tilde{\mathbf{v}}))] - \frac{\delta}{\alpha} \cdot \mathbb{E}_{\tilde{\mathbf{v}}, \mathbf{v}} [\tilde{\mathbf{v}} (\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})}))] \\ &= \frac{1}{1 + \alpha\beta} \cdot \mathbb{E}_{\tilde{\mathbf{v}}} [\tilde{\mathbf{v}} (\text{ALG}(\tilde{\mathbf{v}}))]. \quad \blacksquare \end{aligned}$$

In what follows, we provide an alternative definition of balancedness, in which Property (b) is refined. This definition will be useful for some applications, as exemplified in Section IV.

Definition 2. Let $\alpha > 0, \beta_1, \beta_2 \geq 0$. Given a set of feasible outcomes \mathcal{F} and a valuation profile \mathbf{v} , a pricing rule \mathbf{p} is *weakly* $(\alpha, \beta_1, \beta_2)$ -balanced with respect to allocation rule ALG, an exchange-compatible family of sets $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$, and an indexing of the players $i = 1, \dots, n$ if for all $\mathbf{x} \in \mathcal{F}$

- (a) $\sum_i p_i(x_i \mid \mathbf{x}_{[i-1]}) \geq \frac{1}{\alpha} \cdot (\mathbf{v}(\text{ALG}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})))$,
- (b) for all $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$: $\sum_i p_i(x'_i \mid \mathbf{x}_{[i-1]}) \leq \beta_1 \cdot \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) + \beta_2 \cdot \mathbf{v}(\text{ALG}(\mathbf{v}))$.

The following theorem specifies the refined bound on the welfare that is obtained by weakly $(\alpha, \beta_1, \beta_2)$ -balanced pricing rules. Its proof appears in the full version.

Theorem 7. Suppose that the collection of pricing rules $(\mathbf{p}^{\mathbf{v}})_{\mathbf{v} \in V}$ for feasible outcomes \mathcal{F} and valuation profiles $\mathbf{v} \in V$ is weakly $(\alpha, \beta_1, \beta_2)$ -balanced with respect to allocation ALG and indexing of the players $i = 1, \dots, n$ with $\beta_1 + \beta_2 \geq \frac{1}{\alpha}$. Then for $\delta = \frac{1}{\beta_1 + \max\{2\beta_2, 1/\alpha\}}$ the posted-price mechanism with pricing rule $\delta \mathbf{p}$, where $p_i(x_i \mid \mathbf{y}) = \mathbb{E}_{\tilde{\mathbf{v}}} [p_i^{\tilde{\mathbf{v}}}(x_i \mid \mathbf{y})]$, generates welfare at least $\frac{1}{\alpha(2\beta_1 + 4\beta_2)} \cdot \mathbb{E}_{\mathbf{v}} [\mathbf{v}(\text{ALG}(\mathbf{v}))]$ when approaching players in the order they are indexed.

IV. NEW AND IMPROVED PROPHET INEQUALITIES

We have already argued that our framework unifies and simplifies many of the existing prophet inequality proofs. In this section we show how it can be used to derive new and improved bounds on the approximation ratio that can be obtained via price-based prophet inequalities. We highlight two results: the new poly-time $O(d)$ -approximation for combinatorial auctions with bundle size at most d , and the new poly-time constant-approximation for knapsack problems. Additional results are provided in the full version, and include combinatorial auctions with MPH- k valuations, d -sparse packing integer programs, and multi-dimensional matroids (where the result follows from the Rota exchange theorem [40, Lemma 2.7] and our composition results).

Combinatorial Auctions with Bounded Bundle Size: An existential $O(d)$ -approximate price-based prophet inequality is presented in Section I-A. Combined with the $O(d)$ -approximation greedy algorithm for this setting, it gives a poly-time $O(d^2)$ -approximate price-based prophet inequality (as shown in [7]). In what follows we use the flexibility of our framework to work directly with a relaxation of the allocation problem, thereby improving the approximation of the prophet inequality from $O(d^2)$ to $O(d)$. This is a special case of Theorem 1, which is proved in the full version.

Theorem 8. For combinatorial auctions where every agent can get at most d items, there exist weakly $(1, 1, d - 1)$ -balanced item prices that are static, anonymous, and order oblivious. Moreover, a $(4d - 2 - \epsilon)$ -approximate posted-price mechanism can be computed in $\text{POLY}(n, m, 1/\epsilon)$ demand queries, where ϵ is an additive error due to sampling.

Proof: Consider the canonical fractional relaxation of the combinatorial auction problem: a feasible allocation is described by values $x_{i,S} \in [0, 1]$ for all $i \in N$ and $S \subseteq M$ such that $\sum_S x_{i,S} \leq 1$ for all i and $\sum_{i, S \ni j} x_{i,S} \leq 1$ for all $j \in M$. Take \mathcal{F} to be all such fractional allocations, and $\mathcal{F}_{\mathbf{x}}$ to be the set of fractional allocations \mathbf{y} such that $\sum_{i, S \ni j} (x_{i,S} + y_{i,S}) \leq 1$ for all $j \in M$, and $\sum_S y_{i,S} \leq 1$ for all i . As usual, we think of $\mathcal{F}_{\mathbf{x}}$ as the set of allocations that remain feasible given a partial allocation \mathbf{x} .

Consider the following pricing rule for fractional allocations. Given valuation profile \mathbf{v} , let \mathbf{x}^* be the welfare-maximizing fractional allocation. Then for each item j , set $p_j = \sum_i \sum_{S \ni j} x_{i,S}^* v_i(S)$. We claim that these prices are $(1, 1, d-1)$ -balanced with respect to the optimal allocation rule.

For Property (a), fix some $\mathbf{x} \in \mathcal{F}$. Write $x^j = \sum_{i, S \ni j} x_{i,S}$. Consider the following allocation $\mathbf{y} \in \mathcal{F}_{\mathbf{x}}$: for each S , choose $j_S \in \arg \max_{j \in S} \{x^j\}$. Set $y_{i,S} = (1 - x^{j_S}) \cdot x_{i,S}^*$. We think of \mathbf{y} as the optimal allocation \mathbf{x}^* adjusted downward to lie in $\mathcal{F}_{\mathbf{x}}$. We then have that

$$\begin{aligned} \mathbf{v}(\mathbf{x}^*) - \mathbf{v}(\mathbf{y}) &= \sum_i \sum_S x^{j_S} \cdot x_{i,S}^* \cdot v_i(S) \\ &= \sum_j x^j \sum_{i, S: j=j_S} x_{i,S}^* \cdot v_i(S) \\ &\leq \sum_j x^j \cdot p_j = \sum_i p_i(x_i). \end{aligned}$$

Property (a) follows since $\mathbf{v}(\mathbf{y}) \leq \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))$. For Property (b), fix $\mathbf{x} \in \mathcal{F}$ and $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$. Then

$$\begin{aligned} \sum_i p_i(x'_i) &\leq \sum_j (1 - x^j) p_j \\ &= \sum_j (1 - x^j) \sum_{i, S \ni j} x_{i,S}^* \cdot v_i(S) \\ &= \sum_{i, S} x_{i,S}^* \cdot v_i(S) \sum_{j \in S} (1 - x^j) \\ &= \sum_{i, S} (|S| - 1) x_{i,S}^* \cdot v_i(S) + \\ &\quad \sum_{i, S} x_{i,S}^* \cdot v_i(S) \cdot \left(1 - \sum_{j \in S} x^j\right). \end{aligned}$$

The first expression on the right-hand side is at most $(d-1)\mathbf{v}(\text{OPT}(\mathbf{v}))$, since $|S| \leq d$ whenever $x_{i,S}^* > 0$. For the second expression, note that it is at most the welfare of the allocation \mathbf{y} defined by $y_{i,S} = x_{i,S}^* \cdot (1 - \sum_{j \in S} x^j)^+$. Moreover, this allocation \mathbf{y} is in $\mathcal{F}_{\mathbf{x}}$. So the second expression is at most $\mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))$, giving Property (b).

Theorem 7 therefore yields prices that guarantee a $(4d-2)$ approximation for the fractional allocation problem, and an ϵ -approximation to those prices can be computed via sampling. To complete the proof, note that for every agent i , if all previous agents have selected integral outcomes, then agent i also has a utility-maximizing outcome that is integral. This is because any fractional allocation can be interpreted as a convex combination of integral allocations. These same prices therefore guarantee a $(4d-2-\epsilon)$ approximation even if the mechanism prohibits non-integral allocations from being purchased. ■

The more general Theorem 1 also improves the best-known poly-time prophet inequality for XOS valuations from $2e/(e-1)$ to 2 (which is tight [7]) and for MPH-

k valuations it improves the best known poly-time bounds from $O(k^2)$ to $O(k)$.

Knapsack: In the knapsack allocation problem, there is a single divisible unit of resource and each agent has a private value $v_i \geq 0$ for receiving at least $s_i \geq 0$ units. Assume for now that $s_i \leq 1/2$ for all i . We allow both v_i and s_i to be private information, drawn from a joint distribution. In our notation: $X_i = [0, \frac{1}{2}]$, $\mathcal{F} = \{\mathbf{x} \mid \sum_i x_i \leq 1\}$, and $v_i(x_i) = v_i$ if $x_i \geq s_i$ and $v_i(x_i) = 0$ otherwise. Based on an arbitrary allocation algorithm ALG, we design anonymous, static prices by setting $p_i(x_i \mid \mathbf{y}) = x_i \cdot \mathbf{v}(\text{ALG}(\mathbf{v}))$ if x_i can feasibly be added and ∞ otherwise. The following restates the second half of Theorem 2.

Theorem 9. *For the knapsack allocation problem in which no single agent can request more than half of the total capacity, the prices above are $(1, 2)$ -balanced with respect to ALG. This implies a $(3 + \epsilon)$ -approximate poly-time posted-price mechanism with a single static anonymous per-unit price.*

Proof: The poly-time claim follows from Theorem 6 with ALG set to the classic FPTAS for knapsack [41], so it suffices to prove balancedness. For any $\mathbf{x} \in \mathcal{F}$, let $\mathcal{F}_{\mathbf{x}} = \mathcal{F}$ if $\sum_i x_i < \frac{1}{2}$, and $\mathcal{F}_{\mathbf{x}} = \emptyset$ otherwise. Note that $\mathcal{F}_{\mathbf{x}}$ is exchange compatible with \mathbf{x} since, for any $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$ and any agent k , $x'_k + \sum_i x_i \leq 1$. To establish balancedness with respect to $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x}}$, we consider two cases based on the value of $\sum_i x_i$.

Case 1: $\sum_i x_i < \frac{1}{2}$. Property (a) is trivially fulfilled because $\mathbf{v}(\text{ALG}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) \leq \mathbf{v}(\text{OPT}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) = 0$. For Property (b), note that for any $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$, we have

$$\begin{aligned} \sum_i p_i(x'_i \mid \mathbf{x}_{[i-1]}) &= \sum_i x'_i \cdot \mathbf{v}(\text{ALG}(\mathbf{v})) \leq \mathbf{v}(\text{ALG}(\mathbf{v})) \\ &\leq \mathbf{v}(\text{OPT}(\mathbf{v})) = \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})). \end{aligned}$$

Case 2: $\sum_i x_i \geq \frac{1}{2}$. Property (b) is vacuous since $\mathcal{F}_{\mathbf{x}} = \emptyset$. For Property (a), we have

$$\begin{aligned} \sum_i p_i(x_i \mid \mathbf{x}_{[i-1]}) &= \sum_i x_i \cdot \mathbf{v}(\text{ALG}(\mathbf{v})) \geq \frac{1}{2} \mathbf{v}(\text{ALG}(\mathbf{v})) \\ &= \frac{1}{2} (\mathbf{v}(\text{ALG}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))). \quad \blacksquare \end{aligned}$$

We can remove the restriction that $s_i \leq 1/2$ as follows, completing the proof of Theorem 2. Consider the contribution to the expected optimal welfare separated into welfare from agents with $s_i \leq 1/2$, and agents with $s_i > 1/2$. The posted-price mechanism described above obtains a 3-approximation to the former. For the latter, a mechanism that treats the unit of resource as indivisible, and posts the best take-it-or-leave-it price for the entire unit, is a 2-approximation. This is because at most one agent with $s_i > 1/2$ can win in any realization. Thus, for any

distribution profile, one of these two mechanisms must be a 5-approximation to the unrestricted knapsack problem.⁸ One can therefore obtain a $(5 + \epsilon)$ -approximate price-based prophet inequality by estimating the expected welfare of each pricing scheme (via sampling) and selecting the better of the two. In the full version we show how to generalize the result for the knapsack problem to d -sparse packing integer programs.

Finally, consider the fractional version of the knapsack problem, where agents obtain partial value for receiving a portion of their desired allocation: $v_i(x_i) = v_i \cdot \min\{x_i/s_i, 1\}$. If we restrict allocations x_i to be multiples of some $\delta > 0$, this is a special case of a submodular combinatorial auction with $\lceil 1/\delta \rceil$ identical items. Since Theorem 1 implies that a fixed per-item price yields a 2-approximation for any δ , we can infer by taking the limit as $\delta \rightarrow 0$ that for any $\epsilon > 0$ there is a $(2 + \epsilon)$ -approximate poly-time posted-price mechanism for the fractional knapsack problem, with a single static anonymous per-unit price, even if each agent's size s_i is private and arbitrarily correlated with their value. As mentioned in Section I-D, this improves the previously best-known prophet inequality of ≈ 11.657 due to [18].

V. FROM PRICE OF ANARCHY TO PROPHET INEQUALITIES

In this section we explore the connection between balanced prices and mechanism smoothness. While generally smoothness does not suffice to conclude the existence of a posted-price mechanism with comparable welfare guarantee (see the full version), we will show that this is the case for typical smoothness proofs and present pretty general reductions from the problem of proving prophet inequalities to mechanism smoothness.

We first recall the definition of a smooth mechanism. A (possibly indirect) mechanism \mathcal{M}_π for an allocation problem π is defined by a bid space $B = B_1 \times \dots \times B_n$, an allocation rule $f : B \rightarrow \mathcal{F}$, and a payment rule $P : B \rightarrow \mathbb{R}_{\geq 0}^n$. We focus on first-price mechanisms, where $P_i(\mathbf{b}) = b_i(f(\mathbf{b}))$. Typically, mechanisms are defined for a collection of problems Π , in which case we will simply refer to the mechanism as \mathcal{M} .

Definition 3 (Syrngkanis and Tardos [10]). Mechanism \mathcal{M}_π is (λ, μ) -smooth for $\lambda, \mu \geq 0$ if for any valuation profile $\mathbf{v} \in V$ and any bid profile $\mathbf{b} \in B$ there exists a bid $b'_i(\mathbf{v}, b_i) \in B_i$ for each player $i \in N$ such that

$$\sum_{i \in N} u_i(b'_i, \mathbf{b}_{-i}) \geq \lambda \cdot \mathbf{v}(\text{OPT}(\mathbf{v})) - \mu \cdot \sum_{i \in N} P_i(\mathbf{b}).$$

⁸The worst case is when both mechanisms achieve the same expected welfare, which occurs if 3/5 of the expected welfare is due to agents with $s_i \leq 1/2$. The expected welfare of each mechanism is then $\frac{1}{3} \cdot \frac{3}{5} = \frac{1}{5}$ of the optimum.

A mechanism \mathcal{M} that is (λ, μ) -smooth has a price of anarchy (with respect to correlated and Bayes-Nash equilibria) of at most $\max\{\mu, 1\}/\lambda$ [10].

The following formal notion of a residual market will be useful for our further analysis. For any $\mathbf{x} \in \mathcal{F}$ we define the *contraction* of \mathcal{F} by \mathbf{x} , \mathcal{F}/\mathbf{x} , as follows. Let $N^+(\mathbf{x}) = \{i \in N \mid x_i \neq \emptyset\}$. Then $\mathcal{F}/\mathbf{x} = \{\mathbf{z} = (z_j)_{j \in N \setminus N^+(\mathbf{x})} \mid (\mathbf{z}, \mathbf{x}_{N^+(\mathbf{x})}) \in \mathcal{F}\}$. That is, \mathcal{F}/\mathbf{x} contains allocations to players who were allocated nothing in \mathbf{x} , that remain feasible when combined with the allocations in \mathbf{x} . We think of the contraction by \mathbf{x} as a *subinstance* on players $N \setminus N^+(\mathbf{x})$ with feasibility constraint \mathcal{F}/\mathbf{x} , and refer to it as the subinstance induced by \mathbf{x} . We say that a collection of problems Π is *subinstance closed* if for every $\pi \in \Pi$ with feasible allocations \mathcal{F} and every $\mathbf{x} \in \mathcal{F}$ the subinstance induced by \mathbf{x} is contained in Π . The contraction by \mathbf{x} also naturally leads to an exchange feasible set $\mathcal{F}_\mathbf{x}$ by padding the allocations $\mathbf{z} \in \mathcal{F}/\mathbf{x}$ with null outcomes. We refer to this $\mathcal{F}_\mathbf{x}$ as the *canonical exchange-feasible set*.

A. Warm-up: Binary, Single-Parameter Problems with Monotone Prices

We begin with a simple result that serves to illustrate the connection between balancedness and smoothness. We will show that if a binary, single-parameter problem has the property that the welfare-maximizing mechanism is (λ, μ) -smooth and its critical prices $\tau_i(\cdot \mid \mathbf{y})$ are non-decreasing in \mathbf{y} ,⁹ then there exists a pricing rule that is (α, β) -balanced, where $\alpha\beta = O(\max\{\mu, 1\}/\lambda)$. In particular, this implies that the welfare guarantee due to Theorem 6 is within a constant factor of the price of anarchy of the mechanism implied by smoothness.

Theorem 10. Consider a subinstance-closed collection of binary, single-parameter problems such that the first-price mechanism based on the welfare maximizing allocation rule OPT is (λ, μ) -smooth. If the critical prices $\tau_i(\cdot \mid \mathbf{y})$ are non-decreasing in \mathbf{y} then setting $p_i(1 \mid \mathbf{y}) = \max\{v_i, \tau_i(\mathbf{v}_{-i} \mid \mathbf{y})\}$ and $p_i(0 \mid \mathbf{y}) = 0$ is $(1, \frac{\mu+1+\lambda}{\lambda})$ -balanced with respect to OPT and the canonical exchange-feasible sets $(\mathcal{F}_\mathbf{x})_{\mathbf{x} \in X}$.

Proof: Fix any \mathbf{y} and $\mathbf{x} \in \mathcal{F}_\mathbf{y}$. Observe that by definition of the prices, it holds that

$$p_i(x_i \mid \mathbf{y}) \geq \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{(\emptyset, \mathbf{y}_{-i})})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{(x_i, \mathbf{y}_{-i})})). \quad (5)$$

To see this, first note that both sides of the inequality are equal to 0 if $x_i = 0$. If $x_i = 1$ and $v_i \geq \tau_i(\mathbf{v}_{-i} \mid \mathbf{y})$, then agent i is allocated in $\text{OPT}(\mathbf{v}, \mathcal{F}_{(\emptyset, \mathbf{y}_{-i})})$ and hence both sides of the inequality are equal to v_i . If $x_i = 1$ and $v_i < \tau_i(\mathbf{v}_{-i} \mid \mathbf{y})$, then agent i is not allocated in $\text{OPT}(\mathbf{v}, \mathcal{F}_{(\emptyset, \mathbf{y}_{-i})})$,

⁹The critical price $\tau_i(\mathbf{v}_{-i} \mid \mathbf{y})$ is the infimum of values v_i such that the mechanism allocates 1 to agent i on input (v_i, \mathbf{v}_{-i}) , in the problem subinstance induced by \mathbf{y} .

and hence the right-hand side of the inequality is at most the externality imposed by forcing an allocation to agent i , which is at most $\tau_i(\mathbf{v}_{-i} \mid \mathbf{y}) = p_i(x_i \mid \mathbf{y})$.

We are now ready to prove balancedness. To verify Condition (a), choose $\mathbf{x} \in \mathcal{F}$ and note that

$$\begin{aligned} & \sum_{i=1}^n p_i(x_i \mid \mathbf{x}_{[i-1]}) \\ & \geq \sum_{i=1}^n \left(\mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}_{[i-1]}})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}_{[i]}})) \right) \\ & = \mathbf{v}(\text{OPT}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) \end{aligned}$$

as required, where the inequality follows from Equation (5), and the equality follows by a telescoping sum. For Condition (b), we get

$$\begin{aligned} \sum_{i \in \mathbf{x}'} \tau_i(\mathbf{v}_{-i} \mid \mathbf{x}_{[i-1]}) & \leq \sum_{i \in \mathbf{x}'} \tau_i(\mathbf{v}_{-i} \mid \mathbf{x}) \\ & \leq \frac{\mu + 1}{\lambda} \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})), \end{aligned} \quad (6)$$

where the first inequality follows by the monotonicity of critical prices, and the second inequality follows by a known implication of smoothness [42] (see the full version). Therefore, for any $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$,

$$\begin{aligned} \sum_i p_i(x'_i \mid \mathbf{x}_{[i-1]}) & \leq \sum_{i \in \mathbf{x}'} v_i + \sum_{i \in \mathbf{x}'} \tau_i(\mathbf{v}_{-i} \mid \mathbf{x}_{[i-1]}) \\ & \leq \mathbf{v}(\mathbf{x}') + \frac{\mu + 1}{\lambda} \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) \\ & \leq \frac{\mu + 1 + \lambda}{\lambda} \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})), \end{aligned}$$

where the first inequality follows by replacing the maximum in the definition of the prices by a sum, the second inequality follows by Equation (6), and the last inequality follows by $\mathbf{v}(\mathbf{x}') \leq \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))$, since $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$. ■

B. General Problems and Outcome Smoothness

We proceed to show an implication from smoothness to prices that works in more general settings. It is based on the observation that many smoothness proofs proceed by showing that agent i could bid b'_i to get some target outcome x_i^* . We capture proofs that proceed in this manner through the following notion of outcome smoothness. Similar but different notions were considered in [38], [39].

Definition 4. A mechanism is (λ, μ) -outcome smooth for $\lambda, \mu \geq 0$ if for all valuation profiles $\mathbf{v} \in V$ there exists an outcome $\mathbf{x}'(\mathbf{v}) \in \mathcal{F}$ such that for all bid profiles $\mathbf{b} \in B$,

$$\begin{aligned} \sum_{i \in N} \left(v_i(x'_i) - \inf_{b'_i: f_i(b'_i, \mathbf{b}_{-i}) \geq x'_i} P_i(b'_i, \mathbf{b}_{-i}) \right) \\ \geq \lambda \cdot \mathbf{v}(\text{OPT}(\mathbf{v})) - \mu \cdot \sum_{i \in N} P_i(\mathbf{b}). \end{aligned}$$

We show that if a first-price, declared welfare maximizing mechanism (i.e., a mechanism with allocation rule $f(\mathbf{b}) =$

$\text{OPT}(\mathbf{b})$) is (λ, μ) -outcome smooth and has non-decreasing critical prices, then the critical prices for that mechanism (from the definition of outcome smoothness) can be used as posted prices that yield an $O(\lambda/\mu)$ approximation to the optimal welfare. Recall that these critical prices are different from the first-price payments that make up the mechanism's payment rule. This result has a mild technical caveat: we require that the mechanism continues to be smooth in a modified problem with multiple copies of each bidders. An allocation is feasible in the modified feasibility space \mathcal{F}' if it corresponds to a feasible allocation $\mathbf{x} \in \mathcal{F}$, with each x_i being partitioned between the copies of agent i .

Theorem 11. Fix valuation space V and feasibility space \mathcal{F} , and suppose \mathcal{F}' is an extension of \mathcal{F} as defined above. Suppose that the first-price mechanism based on the declared welfare maximizing allocation rule for valuation space V and feasibility space \mathcal{F}' has non-decreasing critical prices, and is (λ, μ) -outcome smooth for every \mathcal{F}'/\mathbf{z} . Then there is a collection of exchange-feasible sets $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$, and an allocation rule ALG that returns the welfare-maximization allocation with probability λ , such that for every $\mathbf{v} \in V$ there exists a pricing rule that is $(\lambda, \mu/\lambda)$ -balanced with respect to ALG and $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$.

Theorem 11 implies that posting (an appropriately scaled version of) the critical prices from the outcome smooth mechanism yields a welfare approximation of $O(\lambda/\mu)$, matching the price of anarchy guarantee of the original mechanism. The proof of Theorem 11 appears in the full version.

C. Binary, Single-Parameter Problems

We conclude with two general “black-box reductions” for binary single-parameter settings, in which agents can either win or lose, which show how to translate price of anarchy guarantees of $O(\gamma)$ provable via (regular) smoothness into $O(\gamma^2)$ -approximate posted-price mechanisms. The key to both these results is a novel, purely combinatorial implication of smoothness for the greedy allocation rule. Proofs appear in the full version.

Theorem 12. Suppose that the first-price mechanism based on the greedy allocation rule GRD has a price of anarchy of $O(\gamma)$ provable via smoothness, then there exists a $O(\gamma^2)$ -approximate price-based prophet inequality.

Theorem 13. Suppose that the first-price mechanism based on the declared welfare maximizing allocation rule OPT has a price of anarchy of $O(\gamma)$ provable via smoothness, then there exists a $O(\gamma^2)$ -approximate price-based prophet inequality.

We note that Theorem 10 applied to matroids (using known smoothness results for pay-your-bid greedy mechanisms over matroids [19]) implies the existence of $(1, 3)$ -

balanced prices and hence a 4-approximate prophet inequality. A strengthening of Theorem 13 for monotonically increasing critical prices (discussed in the full version) leads to an improved factor of 2, matching the prophet inequality for matroids shown in [5]. This also captures the classic single-item prophet inequality, as a special case.

VI. CONCLUSIONS AND OPEN PROBLEMS

We introduced a general framework for establishing prophet inequalities and posted price mechanisms for multi-dimensional settings. This work leaves many questions open.

A general class of questions is to determine the best approximation guarantee that a prophet inequality can achieve for a particular setting. For example, even for the intersection of two matroids there is a gap between the trivial lower bound of 2 and the upper bound of $4k - 2 = 6$. Similarly, in subadditive combinatorial auctions, the best-known upper bound is logarithmic in the number of items m [7], but again the best-known lower bound is 2, inherited from the case of a single item. Notably, the price of anarchy for simultaneous single-item auctions is known to be constant for subadditive valuations [43], but the proof does not use the smoothness framework and hence our results relating posted prices to smooth mechanisms do not directly apply.

A related question is whether there exist prophet inequalities that cannot be implemented using posted prices. Interestingly, we are not aware of any separation between the two so far. More generally, one could ask about the power of anonymous versus personalized prices, item versus bundle prices, static versus dynamic prices, and so on. For example, to what extent can *static* prices approximate the welfare under a matroid constraint, an intersection of matroids, or an arbitrary downward-closed feasibility constraint?

Regarding the pricing framework itself, it would be interesting to extend the notion of (α, β) -balancedness to allow randomization in a dynamic pricing rule, and to understand the additional power of randomization. One could also generalize beyond feasibility constraints to more general seller-side costs for allocations. For the connection between smoothness and balancedness, we leave open the question of removing the price-monotonicity condition from Theorem 11, or whether the approximation factors can be improved for our single-parameter reductions (Theorems 12 and 13). Finally, recent work has shown that smoothness guarantees often improve as markets grow large [44]; is there a corresponding result for balancedness?

REFERENCES

- [1] U. Krengel and L. Sucheston, "Semiamarts and finite values," *Bulletin of the American Mathematical Society*, vol. 83, pp. 745–747, 1977.
- [2] —, "On semiamarts, amarts, and processes with finite value," *Advances in Probability and Related Topics*, vol. 4, pp. 197–266, 1978.
- [3] E. Samuel-Cahn, "Comparison of threshold stop rules and maximum for independent nonnegative random variables," *Annals of Probability*, vol. 12, pp. 1213–1216, 1984.
- [4] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan, "Multi-parameter mechanism design and sequential posted pricing," in *Proceedings of the 42nd ACM Symposium on Theory of Computing*, 2010, pp. 311–320.
- [5] R. Kleinberg and S. M. Weinberg, "Matroid prophet inequalities," in *Proceedings of the 44th ACM Symposium on Theory of Computing*, 2012, pp. 123–136.
- [6] S. Alaei, "Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers," *SIAM Journal on Computing*, vol. 43, no. 2, pp. 930–972, 2014.
- [7] M. Feldman, N. Gravin, and B. Lucier, "Combinatorial auctions via posted prices," in *Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms*, 2015, pp. 123–135.
- [8] L. Trevisan, "Non-approximability results for optimization problems on bounded degree instances," in *Proceedings of the 33rd ACM Symposium on Theory of Computing*, 2001, pp. 453–461.
- [9] T. Roughgarden, "Intrinsic robustness of the price of anarchy," *Journal of the ACM*, vol. 62, no. 5, p. 32, 2015.
- [10] V. Syrgkanis and É. Tardos, "Composable and efficient mechanisms," in *Proceedings of the 45th ACM Symposium on Theory of Computing*, 2013, pp. 211–220.
- [11] M. Hardt, B. Recht, and Y. Singer, "Train faster, generalize better: Stability of stochastic gradient descent," in *Proceedings of the 33rd International Conference on Machine Learning*, 2016, pp. 1225–1234.
- [12] S. Chawla, J. D. Hartline, and R. D. Kleinberg, "Algorithmic pricing via virtual valuations," in *Proceedings of the 8th ACM Conference on Electronic Commerce*, 2007, pp. 243–251.
- [13] P. Dütting, M. Feldman, T. Kesselheim, and B. Lucier, "Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs," <https://arxiv.org/abs/1612.03161>, 2017.
- [14] B. Lehmann, D. Lehmann, and N. Nisan, "Combinatorial auctions with decreasing marginal utilities," in *Proceedings of the 3rd ACM Conference on Electronic Commerce*, 2001, pp. 18–28.
- [15] S. Dobzinski, "Two randomized mechanisms for combinatorial auctions," in *Proceedings of the 10th/11th APPROX-RANDOM Workshop*, 2007, pp. 89–103.
- [16] K. Bhawalkar and T. Roughgarden, "Welfare guarantees for combinatorial auctions with item bidding," in *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms*, 2011, pp. 700–709.
- [17] U. Feige, M. Feldman, N. Immorlica, R. Izsak, B. Lucier, and V. Syrgkanis, "A unifying hierarchy of valuations with complements and substitutes," in *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 2015, pp. 872–878.

- [18] M. Feldman, O. Svensson, and R. Zenklusen, “A simple $o(\log \log(\text{rank}))$ -competitive algorithm for the matroid secretary problem,” in *Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms*, 2015, pp. 1189–1201.
- [19] B. Lucier and V. Syrgkanis, “Greedy algorithms make efficient mechanisms,” in *Proceedings of the 16th ACM Conference on Economics and Computation*, 2015, pp. 221–238.
- [20] M. Hajiaghayi, R. Kleinberg, and T. W. Sandholm, “Automated mechanism design and prophet inequalities,” in *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, 2007, pp. 58–65.
- [21] P. D. Azar, R. Kleinberg, and S. M. Weinberg, “Prophet inequalities with limited information,” in *Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms*, 2014, pp. 1358–1377.
- [22] P. Dütting and R. Kleinberg, “Polymatroid prophet inequalities,” in *Proceedings of the 23rd European Symposium on Algorithms*, 2015, pp. 437–449.
- [23] S. Alaei, M. Hajiaghayi, and V. Liaghat, “Online prophet-inequality matching with applications to ad allocation,” in *Proceedings of the 13th ACM Conference on Electronic Commerce*, 2012, pp. 18–35.
- [24] —, “The online stochastic generalized assignment problem,” in *Proceedings of the 16th/17th APPROX-RANDOM Workshop*, 2013, pp. 11–25.
- [25] M. Feldman, O. Svensson, and R. Zenklusen, “Online contention resolution schemes,” in *Proceedings of the 27th ACM-SIAM Symposium on Discrete Algorithms*, 2016, pp. 1014–1033.
- [26] A. Rubinstein, “Beyond matroids: Secretary problem and prophet inequality with general constraints,” in *Proceedings of the 48th ACM Symposium on Theory of Computing*, 2016, 324–332.
- [27] S. Dobzinski, N. Nisan, and M. Schapira, “Truthful randomized mechanisms for combinatorial auctions,” in *Proceedings of the 38th ACM Symposium on Theory of Computing*, 2006, pp. 644–652.
- [28] M. Balcan, A. Blum, and Y. Mansour, “Item pricing for revenue maximization,” in *Proceedings 9th ACM Conference on Electronic Commerce*, 2008, pp. 50–59.
- [29] T. Chakraborty, Z. Huang, and S. Khanna, “Dynamic and non-uniform pricing strategies for revenue maximization,” in *Proceedings of the 50th IEEE Symposium on Foundations of Computer Science*, 2009, pp. 495–504.
- [30] S. Alaei, J. D. Hartline, R. Niazadeh, E. Pountourakis, and Y. Yuan, “Optimal auctions vs. anonymous pricing,” in *Proceedings of 56th IEEE Symposium on Foundations of Computer Science*, 2015, pp. 1446–1463.
- [31] V. Cohen-Addad, A. Eden, M. Feldman, and A. Fiat, “The invisible hand of dynamic market pricing,” in *Proceedings of the 17th ACM Conference on Economics and Computation*, 2016, pp. 383–400.
- [32] I. R. Cohen, A. Eden, A. Fiat, and L. Jez, “Pricing online decisions: Beyond auctions,” in *Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms*, 2015, pp. 73–91.
- [33] M. Feldman, A. Fiat, and A. Roytman, “Makespan minimization via posted prices,” in *Working paper*, 2017.
- [34] A. Rubinstein and S. Singla, “Combinatorial prophet inequalities,” in *Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms*, 2017, pp. 1671–1687.
- [35] Y. Cai and M. Zhao, “Simple mechanisms for subadditive buyers via duality,” in *Proceedings of the 49th ACM Symposium on Theory of Computing*, 2017, pp. 170–183.
- [36] J. R. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, and T. Vredeveld, “Posted price mechanisms for a random stream of customers,” in *Proceedings of the 18th ACM Conference on Economics and Computation*, 2017, pp. 169–186.
- [37] M. Abolhassani, S. Ehsani, H. Esfandiari, M. T. Hajiaghayi, R. Kleinberg, and B. Lucier, “Beating $1-1/e$ for ordered prophets,” in *Proceedings of the 49th ACM Symposium on Theory of Computing*, 2017, pp. 61–71.
- [38] C. Daskalakis and V. Syrgkanis, “Learning in auctions: Regret is hard, envy is easy,” in *Proceedings of the 57th IEEE Symposium on Foundations of Computer Science*, 2016, pp. 219–228.
- [39] T. Lykouris, V. Syrgkanis, and É. Tardos, “Learning and efficiency in games with dynamic population,” in *Proceedings of the 27th ACM-SIAM Symposium on Discrete Algorithms*, 2016, pp. 120–129.
- [40] J. Lee, M. Sviridenko, and J. Vondrak, “Submodular maximization over multiple matroids via generalized exchange properties,” *Mathematics of Operations Research*, vol. 35, no. 4, pp. 795–806, 2010.
- [41] O. H. Ibarra and C. E. Kim, “Fast approximation algorithms for the knapsack and sum of subset problems,” *Journal of the ACM*, vol. 22, pp. 463–468, 1975.
- [42] P. Dütting and T. Kesselheim, “Algorithms against anarchy: Understanding non-truthful mechanisms,” in *Proceedings of the 16th ACM Conference on Economics and Computation*, 2015, pp. 239–255.
- [43] M. Feldman, H. Fu, N. Gravin, and B. Lucier, “Simultaneous auctions are (almost) efficient,” in *Proceedings of the 45th ACM Symposium on Theory of Computing*, 2013, pp. 201–210.
- [44] M. Feldman, N. Immorlica, B. Lucier, T. Roughgarden, and V. Syrgkanis, “The price of anarchy in large games,” in *Proceedings of the 48th ACM Symposium on Theory of Computing*, 2016, pp. 963–976.