



Prophet inequalities made easy: stochastic optimization by pricing nonstochastic inputs

LSE Research Online URL for this paper: <http://eprints.lse.ac.uk/104025/>

Version: Accepted Version

Article:

Dütting, Paul, Feldman, Michal, Kesselheim, Thomas and Lucier, Brendan (2020) Prophet inequalities made easy: stochastic optimization by pricing nonstochastic inputs. *SIAM Journal on Computing*, 49 (3). pp. 540-582. ISSN 0097-5397 (In Press)

<https://doi.org/10.1137/20M1323850>

Reuse

Items deposited in LSE Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the LSE Research Online record for the item.

PROPHET INEQUALITIES MADE EASY: STOCHASTIC OPTIMIZATION BY PRICING NONSTOCHASTIC INPUTS*

PAUL DÜTTING[†], MICHAL FELDMAN[‡], THOMAS KESSELHEIM[§],
AND BRENDAN LUCIER[¶]

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), multidimensional 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.

Key words. online algorithms, competitive analysis, optimal stopping, prophet inequalities, posted-price mechanisms, price of anarchy

AMS subject classifications. 68Q25, 68W27, 60G40

DOI. 10.1137/20M1323850

1. 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 willing to attend only 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 inequalities” that compare the performance of an online algorithm to that of an omniscient offline planner. A classic result is that if the

*Received by the editors March 6, 2020; accepted for publication (in revised form) April 1, 2020; published electronically June 2, 2020. An extended abstract appeared in *Proceedings of the 58th IEEE Symposium on Foundations of Computer Science*, Berkeley, CA, 2017.

<https://doi.org/10.1137/20M1323850>

Funding: Part of the work of the first author was done while he was visiting Google Research. The work of the second author was partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement 337122 and by the Israel Science Foundation (grant 317/17). The work of the third author was done while he 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 Simons Institute for the Theory of Computing.

[†]Department of Mathematics, London School of Economics, Houghton Street, WC2A 2AE London, UK (p.d.duetting@lse.ac.uk).

[‡]Blavatnik School of Computer Science, Tel Aviv University, P.O.B. 39040, Ramat Aviv, Tel Aviv, Israel (mfeldman@tau.ac.il).

[§]Institute of Computer Science, University of Bonn, Endenicher Allee 19a, 53115 Bonn, Germany (thomas.kesselheim@uni-bonn.de).

[¶]Microsoft Research New England, Cambridge, MA 02142 (brlucier@microsoft.com).

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 precomputed threshold—yields at least half of the expected maximum value [34, 35, 44]. 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 [13, 33], unit-demand bidders [13, 3], and combinatorial auctions [25]. 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.

1.1. Example: Combinatorial auctions. To introduce our framework we will consider a toy problem of a combinatorial auction with a cardinality constraint, designed to capture settings such as the concert example above. There is a set M of m items (seats) for sale and n buyers. Each buyer i has a valuation function $v_i: 2^M \rightarrow \mathbb{R}_{\geq 0}$ that assigns nonnegative value to every subset of *at most d items*, where $d \geq 1$ is a constant. Valuations are nondecreasing 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 $P \neq NP$ [46]. Our goal will be to match this $O(d)$ approximation as a prophet inequality with static and anonymous posted item prices that can be computed in polynomial time. That is, given distributions over the valuations, compute fixed 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 at those prices, the expected total value is an $O(d)$ approximation to the expected optimum.

Let’s first consider the simpler full information case where all valuations are known in advance. Let us also put computational issues aside for the moment and focus on the *existence* of appropriate item prices; we will revisit the computational version of the problem at the end of section 1.2. This problem is still nontrivial, and in fact there is a lower bound of $\Omega(d)$ on the approximation of any posted item prices. For example, suppose there are d items and two agents. The first agent is unit-demand and has value 1 for any single item; this agent will get the first chance to purchase. 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 will purchase anything. If any item has price 1 or less, 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. The optimal outcome allocates all items to the second agent for a value of d .¹

We wish to show that this bound of $O(d)$ can always be obtained using item prices. 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

¹This example is dependent on tie-breaking in case of indifference, but one can avoid such issues by perturbing the values by an arbitrarily small amount.

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 high enough in the following sense. For each set of items S , the sum of the prices of items in S is at least $1/2d$ of the value of allocations in the optimal allocation \mathbf{x}^* that intersect S , which in turn is an upper bound on the total value that could be lost due to misallocating S . So for any allocation \mathbf{x} , the prices that the players have to pay for that allocation partially offset the value lost due to allocating \mathbf{x} . On the other hand, we will think of prices as low enough if, for any allocation \mathbf{x} , the total price of the leftover items can be upper bounded by a certain fraction of the value of the optimal (re)allocation of the leftover items (called *balanced* below) or all items (called *weakly balanced* below). The prices defined above satisfy the latter: for any allocation \mathbf{x} , the total price of the leftover items is at most $1/2 \cdot \mathbf{v}(\mathbf{x}^*)$.

Let’s first see why these prices lead to the desired approximation guarantee in the example described above. In this example, the welfare-maximizing allocation \mathbf{x}^* allocates all items to the second agent, so $x_1^* = \emptyset$ and $x_2^* = M$. Each item j is therefore assigned a price of $p_j = v_2(x_2^*)/2|x_2^*| = d/(2d) = 1/2$. Under these prices, the first agent will purchase a single item, leading to a welfare of 1 which is indeed an $O(d)$ approximation to the optimal welfare.

To see why these prices yield an $O(d)$ approximation more generally, 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 adding the revenue gives an $O(d)$ approximation.²

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.

1.2. 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 2 and 3), 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 $p_i^{\mathbf{v}}(x_i)$ for every bundle that we can assign to buyer i . For example, the item prices described in section 1.1 define a pricing rule $p_i^{\mathbf{v}}(x_i) = \sum_{j \in x_i} p_j$. Below

²For 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 nonnegative, 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.

we also extend the definition to dynamic prices, i.e., prices that depend on which allocations have already been made.

KEY DEFINITION (SPECIAL CASE) [\$(\alpha, \beta)\$-BALANCED PRICES]. *Let \$\alpha, \beta > 0\$. A pricing rule \$\mathbf{p}^{\mathbf{v}} = (p_1^{\mathbf{v}}, \dots, p_n^{\mathbf{v}})\$ defined by functions \$p_i^{\mathbf{v}} : 2^M \to \mathbb{R}_{\geq 0}\$ is \$(\alpha, \beta)\$-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{v}}(x_i) \ge \frac{1}{\alpha}(\mathbf{v}(\text{OPT}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v} \mid \mathbf{x})))\$,
- (b) \$\sum_i p_i^{\mathbf{v}}(x'_i) \le \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 Weinberg [33]. 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 agent valuations are drawn from a known distribution \$\mathcal{D}\$, which is independent across agents. Suppose that the pricing rule \$\mathbf{p}^{\mathbf{v}}\$ is \$(\alpha, \beta)\$-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 [41, 45] or for establishing algorithmic stability [30]; in that it uses “ghost samples.” It is, however, considerably more involved because of the sequential, online aspect of our problem.

Remark 1.1 (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 \$1/4\alpha\beta\$-approximate prophet inequality. (We actually give a more refined notion, where the prices are upper bounded by \$\beta_1 \mathbf{v}(\text{OPT}(\mathbf{v} \mid \mathbf{x})) + \beta_2 \mathbf{v}(\text{OPT}(\mathbf{v}))\$, and the approximation guarantee is some function of \$\alpha, \beta_1\$, and \$\beta_2\$.) See section 3.2.

Remark 1.2 (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, 25].

Remark 1.3 (static versus dynamic, anonymous versus discriminatory, bundle versus 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, are 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 $p_i(x_i \mid \mathbf{x}_{[i-1]})$ and $p_i(x'_i \mid \mathbf{x}_{[i-1]})$ 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 2 and 3 for details.

Remark 1.4 (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 1.1. 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 1.1, by applying our main theorem to a fractional relaxation of the auction problem. See section 4 for more details.

Remark 1.5 (composition). We also show that balanced prices “compose,” as was shown for mechanism smoothness in [45]. This means that to derive a prophet inequality for a complex setting it often suffices to show balancedness for a simpler problem. See section 5.

1.3. 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 the first example in more detail in section 3.1. The other two examples are covered in Appendices C and E.

Example 1 (classic prophet inequality [34, 35]). 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 [33]). 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 $p_i(x_i \mid \mathbf{y}) = \mathbf{v}(\text{OPT}(\mathbf{v} \mid \mathbf{y})) - \mathbf{v}(\text{OPT}(\mathbf{v} \mid \mathbf{y} \cup x_i))$ 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 [25]). The goal is to assign m goods to n buyers with XOS valuations.³ Let $\mathbf{x}^* = \text{OPT}(\mathbf{v})$ and for every agent $i = 1, \dots, n$, let a_i be the additive supporting function of x_i^* (see footnote 3). Set item prices $p_j = a_i(j)$ for $j \in x_i^*$. This pricing rule is $(1, 1)$ -balanced.

³A 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 , $v(S) = \max_{1 \leq \ell \leq k} a_\ell(S)$. The additive supporting function of a set S is a_ℓ such that $v(S) = a_\ell(S)$. XOS is a generalization of submodular valuations [37].

Example 3 also illustrates the power of our composition results (see Appendix B): the existence of $(1, 1)$ -balanced prices for XOS combinatorial auctions, and hence a 2-approximate prophet inequality, follows directly from the existence of $(1, 1)$ -balanced prices for a single item, despite being significantly more complex. It also yields an $O(\log m)$ -approximate prophet inequality for subadditive valuations by approximating subadditive valuations with XOS valuations [18, 9].

1.4. New and improved prophet inequalities. We also establish several new prophet inequalities using our framework; see Table 1. Our first result is a polytime $(4k - 2)$ -approximate prophet inequality for combinatorial auctions with maximum over positive hypergraphs- k (MPH- k) valuations. The MPH- k hierarchy of valuations [22] is an inclusive hierarchy that captures limited complementarities, where k measures the degree of complementarity. For more details about MPH- k valuations, see Appendix C.

THEOREM 1.6 (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 1.6 improves the polytime result of [25] from $O(k^2)$ to $O(k)$. We note two interesting special cases. First, combinatorial auctions with bundle size d (from section 1.1) belong to MPH- d , so Theorem 1.6 captures the polytime $O(d)$ approximation discussed above. Second, XOS valuations coincide with MPH-1, so Theorem 1.6 improves the previously best-known polytime result of [25] from $2e/(e - 1)$ to 2, matching the existential lower bound. See section 4 and Appendix C.

The second set of new results includes knapsack feasibility constraints and d -sparse packing integer programs (PIPs), for which we obtain a constant- and an $O(d)$ -approximation, respectively. These settings are presented in section 4 and Appendix D, respectively.

TABLE 1

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 2).

Feasibility constraint	Valuation class	Pricing model	Upper bound	Query model
Combinatorial auction	XOS	Static, anonymous item prices	$\frac{2e}{e-1}$ [25] 2 (this work)	XOS, demand
Combinatorial auction	MPH- k	Static, anonymous item prices	$O(k^2)$ [25] $4k - 2$ (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

THEOREM 1.7 (knapsack). *For knapsack constraints, a factor $(5 + \epsilon)$ -approximate posted-price mechanism, with static prices, can be computed in $\text{POLY}(n, 1/\epsilon)$. This improves to a $(3 + \epsilon)$ approximation if no individual demands more than half of the total capacity.*

THEOREM 1.8 (sparse PIPs). *For d -sparse 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 1.7 and 1.8 are the first prophet inequalities for these settings. We note that [27] 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 1.6 (with $k = 1$) is that one can obtain a 2-approximation for a fractional knapsack constraint. Moreover, our prophet inequality is achieved using a static per-unit price and holds even when knapsack weights are private and arbitrarily correlated with buyer values. See section 4 for more details.

Finally, in Appendix E we generalize the matroid prophet inequalities of Kleinberg and Weinberg [33] to settings where players make choices regarding multiple elements of a matroid and have submodular preferences over subsets of elements.

THEOREM 1.9 (multidimensional matroids). *For matroid feasibility constraints and submodular valuations, there is a $(4 + \epsilon)$ -approximate posted-price mechanism, with dynamic prices, that can be computed in $\text{POLY}(n, 1/\epsilon)$ value queries.*

1.5. From price of anarchy to prophet inequalities. The price of anarchy of a class of games is the worst ratio, over all games in the class, of the optimal performance in any outcome and the performance in the worst Nash equilibrium (or another equilibrium concept). A standard technique for obtaining price of anarchy bounds is via the *smoothness* framework [41, 45]. A key step in a smoothness proof is to find an appropriate deviation for each player, such that the utility achieved by the deviation has an appropriate lower bound.

Similar to this approach, in the proof sketch in section 1.1, we derived a lower bound on buyer utility by considering a deviation to a certain purchasing decision. There is a subtle but important difference, however, between our deviation and deviations considered as part of the smoothness framework. 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. Nonetheless, we are able to prove fairly general reductions, whereby the existence of a smooth mechanism of a certain form implies a prophet inequality. We provide such reductions for binary, single parameter settings and for more general settings. Our results for these more general settings go through a natural strengthening of smoothness, which we call outcome smoothness. For more details on the connection between price of anarchy and prophet inequality, see section 6.

Using these results we can, for example, rederive the classic prophet inequality [34, 35] from the smoothness of the first-price single-item auction [45] and the matroid prophet inequality [33] from the smoothness of the pay-your-bid, declared welfare maximizing mechanism for selecting a maximum-weight basis [39]. Our results also allow us to rederive the prophet inequality for XOS combinatorial auctions [25] from smoothness, either via our composition results from the smoothness of the first-price single-item auction [45] or directly from the outcome smoothness of simultaneous first-price auctions [14].

1.6. Further related work. Prophet inequalities and their applicability as posted-price mechanisms were (re-)discovered in theoretical computer science by [29]. Subsequently, threshold-based prophet inequalities and posted-price mechanisms were developed for matroids and matroid intersection [13, 33, 7], polymatroids [21], unit-demand bidders [13, 3], and combinatorial auctions [3, 25].

Not all prophet inequalities in the literature are based on explicit thresholds. Examples include prophet inequalities for the generalized assignment problem [4, 5], for matroids and matroid intersection [28], and for general binary feasibility constraints [42]. On the other hand, many posted-price mechanisms from the literature are constructed either without explicit reference to prophet inequalities or via different techniques. Chawla, Hartline, and Kleinberg [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 [19, 8, 11, 6]. Dynamic posted prices that give optimal welfare for unit-demand buyers were established in [16]. Recently, dynamic posted prices for various online settings have been considered, including k -server on the line and metrical task systems [15] and makespan minimization for scheduling problems [23].

Most recently, and in parallel to this work, combinatorial prophet inequalities were developed in [43] and [10]. The former, among 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, [1] reconsiders the classic prophet inequality setting, but in a large market setting and assuming random or best order.

The notion of smooth games was introduced by Roughgarden [41] 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 [45]. Notions of outcome smoothness were considered in [17, 40].

2. 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 will 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 $p_i(x_i \mid \mathbf{y})$ for the (nonnegative) price assigned to outcome $x_i \in X_i$, offered to agent i , given partial allocation

$\mathbf{y} \in \mathcal{F}$. Define $p_i(x_i) = p_i(x_i \mid \emptyset)$ for convenience. We require that $p_i(x_i \mid \mathbf{y}) = \infty$ for any x_i such that $(x_i, \mathbf{y}_{-i}) \notin \mathcal{F}$. A pricing rule is said to be *monotone nondecreasing* if $p_i(x_i \mid \mathbf{y}) \geq p_i(x_i \mid \mathbf{y}_S)$ 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* [13].

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.

3. 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. It will be useful to define a general notion of feasibility given a partial allocation. 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 $\mathbf{x} \in X$. For example, one might choose $\mathcal{F}_{\mathbf{x}}$ to be a singleton set containing a particular allocation of residual goods not allocated by \mathbf{x} , and in many of our applications we will do just that. But some of our applications require the additional flexibility afforded by this more general definition of exchange compatibility, such as our improved analysis of the knapsack allocation problem in Theorem 4.2.

DEFINITION 3.1. *Let $\alpha > 0, \beta \geq 0$. Given a set of feasible outcomes \mathcal{F} and a valuation profile \mathbf{v} , a pricing rule $\mathbf{p}^{\mathbf{v}}$ 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) $\sum_{i \in N} p_i^{\mathbf{v}}(x_i \mid \mathbf{x}_{[i-1]}) \geq \frac{1}{\alpha} \cdot (\mathbf{v}(\text{ALG}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})))$, and
- (b) for all $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$: $\sum_{i \in N} p_i^{\mathbf{v}}(x'_i \mid \mathbf{x}_{[i-1]}) \leq \beta \cdot \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))$.

The definition provides flexibility in the precise choice of $\mathcal{F}_{\mathbf{x}}$. As the set $\mathcal{F}_{\mathbf{x}}$ 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 3.2. *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 (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 $(\mathbf{p}^{\mathbf{v}})_{\mathbf{v} \in V}$ 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 first obtain an independent sample valuation profile $\mathbf{v}' \sim \mathcal{D}$. 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

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

Summing the previous inequality over all agents we get

$$\begin{aligned} \mathbb{E}_{\mathbf{v}} \left[\sum_{i \in N} u_i(\mathbf{v}) \right] &\geq \mathbb{E}_{\mathbf{v}, \mathbf{v}'} \left[\sum_{i \in N} v'_i \left(x'_i(\mathbf{v}, \mathbf{v}') \right) \right] - \mathbb{E}_{\mathbf{v}, \mathbf{v}'} \left[\sum_{i \in N} \delta \cdot p_i \left(x'_i(\mathbf{v}, \mathbf{v}') \mid \mathbf{x}_{[i-1]}(\mathbf{v}) \right) \right] \\ (3.1) \quad &= \mathbb{E}_{\mathbf{v}, \mathbf{v}'} \left[\mathbf{v}' \left(\text{OPT}(\mathbf{v}', \mathcal{F}_{\mathbf{x}(\mathbf{v})}) \right) \right] - \mathbb{E}_{\mathbf{v}, \mathbf{v}'} \left[\sum_{i \in N} \delta \cdot p_i \left(x'_i(\mathbf{v}, \mathbf{v}') \mid \mathbf{x}_{[i-1]}(\mathbf{v}) \right) \right]. \end{aligned}$$

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

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

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

$$(3.2) \quad \mathbb{E}_{\mathbf{v}, \mathbf{v}'} \left[\sum_{i \in N} \delta \cdot p_i \left(x'_i(\mathbf{v}, \mathbf{v}') \mid \mathbf{x}_{[i-1]}(\mathbf{v}) \right) \right] \leq \delta \beta \cdot \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}} \left[\tilde{\mathbf{v}} \left(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})}) \right) \right].$$

Replacing \mathbf{v}' with $\tilde{\mathbf{v}}$ in inequality (3.1) and combining it with inequality (3.2) we obtain

$$(3.3) \quad \mathbb{E}_{\tilde{\mathbf{v}}} \left[\sum_{i \in N} u_i(\mathbf{v}) \right] \geq (1 - \delta\beta) \cdot \mathbb{E}_{\tilde{\mathbf{v}}, \mathbf{v}} [\tilde{\mathbf{v}}(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})}))].$$

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

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

Taking expectation over \mathbf{v} this shows

$$(3.4) \quad \mathbb{E}_{\tilde{\mathbf{v}}} \left[\sum_{i \in N} \delta \cdot p_i(x_i(\mathbf{v}) \mid \mathbf{x}_{[i-1]}(\mathbf{v})) \right] \geq \frac{\delta}{\alpha} \cdot \mathbb{E}_{\tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{ALG}(\tilde{\mathbf{v}}))] - \frac{\delta}{\alpha} \cdot \mathbb{E}_{\tilde{\mathbf{v}}, \mathbf{v}} [\tilde{\mathbf{v}}(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})}))].$$

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 (3.3) and (3.4), this gives

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

3.1. Classic prophet inequality via our framework. We next show how to rederive the classic prophet inequality via our framework. Specifically, we show the existence of a $(1, 1)$ -balanced pricing rule, which implies the factor 2 approximation by Theorem 3.2.

In the classic setting we have $X_i = \{0, 1\}$ for all i and $\mathcal{F} = \{\mathbf{x} \mid \sum_i x_i \leq 1\}$. We set $p_i(1 \mid \mathbf{x}) = \max_{\ell} v_{\ell}$ if \mathbf{x} does not allocate the item, ∞ otherwise; $p_i(0 \mid \mathbf{x}) = 0$ for all \mathbf{x} . This corresponds to a fixed posted price of $\max_{\ell} v_{\ell}$ on the item.

CLAIM 3.3. *These prices are $(1, 1)$ -balanced with respect to OPT and $\mathcal{F}_{\mathbf{x}}$ defined by $\mathcal{F}_{\mathbf{x}} = \mathcal{F}$ if \mathbf{x} does not allocate the item and $\mathcal{F}_{\mathbf{x}} = \emptyset$ otherwise.*

Proof. Let \mathbf{x} be an arbitrary allocation profile. If \mathbf{x} allocates the item, then $\mathcal{F}_{\mathbf{x}} = \emptyset$ and thus condition (b) is trivially fulfilled. For condition (a), we observe that $\mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) = 0$ and that $\mathbf{v}(\text{OPT}(\mathbf{v})) = \max_{\ell} v_{\ell} = \sum_i p_i(x_i \mid \mathbf{x}_{[i-1]})$, because

exactly one buyer pays $\max_{\ell} v_{\ell}$. If \mathbf{x} does not allocate the item, then $\mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) = \mathbf{v}(\text{OPT}(\mathbf{v}))$, making condition (a) trivial. For condition (b), we use that in \mathbf{x}' at most one buyer is allocated the item. Therefore, $\sum_i p_i(x'_i \mid \mathbf{x}_{[i-1]}) \leq \max_{\ell} v_{\ell} = \mathbf{v}(\text{OPT}(\mathbf{v})) = \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))$. \square

3.2. Weak balancedness. 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 4.

DEFINITION 3.4. *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 $\mathbf{p}^{\mathbf{v}}$ 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) $\sum_{i \in N} p_i^{\mathbf{y}}(x_i \mid \mathbf{x}_{[i-1]}) \geq \frac{1}{\alpha} \cdot (\mathbf{v}(\text{ALG}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})))$, and
- (b) for all $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$: $\sum_{i \in N} p_i^{\mathbf{y}}(x'_i \mid \mathbf{x}_{[i-1]}) \leq \beta_1 \cdot \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) + \beta_2 \cdot \mathbf{v}(\text{ALG}(\mathbf{v}))$.

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 Appendix A.

THEOREM 3.5. *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.*

4. 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 polytime $O(d)$ -approximation for combinatorial auctions with bundle size at most d , and the new polytime constant-approximation for knapsack problems. Additional results include combinatorial auctions with MPH- k valuations (see Appendix C), d -sparse PIPs (see Appendix D), and multidimensional matroids (where the result follows from the Rota exchange theorem [36, Lemma 2.7] and our composition results; see Appendix E).

Combinatorial auctions with bounded bundle size. An existential $O(d)$ -approximate price-based prophet inequality is presented in section 1.1. Combined with the $O(d)$ -approximation greedy algorithm for this setting, it gives a polytime $O(d^2)$ -approximate price-based prophet inequality (as shown in [25]). 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 1.6, which is proved in Appendix C.

THEOREM 4.1. *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. We will define $\mathcal{F}_{\mathbf{x}}$ to be the exchange-compatible

set consisting of all 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 . 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

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

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 3.5 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 nonintegral allocations from being purchased. \square

The more general Theorem 1.6 also improves the best-known polytime prophet inequality for XOS valuations from $2e/(e - 1)$ to 2 (which is tight [25]) and for MPH- k valuations it improves the best-known polytime 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 1.7.

THEOREM 4.2. *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 polytime posted-price mechanism with a single static anonymous per-unit price.*

Proof. The polytime claim follows from Theorem 3.2 with ALG set to the classic FPTAS for knapsack [32], 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

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

We can remove the restriction that $s_i \leq 1/2$ as follows, completing the proof of Theorem 1.7. 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.⁴ One can therefore obtain a $(5 + \epsilon)$ -approximate price-based prophet

⁴The 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.

inequality by estimating the expected welfare of each pricing scheme (via sampling) and selecting the better of the two. In Appendix D we show how to generalize the result for the knapsack problem to d -sparse PIPs.

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 1.6 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 polytime 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 1.4, this improves the previously best-known prophet inequality of ≈ 11.657 due to [27].

5. Composition results. In this section we show that balanced prices are *composable*, in the sense that balanced prices for separate markets remain balanced when the markets are combined. We consider two forms of composition. The first is a composition of preferences: it shows how to extend from a class of valuations V to any maximum over valuations from V . The second shows how to compose allocations across different markets, where agents have additive preferences across markets. Together these two composition results capture XOS composition, in the sense of [45]. Our theorems apply to balancedness with respect to OPT, and they extend to general approximation algorithms ALG under mild conditions. Missing proofs for this section appear in Appendix B.

Closure under maximum. Given an arbitrary valuation space V_i for player i , we consider its extension V_i^{\max} , which contains all functions $v_i^{\max}: X_i \rightarrow \mathbb{R}$ for which there is a finite set $\{v_i^1, \dots, v_i^m\} \in V_i$ such that $v_i^{\max}(x_i) = \max_{\ell} v_i^{\ell}(x_i)$ for all $x_i \in X_i$. We say that a valuation profile $\tilde{\mathbf{v}} \in V$ is a *supporting* valuation profile for allocation \mathbf{x} and valuation profile $\mathbf{v} \in V^{\max}$ if $\tilde{v}_i \leq v_i$ and $\tilde{v}_i(x_i) = v_i(x_i)$ for all i .

DEFINITION 5.1. *Allocation rule ALG is consistent if for every $\mathbf{v} \in V^{\max}$ and corresponding supporting valuation profile $\tilde{\mathbf{v}} \in V$ for $\text{ALG}(\mathbf{v})$, we have $\tilde{\mathbf{v}}(\text{ALG}(\tilde{\mathbf{v}})) \geq \tilde{\mathbf{v}}(\text{ALG}(\mathbf{v}))$.*

LEMMA 5.2. *The optimal allocation rule OPT is consistent.*

Proof. Since $\text{OPT}(\mathbf{v})$ is a feasible outcome under $\tilde{\mathbf{v}}$, it holds that $\tilde{\mathbf{v}}(\text{OPT}(\tilde{\mathbf{v}})) \geq \tilde{\mathbf{v}}(\text{OPT}(\mathbf{v}))$ by optimality. \square

THEOREM 5.3 (closure under maximum). *Suppose that for each $\mathbf{v} \in V$ there exists a pricing rule $\mathbf{p}^{\mathbf{v}}$ that is (α, β) -balanced with respect to \mathbf{v} , consistent allocation rule ALG, and the exchange-compatible family of sets $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$. For each $\mathbf{v} \in V^{\max}$ let $\tilde{\mathbf{v}} \in V$ be a supporting valuation profile for $\text{ALG}(\mathbf{v})$. Then the pricing rule $\mathbf{p}^{\tilde{\mathbf{v}}}$ is (α, β) -balanced with respect to \mathbf{v} , ALG, and $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$.*

Closure under addition. Suppose there are m separate allocation problems, each with feasibility constraint \mathcal{F}^{ℓ} over allocation space $X^{\ell} = X_1^{\ell} \times \dots \times X_n^{\ell}$. The joint problem is then defined over the product allocation space $X = X^1 \times \dots \times X^m$, with feasibility constraint $\mathcal{F} = \mathcal{F}^1 \times \dots \times \mathcal{F}^m$. We say that a valuation $v_i: X \rightarrow \mathbb{R}$ is *additive* if it is defined by a (not necessarily additive) valuation function v_i^{ℓ} for each outcome $x_i^{\ell} \in X_i^{\ell}$, and for an outcome $x_i = (x_i^1, \dots, x_i^m) \in X_i$, the value of x_i is given by $v_i(x_i) = \sum_{\ell=1}^m v_i^{\ell}(x_i^{\ell})$.

THEOREM 5.4 (closure under addition). *Suppose that \mathbf{v} is additive over a set of allocation problems $\mathcal{F}^1, \dots, \mathcal{F}^m$, and for each individual allocation problem there*

exists a pricing rule \mathbf{p}^{ℓ} that is (α, β) -balanced with respect to \mathbf{v}^{ℓ} , allocation rule ALG^{ℓ} on \mathcal{F}^{ℓ} , and the exchange-compatible family of sets $(\mathcal{F}_{\mathbf{x}^{\ell}}^{\ell})_{\mathbf{x}^{\ell} \in X^{\ell}}$. Then the pricing rule $\mathbf{p} = \sum_{\ell=1}^m \mathbf{p}^{\ell}$ is (α, β) -balanced with respect to \mathbf{v} and ALG on \mathcal{F} , where $\text{ALG}(\mathbf{v}) = (\text{ALG}^1(\mathbf{v}^1), \dots, \text{ALG}^m(\mathbf{v}^m))$.

Note that XOS valuations are constructed via maximum and addition compositions. Therefore, using these two composition results, the existence of a factor 2 approximation for combinatorial auctions with XOS valuations via posted prices follows directly from the existence of balanced prices for the classic (single-item) prophet inequality setting discussed in section 3. These composition results are also used to establish prophet inequalities for multidimensional matroid settings, as discussed in section 1 and in Appendix E.

We also note that both Theorem 5.3 and Theorem 5.4 can be generalized so that they apply to weakly balanced pricing rules, under a slightly more stringent notion of consistency. See Appendix B for further details.

6. From price of anarchy to prophet inequalities. In this section we explore the connection between balanced prices and mechanism smoothness. In general, smoothness does not suffice to conclude the existence of a posted-price mechanism with comparable welfare guarantee (see Appendix F). However, we will argue that “typical” smoothness proofs do have implications for prophet inequalities and present 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 6.1 (Syrgkanis and Tardos [45]). *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*

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

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 $\max\{\mu, 1\}/\lambda$ [45].

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*. For example, if our setting is a combinatorial auction and \mathbf{x} is some allocation of items to buyers, then the canonical exchange-feasible set $\mathcal{F}_{\mathbf{x}}$ is itself a combinatorial auction, containing the items that are unallocated in \mathbf{x} and the buyers who receive no items in \mathbf{x} .

6.1. 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 consider a binary, single-parameter setting. That is, players can either “win” or “lose” and have a value $v_i \in \mathbb{R}_{\geq 0}$ for winning. Feasible solutions $\mathbf{x} \in \mathcal{F} \subseteq \{0, 1\}^n$ are subsets of players that can win simultaneously. For ease of notation we identify each vector $\mathbf{x} \in \mathcal{F}$ with the subset of players $i \in N$ for which $x_i = 1$. This lets us write $i \in \mathbf{x}$ if $x_i = 1$ and $i \notin \mathbf{x}$ otherwise. For a player $i \in N$, valuations \mathbf{v}_{-i} for the other players, and an allocation \mathbf{y} , we define the critical price $\tau_i(\mathbf{v}_{-i} \mid \mathbf{y}) := \inf\{v'_i \mid i \in \text{OPT}(v'_i, \mathbf{v}_{-i}) \mid \mathbf{y}\}$ as the infimum of values v_i such that the mechanism allocates $x_i = 1$ to agent i on input (v_i, \mathbf{v}_{-i}) , in the problem subinstance induced by \mathbf{y} .

We will show that if the welfare-maximizing mechanism is (λ, μ) -smooth and its critical prices $\tau_i(\cdot \mid \mathbf{y})$ are nondecreasing in \mathbf{y} , 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 of the prophet inequality that uses these balanced prices as implied by Theorem 3.2 is within a constant factor of the price of anarchy of the mechanism implied by smoothness.

THEOREM 6.2. *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 nondecreasing 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

$$(6.1) \quad p_i(x_i \mid \mathbf{y}) \geq \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{(\emptyset, \mathbf{y}_{-i})})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{(x_i, \mathbf{y}_{-i})})).$$

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 both sides of the inequality are equal to v_i . To see that this is the case for the right-hand side notice that i is allocated in $\text{OPT}(\mathbf{v}, \mathcal{F}_{(\emptyset, \mathbf{y}_{-i})})$ because of the condition on the critical threshold, but then the remaining outcomes are the same as if we contracted on (x_i, \mathbf{y}_{-i}) , since an optimal allocation is also optimal on any contraction. If $x_i = 1$ and $v_i < \tau_i(\mathbf{v}_{-i} \mid \mathbf{y})$, then agent i is not allocated in $\text{OPT}(\mathbf{v}, \mathcal{F}_{(\emptyset, \mathbf{y}_{-i})})$, 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

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

as required, where the inequality follows from (6.1), and the equality follows by a telescoping sum. For condition (b), we get

$$(6.2) \quad \sum_{i \in \mathbf{x}'} \tau_i(\mathbf{v}_{-i} \mid \mathbf{x}_{[i-1]}) \leq \sum_{i \in \mathbf{x}'} \tau_i(\mathbf{v}_{-i} \mid \mathbf{x}) \leq \frac{\mu+1}{\lambda} \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})),$$

where the first inequality follows by the monotonicity of critical prices, and the second inequality follows by a known implication of smoothness [20] (see Appendix G.1).

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 (6.2), and the last inequality follows by $\mathbf{v}(\mathbf{x}') \leq \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))$, since $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$. \square

6.2. General reductions for binary, single-parameter settings. More generally, we can show the following two reductions from smoothness to prophet inequalities for binary, single parameter problems. The first is for mechanisms that use a greedy allocation rule, and the second is for mechanisms that maximize declared welfare. The key to both these results, which we prove in Appendix G, is a novel, purely combinatorial implication of smoothness for the greedy allocation rule proved in Lemma G.6.

THEOREM 6.3. *Consider a subinstance-closed collection of binary, single-parameter problems. 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 6.4. *Consider a subinstance-closed collection of binary, single-parameter problems. 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 6.2 applied to matroids (using known smoothness results for pay-your-bid greedy mechanisms over matroids [39]) implies the existence of (1, 3)-balanced prices and hence a 4-approximate prophet inequality. A strengthening of Theorem 6.4 for monotonically increasing critical prices (discussed in Remark G.13) leads to an improved factor of 2, matching the prophet inequality for matroids shown in Kleinberg and Weinberg [33]. This also captures the classic single-item prophet inequality, as a special case, and the prophet inequality for XOS combinatorial auctions via our composition results.

6.3. Beyond binary, single-parameter settings. We conclude with an implication from smoothness to prices that applies in more general settings. It is based on the observation that many smoothness proofs have a built-in charging scheme, where each agent considers a bid b'_i aimed at getting some target outcome x_i^* , which is attainable as long as its critical price is not too high. We capture proofs that proceed in this manner through the following notion of *outcome smoothness*. Similar but different notions were considered in [17, 40].

To state the definition we need the following notion of *critical prices*: Consider a mechanism $\mathcal{M} = (f, P)$. For player $i \in N$ and allocation $x \in \mathcal{F}$, the critical price $\tau_i(x_i, \mathbf{b}_{-i})$ for bundle x_i given bids \mathbf{b}_{-i} of the other players is $\tau_i(x_i, \mathbf{b}_{-i}) = \inf_{b'_i: f_i(b'_i, \mathbf{b}_{-i}) \succeq x'_i} P_i(b'_i, \mathbf{b}_{-i})$.

DEFINITION 6.5. *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 $\text{ALG}(\mathbf{v}) \in \mathcal{F}$ such that for all bid*

profiles $\mathbf{b} \in B$,

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

So, outcome smoothness requires there to be an *outcome* that could be targeted by a bidder. In the notion of smoothness defined in Definition 6.1, in contrast, there has to be a deviation *bid*. Many smoothness proofs in the sense of Definition 6.1 actually establish the stronger notion of outcome smoothness.

The intuitive connection between outcome smoothness and pricing is that, for an appropriate choice of bids, one can use the critical prices with respect to these bids as posted item prices and use outcome smoothness to establish that these prices are balanced.

We illustrate this connection by considering any mechanism $\mathcal{M} = (f, p)$ for XOS combinatorial auctions with item bidding whose allocation rule f maximizes declared welfare and whose payment rule P satisfies $p_i(f_i(b_i, \mathbf{b}_{-i})) \leq b_i(f_i(b_i, \mathbf{b}_{-i}))$ for all i and $\mathbf{b} = (b_i, \mathbf{b}_{-i})$. We show that if \mathcal{M} is (λ, μ) -outcome smooth with respect to some allocation rule ALG , then (i) ALG is a λ -approximation to OPT and (ii) there exist balanced prices that are $(1, \mu)$ -balanced with respect to ALG .

Our construction (in Appendix H) requires that the mechanism \mathcal{M} remains (λ, μ) -outcome smooth for any number of bidders n and any number of items m , and that ALG satisfies some mild technical conditions, which are satisfied when $\text{ALG} = \text{OPT}$. The first condition is that ALG is *scale invariant*, i.e., for all $\mathbf{v} \in V$ and every $\epsilon > 0$, $\text{ALG}(\mathbf{v}) = \text{ALG}(\epsilon \cdot \mathbf{v})$. The second is that ALG is *onto*, i.e., for every $\mathbf{x} \in \mathcal{F}$ there is a $\mathbf{v} \in V$ such that $\text{ALG}(\mathbf{v}) = \mathbf{x}$.

THEOREM 6.6. *Suppose $\mathcal{M} = (f, P)$ is a mechanism for XOS combinatorial auctions with item bidding that maximizes declared welfare. Suppose that \mathcal{M} is (λ, μ) -outcome smooth with respect to ALG for every number of bidders n and every number of items m , where ALG is scale invariant and onto. Then (i) ALG is a λ -approximation to OPT and (ii) there exists static item prices that are $(1, \mu)$ -balanced with respect to ALG .*

We note that simultaneous first-price auctions for XOS valuations are $(1, 1)$ -outcome smooth with respect to OPT . Thus, Theorem 6.6 provides an alternative way of proving the existence of a 2-approximate prophet inequality for XOS combinatorial auctions.

7. Conclusions and open problems. We introduce a general framework for establishing prophet inequalities and posted price mechanisms for multidimensional 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 [25], 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 [24], 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.

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. Finally, recent work has shown that smoothness guarantees often improve as markets grow large [26]; is there a corresponding result for balancedness?

Appendix A. Proof of Theorem 3.5. Our proof will follow the same steps as the one for Theorem 3.2. As in that proof, denote the exchange-compatible family of sets with respect to which the collection of pricing rules $(\mathbf{p}^{\mathbf{v}})_{\mathbf{v} \in V}$ is balanced by $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$, and let $\mathbf{x}(\mathbf{v})$ be the allocation returned by the posted-price mechanism on input valuation profile \mathbf{v} . Let $\mathbf{x}'(\mathbf{v}, \mathbf{v}') = \text{OPT}(\mathbf{v}', \mathcal{F}_{\mathbf{x}(\mathbf{v})})$ be the allocation that maximizes welfare with respect to valuation profile \mathbf{v}' over feasibility constraint $\mathcal{F}_{\mathbf{x}(\mathbf{v})}$.

Utility bound. Again sample valuations $\mathbf{v}' \sim \mathcal{D}$. By the same reasoning as in the proof of Theorem 3.2, we obtain

$$(A.1) \quad \mathbb{E}_{\mathbf{v}} \left[\sum_{i \in N} u_i(\mathbf{v}) \right] = \mathbb{E}_{\mathbf{v}, \mathbf{v}'} [\mathbf{v}'(\mathbf{x}'(\mathbf{v}, \mathbf{v}'))] - \mathbb{E}_{\mathbf{v}, \mathbf{v}'} \left[\sum_{i \in N} \delta \cdot p_i \left(x'_i(\mathbf{v}, \mathbf{v}') \mid \mathbf{x}_{[i-1]}(\mathbf{v}) \right) \right].$$

We upper-bound the last term in the previous inequality by using property (b). This gives pointwise for any \mathbf{v} and \mathbf{v}'

$$\sum_{i \in N} \delta \cdot p_i \left(x'_i(\mathbf{v}, \mathbf{v}') \mid \mathbf{x}_{[i-1]}(\mathbf{v}) \right) \leq \delta \beta_1 \cdot \mathbb{E}_{\tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})}))] + \delta \beta_2 \cdot \mathbb{E}_{\tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{ALG}(\tilde{\mathbf{v}}))],$$

and therefore also

$$(A.2) \quad \mathbb{E}_{\mathbf{v}, \mathbf{v}'} \left[\sum_{i \in N} \delta \cdot p_i \left(x'_i(\mathbf{v}, \mathbf{v}') \mid \mathbf{x}_{[i-1]}(\mathbf{v}) \right) \right] \leq \delta \beta_1 \cdot \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})}))] + \delta \beta_2 \cdot \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{ALG}(\tilde{\mathbf{v}}))].$$

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

$$(A.3) \quad \mathbb{E}_{\mathbf{v}} \left[\sum_{i \in N} u_i(\mathbf{v}) \right] \geq (1 - \delta \beta_1) \cdot \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})}))] - \delta \beta_2 \cdot \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{ALG}(\tilde{\mathbf{v}}))].$$

Revenue bound. Again, applying property (a) we obtain

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

Combination. To combine our bounds, we distinguish whether $\beta_2 \geq \frac{1}{2\alpha}$.

Case 1. $\beta_2 \geq \frac{1}{2\alpha}$. We use that pointwise for every $i \in N$ we have $u_i(\mathbf{v}) \geq 0$. Therefore, $u_i(\mathbf{v}) \geq \rho u_i(\mathbf{v})$ for all $0 \leq \rho \leq 1$. Using $\delta = \frac{1}{\beta_1 + 2\beta_2}$, $\rho = \frac{1}{2\alpha\beta_2}$, and inequalities (A.3) and (A.4), we get

$$\begin{aligned} \mathbb{E}_{\mathbf{v}} \left[\sum_{i \in N} v_i(x_i(\mathbf{v})) \right] &\geq \rho \cdot \mathbb{E}_{\mathbf{v}} \left[\sum_{i \in N} u_i(\mathbf{v}) \right] + \mathbb{E}_{\mathbf{v}} \left[\sum_{i \in N} \delta \cdot p_i(x_i(\mathbf{v}) \mid \mathbf{x}_{[i-1]}(\mathbf{v})) \right] \\ &\geq \rho(1 - \delta\beta_1) \cdot \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})}))] - \rho\delta\beta_2 \cdot \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{ALG}(\tilde{\mathbf{v}}))] \\ &\quad + \frac{\delta}{\alpha} \cdot \mathbb{E}_{\tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{ALG}(\tilde{\mathbf{v}}))] - \frac{\delta}{\alpha} \cdot \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})}))] \\ &= \frac{1}{\alpha(2\beta_1 + 4\beta_2)} \cdot \mathbb{E}_{\tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{ALG}(\tilde{\mathbf{v}}))]. \end{aligned}$$

Case 2. $\beta_2 < \frac{1}{2\alpha}$. Now, we use $\delta = \frac{1}{\beta_1 + 1/\alpha}$. Inequalities (A.3) and (A.4) yield

$$\begin{aligned} \mathbb{E}_{\mathbf{v}} \left[\sum_{i \in N} v_i(x_i(\mathbf{v})) \right] &\geq \mathbb{E}_{\mathbf{v}} \left[\sum_{i \in N} u_i(\mathbf{v}) \right] + \mathbb{E}_{\mathbf{v}} \left[\sum_{i \in N} \delta \cdot p_i(x_i(\mathbf{v}) \mid \mathbf{x}_{[i-1]}(\mathbf{v})) \right] \\ &\geq (1 - \delta\beta_1) \cdot \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})}))] - \delta\beta_2 \cdot \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{ALG}(\tilde{\mathbf{v}}))] \\ &\quad + \frac{\delta}{\alpha} \cdot \mathbb{E}_{\tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{ALG}(\tilde{\mathbf{v}}))] - \frac{\delta}{\alpha} \cdot \mathbb{E}_{\mathbf{v}, \tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}(\mathbf{v})}))] \\ &= \frac{1 - \alpha\beta_2}{1 + \alpha\beta_1} \cdot \mathbb{E}_{\tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{ALG}(\tilde{\mathbf{v}}))] \\ &\geq \frac{1}{\alpha(2\beta_1 + 4\beta_2)} \mathbb{E}_{\tilde{\mathbf{v}}} [\tilde{\mathbf{v}}(\text{ALG}(\tilde{\mathbf{v}}))], \end{aligned}$$

where the last step uses that $\beta_1 + \beta_2 \geq \frac{1}{\alpha}$ and $\beta_2 < \frac{1}{2\alpha}$.

Appendix B. Missing proofs and extensions from section 5. In this appendix we provide proofs for Theorems 5.3 and 5.4, which we have omitted from section 5, and we show how to extend these results to weakly balanced pricing rules.

Proof of Theorem 5.3. We first establish property (a). To this end, we first use property (a) of the original pricing rules, then the consistency of ALG, and finally the definition of a supporting valuation profile to conclude that

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

Next we establish property (b). By property (b) of the original pricing rule and the definition of a supporting valuation profile,

$$\begin{aligned} \sum_{i \in N} p_i^{\tilde{\mathbf{v}}}(x'_i \mid \mathbf{x}_{[i-1]}) &\leq \beta \cdot \tilde{\mathbf{v}}(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}})) \\ &\leq \beta \cdot \mathbf{v}(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}})) \\ &\leq \beta \cdot \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})). \end{aligned} \quad \square$$

Proof of Theorem 5.4. Set $\mathcal{F}_{\mathbf{x}} = (\prod_{i=1}^m \mathcal{F}_{\mathbf{x}^\ell})_{\mathbf{x} \in X}$. We first verify condition (a). We use the definition of the joint pricing rule and condition (a) of the component pricing rules together with additivity across subproblems to conclude that

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

Next we verify condition (b). We use the definition of the joint pricing rule and condition (b) of the component pricing rules together with additivity across subproblems to conclude that

$$\begin{aligned} \sum_{i \in N} p_i(x'_i \mid \mathbf{x}_{[i-1]}) &= \sum_{\ell=1}^m \sum_{i \in N} p_i^{\mathbf{v}^\ell}((x'_i)^\ell \mid \mathbf{x}_{[i-1]}^\ell) \\ &\leq \sum_{\ell=1}^m \beta \cdot \mathbf{v}^\ell(\text{OPT}(\mathbf{v}^\ell, \mathcal{F}_{\mathbf{x}^\ell})) \\ &= \beta \cdot \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})). \quad \square \end{aligned}$$

We next show how to extend these results so they apply to weakly balanced pricing rules. Recall that for Theorem 5.3, we imposed a requirement that the allocation rule ALG be consistent. For our extension to weakly balanced pricing rules, we will require a stronger notion of consistency.

DEFINITION B.1. *Allocation rule ALG is strongly consistent if for every $\mathbf{v} \in V^{\max}$ and corresponding supporting valuation profile $\tilde{\mathbf{v}} \in V$ for $\text{ALG}(\mathbf{v})$, we have $\tilde{\mathbf{v}}(\text{ALG}(\tilde{\mathbf{v}})) \geq \tilde{\mathbf{v}}(\text{ALG}(\mathbf{v}))$ and $\mathbf{v}(\text{ALG}(\mathbf{v})) \geq \mathbf{v}(\text{ALG}(\tilde{\mathbf{v}}))$.*

We note that Lemma 5.2 from section 5 directly extends to show that the optimal allocation rule OPT is strongly consistent.

We are now ready to state and prove our composition results for weakly balanced pricing rules. All expressions in the theorem statements are as in Theorems 5.3 and 5.4.

THEOREM B.2 (closure under maximum). *Suppose that for each $\mathbf{v} \in V$ there exists a pricing rule $\mathbf{p}^{\mathbf{v}}$ that is weakly $(\alpha, \beta_1, \beta_2)$ -balanced with respect to \mathbf{v} , a strongly consistent allocation rule ALG , and the exchange-compatible family of sets $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$. For each $\mathbf{v} \in V^{\max}$ let $\tilde{\mathbf{v}} \in V$ be a supporting valuation profile for $\text{ALG}(\mathbf{v})$. Then the pricing rule $\mathbf{p}^{\tilde{\mathbf{v}}}$ is weakly $(\alpha, \beta_1, \beta_2)$ -balanced with respect to \mathbf{v} , ALG , and $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$.*

Proof. Property (a) of weak $(\alpha, \beta_1, \beta_2)$ -balancedness is identical to that of (α, β) -balancedness and follows precisely as in Theorem 5.3. To establish property (b) of weak $(\alpha, \beta_1, \beta_2)$ -balancedness, we apply property (b) to the original pricing rule, then use the definition of a supporting valuation profile followed by the strong consistency of both ALG and OPT to conclude that

$$\begin{aligned} \sum_{i \in N} p_i^{\tilde{\mathbf{v}}}(x'_i \mid \mathbf{x}_{[i-1]}) &\leq \beta_1 \cdot \tilde{\mathbf{v}}(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}})) + \beta_2 \cdot \tilde{\mathbf{v}}(\text{ALG}(\tilde{\mathbf{v}})) \\ &\leq \beta_1 \cdot \mathbf{v}(\text{OPT}(\tilde{\mathbf{v}}, \mathcal{F}_{\mathbf{x}})) + \beta_2 \cdot \mathbf{v}(\text{ALG}(\tilde{\mathbf{v}})) \\ &\leq \beta_1 \cdot \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) + \beta_2 \cdot \mathbf{v}(\text{ALG}(\mathbf{v})). \quad \square \end{aligned}$$

THEOREM B.3 (closure under addition). *Suppose that \mathbf{v} is additive over a set of allocation problems $\mathcal{F}^1, \dots, \mathcal{F}^m$, and for each individual allocation problem there exists a pricing rule $\mathbf{p}^{\mathbf{v}^\ell}$ that is weakly $(\alpha, \beta_1, \beta_2)$ -balanced with respect to \mathbf{v}^ℓ , an allocation rule ALG^ℓ on \mathcal{F}^ℓ , and the exchange-compatible family of sets $(\mathcal{F}_{\mathbf{x}^\ell}^\ell)_{\mathbf{x}^\ell \in X^\ell}$. Then the pricing rule $\mathbf{p} = \sum_{\ell=1}^m \mathbf{p}^{\mathbf{v}^\ell}$ is weakly $(\alpha, \beta_1, \beta_2)$ -balanced with respect to \mathbf{v} and ALG on \mathcal{F} , where $\text{ALG}(\mathbf{v}) = (\text{ALG}^1(\mathbf{v}^1), \dots, \text{ALG}^m(\mathbf{v}^m))$.*

Proof. Property (a) of weak $(\alpha, \beta_1, \beta_2)$ -balancedness follows precisely as in Theorem 5.4. To establish property (b) of weak $(\alpha, \beta_1, \beta_2)$ -balancedness, we apply property (b) to the original pricing rule, then use additivity across the subproblems to conclude that

$$\begin{aligned} \sum_{i \in N} p_i(x'_i \mid \mathbf{x}_{[i-1]}) &= \sum_{\ell=1}^m \sum_{i \in N} p_i^{\mathbf{v}^\ell} \left((x'_i)^\ell \mid \mathbf{x}_{[i-1]}^\ell \right) \\ &\leq \sum_{\ell=1}^m (\beta_1 \cdot \mathbf{v}^\ell (\text{OPT}(\mathbf{v}^\ell, \mathcal{F}_{\mathbf{x}^\ell}^\ell)) + \beta_2 \cdot \mathbf{v}^\ell (\text{ALG}(\mathbf{v}^\ell))) \\ &= \beta_1 \cdot \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) + \beta_2 \cdot \mathbf{v}(\text{ALG}(\mathbf{v})). \quad \square \end{aligned}$$

Appendix C. Proof of Theorem 1.6. In this appendix we establish our new prophet inequality results for combinatorial auctions with MPH- k valuations. In particular, we establish the existence of $(1, 1)$ -balanced prices for XOS combinatorial auctions.

Combinatorial auctions with MPH- k valuations. The MPH hierarchy of valuations [22] is an inclusive hierarchy (i.e., it is expressive enough to include all valuations), which subsumes many interesting classes of valuations as special cases.

To formalize this valuation class, we first need a few preliminaries. A hypergraph representation w of valuation function $v: 2^M \rightarrow \mathbb{R}_{\geq 0}$ is a set function that satisfies $v(S) = \sum_{T \subseteq S} w(T)$. Any valuation function v admits a unique hypergraph representation and vice versa. A set S such that $w(S) \neq 0$ is said to be a *hyperedge* of w . Pictorially, the hypergraph representation can be thought as a weighted hypergraph, where every vertex is associated with an item in M , and the weight of each hyperedge $e \subseteq M$ is $w(e)$. Then the value of the function for any set $S \subseteq M$ is the total value of all hyperedges that are contained in S . The *rank* of a hypergraph representation w is the cardinality k of the largest hyperedge. The rank of v is the rank of its corresponding w and we refer to a valuation function v with rank k as a *hypergraph- k valuation*. If the hypergraph representation of v is nonnegative, i.e., for any $S \subseteq M$, $w(S) \geq 0$, then we refer to function v as a *positive hypergraph- k function (PH- k)* [2].

DEFINITION C.1 (maximum over positive hypergraph- k class [22]). *A monotone valuation function $v: 2^M \rightarrow \mathbb{R}_{\geq 0}$ is MPH- k if it can be expressed as a maximum over a set of PH- k functions. That is, there exist PH- k functions $\{v_\ell\}_{\ell \in \mathcal{L}}$ such that for every set $S \subseteq M$,*

$$(C.1) \quad v(S) = \max_{\ell \in \mathcal{L}} v_\ell(S),$$

where \mathcal{L} is an arbitrary index set.

An important special case of MPH- k is XOS valuations, which are defined as the maximum over additive functions and which therefore coincide with MPH-1.

Existential $O(k)$ result. Given an arbitrary allocation algorithm ALG and valuation profile \mathbf{v} , write $\tilde{\mathbf{v}}$ for the supporting hypergraph- k valuations for allocation $\text{ALG}(\mathbf{v})$. Write w_i for the hypergraph representation of \tilde{v}_i , so that $\tilde{v}_i(\text{ALG}(\mathbf{v})) = \sum_{T \subseteq \text{ALG}_i(\mathbf{v})} w_i(T)$.

Then, for each item j , we define

$$(C.2) \quad p^{\mathbf{v}}(\{j\}) = \sum_{\substack{T \ni j \\ T \subseteq \text{ALG}_i(\mathbf{v})}} w_i(T).$$

That is, item prices are determined by adding up the weights of each supporting hyperedge an item is contained in. These prices then extend linearly: for each set of goods x , set $p^{\mathbf{v}}(x) = \sum_{j \in x} p^{\mathbf{v}}(\{j\})$. Finally, for each agent i and allocation x_i , we will have $p_i^{\mathbf{v}}(x_i \mid \mathbf{z}) = p^{\mathbf{v}}(x_i)$ whenever x_i is disjoint from the sets in \mathbf{z} and ∞ otherwise.

THEOREM C.2. *The pricing rule defined in (C.2), extended linearly to sets of items, is weakly $(1, 1, k - 1)$ -balanced with respect to an arbitrary allocation rule ALG and MPH- k valuation profile \mathbf{v} .*

Proof. We will show balancedness with respect to $\mathcal{F}_{\mathbf{x}} = \{\mathbf{y} \in \mathcal{F} : (\bigcup_i y_i) \cap (\bigcup_i x_i) = \emptyset\}$. Observe that we can lower-bound the value of $\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})$ by removing all items that are allocated by \mathbf{x} from the allocation $\text{ALG}(\mathbf{v})$. This gives us

$$(C.3) \quad \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) \geq \sum_{\ell \in N} v_{\ell} \left(\text{ALG}_{\ell}(\mathbf{v}) \setminus \left(\bigcup_i x_i \right) \right) \geq \sum_{\ell} \sum_{\substack{T \subseteq \text{ALG}_{\ell}(\mathbf{v}) \\ \forall i: T \cap x_i = \emptyset}} w_{\ell}(T).$$

Furthermore, note first that $p^{\mathbf{v}}$ is a fixed pricing scheme, meaning that the price of an outcome does not change unless it becomes infeasible. Therefore, for every allocation $\mathbf{y} \in \mathcal{F}_{\mathbf{x}}$, we have

$$(C.4) \quad \begin{aligned} \sum_i p_i(y_i \mid \mathbf{x}_{[i-1]}) &= \sum_i p_i(y_i \mid \mathbf{x}) = \sum_i \sum_{j \in y_i} p^{\mathbf{v}}(\{j\}) \\ &= \sum_i \sum_{\ell} \sum_{j \in \text{ALG}_{\ell}(\mathbf{v}) \cap y_i} p^{\mathbf{v}}(\{j\}) \\ &= \sum_{\ell} \sum_i \sum_{j \in \text{ALG}_{\ell}(\mathbf{v}) \cap y_i} \sum_{\substack{T \ni j \\ T \subseteq \text{ALG}_{\ell}(\mathbf{v})}} w_{\ell}(T). \end{aligned}$$

For condition (a), by (C.4),

$$\begin{aligned} \sum_i p_i^{\mathbf{v}}(x_i \mid \mathbf{x}_{[i-1]}) &\geq \sum_{\ell} \sum_{\substack{T \subseteq \text{ALG}_{\ell}(\mathbf{v}) \\ \exists i: T \cap x_i \neq \emptyset}} w_{\ell}(T) \\ &= \left(\sum_{\ell} \sum_{T \subseteq \text{ALG}_{\ell}(\mathbf{v})} w_{\ell}(T) - \sum_{\ell} \sum_{\substack{T \subseteq \text{ALG}_{\ell}(\mathbf{v}) \\ \forall i: T \cap x_i = \emptyset}} w_{\ell}(T) \right) \\ &\geq (\mathbf{v}(\text{ALG}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))), \end{aligned}$$

where the first inequality follows by changing the order of summation over j and T , and the second inequality is given in (C.3).

For condition (b), note that for all \mathbf{x} and all $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$, by (C.4), after splitting the sum depending on whether the respective set T intersects with any of the bundles in \mathbf{x} or not

$$\begin{aligned} \sum_i p_i(x'_i \mid \mathbf{x}_{[i-1]}) &= \sum_{\ell} \sum_i \sum_{j \in \text{ALG}_{\ell}(\mathbf{v}) \cap x'_i} \sum_{\substack{T \ni j \\ T \subseteq \text{ALG}_{\ell}(\mathbf{v}) \\ T \cap (\bigcup_{i'} x_{i'}) = \emptyset}} w_{\ell}(T) \\ &\quad + \sum_{\ell} \sum_i \sum_{j \in \text{ALG}_{\ell}(\mathbf{v}) \cap x_i} \sum_{\substack{T \ni j \\ T \subseteq \text{ALG}_{\ell}(\mathbf{v}) \\ T \cap (\bigcup_{i'} x_{i'}) \neq \emptyset}} w_{\ell}(T). \end{aligned}$$

Observe that in the first sum for a fixed set T , the term $w_{\ell}(T)$ occurs at most $|T|$ times. In the second sum, it can even occur only $|T| - 1$ times because the intersection of \mathbf{x} and \mathbf{x}' is empty but \mathbf{x} intersects with T . By applying that $|T| \leq k$ whenever $w_{\ell}(T) > 0$, this gives us

$$\begin{aligned} \sum_i p_i(x'_i \mid \mathbf{x}_{[i-1]}) &\leq \sum_{\ell} \sum_{\substack{T \subseteq \text{ALG}_{\ell}(\mathbf{v}) \\ T \cap (\bigcup_{i'} x_{i'}) = \emptyset}} |T| w_{\ell}(T) + \sum_{\ell} \sum_{\substack{T \subseteq \text{ALG}_{\ell}(\mathbf{v}) \\ T \cap (\bigcup_{i'} x_{i'}) \neq \emptyset}} (|T| - 1) w_{\ell}(T) \\ &\leq k \sum_{\ell} \sum_{\substack{T \subseteq \text{ALG}_{\ell}(\mathbf{v}) \\ T \cap (\bigcup_{i'} x_{i'}) = \emptyset}} w_{\ell}(T) + (k - 1) \sum_{\ell} \sum_{\substack{T \subseteq \text{ALG}_{\ell}(\mathbf{v}) \\ T \cap (\bigcup_{i'} x_{i'}) \neq \emptyset}} w_{\ell}(T) \\ &= \sum_{\ell} \sum_{\substack{T \subseteq \text{ALG}_{\ell}(\mathbf{v}) \\ T \cap (\bigcup_{i'} x_{i'}) = \emptyset}} w_{\ell}(T) + (k - 1) \sum_{\ell} \sum_{T \subseteq \text{ALG}_{\ell}(\mathbf{v})} w_{\ell}(T) \\ &\geq \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) + (k - 1) \mathbf{v}(\text{ALG}(\mathbf{v})). \quad \square \end{aligned}$$

Note that Theorem C.2 when specialized to XOS valuations shows the existence of weakly $(1, 1, 0)$ -balanced prices or, equivalently, $(1, 1)$ -balanced prices.

Computational $O(k)$ result. We now show how to obtain a *polytime* $(4k - 2)$ -approximate price-based prophet inequality for MPH- k valuations. For this result we will assume access to the following kind of *MPH- k oracle*. Suppose that valuation function v is MPH- k , with supporting PH- k functions $\{v_{\ell}\}_{\ell \in L}$. The query for v takes as input a set of items S and returns a value oracle to access the PH- k function v_{ℓ} for which $v(S) = v_{\ell}(S)$. That is, for every $T \subseteq M$, we can query the value of $v_{\ell}(T)$ in a second step.

The following linear problem, known as the *configuration LP* for combinatorial auctions, computes a fractional allocation that maximizes the social welfare among all fractional allocations:

$$\begin{aligned} &\max \sum_{i,S} v_i(S) \cdot x_{i,S} \\ &\text{s.t.} \sum_S x_{i,S} \leq 1 \text{ for every } i \in N, \\ &\quad \sum_{i,S:j \in S} x_{i,S} \leq 1 \text{ for every } j \in M, \\ &\quad x_{i,S} \in [0, 1] \text{ for every } i \in N, S \subseteq M \end{aligned}$$

We extend the definition of \mathcal{F} to include all fractional allocations that fulfill the above LP; $\mathcal{F}_{\mathbf{x}}$ is extended to all fractional allocations that use at most a fractional

equivalent of $1 - \sum_{i,S:j \in S} x_{i,S}$ of every item $j \in M$. Our first observation will be that Theorem C.2 holds even for the fractional version of the MPH- k combinatorial auction problem. That is, if the set of feasible outcomes is extended to include all fractional allocations, then an appropriate extension of the prices described above remain weakly $(1, 1, k - 1)$ -balanced. In particular, given a fixed valuation profile \mathbf{v} , we define prices based on the optimal LP solution \mathbf{x}^* . For every agent i and every bundle S , let $w_{i,S}$ be the PH- k representation in the support that maximizes agent i 's valuation of S , which is $v_i(S) = \sum_{T \subseteq S} w_{i,S}(T)$. Then, price item j as follows: $p(\{j\}) = \sum_i \sum_S x_{i,S}^* \sum_{T:j \in T, T \subseteq S} w_{i,S}(T) = \sum_i \sum_S x_{i,S}^* (w_{i,S}(S) - w_{i,S}(S \setminus \{j\}))$. This sum can be computed in polynomial time given oracle access as described above because $x_{i,S}^* \neq 0$ only for polynomially many S . Then, extend this pricing linearly to prices over sets of items and to prices over fractional allocations by simple scaling, i.e., for every $\mathbf{x} \in \mathcal{F}$ and $\mathbf{x}' \in \mathcal{F}_x$, we let $p(x'_i | \mathbf{x}) = \sum_S x'_{i,S} \sum_{j \in S} p(\{j\})$.

It is well known that an optimal fractional solution to the configuration LP can be computed in polynomial time given access to demand queries (using demand queries to implement a separation oracle, as is standard). Therefore, for the space of fractional allocations, prices that are weakly $(1, 1, k - 1)$ -balanced with respect to the optimal fractional solution can be computed in polynomial time, given access to demand and MPH- k oracles.

We can therefore compute prices that give a $(4k - 2)$ -approximation to the optimal fractional allocation, in a posted price mechanism, where agents are free to purchase any fractional allocation at the posted prices. This, however, seems unsatisfactory. After all, we do not wish to allow agents to purchase infeasible sets, and the analysis applies only if every agent can purchase a set in her demand correspondence. The following observation comes to our aid: for every agent i , if all previous agents have selected integral outcomes from the posted price mechanism, then agent i has a utility-maximizing outcome in their demand correspondence that is integral. This is because any fractional allocation can be interpreted as a convex combination of integral allocations. Since our approximation guarantee holds regardless of the demanded set chosen by each agent, it will hold even if we restrict agents to only select integral outcomes in the posted price mechanism. This means that we only need to price integral allocations and the analysis follows. We conclude that the prices computed using the configuration LP, as described above, actually generate a $(4k - 2)$ approximation for the (nonfractional) MPH- k combinatorial auction problem.

Appendix D. Proof of Theorem 1.8. In this appendix we prove our polytime price-based prophet inequalities for sparse linear packing programs. We consider programs with and without integrality constraints. That is, the possible outcomes for agent i are either $[0, 1]$ (fractional solutions) or $\{0, 1\}$ (integral solutions). We assume that valuations are linear, i.e., $v_i(x_i) = v_i \cdot x_i$. In both cases, the feasibility constraints \mathcal{F} are given by a constraint matrix $\mathbf{A} \in \mathbb{R}_{\geq 0}^{m \times n}$ such that $\mathbf{x} \in \mathcal{F}$ if and only if $\mathbf{A} \cdot \mathbf{x} \leq \mathbf{c}$. Without loss of generality, let $c_j = 1$ for all $j \in [m]$.

We assume that $a_{j,i} \leq \frac{1}{2}$ for all i, j and that the column sparsity is bounded by d , meaning that for each i there are at most d choices of j such that $a_{j,i} > 0$.

THEOREM D.1. *For d -sparse linear packing programs with constraint matrix $\mathbf{A} \in \mathbb{R}_{\geq 0}^{m \times n}$ such that $a_{j,i} \leq \frac{1}{2}$ for all i, j and unit capacities, there exist weakly $(1, 0, d)$ - and $(2, 0, d)$ -balanced prices, for fractional and integral solutions, respectively, with respect to arbitrary allocation algorithms ALG. The prices can be computed by running ALG once.*

We will define static prices $p_i(x_i | \mathbf{y})$ as follows. Let $\mathbf{x}^* = \text{ALG}(v)$. For each constraint j , we define a per-unit price ρ_j by setting $\rho_j = \sum_{i \in N: a_{j,i} > 0} v_i x_i^*$. Now, we

define $p_i(x_i \mid \mathbf{y}) = \sum_j a_{j,i} x_i \rho_j$ for every quantity x_i that can feasibly be added to \mathbf{y} . The claim will now follow by the two lemmas below.

LEMMA D.2. *For \mathcal{F} being all fractional solutions, the devised pricing scheme is weakly $(1, 0, d)$ -balanced with respect to ALG.*

Proof. Let $\mathcal{F}_{\mathbf{x}} = \{\mathbf{z} \mid \mathbf{A}(\mathbf{x} + \mathbf{z}) \leq \mathbf{1}\}$. To verify condition (a), we derive a lower bound on $\mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))$ as follows. Given \mathbf{x} , let \mathbf{z} be defined by setting $z_i = x_i^*(1 - \max_{j: a_{j,i} > 0} \sum_{i'} a_{j,i'} x_{i'})$. We have $\mathbf{z} \in \mathcal{F}_{\mathbf{x}}$ because for every constraint j we have

$$\sum_i a_{j,i} x_i + \sum_i a_{j,i} z_i \leq \sum_i a_{j,i} x_i + \sum_i a_{j,i} x_i^* \left(1 - \sum_{i'} a_{j,i'} x_{i'}\right) \leq 1.$$

Note that furthermore, by this definition, $x_{i'}^*(\sum_i a_{j,i} x_i) \geq x_{i'}^* - z_{i'}$ for all i' and all j . Now, it follows that

$$\begin{aligned} \sum_i p_i(x_i \mid x_{[i-1]}) &= \sum_i \sum_j a_{j,i} x_i \rho_j = \sum_i \sum_j a_{j,i} x_i \sum_{i': a_{j,i'} > 0} v_{i'} x_{i'}^* \\ &= \sum_j \sum_{i': a_{j,i'} > 0} v_{i'} x_{i'}^* \left(\sum_i a_{j,i} x_i\right) \geq \sum_j \sum_{i': a_{j,i'} > 0} v_{i'} (x_{i'}^* - z_{i'}) \\ &\geq \sum_{i'} v_{i'} (x_{i'}^* - z_{i'}) \geq \mathbf{v}(\text{ALG}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})). \end{aligned}$$

For condition (b), we simply observe that

$$\begin{aligned} \sum_i p_i(x'_i \mid x_{[i-1]}) &= \sum_i \sum_j a_{j,i} x'_i \rho_j = \sum_j \rho_j \sum_j a_{j,i} x'_i \\ &\leq \sum_j \rho_j = \sum_j \sum_{i \in N: a_{j,i} > 0} v_i x_i^* \\ &= \sum_i \sum_{j: a_{j,i} > 0} v_i x_i^* \leq d \sum_i v_i x_i^* = d \cdot \mathbf{v}(\text{ALG}(\mathbf{v})). \quad \square \end{aligned}$$

LEMMA D.3. *For \mathcal{F} being all integral solutions, the devised pricing scheme is weakly $(2, 0, d)$ -balanced with respect to ALG.*

Proof. To define $\mathcal{F}_{\mathbf{x}}$, let \mathbf{c} be the vector such that $c_j = 1$ if $\sum_i a_{j,i} x_i \leq \frac{1}{2}$ and $c_j = 0$ otherwise. That is, $c_j = 1$ if and only if constraint j still has capacity at least $\frac{1}{2}$ after adding \mathbf{x} . Let $\mathcal{F}_{\mathbf{x}}$ be the set of integral solutions \mathbf{y} that fulfill $\mathbf{A}\mathbf{y} \leq \mathbf{c}$.

Let \mathbf{z} be defined such that $z_i = x_i^*$ if $\sum_{i'} a_{j,i'} x_{i'} \leq \frac{1}{2}$ for all j with $a_{j,i} > 0$ and let $z_i = 0$ otherwise. By this definition, we have $z_i \geq x_i^*(1 - 2 \max_{j: a_{j,i} > 0} \sum_{i'} a_{j,i'} x_{i'})$. For this reason, $x_{i'}^*(\sum_i a_{j,i} x_i) \geq \frac{1}{2}(x_{i'}^* - z_{i'})$ for all i' and all j .

From here on, we can follow the exact same calculations as above. For condition (a), we have

$$\begin{aligned} \sum_i p_i(x_i \mid x_{[i-1]}) &= \sum_i \sum_j a_{j,i} x_i \rho_j = \sum_i \sum_j a_{j,i} x_i \sum_{i': a_{j,i'} > 0} v_{i'} x_{i'}^* \\ &= \sum_j \sum_{i': a_{j,i'} > 0} v_{i'} x_{i'}^* \left(\sum_i a_{j,i} x_i\right) \geq \sum_j \sum_{i': a_{j,i'} > 0} v_{i'} \frac{1}{2}(x_{i'}^* - z_{i'}) \\ &\geq \frac{1}{2} \sum_{i'} v_{i'} (x_{i'}^* - z_{i'}) \geq \frac{1}{2} (\mathbf{v}(\text{ALG}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))). \end{aligned}$$

For condition (b), we observe that again

$$\begin{aligned} \sum_i p_i(x'_i \mid x_{[i-1]}) &= \sum_i \sum_j a_{j,i} x'_i \rho_j = \sum_j \rho_j \sum_i a_{j,i} x'_i \leq \sum_j \rho_j \\ &= \sum_j \sum_{i \in N: a_{j,i} > 0} v_i x_i^* = \sum_i \sum_{j: a_{j,i} > 0} v_i x_i^* \\ &\leq d \sum_i v_i x_i^* = d \cdot \mathbf{v}(\text{ALG}(\mathbf{v})). \quad \square \end{aligned}$$

Pricing integral problems based on fractional solutions. The way we described the pricing schemes above was to use an offline allocation algorithm ALG for the respective problem. This algorithm has to solve an NP-hard problem and as we show any approximation guarantee is preserved in the process. However, it is also possible to use the respective optimal fractional solution instead of ALG . The pricing schemes defined this way then achieve the described approximation guarantees with respect to the fractional optimum. To show this result, we have to slightly extend our framework and allow $\mathcal{F}_{\mathbf{x}}$ to contain distributions over outcome profiles that are exchange compatible.

Specifically, for the integral part of Theorem D.1 we would define $\mathcal{F}_{\mathbf{x}}$ as the set of distributions such that the *expected* consumption in the j th constraint is at most 1 if $\sum_i a_{j,i} x_i \leq \frac{1}{2}$ and 0 otherwise. Note that this extends the definition of $\mathcal{F}_{\mathbf{x}}$ in the proof of Lemma D.3 to distributions. Following the steps above, \mathbf{z} is now randomized based on the optimal fractional solution \mathbf{x}^* . We let $z_i = 1$ with probability x_i^* if $\sum_{i'} a_{j,i'} x_{i'} \leq \frac{1}{2}$ for all j with $a_{j,i} > 0$ and let $z_i = 0$ otherwise. In the remainder, $z_{i'}$ is replaced by its expectation.

Appendix E. Proof of Theorem 1.9. We consider a matroid feasibility constraint $\mathcal{M} = (E, \mathcal{I})$, where E is a ground set of elements and $\mathcal{I} \subseteq 2^E$ is a family of feasible subsets. Set E is partitioned into subsets E_1, \dots, E_n , corresponding to agents, and for each i , the outcome space X_i contains all subsets of E_i . One can think of E_i as the set of elements from which agent i can choose. An allocation \mathbf{x} is feasible if and only if $\bigcup_i x_i \in \mathcal{I}$, that is, if the chosen elements form an independent set.⁵ Given a valuation profile \mathbf{v} , let $\text{OPT}(\mathbf{v}) \in \mathcal{I}$ denote a social welfare maximizing allocation. Furthermore, given $S \in \mathcal{I}$, let $\text{OPT}(\mathbf{v} \mid S)$ denote the set $T \in \mathcal{I}$ maximizing $\mathbf{v}(T)$ such that $S \cap T = \emptyset$ and $S \cup T \in \mathcal{I}$. Theorem E.1 asserts the existence of balanced prices for matroid settings with additive, submodular, and XOS valuations.

THEOREM E.1. *There exist prices for multidimensional matroid settings in which agents have additive, submodular, or XOS valuations that are (1,1)-balanced with respect to OPT.*

As a special case, we obtain (1,1)-balanced prices for the single-dimensional matroid setting of [33], where each agent controls exactly one element of the ground set.

We prove the theorem for additive valuations. That is, there are nonnegative numbers $v_{i,j}$ such that $v_i(x_i) = \sum_{j \in x_i} v_{i,j}$. The proof for submodular and XOS valuations then follows directly from Theorem 5.3 (composition theorem: closure under maximum).

⁵Kleinberg and Weinberg [33] also provide bounds for a multidimensional matroid setting. Their result, which has implications for revenue and welfare, applies when buyers have unit-demand preferences with independent values across elements. Our approach is different and provides a welfare bound for a broader class of valuations.

Additive valuations. Given $S \in \mathcal{I}$, $i \in N$, and a valuation profile \mathbf{v} , define

$$(E.1) \quad p_i^{\mathbf{v}}(x_i | \mathbf{y}) = \begin{cases} \mathbf{v}(\text{OPT}(\mathbf{v} | \bigcup_j y_j)) - \mathbf{v}(\text{OPT}(\mathbf{v} | (\bigcup_j y_j \cup x_i))) & \text{if } (\bigcup_j y_j \cup x_i) \in \mathcal{I}, \\ \infty & \text{otherwise.} \end{cases}$$

The following lemma shows that the prices in (E.1) are monotonically increasing. We then prove in Theorem E.3 that these prices are (1, 1)-balanced.

LEMMA E.2. *Consider the pricing rule $p^{\mathbf{v}}$. For an arbitrary profile \mathbf{v} , feasible outcome profiles $\mathbf{y}, \mathbf{y}' \in \mathcal{F}$ such that $y_j \subseteq y'_j$ for all j , agent i , and allocation x_i ,*

$$p_i^{\mathbf{v}}(x_i | \mathbf{y}) \leq p_i^{\mathbf{v}}(x_i | \mathbf{y}').$$

Proof. Let $S = \bigcup_j y_j$, $T = \bigcup_j y'_j$. If $S \cup x_i \notin \mathcal{I}$, then $T \cup x_i \notin \mathcal{I}$, and so $p_i^{\mathbf{v}}(x_i | \mathbf{y}) = p_i^{\mathbf{v}}(x_i | \mathbf{y}') = \infty$. Otherwise,

$$\begin{aligned} p_i^{\mathbf{v}}(x_i | \mathbf{y}) &= \mathbf{v}(\text{OPT}(\mathbf{v} | S)) - \mathbf{v}(\text{OPT}(\mathbf{v} | (S \cup x_i))) \\ &\leq \mathbf{v}(\text{OPT}(\mathbf{v} | T)) - \mathbf{v}(\text{OPT}(\mathbf{v} | (T \cup x_i))) \\ &\leq p_i^{\mathbf{v}}(x_i | \mathbf{y}'), \end{aligned}$$

where the equation follows by the definition of $p_i^{\mathbf{v}}(x_i | \mathbf{y})$, and the inequality follows from the submodularity of the function $f(U) = \mathbf{v}(\text{OPT}(\mathbf{v} | U))$ (cf. [33, Lemma 3]). \square

THEOREM E.3. *The pricing rule defined in (E.1) is (1, 1)-balanced with respect to the optimal allocation rule OPT.*

Proof. We define $\mathcal{F}_{\mathbf{x}}$ as the set of all outcome profiles $\mathbf{y} \in E_1 \times \dots \times E_n$ such that $(\bigcup_i y_i) \cap (\bigcup_i x_i) = \emptyset$ and $(\bigcup_i y_i) \cup (\bigcup_i x_i) \in \mathcal{I}$. In other words, we consider the matroid contracted by the set $\bigcup_i x_i$.

We first show that condition (a) holds with $\alpha = 1$. For every agent $i \in N$ and every $\mathbf{x} \in \mathcal{F}$ by a telescoping-sum argument,

$$\begin{aligned} \sum_i p_i^{\mathbf{v}}(x_i | \mathbf{x}_{[i-1]}) &= \sum_i \left[\mathbf{v} \left(\text{OPT} \left(\mathbf{v} \mid \bigcup_{j=1}^{i-1} x_j \right) \right) - \mathbf{v} \left(\text{OPT} \left(\mathbf{v} \mid \bigcup_{j=1}^i x_j \right) \right) \right] \\ &= \mathbf{v}(\text{OPT}(\mathbf{v})) - \mathbf{v} \left(\text{OPT} \left(\mathbf{v} \mid \bigcup_j x_j \right) \right) \\ &= \mathbf{v}(\text{OPT}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}})). \end{aligned}$$

We next show that condition (b) holds with $\beta = 1$. Consider some arbitrary $\mathbf{x} \in \mathcal{F}$, $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$. Let $S = \bigcup_i x_i$. Note that $\text{OPT}(\mathbf{v} | S)$ is precisely a maximum-weight basis of the matroid contracted with S . By the generalized Rota exchange theorem [36, Lemma 2.7], for each i there exists some set $R_i \subseteq \text{OPT}(\mathbf{v} | S)$ (which may not be contained in E_i) such that each element of $\text{OPT}(\mathbf{v} | S)$ appears in exactly one R_i , and $(\text{OPT}(\mathbf{v} | S) \setminus R_i) \cup x'_i$ is an independent set in the matroid after contracting S .

We therefore have

$$\begin{aligned}
 \sum_{i \in N} p_i^{\mathbf{v}}(x'_i \mid \mathbf{x}_{[i-1]}) &\leq \sum_{i \in N} p_i^{\mathbf{v}}(x'_i \mid \mathbf{x}) \\
 &= \sum_{i \in N} [\mathbf{v}(\text{OPT}(\mathbf{v} \mid S)) - \mathbf{v}(\text{OPT}(\mathbf{v} \mid S \cup x'_i))] \\
 &\leq \sum_{i \in N} [\mathbf{v}(\text{OPT}(\mathbf{v} \mid S)) - \mathbf{v}(\text{OPT}(\mathbf{v} \mid S) \setminus R_i)] \\
 &= \sum_{i \in N} \mathbf{v}(R_i) \\
 &= \mathbf{v}(\text{OPT}(\mathbf{v} \mid \mathbf{x})),
 \end{aligned}$$

where the first inequality follows from the monotonicity of the prices (Lemma E.2), and the second inequality follows by observing that $(\text{OPT}(\mathbf{v}) \setminus R_i) \cup x'_i$ is feasible in the contracted matroid (see the Rota exchange argument above), and thus $\mathbf{v}(\text{OPT}(\mathbf{v}) \setminus R_i) \leq \mathbf{v}(\text{OPT}(\mathbf{v} \mid S \cup x'_i))$. The final equality follows by additivity. We conclude that the suggested prices are $(1, 1)$ -balanced. \square

Computational aspects. Theorem E.1 establishes the existence of a price-based 2-approximate prophet inequality for additive, submodular, and XOS preferences. For additive valuations the construction is polytime, as the greedy algorithm is optimal. For submodular and XOS valuations computing the optimal allocation is NP-hard.

We claim that the result for submodular valuations can turn into a computational one, by basing the prices on an approximately optimal allocation. Namely, let GRD be the algorithm that allocates items greedily by value, always choosing the item that locally increases $\mathbf{v}(\mathbf{x})$ the most subject to the matroid constraint. We show that GRD is consistent (see Definition 5.1) in this setting.

LEMMA E.4. *The greedy algorithm GRD is consistent for XOS valuations over matroid feasibility constraints.*

Proof. Fix an XOS valuation profile \mathbf{v} , and let $\tilde{\mathbf{v}}$ be a supporting valuation profile for the allocation $\text{GRD}(\mathbf{v})$. Note that as the supporting valuation profile is additive, GRD returns the optimal solution. As $\text{GRD}(\mathbf{v})$ is another feasible solution, we have $\tilde{\mathbf{v}}(\text{GRD}(\tilde{\mathbf{v}})) \geq \tilde{\mathbf{v}}(\text{GRD}(\mathbf{v}))$. \square

Theorem 5.3 (closure under maximum) now implies that the prices given in (E.1) for an additive valuation supporting $\text{GRD}(\mathbf{v})$ are $(1, 1)$ -balanced with respect to the greedy allocation. We can compute these prices in polynomial time by simulating GRD and then determining the supporting additive valuation. Since GRD is a 2-approximation to OPT for submodular valuations, Theorem 3.2 then implies that we can compute prices that yield a 4-approximation to the optimal expected welfare, less an additive sampling error.

Appendix F. A smooth mechanism without good posted prices. In this appendix we show that there are allocation problems that admit constant-factor smooth mechanisms, but for which no posted-price mechanism can guarantee more than a linear fraction of the optimal social welfare.

PROPOSITION F.1. *There exists a downward-closed welfare maximization problem that admits a $(1, 0)$ -smooth mechanism, but for which any posted price mechanism has approximation factor $\Omega(n)$.*

Proof. Let k be a positive integer to be fixed later. In the allocation problem we consider, $X_i = [k]^n \cup \{\emptyset\}$ for each $i \in [n]$. That is, each agent is allocated either \emptyset or a sequence of n integers between 1 and k . For $x_i \in X_i \setminus \{\emptyset\}$, write $x_i = (x_{i1}, \dots, x_{in})$, where each $x_{ij} \in [k]$. The set of feasible allocations is $\mathcal{F} = \{\mathbf{x}: (x_i \neq x_j) \implies ((x_i = \emptyset) \vee (x_j = \emptyset))\}$. That is, all agents who receive a nonempty allocation must receive the same allocation.

Agents have the following valuations. Each agent i has some desired value $z_i \in [k]$. The value of an allocation is $v_i(x_i) = 1$ if $x_{ii} = z_i$ and $v_i(x_i) = 0$ otherwise. Let \mathcal{D}_i be the distribution over such valuations in which z_i is chosen uniformly from $[k]$.

For this feasibility constraint and space of valuations, consider the mechanism that returns the welfare-optimal allocation and charges payments of 0. This mechanism simultaneously satisfies all desires by allocating (z_1, \dots, z_n) to every agent, where z_i is the desired value reported by agent i . This mechanism is $(1, 0)$ -smooth, with truth-telling being the required deviation.

On the other hand, consider any posted-price mechanism. Whichever is the first agent to obtain a non- \emptyset outcome, say, agent i purchasing allocation x_i , each subsequent agent j can obtain positive value only if $z_j = x_{ij}$, which occurs with probability $1/k$. We therefore have that the expected welfare of any posted price mechanism is at most $1 + \frac{n-1}{k}$. Taking $k = n$ and noting that the optimal welfare is n , we conclude that no posted price mechanism can obtain more than an $\Omega(n)$ approximation to the optimal welfare. \square

Appendix G. Proofs of Theorems 6.3 and 6.4. In this appendix we describe how to obtain balanced prices from smooth mechanisms in binary, single-parameter settings.

G.1. Permeable allocation rules. We begin by defining the *permeability* of an allocation rule f and by showing that (λ, μ) -smoothness implies a bound on permeability.

An algorithm f for a binary, single-parameter problem is *monotone* if for every player $i \in N$, any two bids $b'_i \geq b_i$, and any bid vector \mathbf{b}_{-i} ,

$$f_i(b_i, \mathbf{b}_{-i}) = 1 \implies f_i(b'_i, \mathbf{b}_{-i}) = 1.$$

For monotone allocation rules the *critical value* for player i is the smallest bid that ensures that player i wins against bids \mathbf{b}_{-i} . That is, $\tau_i^f(\mathbf{b}_{-i}) = \inf\{b_i \mid f_i(b_i, \mathbf{b}_{-i}) = 1\}$.

DEFINITION G.1 (Dütting and Kesselheim [20]; see also [38, 45, 31]). *A monotone allocation algorithm f for a binary, single-parameter problem \mathcal{F} is γ -permeable if $\gamma \geq 1$ is the smallest multiplier such that for all bid vectors \mathbf{b} and all feasible allocations $\mathbf{x} \in \mathcal{F}$ it holds that*

$$(G.1) \quad \sum_{i \in N: x_i=1} \tau_i^f(\mathbf{b}_{-i}) \leq \gamma \cdot \mathbf{b}(f(\mathbf{b})).$$

THEOREM G.2 (Dütting and Kesselheim [20]). *Suppose that the first-price mechanism \mathcal{M} based on allocation f is (λ, μ) -smooth; then f is γ -permeable with $\gamma \leq (\mu + 1)/\lambda$.*

Proof. Given a bid vector \mathbf{b} , we have to show that $\sum_{i: x_i=1} \tau_i^f(\mathbf{b}_{-i}) \leq \frac{\mu+1}{\lambda} \cdot \mathbf{b}(f(\mathbf{b}))$. To this end, consider fixed $\mathbf{x} \in \mathcal{F}$ and \mathbf{b} . Let $\epsilon > 0$ and let \mathbf{v} be defined by $v_i = \max\{b_i, \tau_i^f(\mathbf{b}_{-i})\}$. By smoothness of \mathcal{M} , there are b'_i such that

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

Observe that if $i \notin f(\mathbf{b})$, then $u_i(b'_i, \mathbf{b}_{-i}) \leq 0$ because $\tau_i^f(\mathbf{b}_{-i}) > b_i$ and this means that $i \notin f(b'_i, \mathbf{b}_{-i})$ unless $b'_i > v_i$. None of these choices results in positive utility. Furthermore, for $i \in f(\mathbf{b})$, we have $u_i(b'_i, \mathbf{b}_{-i}) \leq v_i = b_i$. Therefore $\sum_{i \in N} u_i(b'_i, \mathbf{b}_{-i}) \leq \mathbf{b}(f(\mathbf{b}))$.

Next, we can lower-bound $\mathbf{v}(\text{OPT}(\mathbf{v}))$ by the value of the feasible solution \mathbf{x} , which gives us $\mathbf{v}(\text{OPT}(\mathbf{v})) \geq \sum_{i: x_i=1} v_i \geq \sum_{i: x_i=1} (\tau_i^f(\mathbf{b}_{-i}) - \epsilon) \geq \sum_{i: x_i=1} \tau_i^f(\mathbf{b}_{-i}) - n\epsilon$.

Finally, as \mathcal{M} is first-price, we also have $\sum_{i \in N} P_i(\mathbf{b}) = \mathbf{b}(f(\mathbf{b}))$.

In combination this yields

$$\mathbf{b}(f(\mathbf{b})) \geq \lambda \cdot \left(\sum_{i: x_i=1} \tau_i^f(\mathbf{b}_{-i}) - n\epsilon \right) - \mu \cdot \mathbf{b}(f(\mathbf{b})),$$

which implies

$$\sum_{i: x_i=1} \tau_i^f(\mathbf{b}_{-i}) \leq \frac{\mu + 1}{\lambda} \cdot \mathbf{b}(f(\mathbf{b})) + n\epsilon.$$

As this holds for all $\epsilon > 0$, this shows the claim. □

Applying Theorem G.2 to each problem in a collection of problems Π , we see that if a mechanism \mathcal{M} is (λ, μ) -smooth for Π , then it is $(\mu + 1)/\lambda$ -permeable for Π .

Remark G.3. While the definition of permeability requires γ to be the smallest multiplier for which inequality (G.1) is satisfied, all our results can be derived from any upper bound on this multiplier at the cost of slightly worse guarantees.

G.2. Proof of Theorem 6.3. In this subsection we prove Theorem 6.3, which shows that (λ, μ) -smoothness of the greedy allocation rule for a subinstance-closed collection of binary, single-parameter problems Π implies the existence of a weakly $((\mu + 1)/\lambda, 0, (\mu + 1)/\lambda)$ -balanced pricing rule. By Theorem G.2, in order to show this result, it suffices to show the following theorem.

THEOREM G.4. *Let ALG be any allocation rule. Suppose that the greedy allocation rule GRD is γ -permeable for a subinstance-closed collection of binary, single-parameter feasibility problems Π . Then for every $\mathbf{v} \in V$ there exists a pricing rule that is weakly $(\gamma, 0, \gamma)$ -balanced with respect to ALG and the canonical exchange-feasible sets $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$.*

We first describe the pricing rule that achieves this result. Afterward, we show that this pricing rule has the desirable properties.

G.2.1. Construction of the prices. We set the price $p_i(z_i \mid \mathbf{x})$ for player $i \in N$ and outcome $z_i \in \{0, 1\}$ for arbitrary but fixed valuations \mathbf{v} and allocation $\mathbf{x} \in \mathcal{F}$ through Algorithm G.1. For this let $\text{GRD}(\mathbf{v} \mid \mathbf{x}) \in \mathcal{F}_{\mathbf{x}}$ denote the allocation that results if we go through the players in order of nonincreasing value but only add a player if he is not in \mathbf{x} and feasible together with \mathbf{x} and the previously accepted players.

We generally set $p_i(0 \mid \mathbf{x}) = 0$. That is, the price for losing is always zero. To determine the price $p_i(1 \mid \mathbf{x})$ for winning we first compute a sequence of reference allocations $\mathbf{r}^{(0)} \geq \dots \geq \mathbf{r}^{(n)}$ and a sequence of reference valuations $\mathbf{v}^{(0)} \geq \dots \geq \mathbf{v}^{(n)}$. We then set $p_i(1 \mid \mathbf{x}) = v_i$ if $i \in \mathbf{r}^{(n)}$, i.e., the price of player i is that player's valuation if he is part of the final reference allocation. Otherwise, we set $p_i(1 \mid \mathbf{x}) = \inf\{v'_i : i \in \text{GRD}(v'_i, \mathbf{v}_{-i}^{(n)} \mid \mathbf{x})\}$, i.e., we set the price to the player's critical value against the players in the final reference allocation.

Algorithm G.1. Pricing rule derived from GRD (parametrized by ALG).

Input: $z_i \in \{0, 1\}$, $\mathbf{v}, \mathbf{x} \in \mathcal{F}$
Output: $p_i(z_i \mid \mathbf{x})$
if $z_i = 0$ **then**
 // In this case the price is simply zero
 $p_i(z_i \mid \mathbf{x}) = 0$
else
 // First determine reference allocation and valuations
 $\mathbf{r}^{(0)} \leftarrow \text{ALG}(\mathbf{v})$
 $v_k^{(0)} \leftarrow v_k$ if $k \in \mathbf{r}^{(0)}$ and $v_k^{(0)} \leftarrow 0$ otherwise
 for $j \leftarrow 1$ **to** n **do**
 $\mathbf{r}^{(j)} \leftarrow \text{GRD}(\mathbf{v}^{(j-1)} \mid \mathbf{x}_{[j]})$
 $v_k^{(j)} \leftarrow v_k^{(j-1)} = v_k$ if $k \in \mathbf{r}^{(j)}$ and $v_k^{(j)} \leftarrow 0$ else
 end for
 // Now determine the price
 if $i \in \mathbf{r}^{(n)}$ **then**
 // If player i is part of the reference allocation he pays his valuation
 $p_i(z_i \mid \mathbf{x}) \leftarrow v_i$
 else
 // Otherwise he pays the critical value against the players in the
 // reference allocation
 $p_i(z_i \mid \mathbf{x}) \leftarrow \inf\{v'_i : i \in \text{GRD}(v'_i, \mathbf{v}_{-i}^{(n)} \mid \mathbf{x})\}$
 end if
end if
return $p_i(z_i \mid \mathbf{x})$

While we need to define prices $p_i(z_i \mid \mathbf{x})$ for any possible allocation $\mathbf{x} \in \mathcal{F}$, the prices that player i will actually see are the ones where \mathbf{x} is set to the purchase decisions of the players $j = 1, \dots, i-1$ that precede player i in the ordering. Note that in this case $x_i = x_{i+1} = x_n = 0$ and therefore $\mathbf{r}^{(n)} = \dots = \mathbf{r}^{(i-1)}$ and $\mathbf{v}^{(n)} = \dots = \mathbf{v}^{(i-1)}$. We use the shorthand $\tau_i^{\text{GRD}}(\mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]}) := \inf\{v'_i : i \in \text{GRD}(v'_i, \mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]})\}$.

G.2.2. Proof of Theorem G.4. We prove the theorem in two steps. We first use permeability of the greedy allocation rule to establish condition (a) (in Lemma G.5). We then show condition (b). For this we first prove a novel combinatorial implication of permeability of the greedy allocation rule (in Lemma G.6) by considering valuations that are either zero or one. We then use this property in a careful layering argument to establish condition (b) (in Lemma G.7).

LEMMA G.5. *Let ALG be any allocation rule. Suppose that the greedy allocation rule GRD is γ -permeable for a subinstance-closed collection of binary, single-parameter problems Π . Then the pricing rule described in Algorithm G.1 fulfills condition (a) of Definition 3.4 with $\alpha = \gamma$ with respect to allocation rule ALG and the canonical exchange-feasible sets $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$.*

Proof. Let $\mathbf{x} \in \mathcal{F}$. We will show that $p_i(x_i \mid \mathbf{x}_{[i-1]}) \geq \frac{1}{\gamma} \cdot (\mathbf{v}(\mathbf{r}^{(i-1)}) - \mathbf{v}(\mathbf{r}^{(i)}))$. By a telescoping-sum argument, this then implies

$$\begin{aligned} \sum_{i \in N} p_i(x_i \mid \mathbf{x}_{[i-1]}) &\geq \sum_{i \in N} \frac{1}{\gamma} \cdot \left(\mathbf{v}(\mathbf{r}^{(i-1)}) - \mathbf{v}(\mathbf{r}^{(i)}) \right) \\ &= \frac{1}{\gamma} \cdot \left(\mathbf{v}(\mathbf{r}^{(0)}) - \mathbf{v}(\mathbf{r}^{(n)}) \right) \\ &\geq \frac{1}{\gamma} \cdot (\mathbf{v}(\text{ALG}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))), \end{aligned}$$

where the last step follows from the fact that $\mathbf{r}^{(0)} = \text{ALG}(\mathbf{v})$ and $\mathbf{r}^{(n)} \in \mathcal{F}_{\mathbf{x}}$.

So, it only remains to show $p_i(x_i \mid \mathbf{x}_{[i-1]}) \geq \frac{1}{\gamma} \cdot (\mathbf{v}(\mathbf{r}^{(i-1)}) - \mathbf{v}(\mathbf{r}^{(i)}))$. Observe that if $x_i = 0$, we have $\mathbf{r}^{(i-1)} = \mathbf{r}^{(i)}$ and this claim follows trivially. So, consider an arbitrary player i for which $x_i = 1$. If $i \in \mathbf{r}^{(i-1)}$, then $\mathbf{r}^{(i-1)} \setminus \mathbf{r}^{(i)} = \{i\}$. So $p_i(1 \mid \mathbf{x}_{[i-1]}) = v_i$, while $\mathbf{v}(\mathbf{r}^{(i-1)}) - \mathbf{v}(\mathbf{r}^{(i)}) = v_i$ and the claim is true.

Otherwise, $i \notin \mathbf{r}^{(i-1)}$, and we will first use smoothness with respect to subinstances to bound the size of the set $\mathbf{r}^{(i-1)} \setminus \mathbf{r}^{(i)}$. For a fixed $\epsilon > 0$, define \mathbf{v}' by setting $v'_i = \tau_i^{\text{GRD}}(\mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]}) + \epsilon$, $v'_j = v_j$ for $j \in \mathbf{r}^{(i-1)} \setminus \mathbf{r}^{(i)}$, and $v'_j = 0$ for all other j .

Now player $i \in \text{GRD}(\mathbf{v}' \mid \mathbf{x}_{[i-1]} \cup \mathbf{r}^{(i)})$ by definition of \mathbf{v}' and v'_i in particular, while for each player $j \in \mathbf{r}^{(i-1)} \setminus \mathbf{r}^{(i)}$ we have $j \notin \text{GRD}(\mathbf{v}' \mid \mathbf{x}_{[i-1]} \cup \mathbf{r}^{(i)})$ because it cannot be added to $\mathbf{x}_{[i-1]} \cup \mathbf{r}^{(i)} \cup \{i\} = \mathbf{x}_{[i]} \cup \mathbf{r}^{(i)}$ by definition of $\mathbf{r}^{(i)}$. Hence, the greedy critical values of each player $j \in \mathbf{r}^{(i-1)} \setminus \mathbf{r}^{(i)}$ must be at least $\tau_j^{\text{GRD}}(\mathbf{v}' \mid \mathbf{x}_{[i-1]} \cup \mathbf{r}^{(i)}) \geq v'_i$.

Since both player i and the set of players $\mathbf{r}^{(i-1)} \setminus \mathbf{r}^{(i)}$ are feasible extensions to $\mathbf{x}_{[i-1]} \cup \mathbf{r}^{(i)}$, we can use γ -permeability of the greedy allocation rule in the subinstance in which we hold $\mathbf{x}_{[i-1]} \cup \mathbf{r}^{(i)}$ fixed to obtain

$$\begin{aligned} |\mathbf{r}^{(i-1)} \setminus \mathbf{r}^{(i)}| \cdot v'_i &\leq \sum_{j \in \mathbf{r}^{(i-1)} \setminus \mathbf{r}^{(i)}} \tau_j^{\text{GRD}}(\mathbf{v}' \mid \mathbf{x}_{[i-1]} \cup \mathbf{r}^{(i)}) \\ &\leq \gamma \cdot \mathbf{v}'(\text{GRD}(\mathbf{v}' \mid \mathbf{x}_{[i-1]} \cup \mathbf{r}^{(i)})) \\ &= \gamma \cdot v'_i. \end{aligned}$$

Cancelling v'_i shows that $|\mathbf{r}^{(i-1)} \setminus \mathbf{r}^{(i)}| \leq \gamma$.

To show the claim it now suffices to observe that the greedy critical value of player i in the subinstance where we hold $\mathbf{x}_{[i-1]}$ fixed under the original valuations is the highest value of a player $j \in \mathbf{r}^{(i-1)} \setminus \mathbf{r}^{(i)}$. Namely,

$$p_i(1 \mid \mathbf{x}_{[i-1]}) = \max_{j \in \mathbf{r}^{(i-1)} \setminus \mathbf{r}^{(i)}} v_j \geq \frac{1}{\gamma} \cdot \sum_{j \in \mathbf{r}^{(i-1)} \setminus \mathbf{r}^{(i)}} v_j = \mathbf{v}(\mathbf{r}^{(i-1)}) - \mathbf{v}(\mathbf{r}^{(i)}),$$

which concludes the proof. □

LEMMA G.6. *Suppose that the greedy allocation rule GRD is γ -permeable for a subinstance-closed collection of problems of binary, single-parameter problems Π . Consider any problem $\pi \in \Pi$ with feasibility structure \mathcal{F} . Furthermore, let $B_0 \supseteq B_1 \supseteq \dots \supseteq B_n$ and $A_0 \subseteq A_1 \subseteq \dots \subseteq A_n$ with $B_t \cup A_t \in \mathcal{F}$ for all t . Consider a set C that fulfills $C \cup A_n \in \mathcal{F}$ and for every $i \in C$ there is a $t \in \{0, 1, \dots, n\}$ with $i \in B_t$ or $\{i\} \cup B_t \cup A_t \notin \mathcal{F}$. Then we have $|C| \leq \gamma \cdot |B_0|$.*

Proof. Set $B_{n+1} = \emptyset$ and define $C_{n+1} = C$. Furthermore, define C_t for $0 \leq t \leq n$ recursively as a maximal subset of the players in $C_{t+1} \setminus B_t$ such that $C_t \cup A_t \cup B_t \in \mathcal{F}$.

We will show that for all $t \in \{0, 1, \dots, n\}$,

$$|C_{t+1} \setminus C_t| \leq \gamma \cdot |B_t \setminus B_{t+1}|.$$

Consider some fixed t and define $D := B_t \cap C_{t+1}$.

Now a crucial observation is that the set $C_{t+1} \setminus (C_t \cup D)$ is feasible holding $E := B_{t+1} \cup A_t \cup C_t \cup D$ fixed. This is because

$$(C_{t+1} \setminus (C_t \cup D)) \cup B_{t+1} \cup A_t \cup C_t \cup D = B_{t+1} \cup A_t \cup C_{t+1} \subseteq B_{t+1} \cup A_{t+1} \cup C_{t+1} \in \mathcal{F}.$$

Further note that by the way we have chosen C_t we know that $B_t \setminus (B_{t+1} \cup D)$ is another feasible extension to E because

$$(B_t \setminus (B_{t+1} \cup D)) \cup B_{t+1} \cup A_t \cup C_t \cup D = B_t \cup A_t \cup C_t \in \mathcal{F}.$$

To apply γ -permeability in the subinstance where we hold E fixed, define a valuation profile $\bar{\mathbf{v}}$ by setting $\bar{v}_i = 1$ for $i \in B_t \setminus (B_{t+1} \cup D)$ and 0 otherwise. Now, for every $i \in C_{t+1} \setminus (C_t \cup D)$, we have

$$\tau_i^{\text{GRD}}(\bar{\mathbf{v}} \mid E) = 1.$$

This is due to the maximality of C_t : if for some $i \in C_{t+1} \setminus (C_t \cup D)$ this value is 0, then also $\{i\} \cup (B_t \setminus (B_{t+1} \cup D)) \cup B_{t+1} \cup A_t \cup C_t \cup D \in \mathcal{F}$.

So by permeability,

$$\begin{aligned} |C_{t+1} \setminus (C_t \cup D)| &= \sum_{i \in C_{t+1} \setminus (C_t \cup D)} \tau_i^{\text{GRD}}(\bar{\mathbf{v}} \mid E) \\ &\leq \gamma \cdot \bar{\mathbf{v}}(\text{GRD}(\bar{\mathbf{v}} \mid E)) \\ &= \gamma \cdot |B_t \setminus (B_{t+1} \cup D)|, \end{aligned}$$

and therefore

$$|C_{t+1} \setminus C_t| = |D| + |C_{t+1} \setminus (C_t \cup D)| \leq |D| + \gamma \cdot |B_t \setminus (B_{t+1} \cup D)| \leq \gamma \cdot |B_t \setminus B_{t+1}|.$$

We now obtain the desired bound on the size of the set C by summing the previous inequality over all t and using that it becomes a telescoping sum

$$\begin{aligned} |C| + |C_0| &= |C_{n+1}| + |C_0| = \sum_{t=0}^n |C_{t+1} \setminus C_t| \\ &\leq \gamma \cdot \sum_{t=0}^n |B_t \setminus B_{t+1}| \\ &= \gamma \cdot (|B_0| - |B_{n+1}|) = \gamma \cdot |B_0|. \end{aligned}$$

It remains to show that all players in C will be covered (i.e., that $C_0 = \emptyset$). This follows from the fact that for each player $i \in C$ by the definition of C there exists a t such that $i \in B_t$ or i does not fit into $B_t \cup A_t$ and, thus, in either case $i \notin C_t$. \square

LEMMA G.7. *Let ALG be any allocation rule. Suppose that the greedy allocation rule GRD is γ -permeable for a subinstance-closed collection of binary, single-parameter problems Π . Then the pricing rule described in Algorithm G.1 fulfills condition (b) of Definition 3.4 with $\beta_1 = 0$ and $\beta_2 = \gamma$ with respect to allocation rule ALG and the canonical exchange-feasible sets $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$.*




FIG. 1. Our proof that the pricing rule derived from the greedy allocation rule satisfies condition (b) relies on the fact that we can chop the valuation and price space into discrete layers, which reduces the problem to 0/1-valuations.

Proof. Consider an arbitrary feasible $\mathbf{x} \in \mathcal{F}$ and an arbitrary feasible extension $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$. We want to show that

$$\sum_{i \in N} p_i(x'_i \mid \mathbf{x}_{[i-1]}) \leq \gamma \cdot \mathbf{v}(\text{ALG}(\mathbf{v})).$$

We will prove this claim through a layering argument, as depicted in Figure 1. To this end, let $v_{(j)}$ be the j th highest value of v_1, \dots, v_n ; furthermore $v_{(n+1)} = 0$. For each $j \in [n]$, let S^j denote the set of players with value at least $v_{(j)}$ and let T^j denote the set of players with $x'_i = 1$ that see a price $p_i(1 \mid \mathbf{x}_{[i-1]})$ of at least $v_{(j)}$.

We now apply Lemma G.6 for each $j \in [n]$ by setting $A_t = \mathbf{x}_{[t]}$, $B_t = \mathbf{r}^{(t)} \cap S^j$, $C = T^j$. Note that $C \cup A_n \in \mathcal{F}$ and for every $i \in C$ there is a $t \in \{0, 1, \dots, n\}$ such that $i \in B_t$ or $\{i\} \cup B_t \cup A_t \notin \mathcal{F}$. We obtain $|T^j| \leq \gamma \cdot |\mathbf{r}^{(0)} \cap S^j|$.

We conclude that

$$\begin{aligned} \sum_{i \in N} p_i(x'_i \mid \mathbf{x}_{[i-1]}) &= \sum_{i \in N} \sum_{j=1}^n \mathbb{1}_{i \in T^j} \cdot (v_{(j)} - v_{(j+1)}) = \sum_{j=1}^n |T^j| \cdot (v_{(j)} - v_{(j+1)}) \\ &\leq \sum_{j=1}^n \gamma \cdot |\mathbf{r}^{(0)} \cap S^j| \cdot (v_{(j)} - v_{(j+1)}) = \gamma \cdot \mathbf{v}(\text{ALG}(\mathbf{v})), \end{aligned}$$

where the first equality holds by definition of the sets T^j , the second equality is basic calculus, the inequality follows from Lemma G.6 as argued above, and the final equality holds by definition of $\mathbf{r}^{(0)} = \text{ALG}(\mathbf{v})$ and the sets S^j . \square

G.3. Proof of Theorem 6.4. In this subsection we prove Theorem 6.4, which claims that (λ, μ) -smoothness of the pay-your-bid mechanism based on the welfare-maximizing allocation rule for subinstance-closed collection of binary, single-parameter problems Π implies the existence of a weakly $(1, (\mu + 1)/\lambda)$ -balanced pricing rule. By Theorem G.2 it suffices to show the following theorem.

THEOREM G.8. *Let ALG be any allocation rule. Suppose that the welfare-maximizing allocation rule OPT is γ -permeable for a subinstance-closed collection of binary, single-parameter feasibility problems Π . Then there exists a pricing rule that is weakly $(1, 0, \gamma^2)$ -balanced with respect to ALG and the canonical exchange-feasible sets $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$.*

As in the case of greedy we first describe the construction of the prices, and then we show that these prices are balanced.

Algorithm G.2. Pricing rule derived from OPT (parametrized by ALG).

Input: $z_i \in \{0, 1\}$, \mathbf{v} , $\mathbf{x} \in \mathcal{F}$ **Output:** $p_i(z_i \mid \mathbf{x})$

if $z_i = 0$ **then**
 // In this case the price is simply zero
 $p_i(z_i \mid \mathbf{x}) = 0$
else
 // First determine reference allocation and valuations
 $\mathbf{r}^{(0)} \leftarrow \text{ALG}(\mathbf{v})$
 $v_k^{(0)} \leftarrow v_k$ if $k \in \mathbf{r}^{(0)}$ and $v_k^{(0)} \leftarrow 0$ otherwise
 for $j \leftarrow 1$ **to** n **do**
 $\mathbf{r}^{(j)} \leftarrow \text{OPT}(\mathbf{v}^{(j-1)} \mid \mathbf{x}_{[j]})$
 $v_k^{(j)} \leftarrow v_k^{(j-1)} = v_k$ if $k \in \mathbf{r}^{(j)}$ and $v_k^{(j)} \leftarrow 0$ else
 end for
 // Now determine the price
 if $i \in \mathbf{r}^{(n)}$ **then**
 // If player i is part of the reference allocation he pays his valuation
 $p_i(z_i \mid \mathbf{x}) \leftarrow v_i$
 else
 // Otherwise he pays the critical value against the players in the
 // reference allocation
 $p_i(z_i \mid \mathbf{x}) \leftarrow \inf\{v'_i : i \in \text{OPT}(v'_i, \mathbf{v}_{-i}^{(n)} \mid \mathbf{x})\}$
 end if
end if
return $p_i(z_i \mid \mathbf{x})$

G.3.1. Construction of the prices. We define the price $p_i(z_i \mid \mathbf{x})$ for player $i \in N$ and outcome $z_i \in \{0, 1\}$ and arbitrary but fixed valuation profile \mathbf{v} and allocation $\mathbf{x} \in \mathcal{F}$ through Algorithm G.2. For this section, define $\text{OPT}(\mathbf{v} \mid \mathbf{x})$ as the allocation that results by padding the welfare-maximizing allocation for valuation profile \mathbf{v} over \mathcal{F}/\mathbf{x} with empty allocations.

As in the case of the greedy allocation rule, we again set $p_i(0 \mid \mathbf{x}) = 0$ and we compute $p_i(1 \mid \mathbf{x})$ via reference allocations and reference valuations. We again define the initial reference allocation as $\mathbf{r}^{(0)} = \text{ALG}(\mathbf{v})$ and the initial reference valuations by setting $v_j^{(0)} = v_j$ for $j \in \mathbf{r}^{(0)}$ and $v_j^{(0)} = 0$ otherwise. The subsequent reference allocations and valuations are defined recursively as $\mathbf{r}^{(i)} = \text{OPT}(v^{(i-1)} \mid \mathbf{x}_{[i]})$ and $v_j^{(i)} = v_j^{(i-1)} = v_j$ for $j \in \mathbf{r}^{(i)}$ and $v_j^{(i)} = 0$ otherwise. We then set $p_i(1 \mid \mathbf{x}) = v_i$ if $i \in \mathbf{r}^{(n)}$ and $p_i(1 \mid \mathbf{x}) = \inf\{v'_i \mid i \in \text{OPT}(v'_i, \mathbf{v}_{-i}^{(n)} \mid \mathbf{x})\}$ otherwise. Note that this definition immediately implies that $p_i(x_i \mid \mathbf{x}_{[i-1]}) = \mathbf{v}(\mathbf{r}^{(i-1)}) - \mathbf{v}(\mathbf{r}^{(i)})$.

By substituting all occurrences of n with $i - 1$ we obtain the formula for the price $p_i(1 \mid \mathbf{x}_{[i-1]})$. We use the shorthand $\tau^{\text{OPT}}(\mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]}) := \inf\{v'_i \mid i \in \text{OPT}(v'_i, \mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]})\}$.

G.3.2. Proof of Theorem G.8. We again proceed in two steps. We first show condition (a) (in Lemma G.9 below). Afterward we show that the permeability of OPT provides an upper bound on the permeability of GRD and that the critical prices with respect to OPT are not much higher than those with respect to GRD (in Lemmas G.10 and G.11). This allows us to bound condition (b) using the same machinery that we used in the previous section (in Lemma G.12).

LEMMA G.9. *Let ALG be any allocation rule. Suppose that the welfare-maximizing allocation rule OPT is γ -permeable for a subinstance-closed collection of problems of binary, single-parameter problems Π . Then the pricing rule described in Algorithm G.2 fulfills condition (a) of Definition 3.4 with $\alpha = 1$ with respect to allocation rule ALG and the canonical exchange-feasible sets $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$.*

Proof. Consider $\mathbf{x} \in \mathcal{F}$. We defined prices exactly so that $p_i(x_i \mid \mathbf{x}_{[i-1]}) = \mathbf{v}(\mathbf{r}^{(i-1)}) - \mathbf{v}(\mathbf{r}^{(i)})$. Therefore, using a telescoping-sum argument, we get

$$\sum_{i \in N} p_i(x_i \mid \mathbf{x}_{[i-1]}) = \mathbf{v}(\mathbf{r}^{(0)}) - \mathbf{v}(\mathbf{r}^{(n)}).$$

The claim now follows from the fact that $\mathbf{r}^{(0)} = \text{ALG}(\mathbf{v})$ and $\mathbf{r}^{(n)} \in \mathcal{F}_{\mathbf{x}}$. □

LEMMA G.10. *If the greedy allocation rule GRD is γ^{GRD} -permeable and the welfare-maximizing allocation rule OPT is γ^{OPT} -permeable for a subinstance-closed collection of binary, single-parameter problems Π , then $\gamma^{\text{OPT}} \geq \gamma^{\text{GRD}}$.*

Proof. We only have to show that for all $\mathbf{x} \in \mathcal{F}$, $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$ and all \mathbf{v} , we have

$$\sum_{i \in \mathbf{x}'} \tau_i^{\text{GRD}}(\mathbf{v}_{-i} \mid \mathbf{x}) \leq \gamma^{\text{OPT}} \cdot \mathbf{v}(\text{GRD}(\mathbf{v} \mid \mathbf{x})).$$

Let $\mathbf{y} = \mathbf{x}' \cap \text{GRD}(\mathbf{v} \mid \mathbf{x})$. Observe that because $\mathbf{y} \subseteq \text{GRD}(\mathbf{v} \mid \mathbf{x})$, we have $\text{GRD}(\mathbf{v} \mid \mathbf{x} \cup \mathbf{y}) = \text{GRD}(\mathbf{v} \mid \mathbf{x})$. Define a valuation profile \mathbf{v}' by setting $v'_i = v_i$ for all $i \in \mathbf{x}' \setminus \mathbf{y}$ and $v'_i = 0$ otherwise. By loser independence of greedy, $\text{GRD}(\mathbf{v} \mid \mathbf{x} \cup \mathbf{y}) = \text{GRD}(\mathbf{v}' \mid \mathbf{x} \cup \mathbf{y})$. Furthermore, $\text{OPT}(\mathbf{v}' \mid \mathbf{x} \cup \mathbf{y}) = \text{GRD}(\mathbf{v}' \mid \mathbf{x} \cup \mathbf{y})$.

We claim that for $i \in \mathbf{x}' \setminus \mathbf{y}$ we have

$$\tau_i^{\text{GRD}}(\mathbf{v}_{-i} \mid \mathbf{x}) \leq \tau_i^{\text{GRD}}(\mathbf{v}'_{-i} \mid \mathbf{x} \cup \mathbf{y}) \leq \tau_i^{\text{OPT}}(\mathbf{v}'_{-i} \mid \mathbf{x} \cup \mathbf{y}) .$$

For the first inequality we use that $\tau_i^{\text{GRD}}(\mathbf{v}_{-i} \mid \mathbf{x})$ is the value v_j of some j that has to be outbid by player i . By fixing another set \mathbf{y} , this value can only go up because the options are limited further. Also, when fixing $\mathbf{x} \cup \mathbf{y}$, replacing the valuations by \mathbf{v}' has no influence because the set of players that are selected remains unchanged.

The second inequality holds because under \mathbf{v}' player i in order to win when we hold $\mathbf{x} \cup \mathbf{y}$ fixed has to force some subset $\mathbf{z} \subseteq \text{GRD}(\mathbf{v}'_{-i} \mid \mathbf{x} \cup \mathbf{y}) = \text{OPT}(\mathbf{v}'_{-i} \mid \mathbf{x} \cup \mathbf{y})$ out of the solution. Under OPT his payment is the sum of the respective players' valuations; under GRD it is just the highest valuation of any such player.

On the other hand, for players $i \in \mathbf{y}$, because $\mathbf{y} \subseteq \text{GRD}(\mathbf{v} \mid \mathbf{x})$, the greedy critical value $\tau_i^{\text{GRD}}(\mathbf{v}_{-i} \mid \mathbf{x})$ is at most v_i .

Using these two bounds on $\tau_i^{\text{GRD}}(\mathbf{v}_{-i} \mid \mathbf{x})$ we obtain

$$\begin{aligned} \sum_{i \in \mathbf{x}'} \tau_i^{\text{GRD}}(\mathbf{v}_{-i} \mid \mathbf{x}) &\leq \sum_{i \in \mathbf{y}} v_i + \sum_{i \in \mathbf{x}' \setminus \mathbf{y}} \tau_i^{\text{GRD}}(\mathbf{v}_{-i} \mid \mathbf{x}) \\ &\leq \sum_{i \in \mathbf{y}} v_i + \sum_{i \in \mathbf{x}' \setminus \mathbf{y}} \tau_i^{\text{OPT}}(\mathbf{v}'_{-i} \mid \mathbf{x} \cup \mathbf{y}) \\ &\leq \sum_{i \in \mathbf{y}} v_i + \gamma^{\text{OPT}} \cdot \mathbf{v}'(\text{OPT}(\mathbf{v}' \mid \mathbf{x} \cup \mathbf{y})) \\ &= \sum_{i \in \mathbf{y}} v_i + \gamma^{\text{OPT}} \cdot \mathbf{v}'(\text{GRD}(\mathbf{v}' \mid \mathbf{x} \cup \mathbf{y})) \end{aligned}$$

$$\begin{aligned} &\leq \gamma^{\text{OPT}} \cdot \sum_{i \in \mathbf{y}} v_i + \gamma^{\text{OPT}} \cdot \mathbf{v}'(\text{GRD}(\mathbf{v}' \mid \mathbf{x} \cup \mathbf{y})) \\ &= \gamma^{\text{OPT}} \cdot \mathbf{v}(\text{OPT}(\mathbf{v} \mid \mathbf{x})), \end{aligned}$$

where the third inequality uses γ^{OPT} -permeability of the welfare-maximizing allocation rule OPT in the subinstance in which we hold $\mathbf{x} \cup \mathbf{y}$ fixed, the subsequent equality holds by the definition of \mathbf{v}' , the fourth inequality uses that $\gamma^{\text{OPT}} \geq 1$, and the final equality holds by the definition of \mathbf{y} and \mathbf{v}' . \square

LEMMA G.11. *If the greedy allocation rule GRD is γ^{GRD} -permeable for a subinstance-closed collection of binary, single-parameter problems Π , then*

$$\tau_i^{\text{OPT}}(\mathbf{v}^{(i-1)} \mid \mathbf{x}_{[i-1]}) \leq \gamma^{\text{GRD}} \cdot \tau_i^{\text{GRD}}(\mathbf{v}^{(i-1)} \mid \mathbf{x}_{[i-1]}).$$

Proof. Note that for valuations $\mathbf{v}^{(i-1)}$ both GRD and OPT over \mathcal{F}/\mathbf{x} return the same set of players. The same is true if we drop any player j from $\mathbf{v}^{(i-1)}$. Dropping player i we can define $\mathbf{y} = \text{GRD}(\mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]}) = \text{OPT}(\mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]})$.

Suppose player i bids $b'_i = \tau_i^{\text{GRD}}(\mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]}) + \epsilon$. Then $i \in \text{GRD}(b'_i, \mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]})$. The addition of player i causes the removal of a (possibly empty) subset of players $\mathbf{z} \subseteq \mathbf{y}$. That is, $\text{GRD}(b'_i, \mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]}) = (\mathbf{y} \setminus \mathbf{z}) \cup \{i\}$.

Define valuations \mathbf{v}' by setting $v'_i = b'_i$, $v'_j = v_j^{(i-1)}$ for $j \in \mathbf{z}$, and $v'_j = 0$ for every other player j . Consider the subinstance in which we hold $\mathbf{y} \setminus \mathbf{z}$ fixed. Since both player i and the set of players \mathbf{z} are feasible extensions we can apply γ^{GRD} -permeability of the greedy allocation rule to obtain

$$|\mathbf{z}| \cdot b'_i = \sum_{j \in \mathbf{z}} \tau_j^{\text{GRD}}(\mathbf{v}' \mid \mathbf{y} \setminus \mathbf{z}) \leq \gamma^{\text{GRD}} \cdot \mathbf{v}'(\text{GRD}(\mathbf{v}' \mid \mathbf{y} \setminus \mathbf{z})) = \gamma^{\text{GRD}} \cdot b'_i,$$

and therefore $|\mathbf{z}| \leq \gamma^{\text{GRD}}$.

The final step is now to observe that the critical value $\tau_i^{\text{OPT}}(\mathbf{v}^{(i-1)} \mid \mathbf{x}_{[i-1]})$ of player i under the welfare-maximizing allocation rule is at most $|T|$ times the critical value $p_i^{\text{GRD}}(\mathbf{v}^{(i-1)} \mid \mathbf{x}_{[i-1]})$ of player i under the greedy allocation rule. To see this let $\mathbf{z}' \subseteq \mathbf{y}$ be the set with the smallest value such that $(\mathbf{y} \setminus \mathbf{z}') \cup \{i\} \in \mathcal{F}$. Then,

$$\tau_i^{\text{OPT}}(\mathbf{v}^{(i-1)} \mid \mathbf{x}_{[i-1]}) = \sum_{j \in \mathbf{z}'} v_j \leq \sum_{j \in \mathbf{z}} v_j \leq |T| \cdot \max_{j \in \mathbf{z}} v_j = |\mathbf{z}| \cdot \tau_i^{\text{GRD}}(\mathbf{v}^{(i-1)} \mid \mathbf{x}_{[i-1]}),$$

where we used that \mathbf{z} is some subset of \mathbf{y} such that $(\mathbf{y} \setminus \mathbf{z}) \cup \{i\} \in \mathcal{F}$ and so its combined value can only be larger than that of \mathbf{z}' . \square

LEMMA G.12. *Let ALG be any allocation rule. Suppose that the welfare-maximizing allocation rule OPT is γ -permeable for a subinstance-closed collection of binary, single-parameter problems Π . Then the pricing rule described in Algorithm G.2 fulfills condition (b) of Definition 3.4 with $\beta_1 = 0$ and $\beta_2 = \gamma^2$ with respect to allocation rule ALG and the canonical exchange-feasible sets $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x} \in X}$.*

Proof. Consider an arbitrary allocation $\mathbf{x} \in \mathcal{F}$ and feasible extension $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$. We want to show that

$$\sum_{i \in N} p_i(x'_i \mid \mathbf{x}_{[i-1]}) \leq (\gamma^{\text{OPT}})^2 \cdot \mathbf{v}(\text{ALG}(\mathbf{v})).$$

We will again show this claim through layering. This time, however, we need the layering to be arbitrarily fine-grained. We will specify the granularity by $\epsilon > 0$. For each $j \in \mathbb{N}$, let S^j denote the set of players with value at least $j \cdot \epsilon$. Let T^j denote the set of players with $x'_i = 1$ that see a price $p_i^{\text{OPT}}(x'_i \mid \mathbf{x}_{[i-1]})$ of at least $\gamma^{\text{GRD}} \cdot j \cdot \epsilon$.

For any fixed $j \in \mathbb{N}$ we now bound $|T^j|$ using Lemma G.6. We set $A_t = x_{[t]}$, $B_t = \mathbf{r}^{(t)} \cap S^j$, $C = T^j$. Note that $C \cup A_n \in \mathcal{F}$. We claim that for every $i \in C$ we have $i \in B_t$ or $\{i\} \cup B_t \cup A_t \notin \mathcal{F}$ for $t = i - 1$. To see this, consider the two options for how $p_i^{\text{OPT}}(x'_i \mid \mathbf{x}_{[i-1]})$ can be set. If $i \in \mathbf{r}^{(i-1)}$, then $p_i^{\text{OPT}}(x'_i \mid \mathbf{x}_{[i-1]}) = v_i$. As by definition $p_i^{\text{OPT}}(x'_i \mid \mathbf{x}_{[i-1]}) \geq \gamma^{\text{GRD}} \cdot j \cdot \epsilon \geq j \cdot \epsilon$, this implies $v_i \geq j \cdot \epsilon$ and so $i \in B_{i-1} = \mathbf{r}^{(i-1)} \cap S^j$. Otherwise, if $i \notin \mathbf{r}^{(i-1)}$, then $p_i^{\text{OPT}}(x'_i \mid \mathbf{x}_{[i-1]}) = \tau_i^{\text{OPT}}(\mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]})$. In this case, we can apply Lemma G.11 to get

$$p_i^{\text{OPT}}(x'_i \mid \mathbf{x}_{[i-1]}) = \tau_i^{\text{OPT}}(\mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]}) \leq \gamma^{\text{GRD}} \cdot \tau_i^{\text{GRD}}(\mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]}).$$

So, we know that $\tau_i^{\text{GRD}}(\mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]}) \geq j \cdot \epsilon$ and therefore $\{i\} \cup \mathbf{x}_{[i-1]} \cup (\mathbf{r}^{(i-1)} \cap Y^j) \notin \mathcal{F}$.

So, by Lemma G.6, we get $|T^j| \leq \gamma^{\text{GRD}} \cdot |\mathbf{r}^{(0)} \cap Y^j|$.

We conclude that

$$\begin{aligned} \sum_{i \in \mathbf{x}'} p_i^{\text{OPT}}(x'_i \mid \mathbf{x}_{[i-1]}) &\leq n\epsilon + \sum_{i \in \mathbf{x}'} \sum_{j=1}^{\infty} \mathbb{1}_{i \in T^j} \cdot \gamma^{\text{GRD}} \cdot \epsilon \\ &= n\epsilon + \gamma^{\text{GRD}} \cdot \sum_{j=1}^{\infty} |T^j| \cdot \epsilon \\ &\leq n\epsilon + \gamma^{\text{GRD}} \cdot \sum_{j=1}^{\infty} \gamma^{\text{GRD}} \cdot |\mathbf{r}^{(0)} \cap Y^j| \cdot \epsilon \\ &\leq (\gamma^{\text{GRD}})^2 \cdot \mathbf{v}(\text{ALG}(\mathbf{v})) + 2n\epsilon. \end{aligned}$$

As this argument holds for any $\epsilon > 0$, we also have $\sum_{i \in \mathbf{x}'} p_i^{\text{OPT}}(x'_i \mid \mathbf{x}_{[i-1]}) \leq (\gamma^{\text{GRD}})^2 \cdot \mathbf{v}(\text{ALG}(\mathbf{v}))$. By Lemma G.10, $\gamma^{\text{GRD}} \leq \gamma^{\text{OPT}}$ and the claim follows. \square

Remark G.13. When the prices defined by Algorithm G.2 are nondecreasing then γ -permeability of OPT immediately implies that $\tau_i^{\text{GRD}}(\mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}_{[i-1]}) \leq \tau_i^{\text{GRD}}(\mathbf{v}_{-i}^{(i-1)} \mid \mathbf{x}) \leq \gamma \cdot \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))$ implying that the pricing rule is (α, β) -balanced with $\beta = \gamma$. In this case Theorem G.8 can be strengthened to show the existence of a $(1, \gamma)$ -balanced pricing rule.

Appendix H. Proof of Theorem 6.6. We prove this theorem constructively. The high-level idea behind our construction is to use additive supporting valuations for ALG to define static item prices (generalizing the approach in [25]) and let these prices play the role of the bids in Definition 6.5.

To construct a pricing rule for a setting with n players and m items and known valuations \mathbf{v} we proceed as follows. Let \mathbf{a} denote the additive supporting valuation for $\text{ALG}(\mathbf{v})$. To define $p_i^{\mathbf{y}}(x_i \mid \mathbf{y})$ consider a $(n + 1)$ -first bidder bidding against \mathbf{a} . We set

$$(H.1) \quad p_i^{\mathbf{y}}(x_i \mid \mathbf{y}) = \begin{cases} \tau_{n+1}(x_i, \mathbf{a}) & \text{for } x_i \subseteq [m] \setminus \mathbf{y}, \\ \infty & \text{otherwise} \end{cases}$$

We begin by showing that ALG is a λ -approximation to OPT. We use that for all $\mathbf{v} \in V$ Definition 6.5 with $\mathbf{b} = \mathbf{0} \in B$ implies that

$$\begin{aligned} \mathbf{v}(\text{ALG}(\mathbf{v})) &= \sum_{i \in N} (v_i(\text{ALG}_i(\mathbf{v})) - \tau_i(\text{ALG}_i(\mathbf{v}), \mathbf{b}_{\cdot i})) \\ &\geq \lambda \cdot \mathbf{v}(\text{OPT}(\mathbf{v})) - \mu \cdot \sum_{i \in N} P_i(\mathbf{b}) = \lambda \cdot \mathbf{v}(\text{OPT}(\mathbf{v})). \end{aligned}$$

Next we show that the prices defined in (H.1) are balanced with respect to ALG . We first show that they satisfy property (a) with $\alpha = 1$. To this end fix n, m , and $\mathbf{v} \in V$ and consider any $x \in \mathcal{F}$. Let $\mathbf{a} \in B$ be the additive supporting valuation for $\text{ALG}(\mathbf{v})$ used in the definition of the prices. For each item j define $\hat{a}_j = \max_i a_{i,j}$. Then,

$$\begin{aligned} \sum_{i \in N} p_i^{\mathbf{y}}(x_i \mid \mathbf{x}_{[i-1]}) &= \sum_{i \in N} \tau_{n+1}(x_i, \mathbf{a}) \\ &= \sum_{i \in N} \sum_{j \in x_i} a_j \\ &= \sum_{j \in M} a_j - \sum_{j \notin \mathbf{x}} a_j \geq \mathbf{v}(\text{ALG}(\mathbf{v})) - \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_x)), \end{aligned}$$

where in the last step we used that $(\text{ALG}(\mathbf{v}) \setminus x) \in \mathcal{F}_{\mathbf{x}}$ and $\mathbf{a}(\text{ALG}(\mathbf{v}) \setminus \mathbf{x}) \leq \mathbf{v}(\text{ALG}(\mathbf{v}) \setminus x) \leq \mathbf{v}(\text{OPT}(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))$.

Next we show that the prices satisfy property (b) with $\beta = \mu$. Fix n, m , and $\mathbf{v} \in V$. Consider $\mathbf{x} \in \mathcal{F}$ and $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$. As before let $\mathbf{a} \in B$ denote the additive supporting valuation for $\text{ALG}(\mathbf{v})$. To bound the prices we will invoke Definition 6.5 for a different setting with $n' = 2n$ players and $m' \leq m$ items. For the valuations in this new setting we use that ALG is onto and scale free. Namely, we set the valuations to $\mathbf{v}'' = \epsilon \cdot \mathbf{v}'$ where \mathbf{v}' is such that $\text{ALG}_i(\epsilon \cdot \mathbf{v}') = x'_i$ for $i = 1, \dots, n$ and $\text{ALG}_i(\epsilon \cdot \mathbf{v}') = \emptyset$ for $i = n + 1, \dots, 2n$. The bid of player of player $i = 1, \dots, n$ is $b_i = \mathbf{0}$, and the bid of player $i = n + 1, \dots, 2n$ is the additive supporting valuation of bidder $n - i$ in the original setting, restricted to the items $j \notin \mathbf{x}$.

By outcome smoothness,

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

Since this holds for every $\epsilon > 0$ we can use that $\sum_{i \in N} v_i''(x'_i) \rightarrow 0$ as $\epsilon \rightarrow 0$ and $\mathbf{v}''(\text{OPT}(\mathbf{v}'')) \rightarrow 0$ as $\epsilon \rightarrow 0$ to obtain

$$\sum_{i \in N} \tau_i(\text{ALG}_i(\mathbf{v}''), \mathbf{b}_{\cdot i}) \leq \mu \cdot \sum_{i \in N} P_i(\mathbf{b}).$$

The left-hand side of this inequality is equal to $\sum_{i=1}^n p_i^{\mathbf{y}}(x'_i \mid \mathbf{x}_{[i-1]})$. The right-hand side is at most $\mu \cdot \mathbf{v}(\text{ALG}(\mathbf{v}) \setminus \mathbf{x}) \leq \mu \cdot \mathbf{v}(\text{OPT}(\mathbf{v}), \mathcal{F}_{\mathbf{x}})$. The claim follows.

REFERENCES

- [1] 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.
- [2] I. ABRAHAM, M. BABAILOFF, S. DUGHMI, AND T. ROUGHGARDEN, *Combinatorial auctions with restricted complements*, in Proceedings of the 13th ACM Conference on Electronic Commerce, 2012, pp. 3–16.

- [3] S. ALAEI, *Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers*, SIAM J. Comput., 43 (2014), pp. 930–972.
- [4] 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.
- [5] S. ALAEI, M. HAJIAGHAYI, AND V. LIAGHAT, *The online stochastic generalized assignment problem*, in Proceedings of the 16th/17th APPROX-RANDOM Workshop, 2013, pp. 11–25.
- [6] S. ALAEI, J. D. HARTLINE, R. NIAZADEH, E. POUNTOURAKIS, AND Y. YUAN, *Optimal auctions vs. anonymous pricing*, in Proceedings of the 56th IEEE Symposium on Foundations of Computer Science, 2015, pp. 1446–1463.
- [7] 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.
- [8] M. BALCAN, A. BLUM, AND Y. MANSOUR, *Item pricing for revenue maximization*, in Proceedings of the 9th ACM Conference on Electronic Commerce, 2008, pp. 50–59.
- [9] 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.
- [10] 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.
- [11] 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.
- [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] 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.
- [14] G. CHRISTODOULOU, A. KOVÁCS, AND M. SCHAPIRA, *Bayesian combinatorial auctions*, J. ACM, 63 (2016), p. 11.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] S. DOBZINSKI, *Two randomized mechanisms for combinatorial auctions*, in Proceedings of the 10th/11th APPROX-RANDOM Workshop, 2007, pp. 89–103.
- [19] 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.
- [20] 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.
- [21] P. DÜTTING AND R. KLEINBERG, *Polymatroid prophet inequalities*, in Proceedings of the 23rd European Symposium on Algorithms, 2015, pp. 437–449.
- [22] 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.
- [23] M. FELDMAN, A. FIAT, AND A. ROYTMAN, *Makespan minimization via posted prices*, in Proceedings of the 18th ACM Conference on Economics and Computation, 2017, pp. 405–422.
- [24] 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.
- [25] 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.
- [26] 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.
- [27] M. FELDMAN, O. SVENSSON, AND R. ZENKLUSEN, *A simple $o(\log \log(\text{rank}))$ -competitive algo-*

- rithm for the matroid secretary problem*, in Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms, 2015, pp. 1189–1201.
- [28] 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.
 - [29] 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.
 - [30] 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.
 - [31] J. D. HARTLINE, D. HOY, AND S. TAGGART, *Price of anarchy for auction revenue*, in Proceedings of the 15th ACM Conference on Economics and Computation, 2014, pp. 693–710.
 - [32] O. H. IBARRA AND C. E. KIM, *Fast approximation algorithms for the knapsack and sum of subset problems*, J. ACM, 22 (1975), pp. 463–468.
 - [33] R. KLEINBERG AND S. M. WEINBERG, *Matroid prophet inequalities*, in Proceedings of the 44th ACM Symposium on Theory of Computing, 2012, pp. 123–136.
 - [34] U. KRENGEL AND L. SUCHESTON, *Semiamarts and finite values*, Bull. Amer. Math. Soc., 83 (1977), pp. 745–747.
 - [35] U. KRENGEL AND L. SUCHESTON, *On semiamarts, amarts, and processes with finite value*, Adv. Probab. Relat. Topics, 4 (1978), pp. 197–266.
 - [36] J. LEE, M. SVIRIDENKO, AND J. VONDRAK, *Submodular maximization over multiple matroids via generalized exchange properties*, Math. Oper. Res., 35 (2010), pp. 795–806.
 - [37] B. LEHMANN, D. LEHMANN, AND N. NISAN, *Combinatorial auctions with decreasing marginal utilities*, Games Econom. Behav., 55 (2006), pp. 270–296.
 - [38] B. LUCIER AND A. BORODIN, *Price of anarchy for greedy auctions*, in Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms, 2010, pp. 537–553.
 - [39] 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.
 - [40] 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.
 - [41] T. ROUGHGARDEN, *Intrinsic robustness of the price of anarchy*, J. ACM, 62 (2015), p. 32.
 - [42] A. RUBINSTEIN, *Beyond matroids: Secretary problem and prophet inequality with general constraints*, in Proceedings of the 48th ACM Symposium on Theory of Computing, 2016, pp. 324–332.
 - [43] A. RUBINSTEIN AND S. SINGLA, *Combinatorial prophet inequalities*, in Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms, 2017, pp. 1671–1687.
 - [44] E. SAMUEL-CAHN, *Comparison of threshold stop rules and maximum for independent nonnegative random variables*, Ann. Probab., 12 (1984), pp. 1213–1216.
 - [45] V. SYRGKANIS AND É. TARDOS, *Composable and efficient mechanisms*, in Proceedings of the 45th ACM Symposium on Theory of Computing, 2013, pp. 211–220.
 - [46] 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.