

Extreme-Value Theorems for Optimal Multidimensional Pricing

Yang Cai*
EECS, MIT

Constantinos Daskalakis†
EECS, MIT

November 3, 2013

Abstract

We provide a near-optimal, computationally efficient algorithm for the unit-demand pricing problem, where a seller wants to price n items to optimize revenue against a unit-demand buyer whose values for the items are independently drawn from known distributions. For any chosen accuracy $\epsilon > 0$ and item values bounded in $[0, 1]$, our algorithm achieves revenue that is optimal up to an additive error of at most ϵ , in polynomial time. For values sampled from Monotone Hazard Rate (MHR) distributions, we achieve a $(1 - \epsilon)$ -fraction of the optimal revenue in polynomial time, while for values sampled from regular distributions the same revenue guarantees are achieved in quasi-polynomial time.

Our algorithm for bounded distributions applies probabilistic techniques to understand the statistical properties of revenue distributions, obtaining a reduction in the search space of the algorithm through dynamic programming. Adapting this approach to MHR and regular distributions requires the proof of novel extreme value theorems for such distributions.

As a byproduct, our techniques establish structural properties of approximately-optimal and near-optimal solutions. We show that, when the buyer's values are independently distributed according to MHR distributions, pricing all items at the same price achieves a constant fraction of the optimal revenue. Moreover, for all $\epsilon > 0$, at most $g(1/\epsilon)$ distinct prices suffice to obtain a $(1 - \epsilon)$ -fraction of the optimal revenue, where $g(1/\epsilon)$ is a quadratic function of $1/\epsilon$ that does not depend on the number of items. Similarly, for all $\epsilon > 0$ and $n > 0$, at most $g(1/\epsilon \cdot \log n)$ distinct prices suffice if the values are independently distributed according to regular distributions, where $g(\cdot)$ is a polynomial function. Finally, if the values are i.i.d. from some MHR distribution, we show that, as long as the number of items is a sufficiently large function of $1/\epsilon$, a single price suffices to achieve a $(1 - \epsilon)$ -fraction of the optimal revenue.

*Supported by NSF Awards CCF-0953960 (CAREER) and CCF-1101491.

†Supported by a Sloan Foundation Fellowship, a Microsoft Research Faculty Fellowship, and NSF Awards CCF-0953960 (CAREER) and CCF-1101491.

1 Introduction

We study the following pricing problem. A seller has n items to sell to a buyer who is looking to buy a single item. The seller wants to maximize profit from the sale, leveraging stochastic knowledge she has about the buyer to achieve this goal. In particular, we assume that the seller has access to a distribution \mathcal{F} from which the values (v_1, \dots, v_n) of the buyer for the items are drawn. Given this information, the seller wants to compute prices p_1, \dots, p_n for the items to maximize her revenue, assuming that the buyer is *quasi-linear*—i.e. will buy the item i maximizing $v_i - p_i$, as long as this difference is positive. That is, the seller’s expected revenue from a price vector $P = (p_1, \dots, p_n)$ is

$$\mathcal{R}_P = \sum_{i=1}^n p_i \cdot \Pr \left[(i = \arg \max \{v_j - p_j\}) \wedge (v_i - p_i \geq 0) \right], \quad (1)$$

where we assume that the $\arg \max$ breaks ties in favor of a single item, when there are multiple maximizers. A more sophisticated seller could try to improve her revenue by pricing lotteries over items, that is also price randomized allocations of items [BCKW10], albeit this may be less natural than item pricing, and we will not study it extensively in this paper.

While our problem has a simple statement, it exhibits rich behavior depending on the nature of \mathcal{F} . For example, if \mathcal{F} assigns the same value to all the items with probability 1, i.e. when the buyer always values all items equally, the problem becomes *single-dimensional*. In this setting, it is clear that lotteries do not improve the revenue and that the optimal price vector can assign the same price to all the items. This observation is a special case of the more general, celebrated result of Myerson [Mye81] on optimal mechanism design, i.e. the multi-buyer version of our problem, and generalizations thereof. Myerson’s result provides a closed-form solution to the multi-buyer problem in a single sweep that covers many settings, but only works under the same limiting assumption that every buyer is single-dimensional, i.e. receives the same value from all the items. (More generally, every buyer receives the same value from all outcomes of the mechanism that provide her service.)

Following Myerson, a large body of research in both Economics and Engineering has been devoted to extending his result to the *multi-dimensional setting*, where the buyers’ values come from general distributions. And, while there has been sporadic progress (see survey [MV07] and its references), an optimal multi-dimensional mechanism, generalizing Myerson’s result, does not seem to be in sight. Indeed, there is not even an optimal solution known for the single-buyer item pricing problem. Even the ostensibly easier version of that problem, where the values of the buyer for the items are independent and supported on a set of cardinality 2 is unresolved.¹ Our main contribution in this paper is to develop near-optimal polynomial-time algorithms for this problem, when the buyer’s values for the items are independent.

1.1 Main Results

We partition our results into algorithmic and structural. The former provide efficient algorithmic procedures for computing near-optimal price vectors. The latter shed light into the structure of optimal solutions.

¹Incidentally, the problem is trickier than it originally seems, and various intuitive properties that one would expect from the optimal solution fail to hold. See Appendix A for an interesting example.

Algorithmic Results. Previous work on the item pricing problem has provided constant factor approximation algorithms. The best known polynomial-time algorithm obtains revenue that is at least $1/2$ of the revenue of the optimal price vector [CHK07, CHMS10]. We discuss these approaches in Section 1.3, also noting that they are limited to constant factor approximations. We are aiming instead for item pricing mechanisms that come arbitrarily close to the optimal revenue, obtaining the following results. Their proofs are overviewed in Sections 4 through 9, while complete details are provided in the appendix.

Theorem 1 (Main Algorithmic Result: Additive PTAS for Bounded Distributions). *Suppose that the values of the buyer for n items are independent and normalized to lie in $[0, 1]$. Then, for all $\epsilon > 0$, there exists an algorithm that computes a price vector whose revenue is within an additive ϵ of optimal, and whose running time is polynomial in $n^{\frac{\log^3 1/\epsilon}{\epsilon^4}}$.*

Theorem 2 (General Algorithm). *Suppose that the values of the buyer for n items are independent and supported on some interval $[u_{\min}, r \cdot u_{\min}]$ for some $u_{\min} > 0$ and $r \geq 1$. Then, for all $\epsilon > 0$, there is an algorithm that computes a price vector whose revenue is at least a $(1 - \epsilon)$ -fraction of the optimal revenue, and whose running time is polynomial in $\max \left\{ n^{\log^{11} r \cdot \log \log r}, n^{\frac{\log^3 r \cdot \log \frac{1}{\epsilon}}{\epsilon^8}} \right\}$.²*

Theorem 3 (Multiplicative PTAS for MHR Distributions). *There is a Polynomial-Time Approximation Scheme³ for computing an optimal price vector, when the values of the buyer are independently drawn from Monotone Hazard Rate distributions.⁴*

For any accuracy $\epsilon > 0$, the algorithm runs in time polynomial in $n^{\frac{1}{\epsilon^7}}$, and outputs a price vector whose revenue is at least a $(1 - \epsilon)$ -fraction of the optimal revenue, where n is the number of items.

Theorem 4 (Multiplicative Quasi-PTAS for Regular Distributions). *There is a Quasi-Polynomial-Time Approximation Scheme⁵ for computing an optimal price vector, when the values of the buyer are independent and drawn from regular distributions.⁶*

For any accuracy $\epsilon > 0$, the algorithm runs in time polynomial in $\max \left\{ n^{\log^{11} \frac{n}{\epsilon} \cdot \log \log \frac{n}{\epsilon}}, n^{\frac{\log^3 \frac{n}{\epsilon} \cdot \log \frac{1}{\epsilon}}{\epsilon^8}} \right\}$, and outputs a price vector whose revenue is at least a $(1 - \epsilon)$ -fraction of the optimal revenue, where n is the number of items.

²We note that a natural approach for computing approximately optimal price vectors is to discretize the domain of price vectors and show that searching over the discretized domain suffices for approximating the optimal revenue. However, a straightforward application of the discretizations proposed by Nisan [CHK07] and Hartline and Koltun [HK05] to our problem would result in running time of $(\frac{1}{\epsilon} \log r)^{O(n)}$. The purpose of our theorem is to remove the exponential dependence of the running time on the number of items n .

³A Polynomial-Time Approximation Scheme (PTAS) is a family of algorithms $\{\mathcal{A}_\epsilon\}_\epsilon$, indexed by the accuracy parameter $\epsilon > 0$, such that for every fixed $\epsilon > 0$, \mathcal{A}_ϵ runs in time polynomial in the size of its input. See Section 2 for a formal definition.

⁴Monotone Hazard Rate (MHR) distributions are a commonly studied class of distributions that contain such familiar distributions as the Uniform, Gaussian and Exponential distributions. See Section 2 for a formal definition.

⁵A Quasi-Polynomial-Time Approximation Scheme (Quasi-PTAS) is a family of algorithms $\{\mathcal{A}_\epsilon\}_\epsilon$, indexed by the accuracy parameter $\epsilon > 0$, such that for every fixed $\epsilon > 0$, \mathcal{A}_ϵ runs in time quasi-polynomial in the size of its input. See Section 2 for formal definition.

⁶Regular distributions are another widely studied class of distributions that contain MHR distributions. See Section 2 for a formal definition.

Discussion of Algorithmic Results. Prior to our work there were no optimal algorithms known for multi-dimensional auction problems without special structure. In particular, only constant factor approximation algorithms were known for the item pricing problems addressed by Theorems 1 through 4. (For an extensive discussion of related work, we refer the reader to Section 1.3.) Our results are the first to obtain near-optimal solutions for these problems in polynomial time. We view the main contribution of our results not to be the practicality of our algorithms, but establishing that there is no lingering constant in-approximability results for item pricing. In particular, our results show that, for any desired accuracy $\epsilon > 0$, there are polynomial-time algorithms that compute ϵ -optimal solutions. Complemented with the NP-hardness result for the item pricing problem discussed in Section 1.3, what is left open by our work is obtaining faster near-optimal algorithms.

Structural Results. Our algorithms are obtained by studying the distribution of the optimal revenue, as a function of the buyer’s values (which are random) and the optimal price vector (which is unknown), as overviewed in Section 1.2. As a byproduct of our techniques, we deduce the following structural properties of optimal solutions, whose proofs are given in Appendix I. Theorem 5 states that, when the values are independently distributed according to monotone hazard rate distributions, then pricing all items at the same price guarantees a constant fraction of the optimal revenue. Theorem 6 generalizes this to showing that only the desired approximation ϵ dictates the number of distinct prices that are necessary to achieve a $(1 - \epsilon)$ -fraction of the optimal revenue, and not the number of items or the size of the support of the distributions, as long as they are monotone hazard rate. Theorem 7 generalizes this result to a mild dependence on n for regular distributions.

Theorem 5 (Structural 1 (MHR): Constant Factor Approximation from a Single Price). *If the buyer’s values for the items are independently distributed according to MHR distributions, there exists a price p such that pricing all items at p guarantees a constant fraction of the optimal revenue. Price p can be computed efficiently from the value distributions.*

Theorem 6 (Structural 2 (MHR): A Constant Number of Distinct Prices Suffice for Near-Optimal Revenue). *There exists a quasi-quadratic⁷ function $g(\cdot)$ such that, for all $\epsilon > 0$ and all $n > 0$, $g(1/\epsilon)$ distinct prices suffice to achieve a $(1 - \epsilon)$ -fraction of the optimal revenue, when the buyer’s values for the n items are independently distributed according to MHR distributions. These distinct prices can be computed efficiently from the value distributions.*

Theorem 7 (Structural 3 (Regular): A Polylogarithmic Number of Distinct Prices Suffice for Near-Optimal Revenue). *There exists a polynomial function $g(\cdot)$ such that, for all $\epsilon > 0$ and $n > 0$, $g(1/\epsilon \cdot \log n)$ distinct prices suffice to achieve a $(1 - \epsilon)$ -fraction of the optimal revenue, when the buyer’s values for the n items are independently distributed according to regular distributions. These prices can be computed efficiently from the value distributions.*

Finally, it seems intuitive that, when the value distributions are not widely different, a single price might suffice for extracting a $(1 - \epsilon)$ -fraction of the optimal revenue, as long as there is a sufficient number of items for sale. We show such a result for the case where the buyer’s values are i.i.d. according to a MHR distribution. See Appendix I.1 for the proof of this theorem.

⁷A function $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is *quasi-quadratic* iff it satisfies $g(x) = O(x^2 \log^c x)$, for some absolute constant $c > 0$. For the meaning of the $O(\cdot)$ notation please refer to Section 2.

Theorem 8 (Structural 4 (i.i.d. MHR): A Single Price Suffices for Near-Optimal Revenue). *There is a function $g(\cdot)$ such that, for all $\epsilon > 0$, if the number of items is larger than $g(1/\epsilon)$ then pricing all the items at the same price obtains a $(1 - \epsilon)$ -fraction of the optimal revenue, if the buyer’s values are i.i.d. according to a MHR distribution.*

Extreme Value Theorems. Establishing the above structural properties relies on understanding the tails of MHR and regular distributions. For this purpose, we develop *extreme value theorems* for these classes of distributions. We state our extreme value theorems informally below, referring the reader to Theorems 19 and 21 (in Sections 8 and 9 respectively) for formal statements.

Informal Theorem 9 (Extreme Values of MHR Distributions). *Let X_1, \dots, X_n be a collection of independent random variables whose distributions are MHR, and let $Z = \max_i X_i$. Then, for all ϵ sufficiently small, at least a $(1 - \epsilon)$ -fraction of $\mathbb{E}[Z]$ is contributed to by the event $Z \leq O(\log_2 \frac{1}{\epsilon}) \cdot \mathbb{E}[Z]$.*

Informal Theorem 10 (Extreme Values of Regular Distributions). *Let X_1, \dots, X_n be a collection of independent random variables whose distributions are regular, and let $Z = \max_i X_i$. Then the tails of Z are eventually not fatter than the tails of the equal revenue distribution.⁸*

Bounding the size of the tail of the *maximum* of n independent random variables, which are MHR or regular respectively, is instrumental in establishing the following truncation property: restricting all the item prices into an interval of the form $[\alpha, \text{poly}(1/\epsilon)\alpha]$ in the MHR case, and $[\alpha, \text{poly}(n, 1/\epsilon)\alpha]$ in the regular case, for some α that depends on the value distributions, only loses a fraction of ϵ of the optimal revenue. This is quite remarkable, especially when the value distributions are non-identical or have large tails. Why can we restrict the prices into a bounded interval as above, when each of the underlying value distributions concentrates on different supports, or even worse when they do not exhibit good concentration at all as in a power law distribution?

Besides enabling the aforementioned structural results for our problem, we expect that our extreme value theorems will find applications in future work, and indeed they have already been used in followup research. In [DW12, CDW12a], our theorem is used to convert nearly-optimal multi-item multi-bidder mechanisms for distributions with bounded support to nearly-optimal mechanisms for MHR distributions. For the same setting, [CH13] use our theorem to show that relatively simple auctions can extract near-optimal revenue when the bidders are identical, by showing that the welfare is highly-concentrated. We also note that extreme value theorems have been obtained in Statistics for large classes of distributions [dHF06], and indeed such theorems have been applied to optimal mechanism design prior to our work [BH08]. Nevertheless, the known extreme value theorems are typically asymptotic, only hold for maxima of i.i.d. random variables, and are not known to hold for all MHR or regular distributions. We can instead handle the non-i.i.d. case, maxima of a finite number of random variables, and the full spectrum of MHR and regular distributions.

1.2 Algorithmic Ideas: Covers of Revenue Distributions

We overview our approach for Theorem 1. A natural strategy for reducing the search space for an approximately optimal price vector is to discretize the set of possible prices into a finite set,

⁸The *equal revenue distribution* is supported on $[1, +\infty]$ and has cumulative density function $F(x) = 1 - \frac{1}{x}$. Notice that, if a buyer’s value for a single item is distributed according to this distribution, the buyer’s expected value for the item is $+\infty$. However, if the item is priced at any price x , the expected revenue is 1, hence the name “equal revenue.” The equal revenue distribution is itself a regular distribution. So our theorem says that the tail of the maximum of n regular distributions eventually behaves like the tail of a regular distribution.

whose size scales mildly with the number of items, n , and the approximation accuracy, $1/\epsilon$. Of course, even with discretization the number of possible price vectors is exponential in the number of items, and it is not clear how to search this set efficiently. A natural idea to shortcut the search further is to cluster the value distributions into a small number of buckets, containing distributions with similar statistical properties, and proceed to treat all items in a bucket as essentially identical. However, the expected revenue is not sufficiently smooth for us to perform such bucketing. We do obtain a delicate discretization of the supports of the value distributions (Corollary 16), but cannot discretize the probabilities used by these distributions into a coarse-enough accuracy to allow for polynomial-time solvability of the problem.

Our main algorithmic idea is to shift the focus of attention from the space of *value distributions*, which is inherently exponential in the number of items, to the space of all possible *revenue distributions*, which are *single-dimensional* distributions. The revenue from a given price vector can be viewed as a random variable that depends on the (random) values of the items. So, there is still an exponential number of possible revenue distributions, corresponding to all possible price vectors. Nevertheless, we can exploit the single-dimensional nature of these distributions to construct a polynomial-size δ -cover of the set of all possible revenue distributions under the total variation distance between distributions. That is, for every possible revenue distribution, there exists a distribution in our cover that is within δ in total variation distance from it.

Our cover is implicit, i.e. we do not provide a closed-form description for it. We show instead that it can be constructed efficiently using dynamic programming. Our algorithm iteratively considers prefixes of the items and, for each prefix $1 \dots i$, constructs a cover of all possible revenue distributions from only pricing items $1, \dots, i$. For the next iteration, we show that the cover for items $1, \dots, i + 1$ can be easily computed from the cover for items $1, \dots, i$ and the distribution of v_{i+1} . In the end of our iterations we obtain a polynomial-size δ -cover of all possible revenue distributions, and we argue that only a δ -fraction of revenue is lost if we replace the optimal revenue distribution with its closest one in our cover. And, because the cover has polynomial size, we can exhaustively try every distribution in the cover and its associated price vector to pick the one with the highest expected revenue. A more detailed description of our algorithm is given in Section 4, and complete details are provided in Section 5. Theorem 1 follows then easily in Section 6.

Theorem 2 follows similarly, except we employ a stronger discretization (Theorem 17) before using dynamic programming to obtain a cover of all possible revenue distributions. Finally, our algorithms for MHR and regular distributions (Theorems 3 and 4 respectively) are corollaries of Theorem 2, except they require some extra work for restricting the value distributions into a bounded range. This is achieved in Sections 8 (for MHR distributions) and 9 (for regular distributions) using our extreme value theorems (Theorems 19 and 21). The detailed proofs of Theorems 3 and 4 are provided in Appendix H.

1.3 Related and Future Work

The focus of this paper is the multidimensional item pricing problem for a unit-demand buyer whose values for the items are independent. This problem is related to the celebrated multidimensional mechanism design problem, but it is restricted in two ways. First, there is a single bidder who is unit-demand. Second, we are interested in coming close to the revenue of the optimal *deterministic*—i.e. item pricing—mechanism and not the optimal unrestricted mechanism, which may also price lotteries over items. While it is unclear whether the restriction to deterministic mechanisms should make the problem easier or harder computationally, the restriction to a single unit-demand bidder

should make the problem easier compared to having many bidders with arbitrary valuations.

Despite the apparent simplicity of the item pricing problem, (near-)optimal polynomial-time algorithms for it were not known prior to our work. Chawla et al. [CHK07] provide a 3-approximation algorithm, computing a price vector whose revenue is at least a third of the optimal revenue. Their technique is quite elegant, connecting the item pricing problem to a related, single-dimensional mechanism design problem, which can be analyzed using Myerson’s result [Mye81]. Using the same approach, the approximation factor is improved to 2 in [CHMS10], and the result is generalized to the multi-bidder setting, albeit with a worse approximation factor. Different work [BGGM10, Ala11] obtains polynomial-time constant factor approximations for additive bidders, using convex programming relaxations of the problem.

However, all these approaches are limited to constant factor approximations, as ultimately the attained revenue is compared to the optimal revenue in a related single-dimensional setting [CHK07, CHMS10], or a convex programming relaxation of the problem [BGGM10, Ala11]. In particular, the limitation of these approaches comes from avoiding a direct comparison of the attained revenue to the optimal revenue in the actual problem, comparing it instead to the optimal revenue in a related problem. Our work provides instead near-optimal algorithms, using a direct comparison to the real optimum via covers of revenue distributions.

Our work leaves several directions open for exploration and some have already been studied following the announcement of our results [CD11]. We classify them into three categories discussed below.

- *Unit-demand Bidders:* Can our near-optimal algorithms be improved to be exactly optimal? Recent work has shown that the answer is no, namely that there are no exactly optimal polynomial-time algorithms for product value distributions, unless $P = NP$ [CDP⁺14]. Still there is room for improving the dependence of our running times on the approximation parameter ϵ . E.g., is there an algorithm that runs in time polynomial in n and $1/\epsilon$ when the item values are bounded in $[0, 1]$?

And how about correlated distributions over item values? Here, it had already been known that computing an optimal price vector is highly inapproximable in polynomial-time [BK07]. So there cannot even be a polynomial-time constant factor approximation in this case.

Beyond item pricing, it is important to understand the complexity of optimal randomized mechanisms, which may increase revenue by also pricing lotteries over items [Tha04, BCKW10]. For product distributions, Chawla et al. [CMS10] show that randomization does not increase revenue by more than a factor of 4, thus extending the constant-factor approximation algorithms of [CHK07, CHMS10] to the randomized multi-bidder setting, except with worse approximation guarantees. Is there a polynomial-time optimal mechanism for this setting? No computational lower bound is known at the time of writing of this paper.

For correlated distributions over item values, Cai et al. [CDW12b] obtain near-optimal, randomized mechanisms for multi-bidder multi-item settings with unit-demand bidders. For any desired accuracy $\epsilon > 0$, they compute a mechanism whose revenue is within an additive error of ϵ from optimal in time polynomial in $1/\epsilon$ and the size of the support of each bidder’s distribution over valuations, when these distributions are discrete. (When they are continuous, they are handled via fine enough discretization.) This algorithm is clearly also applicable when every bidder’s values for the items are independent (i.e. the setting discussed in the previous paragraph). However, the dependence of the running time on the support of

the product distribution may be unreasonable computationally. Indeed, a discrete product distribution can be described by specifying all of its marginals, with description complexity logarithmic in the size of its support.

- *Additive Bidders:* Can our algorithms be extended to additive bidders? Here, an optimal mechanism may increase revenue by pricing bundles of items [MV06], or (if randomization is allowed) lotteries over bundles of items. Exploiting our extreme value theorems for MHR distributions, Cai and Huang [CH13] provide near-optimal polynomial-time mechanisms for multiple i.i.d. bidders, whose values for the items are independently distributed according to MHR distributions. Moreover, Daskalakis et al. [DDT14] show that this result cannot be made exact for general product distributions. They show that, subject to widely held complexity theoretic beliefs—in particular that $\text{ZPP} \not\subseteq \text{P}^{\#\text{P}}$,⁹ computing and implementing an exactly optimal mechanism cannot be done computationally efficiently. Indeed, this is true even in the ostensibly simple setting where there is a single additive bidder whose values for the items are independently distributed on two rational numbers with rational probabilities.

For correlated distributions, Cai et al. [CDW12a] obtain (exactly) optimal mechanisms for multi-bidder multi-item settings with additive bidders, in time polynomial in the size of the support of each bidder’s distribution over valuations.

On a different vein, Daskalakis et al. [DDT13] study the structure (rather than the computational complexity) of optimal mechanisms, following earlier work on the topic by Economists, e.g. [Roc85, Arm00, MV06, MV07, Pav11]. They provide a duality framework based on Monge-Kantorovich duality for characterizing the structure of the optimal mechanism of selling multiple items to a single additive bidder.

- *General Settings:* It is important to understand the computational complexity of mechanism design in general settings: multiple bidders, general valuations (beyond unit-demand and additive), general constraints on what allocations of items to bidders are feasible, and general objectives, potentially going beyond the familiar objectives of revenue and welfare.¹⁰ In recent work, Cai et al. [CDW13] provide a computational black-box reduction from *mechanism design* for maximizing an arbitrary concave objective O ¹¹ under arbitrary allocation constraints and an arbitrary family of bidder valuations (e.g. submodular, supermodular, etc.) to *algorithm design* for that same objective O , modified by an additive virtual welfare and virtual revenue term, and under the same allocation constraints and family of allowed valuations. Roughly speaking, they show that, whenever the algorithmic problem is polynomial-time solvable (exactly or approximately), the mechanism design problem also becomes solvable (exactly or approximately) in time polynomial in the size of the support of each bidder’s distribution over valuations. It is important to find applications of this reduction to settings of interest beyond optimizing fractional max-min fairness for additive bidders, which was done in [CDW13].

⁹ $\text{ZPP} \supseteq \text{P}^{\#\text{P}}$ would imply that there are randomized polynomial-time algorithms for NP-complete problems, which is widely believed to be unlikely.

¹⁰A general objective O takes as input the valuations of the bidders \vec{t} and a randomized allocation and price vector (A, p) and outputs a real number $O(\vec{t}, (A, p))$.

¹¹An objective function $O(\vec{t}, (A, p))$ is called *concave* iff, for all bidder valuations \vec{t} , and all (A_1, p_1) and (A_2, p_2) , it holds that $O(\vec{t}, \frac{1}{2}(A_1, p_1) + \frac{1}{2}(A_2, p_2)) \geq \frac{1}{2}O(\vec{t}, (A_1, p_1)) + \frac{1}{2}O(\vec{t}, (A_2, p_2))$, where $\frac{1}{2}(A_1, p_1) + \frac{1}{2}(A_2, p_2)$ denotes uniformly randomizing between (A_1, p_1) and (A_2, p_2) . Clearly, revenue and welfare satisfy this condition with equality, but several other objectives are concave, such as the max-min fairness objective considered in [CDW13].

2 Preliminaries

Computational Problems. We define three variants of the item pricing problem. ADDITIVEPRICE and PRICE are the main computational problems that we aim to solve, but RESTRICTEDPRICE is an auxiliary one that is helpful in the analysis. For the value distributions that we consider, it can be shown that all three problems have finite optimal solutions.

ADDITIVEPRICE: Input: A collection of mutually independent random variables $\{v_i\}_{i=1}^n$, and some $\epsilon > 0$. **Output:** A vector of prices $P = (p_1, \dots, p_n)$ such that the expected revenue \mathcal{R}_P from using P , defined as in Eq. (1), is within an additive ϵ of the optimal revenue achieved by any price vector.

PRICE: Input: A collection of mutually independent random variables $\{v_i\}_{i=1}^n$, and some $\epsilon > 0$. **Output:** A vector of prices $P = (p_1, \dots, p_n)$ such that the expected revenue \mathcal{R}_P from using P , defined as in Eq. (1), is within a $(1 + \epsilon)$ -factor of the optimal revenue achieved by any price vector.

RESTRICTEDPRICE: Input: A collection of mutually independent random variables $\{v_i\}_{i=1}^n$ supported on a common discrete set \mathcal{S} , and a discrete set $\mathcal{P} \subset \mathbb{R}_{\geq 0}$ of possible prices. **Output:** A vector of prices $P = (p_1, \dots, p_n) \in \mathcal{P}^n$ such that the expected revenue \mathcal{R}_P from using P is optimal among all vectors in \mathcal{P}^n .

In Section 3 we describe how these computational problems are interconnected through other results in this paper to establish Theorems 1 through 4.

Computational Efficiency. Throughout the paper we use the standard convention of identifying “computational efficiency” with polynomial-time computation. Namely, we will say that an algorithm is “computationally efficient” iff its running time is polynomial in the number of bits required to describe the input to the algorithm.

Reductions Between Computational Problems. We provide several reductions between different flavors of the item pricing problem. Formally, a (polynomial-time) reduction from a computational problem P_1 (e.g. PRICE) to a computational problem P_2 (e.g. RESTRICTEDPRICE) is a pair of (polynomial-time) algorithms \mathcal{A} and \mathcal{B} satisfying the following properties. For all inputs Π_1 to P_1 :

- $\mathcal{A}(\Pi_1)$ is a valid input to P_2 ;
- if S is a solution to $\mathcal{A}(\Pi_1)$ then $\mathcal{B}(S)$ is a solution to Π_1 .

For example, a polynomial-time reduction from PRICE to RESTRICTEDPRICE would allow us to convert (in polynomial time) any input to PRICE to a valid input to RESTRICTEDPRICE so that, if we found a solution to the latter, we would also be able to compute (in polynomial time) a solution to the former.

Approximation Algorithms. Our algorithmic results use the concept of a *Polynomial-Time Approximation Scheme*, or PTAS. A PTAS for a computational problem such as PRICE is a collection of algorithms $(\mathcal{A}_\epsilon)_\epsilon$, indexed by the approximation parameter ϵ , such that, for all $\epsilon > 0$ and for any given input Π to the algorithm, Algorithm \mathcal{A}_ϵ computes an ϵ -optimal solution to Π in time

$d(\Pi)^{g(1/\epsilon)}$, where $d(\Pi)$ is the number of bits required to describe problem Π and g is some increasing function of $1/\epsilon$, which does not depend on Π or its description complexity. The algorithms in the collection are called polynomial-time because for all fixed ϵ , e.g. $\epsilon = 1/3$, the running time of \mathcal{A}_ϵ is polynomial in the description of the problem. A *quasi-polynomial-time approximation scheme*, or Quasi-PTAS is a similar concept, except that running time is $2^{g(1/\epsilon) \cdot \text{poly}(\log d(\Pi))}$ for some function g instead of $2^{g(1/\epsilon) \cdot \log d(\Pi)}$ as in a PTAS.

Distributions. For a random variable X we denote by $F_X(x)$ the cumulative distribution function of X , and by $f_X(x)$ its probability density function. We also let $u_{min}^X = \sup\{x | F_X(x) = 0\}$ and $u_{max}^X = \inf\{x | F_X(x) = 1\}$. u_{max}^X may be $+\infty$, but we assume that $u_{min}^X \geq 0$, since our distributions represent value distributions. We drop the subscript/superscript of X , if X is clear from context.

We proceed with a precise definition of *Monotone Hazard Rate (MHR)* and *Regular* distributions, which are two commonly studied families of distributions.

Definition 11 (Monotone Hazard Rate Distribution). *We say that a one-dimensional differentiable distribution F has Monotone Hazard Rate if $\frac{f(x)}{1-F(x)}$ is non-decreasing in $[u_{min}, u_{max}]$. We call such F a Monotone Hazard Rate, or MHR, distribution.*

Definition 12 (Regular Distribution). *A one-dimensional differentiable distribution F is called regular if $x - \frac{1-F(x)}{f(x)}$ is non-decreasing in $[u_{min}, u_{max}]$.*

It is worth noting that all MHR distributions are also regular distributions, but there are regular distributions that are not MHR. The family of MHR distributions includes such familiar distributions as the Normal, Exponential, and Uniform distributions. The family of regular distributions contains a broader range of distributions, including fat-tail distributions $f_X(x) \sim x^{-(1+\alpha)}$ for $\alpha \geq 1$ (which are not MHR). In Appendices F.1 and G.1 we establish important properties of MHR and regular distributions. These properties are instrumental in establishing our extreme value theorems for these distributions (Theorems 19 and 21 in Sections 8 and 9).

To formally study the complexity of the item pricing problem, we need to pay attention to how value distributions are described as part of the input to the problem. We discuss this technical issue in Appendix B, entertaining three types of access to a distribution. Maybe we are given an *explicit description* of the distribution, specifying its support and the probabilities assigned to every value in the support. Such explicit description is appropriate if the distribution is discrete and has finite support. Alternatively, we may have more limited access to the distribution. E.g., we may only have *sample access* to the distribution via a procedure that generates independent samples from it. Or, we may have *oracle access* to the cumulative distribution function via a procedure that returns its value at any queried point. We formally discuss these types of access to a distribution in Appendix B, also defining precisely what it means for an algorithm who takes as input a distribution or outputs a distribution to be “computationally efficient” in each case.

Order Notation. Throughout the text we use the $O(\cdot)$, $\Omega(\cdot)$ and $\Theta(\cdot)$ notation. Let $f(x)$, $g(x)$ be two positive functions defined on some infinite subset of \mathbb{R}_+ . Then:

- we write $f(x) = O(g(x))$ iff there exist some positive reals M and x_0 such that $f(x) \leq Mg(x)$, for all $x > x_0$;
- we write $f(x) = \Omega(g(x))$ iff there exist some positive reals m and x_0 such that $f(x) \geq mg(x)$, for all $x > x_0$; and

- we write $f(x) = \Theta(g(x))$ iff $f(x) = O(g(x))$ and $f(x) = \Omega(g(x))$.

Other Notation. Whenever we write $\text{poly}(x)$ in an expression providing a bound to some quantity, we mean that there exists some positive polynomial $p(x)$ which can replace “ $\text{poly}(x)$ ” so that the bound is true. Whenever we write \log in some expression without specifying the base of the logarithm, any constant base that is larger than 1 works. For some positive rational number x , we write $\langle x \rangle$ to denote the bit complexity of x , i.e. the number of bits required to specify the numerator and denominator of an irreducible fraction representing x .

3 Paper Organization

We provide a roadmap to the paper and a high-level description of our approach. We first study `RESTRICTEDPRICE`. Despite its input/output restrictions, it addresses the major algorithmic challenges of the item pricing problem. Our approach to `RESTRICTEDPRICE` is to first design a dynamic programming algorithm that produces a cover of all possible revenue distributions arising from price vectors in \mathcal{P}^n , where \mathcal{P} is the set of possible prices specified in the input to the problem. Using the cover it is then easy to obtain a near-optimal algorithm for `RESTRICTEDPRICE`, which is not necessarily polynomial-time. Section 4 provides an intuitive description of the dynamic programming approach for producing the cover, and Section 5 describes our algorithm for `RESTRICTEDPRICE` formally. This algorithm serves as the main algorithmic tool of this paper, and is at the root of the tree structure of Figure 1, summarizing our results and proofs.

In Section 6, we obtain Theorem 1 by reducing `ADDITIVEPRICE` for value distributions supported on $[0, 1]$ to `RESTRICTEDPRICE`. The reduction is obtained by showing a discretization result, establishing that the supports of the value distributions as well as the candidate prices can be discretized without too much loss in revenue. The reduction is summarized by Corollary 16, which together with our algorithm for `RESTRICTEDPRICE` immediately shows Theorem 1.

In Section 7, we move on to multiplicative approximations, establishing Theorem 2. The approach is similar, reducing `PRICE` to `RESTRICTEDPRICE` by discretizing the supports of the value distributions as well as the set of available prices. However, Theorem 15, the discretization result at the heart of Corollary 16 (our reduction from Section 6), is not strong enough for our purposes here. We establish instead a stronger discretization (Theorem 17) that is sufficiently powerful for our reduction.

In Sections 8 and 9, we establish our algorithms for MHR and regular distributions. In Section 8, we present an extreme value theorem for MHR distributions (Theorem 19). This theorem enables us to obtain a polynomial-time reduction from `PRICE` where the value distributions are MHR to `PRICE` where the value distributions are supported on a common range of the form $[u_{\min}, r \cdot u_{\min}]$, where the multiplier r is independent of the number of items n . Our reduction is summarized by Theorem 18. Theorem 3 follows then as a corollary of Theorem 18 and Theorem 2. Our algorithm for regular distributions follows similarly in Section 9. We show an extreme value theorem for regular distributions (Theorem 21), enabling a reduction from `PRICE` with regular distributions to `PRICE` with distributions supported on a common range of the form $[u_{\min}, r \cdot u_{\min}]$, except that now the multiplier r depends polynomially on the number of items n . Our reduction is summarized in Theorem 20. Theorem 4 follows then as a corollary of Theorem 20 and Theorem 2.




Figure 1: Overview of our results and the proof structure. Arrows are implications.

Reading the paper. Sections 4 through 9 are meant to provide a self-contained overview of the proofs of our algorithmic results, with the appendices containing the complete proof details. Appendix A provides a roadmap to the appendices.

4 Probabilistic Covers of Revenue Distributions

In this section, we discuss our algorithmic approach to `RESTRICTEDPRICE`, postponing the description of our algorithm for it to Section 5. As we have already discussed in Section 3, although seemingly restricted this problem captures the main algorithmic challenges underlying problems `PRICE` and `ADDITIVEPRICE`. In particular, our algorithm for `RESTRICTEDPRICE` will become a

central building block in all our algorithmic results (Theorems 1 through 4).

For convenience, throughout this section we will take F_1, \dots, F_n to be a collection of distributions supported on a discrete set $\mathcal{S} = \{v^{(1)}, v^{(2)}, \dots, v^{(k_1)}\}$, and v_1, \dots, v_n to be a collection of mutually independent random variables distributed according to the F_i 's. We will then assume that the input to RESTRICTEDPRICE comprises the v_i 's and a finite set of prices $\mathcal{P} = \{p^{(1)}, p^{(2)}, \dots, p^{(k_2)}\}$.

The obvious algorithmic challenge in RESTRICTEDPRICE is that, even though the set of possible prices is finite, there are still exponentially many (namely k_2^n) possible price vectors that we need to choose from for an optimal one. If k_2 were a constant and the items were i.i.d., then we could decrease the possible vectors to a polynomial number by exploiting the symmetry of the items.¹² Similarly, we can obtain polynomial-time algorithms for the case where there is only a constant number of possible value distributions and a constant number of possible prices. However, when all the F_i 's may be different, the problem looks inherently exponential, even if both k_1 and k_2 are absolute constants, e.g., even when the value distributions are supported on 2 possible values and there are 2 possible prices available.

Our algorithmic approach is enabled by a shift in perspective, which may be applicable to other problems with a similar structure. To illustrate the approach, let us view our problem in the graphical representation of Figure 2. C is a function that takes as input a price vector $P = (p_1, \dots, p_n)$ and outputs the distribution F_{R_P} of the revenue of the seller under this price vector. Indeed, the revenue of the seller is a random variable R_P that depends on the random variables $\{v_i\}_{i \in [n]}$. So in order to compute the distribution of the revenue, C also takes as input the distributions F_1, \dots, F_n . What we are aiming at maximizing is the expectation \mathcal{R}_P of R_P .




Figure 2: The Revenue Distribution as the output of a function. The inputs to the function are the prices and the value distributions.

Given our restriction of the prices to a finite set $\{p^{(1)}, p^{(2)}, \dots, p^{(k_2)}\}$, there are k_2^n possible inputs to the function, and at most k_2^n possible revenue distributions that the function can output. Our main conceptual idea is the following:

Instead of searching the space of possible price vectors that can be input to C , we search the space of possible outputs of C , i.e. the space of all possible revenue distributions resulting from different price vectors, for one with maximum expectation.

Moreover, to efficiently search the space of all possible revenue distributions, we construct an appropriately small subset of it and only search the distributions in that subset.

The subset we construct is a probabilistic cover (under some appropriate metric) of the space of all

¹²A broader exposition of the role of symmetries in mechanism design can be found in [DW12].

possible revenue distributions.¹³ The properties of our cover that are crucial for our algorithmic applications are the following: (a) the cover has small cardinality, and (b) for any possible revenue distribution that the function may output, there exists a revenue distribution in our cover whose expectation is close.

Constructing the Cover. At a high level, we construct our cover using dynamic programming (henceforth DP for short), whose steps are interleaved with coupling arguments that prune the size of the DP table before proceeding to the next step.

Intuitively, our DP algorithm sweeps the items from 1 through n , maintaining a cover of the revenue distributions produced by all possible price vectors on every prefix of the items. More precisely, for each prefix $1 \dots j$ of the items, our DP table keeps track of all possible feasible collections of $k_1 \times k_2$ probability values, where \Pr_{i_1, i_2} , $i_1 \in [k_1], i_2 \in [k_2]$, denotes the probability that the item with the largest value-minus-price gap (i.e. the item that would have been sold in a sale that only sells items 1 through j) has value $v^{(i_1)}$ for the buyer and is assigned price $p^{(i_2)}$ by the seller. I.e. we store in our DP table all possible (winning-value, winning-price) distributions that can arise from a price vector on every prefix of the items. The reasons we store these distributions are the following:

- First, if we have all possible (winning-value, winning-price) distributions for the full set of items, we can search for the one with the highest expected revenue. For every distribution we will also maintain in our DP table a price vector resulting in that distribution. So, once we have found the distribution with the optimal expected revenue, we will also find the price vector with that optimal revenue.
- Second, we can construct the set of all possible (winning-value, winning-price) distributions for the full set of items, by considering one prefix at a time. In particular, suppose that we have all possible (winning-value, winning-price) distributions for the prefix of items $1 \dots j$. By combining every such distribution with all k_2 possible prices for item $j+1$, we can compute all possible (winning-value, winning-price) distributions for the prefix of items $1 \dots j+1$. That is, if we have these distributions for a prefix of items, we do not need any other information to extend the prefix by one item. For this scheme to work, observe that it is crucial to maintain the joint distribution of both the winning-value and the winning-price, rather than just the distribution of the winning-price.

Clearly, the dynamic programming approach that we just outlined for computing a cover of all possible revenue distributions achieves nothing in terms of reducing the number of distributions. Indeed, there could be one (winning-value, winning-price) distribution for every price vector, so that the total number of distributions that we need to store in our DP table is exponential. To control the size of our cover from exploding, we show that we can be coarse in our bookkeeping of the (winning-value, winning-price) distributions, without sacrificing much revenue. Indeed, it is here where viewing our problem in the “upside-down” manner illustrated in Figure 2 (i.e. targeting a cover of the output of C) is important. We show that we can discretize the probabilities used by the distributions stored in the DP table into multiples of some fraction $\frac{1}{m}$ without losing much revenue. In particular, after a prefix of items is processed by the algorithm, we show that we

¹³A δ -cover of a set of distributions \mathcal{F} with metric d is a subset $\mathcal{F}' \subseteq \mathcal{F}$ such that for all $F \in \mathcal{F}$ there exists some $F' \in \mathcal{F}'$ such that $d(F, F') \leq \delta$.

can discretize the probabilities in all distributions in the table before considering the next item. That the loss due to coarsening the probabilities is not significant follows from coupling arguments interleaved with the steps of dynamic programming.

In the next section we make our ideas precise, obtaining our algorithm for RESTRICTEDPRICE.

5 The Algorithm for the Discrete Problem

In this section, we formalize our ideas from the previous section, describing our main algorithmic result for RESTRICTEDPRICE. We use the same notation as in Section 4, namely we assume that the input distributions are supported on a common set \mathcal{S} of cardinality k_1 and the prices are restricted to a set \mathcal{P} of cardinality k_2 . We also denote by OPT the optimal expected revenue for the input value distributions when the prices are restricted to \mathcal{P} .

The Algorithm. As a first step we discretize the probabilities used by the input distributions. We prove a discretization lemma that provides a polynomial-time reduction from our problem into a new one, where additionally the probabilities that the value distributions assign to each point in their support \mathcal{S} are integer multiples of $1/m$, for some integer m that is a free parameter in our algorithm. We show that the loss in revenue resulting from our reduction is at most an additive $\frac{4k_1n}{m} \max_i \{p^{(i)}\}$ in the following sense: for any price vector P , the expected revenue from the original value distributions $\{F_i\}_i$ and the expected revenue from the discretized distributions $\{\hat{F}_i\}_i$ are within an additive $\frac{4k_1n}{m} \max_i \{p^{(i)}\}$. Moreover, the construction of Lemma 23 is explicit, so from now on we can assume that we know the $\{\hat{F}_i\}_i$ explicitly, regardless of what type of access we have to the $\{F_i\}_i$ (see Appendix B).

The second phase of our algorithm is the dynamic programming algorithm outlined in Section 4. We provide some further details on this now. Our algorithm computes a Boolean function $g(i, \text{Pr})$, whose arguments lie in the following range: $i \in [n]$ and $\text{Pr} = (\text{Pr}_{1,1}, \text{Pr}_{1,2}, \dots, \text{Pr}_{k_1, k_2})$, where each $\text{Pr}_{i_1, i_2} \in [0, 1]$ is an integer multiple of $\frac{1}{m}$. The function g is stored in a table that has one cell for every setting of i and Pr , and the cell contains a 0 or a 1 depending on the value of g at the corresponding input. In the terminology of the previous section, argument i indexes the last item in a prefix of items and Pr defines a (winning-value, winning-price) distribution whose probabilities are integer multiples of $\frac{1}{m}$. If Pr can arise from some pricing of the items $1 \dots i$ (up to discretization of probabilities into multiples of $\frac{1}{m}$), we intend to store $g(i, \text{Pr}) = 1$; otherwise we store $g(i, \text{Pr}) = 0$. For each cell of the table such that $g(i, \text{Pr}) = 1$, we also store a price vector on the corresponding prefix of items $1 \dots i$ consistent with Pr .

For conciseness, we give next a high-level description of the dynamic programming algorithm, postponing its full details to Appendix C.1. The table is filled in a bottom-up fashion from $i = 1$ through n . At the end of the i -th iteration, we have computed all feasible “discretized” (winning-value, winning-price) distributions for the prefix of items $1 \dots i$, where “discretized” means that all probabilities have been rounded into multiples of $1/m$. For the next iteration, we try all possible prices $p^{(j)}$ for item $i + 1$ and compute how each of the feasible discretized (winning-value, winning-price) distributions for the prefix $1 \dots i$ evolves into a discretized distribution for the prefix $1 \dots i + 1$, setting the corresponding cell of layer $g(i + 1, \cdot)$ of the DP table to 1. Notice, in particular, that *we lose accuracy in every step of the dynamic programming algorithm*, as each step involves computing how a discretized distribution for items $1 \dots i$ evolves into a distribution for items $1 \dots i + 1$ and

then rounding the latter back into multiples of $1/m$. We show in the analysis of our algorithm that the error accumulating from these roundings can be controlled via coupling arguments.

After computing the truth-table of function g , we look at all cells such that $g(n, \text{Pr}) = 1$ and evaluate the expected revenue resulting from the distribution Pr , i.e.

$$\mathcal{R}_{\text{Pr}} = \sum_{i_1 \in [k_1], i_2 \in [k_2]} p^{(i_2)} \cdot \text{Pr}_{i_1, i_2} \cdot \mathbb{1}_{v^{(i_1)} \geq p^{(i_2)}}.$$

Having located the cell whose \mathcal{R}_{Pr} is the largest, we output the price vector stored in that cell.

Running Time and Correctness. Next we bound the algorithm's running time and revenue.

Theorem 13. *Given an instance of RESTRICTEDPRICE, where the value distributions are supported on a discrete set \mathcal{S} of cardinality k_1 and the prices are restricted to a discrete set \mathcal{P} of cardinality k_2 , and for any choice of discretization accuracy $m \geq 2k_1$, the algorithm described in this section produces a price vector with expected revenue at least*

$$OPT - \frac{(2nk_1k_2 + 16k_1n)}{m} \cdot \max\{\mathcal{P}\},$$

where $\max\{\mathcal{P}\}$ is the maximum element in \mathcal{P} and OPT the optimal expected revenue. The running time of the algorithm is polynomial in the size of the input and $m^{k_1k_2}$.

The proof of the theorem is given in Appendix C. Intuitively, if we did not perform any rounding of distributions, our algorithm would have been *exact*, outputting an optimal price vector in $\{p^{(1)}, \dots, p^{(k_2)}\}^n$. What we show is that the roundings performed at the steps of the dynamic programming algorithm are fine enough that do not become detrimental to the revenue. To show this, we use coupling arguments, invoking the coupling lemma and the optimal coupling theorem after each step of the algorithm. (See Lemma 22 in Appendix C.2.) This way, we show that the rounded (winning-value, winning-price) distributions maintained by the algorithm for each price vector are close in total variation distance to the corresponding exact distributions arising from these price vectors, culminating in Theorem 13.

6 Additive PTAS for values distributed in $[0, 1]^n$

In this section, we provide a polynomial-time reduction from ADDITIVEPRICE(\mathcal{V}, ϵ), for value distributions $\mathcal{V} = \{v_i\}_i$ supported on $[0, 1]$, to $O(\epsilon)$ -approximating RESTRICTEDPRICE($\hat{\mathcal{V}}, \mathcal{P}$), where $\hat{\mathcal{V}}$ is a collection of mutually independent random variables supported on a common set of cardinality $\text{poly}(1/\epsilon)$ and $|\mathcal{P}| = \text{poly}(1/\epsilon)$. A PTAS for ADDITIVEPRICE then follows from Theorem 13 with an appropriate choice of the discretization m .

As a first step, we reduce ADDITIVEPRICE(\mathcal{V}, ϵ) to ADDITIVEPRICE($\tilde{\mathcal{V}}, O(\epsilon)$), where the random variables $\tilde{\mathcal{V}} = \{\tilde{v}_i\}_i$ are independently distributed in $[O(\epsilon), 1]$. The reduction is quite straightforward, replacing all sampled values that are smaller than some $O(\epsilon)$ with $O(\epsilon)$ and keeping the rest unchanged. We argue that a nearly optimal price vector for the new value distributions is also nearly optimal for the original value distributions. Formally,

Lemma 14. *Let $\mathcal{V} = \{v_i\}_{i \in [n]}$ be a collection of mutually independent random variables supported on $[0, 1]$. For any $\epsilon > 0$, there is a polynomial-time reduction from $\text{ADDITIVEPRICE}(\mathcal{V}, \epsilon)$ to $\text{ADDITIVEPRICE}(\tilde{\mathcal{V}}, \epsilon/3)$, where $\tilde{\mathcal{V}} = \{\tilde{v}_i\}_{i \in [n]}$ is a collection of mutually independent random variables supported on $[\epsilon/6, 1]$.*

The proof can be found in Appendix D.4.

Next we want to discretize the problem $\text{ADDITIVEPRICE}(\tilde{\mathcal{V}}, \epsilon/3)$. As alluded to in Section 1, the expected revenue can be sensitive even to small perturbations of the prices and the probability distributions. So our discretization, summarized in the next theorem, must be done delicately.

Theorem 15 (Price/Value Discretization for Additive Approximation). *Let $\mathcal{V} = \{v_i\}_{i \in [n]}$ be a collection of mutually independent random variables supported on a bounded set $[u_{\min}, u_{\max}] \subset \mathbb{R}_+$, and let $r = \frac{u_{\max}}{u_{\min}} \geq 1$. For any $\epsilon > 0$, there is a reduction from $\text{ADDITIVEPRICE}(\mathcal{V}, \epsilon)$ to approximating $\text{RESTRICTEDPRICE}(\hat{\mathcal{V}}, \mathcal{P})$ to within an additive error of $\frac{\epsilon}{6}$, where*

- $\hat{\mathcal{V}} = \{\hat{v}_i\}_{i \in [n]}$ is a collection of mutually independent random variables that are supported on a common set of cardinality $O\left(\frac{u_{\max}^2 \log r}{\epsilon^2}\right)$;
- $|\mathcal{P}| = O\left(\frac{u_{\max}^2 \log r}{\epsilon^2}\right)$ and $\max_{x \in \mathcal{P}} x \leq \frac{7}{6}u_{\max}$.

Moreover, if u_{\min} and u_{\max} are given explicitly as input to the reduction,¹⁴ the running time of the reduction is polynomial in the description of \mathcal{V} , $\langle u_{\min} \rangle$, $\langle u_{\max} \rangle$, and $1/\epsilon$.

That the prices can be restricted to a discrete set without hurting the revenue too much follows immediately from a discretization lemma attributed to Nisan [CHK07]. (See also [HK05] for a related discretization.) Our price discretization result is summarized in Lemma 25 of Appendix D.2. The discretization of the support of the value distributions is inspired by Nisan’s lemma, and our corresponding discretization result is summarized in Lemma 31.

Combining Lemma 14 and Theorem 15, we complete our reduction from ADDITIVEPRICE to RESTRICTEDPRICE .

Corollary 16. *Let $\mathcal{V} = \{v_i\}_{i \in [n]}$ be a collection of mutually independent random variables supported on $[0, 1]$. For any $\epsilon > 0$, there is a polynomial-time reduction from $\text{ADDITIVEPRICE}(\mathcal{V}, \epsilon)$ to approximating $\text{RESTRICTEDPRICE}(\hat{\mathcal{V}}, \mathcal{P})$ to within an additive error of $\frac{\epsilon}{18}$, where $\hat{\mathcal{V}}$ is a collection of mutually independent random variables supported on a common set of cardinality $O\left(\frac{\log 1/\epsilon}{\epsilon^2}\right)$, $|\mathcal{P}| = O\left(\frac{\log 1/\epsilon}{\epsilon^2}\right)$ and $\max_{x \in \mathcal{P}} x \leq 7/6$.*

We are now ready to prove Theorem 1, using the reduction of Corollary 16 and our algorithm from Section 5.

Proof of Theorem 1: We first perform the reduction of Corollary 16. In the resulting instance of RESTRICTEDPRICE both the cardinality of the support of the value distribution and the number of available prices are $O\left(\frac{\log 1/\epsilon}{\epsilon^2}\right)$. Using $m = O\left(\frac{n \cdot \log^2 1/\epsilon}{\epsilon^5}\right)$ we can solve the resulting instance of RESTRICTEDPRICE to within additive error $O(\epsilon)$ using the algorithm of Theorem 13. The running time of the algorithm is polynomial in the input and $n \frac{\log^3 1/\epsilon}{\epsilon^4}$. \square

¹⁴This requirement is only relevant if we have oracle access to the distributions of the v_i ’s, as if we are given the distributions explicitly we immediately also know u_{\min} and u_{\max} .

7 Multiplicative PTAS

For the remainder of our main exposition, we move on to multiplicative approximations to the item pricing problem, obtaining algorithms for PRICE. In this section, we study the general problem where the values are independently distributed on a bounded range $[u_{min}, u_{max} \equiv r \cdot u_{min}]$ according to arbitrary distributions, proving Theorem 2.

Notice that, using our results from the previous sections, we can already get an algorithm for PRICE. We can first apply our reduction from Theorem 15 to discretize the prices and the supports of the value distributions. Then we can use our algorithm from Theorem 13 to solve the discretized problem. However to convert the additive approximation of this algorithm to a multiplicative one, we need to choose the approximation to be no worse than $\epsilon \cdot u_{min}$. This requirement forces the support of the discretized value distributions to be $\Omega(r^2 \log r / \epsilon^2)$ and the discrete set of prices to also have cardinality $\Omega(r^2 \log r / \epsilon^2)$. Hence, the algorithm has running time polynomial in $n^{r^4 \log^2 r / \epsilon^4}$.

In this section, we present a stronger discretization result, reducing the size of the support of the value distributions and the cardinality of the price set to linear in $\log r$. With this new discretization, we can speed up the running time of our algorithm to $n^{\text{poly}(\log r, 1/\epsilon)}$. Our improved discretization reduction is presented below, and proved in Appendix D.5.

Theorem 17 (Price/Value Distribution Discretization). *Let $\mathcal{V} = \{v_i\}_{i \in [n]}$ be a collection of mutually independent random variables supported on a bounded range $[u_{min}, u_{max}] \subset \mathbb{R}_+$, and let $r = \frac{u_{max}}{u_{min}} \geq 1$. For any $\epsilon \in \left(0, \frac{1}{(4 \lceil \log_2 r \rceil)^{1/6}}\right)$, there is a reduction from $\text{PRICE}(\mathcal{V}, \epsilon)$ to the problem of approximating $\text{RESTRICTEDPRICE}(\hat{\mathcal{V}}, \mathcal{P})$ to within a factor of $(1 - \frac{\epsilon}{16})$, where*

- $\hat{\mathcal{V}} = \{\hat{v}_i\}_{i \in [n]}$ is a collection of mutually independent random variables that are supported on a common set of cardinality $O\left(\frac{\log r}{\epsilon^{16}}\right)$;
- $|\mathcal{P}| = O\left(\frac{\log r}{\epsilon^2}\right)$.

Moreover, if u_{min} and u_{max} are given explicitly as input to the reduction,¹⁵ the running time of the reduction is polynomial in the description of \mathcal{V} , $\langle u_{min} \rangle$, $\langle u_{max} \rangle$, and $1/\epsilon$.

Combining our discretization from Theorem 17 with our algorithm from Theorem 13, it is easy to show Theorem 2. We only sketch the proof here, providing a formal proof in Appendix E.

Proof of Theorem 2: (sketch) We first perform the reduction of Theorem 17 to get an instance of RESTRICTEDPRICE where both the values and the prices come from discrete sets of cardinality $O\left(\frac{\log r}{\text{poly}(\epsilon)}\right)$. Using the algorithm of Theorem 13, we can then approximately solve this instance to within a factor of $1 - O(\epsilon)$ in time polynomial in the input and $n^{\frac{\log^2 r}{\text{poly}(\epsilon)}}$. \square

8 Extreme Values of MHR Distributions

We reduce the problem of finding a near-optimal price vector for value distributions that are MHR to finding a near-optimal price vector for value distributions that are supported on a bounded

¹⁵This requirement is only relevant if we have oracle access to the distributions of the v_i 's, as if we are given the distributions explicitly we immediately also know u_{min} and u_{max} .

range $[u_{\min}, u_{\max}]$, where u_{\max}/u_{\min} is only a function of the desired approximation $\epsilon > 0$. More precisely, we establish the following reduction.

Theorem 18 (From MHR to Bounded Distributions). *Let $\mathcal{V} = \{v_i\}_{i \in [n]}$ be a collection of mutually independent MHR random variables. Then there exists some $\beta = \beta(\mathcal{V}) > 0$ such that for all $\epsilon \in (0, 1/4)$, there is a reduction from $\text{PRICE}(\mathcal{V}, c\epsilon \log_2(\frac{1}{\epsilon}))$ to $\text{PRICE}(\tilde{\mathcal{V}}, \epsilon)$, where $\tilde{\mathcal{V}} := \{\tilde{v}_i\}_i$ is a collection of mutually independent random variables supported on the set $[\frac{\epsilon}{2}\beta, 2\log_2 \frac{1}{\epsilon}\beta]$, and c is some absolute constant.¹⁶*

Moreover, β is efficiently computable from the distributions of the v_i 's, and, for all ϵ , the running time of the reduction is polynomial in the size of the input and $\frac{1}{\epsilon}$.

We discuss the essential elements of our reduction below. Most crucially, the reduction is enabled by the following theorem, characterizing the extreme values of a collection of independent MHR distributions.

Theorem 19 (Extreme Values of MHR distributions). *Let X_1, \dots, X_n be a collection of independent random variables whose distributions are MHR. Then there exists some anchoring point β such that $\Pr[\max_i\{X_i\} \geq \beta/2] \geq 1 - \frac{1}{\sqrt{e}}$ and*

$$\int_{2\beta \log_2 1/\epsilon}^{+\infty} t \cdot f_{\max_i\{X_i\}}(t) dt \leq 36\beta\epsilon \log_2 1/\epsilon, \text{ for all } \epsilon \in (0, 1/4), \quad (2)$$

where $f_{\max_i\{X_i\}}(t)$ is the probability density function of $\max_i\{X_i\}$. Moreover, β is efficiently computable from the distributions of the X_i 's.

Theorem 19, whose proof is given in Appendix F.2, implies that, for ϵ sufficiently small, at least a $(1-\epsilon)$ -fraction of $\mathbb{E}[\max_i X_i]$ is contributed to by values that are no larger than $\mathbb{E}[\max_i X_i] \cdot O(\log_2 \frac{1}{\epsilon})$. Our result is quite surprising, especially for the case of non-identically distributed MHR random variables. Why should most of the contribution to $\mathbb{E}[\max_i X_i]$ come from values that are close (*within a function of ϵ only*) to the expectation, when the underlying random variables X_i may concentrate on widely different supports? To obtain the theorem one needs to understand how the tails of the distributions of a collection of independent MHR random variables contribute to the expectation of their maximum. Our proof technique is intricate, defining a tournament between the tails of the distributions. Each round of the tournament ranks the remaining distributions according to the size of their tails, and eliminates the lightest half of the distributions. The threshold β is then obtained by some side-information that the algorithm records in every round.

Given our understanding of the extreme values of MHR distributions, our reduction of Theorem 18 from MHR to bounded distributions proceeds in the following steps:

- We start with the computation of the threshold β specified by Theorem 19. This computation can be done efficiently, as stated in the statement of the theorem. Given that $\Pr[\max_i\{X_i\} \geq \beta/2]$ is bounded away from 0, the revenue from pricing every item at $\beta/2$ is $\Omega(\beta)$, hence the optimal revenue is also $\Omega(\beta)$. See Appendix F.3.1 for the precise lower bound we obtain. Such lower bound is useful as it implies that, if our transformation loses revenue that is a small fraction of β , this corresponds to a small fraction of optimal revenue lost.

¹⁶Clearly, by plugging $\epsilon = O(\frac{\hat{\epsilon}}{\log_2 1/\hat{\epsilon}})$ into our reduction, we obtain a reduction from $\text{PRICE}(\mathcal{V}, \hat{\epsilon})$ to $\text{PRICE}(\tilde{\mathcal{V}}, O(\frac{\hat{\epsilon}}{\log_2 1/\hat{\epsilon}}))$, for any desired $\hat{\epsilon}$. We phrased our theorem as a reduction from $\text{PRICE}(\mathcal{V}, c\epsilon \log_2(\frac{1}{\epsilon}))$ to $\text{PRICE}(\tilde{\mathcal{V}}, \epsilon)$ only to have better expressions in the supports of the \tilde{v}_i 's.

- Next, using (2) we show that, for all $\epsilon > 0$, if we restrict the prices to lie in the range $[\epsilon \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]$, we only lose a $O(\epsilon \log_2 1/\epsilon)$ fraction of the optimal revenue; this step is detailed in Appendix F.3.2.
- Finally, we show that we can efficiently transform the given MHR random variables $\{v_i\}_{i \in [n]}$ into a new collection of random variables $\{\tilde{v}_i\}_{i \in [n]}$ that take values in $[\frac{\epsilon}{2} \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]$ and satisfy the following: a near-optimal price vector for the setting where the buyer's values are distributed as $\{\tilde{v}_i\}_{i \in [n]}$ can be efficiently transformed into a near-optimal price vector for the original setting, i.e. where the buyer's values are distributed as $\{v_i\}_{i \in [n]}$. This step is detailed in Appendix F.3.3.

Theorem 3 is established by combining the reduction of Theorem 18 with our algorithm for bounded distributions of Theorem 2. See Appendix H.

9 Extreme Values of Regular Distributions

We reduce the problem of finding a near-optimal price vector for value distributions that are regular to finding a near-optimal price vector for value distributions that are supported on a bounded range $[u_{min}, u_{max}]$ satisfying $u_{max}/u_{min} \leq 16n^8/\epsilon^4$, where n is the number of distributions and ϵ is the desired approximation. It is important to notice that our bound on the ratio u_{max}/u_{min} does not depend on the distributions at hand, just their number and the required approximation. We also emphasize that the given regular distributions may be supported on $[0, +\infty)$, so it is a priori not clear if we can truncate these distributions to any finite set without losing substantial revenue. Our reduction is the following.

Theorem 20 (Reduction from Regular to Poly(n)-Bounded Distributions). *Let $\mathcal{V} = \{v_i\}_{i \in [n]}$ be a collection of mutually independent regular random variables. Then there exists some $\alpha = \alpha(\mathcal{V}) > 0$ such that, for any $\epsilon \in (0, 1)$, there is a reduction from $\text{PRICE}(\mathcal{V}, \epsilon)$ to $\text{PRICE}(\tilde{\mathcal{V}}, \epsilon - \Theta(\epsilon/n))$, where $\tilde{\mathcal{V}} = \{\tilde{v}_i\}_{i \in [n]}$ is a collection of mutually independent random variables that are supported on $[\frac{\epsilon\alpha}{4n^4}, \frac{4n^4\alpha}{\epsilon^3}]$.*

Moreover, α is efficiently computable from the distributions of the v_i 's, and, for all ϵ , the running time of the reduction is polynomial in the size of the input and $1/\epsilon$.

Our reduction is based on the following extreme value theorem for regular distributions, whose proof is provided in Appendix G.2. Immediately following the statement of the theorem we sketch how it is used to establish our reduction, whose detailed proof is in Appendix G.3. Section 9.1 gives other example applications of the theorem to illustrate its usefulness in bounding extreme values of regular distributions.

Theorem 21 (Homogenization of the Extreme Values of Regular Distributions). *Let $\{X_i\}_{i \in [n]}$ be a collection of mutually independent regular random variables, where $n \geq 2$. Then there exists some $\alpha = \alpha(\{X_i\}_{i \in [n]})$ such that:*

1. α has the following “anchoring” properties:
 - for all $\ell \geq 1$, $\Pr[X_i \geq \ell\alpha] \leq 2/(\ell n^3)$, for all $i \in [n]$;
 - $\alpha/n^3 \leq c \cdot \max_z (z \cdot \Pr[\max_i\{X_i\} \geq z])$, where c is an absolute constant.

2. For all $\epsilon \in (0, 1)$, the tails beyond $\frac{2n^2\alpha}{\epsilon^2}$ can be “homogenized”, i.e.

- for any integer $m \leq n$, thresholds $t_1, \dots, t_m \geq t \geq \frac{2n^2\alpha}{\epsilon^2}$, and index set $\{a_1, \dots, a_m\} \subseteq [n]$:

$$\sum_{i=1}^m t_i \Pr[X_{a_i} \geq t_i] \leq \left(t - \frac{2\alpha}{\epsilon}\right) \cdot \Pr\left[\max_{i \in [m]} \{X_{a_i}\} \geq t\right] + \frac{7\epsilon}{n} \cdot \left(\frac{2\alpha}{\epsilon} \cdot \Pr\left[\max_{i \in [m]} \{X_{a_i}\} \geq \frac{2\alpha}{\epsilon}\right]\right).$$

Furthermore, α is efficiently computable from the distributions of the X_i 's.

Given our homogenization theorem, our reduction of Theorem 20 is obtained as follows.

- First, we compute the threshold α specified in Theorem 21. This can be done efficiently as stated in Theorem 21. Now given the second anchoring property of α , we obtain an $\Omega(\alpha/n^3)$ lower bound to the optimal revenue. Such a lower bound is useful as it implies that we can ignore prices below some $O(\epsilon\alpha/n^3)$, without losing more than an ϵ -fraction of revenue.
- Next, using the homogenization part of Theorem 21, we show that, if we restrict a price vector to lie in $[\epsilon\alpha/n^4, 2n^2\alpha/\epsilon^2]^n$, we only lose a $O(\frac{\epsilon}{n})$ fraction of the optimal revenue. This step is detailed in Appendix G.3.1.
- Finally, we show that we can efficiently transform the input regular random variables $\{v_i\}_{i \in [n]}$ into a new collection of random variables $\{\tilde{v}_i\}_{i \in [n]}$ that are supported on $[\frac{\epsilon\alpha}{4n^4}, \frac{4n^4\alpha}{\epsilon^3}]$ and satisfy the following: a near-optimal price vector for when the buyer's values are distributed as $\{\tilde{v}_i\}_{i \in [n]}$ can be efficiently transformed into a near-optimal price vector for when the buyer's values are distributed as $\{v_i\}_{i \in [n]}$. This step is detailed in Appendix G.3.2, and Appendix G.3.3 concludes the proof of Theorem 20.

Theorem 4 is established by combining the reduction of Theorem 20 with our algorithm for bounded distributions of Theorem 2. See Appendix H.

9.1 Discussion of Theorem 21

We give a couple of applications of Theorem 21 to gain some intuition about its content:

- Suppose that we set all the t_i 's equal to $t \geq 2n^2\alpha/\epsilon^2$. In this case, the homogenization property of Theorem 21 implies that the union bound is essentially tight for t large enough, as $\frac{14\alpha}{tn} \cdot \Pr\left[\max_{i \in [m]} \{X_{a_i}\} \geq \frac{2\alpha}{\epsilon}\right]$ in the following calculation gets arbitrary close to 0:

$$\begin{aligned} \Pr\left[\max_{i \in [m]} \{X_{a_i}\} \geq t\right] &\leq \left(\sum_{i=1}^m \Pr[X_{a_i} \geq t]\right) \\ &\leq \left(\frac{t - \frac{2\alpha}{\epsilon}}{t}\right) \cdot \Pr\left[\max_{i \in [m]} \{X_{a_i}\} \geq t\right] + \frac{7\epsilon}{tn} \cdot \left(\frac{2\alpha}{\epsilon} \cdot \Pr\left[\max_{i \in [m]} \{X_{a_i}\} \geq \frac{2\alpha}{\epsilon}\right]\right) \\ &\leq \Pr\left[\max_{i \in [m]} \{X_{a_i}\} \geq t\right] + \frac{14\alpha}{tn} \cdot \Pr\left[\max_{i \in [m]} \{X_{a_i}\} \geq \frac{2\alpha}{\epsilon}\right]. \end{aligned}$$

This is not surprising, since for all i , the event $X_{a_i} \geq t$ only happens with tiny probability, by the anchoring property of α .

- Now let's try to set all the t_i 's to the same value $t' > t \geq 2n^2\alpha/\epsilon^2$. The homogenization property can be used to show that the probability of the event $\max_{i \in [m]} \{X_{a_i}\} \geq t'$ scales inverse proportionally with t' . Essentially this says that the tails of $\max_{i \in [m]} \{X_{a_i}\}$ are not fatter than those of the equal revenue distribution.¹⁷

$$\begin{aligned}
\Pr \left[\max_{i \in [m]} \{X_{a_i}\} \geq t' \right] &\leq \sum_{i=1}^m \Pr[X_{a_i} \geq t'] \\
&\leq \left(\frac{t - \frac{2\alpha}{\epsilon}}{t'} \right) \cdot \Pr \left[\max_{i \in [m]} \{X_{a_i}\} \geq t \right] + \frac{7\epsilon}{t'n} \cdot \left(\frac{2\alpha}{\epsilon} \cdot \Pr \left[\max_{i \in [m]} \{X_{a_i}\} \geq \frac{2\alpha}{\epsilon} \right] \right) \\
&\leq \frac{1}{t'} \cdot \left[t \cdot \Pr \left[\max_{i \in [m]} \{X_{a_i}\} \geq t \right] + \frac{7\epsilon}{n} \cdot \left(\frac{2\alpha}{\epsilon} \cdot \Pr \left[\max_{i \in [m]} \{X_{a_i}\} \geq \frac{2\alpha}{\epsilon} \right] \right) \right].
\end{aligned}$$

A similar bound would follow from Markov's inequality, if the expression inside the brackets were within a constant factor of $\mathbb{E}[\max_{i \in [m]} \{X_{a_i}\}]$. The result is interesting as it is possible for that expression to be much smaller than $\mathbb{E}[\max_{i \in [m]} \{X_{a_i}\}]$. For example, if $m = 1$ and X_{a_1} is distributed according to the equal revenue distribution, the expectation of X_{a_1} is $+\infty$, while the expression inside the brackets is $1 + \frac{7\epsilon}{n}$.

¹⁷Recall that the *equal revenue distribution* is supported on $[1, +\infty]$ and has cumulative density function $F(x) = 1 - \frac{1}{x}$.

Appendix

A Roadmap to the Appendix

Appendix B describes several computational models of accessing a value distribution, explaining what it means for an algorithm with each type of access to be “computationally efficient” or “take time polynomial in the input.”

Appendix C contains a formal description and analysis of our dynamic programming approach for RESTRICTEDPRICE, culminating in the proof of Theorem 13.

Appendix D provides several reductions among item pricing problems, whose goal is to discretize some aspect of the problem such as the support of the value distributions, the probabilities they assign to their support, or the set of available prices. The appendix culminates in the reductions of Theorems 15 and 17.

Appendix E provides a proof of Theorem 2, our algorithm for bounded distributions.

The rest of the appendix is dedicated to our treatment of MHR and regular distributions. Appendix F provides the proof of our extreme value theorem for MHR distributions (Theorem 19), as well as our reduction from item pricing problems with MHR distributions to item pricing problems with bounded distributions (Theorem 18). Similarly, Appendix G provides the proof of our extreme value theorem for regular distributions (Theorem 21), as well as our reduction from item pricing problems with regular distributions to item pricing problems with bounded distributions (Theorem 20). The proofs of our algorithmic results for MHR and regular distributions (Theorems 3 and 4) are provided in Appendix H. The proofs of our structural results for independent MHR and regular distributions (Theorems 5, 6 and 7) are provided in Appendix I. Finally, Appendix I.1 contains the proof of our structural result for i.i.d. MHR distributions (Theorem 8).

B Access to Value Distributions, and Computational Complexity

We consider three ways in which a distribution may be input to an algorithm, as well as what it means for the algorithm to run in time “polynomial in the description of the distribution” in each case.

- **Explicitly:** In this case, the distribution has to be discrete, and we are given its support (as a list of numbers), and the probabilities that the distribution places on every element in its support. If a distribution is explicitly input to an algorithm, the algorithm is computationally efficient if it runs in time polynomial in its other inputs and the bit-complexity of the numbers required to specify the distribution, i.e. the numbers in the support of the distribution and the probabilities assigned to them.
- **As an Oracle:** In this case, we are given (potentially black-box) access to a subroutine, called an *oracle*, that answers queries about the value of the cumulative distribution function on a queried point. In particular, a query to the oracle consists of a point x and a precision ϵ , and the oracle outputs a value of bit-complexity polynomial in the bit-complexity of x and ϵ , which is within ϵ from the value of the cumulative distribution function at point x . Moreover, we assume that we are given an *anchoring point* x^* such that the value of the cumulative distribution at that point is between two a priori known absolute constants c_1

and c_2 , such that $0 < c_1 < c_2 < 1$. Having such a point is necessary, as otherwise it would be computationally impossible to find any interesting point in the support of the distribution (i.e. any point where the cumulative is different than 0 or 1).

If a distribution is provided to an algorithm as an oracle, the algorithm is computationally efficient if it runs in time polynomial in its other inputs and the bit complexity of x^* , ignoring the time spent by the oracle to answer queries (since this is not under the algorithm’s control).

If, as it so happens in practice, we have a closed-form description of our input distribution, e.g. if our distribution is $\mathcal{N}(\mu, \sigma^2)$, we think of it as given to us as an oracle, answering queries of the form (x, ϵ) as specified above. In most common cases, such an oracle can be implemented so that it also runs efficiently in the bit-complexity of the query to the oracle.

- **Sample Access:** In this case, our only access to the distribution is our ability to take samples from it. It is easy to see that sample access to a distribution can be reduced to oracle access as follows. Suppose we have an algorithm \mathcal{A} designed to work with oracle access to a distribution, and let B be a bound on the total number of queries that the algorithm may make to the oracle. (B is always upper bounded by the running time of the algorithm.) Suppose now that instead of oracle access we have sample access to the distribution. Here is how we can fix this: For any query (x, ϵ) that \mathcal{A} needs to make to the oracle, we can simply take $\frac{1}{2\epsilon^2} \ln(\frac{2B}{\delta})$ samples from the distribution to estimate the cumulative distribution function at x . By Chernoff bounds, our estimate will have error greater than ϵ with probability at most $\frac{\delta}{B}$. So a union bound shows that all (at most B) queries of the algorithm will have error smaller than ϵ with probability at least $1 - \delta$. (We can tune this probability to be as close to one as we want at a cost of a factor of $\log \frac{1}{\delta}$ in the running time.) It is also easy to find an anchoring point. If we take many samples from the distribution and pick the median as the anchoring point, with very high probability the value of the cumulative distribution at this point is between $1/3$ and $2/3$.

Given the above, whenever we have sample access to a distribution we will pretend to have instead oracle access to it, and we will say that an algorithm is computationally efficient using the same criterion we used for oracle access.

Polynomial-Time Reductions Involving Value Distributions. This paper provides several polynomial-time reductions among item pricing problems. Recall from Section 2 that a reduction contains an algorithm \mathcal{A} that takes as input an instance of the item pricing problem, comprising distributions (and sometimes a restricted set of prices), and outputs another instance of the item pricing problem, comprising potentially different distributions (and prices). But what do we mean when we say that “an algorithm \mathcal{A} outputs a distribution F ?” The algorithm may either output an explicit description of the distribution or an oracle for it.¹⁸ In the former case, \mathcal{A} must enumerate the support of the distribution and specify the probabilities assigned to every point in the support, as required by the first bullet above. In the latter case, \mathcal{A} outputs an oracle for F , i.e. the description of an algorithm that satisfies the requirements of the second bullet above. This oracle may use as subroutines the oracles of the distributions provided in the input to \mathcal{A} , if any. We will then say that “ \mathcal{A} runs in polynomial time” if two properties are satisfied: 1. \mathcal{A} ’s running time is polynomial in its input; and 2. if \mathcal{A} outputs an oracle for some distribution \mathcal{F} , this oracle must

¹⁸Our reductions never output a distribution by providing sample access to it.

run in time polynomial in the description of the oracle and the input (x, ϵ) to the oracle, excluding the time spent in oracles (from the input to \mathcal{A}) that the oracle may use as subroutines.

C The Algorithm for Discrete Distributions

C.1 The Generic DP Step: Add an Item and Discretize Probabilities

In Section 5, we described our intended meaning for the Boolean function $g(i, \text{Pr})$. Here we explain how to compute g using dynamic programming. Our algorithm works bottom-up (i.e. from smaller to larger i 's), filling in g 's table so that the following recursive conditions are met.

- If $i > 1$, we set $g(i, \text{Pr}) = 1$ iff there is a price $p^{(j)}$ and a distribution Pr' so that the following hold:

1. $g(i - 1, \text{Pr}') = 1$.
2. Suppose that P_{i-1} is the price vector stored at cell $(i - 1, \text{Pr}')$ of the table, namely that under price vector P_{i-1} the (winning-value, winning-price) distribution for the prefix $1 \dots i - 1$ of the items is Pr' . What would happen if we assigned price $p^{(j)}$ to the i -th item? If the gap between the winning-value and winning-price among the first $i - 1$ items is larger than the gap between the value and price for the i -th item, the winning-value and winning-price would remain the same. Otherwise, they will become the value and price for the i -th item. Based on this, we can compute the resulting (winning-value, winning-price) distribution $\{\text{Pr}''_{i_1, i_2}\}_{i_1 \in [k_1], i_2 \in [k_2]}$ for the prefix $1 \dots i$ from just Pr'_{i_1, i_2} and the distribution \hat{F}_i of item i . Indeed:

$$\begin{aligned} \text{Pr}''_{i_1, i_2} = & \text{Pr}'_{i_1, i_2} \cdot \Pr_{v_i \sim \hat{F}_i} [v_i - p^{(j)} < v^{(i_1)} - p^{(i_2)}] \\ & + \left(\sum_{\substack{j_1, j_2 \\ \text{s.t. } v^{(j_1)} - p^{(j_2)} \\ \leq v^{(i_1)} - p^{(i_2)}}} \text{Pr}'_{j_1, j_2} \right) \cdot \Pr_{v_i \sim \hat{F}_i} [v_i = v^{(i_1)}] \cdot \mathbb{1}_{p^{(j)} = p^{(i_2)}}. \end{aligned} \quad (3)$$

We require that Pr is a rounded version of Pr'' computed as above, where all the probabilities are integer multiples of $\frac{1}{m}$. The rounding should be of the following canonical form. Setting $\delta_{i_1, i_2} = \text{Pr}''_{i_1, i_2} - \left\lfloor \frac{\text{Pr}''_{i_1, i_2}}{1/m} \right\rfloor \cdot \frac{1}{m}$, and $l = \left(\sum_{i_1 \in [k_1], i_2 \in [k_2]} \delta_{i_1, i_2} \right) / \left(\frac{1}{m} \right)$, we will round the first l probabilities in $\{\text{Pr}''_{i_1, i_2}\}_{i_1 \in [k_1], i_2 \in [k_2]}$ in some fixed lexicographic order up to the closest multiple of $\frac{1}{m}$, and round the rest down to the closest multiple of $\frac{1}{m}$.¹⁹

If Conditions 1 and 2 are met, we also store price vector $(P_{i-1}, p^{(j)})$ in cell $g(i, \text{Pr})$ of the table.

- To fill in the first slice of the table corresponding to $i = 1$, we use the same recursive definition given above, imagining that there is a slice $i = 0$, whose cells are all 0 except for those corresponding to the distributions Pr that satisfy: $\text{Pr}_{i_1, i_2} = 0$, for all i_1, i_2 , except for the lexicographically smallest $(i_1^*, i_2^*) \in \arg \min_{(k_1, k_2)} v^{(k_1)} - p^{(k_2)}$, where $\text{Pr}_{i_1^*, i_2^*} = 1$.

While we described the function g recursively above, we compute it iteratively from $i = 1$ through n .

¹⁹Any rounding would work. We use this one just to make the description of our algorithm explicit.

C.2 Proof of Theorem 13

In this appendix, we prove the correctness and running time of the algorithm presented in Section 5, providing a proof of Theorem 13. Intuitively, if we did not perform any rounding of distributions, our algorithm would have been *exact*, outputting an optimal price vector in $\{p^{(1)}, \dots, p^{(k_2)}\}^n$. We show next that the rounding is fine enough that it does not become detrimental to our revenue. To show this, we use the probabilistic concepts of *total variation distance* and *coupling of random variables*. Recall that the total variation distance between two distributions \mathbb{P} and \mathbb{Q} over a finite set \mathcal{A} is defined as follows

$$\|\mathbb{P} - \mathbb{Q}\|_{TV} = \frac{1}{2} \sum_{\alpha \in \mathcal{A}} |\mathbb{P}(\alpha) - \mathbb{Q}(\alpha)|.$$

Similarly, if X and Y are two random variables ranging over a finite set, their total variation distance, denoted $\|X - Y\|_{TV}$ is defined as the total variation distance between their distributions.

Proceeding to the correctness of our algorithm, let $P = (p_1, p_2, \dots, p_n) \in \{p^{(1)}, \dots, p^{(k_2)}\}^n$ be an arbitrary price vector. We can use this price vector to select n cells of our dynamic programming table, picking one cell per layer. The cells are those that the algorithm would have traversed if it made the decision of assigning price p_i to item i , for all i . Let us call the resulting cells $cell_1, cell_2, \dots, cell_n$.

For all i , we intend to compare the distributions $\left\{ \widehat{\Pr}_{i_1, i_2}^{(i)} \right\}_{i_1 \in [k_1], i_2 \in [k_2]}$ and $\left\{ \Pr_{i_1, i_2}^{(i)} \right\}_{i_1 \in [k_1], i_2 \in [k_2]}$,

which are respectively the (winning-value, winning-price) distribution:

- arising when the prefix $1 \dots i$ of items with distributions $\{\hat{F}_j\}_{j=1, \dots, i}$ is priced according to price vector (p_1, \dots, p_i) ;
- stored in $cell_i$ of the DP table.

The following lemma shows that these distributions have small total variation distance.

Lemma 22. *For all $i \in [n]$, $\|\Pr^{(i)} - \widehat{\Pr}^{(i)}\|_{TV} \leq ik_1k_2/m$.*

Proof. At a high level, our argument shows two properties for every i : (1) if rounding was not performed at step i of the DP algorithm, the distance between $\widehat{\Pr}^{(i)}$ and $\Pr^{(i)}$ would not increase compared to the distance between $\widehat{\Pr}^{(i-1)}$ and $\Pr^{(i-1)}$; (2) after the rounding is performed the distance increases by at most k_1k_2/m . Combining the two properties, we can prove the lemma.

Formally, we prove the lemma by induction. The base case is trivially true as $\Pr^{(1)}$ is just a rounding of $\widehat{\Pr}^{(1)}$ into probabilities that are multiples of $\frac{1}{m}$, whereby the probability of every point in the support is modified by no more than an additive $\frac{1}{m}$.

We proceed to show the inductive step. For convenience, for all i , let X_i be a random variable distributed according to $\Pr^{(i)}$, i.e. $\Pr[X_i = (v^{(i_1)}, p^{(i_2)})] = \Pr_{i_1, i_2}^{(i)}$ for all i_1, i_2 , and let \hat{X}_i be a random variable distributed according to $\widehat{\Pr}^{(i)}$.

Now suppose that the claim is true for i . We want to show that it holds for $i+1$. For this purpose we define an auxiliary random variable Z_{i+1} . Z_{i+1} is a function of the random variable X_i and an independent random variable \hat{v}_{i+1} distributed according to \hat{F}_{i+1} . If $\hat{v}_{i+1} - p_{i+1} \geq X_i(1) - X_i(2)$, we set $Z_{i+1} = (\hat{v}_{i+1}, p_{i+1})$, otherwise we set $Z_{i+1} = X_i$. Clearly, if we replaced X_i by \hat{X}_i in this definition, we would get a random variable with the same distribution as \hat{X}_{i+1} .

Now consider the following coupling of \hat{X}_{i+1} and Z_{i+1} . Use the optimal coupling of \hat{X}_i and X_i . Then generate both \hat{X}_{i+1} and Z_{i+1} using the above procedure with the same sample for \hat{v}_{i+1} . It is clear then that, conditioning on $X_i = \hat{X}_i$, $\hat{X}_{i+1} = Z_{i+1}$ with probability 1. So

$$\|\hat{X}_{i+1} - Z_{i+1}\|_{TV} \leq \Pr[\hat{X}_{i+1} \neq Z_{i+1}] \leq \Pr[X_i \neq \hat{X}_i] = \|X_i - \hat{X}_i\|_{TV}, \quad (4)$$

where the first inequality is true under any coupling, the second inequality is true for our particular coupling, and the last equality is true because we assumed an optimal coupling of X_i and \hat{X}_i .

On the other hand, we know that, if we round the distribution of Z_{i+1} into integer multiples of $1/m$, we will get the distribution of X_{i+1} . Therefore,

$$\|Z_{i+1} - X_{i+1}\|_{TV} \leq k_1 k_2 / m \quad (5)$$

Combining (4) and (5), the triangle inequality implies that $\|X_{i+1} - \hat{X}_{i+1}\|_{TV} \leq \|X_i - \hat{X}_i\|_{TV} + k_1 k_2 / m$, which completes the inductive step. \square

Proof of Theorem 13: Correctness: Let P^* be an optimal price vector for the instance of RESTRICTEDPRICE resulting after the reduction of Lemma 23 is applied to discretize the F_i 's into \hat{F}_i 's. Let $cell^*$ be the cell at layer n of the DP table corresponding to the price vector P^* . Lemma 22 implies that

$$\sum_{i_1 \in [k_1], i_2 \in [k_2]} |\Pr_{i_1, i_2}^{(n)} - \widehat{\Pr}_{i_1, i_2}^{(n)}| \leq n k_1 k_2 / m,$$

where $\widehat{\Pr}^{(n)}$ is the true (winning-value, winning-price) distribution corresponding to price vector P^* and $\Pr^{(n)}$ is the distribution stored in cell $cell^*$. Clearly, the expected revenues \mathcal{R}_{P^*} and \mathcal{R}_{cell^*} from these two distributions are related, as follows

$$|\mathcal{R}_{P^*} - \mathcal{R}_{cell^*}| \leq \sum_{i_1 \in [k_1], i_2 \in [k_2]} |\Pr_{i_1, i_2}^{(n)} - \widehat{\Pr}_{i_1, i_2}^{(n)}| \cdot p^{(i_2)} \leq \frac{n k_1 k_2}{m} \cdot \max_i \{p^{(i)}\}.$$

Now let $cell'$ be the cell at layer n of the DP table that has the highest expected revenue, and let P' be the price vector stored in $cell'$. Using the same notation as above, call $\mathcal{R}_{cell'}$ the revenue from the distribution stored at $cell'$ and $\mathcal{R}_{P'}$ the revenue from price vector P' . Then we have the following:

$$\mathcal{R}_{cell'} \geq \mathcal{R}_{cell^*}; \quad (\text{by the optimality of } cell') \quad (6)$$

$$|\mathcal{R}_{P'} - \mathcal{R}_{cell'}| \leq \frac{n k_1 k_2}{m} \cdot \max_i \{p^{(i)}\}. \quad (\text{using Lemma 22, as above}) \quad (7)$$

Putting all the above together, we obtain that

$$\mathcal{R}_{P'} \geq \mathcal{R}_{P^*} - \frac{2n k_1 k_2}{m} \cdot \max_i \{p^{(i)}\}. \quad (8)$$

Hence, the price vector P' output by our algorithm achieves revenue $\mathcal{R}_{P'}$ that is close to the optimal revenue \mathcal{R}_{P^*} for the discretized distributions $\{\hat{F}_i\}_i$. We now have to relate this revenue to the optimal revenue for the distributions $\{F_i\}_i$. So let us define the following quantities:

- $\mathcal{R}(P^*)$: the revenue achieved by price vector P^* in the original instance $\{F_i\}_i$;

- $\mathcal{R}(P')$: the revenue achieved by price vector P' in the original instance $\{F_i\}_i$.

Using Lemma 23 we easily see the following:

- $\mathcal{R}(P^*) \geq OPT - \frac{8k_1n}{m} \cdot \max_i\{p^{(i)}\}$;
- $\mathcal{R}(P') \geq \mathcal{R}_{P'} - \frac{4k_1n}{m} \cdot \max_i\{p^{(i)}\}$; and
- $\mathcal{R}_{P^*} \geq \mathcal{R}(P^*) - \frac{4k_1n}{m} \cdot \max_i\{p^{(i)}\}$.

Combining these with (8), we get

$$\begin{aligned} \mathcal{R}(P') &\geq \mathcal{R}_{P^*} - \frac{2nk_1k_2}{m} \cdot \max_i\{p^{(i)}\} - \frac{4k_1n}{m} \cdot \max_i\{p^{(i)}\} \\ &\geq \mathcal{R}(P^*) - \frac{2nk_1k_2}{m} \cdot \max_i\{p^{(i)}\} - \frac{8k_1n}{m} \cdot \max_i\{p^{(i)}\} \\ &\geq OPT - \frac{(2nk_1k_2 + 16k_1n)}{m} \cdot \max_i\{p^{(i)}\}. \end{aligned}$$

Running Time: Recall that both the support $\mathcal{S} = \{v^{(1)}, v^{(2)}, \dots, v^{(k_1)}\}$ of the value distributions and the set $\mathcal{P} := \{p^{(1)}, \dots, p^{(k_2)}\}$ of prices are explicitly part of the input to our algorithm.

Given this, the reduction of Lemma 23 (used as the first step of our algorithm) takes time polynomial in the size of the input and $\log m$. After this reduction is carried out, the value distributions $\{\hat{F}_i\}_i$ that are provided as input to the dynamic programming algorithm are known explicitly and the probabilities they assign to every point in \mathcal{S} are integer multiples of $\frac{1}{m}$.

We proceed to bound the running time of the dynamic programming algorithm. First, it is easy to see that its table has at most $n \times (m+1)^{k_1k_2}$ cells, since there are n possible choices for i and $m+1$ possible values for each \Pr_{i_1, i_2} . Our DP computation proceeds iteratively from layer $i = 1$ to layer $i = n$ of the table. For every cell of layer i , there are at most k_2 different prices we can assign to the next item $i+1$. For every such price we need to compute a distribution using Eq. (3) and then round that distribution. Hence, the total work we need to do per cell of layer i is polynomial in the input size and $\log m$, since our computation involves probabilities that are integer multiples of $\frac{1}{m}$. Indeed the probability distributions maintained in the DP table use probabilities that are integer multiples of $\frac{1}{m}$, and recall that the distributions \hat{F}_i also use probabilities in multiples of $\frac{1}{m}$. Hence, the total time we need to spend to fill up the whole table is polynomial in the size of the input and $m^{k_1k_2}$. In the last phase of the algorithm, we exhaustively search for the cell of layer n with the highest expected revenue. This costs time polynomial in the size of the input and $m^{k_1k_2}$, since there are $m^{O(k_1k_2)}$ cells at layer n , and the expected revenue computation for each cell can be done in time polynomial in the input size and $\log m$. Overall, the running time of the algorithm is polynomial in the size of the input and $m^{k_1k_2}$. \square

D Our Discretization Results

This appendix provides various reductions among item pricing problems. All reductions discretize some aspect of a given item pricing problem. This could be the set of allowable prices, the support of the value distributions, or the probabilities assigned by these distributions to the points in their support. We will bound the loss in approximation resulting from each reduction. This bound is useful in telling us how much revenue we are losing if we solve the discretized problem instead of the given problem.

D.1 Discretization of Probabilities

The following lemma allows us to discretize the probabilities assigned by value distributions supported on a discrete set to points in their support.

Lemma 23 (Probability Discretization). *Suppose we are given a collection of mutually independent random variables $\{v_i\}_{i \in [n]}$ supported on a discrete set $S = \{s_1, \dots, s_k\} \subset \mathbb{R}_{\geq 0}$, an interval $[p_{\min}, p_{\max}] \subset \mathbb{R}_{\geq 0}$ of possible prices, and an integer $m \geq 2k$.²⁰ In polynomial-time we can construct another collection of mutually independent random variables $\{v'_i\}_{i \in [n]}$ whose distributions are supported on the same set S but only use probabilities that are integer multiples of $1/m$. The distributions of the v'_i 's are computed explicitly. Moreover, for any price vector $P \in [p_{\min}, p_{\max}]^n$, the difference in expected revenue from the two collections of random variables is upper bounded $\frac{4kn}{m} \cdot p_{\max}$.*

Proof of Lemma 23: If we know the distributions of the v_i 's explicitly, then, for all i , we construct the distribution of v'_i as follows. Let $\pi_{s_j} = \Pr[v_i = s_j]$ and $\pi'_{s_j} = \Pr[v'_i = s_j]$. For all $j \geq 2$, round all π_{s_j} down to the nearest integer multiple of $1/m$ to get π'_{s_j} . We then round π_{s_1} up to get π'_{s_1} to guarantee that π' is still a distribution. We use δ_{s_j} to denote the rounding error at s_j .

As the total variation distance between the distribution of v_i and v'_i is $\frac{1}{2} \sum_{j=1}^m \delta_{s_j} \leq \frac{k}{m}$, we can couple v_i and v'_i so that $\Pr[v_i \neq v'_i] \leq \frac{k}{m}$. Now taking a union bound over all i , the probability that the vector $v = (v_1, v_2, \dots, v_n)$ is different from $v' = (v'_1, v'_2, \dots, v'_n)$ is at most $\frac{kn}{m}$. In other words, with probability at least $1 - \frac{kn}{m}$, $v = v'$. Clearly, for all draws from the distribution such that $v = v'$, the revenues are the same. When $v \neq v'$, the difference between the revenues is at most p_{\max} , since $P \in [p_{\min}, p_{\max}]^n$. And this only happens with probability at most $\frac{kn}{m}$. Therefore, the difference between the expected revenues under the two distributions should be no greater than $\frac{kn}{m} \cdot p_{\max}$.

Clearly, we can compute the distributions of the v'_i 's in time polynomial in $n, k, \log m$ and the description complexity of the distributions of the variables v_i 's, if these distributions are given to us explicitly. If we have oracle access to the distributions of the v_i 's we can query our oracle with high enough precision, say $1/m$, to obtain a function $g_i : S \rightarrow [0, 1]$ that satisfies $\sum_{x \in S} g_i(x) = 1 \pm \frac{k}{m}$. Using the normalized g_i as a proxy for the distribution of v_i we can follow the algorithm outlined above to define the distribution of v'_i . It is not hard to argue that the total variation distance between v_i and v'_i can be bounded by $\frac{4k}{m}$. Hence, we can couple v_i and v'_i so that $\Pr[v_i \neq v'_i] \leq \frac{4k}{m}$ and proceed as above. \square

D.2 Discretization of Prices

In this appendix, we present several easy lemmas that can be used to restrict the search space for a (nearly-)optimal price vector. First, it is straightforward to see that, when the value distributions are supported in some range $[u_{\min}, u_{\max}]$, it is sufficient to only consider prices that lie in the same range, without any sacrifice in revenue.

Lemma 24 (Price Restriction). *In an instance of the item pricing problem, suppose that the values are independently distributed in some range $[u_{\min}, u_{\max}]$. Let also $P = (p_1, \dots, p_n)$ be an arbitrary price vector, and suppose that we modify P into a new price vector P' as follows: for all i , set*

²⁰It is assumed that S, p_{\min}, p_{\max} and m are given explicitly. We may have any access to the value distributions (as discussed in Appendix B).

$p'_i = u_{max}$, if $p_i > u_{max}$; set $p'_i = u_{min}$, if $p_i < u_{min}$; otherwise set $p'_i = p_i$. The expected revenues \mathcal{R}_P and $\mathcal{R}_{P'}$ achieved by the price vectors P and P' respectively satisfy $\mathcal{R}_{P'} \geq \mathcal{R}_P$.

Proof of Lemma 24: Let us do the modification in two steps. We first increase the prices that are below u_{min} to u_{min} , and then decrease the prices that are above u_{max} to u_{max} . We will show that each step will not decrease the expected revenue.

Let us increase the low prices first, and call P'' the resulting price vector. For every sample (v_1, \dots, v_n) from the value distributions, if the buyer makes the same decision under P and P'' , his price under P'' is at least as high as under P . If the buyer makes a different decision under P and P'' , it must be that, under P , the buyer is buying an item priced less than u_{min} and, because the price of that item was increased to u_{min} in P'' , the buyer prefers to buy a different item. In this case, the buyer was paying less than u_{min} under P and is paying at least u_{min} under P'' .

Now let us decrease the prices in P'' that are above u_{max} to u_{max} to obtain P' , and let us compare the price that a buyer will pay under these two price vectors. Whenever the buyer does not buy anything under P' , he is also not buying anything under P'' , as the items under P' are at least as cheap. Also notice that the only items whose value-minus-price gap is different under P'' and P' are those that are priced above u_{max} in P'' and u_{max} in P' , and their gap increases in P' . So if the buyer buys different items under P'' and P' , then the buyer is paying u_{max} under P' and at most u_{max} under P'' . \square

Combining Lemma 24 with a price discretization lemma attributed to Nisan [CHK07], we can restrict the set of prices to a set of cardinality $O(\frac{\log u_{max}/u_{min}}{\epsilon^2})$, if $[u_{min}, u_{max}]$ is the support of the value distributions.

Lemma 25 (Price Discretization). *Suppose that the value distributions in an instance of the item pricing problem are independent and supported on $[u_{min}, u_{max}] \subset \mathbb{R}_+$. For any $\epsilon \in (0, 1/2)$, consider the following finite set of prices:*

$$\mathcal{P}_\epsilon = \left\{ p \mid p = \frac{1 + \epsilon^2 - \epsilon}{(1 - \epsilon^2)^i} \cdot u_{min}, \quad i \in \left[\left\lceil \log_{\frac{1}{(1-\epsilon^2)}} (u_{max}/u_{min}) \right\rceil \right] \right\}.$$

For any price vector $P \in [u_{min}, u_{max}]^n$, there exists a price vector P' such that $p'_i \in \mathcal{P}_\epsilon$ and $p'_i \in [1 - \epsilon, 1 + \epsilon^2 - \epsilon] \cdot p_i$, for all i . The expected revenue achieved by the two price vectors satisfies $\mathcal{R}_{P'} \geq (1 - 2\epsilon)\mathcal{R}_P$.

Proof of Lemma 25: Our proof exploits the following lemma, attributed in [CHK07] to Nisan.

Lemma 26. *Let $\epsilon \in (0, 1)$, and let P, P' be price vectors satisfying $p'_i \in [1 - \epsilon, 1 + \epsilon^2 - \epsilon] \cdot p_i$, for all i . Then the expected revenues achieved by the two price vectors in an instance of the item pricing problem satisfy $\mathcal{R}_{P'} \geq (1 - 2\epsilon)\mathcal{R}_P$.*

To prove Lemma 25, define for every p_i :

$$p'_i = \frac{1 + \epsilon^2 - \epsilon}{(1 - \epsilon^2)^{\lceil \log_{1/(1-\epsilon^2)}(p_i/u_{min}) \rceil}} \cdot u_{min}.$$

Observe that

$$\frac{1}{(1 - \epsilon^2)^{\lceil \log_{1/(1-\epsilon^2)}(p_i/u_{min}) \rceil}} \cdot u_{min} \in [1 - \epsilon^2, 1] \cdot p_i.$$

On the other hand, $(1 - \epsilon^2)(1 + \epsilon^2 - \epsilon) = 1 - \epsilon + \epsilon^3 - \epsilon^4 \geq 1 - \epsilon$. Thus, $p'_i \in [1 - \epsilon, 1 + \epsilon^2 - \epsilon] \cdot p_i$, for all i . Now Lemma 26 implies that $\mathcal{R}_{P'} \geq (1 - 2\epsilon)R_P$. \square

We conclude with a lemma with a similar flavor as Lemma 24.

Lemma 27. *Let $\alpha > 0$ be arbitrary, let P be any price vector, and define P' as follows: set $p'_i = p_i$, if $p_i \geq \alpha$, and $p'_i = \alpha$ otherwise. Then the expected revenues \mathcal{R}_P and $\mathcal{R}_{P'}$ from these price vectors in an instance of the item pricing problem satisfy $\mathcal{R}_{P'} \geq \mathcal{R}_P - \alpha$.*

Proof of Lemma 27: Let $S_{low} = \{i \mid p_i < \alpha\}$ and fix the buyer's values for the items. The only case where the buyer's behavior is different under P and P' is when the buyer is buying some item S_{low} under P , as these are the only items whose value-minus-price gap changed/decreased from P to P' . So the difference in revenue is bounded by the contribution to R_P of items in S_{low} , which is no greater than α . \square

D.3 Discretization of Values

In this appendix, we establish polynomial-time reductions, discretizing the support of the value distributions in the input to the item pricing problem. Our reductions are specialized depending on whether we want to achieve multiplicative (Lemma 30) or additive (Lemma 31) approximations to the optimal revenue. Both reductions are enabled by an extension of Nisan's lemma to value distributions, summarized in Lemma 28.

Lemma 28. *Let $\{v_i\}_{i \in [n]}$ and $\{\hat{v}_i\}_{i \in [n]}$ be two collections of mutually independent random variables, where all v_i 's are supported on a common set $[u_{min}, u_{max}] \subset \mathbb{R}_+$, and let $r = u_{max}/u_{min}$. Let also $\delta \in \left(0, \frac{1}{(4^{\lceil \log_2 r \rceil})^{1/(2a-1)}}\right]$, where $a \in (1/2, 1)$, and suppose that we can couple the two collections of random variables so that, for all $i \in [n]$, $\hat{v}_i \in [1 + \delta - \delta^2, 1 + \delta] \cdot v_i$ with probability 1. Finally, let \mathcal{R}_{OPT} be the optimal expected revenue from any price vector when the buyer's values are $\{v_i\}_{i \in [n]}$. Then, for any price vector $P \in [u_{min}, u_{max}]^n$, such that $\mathcal{R}_P(\{v_i\}_i) \geq \mathcal{R}_{OPT}/2$, it holds that*

$$\mathcal{R}_P(\{\hat{v}_i\}_i) \geq (1 - 3\delta^{1-a})\mathcal{R}_P(\{v_i\}_i),$$

where $\mathcal{R}_P(\{v_i\}_i)$ is the expected revenue under price vector P when the values are $\{v_i\}_{i \in [n]}$, while $\mathcal{R}_P(\{\hat{v}_i\}_i)$ is the revenue under P when the values are $\{\hat{v}_i\}_{i \in [n]}$.

Proof of Lemma 28: For notational convenience, throughout this proof we use $\mathcal{R}_P := \mathcal{R}_P(\{v_i\}_i)$ and $\hat{\mathcal{R}}_P := \mathcal{R}_P(\{\hat{v}_i\}_i)$.

Consider now the joint distribution of $\{v_i\}_{i \in [n]}$ and $\{\hat{v}_i\}_{i \in [n]}$ satisfying $\hat{v}_i \in [1 + \delta - \delta^2, 1 + \delta] \cdot v_i$, for all i , with probability 1. For every point in the support of the joint distribution, we show that the revenue of the seller under price vector P is not much larger in ‘‘Scenario A’’, where the values of the buyer are $\{v_i\}_{i \in [n]}$, than in ‘‘Scenario B’’, where the values are $\{\hat{v}_i\}_{i \in [n]}$. In particular, we argue first that the price paid in ‘‘Scenario A’’ is at most $\delta \cdot u_{max}$ larger than the price paid in ‘‘Scenario B,’’ with probability 1. Indeed, for every point in the support of the joint distribution, we distinguish two cases:

1. The items sold are the same in the two scenarios. In this case, the winning prices are also the same.

2. The items sold are different in the two scenarios. In this case, we show that the winning prices are close. Since \hat{v}_i is greater than v_i for all i , if there is a winner (item) in Scenario A, there is a winner in Scenario B. Let i be the winner in Scenario A, and j be the winner in Scenario B. We have the following two inequalities:

$$\begin{aligned} v_i - p_i &\geq v_j - p_j \\ \hat{v}_j - p_j &\geq \hat{v}_i - p_i \end{aligned}$$

The two inequalities imply that

$$\hat{v}_j - v_j \geq \hat{v}_i - v_i.$$

Since $\hat{v}_j \in [1 + \delta - \delta^2, 1 + \delta] \cdot v_j$, it follows that $\hat{v}_j - v_j \leq \delta \cdot v_j$. Using the same starting condition for i , we can show that $\hat{v}_i - v_i \geq (\delta - \delta^2) \cdot v_i$.

Hence,

$$\delta \cdot v_j \geq (\delta - \delta^2) \cdot v_i.$$

Also we know that

$$p_j \geq p_i + v_j - v_i.$$

Therefore,

$$p_j \geq p_i + v_j - v_i \geq p_i + (1 - \delta) \cdot v_i - v_i = p_i - \delta \cdot v_i. \quad (9)$$

The above establishes that with probability 1 the price paid in ‘‘Scenario A’’ is larger than that paid in ‘‘Scenario B’’ by at most an *additive* δu_{max} . We proceed to convert this additive approximation guarantee into a multiplicative approximation guarantee. Observe that whenever $p_i \geq \delta^a v_i$, $p_i - \delta \cdot v_i \geq (1 - \delta^{1-a}) p_i$. Hence, if we can show that most of the revenue \mathcal{R}_P is contributed by value-price pairs (v_i, p_i) satisfying $p_i \geq \delta^a v_i$, we can convert our additive approximation to a $(1 - \delta^{1-a})$ multiplicative approximation. Indeed, we argue next that when a price vector P satisfies $\mathcal{R}_P \geq \mathcal{R}_{OPT}/2$, the contribution to the revenue from the event

$$S = \{\text{the sold item } k \text{ satisfies } p_k < \delta^a v_k\}$$

is small. More precisely,

Proposition 29. *If $\mathcal{R}_P \geq \mathcal{R}_{OPT}/2$, then the contribution to \mathcal{R}_P from the event S is no greater than $2\delta^{1-a}\mathcal{R}_P$.*

Proof. The proof is by contradiction. For all $i \in [\lceil \log_2 r \rceil]$, define the event

$$S_i = \{(\text{the sold item } k \text{ has price } p_k < \delta^a v_k) \wedge (p_k \in [2^{i-1}u_{min}, 2^i u_{min}))\}.$$

Note that S_i and S_j are disjoint for all $i \neq j$. Let $n_p = \lceil \log_2 r \rceil$ and note that $S = \cup_{i=1}^{n_p} S_i$.²¹ Assuming that the contribution to \mathcal{R}_P from the event S is larger than $2\delta^{1-a}\mathcal{R}_P$, there must exist some i such that the contribution to \mathcal{R}_P from S_i is at least $2\delta^{1-a}\mathcal{R}_P/n_p \geq \delta^{1-a}\mathcal{R}_{OPT}/n_p$. For this i , let us modify the price vector P to P' in the following fashion:

$$p'_k = \begin{cases} +\infty & p_k \notin [2^{i-1}u_{min}, 2^i u_{min}) \\ \frac{2^{i-1}u_{min}}{\delta^a} & \text{otherwise} \end{cases}$$

²¹To be more accurate, replace the set $[2^{i-1}u_{min}, 2^i u_{min})$ by $[2^{i-1}u_{min}, 2^i u_{min}]$ for the definition of the event S_{n_p} .

We claim that for all outcomes $(v_1, v_2, \dots, v_n) \in S_i$, there always exists an item sold under P' . Indeed, let k be the winner under P . Then $p_k < \delta^a v_k$. By the definition of p'_k , we know that

$$p'_k = \frac{2^{i-1} u_{\min}}{\delta^a} \leq p_k / \delta^a < v_k.$$

Thus, an item has to be sold. Moreover, the sold item has price $\frac{2^{i-1} u_{\min}}{\delta^a}$, as all the other prices are set to $+\infty$. Hence, we can lower bound $\mathcal{R}_{P'}$ as follows

$$\mathcal{R}_{P'} \geq \Pr[S_i] \cdot \frac{2^{i-1} u_{\min}}{\delta^a} \geq \frac{\text{Contribution of } S_i \text{ to } \mathcal{R}_P}{2\delta^a} \geq \frac{\delta^{1-a} \mathcal{R}_{OPT}}{2n_p \delta^a}.$$

Given that $\delta \leq (\frac{1}{4n_p})^{1/(2a-1)}$, the above implies $\mathcal{R}_{P'} \geq 2\mathcal{R}_{OPT}$, which is impossible, i.e. we get a contradiction. This concludes the proof of the proposition. \square

Given the proposition, at least $(1 - 2\delta^{1-a})$ fraction of \mathcal{R}_P is contributed by value-price pairs (v_i, p_i) satisfying $p_i \geq \delta^a v_i$. Recalling our earlier discussion, this implies that $\hat{\mathcal{R}}_P \geq (1 - 2\delta^{1-a})(1 - \delta^{1-a})\mathcal{R}_P \geq (1 - 3\delta^{1-a})\mathcal{R}_P$. \square

Lemma 28 enables polynomial-time reductions from value distributions supported on some bounded range $[u_{\min}, u_{\max}]$ to value distributions supported on some discrete set of cardinality $O(\log r)$, where $r = u_{\max}/u_{\min}$. We provide two reductions (Lemmas 30 and 31) depending on whether the approximation to the optimal revenue is intended to be additive or multiplicative. We note that a straightforward extension of Nisan's lemma to value distributions would have resulted in supports of cardinality $O(r^2 \log r)$. The exponential improvement in the size of the support comes from our more intricate extension obtained by Lemma 28.

Lemma 30 (Value Discretization for Multiplicative Approximations). *Let $\{v_i\}_{i \in [n]}$ be a collection of mutually independent random variables supported on a bounded range $[u_{\min}, u_{\max}] \subset \mathbb{R}_+$, and $r = \frac{u_{\max}}{u_{\min}}$. For any $\delta \in (0, \frac{1}{(4 \lceil \log_2 r \rceil)^{4/3}})$, there exists a collection of mutually independent random variables $\{\hat{v}_i\}_{i \in [n]}$, which are supported on a discrete set of cardinality $O(\frac{\log r}{\delta^2})$ and satisfy the following properties.*

1. *The optimal revenue when the buyer's values are $\{\hat{v}_i\}_{i \in [n]}$ is at least a $(1 - 3\delta^{1/8})$ -fraction of the optimal revenue when the values are $\{v_i\}_{i \in [n]}$. I.e. $\hat{\mathcal{R}}_{OPT} \geq (1 - 3\delta^{1/8})\mathcal{R}_{OPT}$, where $\mathcal{R}_{OPT} = \max_P \mathcal{R}_P(\{v_i\}_i)$ and $\hat{\mathcal{R}}_{OPT} = \max_P \mathcal{R}_P(\{\hat{v}_i\}_i)$.*
2. *Moreover, for any constant $\rho \in (0, 1/2)$ and any price vector P such that $\mathcal{R}_P(\{\hat{v}_i\}_i) \geq (1 - \rho)\hat{\mathcal{R}}_{OPT}$, we can construct in time polynomial in the description of P and $1/\delta$ another price vector \tilde{P} such that $\mathcal{R}_{\tilde{P}}(\{v_i\}_i) \geq (1 - 7\delta^{1/8} - \rho)\mathcal{R}_{OPT}$.*

If u_{\min} and u_{\max} are provided explicitly as input to the reduction,²² we can compute the distributions of the \hat{v}_i 's²³ and their support in time polynomial in the description of $\{v_i\}_{i \in [n]}$, $\langle u_{\min} \rangle$, $\langle u_{\max} \rangle$ and $1/\delta$.

²²This requirement is only relevant if we have oracle access to the distributions of the v_i 's, as if we are given the distributions explicitly we immediately also know u_{\min} and u_{\max} .

²³The \hat{v}_i 's will inherit the same type of access that we have to the distributions of the v_i 's, according to Appendix B. In particular, if the v_i 's are specified explicitly in the input to the reduction then the \hat{v}_i 's will also be specified explicitly in the output of the reduction; if the v_i 's are given as oracles then the \hat{v}_i 's will be given as oracles; etc.

Proof of Lemma 30: Informally, our reduction establishes the following properties: (1) Suppose that we transform a buyer with arbitrary valuations (call this buyer “Buyer A”) to a buyer with discrete valuations (called “Buyer B”) by first multiplying each of Buyer A’s values by $(1 + \delta)$ and then rounding them down to the closest real of the form $(1 + \delta)(1 + \xi)^j u_{min}$, for some integer j , where δ is fixed and $\xi = \frac{\delta^2}{1 + \delta - \delta^2}$. We show that the optimal revenue from Buyer B is very close to the optimal revenue from Buyer A by exploiting that Buyer B’s values have been boosted, using Lemma 28. (2) For the reduction to be computationally useful, we also show that given an approximately optimal price vector for Buyer B, if we divide all prices by $(1 + \delta)(1 + \delta - \delta^2)$, the new price vector will be an approximately optimal price vector for Buyer A. Intuitively, scaling down the prices undoes the effect of boosting the values.

We proceed to make the above plan precise, beginning with the description of the random variables $\{\hat{v}_i\}_{i \in [n]}$. We will use $\{F_i\}_{i \in [n]}$ and $\{\hat{F}_i\}_{i \in [n]}$ to denote respectively the cumulative distribution functions of the variables $\{v_i\}_{i \in [n]}$ and $\{\hat{v}_i\}_{i \in [n]}$. Our variables $\{\hat{v}_i\}_{i \in [n]}$ will only be supported on the set

$$\left\{ a_j = (1 + \delta)(1 + \xi)^j u_{min} \mid j \in \left\{ 0, \dots, \lfloor \log_{1+\xi} \frac{u_{max}}{u_{min}} \rfloor \right\} \right\}.$$

Moreover, for all i , \hat{F}_i is defined in terms of F_i as follows:

$$\hat{F}_i(a_j) = F_i(a_j/(1 + \delta - \delta^2)) - F_i(a_j/(1 + \delta)) + \hat{F}_i(a_{j-1}) \mathbb{1}_{j>0}, \quad \forall j.$$

Now, for all i , we couple v_i with \hat{v}_i as follows: If $v_i \in [a_j/(1 + \delta), a_j/(1 + \delta - \delta^2))$, we set $\hat{v}_i = a_j$. Given our definition of the \hat{F}_i ’s, this defines a valid coupling of the collections $\mathcal{V} = \{v_i\}_i$ and $\hat{\mathcal{V}} = \{\hat{v}_i\}_i$. Moreover, by definition, our coupling satisfies

$$\hat{v}_i \in [1 + \delta - \delta^2, 1 + \delta] \cdot v_i, \forall i, \tag{10}$$

with probability 1, and all the \hat{v}_i ’s are supported on $[(1 + \delta)u_{min}, (1 + \delta)u_{max}]$.

We are now ready to establish the first part of the lemma. Using Lemma 28 and the property of our coupling it follows immediately that

$$\mathcal{R}_P(\hat{\mathcal{V}}) \geq (1 - 3\delta^{1/8})\mathcal{R}_P(\mathcal{V}),$$

for any price vector $P \in [u_{min}, u_{max}]^n$ s.t. $\mathcal{R}_P(\mathcal{V}) \geq \frac{1}{2}\mathcal{R}_{OPT}$. Lemma 24 implies that the optimal revenue for \mathcal{V} is achieved by some price vector in $[u_{min}, u_{max}]^n$. Hence, we get from the above that $\hat{\mathcal{R}}_{OPT} \geq (1 - 3\delta^{1/8})\mathcal{R}_{OPT}$.

We proceed to show the second part of the lemma. We do this by defining another collection of random variables $\tilde{\mathcal{V}} = \{\tilde{v}_i\}_{i \in [n]}$. These are defined implicitly via the following coupling between $\{\tilde{v}_i\}_{i \in [n]}$ and $\{\hat{v}_i\}_{i \in [n]}$: for all i , we set

$$\tilde{v}_i = \frac{\hat{v}_i}{(1 + \delta - \delta^2)(1 + \delta)}.$$

It follows that the \tilde{v}_i ’s are supported on $[u_{min}/(1 + \delta - \delta^2), u_{max}/(1 + \delta - \delta^2)]$.

Moreover, for any price vector P , let us construct another price vector \tilde{P} as follows:

$$\tilde{p}_i = \frac{p_i}{(1 + \delta - \delta^2)(1 + \delta)}. \tag{11}$$

Under our coupling between $\{\tilde{v}_i\}_{i \in [n]}$ and $\{\hat{v}_i\}_{i \in [n]}$, it is not hard to see that if we use price vector P when the buyer's values are $\{\hat{v}_i\}_{i \in [n]}$ and price vector \tilde{P} when the buyer's values are $\{\tilde{v}_i\}_{i \in [n]}$, then the index of the item that the buyer buys is the same in the two cases, with probability 1. Hence:

$$\mathcal{R}_{\tilde{P}}(\tilde{\mathcal{V}}) = \frac{\mathcal{R}_P(\hat{\mathcal{V}})}{(1 + \delta - \delta^2)(1 + \delta)}. \quad (12)$$

This follows from the fact that both \tilde{P} and $\{\tilde{v}_i\}_{i \in [n]}$ are the same linear transformations of P and $\{\hat{v}_i\}_{i \in [n]}$ respectively.

Composing the coupling between v_i and \hat{v}_i and the coupling between \hat{v}_i with \tilde{v}_i , we obtain a coupling between v_i and \tilde{v}_i . We show that this coupling satisfies $v_i \in [1 + \delta - \delta^2, 1 + \delta] \cdot \tilde{v}_i$, with probability 1. Since $(1 + \delta - \delta^2)v_i \leq \hat{v}_i \leq (1 + \delta)v_i$, it follows that

$$v_i/(1 + \delta) \leq \hat{v}_i/(1 + \delta - \delta^2)(1 + \delta) = \tilde{v}_i \leq v_i/(1 + \delta - \delta^2).$$

Hence

$$v_i \in [1 + \delta - \delta^2, 1 + \delta] \cdot \tilde{v}_i, \forall i, \quad (13)$$

with probability 1. Now an application of Lemma 28 implies that, for any price vector $\tilde{P} \in [u_{min}/(1 + \delta - \delta^2), u_{max}/(1 + \delta - \delta^2)]^n$ satisfying $\mathcal{R}_{\tilde{P}}(\tilde{\mathcal{V}}) \geq \frac{1}{2}\mathcal{R}_{OPT}(\tilde{\mathcal{V}})$:

$$\mathcal{R}_{\tilde{P}}(\mathcal{V}) \geq (1 - 3\delta^{1/8})\mathcal{R}_{\tilde{P}}(\tilde{\mathcal{V}}). \quad (14)$$

Now let P be a price vector satisfying $\mathcal{R}_P(\hat{\mathcal{V}}) \geq (1 - \rho)\hat{\mathcal{R}}_{OPT}$. Lemma 24 implies that WLOG we can assume that $P \in [(1 + \delta)u_{min}, (1 + \delta)u_{max}]^n$ (as if the given price vector is not in this set, we can efficiently convert it into one that is in this set without losing any revenue). Then the vector \tilde{P} obtained from P via Eq. (11) is in $[u_{min}/(1 + \delta - \delta^2), u_{max}/(1 + \delta - \delta^2)]^n$, and clearly satisfies $\mathcal{R}_{\tilde{P}}(\tilde{\mathcal{V}}) \geq (1 - \rho)\mathcal{R}_{OPT}(\tilde{\mathcal{V}})$, as \tilde{P} and $\tilde{\mathcal{V}}$ are the same linear transformations of P and $\hat{\mathcal{V}}$ respectively. Hence, Equations (12) and (14) give

$$\begin{aligned} \mathcal{R}_{\tilde{P}}(\mathcal{V}) &\geq \left((1 - 3\delta^{1/8}) / (1 + \delta)(1 + \delta - \delta^2) \right) \mathcal{R}_P(\hat{\mathcal{V}}) \\ &\geq (1 - 3\delta^{1/8})(1 - 2\delta)\mathcal{R}_P(\hat{\mathcal{V}}) \\ &\geq (1 - 4\delta^{1/8})\mathcal{R}_P(\hat{\mathcal{V}}) \\ &\geq (1 - 4\delta^{1/8})(1 - \rho)\hat{\mathcal{R}}_{OPT} \\ &\geq (1 - 4\delta^{1/8})(1 - \rho)(1 - 3\delta^{1/8})\mathcal{R}_{OPT} \quad (\text{using the first part of the theorem}) \\ &\geq (1 - 7\delta^{1/8} - \rho)\mathcal{R}_{OPT}. \end{aligned}$$

□

Lemma 31 (Value Discretization for Additive Approximations). *Let $\{v_i\}_{i \in [n]}$ be a collection of mutually independent random variables supported on a bounded range $[u_{min}, u_{max}] \subset \mathbb{R}_+$, and $r = \frac{u_{max}}{u_{min}}$. For any $\delta > 0$, there exists another collection of mutually independent random variables $\{\hat{v}_i\}_{i \in [n]}$, which are supported on a discrete set of cardinality $O\left(\frac{\log r}{\delta^2}\right)$ and satisfy the following properties.*

1. The optimal revenue when the buyer's values are $\{\hat{v}_i\}_{i \in [n]}$ is at most δu_{\max} smaller than the optimal revenue when the values are $\{v_i\}_{i \in [n]}$. I.e. $\hat{\mathcal{R}}_{OPT} \geq \mathcal{R}_{OPT} - \delta u_{\max}$, where $\mathcal{R}_{OPT} = \max_P \mathcal{R}_P(\{v_i\}_i)$ and $\hat{\mathcal{R}}_{OPT} = \max_P \mathcal{R}_P(\{\hat{v}_i\}_i)$.
2. Moreover, for any constant $\rho > 0$ and any price vector P such that $\mathcal{R}_P(\{\hat{v}_i\}_i) \geq \hat{\mathcal{R}}_{OPT} - \rho$, we can construct in time polynomial in the description of P and $1/\delta$ another price vector \tilde{P} such that $\mathcal{R}_{\tilde{P}}(\{v_i\}_i) \geq \mathcal{R}_{OPT} - 4\delta u_{\max} - \rho$.

If u_{\min} and u_{\max} are provided explicitly as input to the reduction,²⁴ we can compute the distributions of the \hat{v}_i 's²⁵ and their support in time polynomial in the description of $\{v_i\}_{i \in [n]}$, $\langle u_{\min} \rangle$, $\langle u_{\max} \rangle$, and $1/\delta$.

Proof. The proof is very similar to the proof of Lemma 30. In particular, let $\hat{\mathcal{V}} = \{\hat{v}_i\}_{i \in [n]}$ and $\tilde{\mathcal{V}} = \{\tilde{v}_i\}_{i \in [n]}$ be defined in the same way as in that lemma. So, with probability 1, (10) and (13) are satisfied. So Eq. (9) of Lemma 28 implies that $\hat{\mathcal{R}}_{OPT} \geq \mathcal{R}_{OPT} - \delta u_{\max}$.

Now let P be a price vector satisfying $\mathcal{R}_P(\hat{\mathcal{V}}) \geq \hat{\mathcal{R}}_{OPT} - \rho$. Lemma 24 implies that WLOG we can assume that $P \in [(1 + \delta)u_{\min}, (1 + \delta)u_{\max}]^n$. Then the vector \tilde{P} obtained from P via Eq. (11) is in $[u_{\min}/(1 + \delta - \delta^2), u_{\max}/(1 + \delta - \delta^2)]^n$, and satisfies $\mathcal{R}_{\tilde{P}}(\tilde{\mathcal{V}}) = \frac{\mathcal{R}_P(\hat{\mathcal{V}})}{(1 + \delta)(1 + \delta - \delta^2)}$, as \tilde{P} and $\tilde{\mathcal{V}}$ are the same linear transformations of P and $\hat{\mathcal{V}}$ respectively. Hence,

$$\begin{aligned}
\mathcal{R}_{\tilde{P}}(\mathcal{V}) &\geq \mathcal{R}_{\tilde{P}}(\tilde{\mathcal{V}}) - \delta u_{\max} && \text{(by Eq. (9) of Lemma 28 given (13))} \\
&= \frac{\mathcal{R}_P(\hat{\mathcal{V}})}{(1 + \delta)(1 + \delta - \delta^2)} - \delta u_{\max} \\
&\geq (1 - 2\delta)\mathcal{R}_P(\hat{\mathcal{V}}) - \delta u_{\max} \\
&\geq \hat{\mathcal{R}}_{OPT} - 3\delta u_{\max} - \rho \\
&\geq \mathcal{R}_{OPT} - 4\delta u_{\max} - \rho.
\end{aligned}$$

□

D.4 Omitted Details from Section 6

Proof of Lemma 14: Let OPT be the optimal revenue under \mathcal{V} . Also, let $\epsilon' = \epsilon/3$.

By Lemma 24, we only need to consider price vectors in $[0, 1]^n$ for an optimal one. Moreover, it follows from Lemma 27 that, if we restrict all prices to be higher than ϵ' , we lose at most an additive ϵ' in revenue. So there exists a price vector $\bar{P} \in [\epsilon', 1]^n$ such that $\mathcal{R}_{\bar{P}}(\mathcal{V}) \geq OPT - \epsilon'$.

Now, let us define a new collection of random variables $\tilde{\mathcal{V}} = \{\tilde{v}_i\}_{i \in [n]}$ via the following coupling: for all $i \in [n]$, set $\tilde{v}_i = \frac{\epsilon'}{2}$ if $v_i < \epsilon'$, and $\tilde{v}_i = v_i$ otherwise. We claim the following:

Claim 32. For any price vector P in $[\epsilon', 1]^n$, $\mathcal{R}_P(\tilde{\mathcal{V}}) = \mathcal{R}_P(\mathcal{V})$.

²⁴This requirement is only relevant if we have oracle access to the distributions of the v_i 's, as if we are given the distributions explicitly we immediately also know u_{\min} and u_{\max} .

²⁵The \hat{v}_i 's will inherit the same type of access that we have to the distributions of the v_i 's, according to Appendix B. In particular, if the v_i 's are specified explicitly in the input to the reduction then the \hat{v}_i 's will also be specified explicitly in the output of the reduction; if the v_i 's are given as oracles then the \hat{v}_i 's will be given as oracles; etc.

Proof of Claim 32: Recall that the variables $\{\tilde{v}_i\}_i$ are defined via a coupling with the v_i 's. Under the same coupling, sample values from the v_i 's and the \tilde{v}_i 's. For all items i such that $v_i \neq \tilde{v}_i$, the price of item i is higher than the value of item i in both cases, so item i will not be purchased in both cases. That means the buyer will make the same decision in both cases, as she will only consider items whose values are the same. So the revenues under \mathcal{V} and $\tilde{\mathcal{V}}$ are pointwise equal. \square

We proceed to show that an approximately optimal solution for $\tilde{\mathcal{V}}$ provides an approximately optimal solution for \mathcal{V} . Suppose that a price vector \tilde{P} satisfies $\mathcal{R}_{\tilde{P}}(\tilde{\mathcal{V}}) \geq \mathcal{R}_{OPT}(\tilde{\mathcal{V}}) - \epsilon'$. By Lemmas 24 and 27, we can efficiently convert \tilde{P} to $P' \in [\epsilon', 1]^n$, such that $\mathcal{R}_{P'}(\tilde{\mathcal{V}}) \geq \mathcal{R}_{\tilde{P}}(\tilde{\mathcal{V}}) - \epsilon'$.

Combining the inequalities above, we have

$$\mathcal{R}_{OPT}(\tilde{\mathcal{V}}) \geq \mathcal{R}_{\tilde{P}}(\tilde{\mathcal{V}}) = \mathcal{R}_{\tilde{P}}(\mathcal{V}) \geq OPT - \epsilon',$$

and

$$\mathcal{R}_{P'}(\mathcal{V}) = \mathcal{R}_{P'}(\tilde{\mathcal{V}}) \geq \mathcal{R}_{\tilde{P}}(\tilde{\mathcal{V}}) - \epsilon' \geq \mathcal{R}_{OPT}(\tilde{\mathcal{V}}) - 2\epsilon'.$$

Thus,

$$\mathcal{R}_{P'}(\mathcal{V}) \geq OPT - 3\epsilon'.$$

\square

Proof of Theorem 15: Lemma 31 implies that we can reduce the problem $\text{ADDITIVEPRICE}(\mathcal{V}, \epsilon)$ to the problem $\text{ADDITIVEPRICE}(\hat{\mathcal{V}}, \frac{\epsilon}{3})$, where $\hat{\mathcal{V}} = \{\hat{v}_i\}_i$ is a collection of mutually independent random variables supported on a common discrete set $\mathcal{S} = \{s^{(1)}, \dots, s^{(k_1)}\} \subset [(1 + \frac{\epsilon}{6u_{max}})u_{min}, (1 + \frac{\epsilon}{6u_{max}})u_{max}]$ of cardinality $k_1 = O(\frac{u_{max}^2 \log r}{\epsilon^2})$. Now, Lemmas 24 and 25 imply that we can reduce the problem $\text{ADDITIVEPRICE}(\hat{\mathcal{V}}, \frac{\epsilon}{3})$ to the problem of approximating $\text{RESRTICTEDPRICE}(\hat{\mathcal{V}}, \mathcal{P})$ to within an additive $\frac{\epsilon}{6}$, where \mathcal{P} is a discrete set of prices of cardinality $O(\frac{u_{max}^2 \log r}{\epsilon^2})$, satisfying $\max_{x \in \mathcal{P}} x \leq \frac{7}{6}u_{max}$. \square

D.5 Proof of Theorem 17

Proof of Theorem 17: Lemma 30 implies that we can reduce the problem $\text{PRICE}(\mathcal{V}, \epsilon)$ to the problem $\text{PRICE}(\hat{\mathcal{V}}, \frac{\epsilon}{8})$, where $\hat{\mathcal{V}} = \{\hat{v}_i\}_i$ is a collection of mutually independent random variables supported on a common discrete set $\mathcal{S} = \{s^{(1)}, \dots, s^{(k_1)}\} \subset [(1 + (\epsilon/8)^8)u_{min}, (1 + (\epsilon/8)^8)u_{max}]$ of cardinality $k_1 = O(\frac{\log r}{\epsilon^{16}})$. Now, Lemmas 24 and 25 imply that we can reduce the problem $\text{PRICE}(\hat{\mathcal{V}}, \frac{\epsilon}{8})$ to the problem of approximating $\text{RESRTICTEDPRICE}(\hat{\mathcal{V}}, \mathcal{P})$ to within a multiplicative factor of $1 - \frac{\epsilon}{16}$, where \mathcal{P} is a discrete set of prices of cardinality $O(\frac{\log r}{\epsilon^2})$. \square

E Proof of Theorem 2

We restate and prove Theorem 2. We note that we have not tried to carefully optimize the constants in the running time. There may be room for improvement with a more careful analysis.

Theorem 2 [Restated] Let $\{F_i\}_{i \in [n]}$ be a collection of distributions that are supported on a bounded set $[u_{min}, u_{max}] \subset \mathbb{R}_+$, where u_{min} and u_{max} are specified as part of the input,²⁶ and let

²⁶This requirement is only relevant if we have oracle access to the F_i 's, as if we are given the distributions explicitly we immediately also know u_{min} and u_{max} .

$r := u_{max}/u_{min}$. Then, for any constant $\epsilon > 0$, there is an algorithm that runs in time polynomial in the size of the input and $\max \left\{ n^{\log^{11} r \cdot \log \log r}, n^{\frac{\log^3 r \cdot \log \frac{1}{\epsilon}}{\epsilon^8}} \right\}$ and computes a price vector P such that

$$\mathcal{R}_P \geq (1 - \epsilon)OPT,$$

where \mathcal{R}_P is the expected revenue under price vector P when the buyer's values for the items are independent draws from the distributions $\{F_i\}_i$ and OPT is the optimal revenue.

Proof of Theorem 2: First set $\hat{\epsilon} = \min \left\{ \epsilon, \frac{1}{(4^{\lceil \log_2 r \rceil})^{1/6}} \right\}$. Clearly, it suffices to find a price vector with expected revenue $(1 - \hat{\epsilon})OPT$. Now, let us invoke the reduction of Theorem 17, reducing this task to approximating $\text{RESTRICTEDPRICE}(\{\hat{F}_i\}_i, \mathcal{P})$ to within a factor of $(1 - \frac{\hat{\epsilon}}{16})$, where the distributions $\{\hat{F}_i\}_i$ are supported on a discrete set $\mathcal{S} = \{s^{(1)}, \dots, s^{(k_1)}\}$ of cardinality $k_1 = O(\log r / \hat{\epsilon}^{16})$ and the prices are also restricted to a discrete set $\mathcal{P} = \{p^{(1)}, \dots, p^{(k_2)}\}$ of cardinality $k_2 = O(\log r / \hat{\epsilon}^2)$. It is important to note that $\mathcal{S} \subset [(1 + (\hat{\epsilon}/8)^8)u_{min}, (1 + (\hat{\epsilon}/8)^8)u_{max}]$ and $\min_i \{p^{(i)}\} \leq \min_i \{s^{(i)}\}$ (this can be checked by a careful study of the proof of Theorem 17). Hence, if \widehat{OPT} is the optimal revenue of the discrete instance resulting from the reduction, we have $\widehat{OPT} \geq \min_i \{p^{(i)}\}$. It is our goal to achieve revenue at least $(1 - \frac{\hat{\epsilon}}{16})\widehat{OPT}$.

To do this, we invoke the algorithm of Theorem 13 with a choice of $m = \Theta(\frac{nrk_1k_2}{\hat{\epsilon}}) = O(\frac{nr \log^2 r}{\hat{\epsilon}^{19}})$, obtaining a price vector with revenue at least:

$$\widehat{OPT} - O\left(\frac{nk_1k_2}{m} \max_i \{p^{(i)}\}\right) = \widehat{OPT} - O(\hat{\epsilon}) \cdot \min_i \{p^{(i)}\} \geq \widehat{OPT} (1 - O(\hat{\epsilon})), \quad (15)$$

as we wanted. The running time of the algorithm in this case is polynomial in the input and $m^{\frac{\log^2 r}{\hat{\epsilon}^{18}}}$, that is polynomial in the input and $\max \left\{ n^{\log^6 r \cdot \log \log r}, n^{\frac{\log^3 r \cdot \log \frac{1}{\epsilon}}{\epsilon^{18}}} \right\}$.

Being a bit more careful in the application of our discretization lemmas we can obtain running time polynomial in the input and $\max \left\{ n^{\log^{11} r \cdot \log \log r}, n^{\frac{\log^3 r \cdot \log \frac{1}{\epsilon}}{\epsilon^8}} \right\}$. Recall that to establish our reduction in Theorem 17 we employed Lemma 30, which in turn made use of Lemma 28, setting $a = \frac{7}{8}$. Setting instead $a = \frac{2}{3}$ would result in a stronger Lemma 30 and Theorem 17, improving our running time here. \square

F Details of Section 8: MHR to Bounded Distributions

F.1 Basic Properties of MHR Distributions

Definition 33. For a random variable X , we define $\alpha_1 = u_{min}$, and for every real number $p \in (1, +\infty)$, we define $\alpha_p = \inf \left\{ x \mid F(x) \geq 1 - \frac{1}{p} \right\}$.

The following lemma establishes an interesting property of MHR distributions. Intuitively, the lemma provides a lower bound on the speed of the decay of the tail of a MHR distribution. We prove the lemma by showing that the function $\log_e (1 - F(x))$ is concave if F is MHR, and exploiting this concavity (see Appendix F.1.1).

Lemma 34. *If the distribution of a random variable X satisfies MHR, $m \geq 1$ and $d \geq 1$, $d \cdot \alpha_m \geq \alpha_{m^d}$.*

Next we study the expectation of a random variable that satisfies MHR. We show that the contribution to the expectation from values $\geq m$, is $O(m \cdot \Pr[X \geq m])$. We start with a definition.

Definition 35. *For a random variable X , let $\text{Con}[X \geq x] = \mathbb{E}[X|X \geq x] \cdot \Pr\{X \geq x\}$ be the contribution to expectation of X from values which are no smaller than x , i.e.*

$$\text{Con}[X \geq x] = \int_x^{+\infty} t \cdot f(t) dt.$$

It is an obvious fact that for any random variable X and any two points $x_1 \leq x_2$, $\text{Con}[X \geq x_1] \geq \text{Con}[X \geq x_2]$. Using the bound on the tail of a MHR distribution obtained in Lemma 34, we bound the contribution to the expectation of X by the values at the tail of the distribution. The proof is given in Appendix F.1.

Lemma 36. *Let X be a random variable whose distribution satisfies MHR. For all $m \geq 2$, $\text{Con}[X \geq \alpha_m] \leq 6\alpha_m/m$.*

F.1.1 Proofs Omitted from Appendix F.1

Proof of Lemma 34: It is not hard to see that $f(x) > 0$, for all $x \in (u_{\min}, u_{\max})$. For a contradiction, assume this is not true, that is, for some $x' \in (u_{\min}, u_{\max})$, $f(x') = 0$. We know $1 - F(x') > 0$. Thus $\frac{f(x')}{1 - F(x')} = 0$. Since the distribution satisfies MHR and $1 - F(x)$ is positive for all $x \in (u_{\min}, x')$, $f(x) = 0$ in this interval. Hence, it must also be that $F(x) = 0$ in $[u_{\min}, x')$. Since $x' > u_{\min}$, it follows that $u_{\min} \neq \sup\{x|F(x) = 0\}$, a contradiction.

Since $f(x) > 0$ in (u_{\min}, u_{\max}) , $F(x)$ is monotone in (u_{\min}, u_{\max}) . So we can define the inverse $F^{-1}(x)$ in (u_{\min}, u_{\max}) . It is not hard to see that for any $p \in [1, +\infty)$, $F(\alpha_p) = 1 - 1/p$ and $\alpha_p = F^{-1}(1 - 1/p)$.

Now let $G(x) = \log_e(1 - F(x))$. We will show that $G(x)$ is a concave function.

Let us consider the derivative of $G(x)$. By the definition of MHR, $G'(x) = \frac{-f(x)}{1 - F(x)}$ is monotonically non-increasing. Therefore, $G(x)$ is concave. It follows that, for every m , by the concavity of $G(x)$, the following inequality holds:

$$G\left(\frac{d-1}{d} \cdot \alpha_1 + \frac{1}{d} \cdot \alpha_{m^d}\right) \geq \frac{d-1}{d} G(\alpha_1) + \frac{1}{d} G(\alpha_{m^d}).$$

Let us rewrite the RHS as follows

$$\frac{d-1}{d} G(\alpha_1) + \frac{1}{d} G(\alpha_{m^d}) = \frac{d-1}{d} \log_e 1 + \frac{1}{d} \log_e(1 - F(\alpha_{m^d})) = \frac{1}{d} \log_e\left(\frac{1}{m^d}\right) = \log_e\left(\frac{1}{m}\right)$$

Hence, we have the following:

$$\begin{aligned}
& G\left(\frac{d-1}{d} \cdot \alpha_1 + \frac{1}{d} \cdot \alpha_{m^d}\right) \geq \log_e\left(\frac{1}{m}\right) \\
\implies \log_e\left(1 - F\left(\frac{d-1}{d} \cdot \alpha_1 + \frac{1}{d} \cdot \alpha_{m^d}\right)\right) & \geq \log_e\left(\frac{1}{m}\right) \\
\implies 1 - F\left(\frac{d-1}{d} \cdot \alpha_1 + \frac{1}{d} \cdot \alpha_{m^d}\right) & \geq \frac{1}{m} \\
\implies 1 - F\left(\frac{d-1}{d} \cdot \alpha_1 + \frac{1}{d} \cdot \alpha_{m^d}\right) & \geq 1 - F(\alpha_m) \\
\implies F(\alpha_m) & \geq F\left(\frac{d-1}{d} \cdot \alpha_1 + \frac{1}{d} \cdot \alpha_{m^d}\right) \\
\implies \alpha_m & \geq \frac{d-1}{d} \cdot \alpha_1 + \frac{1}{d} \cdot \alpha_{m^d} \quad (F \text{ is monotone increasing}) \\
\implies \alpha_m & \geq \frac{1}{d} \cdot \alpha_{m^d} \quad (u_{min} \geq 0) \\
\implies d \cdot \alpha_m & \geq \alpha_{m^d}.
\end{aligned}$$

□

Proof of Lemma 36: Let $S = \text{Con}[X \geq \alpha_m]$, and consider the sequence $\{\beta_i := \alpha_{m^{(2^i)}}\}$, defined for all non-negative integers i . It can easily be seen that $\lim_{i \rightarrow +\infty} \alpha_{m^{(2^i)}} = u_{max}$; hence, $\lim_{i \rightarrow +\infty} \beta_i = u_{max}$ and by continuity $\lim_{i \rightarrow +\infty} F(\beta_i) = F(u_{max}) = 1$.

Also,

$$\int_{\beta_i}^{\beta_{i+1}} x \cdot f(x) dx \leq \beta_{i+1}(1 - F(\beta_i)) = \beta_{i+1}/m^{(2^i)}.$$

Moreover, Lemma 34 implies that $\beta_i \leq 2\beta_{i-1}$; thus, $\beta_i \leq 2^i \beta_0 \leq 2^i \alpha_m$. Hence, we have the following:

$$\begin{aligned}
S &= \int_{\alpha_m}^{u_{max}} x \cdot f(x) dx \leq \sum_{i=0}^{+\infty} \frac{\beta_{i+1}}{m^{(2^i)}} \leq \sum_{i=0}^{+\infty} \frac{2^{i+1} \alpha_m}{m^{(2^i)}} \\
&\leq \frac{2\alpha_m}{m} + \sum_{i=1}^{+\infty} \frac{2^{(i+1)} \alpha_m}{m^{(2^i)}} = \frac{2\alpha_m}{m} + \frac{4\alpha_m}{m^2} \sum_{i=0}^{+\infty} \left(\frac{2}{m^2}\right)^i \\
&= \frac{2\alpha_m}{m} + \frac{4\alpha_m}{m^2} \cdot \frac{1}{1 - 2/m^2} \\
&\leq \frac{2\alpha_m}{m} + \frac{4\alpha_m}{m} \\
&\leq \frac{6\alpha_m}{m}.
\end{aligned}$$

□

F.2 Proof of Theorem 19: Extreme Value Theorem for MHR Distributions

We start with some useful notation. For all $i = 1, \dots, n$, we denote by F_i the distribution of variable X_i . We also let $\alpha_m^{(i)} := \inf \{x | F_i(x) \geq 1 - \frac{1}{m}\}$, for all $m \geq 1$. Moreover, we assume that

n is a power of 2. If not, we can always include at most n additional random variables that are detremistically 0, making the total number of variables a power of 2.

We proceed with the proof of Theorem 19. The threshold β is computed by an algorithm. At a high level, the algorithm proceeds in $O(\log n)$ rounds, indexed by $t \in \{0, \dots, \log_2 n\}$, eliminating half of the variables at each round. The way the elimination works is as follows. In round t , we compute for each of the variables that have survived so far the threshold $\alpha_{n/2^t}$ beyond which the size of the tail of their distribution becomes smaller than $\frac{2^t}{n}$. We then sort these thresholds and eliminate the bottom half of the variables, recording the threshold of the last variable that survived this round. The maximum of these records among the $\log_2 n$ rounds of the algorithm is our β . The pseudocode of the algorithm is given below. Given that we may only be given oracle access to the distributions $\{F_i\}_{i \in [n]}$, we allow some slack $\eta \leq \frac{1}{2}$ in the computation of our thresholds so that the computation is efficient. If we know the distributions explicitly, the description of the algorithm simplifies to the case $\eta = 0$.

Algorithm 1 Algorithm for finding β

- 1: Define the permutation of the variables $\pi_0(i) = i, \forall i \in [n]$, and the set of remaining variables $Q_0 = [n]$.
 - 2: **for** $t := 0$ to $\log_2 n - 1$ **do**
 - 3: For all $j \in [n/2^t]$, compute some $x_{n/2^t}^{(\pi_t(j))} \in [1 - \eta, 1 + \eta] \cdot \alpha_{n/2^t}^{(\pi_t(j))}$, for a small constant $\eta \in [0, 1/2)$
 - 4: Sort these $n/2^t$ numbers in decreasing order π_{t+1} such that $x_{n/2^t}^{(\pi_{t+1}(1))} \geq x_{n/2^t}^{(\pi_{t+1}(2))} \geq \dots \geq x_{n/2^t}^{(\pi_{t+1}(n/2^t))}$
 - 5: $Q_{t+1} := \{ \pi_{t+1}(i) \mid i \leq n/2^{t+1} \}$
 - 6: $\beta_t := x_{n/2^t}^{(\pi_{t+1}(n/2^{t+1}))}$
 - 7: **end for**
 - 8: Compute $x_2^{(\pi_{\log_2 n}(1))} \in [1 - \eta, 1 + \eta] \cdot \alpha_2^{(\pi_{\log_2 n}(1))}$
 - 9: Set $\beta_{\log_2 n} := x_2^{(\pi_{\log_2 n}(1))}$
 - 10: Output $\beta := \max_t \beta_t$
-

Crucial in the proof of the theorem is the following lemma.

Lemma 37. *For all $i \in [n]$ and $\epsilon \in (0, 1/4)$, let $S_i = \text{Con}[X_i \geq 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]$, where $\text{Con}[\cdot]$ is defined as in Definition 35. Then*

$$\sum_{i=1}^n S_i \leq 36 \log_2(1/\epsilon) \epsilon \cdot \beta, \quad \text{for all } \epsilon \in (0, 1/4).$$

Proof. Let $d = \log_2(\frac{1}{\epsilon})$ and notice that $d \geq 2$. It is not hard to see that we can divide $[n]$ into $(\log_2 n) + 1$ different groups $\{G_t\}_{t \in \{0, \dots, \log_2 n\}}$ based on the sets Q_t maintained by the algorithm, as follows. For $t \in \{0, \dots, \log_2 n\}$, set

$$G_t = \begin{cases} Q_t \setminus Q_{t+1} & t < \log_2 n \\ Q_{\log_2 n} & t = \log_2 n \end{cases}$$

Now, it is not hard to see that, for all $t < \log_2 n$ and all $i \in G_t$, $S_i \leq \text{Con}[X_i \geq 2d \cdot \beta_t]$, since $\beta_t \leq \beta$. Also for any $i \in G_t$, there must exist some $k \in (n/2^{t+1}, n/2^t]$, such that $i = \pi_{t+1}(k)$. Then

by the definition of the algorithm, we know that

$$(1 - \eta)\alpha_{n/2^t}^{(i)} \leq x_{n/2^t}^{(i)} \leq x_{n/2^t}^{(\pi_{t+1}(n/2^{t+1}))} = \beta_t.$$

Recall that η is chosen to satisfy $2 \geq 1/(1 - \eta)$. Then $d \cdot \alpha_{n/2^t}^{(i)} \leq 2d \cdot \beta_t$. But Lemma 34 gives $d \cdot \alpha_{n/2^t}^{(i)} \geq \alpha_{(n/2^t)^d}^{(i)}$. Hence,

$$2d \cdot \beta_t \geq d \cdot \alpha_{n/2^t}^{(i)} \geq \alpha_{(n/2^t)^d}^{(i)},$$

which implies that

$$\text{Con}[v_i \geq 2d \cdot \beta_t] \leq \text{Con}[v_i \geq \alpha_{(n/2^t)^d}^{(i)}].$$

Using Lemma 36, we know that

$$\text{Con}[v_i \geq \alpha_{(n/2^t)^d}^{(i)}] \leq 6\alpha_{(n/2^t)^d}^{(i)}(2^t/n)^d \leq 12d\beta_t(2^t/n)^d.$$

Now, since $|G_t| = n/2^{t+1}$,

$$\sum_{i \in G_t} S_i \leq 12d\beta_t(2^t/n)^d \times n/2^{t+1} = 6d \cdot \beta_t(2^t/n)^{d-1} = \frac{6d \cdot \beta_t}{n^{d-1}}(2^{d-1})^t.$$

Thus,

$$\begin{aligned} \sum_{i \in [n] \setminus G_{\log_2 n}} S_i &\leq \sum_{t=0}^{(\log_2 n)-1} \frac{6d \cdot \beta_t}{n^{d-1}}(2^{d-1})^t \\ &\leq \frac{6d \cdot \beta}{n^{d-1}} \cdot \frac{(2^{d-1})^{\log_2 n} - 1}{2^{d-1} - 1} \\ &= \frac{6d \cdot \beta}{n^{d-1}} \cdot \frac{n^{d-1} - 1}{2^{d-1} - 1} \\ &\leq \frac{12d \cdot \beta}{2^d - 2} \\ &\leq \frac{24d \cdot \beta}{2^d} \\ &= 24 \log_2(1/\epsilon)\epsilon \cdot \beta \end{aligned}$$

Let i be the unique element in $G_{\log_2 n}$. Then $\beta_{\log_2 n} = x_2^{(i)}$. Using Lemma 34 and the definition of $x_2^{(i)}$, we obtain

$$2d \cdot \beta \geq 2d \cdot \beta_{\log_2 n} \geq 2d \cdot x_2^{(i)} \geq 2(1 - \eta)d \cdot \alpha_2^{(i)} \geq d \cdot \alpha_2^{(i)} \geq \alpha_{2^d}^{(i)} = \alpha_{1/\epsilon}^{(i)}.$$

Using the above and Lemma 36 we get

$$S_i \leq \text{Con}[v_i \geq \alpha_{1/\epsilon}^{(i)}] \leq 6\epsilon \cdot \alpha_{1/\epsilon}^{(i)} \leq 12\epsilon d \cdot \beta.$$

Putting everything together,

$$\sum_{i=1}^n S_i \leq 36 \log_2(1/\epsilon)\epsilon \cdot \beta.$$

□

Using Lemma 37, we obtain

$$\int_{2\beta \log_2 1/\epsilon}^{+\infty} t \cdot f_{\max_i \{X_i\}}(t) dt \leq \sum_{i=1}^n S_i \leq 36 \log_2(1/\epsilon) \epsilon \cdot \beta.$$

It remains to show that

$$\Pr[\max_i \{X_i\} \geq \beta/2] \geq 1 - \frac{1}{e^{1/2}}. \quad (16)$$

We show that, for all t , $\Pr[\max_i \{X_i\} \geq \frac{\beta_t}{1+\eta}] \geq 1 - \frac{1}{e^{1/2}}$, where η is the parameter used in Algorithm 1. This is sufficient to imply (16), as $\eta \leq 1/2$. Observe that for all $i \in [n/2^{t+1}]$,

$$(1 + \eta) \cdot \alpha_{n/2^t}^{(\pi_{t+1}(i))} \geq x_{n/2^t}^{(\pi_{t+1}(i))} \geq \beta_t,$$

where π_{t+1} is the permutation constructed in the t -th round of the algorithm. This implies

$$\alpha_{n/2^t}^{(\pi_{t+1}(i))} \geq \frac{\beta_t}{1 + \eta}.$$

Hence, for all $i \in [n/2^{t+1}]$, $\Pr[X_{\pi_{t+1}(i)} \leq \frac{\beta_t}{1+\eta}] \leq 1 - 2^t/n$. Thus,

$$\begin{aligned} \Pr\left[\max_i \{X_i\} \geq \frac{\beta_t}{1 + \eta}\right] &\geq \Pr\left[\exists i \in [n/2^{t+1}], X_{\pi_{t+1}(i)} \geq \frac{\beta_t}{1 + \eta}\right] \\ &\geq 1 - (1 - 2^t/n)^{n/2^{t+1}} \\ &\geq 1 - \frac{1}{e^{1/2}}. \end{aligned}$$

Eq. (16) now follows.

F.3 Proof of Theorem 18: Reduction from MHR to Bounded Distributions

Recall that we represent by $\{v_i\}_{i \in [n]}$ the values of the buyer for the items. We will denote their distributions by $\{F_i\}_{i \in [n]}$ throughout this appendix.

F.3.1 Relating OPT to β

We demonstrate that the anchoring point β of Theorem 19 provides a lower bound to the optimal revenue. In particular, we show that the optimal revenue satisfies $OPT = \Omega(\beta)$. This lemma justifies the relevance of β .

Lemma 38. *If β is the anchoring point of Theorem 19, then $OPT \geq \left(1 - \frac{1}{\sqrt{e}}\right) \frac{\beta}{2}$.*

Proof of Lemma 38: Suppose we priced all items at $\frac{\beta}{2}$. The revenue we would get from such price vector would be at least

$$\frac{\beta}{2} \Pr\left[\max\{v_i\} \geq \frac{\beta}{2}\right] \geq \frac{\beta}{2} \left(1 - \frac{1}{\sqrt{e}}\right),$$

where we used Theorem 19. Hence, $OPT \geq \left(1 - \frac{1}{\sqrt{e}}\right) \frac{\beta}{2}$. \square

For simplicity, we set $c_1 := \frac{1}{2} \left(1 - \frac{1}{\sqrt{e}}\right)$ for the next appendices, keeping in mind that c_1 is an absolute constant.

F.3.2 Restricting the Prices

This appendix culminates in Lemma 40 (given below), which states that we can constrain our prices to the set $[\epsilon \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]$ without hurting the revenue by more than a fraction of $\frac{\epsilon + c_2(\epsilon)}{c_1}$, where $c_2(\epsilon) := 36 \log_2(\frac{1}{\epsilon}) \epsilon$ and c_1 is the constant defined in Appendix F.3.1. We prove this in two steps. First, exploiting our extreme value theorem for MHR distributions (Theorem 19), we show that for a given price vector, if we lower the prices that are above $2 \log_2(\frac{1}{\epsilon}) \cdot \beta$ to $2 \log_2(\frac{1}{\epsilon}) \cdot \beta$, the loss in revenue is bounded by $c_2(\epsilon) \cdot \beta$, namely

Lemma 39. *Fix an arbitrary $\epsilon \in (0, 1/4)$. Given a price vector P , define P' as follows: set $p'_i = p_i$, if $p_i \leq 2 \log_2(\frac{1}{\epsilon}) \cdot \beta$, and $p'_i = 2 \log_2(\frac{1}{\epsilon}) \cdot \beta$ otherwise. Then the expected revenues \mathcal{R}_P and $\mathcal{R}_{P'}$ achieved by price vectors P and P' respectively satisfy: $\mathcal{R}_{P'} \geq \mathcal{R}_P - c_2(\epsilon) \cdot \beta$.*

Using Lemma 39, we obtain our main result for this appendix. Observe that we can make the loss in revenue arbitrarily small by taking ϵ sufficiently small.

Lemma 40. *For all $\epsilon \in (0, 1/4)$, there exists a price vector $P^* \in [\epsilon \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]^n$, such that the revenue from this price vector satisfies $\mathcal{R}_{P^*} \geq \left(1 - \frac{\epsilon + c_2(\epsilon)}{c_1}\right) OPT$, where OPT is the optimal revenue under any price vector.*

All proofs of this appendix can be found in Appendix F.4.1.

F.3.3 Truncating the Value Distributions

Exploiting Lemma 40, i.e. that we can constrain the prices to $[\epsilon \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]$ without hurting the revenue, we show Theorem 18, i.e. that we can also constrain the support of the value distributions into a bounded range. In particular, we show that we can “truncate” the value distributions to the range $[\frac{\epsilon}{2} \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]$, where for our purposes “truncating” means this: for every distribution F_i , we shift all probability mass from $(2 \log_2(\frac{1}{\epsilon}) \cdot \beta, +\infty)$ to the point $2 \log_2(\frac{1}{\epsilon}) \cdot \beta$, and all probability mass from $(-\infty, \epsilon \cdot \beta)$ to $\frac{\epsilon}{2} \cdot \beta$. We show that our modification does not hurt the revenue. That is, we establish a polynomial-time reduction from the problem of computing a near-optimal price vector when the buyer’s value distributions are arbitrary MHR distributions to the case where the buyer’s value distributions are supported on a bounded interval $[u_{min}, c \cdot u_{min}]$, where $c = c(\epsilon) = 4 \frac{1}{\epsilon} \log_2(\frac{1}{\epsilon})$ is a constant that only depends on the desired approximation ϵ . The proof of Theorem 41 is given in Appendix F.4.2.

Theorem 41 (Reduction from MHR to Bounded Distributions). *Given $\epsilon \in (0, 1/4)$ and a collection of mutually independent random variables $\{v_i\}_i$ that are MHR, let us define a new collection of random variables $\{\tilde{v}_i\}_i$ via the following coupling: for all $i \in [n]$, set $\tilde{v}_i = \frac{\epsilon}{2} \cdot \beta$ if $v_i < \epsilon \cdot \beta$, set $\tilde{v}_i = 2 \log_2(\frac{1}{\epsilon}) \cdot \beta$ if $v_i \geq 2 \log_2(\frac{1}{\epsilon}) \cdot \beta$, and set $\tilde{v}_i = v_i$ otherwise, where $\beta = \beta(\{v_i\}_i)$ is the anchoring point of Theorem 19 computed from the distributions of the variables $\{v_i\}_i$. Let also \widetilde{OPT} be the optimal revenue of the seller when the buyer’s values are distributed as $\{\tilde{v}_i\}_{i \in [n]}$ and OPT the optimal revenue when the buyer’s values are distributed as $\{v_i\}_{i \in [n]}$. Then given a price vector that achieves revenue $(1 - \delta) \cdot \widetilde{OPT}$ when the buyer’s values are distributed as $\{\tilde{v}_i\}_{i \in [n]}$, we can efficiently compute a price vector with revenue*

$$\left(1 - \delta - \frac{2\epsilon + 3c_2(\epsilon)}{c_1}\right) OPT$$

when the buyer’s values are distributed as $\{v_i\}_{i \in [n]}$.

Theorem 18 follows from Theorem 41.

F.4 Proofs Omitted from Appendix F.3

F.4.1 Restricting the Price Range for MHR Distributions: the Proofs

Proof of Lemma 39: We will show a slightly more general result. Given a price vector, if we make all prices that are above α equal to α , then the loss in revenue can be bounded by the sum, over all items whose price was above α , of the contribution to this item's expected value by points above α . Formally,

Lemma 42. *Let $\alpha > 0$ and $S(\alpha) = \text{Con}[\max_i v_i \geq \alpha]$, where $\text{Con}[\cdot]$ is defined as in Definition 35. Moreover, for a given price vector P , define P' as follows: set $p'_i = p_i$, if $p_i < \alpha$, and $p'_i = \alpha$, otherwise. Then the expected revenues \mathcal{R}_P and $\mathcal{R}_{P'}$ from P and P' respectively satisfy*

$$\mathcal{R}_{P'} \geq \mathcal{R}_P - S(\alpha).$$

Proof. Let $S_{exp} = \{ i \mid p_i > \alpha \}$ be the set of expensive items under P , and let P'' be a new price vector obtained from P by changing the price of all items $i \in S_{exp}$ from p_i to $+\infty$. If we switch from P to P'' , the only case where the buyer makes a different decision is when she used to buy some item from S_{exp} under P . So the decrease in revenue can be bounded by the contribution to \mathcal{R}_P from the items in S_{exp} . Clearly, this contribution is at most $\text{Con}[\max_{i \in S_{exp}} v_i \geq \alpha] \leq S(\alpha)$.

We proceed to argue that $\mathcal{R}_{P'} \geq \mathcal{R}_{P''}$. If we switch from P'' to P' , the buyer will pointwise either make the same decision or switch to buy some item in S_{exp} paying α . Since α is larger than any finite price in P'' , the revenue does not decrease.

Combining the two inequalities, we have $\mathcal{R}_{P'} \geq \mathcal{R}_{P''} \geq \mathcal{R}_P - S(\alpha)$. \square

Combining Theorem 19 with Lemma 42, it is easy to argue that if we truncate a price vector P at value $2 \log_2(\frac{1}{\epsilon}) \cdot \beta$ to obtain a new price vector P' the change in revenue can be bounded as follows: $\mathcal{R}_{P'} \geq \mathcal{R}_P - c_2(\epsilon) \cdot \beta$. \square

Proof of Lemma 40: Lemma 39 implies that, if we start from any price vector P , we can modify it into another price vector P' that does not use any price above $2 \log_2(\frac{1}{\epsilon}) \cdot \beta$, and satisfies $\mathcal{R}_{P'} \geq \mathcal{R}_P - c_2(\epsilon) \cdot \beta$. Then Lemma 27 implies that we can change P' into another vector $P'' \in [\epsilon \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]^n$, such that $\mathcal{R}_{P''} \geq \mathcal{R}_{P'} - \epsilon \cdot \beta$.

By Lemma 38, we know that $OPT \geq c_1 \cdot \beta$. Hence, if we start with the optimal price vector P and apply the above transformations, we will obtain a price vector $P^* \in [\epsilon \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]^n$ such that

$$\mathcal{R}_{P^*} \geq OPT - (\epsilon + c_2(\epsilon)) \cdot \beta \geq \left(1 - \frac{\epsilon + c_2(\epsilon)}{c_1}\right) OPT.$$

\square

F.4.2 Bounding the Support of the Distributions: the Proofs

To establish Theorem 41 we show that we can transform $\{v_i\}_{i \in [n]}$ into $\{\tilde{v}_i\}_{i \in [n]}$ such that, for all i , \tilde{v}_i only takes values in $[\frac{\epsilon}{2} \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]$, and for any price vector $P \in [\epsilon \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]^n$, $|\tilde{\mathcal{R}}_P - \mathcal{R}_P| \leq c_2(\epsilon) \cdot \beta$, where \mathcal{R}_P and $\tilde{\mathcal{R}}_P$ are respectively the revenues of the seller when the buyer's values are distributed as $\{v_i\}_{i \in [n]}$ and $\{\tilde{v}_i\}_{i \in [n]}$. We first show that one side of our truncation works.

Lemma 43. *Given $\epsilon \in (0, 1/4)$ and a collection of random variables $\{v_i\}_i$ that are MHR, let us define a new collection of random variables $\{\hat{v}_i\}_i$ via the following coupling: for all $i \in [n]$, if $v_i \leq 2 \log_2(\frac{1}{\epsilon}) \cdot \beta$, set $\hat{v}_i = v_i$, otherwise set $\hat{v}_i = 2 \log_2(\frac{1}{\epsilon}) \cdot \beta$, where $\beta = \beta(\{v_i\}_i)$ is the anchoring point of Theorem 19 computed from the distributions of the variables $\{v_i\}_i$. Then, for any price vector $P \in [\epsilon \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]^n$, $|\mathcal{R}_P - \hat{\mathcal{R}}_P| \leq c_2(\epsilon) \cdot \beta$, where \mathcal{R}_P and $\hat{\mathcal{R}}_P$ are respectively the revenues of the seller when the buyer's values are distributed as $\{v_i\}_{i \in [n]}$ and as $\{\hat{v}_i\}_{i \in [n]}$.*

Proof. For convenience let $d = \log_2(\frac{1}{\epsilon})$, and let R_P and \hat{R}_P be random variables representing the revenue when the buyer's values are $\{v_i\}_{i \in [n]}$ and $\{\hat{v}_i\}_{i \in [n]}$ respectively. Recall that $\{v_i\}_i$ and $\{\hat{v}_i\}_i$ are defined via a coupling, so $R_P - \hat{R}_P \neq 0$ only in the event $v_i \neq \hat{v}_i$, for some i . Notice that the probability of this event is $\Pr[\exists i, v_i > 2d \cdot \beta]$, and that $0 \leq R_P, \hat{R}_P \leq 2d \cdot \beta$ since the maximum price of any item is $2d \cdot \beta$. Hence, we can bound $|\mathbb{E}[R_P] - \mathbb{E}[\hat{R}_P]|$ using Theorem 19 as follows:

$$\begin{aligned} |\mathbb{E}[R_P] - \mathbb{E}[\hat{R}_P]| &\leq 2d \cdot \beta \cdot \Pr[\exists i, v_i > 2d \cdot \beta] \\ &= 2d \cdot \beta \cdot \Pr[\max_i v_i \geq 2d \cdot \beta] \leq \text{Con}[\max_i v_i \geq 2d \cdot \beta] \leq c_2(\epsilon) \cdot \beta. \end{aligned}$$

□

Next we show that the other side of the truncation works.

Lemma 44. *Given $\epsilon, \beta > 0$ and a collection of random variables $\{\hat{v}_i\}_i$, let us define a new collection of random variables $\{\tilde{v}_i\}_i$ via the following coupling: for all $i \in [n]$, if $\hat{v}_i \geq \epsilon \cdot \beta$, set $\tilde{v}_i = \hat{v}_i$, otherwise set $\tilde{v}_i = \frac{\epsilon}{2} \cdot \beta$. Then, for any price vector $P \in [\epsilon \cdot \beta, +\infty)^n$, $\mathcal{R}_P = \tilde{\mathcal{R}}_P$, where $\hat{\mathcal{R}}_P$ and $\tilde{\mathcal{R}}_P$ are respectively the revenues of the seller when the buyer's values are distributed as $\{\hat{v}_i\}_{i \in [n]}$ and as $\{\tilde{v}_i\}_{i \in [n]}$.*

Proof. Recall that the variables $\{\tilde{v}_i\}_i$ are defined via a coupling with the \hat{v}_i 's. Under the same coupling, sample values from the \hat{v}_i 's and the \tilde{v}_i 's. For all items i such that $\hat{v}_i \neq \tilde{v}_i$, the price of item i is higher than the value of item i in both cases, so item i will not be purchased in both cases. That means the buyer will make the same decision in both cases, as she will only consider items whose values are the same. So the revenues are pointwise equal. □

Putting these lemmas together we obtain our reduction.

Proof of Theorem 41: Let P be a near-optimal price vector when the values of the buyer are distributed as $\{\tilde{v}_i\}_{i \in [n]}$, i.e. one that satisfies

$$\tilde{\mathcal{R}}_P \geq (1 - \delta) \cdot \widetilde{\text{OPT}},$$

where $\tilde{\mathcal{R}}_P$ denotes the expected revenue of the seller under price vector P when the buyer's values are $\{\tilde{v}_i\}_{i \in [n]}$. Given that each \tilde{v}_i lies in $[\frac{\epsilon}{2} \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]$, it follows from Lemma 24 that we can (efficiently) transform P into another vector $P' \in [\frac{\epsilon}{2} \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]^n$, such that $\tilde{\mathcal{R}}_P \leq \tilde{\mathcal{R}}_{P'}$.

We can then apply the following efficient transformation from P' to P'' : For any i , if $p'_i < \epsilon \cdot \beta$, set $p''_i = \epsilon \cdot \beta$, and set $p''_i = p'_i$ otherwise. By Lemma 27, we know that, $\tilde{\mathcal{R}}_{P''} \geq \tilde{\mathcal{R}}_{P'} - \epsilon \cdot \beta$.

Now, since P'' is a price vector in $[\epsilon \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]^n$, by Lemmas 43 and 44, we get $\mathcal{R}_{P''} \geq \tilde{\mathcal{R}}_{P''} - c_2(\epsilon) \cdot \beta$, where $\mathcal{R}_{P''}$ is the expected revenue of the seller under price vector P'' when the values of the buyers are $\{v_i\}_i$.

On the other hand, suppose that P^* is the optimal price vector in $[\epsilon \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]^n$ for values $\{v_i\}_{i \in [n]}$. By Lemma 40, we know that $\mathcal{R}_{P^*} \geq \left(1 - \frac{\epsilon + c_2(\epsilon)}{c_1}\right) OPT$. Now Lemmas 43 and 44 give

$$\tilde{\mathcal{R}}_{P^*} \geq \mathcal{R}_{P^*} - c_2(\epsilon) \cdot \beta \geq \left(1 - \frac{\epsilon + c_2(\epsilon)}{c_1}\right) OPT - c_2(\epsilon) \cdot \beta \geq \left(1 - \frac{\epsilon + 2c_2(\epsilon)}{c_1}\right) OPT,$$

where we used that $OPT \geq c_1 \cdot \beta$, by Lemma 38.

Since $\widetilde{OPT} \geq \tilde{\mathcal{R}}_{P^*}$,

$$\tilde{\mathcal{R}}_{P'} \geq \tilde{\mathcal{R}}_P \geq (1 - \delta) \widetilde{OPT} \geq (1 - \delta) \left(1 - \frac{\epsilon + 2c_2(\epsilon)}{c_1}\right) OPT \geq \left(1 - \delta - \frac{\epsilon + 2c_2(\epsilon)}{c_1}\right) OPT.$$

Recall that $\mathcal{R}_{P''} \geq \tilde{\mathcal{R}}_{P''} - c_2(\epsilon) \cdot \beta \geq \tilde{\mathcal{R}}_{P'} - \epsilon \cdot \beta - c_2(\epsilon) \cdot \beta$. Therefore,

$$\mathcal{R}_{P''} \geq \left(1 - \delta - \frac{\epsilon + 2c_2(\epsilon)}{c_1}\right) OPT - \epsilon \cdot \beta - c_2(\epsilon) \cdot \beta \geq \left(1 - \delta - \frac{2\epsilon + 3c_2(\epsilon)}{c_1}\right) OPT.$$

So given a near-optimal price vector P for $\{\tilde{v}_i\}_{i \in [n]}$, we can construct a near-optimal price vector P'' for $\{v_i\}_{i \in [n]}$ in polynomial time. \square

G Details of Section 9

We start by establishing some useful properties of regular distributions, and proceed to show our extreme value theorem (Theorem 21), and our reduction from item pricing problems with regular distributions to item pricing problems with bounded distributions (Theorem 20).

G.1 Basic Properties of Regular Distributions

If F is a differentiable continuous regular distribution, it is not hard to see the following: if $f(x) = 0$ for some $x \in (u_{min}^X, u_{max}^X)$, then $f(x') = 0$ for all $x' < x$ (as otherwise the definition of regularity would be violated.) Hence, if X is a random variable distributed according to F , it must be that $f(x) > 0$ for $x \in [u_{min}^X, u_{max}^X]$. So we can define F^{-1} on $[u_{min}^X, u_{max}^X]$, and it will be differentiable, since F is differentiable and f is non-zero. Now we can make the following definition, capturing the revenue of a seller who prices an item with value distribution F , so that the item is bought with probability exactly q .

Definition 45 (Revenue Curve). *For a differentiable continuous regular distribution F , define $R_F : [0, 1] \rightarrow \mathbb{R}$ as follows*

$$R_F(q) = q \cdot F^{-1}(1 - q).$$

The following is well-known. We include its short proof for completeness.

Lemma 46. *If F is regular, $R_F(q)$ is a concave function on $(0, 1]$.*

Proof. The derivative of $R_F(q)$ is

$$R'_F(q) = F^{-1}(1 - q) - \frac{q}{f(F^{-1}(1 - q))}.$$

Notice that $F^{-1}(1 - q)$ is monotonically non-increasing in q . This observation and the regularity of F imply that $R'_F(q)$ is monotonically non-increasing in q . (To see this try the change of variable $x(q) = F^{-1}(1 - q)$.) This implies that $R_F(q)$ is concave. \square

Lemma 47. For any regular distribution F , if $0 < \tilde{q} \leq q \leq p < 1$, then

$$R_F(\tilde{q}) \leq \frac{1}{1-p} R_F(q).$$

Proof. Since $q \in [\tilde{q}, 1)$, there exists a $\lambda \in (0, 1]$, such that

$$\lambda \cdot \tilde{q} + (1 - \lambda) \cdot 1 = q.$$

Hence: $\lambda = \frac{1-q}{1-\tilde{q}} \geq \frac{1-p}{1} = 1 - p$. Now, from Lemma 46, we have that $R_F(x)$ is concave. Thus

$$R_F(q) = R_F(\lambda \cdot \tilde{q} + (1 - \lambda) \cdot 1) \geq \lambda \cdot R_F(\tilde{q}) + (1 - \lambda) \cdot R_F(1).$$

Since $R_F(1) \geq 0$, $R_F(q) \geq \lambda \cdot R_F(\tilde{q}) \geq (1 - p)R_F(\tilde{q})$. Thus, $R_F(\tilde{q}) \leq \frac{1}{1-p} R_F(q)$. \square

Corollary 48. For any regular distribution F , if $\tilde{q} \leq q \leq \frac{1}{n^3}$, then

$$R_F(\tilde{q}) \leq \frac{n^3}{n^3 - 1} R_F(q).$$

G.2 Proof of Theorem 21: Extreme Value Theorem for Regular Distributions

We define α explicitly from the distributions $\{F_i\}_i$ of the variables $\{X_i\}_i$. We first need a definition.

Definition 49. A point x is a (c_1, c_2) -anchoring point of a distribution F , if $F(x) \in [c_1, c_2]$.

Now fix two arbitrary constants $0 < c_1 < c_2 \leq \frac{7}{8}$, and let, for all i , α_i be a (c_1, c_2) -anchoring point of the distribution F_i . Then define

$$\alpha = \frac{n^3}{c_1} \cdot \max_i \left[\alpha_i \cdot (1 - F_i(\alpha_i)) \right].$$

Clearly, a collection $\alpha_1, \dots, \alpha_n$ of (c_1, c_2) -anchoring points can be computed efficiently from the F_i 's. Hence, an α as above can be computed efficiently. We proceed to establish anchoring properties satisfied by α .

Proposition 50. $\alpha \geq \max_i \alpha_{n^3}^{(i)}$, where $\alpha_p^{(i)} = \inf \left\{ x \mid F_i(x) \geq 1 - \frac{1}{p} \right\}$ as in Definition 33.

Proof. Because $1/n^3 \leq 1 - c_2 \leq 1 - F(\alpha_i) \leq 1 - c_1$, it follows from Lemma 47 that

$$\frac{1}{c_1} \cdot \alpha_i \cdot (1 - F_i(\alpha_i)) \geq \alpha_{n^3}^{(i)}/n^3.$$

Hence: $\alpha \geq \frac{n^3}{c_1} \cdot \left[\alpha_i \cdot (1 - F_i(\alpha_i)) \right] \geq \alpha_{n^3}^{(i)}$. This is true for all i , hence the theorem. \square

Proof of Theorem 21: We first show that $\Pr[X_i \geq \ell\alpha] \leq 2/(\ell n^3)$, for any $\ell \geq 1$. By Proposition 50 and Corollary 48, we have that $(\ell\alpha) \Pr[X_i \geq \ell\alpha] \leq \frac{n^3}{n^3 - 1} \alpha \Pr[X_i \geq \alpha]$. Thus

$$\Pr[X_i \geq \ell\alpha] \leq \frac{n^3}{n^3 - 1} \cdot \frac{1}{\ell} \cdot \Pr[X_i \geq \alpha] \leq 2/(\ell n^3), \quad (17)$$

which establishes the first anchoring property satisfied by α .

Moreover, we have that

$$\alpha/n^3 = \frac{1}{c_1} \cdot \max_i \left[a_i \cdot (1 - F_i(a_i)) \right] \leq \frac{1}{c_1} \max_z (z \cdot \Pr[\max_i \{X_i\} \geq z]),$$

which establishes the second anchoring property of α .

Finally, we demonstrate the homogenization property of α . We want to show that, for any integer $m \leq n$, thresholds $t_1, \dots, t_m \geq t \geq \frac{2n^2\alpha}{\epsilon^2}$, index set $S = \{a_1, \dots, a_m\} \subseteq [n]$, and $\epsilon \in (0, 1)$:

$$\sum_{i=1}^m t_i \Pr[X_{a_i} \geq t_i] \leq \left(t - \frac{2\alpha}{\epsilon} \right) \cdot \Pr \left[\max_i \{X_{a_i}\} \geq t \right] + \frac{7\epsilon \cdot (2\alpha/\epsilon \cdot \Pr[\max_i \{X_{a_i}\} \geq 2\alpha/\epsilon])}{n}. \quad (18)$$

For notational simplicity, we define $f_i(z_i) = z_i \cdot \Pr[X_{a_i} \geq z_i]$ and $f_{max}^{(S)}(z) = z \cdot \Pr[\max_i \{X_{a_i}\} \geq z]$. Notice that for any $t_i \geq t \geq 2\alpha/\epsilon$, a double application of Proposition 50, Lemma 47 and Equation (17) gives

$$f_i(t_i) \leq \frac{(n^3/\epsilon)}{(n^3/\epsilon) - 1} f_i(t) \leq \frac{2(n^3/\epsilon)}{(n^3/\epsilon) - 1} f_i \left(\frac{2\alpha}{\epsilon} \right). \quad (19)$$

Thus,

$$\begin{aligned} \text{LHS of (18)} &\leq \sum_{i=1}^m f_i(t) + \frac{1}{(n^3/\epsilon) - 1} \sum_{i=1}^m f_i(t) \\ &\leq \sum_{i=1}^m f_i(t) + \frac{2}{(n^3/\epsilon) - 1} \sum_{i=1}^m f_i \left(\frac{2\alpha}{\epsilon} \right) \\ &\leq \sum_{i=1}^m f_i(t) + \frac{2n}{(n^3/\epsilon) - 1} f_{max}^{(S)} \left(\frac{2\alpha}{\epsilon} \right) \\ &\leq \sum_{i=1}^m f_i(t) + \frac{2\epsilon}{n} f_{max}^{(S)} \left(\frac{2\alpha}{\epsilon} \right). \end{aligned}$$

On the other hand, for any $t \geq 2\alpha/\epsilon$: $\Pr[X_{a_i} \geq t] \leq \Pr[X_{a_i} \geq 2\alpha/\epsilon] \leq \epsilon/n^3$ (using (17)). Thus:

$$\sum_i \Pr[X_{a_i} \geq t] \geq \Pr[\max_i \{X_{a_i}\} \geq t] \geq (1 - \epsilon/n^2) \sum_i \Pr[X_{a_i} \geq t], \quad (20)$$

where the last inequality follows from the fact that, for all i , the probability that $X_{a_i} \geq t$, while $X_{a_j} < t$ for all $j \in S \setminus \{i\}$ is at least $\Pr[X_{a_i} \geq t](1 - \epsilon/n^3)^{m-1} \geq \Pr[X_{a_i} \geq t](1 - \epsilon/n^2)$. Therefore,

continuing our upper-bounding from above:

$$\begin{aligned}
\text{LHS of (18)} &\leq \sum_{i=1}^m f_i(t) + \frac{2\epsilon}{n} f_{\max}^{(S)}\left(\frac{2\alpha}{\epsilon}\right) \\
&\leq (t - 2\alpha/\epsilon) \Pr[\max_i \{X_{a_i}\} \geq t] + (2\alpha/\epsilon) \Pr[\max_i \{X_{a_i}\} \geq t] + (\epsilon/n^2) \sum_{i=1}^m f_i(t) + \frac{2\epsilon}{n} f_{\max}^{(S)}\left(\frac{2\alpha}{\epsilon}\right) \\
&\leq (t - 2\alpha/\epsilon) \Pr[\max_i \{X_{a_i}\} \geq t] + (2\alpha/\epsilon t) \sum_{i=1}^m f_i(t) + (2\epsilon/n^2) \sum_{i=1}^m f_i\left(\frac{2\alpha}{\epsilon}\right) + \frac{2\epsilon}{n} f_{\max}^{(S)}\left(\frac{2\alpha}{\epsilon}\right) \\
&\leq (t - 2\alpha/\epsilon) \Pr[\max_i \{X_{a_i}\} \geq t] + (\epsilon/n^2) \sum_{i=1}^m f_i(t) + \frac{4\epsilon}{n} f_{\max}^{(S)}\left(\frac{2\alpha}{\epsilon}\right) \\
&\leq (t - 2\alpha/\epsilon) \Pr[\max_i \{X_{a_i}\} \geq t] + \frac{6\epsilon}{n} f_{\max}^{(S)}\left(\frac{2\alpha}{\epsilon}\right),
\end{aligned}$$

where we got the third inequality by invoking (19) and (20), the fourth inequality by invoking (20) with $t = 2\alpha/\epsilon$, and the fifth inequality by invoking (19) and then (20) with $t = 2\alpha/\epsilon$. This concludes the proof of Theorem 21. \square

G.3 Proof of Theorem 20: Reduction from Regular to Bounded Distributions

G.3.1 Restricting the Prices for the Input Regular Distributions

Lemma 51. *Let $\mathcal{V} = \{v_i\}_{i \in [n]}$ be a collection of independent regular value distributions, $\epsilon \in (0, 1)$, and c the absolute constant in the statement of Theorem 21. For any price vector P , we can construct a new price vector $\hat{P} \in [\epsilon\alpha/n^4, 2n^2\alpha/\epsilon^2]^n$, such that $\mathcal{R}_{\hat{P}} \geq \mathcal{R}_P - \frac{(c+9)\epsilon\mathcal{R}_{OPT}}{n}$, where \mathcal{R}_P and $\mathcal{R}_{\hat{P}}$ are respectively the expected revenues under price vectors P and \hat{P} , and \mathcal{R}_{OPT} is the optimal expected revenue for \mathcal{V} .*

Proof. First step: We first construct a price vector $P' \in [0, 2n^2\alpha/\epsilon^2]^n$ based on P , such that the revenue under P' is at most an additive $O(\frac{\epsilon\mathcal{R}_{OPT}}{n})$ smaller than the revenue under P .

We define P' as follows. Let $S = \{i \mid p_i > 2n^2\alpha/\epsilon^2\}$. For any $i \in S$ set $p'_i = \frac{2(n^2/\epsilon-1)\alpha}{\epsilon}$, while if $i \notin S$ set $p'_i = p_i$. Now assume $|S| = m$. For notational convenience we assume that $S = \{a_i \mid i \in [m]\}$, and set $X_{a_i} = v_{a_i}$. Moreover, let $t = \frac{2n^2\alpha}{\epsilon^2}$ and $t_i = p_{a_i}$.

Clearly, the contribution to \mathcal{R}_P from items in S is upper bounded by $\sum_{i=1}^m t_i \Pr[X_{a_i} \geq t_i]$. We proceed to analyze the contribution to revenue \mathcal{R}'_P from items in S . Notice that, when $\max_{i \in S} \{v_i\} = \max_i \{X_{a_i}\} \geq t$, the largest value-minus-price gap for items in S is at least $2\alpha/\epsilon$ (given our subtle choice of prices for items in S above). Hence, for the item of S achieving this gap not to be the winner, it must be that some item in $[n] \setminus S$ has a larger value-minus-price gap. For this to happen, the value for this item has to be higher than $2\alpha/\epsilon$. However, the probability that there exists an item in $[n] \setminus S$ with value greater than $2\alpha/\epsilon$ is smaller than $n \cdot \epsilon/n^3 = \epsilon/n^2$ (by Theorem 21). Thus, when $\max_i \{X_{a_i}\} \geq t$, then with probability at least $1 - \epsilon/n^2$, the item in S achieving the largest value-minus-price gap is the item bought by the buyer. So when the price vector is P' , the revenue from the items in S is lower bounded by $(t - 2\alpha/\epsilon) \Pr[\max_i \{X_{a_i}\} \geq t](1 - \epsilon/n^2)$ (where we used independence and the fact that $p'_i = t - 2\alpha/\epsilon$ for all $i \in S$).

Clearly, $(t - 2\alpha/\epsilon) \Pr[\max_i \{X_{a_i}\} \geq t] \leq t \Pr[\max_i \{X_{a_i}\} \geq t] \leq \mathcal{R}_{OPT}$. To see this, notice that the first inequality is obvious and the second follows from the observation that we could set the

prices of all items in S to t and of all other items to $+\infty$ to achieve revenue $t \Pr[\max_i \{X_{a_i}\} \geq t]$. So \mathcal{R}_{OPT} should be larger than this revenue. Similarly, we see that $2\alpha/\epsilon \cdot \Pr[\max_i X_{a_i} \geq 2\alpha/\epsilon] \leq \mathcal{R}_{OPT}$. Using these observations and Theorem 21 we get

$$\begin{aligned} & (t - 2\alpha/\epsilon) \Pr[\max_i \{X_{a_i}\} \geq t] (1 - \epsilon/n^2) + \frac{8\epsilon \cdot \mathcal{R}_{OPT}}{n} \\ & \geq (t - 2\alpha/\epsilon) \Pr[\max_i \{X_{a_i}\} \geq t] + \frac{7\epsilon \cdot (2\alpha/\epsilon \cdot \Pr[\max_i \{X_{a_i}\} \geq 2\alpha/\epsilon])}{n} \\ & \geq \sum_{i=1}^m t_i \Pr[X_{a_i} \geq t_i]. \end{aligned}$$

The above imply that the contribution to $\mathcal{R}_{P'}$ from the items in S is at most an additive $\frac{8\epsilon \cdot \mathcal{R}_{OPT}}{n}$ smaller than the contribution to \mathcal{R}_P from the items in S .

We proceed to compare the contributions from the items in $[n] \setminus S$ to \mathcal{R}_P and $\mathcal{R}_{P'}$. We start with \mathcal{R}_P . The contribution from the items in $[n] \setminus S$ is no greater than the total revenue when we ignore the existence of the items in S (e.g. by setting the prices of these items to $+\infty$), since this only boosts the winning probabilities of each item in $[n] \setminus S$.

Under price vector P' , $\forall i \in S$, $\Pr[v_i \geq p'_i] \leq \frac{\epsilon}{n^3}$ (Theorem 21). So with probability at least $1 - \frac{\epsilon}{n^2}$, no item in S has a positive value-minus-price gap and the item that has the largest positive gap among the items in $[n] - S$ is the item that is bought by the buyer. Hence, by independence the contribution to $\mathcal{R}_{P'}$ from the items in $[n] - S$ is at least a $1 - \frac{\epsilon}{n^2}$ fraction of the revenue when the items of S are ignored.

By the above discussion, the contribution to $\mathcal{R}_{P'}$ from the items in $[n] - S$ is at most an additive $\frac{\epsilon \mathcal{R}_{OPT}}{n^2}$ smaller than the contribution to \mathcal{R}_P from the items in $[n] - S$.

Putting everything together, we get that $\mathcal{R}_{P'} \geq \mathcal{R}_P - \frac{9\epsilon \mathcal{R}_{OPT}}{n}$.

Second step: To truncate the lower prices, we invoke Lemma 27. This implies that we can set all the prices below $\epsilon\alpha/n^4$ to $\epsilon\alpha/n^4$, only hurting our revenue by an additive $\epsilon\alpha/n^4 \leq \frac{\epsilon}{n} \cdot \max_z (z \cdot \Pr[\max_i \{X_i\} \geq z]) \leq c\epsilon \mathcal{R}_{OPT}/n$ (where we used Theorem 21 for the first inequality).

Hence, we can define \hat{P} as follows: if $p'_i \leq \epsilon\alpha/n^4$, set $\hat{p}_i = \epsilon\alpha/n^4$, otherwise set $\hat{p}_i = p'_i$. It follows from the above that $\mathcal{R}_{\hat{P}} \geq \mathcal{R}_P - \frac{(c+9)\epsilon \mathcal{R}_{OPT}}{n}$. \square

Thus, we have reduced the problem of finding a near-optimal price vector in $[0, +\infty]^n$ to the problem of finding a near-optimal price vector in the set $[\epsilon\alpha/n^4, 2n^2\alpha/\epsilon^2]^n$.

G.3.2 Truncating the Support of the Input Regular Distributions

We show that we can truncate the support of the distributions if the price vectors are restricted. Namely

Lemma 52. *Given a collection of independent regular random variables $\mathcal{V} = \{v_i\}_{i \in [n]}$ and any $\epsilon \in (0, 1)$, let us define a new collection of random variables $\tilde{\mathcal{V}} = \{\tilde{v}_i\}_{i \in [n]}$ via the following coupling: for all $i \in [n]$, set $\tilde{v}_i = \frac{\epsilon\alpha}{4n^4}$ if $v_i < \frac{\epsilon\alpha}{2n^4}$, set $\tilde{v}_i = 4n^4\alpha/\epsilon^3$, if $v_i \geq 4n^4\alpha/\epsilon^3$, and $\tilde{v}_i = v_i$ otherwise. Also, let c be the absolute constant defined in Theorem 21. For any price vector $P \in [\epsilon\alpha/n^4, 2n^2\alpha/\epsilon^2]^n$, $|\mathcal{R}_P(\mathcal{V}) - \mathcal{R}_P(\tilde{\mathcal{V}})| \leq \frac{c\epsilon \mathcal{R}_{OPT}(\mathcal{V})}{n}$, where $\mathcal{R}_P(\mathcal{V})$ and $\mathcal{R}_P(\tilde{\mathcal{V}})$ are respectively the revenues of the seller under price vector P when the values of the buyer are \mathcal{V} and $\tilde{\mathcal{V}}$.*

Proof. First, let us define another collection of mutually independent random variables $\hat{\mathcal{V}} = \{\hat{v}_i\}_{i \in [n]}$ via the following coupling: for all $i \in [n]$ set $\hat{v}_i = 4n^4\alpha/\epsilon^3$ if $v_i \geq 4n^4\alpha/\epsilon^3$, and set $\hat{v}_i = v_i$ otherwise.

By Theorem 21, we know that for every i , $\Pr[v_i \geq 4n^4\alpha/\epsilon^3] \leq \frac{\epsilon^3}{2n^7}$. Hence, the probability of the event that there exists an i such that $v_i \geq 4n^4\alpha/\epsilon^3$ is no greater than $n \times \epsilon^3/2n^7 = \epsilon^3/2n^6$. Thus the difference between the contributions of this event to the revenues $\mathcal{R}_P(\mathcal{V})$ and $\mathcal{R}_P(\hat{\mathcal{V}})$ is no greater than $2n^2\alpha/\epsilon^2 \cdot (\epsilon^3/2n^6) = \frac{\epsilon\alpha}{n^4} \leq \frac{c\epsilon}{n} \cdot \max_z(z \cdot \Pr[\max_i\{v_i\} \geq z]) \leq \frac{c\epsilon\mathcal{R}_{OPT}}{n}$, given that the largest price is at most $2n^2\alpha/\epsilon^2$.

Now let us consider the event: $v_i \leq 4n^4\alpha/\epsilon^3$, for all i . In this case $\hat{v}_i = v_i$ for all i . So the contribution of this event to the revenues $\mathcal{R}_P(\mathcal{V})$ and $\mathcal{R}_P(\hat{\mathcal{V}})$ is the same.

Thus, $|\mathcal{R}_P(\mathcal{V}) - \mathcal{R}_P(\hat{\mathcal{V}})| \leq \frac{c\epsilon\mathcal{R}_{OPT}}{n}$.

Now it follows from Lemma 44 that the seller's revenue under any price vector in $[\epsilon\alpha/n^4, 2n^2\alpha/\epsilon^2]^n$ is the same when the buyer's value distributions are $\hat{\mathcal{V}}$ and $\tilde{\mathcal{V}}$. \square

The above lemma shows that we can reduce the problem of finding a near-optimal price vector in $[\epsilon\alpha/n^4, 2n^2\alpha/\epsilon^2]^n$ for the original value distributions \mathcal{V} to the problem of finding a near-optimal price vector in the set $[\epsilon\alpha/n^4, 2n^2\alpha/\epsilon^2]^n$ for a collection of value distributions $\tilde{\mathcal{V}}$ supported on the set $[\frac{\epsilon\alpha}{4n^4}, 4n^4\alpha/\epsilon^3]$. Next, we establish that the latter problem can be reduced to finding any (i.e. not necessarily restricted) near-optimal price vector for the distributions $\tilde{\mathcal{V}}$.

Lemma 53. *Given a collection of independent regular random variables $\mathcal{V} = \{v_i\}_{i \in [n]}$ and any $\epsilon \in (0, 1)$, let us define a new collection of random variables $\tilde{\mathcal{V}} = \{\tilde{v}_i\}_{i \in [n]}$ via the following coupling: for all $i \in [n]$, set $\tilde{v}_i = \frac{\epsilon\alpha}{4n^4}$ if $v_i < \frac{\epsilon\alpha}{2n^4}$, set $\tilde{v}_i = 4n^4\alpha/\epsilon^3$ if $v_i \geq 4n^4\alpha/\epsilon^3$, and set $\tilde{v}_i = v_i$ otherwise. Let also c be the absolute constant defined in Theorem 21. For any price vector P , we can efficiently construct a new price vector $\hat{P} \in [\epsilon\alpha/n^4, 2n^2\alpha/\epsilon^2]^n$, such that $\mathcal{R}_{\hat{P}}(\tilde{\mathcal{V}}) \geq \mathcal{R}_P(\tilde{\mathcal{V}}) - \frac{(c+9)\epsilon\mathcal{R}_{OPT}(\tilde{\mathcal{V}})}{n}$.*

The proof is essentially the same as the proof of Lemma 51 and we skip it. Combining Lemmas 51, 52 and 53 we obtain Theorem 20. The proof is given in the next appendix.

G.3.3 Finishing the Reduction

Proof of Theorem 20: We start with computing α . This can be done efficiently as specified in the statement of Theorem 21. Now let us define $\tilde{\mathcal{V}}$ via the following coupling: for all $i \in [n]$, set $\tilde{v}_i = \frac{\epsilon\alpha}{4n^4}$ if $v_i < \frac{\epsilon\alpha}{2n^4}$, set $\tilde{v}_i = 4n^4\alpha/\epsilon^3$ if $v_i \geq 4n^4\alpha/\epsilon^3$, and set $\tilde{v}_i = v_i$ otherwise.

Now let P be a price vector such that $\mathcal{R}_P(\tilde{\mathcal{V}}) \geq (1 - \epsilon + \frac{(4c+19)\epsilon}{n}) \cdot \mathcal{R}_{OPT}(\tilde{\mathcal{V}})$. It follows from Lemma 53 that we can efficiently construct a price vector $P' \in [\epsilon\alpha/n^4, 2n^2\alpha/\epsilon^2]^n$, such that

$$\mathcal{R}_{P'}(\tilde{\mathcal{V}}) \geq \left(1 - \epsilon + \frac{(4c+19)\epsilon}{n}\right) \cdot \mathcal{R}_{OPT}(\tilde{\mathcal{V}}) - \frac{(c+9)\epsilon}{n} \mathcal{R}_{OPT}(\tilde{\mathcal{V}}) \geq \left(1 - \epsilon + \frac{(3c+10)\epsilon}{n}\right) \cdot \mathcal{R}_{OPT}(\tilde{\mathcal{V}}).$$

Lemma 51 implies that there exists a price vector $\hat{P} \in [\epsilon\alpha/n^4, 2n^2\alpha/\epsilon^2]^n$, such that $\mathcal{R}_{\hat{P}}(\mathcal{V}) \geq (1 - \frac{(c+9)\epsilon}{n}) \cdot \mathcal{R}_{OPT}(\mathcal{V})$. By Lemma 52, we know that

$$\mathcal{R}_{OPT}(\tilde{\mathcal{V}}) \geq \mathcal{R}_{\hat{P}}(\tilde{\mathcal{V}}) \geq \mathcal{R}_{\hat{P}}(\mathcal{V}) - \frac{c\epsilon}{n} \mathcal{R}_{OPT}(\mathcal{V}) \geq \left(1 - \frac{(2c+9)\epsilon}{n}\right) \cdot \mathcal{R}_{OPT}(\mathcal{V}).$$

So $\mathcal{R}_{P'}(\tilde{\mathcal{V}}) \geq (1 - \epsilon + \frac{c\epsilon}{n}) \cdot \mathcal{R}_{OPT}(\mathcal{V})$. We can now apply Lemma 52 again, and get

$$\mathcal{R}_{P'}(\mathcal{V}) \geq \mathcal{R}_{P'}(\tilde{\mathcal{V}}) - \frac{c\epsilon}{n} \mathcal{R}_{OPT}(\mathcal{V}) \geq (1 - \epsilon) \cdot \mathcal{R}_{OPT}(\mathcal{V}).$$

\square

H Algorithmic Results for MHR and Regular Distributions

The proofs of Theorems 3 and 4 follow immediately from Theorem 2 using our reductions to bounded distributions (Theorems 18 and 20 of Sections 8 and 9 respectively). We restate the theorems and prove them.

Theorem 3 [Restated] Suppose we are given a collection of MHR distributions $\{F_i\}_{i \in [n]}$. Then, for any constant $\epsilon > 0$, there is an algorithm that runs in time polynomial in the input and $n^{\frac{1}{\epsilon^7}}$ and computes a price vector P such that

$$\mathcal{R}_P \geq (1 - \epsilon)\mathcal{R}_{OPT},$$

where \mathcal{R}_P is the expected revenue under price vector P when the buyer's values for the items are independently distributed according to the distributions $\{F_i\}_i$ and \mathcal{R}_{OPT} is the revenue achieved by the optimal price vector.

Proof of Theorem 3: We apply Theorem 18 to reduce the item pricing problem for MHR distributions to the item pricing problem for bounded distributions. Then we use our algorithm from Theorem 2 for bounded distributions. The resulting running time is polynomial in the input and $n^{\frac{\log^4 \frac{1}{\epsilon}}{\epsilon^8}}$. Being a bit more careful in the application of our discretization lemmas we obtain running time polynomial in the input and n^{1/ϵ^7} . Recall that in the algorithm of Theorem 2 we employed the reduction of Theorem 17 to discretize supports and prices into sets of bounded cardinalities. To establish our reduction in Theorem 17 we employed Lemma 30, which in turn made use of Lemma 28, where we set $a = \frac{2}{3}$. Setting instead $a \approx \frac{1}{2}$ would result in a different tradeoff of parameters, improving our running time here. \square

Theorem 4 [Restated] Suppose we are given a collection of regular distributions $\{F_i\}_{i \in [n]}$. Then, for any constant $\epsilon > 0$, there exists an algorithm that runs in time polynomial in the input and $\max \left\{ n^{\log^{11} \frac{n}{\epsilon} \cdot \log \log \frac{n}{\epsilon}}, n^{\frac{\log^3 \frac{n}{\epsilon} \cdot \log \frac{1}{\epsilon}}{\epsilon^8}} \right\}$ and computes a price vector P such that

$$\mathcal{R}_P \geq (1 - \epsilon)\mathcal{R}_{OPT},$$

where \mathcal{R}_P is the expected revenue under price vector P when the buyer's values for the items are independently distributed according to the distributions $\{F_i\}_i$ and \mathcal{R}_{OPT} is the revenue achieved by the optimal price vector.

Proof of Theorem 4: We apply Theorem 20 to reduce the item pricing problem for regular distributions to the item pricing problem for bounded distributions. Then we use our algorithm from Theorem 2 for bounded distributions. The resulting running time is polynomial in the input and $\max \left\{ n^{\log^{11} \frac{n}{\epsilon} \cdot \log \log \frac{n}{\epsilon}}, n^{\frac{\log^3 \frac{n}{\epsilon} \cdot \log \frac{1}{\epsilon}}{\epsilon^8}} \right\}$. \square

I Proofs of Structural Results

Proof of Theorem 5: Let β be the anchoring point of Theorem 19. It follows from the properties of the anchoring point that pricing all the items at price $\beta/2$ achieves revenue

$$\frac{\beta}{2} \cdot \Pr[\max_i \{X_i\} \geq \beta/2] \geq \frac{\beta}{2} \cdot \left(1 - \frac{1}{\sqrt{e}}\right) = \beta \cdot c_1,$$

where $c_1 = \frac{1}{2} \left(1 - \frac{1}{\sqrt{e}}\right)$. On the other hand, Lemma 39 shows that the optimal revenue is upper bounded by

$$\beta \cdot \min_{\epsilon \in (0, \frac{1}{4})} \left(2 \log_2 \frac{1}{\epsilon} + c_2(\epsilon) \right),$$

where $c_2(\epsilon) = 36\epsilon \log_2(\frac{1}{\epsilon})$. So pricing all items at $\beta/2$ achieves a constant factor approximation to the optimal revenue. \square

Proof of Theorem 6: Suppose that the buyer's values are $\{v_i\}_{i \in [n]}$ where the v_i 's are mutually independent, MHR random variables. We can apply Lemma 40 to restrict the price-vectors to $[\epsilon \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]^n$, where $\beta = \beta(\{v_i\}_i)$ is the anchoring point of Theorem 19 computed from the distributions of the variables $\{v_i\}_i$. The loss in revenue from this restriction is bounded by Lemma 40. Having this restriction in place, we may now modify the variables $\{v_i\}_{i \in [n]}$ into a new collection of random variables $\{\tilde{v}_i\}_i$ as follows: for all $i \in [n]$, set $\tilde{v}_i = \frac{\epsilon}{2} \cdot \beta$ if $v_i < \epsilon \cdot \beta$, set $\tilde{v}_i = 2 \log_2(\frac{1}{\epsilon}) \cdot \beta$ if $v_i \geq 2 \log_2(\frac{1}{\epsilon}) \cdot \beta$, and set $\tilde{v}_i = v_i$ otherwise. Lemmas 43 and 44 show that the expected revenue of any price vector $P \in [\epsilon \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]^n$ is approximately the same for $\{v_i\}_{i \in [n]}$ and for $\{\tilde{v}_i\}_i$. Now we can apply Lemma 25 to discretize $[\epsilon \cdot \beta, 2 \log_2(\frac{1}{\epsilon}) \cdot \beta]^n$. The chain of reductions we used guarantees that a nearly-optimal among discretized prize-vectors for $\{\tilde{v}_i\}_i$ is also nearly-optimal among all possible price-vectors for $\{v_i\}_i$. \square

Proof of Theorem 7: Suppose that the buyer's values are $\{v_i\}_{i \in [n]}$ where the v_i 's are mutually independent, regular random variables. We can apply Lemma 51 to restrict the price-vectors to $[\epsilon\alpha/n^4, 2n^2\alpha/\epsilon^2]^n$ where α is chosen as in Appendix G.2. The loss in revenue from this restriction is bounded by Lemma 51. Having this restriction in place, we may now modify the variables $\{v_i\}_{i \in [n]}$ into a new collection of random variables $\{\tilde{v}_i\}_i$ as follows: for all $i \in [n]$, set $\tilde{v}_i = \frac{\epsilon\alpha}{4n^4}$ if $v_i < \frac{\epsilon\alpha}{2n^4}$, set $\tilde{v}_i = 4n^4\alpha/\epsilon^3$, if $v_i \geq 4n^4\alpha/\epsilon^3$, and $\tilde{v}_i = v_i$ otherwise. Lemma 52 shows that the expected revenue of any price vector $P \in [\epsilon\alpha/n^4, 2n^2\alpha/\epsilon^2]^n$ is approximately the same for $\{v_i\}_{i \in [n]}$ and for $\{\tilde{v}_i\}_i$. Now we can apply Lemma 25 to discretize $[\epsilon\alpha/n^4, 2n^2\alpha/\epsilon^2]^n$. The chain of reductions we used guarantees that a nearly-optimal among discretized prize-vectors for $\{\tilde{v}_i\}_i$ is also nearly-optimal among all possible price-vectors for $\{v_i\}_i$. \square

I.1 Proof of Theorem 8: A Single Price Suffices for I.I.D. MHR Distributions

We improve the running time of Theorem 3 for when the buyer's values are i.i.d. according to some MHR distribution. The main technical idea that goes into the algorithm is establishing our structural result for i.i.d. MHR distributions described by Theorem 8. In particular, we show that, if the number of items is a sufficiently large function of $1/\epsilon$, then using a single price suffices to get an $(1 - \epsilon)$ -fraction of the optimal revenue. Theorem 56 below summarizes the improvement on the running time as well as the structural result for i.i.d. MHR distributions.

We proceed to the details of our algorithm. To simplify our notation, let us assume that all the v_i 's are independent copies of the random variable v , and denote the cumulative distribution function of v by F . Moreover, let $\alpha_n = \inf \{x | F(x) \geq 1 - \frac{1}{n}\}$ (as in Definition 33). We start by showing an analogue of Lemma 37.

Lemma 54. *If $S = \text{Con}[v \geq (1 + \epsilon)\alpha_n]$, then $S \leq \frac{6(1+\epsilon)\alpha_n}{n^{1+\epsilon}}$.*

Proof. By Lemma 34, we know that $(1 + \epsilon)\alpha_n \geq \alpha_{n^{1+\epsilon}}$. Thus, $S \leq \text{Con}[v \geq \alpha_{n^{1+\epsilon}}]$. But Lemma 36 gives $\text{Con}[v \geq \alpha_{n^{1+\epsilon}}] \leq 6\alpha_{n^{1+\epsilon}}/n^{1+\epsilon}$. Hence,

$$S \leq \frac{6\alpha_{n^{1+\epsilon}}}{n^{1+\epsilon}} \leq \frac{6(1 + \epsilon)\alpha_n}{n^{1+\epsilon}}.$$

□

Using Lemma 54 and Lemma 42, we deduce that if we constrain our prices to be $\leq (1 + \epsilon)\alpha_n$, we lose no more than $\frac{6(1+\epsilon)\alpha_n}{n^\epsilon}$ revenue. Given that the optimal revenue with the restriction that all prices be $\leq (1 + \epsilon)\alpha_n$ is at most $(1 + \epsilon)\alpha_n$, it follows that the optimal revenue without the restriction is at most $(1 + \epsilon)\alpha_n + \frac{6(1+\epsilon)\alpha_n}{n^\epsilon} = (1 + \epsilon)(1 + \frac{6}{n^\epsilon})\alpha_n$. This is very close to α_n if n is a sufficiently large function of ϵ . If that's the case, it suffices to find a price vector achieving revenue close to α_n .

Lemma 55. *If we use the price vector $P = ((1 - \epsilon)\alpha_n, (1 - \epsilon)\alpha_n, \dots, (1 - \epsilon)\alpha_n)$, we receive revenue at least $(1 - e^{(-n^\epsilon)} - \epsilon)\alpha_n$.*

Proof. Let $p = (1 - \epsilon)\alpha_n$. By Lemma 34, we know that $\frac{\alpha_{n^{1-\epsilon}}}{(1-\epsilon)} \geq \alpha_n$. Hence, for all i ,

$$\Pr[v_i < p] \leq \Pr[v_i < \alpha_{n^{1-\epsilon}}] \leq 1 - \frac{1}{n^{1-\epsilon}}.$$

It follows that

$$\Pr[\exists i, v_i \geq p] \geq 1 - \left(1 - \frac{1}{n^{1-\epsilon}}\right)^n \geq 1 - e^{(-n^\epsilon)}.$$

Hence, with probability at least $1 - e^{(-n^\epsilon)}$, the buyer will purchase an item and will pay p . Hence, the revenue is at least $(1 - e^{(-n^\epsilon)})(1 - \epsilon)\alpha_n \geq (1 - e^{(-n^\epsilon)} - \epsilon)\alpha_n$. □

Notice that, when $n \geq (1/\epsilon)^{1/\epsilon}$, $n^\epsilon \geq 1/\epsilon$. In this case, we have shown that $OPT \leq (1 + \epsilon)(1 + 6\epsilon)\alpha_n \leq (1 + 8\epsilon)\alpha_n$. On the other hand, Lemma 55, says that we can achieve revenue at least $(1 - \frac{1}{e^{1/\epsilon}} - \epsilon)\alpha_n$ using a single price. Since $e^{1/\epsilon} \geq 1/\epsilon$, this revenue is at least $(1 - 2\epsilon)\alpha_n$. Given that $(1 + 8\epsilon)(1 - 10\epsilon) \leq (1 - 2\epsilon)$, we have $(1 - 2\epsilon)\alpha_n \geq (1 - 10\epsilon)OPT$. So if we set the price for every item to be $(1 - \epsilon)\alpha_n$, we achieve a revenue that is at least $(1 - 10\epsilon)OPT$.

Theorem 56. *If the values of the buyer are i.i.d. according to a MHR distribution, there is a PTAS for finding a price vector that achieves a $(1 - \epsilon)$ -fraction of the optimal revenue. The algorithm runs in time polynomial in $\log(\frac{\log n}{\epsilon})$, $2^{\frac{\log(1/\epsilon)}{\epsilon^8}}$ and the size of the input. Moreover, if $n \geq (12/\epsilon)^{12/\epsilon}$, there exists an efficiently computable price such that, if all items are priced at this price, the resulting revenue is at least $(1 - \epsilon)OPT$.*

Proof. Let $\epsilon' = \epsilon/12$. Depending on the value of n our algorithm proceeds in one of the following ways:

- If $n \geq (1/\epsilon')^{1/\epsilon'}$, we do binary search starting at an anchoring point of the distribution (see Appendix B) to find some $p \in [1 - \epsilon', 1 + \epsilon']\alpha_n$. This takes time polynomial in $O(\log(\frac{\log n}{\epsilon}))$ and the size of the input, since $\alpha_n \leq \alpha_2 \cdot \log_2 n$. We then set every item's price to $(1 - 2\epsilon')p$. Since $(1 - 2\epsilon')p \leq (1 - \epsilon')\alpha_n$,

$$\Pr[\exists i, v_i \geq (1 - 2\epsilon')p] \geq \Pr[\exists i, v_i \geq (1 - \epsilon')\alpha_n].$$

On the other hand, $(1 - 2\epsilon')p \geq (1 - 2\epsilon')(1 - \epsilon')\alpha_n$. Thus, the revenue we obtain if we price all items at $(1 - 2\epsilon')p$ is at least $(1 - 2\epsilon')$ times the revenue under price vector $P = ((1 - \epsilon')\alpha_n, (1 - \epsilon')\alpha_n, \dots, (1 - \epsilon')\alpha_n)$. Hence, the revenue is at least $(1 - 12\epsilon')OPT = (1 - \epsilon)OPT$.

- If $n < (1/\epsilon')^{1/\epsilon'}$, we simply use the algorithm for the non-i.i.d. case (Theorem 3).

□

References

- [Ala11] Saeed Alaei. Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers. In *the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2011.
- [Arm00] Mark Armstrong. Optimal Multi-Object Auctions. *The Review of Economic Studies*, 67(3):455–481, 2000.
- [BCKW10] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. Pricing Randomized Allocations. In *the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2010.
- [BGGM10] Sayan Bhattacharya, Gagan Goel, Sreenivas Gollapudi, and Kamesh Munagala. Budget Constrained Auctions with Heterogeneous Items. In *the 42nd ACM Symposium on Theory of Computing (STOC)*, 2010.
- [BH08] Liad Blumrosen and Thomas Holenstein. Posted Prices vs. Negotiations: An Asymptotic Analysis. In *the 9th ACM Conference on Electronic Commerce (EC)*, 2008.
- [BK07] Patrick Briest and Piotr Krysta. Buying Cheap is Expensive: Hardness of Non-Parametric Multi-Product Pricing. In *the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2007.
- [CD11] Yang Cai and Constantinos Daskalakis. Extreme-Value Theorems for Optimal Multidimensional Pricing. In *the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2011.
- [CDP⁺14] Xi Chen, Ilias Diakonikolas, Dimitris Paparas, Xiaorun Sun, and Mihalis Yannakakis. The Complexity of Optimal Multidimensional Pricing. In *the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2014.
- [CDW12a] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. An Algorithmic Characterization of Multi-Dimensional Mechanisms. In *the 43rd Annual ACM Symposium on Theory of Computing (STOC)*, 2012.
- [CDW12b] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization. In *the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2012.

- [CDW13] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Understanding Incentives: Mechanism Design becomes Algorithm Design. In *54th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2013.
- [CH13] Yang Cai and Zhiyi Huang. Simple and Nearly Optimal Multi-Item Auctions. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2013.
- [CHK07] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic Pricing via Virtual Valuations. In *the 8th ACM Conference on Electronic Commerce (EC)*, 2007.
- [CHMS10] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*, 2010.
- [CMS10] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. The Power of Randomness in Bayesian Optimal Mechanism Design. In *the 11th ACM Conference on Electronic Commerce (EC)*, 2010.
- [DDT13] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism Design via Optimal Transport. In *the 14th ACM Conference on Electronic Commerce (EC)*, 2013.
- [DDT14] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. The Complexity of Optimal Mechanism Design. In *25th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2014.
- [dHF06] Laurens de Haan and Ana Ferreira. *Extreme Value Theory: An Introduction*. Springer Series in Operations Research and Financial Engineering, 2006.
- [DW12] Constantinos Daskalakis and S. Matthew Weinberg. Symmetries and Optimal Multi-Dimensional Mechanism Design. In *the 13th ACM Conference on Electronic Commerce (EC)*, 2012.
- [HK05] Jason D. Hartline and Vladlen Koltun. Near-Optimal Pricing in Near-Linear Time. In *the 9th International Workshop on Algorithms and Data Structures (WADS)*, 2005.
- [MV06] Alejandro M. Manelli and Daniel R. Vincent. Bundling as an Optimal Selling Mechanism for a Multiple-Good Monopolist. *Journal of Economic Theory*, 127(1):1–35, 2006.
- [MV07] Alejandro M. Manelli and Daniel R. Vincent. Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly. *Journal of Economic Theory*, 137(1):153–185, 2007.
- [Mye81] Roger B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [Pav11] Gregory Pavlov. Optimal Mechanism for Selling Two Goods. *The BE Journal of Theoretical Economics*, 11(1), 2011.

[Roc85] Jean-Charles Rochet. The Taxation Principle and Multi-time Hamilton-Jacobi Equations. *Journal of Mathematical Economics*, 14(2):113–128, 1985.

[Tha04] John Thanassoulis. Hagglng Over Substitutes. *Journal of Economic Theory*, 117(2):217–245, 2004.

A An interesting example

A natural property than one would expect to hold is that, when the value distributions are discrete, there always exists an optimal solution that uses prices from the support of the value distributions. It turns out that this is not true. Here is an example:

Suppose that the seller has two items to sell, and the buyer's values for the items are v_1 , which is uniform on $\{1, 5\}$, and v_2 , which is uniform on $\{3, 3.5\}$. Moreover, assume that, if there is a tie between the value-minus-price gap for the two items, the buyer tie-breaks in favor of item 1. We claim that in this case the price vector $P = (4.5, 3)$ achieves higher revenue than any price vector that uses prices from the set $\{1, 3, 3.5, 5\}$ (where the values are drawn from.) Let us do the calculation. All our calculations are written in the form

$$\mathcal{R}_P = p_1 \times \Pr[\textit{item 1 is the winner}] + p_2 \times \Pr[\textit{item 2 is the winner}].$$

1. When $P = (4.5, 3)$

$$\mathcal{R}_P = 4.5 \times (1/2 \times 1) + 3 \times (1/2 \times 1) = 30/8$$

2. When $P \in \{1, 3, 3.5, 5\}^2$:

- If $P = (5, 3.5)$ then

$$\mathcal{R}_P = 5 \times (1/2 \times 1) + 3.5 \times (1/2 \times 1/2) = 27/8 < 30/8$$

- If $P = (5, 3)$ then

$$\mathcal{R}_P = 5 \times (1/2 \times 1/2) + 3 \times (1 \times 1/2 + 1/2 \times 1/2) = 28/8 < 30/8$$

- For any other price vector, the maximum revenue is bounded by $3.5 = 28/8 < 30/8$.