Computing Simple Mechanisms: Lift-and-Round over Marginal Reduced Forms

Yang Cai*
Computer Science Department
Yale University
USA
yang.cai@yale.edu

Argyris Oikonomou†
Computer Science Department
Yale University
USA
argyris.oikonomou@yale.edu

Mingfei Zhao‡
Google Research Mountain View
USA
mingfei@google.com

ABSTRACT

We study revenue maximization in multi-item multi-bidder auctions under the natural *item-independence* assumption – a classical problem in Multi-Dimensional Bayesian Mechanism Design. One of the biggest challenges in this area is developing algorithms to compute (approximately) optimal mechanisms that are not brute-force in the size of the bidder type space, which is usually exponential in the number of items in multi-item auctions. Unfortunately, such algorithms were only known for basic settings of our problem when bidders have unit-demand or additive valuations.

In this paper, we significantly improve the previous results and design the first algorithm that runs in time *polynomial in the number of items and the number of bidders* to compute mechanisms that are $O(1)$-approximations to the optimal revenue when bidders have XOS valuations, resolving an open problem raised by Chawla, Miller and Cai, Zhao. Moreover, the computed mechanism has a simple structure: It is either a posted price mechanism or a two-part tariff mechanism. As a corollary of our result, we show how to compute an approximately optimal and simple mechanism efficiently using *only sample access* to the bidders' value distributions. Our algorithm builds on two innovations that allow us to search over the space of mechanisms efficiently: (i) a new type of succinct representation of mechanisms – the *marginal reduced forms*, and (ii) a novel *Lift-and-Round procedure* that concavifies the problem.

CCS CONCEPTS

• **Theory of computation → Algorithmic mechanism design**.

KEYWORDS

Revenue Maximization, Simple and Approximately Optimal Auctions, Computational Complexity, XOS valuations, Linear Program Lifting

*Supported by a Sloan Foundation Research Fellowship and the NSF Award CCF-1942583 (CAREER).
†Supported by a Sloan Foundation Research Fellowship and the NSF Award CCF-1942583 (CAREER).
‡Work done in part while the author was studying at Yale University, supported by the NSF Award CCF-1942583 (CAREER).

STOC '22, June 20–24, 2022, Rome, Italy

ACM ISBN 978-1-4503-9264-8/22/06.
https://doi.org/10.1145/3519935.3520029

ACM Reference Format:
Yang Cai, Argyris Oikonomou, and Mingfei Zhao. 2022. Computing Simple Mechanisms: Lift-and-Round over Marginal Reduced Forms. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC '22), June 20–24, 2022, Rome, Italy*. ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3519935.3520029

1 INTRODUCTION

Revenue-maximization in multi-item auctions has been recognized as a central problem in Economics and more recently in Computer Science. While Myerson's celebrated work showed that a simple mechanism is optimal in single-item settings [32], the optimal multi-item mechanism is known to be prohibitively complex and notoriously difficult to characterize even in basic settings. Facing the challenge, a major research effort has been dedicated to understanding the computational complexity for finding an approximately revenue-optimal mechanism in multi-item settings. Despite significant progress, there is still a substantial gap in our understanding of the problem, for example, in the natural and extensively studied *item-independent* setting, first introduced in the influential paper by Chawla, Hartline, and Kleinberg [19].

Formally, the item-independent setting is defined as follows: A seller is selling m heterogeneous items to n bidders, where the i-th bidder's type is drawn independently from an m-dimensional product distribution $D_i = \bigtimes_{j\in[m]} D_{ij}$.[1] We only understand the computational complexity of finding the revenue-optimal mechanism in the item-independent setting for the two most basic valuations: unit-demand and additive valuations. First, we know that finding an exactly optimal mechanism is computationally intractable even for a single bidder with either unit-demand [23] or additive valuation [24]. Second, there exists a polynomial time algorithm that computes a mechanism whose revenue is at least a constant fraction of the optimal revenue when bidders have unit-demand [20, 21] or additive valuations [35]. However, unit-demand and additive valuations are only two extremes within a broader class of value functions known as the constrained additive valuations, where the bidder's value is additive subject to a downward-closed feasibility constraint.[2] Furthermore, all constrained additive valuations are contained in an even more general class known as the XOS valuations. Beyond unit-demand and additive valuations, our understanding was limited, and we only knew how to compute an

[1] $[m]$ denotes $\{1, 2, ..., m\}$. D_{ij} is the distribution of bidder i's value for item j. The definition is extended to XOS in Section 2.
[2] A bidder has constrained-additive valuation if the bidder's value for a bundle S is defined as $\max_{V\in 2^S\cap\mathcal{I}} \sum_{j\in V} t_j$, where t_j is the bidder's value for item j, and $\mathcal{I}$ is a downward-closed set system over the items specifying the feasible bundles. Note that constrained-additive valuations contain familiar valuations such as additive, unit-demand, or matroid-rank valuations.

approximately optimal mechanism when bidders are symmetric, i.e., all D_i's are identical [17, 22]. Finding a polynomial time algorithm for asymmetric bidders was thus raised as a major open problem in both [17, 22]. In this paper, we resolve this open problem.

Result I: For the item-independent setting with (asymmetric) XOS bidders, there exists an algorithm that computes a Dominant Strategy Incentive Compatible (DSIC) and Individually Rational (IR) mechanism whose revenue is at least $c \cdot \text{OPT}$ for some absolute constant $c > 0$, where OPT is the optimal revenue achievable by any Bayesian Incentive Compatible (BIC) and IR mechanism. Our algorithm has running time polynomial in $\sum_{i\in[n],j\in[m]} |\mathcal{T}_{ij}|$, where $\mathcal{T}_{ij}$ is the support of D_{ij}. See Theorem 1 for the formal statement.

Computing Approximately Optimal Mechanisms under Structured Distributions. When the bidders' types are drawn from arbitrary distributions, a line of works provide algorithms for finding almost revenue-optimal mechanisms in multi-item settings in time polynomial in the total number of types, i.e., $\sum_{i\in[n]} |\text{Supp}(D_i)|$ ($\text{Supp}(D_i)$ denotes the support of D_i) [2, 8, 9, 11, 12, 16].

However, the total number of types could be exponential in the number of items, e.g., there are $\sum_{i\in[n]} \left(\prod_{j\in[m]} |\mathcal{T}_{ij}|\right)$ types in the item-independent case, making these algorithms unsuitable. For unstructured type distributions, such dependence is unavoidable as even describing the distributions requires time $\Omega\left(\sum_{i\in[n]} |\text{Supp}(D_i)|\right)$. What if the type distributions are *structured and permit a more succinct description*, e.g., product measures? Arguably, high-dimensional distributions that arise in practice (such as bidders' type distributions in multi-item auctions) are rarely arbitrary, as arbitrary high-dimensional distributions cannot be represented or learned efficiently; see e.g. [25] for a discussion. Indeed, one of the biggest challenges in Bayesian Algorithmic Mechanism Design is designing algorithms to compute (approximately) optimal mechanisms that are not brute-force in the size of the bidder type space when the type distributions are structured. In this paper, we develop computational tools to exploit the item-independence to obtain an exponential speed-up in running time.

Simple vs. Optimal. An additional feature of our algorithm is that the mechanisms computed have a simple structure. It is either a *posted price mechanism* or a *two-part tariff mechanism*. Given the description of the two mechanisms, it is clear that both of them are DSIC and IR.

Rationed Posted Price Mechanism (RPP). There is a price p_{ij} for bidder i to purchase item j. The bidders arrive in some arbitrary order, and each bidder can purchase *at most one* item among the available ones at the given price.[3]

Two-part Tariff Mechanism (TPT). All bidders face the same set of prices $\{p_j\}_{j\in[m]}$. Bidders arrive in some arbitrary order. For each bidder, we show her the available items and the associated price for each item, then ask her to pay an entry fee depending on the bidder's identity and the available items. If the bidder accepts the entry fee, she proceeds to purchase any of the available items at the given prices; if she rejects the entry fee, then she pays nothing and receives nothing.

A recent line of works focus on designing simple and approximately optimal mechanisms [1, 3, 13, 14, 17, 19, 20, 22, 29, 30, 34, 35]. The main takeaway of these results is that in the item-independent setting, there exists a simple mechanism that achieves a constant fraction of the optimal revenue. The most general setting where such a simple $O(1)$-approximation is known is exactly the setting in **Result I**, where bidders have XOS valuations [17]. More specifically, [17] show that there is a RPP or TPT that achieves a constant fraction of the optimal revenue, however their result is purely existential and does not suggest how to compute these simple mechanisms. Our result makes their existential result constructive.

Finally, combining our result with the learnability result for multi-item auctions in [6], we can extend our algorithm to the case when we only have sample access to the distributions.

Result II: For constrained-additive bidders, there exists an algorithm that computes a simple, DSIC, and IR mechanism whose revenue is at least $c \cdot \text{OPT} - O(\varepsilon \cdot \text{poly}(n, m))$ for some absolute constant $c > 0$ in time polynomial in n, m, and $1/\varepsilon$, given sample access to bidders' type distributions, and assuming each bidder's value for each item lies in $[0, 1]$. See Theorem 4 for the formal statement.

Due to space limit, we only include the formal statements of our results and proof sketches. All details can be found in the full version of the paper on arXiv https://arxiv.org/abs/2111.03962.

1.1 Our Approach and Techniques

Our main technical contribution is *a novel relaxation of the revenue optimization problem that can be solved approximately in polynomial time* and *an accompanying rounding scheme that converts the solution to a simple and approximately optimal mechanism.*[4] Our first step is to replace the objective of revenue with a duality-based benchmark of the revenue proposed in [17]. One can view the new objective as maximizing the virtual welfare, similar to Myerson's elegant solution for the single-item case. The main difference is that, while one can use *a fixed set of virtual valuations for any allocation* in the single-item case, due to the multi-dimensionality of our problem, *the virtual valuations must depend on the allocation, causing the virtual welfare to be a non-concave function in the allocation.* In this paper, we develop algorithmic tools to concavify and approximately optimize the virtual welfare maximization problem. We believe our techniques will be useful to address other similar challenges in Multi-Dimensional Mechanism Design.

More specifically, for every BIC and IR mechanism $\mathcal{M}$ with allocation rule σ and payment rule p, one can choose a set of *dual parameters* $\theta(\sigma)$ based on σ to construct an upper bound $U(\sigma, \theta(\sigma))$ for the revenue of $\mathcal{M}$. We refer to θ as the dual parameters because θ corresponds to a set of "canonical" dual variables, which can be used to derive the virtual valuations via the Cai-Devanur-Weinberg duality framework [13]. The upper bound $U(\sigma, \theta(\sigma))$ is then simply the corresponding virtual welfare. The computational problem is

[3]Usually, posted price mechanisms do not restrict the maximum number of items a bidder can buy. We consider a rationed version of posted price mechanism to make the computational task easy.

[4]An influential framework known as the ex-ante relaxation has been widely used in Mechanism Design, but is insufficient for our problem. See Appendix B.2 for a detailed discussion.

to find an allocation σ that (approximately) maximizes $U(\sigma, \theta(\sigma))$. With such a σ, we could use the result in [17] to convert it to a simple and approximately optimal mechanism. Unfortunately, the function $U(\sigma, \theta(\sigma))$ is highly non-concave in σ,[5] and thus hard to maximize efficiently. See Section 3.1 for a detailed discussion.

LP Relaxation via Lifting. We further relax our objective, i.e., $U(\sigma, \theta(\sigma))$, to obtain a computationally tractable problem. One specific difficulty in optimizing $U(\sigma, \theta(\sigma))$ comes from the fact that $\theta(\sigma)$ is highly non-linear in σ. We address this difficulty in two steps. In the first step of our relaxation, we flip the dependence of σ and θ by relaxing the problem to the following two-stage optimization problem (Figure 1):

- Stage I: Maximize $H(\theta)$ subject to some constraints. $H(\theta)$ is the optimal value of the Stage II problem.

- Stage II: Maximize an LP over σ with θ-dependent constraints.

This makes the problem much more structured and significantly disentangles the complex dependence between σ and θ. Yet we still do not know how to solve it efficiently. In the second step of our relaxation, we merge the two-stage optimization into a single LP. In particular, we *lift the problem to a higher dimensional space and optimize over joint distributions of the allocation σ and the dual parameters θ via an LP (Figure 3).* Since the number of dual parameters is already exponential in the number of bidders and the number of items, it is too expensive to represent such a joint distribution explicitly. We show it is unnecessary to search over all joint distributions. By leveraging the independence across bidders and items, it suffices for us to consider a set of succinctly representable distributions – the ones whose marginals over the dual parameters are product measures. See Section 3.1 for a more detailed discussion on the development of our relaxation.

"Rounding" any Feasible Solution to a Simple Mechanism. Can we still approximate the optimal solution of the LP relaxation using a simple mechanism? Unfortunately, the result from [17] no longer applies. We provide a *generalization of [17]*, that is, given any feasible solution of our LP relaxation, we can construct in polynomial time a simple mechanism whose revenue is at least a constant fraction of the objective value of the feasible solution (Theorem 3). Our proof provides several novel ideas to handle the new challenges due to the relaxation, which may be of independent interest.

Marginal Reduced Forms. We deliberately postpone the discussion on how we represent the allocation of a mechanism until now. A widely used succinct representation a mechanism $\mathcal{M}$ is known as the reduced form or the interim allocation rule: $\{r_{ij}(t_i)\}_{i\in[n],j\in[m],t_i\in\times_{j\in[m]}\mathcal{T}_{ij}}$ where $r_{ij}(t_i)$ is the probability for bidder i to receive item j when *her type is* $t_i = (t_{i1}, \ldots, t_{im})$ [8]. Despite being more succinct than the ex-post allocation rule, the reduced form is still too expensive to store in our setting, as its size is exponential in m. A key innovation in our relaxation is the introduction of an even more succinct representation – the *marginal reduced forms* and a *multiplicative approximation* to the polytope of all feasible marginal reduced forms. Although this is a natural concept, to the best of our knowledge, we are the first to introduce and make use of it. The marginal reduced form is represented as : $\{w_{ij}(t_{ij})\}_{i\in[n],j\in[m],t_{ij}\in\mathcal{T}_{ij}}$, where $w_{ij}(t_{ij})$ is the probability for bidder i to receive item j in $\mathcal{M}$ and *her value for item j is t_{ij}*.[6] Importantly, the size of a marginal reduced form is polynomial in the input size of our problem. As our LP relaxation uses marginal reduced forms as decision variables, it is crucial for us to be able to optimize over the polytope P that contains all *feasible marginal reduced forms.* To the best of our knowledge, P does not have a succinct explicit description or an efficient separation oracle. To overcome the obstacle, we provide *an efficient separation oracle for a different polytope Q that is a multiplicative approximation to P, i.e., $c \cdot P \subseteq Q \subseteq P$ for some absolute constant $c \in (0, 1)$* (Theorem 2). Using the separation oracle for Q, we can find a c-approximation to the optimum of the LP relaxation efficiently. Note that a sampling technique was developed in [9] to approximate the polytope of feasible reduced forms. However, their technique only provides an "additive approximation to the polytope", which is insufficient for our purpose. Indeed, our multiplicative approximation holds for a wide class of polytopes that frequently appear in Mechanism Design. We believe our technique has further applications, for example, to convert the additive FPRAS of Cai-Daskalakis-Weinberg [8, 9, 11, 12] to a multiplicative FPRAS.

1.2 Related Work

Simple vs. Optimal. We provide an algorithm for the most general setting where an $O(1)$-approximation to the optimal revenue is known using simple mechanisms. It is worth mentioning that a recent result by Dütting et al. [27] shows that simple mechanisms can be used to obtain a $O(\log\log m)$-approximation to the optimal revenue even when the bidders have subadditive valuations. We leave it as an interesting open problem to extend our algorithm to bidders with subadditive valuations.

$(1-\varepsilon)$-Approximation in Item-Independent Settings. We focus on constant factor approximations for general valuations. For more specialized valuations, e.g., unit-demand/additive, there are several interesting results for finding $(1-\varepsilon)$-approximation to the "optimal mechanism". For example, PTASes are known if we restrict our attention to finding the optimal simple mechanism for a single bidder, e.g., item-pricing [5] or partition mechanisms [33]. For multiple bidders, PTASes are known for bidders with additive valuations under extra assumptions on distributions (such as i.i.d., MHR,[7] etc.) [14, 26]. The only result that does not require simplicity of the mechanism or extra assumptions on the distribution is [31], but their algorithm is only a quasi-polynomial time approximation scheme (QPTAS) and computes a $(1-\varepsilon)$-approximation to the optimal revenue for a single unit-demand bidder.

Structured Distributions beyond Item-Independence. When the type distributions can be represented as other structured distributions such as Bayesian networks, Markov Random Fields, or Topic Models, recent results show how to utilize the structure to improve the learnability, approximability, and communication complexity

[5] See Appendix B.1 for an example of the non-concavity of the function.

[6] We refer to $\{w_{ij}(t_{ij})\}_{i\in[n],j\in[m],t_{ij}\in\mathcal{T}_{ij}}$ as the marginal reduced forms as they are the marginals of the reduced forms multiplied by the probability that t_{ij} is bidder i's value for item j, i.e., $\frac{w_{ij}(t_{ij})}{\Pr_{D_{ij}}[t_{ij}]} = \mathbb{E}_{t_{i,-j}\sim\times_{\ell\neq j}D_{i\ell}}\left[r_{ij}\left(t_{ij}, t_{i,-j}\right)\right]$.

[7] That is, $f_{ij}(v)/1 - F_{ij}(v)$ is monotone non-decreasing (MHR) for each i, j, where f_{ij} is the pdf and F_{ij} is the cdf.

of multi-item auctions [4, 7, 15]. We believe that tools developed in this work would be useful to obtain similar improvement in terms of the computational complexity for computing approximately optimal mechanisms for structured distributions beyond item-independence.

2 PRELIMINARIES

We focus on revenue maximization in the combinatorial auction with n independent bidders and m heterogeneous items. We denote bidder i's type t_i as $\{t_{ij}\}_{j\in[m]}$, where t_{ij} is bidder i's private information about item j. For each i, j, we assume t_{ij} is drawn independently from the distribution D_{ij}. Let $D_i = \times_{j=1}^m D_{ij}$ be the distribution of bidder i's type and $D = \times_{i=1}^n D_i$ be the distribution of the type profile. We only consider discrete distributions in this paper. We use $\mathcal{T}_{ij}$ (or $\mathcal{T}_i, \mathcal{T}$) and f_{ij} (or f_i, f) to denote the support and the probability mass function of D_{ij} (or D_i, D). For notational convenience, we let t_{-i} to be the types of all bidders except i and $t_{<i}$ (or $t_{\leq i}$) to be the types of the first $i-1$ (or i) bidders. Similarly, we define D_{-i}, $\mathcal{T}_{-i}$ and f_{-i} for the corresponding distribution, support of the distribution, and probability mass function.

Valuation Functions. For every bidder i, denote her valuation function as $v_i(\cdot,\cdot) : \mathcal{T}_i \times 2^{[m]} \to \mathbb{R}_+$. For every $t_i \in \mathcal{T}_i, S \subseteq 2^{[m]}$, $v_i(t_i, S)$ is bidder i's value for receiving a set S of items, when her type is t_i. In the paper, we are interested in constrained-additive and XOS valuations. For every $i \in [n]$, bidder i's valuation $v_i(\cdot,\cdot)$ is *constrained-additive* if the bidder can receive a set of items subject to some downward-closed feasibility constraint $\mathcal{F}_i$. Formally, $v_i(t_i, S) = \max_{R\in 2^S\cap\mathcal{F}} \sum_{j\in R} t_{ij}$ for every type t_i and set S. It contains classic valuations such as additive ($\mathcal{F}_i = 2^{[m]}$) and unit-demand ($\mathcal{F}_i = \cup_{j\in[m]}\{j\}$). For constrained-additive valuations, we use t_{ij} to denote bidder i's value for item j. For every $i \in [n]$, bidder i's valuation $v_i(\cdot,\cdot)$ is *XOS (or fractionally-subadditive)* if each t_{ij} represents a set of K non-negative numbers $\{\alpha_{ij}^{(k)}(t_{ij})\}_{k\in[K]}$, for some integer K, and $v_i(t_i, S) = \max_{k\in[K]} \sum_{j\in S} \alpha_{ij}^{(k)}(t_{ij})$, for every type t_i and set S. We denote by $V_{ij}(t_i) = v_i(t_i, \{j\})$ the value for a single item j. Since the value of the bidder for item j only depends on t_{ij}, we denote $V_{ij}(t_{ij})$ as the singleton value.

Mechanisms. A mechanism $\mathcal{M}$ can be described as a tuple (σ, p), where σ is the *interim allocation* rule of $\mathcal{M}$ and p stands for the payment rule. Formally, for every bidder i, type t_i and set S, $\sigma_{iS}(t_i)$ is the interim probability that bidder i with type t_i receives exact bundle S. We use standard concepts of BIC, DSIC and IR for mechanisms. See Appendix A for the formal definitions. For any BIC and IR mechanism $\mathcal{M}$, denote $\text{Rev}(\mathcal{M})$ the revenue of $\mathcal{M}$. Denote OPT the optimal revenue among all BIC and IR mechanisms. Throughout this paper, the two classes of simple mechanisms we focus on are *rationed posted price* (RPP) mechanisms and *two-part tariff* (TPT) mechanisms, which are both described in Section 1. We denote PRev the optimum revenue achievable among all RPP mechanisms.

Access to the Bidders' Valuations. We define several ways to access a bidder's valuation.

Definition 1 (Value and Demand Oracle). *A* value oracle *for a valuation $v(\cdot,\cdot)$ takes a type t and a set of items $S \subseteq [m]$ as input, and returns the bidder's value $v(t, S)$ for the bundle S. A* demand oracle *for a valuation $v(\cdot,\cdot)$ takes a type t and a collection of non-negative prices $\{p_j\}_{j\in[m]}$ as input, and returns a utility-maximizing bundle, i.e. $S^* \in \arg\max_{S\subseteq[m]} \left(v(t, S) - \sum_{j\in S} p_j\right)$. In this paper, we use $\text{Dem}_i(\cdot,\cdot)$ to denote the demand oracle for bidder i's valuation $v_i(\cdot,\cdot)$.*

For constrained-additive valuations, our result only requires query access to a value oracle and a *demand oracle* for every bidder i's valuation $v_i(\cdot,\cdot)$. For XOS valuations, we need a stronger demand oracle that allows "scaled types" as input. We refer to the stronger oracle as the *adjustable demand oracle.*

Definition 2 (Adjustable Demand Oracle). *An* adjustable demand oracle *for bidder i's XOS valuation $v_i(\cdot,\cdot)$ takes a type t, a collection of non-negative coefficients $\{b_j\}_{j\in[m]}$, and a collection of non-negative prices $\{p_j\}_{j\in[m]}$ as input. For every item j, b_j is a scaling factor for t_{ij}, meaning that each of the K numbers $\{\alpha_{ij}^{(k)}(t_{ij})\}_{k\in[K]}$, i.e. the contribution of item j under each additive function, is multiplied by b_j. The oracle outputs a favorite bundle S^* with respect to the adjusted contributions and the prices $\{p_j\}_{j\in[m]}$, as well as the additive function $\{\alpha_{ij}^{(k^*)}(t_{ij})\}_{j\in[m]}$ for some $k^* \in [K]$ that achieves the highest value on S^*. Formally, $(S^*, k^*) \in \arg\max_{S\subseteq[m],k\in[K]} \left\{\sum_{j\in S} b_j\alpha_{ij}^{(k)}(t_{ij}) - \sum_{j\in S} p_j\right\}$. We use $ADem_i(\cdot,\cdot,\cdot)$ to denote the adjustable demand oracle for bidder i's XOS valuation $v_i(\cdot,\cdot)$.*

The adjustable demand oracle can be viewed as a generalization of the demand oracle for XOS valuations. In the above definition, if every coefficient b_j is 1, then the adjustable demand oracle outputs the utility-maximizing bundle S^* (as in the demand oracle) and the additive function that achieves the value for this set. For general b_j's, the adjustable demand oracle scales item j's contribution to bidder i's value by a b_j factor. The output bundle S^* maximizes the adjusted utility.[8]

Definition 3 (Bit Complexity of an Instance). *Given any instance of our problem represented as the tuple $(\mathcal{T}, D, v = \{v_i(\cdot,\cdot)\}_{i\in[n]})$. Denote as b_f the bit complexity of elements in $\{f_{ij}(t_{ij})\}_{i\in[n],j\in[m],t_{ij}\in\mathcal{T}_{ij}}$. For constrained-additive valuations, denote as b_v the bit complexity of elements in $\{t_{ij}\}_{i\in[n],j\in[m],t_{ij}\in\mathcal{T}_{ij}}$. For XOS valuations, denote as b_v the bit complexity of elements in $\{\alpha_{ij}^{(k)}(t_{ij})\}_{i\in[n],j\in[m],t_{ij}\in\mathcal{T}_{ij},k\in[K]}$. We define the value $\max(b_v, b_f)$ to be the bit complexity of the instance.*

3 LINEAR PROGRAM RELAXATION VIA LIFTING

In this section, we present the linear program relaxation for computing an approximately optimal simple mechanism. Due to the space limit, all proofs can be found in the arXiv version of the paper.

The main result of our paper is as follows:

[8]Note that for every collection of scaling factors, the query to the adjusted demand oracle is simply a demand query for a different XOS valuation. If all additive functions of t_i are explicitly given, then the adjusted demand oracle can be simulated in time $O(mK)$.

Theorem 1. *Let $T = \sum_{i,j} |\mathcal{T}_{ij}|$ and b be the bit complexity of the problem instance (Definition 3). For any $\delta > 0$, there exists an algorithm that computes a RPP mechanism or a TPT mechanism, such that the revenue of the mechanism is at least $c \cdot OPT$ for some absolute constant $c > 0$ with probability $1 - \delta - \frac{2}{nm}$. For constrained-additive valuations, our algorithm assumes query access to a value oracle and a demand oracle of bidders' valuations. For XOS valuations, our algorithm assumes query access to a value oracle and an adjustable demand oracle. The algorithm has running time $poly(n, m, T, b, \log(1/\delta))$.*

For any matroid-rank valuation, i.e., the downward-closed feasibility constraint is a matroid, the value and demand oracle can be simulated in polynomial time using greedy algorithms. For more general constraints, it is standard to assume access to the value and demand oracle. We also show that the adjustable demand oracle (rather than a demand oracle) is necessary to obtain our XOS result. In Theorem 6, we show that (even an approximation of) ADem_i can not be implemented in polynomial time, given access to the value oracle, demand oracle, and XOS oracle.

As most of the technical barriers already exist in the constrained-additive case, for exposition purposes, we focus on constrained-additive valuations in the main body (unless explicitly stated).[9] Before stating our LP, we first provide a brief recap of the existential result by Cai and Zhao [17] summarized in Lemma 1.[10]

Definition 4. *For any $i \in [n], j \in [m]$, and any feasible interim allocation σ,[11] and non-negative numbers $\tilde{\boldsymbol{\beta}} = \{\tilde{\beta}_{ij} \in \mathcal{T}_{ij}\}_{i\in[n],j\in[m]}$, $\mathbf{c} = \{c_i\}_{i\in[n]}$ and $\mathbf{r} = \{r_{ij}\}_{i\in[n],j\in[m]} \in [0,1]^{nm}$ (referred to as the dual parameters), let $\text{Core}(\sigma, \tilde{\boldsymbol{\beta}}, \mathbf{c}, \mathbf{r})$ be the welfare under allocation σ truncated at $\tilde{\beta}_{ij} + c_i$ for every i, j. Formally,*

$$\text{Core}(\sigma, \tilde{\boldsymbol{\beta}}, \mathbf{c}, \mathbf{r}) = \sum_i \sum_{t_i} f_i(t_i) \cdot \sum_{S \subseteq [m]} \sigma_{iS}(t_i) \cdot \sum_{j \in S} t_{ij} \cdot \left(\mathbb{1}[t_{ij} < \tilde{\beta}_{ij} + c_i] + r_{ij} \cdot \mathbb{1}[t_{ij} = \tilde{\beta}_{ij} + c_i] \right).$$

Lemma 1. *[17] Given any BIC and IR mechanism $\mathcal{M}$ with interim allocation σ, where $\sigma_{iS}(t_i)$ is the interim probability for bidder i to receive exactly bundle S when her type is t_i, there exist non-negative numbers $\tilde{\boldsymbol{\beta}}^{(\sigma)} = \{\tilde{\beta}_{ij}^{(\sigma)} \in \mathcal{T}_{ij}\}_{i\in[n],j\in[m]}$, $\mathbf{c}^{(\sigma)} = \{c_i^{(\sigma)}\}_{i\in[n]}$ and $\mathbf{r}^{(\sigma)} \in [0,1]^{nm}$ that satisfy*[12]

(1) $\sum_{i\in[n]} \left(\Pr_{t_{ij}}[t_{ij} > \tilde{\beta}_{ij}^{(\sigma)}] + r_{ij}^{(\sigma)} \cdot \Pr_{t_{ij}}[t_{ij} = \tilde{\beta}_{ij}^{(\sigma)}] \right) \le \frac{1}{2}, \forall j$

(2) $\frac{1}{2} \cdot \sum_{t_i \in \mathcal{T}_i} f_i(t_i) \cdot \sum_{S: j \in S} \sigma_{iS}(t_i) \le \Pr_{t_{ij}}[t_{ij} > \tilde{\beta}_{ij}^{(\sigma)}] + r_{ij}^{(\sigma)} \cdot \Pr_{t_{ij}}[t_{ij} = \tilde{\beta}_{ij}^{(\sigma)}], \forall i, j$

(3) $\sum_{i\in[n]} c_i^{(\sigma)} \le 8 \cdot \text{PRev}$

and the corresponding $\text{Core}(\sigma, \tilde{\boldsymbol{\beta}}^{(\sigma)}, \mathbf{c}^{(\sigma)}, \mathbf{r}^{(\sigma)})$ satisfies the following inequalities:

(4) $\text{Rev}(\mathcal{M}) \le 28 \cdot \text{PRev} + 4 \cdot \text{Core}(\sigma, \tilde{\boldsymbol{\beta}}^{(\sigma)}, \mathbf{c}^{(\sigma)}, \mathbf{r}^{(\sigma)})$,

(5) $\text{Core}(\sigma, \tilde{\boldsymbol{\beta}}^{(\sigma)}, \mathbf{c}^{(\sigma)}, \mathbf{r}^{(\sigma)}) \le 64 \cdot \text{PRev} + 8 \cdot \text{Rev}(\mathcal{M}_1^{(\sigma)})$, *where $\mathcal{M}_1^{(\sigma)}$ is some TPT mechanism.*

Remark 1. *For continuous type distributions, there exists $\tilde{\boldsymbol{\beta}}^{(\sigma)}$ that satisfy both Property 1 and 2 of Lemma 1 with $r_{ij}^{(\sigma)} = 1, \forall i, j$ for every σ. For discrete distributions, such a $\tilde{\boldsymbol{\beta}}^{(\sigma)}$ may not exist. This is simply a tie-breaking issue, and the role of $\mathbf{r}^{(\sigma)}$ is to fix it. Roughly speaking, $r_{ij}^{(\sigma)}$ is the probability that bidder i wins item j, when she is indifferent between purchasing or not. Readers can treat $\mathbf{r}^{(\sigma)}$ as the all-one vector to get the intuition behind our approach.*

By combining Property 4 and 5 of Lemma 1, Cai and Zhao [17] proved that the revenue of any BIC, IR mechanism $\mathcal{M}$ is bounded by a constant number of PRev and the revenue of some TPT mechanism. Recall that PRev is the optimal revenue achieved by an RPP mechanism, which is exactly the Sequential Posted Price mechanism if we restrict the bidders' valuations to unit-demand. Thus we can compute a set of posted prices that approximates PRev by Chawla et al. [18].

3.1 Tour to Our Relaxation

To facilitate our discussion about the key components and the intuition behind the relaxation, we present the development of our relaxation and along the way examine several failed attempts. In Theorem 3, we show that the optimal solution of the relaxed problem can indeed be approximated by simple mechanisms. Due to space limitations, we do not include details on the approximation analysis in this section, but focus on our intuition behind each step of our relaxation. We also assume r_{ij} to be 1 for every i and j to keep the notation light.

Step 0: Replace Revenue with the Duality-Based Benchmark. Instead of optimizing the revenue, we optimize the upper bound of revenue. As guaranteed by Lemma 1, for any BIC and IR mechanism $\mathcal{M} = (\sigma, p)$, its revenue is upper bounded by $O(\text{PRev} + \text{Core}(\sigma, \theta(\sigma)))$, where we use $\theta(\sigma)$ to denote the set of dual parameters $(\tilde{\boldsymbol{\beta}}^{(\sigma)}, \mathbf{c}^{(\sigma)})$ guaranteed to exist by Lemma 1. Since we can approximate PRev, it suffices to first approximately maximize $\text{Core}(\sigma, \theta(\sigma))$ over all feasible interim allocations σ, then compute the TPT in Lemma 1 based on the computed σ. $\text{Core}(\sigma, \theta(\sigma))$ is the truncated welfare, but the truncation depends on σ in a complex way, causing the function to be highly non-concave in σ (Example 1).

Step 1: Two-Stage Optimization. To overcome the barrier mentioned above, we consider a two-stage optimization problem (Figure 1) by *switching the order of dependence* between the interim allocation σ and dual parameters $\theta = (\boldsymbol{\beta}, \mathbf{c})$. In Stage I, we optimize some function H over the dual parameters $\theta = (\boldsymbol{\beta}, \mathbf{c})$, where $H(\boldsymbol{\beta}, \mathbf{c})$ is the optimum of the Stage II problem for every fixed set of parameters $(\boldsymbol{\beta}, \mathbf{c})$. Constraint (1) and (2) in the Stage I problem are due to Property 1 and 3 of Lemma 1 respectively. In Stage II, for any fixed set of parameters $\theta = (\boldsymbol{\beta}, \mathbf{c})$, we optimize $\text{Core}(\sigma, \theta)$ over all feasible σ such that the tuple $(\sigma, \boldsymbol{\beta}, \mathbf{c})$ satisfy Property 1,

[9] The linear program for XOS valuations can be found in Figure 5 in Appendix C.2.
[10] The statement is for constrained-additive bidders. The statement for XOS bidders can be found in the arxiv version of the paper.
[11] For constrained-additive bidders, an interim allocation σ is feasible if it can be implemented by a mechanism whose allocation rule always respects all bidders' feasibility constraints. It is without loss of generality to consider feasible interim allocations.
[12] [17] provides an explicit way to calculate $\tilde{\boldsymbol{\beta}}^{(\sigma)}, \mathbf{c}^{(\sigma)}, \mathbf{r}^{(\sigma)}$. We only include the crucial properties of these parameters here.

2, and 3 of Lemma 1. We choose the interim allocation σ as the variables, $\textsc{Core}(\sigma, \boldsymbol{\beta}, \mathbf{c})$ as the objective, and include Constraint (4), which corresponds to Property 2 of Lemma 1. Why is the two-stage optimization a relaxation? For any interim allocation σ, (i) the corresponding set of dual parameters $\theta(\sigma)$ is a feasible solution of the first-stage optimization problem, and (ii) σ is feasible in the second-stage optimization w.r.t. $\theta(\sigma)$, so $(\theta(\sigma), \sigma)$ is a feasible solution of the two-stage optimization problem.

Stage I:

$$\begin{aligned} &\max H(\boldsymbol{\beta}, \mathbf{c}) \\ \textbf{s.t.}\quad &(1)\quad \sum_{i\in[n]} \Pr_{t_{ij}}[t_{ij} \geq \beta_{ij}] \leq \frac{1}{2} \qquad \forall j \\ &(2)\quad \sum_{i\in[n]} c_i \leq 8 \cdot \textsc{PRev} \end{aligned}$$

Stage II:

$$\begin{aligned} H(\boldsymbol{\beta}, \mathbf{c}) = &\max \sum_{i\in[n]}\sum_{t_i\in\mathcal{T}_i} f_i(t_i)\cdot \sum_{S\subseteq[m]} \sigma_{iS}(t_i) \sum_{j\in S} t_{ij}\cdot \mathbb{1}[t_{ij} \leq \beta_{ij}+c_i] \\ \textbf{s.t.}\quad &(3)\quad \sigma \text{ is feasible} \\ &(4)\quad \frac{1}{2}\sum_{t_i\in\mathcal{T}_i} f_i(t_i)\cdot \sum_{S:j\in S} \sigma_{iS}(t_i) \leq \Pr_{t_{ij}}[t_{ij} \geq \beta_{ij}] \qquad \forall i, j \end{aligned}$$

Figure 1: Two-stage Optimization over $\theta = (\boldsymbol{\beta}, \mathbf{c})$ and the allocation σ

We now focus on the Stage II problem and try to solve it efficiently for a fixed set of parameters θ. The objective is a linear function of the variables σ, yet the set of variables $\sigma = \{\sigma_{iS}(t_i)\}_{i\in[n],S\subseteq[m],t_i\in\mathcal{T}_i}$ has exponential size. Luckily, the problem can be expressed more succinctly. For any interim allocation σ and dual parameters $\theta = (\boldsymbol{\beta}, \mathbf{c})$, the objective ($\textsc{Core}(\sigma, \theta)$) can be simplified as follows:

$$\begin{aligned} \textsc{Core}(\sigma, \theta) &= \sum_{\substack{i\in[n]\\ t_i\in\mathcal{T}_i}} f_i(t_i) \sum_{S\subseteq[m]} \sigma_{iS}(t_i) \sum_{j\in S} t_{ij}\cdot \mathbb{1}[t_{ij} \leq \beta_{ij}+c_i] \\ &= \sum_{\substack{i\in[n],j\in[m]\\ t_{ij}\in\mathcal{T}_{ij}}} \widehat{w}_{ij}(t_{ij}) t_{ij}\cdot \mathbb{1}[t_{ij} \leq \beta_{ij}+c_i], \end{aligned} \tag{1}$$

where $\widehat{w}_{ij}(t_{ij}) = f_{ij}(t_{ij}) \cdot \sum_{t_{i,-j}} f_{i,-j}(t_{i,-j}) \cdot \sum_{S:j\in S} \sigma_{iS}(t_{ij}, t_{i,-j})$ for every $i \in [n], j \in [m], t_{ij} \in \mathcal{T}_{ij}$. We refer to $\{\widehat{w}_{ij}(t_{ij})\}_{i\in[n],j\in[m],t_{ij}\in\mathcal{T}_{ij}}$ as the **marginal reduced form** of the interim allocation rule σ. $\widehat{w}_{ij}(t_{ij})$ represents the probability that bidder i's value for item j is t_{ij} and she receives item j, and the probability is taken over the randomness of the allocation, other bidders' types, as well as her own values for all the other items. Now for every fixed dual parameters θ, $\textsc{Core}$ is expressed as a linear function of the much more succinct representation $\widehat{w} = \{\widehat{w}_{ij}(t_{ij})\}_{i,j,t_{ij}}$ that has polynomial description size. We rewrite the Stage II problem as an LP using the variables $\widehat{w}$. Denote $\textsc{Core}(\hat{w}, \theta)$ the last term of Equation (1), which is the objective of the problem. By the definition of $\widehat{w}$, Constraint (4) is equivalent to

$$\frac{1}{2}\cdot \sum_{t_{ij}\in\mathcal{T}_{ij}} \widehat{w}_{ij}(t_{ij}) \leq \Pr_{t_{ij}}[t_{ij} \geq \beta_{ij}], \qquad \forall i, j \tag{2}$$

which is a linear constraint on $\widehat{w}$. Let $\mathcal{P}_1$ be the convex polytope that contains all marginal reduced forms $\widehat{w}$ that can be implemented by some feasible allocation σ (corresponds to Constraint (3)) and $\mathcal{P}_2$ be the set of all $\widehat{w}$ that satisfy all constraints in Equation (2). The Stage II problem is equivalent to the LP $\max_{\widehat{w}\in\mathcal{P}_1\cap\mathcal{P}_2} \textsc{Core}(\widehat{w}, \theta)$. Unfortunately, since $\mathcal{P}_1$ does not have an explicit succinct description or an efficient separation oracle, it is unclear if the problem can be solved efficiently.

Step 2: Marginal Reduced Form Relaxation. To overcome this barrier, we consider a relaxation of $\mathcal{P}_1$, where the feasibility constraint is only enforced on each bidder separately. We refer to this step as the *marginal reduced form relaxation.* We use $\widehat{w}_i = \{\widehat{w}_{ij}(t_{ij})\}_{j\in[m],t_{ij}\in\mathcal{T}_{ij}}$ to denote a feasible *single-bidder marginal reduced form* for bidder i. Formally, we define the feasible region W_i of $\widehat{w}_i$ in Definition 5.

Definition 5 (Constrained-additive valuations: single-bidder marginal reduced form polytope). *For every $i \in [n]$, suppose bidder i has a constrained-additive valuation with feasibility constraint $\mathcal{F}_i$. Bidder i's single-bidder marginal reduced form polytope $W_i \subseteq [0,1]^{\sum_{j\in[m]}|\mathcal{T}_{ij}|}$ is defined as follows: $\widehat{w}_i \in W_i$ if and only if there exists an allocation rule $\{\sigma_S(t_i)\}_{t_i\in\mathcal{T}_i,S\in\mathcal{F}_i}$, i.e., $\sigma_S(t_i)$ is the probability that i receives set S when her type is t_i, such that **(i)** $\sum_{S\in\mathcal{F}_i} \sigma_S(t_i) \leq 1$, $\forall t_i \in \mathcal{T}_i$, and **(ii)** $\widehat{w}_{ij}(t_{ij}) = f_{ij}(t_{ij}) \cdot \sum_{t_{i,-j}} f_{i,-j}(t_{i,-j}) \cdot \sum_{S:j\in S} \sigma_S(t_i)$, for all $j \in [m]$ and $t_{ij} \in \mathcal{T}_{ij}$.*

Throughout this section, we assume access to a separation oracle of W_i for every bidder i. In Theorem 2, we present an efficient separation oracle for another polytope $\widehat{W}_i$ that is a multiplicative approximation to W_i, i.e., $\widehat{W}_i$ is sandwiched between $c \cdot W_i$ and W_i for some absolute constant $c \in (0,1)$, using only queries to bidder i's demand oracle. We will argue later that we can efficiently approximate our problem with the separation oracle for $\widehat{W}_i$.

Here is our relaxation to the (rewritten) Stage II problem: Instead of forcing $\widehat{w}$ to be *implementable jointly* ($\widehat{w} \in \mathcal{P}_1$), we consider the relaxed region $\mathcal{P}' \supseteq \mathcal{P}_1$: $\widehat{w} \in \mathcal{P}'$ if and only if: (i) $\widehat{w}_i \in W_i$, for all bidder $i \in [n]$, and (ii) $\sum_i \sum_{t_{ij}} \widehat{w}_{ij}(t_{ij}) \leq 1, \forall j \in [m]$. In other words, $\mathcal{P}'$ guarantees that, for every bidder i, $\widehat{w}_i$ is a feasible single-bidder marginal reduced form for i, and the supply constraint is met in terms of marginal reduced forms (rather than ex-post allocations).

Relaxed Stage II:

$$\begin{aligned} H(\boldsymbol{\beta}, \mathbf{c}) = &\max \sum_{i\in[n]}\sum_{j\in[m]}\sum_{t_{ij}\in\mathcal{T}_{ij}} \widehat{w}_{ij}(t_{ij})\cdot t_{ij}\cdot \mathbb{1}[t_{ij} \leq \beta_{ij}+c_i] \\ \textbf{s.t.}\quad &(3)\quad \widehat{w}_i \in W_i \qquad \forall i \\ &(4)\quad \frac{1}{2}\sum_{t_{ij}} \widehat{w}_{ij}(t_{ij}) \leq \Pr_{t_{ij}}[t_{ij} \geq \beta_{ij}] \qquad \forall i, j \\ &\qquad \widehat{w}_{ij}(t_{ij}) \geq 0 \qquad \forall i, j, t_{ij} \end{aligned}$$

Figure 2: The Relaxed Stage II Problem over the Marginal Reduced Forms

The main benefit of this relaxation is computational. Without the relaxation, we need a multiplicative approximation of $\mathcal{P}_1$. Theorem 2 provides such an approximation if we can exactly maximizes the social welfare – a computational task that is substantially harder than answering demand queries. Indeed, we are not aware of any efficient algorithm that exactly maximizes the social welfare with only access to demand oracles of every bidder. The relaxed problem $\max_{\widehat{w}\in\mathcal{P}'\cap\mathcal{P}_2} \text{Core}(\widehat{w},\theta)$ is captured by the LP in Figure 2.[13]

Consider the two-stage optimization with the relaxed Stage II problem. For every fixed parameters θ, the relaxed Stage II problem can be solved efficiently (assuming a separation oracle of W_i for every i). Unfortunately, we do not know how to solve the two-stage optimization problem efficiently, as the number of different dual parameters is exponential in n and m, and enumerating through all possible choices of dual parameters is not an option. To overcome this obstacle, we need ideas explained in the following step.

Step 3: Lifting the problem to a higher dimensional space. Instead of enumerating all possible dual parameters θ, we optimize over *distributions of the parameters*. To guarantee that the number of decision variables in our program is polynomial, we focus on *product distributions* over the parameters. Formally, for every i, j, let C_{ij} be a distribution over $\mathcal{V}_{ij}\times\Delta$, where $\mathcal{V}_{ij}$ and Δ are the set of possible values of β_{ij} and c_i accordingly, after discretization (See Footnote *a* in Figure 3 for a formal definition). All C_{ij}'s are independent. In our program, we use decision variables $\{\hat{\lambda}_{ij}(\beta_{ij},\delta_{ij})\}_{i\in[n],j\in[m],\beta_{ij}\in\mathcal{V}_{ij},\delta_{ij}\in\Delta}$ to represent the distribution C_{ij}, i.e., $\hat{\lambda}_{ij}(a,b)=\Pr_{(\beta_{ij},\delta_{ij})\sim C_{ij}}[\beta_{ij}=a\wedge\delta_{ij}=b]$. Notice that if the parameters are drawn from a product distribution, the parameter "c_i" may be different for each item j. To distinguish them, we use δ_{ij} to replace the original parameter c_i in our program.

Now we maximize the expected value of the Core function over all product distributions $\times_{i,j}C_{ij}$ (represented by decision variables $\hat{\lambda}$) and the allocations (represented by the marginal reduced form $\widehat{w}$). If the parameters θ and allocation $\widehat{w}$ are generated independently, the expected Core is not a linear objective, since the contributed truncated welfare in Core is $\widehat{w}_{ij}(t_{ij})\cdot\hat{\lambda}_{ij}(\beta_{ij},\delta_{ij})\cdot t_{ij}\cdot\mathbb{1}[t_{ij}\le\beta_{ij}+\delta_{ij}]$ for every $t_{ij},\beta_{ij},\delta_{ij}$. *To linearize the objective, we lift the problem to a higher dimensional space and consider joint distributions over the parameters and allocations.* We do not consider arbitrary joint distributions, and only focus on the ones that correspond to the following generative process: first draw (β,δ) from a product distribution (according to $\hat{\lambda}$), then choose a feasible allocation $\widehat{w}^{(\beta,\delta)}=\{\widehat{w}_{ij}^{(\beta,\delta)}(t_{ij})\}_{i,j,t_{ij}}$ conditioned on (β,δ). Since there are too many parameters (β,δ), we certainly cannot afford to store all $\widehat{w}^{(\beta,\delta)}$'s explicitly. Instead, for each bidder i and item j we introduce a new set of decision variables $\{\lambda_{ij}(t_{ij},\beta_{ij},\delta_{ij})\}_{t_{ij}\in\mathcal{T}_{ij},\beta_{ij}\in\mathcal{V}_{ij},\delta_{ij}\in\Delta}$, where $\lambda_{ij}(t_{ij},\beta_{ij},\delta_{ij})$ is the marginal probability for the following three events to happen simultaneously in our generative process: **(a)** (β_{ij},δ_{ij}) are the parameters for i and j. **(b)** Bidder i receives item j. **(c)** Bidder i's value for item j is t_{ij}. Formally,

$$\begin{aligned}\lambda_{ij}(t_{ij},\beta_{ij},\delta_{ij}) =&\hat{\lambda}_{ij}(\beta_{ij},\delta_{ij})\\ &\cdot\sum_{\{(\beta_{i'j'},\delta_{i'j'})\}_{(i',j')\neq(i,j)}}\left(\widehat{w}_{ij}^{(\beta,\delta)}(t_{ij})/f_{ij}(t_{ij})\right)\\ &\cdot\prod_{(i',j')\neq(i,j)}\hat{\lambda}_{i'j'}(\beta_{i'j'},\delta_{i'j'})\end{aligned}\tag{3}$$

With the new variables $\lambda_{ij}(t_{ij},\beta_{ij},\delta_{ij})$'s, we can express the objective as an linear function:

$$\sum_{i\in[n]}\sum_{j\in[m]}\sum_{t_{ij}\in\mathcal{T}_{ij}}f_{ij}(t_{ij})\cdot t_{ij}\cdot\sum_{\beta_{ij}\in\mathcal{V}_{ij},\delta_{ij}\in\Delta}\lambda_{ij}(t_{ij},\beta_{ij},\delta_{ij})\cdot\mathbb{1}[t_{ij}\le\beta_{ij}+\delta_{ij}].$$

Our program can be viewed as an "expected version" of the the two-stage optimization, when the parameters $\theta=(\beta,\delta)\sim\bigtimes_{i,j}C_{ij}$. In other words, we only require the constraints to be satisfied in expectation. We discuss our relaxation in more details in Section 3.2.

3.2 Our LP and a Sketch of the Proof of Theorem 1

We present a sketch of the proof of Theorem 1 for constrained-additive bidders and our main linear program (Figure 3). Although the LP has many constraints and may seem intimidating at first, all constraints follow quite naturally from our derivation in Section 3.1. See Section 3.3 for more details.

The first step of our proof is to estimate PRev using Lemma 2 from [18].

Lemma 2 (Theorem 14 and Appendix F in [18]). *There exists an algorithm that with probability at least $1-\frac{2}{nm}$, computes a Rationed Posted Price mechanism $\mathcal{M}$ such that $\frac{1}{6.75}(1-\frac{1}{nm})\cdot$ PRev. The algorithm runs in time $poly(n,m,\sum_{i,j}|\mathcal{T}_{ij}|)$.*

Denote $\mathcal{E}$ the event that an RPP in Lemma 2 is computed successfully. For simplicity, we will condition on the event $\mathcal{E}$ for the rest of this section. Let $\widetilde{\text{PRev}}$ be the revenue of the RPP mechanism found in Lemma 2.

Next, we argue that the LP in Figure 3 (or Figure 5 when the valuations are XOS) can be solved efficiently. Note that there are $\text{poly}(n,m,\sum_{i,j}|\mathcal{T}_{ij}|)$ constraints except for Constraint (1), where we need to enforce the feasibility of single-bidder marginal reduced forms. It suffices to construct an efficient separation oracle for W_i for every i. However, to the best of our knowledge, W_i does not have a succinct explicit description or an efficient separation oracle. For constrained-additive valuations, we construct another polytope $\widehat{W}_i$ such that: (i) $\widehat{W}_i$ is a multiplicative approximation of W_i, i.e., $c\cdot W_i\subseteq\widehat{W}_i\subseteq W_i$ for some absolute constant $c\in(0,1)$, and (ii) There exists an efficient separation oracle for $\widehat{W}_i$ given access to the demand oracle.

Theorem 2. *Let $T=\sum_{i,j}|\mathcal{T}_{ij}|$ and b be the bit complexity of the problem instance (Definition 3). For any $i\in[n]$ and $\delta\in(0,1)$, there is an algorithm that constructs a convex polytope $\widehat{W}_i\in[0,1]^{\sum_{j\in[m]}|\mathcal{T}_{ij}|}$ using $poly(n,m,T,\log(1/\delta))$ samples from D_i, such that with probability at least $1-\delta$,*

[13] We omit the supply constraint $\sum_i\sum_{t_{ij}}\widehat{w}_{ij}(t_{ij})\le 1$ as it is implied by Constraint (1) in the Stage I problem and Constraint (4).

(1) $\frac{1}{12} \cdot W_i \subseteq \widehat{W}_i \subseteq W_i$, *and the vertex-complexity (Definition 11) of* $\widehat{W}_i$ *is* $poly(n, m, T, b, \log(1/\delta))$.

(2) There exists a separation oracle SO for $\widehat{W}_i$, *given access to the demand oracle for bidder i's valuation. The running time of SO on any input with bit complexity* b' *is* $poly(n, m, T, b, b', \log(1/\delta))$ *and makes* $poly(n, m, T, b, b', \log(1/\delta))$ *queries to the demand oracle.*

The algorithm constructs the polytope and the separation oracle SO in time $poly(n, m, T, b, \log(1/\delta))$.

Indeed, we prove a more general result regarding polytopes that can be expressed as a "mixture of polytopes", which can be viewed as a generalization of the technique developed in [10] for approximating the polytope of all feasible reduced forms.

To solve the LP relaxation, we replace W_i by $\widehat{W}_i$ in the LP in Figure 3 for every $i \in [n]$, and solve the LP in polynomial time using the ellipsoid method. Clearly, this solution is also feasible for the original LP in Figure 3. Moreover, since $\widehat{W}_i$ contains $c \cdot W_i$, we can show that **the objective value of the solution we computed is at least** $c \cdot \textbf{OPT}_{\textbf{LP}}$, where OPT_{LP} the optimum of the LP in Figure 3. Our proof of Theorem 2 heavily relies on the fact that W_i is a down-monotone polytope,[14] which does not hold in the XOS case. For XOS valuations, we construct the polytope $\widehat{W}_i$ with a weaker guarantee: For every vector x in W_i, there exists another vector x' in $\widehat{W}_i$ such that for every coordinate j, $x_j/x'_j \in [a, b]$ for some absolute constant $0 < a < b$, and vice versa.

Next, we argue that the LP optimum can be approximated by simple mechanisms. [17] shows that for any BIC and IR mechanism $\mathcal{M}$, $\textsc{Core}(\sigma, \tilde{\boldsymbol{\beta}}^{(\sigma)}, \mathbf{c}^{(\sigma)}, \mathbf{r}^{(\sigma)})$ (as stated in Lemma 1) can be bounded by a constant number of $\textsc{PRev}$ and the revenue of a TPT (see Property 5 of Lemma 1). We generalize their result by proving that for any feasible solution of the LP, its objective can be bounded by (a constant times) the revenue of a RPP or TPT mechanism, and the mechanism can be computed efficiently given the feasible solution.

Definition 6. *Let* $(w, \lambda, \hat{\lambda}, \boldsymbol{d} = (d_i)_{i\in[n]})$ *be any feasible solution of the LP in Figure 3. For every* $j \in [m]$, *define* $Q_j = \frac{1}{2} \cdot \sum_{i\in[n]} \sum_{t_{ij}\in\mathcal{T}_{ij}} f_{ij}(t_{ij}) \cdot t_{ij} \cdot \sum_{\beta_{ij}\in\mathcal{V}_{ij}, \delta_{ij}\in\Delta} \lambda_{ij}(t_{ij}, \beta_{ij}, \delta_{ij}) \cdot \mathbb{1}[t_{ij} \le \beta_{ij} + \delta_{ij}]$.[15]

Clearly, for any feasible solution of the LP, the objective function is $2 \cdot \sum_{j\in[m]} Q_j$. We prove in Theorem 3 that $2 \cdot \sum_{j\in[m]} Q_j$ can be bounded by the revenue of $\mathcal{M}_{\text{TPT}}$ (Mechanism 1) and the RPP $\mathcal{M}_{\text{PP}}$ (Lemma 2). As we can efficiently compute a feasible solution whose objective is $\Omega(\text{OPT}_{\text{LP}})$, Theorem 3 implies that we can compute in polynomial time a simple mechanism whose revenue is at least $\Omega(\text{OPT}_{\text{LP}} + \textsc{PRev})$.

Theorem 3. *Let* $(w, \lambda, \hat{\lambda}, \boldsymbol{d})$ *be any feasible solution of the LP in Figure 3. Let* $\mathcal{M}_{PP}$ *be the rationed posted price mechanism computed in Lemma 2. Let* $\mathcal{M}_{TPT}$ *be the two-part tariff mechanism shown in Mechanism 1 with prices* $\{Q_j\}_{j\in[m]}$. *Then the objective function of the solution* $2 \cdot \sum_{j\in[m]} Q_j$ *is bounded by* $c_1 \cdot \textsc{Rev}(\mathcal{M}_{PP}) + c_2 \cdot \textsc{Rev}(\mathcal{M}_{TPT})$, *for some absolute constant* $c_1, c_2 > 0$. *Moreover, both* $\mathcal{M}_{PP}$ *and* $\mathcal{M}_{TPT}$ *can be computed in time* $poly(n, m, \sum_{i,j} |\mathcal{T}_{ij}|)$ *with access to the demand oracle for the bidders' valuations.*

The proof of Theorem 3 combines the "shifted $\textsc{Core}$" technique by Cai and Zhao [17] with several novel ideas to handle the new challenges due to the relaxation.

$$\textbf{max} \sum_{i\in[n]} \sum_{j\in[m]} \sum_{t_{ij}\in\mathcal{T}_{ij}} f_{ij}(t_{ij}) \cdot t_{ij} \cdot \sum_{\beta_{ij}\in\mathcal{V}_{ij}, \delta_{ij}\in\Delta} \lambda_{ij}(t_{ij}, \beta_{ij}, \delta_{ij}) \cdot \mathbb{1}[t_{ij} \le \beta_{ij} + \delta_{ij}]$$

s.t.

Allocation Feasibility Constraints:

$$(1) \quad w_i \in W_i, \forall i$$

$$(2) \quad \sum_i \sum_{t_{ij}\in\mathcal{T}_{ij}} w_{ij}(t_{ij}) \le 1, \forall j$$

Natural Feasibility Constraints:

$$(3) \quad f_{ij}(t_{ij}) \cdot \sum_{\beta_{ij}\in\mathcal{V}_{ij}} \sum_{\delta_{ij}\in\Delta} \lambda_{ij}(t_{ij}, \beta_{ij}, \delta_{ij}) = w_{ij}(t_{ij}), \quad \forall i, j, t_{ij} \in \mathcal{T}_{ij}$$

$$(4) \quad \lambda_{ij}(t_{ij}, \beta_{ij}, \delta_{ij}) \le \hat{\lambda}_{ij}(\beta_{ij}, \delta_{ij}), \forall i, j, t_{ij}, \beta_{ij} \in \mathcal{V}_{ij}, \delta_{ij}$$

$$(5) \quad \sum_{\beta_{ij}\in\mathcal{V}_{ij}, \delta_{ij}\in\Delta} \hat{\lambda}_{ij}(\beta_{ij}, \delta_{ij}) = 1, \forall i, j$$

Problem Specific Constraints:

$$(6) \quad \sum_{i\in[n]} \sum_{\beta_{ij}\in\mathcal{V}_{ij}} \sum_{\delta_{ij}\in\Delta} \hat{\lambda}_{ij}(\beta_{ij}, \delta_{ij}) \cdot \Pr_{t_{ij}\sim D_{ij}}[t_{ij} \ge \beta_{ij}] \le \frac{1}{2}, \forall j$$

$$(7) \quad \frac{1}{2} \sum_{t_{ij}\in\mathcal{T}_{ij}} f_{ij}(t_{ij}) \left(\lambda_{ij}(t_{ij}, \beta_{ij}, \delta_{ij}) + \lambda_{ij}(t_{ij}, \beta^+_{ij}, \delta_{ij})\right) \le \hat{\lambda}_{ij}(\beta_{ij}, \delta_{ij}) \cdot \Pr_{t_{ij}}[t_{ij} \ge \beta_{ij}] + \hat{\lambda}_{ij}(\beta^+_{ij}, \delta_{ij}) \cdot \Pr_{t_{ij}}[t_{ij} \ge \beta^+_{ij}], \quad \forall i, j, \beta_{ij} \in \mathcal{V}^0_{ij}, \delta_{ij} \in \Delta$$

$$(8) \quad \sum_{\beta_{ij}\in\mathcal{V}_{ij}, \delta_{ij}\in\Delta} \delta_{ij} \cdot \hat{\lambda}_{ij}(\beta_{ij}, \delta_{ij}) \le d_i, \forall i, j$$

$$(9) \quad \sum_{i\in[n]} d_i \le 111 \cdot \widetilde{\textsc{PRev}}$$

$$\lambda_{ij}(t_{ij}, \beta_{ij}, \delta_{ij}) \ge 0, \hat{\lambda}_{ij}(\beta_{ij}, \delta_{ij}) \ge 0, \quad w_{ij}(t_{ij}) \ge 0, d_i \ge 0 \quad \forall i, j, t_{ij}, \beta_{ij} \in \mathcal{V}_{ij}, \delta_{ij}$$

Figure 3: LP Relaxation for Constrained-Additive Bidders

Remark 2. *Theorem 3 indeed holds even if the bidders arrive in an arbitrary order in* $\mathcal{M}_{TPT}$. *We choose the lexicographical order only to keep the notation light.*

We complete the last step of our proof by showing $\text{OPT} = O(\text{OPT}_{\text{LP}} + \textsc{PRev})$ in Lemma 3. More specifically, we show that for any mechanism $\mathcal{M} = (\sigma, p)$, the tuple $(\sigma, \tilde{\boldsymbol{\beta}}^{(\sigma)}, \mathbf{c}^{(\sigma)}, r^{(\sigma)})$ stated in Lemma 1 corresponds to a feasible solution of the LP in Figure 3 whose objective is at least $\textsc{Core}(\sigma, \tilde{\boldsymbol{\beta}}^{(\sigma)}, \mathbf{c}^{(\sigma)}, r^{(\sigma)})$. Hence, the revenue of $\mathcal{M}$ is upper bounded by $\textsc{PRev}$ and OPT_{LP}.

[14] A polytope $\mathcal{P} \subseteq [0, 1]^d$ is down-monotone if and only if for every $\boldsymbol{x} \in \mathcal{P}$ and $\boldsymbol{0} \le \boldsymbol{x}' \le \boldsymbol{x}$, we have $\boldsymbol{x}' \in \mathcal{P}$.

[15] Recall that $\lambda_{ij}(t_{ij}, \beta_{ij}, \delta_{ij})$ is introduced in Step 3 of Section 3.1. See Figure 3 for the formal definition.

Variables: [a]

- $\lambda_{ij}(t_{ij}, \beta_{ij}, \delta_{ij})$, for all i, j and $t_{ij} \in \mathcal{T}_{ij}$, $\beta_{ij} \in \mathcal{V}_{ij}$, $\delta_{ij} \in \Delta$. See Step 3 of Section 3.1 for an explanation of this variable.

- $\hat{\lambda}_{ij}(\beta_{ij}, \delta_{ij})$, for all i, j, $\beta_{ij} \in \mathcal{V}_{ij}, \delta_{ij} \in \Delta$, denoting the distribution C_{ij} over $\mathcal{V}_{ij} \times \Delta$.

- $w_{ij}(t_{ij})$, for all $i \in [n], j \in [m], t_{ij} \in \mathcal{T}_{ij}$, denoting the expected marginal reduced form. We denote $w_i = \{w_{ij}(t_{ij})\}_{j\in[m],t_{ij}\in\mathcal{T}_{ij}}$ the vector of all variables associated with bidder i.

- d_i, for all $i \in [n]$, denoting an upper bound of the expectation of δ_{ij} over distribution C_{ij} for all j.

[a] For every i, j, let $\mathcal{V}^0_{ij} = \mathcal{T}_{ij}$ be the set of all possible values of t_{ij}. To address the tie-breaking issue in Remark 1, let $\varepsilon_r > 0$ be an arbitrarily small number,and define $\mathcal{V}^+_{ij} = \{t_{ij} + \varepsilon_r : t_{ij} \in \mathcal{T}_{ij}\}$ and $\mathcal{V}_{ij} = \mathcal{V}^0_{ij} \cup \mathcal{V}^+_{ij}$. Let Δ be a geometric discretization of range $[\widetilde{\text{PRev}}/n, 55 \cdot \widetilde{\text{PRev}}]$. Formally, $\delta \in \Delta$ if and only if $\delta = \frac{2^x}{n} \cdot \widetilde{\text{PRev}}$ for some integer x such that $0 \le x \le \lceil \log(55n) \rceil$. Finally, for each $\beta \in \mathcal{V}^0_{ij}$, let $\beta^+ = \beta + \varepsilon_r \in \mathcal{V}^+_{ij}$. Note that the LP (or the LP in Figure 5) do not depend on the choice of ε_r, so we can choose ε_r to be sufficiently small. In fact, let b be an upper bound of the bit complexity of the problem instance, and the bit complexity of any feasible solution of our LP. Our proof works as long as $\varepsilon_r < \min\{\frac{1}{2^{\text{poly}(b)}}, \frac{\text{PRev}}{\sum_{i,j} |\mathcal{T}_{ij}|}\}$.

Figure 4: Interpretation of the variables of the LP in Figure 3.

Mechanism 1 Two-part Tariff Mechanism $\mathcal{M}_{\text{TPT}}$

0: Before the mechanism starts, the seller computes the *price* Q_j (Definition 6) for every item j.
1: Bidders arrive sequentially in the lexicographical order.
2: When bidder i arrive, the seller shows her the set of available items $S_i(t_{<i}) \subseteq [m]$, as well as their prices. Note that $S_i(t_{<i})$ is the set of items that are not purchased by the first $i-1$ bidders, which depends on $t_{<i}$. We use $S_1(t_{<1})$ to denote $[m]$.
3: Bidder i is asked to pay an *entry fee*. The seller samples a type $t'_i \sim D_i$, and sets the entry fee as: $\xi_i(S_i(t_{<i}), t'_i) = \max_{S' \subseteq S_i(t_{<i})} \left(v_i(t'_i, S') - \sum_{j\in S'} Q_j\right)$. The entry fee is bidder i's utility for her favorite set under prices Q_j's if her type is t'_i.
4: If bidder i (with type t_i) agrees to pay the entry fee $\xi_i(S_i(t_{<i}), t'_i)$, then she can enter the mechanism and take her favorite set $S^* \in \arg\max_{S' \subseteq S_i(t_{<i})} \left(v_i(t_i, S') - \sum_{j\in S'} Q_j\right)$, by paying $\sum_{j\in S^*} Q_j$. Otherwise, the bidder gets nothing and pays 0.

Indeed, for both Theorem 3 and Lemma 3, we prove a general statement that applies to XOS valuations, which requires a generalized LP and definitions. See the arXiv version for details.

Lemma 3. *For any BIC and IR mechanism $\mathcal{M}$, $\text{Rev}(\mathcal{M}) \le 28 \cdot \text{PRev} + 4 \cdot \text{OPT}_{LP}$.*

3.3 Interpretation of Our LP in Figure 3.

We explain our LP in this section. The objective is the expected Core, as explained in Step 3 in Section 3.1. According to our definition of $\lambda_{ij}(t_{ij}, \beta_{ij}, \delta_{ij})$ and Constraint (3), $\{w_{ij}(t_{ij})\}_{i,j,t_{ij}}$ corresponds to the expected marginal reduced form, that is, $w_{ij}(t_{ij})$ is the expected probability for bidder i to receive item j and her value for item j is t_{ij}. Constraints (1) and (2) simply sets the feasible region of the expected marginal reduced form w. They follow directly from the fact that every realized $\widehat{w}^{(\beta,\delta)}$ is feasible (see Step 3 in Section 3.1). Constraint (4) follows from Equation (3) and the fact that every $\widehat{w}^{(\beta,\delta)}_{ij}(t_{ij})/f_{ij}(t_{ij})$ is in $[0, 1]$. Constraint (5) implies that $\{\hat{\lambda}_{ij}(\beta_{ij}, \delta_{ij})\}_{\beta_{ij},\delta_{ij}}$ correspond to a distribution C_{ij}.

Constraints (6) - (9) are specialized for our problem, which guarantees that the LP optimum can be bounded by simple mechanisms. Constraint (6) follows from taking expectation on both sides of Constraint (1) in Figure 1, over the randomness of β_{ij}. Constraints (8) and (9) correspond to Constraint (2) in Figure 1. Here we bound the expectation of δ_{ij} by a unified upper bound d_i for every j.[16] It is worth emphasizing Constraint (7), which corresponds to Constraint (4) in Figure 1 (and Figure 2). Instead of taking expectations over the dual parameters, we force the constraint to hold for every β_{ij} and δ_{ij}. This is an important property that is crucial in our analysis. Readers may notice that Constraint (2) is implied by Constraints (3), (6) and (7). We keep Constraint (2) so that it is clear that the supply constraint is enforced over the (expected) marginal reduced form. The Problem Specific Constraints (6)-(9) in the LP in Figure 3 are expected versions of the constraints in the two-stage optimization problem in Figure 1, which are directly inspired by the Properties 1, 2, and 3 in Lemma 1. They are crucial to guarantee that optimal value of the LP in Figure 3 is still approximable by simple mechanisms.

4 SAMPLE ACCESS TO THE DISTRIBUTION

In this section, we focus on the case where we have only access to the bidders' distribution. Our goal is again to compute an approximately optimal mechanism. Our plan is as follows: (i) for each $i \in [n]$ and $j \in [m]$, take $O(\log(1/\delta)/\varepsilon^2)$ samples from D_{ij}, and let $\widehat{D}$ be the uniform distribution over the samples. By the DKW inequality [28], $\widehat{D}_{ij}$ and D_{ij} have Kolmogorov distance (Definition 7) no more than ε with probability at least $1-\delta$. (ii) We then apply our algorithm in Theorem 1 to compute an RPP or TPT that is approximately optimal w.r.t. distributions $\{\widehat{D}_{ij}\}_{i\in[n],j\in[m]}$. We show that the computed simple mechanism is approximately optimal for the true distributions as well.

Definition 7. *The Kolmogorov distance between two distributions $\mathcal{D}$ and $\widehat{\mathcal{D}}$ supported on $\mathbb{R}$ is defined as*

$$d_K(\mathcal{D}, \widehat{\mathcal{D}}) = \sup_{z\in\mathbb{R}} \left| \Pr_{t\sim\mathcal{D}}[t \le z] - \Pr_{\hat{t}\sim\widehat{\mathcal{D}}}[\hat{t} \le z] \right|$$

Theorem 4. *Suppose all bidders' valuations are constrained additive. If for each $i \in [n]$ and $j \in [m]$, D_{ij} is supported on numbers in $[0, 1]$ with bit complexity no more than b, then for any $\varepsilon > 0$ and $\delta > 0$, with probability $1-\delta$, we can compute in time*

[16] This corresponds to the fact that in Lemma 1 (and Figure 1), there is a single c_i that represents δ_{ij}.

$poly(n, m, 1/\varepsilon, \log(1/\delta), b)$ a rationed posted price mechanism or a two-part tariff mechanism, whose revenue is at least $c \cdot OPT - O(nm^2\varepsilon)$ for some absolute constant c. The algorithm takes $O\left(\frac{\log(nm/\delta)}{\varepsilon^2}\right)$ samples from each D_{ij} and assumes query access to a demand oracle for each bidder.

A ADDITIONAL PRELIMINARIES

Definition 8. *[34] Let $\mathcal{D}_i$ be the type distribution of bidder i and denote by $\mathcal{V}_i$ her distribution over valuations $v_i(t_i, \cdot)$ where $t_i \sim \mathcal{D}_i$. We say that $\mathcal{V}_i$ is subadditiver over independent items if*

- *$v_i(\cdot,\cdot)$ has no externalities, that is for any $S \subseteq [m]$, $t_i, t_i' \in \mathcal{T}_i$ such that $t_{ij} = t'_{ij}$ for $j \in S$, then $v_i(t_i, S) = v_i(t_i', S)$.*
- *$v_i(\cdot,\cdot)$ is monotone, that is for each $t_i \in \mathcal{T}_i$ and $S \subseteq T \subseteq [m]$, $v_i(t_i, S) \leq v_i(t_i, T)$*
- *$v_i(\cdot,\cdot)$ is subadditive function, that is for all $t_i \in \mathcal{T}_i$ and $S_1, S_2 \subseteq [m]$, $v_i(t_i, S_1 \cup S_2) \leq v_i(t_i, S_1) + v_i(t_i, S_2)$*

Mechanisms: A mechanism M in multi-item auctions can be described as a tuple (x, p). For every type profile t, buyer i and bundle $S \subseteq [m]$, $x_{iS}(t)$ is the probability of buyer i receiving the exact bundle S at profile t, $p_i(t)$ is the payment for buyer i at the same type profile. To ease notations, for every buyer i and types t_i, we use $p_i(t_i) = \mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})]$ as the interim price paid by buyer i and $\sigma_{iS}(t_i) = \mathbb{E}_{t_{-i}}[x_{iS}(t_i, t_{-i})]$ as the interim probability of receiving the exact bundle S.

IC and IR constraints: A mechanism $M = (x, p)$ is BIC if:

$$\sum_{S\subseteq[m]} \sigma_{iS}(t_i)\cdot v_i(t_i, S) - p_i(t_i) \geq \sum_{S\subseteq[m]} \sigma_{iS}(t_i') \cdot v_i(t_i, S) - p_i(t_i'), \forall i, t_i, t_i' \in \mathcal{T}_i.$$

The mechanism is DSIC if $\forall i, t_i, t_i' \in \mathcal{T}_i, t_{-i} \in \mathcal{T}_{-i}$:

$$\sum_{S\subseteq[m]} x_{iS}(t_i, t_{-i}) \cdot v_i(t_i, S) - p_i(t_i, t_{-i}) \geq \sum_{S\subseteq[m]} x_{iS}(t_i', t_{-i}) \cdot v_i(t_i, S) - p_i(t_i', t_{-i}).$$

The mechanism is (interim) IR if:

$$\sum_{S\subseteq[m]} \sigma_{iS}(t_i) \cdot v_i(t_i, S) - p_i(t_i) \geq 0, \forall i, t_i \in \mathcal{T}_i.$$

The mechanism is ex-post IR if:

$$\sum_{S\subseteq[m]} x_{iS}(t_i, t_{-i}) \cdot v_i(t_i, S) - p_i(t_i, t_{-i}) \geq 0, \forall i, t_i \in \mathcal{T}_i, t_{-i} \in \mathcal{T}_{-i}.$$

Definition 9 (Separation Oracle for Convex Polytope $\mathcal{P}$). *A Separation Oracle SO for a convex polytope $\mathcal{P} \subseteq \mathbb{R}^d$, takes as input a point $\boldsymbol{x} \in \mathbb{R}^d$ and if $\boldsymbol{x} \in \mathcal{P}$, then the oracle says that the point is in the polytope. If $\boldsymbol{x} \notin \mathcal{P}$, then the oracle output a separating hyperplane, that is it outputs a vector $\boldsymbol{y} \in \mathbb{R}^d$ and $c \in \mathbb{R}$ such that $\boldsymbol{y}^T x \leq c$, but for $\boldsymbol{z} \in \mathcal{P}$, $\boldsymbol{y}^T \boldsymbol{z} > c$.*

Definition 10 (Polytopes and Facet-Complexiy). *We say P has* facet-complexity *at most b if it can be written as $P := \{\vec{x} \mid \vec{x} \cdot \vec{w}^{(i)} \leq c_i,\ \forall i \in \mathcal{I}\}$, where each $\vec{w}^{(i)}$ and c_i has bit complexity at most b for all $i \in \mathcal{I}$. We use the term* convex polytope *to refer to a set of points that is closed, convex, bounded,[17] and has finite facet-complexity.*

Definition 11 (Vertex-Complexity). *We use the term* corner *to refer to non-degenerate extreme points of a convex polytope. In other words, $\vec{y}$ is a corner of the d-dimensional convex polytope P if $\vec{y} \in P$ and there exist d linearly independent directions $\vec{w}^{(1)}, \ldots, \vec{w}^{(d)}$ such that $\vec{x} \cdot \vec{w}^{(i)} \leq \vec{y} \cdot \vec{w}^{(i)}$ for all $\vec{x} \in P, 1 \leq i \leq d$. We use **Corner**$(P)$ to denote the set of corners of a convex polytope P. We say P has vertex-complexity at most b if all vectors in Corner(P) have bit complexity no more than b.*

B SOME EXAMPLES

B.1 Non-Concavity of Core

In this section, we show that the $\textsc{Core}(\sigma, \theta(\sigma))$ function is non-concave in the interim allocation rule σ. We first provide the formal definition of $\theta(\sigma)$ for a single-bidder two-item instance, and we use $\textsc{Core}^{CZ}(\sigma)$ to denote $\textsc{Core}(\sigma, \theta(\sigma))$.

Definition 12 (Core for a single additive bidder over two items with continuous distributions - [17]). *Consider a single bidder interested in two items, whose value is sampled from continuous distribution D with support $T = \textsc{Supp}(D)$ and density function $f(t)$ for $t \in T$. Consider a feasible interim allocation $\sigma = \{\sigma_1(t), \sigma_2(t)\}_{t\in \textsc{Supp}(D)}$, that is $\sigma_1(t)$ ($\sigma_2(t)$ resp.) is the probability that the allocation rule awards item 1 (item 2 resp.) to a bidder with type t. Define*

$$\beta_1(\sigma) = \arg\min_{\beta\geq 0} \left[\Pr_{t_1\sim D_1}\left[t_1 \geq \beta\right] = \mathbb{E}_{t\sim D}\left[\sigma_1(t)\right]\right],$$

$$\beta_2(\sigma) = \arg\min_{\beta\geq 0} \left[\Pr_{t_2\sim D_2}\left[t_2 \geq \beta\right] = \mathbb{E}_{t\sim D}\left[\sigma_2(t)\right]\right]$$

and

$$c(\sigma) = \arg\min_{a\geq 0}\left\{\Pr_{t\sim D}\left[t_1 \leq \beta_1(\sigma) + a\right] + \Pr_{t\sim D}\left[t_2 \leq \beta_2(\sigma) + a\right] \geq \frac{1}{2}\right\}$$

The term $\textsc{Core}^{CZ}$ for interim allocation σ is defined as follows:

$$\textsc{Core}^{CZ}(\sigma) = \mathbb{E}_{t\sim D}\left[\sigma_1(t) t_1 \cdot \mathbb{1}\left[t_1 \leq \beta_1(\sigma) + c(\sigma)\right]\right] + \mathbb{E}_{t\sim D}\left[\sigma_2(t) t_2 \cdot \mathbb{1}\left[t_2 \leq \beta_2(\sigma) + c(\sigma)\right]\right]$$

In Example 1 we show that $\textsc{Core}^{CZ}(\sigma)$ is a non-concave function even in the setting with a single bidder and two items. The reason for the $\textsc{Core}^{CZ}$ being non-concave lies in the fact that the interval which we truncate depends on the interim allocation σ. Computing the concave hull of $\textsc{Core}^{CZ}(\sigma)$ in the worst case requires exponential time in the dimension of the space, which is m is our case.

Example 1. *Consider a single additive bidder interested in two items whose values are both drawn from the uniform distribution $U[0, 1]$. Consider two interim allocation rules σ and σ':*

- *σ: Award the first item to the buyer if her value for it lies in the interval $[0, 1/2]$ and never award the second item to the buyer.*

[17] $P \subseteq \mathbb{R}^d$ is bounded if it is contained in $[-x, x]^d$ for some $x \in \mathbb{R}$.

- σ': *Always award the first item to the buyer and never award the second item to the buyer.*

According to Definition 13, for allocation rule σ, the dual parameters are $\beta_1(\sigma) = 1/2$, $\beta_2(\sigma) = 1$ and $c(\sigma) = 0$, which implies $\text{Core}^{CZ}(\sigma) = 1/8$. Similarly for allocation rule σ' we have $\beta_1(\sigma') = 0$, $\beta_2(\sigma') = 1$ and $c(\sigma') = 0$, which implies that $\text{Core}^{CZ}(\sigma') = 0$.

Consider the interim allocation σ'' that uses allocation rule σ with probability 50% and σ' with 50%. Note that σ'' is in the convex combination of σ and σ' and more specifically $\sigma'' = \frac{\sigma+\sigma'}{2}$. For interim allocation σ'' we have that $\beta_1(\sigma'') = 1/4$, $\beta_2(\sigma'') = 1$ and $c(\sigma'') = 0$, which implies that $\text{Core}^{CZ}(\sigma'') = \frac{1}{32}$. We notice that the second item contributes nothing to the Core^{CZ}, but it ensures that $c = 0$ regardless of the allocation of the first item. Thus $\text{Core}^{CZ}(\sigma'') < \frac{1}{2}(\text{Core}^{CZ}(\sigma) + \text{Core}^{CZ}(\sigma'))$, which implies that $\text{Core}^{CZ}(\cdot)$ is not a concave function.

B.2 Why can't we use the Ex-Ante Relaxation?

An influential framework known as the ex-ante relaxation has been widely used in Mechanism Design, but is insufficient for our problem. Informally speaking, ex-ante relaxation reduces a multi-bidder objective to the sum of single-bidder objectives subject to some global supply constraints over ex-ante allocation probabilities. To solve the ex-ante relaxation program efficiently, the single-bidder objective has to be concave and efficiently computable given the ex-ante probabilities [1].

In revenue maximization, the single-bidder objective – the optimal revenue subject to ex-ante probabilities – is indeed a concave function. However, we do not have a polynomial time algorithm to even compute the single-bidder objective given a set of fixed ex-ante probabilities. [18] To fix this issue, one can try to find a concave function that is always a good approximation to the single-bidder objective for any ex-ante probabilities. To the best of our knowledge, such a concave function only exists for unit-demand bidders via the copies setting technique [20]. Alternatively, one can replace the global objective – optimal revenue by the upper bound of revenue proposed in [17]. Yet the corresponding single-bidder objective for one term Core in the upper bound is highly non-concave, which makes the ex-ante relaxation not applicable.

Although the term Core was originally defined for interim allocation rules (as in Definition 12), it can also be defined for ex-ante probabilities. We only define it for the single-bidder two-item case. Let $q = \{q_1, q_2\} \in [0,1]^2$, and $\text{Max-Core} = \max_{\sigma\in\Sigma(q)} \text{Core}^{CZ}(\sigma)$, where $\Sigma(q)$ is the set of feasible interim allocations that awards the first item with probability at most q_1 and the second item with probability at most q_2. Example 2 also shows that $\text{Max-Core}(\cdot)$ is a non-concave function by observing that $\sigma \in \Sigma(1/2, 0)$, $\sigma' \in \Sigma(1, 0)$ and $\sigma'' \in \Sigma(3/4, 0)$.

Definition 13 (Core for a single additive bidder over two items - [17]). *Consider a single bidder interested in two items, whose value is sampled from $D_1 \times D_2$. Consider a supply constraints $q_1, q_2 \in [0,1]$. Note that q_1 (or q_2) is the probability that a mechanism awards the first item (or the second item) to the bidder. Define*

$$\beta_1 = \arg\min_{\beta\geq 0}\left[\Pr_{t_1\sim D_1}\left[t_1 \geq \beta\right] = q_1\right],$$

$$\beta_2 = \arg\min_{\beta\geq 0}\left[\Pr_{t_2\sim D_2}\left[t_2 \geq \beta\right] = q_2\right]$$

and

$$c = \arg\min_{a\geq 0}\left\{\Pr_{t_1\sim D_1}\left[t_1 \leq \beta_1 + a\right] + \Pr_{t_2\sim D_2}\left[t_2 \leq \beta_2 + a\right] \geq \frac{1}{2}\right\}$$

The term Max-Core is defined as follows:

$$\begin{aligned}\text{Max-Core}(q) = &\max_{\substack{x_1:\mathcal{T}_1\to[0,1]\\ \sum_{t_1\in\mathcal{T}_1} f_1(t_1)x_1(t_1)=q_1}} \sum_{\substack{t_1\in\mathcal{T}_1\\ t_1\leq\beta_1+c}} f_1(t_1)\cdot t_1\cdot x_1(t_1)\\ &+ \max_{\substack{x_2:\mathcal{T}_2\to[0,1]\\ \sum_{t_2\in\mathcal{T}_2} f_2(t_2)x_2(t_2)=q_2}} \sum_{\substack{t_2\in\mathcal{T}_2\\ t_2\leq\beta_2+c}} f_2(t_2)\cdot t_2\cdot x_2(t_2)\end{aligned}$$

In Example 2 we show that $\text{Max-Core}(q)$ is a non-concave function even in the setting with a single bidder and two items. The reason for the Max-Core being non-concave lies in the fact that the interval which we truncate depends on the supply constraints q. Computing the concave hull of $\text{Max-Core}(q)$ in the worst case requires exponential time in the dimension of the space, which is m is our case. These facts make the ex-ante relaxation approach not applicable to solve our problem.

Example 2. *Consider a single additive bidder interested in two items whose values are both drawn from the uniform distribution $U[0,1]$. Consider the values $q = (1/2, 0)$ and $q' = (1, 0)$. According to Definition 13, for q we have that $\beta_1^{(q)} = 1/2$, $\beta_2^{(q)} = 1$ and $c^{(q)} = 0$ and for q' we have $\beta_1^{(q')} = 0$, $\beta_2^{(q')} = 1$ and $c^{(q')} = 0$. We notice that the second item contributes nothing to the Max-Core, but it ensures that $c = 0$ regardless of the supply demand for the first item. Observe that $\text{Max-Core}(q) = 1/8$ and $\text{Max-Core}(q') = 0$. Let $q'' = (q+q')/2 = (3/4, 0)$. For q'', observe that $\beta_1^{(q'')} = 1/4$, $\beta_2^{(q'')} = 1$ and $c^{(q'')} = 0$. We have $\text{Max-Core}(q'') = 1/32$. Thus $\text{Max-Core}(q'') < \frac{1}{2}(\text{Max-Core}(q) + \text{Max-Core}(q'))$, which implies that $\text{Max-Core}(\cdot)$ is not a concave function.*

C XOS VALUATIONS

We show a generalization of Theorem 3 that works for XOS buyers (Theorem 5), with the generalized of the single-bidder marginal reduced form polytope Definition 14 and a generalized LP (Figure 5).

Theorem 5. *Let $(w, \lambda, \hat{\lambda}, \boldsymbol{d})$ (or $(\pi, w, \lambda, \hat{\lambda}, \boldsymbol{d})$) be any feasible solution of the LP in Figure 3 (or Figure 5). Let $\mathcal{M}_{PP}$ be the rationed posted price mechanism computed in Lemma 2.*

Let

$$Q_j = \frac{1}{2}\cdot\sum_{i\in[n]}\sum_{t_{ij}\in\mathcal{T}_{ij}} f_{ij}(t_{ij})\cdot V_{ij}(t_{ij})\cdot \sum_{\substack{\beta_{ij}\in\mathcal{V}_{ij}\\ \delta_{ij}\in\Delta}} \lambda_{ij}(t_{ij}, \beta_{ij}, \delta_{ij}) \cdot \mathbb{1}[V_{ij}(t_{ij}) \leq \beta_{ij} + \delta_{ij}],$$

and $\mathcal{M}_{TPT}$ be the two-part tariff mechanism shown in Mechanism 1 with prices $\{Q_j\}_{j\in[m]}$. Then the objective function of the solution

[18] The closest thing we know is a QPTAS for a unit-demand bidder. See Section 1.2.

$2 \cdot \sum_{j\in[m]} Q_j$ is bounded by $c_1 \cdot \textsc{Rev}(\mathcal{M}_{PP}) + c_2 \cdot \textsc{Rev}(\mathcal{M}_{TPT})$, for some constant $c_1, c_2 > 0$. Moreover, both $\mathcal{M}_{PP}$ and $\mathcal{M}_{TPT}$ can be computed in time $poly(n, m, \sum_{i,j} |\mathcal{T}_{ij}|)$, with access to the demand oracle for the buyers' valuations.

C.1 Single-Bidder Marginal Reduced Form Polytope for XOS Valuations

In Definition 14 we define the single-bidder marginal reduced form polytope W_i for XOS buyers, which differs from the single-bidder marginal reduced form polytope for constrained-additive valuation is several ways. In Definition 14, we define a distribution σ_S^k over all possible subset of items $S \subseteq [m]$ and over the finite number $k \in [K]$ over additive functions that can be chosen when we evaluate the value that the buyer has for a set of items. In Definition 5, the distribution σ_S was only over sets in the set of feasible allocations.

Similar to Definition 5, $\pi_{ij}(t_{ij})$ is equal to $f_{ij}(t_{ij})$ times the probability that the i-th buyer receives the j-th item. In contrast to Definition 5, in Definition 14, the value of $w_{ij}(t_{ij})$ is $\frac{f_{ij}(t_{ij})}{V_{ij}(t_{ij})}$ times the expected value that the buyer has for the item when we are allowed to choose which additive functions in $k \in [K]$ we count the value of the buyer, or we are even allowed to allocate an item to the buyer but count zero value for it (that is equivalent to just throwing away the item).

Definition 14 (XOS valuations: single-bidder marginal reduced form polytope). *For every $i \in [n]$, the single-bidder marginal reduced form polytope $W_i \subseteq [0,1]^{2\cdot\sum_j |\mathcal{T}_{ij}|}$ is defined as follows. Let $\pi_i = (\pi_{ij}(t_{ij}))_{j,t_{ij}\in\mathcal{T}_{ij}}$ and $w_i = (w_{ij}(t_{ij}))_{j,t_{ij}\in\mathcal{T}_{ij}}$. Then $(\pi_i, w_i) \in W_i$ if and only if there exist a number $\sigma_S^{(k)}(t_i) \in [0,1]$ for every $t_i \in \mathcal{T}_i, S \subseteq [m], k \in [K]$, such that*

(1) $\sum_{S,k} \sigma_S^{(k)}(t_i) \le 1, \forall t_i \in \mathcal{T}_i$.
(2) $\pi_{ij}(t_{ij}) = f_{ij}(t_{ij}) \cdot \sum_{t_{i,-j}} f_{i,-j}(t_{i,-j}) \cdot \sum_{S:j\in S}\sum_{k\in[K]} \sigma_S^{(k)}(t_{ij}, t_{i,-j})$, for all $i, j, t_{ij} \in \mathcal{T}_{ij}$.
(3) $w_{ij}(t_{ij}) \le f_{ij}(t_{ij}) \cdot \sum_{t_{i,-j}} f_{i,-j}(t_{i,-j}) \cdot \sum_{S:j\in S}\sum_{k\in[K]} \sigma_S^{(k)}(t_{ij}, t_{i,-j}) \cdot \frac{\alpha_{ij}^{(k)}(t_{ij})}{V_{ij}(t_{ij})}$, for all $i, j, t_{ij} \in \mathcal{T}_{ij}$.

C.2 The Linear Program for XOS valuations

The LP for XOS valuations can be found in Figure 5. Here $V_{ij}^0 = \{V_{ij}(t_{ij}) : t_{ij} \in \mathcal{T}_{ij}\}$, $\mathcal{V}_{ij}^+ = \{V_{ij}(t_{ij}) + \varepsilon_r : t_{ij} \in \mathcal{T}_{ij}\}$ and $\mathcal{V}_{ij} = \mathcal{V}_{ij}^0 \cup \mathcal{V}_{ij}^+$. We notice that this is consistent with our LP for constrained-additive buyers (Figure 3), as $V_{ij}(t_{ij}) = t_{ij}$ for constrained-additive buyers.

Denote OPT_{LP} the optimum objective of the LP in Figure 5. Similar to the constrained-additive case, we have the following lemma.

Lemma 4. *When buyers have XOS valuations, for any BIC and IR mechanism $\mathcal{M}$, $\textsc{Rev}(\mathcal{M}) \le 28 \cdot \textsc{PRev} + 4 \cdot OPT_{LP}$.*

$$\max \sum_{i\in[n]}\sum_{j\in[m]}\sum_{t_{ij}\in\mathcal{T}_{ij}} f_{ij}(t_{ij}) \cdot V_{ij}(t_{ij}) \cdot \sum_{\substack{\beta_{ij}\in\mathcal{V}_{ij}\\ \delta_{ij}\in\Delta}} \lambda_{ij}(t_{ij},\beta_{ij},\delta_{ij}) \cdot \mathbb{1}[V_{ij}(t_{ij}) \le \beta_{ij} + \delta_{ij}]$$

s.t.

Allocation Feasibility Constraints:

(1) $(\pi_i, w_i) \in W_i, \forall i$

(2) $\sum_i \sum_{t_{ij}\in\mathcal{T}_{ij}} \pi_{ij}(t_{ij}) \le 1, \forall j$

Natural Feasibility Constraints:

(3) $f_{ij}(t_{ij}) \cdot \sum_{\beta_{ij}\in\mathcal{V}_{ij}}\sum_{\delta_{ij}\in\Delta} \lambda_{ij}(t_{ij},\beta_{ij},\delta_{ij}) = w_{ij}(t_{ij}, \forall i, j, t_{ij} \in \mathcal{T}_{ij}$

(4) $\lambda_{ij}(t_{ij},\beta_{ij},\delta_{ij}) \le \hat\lambda_{ij}(\beta_{ij},\delta_{ij}), \forall i, j, t_{ij}, \beta_{ij} \in \mathcal{V}_{ij}, \delta_{ij}$

(5) $\sum_{\substack{\beta_{ij}\in\mathcal{V}_{ij}\\ \delta_{ij}\in\Delta}} \hat\lambda_{ij}(\beta_{ij},\delta_{ij}) = 1, \forall i, j$

Problem Specific Constraints:

(6) $\sum_{i\in[n]}\sum_{\beta_{ij}\in\mathcal{V}_{ij}}\sum_{\delta_{ij}\in\Delta} \hat\lambda_{ij}(\beta_{ij},\delta_{ij}) \cdot \Pr_{t_{ij}\sim D_{ij}}[V_{ij}(t_{ij}) \ge \beta_{ij}] \le \frac{1}{2}, \forall j$

(7) $\frac{1}{2}\sum_{t_{ij}\in\mathcal{T}_{ij}} f_{ij}(t_{ij}) \left(\lambda_{ij}(t_{ij},\beta_{ij},\delta_{ij}) + \lambda_{ij}(t_{ij},\beta_{ij}^+,\delta_{ij})\right) \le \hat\lambda_{ij}(\beta_{ij},\delta_{ij}) \cdot \Pr_{t_{ij}}[V_{ij}(t_{ij}) \ge \beta_{ij}] + \hat\lambda_{ij}(\beta_{ij}^+,\delta_{ij}) \cdot \Pr_{t_{ij}}[V_{ij}(t_{ij}) \ge \beta_{ij}^+]$,
$\forall i, j, \beta_{ij} \in \mathcal{V}_{ij}^0, \delta_{ij} \in \Delta$

(8) $\sum_{\substack{\beta_{ij}\in\mathcal{V}_{ij}\\ \delta_{ij}\in\Delta}} \delta_{ij} \cdot \hat\lambda_{ij}(\beta_{ij},\delta_{ij}) \le d_i, \forall i, j$

(9) $\sum_{i\in[n]} d_i \le 111 \cdot \widetilde{\textsc{PRev}}$

$\lambda_{ij}(t_{ij},\beta_{ij},\delta_{ij}) \ge 0, \hat\lambda_{ij}(\beta_{ij},\delta_{ij}) \ge 0, \pi_{ij}(t_{ij}) \ge 0,$
$w_{ij}(t_{ij}) \ge 0, d_i \ge 0, \forall i, j, t_{ij}, \beta_{ij} \in \mathcal{V}_{ij}, \delta_{ij}$

Figure 5: LP for XOS Valuations

D COUNTEREXAMPLE FOR ADJUSTABLE DEMAND ORACLE

For XOS valuations, our algorithm for constructing the simple mechanism requires access to a special adjustable demand oracle $\textsc{ADem}_i(\cdot,\cdot,\cdot)$. Readers may wonder if this enhanced oracle (rather than a demand oracle) is necessary to prove our result. In this section we show that (even an approximation of) $\textsc{ADem}_i$ can not be implemented using polynomial number of queries from the value oracle, demand oracle and a classic XOS oracle. All the oracles are defined as follows. Throughout this section, we only consider a single buyer and thus drop the subscript i. Recall that the XOS valuation $v(\cdot)$ satisfies that $v(S) = \max_{k\in[K]}\left\{\sum_{j\in S}\alpha_j^{(k)}\right\}$ for every set S, where $\{\alpha_j^{(k)}\}_{j\in[m]}$ is the k-th additive function.

- Demand Oracle ($\textsc{Dem}$): takes a price vector $p \in \mathbb{R}^m$ as input, and outputs
  $S^* \in \arg\max_{S\subseteq[m]}\left(v(S) - \sum_{j\in S} p_j\right)$.

- XOS Oracle (Xos): takes a set $S \subseteq [m]$ as input, and outputs the k^*-th additive function $\{\alpha_j^{(k^*)}\}_{j\in[m]}$, where $k^* \in \arg\max_{k\in[K]} \left\{\sum_{j\in S} \alpha_j^{(k)}\right\}$.
- Value Oracle: takes a set $S \subseteq [m]$ as input, and outputs $v(S)$. We notice that a value oracle can be easily simulated with an XOS oracle. Thus we focus on XOS oracles for the rest of this section.
- Adjustable Demand Oracle (ADem): takes a coefficient vector $b \in \mathbb{R}^m$ and a price vector $p \in \mathbb{R}^m$ as inputs, and outputs $(S^*, \{\alpha_j^{(k^*)}\}_{j\in[m]})$ where $(S^*, k^*) \in \arg\max_{S\subseteq[m],k\in[K]} \left\{\sum_{j\in S} b_j\alpha_j^{(k)} - \sum_{j\in S} p_j\right\}$.

An (approximate) implementation of ADem is an algorithm that takes inputs $b, p \in \mathbb{R}^m$, and outputs a set $S \subseteq [m]$ and $k \in [K]$. The algorithm has access to the demand oracle and XOS oracle of v. We denote $\text{ALG}(v, b, p)$ the output of the algorithm. For any $\alpha > 1$, ALG is an α-approximation to ADem if for every XOS valuation v and every $b, p \in \mathbb{R}^m$, the algorithm outputs (S', k') that satisfies:

$$\max_{S\subseteq[m],k\in[K]} \left\{\sum_{j\in S} b_j\alpha_j^{(k)} - \sum_{j\in S} p_j\right\} \leq \alpha \cdot \left(\sum_{j\in S} b_j\alpha_j^{(k')} - \sum_{j\in S'} p_j\right)$$

In Theorem 6 we show that we cannot approximate the output of an Adjustable Demand Oracle within any finite factor, if we are permitted to query polynomial many times the XOS, Value and Demand Oracle.

Theorem 6. *Given any $\alpha > 1$, there* ***does not exist*** *an implementation of ADem (denoted as ALG) that satisfies both of the following properties:*

(1) For any XOS valuation v over m items, ALG makes $poly(m)$ queries to the value oracle, the demand oracle and XOS oracle of v, and runs in time $poly(m, b)$. Here b is the bit complexity of the problem instance (See Definition 3).
(2) ALG is an α-approximation to ADem.

REFERENCES

[1] Saeed Alaei. 2011. Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers. In *the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
[2] Saeed Alaei, Hu Fu, Nima Haghpanah, Jason Hartline, and Azarakhsh Malekian. 2012. Bayesian Optimal Auctions via Multi- to Single-agent Reduction. In *the 13th ACM Conference on Electronic Commerce (EC)*.
[3] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. 2014. A Simple and Approximately Optimal Mechanism for an Additive Buyer. In *the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
[4] Johannes Brustle, Yang Cai, and Constantinos Daskalakis. 2020. Multi-Item Mechanisms without Item-Independence: Learnability via Robustness *(EC '20)*. Association for Computing Machinery, New York, NY, USA, 715–761. https://doi.org/10.1145/3391403.3399541
[5] Yang Cai and Constantinos Daskalakis. 2011. Extreme-Value Theorems for Optimal Multidimensional Pricing. In *the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
[6] Yang Cai and Constantinos Daskalakis. 2017. Learning Multi-item Auctions with (or without) Samples. In *the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
[7] Yang Cai and Constantinos Daskalakis. 2021. Recommender Systems meet Mechanism Design. *arXiv preprint arXiv:2110.12558* (2021).
[8] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 2012. An Algorithmic Characterization of Multi-Dimensional Mechanisms. In *the 44th Annual ACM Symposium on Theory of Computing (STOC)*.
[9] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 2012. Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization. In *the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
[10] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 2012. Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization. *arxiv report* (2012). http://arxiv.org/abs/1207.5518.
[11] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 2013. Reducing Revenue to Welfare Maximization : Approximation Algorithms and other Generalizations. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
[12] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 2013. Understanding Incentives: Mechanism Design becomes Algorithm Design. In *the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
[13] Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. 2016. A Duality Based Unified Approach to Bayesian Mechanism Design. In *the 48th Annual ACM Symposium on Theory of Computing (STOC)*.
[14] Yang Cai and Zhiyi Huang. 2013. Simple and Nearly Optimal Multi-Item Auctions. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
[15] Yang Cai and Argyris Oikonomou. 2021. On Simple Mechanisms for Dependent Items. In *Proceedings of the 22nd ACM Conference on Economics and Computation* (Budapest, Hungary) *(EC '21)*. Association for Computing Machinery, New York, NY, USA, 242–262. https://doi.org/10.1145/3465456.3467643
[16] Yang Cai, Argyris Oikonomou, Grigoris Velegkas, and Mingfei Zhao. 2021. An Efficient ε-BIC to BIC Transformation and Its Application to Black-Box Reduction in Revenue Maximization. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1337–1356.
[17] Yang Cai and Mingfei Zhao. 2017. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*. 170–183. https://doi.org/10.1145/3055399.3055465
[18] Shuchi Chawla, Jason Hartline, David Malec, and Balasubramanian Sivan. 2010. Sequential Posted Pricing and Multi-parameter Mechanism Design. *CoRR* abs/0907.2435 (2010). arXiv:0907.2435 [cs.GT]
[19] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. 2007. Algorithmic Pricing via Virtual Valuations. In *the 8th ACM Conference on Electronic Commerce (EC)*.
[20] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. 2010. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*.
[21] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. 2015. The power of randomness in Bayesian optimal mechanism design. *Games and Economic Behavior* 91 (2015), 297–317. https://doi.org/10.1016/j.geb.2012.08.010
[22] Shuchi Chawla and J. Benjamin Miller. 2016. Mechanism Design for Subadditive Agents via an Ex-Ante Relaxation, In Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016. *CoRR*, 579–596. https://doi.org/10.1145/2940716.2940756
[23] Xi Chen, Ilias Diakonikolas, Anthi Orfanou, Dimitris Paparas, Xiaorui Sun, and Mihalis Yannakakis. 2015. On the complexity of optimal lottery pricing and randomized mechanisms. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*. IEEE, 1464–1479.
[24] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2014. The Complexity of Optimal Mechanism Design. In *the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
[25] Constantinos Daskalakis, Nishanth Dikkala, and Gautam Kamath. 2018. Testing ising models. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 1989–2007.
[26] Constantinos Daskalakis and S. Matthew Weinberg. 2012. Symmetries and Optimal Multi-Dimensional Mechanism Design. In *the 13th ACM Conference on Electronic Commerce (EC)*.
[27] Paul Dütting, Thomas Kesselheim, and Brendan Lucier. 2020. An O(log log m) Prophet Inequality for Subadditive Combinatorial Auctions. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*. IEEE, 306–317. https://doi.org/10.1109/FOCS46700.2020.00037
[28] Aryeh Dvoretzky, Jack Kiefer, and Jacob Wolfowitz. 1956. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics* (1956), 642–669.
[29] Sergiu Hart and Noam Nisan. 2012. Approximate Revenue Maximization with Multiple Items. In *the 13th ACM Conference on Electronic Commerce (EC)*.
[30] Robert Kleinberg and S. Matthew Weinberg. 2012. Matroid Prophet Inequalities. In *the 44th Annual ACM Symposium on Theory of Computing (STOC)*.
[31] Pravesh Kothari, Sahil Singla, Divyarthi Mohan, Ariel Schvartzman, and S Matthew Weinberg. 2019. Approximation schemes for a unit-demand buyer with independent items via symmetries. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 220–232.

[32] Roger B. Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
[33] Aviad Rubinstein. 2016. On the computational complexity of optimal simple mechanisms. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*. 21–28.
[34] Aviad Rubinstein and S. Matthew Weinberg. 2015. Simple Mechanisms for a Subadditive Buyer and Applications to Revenue Monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*. 377–394. https://doi.org/10.1145/2764468.2764510
[35] Andrew Chi-Chih Yao. 2015. An n-to-1 Bidder Reduction for Multi-item Auctions and its Applications, In SODA. *CoRR* (2015). http://arxiv.org/abs/1406.3278