

Simple and Nearly Optimal Multi-Item Auctions

Yang Cai*
EECS, MIT
ycai@csail.mit.edu

Zhiyi Huang†
University of Pennsylvania
hzhiyi@cis.upenn.edu

Abstract

We provide a Polynomial Time Approximation Scheme (PTAS) for the Bayesian optimal multi-item multi-bidder auction problem under two conditions. First, bidders are independent, have additive valuations and are from the same population. Second, every bidder’s value distributions of items are independent but not necessarily identical monotone hazard rate (MHR) distributions. For non-i.i.d. bidders, we also provide a PTAS when the number of bidders is small. Prior to our work, even for a single bidder, only constant factor approximations are known.

Another appealing feature of our mechanism is the simple allocation rule. Indeed, the mechanism we use is either the second-price auction with reserve price on every item individually, or VCG allocation with a few outlying items that requires additional treatments. It is surprising that such simple allocation rules suffice to obtain nearly optimal revenue.

1 Introduction

The multi-dimensional mechanism design problem has been widely studied in Economics, and recently in the theory of computation community. Consider a seller who has a limited supply of several distinguishable items and many interested bidders. The goal for the seller is to design an auction that will incentivize the bidders to truthfully report their private valuations and maximize her revenue. Unfortunately, optimal mechanism is not even well-defined in the worst-case analysis, as no truthful mechanism can be universally optimal for all possible valuation profiles. Economists have taken the Bayesian approach to cope with this impossibility, where the valuations of the bidders are assumed to be drawn from some publicly known distributions. Given such prior distributions, the optimal mechanism is de-

finer as the one that maximizes the expected revenue among all (possibly randomized) truthful and individual rational mechanisms. In this paper, the notion of truthfulness we will focus on is Bayesian incentive compatibility (BIC), while we will also consider other notions of truthfulness such as incentive compatibility and deterministic truthfulness. Informally, a mechanism is BIC if each bidder maximizes her expected utility by truth-telling assuming other bidders are also truthful, where the expectation is over the randomness of the mechanisms and random realizations of other agents’ valuations.

When there is only a single item for sale, the structure of the optimal mechanism is very well-understood. Myerson [18] provides an elegant solution to the optimal single-item auction problem. However, Myerson’s result does not extend to the more general *multi-item* setting. Following Myerson’s work, a large body of research in Economics has been devoted to extending his result to the multi-item setting (see survey [17] and the references therein).

The theory of computation community has also studied this problem during the past decade, with an eye on the computational efficiency of the mechanism. There has been lots of success in obtaining constant factor approximations in various settings (e.g., [11, 12, 6, 1]). Lately, attention has been shifted to getting nearly optimal revenue and such mechanisms have been proposed for several cases (e.g., [13, 8, 2]).

In a very recent paper [9], Cai et al. consider a very general setting. In their setting, bidders are additive with arbitrary combinatorial feasibility constraints. They show how to design revenue-optimal auctions by reducing the revenue optimization to welfare optimization under the same constraints. Their algorithm has runtime polynomial in the total number of bidder types¹. This is the natural description size for the problem if we allow items to have correlated values. However, when items are independent, the natural de-

*Supported by NSF Award CCF-0953960 (CAREER) and CCF-1101491. Part of this work was done while the author was visiting Microsoft Research, Redmond.

†This work was supported in part by an ONR MURI Grant N000140710907. Part of this work was done while the author was visiting Microsoft Research, Redmond.

¹More precisely, these algorithm is polynomial in $\sum_i |S_i|$, where S_i is the support of the joint value distribution for bidder i .

description is much more succinct. Making their algorithm inefficient (exponential in the input size). Moreover, to handle such a broad setting, their solution has to be relatively complicated, which might sometimes makes it hard to implement in reality. The above drawbacks motivate the research in this paper, that is, *designing simple, computational efficient, and nearly optimal auctions*.

1.1 Main Results In this paper, we will focus on a very important and fundamental case: Bidders have independent and additive valuations, and items values are independent. Our goal is to obtain an algorithm whose runtime is polynomial in the succinct input-size and propose much simpler revenue-optimal auctions.

More concretely, let there be m bidders² and n heterogeneous items (unit-supply). Let there be no feasibility constraints on the allocations. We will assume the bidders' valuations are additive and the values are drawn from independent but not necessarily identical distributions subject to the standard monotone hazard rate (MHR) assumption. Roughly speaking, MHR distributions are those whose tails are "thinner" than exponential distributions. The formal definition of MHR is deferred to Section 2. We want to efficiently find a mechanism whose expected revenue is optimal relative to any (possibly randomized) truthful and individual rational mechanism. Prior to our work, even the case of a single bidder is elusive in the presence of many independent but not necessarily identical items.

Our main results are the following two theorems.

THEOREM 1.1. *Let there be n heterogeneous items, m additive bidders, and $\{\mathcal{F}_j\}_{j \in [n]}$ be a collection of independent but not necessarily identical MHR distributions. Suppose for each bidder i , her value for item j is drawn independently from \mathcal{F}_j . Then, there is a Polynomial Time Approximation Scheme³ (PTAS) for computing the revenue-optimal truthful mechanism.*

In the above theorem, we consider the case when the bidders are from the same population. So any two bidders have the same value distributions for any particular item. This is a realistic assumption, as to tell which demographic group the bidder is from, the seller needs to collect lots of information, e.g., her occupation,

²We will sometimes use k to denote the number of bidders when this number is an absolute constant.

³Recall that a *Polynomial Time Approximation Scheme* (PTAS) is a family of algorithms $\{\mathcal{A}_\epsilon\}_\epsilon$, indexed by a parameter $\epsilon > 0$, such that for every fixed $\epsilon > 0$, \mathcal{A}_ϵ runs in poly-time. In particular, for any constant $\epsilon > 0$, the PTAS constructs an auction whose expected revenue is a $(1 + \epsilon)$ factor approximation to the optimal, in time polynomial in n and m .

income, marital status etc. which is usually infeasible in practice, especially when the number of bidders is huge. When there are only a handful of bidders, however, the seller might have enough knowledge to distinguish different bidders. We develop the following theorem to address this case.

THEOREM 1.2. *Let there be n heterogeneous items, k additive bidders (consider k as an absolute constant), and $\{\mathcal{F}_{ij}\}_{i \in [k], j \in [n]}$ be a collection of independent but not necessarily identical MHR Distributions. For any bidder i and item j , her value for the item is drawn from \mathcal{F}_{ij} . There is a PTAS for computing the revenue-optimal BIC mechanism.*

Although the above theorems are stated only for BIC mechanism here, our techniques can be extended to other solution concepts as well, such as IC and deterministic truthfulness. We will elaborate these theorems in the corresponding sections and explain the results for various solution concepts.

Besides achieving nearly optimal revenue, our mechanisms in Theorem 1.1 and Theorem 1.2 have an additional appealing feature of using very simple allocation rules. In fact, all of our mechanisms essentially has one of the two following simple forms: 1) Run a second price auction with reserve price on every item individually. 2) Use the VCG allocation with a threshold welfare whose role is similar to the reserve price, except for a few outlying items which we need to handle separately. It is surprising that such simple allocation rules can actually obtain nearly optimal revenue.

1.2 Overview of Techniques First let us explain by example why the obvious attempt of running Myerson's auction on every item individually fails. Consider a single bidder and n items whose values are i.i.d. and uniformly drawn from $[0, 1]$. On the one hand, Myerson's optimal auction only gets $\frac{1}{4}$ revenue per item. On the other hand, if n is large, the total value of the grand-bundle concentrated at $\frac{n}{2}$. So a simple grand-bundle-reserve-price auction (e.g., [3]) can get almost $\frac{n}{2}$ revenue.

One might also argue that when the bidders' values are additive, the overall values will be concentrated and thus it is easy to find the optimal. But as items are non-i.i.d., we may not have such a concentration phenomena in some cases.⁴

Instead, our first technical contribution is by understanding the probabilistic structure to prove the following structural lemma which we will use heavily:

⁴For instance, consider an item whose value is uniformly drawn from $[1/2, 1]$ and $n - 1$ items whose values are i.i.d. and uniformly drawn from $[0, \frac{1}{n^2}]$.

PARTITIONING LEMMA (INFORMAL). *Assuming MHR distributions, then we can partition the items into two sets, where the first set contains only a constant number of items, and the second set has many items but the social welfare of which highly concentrates.*

Based on this lemma, we manage to reduce the problem of finding nearly-optimal mechanisms for many independent items into two simpler sub-problems: Designing nearly-optimal mechanisms for a constant number of independent items, and designing nearly-optimal mechanisms when the total value of the items concentrates. The formal statement of the partition lemma and its proof will be given in Section 3.

Constant Number of Bidders In this case, designing nearly-optimal mechanisms for the sub-problem with only a constant number of items is almost folklore and we sketch these mechanisms in Section 3.1. In order to handle the second sub-problem, we propose a novel mechanism that falls into the VCG family, which we shall introduce as the *reserve welfare mechanism* in Section 3.2. The reserve welfare mechanism allocates items to the bidders only if the social welfare exceeds a certain reserve welfare, in which case it will use the welfare-maximizing allocation. We show that with the proper pricing scheme, the reserve welfare mechanism is deterministically truthful and solves the welfare-concentrated case nearly optimally. The proof of Theorem 1.2 follows by combining these technical ingredients.

Many I.I.D. Bidders The key observation in this case is that when the number of bidders is sufficiently large, simply running second price auction with a properly chosen reserve price for each item suffices to guarantee nearly optimal revenue. More concretely, inspired by Theorem 7 in [7], we can argue that for any constant $\epsilon > 0$, if the number of bidders is larger than an absolute constant that only depends on ϵ , then for every item there is a second price auction with reserve price that achieves revenue at least a $(1 - \epsilon)$ fraction of the social welfare. In [5], Bhargat and Khanna have independently provided similar insights when there are sufficiently many i.i.d. bidders. On the other hand, if the number of bidders is smaller than this absolute constant, then we can reduce the problem to Theorem 1.2.

1.3 Related Work The theory of computation community has contributed many computational efficient solutions to various special cases of the multi-dimensional mechanism design problem. Chawla et.al [11] consider the case of a single unit-demand bidder, and propose an item pricing mechanism that achieves a constant factor approximation of the optimal. Their result is based on an elegant reduction to Myer-

son’s optimal auction in the single-dimensional setting. For the same problem, Cai and Daskalakis [7] propose a PTAS for optimal item-pricing, thus close the constant approximation gap.

In the multi-bidder setting, [12, 6, 1] provide efficient constant factor approximations for cases when the bidders are additive or unit-demand. More recently, near-optimal solutions have been obtained for several cases. Daskalakis and Weinberg [13] solve the case where there are few bidders with symmetric⁵ items or symmetric bidders with few items.

For asymmetric distributions, Cai et al. [8] give the optimal solution to the many-bidder and many-item setting. Alaei et al. [2] consider serving many copies of an item with a matroid feasibility constraint on which bidders can be served an item simultaneously, and obtain the optimal solution. In [9], Cai et al. provide the optimal solution for a much more general setting where bidders are additive, and with (possibly) arbitrary feasibility constraints, by reducing the revenue optimization to welfare optimization. Their reduction provides a poly-time solution to the optimal mechanism design problem in all auction settings where welfare optimization can be solved efficiently. However, it is fragile to approximation, as the reduction requires an exact solution for the welfare optimization problem. In [10], the same group of authors show that even when the welfare optimization problem is only approximately solvable, they can still carry over the reduction while preserving the approximation factor. All of these algorithms allow correlation among items, so the total number of bidder types is the natural input size. However, for independent but not necessarily identical items, even when support size of every value distribution is only 2 and bidders are i.i.d., the total number of bidder types could still be as large as 2^n , making their algorithm highly inefficient in our setting. Nonetheless, for symmetric items, Cai et al. [8] show how to reduce the “effective number” of types by utilizing the symmetric structure of the items, yielding mechanisms that are polynomial in both n and m . But designing nearly optimal auctions for asymmetric items remains open prior to our work even for a single bidder.

Our result can be viewed as an improvement of [13] and a complement to [8]. Although the results are related, the techniques are orthogonal. The approaches in [13] and [8] are LP-based, and they use symmetry to reduce the size of the LP. We take a different path. By understanding the probabilistic structure, we argue that the social welfare of most of the items are highly concentrated, and can be easily extracted by the seller

⁵See [13] for a formal definition of symmetric distributions. E.g., i.i.d. distributions is symmetric, but general independent distributions are not.

using a modified VCG mechanism. Further, for the other constant number of items, as there are only a small number of possible types, they can be easily handled by previous results (e.g., [13]).

2 Preliminaries

2.1 Model Formally, in an *multi-item auction*, a seller has n heterogeneous items that she wants to auction to m quasi-linear risk-neutral bidders. Each bidder i has a private valuation profile $v_i = (v_{i1}, \dots, v_{in}) \in \mathbb{R}^n$, where v_{ij} is bidder i 's value for item j . v_i is sometimes referred to as the type of the bidder i . We will assume the valuation function to be additive, that is, $v_i(S) = \sum_{j \in S} v_{ij}$ for any $S \subseteq [n]$. We will let v_{-i} denote the type profile of every bidder except i .

A *mechanism* M consists of two parts: An allocation rule $x(\cdot)$ and a payment rule $p(\cdot)$.

The *allocation rule* $x(\cdot)$ maps a type profile \mathbf{v} to a feasible allocation $x(\mathbf{v}) = \{x(\mathbf{v})_{ij}\}_{i \in [m], j \in [n]}$, where $x(\mathbf{v})_{ij}$ is the probability for bidder i to receive item j when the type profile is \mathbf{v} . For deterministic mechanisms, we will let $x(\mathbf{v})_{ij}$ to be either 0 or 1. We will let $x(\mathbf{v})_i$ denote the n dimensional vector $(x(\mathbf{v})_{i1}, \dots, x(\mathbf{v})_{in})$.

The *payment rule* maps a type profile \mathbf{v} to a m -dimensional real vector $p(\mathbf{v}) = (p_1(\mathbf{v}), \dots, p_m(\mathbf{v}))$, where $p_i(\mathbf{v})$ is the price charged to bidder i .

Since the valuations are private information of the bidders, the mechanism needs to retrieve these information from the bidders, who may or may not manipulate the information. We will let $\mathbf{b} = (b_1, \dots, b_m)$ denote the *bids* of the bidders. Given the bids, the allocation $x(\mathbf{b})$, and the payments $p(\mathbf{b})$, we will assume the bidders are utility maximizers w.r.t. the standard notion of *quasi-linear utility*:

$$u_i(v_i, x(\mathbf{b}), p(\mathbf{b})) = v_i \cdot x_i(\mathbf{b}) - p_i(\mathbf{b}) .$$

We will consider the *Bayesian setting*. Namely, we will assume that the valuations v_{ij} , $i \in [m]$ and $j \in [n]$, are drawn from some publicly known independent (but not necessarily identical) distributions \mathcal{F}_{ij} . When bidders are from the same population, we will use \mathcal{F}_j to denote the value distribution for item j , and omit subscript i . We will let $F_{ij}(x)$ and $f_{ij}(x)$ denote the cumulative distribution function and probability density function of \mathcal{F}_{ij} respectively.

Next, we formally define how the distributions \mathcal{F}_{ij} are specified to the mechanism. We shall consider two different models. The first one is the *discrete explicit access* model, where the support of each \mathcal{F}_{ij} is discrete and explicitly given, and so is the probability of each value in the support being chosen. The second one is the *continuous oracle access* model, where the support

of \mathcal{F}_{ij} could be continuous and even unbounded. In the latter, we will assume there is an oracle sampler such that each access to the oracle returns a random value drawn from \mathcal{F}_{ij} . In the former case, the running time of our mechanisms shall be polynomial in the sum of the support sizes of \mathcal{F}_{ij} for all $i \in [m]$ and $j \in [n]$. In the latter case, both the running time and the number of accesses to the oracle of our mechanisms shall be polynomial in m and n .

2.2 Solution Concepts We will consider the following standard game-theoretic solution concepts:

DEFINITION 1. A *deterministic mechanism* M is *deterministically truthful (DT)*, if *truth-telling is a utility-maximizing strategy*, i.e.,

$$\forall b_{-i} : v_i \in \arg \max_{b_i} \{v_i \cdot x_i(b_i, b_{-i}) - p(b_i, b_{-i})\} .$$

DEFINITION 2. A *randomized mechanism* M is *truthful-in-expectation or incentive compatible (IC)* if *truth-telling maximizes the expected utility*, i.e.,

$$\forall b_{-i} : v_i \in \arg \max_{b_i} \{\mathbf{E}[v_i \cdot x_i(b_i, b_{-i}) - p(b_i, b_{-i})]\} .$$

where *expectation is over random coin-flips of the mechanism*.

DEFINITION 3. A (*randomized*) *mechanism* M is *Bayesian-incentive-compatible (BIC)* if *truth-telling maximizes the expected utility*, i.e.,

$$\forall b_{-i} : v_i \in \arg \max_{b_i} \{ \mathbf{E}_{b_{-i} \sim \mathcal{F}_{-i}} [v_i \cdot x_i(b_i, b_{-i}) - p(b_i, b_{-i})] \} ,$$

where *expectation is over random coin-flips of the mechanism and random realization of the valuations of other bidders*.

We will also consider the following relaxed notions of deterministic truthfulness.

DEFINITION 4. A *deterministic mechanism* M is *ϵ -deterministically truthful (ϵ -DT)*, if

$$\begin{aligned} \forall b_i, b_{-i} : & v_i \cdot x_i(b_i, b_{-i}) - p(b_i, b_{-i}) \\ & \leq v_i \cdot x_i(v_i, b_{-i}) - p(v_i, b_{-i}) + \epsilon . \end{aligned}$$

The notions of ϵ -IC and ϵ -BIC are defined similarly.

Further, it is very important not to overcharge the bidders, especially when we are aiming for revenue.

DEFINITION 5. A *mechanism* M is *individually rational (IR)* if *the utility of any bidder in any outcome is always non-negative*, i.e.,

$$\forall b_{-i} : v_i \cdot x_i(v_i, b_{-i}) - p(v_i, b_{-i}) \geq 0 .$$

The following *taxation principle* is a well known characterization for truthful mechanisms (e.g., see [14, 16]) which will be useful for our discussion.

THEOREM 2.1. (TAXATION PRINCIPLE) *A mechanism is DT/IC if and only if each bidder i is presented a menu of bundles/lotteries of items such that the prices of the bundles/lotteries only depends on the other bidders' valuations v_{-i} , and bidder i always gets one of the utility maximizing bundles/lotteries.*

In particular, if there is only one bidder, then such menus are fixed regardless of the reported value. So any DT/IC mechanism can be viewed as a bundle-pricing/lottery-pricing of the items.

2.3 Extreme Value Theorem Throughout this paper, we will consider distributions that have *monotone hazard rate* (MHR):

DEFINITION 6. *A distribution \mathcal{F} has monotone hazard rate if $\frac{f(x)}{1-F(x)}$ is non-decreasing in the support of \mathcal{F} .*

The MHR distributions is a commonly studied family of distributions in Economics and recently in the algorithmic game theory community. It includes familiar distributions such as the Normal, Exponential, and Uniform distributions. Intuitively, a distribution has monotone hazard rate if its tail is at most as large as that of an Exponential distribution. We note that in our results the MHR assumption can be replaced by the following assumption: There exists a constant C such that the for each bidder i and item j the support of \mathcal{F}_{ij} is an interval whose upper and lower bounds differs by at most a C multiplicative factor. In other words, our algorithms work well as long as we have a rough idea on each bidder i 's value on each item j . In fact, this is the alternative assumption we will use for the discrete explicit access model.

We will use the following *extreme value theorem* for MHR distributions developed in [7] as an important technical tool in our proofs. Readers are referred to [7] for the proof of the theorem.

THEOREM 2.2. (EXTREME VALUE THEOREM [7]) *Suppose X_1, \dots, X_n are a collection of independent (but not necessarily identically distributed) random variables whose distributions are MHR, and $f_{\max_i\{X_i\}}$ is the probability density function of the random variable $\max_i X_i$. Then, for all $\epsilon \in (0, \frac{1}{4})$, there exists some anchoring point β such that $\Pr[\max_i X_i \geq \frac{\beta}{2}] \geq 1 - \frac{1}{\sqrt{\epsilon}}$ and*

$$\int_{2\beta \log(\frac{1}{\epsilon})}^{+\infty} t \cdot f_{\max_i\{X_i\}}(t) dt \leq 36\beta\epsilon \log\left(\frac{1}{\epsilon}\right).$$

Moreover, β is efficiently computable from the distributions of the X_i 's.

Based on the above extreme value theorem and an additional probabilistic argument, we will show in Section 3 that the social welfare of some carefully chosen subset of items highly concentrates. In particular, we will consider the followings notion of concentration.

DEFINITION 7. *A random variable X is (ϵ, δ) -concentrated if $X \in (1 - \epsilon, 1 + \epsilon) \mathbf{E}[X]$ with probability at least $1 - \delta$.*

3 Nearly Optimal Mechanism for Constant Number of Bidders

In this section, we will consider the case when there are only k bidders, where k is an absolute constant.⁶ We will prove that for various solution concepts, the problem of finding revenue-optimal truthful mechanisms can be solved under a unified framework. Formally, our results can be summarized as the following theorem.

THEOREM 3.1. (THM. 1.2 ELABORATED) *Suppose the number of bidders is a constant. Then, there is a PTAS (polynomial in n) for finding revenue-optimal mechanisms among (all settings require IR):*

- IC/BIC mechanisms with discrete explicit access.
- DT mechanisms with discrete explicit access.¹
- DT/IC mechanisms with continuous oracle access for a single bidder.
- BIC mechanisms with continuous oracle access.
- DT/IC mechanisms with continuous oracle access.¹

¹Our mechanisms in these cases are only ϵ -deterministically truthful and ϵ -IC.

The general proof strategy of Theorem 3.1 is to reduce the problem of designing almost optimal mechanisms for the multi-item auction problem into two easier sub-problems (assuming MHR distributions). More precisely, we will prove that if there are PTAS for the special cases in the next two lemmas, then it is possible to combine the nearly optimal mechanisms for these two cases to derive the PTAS in Theorem 3.1.

⁶We note that our mechanisms can be extended to the case of $O(\log^c n)$ bidders for sufficiently small constant c via almost identical proofs, where c depends on the solution concept. However, we feel such extension is not very insightful. So we will only present the case for a constant number of bidders in this extended abstract for the sake of presentation.

LEMMA 3.1. (FEW-ITEM CASE) *Theorem 3.1 holds if both the number of items and the number of bidders are constants.*

LEMMA 3.2. (CONCENTRATED CASE) *Suppose the optimal social welfare is (ϵ, δ) -concentrated, and the number of bidders is a constant. Then, there is a polynomial-time and deterministically truthful mechanism whose expected revenue is at least $(1 - f(\epsilon, \delta))$ fraction of the expected optimal social welfare, where $f(\epsilon, \delta)$ goes to zero as ϵ and δ goes to zero.*

At this point, we will focus on how to combine the mechanisms obtained from the above lemmas to derive the proof of Theorem 3.1. The proofs of Lemma 3.1 and Lemma 3.2 are deferred to Section 3.1 and Section 3.2 respectively.

Proof Outline of Theorem 3.1 Before getting into the technical details, let us first sketch the roadmap of our proof. First of all, we notice that under the MHR assumption, it is easy to achieve expected revenue that is at least a constant fraction of the expected social welfare. This follows easily from previous work (e.g. see [4]) and we will formally state it as Lemma 3.3. By this result, we know that we can throw away items whose contribution to the expected social welfare is tiny without overhurting the optimal revenue. Next, we proceed by proving a structural result saying that we can partition the items into three groups: a small group of items with large variance, which we shall handle with the mechanism from Lemma 3.1 (Section 3.1);⁷ a group of items whose contribution to the social welfare concentrates, which we will handle with the mechanism from Lemma 3.2 (Section 3.2); and finally a group of items whose total contributions to the expected social welfare is tiny, which we will simply ignore (never allocate them to any bidder). This result is formally stated and proved in Lemma 3.4. At last, in order to show this approach is a PTAS for the multi-item auction problem, we need to show that the optimal revenue of the problem is upper bounded by the optimal

⁷The philosophy of selecting a few distinguished items to reduce the size of the problem and solve it nearly optimally may look similar to that of the k -lookahead auction (e.g. see [19, 15]), where we choose a few distinguished bidders and design nearly optimal mechanism for them (based on the bids of the other bidders). However, there are a few crucial differences. First, in our approach the small set of items are chosen without knowing the bids while in the k -lookahead auction the set of bidders are chosen based on the bids. Further, in our approach we also derive good revenue from the rest of the items while the k -lookahead auction never derive revenue directly from the rest of the bidders. Finally, as a result of the previous point, our approach admits nearly optimal revenue while the k -lookahead auction only guarantees constant-factor approximation so far.

revenue when only a subset of items present (the items with large variance) plus the expected social welfare of the remaining items (the concentrated group of items). Indeed, we will prove this claim as Lemma 3.5. The complete proof of Theorem 3.1 is given in Appendix A.

Next, let us formally state and prove the technical lemmas mentioned in the proof sketch. The following lemma is folklore from previous work.

LEMMA 3.3. (E.G., COROLLARY 3.7 OF [6]) *For any multi-item auction with MHR bidders, the optimal expected revenue is at least a constant fraction of the expected social welfare.*

Now let us consider the partition lemma. For presentation purpose, we will only show a weaker version of the partition lemma, under the additional assumption that the upper and lower bounds of the value range of $X_j = \max_i v_{ij}$ only differ by at most a constant factor c for every item j . Note that Lemma 2.2 implies that at least $1 - O(\epsilon \log \frac{1}{\epsilon})$ fraction of the contribution to $\mathbf{E}[X_j]$ comes from a range whose upper and lower bounds differ by at most an $O(\frac{1}{\epsilon} \log(\frac{1}{\epsilon}))$ factor. It is easy to see that by choosing $c = \Theta(\frac{1}{\epsilon} \log(\frac{1}{\epsilon}))$ and taking into account the fact that the contribution outside the range is tiny, we can prove the partition lemma without the additional assumption. We omit the details here.

LEMMA 3.4. *Suppose X_1, \dots, X_n are n non-negative independent random variables, where $[\alpha_j, \beta_j]$ is the range of X_j , $1 \leq j \leq n$, such that $c = \max_j \frac{\beta_j}{\alpha_j}$ is a constant. Suppose $\epsilon > 0$ and $\frac{1}{8} > \delta > 0$ are small constants. Then, we can partition X_1, \dots, X_n into three groups R , S , and T in polynomial time, such that:*

1. *The size of R is small: $|R| \leq \frac{16c^2}{\epsilon^3} \ln(\frac{2}{\delta})$.*
2. *The sum in S , $\sum_{X_j \in S} X_j$, is (ϵ, δ) -concentrated.*
3. *The contribution from group T is tiny: $\sum_{X_j \in T} \mathbf{E}[X_j] \leq \epsilon \sum_{j=1}^n \mathbf{E}[X_j]$.*

Proof. Let $s = \sum_{j=1}^n \mathbf{E}[X_j]$ be the sum of the expectation of these random variables. Note that s can be estimated up to a constant factor in polynomial time and such estimated value is sufficient for our purpose. For the sake of presentation, we will assume that we know the value of s .

We will first partition the random variables into $\Theta(\log n)$ buckets $B_1, B_2, \dots, B_{\log n + \log(\frac{2}{\epsilon})}$ according to their expectations. If $\mathbf{E}[X_j] \in [\frac{s}{2^\ell}, \frac{s}{2^{\ell-1}}]$ for $1 \leq \ell \leq \log n + \log(\frac{2}{\epsilon})$, then we put X_j into bucket B_ℓ . If $\mathbf{E}[X_j] \leq \frac{\epsilon s}{2n}$, then its contribution to the social welfare is negligible and we will put X_j into T .

Briefly speaking, we will proceed as follows. First pick a small threshold index ℓ^* (the value of ℓ^* will be defined later); then for each bucket B_ℓ such that $\ell \leq \ell^*$, we put all random variables in B_ℓ into R ; for each bucket B_ℓ such that $\ell > \ell^*$, we will show that either $\sum_{X_j \in B_\ell} X_j$ concentrates with high probability, or the contribution of $\sum_{X_j \in B_\ell} \mathbf{E}[X_j]$ is tiny. In the former case, we will put the variables in B_ℓ into S ; in the latter case, we will put the variables in B_ℓ into T .

More precisely, for each bucket B_ℓ where $\ell > \ell^*$, if $|B_\ell| \geq \frac{2c^2}{\epsilon^2} (\ln(\frac{2}{\delta}) + \ell - \ell^*)$, we will put these random variables into S . Note that for every $X_j \in B_\ell$, we have that $\beta_j \leq c\alpha_j \leq c\mathbf{E}[X_j] \leq \frac{c \cdot s}{2^{\ell-1}}$ and $\mathbf{E}[X_j] \geq \frac{s}{2^\ell}$. By Chernoff-Hoeffding bound, we get that

$$\begin{aligned} \Pr \left[\left| \sum_{X_j \in B_\ell} X_j - \sum_{X_j \in B_\ell} \mathbf{E}[X_j] \right| > \epsilon \sum_{X_j \in B_\ell} \mathbf{E}[X_j] \right] \\ \leq 2 \exp \left(-\ln \left(\frac{2}{\delta} \right) - \ell + \ell^* \right) = \delta \exp(-\ell + \ell^*) . \end{aligned}$$

Now consider all the buckets that we put into S . By union bound, the probability that the sum of the random variables in any of these buckets does not concentrate is at most $\sum_{\ell > \ell^*} \delta \exp(-\ell + \ell^*) < \delta$. Thus, we have proved that S satisfies the desired property in the lemma.

If $|B_\ell| < \frac{2c^2}{\epsilon^2} (\ln(\frac{2}{\delta}) + \ell - \ell^*)$, we shall put all variables in B_ℓ into T . In this case, we have $\sum_{X_j \in B_\ell} \mathbf{E}[X_j] \leq \frac{s}{2^{\ell-1}} \frac{2c^2}{\epsilon^2} (\ln(\frac{2}{\delta}) + \ell - \ell^*)$. Therefore, the sum of the expected values of the random variables in T is at most

$$\begin{aligned} & \sum_{\ell > \ell^*} \frac{s}{2^{\ell-1}} \frac{2c^2}{\epsilon^2} \left(\ln \left(\frac{2}{\delta} \right) + \ell - \ell^* \right) + \sum_{X_j: \mathbf{E}[X_j] \leq \frac{\epsilon s}{2n}} \frac{\epsilon s}{2n} \\ & \leq \sum_{\ell > \ell^*} \frac{s}{2^{\ell-1}} \frac{2c^2}{\epsilon^2} \ln \left(\frac{2}{\delta} \right) + \sum_{\ell > \ell^*} \frac{s}{2^{\ell-1}} \frac{2c^2}{\epsilon^2} (\ell - \ell^*) + \frac{\epsilon s}{2} \\ & = \frac{s}{2^{\ell^*-1}} \frac{2c^2}{\epsilon^2} \ln \left(\frac{2}{\delta} \right) + \frac{s}{2^{\ell^*-2}} \frac{2c^2}{\epsilon^2} + \frac{\epsilon s}{2} . \end{aligned}$$

In order to guarantee that $\sum_{X_j \in T} \mathbf{E}[X_j] \leq \epsilon s$, it suffices to choose ℓ^* such that

$$\frac{1}{2^{\ell^*-1}} \frac{2c^2}{\epsilon^2} \ln \left(\frac{2}{\delta} \right) + \frac{1}{2^{\ell^*-2}} \frac{2c^2}{\epsilon^2} \leq \frac{\epsilon}{2} .$$

Note that $\delta < \frac{1}{2\epsilon^2}$ implies that

$$\frac{1}{2^{\ell^*-1}} \frac{2c^2}{\epsilon^2} \ln \left(\frac{2}{\delta} \right) + \frac{1}{2^{\ell^*-2}} \frac{2c^2}{\epsilon^2} \leq \frac{1}{2^{\ell^*-2}} \frac{2c^2}{\epsilon^2} \ln \left(\frac{2}{\delta} \right) .$$

We shall let $\ell^* = \log \left(\frac{16c^2}{\epsilon^3} \ln \left(\frac{2}{\delta} \right) \right)$ and conclude that T satisfies the claimed property.

Finally, we note that for any $X_j \in R$, we have $\mathbf{E}[X_j] \geq \frac{s}{2^{\ell^*}}$. So by $\sum_{X_j \in R} \mathbf{E}[X_j] \leq s$ we get that the size of R is at most $2^{\ell^*} = \frac{16c^2}{\epsilon^3} \ln \left(\frac{2}{\delta} \right)$.

At last, we will show that by decomposing the problem into two sub-problems we do not hurt the optimal revenue by too much. Concretely, for any $S \subseteq [n]$, we let $\text{opt}^{\text{DT}}(S)$, $\text{opt}^{\text{IC}}(S)$, and $\text{opt}^{\text{BIC}}(S)$ denote the optimal revenue by deterministically truthful/IC/BIC mechanisms respectively when only the items in S are available on the market (value distributions are the same). We have

LEMMA 3.5. *For any $S \subseteq [n]$, we have*

$$\text{opt}^{\text{truthful}}([n]) \leq \text{opt}^{\text{truthful}}(S) + \sum_{j \notin S} \mathbf{E}[\max_i v_{ij}] ,$$

where *truthful* can be instantiated with *deterministically truthful (DT)*, or *IC*, or *BIC*.

Proof. Suppose M is the truthful (under the instantiated solution concept) mechanism that achieves optimal revenue. Let us construct a truthful mechanism M_S for the market when only the items in S is presented. The revenue of the mechanism shall be at least $\text{opt}^{\text{truthful}}([n]) - \sum_{j \notin S} \mathbf{E}[\max_i v_{ij}]$:

1. Let bidders submit their bids $\mathbf{b}_{1,S}, \dots, \mathbf{b}_{k,S}$.
2. Sample values $\mathbf{v}_{1,-S} \sim \mathcal{F}_{1,-S}, \dots, \mathbf{v}_{k,-S} \sim \mathcal{F}_{k,-S}$ for items not in S .
3. Run M on bids $(\mathbf{b}_{1,S}, \mathbf{v}_{1,-S}), \dots, (\mathbf{b}_{k,S}, \mathbf{v}_{k,-S})$. Let \mathbf{S} and \mathbf{p} denote the allocation and prices.
4. Give bidder i the items in $S_i \cap S$ and charge her $p_i - \sum_{j \in S_i \setminus S} v_{ij}$.

First, let us analyze the revenue achieved by M_S assuming the bidders bid truthfully: $\mathbf{b}_{i,S} = \mathbf{v}_{i,S}$ for $1 \leq i \leq k$. The revenue by M_S is $\sum_{i=1}^k \mathbf{E} \left[p_i - \sum_{j \in S_i \setminus S} v_{ij} \right]$. By linearity of expectation, this can be divided into two parts: $\sum_{i=1}^k \mathbf{E}[p_i] - \sum_{i=1}^k \mathbf{E}[\sum_{j \in S_i \setminus S} v_{ij}]$. The first part $\sum_{i=1}^k \mathbf{E}[p_i]$ is the expected revenue $\text{opt}^{\text{truthful}}([n])$ achieved by M . The second part $\sum_{i=1}^k \mathbf{E} \left[\sum_{j \in S_i \setminus S} v_{ij} \right]$ is social welfare from items outside S , which is upper bounded by the optimal social welfare. Note that the latter part is upper bounded by the $\sum_{j \notin S} \mathbf{E}[\max_i v_{ij}]$. Therefore, the revenue by M_S is at least

$$\text{opt}^{\text{truthful}}([n]) - \sum_{j \notin S} \mathbf{E}[\max_i v_{ij}] .$$

Now let us explain why mechanism M_S is indeed truthful with respect to the corresponding solution concepts. Note that bidder i 's utility is

$$\sum_{j \in S_i \cap S} v_{ij} - (p_i - \sum_{j \in S_i \setminus S} v_{ij}) = \sum_{j \in S_i} v_{ij} - p_i ,$$

which is exactly the utility of a virtual bidder whose values are v_i and bids $(b_{i,S}, v_{i,-S})$. So if M is IC/BIC, then M_S is also IC/BIC. Finally, if M is a deterministically truthful mechanism then M_S is uniformly truthful. We further note that there is no performance gap between optimal uniformly truthful mechanisms and optimal deterministically truthful mechanisms in the Bayesian setting. So Lemma 3.5 follows.

3.1 Nearly Optimal Mechanism for Constant Number of Items and bidders The mechanisms for constant number of items and constant number of bidders mostly follow directly from previous work. The general approach is to brute-force search with the hope that the search space would be small since both the number of items and the number of bidders are small. However, the strategy spaces for mechanism design problems are typically infinite. Hence, appropriate discretization is needed in order to reduce the size of the search space. We will briefly describe these mechanisms and thus prove Lemma 3.1 in the Appendix B for self-containness.

3.2 Nearly Optimal Mechanism When Social Welfare Concentrates In this section, we will prove Lemma 3.2 by demonstrating how to design nearly optimal mechanisms, when the social welfare concentrates near its expectation as the bidders' values are drawn from the corresponding distributions.

3.2.1 Single-Bidder Case As a warm-up, let us first consider the single-bidder case. This case is quite straight-forward. We note that a *grand-bundle-reserve-price auction* (e.g. see Armstrong [3]) shall suffice. More precisely, the auction will offer the bidder the grand bundle with a take-it-or-leave-it price

$$r^* = (1 - \epsilon) \mathbf{E}[\sum_j v_j] .$$

If the bidder values the grand bundle above r^* , she will take the grand bundle and pay r^* ; no item is allocated otherwise and the bidder pays nothing. The proof of the next theorem follows straightforwardly from the definition of the mechanism and (ϵ, δ) -concentrated. So we will omit the tedious details.

THEOREM 3.2. *The grand-bundle-reserve-price auction is deterministically truthful, individually rational, and its expected revenue is at least $(1 - \epsilon)(1 - \delta) \mathbf{E}[\sum_j v_j]$ if the social welfare is (ϵ, δ) -concentrated.*

3.2.2 Constant Number of Bidders Now we show a similar result for multiple bidders. As a natural first

attempt, it might be tempting to think there exists a *reserve-revenue mechanism* with reserve revenue $r^* = (1 - \epsilon) \mathbf{E}[\sum_j \max_i v_{ij}]$ such that the mechanism offers the grand-bundle to all the bidders at a reserve price r^* and let the bidders discuss and decide whether to accept this offer and how to share the items and the costs if they decide to accept. Of course, the last step in the above procedure is not well-defined. The hope is that there is a truthful way for the bidders to come to a consensus of accepting the offer whenever the optimal social welfare is greater than r^* , since in such cases the bidders as a whole has positive surplus when buying the grand-bundle at the reserve price r^* . It is easy to see that this mechanism (if implementable) achieves a revenue of at least $(1 - \epsilon)(1 - \delta) \mathbf{E}[\sum_j \max_i v_{ij}]$. Unfortunately, we show that such mechanisms cannot be implemented in a truthful and IR manner. We will defer the discussion of this impossibility result to Appendix F.

Reserve-Welfare Mechanism In order to handle the multiple-bidder case, we will propose a novel mechanism in the VCG-family, which we shall refer to as the *reserve-welfare mechanism*. The idea is the following: it might be too aggressive to ask for a certain reserve revenue whenever the social welfare is above this reserve revenue; but it suffices to aim for the weaker goal of getting good revenue only when the social welfare is closed to its expectation because the social welfare concentrates by our assumption. Concretely, the *reserve-welfare mechanism* is defined in Figure 1.

Notice when $k = 1$, this mechanism indeed becomes the grand-bundle-reserve-price auction. So the reserve-welfare mechanism can be viewed as a generalization of the grand-bundle-reserve-price auction. We shall prove that this mechanism satisfies the desired properties.

THEOREM 3.3. *The reserve-welfare mechanism is deterministically truthful, individually rational, and its expected revenue is at least $(1 - k\epsilon - k\delta) \mathbf{E}[\sum_j \max_i v_{ij}]$ if the social welfare is (ϵ, δ) -concentrated for constants $\frac{1}{3} > \epsilon > 0$ and $1 > \delta > 0$.*

Briefly speaking, the proof goes as follows. By our definition of the payments, each bidder pays almost up to her value on the subset she gets when if social welfare is near the reserve welfare \hat{s} . Further, the social welfare will be near the reserve welfare \hat{s} almost for sure by our choice of \hat{s} and that the social welfare is (ϵ, δ) -concentrated. The only catch is the prices in the reserve-welfare mechanism might be negative when the values of the bidders are very large. We manage to show that the contribution of the negative prices can be bounded as well. So the expected revenue almost matches the expected social welfare. Below let us present the formal argument.

1. Let $\hat{s} = (1 - \epsilon) \mathbf{E} \left[\sum_j \max_i v_{ij} \right]$ be the reserve welfare.
2. If the optimal social welfare according to the bids, $\sum_j \max_i b_{ij}$, is smaller the reserve welfare \hat{s} , then no item is allocated and the bidders pay nothing.
3. Otherwise, allocate items according to an allocation $\mathbf{S} = (S_1, \dots, S_k)$ that maximizes the social welfare.
4. Charge bidder i price $p_i = \hat{s} - \sum_{\ell \neq i} \sum_{j \in S_\ell} b_{\ell j}$.

Figure 1: The reserve-welfare mechanism

Proof. (Theorem 3.3) If we omit step 2 and always allocate items according to the social-welfare-maximizing allocation, then the mechanism falls into the VCG family except that we are using the reserve welfare \hat{s} as our pivot instead of the Clarke pivot. So this variance of the reserve-welfare mechanism is deterministically truthful. Yet it is not individually rational. The reason of doing step 2 is exactly to fix the individual rationality.

Formally, for each bidder i , suppose her true valuations are \mathbf{v}_i and she bids \mathbf{b}_i . If reporting \mathbf{v}_i the items will not be allocated, then she should not lie and get the items allocated, since in the former case, her utility is 0, while in the latter case her utility is negative. Now assuming the items are allocated, her utility is

$$\sum_{j \in S_i} v_{ij} - p_i = \sum_{j \in S_i} v_{ij} + \sum_{\ell \neq i} \sum_{j \in S_\ell} b_{\ell j} - \hat{s} .$$

Note that the mechanism chooses the allocation that maximizes $\sum_i \sum_{j \in S_i} b_{ij}$. So by reporting her value truthfully the bidder maximizes her utility. Thus, the mechanism is deterministically truthful. Moreover, step 2 guarantees that allocation will be made only if $\sum_i \sum_{j \in S_i} b_{ij} \geq \hat{s}$. Therefore, the mechanism is individually rational.

Finally, let us analyze the revenue achieved by the reserve-welfare mechanism. We will let

$$s^* = \mathbf{E} \left[\sum_j \max_i v_{ij} \right]$$

denote the optimal expected social welfare, and recall that $\hat{s} = (1 - \epsilon)s^*$ is the reserve welfare. Assuming the bidders bid truthfully, the revenue is zero if the social welfare is less than \hat{s} , and is the following otherwise:

$$\begin{aligned} \sum_i p_i &= \sum_i \left(\hat{s} - \sum_{\ell \neq i} \sum_{j \in S_\ell} v_{\ell j} \right) \\ &= k\hat{s} - \sum_{\ell} \sum_{j \in S_\ell} (k-1)v_{\ell j} \\ &= k\hat{s} - (k-1) \sum_{j \in [n]} \max_{\ell} v_{\ell j} . \end{aligned}$$

By our assumption that the social welfare is (ϵ, δ) -concentrated, the expected revenue is at least

$$(3.1) \quad (1 - \delta)(1 - \epsilon)ks^* - (1 - \delta)(k - 1) \cdot \mathbf{E} \left[\sum_{j \in [n]} \max_{\ell} v_{\ell j} \mid \sum_{j \in [n]} \max_{\ell} v_{\ell j} \geq \hat{s} \right] .$$

Note that

$$(3.2) \quad \begin{aligned} s^* &\geq \Pr \left[\sum_{j \in [n]} \max_{\ell} v_{\ell j} \geq \hat{s} \right] \\ &\quad \mathbf{E} \left[\sum_{j \in [n]} \max_{\ell} v_{\ell j} \mid \sum_{j \in [n]} \max_{\ell} v_{\ell j} \geq \hat{s} \right] \\ &\geq (1 - \delta) \mathbf{E} \left[\sum_j \max_{\ell} v_{\ell j} \mid \sum_j \max_{\ell} v_{\ell j} \geq \hat{s} \right] . \end{aligned}$$

Combining (3.1) and (3.2) we get that the expected revenue of the reserve-welfare mechanism is at least

$$(1 - \delta)(1 - \epsilon)ks^* - (k - 1)s^* \geq (1 - k\delta - k\epsilon)s^* .$$

This proves the desired revenue guarantee.

Since the number of bidders is an absolute constant, Theorem 3.3 implies Lemma 3.2.

4 Many Bidders From the Same Population

As a natural restriction of the general multi-item auction problem, we will consider multi-item auctions with arbitrary number of items and bidders under the assumption that they are from the same population. Formally, for every item j , the value distributions \mathcal{F}_{ij} are identical for every bidder i . In this case, we manage to design nearly-optimal mechanisms based on our results for the few-bidder case.

THEOREM 4.1. (THM. 1.1 ELABORATED) *Suppose the bidders are from the same population, then there is a PTAS (polynomial in both n and m) for finding revenue-optimal mechanisms among (all settings require the mechanism to be IR):*

- IC/BIC mechanisms with discrete explicit access.
- DT mechanisms with discrete explicit access.¹
- BIC mechanisms with continuous oracle access.
- DT/IC mechanisms with continuous oracle access.¹

¹ Our mechanisms in these cases are only ϵ -deterministically truthful and ϵ -IC.

We will need the following lemma for i.i.d. MHR distributions following Theorem 7 in Cai and Daskalakis [7] and the proof therein.

LEMMA 4.1. *Suppose v_1, \dots, v_k are i.i.d. according to a MHR distribution, and $k \geq (12/\epsilon)^{12/\epsilon}$. Then there is a threshold r^* such that*

$$\Pr \left[\max_i v_i \geq r^* \right] \cdot r^* \geq (1 - \epsilon) \mathbf{E} \left[\max_i v_i \right] .$$

Moreover, one can efficiently find such a threshold r^* in polynomial time.

Roughly speaking, Cai and Daskalakis managed to improve their extreme value theorem when the bidders are i.i.d. so that consider the expectation of the random variable $\max_i v_i$, we only need to focus on the contribution from a small interval whose upper and lower bounds only differ by a $(1 + \epsilon)$ factor. As a simple corollary of this stronger extreme value theorem, we have the above lemma.

Equipped with this lemma, we are now ready to solve the case of arbitrary number of i.i.d. bidders.

Proof. (Theorem 4.1) Note that for each item j , the bidders' valuations for this item v_{1j}, \dots, v_{kj} are i.i.d. random variable according to a MHR distribution. Therefore, if the number of bidders k is greater than $(12/\epsilon)^{12/\epsilon}$, then by Lemma 4.1, we can find in polynomial time a threshold r_j^* for each item j such that

$$\Pr[\max_i v_{ij} \geq r_j^*] \cdot r_j^* \geq (1 - \epsilon) \mathbf{E}[\max_i v_{ij}] .$$

Therefore, if we run the second price auction with reserve prices r_j^* for each item j , then the expected revenue is at least

$$\sum_j \Pr[\max_i v_{ij} \geq r_j^*] \cdot r_j^* \geq (1 - \epsilon) \sum_j \mathbf{E}[\max_i v_{ij}] .$$

Note that the right-hand-side of the above inequality is the optimal expected social welfare and therefore

is an upper bound on the optimal revenue. So in the case when there are at least $(12/\epsilon)^{12/\epsilon}$ bidders, a simple reserve-price auction suffices to obtain a $(1 - \epsilon)$ fraction of the optimal revenue. Note that this mechanism is deterministic truthful and thus satisfies all our definitions of truthfulness.

So it suffices to solve the case when the number of bidders are smaller than $(12/\epsilon)^{12/\epsilon}$. But this falls into the case of constant number of bidders for any constant $\epsilon > 0$. So we could use the mechanism in Theorem 1.2 to solve the few-bidder case. In sum, we have proved the theorem.

References

- [1] Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *FOCS*, 2011.
- [2] Saeed Alaei, Hu Fu, Nima Haghpanah, Jason Hartline, and Azarakhsh Malekian. Bayesian optimal auctions via multi- to single-agent reduction. In *ArXiv*, 2012.
- [3] Mark Armstrong. Price discrimination by a many-product firm. *Review of Economic Studies*, 66(1):151–168, 1999.
- [4] Maria-Florina Balcan, Avrim Blum, Jason D. Hartline, and Yishay Mansour. Mechanism design via machine learning. In *FOCS*, pages 605–614. IEEE, 2005.
- [5] Anand Bhargat and Sanjeev Khanna. Personal communication.
- [6] Sayan Bhattacharya, Gagan Goel, Sreenivas Gollapudi, and Kamesh Munagala. Budget constrained auctions with heterogeneous items. In *STOC*, pages 379–388. ACM, 2010.
- [7] Yang Cai and Constantinos Daskalakis. Extreme-value theorems for optimal multidimensional pricing. In *FOCS*, 2011.
- [8] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *STOC*, 2012.
- [9] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue maximization to VCG. In *FOCS*, 2012, to appear.
- [10] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Reducing revenue to welfare maximization: Approximation algorithms and other generalizations. In *SODA*, 2013, to appear.
- [11] Shuchi Chawla, Jason D. Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In *EC*, pages 243–251. ACM, 2007.
- [12] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *STOC*, pages 311–320. ACM, 2010.

- [13] Constantinos Daskalakis and S. Matthew Weinberg. Symmetries and optimal mechanism design. In *EC*, 2012.
- [14] Shahar Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. In *STOC*, pages 139–148. ACM, 2011.
- [15] Shahar Dobzinski, Hu Fu, and Robert Kleinberg. Optimal auctions with correlated bidders are easy. In *STOC*, pages 129–138. ACM, 2011.
- [16] Shahar Dughmi and Jan Vondrák. Limitations of randomized mechanisms for combinatorial auctions. In *FOCS*, 2011.
- [17] Alejandro M. Manelli and Daniel R. Vincent. Multi-dimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1):153–185, 2007.
- [18] Roger B. Myerson. Optimal auction design. *Mathematics of operations research*, pages 58–73, 1981.
- [19] Amir Ronen. On approximating optimal auctions. In *EC*, pages 11–17. ACM, 2001.

A Omitted Proofs in Section 3

Proof. (Theorem 3.1) First, we will use Lemma 2.2 to truncate the random variables $X_j = \max_i v_{ij}$ and get \hat{X}_j , so that every \hat{X}_j lies in an interval where upper bound and lower bound are only $(\frac{1}{\epsilon} \log(\frac{1}{\epsilon}))$ factor away and $\mathbf{E}[\hat{X}_i] \geq (1 - O(\epsilon \log(\frac{1}{\epsilon}))) \mathbf{E}[X_i]$.

By Lemma 3.4, we partition $[n]$ into three sets R , S and T using the same ϵ . Let S_1 , S_2 and S_3 be the sets of items whose max value are in R , S and T , respectively. Then, the size of S_1 is a constant that only depends on ϵ and the following is true.

$$\begin{aligned}
 \sum_{j \in S_3} \mathbf{E}[\max_i v_{ij}] &\leq \frac{\sum_{\hat{X}_i \in T} \mathbf{E}[\hat{X}_i]}{1 - O(\epsilon \log(\frac{1}{\epsilon}))} \\
 &\leq \frac{\epsilon \sum_{i=1}^n \mathbf{E}[X_i]}{1 - O(\epsilon \log(\frac{1}{\epsilon}))} \\
 \text{(A.1)} \quad &\leq \frac{O(\epsilon) \text{opt}([n])}{1 - O(\epsilon \log(\frac{1}{\epsilon}))},
 \end{aligned}$$

where the last inequality follows from Lemma 3.3.

Let M_1 and M_2 be the $(1 - \epsilon)$ -approximate mechanisms from Lemma 3.1 and Lemma 3.2 respectively. Consider the following mechanism M for $[n]$:

1. Let the bidders submit their bids $\mathbf{b}_{1,[n]}, \dots, \mathbf{b}_{k,[n]}$.
2. Run M_1 on bids $\mathbf{b}_{1,S_1}, \mathbf{b}_{2,S_1}, \dots, \mathbf{b}_{k,S_1}$ and M_2 on bids $\mathbf{b}_{1,S_2}, \mathbf{b}_{2,S_2}, \dots, \mathbf{b}_{k,S_2}$.
3. Let \mathbf{S}' , \mathbf{p}' and \mathbf{S}'' , \mathbf{p}'' be the corresponding allocation and prices for M_1 and M_2 , give items in $\mathbf{S}'_i \cup \mathbf{S}''_i$ to bidder i , and charge him $p'_i + p''_i$.

Let $\mathcal{R}(N)$ be the revenue for mechanism N , then $\mathcal{R}(M) = \mathcal{R}(M_1) + \mathcal{R}(M_2)$. By Lemma 3.5, we know

$$\text{opt}([n]) \leq \text{opt}(S_1) + \max_{j \in S_2} \mathbf{E}[\max_i v_{ij}] + \max_{j \in S_3} \mathbf{E}[\max_i v_{ij}].$$

First of all, the contribution of the last term is small according to (A.1). So it suffices to obtain revenue close to $\text{opt}(S_1) + \max_{j \in S_2} \mathbf{E}[\max_i v_{ij}]$. Further, we know that $\mathcal{R}(M_1) \geq (1 - \epsilon) \text{opt}(S_1)$ and $\mathcal{R}(M_2) \geq (1 - \epsilon) \sum_{j \in S_2} \mathbf{E}[\max_i v_{ij}]$. Therefore, we have $\mathcal{R}(M) \geq (1 - O(\epsilon)) \text{opt}([n])$.

Finally, the truthfulness (with respect to the corresponding solution concept) of M follows straightforwardly from the truthfulness of M_1 and M_2 . So we have proved the theorem.

B Nearly Optimal Mechanism for Constant Number of Items and Bidders

B.1 Discrete Explicit Access Model In this setting, the problem of optimal mechanism design for revenue among IC and IR mechanisms or among BIC and IR mechanisms can be written as polynomial-size linear programs (each bidder might have many different values for an item). Therefore, we can efficiently find the optimal mechanism in these two settings. Since the LPs we used are very standard (e.g. see [4]), we will defer the discussion of these LPs to Appendix C.

For the problem of optimal mechanism design among deterministically truthful mechanisms, however, we need to solve the integer program version of the LP of optimal IC mechanisms. In order to do so, we need to reduce the size of the integer program from polynomial to constant. We will take the standard approach of rounding down each bidder's value to the nearest multiple of ϵ . As a result, for each bidder-item pair we only need to consider a constant number of possible valuations. Recall there are only a constant number of items and bidders, we can solve the constant-size integer program for this coarsened support set efficiently. As a result of the coarsening, however, we only get ϵ -deterministically truthful instead of perfect truthfulness.

B.2 Continuous Oracle Access Model

B.2.1 DT Mechanism for a Single Bidder By the taxation principle, any deterministically truthful mechanism can be interpreted as a bundle-pricing mechanism: the bidder is given a menu of bundles of items such that the prices of the bundles are independent on the reported values; moreover, the bidder always gets one of the utility-maximizing bundles. In other words, it suffices to find the nearly optimal bundle prices. In order to do so, we first show that in order to obtain nearly

optimal revenue it suffice to consider a finite number of prices for each bundle via a standard price discretization lemma attributed to Nisan (e.g. see [11]). Then, we can search over all possible bundle-pricings within the discretized price set and choose the optimal one. Since there are only constant number of items and thus constant number of bundles, such brute-force search can be done efficiently. For completeness we include a formal statement and the proof of the price discretization lemma in Appendix D.

B.2.2 IC Mechanism for a Single Bidder In this case, our starting point is again the taxation principle. Any IC mechanism can be interpreted as a lottery-pricing mechanism: the bidder is given a (not necessarily finite) menu of lotteries, each of which is represented by a vector of the probabilities of getting each item, such that the prices of the lotteries are independent on the reported values; moreover, the bidder always gets the utility maximizing lottery. By the same price discretization lemma, we only need to consider a finite number of prices for each lottery. However, there is an infinite number of possible lotteries. We settle this problem by showing the lottery space can be discretized as well. Concretely, we prove that in order to obtain $1 - O(\epsilon)$ of the optimal revenue, it suffices to consider lotteries in which the probabilities of getting each item are powers of $(1 + \epsilon^2)$ and are greater than ϵ^2 . As a result, we can combine the lottery discretization lemma and the pricing discretization lemma to show that in order to get nearly optimal lottery pricing it suffices to search over constant number of lottery-pricing mechanisms and choose the best one. The proof of the lottery discretization lemma is deferred to Appendix E.

B.2.3 DT/IC/BIC Mechanisms for Multiple Bidders In order to solve the problem for multiple bidders, we use a reduction to the discrete case: discretize the prior distributions by rounding each sampled value to the closest powers of $(1 + \epsilon)$ and truncate values that are too large or too small according to the extreme value theorem in [7]. We then find the nearly optimal mechanism for the coarsened problem via the integer programming/linear programming approach for the discrete case. Finally, we will round the bids of the bidders to the closest powers of $(1 + \epsilon)$, run the above mechanism on the coarsened bids, and use the allocation and prices chosen by the mechanism. As a result of the rounding, the mechanisms we obtain are only ϵ -truthful with respect to the corresponding solution concepts. Nonetheless, in the BIC case, we can use the technique recently developed by Daskalakis and Weinberg [13] to convert our ϵ -BIC mechanism into a BIC one with only a small

additional loss in the expected revenue.

C Linear Programs for Multi-Item Auctions

It has long been known that if the support set is finite, then the problem of designing truthful (IC/BIC) mechanism that achieves optimal revenue can be characterized by a linear problem. For completeness, we will describe the standard linear programs for the multi-item auction.

For any type profile \mathbf{v} , any bidder i , and any item j , we let $x(\mathbf{v})_{ij}$ denote the probability that bidder i gets item j when the valuations are \mathbf{v} , and let $p(\mathbf{v})_i$ denote the expected payment of bidder i . The problem of optimal multi-item auction among IC mechanisms has the following exact LP characterization:

$$\begin{aligned} & \text{Maximize} && \sum_{\mathbf{v}} \Pr[\mathbf{v}] \sum_{i=1}^k p(\mathbf{v})_i && \text{s.t.} \\ & \forall j, \mathbf{v} && \sum_{i=1}^k x(\mathbf{v})_{ij} \leq 1 \\ & \forall i, \mathbf{v}, v'_i && \sum_{j=1}^n x(\mathbf{v})_{ij} v_{ij} - p(\mathbf{v})_i \geq \\ & && \sum_{j=1}^n x(v'_i, v_{-i})_{ij} v_{ij} - p(v'_i, v_{-i})_i \\ & \forall i, \mathbf{v} && \sum_{j=1}^n x(\mathbf{v})_{ij} v_{ij} - p(\mathbf{v})_i \geq 0 \end{aligned}$$

The LP characterization of the problem of optimal multi-item auction among BIC mechanisms is almost the same, except for replacing the IC constraints with the following BIC constraints for all i , \mathbf{v} , and v'_i :

$$\begin{aligned} & \sum_{v_{-i}} \Pr[v_{-i}] \left(\sum_{j=1}^n x(v_i, v_{-i})_{ij} v_{ij} - p(v_i, v_{-i})_i \right) \\ & \geq \sum_{v_{-i}} \Pr[v_{-i}] \left(\sum_{j=1}^n x(v'_i, v_{-i})_{ij} v_{ij} - p(v'_i, v_{-i})_i \right). \end{aligned}$$

D Price Discretization Lemma

The following price discretization lemma is attributed to Nisan (e.g., [4, 11]):

LEMMA D.1. *For $\epsilon \in (0, 1)$, let \mathbf{p} and \mathbf{p}' be two bundle pricing schemes such that for any bundle i , $p_i \in [1 - \epsilon, 1 - \epsilon + \epsilon^2] p'_i$. Suppose the bidder buys bundle j when \mathbf{p} are the prices and buys bundle ℓ when \mathbf{p}' are the prices, then $p_j \geq (1 - 2\epsilon) p'_\ell$.*

Proof. By our assumption, we have $v_j - p_j \geq v_\ell - p_\ell$ and $v_\ell - p'_\ell \geq v_j - p'_j$. Summing up the two inequalities and

cancelling the common terms, we have $p'_j - p_j \geq p'_\ell - p_\ell$. Note that by our assumption $p'_j - p_j \leq p'_j - (1 - \epsilon)p'_j = \epsilon p'_j$, and $p'_\ell - p_\ell \geq p'_\ell - (1 - \epsilon + \epsilon^2)p'_\ell = (\epsilon - \epsilon^2)p'_\ell$. So we have $p'_j \geq (1 - \epsilon)p'_\ell$. Finally, $p_j \geq (1 - \epsilon)p'_j$. So $p_j \geq (1 - \epsilon)^2 p'_\ell \geq (1 - 2\epsilon)p'_\ell$.

By Lemma D.1 we know that it suffices to consider prices that are powers of $(1 + \epsilon^2)$ in order to get $(1 - 2\epsilon)$ of the optimal revenue. Of course, we still have infinite number of prices to consider. In order to settle this problem, we will use the extreme value theorem in [7] to conclude that for each bundle it suffices to consider prices that are in a range whose upper and lower bounds differ by at most an $O(\frac{1}{\epsilon} \log(\frac{1}{\epsilon}))$ factor (this range may be different for different bundles). Therefore, we only need to consider $O(\log_{1+\epsilon}(\frac{1}{\epsilon} \log(\frac{1}{\epsilon}))) = O(\frac{1}{\epsilon} \log(\frac{1}{\epsilon}))$ number of prices per bundle.

E Lottery Discretization Lemma

The following lottery discretization lemma is inspired by the idea in the price discretization lemma. First, let us define some notations. We will use a n -dimensional vector $\mathbf{q} = (q_1, \dots, q_n)$ to denote a lottery where q_j is the probability of getting item j . A lottery menu is a collection (may or may not be finite) of lottery-price pairs: $\{(\mathbf{q}_1, p_1), (\mathbf{q}_2, p_2), \dots\}$.

LEMMA E.1. *Suppose $\epsilon \in (0, 1)$. Consider the optimal lottery menu L and another lottery menu L' obtained by rounding up probabilities of each lottery $(\mathbf{q}_i, p_i) \in L$ into (\mathbf{q}'_i, p_i) such that for all $j \in [n]$, $q'_{ij} \in [1 + \epsilon - \epsilon^2, 1 + \epsilon]q_{ij}$. Then, the expected revenue from menu L' is at least a $(1 - O(\epsilon))$ fraction of that from menu L .*

Proof. Suppose the type profile is \mathbf{v} . Further, let us assume the bidder buys lottery (\mathbf{q}_j, p_j) when L is presented and $(\mathbf{q}'_\ell, p_\ell)$ when L' is presented. We have

$$(E.2) \quad \begin{aligned} \mathbf{v} \cdot \mathbf{q}_j - p_j &\geq \mathbf{v} \cdot \mathbf{q}_\ell - p_\ell \\ \mathbf{v} \cdot \mathbf{q}'_\ell - p_\ell &\geq \mathbf{v} \cdot \mathbf{q}'_j - p_j \end{aligned}$$

By summing up these two inequalities and cancelling the common terms, we have

$$\mathbf{v} \cdot \mathbf{q}'_\ell - \mathbf{v} \cdot \mathbf{q}_\ell \geq \mathbf{v} \cdot \mathbf{q}'_j - \mathbf{v} \cdot \mathbf{q}_j .$$

By our assumption, we further have

$$\mathbf{v} \cdot \mathbf{q}'_j - \mathbf{v} \cdot \mathbf{q}_j \geq \mathbf{v} \cdot (1 + \epsilon - \epsilon^2)\mathbf{q}_j - \mathbf{v} \cdot \mathbf{q}_j = (\epsilon - \epsilon^2)\mathbf{v} \cdot \mathbf{q}_j ,$$

and

$$\mathbf{v} \cdot \mathbf{q}'_\ell - \mathbf{v} \cdot \mathbf{q}_\ell \leq \mathbf{v} \cdot (1 + \epsilon)\mathbf{q}_\ell - \mathbf{v} \cdot \mathbf{q}_\ell = \epsilon\mathbf{v} \cdot \mathbf{q}_\ell .$$

Therefore, we have $\mathbf{v} \cdot \mathbf{q}_\ell \geq (1 - \epsilon)\mathbf{v} \cdot \mathbf{q}_j$. By this inequality and (E.2), we have

$$p_\ell \geq p_j + \mathbf{v} \cdot \mathbf{q}_\ell - \mathbf{v} \cdot \mathbf{q}_j \geq p_j - \epsilon\mathbf{v} \cdot \mathbf{q}_j .$$

Hence, if we compare the expected revenue from L' , $\mathbf{E}[p_\ell]$, and the expected revenue of L , $\mathbf{E}[p_j]$, then the former is worse than the latter by no more than an ϵ fraction of the social welfare by L . We further note that the optimal social welfare and the optimal revenue differ by at most a constant factor. Thus, we have proved the lemma.

By Lemma E.1, we can round up the probabilities in each lottery to some powers of $(1 + \epsilon^2)$ in order to get $1 - O(\epsilon)$ of the optimal revenue. There is only one catch in this argument: by rounding up the probabilities, some of them may exceed 1 and therefore become infeasible. We resolve this problem by rounding down the probabilities as well as the prices of the resulting discretized lotteries by a factor of $1 - \epsilon$. By doing so, we retain feasibility with the extra cost of a $1 - \epsilon$ factor, but we still gets $1 - O(\epsilon)$ of the optimal revenue.

Lemma E.1 reduces the number of lotteries from uncountably infinite to countably infinite. We observe that we can further reduce this number to finite by dropping invaluable lotteries and the negligible entries in the valuable lotteries. Concretely, if the expected value of a lottery is at most an ϵ fraction of the expected welfare, then we can ignore this lottery because the total revenue from such lotteries is at most an $O(\epsilon)$ fraction of the optimal. Next, suppose we have a lottery whose expected value is at most an ϵ fraction of the expected social welfare. Then, any entry smaller than ϵ^2 contributes at most an $O(\epsilon)$ fraction to the expected value of this lottery, and hence can be dropped.

F Impossibility of Truthful Reserve-Revenue Mechanism

In this section, we will show that the reserve-revenue mechanisms can not be implemented in a truthful and individually rational manner. First of all, let us formally define the family of reserve-revenue mechanisms.

DEFINITION 8. *A mechanism M is a reserve-revenue mechanism if there is a reserve revenue $r^* > 0$ and a threshold social welfare $s^* > 0$ (presumably $s^* \geq r^*$), such that M achieves revenue at least r^* whenever the social welfare (according to the bids) is at least s^* .*

We have the following negative result.

THEOREM F.1. *If there are at least $k \geq 2$ bidders and $n \geq 2$ items, then there are no truthful and individually rational reserve-revenue mechanism.*

Proof. It suffices to prove the theorem for the case of $k = n = 2$. Assume for contradiction that M is a truthful and individually rational reserve-revenue

mechanism with reserve revenue r^* and threshold social welfare s^* .

Let us consider what happens when $v_{11} = s^* - \frac{r^*}{3}$, $v_{12} = 0$, $v_{21} = 0$, $v_{22} = s^* - \frac{r^*}{3}$. We claim that in this case $p_1, p_2 \leq \frac{r^*}{3}$.

Consider the alternative type profile in which bidder 1's values are $v'_{11} = \frac{r^*}{3}$ and $v'_{12} = 0$, and bidder 2's values are still the same. Note that the social welfare for this type profile is exactly s^* . So M shall achieve revenue at least r^* . It is obvious that the only allocation that could achieve this level of revenue in an individually rational fashion is to give bidder 1 item 1 and to give bidder 2 item 2. The price for bidder 1 in this case is at most $\frac{r^*}{3}$. By the taxation principle, from bidder 1's viewpoint any truthful mechanism should look like a menu of lotteries over possible outcomes with prices that do not depend on the value of bidder 1. Moreover, bidder 1 should always get one of the utility maximizing lottery. Therefore, we know that the lottery that corresponds to getting item 1 and not getting item 2 with probability 1 is available to bidder 1 with price at most $\frac{r^*}{3}$ when bidder 2 bids v_2 and it is bidder 1's utility-maximizing lottery when her valuation is v'_1 . Note that the only difference between v_1 and v'_1 is the value for item 1 increases. So we conclude that when the type profile is v_1, v_2 , bidder 1 should purchase the same lottery with the same price. Hence, we have proved that $p_1 \leq \frac{r^*}{3}$.

Similarly, we can show that $p_2 \leq \frac{r^*}{3}$. Now we get that the revenue achieved by M when the values are v_1 and v_2 is at most $p_1 + p_2 \leq \frac{2r^*}{3} < r^*$. Thus, we have obtained a contradiction.

REMARK F.1. *The conditions in Theorem F.1 cannot be relaxed for that if $k = 1$, then the grand-bundle-reserve-price auction is a reserve-revenue mechanism, and if $n = 1$, then the standard reserve price auction is a reserve-revenue mechanism.*