On the Complexity of Simple and Optimal Deterministic Mechanisms for an Additive Buyer

Xi Chen[*]   George Matikas[†]   Dimitris Paparas[‡]   Mihalis Yannakakis[§]

Abstract

We show that the Revenue-Optimal Deterministic Mechanism Design problem for a single additive buyer is #P-hard, even when the distributions have support size 2 for each item and, more importantly, even when the optimal solution is guaranteed to be of a very simple kind: the seller picks a price for each individual item and a price for the grand bundle of all the items; the buyer can purchase either the grand bundle at its given price or any subset of items at their total individual prices. The following problems are also #P-hard, as immediate corollaries of the proof:

1. determining if individual item pricing is optimal for a given instance,

2. determining if grand bundle pricing is optimal, and

3. computing the optimal (deterministic) revenue.

On the positive side, we show that when the distributions are i.i.d. with support size 2, the optimal revenue obtainable by any mechanism, even a randomized one, can be achieved by a simple solution of the above kind (individual item pricing with a discounted price for the grand bundle) and furthermore, it can be computed in polynomial time. The problem can be solved in polynomial time too when the number of items is constant.

1 Introduction

Consider the following natural scenario: A customer walks in a grocery store with the intention of buying some items. The store owner has statistical information from past customers that reveals how much a typical customer values each item. Her goal is to assign prices for the items and offer discounts for bundles of them to encourage the customer to spend more money in a way that maximizes her expected revenue.

In this paper we formally study practices like the above, which we refer to as the *optimal bundle-pricing* problem, under the setting where a single *additive* buyer is interested in n *heterogeneous* items offered by a seller. While the buyer's values for the items are unknown, the seller is given as input a product distribution $\mathcal{F} = \mathcal{F}_1 \times \cdots \times \mathcal{F}_n$ from which the valuations $\mathbf{v} = (v_1, \ldots, v_n)$ of the buyer for the n items are drawn, where each $\mathcal{F}_i$ is a discrete distribution given explicitly (by listing its support and probabilities). The seller offers a finite *menu* M of *bundles* to the buyer (or a *bundle-pricing*), with each entry of the menu consisting of a subset (bundle) $T \subseteq [n]$ of items and the price π_T at which it is sold. Given a menu M, the buyer draws her valuations $\mathbf{v}$ from $\mathcal{F}$ and then either buys a bundle T from M that maximizes her utility $\sum_{i \in T} v_i - \pi_T$ or nothing if the utility of every bundle in M is negative[1]; the price π_T of the bundle bought by the buyer is the *revenue* of the seller. The goal of the seller is to find a menu that maximizes her expected revenue (i.e., the expected price π_T that the buyer pays), which is known to be equivalent to the problem of *Revenue Optimal Deterministic Mechanism Design*. If we extend bundle-pricings to allow the seller to offer a finite menu of lotteries (or a *lottery-pricing*), where a lottery is a pair $((x_1, \ldots, x_n), \pi)$ with π being its price and $x_i \in [0, 1]$ being the probability of the buyer getting item $i \in [n]$ if this lottery is purchased, we obtain the *optimal lottery-pricing problem*, also known in the literature as the problem of *Revenue Optimal Randomized Mechanism Design*. We define these problems formally in Section 2.[2]

It is worth pointing out that bundle-pricing schemes commonly used in practice do not necessarily list explic-

---

[1]Ties in utility are broken in favor of a bundle with higher price (and value).

[2]We choose to follow the language of bundle-pricings and lottery-pricings in this paper, instead of deterministic and randomized mechanisms, mainly because they are conceptually closer to common practices seen in the real world and are easier to understand for readers that are not familiar with mechanism design (e.g., there is no need to introduce the notion of truthfulness).

itly the bundles offered in the menu, but may specify them implicitly in a succinct manner. For example, in the case of offering a simple *item-pricing* without any discounts for bundles, the seller needs only to specify a price for each item (n numbers in total); the induced menu consists of all 2^n subsets of items, each priced at the sum of prices of items in the subset, and has the desired property that the buyer's problem, i.e., finding an optimal bundle in the menu given $\mathbf{v}$, is easy to solve. Because of this, it would not be appropriate to require the output of the bundle-pricing problem to be an explicit list of bundles in an optimal menu, but rather it can be represented in a reasonably succinct way. The exact representation of the output, however, will not affect our main results, as we explain when we describe them later in this section.

Both optimal deterministic and randomized mechanism design problems have been studied intensively during the past decade [Tha04, GHK$^+$05, Bri08, CHK07, CHMS10, CD11, CMS10, BCKW10, Pav10, WT14, BILW14, DDT14, DDT12, DDT13, DDT15, HN12, MV06, Rub16, LY13, Yao15]. For some instances, randomized mechanisms can achieve strictly higher revenue than deterministic mechanisms. However, deterministic mechanisms (bundle-pricings) are much more widely used in practice (especially "simple" pricing schemes); we will focus on deterministic mechanisms in this paper. Recently, much effort has been devoted to understanding the power and limitations of simple pricing schemes, that is, menus that can be described succinctly in a natural way and at the same time induce an easy-to-solve buyer's problem. Some of the examples include (i) selling all items separately (item-pricing), (ii) selling only the grand bundle that consists of all items (grand-bundle pricing), and (iii) partition mechanisms, where one partitions the items into disjoint groups, each with its own price, and sells the groups separately. While it is known that none of these solutions is optimal in general among bundle-pricings, there has been substantial work studying basic questions for each of these simple solutions, including the following: How does the revenue achievable by these solutions compare with optimal revenues achievable by bundle or lottery pricings? What are conditions under which these solutions are optimal? Can we compute an optimal solution of each type?

In case (i) of selling the items separately, we know how to compute efficiently an optimal item-pricing: each item is assigned separately its optimal price following Myerson's theory [Mye81]. In both cases (ii) and (iii) of the grand bundle and partition mechanisms, the problem of finding an optimal solution is intractable (#P-hard [DDT12] and NP-hard [Rub16], respectively). However, the fact that it is hard to find an optimal solution of a certain type (grand bundle or partition mechanisms) does not mean that one cannot easily find a solution that is not of this type and has higher revenue (for example, by selling also individual items), or possibly even find a solution that is optimal among all bundle-pricings. Thus, two central questions remain concerning the bundle-pricing (or optimal deterministic mechanism design) problem:

1. Is there an efficient algorithm that finds an optimal bundle-pricing?

2. If the problem above is hard in general, is there such an algorithm when the instance is promised to have a "simple" optimal bundle-pricing?

Our results resolve both questions in the negative by showing that the problem is #P-hard, even when the distributions have support size 2 for each item and, more importantly, even when the instance is promised to have a unique optimal bundle-pricing that is of a very simple kind, which we call a *discounted item-pricing*: the seller picks a price for each individual item and a price for the grand bundle of all the items; the buyer can purchase either the grand bundle at its given price or any subset of items at their total individual prices. Such a solution can be described using $n + 1$ numbers and the buyer's problem is also easy to solve. This is the reason why the exact output format of the problem does not affect our hardness result.

This result tells us that the bundle-pricing (deterministic mechanism design) problem is inherently computationally hard, and furthermore the difficulty is not (only) due to the fact that the optimal solution can be very complex, of the kind that one would not use in practice anyway; the problem is hard even when the optimal solution is extremely simple: standard item pricing with a discount for the grand bundle.

As a by-product of the proof, we also resolve in the negative the question of whether there is a 'nice' characterization (i.e., an easy-to-check necessary and sufficient condition) of when item-pricings are optimal, i.e., whether an item-pricing can achieve the optimal revenue achievable by bundle pricings or whether bundling helps. The same applies to grand-bundle pricings (and partition mechanisms), i.e., there is no easy-to-check characterization for the optimality of grand-bundle pricings under standard complexity-theoretic assumptions.

On the positive side, we show that when $\mathcal{F}_1, \ldots, \mathcal{F}_n$ are i.i.d. with support size 2, the optimal revenue achievable by any pricing scheme, even a lottery one, can be achieved by a discounted item-pricing which, furthermore, can be computed in polynomial time. We discuss our results in detail below in Section 1.1.

1.1 Our Results. Given an input distribution $\mathcal{F}$ we use $\mathrm{BReV}(\mathcal{F}), \mathrm{SReV}(\mathcal{F}), \mathrm{DReV}(\mathcal{F})$ and $\mathrm{ReV}(\mathcal{F})$ to denote the optimal expected revenues achievable by a grand-bundle pricing (i.e., selling the grand bundle only), an item-pricing (i.e., selling all items separately), a bundle-pricing, and a lottery-pricing, respectively.

First we state our positive result for i.i.d. distributions with support size 2:

THEOREM 1.1. *When $\mathcal{F}_1, \ldots, \mathcal{F}_n$ are i.i.d. with support size 2 ($\{a, b\}$ with $a < b$), $\mathrm{ReV}(\mathcal{F})$ can always be achieved by a discounted item-pricing where the grand bundle is priced at $kb + (n - k)a$ for some $k \in [0 : n]$ and each item is priced at b. Moreover, the parameter k can be found in polynomial time.*

Our main result addresses the two questions from the introduction in the negative: we show that it is #P-hard to find an optimal bundle-pricing for a single additive buyer, even when $\mathcal{F}$ is a product distribution and each $\mathcal{F}_i$ has support size 2. Although in general an optimal solution can be highly complex and consist of exponentially many bundles without a succinct description, our hardness result is established on instances that are guaranteed to have a unique and very simple optimal solution, namely, the discounted item-pricing that we defined earlier. Such a pricing scheme corresponds to the ubiquitous practice of offering an individual price π_i for each item i and also the grand bundle of all items at a discounted price π, for example a combo of a toothpaste, a toothbrush, dental floss, and mouth wash offered at a 15% discount as compared to the cost of buying them separately. The buyer can choose to buy the grand bundle at π or any subset T of items at $\sum_{i \in T} \pi_i$, whichever brings the highest (nonnegative) utility (note that in the latter case the buyer will obviously buy the set T of all items whose price is less than or equal to the buyer's value). While a discounted item-pricing offers exponentially many bundles, it has a succinct representation by $n + 1$ numbers and is easy to implement in practice. We state our main hardness result in Theorem 1.2:

THEOREM 1.2. *The optimal bundle-pricing problem is #P-hard even when (1) all distributions have support size 2 and (2) the instance is promised to have a unique optimal solution that is a discounted item-pricing.*

Indeed, the hard instances constructed in the proof of Theorem 1.2 have the property that either

  (i) the grand-bundle pricing[3] that offers the grand bundle at the sum of low values of all items; or

  (ii) the discounted item-pricing that offers each individual item at its high value and the grand bundle at a specific value that can be computed from the instance in polynomial time,

is guaranteed to be optimal among all bundle-pricings, but it is #P-hard to determine which one is better. Note that (i) can be equivalently described as an item-pricing, with each item priced at its low value, and the revenue can be computed in polynomial time. These observations together lead to a number of corollaries.

COROLLARY 1.1. *The following problems are #P-hard:*

  1. *Given a product distribution $\mathcal{F}$, decide whether $\mathrm{DReV}(\mathcal{F}) = \mathrm{SReV}(\mathcal{F})$, i.e., whether an item-pricing is optimal among all bundle-pricings.*

  2. *Given a product distribution $\mathcal{F}$, decide whether $\mathrm{DReV}(\mathcal{F}) = \mathrm{BReV}(\mathcal{F})$, i.e., whether a grand-bundle pricing is optimal among all bundle-pricings.*

COROLLARY 1.2. *The following problems are #P-hard:*

  1. *Given a product distribution $\mathcal{F}$, compute $\mathrm{DReV}(\mathcal{F})$.*

  2. *Given a product distribution $\mathcal{F}$ and a valuation $\mathbf{v}$, compute the bundle bought at $\mathbf{v}$ in any optimal bundle-pricing.*

We remark finally that all the hardness results hold if the number of items is unbounded. For a constant number of items, we obtain a polynomial-time algorithm (though the dependency of its running time on the number of items is exponential).

THEOREM 1.3. *When the number of items is constant, an optimal bundle-pricing can be computed in polynomial time.*

1.2 Related Work. The seminal work of Myerson [Mye81] completely settles the case of selling a single item, by giving a computationally efficient and deterministic mechanism (i.e., a pricing of the item) that maximizes the expected revenue among all possible, randomized or deterministic, mechanisms. The more general multi-dimensional setting, however, turns out to be inherently more difficult. Unlike Myerson's setting, randomization in general improves the revenue when there are many items for sale, even if there is a single unit-demand buyer [Tha04] (i.e. the buyer wants to buy only one item) or an additive buyer [MV06]. It is also known that the optimal menu of lotteries may have exponential size [CDO+15, HN13]. Moreover, under standard complexity-theoretic assumptions,

recent results rule out the existence of computationally efficient algorithms that find a revenue-optimal deterministic or randomized mechanism for a unit-demand buyer [CDP+14, CDO+15, GHK+05, Bri08], or a randomized mechanism for an additive buyer [DDT14]. However, hardness results for the optimal deterministic mechanism design problem with an additive buyer are limited. Rubinstein [Rub16] proved that finding an optimal partition mechanism is strongly NP-hard; Daskalakis et al. [DDT12] proved that finding an optimal price for selling the grand bundle is #P-hard. These results are for restrictions of the problem that impose a specific menu structure, and the original problem remained open before this work.

Among these hardness results, the one that is most relevant to ours is that of Daskalakis, Deckelbaum, and Tzamos [DDT14]. They construct instances $\mathcal{F}$ with support size 2 for each $\mathcal{F}_i$ and show that the problem of finding an optimal lottery-pricing (or randomized mechanism) is #P-hard. This, however, does not have any consequences for the bundle-pricing problem for two reasons. First, the deterministic mechanism design problem is not necessarily harder than the randomized one. In fact, in the setting of a unit-demand buyer, the deterministic problem is provably "easier": the randomized problem is #P-hard [CDO+15] while the deterministic one is in NP [CDP+14]. Second, for the construction of [DDT14] to work for bundle-pricings, it would need to be the case that optimal menus of lotteries of their instances are deterministic and consist of bundles only. However, this is not the case: the solution in [DDT14] makes essential use of the randomization feature and the optimal menu for $\mathcal{F}$ contains a large number of lotteries (with probabilities in $(0, 1)$) for valuations in a certain "critical" region. Compared to techniques used in [DDT14], ours are different in the following two aspects: (1) Since $\mathrm{DREV}(\mathcal{F})$ is captured by an integer program (instead of a linear program for $\mathrm{REV}(\mathcal{F})$; see Section 2), we cannot use LP duality but have to rely on more discrete and combinatorial arguments to identify its optimal integer solutions; (2) An important step in both proofs is to relax the integer (or linear) program that captures $\mathrm{DREV}(\mathcal{F})$ (or $\mathrm{REV}(\mathcal{F})$). Our relaxation is significantly different from the LP relaxation of [DDT14]. We need to keep a large set of global constraints from the original IP while local constraints suffice for the purpose of [DDT14].

Most of the work on the deterministic mechanism design problem for an additive buyer so far focuses on approximation. Hart and Nisan [HN12] studied two simple deterministic mechanisms for product distributions: selling items separately or selling the grand bundle only. They showed that selling items separately and grand

bundling are respectively $\Omega(1/\log^2 n)$ and $\Omega(1/\log n)$ approximations of the optimal revenue achievable by any (possibly randomized) mechanism (later improved by Li and Yao [LY13] to $\Omega(1/\log n)$ for both schemes, which is known to be tight [HN12]). While neither of these two schemes can achieve by itself a constant factor approximation, Babaioff et al. [BILW14] showed that the *better* of the two gives a $(1/6)$-approximation. Recently, Daskalakis et al. [DDT13, DDT15] studied conditions for grand-bundling mechanisms to be optimal (for continuous distributions), and showed that this happens if and only if two stochastic dominance conditions hold. Rubinstein [Rub16] worked on partition mechanisms and obtained a polynomial-time approximation scheme (PTAS) for a revenue maximizing partition mechanism. A number of other results [DW12, CH13] obtained approximation schemes for i.i.d. distributions with the MHR property. Giannakopoulos and Koutsoupias [GK14] obtained optimal mechanisms for i.i.d. uniform distributions with up to six items. Finally Yao [Yao15] introduced a new approach for reducing the k-item n-bidder problem to the k-item 1-bidder setting and gave a deterministic mechanism that yields at least a constant fraction of the optimal revenue for the more general k-item n-bidder setting.

We also note that there is extensive work studying unit-demand buyers (e.g., [CDP+14, CDO+15, GHK+05, Bri08, CHK07, CHMS10, CD11, CMS10, BCKW10, Pav10, WT14]). Besides the papers cited earlier that address the complexity of an optimal mechanism in that context, the rest of the work, which mostly concerns special cases or approximation, is not directly related to the topic of the present paper.

2 Preliminaries

Let D_i be the support of $\mathcal{F}_i$, and $D = D_1 \times \cdots \times D_n$ be the set of valuation vectors. For each $\mathbf{v} \in D$, let

$$\Pr[\mathbf{v}] = \Pr_{\mathcal{F}_1}[v_1] \times \cdots \times \Pr_{\mathcal{F}_n}[v_n]$$

denote the probability of $\mathbf{v}$ drawn from $\mathcal{F}$.

We first define $\mathrm{DREV}(\mathcal{F})$, the optimal expected revenue obtainable by a bundle-pricing, by formulating it using an integer program with $n + 1$ variables associated with each valuation $\mathbf{v} \in D$: $x_{\mathbf{v},1}, \ldots, x_{\mathbf{v},n}$ and $\pi_{\mathbf{v}}$, where $x_{\mathbf{v},i} \in \{0, 1\}$ indicates whether item i is included in the bundle the buyer chooses from the menu (with $x_{\mathbf{v},i} = 1$ if item i is included) when her valuation is $\mathbf{v}$ and $\pi_{\mathbf{v}}$ denotes the price of the bundle. We also write $\mathbf{x}_{\mathbf{v}} = (x_{\mathbf{v},1}, \ldots, x_{\mathbf{v},n}) \in \{0, 1\}^n$ to denote the allocation vector for valuation $\mathbf{v}$. The integer program then maximizes the expected revenue: $\sum_{\mathbf{v} \in D} \pi_{\mathbf{v}} \cdot \Pr[\mathbf{v}]$ subject to the following constraints:

1. $x_{\mathbf{v},i} \in \{0,1\}$ for all $\mathbf{v} \in D$;

2. For each $\mathbf{v} \in D$, the utility is nonnegative:
$$\sum_{i \in [n]} v_i \cdot x_{\mathbf{v},i} - \pi_{\mathbf{v}} \geq 0;$$

3. For all $\mathbf{w}, \mathbf{v} \in D$, $\mathbf{w}$ does not envy the bundle of $\mathbf{v}$:
$$\sum_{i \in [n]} w_i \cdot x_{\mathbf{w},i} - \pi_{\mathbf{w}} \geq \sum_{i \in [n]} w_i \cdot x_{\mathbf{v},i} - \pi_{\mathbf{v}}.$$

We refer to this integer program as the *standard IP* for $\mathrm{DRev}(\mathcal{F})$ and the goal of the optimal bundle-pricing problem is to find an optimal solution to the standard IP. As discussed earlier, the exact way of defining the output of the problem does not affect our main results. (For example, one can adopt the model used in [DDT14, CDO$^+$15], where a polynomial-time algorithm A for the optimal bundle-pricing problem takes as input a distribution $\mathcal{F}$ and a valuation $\mathbf{v} \in D$ and outputs a bundle $A(\mathcal{F}, \mathbf{v})$ such that $\{A(\mathcal{F}, \mathbf{v}) : \mathbf{v} \in D\}$ is an optimal solution to the standard IP for $\mathrm{DRev}(\mathcal{F})$. Under this formulation Theorem 1.2 implies that there cannot be any such polynomial-time algorithm unless #P can be solved in polynomial time.)

The equivalence between the optimal bundle-pricing problem and deterministic mechanism design follows from the observation that any feasible solution $\{\mathbf{x}_{\mathbf{v}}, \pi_{\mathbf{v}} : \mathbf{v} \in D\}$ to the standard IP for $\mathrm{DRev}(\mathcal{F})$ can be equivalently viewed as a deterministic mechanism that is both individually rational and truthful, and vice versa: the mechanism, upon $\mathbf{v}$ reported by the buyer, assigns items $\mathbf{x}_{\mathbf{v}}$ to the buyer and charges her $\pi_{\mathbf{v}}$.

Sometimes (e.g., in Section 4), it is more convenient to replace $\pi_{\mathbf{v}}$ by a nonnegative utility variable $u_{\mathbf{v}}$. The standard IP maximizes the same expected revenue:
$$\sum_{\mathbf{v} \in D} \left(\sum_{i \in [n]} v_i \cdot x_{\mathbf{v},i} - u_{\mathbf{v}} \right) \cdot \Pr[\mathbf{v}]$$

subject to the following (slightly simpler) constraints:

1. $x_{\mathbf{v},i} \in \{0,1\}$ and $u_{\mathbf{v}} \geq 0$ for all $\mathbf{v} \in D$;

2. For all $\mathbf{w}, \mathbf{v} \in D$, $\mathbf{w}$ does not envy the bundle of $\mathbf{v}$:
$$u_{\mathbf{w}} \geq \sum_{i \in [n]} w_i \cdot x_{\mathbf{v},i} - \left(\sum_{i \in [n]} v_i \cdot x_{\mathbf{v},i} - u_{\mathbf{v}} \right)$$
$$= u_{\mathbf{v}} + \sum_{i \in [n]} (w_i - v_i) \cdot x_{\mathbf{v},i}.$$

We refer to this IP as the standard IP (utility version) for $\mathrm{DRev}(\mathcal{F})$.

On the other hand, the optimal revenue $\mathrm{Rev}(\mathcal{F})$ obtainable by a lottery-pricing is captured by the same objective function and linear constraints, except that $x_{\mathbf{v},i}$ takes values in $[0,1]$ instead of $\{0,1\}$. We refer to this linear program as the standard LP for $\mathrm{Rev}(\mathcal{F})$.

3 IID with Support Size 2

We establish Theorem 1.1 in this section.

Let $\mathcal{F}_1, \ldots, \mathcal{F}_n$ be i.i.d. distributions with support size 2. Without loss of generality we can assume that the support is $\{1, b\}$ with $b > 1$. (If the support is $\{0, b\}$ the problem is trivial: the optimal revenue can be achieved by offering every item at price b; if the support is $\{a, b\}$ with $0 < a < b$, then we can equivalently rescale it to $\{1, b/a\}$.) Let $p \in (0, 1)$ be the probability that an item takes value b, and $1 - p$ that it takes value 1. We let $\mathcal{P}_i$ denote the probability of $\mathbf{v} \sim \mathcal{F}$ having i items at value b and $n - i$ at 1, for each $i \in [0 : n]$. That is,
$$\mathcal{P}_i = \binom{n}{i} \cdot p^i \cdot (1 - p)^{n-i}.$$

The following lemma for $\mathcal{P}_i$'s is crucial. Its proof can be found in the full version of the paper [CMPY17].

LEMMA 3.1. *There exists a $k \in [0 : n]$ such that*
$$(3.1) \qquad (n - i)\mathcal{P}_i - (b - 1)(\mathcal{P}_{i+1} + \cdots + \mathcal{P}_n)$$

is negative for all $i : 0 \leq i < k$ and is nonnegative for all $i : k \leq i \leq n$.

Let $k \in [0 : n]$ be an integer that satisfies Lemma 3.1, which is unique and can be computed in polynomial time. We write S^* to denote the following discounted item-pricing:

> The grand bundle $[n]$ is offered at $kb + n - k$ and each item is offered individually at b (the latter means that the buyer can buy any bundle $T \subseteq [n]$ at price $|T|b$).

Given S^*, the behavior of the buyer is as follows. If a valuation vector has k or more items at b then the buyer buys the grand bundle at $kb + n - k$; otherwise it buys all the items that have value b. The expected revenue R^* of the discounted item-pricing S^* is then
$$R^* = \sum_{1 \leq i < k} bi \cdot \mathcal{P}_i + (kb + n - k) \sum_{k \leq i \leq n} \mathcal{P}_i.$$

It is clear that given k, R^* can be computed in polynomial time.

To finish the proof of Theorem 1.1, we show that S^* achieves the optimal revenue $\mathrm{Rev}(\mathcal{F})$.

LEMMA 3.2. $R^* = \mathsf{Rev}(\mathcal{F})$ when $\mathcal{F}_1, \ldots, \mathcal{F}_n$ are i.i.d. with support size 2 and k satisfies Lemma 3.1.

We start with some preparation. First recall that when distributions are i.i.d., Daskalakis and Weinberg [DW12] showed that there always exists an optimal solution to the standard LP for $\mathsf{Rev}(\mathcal{F})$ (we use the price version in this section) that is "symmetric": For any permutation σ over $[n]$ with $\sigma(\mathbf{v}) = \mathbf{w}$ (i.e. $v_{\sigma(i)} = w_i$ for all $i \in [n]$), we always have $\sigma(\mathbf{x_v}) = \mathbf{x_w}$ and $\pi_{\mathbf{v}} = \pi_{\mathbf{w}}$, i.e., the lotteries bought at $\mathbf{v}$ and $\mathbf{w}$ are the same under the permutation σ. Based on that, one can significantly reduce the number of variables for the i.i.d. support-size 2 case, and we refer to the new LP described below as the symmetric LP for $\mathsf{Rev}(\mathcal{F})$.

The symmetric LP has $3n + 1$ variables: x_i, for $i = 1, \ldots, n$, is the probability of getting an item with value b in $\mathbf{v}$ when the valuation $\mathbf{v}$ has i items at b (and $n - i$ items at 1); y_i, for $i = 0, 1, \ldots, n - 1$, is the probability of getting an item with value 1 when the valuation has i items at b; finally, π_i for $i = 0, 1, \ldots, n$ is the price of the lottery for a valuation with i items at b. The symmetric LP maximizes the expected revenue: $\sum_{i=0}^{n} \pi_i \cdot \mathcal{P}_i$ subject to the same constraints of the standard LP after replacing $\pi_{\mathbf{v}}$ by π_ℓ when $\mathbf{v}$ has ℓ items at b and $x_{\mathbf{v},i}$ by x_ℓ if $v_i = b$ and by y_ℓ if $v_i = 1$. It is not hard to see that the number of distinct constraints left after the replacement is polynomial in n and thus, the symmetric LP can be solved exactly in polynomial time. By [DW12], the optimal value of the symmetric LP is $\mathsf{Rev}(\mathcal{F})$.

We are now ready to prove Lemma 3.2.

Proof. [Proof of Lemma 3.2] Since R^* is the expected revenue of S^*, it suffices to show $\mathsf{Rev}(\mathcal{F}) \le R^*$.

For this purpose we will relax the symmetric LP for $\mathsf{Rev}(\mathcal{F})$ and show that its optimal value is at most R^*. In the relaxed LP we only keep the following constraints of the symmetric LP:

1. $0 \le x_i \le 1$ for each $i \in [n]$ and $0 \le y_i \le 1$ for each $i \in [0 : n - 1]$.

2. $\pi_0 \le n y_0$ (i.e., the utility at the all-1 vector is nonnegative);

3. For each $i \in [n]$, the constraint that the valuation $\mathbf{w}$ with $w_j = b$ for $j \in [i]$ and $w_j = 1$ for $j > i$ does not envy the lottery of $\mathbf{v}$ with $v_j = b$ for $j \in [i - 1]$ and $v_j = 1$ for $j > i - 1$:

$$bix_i + (n - i)y_i - \pi_i$$
$$(3.2) \qquad \ge b(i - 1)x_{i-1} + (n - i + b)y_{i-1} - \pi_{i-1}.$$

Note that when $i = 1$, x_0 appears on the RHS with coefficient 0; when $i = n$, y_n appears on the

LHS with coefficient 0. For convenience we introduce $x_0 = y_n = 0$ as dummy variables that never appear in the relaxed LP but help simplify the presentation of these constraints.

Since all of them are part of the symmetric LP, the optimal value of the relaxed LP is at least $\mathsf{Rev}(\mathcal{F})$. In the rest of the proof we show that the optimal value of the relaxed LP is at most R^*.

To this end we use the constraints above to upper-bound each π_i using x and y variables. For $i = 0$ we use $\pi_0 \le n y_0$. For each $j \in [n]$ we have from constraints (2) in the relaxed LP that:

$$\pi_j \le \pi_{j-1} + bjx_j + (n - j)y_j$$
$$- b(j - 1)x_{j-1} - (n - j + b)y_{j-1}.$$

Summing these inequalities for all $j = 1, \ldots, i$, we get after some cancellations:

$$\pi_i \le \pi_0 + bix_i + (n - i)y_i$$
$$- (b - 1)(y_{i-1} + y_{i-2} + \cdots + y_1) - (n + b - 1)y_0.$$

Plugging in $\pi_0 \le n y_0$, we have for each $i \in [n]$:

$$\pi_i \le bix_i + (n - i)y_i - (b - 1)(y_{i-1} + \cdots + y_1 + y_0).$$

Replacing in the objective function $\sum_i \mathcal{P}_i \cdot \pi_i$ each π_i by its upper bound, we get a linear form in the x_i's, $i \in [n]$, and y_i's, $i \in [0 : n - 1]$, which upperbounds the value of the relaxed LP (note that x_0 and y_n are dummy variables that do not really appear in any constraint). For each $i \in [n]$, the coefficient of x_i is $bi \cdot \mathcal{P}_i$, thus this term is maximized if we set $x_i = 1$ for each $i \in [n]$. The coefficient of y_0 is $n\mathcal{P}_0 - (b - 1)(\mathcal{P}_1 + \cdots + \mathcal{P}_n)$ and the coefficient of y_i for each $i \in [n - 1]$ is $(n - i)\mathcal{P}_i - (b - 1)(\mathcal{P}_{i+1} + \cdots + \mathcal{P}_n)$. From the choice of $k \in [0 : n]$ and Lemma 3.1, we have for all $i \in [0 : n - 1]$: The coefficient of y_i is negative if $i < k$, and is nonnegative if $i \ge k$. Therefore, the linear form is maximized when we set $y_i = 0$ for all $i < k$ and $y_i = 1$ for all $i \ge k$. Applying these substitutions in the linear form, the upper bound on the value of the LP becomes (note that $(n - i)\mathcal{P}_i - (b - 1)(\mathcal{P}_{i+1} + \cdots + \mathcal{P}_n)$ is 0 when $i = n$):

$$\sum_{i=1}^{n} bi \cdot \mathcal{P}_i + \sum_{i=k}^{n-1} \left[(n - i)\mathcal{P}_i - (b - 1)(\mathcal{P}_{i+1} + \cdots + \mathcal{P}_n) \right]$$
$$= \sum_{i=1}^{n} bi \cdot \mathcal{P}_i + \sum_{i=k}^{n} \left[(n - i)\mathcal{P}_i - (b - 1)(\mathcal{P}_{i+1} + \cdots + \mathcal{P}_n) \right]$$
$$= \sum_{i=1}^{n} bi \cdot \mathcal{P}_i + \sum_{i=k}^{n} \mathcal{P}_i \cdot \left[(n - i) - (b - 1)(i - k) \right] = R^*.$$

This finishes the proof of the lemma. $\square$

4 Hardness of Revenue-Optimal Deterministic Mechanism Design

We prove Theorem 1.2 in this section. The plan is to reduce from the following #P-hard decision problem called COMP introduced in [CDO+15]. The input consists of three parts: 1) a set B of n nonnegative integers $B = \{b_1, \ldots, b_n\}$ between 0 and 2^n (we assume without loss of generality that $b_1 \leq \cdots \leq b_n$); 2) a subset $W \subset [n]$ of size $|W| = n/2$ (assume without loss of generality that n is even) and we use w to denote $\sum_{i \in W} b_i$; and 3) an integer t. The question is to decide whether the number of $S \subset [n]$ such that $|S| = n/2$ and $\sum_{i \in S} b_i \geq w$ is at least t or at most $t - 1$. While COMP was shown to be #P-hard in [CDO+15], we need here the same problem with the following two extra conditions on the two input sets B and W, which we will refer to as COMP*:

1. Every $(n/2)$-subset $S \subset [n]$ with $b_n \in S$ satisfies $\sum_{i \in S} b_i \geq w$, i.e. $b_1 + \cdots + b_{n/2-1} + b_n \geq w$.

2. Every $(n/2)$-subset $S \subset [n]$ that does not contain b_n but contains either b_1 or b_2 must satisfy $\sum_{i \in S} b_i < w$, i.e. $b_2 + b_{n/2+1} + \cdots + b_{n-1} < w$.

These conditions will come in handy in the reduction below. The #P-hardness proof of the problem COMP* can be found in the full version [CMPY17].

4.1 The Reduction.
We now present the reduction from COMP* to the optimal bundle-pricing problem.

Given an input instance (B, W, t) of COMP* (with B and W satisfying the extra conditions) we define an instance $\mathcal{F}$ of optimal bundle-pricing with $n + 1$ items and support size 2. We shall refer to the first n items as item i for $i \in [n]$ and refer to the last item as the *special* item. We use the following parameters:

$$h = 2^{2n}, \quad p = \frac{1}{2(h+1)}, \quad \delta = \frac{1}{2^{3n}},$$

$$a_i = b_i \delta \quad \text{and} \quad h_i = h + a_i, \quad \text{for each } i \in [n]$$

(so $a_i \in [0, 2^n \delta]$ and is an integer multiple of δ). Then item i is supported on $\{1, h_i + 1\}$. The probability of $h_i + 1$ is p and the probability of 1 is $1 - p$. Let

$$c = w\delta, \quad \alpha = \frac{nh}{2} + c, \quad \sigma = \frac{1}{p^n} \quad \text{and} \quad \tau = \frac{\sigma}{\sigma + \alpha}$$

(hence $c \in [0, (n/2)2^n \delta]$ and is an integer multiple of δ). The special item is supported on $\{\sigma, \sigma + \alpha\}$. The probability of $\sigma + \alpha$ is $\tau - \epsilon$ and the probability of σ is $1 - \tau + \epsilon = (\alpha/(\sigma + \alpha)) + \epsilon$ for some ϵ (which is not necessarily positive) to be specified at the end of the proof; for now we only require that $|\epsilon| = o(1/\sigma)$. Note

that since $\sigma \gg \alpha$, the probability $\tau - \epsilon$ of $\sigma + \alpha$ is very close to 1, and the probability $1 - \tau + \epsilon$ of σ is positive but very close to 0. This finishes the description of the bundle pricing instance $\mathcal{F}$ (except the choice of the parameter ϵ which we will set at the end).

4.2 Plan of the Proof.
Our plan for the proof is the following. In Section 4.3, we introduce some notation and define two simple bundle-pricings (Solution 1 and 2), as feasible solutions to the standard IP for $\mathrm{DRev}(\mathcal{F})$, both of which are discounted item-pricings. Most of the work lies in Section 4.4, where we show that whenever $|\epsilon| = o(1/\sigma)$, one of these two solutions is the *unique* optimal solution to the standard IP and achieves $\mathrm{DRev}(\mathcal{F})$. This is done by relaxing the standard IP and showing that one of these two solutions is the unique optimal solution to the relaxed IP. As they are both feasible to the standard IP, we conclude that one of them is uniquely optimal for the standard IP.

Finally, we set ϵ carefully (with $|\epsilon| = o(1/\sigma)$ as promised) in Section 4.5 to show that Solution 2 is strictly better if the (B, W, t) used in the construction of $\mathcal{F}$ is a yes-instance of problem COMP*, and Solution 1 is strictly better if it is a no-instance. This finishes the proof of Theorem 1.2.

4.3 Notation and two simple solutions.
For convenience, we will use a subset $S \subseteq [n]$ to denote a valuation vector over the first n items, where $i \in S$ (or $i \notin S$) means that item i takes the high value $h_i + 1$ (or low value 1). We will also use $(S, \sigma + \alpha)$ (or (S, σ)) to denote a full valuation vector over all the $n + 1$ items in which the special item has the high value $\sigma + \alpha$ (or low value σ).

Given $S \subseteq [n]$, we write $\mathrm{Pr}[S]$ to denote $p^{|S|}(1 - p)^{n-|S|}$; given an integer $i \in [0 : n]$, we write

$$\mathrm{Pr}[i] = \binom{n}{i} \cdot p^i (1-p)^{n-i},$$

use $\mathrm{Pr}[i \geq k]$ to denote $\sum_{i=k}^{n} \mathrm{Pr}[i]$, and $\mathrm{Pr}[i > k]$ to denote $\sum_{i=k+1}^{n} \mathrm{Pr}[i]$. We write $\mathrm{Pr}[S, \sigma]$ and $\mathrm{Pr}[S, \sigma + \alpha]$ to denote the probabilities of (S, σ) and $(S, \sigma + \alpha)$:

$$\mathrm{Pr}[S, \sigma] = \mathrm{Pr}[S] \cdot (1 - \tau + \epsilon) \quad \text{and}$$

$$\mathrm{Pr}[S, \sigma + \alpha] = \mathrm{Pr}[S] \cdot (\tau - \epsilon).$$

We use the standard IP for $\mathrm{DRev}(\mathcal{F})$ (the utility version) but rename the variables as follows. For each $S \subseteq [n]$ we use $x_{S,i} \in \{0, 1\}$ to denote the variable for item i in valuation (S, σ) for each $i \in [n]$, $z_S \in \{0, 1\}$ to denote the variable for the special item, and $u_S \geq 0$ to denote the utility. For each $S \subseteq [n]$, we use $x'_{S,i} \in \{0, 1\}$ to denote the variable for item i in valuation $(S, \sigma + \alpha)$,

$z'_S \in \{0,1\}$ to denote the variable for the special item, and $u'_S \geq 0$ to denote the utility.

The two simple bundle-pricings we are interested in are the following:

Solution 1: Offer the grand bundle at $\sigma + n$, or (equivalently) offer each item at its low value.

Solution 2: The discounted item-pricing where the grand bundle is offered at $\sigma + \alpha + n$ and each individual item is offered at its high value. As discussed in the introduction, this means that the buyer can buy either the grand bundle at $\sigma + \alpha + n$ or any bundle of items at the sum of their high values. When this menu is offered, the buyer buys the grand bundle if its utility is positive (since her utility from buying items priced at their high values can never be positive) and buys the bundle of items at their high values (if any) if the utility from the grand bundle is negative. For the case when the utility from the grand bundle is 0, the buyer still gets the grand bundle since it always gives a higher revenue.

These two bundle-pricings induce two feasible solutions to the standard IP:

1. In Solution 1, every valuation (S, σ) and $(S, \sigma + \alpha)$ buys the grand bundle. So we have for each $S \subseteq [n]$ and $i \in [n]$, $x_{S,i} = z_S = x'_{S,i} = z'_S = 1$. Regarding the utilities we have

$$u_S = \sum_{i \in S} h_i \quad \text{and} \quad u'_S = \alpha + \sum_{i \in S} h_i.$$

2. In Solution 2 we have 1) for each $S \subseteq [n]$ and $i \in [n]$, $x'_{S,i} = z'_S = 1$ and $u'_S = \sum_{i \in S} h_i$; 2) for each $S \subseteq [n]$ with $\sum_{i \in S} h_i \geq \alpha$, we have $x_{S,i} = z_S = 1$ for all $i \in [n]$, $u_S = \sum_{i \in S} h_i - \alpha$; 3) for each $S \subseteq [n]$ with $\sum_{i \in S} h_i < \alpha$, we have $x_{S,i} = 1$ for each $i \in S$, $x_{S,i} = 0$ for each $i \notin S$, $z_S = 0$, and $u_S = 0$. Given our choice of parameters (i.e. $h \gg a_i, c$), every S with $|S| > n/2$ satisfies case 2) and every S with $|S| < n/2$ satisfies case 3). A set S with $|S| = n/2$ satisfies case 2) if we have $\sum_{i \in S} a_i \geq c$ (equivalently, $\sum_{i \in S} b_i \geq w$) and satisfies case 3) otherwise. This is the connection with COMP* that we will explore in the reduction.

As discussed in the plan, we introduce a relaxation of the standard IP that only contains a subset of its constraints and refer to it as the *relaxed* IP. It contains the following constraints:

1. $x_{S,i}, z_S, x'_{S,i}, z'_S \in \{0,1\}$ and $u_S, u'_S \geq 0$, for all $S \subseteq [n]$ and $i \in [n]$;

2. For each $S \neq \emptyset$, $(S, \sigma + \alpha)$ does not envy $(\emptyset, \sigma + \alpha)$:
$$u'_S \geq u'_\emptyset + \sum_{i \in S} h_i \cdot x'_{\emptyset,i}.$$

3. For each $S \subseteq [n]$, $(\emptyset, \sigma + \alpha)$ does not envy (S, σ):
$$u'_\emptyset \geq u_S - \sum_{i \in S} h_i \cdot x_{S,i} + \alpha \cdot z_S.$$

4. For each $S \subseteq [n]$, (S, σ) does not envy $(\emptyset, \sigma + \alpha)$:
$$u_S \geq u'_\emptyset + \sum_{i \in S} h_i \cdot x'_{\emptyset,i} - \alpha \cdot z'_\emptyset.$$

5. For each pair of T, S with $T \subset S \subseteq [n]$, (S, σ) does not envy (T, σ):
$$u_S \geq u_T + \sum_{i \in S \setminus T} h_i \cdot x_{T,i}.$$

The objective function of the relaxed IP is the same expected revenue, which is the sum of the following over all $S \subseteq [n]$:

$$\left(\sum_{i \in S} (h_i + 1) x'_{S,i} + \sum_{i \notin S} x'_{S,i} + (\sigma + \alpha) z'_S - u'_S \right) \Pr[S, \sigma + \alpha]$$

$$+ \left(\sum_{i \in S} (h_i + 1) x_{S,i} + \sum_{i \notin S} x_{S,i} + \sigma z_S - u_S \right) \Pr[S, \sigma].$$

4.4 One of the two simple solutions is optimal. Our goal is the following lemma about optimal solutions to the relaxed IP:

LEMMA 4.1. *In an optimal solution to the relaxed IP, if $z_\emptyset = 1$, then it must be Solution 1; if $z_\emptyset = 0$, then it must be Solution 2.*

As both solutions are feasible for the standard IP, we have the following corollary:

COROLLARY 4.1. *Either Solution 1 or 2 is the unique optimal solution to the standard IP.*

We start the proof of Lemma 4.1 with a few simple observations.

LEMMA 4.2. *In any optimal solution to the relaxed IP, we have $x_{S,i} = 1$ for all $S \subseteq [n]$ and $i \in S$, $x'_{S,i} = 1$ for all $S \neq \emptyset$ and $i \in [n]$, and $z'_S = 1$ for all $S \subseteq [n]$.*

Proof. The part of $x'_{S,i} = 1$ for all $S \neq \emptyset$ and $i \in [n]$ and $z'_S = 1$ for all $S \neq \emptyset$ is trivial as they do not appear in the constraints of the relaxed IP but appear with a positive coefficient in the objective function. For $z'_\emptyset$, note that it only appears on the right hand side with a negative coefficient. Thus, if $z'_\emptyset = 0$ in a feasible solution, we can switch it to 1 and the new solution remains feasible but the expected revenue goes up strictly. The same argument works for $x_{S,i}$ for all $S \subseteq [n]$ and $i \in S$. $\square$

Next we show that $x'_{\emptyset,i} = 1$ for all $i \in [n]$ in any optimal solution to the relaxed IP.

Lemma 4.3. *In an optimal solution to the relaxed IP, $x'_{\emptyset,i} = 1$ for all $i \in [n]$ and $u'_S = u'_\emptyset + \sum_{i \in S} h_i$ for all nonempty $S \subseteq [n]$.*

Proof. Assume for contradiction that $x'_{\emptyset,i} = 0$ for $i \in I$ and I is nonempty. Then we make the following changes to obtain a new solution:

i) We change $x'_{\emptyset,i}$ from 0 to 1 for each $i \in I$;

ii) For each $S \neq \emptyset$, we increase u'_S by $\sum_{i \in S \cap I} h_i$;

iii) For each $S \subseteq [n]$, we replace u_S by the maximum of the original u_S and $u'_\emptyset + \sum_{i \in S} h_i - \alpha$ (by doing this the new u_S can go up by at most $\sum_{i \in S \cap I} h_i$ because of constraints (4)).

We first verify that the new solution remains feasible and then show that its expected revenue is strictly better than that of the original solution. For its feasibility, constraints in (2) and (4) are trivial. (3) for S holds trivially if u_S remains the same. Otherwise $u_S = u'_\emptyset + \sum_{i \in S} h_i - \alpha$ and thus,

$$u_S - \sum_{i \in S} h_i \cdot x_{S,i} + \alpha \cdot z_S = u_S - \sum_{i \in S} h_i + \alpha \cdot z_S$$

$$\leq u_S - \sum_{i \in S} h_i + \alpha = u'_\emptyset.$$

For 5), the constraint for $T \subset S$ holds trivially if u_T remains the same. If u_T goes up, we have

$$u_S \geq u'_\emptyset + \sum_{i \in S} h_i - \alpha \quad \text{and}$$

$$u_T = u'_\emptyset + \sum_{i \in T} h_i - \alpha$$

and thus, $u_S - u_T \geq \sum_{i \in S \setminus T} h_i \geq \sum_{i \in S \setminus T} h_i \cdot x_{T,i}$. This shows that the new solution is also feasible.

Finally, comparing the two solutions, the net gain of expected revenue in the new one is at least

$$|I| \cdot \Pr[\emptyset, \sigma + \alpha] - \sum_{S \subseteq [n]} \big(\Pr[S, \sigma + \alpha] + \Pr[S, \sigma] \big) \sum_{i \in S \cap I} h_i$$

$$= |I| \cdot (\tau - \epsilon) \cdot (1 - p)^n - \sum_{i \in I} p h_i > 0$$

as by our choice of parameters the first term is close to $|I|$ and the second term is close to $|I|/2$.

The second part of the lemma follows trivially since u'_S appears only in the LHS of constraints (2) when $S \neq \emptyset$ and appears in the objective function with a negative coefficient. $\square$

So far we have shown in an optimal solution to the relaxed IP that every entry in $\mathbf{x}'$ and $\mathbf{z}'$ is 1 and the only unsettled variable in $\mathbf{u}'$ is $u'_\emptyset$, with $u'_S = u'_\emptyset + \sum_{i \in S} h_i$ for all $S \neq \emptyset$. (For a sanity check, note that these conditions hold in Solution 1 and 2.) Next we prove the first case of Lemma 4.1.

Lemma 4.4. *In an optimal solution to the relaxed IP, if $z_\emptyset = 1$, then it must be Solution 1.*

Proof. Assuming $z_\emptyset = 1$, we have $u'_\emptyset \geq u_\emptyset + \alpha$ from constraints (3). Next we show that $x_{\emptyset,i} = 1$ for all $i \in [n]$. Assume for contradiction that $x_{\emptyset,i} = 0$ for $i \in I$ and I is nonempty. Then we make the following changes to obtain a new solution: i) We change $x_{\emptyset,i}$ from 0 to 1 for every $i \in I$; and ii) For each $S \neq \emptyset$, we change u_S to be the maximum of the original u_S and $u_\emptyset + \sum_{i \in S} h_i$ (by doing this u_S can go up by at most $\sum_{i \in S \cap I} h_i$ due to constraints (5) for $T = \emptyset$).

Following the same argument used in the previous lemma, we verify that the new solution is feasible and then show that its expected revenue is strictly higher; the details can be found in the full version [CMPY17]. This allows us to conclude that $x_{\emptyset,i} = 1$ for all $i \in [n]$.

Finally, given that $x_{\emptyset,i} = 1$ for all $i \in [n]$ and $u_\emptyset \geq 0$, we have $u_S \geq \sum_{i \in S} h_i$, $u'_\emptyset \geq \alpha$ and

$$u'_S = u'_\emptyset + \sum_{i \in S} h_i \geq \alpha + \sum_{i \in S} h_i.$$

Comparing it with Solution 1, a) the utility of every valuation of this solution is at least as large as that of Solution 1; and b) every allocation variable $x_{S,i}$, $x'_{S,i}$, z_S and z'_S in Solution 1 is 1 (the maximum possible). Thus, its revenue is no more than that of Solution 1 and in order to be optimal, it must be exactly the same as Solution 1. This finishes the proof of the lemma. $\square$

In the rest of this subsection, we consider the more challenging case when $z_\emptyset = 0$, and prove the second part of Lemma 4.1 using a sequence of lemmas.

First we show that $u'_\emptyset = 0$ when $z_\emptyset = 0$.

Lemma 4.5. *In any optimal solution to the relaxed IP, if $z_\emptyset = 0$ then we have $u'_\emptyset = 0$.*

Proof. Assume for contradiction that $u'_\emptyset > 0$. We show first that $u'_\emptyset > 0$ implies that $u'_\emptyset \geq \delta$.

For this purpose we fix all the allocation variables to their $\{0, 1\}$ values in the optimal solution and replace each u'_S, $S \neq \emptyset$, using $u'_S = u'_\emptyset + \sum_{i \in S} h_i$ (by Lemma 4.4) to obtain a linear program over the rest of utility variables u_S, $S \subseteq [n]$, and $u'_\emptyset$. The linear program maximizes the objective function (hence these utility variables must have the minimum possible values),

subject to the constraints (1), (3), (4), (5) of the relaxed IP (note that constraints (2) are already satisfied). All the constraints (3), (4), (5) of the relaxed IP have the form $v \geq v' + b$ where v, v' are two variables from the set $V = \{u_S : S \subseteq [n]\} \cup \{u'_\emptyset\}$ and b has the form $c_1 h + c_2 \delta$, where c_1, c_2 are integers and $|c_2| \leq O(n2^n)$ (note that all h_i have this form with $c_2 \in [0, 2^n]$, α also has this form with $c_2 \in [0, (n/2)2^n]$, and the allocation variables have value 0 or 1). Recall that $h = 2^{2n}$ and $\delta = 1/2^{3n}$. The problem of finding the optimal utilities becomes a shortest-path problem over $2^n + 2$ vertices: the set V plus an extra vertex 0 as the source vertex. The distance from 0 to each vertex in V is set to 0 and the distance from v' to v is set to $-b$ if $v \geq v' + b$ is a constraint in the relaxed IP. The optimal value of $u'_\emptyset$ is then the length of the shortest path from vertex 0 to vertex $u'_\emptyset$ with the sign changed. Hence $u'_\emptyset$ is a sum of at most $2^n + 1$ edge weights $(-b)$, with each of them being an integer multiple c_1 of h plus an integer multiple c_2 of δ (with $|c_2| \leq O(n2^n)$). As a result, $u'_\emptyset$ is of the form $d_1 h + d_2 \delta$ where d_1 and d_2 are both integers and $|d_2| = O(n2^{2n})$. For $u'_\emptyset$ to be positive, we have either $d_1 > 0$, in which case $u'_\emptyset = \Omega(h)$ (as $\delta = 1/2^{3n}$), or $d_1 = 0$ and $d_2 > 0$, in which case $u'_\emptyset \geq \delta$.

Now we modify the given solution to obtain a new solution. Let $\mathcal{D}$ denote the set of all subsets $S \subset [n]$ with $\sum_{i \in S} h_i < \alpha$, and $\mathcal{U}$ denote the set of all $S \subseteq [n]$ with $\sum_{i \in S} h_i \geq \alpha$. We modify the solution as follows:

i) We set $x_{S,i} = 0$ and $z_S = 0$ for all $S \in \mathcal{D}$ and $i \notin S$;

ii) We change $u'_\emptyset$ to 0 and u'_S to $\sum_{i \in S} h_i$;

iii) For each $S \in \mathcal{D}$ we change u_S to 0; For each $S \in \mathcal{U}$ we change u_S to $\sum_{i \in S} h_i - \alpha$ (which is nonnegative by the definition of $\mathcal{U}$ and can only go down compared to the old solution by constraints (4) using $x'_{\emptyset,i} = 1$ from Lemma 4.3 and $z'_\emptyset = 1$ from Lemma 4.2).

Every other entry remains the same as in the original solution. As before, we first show that the new solution is feasible and then show that its expected revenue is strictly higher. Details of the analysis can be found in the full version [CMPY17]. This finishes the proof of the lemma. $\square$

Let $\mathcal{D}$ and $\mathcal{U}$ be the sets defined in the proof above. We have the following simple corollary:

COROLLARY 4.2. *In any optimal solution to the relaxed IP, if $z_\emptyset = 0$ then $z_S = 0$ for all $S \in \mathcal{D}$.*

Proof. By constraints (3), we have for $S \in \mathcal{D}$:

$$0 \geq u_S - \sum_{i \in S} h_i + \alpha \cdot z_S.$$

As a result, $\alpha \cdot z_S \leq \sum_{i \in S} h_i < \alpha$.

Next we show that if $z_\emptyset = 0$ in an optimal solution, then $x_{S,i} = 0$ for all $S \in \mathcal{D}$ and $i \notin S$.

LEMMA 4.6. *In an optimal solution, if $z_\emptyset = 0$ then $x_{S,i} = 0$ for all $S \in \mathcal{D}$ and $i \notin S$.*

Proof. Assume for contradiction that we do not have $x_{T,i} = 0$ for all sets $T \in \mathcal{D}$ and $i \notin T$. (For reasons that will become clear later we need special treatments for not only those T at level $n/2$ but those at level $(n/2) - 1$ as well.) We use $\mathcal{G}$ to denote the set of pairs (T, k) such that $|T| < (n/2) - 1$, $k \notin T$ and $x_{T,k} = 1$; we use $\mathcal{G}^*$ to denote the set of pairs (T, k) such that $|T| = (n/2) - 1$, $k \notin T$, $T \cup \{k\} \in \mathcal{U}$ and $x_{T,k} = 1$; we use $\mathcal{G}_1^\dagger$ to denote the set of pairs (T, k) such that $|T| = (n/2) - 1$, $k \notin T$, $T \cup \{k\} \in \mathcal{D}$ (so it follows from the first condition that B and W satisfy that $n \notin T \cup \{k\}$) and $x_{T,k} = 1$; we use $\mathcal{G}_2^\dagger$ to denote the set of pairs (T, k) such that $|T| = n/2$, $T \in \mathcal{D}$, $k \notin T$ and $x_{T,k} = 1$. Then $\mathcal{G} \cup \mathcal{G}^* \cup \mathcal{G}_1^\dagger \cup \mathcal{G}_2^\dagger$ is nonempty. We next use $\mathcal{G}, \mathcal{G}^*, \mathcal{G}_1^\dagger$ and $\mathcal{G}_2^\dagger$ to define two sets $\mathcal{E}, \mathcal{E}^\dagger$, where $\mathcal{E}$ consists of subsets of size $n/2$ and $\mathcal{E}^\dagger$ consists of subsets of size $n/2 + 1$ so they are disjoint.

1. $S \subset [n]$ is in $\mathcal{E}$ if $|S| = n/2$, $S \in \mathcal{U}$ and $T \cup \{k\} \subseteq S$ for some pair $(T, k) \in \mathcal{G} \cup \mathcal{G}^*$.

2. $S \subset [n]$ is in $\mathcal{E}^\dagger$ if $|S| = (n/2) + 1$ and satisfies either $S = T \cup \{k\}$ for some pair $(T, k) \in \mathcal{G}_2^\dagger$ or $S = T \cup \{k, r\}$ for some pair $(T, k) \in \mathcal{G}_1^\dagger$ with r being the smallest index not in $T \cup \{k\}$.

We need the following simple claim about z_S, $S \in \mathcal{E} \cup \mathcal{E}^\dagger$, in the original solution. The proof can be found in the full version [CMPY17].

CLAIM 4.1. *We have $z_S = 0$ for every $S \in \mathcal{E} \cup \mathcal{E}^\dagger$ in the original solution.*

Claim 4.1 inspires us to derive a new solution by making the following changes in the old one:

i) For all $(T, k) \in \mathcal{G} \cup \mathcal{G}^* \cup \mathcal{G}_1^\dagger \cup \mathcal{G}_2^\dagger$, change $x_{T,k}$ from 1 to 0 (so that $x_{T,k} = 0$ in the new solution for all $T \in \mathcal{D}$ and $k \notin T$);

ii) For all $S \in \mathcal{E} \cup \mathcal{E}^\dagger$, change z_S from 0 to 1;

iii) For all $S \in \mathcal{D}$, change u_S to 0; For all $S \in \mathcal{U}$, change u_S to $\sum_{i \in S} h_i - \alpha$ (note that the new u_S is nonnegative as $S \in \mathcal{U}$ and can only go down from the original by constraint (4)).

All other entries remain the same in the new solution.

We first verify that the new solution is feasible and then show that it is strictly better than the old solution.

For the feasibility, constraints (2) are trivial. For (3), the constraint is trivial if $S \in \mathcal{D}$ (since $u_S = z_S = 0$); if $S \in \mathcal{U}$, we have $u_S = \sum_{i \in S} h_i - \alpha$ and thus, the RHS is at most 0. For (4), the constraint is trivial if $S \in \mathcal{D}$ (since the RHS is negative); if $S \in \mathcal{U}$, the LHS and RHS are the same. For (5), the constraint is trivial if $T \in \mathcal{D}$ (since the RHS is 0); otherwise we have both S and T are in $\mathcal{U}$ and the constraint follows from $u_S = \sum_{i \in S} h_i - \alpha$ and $u_T = \sum_{i \in T} h_i - \alpha$ in the new solution. This finishes the proof of feasibility of the new solution.

Note that each utility variable in the new solution can only go down from that in the old solution. As a result, the net gain of expected revenue in the new solution is at least

$$\sum_{S \in \mathcal{E} \cup \mathcal{E}^\dagger} \sigma \cdot \Pr[S, \sigma] - \sum_{(T,k) \in \mathcal{G} \cup \mathcal{G}^* \cup \mathcal{G}_1^\dagger \cup \mathcal{G}_2^\dagger} \Pr[T, \sigma].$$

Ignoring the common factor of $1 - \tau + \epsilon$ and rearranging the terms, we obtain

$$\left(\sigma \sum_{S \in \mathcal{E}} \Pr[S] - \sum_{(T,k) \in \mathcal{G} \cup \mathcal{G}^*} \Pr[T] \right)$$

$$+ \left(\sigma \sum_{S \in \mathcal{E}^\dagger} \Pr[S] - \sum_{(T,k) \in \mathcal{G}_1^\dagger \cup \mathcal{G}_2^\dagger} \Pr[T] \right).$$

The rest of the proof can be found in the full version [CMPY17], where we show that the first term is positive if $\mathcal{G} \cup \mathcal{G}^*$ is nonempty and the second term is positive if $\mathcal{G}_1^\dagger \cup \mathcal{G}_2^\dagger$ is nonempty. This shows that the expected revenue goes up strictly in the new solution. $\square$

Finally we prove the second part of Lemma 4.1. Lemma 4.1 follows from Lemma 4.4 and 4.7.

LEMMA 4.7. *In an optimal solution to the relaxed IP, if $z_\emptyset = 0$, then it must be Solution 2.*

Proof. We show that, when $z_\emptyset = 0$, every utility variable in the optimal solution is at least as large as that in Solution 2, and every allocation variable is at most as large as that in Solution 2. So for it to be optimal, it must be exactly the same as Solution 2.

For utilities we first note that u'_S is the same in both solutions for all $S \subseteq [n]$ since in both we have $u'_\emptyset = 0$ and $u'_S = \sum_{i \in S} h_i$. Next for each $S \in \mathcal{D}$, we have $u_S = 0$ in Solution 2. Finally, for each $S \in \mathcal{U}$ we have from constraint (4) that $u_S \geq u'_\emptyset + \sum_{i \in S} h_i - \alpha = \sum_{i \in S} h_i - \alpha$ in the optimal solution, which is at least as large as u_S in Solution 2. For allocation variables, we first have $x'_{S,i} = z'_S = 1$ in both solutions for all $S \subseteq [n]$ and $i \in [n]$. Next we have $x_{S,i} = 1$ in both solutions for

all $S \subseteq [n]$ and $i \in S$. For each $S \in \mathcal{D}$, we have $x_{S,i} = z_S = 0$ in both solutions for all $i \notin S$. Finally for each $S \in \mathcal{U}$, we have $x_{S,i} = z_S = 1$ in Solution 2 for all $i \notin S$. This finishes the proof of the lemma. $\square$

4.5 Finishing the reduction from COMP*.
Finally, we show that with an appropriate choice of ϵ (with $|\epsilon| = o(1/\sigma)$ as promised) that can be computed in polynomial time, we have 1) Solution 2 is strictly better than Solution 1 if (B, W, t) is a yes-instance of COMP*; and 2) Solution 1 is strictly better than Solution 2 if it is a no-instance.

The expected revenue of Solution 1 is $\mathsf{Rev}_1 = n + \sigma$. Let t^* be the number of $(n/2)$-sets $S \subset [n]$ with $\sum_{i \in S} h_i \geq \alpha$ (recall that (B, W, t) is a yes-instance if $t^* \geq t$ and is a no-instance if $t^* \leq t - 1$) and let $\mathcal{R}$ be the set of all sets $S \subset [n]$ with $\sum_{i \in S} h_i < \alpha$. Then Rev_2 of Solution 2 is

$$(n + \sigma + \alpha)\Big((\tau - \epsilon) + (1 - \tau + \epsilon)$$
$$\times \Big(\Pr[i > n/2] + t^* \cdot p^{n/2} \cdot (1 - p)^{n/2}\Big)\Big)$$
$$+ (1 - \tau + \epsilon) \sum_{S \in \mathcal{R}} \Pr[S] \cdot \sum_{i \in S}(h_i + 1).$$

Below we use $a = b \pm c$ to denote $|a - b| \leq c$. The sum $\sum_{S \in \mathcal{R}} \Pr[S] \cdot \sum_{i \in S}(h_i + 1)$ is equal to

$$\sum_{i \in [n]}(h_i + 1)p - \sum_{S \notin \mathcal{R}} \Pr[S] \sum_{i \in S}(h_i + 1)$$
$$= \sum_{i \in [n]}(h_i + 1)p \pm O(2^n p^{n/2} nh).$$

Then $\mathsf{Rev}_2 = B - A\epsilon$, where $A = A' \pm O(\sigma 2^n p^{n/2})$ with

$$A' = (n + \sigma + \alpha)\big(1 - \Pr[i > n/2]\big) - \sum_{i \in [n]}(h_i + 1)p$$

and $A' = \Theta(\sigma)$ can be computed efficiently. On the other hand, we note for B that $\Pr[i > n/2] \leq 2^n p^{n/2+1}$ and $(n + \sigma + \alpha) \cdot (1 - \tau) = O(\alpha)$. As a result, we have

$$B = (n + \sigma + \alpha)\left(\tau + (1 - \tau) \cdot t^* \cdot p^{n/2} \cdot (1 - p)^{n/2}\right)$$
$$+ (1 - \tau) \sum_{i \in [n]}(h_i + 1)p \pm O(\alpha 2^n p^{n/2+1}).$$

For convenience we write B as

$$B = B' + C' \cdot t^* \pm O(\alpha 2^n p^{n/2+1}),$$

where

$$B' = (n + \sigma + \alpha)\tau + (1 - \tau) \sum_{i \in [n]}(h_i + 1)p \quad \text{and}$$

$$C' = (n + \sigma + \alpha) \cdot (1 - \tau) \cdot p^{n/2} \cdot (1 - p)^{n/2}$$

can be computed efficiently. Plugging in $\tau = \sigma/(\sigma + \alpha)$ and $h_i = h + a_i$, we have (after simplification)

$$B' = \sigma + n - \frac{\alpha n}{2(\sigma + \alpha)} + (1 - \tau)p \sum_{i \in [n]} a_i.$$

Finally we choose ϵ to be (recall t is between 1 and 2^n; otherwise the problem is trivial)

$$\epsilon = \frac{1}{A'} \cdot \left(C'\big(t - (1/2)\big) - \frac{\alpha n}{2(\sigma + \alpha)} + (1 - \tau)p \sum_{i \in [n]} a_i \right)$$

which can be computed efficiently and (by $|A'| = \Theta(\sigma)$)

$$|\epsilon| \leq (1/A') \cdot \big(O(\alpha p^{n/2} 2^n) + O(\alpha n/\sigma) \big)$$

$$= O(\alpha p^{n/2} 2^n/\sigma) = o(1/\sigma).$$

Plugging in our choice of ϵ, we finally have

$$\mathsf{Rev}_2 - \mathsf{Rev}_1 = C'(t^* - t + 1/2) \pm O(\alpha 2^n p^{n/2+1}).$$

Note that $C' = \Omega(\alpha p^{n/2}) \gg O(\alpha 2^n p^{n/2+1})$. If $t^* \geq t$, Solution 2 is strictly better than Solution 1; if $t^* \leq t - 1$, Solution 1 is strictly better. This finishes the proof of Theorem 1.2.

5 Constant Number of Items

In this section we prove Theorem 1.3. Let $\mathcal{F} = \mathcal{F}_1 \times \cdots \times \mathcal{F}_k$ be an instance of the bundle-pricing problem for some constant number of items k and assume without loss of generality that $|\mathrm{SUPPORT}(\mathcal{F}_i)| = m$ for all $i \in [k]$. In this case, there are m^k possible valuation vectors (a polynomial number), and $d = 2^k$ possible distinct bundles (a constant number). The standard IP in this case has a polynomial number of variables and constraints. However, Integer Programming is NP-hard, so we will use a different method to solve the problem in polynomial time. For the rest of this section, we assume two arbitrary orderings, one for the valuation vectors and one for the bundles, and we will use v_i to denote the ith valuation vector, and $B(j)$ to denote the jth bundle and p_j to denote its price.

We will argue that we can generate in polynomial time a set of price vectors p that includes an optimal one; we can then compute the expected revenue for each of these vectors and pick the best one. To this end, we consider a partitioning of the d-dimensional space of possible price vectors p into cells, such that for all p in the same cell, the buyer has the same behavior for every v_i, i.e., buys the same bundle, if any. Consider the following set H of hyperplanes over p.

1. For each valuation v_i and bundle $B(j)$, the set H includes $\sum_{\ell \in B(j)} v_{i,\ell} - p_j = 0$. (If the price vector p is below the hyperplane, the buyer will not consider bundle $B(j)$ for valuation v_i.)

2. For each valuation v_i and each pair of bundles $B(j)$ and $B(j')$, the set H includes the hyperplane

$$\sum_{\ell \in B(j)} v_{i,\ell} - p_j = \sum_{\ell \in B(j')} v_{i,\ell} - p_{j'}.$$

Note that for v_i, the buyer prefers $B(j)$ to $B(j')$ if p is on one side of the hyperplane, she prefers $B(j')$ to $B(j)$ if p is on the other side, and if p lies on the hyperplane itself then it depends on the order between the prices $p_j, p_{j'}$.

3. For each pair of bundles $B(j)$ and $B(j')$, the set H includes the hyperplane $p_j = p_{j'}$.

These hyperplanes partition the space of prices into cells, where a cell consists of all price vectors that have the same relation to each of these hyperplanes, i.e., lie in the same open half-space or on the hyperplane. We can assume that for a valuation v_i and price vector p, if there is a tie both in utility and in the price between some bundles, then the buyer selects a bundle according to some fixed tie-breaking rule, for example she chooses among the tied bundles the one with the smallest index (the rule does not matter for the revenue). It follows from the definition of the set H of hyperplanes, that for every cell C and every valuation v_i there is $k_i \in [d]$ such that the buyer selects the same bundle $B(k_i)$ for every price vector p in C, or buys no bundle (if they all have negative utility). Let $V_C(j)$ be the set of valuations for which the buyer selects bundle $B(j)$ if the price vector $p \in C$, and let

$$Q_C(j) = \Pr\big[V_C(j)\big] = \sum_{v_i \in V_C(j)} \Pr[v_i]$$

be the probability that the buyer selects bundle $B(j)$. The supremum revenue that the seller can extract for a price vector p in the cell C, can be computed by solving the LP of maximizing $\sum_j Q_C(j) \cdot p_j$, subject to p belonging to the closure of the cell C, i.e., p satisfying all the weak inequalities corresponding to the bounding hyperplanes of C. By LP theory, the maximum value of the LP is achieved at some vertex; even if the vertex does not belong to C but is in the closure, the corresponding price vector achieves this expected revenue (by the maximum price tie breaking rule). The maximum over all cells C gives the supremum revenue that can be achieved by any price vector. Thus, the supremum revenue is achieved at some vertex, i.e., at the intersection of some d hyperplanes of the set H.

Therefore, we can compute an optimal solution by generating all vertices and picking the best one. For every subset of $d = 2^k$ hyperplanes of H, we solve the corresponding linear system of equations to check if the

hyperplanes intersect at a unique point p, and if p is nonnegative (if a price is negative then p clearly cannot be optimal). If so, compute the expected revenue of p by examining each valuation v_i and determining the bundle selected for v_i, if any. Choose among these price vectors p the one that yields the maximum revenue. Since the set H has a polynomial number of hyperplanes, and the dimension $d = 2^k$ is constant, we only need to consider a polynomial number of subsets to generate the set of price vectors p. Because the number of valuations is also polynomial, it takes polynomial time to compute the expected revenue of each vector p. Hence the total time is polynomial. This finishes the proof of Theorem 1.3.

6 Conclusions

In this work we studied the optimal bundle-pricing problem (or equivalently, the Revenue-Optimal Deterministic Mechanism Design problem). We showed that the problem is intractable (#P-hard) even when the (independent) item distributions have support size 2 and the optimal solution has a very simple form of discounted item-pricing: the seller prices the individual items and offers also the grand bundle at a (possibly) discounted price. Another consequence of the results is that there is no 'nice' (easy-to-check) characterization of when separate item pricing, or grand bundling extracts the maximum revenue $\mathrm{DRev}(\mathcal{F})$ achievable by any bundle-pricing. On the positive side, we showed that for i.i.d. distributions with support size 2, the maximum revenue $\mathrm{Rev}(\mathcal{F})$ achievable by any lottery pricing can always be achieved by a discounted item-pricing, and we can compute it in polynomial time. The problem can be also solved in polynomial time for a constant number of items.

A number of interesting problems present themselves. First, we know from Babaioff et al. [BILW14] that discounted item-pricing always achieves a constant fraction (at least 1/6th) of the maximum revenue; what is the constant that can always be guaranteed with respect to the deterministic and randomized maximum revenue? Second, we know that we can compute efficiently an optimal item pricing, and it can be shown that we can also compute an $(1 - \epsilon)$-approximately optimal grand bundle price; can we compute efficiently an $(1 - \epsilon)$-approximately optimal discounted item pricing? (We believe this is the case.) Third, besides extending simple item-pricing with the grand bundle, it is more generally natural to offer discounts on disjoint groups of items, as in partition mechanisms. How powerful are such partitioned discounted item-pricings, and can we compute efficiently an $(1 - \epsilon)$-approximately optimal solution of this type? Finally, regarding i.i.d. distributions, we know that randomization can increase the

revenue for support size 3 in some cases (an example is given by Hart and Nisan [HN12]). Are simple schemes able to extract (approximately) the maximum revenue $\mathrm{DRev}(\mathcal{F})$ achievable by any bundle pricing for general i.i.d. distributions?

References

[BCKW10] P. Briest, S. Chawla, R. Kleinberg, and S.M. Weinberg, *Pricing randomized allocations*, Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms, 2010, pp. 585–597.

[BILW14] M. Babaioff, N. Immorlica, B. Lucier, and S.M. Weinberg, *A simple and approximately optimal mechanism for an additive buyer*, 55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014, 2014, pp. 21–30.

[Bri08] P. Briest, *Uniform budgets and the envy-free pricing problem*, Proceedings of the 35th International Colloquium on Automata, Languages and Programming, 2008, pp. 808–819.

[CD11] Y. Cai and C. Daskalakis, *Extreme-value theorems for optimal multidimensional pricing*, Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science, 2011, pp. 522–531.

[CDO+15] X. Chen, I. Diakonikolas, A. Orfanou, D. Paparas, X. Sun, and M. Yannakakis, *On the complexity of optimal lottery pricing and randomized mechanisms*, Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Sciences, 2015.

[CDP+14] X. Chen, I. Diakonikolas, D. Paparas, X. Sun, and M. Yannakakis, *The complexity of optimal multidimensional pricing*, Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms, 2014.

[CH13] Y. Cai and Z. Huang, *Simple and nearly optimal multi-item auctions*, Proceedings of the 24th ACM-SIAM Symposium on Discrete Algorithms, 2013, pp. 564–577.

[CHK07] S. Chawla, J.D. Hartline, and R.D. Kleinberg, *Algorithmic pricing via virtual valuations*, Proceedings of the 8th ACM Conference on Electronic Commerce, 2007, pp. 243–251.

[CHMS10] S. Chawla, J.D. Hartline, D.L. Malec, and B. Sivan, *Multi-parameter mechanism design and sequential posted pricing*, Proceedings of the 42nd Annual ACM Symposium on Theory of Computing, 2010, pp. 311–320.

[CMPY17] X. Chen, G. Matikas, D. Paparas, and M. Yannakakis, *On the complexity of bundle-pricing and simple mechanisms*, CoRR **abs/1702.07032** (2017).

[CMS10] S. Chawla, D.L. Malec, and B. Sivan, *The power of randomness in bayesian optimal mechanism design*,

Proceedings of the 11th ACM Conference on
Electronic Commerce, 2010, pp. 149–158.

[DDT12] C. Daskalakis, A. Deckelbaum, and C. Tzamos, *Optimal pricing is hard*, Proceedings of the 8th Workshop on Internet and Network Economics, 2012, pp. 298–308.

[DDT13] _____, *Mechanism design via optimal transport*, Proceedings of the 14th ACM conference on Electronic commerce, 2013, pp. 269–286.

[DDT14] _____, *The complexity of optimal mechanism design*, Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms, 2014, pp. 1302–1318.

[DDT15] C. Daskalakis, A. Deckelbaum, and C. Tzamos, *Strong duality for a multiple-good monopolist*, Proceedings of the Sixteenth ACM Conference on Economics and Computation (New York, NY, USA), EC '15, 2015.

[DW12] C. Daskalakis and S.M. Weinberg, *Symmetries and optimal multi-dimensional mechanism design*, Proceedings of the 13th ACM Conference on Electronic Commerce, 2012, pp. 370–387.

[GHK$^+$05] V. Guruswami, J.D. Hartline, A.R. Karlin, D. Kempe, C. Kenyon, and F. McSherry, *On profit-maximizing envy-free pricing*, Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms, 2005, pp. 1164–1173.

[GK14] Y. Giannakopoulos and E. Koutsoupias, *Duality and optimality of auctions for uniform distributions*, Proceedings of the 14th ACM Conference on Electronic Commerce, 2014, pp. 259–276.

[HN12] S. Hart and N. Nisan, *Approximate revenue maximization with multiple items*, Proceedings of the 13th ACM Conference on Electronic Commerce, 2012, pp. 656–656.

[HN13] S. Hart and N. Nisan, *The menu-size complexity of auctions*, Proceedings of the 14th ACM Conference on Electronic Commerce, 2013, pp. 565–566.

[LY13] X. Li and A.C.-C. Yao, *On revenue maximization for selling multiple independently distributed items*, Proceedings of the National Academy of Sciences of the United States of America, vol. 110, 2013, pp. 11232–11237.

[MV06] A.M. Manelli and D.R. Vincent, *Bundling as an optimal selling mechanism for a multiple-good monopolist*, Journal of Economic Theory **127** (2006), no. 1, 1–35.

[Mye81] R.B. Myerson, *Optimal auction design*, Mathematics of Operations Research **6** (1981), no. 1, 58–73.

[Pav10] G. Pavlov, *Optimal Mechanism for Selling Two Goods*, UWO Department of Economics Working Papers 20103, University of Western Ontario, Department of Economics, 2010.

[Rub16] A. Rubinstein, *On the computational complexity of optimal simple mechanisms*, Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January

14-16, 2016, 2016, pp. 21–28.

[Tha04] J. Thanassoulis, *Haggling over substitutes*, Journal of Economic Theory **117** (2004), no. 2, 217–245.

[WT14] Z. Wang and P. Tang, *Optimal mechanisms with simple menus*, Proceedings of the 15th ACM Conference on Economics and Computation, 2014, pp. 227–240.

[Yao15] A.C.-C. Yao, *An n-to-1 bidder reduction for multi-item auctions and its applications*, Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms, 2015, pp. 92–109.