2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS) | 979-8-3315-1674-1/24/$31.00 ©2024 IEEE | DOI: 10.1109/FOCS61266.2024.00082

Benchmark-Tight Approximation Ratio of Simple Mechanism for a Unit-Demand Buyer

1st Yaonan Jin
Huawei TCS Lab
Shanghai, China
jinyaonan@huawei.com

2nd Pinyan Lu
Shanghai University of Finance and Economic
Laboratory of Interdisciplinary Research of Computation and Economics (SUFE)
Huawei TCS Lab
Shanghai, China
lu.pinyan@mail.shufe.edu.cn

***Abstract*—We study revenue maximization in the unit-demand single-buyer setting. Our main result is that Uniform-Ironed-Virtual-Value Item Pricing guarantees a *tight* 3-approximation to the Duality Relaxation Benchmark [Chawla-Malec-Sivan, EC'10/GEB'15; Cai-Devanur-Weinberg, STOC'16/ SICOMP'21], breaking the barrier of 4 since [Chawla-Hartline-Malec-Sivan, STOC'10; Chawla-Malec-Sivan, EC'10/GEB'15]. To our knowledge, this is the first *benchmark-tight* revenue guarantee of any simple multi-item mechanism.**

Technically, all previous works employ Myerson Auction as an intermediary. The barrier of 4 follows as Uniform-Ironed-Virtual-Value Item Pricing achieves a *tight* 2-approximation to Myerson Auction, which then achieves a *tight* 2-approximation to Duality Relaxation Benchmark. Instead, our new approach avoids Myerson Auction, thus enabling the improvement. Central to our work are a *benchmark-based* 3-competitive prophet inequality and its *fully constructive* proof. Such variant prophet inequalities shall find future applications, e.g., to Multi-Item Mechanism Design where optimal revenues are relaxed to various more accessible benchmarks.

We complement our benchmark-tight ratio with an impossibility result. All previous works and ours follow the *single-dimensional representative* approach introduced by [Chawla-Hartline-Kleinberg, EC'07]. Against Duality Relaxation Benchmark, it turns out that this approach cannot beat our bound of 3 for a large class of Item Pricing's.

***Index Terms*—Mechanism Design, Revenue Maximization, Approximation Algorithms, Prophet Inequalities**

I. Introduction

The TCS community has tremendously advanced on Multi-Item Mechanism Design in the last two decades. Once it became clear that, even only with a single buyer, optimal multi-item mechanisms are intolerably complicated [Tha04], [BCKW15], [HN19], [DDT14], [CDP+18], [CDO+22], [CMPY18], the community turned to examine the problem through the lenses of *approximation*. Hitherto there is a long line of works on the design and analysis of *simple and approximately optimal* mechanisms [CHK07], [CHMS10], [CMS15], [CD15], [KW19], [HN17], [LY13], [BILW20], [Yao15], [HH15], [RW18], [CM16], [CDW21], [CZ17], [Yao18], [JLQ+19b], [DFL+22], [MS21], [DKL20], [CC23], [CCW23, and the references therein].

This rich literature dates back to the seminal work of Chawla, Hartline, and Kleinberg [CHK07], who studied a revenue-maximizing seller (she) with $n \geq 2$ items facing a *unit-demand* buyer (he) with item-wise values v_i for $i \in [n]$ and bundle-wise values $v_S = \max_{i \in S}\{v_i\}$ for $S \subseteq [n]$. Namely, item-wise values $v_i \sim F_i$ are *independently distributed*; the seller knows the priors $\boldsymbol{F} = \{F_i\}_{i \in [n]}$ but does not know the outcomes $\boldsymbol{v} = (v_i)_{i \in [n]}$, thus aiming to maximize her expected revenue.

This revenue maximization problem can be formulated as an (infinite-dimensional) linear program[1] and a feasible solution can be interpreted as follows: The seller offers the buyer a menu of lotteries $\{(\boldsymbol{x}, \ell)\}$; each lottery includes its item-wise allocation probabilities $\boldsymbol{x} = (x_i)_{i \in [n]}$ and its price ℓ.[2] The buyer will purchase the utility-maximizing lottery $\mathrm{argmax}_{(\boldsymbol{x}, \ell)}\{\boldsymbol{v} \cdot \boldsymbol{x} - \ell\}$, or nothing if no lottery induces a nonnegative expected utility. The revenue-maximizing menu is called Optimal Lottery Pricing (OLP). Unfortunately, the associated computation problem is intractable [CDO+22]: First, $\Omega(2^n)$ many lotteries can be necessary even for binary-supported distributions $|\mathrm{supp}(F_i)| = 2$, hence an exponential *menu complexity* [HN13]. Second, no randomized polynomial-time algorithm can implement Optimal Lottery Pricing, unless $\mathbf{P}^{\#\mathbf{P}} = \mathbf{P}^{\mathbf{NP}}$, even for ternary-supported distributions $|\mathrm{supp}(F_i)| = 3$.

An equally fascinating problem, revenue maximization among deterministic mechanisms, can be formulated as an (infinite-dimensional) mixed-integer program[1] and a feasible solution can be interpreted as follows: The seller posts item-wise prices $\boldsymbol{p} = (p_i)_{i \in [n]}$ and, then, the buyer purchases his favorite item $\mathrm{argmax}_{i \in [n]}\{v_i - p_i\}$, or nothing if every price p_i is above value v_i. The revenue-maximizing deterministic mechanism is called Optimal Item Pricing (OIP). Unfortunately, the associated computation problem is $\mathbf{NP}$-hard even for ternary-supported distributions $|\mathrm{supp}(F_i)| = 3$, or even for i.i.d. distributions $\boldsymbol{F} = \{F\}^{\otimes n}$ of support size $|\mathrm{supp}(F)| = \mathrm{poly}(n)$ [CDP+18].

Given the above hardness results, it is of particular interest

[1] More precisely, this linear or mixed-integer program has a size polynomial in $\prod_{i \in [n]} |\mathrm{supp}(F_i)|$, *the support size of the joint value distribution*. Thus, even for binary-supported distributions $|\mathrm{supp}(F_i)| = 2$, it has a size exponential in the number of items $n \geq 2$.

[2] Since the buyer is unit-demand, without loss of generality, a lottery $(\boldsymbol{x}, \ell)$ can have a total allocation probability of at most one $\sum_{i \in [n]} x_i \leq 1$, i.e., exclusively allocating each item $i \in [n]$ with probability x_i.

2575-8454/24/$31.00 ©2024 IEEE
DOI 10.1109/FOCS61266.2024.00082

to investigate simple *deterministic polynomial-time* mechanisms and to analyze their revenue guarantees against Optimal Item Pricing and Optimal Lottery Pricing.[3] Among all the simple mechanisms studied in the literature, only one unconditionally achieves a constant-factor revenue guarantee against either optimum.[4] It is called Uniform-Ironed-Virtual-Value Item Pricing, or simply UIVV Item Pricing (UIVVIP); cf. Figure 1 for a diagram of its revenue guarantees against either Optimal Item Pricing or Optimal Lottery Pricing, as well as other revenue gaps.

UIVV Item Pricing was proposed in Chawla, Hartline, and Kleinberg's original work [CHK07]. In particular, they proved that it achieves a 3-approximation to Optimal Item Pricing. Afterward, Chawla, Hartline, Malec, and Sivan [CHMS10] improved this upper bound to 2, which remains the state of the art.

In a breakthrough, Chawla, Malec, and Sivan [CMS15] (the conference version was in EC'10) established that UIVV Item Pricing achieves a 4-approximation to Optimal Lottery Pricing. Indeed, since the $\mathbf{P}^{\#\mathbf{P}}$-hard Optimal Lottery Pricing is prohibitively difficult to deal with, [CMS15] instead relaxed it to a more accessible benchmark, Duality Relaxation Benchmark (DRB).[5] I.e., they proved that this benchmark revenue-surpasses Optimal Lottery Pricing and that UIVV Item Pricing even ensures a 4-approximation to it. The follow-up by Cai, Devanur, and Weinberg [CDW21] presented an alternative and possibly more "universal" proof, including a *duality-based* new interpretation for the benchmark. Nonetheless, for nearly fifteen years, the original bound of 4 from [CMS15] remains state of the art and Duality Relaxation Benchmark remains the only known useful relaxation of Optimal Lottery Pricing.[6]

Exploring "Optimal Lottery Pricing vs. Optimal Item Pricing" is also interesting, i.e., the gap between randomized and deterministic revenues, regardless of computational constraints. Unfortunately, mainly because both mechanisms are computationally intractable, nothing is known except an upper bound of 4 (by implication) and a best-known lower bound of $1 + \frac{\ln 2}{5} \approx 1.1386$ [CMS15].

II. Our contributions

In this work, we show the *tight* approximation ratio $\mathcal{C}_{\mathsf{DRB/UIVVIP}} = 3$ of UIVV Item Pricing to Duality Relaxation Benchmark. To our knowledge, this is the first benchmark-tight approximation ratio of any simple multi-item mechanism – within or beyond the unit-demand single-buyer setting.

The worst-case instances for $\mathcal{C}_{\mathsf{DRB/UIVVIP}} = 3$ further inspire us to improve the lower bounds on the approximation ratios of UIVV Item Pricing against either Optimal Lottery Pricing or Optimal Item Pricing. Namely, we show a family of n-item instances by which $\mathcal{C}_{\mathsf{OLP/UIVVIP}} \geq \mathcal{C}_{\mathsf{OIP/UIVVIP}} \geq \frac{2}{1+1/n}$. This lower bound can be arbitrarily close to 2, when the number of items is large enough $n \gg 2$. Also, it matches the upper bound $\mathcal{C}_{\mathsf{OIP/UIVVIP}} \leq 2$ from [CHMS10], thus closing this approximation ratio $\mathcal{C}_{\mathsf{OIP/UIVVIP}} = 2$.

The following theorem summarizes our improved upper and lower bounds on the revenue guarantees of UIVV Item Pricing; cf. Figure 1 for a diagram. (The upper-bound part of Item (iii) is credited to [CHK07], [CHMS10].)

Theorem 1 (Revenue Guarantees of UIVVIP)**.** *Uniform-Ironed-Virtual-Value Item Pricing achieves*
(i) a tight $\mathcal{C}_{\mathsf{DRB/UIVVIP}} = 3$ approximation to Duality Relaxation Benchmark,
(ii) a $\mathcal{C}_{\mathsf{OLP/UIVVIP}} \in [2, 3]$ approximation to Optimal Lottery Pricing, and
(iii) a tight $\mathcal{C}_{\mathsf{OIP/UIVVIP}} = 2$ approximation to Optimal Item Pricing.

By implication, we also improve the upper bounds on the revenue guarantees of Optimal Item Pricing against either Duality Relaxation Benchmark or Optimal Lottery Pricing, both from 4 to 3. (The previous upper bounds also are implications of $\mathcal{C}_{\mathsf{DRB/UIVVIP}} \leq 4$ [CHK07], [CHMS10], [CMS15], [CDW21].) In this regard, we construct another 2-approximation lower-bound instance for "Duality Relaxation Benchmark vs. Optimal Item Pricing".

The following corollary concludes our improved upper and lower bounds on the revenue guarantees of Optimal Item Pricing; cf. Figure 1 for a diagram. (The lower-bound part of Item (ii) is credited to [CMS15].)

Corollary 2 (Revenue Guarantees of OIP)**.** *Optimal Item Pricing achieves*
(i) a $\mathcal{C}_{\mathsf{DRB/OIP}} \in [2, 3]$ approximation to Duality Relaxation Benchmark and
(ii) a $\mathcal{C}_{\mathsf{OLP/OIP}} \in [1 + \frac{\ln 2}{5} \approx 1.1386,\ 3]$ approximation to Optimal Lottery Pricing.

Limitations on Further Improvements. Given our improved upper bounds of 3, one might be optimistic with better revenue guarantees of UIVV Item Pricing, Optimal Item Pricing, or other simple deterministic mechanisms in between. Unfortunately, en route to Theorem 1, we also show that beating

[3] It is equally interesting to design and analyze simple *randomized* mechanisms. In this regard, however, the only positive progress is a quasi-polynomial time approximation scheme ($\mathbf{QPTAS}$) for Optimal Lottery Pricing [KSM+19].

[4] But if we impose mild assumptions on the value distributions $\boldsymbol{F} = \{F_i\}_{i \in [n]}$, more simple mechanisms also can achieve constant-factor revenue guarantees. E.g., if the value distributions satisfy the *regularity* condition, a standard condition from Myerson's semial work [Mye81], Optimal Item Pricing admits a $\mathbf{QPTAS}$ [CD15] and even the simplest mechanism Uniform Pricing, i.e., posting an optimal uniform price p for all items, achieves a tight $\mathcal{C}_{\mathsf{OIP/UP}} \approx 2.6202$-approximation to Optimal Item Pricing and a $2\mathcal{C}_{\mathsf{OIP/UP}} \approx 5.2404$-approximation to Optimal Lottery Pricing [JLQ+19b].

[5] [CMS15] did not name this benchmark by Duality Relaxation Benchmark. Instead, the follow-ups [HH15], [CDW21] gave it a *duality-based* new interpretation. Thus, without ambiguity, we call it Duality Relaxation Benchmark.

[6] More precisely, [EFF+17], [BW19] considered another duality-based benchmark (slightly different from the one by [CMS15], [HH15], [CDW21]). Nonetheless, UIVV Item Pricing turns out to have a worse or the same revenue guarantee ≥ 3 against the [EFF+17], [BW19] benchmark, which means that benchmark is not so useful for our purpose. (Our lower-bound instances (see the full version) are *symmetric instances*, by which both benchmarks have the same revenue, so the [EFF+17], [BW19] benchmark can only establish a worse or the same revenue guarantee ≥ 3.)

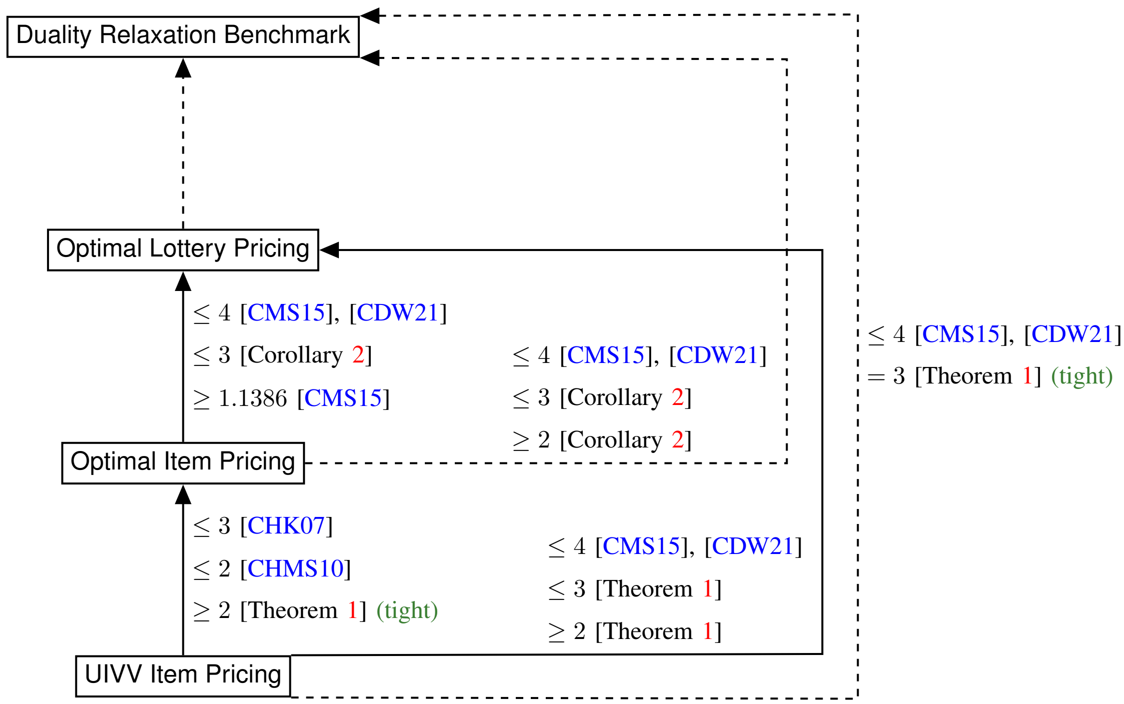

Fig. 1. A diagram of the previous results and our new results. The four mechanisms constitute a revenue hierarchy: Duality Relaxation Benchmark $\succeq$ Optimal Lottery Pricing $\succeq$ Optimal Item Pricing $\succeq$ UIVV Item Pricing. (Duality Relaxation Benchmark itself may lack a concrete economic meaning but remains the only known useful relaxation of Optimal Lottery Pricing.) One solid arrow refers to the revenue gap between two concrete mechanisms, and one dashed arrow refers to the revenue gap between one concrete mechanism and Duality Relaxation Benchmark.
We emphasize that, mainly because the $\mathbf{P}^{\#\mathbf{P}}$-hard Optimal Lottery Pricing and the $\mathbf{NP}$-hard Optimal Item Pricing both are sophisticated, actual positive progress has been made only for "Duality Relaxation Benchmark (resp. Optimal Item Pricing) vs. UIVV Item Pricing".
For the other revenue gaps, all the previous upper bounds and our new upper bounds are implications from "Duality Relaxation Benchmark vs. UIVV Item Pricing".

our bound of 3 transcends the scope of known methods; see Section III for details. Roughly speaking, this is because both the previous works [CHK07], [CHMS10], [CMS15], [CDW21] and this paper (possibly implicitly) apply the *single-dimensional representative* approach initiated by [CHK07]. We show that, against Duality Relaxation Benchmark, this approach cannot establish an approximation ratio better than 3 for any *ironed-virtual-value-based* Item Pricing. We note that all the mentioned works (and many others beyond the unit-demand single-buyer setting) restricted their attention to this family of mechanisms. On the other hand, direct analysis of the $\mathbf{P}^{\#\mathbf{P}}$-hard Optimal Lottery Pricing seems out of reach.

III. Technical overview

Below, we sketch our new approach for the tight revenue guarantee $\mathcal{C}_{\mathsf{DRB/UIVVIP}} = 3$ of UIVV Item Pricing against Duality Relaxation Benchmark, assuming that the reader has basic familiarity with the previous works, especially Myerson's virtual value theory [Mye81].

To begin with, recall that the previous approach [CHK07], [CHMS10], [CMS15], [CDW21] for the weaker bound $\mathcal{C}_{\mathsf{DRB/UIVVIP}} \leq 4$ takes three steps.[7]

$$\text{UIVV Item Pricing} \geq \text{UIVV Sequential Posted Pricing} \tag{1}$$

▷ the single-dimensional representative approach (tight)

$$\geq \tfrac{1}{2} \cdot \text{Myerson Auction} \tag{2}$$

▷ the order-oblivious prophet inequalities (tight)

$$\geq \tfrac{1}{4} \cdot \text{Duality Relaxation Benchmark}. \tag{3}$$

▷ benchmark decomposition (tight)

Step (1) applies the single-dimensional representative approach [CHK07], [CHMS10], which relates the (less understood) unit-demand single-buyer setting to the (better understood) single-item multi-buyer setting. The former single buyer's value and distribution $v_i \sim F_i$ for each former item $i \in [n]$ is reinterpreted as each latter buyer i's value and distribution $v_i \sim F_i$ for the latter single item. Therefore, UIVV

[7]Of course, if we consider Optimal Item Pricing and/or Optimal Lottery Pricing instead, there are two other steps:
Optimal Item Pricing $\geq$ UIVV Item Pricing (obvious) and $\frac{1}{4} \cdot$ Duality Relaxation Benchmark $\geq \frac{1}{4} \cdot$ Optimal Lottery Pricing (proved in [CMS15], [CDW21]). Both equalities simultaneously hold for (say) the first instance in Footnote 8.

Item Pricing shall be replaced with its single-item multi-buyer counterpart, UIVV Sequential Posted Pricing (UIVVSPP) in the *order-oblivious* model (elaborated soon).

Step (2), due to well-known reductions [CHMS10], [CFPV19] based on Myerson's virtual value theory, is equivalent to the order-oblivious prophet inequalities [KS78], [SC84], [KW19].

Step (3) is evident once we write down the revenue formula of Duality Relaxation Benchmark. It is basically a "mixture" of the revenues from Myerson Auction (MA) and Second Price Auction (SPA), thus being upper-bounded by the latter two's total revenue.

We emphasize that all individual steps are tight, i.e., the equalities hold for some step-specific instances,[8] which may explain why the bound of 4 is long-standing. In contrast, our approach will enable the improvement by unifying and then refining Steps (2) and (3).

Mechanisms and Their Relation. First, we shall formalize Item Pricing and Sequential Posted Pricing and recap how to relate them via the single-dimensional representative approach.

Definition III.1 (Item Pricing)**.** For a unit-demand single-buyer instance $\boldsymbol{F} = \{F_i\}_{i\in[n]}$,

- Item Pricing: The seller posts item-wise prices $\boldsymbol{p} = (p_i)_{i\in[n]}$ on items $i \in [n]$. Then the buyer purchases the utility-maximizing item $\text{argmax}_{i\in[n]}\{v_i - p_i\}$, breaking ties in favor of (one of) the highest-priced item(s),[9] or nothing if every price is above the value $p_i > v_i$, $\forall i \in [n]$.
  Denote by $\mathsf{IP}(\boldsymbol{F},\, \boldsymbol{p})$ the resulting revenue.

Definition III.2 (Sequential Posted Pricing)**.** For a single-item multi-buyer instance $\boldsymbol{F} = \{F_i\}_{i\in[n]}$,

- Sequential Posted Pricing: The seller posts buyer-wise prices $\boldsymbol{p} = \{p_i\}_{i\in[n]}$ for buyers $i \in [n]$. Then the buyers arrive one by one $\pi(1), \pi(2), \ldots, \pi(n)$ in a specific order $\pi \in \Pi_n$; the first arrival buyer $\pi(i)$ who accepts her price $v_{\pi(i)} \geq p_{\pi(i)}$ gets the item.
  Denote by $\mathsf{SPP}(\boldsymbol{F},\, \boldsymbol{p},\, \pi)$ the resulting revenue.
- The *order-oblivious* model: Once the seller determines the buyer-wise prices $\boldsymbol{p} = \{p_i\}_{i\in[n]}$, a hypothetical adversary chooses (one of) the revenue-worst arrival order(s) $\pi^* \in \Pi_n$ for the seller.
  Denote by $\mathsf{SPP}^*(\boldsymbol{F},\, \boldsymbol{p}) := \min_{\pi\in\Pi_n}\{\mathsf{SPP}(\boldsymbol{F},\, \boldsymbol{p},\, \pi)\}$ the resulting revenue.

[8] Here are three "tight" instances, respectively; for brevity, we omit calculation of revenues (elementary algebra).
The equality in Step (1): $F_1(v) = F_2(v) = \cdots = F_n(v) := \mathbb{I}(v \geq 1)$, i.e., all values are deterministic $\equiv 1$.
The equality in Step (2): $F_1(v) := \mathbb{I}(v \geq 1)$ and $F_2(v) := 1 - 1/v$ for $v \geq 1$ (a.k.a. the equal-revenue distribution).
The equality in Step (3): $F_1(v) = F_2(v) = \cdots = F_n(v) := 1 - 1/v$ for $v \geq 1$ and $n \geq 2$ (see more details in the full version).

[9] As shown in [CDP+18, Section 2.2], this tie-breaking rule is without loss of generality.

The single-dimensional representative approach [CHMS10, Theorem 4] asserts that, under the same *deterministic* prices $\boldsymbol{p}$, Item Pricing revenue-surpasses Sequential Posted Pricing in the order-oblivious model. (We note that, to validate Proposition III.3 as a black-box reduction, the involved prices $\boldsymbol{p}$ must be deterministic; see the full version for more details. The previous work [CHMS10] did not emphasize this issue.)

Proposition III.3 (The Single-Dimensional Representative Approach [CHMS10])**.** *Given the same instance $\boldsymbol{F}$ and the same* ***deterministic*** *prices $\mathbf{p}$, Item Pricing revenue-surpasses Sequential Posted Pricing in the order-oblivious model* $\mathsf{IP}(\boldsymbol{F},\, \boldsymbol{p}) \geq \mathsf{SPP}^*(\boldsymbol{F},\, \boldsymbol{p})$.

Specifically, [CHK07], [CHMS10] designed the uniform-ironed-virtual-value prices $\boldsymbol{p}$ as follows.

Definition III.4 (Uniform-Ironed-Virtual-Value Prices)**.** For a single-item multi-buyer instance $\boldsymbol{F} = \{F_i\}_{i\in[n]}$, let $\bar{\boldsymbol{\varphi}} = \{\bar{\varphi}_i\}_{i\in[n]}$ be the *increasing* ironed virtual value functions. Given a uniform ironed-virtual-value threshold (UIVV-threshold) $T > 0$, define the item-wise prices as $p_i := \sup\{v \in [0, +\infty] \mid \bar{\varphi}_i(v) < T\} = \inf\{v \in [0, +\infty] \mid \bar{\varphi}_i(v) \geq T\}$ for each $i \in [n]$,[10]
where the equality holds since each $\bar{\varphi}_i$ is an increasing function.

Now we are ready to elaborate on how to unify and refine Steps (2) and (3), thus obtaining our benchmark-tight revenue guarantee $\mathcal{C}_{\mathsf{DRB/UIVVIP}} = 3$.

Unifying Steps (2) and (3). To clarify the differences between our approach and the previous works [CHK07], [CHMS10], let us recap the previous proof of Step (2), as follows.

Theorem 3 (Revenue Guarantees of UIVVSPP [CHK07], [CHMS10])**.** *Against Myerson Auction:*
(i) Uniform-Ironed-Virtual-Value Sequential Posted Pricing with a UIVV-threshold $T = \frac{1}{2} \cdot \mathsf{MA}$ of one half of the Myerson Auction revenue achieves a $\mathcal{C}_{\mathsf{MA/UIVVSPP}} = 2$ approximation.
(ii) This ratio is optimal for ***all*** *stopping rules in the order-oblivious model.*

Proof (The Upper-Bound Part). Given a specific threshold $T > 0$ and an arbitrary arrival order $\pi \in \Pi_n$ for Sequential Posted Pricing, the prices $\boldsymbol{p} = \{p_i\}_{i\in[n]}$ from Definition III.4 satisfy that "each buyer $i \in [n]$ gets allocated with the same probability, under any two values v_i and $\tilde{v}_i$ with the same ironed virtual value $\bar{\varphi}_i(v_i) = \bar{\varphi}_i(\tilde{v}_i)$." Thus, we can apply the revenue equivalence [Mye81], [CDW21] (see the full version for a formal statement). I.e., the UIVV Sequential Posted Pricing revenue $=$ its ironed virtual welfare, and its revenue guarantee reduces to the order-oblivious prophet inequalities [KS78], [SC84], [KW19]. More precisely, a threshold of

[10] We can alternatively define the prices as $p_i := \sup\{v \in [0, +\infty] \mid \bar{\varphi}_i(v) \leq T\} = \inf\{v \in [0, +\infty] \mid \bar{\varphi}_i(v) > T\}$. Both definitions will induce different prices $\boldsymbol{p} = \{p_i\}_{i\in[n]}$ but the same bounds $\mathcal{C}_{\mathsf{MA/UIVVIP}} = 2$ and $\mathcal{C}_{\mathsf{DRB/UIVVIP}} = 3$.

$T = \frac{1}{2} \cdot \mathsf{MA}$ is well-known to 2-approximate the optimal ironed virtual welfare = the Myerson Auction revenue. □

Theorem 3 in combination with "Myerson Auction $\geq \frac{1}{2}$ · Duality Relaxation Benchmark" [CMS15], [CDW21] gives the previous revenue guarantee $\mathcal{C}_{\mathsf{DRB/UIVVSPP}} \leq 4$ of UIVV Sequential Posted Pricing against Duality Relaxation Benchmark.

An astute reader may already note that, for our sake of the revenue gap $\mathcal{C}_{\mathsf{DRB/UIVVSPP}}$ between Duality Relaxation Benchmark and UIVV Sequential Posted Pricing, the "right" approach shall avoid the intermediary, Myerson Auction, and instead attack this revenue gap directly. This is precisely our main technical ingredient, which we call *benchmark-based* prophet inequalities.

Theorem 4 (Revenue Guarantees of UIVVSPP)**.** *Against Duality Relaxation Benchmark:*
(i) Uniform-Ironed-Virtual-Value Sequential Posted Pricing with a UIVV-threshold $T = \frac{1}{3} \cdot \mathsf{DRB}$ of one third of the Duality Relaxation Benchmark revenue achieves a $\mathcal{C}_{\mathsf{DRB/UIVVSPP}} = 3$ approximation.
*(ii) This ratio is optimal for **deterministic ironed-virtual-value-based** and/or **(arbitrary) uniform-ironed-virtual-value** stopping rules in the order-oblivious model.*[11]

As mentioned, Duality Relaxation Benchmark is a mixture of the revenues from Myerson Auction and Second Price Auction. This means, in our approach to Theorem 4, we must frequently switch between the value space and the ironed virtual value space. To prove Theorem 3, in contrast, the revenue equivalence allows us to concentrate solely on the ironed virtual value space and thus apply the classical prophet inequalities. This difference incurs many technical difficulties (and may explain why the previous works [CMS15], [CDW21] did not investigate this revenue gap directly but leveraged Myerson Auction as an intermediary).

Nonetheless, we develop many new techniques here and give a *fully constructive* proof of Theorem 4. More concretely, we will prove more and more necessary conditions for the worst-case instances, thus gradually shrinking the optimization space and finally deriving the exact worst-case instances and the tight bound $\mathcal{C}_{\mathsf{DRB/UIVVSPP}} = 3$. Such a fully constructive proof might be more structure-revealing than the nonconstructive proofs of the classical prophet inequalities.[12]

Such constructive proofs have made many success stories in Single-Item Mechanism Design. There is a long line of works on proving tight revenue/welfare guarantees of both truthful and nontruthful single-item mechanisms [CGL14], [AHN+19], [CFH+21], [AB20], [JLTX20], [JLQ+19b], [JLQ19a], [HJL20], [JJLZ22], [CST18], [BMTT19], [JL22], [JL23]. To our knowledge, however, our work is the first to adopt such techniques to Multi-Item Mechanism Design. In this regard, we look forward to more future works that technically bridge these two highly relevant areas.

Given Proposition III.3 and the lower-bound part of Theorem 4, we can also conclude that the single-dimensional representative approach cannot establish a bound better than 3 for any *ironed-virtual-value-based* Item Pricing. (As mentioned, all the previous works [CHK07], [CHMS10], [CMS15], [CDW21] and many others beyond the unit-demand single-buyer setting restricted their attention to this family of mechanisms.)

Further, we remark that the Duality Relaxation Benchmark considered here actually is the instantiation of the benchmark from [CMS15], [HH15], [CDW21] to the unit-demand single-buyer setting – the original benchmark accommodates much more general settings. It is conceivable that our main technical ingredient – the benchmark-based prophet inequality (Theorem 4) – also leaves enough room for future generalizations, which shall give improved or even benchmark-tight approximation ratios of other simple multi-item mechanisms.

A. Proof overview

In the section, we provide an outline of the proof for our main theorem.

One of our first key observation is that the Duality Relaxation Benchmark for a single unit-demand buyer with independent values is monotone with respect to the stochastic dominance of value distributions; see the full version for more details. This property turns out to be very useful in our proof. Actually, one of the reasons why Multi-Item Mechanism Design is much more difficult than Single-Item Mechanism Design is that the optimal mechanisms (in general) no longer satisfy revenue monotonicity [HR15]. It is very difficult to compare to a benchmark which itself is not even monotone. The proof for the monotonicity of Duality Relaxation Benchmark is not that difficult but also nontrivial since it is not point-wise dominance. To the best of our knowledge, this monotonicity has not been pointed out in literature before.[13] We believe that this observation may find its applications in other problems.

[11]By a **deterministic ironed-virtual-value-based** stopping rule, we mean: First, the seller chooses a deterministic set of *acceptable* ironed virtual values $\Phi_i \subseteq \mathbb{R}$ for every buyer $i \in [n]$, i.e., different possible values $v_i \sim F_i$ with the same ironed virtual value either all are acceptable $\bar{\varphi}_i(v_i) \in \Phi_i$ or all are unacceptable $\bar{\varphi}_i(v_i) \notin \Phi_i$. (So a buyer $i \in [n]$ on his arrival takes the item iff $\bar{\varphi}_i(v_i) \in \Phi_i$ and no earlier buyer had taken the item.) Then, the adversary chooses a worst-case arrival order $\pi^* \in \Pi_n$ against this stopping rule, i.e., the acceptance sets $\{\Phi_i\}_{i \in [n]}$.
By an **arbitrary uniform-ironed-virtual-value**, we mean: First, the seller chooses an arbitrary ironed-virtual-value threshold T (which can be randomized) and posted prices $\boldsymbol{p} = (p_i)_{i \in [n]}$ such that $\bar{\varphi}_i(p_i) = T$ (which can be randomized even conditioned on T). Then, the adversary chooses a worst-case arrival order $\pi^* \in \Pi_n$ against the distributional information of $\boldsymbol{p}$.

[12]I.e., [KS78], [SC84], [KW19] set the threshold as $T = \frac{1}{2} \cdot \mathsf{MA}$ (say) constructively, but proved the upper bound $\mathcal{C} \leq 2$ non-constructively using the probabilistic method and gave the matching lower-bound instances separately.

[13]Rather, approximate versions of the revenue monotonicity of (optimal) multi-item mechanisms have been established in [BILW20], [RW18], [Yao18], [CCW23].

Based on this revenue monotonicity, we can do a few transformations of the given instance to get instances with worse revenue gap. Our first reduction Iron that transforms a given instance to a regular instance is standard and natural, as UIVV Sequential Posted Pricing uses *ironed* virtual values at the first place. After that, we simply assume that all value distributions are regular. The next two transformations are crucial and also quite intuitive. They preserve the revenue of UIVV Sequential Posted Pricing (UIVVSPP) but increase the revenue of Duality Relaxation Benchmark by its revenue monotonicity, thus a larger revenue gap after the transformation. Since the revenue of UIVVSPP only depends on the probability of one's value is below or above the given price but not on how much, we can modify the instance so that their virtual values are either slightly above or slightly below the given threshold. In particular, the Truncate reduction makes the not-buying part (below the price) only slightly below the given threshold. This will not change the revenue of UIVVSPP but will increase the Duality Relaxation Benchmark since the distribution after the modification stochastically dominates the original one. The Extend reduction modifies the buying part (above the price) so that the virtual value only slightly above the threshold for most probability and at a very high value ($\to \infty$) for a very small probability. Again, this will not change the revenue of UIVVSPP. The fact that the distribution after the modification dominates the previous one seems less obvious than the Truncate reduction if one just reads the above description. However, it *is* obvious once one looks at the revenue-quantile curve before and after the modification as the later revenue-quantile curve is point-wise above the original one. Such instances after the modification are called *semi-linear* instances (see the definition in the full version) in our paper as their revenue-quantile curves are almost linear.

We note that the revenue-quantile curve is the most convenient representation for a value distribution in our proof and we use it extensively [Har13, Chapter 3]. This is because we need to go back and forth between the value space and the virtual value space; fortunately, both values and virtual values have intuitive geometry representation in revenue-quantile curves. If we use value CDF (say) instead, it is inconvenient and less intuitive to represent virtual values. Another notational advantage of using revenue-quantile curves rather than value distributions directly is that we can implicitly introduce the value of ∞, which is both mathematical sound and notational simple. This is achieved by allowing $R(0) > 0$, namely we have strictly positive revenue by a sale probability of 0, which means that the price (and the value) is ∞. However, this convention may cause some confusion if the reader is not so familiar with it. Thus, let us point it out in this outline to avoid potential confusion since we use it very often in our proof. If one does not like infinite values, this can always be interpreted as a sufficiently high value $H \to \infty$ with a sufficiently small probability $R(0)/H \to 0^+$. Once the reader gets familiar with this, it is convenient to use revenue-quantile curves with $R(0) > 0$.

These transforms and the above arguments omit a subtlety; they assume that the threshold remains the same after the modifications of the instances. This is inaccurate. Actually, the threshold is selected after the instance is given, which is set to be $T = \frac{1}{3} \cdot \mathsf{DRB}$ in our final Sequential Posted Pricing. To address this subtlety, we introduced the Scale transform. This reduction is controlled by a continuous parameter $\gamma \in [0, 1]$ which will reduce the benchmark revenue and the UIVVSPP revenue simultaneously and continuously. But by the mean value theorem, we prove that there exist a $\gamma^* \in [0, 1]$ such that the instance after the Scale reduction has the same benchmark revenue as the most original instance, thus the same threshold $T = \frac{1}{3} \cdot \mathsf{DRB}$. This reduction remedies the mentioned issue. Thus, compared with the most original instance, the new instance (after Scale) has the same benchmark revenue but a lower UIVVSPP revenue, thus a larger revenue gap.

Our last reduction Perturb transforms a semi-linear instance into a *linear instance* (see the definition in the full version). To achieve this, we can take a different but equivalent viewpoint of the above semi-linear instances. Namely, their virtual values are exactly equal to the threshold, rather than slightly above or slightly below it. The take-it-or-leave-it decisions in UIVV Sequential Posted Pricing are determined by a tie-breaking rule that yields the same posted prices and take-it-or-leave-it decisions, hence the same revenue as before. From this viewpoint, it is already a linear instance. These two different viewpoints do not change the benchmark revenue, as the tie-breaking rule only effects the UIVVSPP revenue. The key observation here is to consider two extreme-case tie-breaking rules, namely "always take it" vs. "always leave it" in ties. Regarding the UIVVSPP revenues, the worse one between these two tie-breaking rule *is* the worst one among all possible tie-breaks. Thus, it is sufficient to only consider these two extreme cases. For these two families of (extreme-case) linear instances, we can easily obtain the formulae for both the benchmark revenues and the UIVVSPP revenues. Thus, we can derive an upper bound on the revenue gap $\mathcal{C}_{\mathsf{DRB/UIVVSPP}}$ after optimizing the parameters.

We notice that, for revenue maximization in *uniform pricing* mechanisms, breaking ties in favor of buying (trivially) is always good. But this is not the case for UIVV Item Pricing, since the same virtual value may correspond to many different values/prices. (Our linear instances are typical for this phenomenon.) The so-called "tie-breaking rule" throughout our proof is not an actual tie-breaking rule. Rather, it refers to different choices (as the UIVV posted prices) of different values that have the same ironed virtual value. It is unclear which choice, "higher prices with lower sale probabilities" vs. "lower prices with higher sale probabilities", is better for revenue maximization. To our knowledge, this subtlety about different choices of UIVV posted prices (or, generally, *ironed-value-value-based* posted prices) has not been pointed out in the literature.

Whether this subtlety has effects on revenue maximization is an important question to study. For the *upper bounds*, we remark that all mentioned revenue guarantees hold for any choice of UIVV posted prices (i.e., even the worst-case UIVV posted prices). For the *lower bounds*, however, this subtlety seems to influence Item Pricing more than Sequential Posted Pricing. (i) The tight revenue guarantee of 3 for Sequential Posted Pricing is quite robust to different choices of UIVV posted prices; see the statement of Theorem 4, Footnote 11, and the full version for more details. (ii) The revenue guarantee for Item Pricing is trickier. On the one hand, if we insist on the "naive" tie-breaking rule in Definition III.4 for UIVV posted prices, namely the smallest/largest values that correspond to the given UIVV threshold T, then the bound of 3 *is* tight. (Also, if we follow the single-dimensional representative approach, this bound of 3 is the best possible.) On the other hand, if we instead adopt (say) the "best-UIVV" posted prices, then we cannot find a matching lower-bound instance. In this regard, it is interesting for future works to study the revenue guarantees of the "best-UIVV" Item Pricing; see the full version for more details.

IV. Further related works

Beyond the paradigmatic unit-demand single-buyer setting, Multi-Item Mechanism Design has been systematically studied in other parallel and more general settings, which can be categorized as the single- vs. multi-buyer settings, and based on the underlying valuation families:

$$\text{unit-demand} \;\&\; \text{additive} \subseteq \text{matroid-rank} \subseteq \text{constrained-additive} \subseteq \text{XOS} \subseteq \text{subadditive}.$$

See, e.g., [CZ17] for their formal definitions. In particular, "unit-demand" and "additive" are the two most basic valuation families; one is neither a subset nor a superset of the other.

This flexibility in modeling has made Multi-Item Mechanism Design a centerpiece in the intersection of Theoretical Computer Science and Mathematical Economics. As we quote from Cai and Papadimitriou [CP14]: "... *multi-item mechanisms have played a role akin to that enjoyed by the Traveling Salesman Problem in combinatorial optimization: A paradigmatic hard nut on which all new ideas must be tried* ..." It is impossible to review such an extensive literature here thoroughly. Instead, we list some of these works in Table I for guidance.

Very recently, it was shown in [CC23], [CCW23] that, even in the (most general) subadditive multi-buyer setting, simple deterministic mechanisms are constant-factor approximately revenue-optimal, for which prophet inequalities again play a pivotal role. Hence, it is now a natural time to tighten these bounds. For this sake, our contributions are threefold: (i) the benchmark-tight ratio of UIVV Item Pricing, (ii) a new technique with enough room for future generalizations, benchmark-based prophet inequalities, and (iii) showing the limitations of the single-dimensional representative approach. We hope all our positive and negative results will inspire future works.

Acknowledgment

We would like to thank Yingkai Li for invaluable discussions and Dongxu Wang for help with mathematical experiments.

References

[AB20] Amine Allouah and Omar Besbes. Prior-independent optimal auctions. *Manag. Sci.*, 66(10):4417–4432, 2020. doi:10.1287/mnsc.2019.3459.

[AHN+19] Saeed Alaei, Jason D. Hartline, Rad Niazadeh, Emmanouil Pountourakis, and Yang Yuan. Optimal auctions vs. anonymous pricing. *Games Econ. Behav.*, 118:494–510, 2019. doi:10.1016/j.geb.2018.08.003.

[BCKW15] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. Pricing lotteries. *J. Econ. Theory*, 156:144–174, 2015. doi:10.1016/j.jet.2014.04.011.

[BILW20] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. *J. ACM*, 67(4):24:1–24:40, 2020. doi:10.1145/3398745.

[BMTT19] Georgios Birmpas, Evangelos Markakis, Orestis Telelis, and Artem Tsikiridis. Tight welfare guarantees for pure nash equilibria of the uniform price auction. *Theory Comput. Syst.*, 63(7):1451–1469, 2019. doi:10.1007/s00224-018-9889-7.

[BW19] Hedyeh Beyhaghi and S. Matthew Weinberg. Optimal (and benchmark-optimal) competition complexity for additive buyers over independent items. In Moses Charikar and Edith Cohen, editors, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 686–696. ACM, 2019. doi:10.1145/3313276.3316405.

[CC23] José R. Correa and Andrés Cristi. A constant factor prophet inequality for online combinatorial auctions. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 686–697. ACM, 2023. doi:10.1145/3564246.3585151.

[CCW23] Yang Cai, Ziyun Chen, and Jinzhao Wu. Simultaneous auctions are approximately revenue-optimal for subadditive bidders. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023, Santa Cruz, CA, USA, November 6 - November 9, 2023*. IEEE, 2023.

[CD15] Yang Cai and Constantinos Daskalakis. Extreme value theorems for optimal multidimensional pricing. *Games Econ. Behav.*, 92:266–305, 2015. doi:10.1016/j.geb.2015.02.003.

[CDO+22] Xi Chen, Ilias Diakonikolas, Anthi Orfanou, Dimitris Paparas, Xiaorui Sun, and Mihalis Yannakakis. On the complexity of optimal lottery pricing and randomized mechanisms for a unit-demand buyer. *SIAM J. Comput.*, 51(3):492–548, 2022. doi:10.1137/17m1136481.

[CDP+18] Xi Chen, Ilias Diakonikolas, Dimitris Paparas, Xiaorui Sun, and Mihalis Yannakakis. The complexity of optimal multidimensional pricing for a unit-demand buyer. *Games Econ. Behav.*, 110:139–164, 2018. doi:10.1016/j.geb.2018.03.016.

[CDW21] Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. A duality-based unified approach to bayesian mechanism design. *SIAM J. Comput.*, 50(3), 2021. doi:10.1137/16M1100113.

[CFH+21] José R. Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Posted price mechanisms and optimal threshold strategies for random arrivals. *Math. Oper. Res.*, 46(4):1452–1478, 2021. doi:10.1287/moor.2020.1105.

[CFPV19] José R. Correa, Patricio Foncea, Dana Pizarro, and Victor Verdugo. From pricing to prophets, and back! *Oper. Res. Lett.*, 47(1):25–29, 2019. doi:10.1016/j.orl.2018.11.010.

[CGL14] Ning Chen, Nick Gravin, and Pinyan Lu. Optimal competitive auctions. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 253–262. ACM, 2014. doi:10.1145/2591796.2591855.

[CHK07] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic pricing via virtual valuations. In Jeffrey K. MacKie-Mason, David C. Parkes, and Paul Resnick, editors, *Proceedings 8th ACM Conference on Electronic Commerce (EC-2007), San Diego, California, USA, June 11-15, 2007*,

	single-buyer	multi-buyer
unit-demand	[CHK07], [CHMS10], [CMS15] [CD15], [HH15], [CDW21], [JLQ+19b]	[CMS15], [KW19], [HH15], [CDW21]
additive	[HN17], [LY13], [BILW20] [HH15], [CDW21], [MS21]	[Yao15], [HH15], [CDW21], [DFL+22]
matroid-rank	[CM16], [CZ17]	[CM16], [CZ17]
constrained-additive	[CM16], [CZ17]	[CZ17]
XOS	[RW18], [CZ17]	[CZ17], [Yao18], [COZ22]
subadditive	[RW18], [CZ17]	[CZ17], [DKL20], [CC23], [CCW23]

TABLE I
A SUMMARY OF SOME PREVIOUS WORKS ON MULTI-ITEM MECHANISM DESIGN.

pages 243–251. ACM, 2007. doi:10.1145/1250910.1250946.

[CHMS10] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 311–320. ACM, 2010. doi:10.1145/1806689.1806733.

[CM16] Shuchi Chawla and J. Benjamin Miller. Mechanism design for subadditive agents via an ex ante relaxation. In Vincent Conitzer, Dirk Bergemann, and Yiling Chen, editors, *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 579–596. ACM, 2016. doi:10.1145/2940716.2940756.

[CMPY18] Xi Chen, George Matikas, Dimitris Paparas, and Mihalis Yannakakis. On the complexity of simple and optimal deterministic mechanisms for an additive buyer. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 2036–2049. SIAM, 2018. doi:10.1137/1.9781611975031.133.

[CMS15] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games Econ. Behav.*, 91:297–317, 2015. doi:10.1016/j.geb.2012.08.010.

[COZ22] Yang Cai, Argyris Oikonomou, and Mingfei Zhao. Computing simple mechanisms: Lift-and-round over marginal reduced forms. In Stefano Leonardi and Anupam Gupta, editors, *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 704–717. ACM, 2022. doi:10.1145/3519935.3520029.

[CP14] Yang Cai and Christos H. Papadimitriou. Simultaneous bayesian auctions and computational complexity. In Moshe Babaioff, Vincent Conitzer, and David A. Easley, editors, *ACM Conference on Economics and Computation, EC '14, Stanford , CA, USA, June 8-12, 2014*, pages 895–910. ACM, 2014. doi:10.1145/2600057.2602877.

[CST18] George Christodoulou, Alkmini Sgouritsa, and Bo Tang. On the efficiency of all-pay mechanisms. *Algorithmica*, 80(4):1115–1145, 2018. doi:10.1007/s00453-017-0296-2.

[CZ17] Yang Cai and Mingfei Zhao. Simple mechanisms for subadditive buyers via duality. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 170–183. ACM, 2017. doi:10.1145/3055399.3055465.

[DDT14] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. The complexity of optimal mechanism design. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1302–1318. SIAM, 2014. doi:10.1137/1.9781611973402.96.

[DFL+22] Constantinos Daskalakis, Maxwell Fishelson, Brendan Lucier, Vasilis Syrgkanis, and Santhoshini Velusamy. Multi-item nontruthful auctions achieve good revenue. *SIAM J. Comput.*, 51(6):1796–1838, 2022. doi:10.1137/22m1471742.

[DKL20] Paul Dütting, Thomas Kesselheim, and Brendan Lucier. An o(log log m) prophet inequality for subadditive combinatorial auctions. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 306–317. IEEE, 2020. doi:10.1109/FOCS46700.2020.00037.

[EFF+17] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. The competition complexity of auctions: A bulow-klemperer result for multi-dimensional bidders. In Constantinos Daskalakis, Moshe Babaioff, and Hervé Moulin, editors, *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, page 343. ACM, 2017. doi:10.1145/3033274.3085115.

[Har13] Jason D Hartline. Mechanism design and approximation. *Book draft. October*, 122(1), 2013.

[HH15] Nima Haghpanah and Jason D. Hartline. Reverse mechanism design. In Tim Roughgarden, Michal Feldman, and Michael Schwarz, editors, *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 757–758. ACM, 2015. doi:10.1145/2764468.2764498.

[HJL20] Jason D. Hartline, Aleck C. Johnsen, and Yingkai Li. Benchmark design and prior-independent optimization. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 294–305. IEEE, 2020. doi:10.1109/FOCS46700.2020.00036.

[HN13] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In Michael J. Kearns, R. Preston McAfee, and Éva Tardos, editors, *Proceedings of the fourteenth ACM Conference on Electronic Commerce, EC 2013, Philadelphia, PA, USA, June 16-20, 2013*, pages 565–566. ACM, 2013. doi:10.1145/2492002.2482544.

[HN17] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. *J. Econ. Theory*, 172:313–347, 2017. doi:10.1016/j.jet.2017.09.001.

[HN19] Sergiu Hart and Noam Nisan. Selling multiple correlated goods: Revenue maximization and menu-size complexity. *J. Econ. Theory*, 183:991–1029, 2019. doi:10.1016/j.jet.2019.07.006.

[HR15] Sergiu Hart and Philip J Reny. Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Theoretical Economics*, 10(3):893–922, 2015.

[JJLZ22] Yaonan Jin, Shunhua Jiang, Pinyan Lu, and Hengjie Zhang. Tight revenue gaps among multiunit mechanisms. *SIAM J. Comput.*, 51(5):1535–1579, 2022. doi:10.1137/21m1456364.

[JL22] Yaonan Jin and Pinyan Lu. First price auction is 1 - 1 /e^2 efficient. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 179–187. IEEE, 2022. doi:10.1109/FOCS54457.2022.00024.

[JL23] Yaonan Jin and Pinyan Lu. The price of stability for first price auction. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 332–352. SIAM, 2023. doi:10.1137/1.9781611977554.ch14.

[JLQ19a] Yaonan Jin, Weian Li, and Qi Qi. On the approximability of simple mechanisms for MHR distributions. In Ioannis Caragiannis, Vahab S. Mirrokni, and Evdokia Nikolova, editors, *Web and Internet Economics - 15th International Conference, WINE 2019, New York, NY, USA, December 10-12, 2019, Proceedings*, volume 11920 of *Lecture Notes in Computer Science*, pages 228–240. Springer, 2019. doi:10.1007/978-3-030-35389-6_17.

[JLQ$^+$19b] Yaonan Jin, Pinyan Lu, Qi Qi, Zhihao Gavin Tang, and Tao Xiao. Tight approximation ratio of anonymous pricing. In Moses Charikar and Edith Cohen, editors, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 674–685. ACM, 2019. doi:10.1145/3313276.3316331.

[JLTX20] Yaonan Jin, Pinyan Lu, Zhihao Gavin Tang, and Tao Xiao. Tight revenue gaps among simple mechanisms. *SIAM J. Comput.*, 49(5):927–958, 2020. doi:10.1137/19M126178X.

[KS78] Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, 4:197–266, 1978.

[KSM$^+$19] Pravesh Kothari, Sahil Singla, Divyarthi Mohan, Ariel Schvartzman, and S. Matthew Weinberg. Approximation schemes for a unit-demand buyer with independent items via symmetries. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 220–232. IEEE Computer Society, 2019. doi:10.1109/FOCS.2019.00023.

[KW19] Robert Kleinberg and S. Matthew Weinberg. Matroid prophet inequalities and applications to multi-dimensional mechanism design. *Games Econ. Behav.*, 113:97–115, 2019. doi:10.1016/j.geb.2014.11.002.

[LY13] Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. *Proc. Natl. Acad. Sci. USA*, 110(28):11232–11237, 2013. doi:10.1073/pnas.1309533110.

[MS21] Will Ma and David Simchi-Levi. Reaping the benefits of bundling under high production costs. In Arindam Banerjee and Kenji Fukumizu, editors, *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event*, volume 130 of *Proceedings of Machine Learning Research*, pages 1342–1350. PMLR, 2021. URL: http://proceedings.mlr.press/v130/ma21b.html.

[Mye81] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[RW18] Aviad Rubinstein and S. Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. *ACM Trans. Economics and Comput.*, 6(3-4):19:1–19:25, 2018. doi:10.1145/3105448.

[SC84] Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *the Annals of Probability*, pages 1213–1216, 1984.

[Tha04] John Thanassoulis. Haggling over substitutes. *J. Econ. Theory*, 117(2):217–245, 2004. doi:10.1016/j.jet.2003.09.002.

[Yao15] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In Piotr Indyk, editor, *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 92–109. SIAM, 2015. doi:10.1137/1.9781611973730.8.

[Yao18] Andrew Chi-Chih Yao. On revenue monotonicity in combinatorial auctions. In Xiaotie Deng, editor, *Algorithmic Game Theory - 11th International Symposium, SAGT 2018, Beijing, China, September 11-14, 2018, Proceedings*, volume 11059 of *Lecture Notes in Computer Science*, pages 1–11. Springer, 2018. doi:10.1007/978-3-319-99660-8_1.