On Infinite Separations Between Simple and Optimal Mechanisms

Alexandros Psomas [*]   Ariel Schvartzman [†]   S. Matthew Weinberg [‡]

Abstract

We consider a revenue-maximizing seller with k heterogeneous items for sale to a single additive buyer, whose values are drawn from a known, possibly correlated prior $\mathcal{D}$. It is known that there exist priors $\mathcal{D}$ such that simple mechanisms — those with bounded menu complexity — extract an arbitrarily small fraction of the optimal revenue ([BCKW15, HN19]). This paper considers the opposite direction: given a correlated distribution $\mathcal{D}$ witnessing an infinite separation between simple and optimal mechanisms, what can be said about $\mathcal{D}$?

[HN19] provides a framework for constructing such $\mathcal{D}$: it takes as input a sequence of k-dimensional vectors satisfying some geometric property, and produces a $\mathcal{D}$ witnessing an infinite gap. Our first main result establishes that this framework is without loss: *every* $\mathcal{D}$ witnessing an infinite separation could have resulted from this framework. An earlier version of their work provided a more streamlined framework [HN13]. Our second main result establishes that this restrictive framework is *not* tight. That is, we provide an instance $\mathcal{D}$ witnessing an infinite gap, but which provably could not have resulted from the restrictive framework.

As a corollary, we discover a new kind of mechanism which can witness these infinite separations on instances where the previous "aligned" mechanisms do not.

1 Introduction

Consider a revenue-maximizing seller with k items for sale to a single additive buyer, whose values for the k items are drawn from a known distribution $\mathcal{D}$. When $k = 1$, Myerson's seminal work provides a closed-form solution to the revenue-optimal mechanism, and it has a particularly simple form: simply post a price $p := \arg\max_p \{p \cdot \Pr_{v \leftarrow \mathcal{D}}[v \geq p]\}$, and let the buyer purchase the item if they please ([Mye81]). For $k > 1$, however, this *multi-dimensional mechanism design* problem remains an active research agenda forty years later.

While simple, constant-factor approximations are known in quite general settings when $\mathcal{D}$ is a product distribution ([CHK07, CHMS10, CMS15, KW12, LY13, BILW14, Yao15, RW15, CDW16, CM16, CZ17]), there may be an *infinite* gap between the revenue-optimal auction and any simple counterpart when values are correlated ([BCKW15, HN19]). More specifically: one simple way to sell k items is to treat the grand bundle of all items as if it were a single item, and sell it using Myerson's optimal

---

[*]Department of Computer Science, Purdue University, apsomas@cs.purdue.edu. Supported in part by an NSF CAREER award CCF-2144208, a Google Research Scholar Award, and by the Algorand Centres of Excellence program managed by Algorand Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of Algorand Foundation.

[†]Center for Discrete Mathematics and Theoretical Computer Science (DIMACS), schvartman.ariel@gmail.com. Research conducted while at DIMACS, Rutgers University and supported by the National Science Foundation under grant number CCF-1445755.

[‡]Department of Computer Science, Princeton University, smweinberg@princeton.edu. Supported by NSF CCF-1717899

auction (which sets price $\arg\max_p\{p \cdot \Pr_{\vec{v}\leftarrow\mathcal{D}}[\sum_i v_i \geq p]\}$). Letting $\mathrm{BRev}(\mathcal{D})$ denote the revenue achieved by this simple scheme, [HN19] further show a connection between $\mathrm{BRev}(\mathcal{D})$ and *any* simple mechanism through the lens of menu complexity: any mechanism with menu complexity at most C generates expected revenue at most $C \cdot \mathrm{BRev}(\mathcal{D})$.

These works establish a strong separation between simple and optimal auctions: *even when $k = 2$, there exist distributions $\mathcal{D}$ such that $\mathrm{Rev}(\mathcal{D}) = \infty$ (the optimal revenue) while $\mathrm{BRev}(\mathcal{D}) = 1$.* The fact that $\mathrm{Rev}(\mathcal{D}) = \infty$ does not on its own suggest that $\mathcal{D}$ must be "weird" (the one-dimensional distribution with CDF $1 - 1/\sqrt{x}$ has this property). The "weird" property is that $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$, which can never occur for a one-dimensional distribution.

[BCKW15] and [HN19] establish sufficient conditions for a distribution $\mathcal{D}$ to satisfy $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$. Simply put, the goal of this paper is to study *necessary* conditions for a distribution to satisfy $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$. We provide two main results. The first establishes that the sufficient condition presented in [HN19] *is in fact necessary* for $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$ (Theorem 6). The second establishes that the sufficient condition used in an earlier version of that work [HN13] and [BCKW15] is *not necessary* (Theorem 9). In establishing Theorem 9, we also construct a distribution $\mathcal{D}$ such that $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$ yet the mechanism witnessing this provably falls outside the scope of any previous constructions (Corollary 11).

Proper context and formal statements of our results require precise definitions, which we provide in Section 2 immediately below. Section 3 provides formal statements of our results, along with context alongside related work. Subsequent sections provide proofs.

2 Preliminaries

We consider an auction design setting with a single buyer, single seller, and k heterogeneous items. Note that our positive results hold for arbitrary k, while our constructions use only $k = 2$ (and $k = 1$ is not possible). We use $\mathcal{D}$ to denote a distribution over $\mathbb{R}^k_{\geq 0}$, the (possibly correlated) distribution over the buyer's values for the k items. The buyer is additive, meaning that their value for a set of items S is equal to $\sum_{i \in S} v_i$. We use $\mathrm{Rev}(\mathcal{D})$ to denote the optimal expected revenue achievable by any incentive-compatible mechanism (formally, the supremum of expected revenues, or ∞ if the supremum is undefined), and let $\mathrm{BRev}(\mathcal{D})$ denote the revenue achieved by selling the grand bundle as a single item using Myerson's auction.[4] Finally, a *mechanism M* is a set $\{(\vec{q}_i, p_i)\}_i$, where each $\vec{q}_i \in [0,1]^k$ denotes a vector of probabilities, and $p_i \in \mathbb{R}$ denotes a price. When the buyer's valuation is $\vec{v}$, they pay the auctioneer $p_{i(\vec{v})}$, where $i(\vec{v}) := \arg\max_i\{\vec{v} \cdot \vec{q}_i - p_i\}$.[5] $\mathrm{Rev}(\mathcal{D}, M)$ denotes the expected revenue achieved by a particular mechanism M on distribution $\mathcal{D}$. We will also use the shorthand $\vec{q}^M(\vec{v})$ to denote the allocation vector purchased by a vector $\vec{v}$, and $p^M(\vec{v})$ to denote the price paid (we may drop the superscript of M if the mechanism is clear from context).

Brief Overview of [HN19]. Below, we formally define two geometric properties of sequences of points, which are the focus of this paper. Many of the ideas below appear in both [BCKW15] and [HN19], but we will use the formal definitions from [HN19] (and an earlier published version [HN13]). Below, morally the vectors $\vec{x}_i$ correspond to possible (scaled) valuation vectors, and the vectors $\vec{q}_i$ correspond to possible vectors of allocation probabilities.

Definition 1. *Let $X = (\vec{x}_i)_{i=1}^N$ be an ordered sequence of N points (N may be finite, or equal to $+\infty$), with each $\vec{x}_i \in \mathbb{R}^k_{\geq 0}$. Let $Q = (\vec{q}_i)_{i=0}^N$ be another ordered sequence of N points, with each $\vec{q}_i \in [0,1]^k$, and starting with $\vec{q}_0 = (0,...,0)$. Define the following:*

$$\mathrm{gap}_i^{X,Q} := \min_{0 \leq j < i}(\vec{q}_i - \vec{q}_j) \cdot \vec{x}_i \qquad \mathrm{MenuGap}(X, Q) := \sum_{i=1}^N \frac{\mathrm{gap}_i^{X,Q}}{||\vec{x}_i||_1}.$$

---

[4] We briefly remind the reader that $\mathrm{BRev}(\mathcal{D})$ serves as a proxy for the achievable revenue by any simple mechanism, especially when focusing on the gap between infinite and finite. For example, the revenue achieved by selling the items separately is at most $k\mathrm{BRev}(\mathcal{D})$, the revenue achieved by any deterministic mechanism is at most $2^k\mathrm{BRev}(\mathcal{D})$, and, more generally, the revenue achieved by any mechanism which offers at most m distinct options is at most $m\mathrm{BRev}(\mathcal{D})$ [HN19].

[5] How ties are broken is irrelevant to our results — all results hold for arbitrary tie-breaking. Also, all mechanisms include an all-zero pair $\vec{q}_0 = (0, \dots, 0)$, $p_0 = 0$ to insure individual rationality.

We will also slightly abuse notation and define $\mathrm{MenuGap}(X) := \sup_Q\{\mathrm{MenuGap}(X, Q)\}$.[6]

$\mathrm{MenuGap}(X)$ is some measure of how distinct the angles of points in X are. To get intuition for this, one might try to write a short proof that when $k = 1$, $\mathrm{MenuGap}(X) = 1$ for all X (or that $\mathrm{MenuGap}(X) = 1$ whenever all $\vec{x}_i \in X$ are parallel). We provide such a proof in Appendix A. $\mathrm{MenuGap}(X, Q)$ is just some geometric measure with no obvious intuition for why this quantity should be of interest to auction designers. However, one key result of [HN19] shows that this quantity has connections to simplicity vs. optimality gaps. Specifically, they show:

Theorem 2 ([HN19], Proposition 7.1). *For every pair of sequences* $X = (\vec{x}_i)_{i=1}^N, Q = (\vec{q}_i)_{i=0}^N$ *starting with* $\vec{q}_0 = (0, ..., 0)$, *and all* $\varepsilon > 0$, *there exists a distribution* $\mathcal{D}$ *and mechanism* M *s.t.:*

$$\frac{\mathrm{Rev}(\mathcal{D}, M)}{\mathrm{BRev}(\mathcal{D})} \geq (1 - \varepsilon) \cdot \mathrm{MenuGap}(X, Q).$$

Moreover, for all $i \in [N]$, *the support of* $\mathcal{D}$ *contains a single point of the form* $c_i \vec{x}_i$, *for some* $c_i \in \mathbb{R}_+$ *(and no other points). Additionally,* $\vec{q}^M(c_i \vec{x}_i) = \vec{q}_i$.

The "Moreover,..." portion of Theorem 2 gives some insight to their construction. Further insight can be deduced by observing that the constraint "$\mathrm{gap}_i^{X,Q} \leq \vec{x}_i \cdot (\vec{q}_i - \vec{q}_j)$" looks similar (but far from identical) to an incentive compatibility constraint involving a valuation vector $\vec{x}_i$ and two allocation vectors $\vec{q}_i, \vec{q}_j$. We refer the reader to [HN19] for further details and intuition for this connection. Theorem 2 gives a framework for proving simplicity vs. optimality gaps, but leaves open the question of actually finding a pair of sequences X, Q. They approach this through the following observation:

Definition 3. *Given a sequence of points* $X = (\vec{x}_i)_{i=0}^N \in [0, 1]^k$ *starting with* $\vec{x}_0 = (0, ..., 0)$, *define* $\mathrm{SupGap}(X)$ *(read "support gap of* X") *as* $\mathrm{SupGap}(X) := \mathrm{MenuGap}(X, X)$.

Observation 4. *For all* X, $\mathrm{MenuGap}(X) \geq \mathrm{SupGap}(X)$.

Finally, [HN13] propose an explicit construction of a sequence X with infinite SupGap.

Theorem 5 ([HN13]). *There exists an infinite sequence* X *of points in* $[0, 1]^2$ *s.t.* $\mathrm{SupGap}(X) = \infty$.

Theorems 2 and 5 together yield a two-dimensional distribution with $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$. It is also worth noting that all prior constructions follow this approach as well. For example, [BCKW15] provides an infinite sequence X of points in $[0, 1]^3$ such that $\mathrm{SupGap}(X) = \infty$.

3 Our Results

Independent of Observation 4 and Theorem 5, Theorem 2 *alone* provides a framework for constructing distributions $\mathcal{D}$ so that $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D})$ is large: find sequences X so that $\mathrm{MenuGap}(X)$ is large. Our goal is to understand to what extent this framework is *complete* for constructing such instances. Our first main result establishes that *any distribution with* $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$ *could have resulted from the framework induced by Theorem 2*. Specifically:

Theorem 6. *For any distribution* $\mathcal{D}$ *over* k *items, there exists a sequence of* N *points (*N *can be finite, or equal to* $+\infty$*)* $X = (\vec{x}_i)_{i=1}^N$, *with each* $\vec{x}_i \in supp(\mathcal{D})$, *such that*

$$\mathrm{MenuGap}(X) \geq \frac{\mathrm{Rev}(\mathcal{D})}{9\mathrm{BRev}(\mathcal{D})}.$$

In particular, if $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$, *then* $\mathrm{MenuGap}(X) = \infty$ *as well.*

A complete proof of Theorem 6 appears in Section 4. Observe also that because $\mathrm{MenuGap}(X)$ is monotone (in the sense that adding points to X, anywhere, cannot possibly decrease $\mathrm{MenuGap}(X)$), the fact that X is a subset of the support of $\mathcal{D}$ (rather than the entire support) is immaterial.[7] Put

---

[6] Observe that $\mathrm{gap}_i^{X,Q}$ and $\mathrm{MenuGap}(X, Q)$ might be negative. Any claims made throughout this paper regarding $\mathrm{MenuGap}(X, Q)$ are vacuously true when $\mathrm{MenuGap}(X, Q) < 0$. We allow $\mathrm{gap}_i^{X,Q}$ to be negative to match the definition of [HN19] verbatim (although our work will also show that this peculiarity of their definition is not significant).

[7] Of course, if the support of $\mathcal{D}$ is uncountable, then clearly the entire support of $\mathcal{D}$ cannot be included in X.

another way, the important aspect in constructing X is how elements in the support of $\mathcal{D}$ are ordered, rather than which points are included.

Observation 4 further provides a framework to construct sequences so that $\mathrm{MenuGap}(X)$ is large: construct sequences X so that $\mathrm{SupGap}(X)$ is large. One may then wonder if $\mathrm{SupGap}(X)$ and $\mathrm{MenuGap}(X)$ are approximately related, for all X. For this specific question, the answer is trivially no, due to incompatibility with scaling (multiplying every point in X by $1/2$ will decrease $\mathrm{SupGap}(X)$ by a factor of 2, but not $\mathrm{MenuGap}(X)$). Therefore, not much insight is gained by studying this precise question.

Instead, we observe that the interesting aspect of constructions resulting through $\mathrm{SupGap}(X)$ is that $\vec{x}_i$ and $\vec{q}_i$ are *aligned* (that is, $\vec{x}_i = c_i \cdot \vec{q}_i$ for some $c_i \in \mathbb{R}_{\geq 0}$). Specifically, even if $\vec{q}_i = \vec{x}_i$, this equality is not maintained by the construction of Theorem 2. However, if $\vec{q}_i$ and $\vec{x}_i$ are aligned, this alignment property *is* maintained by the construction. We therefore propose the following definition, which captures the maximum value achievable by $\mathrm{MenuGap}(X, Q)$ when X, Q are aligned.

Definition 7. *Let $X = (\vec{x}_i)_{i=1}^N$ be an ordered sequence of N points in $[0, 1]^k$ (N may be finite, or equal to $+\infty$). Let also $C = (c_i)_{i=0}^N$ be an ordered sequence of numbers, with each $c_i \in [0, 1/||\vec{x}_i||_\infty]$, starting with $c_0 = 0$. Define:*

$$\mathrm{sgap}_i^{X,C} := \min_{j<i} \vec{x}_i \cdot (c_i \vec{x}_i - c_j \vec{x}_j), \qquad and \qquad \mathrm{AlignGap}(X, C) := \sum_{i=1}^N \frac{\max\{0, \mathrm{sgap}_i^{X,C}\}}{||\vec{x}_i||_1}.$$

We will also slightly abuse notation and denote by $\mathrm{AlignGap}(X) := \sup_C \{\mathrm{AlignGap}(X, C)\}$.

Recall that we have chosen to let c_i range in $[0, 1/||\vec{x}_i||_\infty]$ (rather than be fixed at 1, or $1/||\vec{x}_i||_\infty$) to give potential constructions flexibility in scaling $\vec{q}_i$. Additionally, by ensuring that the contribution of each $\mathrm{sgap}_i^{X,C}$ is non-negative, we give potential constructions flexibility to ignore points in the sequence. That is, any construction using MenuGap directly can always set $\vec{q}_i := \arg\max_{j<i}\{\vec{q}_j \cdot \vec{x}_i\}$, which effectively just drops $\vec{x}_i$ from the sequence. Counting $\max\{0, \mathrm{sgap}_i^{X,C}\}$ towards the objective (rather than just $\mathrm{sgap}_i^{X,C}$) gives constructions that arise through AlignGap the same flexibility.

Lemma 8. *For all X, $\mathrm{AlignGap}(X) \leq \mathrm{MenuGap}(X)$.*

The proof of Lemma 8 is in Appendix B. Lemma 8 induces a framework to design sequences with large MenuGap: design sequences with large AlignGap. Our second main result establishes that this framework *is not* without loss of generality, even for $k = 2$. Specifically:

Theorem 9. *There exist sequences $X = (\vec{x}_i)_{i=1}^\infty \in [0, 1]^2$ such that:*

$$\mathrm{AlignGap}(X) \leq 6 \qquad but \qquad \mathrm{MenuGap}(X) = \infty.$$

A complete proof of Theorem 9 appears in Section 5. By the discussion following Definition 7, the source of this gap is entirely due to the requirement that the sequence Q be aligned with X (and is not due to inability to scale, or inability to ignore difficult points in X). We make this crisp with the following definition and corollary, which construct a novel distribution witnessing $\mathrm{Rev}(D)/\mathrm{BRev}(D) = \infty$ that is provably distinct from all previous approaches.

Definition 10. *For a distribution $\mathcal{D}$ and mechanism M, define the* Aligned Revenue *of M on $\mathcal{D}$:*

$$\mathrm{ARev}(\mathcal{D}, M) := \mathbb{E}_{\vec{v} \leftarrow \mathcal{D}}[p^M(\vec{v}) \cdot I(\vec{v} \text{ is parallel to } \vec{q}^M(\vec{v}))], \qquad \mathrm{ARev}(\mathcal{D}) := \sup_M \{\mathrm{ARev}(\mathcal{D}, M)\}.$$

Corollary 11. *There exist distributions $\mathcal{D}$ over two items such that $\mathrm{Rev}(\mathcal{D})/\mathrm{ARev}(\mathcal{D}) = \infty$.*

A proof of Corollary 11 appears in Appendix C. It is worth noting that *all* previous constructions establishing $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$ proceeded by producing an X such that $\mathrm{SupGap}(X) = \infty$. Indeed, the [BCKW15] construction provides such an X when $k = 3$, the [HN13] construction provides an X when $k = 2$, and the [PSW19] construction adapts parameters in that of [HN13]. By Theorem 2, this implies not only that $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$, but also that $\mathrm{ARev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$. Corollary 11 establishes the existence of a fundamentally different construction,[8] as our $\mathcal{D}$ has a finite ratio between $\mathrm{ARev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D})$, yet still maintains an infinite ratio between $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D})$.

---

[8]On a technical level, our construction certainly borrows several ideas from previous ones, however.

Additional Related Work.   We've already discussed the most related work to ours, which is that of [HN19, BCKW15]. There is also a large body of work studying *product distributions* specifically, and establishes that simple mechanisms can achieve constant factor approximations in quite general settings [CHK07, CHMS10, CMS15, KW12, LY13, BILW14, Yao15, RW15, CDW16, CM16, CZ17]). Recent works have made progress in obtaining arbitrary approximations ([BGN17, KMS$^+$19]), which again rely on the assumption that $\mathcal{D}$ is a product distribution.

Three recent lines of work address the [BCKW15, HN19] constructions in a different manner. First, [CTT19, CTT20] consider the related *buy-many* model (where the auctioneer cannot prevent the buyer from interacting multiple times with the auction). [CTT19] establishes that selling separately achieves an $O(\log k)$-approximation to the optimal buy-many mechanism in *quite* general settings (including the settings considered in this paper). In a different direction, [PSW19] uses the lens of smoothed analysis ([ST04]) to reason about the robustness of the [HN19] constructions. Finally, [Car17] considers a correlation robust framework in which the valuation profile is drawn from a correlated distribution that is not completely known to the seller; the goal is to design a mechanism that maximizes the worst-case (over correlations) seller revenue, when only the items' marginal distributions are known. [Car17] shows that selling each item separately is optimal; see [BGLT19, GL18] for further work in this model.

4   A Converse to Theorem 2: The [HN19] Framework is WLOG

In this section, we prove Theorem 6. Our proof has two main parts. First, we will take the optimal auction for $\mathcal{D}$ (or one that is arbitrarily close to optimal) and repeatedly simplify it through a sequence of lemmas, at the cost of small constant-factors of revenue. The second part takes this simple menu and draws a connection to MenuGap. Missing proofs can be found in Appendix B.

Simplifying the Optimal Mechanism.   We show that for every $\mathcal{D}$, there exists an approximately-optimal mechanism which satisfies some useful properties. First, we argue that we may ignore menu options with low prices.

Definition 12. *We say that a mechanism M is c-expensive if every option has price at least c.*

Claim 13. *For all $c \in \mathbb{R}_{\geq 0}$, all distributions $\mathcal{D}$, and all mechanisms M, there exists a c-expensive mechanism M' satisfying $\mathrm{Rev}(\mathcal{D}, M') \geq \mathrm{Rev}(\mathcal{D}, M) - c$.*

Our next step will show that we can assume further structure on the prices charged, at the cost of a factor of two.

Definition 14. *A c-expensive mechanism M is oddly-priced (respectively, evenly-priced) if for all $\vec{v}$, there exists an odd (respectively, even) integer i such that $p^M(\vec{v}) \in [c \cdot 2^i, c \cdot 2^{i+1})$.*

Claim 15. *For all c-expensive mechanisms M, and all $\mathcal{D}$, there exists either an oddly-priced or evenly-priced c-expensive mechanism M' satisfying $\mathrm{Rev}(\mathcal{D}, M') \geq \mathrm{Rev}(\mathcal{D}, M)/2$.*

This concludes our simplification of the mechanism. In the subsequent section, we draw a connection between MenuGap and the revenue of oddly- or evenly-price c-expensive mechanisms.

Connecting Structured Mechanisms to MenuGap.   We begin with the following definition, which describes our proposed X, Q based on a structured mechanism for $\mathcal{D}$.

Definition 16 (Representative Sequences)**.** *Let M be a c-expensive mechanism which is oddly-priced or evenly-priced, and let $\mathcal{D}$ be any distribution. An ε-representative sequence for $M, \mathcal{D}$ is the following:*

- *Define $\vec{q}_0(\mathcal{D}, M) := (0, \ldots, 0)$.*

- *Define offset a to be 1 if M is oddly-priced, and 0 if M is evenly-priced.*

- *For all $j \in \mathbb{N}_+$, define $B_j := \{\vec{v} \in supp(\mathcal{D}) : p^M(\vec{v}) \in [c \cdot 2^{2(j-1)+a}, c \cdot 2^{2(j-1)+a+1})\}$.*

- *For all $j \in \mathbb{N}_+$, let $\vec{x}_j \in B_j$ be such that $||\vec{x}_j||_1 \leq ||\vec{v}||_1 \cdot (1 + \varepsilon)$ for all $\vec{v} \in B_j$.*[9]

---

[9]Note that for all $\varepsilon > 0$, such an $\vec{x}_j$ exists, even if B_j is not closed (as long as B_j is non-empty). In particular, because M is c-expensive, we know that $||\vec{v}||_1 \geq c > 0$ for all $\vec{v}$ who pay $\geq c$. If B_j is empty, instead omit $\vec{x}_j, \vec{q}_j$ from both lists (i.e. decrease all future indices by one).

- *For all $j \in \mathbb{N}_+$, let $\vec{q}_j := \vec{q}^M(\vec{x}_j)$.*

Proposition 17. *Let M be a c-expensive, and oddly- or evenly-priced. Let (X, Q) be an ε-representative sequence for $M, \mathcal{D}$. Then:* $\mathrm{MenuGap}(X, Q) \geq \frac{\mathrm{Rev}(\mathcal{D}, M)}{4(1+\varepsilon)\mathrm{BRev}(\mathcal{D})}$.

Proof. The proof will follow from two technical claims. The first claim relates $||\vec{x}_i||_1$ and $\mathrm{BRev}(\mathcal{D})$. The second claim relates $\mathrm{Rev}(\mathcal{D}, M)$ and $\mathrm{MenuGap}(X, Q)$. Crucially, this claim uses the fact that the mechanism is either oddly-priced or evenly-priced (and therefore $p^M(\vec{x}_i), p^M(\vec{x}_j)$ differ by at least a factor of 2, for any $i \neq j$).

Claim 18. $(1 + \varepsilon) \cdot \mathrm{BRev}(\mathcal{D}) \geq ||\vec{x}_i||_1 \cdot \Pr_{\vec{v} \sim \mathcal{D}}[\vec{v} \in B_i]$.

Claim 19. *If X, Q is an ε-representative sequence for $M, \mathcal{D}$, then $\mathrm{gap}_i^{X,Q} \geq p^M(\vec{x}_i)/2$.*

With these two claims, we can complete the proof of Proposition 17:

$$
\begin{aligned}
\mathrm{MenuGap}(X, Q) &:= \sum_{i=1}^{\infty} \frac{\mathrm{gap}_i^{X,Q}}{||\vec{x}_i||_1} \\
&\geq^{(Claim\ 19)} \sum_{i=1}^{\infty} \frac{p^M(\vec{x}_i)}{2 \cdot ||\vec{x}_i||_1} \\
&\geq^{(Claim\ 18)} \sum_{i=1}^{\infty} \frac{p^M(\vec{x}_i) \cdot \Pr_{\vec{v} \sim \mathcal{D}}[\vec{v} \in B_i]}{2 \cdot (1+\varepsilon)\mathrm{BRev}(\mathcal{D})} \\
&\geq \sum_{i=1}^{\infty} \frac{\mathbb{E}_{\vec{v} \sim \mathcal{D}|\vec{v} \in B_i}[p^M(\vec{v})] \cdot \Pr_{\vec{v} \sim \mathcal{D}}[\vec{v} \in B_i]}{4 \cdot (1+\varepsilon)\mathrm{BRev}(\mathcal{D})} \\
&\geq \frac{\mathrm{Rev}(\mathcal{D}, M)}{4 \cdot (1+\varepsilon)\mathrm{BRev}(\mathcal{D})}.
\end{aligned}
$$

Above, the first line is the definition of MenuGap. The fourth line uses the fact that for all $\vec{v} \in B_i$, $p^M(\vec{v}) \leq 2^{2(i-1)+a+1} = 2 \cdot 2^{2(i-1)+a} \leq 2p^M(\vec{x}_i)$. The final line is just rewriting the definition of $\mathrm{Rev}(\mathcal{D}, M)$, and using the fact that every $\vec{v} \in \mathrm{supp}(\mathcal{D})$ with $p(\vec{v}) > 0$ is in some B_i. $\square$

Wrapping Up. Proposition 17, together with the fact that an oddly- or evenly-priced mechanism is guaranteed to get a good approximation to the optimum, now suffices to prove Theorem 6.

Proof of Theorem 6. Set $c = \mathrm{Rev}(\mathcal{D})/100$, and $\varepsilon = 1/100$. Note that ε-representative sequences are guaranteed to exist for any $M, \mathcal{D}$, as $\varepsilon > 0$. Let M' denote the c-expensive mechanism promised by Claim 13. Let then M'' denote the oddly- or evenly-priced mechanism promised by Claim 15. Finally, let X, Q denote the ε-representative sequence for $M'', \mathcal{D}$. We get:

$$
\begin{aligned}
\mathrm{MenuGap}(X, Q) &\geq^{(\text{Proposition } 17)} \frac{\mathrm{Rev}(\mathcal{D}, M'')}{4(1+\varepsilon)\mathrm{BRev}(\mathcal{D})} \\
&\geq^{(\text{Claim } 15)} \frac{\mathrm{Rev}(\mathcal{D}, M')}{8(1+\varepsilon)\mathrm{BRev}(\mathcal{D})} \\
&\geq^{(\text{Claim } 13)} \frac{\mathrm{Rev}(\mathcal{D}) - c}{8(1+\varepsilon)\mathrm{BRev}(\mathcal{D})} \\
&\geq \frac{\mathrm{Rev}(\mathcal{D})}{9\mathrm{BRev}(\mathcal{D})}.
\end{aligned}
$$
$\square$

5 No Converse to Lemma 8: Separating MenuGap and AlignGap

In this section we prove our second main result: Lemma 8 does *not* admit a converse, even approximately. We briefly remind the reader that *all* previous constructions witnessing $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$ arose by establishing sequences X with $\mathrm{AlignGap}(X) = \infty$ (in fact, even $\mathrm{SupGap}(X) = \infty$). Theorem 9 establishes that constructions exist outside of this more restrictive framework. Missing proofs can be found in Appendix B.

Figure 1: An illustration of two layers of our construction. The number of points in each layer increases with ℓ, but they are always evenly spaced between $(1,0)$ and $(0,1)$. The direction in which the points are placed alternates between clockwise and counterclockwise.

Description of our construction. Our infinite sequence X consists of consecutive *layers* of points. For $\ell = 2$ to ∞, layer ℓ has $n_\ell := \ell\lceil\ln^2(\ell)\rceil + 1$ points. These points/vectors have ℓ_2 norm equal to one, and are evenly spaced (in terms of their angle) between $(1,0)$ and $(0,1)$. If ℓ is even, they go counterclockwise from $(1,0)$ to $(0,1)$. If ℓ is odd, they go clockwise from $(0,1)$ to $(1,0)$:

- Define $n_\ell := \ell\lceil\ln^2(\ell)\rceil + 1$. Define $\theta_\ell := \frac{\pi}{2(n_\ell-1)}$.

- Point $\vec{x}_{\ell,j}$ is the j^{th} point in the ℓ^{th} layer, and is $(\cos(j\theta_\ell), \sin(j\theta_\ell))$ when ℓ is even, or $(\sin(j\theta_\ell), \cos(j\theta_\ell))$ when ℓ is odd.

- The infinite sequence X is $((\vec{x}_{\ell,j})_{j=0}^{n_\ell-1})_{\ell=2}^\infty$.

Figure 1 demonstrates two layers of our construction. In the remainder of this section, we may refer to the sequence of points $\vec{x}_{\ell,j}$ by a single indexed sequence $\vec{x}_i$. The latter is the same as the former where points are ordered lexicographically with respect to the original indexing.

Upper Bounding $\mathrm{AlignGap}(X)$ via Lagrangian duality. We first upper bound $\mathrm{AlignGap}(X)$ and establish that it's finite. To this end, observe that for any sequence X, $\mathrm{AlignGap}(X)$ is the solution to the following (infinite, if X is infinite) mathematical program, where the variables are sgap_i, c_i (the sequence $\vec{x}_i$ is fixed, as we're aiming to compute $\mathrm{AlignGap}(X)$):

$$\mathrm{AlignGap}(X) : \begin{cases} \max \sum_i \max\{0, \mathrm{sgap}_i\}/||\vec{x}_i||_1 \\ \text{subject to:} \quad \forall i, j < i : \mathrm{sgap}_i \leq \vec{x}_i \cdot (c_i\vec{x}_i - c_j\vec{x}_j) \\ \forall i : 0 \leq c_i \leq 1/||\vec{x}_i||_\infty \end{cases}$$

We next proceed with a series of relaxations of this program. Some steps are specific to our choice of X from Section 5, while others hold for arbitrary X. Our first step is specific to this construction, and simply bounds $||\vec{x}||_1$. Consider the following mathematical program:

$$\mathrm{AlignGap}'(X) : \begin{cases} \max \sum_i \max\{0, \mathrm{sgap}_i\} \\ \text{subject to:} \quad \forall i, j < i : \mathrm{sgap}_i \leq \vec{x}_i \cdot (c_i\vec{x}_i - c_j\vec{x}_j) \\ \forall i : 0 \leq c_i \leq \sqrt{2}. \end{cases}$$

Observation 20. *For the sequence X defined in Section 5, $\mathrm{AlignGap}'(X) \geq \mathrm{AlignGap}(X)$.*

Proof. Every $\vec{x}_i$ in the construction has $||\vec{x}_i||_1 \geq ||\vec{x}_i||_2 = 1$. Therefore, the new objective function is only larger. Moreover, every $\vec{x}_i$ in the construction has $||\vec{x}_i||_\infty \geq 1/\sqrt{2}$ (because the ℓ_2 norm is 1), so this is relaxing the upper bound on c_i. $\square$

We proceed to upper bound $\mathrm{AlignGap}'(X)$ via a Lagrangian relaxation of the formulation above. Consider the following Lagrangian relaxation. We put a Lagrangian multiplier of 1 on every constraint of the form $\mathrm{sgap}_i \leq \vec{x}_i \cdot (c_i\vec{x}_i - c_{i-1}\vec{x}_{i-1})$, for all $i > 1$, and a Lagrangian multiplier of 0 on all

other constraints involving sgap. We will *not* put a Lagrangian multiplier on constraints binding c_i to $[0, \sqrt{2}]$, and keep those in the program. This yields the following Lagrangian relaxation (for simplicity of notation below, define $c_0 := 0$, and define $\vec{x}_0 = \vec{0}$):

$$\text{LagRel}_1(\text{X}) : \begin{cases} \max \sum_i \max\{0, \text{sgap}_i\} + \vec{x}_i \cdot (c_i \vec{x}_i - c_{i-1} \vec{x}_{i-1}) - \text{sgap}_i \\ \text{subject to:} \quad \forall i : 0 \leq c_i \leq \sqrt{2} \end{cases}$$

Observation 21. *For all X,* $\text{AlignGap}'(X) \leq \text{LagRel}_1(X)$.

We now proceed to further simplify $\text{LagRel}_1(X)$. The next step is defined below:

$$\text{LagRel}_2(\text{X}) : \begin{cases} \max \sum_i \vec{x}_i \cdot (c_i \vec{x}_i - c_{i-1} \vec{x}_{i-1}) \\ \text{subject to:} \quad \forall i : 0 \leq c_i \leq \sqrt{2} \end{cases}$$

Observation 22. *For all X,* $\text{LagRel}_1(X) = \text{LagRel}_2(X)$.

Finally, we can rewrite the objective function to group all coefficients of c_i. For ease of notation below, define $\vec{x}_{N+1} := \vec{0}$ (if N is finite. If $N = \infty$, there are no notational issues).

$$\text{LagRel}(\text{X}) : \begin{cases} \max \sum_i c_i \vec{x}_i \cdot (\vec{x}_i - \vec{x}_{i+1}) \\ \text{subject to:} \quad \forall i : 0 \leq c_i \leq \sqrt{2} \end{cases}$$

Observation 23. *For all X,* $\text{AlignGap}'(X) \leq \text{LagRel}_1(X) = \text{LagRel}_2(X) = \text{LagRel}(X)$.

Now, we move to analyze $\text{LagRel}(X)$ for our particular sequence X.

Claim 24. *For the sequence X defined in Section 5,* $\text{LagRel}(X) = \sqrt{2} \cdot \sum_i 1 - \vec{x}_i \cdot \vec{x}_{i+1}$.

Proof. For all i, $\vec{x}_i \cdot (\vec{x}_i - \vec{x}_{i+1}) \geq 0$, as each $\vec{x}_i$ has ℓ_2 norm exactly one. Thus, the optimal solution for $\text{LagRel}(X)$ sets each $c_i := \sqrt{2}$. Recalling that $\vec{x}_i \cdot \vec{x}_i = 1$ for all i concludes the claim. □

Proposition 25. *For the sequence X defined in Section 5,* $\text{LagRel}(X) \leq 6$.

Proof. Let us first observe that if $\vec{x}_i$ is the last point in a layer, then in fact $\vec{x}_{i+1} = \vec{x}_i$, and therefore $1 - \vec{x}_i \cdot \vec{x}_{i+1} = 0$. Therefore, these terms do not contribute to the sum. We can then rewrite the term to sum over all layers as follows:

$$\begin{aligned}
\text{LagRel}(X) &= \sqrt{2} \cdot \sum_\ell \sum_{j=0}^{n_\ell - 2} 1 - \vec{x}_{\ell,j} \cdot \vec{x}_{\ell,j+1} \\
&= \sqrt{2} \cdot \sum_\ell (n_\ell - 1) \cdot (1 - \cos(\theta_\ell)) \\
&\leq \sqrt{2} \cdot \sum_\ell (n_\ell - 1) \cdot \theta_\ell^2/2 \\
&= \sqrt{2} \cdot \sum_\ell \frac{\pi^2}{8(n_\ell - 1)} \\
&= \sqrt{2} \cdot \sum_\ell \frac{\pi^2}{8\ell \lceil \ln^2(\ell) \rceil} \\
&\leq 6.
\end{aligned}$$

Above, the first line follows by the reasoning in the first paragraph. The second line follows by observing that the angle between any two points in layer ℓ is exactly θ_ℓ (and the two points in question of ℓ_2 norm equal to one). The third line follows as $\cos(\theta_\ell) \geq 1 - \theta_\ell^2/2$ for any $\theta_\ell \in [0, \pi/2]$ (and all θ_ℓ are indeed in $[0, \pi/2]$). The fourth line follows by substituting the definition of θ_ℓ as a function of n_ℓ. The fifth line follows by definition of n_ℓ. The final line is just calculation for this particular infinite series (this follows as $\sum_{\ell=2}^\infty \frac{1}{\ell \ln^2(\ell)} \leq 3$). □

Observation 20, Observation 23, and Proposition 25 yield the main result of this section:

Proposition 26. *For the sequence X defined in Section 5,* $\text{AlignGap}(X) \leq 6$.

Step 3: Picking a Q to Lower Bound $\mathrm{MenuGap}(X)$. Finally, we propose a sequence Q and show that $\mathrm{MenuGap}(X, Q) = \infty$. We describe the sequence again in layers, to match our description of X (that is, the vector $\vec{q}_{\ell,j}$ corresponds to the vector $\vec{x}_{\ell,j}$). In particular, for each even layer ℓ, the vectors $\vec{q}_{\ell,j}$ will have a fixed x-coordinate, and the y-coordinate will increase with j. For each odd layer, we will introduce no new vectors (i.e. we will just let $\mathrm{gap}_i^{X,Q} = 0$ for all i in an odd layer). Specifically, the construction is as follows:

- Define $\alpha := \sum_{\ell=2}^{\infty} \frac{1}{\ell \ln^2(\ell)}$ (and note that $\alpha < \infty$).

- Define $z_\ell := \frac{1}{\alpha} \sum_{j=2}^{\ell} \frac{1}{j \ln^2(j)}$.

- Define $\delta_\ell := z_\ell - z_{\ell-1} = \frac{1}{\alpha \ell \ln^2(\ell)}$.

- For $j < n_\ell - 1$, define $z_{\ell,j} := 1 - \delta_\ell \cot((j+1)\theta_\ell)$. For $j = n_\ell - 1$, define $z_{\ell,j} := 1$.[10]

- For all even ℓ, and all j, set $\vec{q}_{\ell,j} := (z_\ell, z_{\ell,j})$.

- For all odd ℓ, and all j, set $\vec{q}_{\ell,j} := \arg\max_{\ell' < \ell, j'} \{ \vec{x}_{\ell,j} \cdot \vec{q}_{\ell',j'} \}$.

Proposition 27. $\mathrm{MenuGap}(X, Q) = \infty$.

In our proof (Appendix B) we first analyze which point sets the gap for $\vec{x}_{\ell,j}$, and observe that it must either be $\vec{q}_{\ell,j-1}$ or $\vec{q}_{\ell-2,n_{\ell-2}-1}$ (that is, it must be the previous point in the same layer, or the final point in the previous even layer). We can then bound the gap of each layer exactly, and then take a sum over layers: for any even layers we have that the sum of gaps is at least $\sum_{j=0}^{n_\ell-1} \mathrm{gap}_{\ell,j}^{X,Q} \geq \delta_\ell \cdot \ln(n_\ell)/2$, where $\mathrm{gap}_{\ell,j}^{X,Q} := \mathrm{gap}_i^{X,Q}$, where $\vec{x}_i := \vec{x}_{\ell,j}$ ($\vec{x}_i$ is the j^{th} point on layer ℓ). Recall that $\delta_\ell := \frac{1}{\alpha n_\ell} = \frac{1}{\alpha \ell \ln^2(\ell)}$ we conclude that:

$$\sum_{\ell \text{ even}} \sum_{j=0}^{n_\ell-1} \mathrm{gap}_{\ell,j}^{X,Q} \geq \sum_{\ell \text{ even}} \delta_\ell \cdot \ln(n_\ell)/2 = \sum_{\ell \text{ even}} \frac{1}{2\alpha \ell \ln(\ell)} = \infty.$$

This completes the proof of Theorem 9: Proposition 26 establishes that $\mathrm{AlignGap}(X) \leq 6$, while Proposition 27 establishes that $\mathrm{MenuGap}(X) = \infty$.

6 Conclusion

We study the nature of distributions $\mathcal{D}$ with $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$. Prior work established a framework to construct such distributions, and therefore established sufficient conditions [BCKW15, HN19]. Our first main result establishes that the most general of these frameworks is in fact complete (Theorem 6). Our second main result establishes that the more restrictive framework, through which all previous constructions arose, is not complete (Theorem 9). Finally, we build upon our main construction to develop a novel distribution $\mathcal{D}$ witnessing $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$, but for which none of the "aligned" mechanisms of prior work can possibly witness this (Corollary 11).

In terms of future work, it remains open as to whether there is an alternative definition for "sequences with $\mathrm{MenuGap}(X) = \infty$" which is easier to parse. Our work establishes that understanding such sequences is necessary and sufficient to understand distributions with $\mathrm{Rev}(\mathcal{D})/\mathrm{BRev}(\mathcal{D}) = \infty$, and establishes that "sequences with $\mathrm{MenuGap}(X) = \infty$" is not equivalent to "sequences with $\mathrm{AlignGap}(X) = \infty$." But it would be exciting for future work to better understand sequences with $\mathrm{MenuGap}(X) = \infty$.

Interestingly, ongoing work by [AS22] uses the framework developed in this paper in order to prove *approximation results* for so-called fine-grained buy-many mechanisms, a class of mechanisms which interpolates between buy-one and the recently introduced buy-many mechanisms [CTT19, CTT20]. In general, our techniques make it possible for future work to explore the gap between various simple versus optimal benchmarks without having to reason directly about the underlying mechanisms.

---

[10]To see that $z_{\ell,j} \geq 0$ observe that $\cot(x) \leq 1/x$ for all $x \in [0, \pi/2]$. Therefore we get $\delta_\ell \cot((j+1)\theta_\ell) \leq \delta_\ell \frac{1}{(j+1)\theta_\ell} \leq \delta_\ell \frac{2(n_\ell-1)}{\pi(j+1)} \leq \frac{2}{\pi} \delta_\ell n_\ell = \frac{2\ell\lceil \log^2 \ell \rceil}{\pi \alpha \ell \log^2 \ell} \leq 1$, for all ℓ, j ($\alpha > 1.9$). To see that $z_{\ell,j} \leq 1$ simply note that the term we subtract can't be negative.

References

[AS22] Sepehr Assadi and Ariel Schvartzman. Fine-grained buy-many mechanisms are not much better than bundling. *Manuscript from author's website.*, 2022.

[BCKW15] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S Matthew Weinberg. Pricing lotteries. *Journal of Economic Theory*, 156:144–174, 2015.

[BGLT19] Xiaohui Bei, Nick Gravin, Pinyan Lu, and Zhihao Gavin Tang. Correlation-robust analysis of single item auction. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 193–208. SIAM, 2019.

[BGN17] Moshe Babaioff, Yannai A. Gonczarowski, and Noam Nisan. The menu-size complexity of revenue approximation. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 869–877, 2017.

[BILW14] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 21–30. IEEE, 2014.

[Car17] Gabriel Carroll. Robustness and separation in multidimensional screening. *Econometrica*, 85(2):453–488, 2017.

[CDW16] Yang Cai, Nikhil R Devanur, and S Matthew Weinberg. A duality-based unified approach to bayesian mechanism design. *ACM SIGecom Exchanges*, 15(1):71–77, 2016.

[CHK07] Shuchi Chawla, Jason D. Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, EC '07, pages 243–251, New York, NY, USA, 2007. ACM.

[CHMS10] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320. ACM, 2010.

[CM16] Shuchi Chawla and J Benjamin Miller. Mechanism design for subadditive agents via an ex ante relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 579–596. ACM, 2016.

[CMS15] Shuchi Chawla, David Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015.

[CTT19] Shuchi Chawla, Yifeng Teng, and Christos Tzamos. Buy-many mechanisms are not much better than item pricing. In Anna Karlin, Nicole Immorlica, and Ramesh Johari, editors, *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*, pages 237–238. ACM, 2019.

[CTT20] Shuchi Chawla, Yifeng Teng, and Christos Tzamos. Menu-size complexity and revenue continuity of buy-many mechanisms. *CoRR*, abs/2003.10636, 2020.

[CZ17] Yang Cai and Mingfei Zhao. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 170–183. ACM, 2017.

[GL18] Nick Gravin and Pinyan Lu. Separation in correlation-robust monopolist problem with budget. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2069–2080. SIAM, 2018.

[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.

[HN19] Sergiu Hart and Noam Nisan. Selling multiple correlated goods: Revenue maximization and menu-size complexity. *Journal of Economic Theory*, 183:991–1029, 2019.

[KMS+19] Pravesh Kothari, Divyarthi Mohan, Sahil Singla, 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.

[KW12] Robert Kleinberg and Seth Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 123–136. ACM, 2012.

[LY13] Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28):11232–11237, 2013.

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

[PSW19] Alexandros Psomas, Ariel Schvartzman, and S Matthew Weinberg. Smoothed analysis of multi-item auctions with correlated values. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 417–418. ACM, 2019.

[RW15] Aviad Rubinstein and S Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 377–394. ACM, 2015.

[ST04] Daniel A Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 81–90. ACM, 2004.

[Yao15] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 92–109. Society for Industrial and Applied Mathematics, 2015.

A MenuGap$(X) = 1$ when $k = 1$

In this brief section we prove that when $k = 1$, for any sequence of $x_i \in \mathbb{R}^+_{\geq 0}$, MenuGap$(X) = 1$.

Claim 28. *When $k = 1$, for any $X = \{x_i\}_{i=1}^N$, $x_i \in \mathbb{R}^+_{\geq 0}$, MenuGap$(X) = 1$.*

Proof. Note that when $k = 1$, $||x_i||_1 = x_i$. Therefore, MenuGap$(X, Q) = \sum_i \min_{j<i}(q_i - q_j)$. We make the following observation which allows us to look at structured optimal solutions.

Observation 29. *Any optimal solution Q to MenuGap(X) is monotone non-decreasing.*

Proof. For the sake of contradiction, suppose we are given an optimal solution Q that is not monotone non-decreasing. Let i be the smallest index for which $q_i < q_{i-1}$. Then $\text{gap}_i^{X,Q} = (q_i - q_{i-1})x_i < 0$. Consider instead a solution Q' where $q'_j = q_j$ for all $j \neq i$ and $q'_i = q_{i-1}$. Now, $\text{gap}_i^{X,Q'} = 0$. Since $Q_{\leq i-1} = Q'_{\leq i-1}$, $\text{gap}_j^{X,Q} = \text{gap}_j^{X,Q'}$ for all $j < i$. Since $q_{i-1} > q_i$, for any $j > i$ it holds that $(q_j - q_{i-1}) < (q_j - q_i)$. Therefore, q_i is not "setting the gap" for any point after it. Hence it also holds that $\text{gap}_j^{X,Q} = \text{gap}_j^{X,Q'}$ for all $j > i$. Putting everything together we get that MenuGap$(X, Q') - $ MenuGap$(X, Q) = \text{gap}_i^{X,Q'} - \text{gap}_i^{X,Q} > 0$ contradicting the optimality of Q. □

With this observation in hand, since the q_i are monotone non-decreasing, without loss of generality it holds that $\text{gap}_i^{X,Q} = \min_{j<i} q_i - q_j = q_i - q_{i-1}$ ($q_{i-1} \geq q_j$ for all $j < i$). Therefore, we get MenuGap$(X, Q) = \sum_i q_i - q_{i-1} = q_N - q_0$. Since $q_0 = 0$ and $0 \leq q_N \leq 1$, we get that MenuGap$(X, Q) \leq 1$.

Finally, note that for any X, we can set $q_N = 1$ and $q_i = 0$ for all other i, proving that MenuGap$(X) \geq 1$.

□

B Omitted Proofs

Proof of Lemma 8. We prove that for all X, C, AlignGap$(X, C) \leq $ MenuGap(X), which implies the lemma. For a given X, C, define:

- $\vec{q}_i := c_i \cdot \vec{x}_i$, if $\text{sgap}_i^{X,C} > 0$.

- $\vec{q}_i := \arg \max_{j<i}\{c_j \cdot \vec{x}_j\}$, if $\text{sgap}_i^{X,C} \leq 0$.

Observe first that each $\vec{q}_i \in [0, 1]^k$, as each $c_i \vec{x}_i \in [0, 1]^k$ (this follows because each component of $\vec{x}_i$ is at most $||\vec{x}_i||_\infty$, and each c_i is at most $1/||\vec{x}_i||_\infty$). Next, observe that if $\text{sgap}_i^{X,C} \leq 0$, then $\text{gap}_i^{X,Q} = 0$. This is by definition in bullet two above. Finally, observe that if $\text{sgap}_i^{X,C} > 0$, then $\text{gap}_i^{X,Q} \geq \text{sgap}_i^{X,C}$. This is because the set of $\{\vec{q}_j\}_{j<i}$ is a subset of $\{c_j \vec{x}_j\}_{j<i}$, and because $\vec{q}_i := c_i \cdot \vec{x}_i$ by bullet one. Therefore, $\text{gap}_i^{X,Q} \geq \max\{0, \text{sgap}_i^{X,C}\}$ for all i and the lemma follows. □

Proof of Claim 13. Take M' to be exactly the same as M, except having removed all entries with price $< c$. For every value in the support of $\mathcal{D}$ with $p^M(\vec{v}) \geq c$ in M, we still have $p^{M'}(\vec{v}) \geq c$. This is simply $\vec{v}$'s favorite option in M is still available in M', and all options in M' were also available in M. For any value with $p^M(\vec{v}) < c$, we clearly have $p^{M'}(\vec{v}) \geq 0$. So for all $\vec{v}$, we have $p^{M'}(\vec{v}) \geq p^M(\vec{v}) - c$, and the claim follows by taking an expectation with respect to $\vec{v}$. □

Proof of Claim 15. Simply let M_1 denote the set of menu options from M whose price lies in $[c \cdot 2^i, c \cdot 2^{i+1})$ for an odd integer i, and M_2 denote the remaining menu options (which lie in $[c \cdot 2^i, c \cdot 2^{i+1})$ for an even power of i). It is easy to see that M_1 is oddly-priced and M_2 is evenly-priced. Then for all $\vec{v}$, we must have $p^{M_1}(\vec{v}) + p^{M_2}(\vec{v}) \geq p^M(\vec{v})$. This is because $\vec{v}$'s favorite menu

option from M appears in one of the two menus, and is necessarily $\vec{v}$'s favorite option on that menu (and they pay non-zero from the other menu). Taking an expectation with respect to $\vec{v}$ yields that $\mathrm{Rev}(\mathcal{D}, M_1) + \mathrm{Rev}(\mathcal{D}, M_2) \geq \mathrm{Rev}(\mathcal{D}, M)$, completing the proof. $\qquad\square$

Proof of Claim 18. Recall that $(1 + \varepsilon) \cdot ||\vec{v}||_1 \geq ||\vec{x}_i||_1$ for all $\vec{v} \in B_i$. Therefore, if we set a price of $||\vec{x}_i||_1 / (1 + \varepsilon)$ for the grand bundle, every $\vec{v} \in B_i$ would choose to purchase the grand bundle. This immediately implies the claim, as: $\mathrm{BRev}(\mathcal{D}) \geq \frac{||\vec{x}_i||_1}{1+\varepsilon} \cdot \mathrm{Pr}_{\vec{v} \sim \mathcal{D}} \left[||\vec{v}||_1 \geq \frac{||\vec{x}_i||_1}{1+\varepsilon} \right] \geq \frac{||\vec{x}_i||_1}{1+\varepsilon} \cdot \mathrm{Pr}_{\vec{v} \sim \mathcal{D}}[\vec{v} \in B_i]$. $\qquad\square$

Proof of Claim 19. Recall that $\mathrm{gap}_i^{X,Q} := \min_{j<i}\{\vec{x}_i \cdot (\vec{q}_i - \vec{q}_j)\}$, and that $\vec{q}_i := \vec{q}^M(\vec{x}_i)$. For any fixed $j < i$, recall that because M was a truthful mechanism, we must have:

$$\vec{x}_i \cdot \vec{q}^M(\vec{x}_i) - p^M(\vec{x}_i) \geq \vec{x}_i \cdot \vec{q}^M(\vec{x}_j) - p^M(\vec{x}_j)$$
$$\Rightarrow \vec{x}_i \cdot (\vec{q}_i - \vec{q}_j) \geq p^M(\vec{x}_i) - p^M(\vec{x}_j)$$
$$\Rightarrow \vec{x}_i \cdot (\vec{q}_i - \vec{q}_j) \geq p^M(\vec{x}_i)/2.$$

The first line is simply restating incentive compatibility. The second line is basic algebra, and substituting $\vec{q}_i := \vec{q}^M(\vec{x}_i)$. The third line invokes the fact that $p^M(\vec{x}_i) \geq 2^{2(i-1)+a}$, while $p^M(\vec{x}_j) < 2^{2(j-1)+a+1} \leq 2^{2(i-1)+a-1}$. $\qquad\square$

Proof of Observation 21. This follows immediately from weak Lagrangian duality. For a quick refresher on weak Lagrangian duality, observe that for any feasible solution to the LP defining $\mathrm{AlignGap}'(X)$ we must have $\vec{x}_i \cdot (c_i\vec{x}_i - c_{i-1}\vec{x}_{i-1}) - \mathrm{sgap}_i \geq 0$. Therefore, for any feasible solution to the original LP, that solution is also feasible for $\mathrm{LagRel}_1(X)$, and the objective is only larger. Therefore, the optimal solution to $\mathrm{LagRel}_1(X)$ must be at least as large as $\mathrm{AlignGap}'(X)$. $\qquad\square$

Proof of Observation 22. For all i, $\max\{0, \mathrm{sgap}_i\} - \mathrm{sgap}_i \leq 0$. When $\mathrm{sgap}_i = 0$, the maximum is achieved (and $\mathrm{sgap}_i := 0$ is feasible). Substituting $\max\{0, \mathrm{sgap}_i\} - \mathrm{sgap}_i = 0$ for all i concludes the proof. $\qquad\square$

Proof of Proposition 27. To ease notation throughout the proof, we'll use the notation $\mathrm{gap}_{\ell,j}^{X,Q} := \mathrm{gap}_i^{X,Q}$, where $\vec{x}_i := \vec{x}_{\ell,j}$ ($\vec{x}_i$ is the j^{th} point on layer ℓ). We will also use the notation $(\ell', j') < (\ell, j)$ if $\ell' < \ell$, or $\ell' = \ell$ and $j' < j$ (that is, if the j'^{th} point in the ℓ'^{th} layer comes before the j^{th} point in the ℓ^{th} layer). To understand $\mathrm{gap}_{\ell,j}^{X,Q}$, we need to understand which point "sets the gap" for $\vec{x}_{\ell,j}$, that is, which $(\ell', j') := \arg\min_{(\ell',j')<(\ell,j)}\{(\vec{q}_{\ell,j} - \vec{q}_{\ell',j'}) \cdot \vec{x}_i\}$.

We first analyze which point sets the gap for $\vec{x}_{\ell,j}$ (for even ℓ; for odd ℓ the gap is zero and we don't care which point sets it), and observe that it must either be $\vec{q}_{\ell,j-1}$ or $\vec{q}_{\ell-2,n_{\ell-2}-1}$ (that is, it must be the previous point in the same layer, or the final point in the previous even layer).

Claim 30. *For all j, and all even ℓ, $\mathrm{gap}_{\ell,j}^{X,Q} = \vec{x}_{\ell,j} \cdot \vec{q}_{\ell,j} - \max\{\vec{x}_{\ell,j} \cdot \vec{q}_{\ell-2,n_{\ell-2}-1}, \vec{x}_{\ell,j} \cdot \vec{q}_{\ell,j-1}\}$.*[11]

Proof of Claim 30. First, note that $\mathrm{gap}_{\ell,j}^{X,Q} := \min_{(\ell',j')<(\ell,j)}\{\vec{x}_{\ell,j} \cdot (\vec{q}_{\ell,j} - \vec{q}_{\ell',j'})\} = \vec{x}_{\ell,j} \cdot \vec{q}_{\ell,j} - \max_{(\ell',j')<(\ell,j)}\{\vec{x}_{\ell,j} \cdot \vec{q}_{\ell',j'}\}$. To conclude the proof, simply observe that the first component of $\vec{q}_{\ell',j'}$ is monotone increasing in ℓ' (for fixed j'), and the second component is monotone increasing in j' (for fixed ℓ'). Moreover, the second component of $\vec{q}_{\ell',n_{\ell'}-1}$ is 1, and this is the maximum possible. Also, both components of $\vec{x}_{\ell,j}$ are non-negative, and therefore we conclude that $\vec{x}_{\ell,j} \cdot \vec{q}_{\ell-2,n_{\ell-2}-1} \geq \vec{x}_{\ell,j} \cdot \vec{q}_{\ell',j'}$ whenever $(\ell', j') \leq (\ell-2, n_{\ell-2}-1)$ (in fact, this extends even to $(\ell', j') \leq (\ell-1, n_{\ell-1}-1)$ as no new $\vec{q}$ are introduced in layer $\ell-1$). Also, $\vec{x}_{\ell,j} \cdot \vec{q}_{\ell,j-1} \geq \vec{x}_{\ell,j} \cdot \vec{q}_{\ell,j'}$ whenever $j' \leq j-1$. $\qquad\square$

Now that we know that the gap is set either by the last point in the previous layer, or the previous point in the current layer, we can nail down $\mathrm{gap}_{\ell,j}^{X,Q}$ exactly.

Lemma 31. *For all even $\ell > 2$, and all $j \in [0, n_\ell - 1]$: $\mathrm{gap}_{\ell,j}^{X,Q} \geq \delta_\ell \frac{\sin(\theta_\ell)}{\sin((j+1)\theta_\ell)}$.*

---

[11]For simplicity of notation, define $\vec{q}_{0,j} = \vec{0} = \vec{q}_{\ell,-1}$ for all ℓ, j.

Proof of Lemma 31. To prove the lemma, we simply compute the inner product of $\vec{x}_{\ell,j}$ with the three relevant vectors $\vec{q}_{\ell,j}, \vec{q}_{\ell-2,n_{\ell-2}-1}, \vec{q}_{\ell,j-1}$. To this end, recall that:

$$\vec{q}_{\ell,j} = (z_\ell, 1 - \delta_\ell \cot((j+1)\theta_\ell)),$$
$$\vec{q}_{\ell,j-1} = (z_\ell, 1 - \delta_\ell \cot(j\theta_\ell)),$$
$$\vec{q}_{\ell-2,n_{\ell-2}-1} = (z_{\ell-2}, 1).$$

Therefore, observe that

$$
\begin{aligned}
\vec{x}_{\ell,j} \cdot (\vec{q}_{\ell,j} - \vec{q}_{\ell,j-1}) &= \sin(j\theta_\ell) \cdot \delta_\ell \cdot (\cot(j\theta_\ell) - \cot((j+1)\theta_\ell)) \\
&= \sin(j\theta_\ell) \cdot \delta_\ell \cdot \left(\frac{\cos(j\theta_\ell)}{\sin(j\theta_\ell)} - \frac{\cos((j+1)\theta_\ell)}{\sin((j+1)\theta_\ell)} \right) \\
&= \delta_\ell \cdot \frac{\cos(j\theta_\ell)\sin((j+1)\theta_\ell) - \sin(j\theta_\ell)\cos((j+1)\theta_\ell)}{\sin((j+1)\theta_\ell)} \\
&= \delta_\ell \cdot \frac{\sin(\theta_\ell)}{\sin((j+1)\theta_\ell)}.
\end{aligned}
$$

Similarly,

$$
\begin{aligned}
\vec{x}_{\ell,j} \cdot (\vec{q}_{\ell,j} - \vec{q}_{\ell-2,n_{\ell-2}-1}) &= (\delta_\ell + \delta_{\ell-1}) \cdot \cos(j\theta_\ell) - \delta_\ell \cot((j+1)\theta_\ell) \cdot \sin(j\theta_\ell) \\
&\geq \delta_\ell \cdot \cos(j\theta_\ell) - \delta_\ell \cot((j+1)\theta_\ell) \cdot \sin(j\theta_\ell) \\
&= \frac{\delta_\ell}{\sin((j+1)\theta_\ell)} (\sin((j+1)\theta_\ell)\cos(j\theta_\ell) - \sin(j\theta_\ell)\cos((j+1)\theta_\ell)) \\
&= \delta_\ell \frac{\sin(\theta_\ell)}{\sin((j+1)\theta_\ell)}.
\end{aligned}
$$

This means that no matter which point sets the gap (or if one of the points does not exist), the gap is at least $\delta_\ell \frac{\sin(\theta_\ell)}{\sin((j+1)\theta_\ell)}$. $\qquad\square$

Finally, we need to sum over each even layer.

Corollary 32. *For any even $\ell > 2$, $\sum_{j=0}^{n_\ell-1} \text{gap}_{\ell,j}^{X,Q} \geq \delta_\ell \cdot \ln(n_\ell)/2$.*

Proof of Corollary 32. Consider the following sequence of calculations:

$$
\begin{aligned}
\sum_{j=0}^{n_\ell-1} \text{gap}_{\ell,j}^{X,Q} &\geq \sum_{j=0}^{n_\ell-1} \delta_\ell \frac{\sin(\theta_\ell)}{\sin((j+1)\theta_\ell)} \\
&\geq \delta_\ell \cdot (\theta_\ell - \theta_\ell^3/6) \cdot \sum_{j=0}^{n_\ell-1} \frac{1}{(j+1)\theta_\ell} \\
&\geq \delta_\ell \cdot (1 - \theta_\ell^2/6) \cdot \ln(n_\ell) \\
&\geq \delta_\ell \cdot \ln(n_\ell)/2
\end{aligned}
$$

Above, the first line follows from Lemma 31. The second line uses the fact that $\theta_\ell - \theta_\ell^3/6 \leq \sin(\theta_\ell) \leq \theta_\ell$, because $\theta_\ell \in [0, \pi/2]$. The third line follows as the n^{th} harmonic sum is at least $\ln(n)$. The final line follows as $\theta_\ell^2/6 = \pi^2/(24(n_\ell - 1)^2) \leq 1/2$. $\qquad\square$

And finally, we can wrap up the proof of the proposition. Here, we just need to recall that $\delta_\ell := \frac{1}{\alpha n_\ell} = \frac{1}{\alpha \ell \ln^2(\ell)}$. Therefore, we conclude that:

$$
\sum_{\ell \text{ even}} \sum_{j=0}^{n_\ell-1} \text{gap}_{\ell,j}^{X,Q} \geq \sum_{\ell \text{ even}} \delta_\ell \cdot \ln(n_\ell)/2 = \sum_{\ell \text{ even}} \frac{1}{2\alpha \ell \ln(\ell)} = \infty. \qquad\square
$$

C Proof of Corollary 11

We prove Corollary 11 by making use of Theorem 2 combined with the sequence X from Section 5. The only task is to confirm that $\mathrm{ARev}(D) < \infty$ for the resulting $\mathcal{D}$, which essentially requires that we execute and analyze the construction fully. Let us quickly review the [HN19] construction, given as input a sequence X:

- Let B be a very large constant, to be defined later.
- Let $\vec{v}_i := B^{2^i} \cdot \vec{x}_i / ||\vec{x}_i||_1$ (for all i).
- Let $\mathcal{D}$ sample $\vec{v}_i$ with probability $1/B^{2^i}$ (for all i).
- Let $\mathcal{D}$ sample $\vec{0}$ with probability $1 - \sum_{i \geq 1} 1/B^{2^i}$.

[HN19] establishes that the above construction yields Theorem 2 (for sufficiently large B, as a function of ε). To complete the proof of Corollary 11, we just need to relate $\mathrm{ARev}(D)$ for this construction to $\mathrm{AlignGap}(X)$.

Proposition 33. *The construction above yields a $\mathcal{D}$ satisfying $\mathrm{ARev}(\mathcal{D}) \leq \mathrm{AlignGap}(X) + 1/B$.*

Proof. Consider any mechanism M. We show that $\mathrm{AlignGap}(X) \geq \mathrm{ARev}(\mathcal{D}, M) - 1/B$. To see this, consider the following choice of C:

- If $\vec{v}_i$ is parallel to $\vec{q}^M(\vec{v}_i)$, set $c_i := ||\vec{q}^M(\vec{v}_i)||_2 / ||\vec{x}_i||_2$.
- If $\vec{v}_i$ is not parallel to $\vec{q}^M(\vec{v}_i)$, set $c_i := 0$.

We now need to lower bound $\mathrm{sgap}_i^{X,C}$, when i satisfies the first bullet. Observe that, because M is truthful, we must have, for all $j < i$:

$$\vec{v}_i \cdot \vec{q}^M(\vec{v}_i) - p^M(\vec{v}_i) \geq \vec{v}_i \cdot \vec{q}^M(\vec{v}_j) - p^M(\vec{v}_j)$$
$$\Rightarrow p^M(\vec{v}_i) \leq p^M(\vec{v}_j) + B^{2^i} \vec{x}_i \cdot (c_i \vec{x}_i - c_j \vec{x}_j) / ||\vec{x}||_1$$
$$\Rightarrow p^M(\vec{v}_i) \leq 2B^{2^{i-1}} + B^{2^i} \mathrm{sgap}_i^{X,C} / ||\vec{x}_i||_1$$

Above, the first line follows from incentive compatibility. The second line follows as $\vec{q}^M(\vec{v}_i) = c_i \vec{x}_i$ for all i in the first bullet, and either $\vec{q}^M(\vec{v}_j) = c_j \vec{x}_j$, or $c_j = 0$. The final line follows by taking $j := \arg\min_{j<i}\{\vec{v}_i \cdot (c_i \vec{v}_i - c_j \vec{v}_j)\}$, and by observing that $\vec{v}_j$ cannot possibly pay more than their value for the grand bundle.

We can then conclude that:

$$\mathrm{ARev}(\mathcal{D}, M) \leq \sum_i (2B^{2^{i-1}} + B^{2^i} \mathrm{sgap}_i^{X,C} / ||\vec{x}_i||_1) / B^{2^i}$$
$$\leq \sum_i 2/B^{2^{i-1}} + \mathrm{AlignGap}(X)$$
$$\leq \mathrm{AlignGap}(X) + 1/B.$$

$\square$

Because we can take B as large as we like, we can construct a $\mathcal{D}$ such that $\mathrm{ARev}(\mathcal{D})$ is arbitrarily close to $\mathrm{AlignGap}(X)$, while also maintaining that $\mathrm{Rev}(\mathcal{D})$ is arbitrarily close to $\mathrm{MenuGap}(X)$. Because Theorem 9 provides a construction X such that $\mathrm{MenuGap}(X)/\mathrm{AlignGap}(X) = \infty$, the [HN19] construction, with sufficiently large B, yields a $\mathcal{D}$ with $\mathrm{Rev}(\mathcal{D})/\mathrm{ARev}(\mathcal{D}) = \infty$, completing the proof of Corollary 11.