

Smoothed Analysis of Multi-Item Auctions with Correlated Values

Christos-Alexandros Psomas^{*} Ariel Schwartzman[†] S. Matthew Weinberg[‡]

Abstract

Consider a seller with m heterogeneous items for sale to a single additive buyer whose values for the items are arbitrarily correlated. It was previously shown that, in such settings, distributions exist for which the seller’s optimal revenue is infinite, but the best “simple” mechanism achieves revenue at most one (Briest et al. [5], Hart and Nisan [22]), even when $m = 2$. This result has long served as a cautionary tale discouraging the study of multi-item auctions without some notion of “independent items”.

In this work we initiate a smoothed analysis of such multi-item auction settings. We consider a buyer whose item values are drawn from an arbitrarily correlated multi-dimensional distribution then randomly perturbed with magnitude δ under several natural perturbation models. On one hand, we prove that the [5, 22] construction is surprisingly robust to certain natural perturbations of this form, and the infinite gap remains.

On the other hand, we provide a smoothed model such that the approximation guarantee of simple mechanisms is smoothed-finite. We show that when the perturbation has magnitude δ , pricing only the grand bundle guarantees an $O(1/\delta)$ -approximation to the optimal revenue. That is, no matter the (worst-case) initially correlated distribution, these tiny perturbations suffice to bring the gap down from infinite to finite. We further show that the same guarantees hold when n buyers have values drawn from an arbitrarily correlated mn -dimensional distribution (without any dependence on n).

Taken together, these analyses further pin down key properties of correlated distributions that result in large gaps between simplicity and optimality.

^{*}Computer Science Department, Carnegie Mellon University, cpsomas@cs.cmu.edu.

[†]Department of Computer Science, Princeton University, acohenca@cs.princeton.edu.

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

1 Introduction

How should a revenue-maximizing seller sell heterogeneous goods to interested buyers? This problem has been extensively studied by economists and computer scientists alike, from a variety of perspectives. One major highlight from these works is numerous impossibility results, essentially proving that one cannot hope to find a mechanism that is simple, yet optimal in general settings [5, 22, 23, 16]. A major highlight on the flip side are numerous approximation results, where simple mechanisms are now known to be approximately optimal in quite general (but still structured) settings [10, 24, 11, 21, 2].

Perhaps the clearest example demonstrating the interaction of these two lines of work is the following. Consider a single additive buyer whose values for the two items are drawn jointly from a two-dimensional distribution \mathcal{D} ¹. Then there exist correlated distributions \mathcal{D} such that the revenue-optimal mechanism achieves infinite revenue, yet the revenue of any simple mechanism is at most 1 [5, 22]. Without getting into details of exactly what “simple” means², this impossibility result rules out any hope of simple mechanisms that are even approximately optimal for all two-dimensional \mathcal{D} . Still, a fantastic complementary result shows that if the two item values are drawn independently, selling separately (post a price p_i for each item and let the buyer purchase a single item i for p_i , or both for $p_1 + p_2$) guarantees a 2-approximation [21]. So while the impossibility results for arbitrary \mathcal{D} are quite strong, compelling positive results still exist under natural assumptions.

This avenue turned out to be quite productive: A long line of work recently culminated in a simple and approximately optimal mechanism for any number of buyers with subadditive valuations over any number of independent items [10, 11, 12, 25, 21, 27, 2, 35, 32, 7, 9, 6]. Modeling assumptions aside, this body of works constitutes a major contribution to the theory of optimal auction design.

The impact of these works notwithstanding, one key direction is left largely unaddressed: even as these works generalized in various directions (arbitrary feasibility constraints, combinatorial valuations, etc.), the “independent items” assumption remained. Even the few works that pose models of limited correlation have some underlying notion of independence (e.g. there are “independent features,” and item values depend linearly on features) [12, 3]. While “independent items” is a perfectly natural assumption (and we have greatly deepened our understanding of mechanism design under it), it was never intended to be ubiquitous in all future works. This is especially true due to the nature of the motivating impossibility result: the \mathcal{D} witnessing these impossibilities is so carefully crafted (we overview the construction in Section 3) that it is far removed from any “real-world” motivation. That is, the constructions provided in [5, 22] require carefully building \mathcal{D} by perfectly placing the infinitely many points in its support just so, and even a tiny deviation in the construction would cause the entire argument to collapse. The thoughtful reader may at this point be thinking: if this construction is so fragile that even a tiny deviation breaks it, perhaps a smoothed analysis might prove insightful. Indeed, this is the focus of this paper.

So, what might a “smoothed” distribution look like? Given an arbitrarily correlated distribution \mathcal{D} , its smoothing $\hat{\mathcal{D}}$ first draws a valuation vector \mathbf{v} from \mathcal{D} , and then randomly perturbs \mathbf{v} to $\hat{\mathbf{v}}$ (where the size of the perturbation is parameterized by some $\delta > 0$). This makes sense for the same reason that it makes sense in all other applications of smoothed analysis: these distributions come from somewhere (e.g. past data), inevitably in the presence of tiny noise.

¹That is, $(v_1, v_2) \sim \mathcal{D}$, and the buyer’s value for receiving both items is $v_1 + v_2$ (and receiving just item i is v_i).

²Formally, these results show lower bounds on a measure of simplicity called the menu-size complexity.

1.1 What Makes a Good Smoothed Model?

Once we’ve decided to do smoothed analysis, we also need to pick a model of random perturbation. We’ll be interested in balancing relevance (e.g. how natural a model is) with transparent analysis (e.g. what insights can we derive?). To help illustrate this point, consider the following perturbation proposal: assume that \mathcal{D} is supported on $[0, v_{\max}]^2$,³ and perturb \mathbf{v} by adding a uniformly random vector from $[0, \delta v_{\max}]^2$.

This is indeed a natural starting point. Unfortunately, the model lends itself to a trivial analysis. Specifically, we can first conclude that the optimal revenue for \mathcal{D} (respectively $\hat{\mathcal{D}}$) is at most $2v_{\max}$ (respectively $2v_{\max}(1 + \delta)$). Moreover, for $\hat{\mathcal{D}}$, we can set a price of δv_{\max} on the grand bundle (that is, allow the buyer to pay δv_{\max} to receive all items or 0 to receive nothing), which sells with probability at least $1/2$. Therefore, the revenue achieved by bundling all items together (henceforth, $\text{BREV}(\hat{\mathcal{D}})$) is at least $\delta v_{\max}/2$, immediately guaranteeing an $O(1/\delta)$ -approximation. For the sake of completeness, we include an improved analysis in Appendix E showing that $\text{BREV}(\hat{\mathcal{D}})$ is in fact a $O(\ln(1/\delta))$ -approximation in this model (and that this is nearly tight). So, while technically there’s a “positive result” here, we simply don’t learn much from this exercise. For readers especially interested in the modeling aspect, Appendix F overviews similar issues with alternative models.

From here, we consider two natural modifications. We call the first *Rectangle-Shift*, which essentially replaces the additive shift of the previous model with a multiplicative shift. In this model, values (v, w) are perturbed to $(v + \varepsilon_1 v, w + \varepsilon_2 w)$, where each ε_i is drawn independently from $U[0, \delta]$. Our first main result proves that in fact the infinite gap persists in this model! That is, for all $\delta < 1/2$, there exists a correlated \mathcal{D} such that the corresponding $\hat{\mathcal{D}}$ satisfies $\text{REV}(\hat{\mathcal{D}}) = \infty$ but $\text{BREV}(\hat{\mathcal{D}}) \in O(1)$ (which implies that any “simple” mechanism achieves finite revenue as well).

Theorem 1.1. *For all $\delta < 1/2$, there exists a bivariate distribution \mathcal{D} such that for its corresponding perturbed distribution $\hat{\mathcal{D}}$ in the Rectangle-Shift model, $\text{BREV}(\hat{\mathcal{D}}) \in O(1)$ and $\text{REV}(\hat{\mathcal{D}}) = \infty$.*

We call our second model *Square-Shift*, which essentially replaces the scale v_{\max} in the initial model with $\max\{v, w\}$. In this model, values (v, w) are perturbed to $(v + \varepsilon_1 \max\{v, w\}, w + \varepsilon_2 \max\{v, w\})$, where each ε_i is drawn independently from $U[0, \delta]$. Our second main result shows that $\text{BREV}(\hat{\mathcal{D}}) \in \Omega(\delta \cdot \text{REV}(\hat{\mathcal{D}}))$ (and this is nearly tight, see Appendix C.2).

Theorem 1.2. *In the Square-Shift model, for a single additive buyer and 2 items*

$$\text{REV}(\hat{\mathcal{D}}) \leq \frac{\sqrt{2}\pi(1 + \delta)^3 \log(1 + \delta)}{\delta^2} \text{BREV}(\hat{\mathcal{D}}) \in O\left(\frac{1}{\delta} \cdot \text{BREV}(\hat{\mathcal{D}})\right).$$

Before continuing, we briefly share the distinguishing feature causing Rectangle-Shift to admit an infinite gap, yet Square-Shift to admit finite approximations: the *angle* by which a buyer’s valuation vector may be perturbed. Observe that in the Rectangle-Shift model, valuation vectors extremely close to an axis remain extremely close to that axis after perturbation. This fact turns out to be crucial in enabling a lower bound construction. On the other hand, valuation vectors in the Square-Shift model are likely to have their angle non-trivially perturbed. This property turns out to be crucial in establishing our approximation guarantee. Of course, both results require much more than this simple observation, but this property is the distinguishing factor causing results in the two models to diverge.

We conclude with two additional technical observations. First, note that in both models $\text{BREV}(\mathcal{D}) \approx \text{BREV}(\hat{\mathcal{D}})$ (unlike the initial additive model). This means that the entirety of our

³The constructions of [5, 22] work subject to this, but one has to replace “infinite gap” with “unbounded gap”.

analysis rests on studying $\text{REV}(\hat{\mathcal{D}})$, and determining whether the perturbations guarantee that it’s finite. Second, note that in both models $\hat{\mathcal{D}}$ *stochastically dominates* \mathcal{D} . Therefore, when we prove in Theorem 1.2 that $\text{REV}(\hat{\mathcal{D}})$ is bounded, this is *not* because we perturbed the buyer into valuing the items less, but really because this perturbation negates whatever bizarre properties of the original \mathcal{D} led to $\text{REV}(\mathcal{D}) = \infty$. This serves as another example of revenue non-monotonicity [23].

Brief Discussion of Models. The refuted additive-noise model perhaps seems most natural from a mathematical perspective (indeed, it is the most similar to models used in prior smoothed analysis). From an economic perspective, this is somewhat natural as well (each consumer in the population makes errors on the scale of, e.g., \$10, while no consumer values any item more than \$1 million). The “Rectangle-Shift” model is also uncontroversially natural from both perspectives: a consumer’s value for item j is inaccurately measured proportional to her value for item j . The “Square-Shift” model requires a touch more thought. From a mathematical perspective, it is simply replacing the universal \$1 million scale for all consumers in the population with a scale proportional to that consumer’s value for the items. From an economic perspective, it may initially seem unconvincing that a high value for one item may cause larger error in estimating the value of the other item. However, numerous works in the behavioral economics literature suggest that bidders indeed value items differently in the presence or absence of more/fewer valuable items [1, 30, 26, 28]. For example, in an experiment of [30], the authors observed that a buyer’s willingness to pay for a cheap item (a CD) was (statistically significantly) higher when a second item (a sweater) was being sold at 80\$ than when it was being sold at 10\$.

1.2 Extensions: Many Buyers and Many Items

After resolving the single-buyer, two-item case, the natural question is whether a similar analysis extends to multiple buyers or multiple items. For the Rectangle-Shift model, the impossibility results extend to the multi-item case, so we focus on the Square-Shift model. For n buyers and two items, we consider an arbitrarily correlated $2n$ -dimensional distribution \mathcal{D} (denoting n buyers’ values for two items). Our perturbation then draws ε_{ij} independently and uniformly from $[0, \delta]$ for all buyers i , and $j \in \{1, 2\}$, and maps v_{ij} to $v_{ij} + \varepsilon_{ij} \max_{j' \in \{1, 2\}} v_{ij'}$. In other words, while the buyers’ values are correlated, the perturbations are independent (and only depend on that buyer’s values).

In this model, we’re again able to prove that $\text{RONENREV}(\hat{\mathcal{D}}) \in \Omega(\delta \cdot \text{DSICREV}(\hat{\mathcal{D}}))$ (and this is again nearly tight). Here, $\text{DSICREV}(\hat{\mathcal{D}})$ denotes the revenue of the optimal dominant-strategy truthful mechanism⁴, and $\text{RONENREV}(\hat{\mathcal{D}})$ denotes the revenue achievable by running Ronen’s simple single-item auction [31] for the grand bundle, i.e. treat the grand bundle as a single-item auction and run Ronen’s auction. See Section 2 for the definition of Ronen’s auction — it is a second-price auction with reserve, but the reserve depends on the other buyers’ bids. If \mathcal{D} happens to be independent across buyers, but still correlated across items, we further get $\text{RONENREV}(\hat{\mathcal{D}}) \in \Omega(\delta \cdot \text{BICREV}(\hat{\mathcal{D}}))$, where $\text{BICREV}(\hat{\mathcal{D}})$ is the revenue of the optimal BIC mechanism. Note that the guarantee does not depend on n .

Finally, we consider the single-buyer, multi-item case. The extension of the Square-Shift model is the obvious one: perturb each value independently by $U([0, \delta]) \cdot \max_j v_j$. Here, we show that our techniques extend, but give $\text{BREV}(\hat{\mathcal{D}}) \in \tilde{\Omega}(\delta^m \cdot \text{REV}(\hat{\mathcal{D}}))$. Furthermore, the exponential dependence

⁴That is, it is in each buyer’s interest to report their true value, no matter the behavior of the other buyers. Contrast this with Bayesian truthful, where it is in each buyer’s interest to report their true value as long as the other buyers do so as well. Note that we cannot hope to replace $\text{DSICREV}(\hat{\mathcal{D}})$ with $\text{BICREV}(\hat{\mathcal{D}})$ here due to [18] - see Section 2 for further discussion.

on m is unavoidable: a simple counterexample provides a \mathcal{D} such that for $\delta \in O(1)$, $\text{BREV}(\hat{\mathcal{D}}) < \frac{m}{2^m} \cdot \text{REV}(\hat{\mathcal{D}})$ (Appendix C.3). Our analysis extends to the m -item n -buyer case, again obtaining an approximation guarantee exponential in m .

Theorem 1.3. *In the Square-Shift model, for a single additive buyer and m items*

$$\text{REV}(\hat{\mathcal{D}}) \leq \left(\sqrt{\frac{\pi e (1 + \delta)}{2 \delta}} \right)^m (1 + \delta) \log(1 + \delta) m \sqrt{m} \cdot \text{BREV}(\hat{\mathcal{D}}).$$

Theorem 1.4. *For the n additive buyer, m item, Square-Shift model,*

$$\text{REV}(\hat{\mathcal{D}}) \leq 4 \left(\sqrt{\frac{\pi e (1 + \delta)}{2 \delta}} \right)^m (1 + \delta) \log(1 + \delta) m \sqrt{m} \cdot \text{RONENREV}(\hat{\mathcal{D}}),$$

where $\text{REV}(\hat{\mathcal{D}}) = \text{DSICREV}(\hat{\mathcal{D}})$ for correlated buyers, and $\text{REV}(\hat{\mathcal{D}}) = \text{BICREV}(\hat{\mathcal{D}})$ for independent buyers.

Our analysis further extends to settings beyond those where the buyers’ valuation is additive across items. For example, all of our positive results hold for the general class of valuations which are *additive subject to downwards-closed constraints*, at the cost of an additional factor of m ; we discuss this in detail in Section 5.

1.3 Related Work and Roadmap

The present paper is the first to explore smoothed analysis in auction design. Other hybrid worst-case/average-case guarantees have been studied, for instance, in the digital goods setting (e.g. [20, 19, 13]), where at a high level, auctions compete instance-by-instance against any auction that could potentially be optimal in the average-case. Work of Carroll [8] on robust mechanism design is also thematically related, and argues that simple mechanisms are optimal if the auctioneer wishes to obtain a worst-case guarantee against all value distributions consistent with given marginals (but neither directly nor indirectly addresses the concepts in the constructions of [5, 22]).

There is a growing body of related literature on multi-item auction design, largely proving impossibility results for arbitrarily correlated distributions (e.g. [5, 22]), or approximation results for distributions with “independent items” (e.g. [10, 11, 12, 25, 21, 27, 2, 35, 32, 7, 9, 6]). Limited work exists on models with limited correlation, but such models make use of the same tools, essentially replacing “independent items” with “independent features” and items that are linear combinations of features [12, 3]. The present paper is the first tractable multi-item model of limited correlation that doesn’t rely on these tools.

Smoothed analysis is a popular framework for analyzing algorithms on “real-world worst-case” inputs. Smoothed analysis most commonly refers to smoothed computational complexity (e.g. an algorithm might run in exponential time in the worst case, but in polynomial time if the worst-case inputs are randomly perturbed), and has been an extremely influential paradigm [34]. For instance, the Simplex Method for solving LPs is known to take exponential time in the worst case, but has smoothed-polynomial computational complexity [33]. More similar in spirit to the present paper is prior work that considers the smoothed competitive ratio of online algorithms [4] or smoothed approximation ratio of mechanisms [17]. The motivation for considering smoothed analysis in these works is, of course, similar to ours, but there is no similarity in techniques: the process of proving smoothed guarantees is a domain-specific process.

Roadmap: Section 2 poses our model and some preliminaries. Section 3 presents our lower bound for the Rectangle-Shift model. In Section 4 we present our results in the Square-Shift model,

including its extension to multiple buyers and multiple items. Appendix A contains a detailed summary of our results in the Square-Shift model, as well as comparisons with known results for independent items. Extensions beyond additive buyers are presented in Section 5. Nearly tight bounds for the additive noise model can be found in Appendix E.

2 Preliminaries

Throughout this paper we study auctions with n buyers who are bidding for $m \geq 2$ items. Each buyer i draws its valuation vector \mathbf{v}_i for the items from an m -dimensional distribution \mathcal{D}_i with density f_i , and we refer to \mathcal{D} as the joint nm -dimensional distribution. The buyer has value $v_{i,j}$ for the j -th item, with marginal distribution $\mathcal{D}_{i,j}$. Our results do not assume any kind of independence between buyers or items. That is, $\mathcal{D}_{i,j}$ can be correlated with $\mathcal{D}_{i,j'}$ (same buyer different item), as well as $\mathcal{D}_{i',k}$ (different buyer and possibly different item). For now, all buyers considered are additive, quasi-linear, and risk-neutral (or expected utility maximizers). See Section 5 for extensions to non-additive valuation functions.

Mechanisms and Benchmarks. As with many results in approximate mechanism design we compare the optimal revenue attainable to benchmarks that result from simple mechanisms. For a single buyer, let $\text{REV}(\mathcal{D})$ be the optimal revenue that can be extracted by a truthful mechanism when valuations are sampled according to \mathcal{D} . Let $\text{BREV}(\mathcal{D})$ be the optimal revenue that can be attained by selling the grand bundle as a single item — that is to say, the mechanism posts a single price and allocates all the items if the bidder is willing to pay that price, otherwise it charges nothing and allocates nothing. For multiple buyers, let $\text{BICREV}(\mathcal{D})$ denote the optimal revenue that can be extracted by a BIC mechanism (see Appendix B for a formal definition for correlated buyers), $\text{DSICREV}(\mathcal{D})$ denote the optimal revenue that can be extracted by a DSIC mechanism, and $\text{RONENREV}(\mathcal{D})$ denote the revenue extracted by Ronen’s single-item auction (treating the grand bundle as a single item). Ronen’s auction will always award the item to the highest bidder (or no one), and the price charged is the maximum of the second-highest bid and a per-buyer reserve (that depends on the other buyers’ bids). Ronen shows that for any single-item setting, $2\text{RONENREV}(\mathcal{D}) \geq \text{DSICREV}(\mathcal{D})$. Fu et al. [18] show that $\text{RONENREV}(\mathcal{D})$ does not generally guarantee any finite approximation to $\text{BICREV}(\mathcal{D})$, but that under some assumptions it attains a 5-approximation.⁵

Models. We consider two different smoothing models and refer to the resulting distributions as $\hat{\mathcal{D}}_i$. The magnitude of the perturbation depends on a parameter δ . We write \hat{f}_i for the density of $\hat{\mathcal{D}}_i$. $R(\hat{\mathbf{x}}, \delta) = \{\mathbf{x} | \mathbf{x} \text{ can get mapped to } \hat{\mathbf{x}}\}$ is the set of points that could map to $\hat{\mathbf{x}}$ under a certain model with parameter δ ; the model will always be clear from the context and is omitted from notation. We also drop δ when clear from the context, and write $R^{-1}(\mathbf{x}, \delta)$ for the set of points that \mathbf{x} maps to. The models we consider are (see Figures 1a, 1b in Appendix A for illustrations):

- **Rectangle-Shift:** buyer i ’s value \mathbf{v}_i is replaced by $\hat{\mathbf{v}}_i$ with $\hat{v}_{ij} = v_{ij} + \delta_j v_{ij}$ where each δ_j is sampled independently from $U[0, \delta]$. Intuitively, this spreads the mass $f(\mathbf{v})$ uniformly on the m -parallelepiped with side lengths $\delta \mathbf{v}$ and \mathbf{v} as its smallest vertex.

⁵All of these claims assume that the mechanism is required to be ex-post individually rational, due to seminal work of Cremer and McLean [14]: see Appendix B for further discussion.

- **Square-Shift:** buyer i 's value \mathbf{v}_i is replaced by $\hat{\mathbf{v}}_i$ with $\hat{v}_{ij} = v_{ij} + \delta_j \max_k v_{ik}$ where each δ_j is sampled independently from $U[0, \delta]$. Intuitively, this spreads the mass $f(\mathbf{v})$ uniformly on the m -dimensional cube with side length $\delta \max_k v_{ik}$ and \mathbf{v} as its smallest vertex.

As mentioned earlier, it is known that even for the single buyer, two items case, the gap between $\text{REV}(\mathcal{D})$ and $\text{BREV}(\mathcal{D})$ is unbounded when there is correlation between the items [5, 22]. In this paper we compare $\text{REV}(\hat{\mathcal{D}})$ and $\text{BREV}(\hat{\mathcal{D}})$ as a function of δ .

3 Persistence of Infinite Gap in the Rectangle-Shift Model

In this section we first overview the key ideas from [5, 22], and then present our construction witnessing an infinite gap in the Rectangle-Shift model. Proofs missing from this section can be found in Appendix C.1. Recall this section's main result:

Theorem 1.1. *For all $\delta < 1/2$, there exists a bivariate distribution \mathcal{D} such that for its corresponding perturbed distribution $\hat{\mathcal{D}}$ in the Rectangle-Shift model, $\text{BREV}(\hat{\mathcal{D}}) \in O(1)$ and $\text{REV}(\hat{\mathcal{D}}) = \infty$.*

The key insight behind the [5, 22] construction is the following: assume that you have k points in the positive orthant of the unit hypercube, $\mathbf{x}_1, \dots, \mathbf{x}_k$, such that $\mathbf{x}_i \cdot \mathbf{x}_j < \mathbf{x}_i \cdot \mathbf{x}_i - \varepsilon$ for all i, j . If the valuation vectors in the support of our distribution are exactly $\mathbf{x}_1, \dots, \mathbf{x}_k$, we can design a mechanism such that when the bidder reports a valuation of \mathbf{x}_i , the mechanism offers a randomized allocation (or *lottery*) of \mathbf{x}_i for price ε . When the buyer has valuation vector \mathbf{x}_i , they get utility $\mathbf{x}_i \cdot \mathbf{x}_i - \varepsilon$ for the lottery designed for them, which is greater than the utility of reporting some other valuation \mathbf{x}_j , $\mathbf{x}_i \cdot \mathbf{x}_j$. Therefore, their utility for purchasing the lottery tailored for them exceeds their *value* for any other lottery (and therefore their utility as well). This guarantees that the buyer will always purchase the lottery designed for them. Observe that simply selling the grand bundle at price ε achieves just as much revenue, so this idea is just the first building block.

From here, [5, 22] modify this construction to increase REV while keeping BREV small. The second insight is that if the valuation vectors in the support are actually $2^i \mathbf{x}_i$, then we can offer the lottery \mathbf{x}_i for price $\varepsilon \cdot 2^i$ and the buyer will still always prefer their tailored lottery to any other one, because their utility will be $2^i(\mathbf{x}_i \cdot \mathbf{x}_i - \varepsilon) > 2^i \mathbf{x}_i \cdot \mathbf{x}_j$, which is their value for the lottery \mathbf{x}_j . Therefore, if the buyer has value $2^i \cdot \mathbf{x}_i$ with probability 2^{-i} , REV can be quite large, while BREV will remain small. Of course, there are still some details left to work out, such as exactly how to analyze REV and BREV , and we've also simplified the key ideas at the cost of technical accuracy, but these are the main ideas we'll borrow from previous work: the point is that these constructions are packing points inside the unit hypercube with small pairwise dot products. Our construction is a "perturbation robust" version of that of [22], as we have to show not only that \mathbf{x}_i will prefer the intended option, but that any smoothing of \mathbf{x}_i will prefer it as well.

The intuition of the preceding paragraphs is captured by Lemma 1 below. For a set of points $\mathbf{x}_1, \mathbf{x}_2, \dots \in [0, 1]^2$ let $\text{gap}_\delta^i = \min_{j < i, \hat{\mathbf{x}}_i \in R^{-1}(\mathbf{x}_i, \delta)} \hat{\mathbf{x}}_i \cdot (\mathbf{x}_i - \mathbf{x}_j)$. That is, gap_δ^i roughly captures the minimum angle between any smoothing of \mathbf{x}_i and any \mathbf{x}_j , $j < i$. First, we show that given any such set of points, there exists a distribution \mathcal{D} such that $\text{BREV}(\hat{\mathcal{D}})$ is at most a constant, but there exists an auction that extracts revenue $\sum_{i=1}^{\infty} \text{gap}_\delta^i$. Roughly speaking, \mathcal{D} has a type \mathbf{v}_i for every point \mathbf{x}_i , such that all post-perturbation types $\hat{\mathbf{v}}_i$ prefer randomized allocation \mathbf{x}_i at price $q_i = \text{gap}_\delta^i$ over any other (allocation, price) pair (\mathbf{x}_j, q_j) . The proof of Lemma 1 is similar to Proposition 9 of [22]⁶ after adjusting for our perturbations.

⁶This is clearer when compared to Proposition 5.1 in the arXiv version of [22] uploaded April 22nd, 2013.

Lemma 1. *In the Rectangle-Shift model, for all $\delta < 1/2$, and sequences of points $\mathbf{x}_1, \mathbf{x}_2, \dots \in [0, 1]^2$, there exists a bivariate distribution \mathcal{D} , such that for its corresponding perturbed distribution $\hat{\mathcal{D}}$ it holds that $\text{BREV}(\hat{\mathcal{D}}) \in O(1)$ and $\text{REV}(\hat{\mathcal{D}}) \geq \sum_{i=1}^{\infty} \text{gap}_{\delta}^i$.*

Once we have Lemma 1, the crux of the proof becomes constructing a set of points such that $\sum_{i=1}^{\infty} \text{gap}_{\delta}^i$ diverges. We indeed follow an outline similar to the construction of [22], but the requirement that any *smoothing* of \mathbf{x}_i has sufficiently large angle with all \mathbf{x}_j (versus just \mathbf{x}_i itself) poses additional challenges. We include a fairly detailed overview below, with proofs of the key steps.

Lemma 2. *For all $\delta < 1/2$, there exist a countably infinite set of points $\mathbf{x}_1, \mathbf{x}_2, \dots \in [0, 1]^2$, such that $\sum_i \text{gap}_{\delta}^i$ diverges.*

Proof. The points will lie on a sequence of spherical shells restricted to the non-negative orthant, with increasing radii. Specifically, shell N will contain a consecutive subsequence of points, all with the same magnitude, ℓ_N . We first make sure that the shells are sufficiently far apart from each other (i.e. that $\ell_N - \ell_{N-1}$ is sufficiently large) so that the distance between \mathbf{x}_i on shell N is large when compared to any \mathbf{x}_j on shell $< N$. Specifically, we define $\ell_N = \frac{1}{\alpha} \sum_{k=2}^N c_k$, where $c_k = \frac{1}{k \log^2 k}$, and $\alpha = \sum_{k=2}^{\infty} c_k$. Note crucially that α is finite (by the integral test), and therefore ℓ_N is well-defined and bounded above by 1 for all N .

Now for each shell N , we'll place points with magnitude ℓ_N starting from angle $\pi/2$ and going down to angle 0 such that *within each shell*, the points are sufficiently far apart from each other. Specifically, we'll place points so that the angles form a geometric progression: the k^{th} point on shell N will have angle $\frac{\pi}{2} \cdot (1 - 3\delta)^{k-1}$, and we'll stop placing points on shell N once the angle drops below $\sqrt{c_N}$ (we'll determine later exactly what the maximum k is as a function of N).

From here we have two tasks of analysis. First, we must figure out what is gap_{δ}^i for some point \mathbf{x}_i on shell N . Second, we must figure out how many points we can pack in each shell with the above construction. We begin by analyzing the gap for a point on shell N with the following lemma.

Lemma 3. *Let \mathbf{x}_i be a point on shell N . Then $\text{gap}_{\delta}^i \in \Omega(c_N)$ (recall $c_N = \frac{1}{N \log^2 N}$).*

Proof. In order to evaluate gap_{δ}^i for a point in shell N , we'll compare it to all points \mathbf{x}_j on a shell $\leq N$. Note that technically we don't need to compare to every point on shell N , but the analysis is more straightforward this way. For every point \mathbf{x}_j we also need to find the point $\hat{\mathbf{x}}_i \in R^{-1}(\mathbf{x}_i, \delta)$ such that gap_{δ}^i is minimized.

First, consider \mathbf{x}_i and the set of points Θ_i that have the same angle as \mathbf{x}_i but are on earlier shells. Then, certainly the point \mathbf{x}_j in Θ_i that minimizes $\hat{\mathbf{x}}_i \cdot (\mathbf{x}_i - \mathbf{x}_j)$ is the one on shell $N-1$. Furthermore, the worst possible perturbation for that point is exactly $\hat{\mathbf{x}}_i = \mathbf{x}_i$, since \mathbf{x}_i has both coordinates greater than \mathbf{x}_j . Therefore, the gap between these points is $\mathbf{x}_i \cdot (\mathbf{x}_i - \mathbf{x}_j) = |\mathbf{x}_i|^2 - |\mathbf{x}_i| |\mathbf{x}_j| = \ell_N^2 - \ell_N \ell_{N-1} = \frac{1}{\alpha} \ell_N c_N \in \Omega(c_N)$ (the final relation is because $\ell_N \geq 1/4$ for all N). Therefore, when comparing to any \mathbf{x}_j with the same angle as \mathbf{x}_i , the gap is indeed $\Omega(c_N)$.

Now, consider some \mathbf{x}_j with a different angle. By construction, for any angle θ such that there exists a point on a shell $\leq N$ with angle θ , there also exists a point \mathbf{x}_j on shell N with angle θ . Therefore, such a \mathbf{x}_j inducing the smallest gap is clearly on shell N . Moreover, it will either be the point immediately clockwise of \mathbf{x}_i , or immediately counterclockwise (so as to maximize the dot product and minimize the gap).

Let θ_i be the angle of \mathbf{x}_i on the N -th shell, θ_j be the angle of the \mathbf{x}_j we're comparing to (also on the N -th shell), and $\hat{\theta}_i$ the angle of a point $\hat{\mathbf{x}}_i \in R^{-1}(\mathbf{x}_i, \delta)$. Our plan is to show first that θ_i must be far from θ_j (Claim 1), that $\hat{\theta}_i$ must be close to θ_i (Claim 2), and finally that conditioned on these facts, $\hat{\mathbf{x}}_i \cdot (\mathbf{x}_i - \mathbf{x}_{i-1})$ must be large (Claim 3).

Claim 1. $|\theta_i - \theta_j| \geq 3\delta\theta_i$.

Claim 2. $|\hat{\theta}_i - \theta_i| \leq \delta\theta_i$.

Claim 3. Let \mathbf{x}_i and \mathbf{x}_j both lie on shell N , and $\hat{\mathbf{x}}_i \in R^{-1}(\mathbf{x}_i, \delta)$. Then $\hat{\mathbf{x}}_i \cdot (\mathbf{x}_i - \mathbf{x}_{j-1}) \geq \frac{2}{\pi^2} \ell_N^2 \delta^2 \theta_i^2$.

Given Claim 3, we simply need to observe that most of these terms are in fact independent of N (up to constant factors). Specifically, $\frac{2\delta^2}{\pi^2}$ is obviously independent of N . Moreover, $\ell_N \in (1/4, 1)$ for all N . So Claim 3 indeed concludes that $\hat{\mathbf{x}}_i \cdot (\mathbf{x}_i - \mathbf{x}_j) \in \Omega(\theta_i^2)$.

This completes the proof of Lemma 3: we have shown that the gap to any point on a lower shell is $\Omega(c_N)$, and the gap to any point on the same shell is $\Omega(\theta_i^2) = \Omega(c_N)$. The final equality is simply because our construction stops placing points on shell N once the angle drops below $\sqrt{c_N}$. \square

Now that we've computed gap_δ^i for any point on shell N , we need to compute the number of points on shell N and then we can analyze $\sum_i \text{gap}_\delta^i$.

Lemma 4. The number of points on shell N is $\Theta(\log(1/c_N)) = \Theta(\log(N))$.

Proof. Here, we simply need to figure out the largest value of k such that the angle $\theta_k \geq \sqrt{c_N}$. So we seek the maximum k such that $\pi/2(1-3\delta)^{k-1} \geq \sqrt{c_N}$. Taking logs of both sides and rearranging yields (note that $\log(1-3\delta) < 0$): $k \leq \frac{\log(c_N)/2 - \ln(\pi/2)}{\ln(1-3\delta)} + 1 \in \Theta(\log(1/c_N)) = \Theta(\log(N))$. \square

We can now complete the proof of Lemma 2. We've just shown that our construction has $\Theta(\log N)$ points on shell N (Lemma 4), and each such point has gap at least $\Omega(c_N)$ (Lemma 3). This means that:

$$\sum_i \text{gap}_\delta^i \in \Omega\left(\sum_N c_N \log N\right) = \Omega\left(\sum_N \frac{1}{N \log N}\right).$$

$\sum_N \frac{1}{N \log N}$ diverges by the integral test, meaning that $\sum_i \text{gap}_\delta^i$ also diverges. \square

Theorem 1.1 follows from combining Lemma 1 and Lemma 2. A technical takeaway from this section is that the ability to construct bad examples is heavily tied to the ability to pack points in a sphere ‘‘sufficiently far away’’ from each other. Without any smoothing, it's possible to pack infinitely many such points [22], because the points claim none of the region around them. With rectangle smoothing, the points now claim some of the region around them, and constructions become trickier, but still exist (note, in particular, that because our choice of angles form a geometric progression, our sequence of points is heavily concentrated near the x-axis, where the angle is barely perturbed). On the other hand, in the Square-Shift model, the points now claim a large region around them and it becomes unclear how to pack so many points. Of course, this is just intuition for why these specific constructions no longer work; the following sections show that the Square-Shift model indeed allows for smoothed-finite approximations.

4 Square-Shift Upper Bounds

In this section we present our positive results for the Square-Shift model. We start with a complete proof in a toy ‘‘Angle-Shift’’ model. This will help us highlight some of the key insights without the technical barriers. In this model a buyer's value $(x, y) = (r \sin \theta, r \cos \theta)$ is perturbed to a point $(r \sin(\theta + \varepsilon), r \cos(\theta + \varepsilon))$, where ε is drawn from $U[-\delta, \delta]$. In other words, we output a vector (\hat{x}, \hat{y}) with the same length as (x, y) , i.e. $x^2 + y^2 = \hat{x}^2 + \hat{y}^2$, whose angle is uniformly distributed in the interval $[\max(0, \theta - \delta), \min(\pi/2, \theta + \delta)]$, where θ is the angle of (x, y) (See Figure 1c in Appendix A).

In Section 4.2 we restate all our results for the Square-Shift model. The proof for a single buyer and two items can be found in Section 4.3. We present the proof for multiple buyers and two items in Section 4.4. We include the single buyer, multi-item case in Section 4.5, and the multi-buyer, multi-item case (with arbitrary correlation across everything) in Section 4.6. Throughout this section we assume familiarity with polar coordinates. For a brief review, see Appendix D.

4.1 Angle-Shift Upper Bounds

In this section we describe our approach for the Angle-Shift model. We show that for this perturbation model, bundling recovers a fraction of the optimal revenue for a single buyer interested in two correlated items. This proof highlights some key insights behind Theorem 1.2 without most of the technical barriers.

Theorem 4.1. *For the Two-Dimensional Angle-Shift, $\text{REV}(\hat{\mathcal{D}}) \leq \frac{\pi}{2\delta} \text{BREV}(\hat{\mathcal{D}})$.*

Proof. We break the proof into two big steps. The first step will be almost identical for all our positive results. In the second step we use model-specific facts.

Step one. We begin by writing the revenue of $\hat{\mathcal{D}}$ in polar coordinates. Below, g denotes the density of \mathcal{D} in polar coordinates, and \hat{g} denotes the density of $\hat{\mathcal{D}}$ in polar coordinates. Recall that g is *not* the same as f . We also use the notation $g(\theta)$ to denote $\int_0^\infty g(w, \theta) dw$, and $g(r|\theta)$ to denote $g(r, \theta)/g(\theta)$ (note that $g(r|\theta)$ is the density of a one-dimensional distribution, which is the continuous analog of conditioning on θ). $p^*(r, \theta)$ denotes the payment of a buyer with values $(r \cos(\theta), r \sin(\theta))$ in the optimal mechanism.

$$\text{REV}(\hat{\mathcal{D}}) = \int_{\theta=0}^{\frac{\pi}{2}} \int_{r=0}^{\infty} p^*(r, \theta) \hat{g}(r, \theta) dr d\theta = \int_{\theta=0}^{\frac{\pi}{2}} \hat{g}(\theta) \int_{r=0}^{\infty} p^*(r, \theta) \hat{g}(r|\theta) dr d\theta.$$

Let $\hat{\mathcal{D}}_\theta$ be the single parameter distribution from which we first sample r according to the density function $\hat{g}(r|\theta)$, and then output the vector $(r \cos \theta, r \sin \theta)$. Then, notice that the inner most integral is the expected payments according to p^* for $\hat{\mathcal{D}}_\theta$. Observe also that if our optimal mechanism is truthful on the entire domain \mathbb{R}^2 , it is certainly also truthful on the domain restricted to θ . Therefore, this expected payment is upper bounded by $\text{REV}(\hat{\mathcal{D}}_\theta)$, the revenue of the optimal truthful mechanism for the same distribution. Furthermore, since this is a single parameter distribution, its revenue is equal to the revenue of a posted price auction (with some reserve r_θ) [29].

$$\begin{aligned} \text{REV}(\hat{\mathcal{D}}) &= \int_{\theta=0}^{\frac{\pi}{2}} \hat{g}(\theta) \int_{r=0}^{\infty} p^*(r, \theta) \hat{g}(r|\theta) dr d\theta \leq \int_{\theta=0}^{\frac{\pi}{2}} \hat{g}(\theta) \text{REV}(\hat{\mathcal{D}}_\theta) d\theta \\ &= \int_{\theta=0}^{\frac{\pi}{2}} \hat{g}(\theta) \left(r_\theta \int_{r=\frac{r_\theta}{\sin \theta + \cos \theta}}^{\infty} \hat{g}(r|\theta) dr \right) d\theta = \int_{\theta=0}^{\frac{\pi}{2}} r_\theta \int_{r=\frac{r_\theta}{\sin \theta + \cos \theta}}^{\infty} \hat{g}(r, \theta) dr d\theta. \end{aligned} \quad (1)$$

In going from line one to two above, we are observing that $\text{REV}(\hat{\mathcal{D}}_\theta)$ integrates the density of $\hat{\mathcal{D}}_\theta$ over all r with $r(\sin \theta + \cos \theta) \geq r_\theta$. The final step is simply rearranging for convenience later. This concludes step one, which is simply to upper bound the optimal revenue by an integral over one-dimensional revenues. Step Two follows, which appeals to the model at hand to unpack $\hat{g}(r, \theta)$.

Step two: model specific analysis. For the Angle-Shift model we have that

$$\hat{g}(r, \theta) = \int_{\phi=\max(0, \theta-\delta)}^{\min(\pi/2, \theta+\delta)} \frac{g(r, \phi)}{\min(\pi/2, \theta+\delta) - \max(0, \theta-\delta)} d\phi \leq \int_{\phi=\max(0, \theta-\delta)}^{\min(\pi/2, \theta+\delta)} \frac{g(r, \phi)}{\delta} d\phi \leq \frac{g(r)}{\delta}.$$

Above, the equality is simply because each (r, ϕ) when sampled from g sends a $\Theta(1/\delta)$ fraction of its density to (r, θ) . The first inequality is just observing that the fraction of (r, ϕ) 's density that goes to (r, θ) is at most $1/\delta$ (it is exactly $1/\delta$ unless ϕ is near an axis). The final inequality just observes that the third term is integrating over a proper subset of possible ϕ , so the integral is certainly upper bounded by integrating the entire region from 0 to $\pi/2$ (similarly to above, we let $g(r) = \int_0^{\pi/2} g(r, \theta) d\theta$). This concludes step two, which is simply to use the specifics of the model to upper bound the density of the smoothed distribution by that of the original distribution.

We combine steps one and two directly to continue our derivation:

$$\begin{aligned}
\text{REV}(\hat{\mathcal{D}}) &\leq \frac{1}{\delta} \int_{\theta=0}^{\frac{\pi}{2}} r_{\theta} \int_{r=\frac{r_{\theta}}{\sin \theta + \cos \theta}}^{\infty} g(r) dr d\theta \\
&\leq \frac{1}{\delta} \int_{\theta=0}^{\frac{\pi}{2}} r_{\theta} Pr \left[\text{a sample } \mathbf{v} \sim \mathcal{D} \text{ has } \|\mathbf{v}\|_2 \geq \frac{r_{\theta}}{\sin \theta + \cos \theta} \right] d\theta \\
&= \frac{1}{\delta} \int_{\theta=0}^{\frac{\pi}{2}} r_{\theta} Pr \left[\text{a sample } \mathbf{v} \sim \hat{\mathcal{D}} \text{ has } \|\mathbf{v}\|_2 \geq \frac{r_{\theta}}{\sin \theta + \cos \theta} \right] d\theta \\
&= \frac{1}{\delta} \int_{\theta=0}^{\frac{\pi}{2}} r_{\theta} Pr \left[\text{a sample } \mathbf{v} \sim \hat{\mathcal{D}} \text{ has } \|\mathbf{v}\|_1 \geq r_{\theta} \right] d\theta \\
&\leq \frac{1}{\delta} \int_{\theta=0}^{\frac{\pi}{2}} \text{BREV}(\hat{\mathcal{D}}) d\theta = \frac{\pi}{2\delta} \text{BREV}(\hat{\mathcal{D}}). \quad \square
\end{aligned}$$

This completes the proof for the Angle-Shift model. Note that above the first inequality simply combines steps one and two. The second observes that integrating the density of a distribution yields (one minus) its CDF. The next line observes that the described event occurs with equal probability for samples from \mathcal{D} and $\hat{\mathcal{D}}$ (this is the step where angle-shifting saves some technical work over square-shifting). The fourth is basic geometry. The final inequality notes that the expression we are integrating over is the revenue of a posted-price mechanism for the bundle (specifically, with reserve r_{θ}) and in particular can be no better than $\text{BREV}(\hat{\mathcal{D}})$ itself.

Let us now highlight the key step of the proof where we make use of the smoothed model. Step one applies for any distribution (even unsmoothed), and is in general *very* loose. In particular, it could be as high as the full welfare for some distributions. Smoothing gets us mileage in the first half of step two where we transition from terms that depend on \hat{g} to terms that depend on g . Mathematically, we capitalize on the following: if there were no angle-shifting, then we would simply have had $\hat{g}(r, \theta) = g(r, \theta)$ instead of the integral with respect to ϕ . This is important because integrals over density functions are probabilities, but density functions themselves are not probabilities! That is, if we concentrate on the first part of step two, this is claiming that an integral of a non-negative function over a smaller range is upper bounded by the integral of the same function over a larger range. Without smoothing, we would instead just have a single point of that function, which *can't* generically be upper bounded by an integral over any region.

4.2 Square-Shift: Our Results

We remind the reader of our results for the Square-Shift model below. Section 4.1 captures the intuition for one core step which appears in each of the proofs, but in order to prove our main positive results we still need to overcome a number of technical obstacles.

Theorem 1.2. *In the Square-Shift model, for a single additive buyer and 2 items*

$$\text{REV}(\hat{\mathcal{D}}) \leq \frac{\sqrt{2}\pi(1+\delta)^3 \log(1+\delta)}{\delta^2} \text{BREV}(\hat{\mathcal{D}}) \in O\left(\frac{1}{\delta} \cdot \text{BREV}(\hat{\mathcal{D}})\right).$$

Theorem 4.2. *In the Square-Shift model, for n additive buyers and 2 items:*

$$\begin{aligned} \text{DSICREV}(\hat{\mathcal{D}}) &\leq \frac{9\sqrt{2}(1+\delta)^3 \log(1+\delta)}{\delta^2} \text{RONENREV}(\hat{\mathcal{D}}) \in O(\text{RONENREV}(\hat{\mathcal{D}})/\delta) \text{ (if buyers are correlated),} \\ \text{BICREV}(\hat{\mathcal{D}}) &\leq \frac{9\sqrt{2}(1+\delta)^3 \log(1+\delta)}{\delta^2} \text{RONENREV}(\hat{\mathcal{D}}) \in O(\text{RONENREV}(\hat{\mathcal{D}})/\delta) \text{ (if buyers are independent).} \end{aligned}$$

Theorem 1.3. *In the Square-Shift model, for a single additive buyer and m items*

$$\text{REV}(\hat{\mathcal{D}}) \leq \left(\sqrt{\frac{\pi e (1+\delta)}{2}} \frac{1+\delta}{\delta} \right)^m (1+\delta) \log(1+\delta) m \sqrt{m} \cdot \text{BREV}(\hat{\mathcal{D}}).$$

Theorem 1.4. *For the n additive buyer, m item, Square-Shift model,*

$$\text{REV}(\hat{\mathcal{D}}) \leq 4 \left(\sqrt{\frac{\pi e (1+\delta)}{2}} \frac{1+\delta}{\delta} \right)^m (1+\delta) \log(1+\delta) m \sqrt{m} \cdot \text{RONENREV}(\hat{\mathcal{D}}),$$

where $\text{REV}(\hat{\mathcal{D}}) = \text{DSICREV}(\hat{\mathcal{D}})$ for correlated buyers, and $\text{REV}(\hat{\mathcal{D}}) = \text{BICREV}(\hat{\mathcal{D}})$ for independent buyers.

4.3 Square-Shift: One buyer, Two Items.

Proof of Theorem 1.2. Our first few steps are identical to the Angle-Shift model. Recall that the first step on the Angle-Shift model was simply to upper bound the optimal multi-dimensional revenue by that of multiple single-dimensional distributions. We then write the revenue of a single-dimensional distribution as the revenue of a posted-price.

$$\begin{aligned} \text{REV}(\hat{\mathcal{D}}) &= \int_{\theta=0}^{\pi/2} \int_{r=0}^{\infty} p^*(r, \theta) \hat{g}(r, \theta) dr d\theta = \int_{\theta=0}^{\pi/2} \int_{r=0}^{\infty} p^*(r, \theta) \hat{g}(\theta) \hat{g}(r|\theta) dr d\theta \\ &\stackrel{\text{Lemma 6}}{\leq} \int_{\theta=0}^{\pi/2} \hat{g}(\theta) \text{REV}(\hat{\mathcal{D}}_{\theta}) d\theta = \int_{\theta=0}^{\pi/2} \hat{g}(\theta) \left(r_{\theta} \int_{r=\frac{r_{\theta}}{\cos \theta + \sin \theta}}^{\infty} \hat{g}(r|\theta) dr \right) d\theta \\ &= \int_{\theta=0}^{\pi/2} r_{\theta} \int_{r=\frac{r_{\theta}}{\cos \theta + \sin \theta}}^{\infty} \hat{g}(r, \theta) dr d\theta. \end{aligned} \tag{2}$$

We now want to proceed with the analog of the second step: upper bound $\int_{r=\frac{r_{\theta}}{\cos \theta + \sin \theta}}^{\infty} \hat{g}(r, \theta) dr$ in terms of the original distribution. We start by moving from polar coordinates to cartesian coordinates, using the transformation $\hat{g}(r, \theta) = r \hat{f}(r \cos \theta, r \sin \theta)$. Then, we can connect \hat{f} , the density of $\hat{\mathcal{D}}$, to f , the density of \mathcal{D} , using the following:

$$\begin{aligned} \hat{f}(x, y) &= \int_0^{\infty} \int_0^{\infty} \frac{f(a, b) I\{(a, b) \in R(x, y)\} da db}{\delta^2 \max\{a, b\}^2} \\ &\leq \frac{(1+\delta)^2}{\delta^2 \max\{x, y\}^2} \int_{a=0}^{\infty} \int_{b=0}^{\infty} f(a, b) I\{(a, b) \in R(x, y)\} db da, \end{aligned} \tag{3}$$

where $I\{(a, b) \in R(x, y)\}$ is an indicator for the event that (a, b) belongs in $R(x, y)$. The inequality follows from the fact that for $I\{(a, b) \in R(x, y)\}$ to be 1 we must have that $\max\{a, b\} \geq \frac{\max\{x, y\}}{1+\delta}$.

$$\int_{r=\frac{r_{\theta}}{\cos \theta + \sin \theta}}^{\infty} \hat{g}(r, \theta) dr = \int_{r=\frac{r_{\theta}}{\cos \theta + \sin \theta}}^{\infty} r \hat{f}(r \cos \theta, r \sin \theta) dr$$

$$\begin{aligned}
&\leq \frac{(1+\delta)^2}{\delta^2} \int_{r=\frac{r_\theta}{\cos\theta+\sin\theta}}^{\infty} \frac{r}{r^2 \max\{\cos\theta, \sin\theta\}^2} \int_{a=0}^{\infty} \int_{b=0}^{\infty} f(a,b) I\{(a,b) \in R(r,\theta)\} db da dr \\
&\leq \frac{2(1+\delta)^2}{\delta^2} \int_{r=\frac{r_\theta}{\cos\theta+\sin\theta}}^{\infty} \frac{1}{r} \int_{a=0}^{\infty} \int_{b=0}^{\infty} f(a,b) I\{(a,b) \in R(r,\theta)\} db da dr. \quad \left[\max\{\cos\theta, \sin\theta\}^2 \geq \frac{1}{2} \right]
\end{aligned}$$

This is starting to look familiar to the expression we had at the end of step two for Angle-Shift. The goal of the next few steps is to replace the innermost integral by the integral of the original density on some controlled region and get rid of the indicator function. At this end of this process, we hope to get a double integral integral (with respect to a and b) that we can interpret as a probability. Our first step is to apply Fubini's Theorem to swap the order of integration:

$$\int_{r=\frac{r_\theta}{\cos\theta+\sin\theta}}^{\infty} \hat{g}(r,\theta) dr \leq \frac{2(1+\delta)^2}{\delta^2} \int_{a=0}^{\infty} \int_{b=0}^{\infty} f(a,b) \int_{r=\frac{r_\theta}{\cos\theta+\sin\theta}}^{\infty} \frac{I\{(a,b) \in R(r,\theta)\}}{r} dr db da.$$

Now we argue about values of a and b where the indicator $I\{(a,b) \in R(r,\theta)\}$ is non-zero.

Claim 4. $\forall (a,b) \in R(r,\theta)$, if $r \geq \frac{r_\theta}{\cos\theta+\sin\theta}$ then $a+b \geq \frac{r_\theta}{(1+\delta)(\cos\theta+\sin\theta)}$.

Proof. Recall that $(a,b) \in R(r,\theta)$ means that (a,b) that could map to a point with length r and angle θ . The length $\ell_{(a,b)}$ of a point (a,b) that could map to (r,θ) via our perturbing process must be at least $\frac{r}{1+\delta}$. Furthermore, this length $\ell_{(a,b)}$ is at most $a+b$. Combining we get $a+b \geq \ell_{(a,b)} \geq \frac{r}{1+\delta} \geq \frac{r_\theta}{(1+\delta)(\cos\theta+\sin\theta)}$. \square

We apply Claim 4 to change the limit of integration:

$$\begin{aligned}
&\int_{r=\frac{r_\theta}{\cos\theta+\sin\theta}}^{\infty} \hat{g}(r,\theta) dr \leq \\
&\frac{2(1+\delta)^2}{\delta^2} \int_{a=0}^{\infty} \int_{b=\frac{r_\theta}{(1+\delta)(\cos\theta+\sin\theta)}-a}^{\infty} f(a,b) \int_{r=\frac{r_\theta}{\cos\theta+\sin\theta}}^{\infty} \frac{I\{(a,b) \in R(r,\theta)\}}{r} dr db da.
\end{aligned}$$

We can now upper bound the value of the inner most integral.

Claim 5. For the current Square-Shift model, for any $(a,b) \in \mathbb{R}^2$,

$$\int_{r=\frac{r_\theta}{\cos\theta+\sin\theta}}^{\infty} \frac{1}{r} I\{(a,b) \in R(r,\theta)\} dr \leq \log(1+\delta).$$

Proof. For a fixed angle θ and point (a,b) , in order to find out where the indicator is non-zero we need to find the set of points on the line with slope θ that intersect the set of points to which (a,b) can map to. The line will first intersect the square defined by (a,b) at some point (x_1, y_1) with length ℓ_1 (if it ever intersects), and leave the square at some point (x_2, y_2) with length ℓ_2 . The ratio between ℓ_1 and ℓ_2 is at most $(1+\delta)$:

$$\int_{r=\frac{r_\theta}{\cos\theta+\sin\theta}}^{\infty} \frac{1}{r} I\{(a,b) \in R(r,\theta)\} dr \leq \int_{r=\ell_1}^{(1+\delta)\ell_1} \frac{1}{r} dr.$$

The claim follows from integrating the simplified upper bound. \square

Once the indicator has been dealt with, the remaining upper bound looks similar to that at the end of step two in the Angle-Shift model. The next couple of lines interpret the double integral as the probability that the buyer drew a sample whose bundle value was above the reserve price of that distribution. This is the moral analog of final step in the Angle-Shift model. Applying Claim 5 we have

$$\begin{aligned} \int_{r=\frac{r_\theta}{\cos\theta+\sin\theta}}^{\infty} \hat{g}(r, \theta) dr &\leq \frac{2(1+\delta)^2 \log(1+\delta)}{\delta^2} \int_{a=0}^{\infty} \int_{b=\frac{r_\theta}{(1+\delta)(\cos\theta+\sin\theta)}-a}^{\infty} f(a, b) db da \\ &= \frac{2(1+\delta)^2 \log(1+\delta)}{\delta^2} Pr \left[\text{sample from } \mathcal{D} \text{ has sum of values at least } \frac{r_\theta}{(1+\delta)(\cos\theta+\sin\theta)} \right] \\ &\leq \frac{2(1+\delta)^2 \log(1+\delta)}{\delta^2} Pr \left[\text{sample from } \hat{\mathcal{D}} \text{ has sum of values at least } \frac{r_\theta}{(1+\delta)(\cos\theta+\sin\theta)} \right], \end{aligned}$$

where we used the fact that perturbing only increases the value for an item.

The remaining step is to combine steps one and two and interpret the right hand side as a possible bundling mechanism. This can be further upper bounded by the revenue of the optimal bundling scheme. So plugging back in 2:

$$\begin{aligned} \text{REV}(\hat{\mathcal{D}}) &\leq \frac{2(1+\delta)^2 \log(1+\delta)}{\delta^2} \int_{\theta=0}^{\pi/2} r_\theta Pr \left[\hat{v}_1 + \hat{v}_2 \geq \frac{r_\theta}{(1+\delta)(\cos\theta+\sin\theta)} \right] d\theta \\ &\leq \frac{2(1+\delta)^2 \log(1+\delta)}{\delta^2} \int_{\theta=0}^{\pi/2} r_\theta Pr \left[\hat{v}_1 + \hat{v}_2 \geq \frac{r_\theta}{(1+\delta)\sqrt{2}} \right] d\theta \\ &= \frac{2\sqrt{2}(1+\delta)^3 \log(1+\delta)}{\delta^2} \int_{\theta=0}^{\pi/2} \frac{r_\theta}{(1+\delta)\sqrt{2}} Pr \left[\hat{v}_1 + \hat{v}_2 \geq \frac{r_\theta}{(1+\delta)\sqrt{2}} \right] d\theta. \end{aligned}$$

Claim 6. For any price r , $\text{BREV}(\hat{\mathcal{D}}) \geq r Pr[\hat{v}_1 + \hat{v}_2 \geq r]$.

Proof. $\text{BREV}(\hat{\mathcal{D}})$ is the revenue maximizing auction among those that sell both items (with allocation probability 1) whenever the allocation probability is non-zero. $r Pr[\hat{v}_1 + \hat{v}_2 \geq r]$ is the revenue of setting a price r for the grand bundle. \square

Applying Claim 6 completes the proof of Theorem 1.2:

$$\text{REV}(\hat{\mathcal{D}}) \leq \frac{2\sqrt{2}(1+\delta)^3 \log(1+\delta)}{\delta^2} \int_{\theta=0}^{\pi/2} \text{BREV}(\hat{\mathcal{D}}) d\theta = \frac{\sqrt{2}\pi(1+\delta)^3 \log(1+\delta)}{\delta^2} \text{BREV}(\hat{\mathcal{D}}). \quad \square$$

4.4 Square-Shift: Multiple Buyers, Two Items.

We study the multi-buyer, two-item scenario. The main difference to the single-buyer case is that interaction between different buyers in optimal auctions (even optimal single-item auctions) can be very complex. Similar to prior work of [35, 7], the main insight (beyond the single-buyer setting) here is to make use of Ronen's auction instead of the optimal auction [31].

Theorem 4.2. In the Square-Shift model, for n additive buyers and 2 items:

$$\begin{aligned} \text{DSICREV}(\hat{\mathcal{D}}) &\leq \frac{9\sqrt{2}(1+\delta)^3 \log(1+\delta)}{\delta^2} \text{RONENREV}(\hat{\mathcal{D}}) \in O(\text{RONENREV}(\hat{\mathcal{D}})/\delta) \text{ (if buyers are correlated),} \\ \text{BICREV}(\hat{\mathcal{D}}) &\leq \frac{9\sqrt{2}(1+\delta)^3 \log(1+\delta)}{\delta^2} \text{RONENREV}(\hat{\mathcal{D}}) \in O(\text{RONENREV}(\hat{\mathcal{D}})/\delta) \text{ (if buyers are independent).} \end{aligned}$$

Proof. The proof is broken into two steps, similarly to the proofs of Theorem 4.1 and Theorem 1.2. In the first step, we upper bound $\text{REV}(\hat{\mathcal{D}})$ by an expression that involves $\text{REV}(\hat{\mathcal{D}}_\theta)$, where $\hat{\mathcal{D}}_\theta$ denotes the n -dimensional distribution \mathbf{r} conditioned on angles θ (so below, \mathbf{r} denotes the vector of magnitudes for each buyer, and θ their angles).

$$\text{DSICREV}(\hat{\mathcal{D}}) = \int_{\theta} \hat{g}(\theta) \int_{\mathbf{r}} \hat{g}(\mathbf{r}|\theta) \sum_{i=1}^n p_i^*(\mathbf{r}, \theta) d\mathbf{r} d\theta \leq \int_{\theta} \hat{g}(\theta) \text{DSICREV}(\hat{\mathcal{D}}_\theta) d\theta,$$

where in the last inequality we used a generalization of Equation 1. See Lemma 6 in Appendix D for a full re-derivation. Again observe that $\hat{\mathcal{D}}_\theta$ is now a single-parameter instance.

Main intuition. From here, we wish to mimic the analysis of the previous section. Unfortunately, optimal multi-buyer mechanisms (even Myerson's [29]) don't take a posted-price format, so we use Ronen's [31] instead. Fixing angles θ for all buyers, and lengths r_{-i} for all buyers but i , Ronen's mechanism (for the possibly correlated, single parameter distribution $\hat{\mathcal{D}}_\theta$) makes a take it or leave it offer $r_i^*(\theta, \mathbf{r}_{-i})$ to buyer i , which is larger than the value of any other buyer, i.e. at least $\max_{j \neq i} (r_j(\cos \theta_j + \sin \theta_j))$. Ronen proved that $2\text{RONENREV}(\hat{\mathcal{D}}_\theta) \geq \text{DSICREV}(\hat{\mathcal{D}}_\theta)$; for independent buyers, $2\text{RONENREV}(\hat{\mathcal{D}}_\theta) \geq \text{BICREV}(\hat{\mathcal{D}}_\theta)$ (and this is the only difference between the correlated and independent buyer proofs).

Let $w_{-i}^*(\theta_{-i}, \mathbf{r}_{-i}) = \max_{j \neq i} (r_j(\cos \theta_j + \sin \theta_j))$ be the highest valuation among all buyers excluding i . Then Ronen's mechanism gives buyer i the option to purchase the item at price $r_i^*(\theta, \mathbf{r}_{-i}) = \arg \max_{w \geq w_{-i}^*(\theta_{-i}, \mathbf{r}_{-i})} (w \cdot \Pr[r_i(\cos \theta_i + \sin \theta_i) \geq w])$. Using this, we get

$$\begin{aligned} \text{DSICREV}(\hat{\mathcal{D}}) &\leq 2 \int_{\theta} \hat{g}(\theta) \left(\int_{\mathbf{r}} \hat{g}(\mathbf{r}|\theta) \sum_{i=1}^n r_i^*(\theta, \mathbf{r}_{-i}) I\{r_i(\cos \theta_i + \sin \theta_i) \geq r_i^*(\theta, \mathbf{r}_{-i})\} d\mathbf{r} \right) d\theta \\ &= 2 \sum_{i=1}^n \int_{\theta_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \theta_{-i}) \int_{\theta_i} r_i^*(\theta, \mathbf{r}_{-i}) \int_{r_i = \frac{r_i^*(\theta, \mathbf{r}_{-i})}{\cos \theta_i + \sin \theta_i}}^{\infty} \hat{g}(r_i, \theta_i | \mathbf{r}_{-i}, \theta_{-i}) dr_i d\theta_i d\mathbf{r}_{-i} d\theta_{-i}. \end{aligned} \quad (4)$$

Above, the first line simply replaces $\text{DSICREV}(\hat{\mathcal{D}}_\theta)$ with a formula for the expected revenue of Ronen's mechanism. The second line performs a few operations: the inner sum is pulled outside the integral, some \hat{g} terms are lumped together for cleanliness, and the integral over an indicator is replaced by an integral over the region where the indicator is one (and the indicator is dropped).

The brief derivation above captures the main intuition: it's crucial to upper bound the optimal revenue as an integral over posted-price revenues. Of course, as we see next, wrapping up the multi-buyer case has unique subtleties.

The formal argument For simplicity, we write $\text{REV}(\hat{\mathcal{D}})$, and don't separate between correlated and independent buyers. Recall though, that for independent buyers $\text{REV}(\hat{\mathcal{D}}) = \text{BICREV}(\hat{\mathcal{D}})$ and for correlated buyers, $\text{REV}(\hat{\mathcal{D}}) = \text{DSICREV}(\hat{\mathcal{D}})$.

$$\begin{aligned} \text{REV}(\hat{\mathcal{D}}) &\leq 2 \int_{\theta} \hat{g}(\theta) \left(\int_{\mathbf{r}} \hat{g}(\mathbf{r}|\theta) \sum_{i=1}^n r_i^*(\theta, \mathbf{r}_{-i}) I\{r_i(\cos \theta_i + \sin \theta_i) \geq r_i^*(\theta, \mathbf{r}_{-i})\} d\mathbf{r} \right) d\theta \\ &= 2 \sum_{i=1}^n \int_{\theta_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \theta_{-i}) \int_{\theta_i} r_i^*(\theta, \mathbf{r}_{-i}) \int_{r_i = \frac{r_i^*(\theta, \mathbf{r}_{-i})}{\cos \theta_i + \sin \theta_i}}^{\infty} \hat{g}(r_i, \theta_i | \mathbf{r}_{-i}, \theta_{-i}) dr_i d\theta_i d\mathbf{r}_{-i} d\theta_{-i} \end{aligned}$$

$$\begin{aligned}
&= 2 \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{\theta_i} I\{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) \leq \sqrt{2}(1+\delta)w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})\} r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) \\
&\quad \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\cos \theta_i + \sin \theta_i}}^{\infty} \hat{g}(r_i, \theta_i | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) dr_i d\theta_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i}
\end{aligned} \tag{5}$$

$$\begin{aligned}
&+ 2 \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{\theta_i} I\{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) > \sqrt{2}(1+\delta)w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})\} r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) \\
&\quad \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\cos \theta_i + \sin \theta_i}}^{\infty} \hat{g}(r_i, \theta_i | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) dr_i d\theta_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i}.
\end{aligned} \tag{6}$$

Recall that $w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}) = \max_{j \neq i} (r_j(\cos \theta_j + \sin \theta_j))$ (the highest valuation among all buyers excluding i) and $r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) = \arg \max_{w \geq w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})} (w \cdot \Pr[r_i(\cos \theta_i + \sin \theta_i) \geq w])$.

Bounding 5. We bound 5 by the revenue of a second price auction for the grand bundle. Notice that replacing $r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})$ with $w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})$ makes the probability of sale increase, and the price decreases by a factor of at most $\frac{1}{\sqrt{2}(1+\delta)}$:

$$\begin{aligned}
5 &\leq 2 \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{\theta_i} (1+\delta)w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}) \int_{r_i = \frac{w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})}{\cos \theta_i + \sin \theta_i}}^{\infty} \hat{g}(r_i, \theta_i | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) dr_i d\theta_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i} \\
&= 2\sqrt{2}(1+\delta) \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}) \int_{\theta_i} \int_{r_i = \frac{w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})}{\cos \theta_i + \sin \theta_i}}^{\infty} \hat{g}(r_i, \theta_i | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) dr_i d\theta_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i} \\
&= 2\sqrt{2}(1+\delta) \text{SECONDPRICE}(\hat{\mathcal{D}}) \leq 2\sqrt{2}(1+\delta) \text{RONENREV}(\hat{\mathcal{D}}),
\end{aligned}$$

where $\text{SECONDPRICE}(\hat{\mathcal{D}})$ is the optimal revenue that can be attained from using a second price auction for the grand bundle, i.e. the auction where each buyer bids her value for the grand bundle, the buyer with the highest bid wins, and the payment is the second highest bid.

Bounding 6. Our approach for bounding 6 will be similar to the single buyer case. Specifically, we first upper bound the inner most integral, for fixed $\boldsymbol{\theta}$ and \mathbf{r}_{-i} , by switching to cartesian coordinates, applying the formula to go from \hat{f} to f for our perturbed model, changing order of integration using Fubini's theorem, and then arguing about the size of the area that gets mapped to the point $(r_i \cos \theta_i, r_i \sin \theta_i)$:

$$\begin{aligned}
&\int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\cos \theta_i + \sin \theta_i}}^{\infty} \hat{g}(r_i, \theta_i | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) dr_i = \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\cos \theta_i + \sin \theta_i}}^{\infty} r_i \hat{f}(r_i \cos \theta_i, r_i \sin \theta_i | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) dr_i \\
&\leq \frac{2(1+\delta)^2}{\delta^2} \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\cos \theta_i + \sin \theta_i}}^{\infty} \frac{1}{r_i} \int_{a=0}^{\infty} \int_{b=0}^{\infty} f(a, b | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) I\{(a, b) \in R(r_i, \theta_i)\} db da dr_i \\
&\leq \frac{2(1+\delta)^2}{\delta^2} \int_{a=0}^{\infty} \int_{b=0}^{\infty} f(a, b | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\cos \theta_i + \sin \theta_i}}^{\infty} \frac{I\{(a, b) \in R(r_i, \theta_i)\}}{r} dr db da \\
&\leq \frac{2(1+\delta)^2}{\delta^2} \int_{a=0}^{\infty} \int_{b = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)(\cos \theta_i + \sin \theta_i)} - a}^{\infty} f(a, b | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\cos \theta_i + \sin \theta_i}}^{\infty} \frac{I\{(a, b) \in R(r_i, \theta_i)\}}{r} dr_i db da \\
&\leq \frac{2(1+\delta)^2 \log(1+\delta)}{\delta^2} \int_{a=0}^{\infty} \int_{b = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)(\cos \theta_i + \sin \theta_i)} - a}^{\infty} f(a, b | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) db da \\
&= \frac{2(1+\delta)^2 \log(1+\delta)}{\delta^2} Pr \left[x \sim \mathcal{D} \text{ has } \|x\|_1 \geq \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)(\cos \theta_i + \sin \theta_i)} \mid \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i} \right] \\
&\leq \frac{2(1+\delta)^2 \log(1+\delta)}{\delta^2} Pr \left[x \sim \hat{\mathcal{D}} \text{ has } \|x\|_1 \geq \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{2}} \mid \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i} \right].
\end{aligned}$$

We can now bound 6:

$$\begin{aligned}
6 &\leq \frac{4(1+\delta)^2 \log(1+\delta)}{\delta^2} \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{\theta_i} I\left\{ \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{2}} > w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}) \right\} \\
&\quad r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) Pr \left[x_i \sim \hat{\mathcal{D}}_i \text{ has } \|x_i\|_1 \geq \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{2}} \mid \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i} \right] d\theta_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i} \\
&= \frac{4\sqrt{2}(1+\delta)^3 \log(1+\delta)}{\delta^2} \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{\theta_i} I\left\{ \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{2}} > w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}) \right\} \\
&\quad \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{2}} Pr \left[x_i \sim \hat{\mathcal{D}}_i \text{ has } \|x_i\|_1 \geq \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{2}} \mid \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i} \right] d\theta_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i} \\
&\leq \frac{2\sqrt{2}\pi(1+\delta)^3 \log(1+\delta)}{\delta^2} \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) I\left\{ \frac{r_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{2}} > w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}) \right\} \\
&\quad \frac{r_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{2}} Pr \left[x_i \sim \hat{\mathcal{D}}_i \text{ has } \|x_i\|_1 \geq \frac{r_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{2}} \mid \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i} \right] d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i},
\end{aligned}$$

where the last inequality follows from switching from $r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})$ to $r_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})$ and paying an additional factor of $\pi/2$ (to remove the integral with respect to θ_i).

Let \mathcal{A} be the set of auctions that sells the grand bundle as a whole, and only sells to the highest bidder. The auction in the RHS offers the grand bundle to buyer i for some price q_i , such that $q_i \geq w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})$ (and therefore it sells to at most one buyer), and is an auction in \mathcal{A} . $\text{RONENREV}(\hat{\mathcal{D}})$ is the revenue maximizing auction in \mathcal{A} , therefore:

$$6 \leq \frac{2\sqrt{2}\pi(1+\delta)^3 \log(1+\delta)}{\delta^2} \text{RONENREV}(\hat{\mathcal{D}}).$$

Putting it all together. Combining our bound for 5 and 6 we get:

$$\begin{aligned}
\text{REV}(\hat{\mathcal{D}}) &\leq \left(2\sqrt{2}(1+\delta) + \frac{2\sqrt{2}\pi(1+\delta)^3 \log(1+\delta)}{\delta^2} \right) \text{RONENREV}(\hat{\mathcal{D}}) \\
&= \frac{2\sqrt{2}(1+\delta)}{\delta^2} (\delta^2 + \pi(1+\delta)^2 \log(1+\delta)) \text{RONENREV}(\hat{\mathcal{D}}) \\
&\leq \frac{9\sqrt{2}(1+\delta)^3 \log(1+\delta)}{\delta^2} \text{RONENREV}(\hat{\mathcal{D}}). \quad \square
\end{aligned}$$

4.5 Square-Shift: One buyer, Many Items.

Theorem 1.3. *In the Square-Shift model, for a single additive buyer and m items*

$$\text{REV}(\hat{\mathcal{D}}) \leq \left(\sqrt{\frac{\pi e}{2}} \frac{(1+\delta)}{\delta} \right)^m (1+\delta) \log(1+\delta) m \sqrt{m} \cdot \text{BREV}(\hat{\mathcal{D}}).$$

Proof. Our first few steps are the same. In order to write revenue in polar coordinates we need $m-1$ angles $\boldsymbol{\theta} = (\theta_1, \dots, \theta_{m-1})$ and a length r :

$$\begin{aligned}
\text{REV}(\hat{\mathcal{D}}) &= \int_{\boldsymbol{\theta} \in [0, \pi/2]^{m-1}} \int_{r=0}^{\infty} p^*(r, \boldsymbol{\theta}) \hat{g}(r, \boldsymbol{\theta}) dr d\boldsymbol{\theta} = \int_{\boldsymbol{\theta} \in [0, \pi/2]^{m-1}} \int_{r=0}^{\infty} p^*(r, \boldsymbol{\theta}) \hat{g}(\boldsymbol{\theta}) \hat{g}(r|\boldsymbol{\theta}) dr d\boldsymbol{\theta} \\
&\stackrel{\text{Lemma 6}}{\leq} \int_{\boldsymbol{\theta} \in [0, \pi/2]^{m-1}} \hat{g}(\boldsymbol{\theta}) \text{REV}(\hat{\mathcal{D}}_{\boldsymbol{\theta}}) d\boldsymbol{\theta} = \int_{\boldsymbol{\theta} \in [0, \pi/2]^{m-1}} \hat{g}(\boldsymbol{\theta}) \left(r_{\boldsymbol{\theta}} \int_{r=\frac{r_{\boldsymbol{\theta}}}{\sum_{j=1}^m \text{trig}_j}}^{\infty} \hat{g}(r|\boldsymbol{\theta}) dr \right) d\boldsymbol{\theta}
\end{aligned}$$

$$= \int_{\boldsymbol{\theta} \in [0, \pi/2]^{m-1}} r_{\boldsymbol{\theta}} \int_{r=\frac{r_{\boldsymbol{\theta}}}{\sum_{j=1}^m \text{trig}_j}}^{\infty} \hat{g}(r, \boldsymbol{\theta}) dr d\boldsymbol{\theta}. \quad (7)$$

Next, we bound the integral $\int_{r=\frac{r_{\boldsymbol{\theta}}}{\sum_{j=1}^m \text{trig}_j}}^{\infty} \hat{g}(r, \boldsymbol{\theta}) dr$. Our first step is to go from polar to cartesian coordinates. This time, the transformation is a bit more complex. More specifically, given a point $(w, \phi_1, \dots, \phi_{m-1})$ in m dimensions expressed in polar coordinates, the corresponding point in Cartesian coordinates is \mathbf{x} , where $x_i = w \cos \phi_i \prod_{j=1}^{i-1} \sin \phi_j$. The transformation from \hat{g} to \hat{f} is given in Claim 13. Furthermore, we need to upper bound \hat{f} by f . For the square model the two are connected by the following equations:

$$\hat{f}(\hat{\mathbf{x}}) = \int_{\mathbf{x} \in [0, \infty]^m} \frac{f(\mathbf{x}) I\{\mathbf{x} \in R(\hat{\mathbf{x}})\}}{(\delta \max_{j \in [1, m]} x_j)^m} d\mathbf{x} \leq \frac{(1 + \delta)^m}{(\delta \max_{j \in [1, m]} \hat{x}_j)^m} \int_{\mathbf{x} \in [0, \infty]^m} f(\mathbf{x}) I\{\mathbf{x} \in R(\hat{\mathbf{x}})\} d\mathbf{x}, \quad (8)$$

where $I\{\mathbf{x} \in R(\hat{\mathbf{x}})\}$ is an indicator for the event that \mathbf{x} belongs in $R(\hat{\mathbf{x}})$. The inequality follows from the fact that for $I\{\mathbf{x} \in R(\hat{\mathbf{x}})\}$ to be 1 we must have that $\max_{j \in [1, m]} x_j \geq \frac{\max_{j \in [1, m]} \hat{x}_j}{1 + \delta}$.

$$\begin{aligned} \int_{r=\frac{r_{\boldsymbol{\theta}}}{\sum_{j=1}^m \text{trig}_j}}^{\infty} \hat{g}(r, \boldsymbol{\theta}) dr &= \int_{r=\frac{r_{\boldsymbol{\theta}}}{\sum_{j=1}^m \text{trig}_j}}^{\infty} r^{m-1} \prod_{j=1}^{m-2} (\sin \theta_j)^{m-j-1} \hat{f}(r \cdot \text{trig}) dr \\ &\leq \int_{r=\frac{r_{\boldsymbol{\theta}}}{\sum_{j=1}^m \text{trig}_j}}^{\infty} r^{m-1} \prod_{j=1}^{m-2} (\sin \theta_j)^{m-j-1} \\ &\quad \left(\frac{(1 + \delta)^m}{(\delta r \max_{j \in [1, m]} \text{trig}_j)^m} \int_{\mathbf{x} \in [0, \infty]^m} f(\mathbf{x}) I\{\mathbf{x} \in R(r \cdot \text{trig})\} d\mathbf{x} \right) dr \\ &\stackrel{\text{Cl. 12}}{\leq} \left(\frac{(1 + \delta) \sqrt{m}}{\delta} \right)^m \prod_{j=1}^{m-2} (\sin \theta_j)^{m-j-1} \int_{r=\frac{r_{\boldsymbol{\theta}}}{\sum_{j=1}^m \text{trig}_j}}^{\infty} \int_{\mathbf{x} \in [0, \infty]^m} \frac{1}{r} f(\mathbf{x}) I\{\mathbf{x} \in R(r \cdot \text{trig})\} d\mathbf{x} dr \\ &= \left(\frac{(1 + \delta) \sqrt{m}}{\delta} \right)^m \prod_{j=1}^{m-2} (\sin \theta_j)^{m-j-1} \int_{\mathbf{x} \in [0, \infty]^m} f(\mathbf{x}) \int_{r=\frac{r_{\boldsymbol{\theta}}}{\sum_{j=1}^m \text{trig}_j}}^{\infty} \frac{1}{r} I\{\mathbf{x} \in R(r \cdot \text{trig})\} dr d\mathbf{x}. \end{aligned}$$

Next, we argue about values of \mathbf{x} such that the indicator $I\{\mathbf{x} \in R(r \cdot \text{trig})\}$ is non-zero.

Claim 7. $\forall \mathbf{x} \in R(r \cdot \text{trig})$, if $r \geq \frac{r_{\boldsymbol{\theta}}}{\sum_{j=1}^m \text{trig}_j}$ then $\sum_{j=1}^m x_j \geq \frac{r_{\boldsymbol{\theta}}}{(1 + \delta)(\sum_{j=1}^m \text{trig}_j)}$.

Proof. Similarly to Claim 4, note that for any vector $\mathbf{x} \in \mathbb{R}_m^+$, $\sum_{i=1}^m x_i \geq r_{\mathbf{x}} = \sqrt{\sum x_i^2}$. Moreover, any vector \mathbf{x} that can map to $(r, \boldsymbol{\theta})$ must have its own length, $r_{\mathbf{x}}$ be at least $\frac{1}{1 + \delta} r$. Combining these two claims with that given on the statement completes the proof. \square

We apply Claim 7 to change the limit of integration:

$$\begin{aligned} \int_{r=\frac{r_{\boldsymbol{\theta}}}{\sum_{j=1}^m \text{trig}_j}}^{\infty} \hat{g}(r, \boldsymbol{\theta}) dr &\leq \left(\frac{(1 + \delta) \sqrt{m}}{\delta} \right)^m \prod_{j=1}^{m-2} (\sin \theta_j)^{m-j-1} \\ &\quad \cdot \int_{\mathbf{x}: \|\mathbf{x}\|_1 \geq \frac{r_{\boldsymbol{\theta}}}{(1 + \delta)(\sum_{j=1}^m \text{trig}_j)}} f(\mathbf{x}) \int_{r=\frac{r_{\boldsymbol{\theta}}}{\sum_{j=1}^m \text{trig}_j}}^{\infty} \frac{1}{r} I\{\mathbf{x} \in R(r \cdot \text{trig})\} dr d\mathbf{x}. \end{aligned}$$

We can now upper bound the value of the inner most integral.

Claim 8. For the current perturbation model, for any $\forall \mathbf{x} \in R(r \cdot \text{trig})$,

$$\int_{r=\frac{r_{\boldsymbol{\theta}}}{\sum_{j=1}^m \text{trig}_j}}^{\infty} \frac{1}{r} I\{\mathbf{x} \in R(r \cdot \text{trig})\} dr \leq \log(1 + \delta).$$

Proof. The points that \mathbf{x} can map to form a hypercube, and the indicator $I\{\mathbf{x} \in R(r \cdot \text{trig})\}$ is non-zero for a length r , when the point $(r, \boldsymbol{\theta})$ intersects that hypercube. Let r_{\min} be the smallest length among all points $(r, \boldsymbol{\theta})$ in the hypercube defined by \mathbf{x} . Then, the maximum length cannot be larger than $r_{\min}(1 + \delta)$. Therefore

$$\int_{r=\frac{r_{\boldsymbol{\theta}}}{\sum_{j=1}^m \text{trig}_j}}^{\infty} \frac{1}{r} I\{\mathbf{x} \in R(r \cdot \text{trig})\} dr \leq \int_{r=r_{\min}}^{(1+\delta)r_{\min}} \frac{1}{r} dr.$$

Integrating the latter integral proves the Claim. \square

We apply Claim 8 to continue our derivation,

$$\begin{aligned} \int_{r=\frac{r_{\boldsymbol{\theta}}}{\sum_{j=1}^m \text{trig}_j}}^{\infty} \hat{g}(r, \boldsymbol{\theta}) dr &\leq \left(\frac{(1 + \delta)\sqrt{m}}{\delta} \right)^m \log(1 + \delta) \prod_{j=1}^{m-2} (\sin \theta_j)^{m-j-1} \int_{\substack{\mathbf{x} \in [0, \infty)^m \text{ such that} \\ \|\mathbf{x}\|_1 \geq \frac{r_{\boldsymbol{\theta}}}{(1+\delta)(\sum_{j=1}^m \text{trig}_j)}}} f(\mathbf{x}) d\mathbf{x} \\ &= \left(\frac{(1 + \delta)\sqrt{m}}{\delta} \right)^m \log(1 + \delta) \prod_{j=1}^{m-2} (\sin \theta_j)^{m-j-1} Pr \left[\mathbf{x} \sim \mathcal{D} \text{ has } \|\mathbf{x}\|_1 \geq \frac{r_{\boldsymbol{\theta}}}{(1 + \delta)(\sum_{j=1}^m \text{trig}_j)} \right] \\ &\leq \left(\frac{(1 + \delta)\sqrt{m}}{\delta} \right)^m \log(1 + \delta) \prod_{j=1}^{m-2} (\sin \theta_j)^{m-j-1} Pr \left[\mathbf{x} \sim \hat{\mathcal{D}} \text{ has } \|\mathbf{x}\|_1 \geq \frac{r_{\boldsymbol{\theta}}}{(1 + \delta)(\sum_{j=1}^m \text{trig}_j)} \right] \\ &\leq \left(\frac{(1 + \delta)\sqrt{m}}{\delta} \right)^m \log(1 + \delta) \prod_{j=1}^{m-2} (\sin \theta_j)^{m-j-1} Pr \left[\mathbf{x} \sim \hat{\mathcal{D}} \text{ has } \|\mathbf{x}\|_1 \geq \frac{r_{\boldsymbol{\theta}}}{(1 + \delta)\sqrt{m}} \right], \end{aligned}$$

where we used the facts that perturbing only increases values, and that $\sum_{j=1}^m \text{trig}_j \leq \sqrt{m}$. Plugging back in Equation 7:

$$\begin{aligned} \text{REV}(\hat{\mathcal{D}}) &\leq \int_{\boldsymbol{\theta} \in [0, \pi/2]^{m-1}} r_{\boldsymbol{\theta}} \cdot \\ &\left(\left(\frac{(1 + \delta)\sqrt{m}}{\delta} \right)^m \log(1 + \delta) \prod_{j=1}^{m-2} (\sin \theta_j)^{m-j-1} Pr \left[\text{sample from } \hat{\mathcal{D}} \text{ has } \|\mathbf{x}\|_1 \geq \frac{r_{\boldsymbol{\theta}}}{(1 + \delta)\sqrt{m}} \right] \right) d\boldsymbol{\theta} \\ &\leq \left(\frac{(1 + \delta)\sqrt{m}}{\delta} \right)^m \log(1 + \delta) (1 + \delta) \sqrt{m} \int_{\boldsymbol{\theta} \in [0, \pi/2]^{m-1}} \prod_{j=1}^{m-2} (\sin \theta_j)^{m-j-1} \text{BREV}(\hat{\mathcal{D}}) d\boldsymbol{\theta} \\ &= \text{BREV}(\hat{\mathcal{D}}) \left(\frac{(1 + \delta)\sqrt{m}}{\delta} \right)^m (1 + \delta) \log(1 + \delta) \sqrt{m} \int_{\boldsymbol{\theta} \in [0, \pi/2]^{m-1}} \prod_{j=1}^{m-2} (\sin \theta_j)^{m-j-1} d\boldsymbol{\theta}, \end{aligned}$$

where we used the fact that $\text{BREV}(\hat{\mathcal{D}})$ is at least $r \cdot Pr \left[\sum_{j=1}^m x_j \geq r \right]$ for all r , where (x_1, \dots, x_m) is a sample from $\hat{\mathcal{D}}$. The next Claim is used to simplify the integral over the product of sine functions:

Claim 9.

$$\int_{\boldsymbol{\theta} \in [0, \pi/2]^{m-1}} \prod_{j=1}^{m-2} (\sin \theta_j)^{m-j-1} d\boldsymbol{\theta} \leq \left(\sqrt{\frac{\pi e}{2}} \right)^m \frac{m}{(\sqrt{m})^m}.$$

Proof. We repeatedly use the fact that $\int_{\theta=0}^{\pi/2} (\sin \theta)^k d\theta = \frac{\sqrt{\pi}}{2} \frac{\Gamma(\frac{k+1}{2})}{\Gamma(\frac{k}{2}+1)}$.

$$\begin{aligned} & \int_{\theta_1=0}^{\pi/2} (\sin \theta_1)^{m-2} \int_{\theta_2=0}^{\pi/2} (\sin \theta_2)^{m-3} \int_{\theta_3=0}^{\pi/2} \cdots \int_{\theta_{m-2}=0}^{\pi/2} \sin \theta_{m-2} \left(\int_{\theta_{m-1}=0}^{\pi/2} 1 d\theta_{m-1} \right) d\theta_{m-2} \cdots d\theta_1 \\ &= \frac{\pi}{2} \int_{\theta_1=0}^{\pi/2} (\sin \theta_1)^{m-2} \int_{\theta_2=0}^{\pi/2} (\sin \theta_2)^{m-3} \int_{\theta_3=0}^{\pi/2} \cdots \left(\int_{\theta_{m-2}=0}^{\pi/2} \sin \theta_{m-2} d\theta_{m-2} \right) \cdots d\theta_1 \\ &= \frac{\pi}{2} \frac{\sqrt{\pi}}{2} \frac{\Gamma(1)}{\Gamma(\frac{1}{2}+1)} \int_{\theta_1=0}^{\pi/2} (\sin \theta_1)^{m-2} \int_{\theta_2=0}^{\pi/2} (\sin \theta_2)^{m-3} \int_{\theta_3=0}^{\pi/2} \cdots \left(\int_{\theta_{m-3}=0}^{\pi/2} (\sin \theta_{m-3})^2 d\theta_{m-3} \right) \cdots d\theta_1 \\ &= \frac{\pi}{2} \left(\frac{\sqrt{\pi}}{2} \right)^2 \frac{\Gamma(1)}{\Gamma(\frac{1}{2}+1)} \frac{\Gamma(\frac{1}{2}+1)}{\Gamma(2)} \int_{\theta_1=0}^{\pi/2} (\sin \theta_1)^{m-2} \int_{\theta_2=0}^{\pi/2} \cdots \left(\int_{\theta_{m-4}=0}^{\pi/2} (\sin \theta_{m-4})^3 d\theta_{m-4} \right) \cdots d\theta_1 \\ &= \frac{\pi}{2} \left(\frac{\sqrt{\pi}}{2} \right)^{m-2} \frac{\Gamma(1)}{\Gamma(\frac{m-2}{2}+1)} \\ &= \frac{\pi}{2} \left(\frac{\sqrt{\pi}}{2} \right)^{m-2} \frac{1}{\Gamma(\frac{m}{2})}. \end{aligned}$$

Recall that for integer n , $\Gamma(n+1) = n!$, therefore $\Gamma(\frac{m}{2}) = (\frac{m}{2}-1)!$. Using the fact that $k! \geq (\frac{k}{e})^k$ we get:

$$\begin{aligned} \frac{\pi}{2} \left(\frac{\sqrt{\pi}}{2} \right)^{m-2} \frac{1}{\Gamma(\frac{m}{2})} &= \frac{\pi}{2} \left(\frac{\sqrt{\pi}}{2} \right)^{m-2} \frac{1}{(\frac{m}{2}-1)!} \\ &= \frac{\pi}{2} \left(\frac{\sqrt{\pi}}{2} \right)^{m-2} \frac{\frac{m}{2}}{(\frac{m}{2})!} \\ &\leq \frac{\pi}{4} \left(\frac{\sqrt{\pi}}{2} \right)^{m-2} \frac{m}{(\frac{m}{2e})^{m/2}} \\ &= \left(\frac{\sqrt{\pi}}{2} \right)^m (2e)^{m/2} \frac{m}{(\sqrt{m})^m} \\ &= \left(\sqrt{\frac{\pi e}{2}} \right)^m \frac{m}{(\sqrt{m})^m}, \end{aligned}$$

which concludes the proof of the Claim. \square

Applying Claim 9 completes the proof of Theorem 1.3:

$$\begin{aligned} \text{REV}(\hat{\mathcal{D}}) &\leq \text{BREV}(\hat{\mathcal{D}}) \left(\frac{(1+\delta)\sqrt{m}}{\delta} \right)^m (1+\delta) \log(1+\delta) \sqrt{m} \left(\sqrt{\frac{\pi e}{2}} \right)^m \frac{m}{(\sqrt{m})^m} \\ &= \text{BREV}(\hat{\mathcal{D}}) \left(\sqrt{\frac{\pi e}{2}} \frac{(1+\delta)}{\delta} \right)^m (1+\delta) \log(1+\delta) m \sqrt{m}. \end{aligned} \quad \square$$

4.6 Square-Shift: Many buyers, Many Items.

In this section we prove upper bounds for the case of multiple (possibly correlated) buyers and multiple items, in the Square-Shift model.

Theorem 1.4. *For the n additive buyer, m item, Square-Shift model,*

$$\text{REV}(\hat{\mathcal{D}}) \leq 4 \left(\sqrt{\frac{\pi e}{2}} \frac{(1 + \delta)}{\delta} \right)^m (1 + \delta) \log(1 + \delta) m \sqrt{m} \cdot \text{RONENREV}(\hat{\mathcal{D}}),$$

where $\text{REV}(\hat{\mathcal{D}}) = \text{DSICREV}(\hat{\mathcal{D}})$ for correlated buyers, and $\text{REV}(\hat{\mathcal{D}}) = \text{BICREV}(\hat{\mathcal{D}})$ for independent buyers.

Proof. We write $\boldsymbol{\theta}$ for the vector $(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_n)$, where $\boldsymbol{\theta}_i = (\theta_{i,1}, \dots, \theta_{i,m-1})$ is the vector of angles for buyer i :

$$\text{REV}(\hat{\mathcal{D}}) = \int_{\boldsymbol{\theta}} \hat{g}(\boldsymbol{\theta}) \int_{\mathbf{r}} \hat{g}(\mathbf{r}|\boldsymbol{\theta}) \sum_{i=1}^n p_i^*(\mathbf{r}, \boldsymbol{\theta}) d\mathbf{r} d\boldsymbol{\theta} \leq \int_{\boldsymbol{\theta}} \hat{g}(\boldsymbol{\theta}) \text{REV}(\hat{\mathcal{D}}_{\boldsymbol{\theta}}) d\boldsymbol{\theta},$$

where $\hat{\mathcal{D}}_{\boldsymbol{\theta}}$ is a correlated, single parameter distribution where we draw a value vector $\mathbf{r} = (r_1, \dots, r_n)$ according to the density function $\hat{g}(\mathbf{r}|\boldsymbol{\theta})$, and give each buyer i value $r_i \cdot \text{trig}_k(\boldsymbol{\theta}_i)$ for item k . Ronen's mechanism [31] for this distribution makes a take it or leave it offer $r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})$ to buyer i for the grand bundle; $r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})$ is larger than the value of any other buyer, i.e. at least $\max_{j \neq i} (r_j \sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_j))$. This mechanism is a 2 approximation to the optimal mechanism for $\hat{\mathcal{D}}_{\boldsymbol{\theta}}$. Let $w_{-i}^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}) = \max_{j \neq i} (r_j \sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_j))$ be the second highest valuation among all buyers excluding i , and $r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) = \arg \max_{w \geq w_{-i}^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})} (w \cdot \Pr[r_i \sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i) \geq w])$ be the price offered by Ronen's mechanism.

$$\begin{aligned} \text{REV}(\hat{\mathcal{D}}) &\leq 2 \int_{\boldsymbol{\theta}} \hat{g}(\boldsymbol{\theta}) \left(\int_{\mathbf{r}} \hat{g}(\mathbf{r}|\boldsymbol{\theta}) \sum_{i=1}^n r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) I\{r_i \sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i) \geq r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})\} d\mathbf{r} \right) d\boldsymbol{\theta} \\ &= 2 \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{\boldsymbol{\theta}_i} r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)}}^{\infty} \hat{g}(r_i, \boldsymbol{\theta}_i | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) dr_i d\boldsymbol{\theta}_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i} \quad (9) \end{aligned}$$

$$\begin{aligned} &= 2 \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{\boldsymbol{\theta}_i} I\{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) \leq \sqrt{m}(1 + \delta) w_{-i}^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})\} r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) \\ &\quad \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)}}^{\infty} \hat{g}(r_i, \boldsymbol{\theta}_i | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) dr_i d\boldsymbol{\theta}_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i} \quad (10) \end{aligned}$$

$$\begin{aligned} &+ 2 \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{\boldsymbol{\theta}_i} I\{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) > \sqrt{m}(1 + \delta) w_{-i}^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})\} r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) \\ &\quad \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)}}^{\infty} \hat{g}(r_i, \boldsymbol{\theta}_i | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) dr_i d\boldsymbol{\theta}_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i}. \quad (11) \end{aligned}$$

Similarly to the two item case, bounding 9 is problematic. To go around the issue we separate 9 cases depending on whether $r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})$ is almost the second highest value, or at least $\sqrt{m}(1 + \delta)$ times the second highest value. We get terms 10 and 11 which we bound separately.

Bounding 10. We bound 10 by the revenue of a second price auction for the grand bundle. Notice that replacing $r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})$ with $w_{-i}^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})$ makes the probability of sale increase, and the price decreases by a factor of at most $\frac{1}{\sqrt{m}(1 + \delta)}$:

$$10 \leq 2 \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{\boldsymbol{\theta}_i} (1 + \delta) w_{-i}^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})$$

$$\begin{aligned}
& \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})}{\sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)}}^{\infty} \frac{w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})}{\sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)} \hat{g}(r_i, \boldsymbol{\theta}_i | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) dr_i d\boldsymbol{\theta}_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i} \\
&= 2\sqrt{m}(1+\delta) \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}) \\
& \quad \int_{\boldsymbol{\theta}_i} \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})}{\sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)}}^{\infty} \frac{w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})}{\sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)} \hat{g}(r_i, \boldsymbol{\theta}_i | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) dr_i d\boldsymbol{\theta}_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i} \\
&= 2\sqrt{m}(1+\delta) \text{SECONDPRICE}(\hat{\mathcal{D}}) \leq 2\sqrt{m}(1+\delta) \text{RONENREV}(\hat{\mathcal{D}}).
\end{aligned}$$

Bounding 11. Our approach for bounding 11 will be similar to the single buyer case. Specifically, we first upper bound the inner most integral, for fixed $\boldsymbol{\theta}$ and \mathbf{r}_{-i} , by switching to cartesian coordinates, applying the formula to go from \hat{f} to f for our perturbation model, changing order of integration using Fubini's theorem, and then arguing about the size of the area that gets mapped to the point $r_i \sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)$.

$$\begin{aligned}
& \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)}}^{\infty} \hat{g}(r_i, \boldsymbol{\theta}_i | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) dr_i \\
&= \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)}}^{\infty} r_i^{m-1} \prod_{k=1}^{m-2} (\sin \boldsymbol{\theta}_{i,k})^{m-k-1} \hat{f}(r_i \sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i) | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) dr_i \\
&\leq \left(\frac{(1+\delta)\sqrt{m}}{\delta} \right)^m \prod_{k=1}^{m-2} (\sin \boldsymbol{\theta}_{i,k})^{m-k-1} \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)}}^{\infty} \frac{1}{r_i} \int_{\mathbf{x}} f(\mathbf{x} | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) I\{\mathbf{x} \in R(r_i, \boldsymbol{\theta}_i)\} d\mathbf{x} dr_i \\
&\leq \left(\frac{(1+\delta)\sqrt{m}}{\delta} \right)^m \prod_{k=1}^{m-2} (\sin \boldsymbol{\theta}_{i,k})^{m-k-1} \int_{\mathbf{x}} f(\mathbf{x} | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)}}^{\infty} \frac{I\{\mathbf{x} \in R(r_i, \boldsymbol{\theta}_i)\}}{r_i} dr_i d\mathbf{x} \\
&\leq \left(\frac{(1+\delta)\sqrt{m}}{\delta} \right)^m \prod_{k=1}^{m-2} (\sin \boldsymbol{\theta}_{i,k})^{m-k-1} \\
& \quad \int_{\mathbf{x}: \|\mathbf{x}\| \geq \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta) \sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)}} f(\mathbf{x} | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{r_i = \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{\sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)}}^{\infty} \frac{I\{\mathbf{x} \in R(r_i, \boldsymbol{\theta}_i)\}}{r_i} dr_i d\mathbf{x} \\
&\leq \left(\frac{(1+\delta)\sqrt{m}}{\delta} \right)^m \prod_{k=1}^{m-2} (\sin \boldsymbol{\theta}_{i,k})^{m-k-1} \log(1+\delta) \int_{\mathbf{x}: \|\mathbf{x}\| \geq \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta) \sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)}} f(\mathbf{x} | \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \\
&\leq \left(\frac{(1+\delta)\sqrt{m}}{\delta} \right)^m \prod_{k=1}^{m-2} (\sin \boldsymbol{\theta}_{i,k})^{m-k-1} \log(1+\delta) Pr \left[\mathbf{x} \sim \mathcal{D} \text{ has } \|\mathbf{x}\|_1 \geq \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta) \sum_{k=1}^m \text{trig}_k(\boldsymbol{\theta}_i)} \mid \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i} \right] \\
&\leq \left(\frac{(1+\delta)\sqrt{m}}{\delta} \right)^m \prod_{k=1}^{m-2} (\sin \boldsymbol{\theta}_{i,k})^{m-k-1} \log(1+\delta) Pr \left[\mathbf{x} \sim \hat{\mathcal{D}} \text{ has } \|\mathbf{x}\|_1 \geq \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{m}} \mid \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i} \right].
\end{aligned}$$

We can now bound 11:

$$\begin{aligned}
11 &\leq 2 \left(\frac{(1+\delta)\sqrt{m}}{\delta} \right)^m \log(1+\delta) \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{\boldsymbol{\theta}_i} \prod_{k=1}^{m-2} (\sin \boldsymbol{\theta}_{i,k})^{m-k-1} \\
& \quad I\left\{ \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{m}} > w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}) \right\} r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) Pr \left[\mathbf{x} \sim \hat{\mathcal{D}} \text{ has } \|\mathbf{x}\|_1 \geq \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{m}} \mid \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i} \right] d\boldsymbol{\theta}_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i} \\
&= 2\sqrt{m}(1+\delta) \left(\frac{(1+\delta)\sqrt{m}}{\delta} \right)^m \log(1+\delta) \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{\boldsymbol{\theta}_i} \prod_{k=1}^{m-2} (\sin \boldsymbol{\theta}_{i,k})^{m-k-1} \\
& \quad I\left\{ \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{m}} > w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}) \right\} \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{m}} Pr \left[\mathbf{x} \sim \hat{\mathcal{D}} \text{ has } \|\mathbf{x}\|_1 \geq \frac{r_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i})}{(1+\delta)\sqrt{m}} \mid \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i} \right] d\boldsymbol{\theta}_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i} \\
&= 2\sqrt{m}(1+\delta) \left(\frac{(1+\delta)\sqrt{m}}{\delta} \right)^m \log(1+\delta) \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{\boldsymbol{\theta}_i} \prod_{k=1}^{m-2} (\sin \boldsymbol{\theta}_{i,k})^{m-k-1} \\
& \quad I\{q_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) > w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})\} q_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) Pr \left[\mathbf{x} \sim \hat{\mathcal{D}} : \|\mathbf{x}\|_1 \geq q_i^*(\boldsymbol{\theta}, \mathbf{r}_{-i}) \mid \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i} \right] d\boldsymbol{\theta}_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i} \\
&\leq 2\sqrt{m}(1+\delta) \left(\frac{(1+\delta)\sqrt{m}}{\delta} \right)^m \log(1+\delta) \sum_{i=1}^n \int_{\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}} \hat{g}(\mathbf{r}_{-i}, \boldsymbol{\theta}_{-i}) \int_{\boldsymbol{\theta}_i} \prod_{k=1}^{m-2} (\sin \boldsymbol{\theta}_{i,k})^{m-k-1} \\
& \quad I\{q_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}) > w_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i})\} q_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}) Pr \left[\mathbf{x} \sim \hat{\mathcal{D}} : \|\mathbf{x}\|_1 \geq q_i^*(\boldsymbol{\theta}_{-i}, \mathbf{r}_{-i}) \mid \mathbf{r}_{-i}, \boldsymbol{\theta}_{-i} \right] d\boldsymbol{\theta}_i d\mathbf{r}_{-i} d\boldsymbol{\theta}_{-i}
\end{aligned}$$

$$\begin{aligned}
&\leq 2\sqrt{m}(1+\delta) \left(\frac{(1+\delta)\sqrt{m}}{\delta}\right)^m \log(1+\delta) \cdot \left(\left(\sqrt{\frac{\pi e}{2}}\right)^m \frac{m}{(\sqrt{m})^m}\right) \cdot \text{RONENREV}(\hat{\mathcal{D}}) \\
&= 2 \left(\sqrt{\frac{\pi e}{2}} \frac{(1+\delta)}{\delta}\right)^m (1+\delta) \log(1+\delta) m \sqrt{m} \cdot \text{RONENREV}(\hat{\mathcal{D}}).
\end{aligned}$$

Putting it all together. Combining our bound for 10 and 11 we get:

$$\text{REV}(\hat{\mathcal{D}}) \leq 4 \left(\sqrt{\frac{\pi e}{2}} \frac{(1+\delta)}{\delta}\right)^m (1+\delta) \log(1+\delta) m \sqrt{m} \cdot \text{RONENREV}(\hat{\mathcal{D}}). \quad \square$$

5 Extensions

It is straightforward to extend our results beyond additive valuations.

Unit-Demand Valuations

Our bounds for unit-demand valuations are the same as the ones for additive valuations, in both the Angle-Shift and the Square-Shift models. Starting from our toy Angle-Shift model, notice that in the proof of Theorem 4.1 Lemma 6 can still be applied, but $\hat{\mathcal{D}}_\theta$ can be replaced by the single parameter distribution where we sample r as before and output $\max\{r \cos \theta, r \sin \theta\}$ (instead of $r(\cos \theta + \sin \theta)$). This would give a term $\int_{r=\frac{r_\theta}{\max\{\sin \theta, \cos \theta\}}}^{\infty} g(r) dr$, which is equal to the probability that a sample from \mathcal{D} has length at least $\frac{r_\theta}{\max\{\sin \theta, \cos \theta\}}$, which is equal to the probability of the same event for $\hat{\mathcal{D}}$, which is equal to $\Pr\left[(x, y) \sim \hat{\mathcal{D}} \text{ has } \max\{x, y\} \geq r_\theta\right]$. The latter probability multiplied by r_θ is upper bounded by $\text{BREV}(\hat{\mathcal{D}})$. Therefore, for this setting, we overall get the same bound of $\frac{\pi\sqrt{2}}{2\delta} \text{BREV}(\hat{\mathcal{D}})$ for unit-demand valuations.

In the Square-Shift model, starting from a single buyer and two items, the same observation about $\hat{\mathcal{D}}_\theta$ holds. When bounding $\int_{r \max\{\cos \theta, \sin \theta\}=r_\theta}^{\infty} \hat{g}(r, \theta) dr$, Claim 4 is replaced by:

Claim 10. $\forall (a, b) \in R(r, \theta)$, if $r \geq \frac{r_\theta}{\max\{\cos \theta, \sin \theta\}}$ then $\max\{a, b\} \geq \frac{r_\theta}{1+\delta}$.

Proof. Recall that $(a, b) \in R(r, \theta)$ means that (a, b) that could map to a point with length r and angle θ . Then, $(1+\delta) \max\{a, b\} \geq \max\{r \cos \theta, r \sin \theta\} \geq r_\theta$. The Claim follows. \square

Claim 5 remains unchanged. Overall we get:

$$\begin{aligned}
\int_{r=\frac{r_\theta}{\max\{\cos \theta, \sin \theta\}}}^{\infty} \hat{g}(r, \theta) dr &\leq \frac{2(1+\delta)^2 \log(1+\delta)}{\delta^2} \int_{a=0}^{\infty} \int_{b=\frac{r_\theta}{1+\delta}-a}^{\infty} f(a, b) db da \\
&= \frac{2(1+\delta)^2 \log(1+\delta)}{\delta^2} \Pr \left[\text{sample from } \mathcal{D} \text{ has maximum value at least } \frac{r_\theta}{1+\delta} \right] \\
&\leq \frac{2(1+\delta)^2 \log(1+\delta)}{\delta^2} \Pr \left[\text{sample from } \hat{\mathcal{D}} \text{ has maximum value at least } \frac{r_\theta}{1+\delta} \right].
\end{aligned}$$

Therefore, the bound $\text{REV}(\hat{\mathcal{D}}) \leq \frac{\sqrt{2}\pi(1+\delta)^3 \log(1+\delta)}{\delta^2} \text{BREV}(\hat{\mathcal{D}})$ is valid a unit demand buyer and two items in the Square-Shift model. For a single buyer and m items we reach the same conclusion by replacing Claim 7 appropriately. Similarly for the many buyer case two item, and many buyer multi-item cases: 5 is bounded by a second price auction for the grand bundle (where this time buyer i 's value for the grand bundle is $\max\{v_{i,1}, v_{i,2}\}$ and not $v_{i,1} + v_{i,2}$), and the second term (6) is bounded similarly to the single buyer case.

Additive subject to downward-closed constraints \mathcal{I}

Let $\text{REV}(\mathcal{D})$ be the revenue of the optimal truthful mechanism on an m -dimensional, n (correlated) buyer distribution \mathcal{D} , where the utility of buyer i for a subset of the items S is equal $v_i(S) = \max_{T \subseteq S, T \in \mathcal{I}} \{\sum_{j \in T} v_{i,j}\}$. Also, let $\text{ADDITIVEREV}(\mathcal{D})$ be the optimal revenue of the same distribution, but this time buyers have additive valuation functions.

Claim 11. $\text{REV}(\mathcal{D}) \leq \text{ADDITIVEREV}(\mathcal{D})$.

Therefore, we can apply our approximations for the additive setting. It remains to observe that the revenue of bundling under additive valuations is at most m times the revenue of bundling under downward-closed constraints \mathcal{I} , i.e. we lose a factor m in the approximation guarantee.

6 Conclusion and Future Directions

We initiate the smoothed analysis of multi-item auctions with arbitrarily correlated buyers. We present two main results: first, the [5, 22] construction is surprisingly robust, and the infinite gap persists in our “Rectangle-Shift” model. Second, we show that a smoothed-finite approximation is indeed possible in our “Square-Shift” model, with extensions to multiple buyers and multiple items. As a bonus technical insight, we learn that whether or not an infinite gap exists between simple and optimal for some distribution \mathcal{D} is intimately connected to what \mathcal{D} looks like in small cones.

Our work takes the view of studying truly simple mechanisms whose approximation ratios degrade with the smoothing parameter. An alternative approach would be to study truly good approximation ratios via mechanisms whose complexity degrades with the smoothing parameter. To be a little more concrete, an interesting open question is the following: in the Square-Shift model, what is the required *menu complexity* of an auction (that is, the number of possible outcomes based on the buyers’ input) to guarantee a $(1 - \varepsilon)$ -approximation if $\hat{\mathcal{D}}$ is smoothed with parameter δ ?

A general direction for future work is exploring alternative smoothed models. As an example, for many items our guarantees are exponential in m , and this is tight for our model. Would an alternative natural smoothed model provide an improved guarantee, or is there some inherent barrier?

The “independent items” paradigm led to numerous developments which deepened our understanding of multi-item auctions. This work shows that smoothed analysis can help guide future work to better understand worst-case distributions beyond the impossibility results of [5, 22].

A Results and Figures

	Independent	Correlated	Square-Shift
$n = 1, m = 2$	$\frac{\text{REV}(\mathcal{D})}{\text{BREV}(\mathcal{D})} \leq 2$ [21]	$\frac{\text{REV}(\mathcal{D})}{\text{BREV}(\mathcal{D})}$ unbounded [22]	$\frac{\text{REV}(\mathcal{D})}{\text{BREV}(\mathcal{D})} \leq \frac{\sqrt{2\pi(1+\delta)^3 \log(1+\delta)}}{\delta^2}$
$n = 1, \text{ general } m$	$\frac{\text{REV}(\mathcal{D})}{\max\{\text{SREV}(\mathcal{D}), \text{BREV}(\mathcal{D})\}} \leq 6$ [2]		$\frac{2^m}{m} \leq \frac{\text{REV}(\mathcal{D})}{\text{BREV}(\mathcal{D})} \leq \left(\sqrt{\frac{\pi e(1+\delta)}{2}}\right)^m (1+\delta) \log(1+\delta) m \sqrt{m}$
general $n, m = 2$	$\frac{\text{DSICREV}(\mathcal{D})}{\text{SREV}(\mathcal{D})} \in O(1)$ [2, 35]		$\frac{\text{DSICREV}(\mathcal{D})}{\text{RONENREV}(\mathcal{D})} \leq \frac{9\sqrt{2}(1+\delta)^3 \log(1+\delta)}{\delta^2}$
general n, m	$\frac{\text{BICREV}(\mathcal{D})}{\max\{\text{SREV}(\mathcal{D}), \text{BREV}(\mathcal{D})\}} \leq 8$ [35, 7]		$\frac{\text{DSICREV}(\mathcal{D})}{\text{RONENREV}(\mathcal{D})} \in O\left(\frac{(1+\delta)}{\delta}\right)^m (1+\delta) \log(1+\delta) m \sqrt{m}$




Figure 1: Our models.

B Extended Preliminaries

For the sake of completion, we provide a brief discussion on the BIC solution concept for correlated buyers. Note that when buyers are correlated, your prior on the other buyers' values changes depending on your own prior. So Bayesian IC truly means “you prefer to tell the truth, assuming that all other buyers tell the truth (even taking your updated prior into account).” Formally, one needs to define $\pi_{ij}(\mathbf{v}_i, \mathbf{v}'_i)$ to be the probability that buyer i receives item j conditioned on reporting \mathbf{v}_i while their true type is \mathbf{v}'_i . When buyers are independent the true type doesn't affect the conditioning. And then a mechanism is BIC as long as:

$$\mathbf{v}_i \cdot \bar{\pi}_i(\mathbf{v}_i, \mathbf{v}_i) - p_i(\mathbf{v}_i) \geq \mathbf{v}_i \cdot \pi_i(\mathbf{v}_i, \mathbf{v}'_i) - p_i(\mathbf{v}'_i).$$

Also, while the difference between ex-post IR (always receive non-negative utility for reporting the truth) and interim IR (receive non-negative utility for reporting the truth in expectation) is unimportant for auctions with independent buyers (due to a very simple reduction from interim IR to ex-post IR, see e.g. [15]). There's a fundamental difference for correlated buyers. This is because seminal work of Cremer and McLean [14] shows that for any non-generic correlated distribution, it is possible for the seller to extract the full welfare with a DSIC, interim IR mechanism (but this is not possible with a DSIC, ex-post IR mechanism). So all of our results stated for correlated buyers necessarily reference ex-post IR mechanisms (and this is explicitly stated in the body).

C Lower bounds

C.1 Rectangle-Shift lower bounds

Proof of Lemma 1. Let $M_i = \frac{4^i}{\prod_{j < i} \text{gap}_j^2}$. Let \mathcal{D} be the following distribution: the buyer has value $\mathbf{v}_i = M_i \mathbf{x}_i$ with probability $\frac{1}{M_i}$. With the remaining probability $1 - \sum_{i=1}^{\infty} \frac{1}{M_i}$ the value is zero. We

sample from $\hat{\mathcal{D}}$ by first sampling \mathbf{x}_i w.p. $\frac{1}{M_i}$, perturbing with parameter δ to $\hat{\mathbf{x}}_i$, and outputting value $\mathbf{v} = M_i \hat{\mathbf{x}}_i$.

Consider the following menu. The i -th menu item costs $M_i \text{gap}_\delta^i$ for an allocation of \mathbf{x}_i . We first show that every buyer $\hat{\mathbf{x}}_i \in R^{-1}(\mathbf{x}_i, \delta)$ prefers the i -th menu item. First, notice that the utility of such a buyer for the i -th item is $M_i \hat{\mathbf{x}}_i \mathbf{x}_i - M_i \text{gap}_\delta^i \geq M_i \hat{\mathbf{x}}_i \mathbf{x}_j$, for all $j < i$ (from the definition of gap_δ^i), which is the utility from getting j -th menu item without any payment. For $j > i$, the utility of buyer $\hat{\mathbf{x}}_i$ for the j -th menu item is: $M_i \hat{\mathbf{x}}_i \mathbf{x}_j - M_j \text{gap}_\delta^j \leq 2(1 + \delta)M_i - 4M_{j-1} < 0$, since $j - 1 \geq i$ and $\delta < 1/2$. Finally, the revenue of this auction is $\sum_{i=1}^{\infty} M_i \text{gap}_\delta^i \frac{1}{M_i} = \sum_{i=1}^{\infty} \text{gap}_\delta^i$.

It remains to show that $\text{BREV}(\hat{\mathcal{D}}) \in O(1)$. For any point $\hat{\mathbf{x}}_i$, the buyer's value (and seller's revenue) for the grand bundle is at most $2M_i(1 + \delta)$ (post perturbing). Since $2M_{i-1}(1 + \delta) < M_i$, the only types willing to buy the bundle at price $M_i \hat{\mathbf{x}}_i$ are those $j \geq i$. Then the probability of selling is at most $\sum_{j \geq i} \frac{1}{M_j} \leq \frac{2}{M_i}$ (since the series of M_i is dominated by a geometric sequence with rate of growth of 2). Therefore the revenue of selling the grand bundle at any price $M_i \hat{\mathbf{x}}_i$ is at most $\frac{2}{M_i} 2M_i(1 + \delta) = 4(1 + \delta)$. \square

Proof of Claim 1. Recall that the k^{th} point on shell N has angle $\frac{\pi}{2} \cdot (1 - 3\delta)^{k-1}$. Observe that for any $j > k$, we have $\theta_j = (1 - 3\delta)^{j-k} \cdot \theta_k \leq (1 - 3\delta)\theta_k$. In particular, this implies that $\theta_k - \theta_j \geq 3\delta\theta_k$.

For the other direction, if $j < k$, we have just shown that $|\theta_k - \theta_j| \geq 3\delta\theta_j \geq 3\delta\theta_k$, since θ_k is a decreasing sequence. \square

Proof of Claim 2. The most extreme angles $\hat{\theta}_i$ a point p_i can be mapped to are $\arctan(\tan(\theta_i(1 + \delta)))$ and $\arctan(\tan(\frac{\theta_i}{1 + \delta}))$. We now bound how big the gap can be, as a function of θ_i . We do so by examining the integral of the derivative of the function $\arctan((1 + x)\tan(\theta_i))$ with respect to x .

$$\begin{aligned} \hat{\theta}_i - \theta_i &\leq \arctan((1 + \delta)\tan(\theta_i)) - \arctan(\tan(\theta_i)) \\ &\leq \int_{x=0}^{x=\delta} \frac{\tan(\theta_i)}{(\tan(\theta_i) + x\tan(\theta_i))^2 + 1} dx \\ &\leq \delta \frac{\tan(\theta_i)}{\tan(\theta_i)^2 + 1} \\ &= \delta \sin \theta_i \cos \theta_i \leq \delta \theta_i. \end{aligned} \quad \square$$

Similarly, we have (by examining the function $\arctan((1 + x)\tan(\theta_i)/(1 + \delta))$):

$$\begin{aligned} \theta_i - \hat{\theta}_i &\leq \arctan(\tan(\theta_i)) - \arctan(\tan(\theta_i)/(1 + \delta)) \\ &\leq \int_{x=0}^{x=\delta} \frac{\tan(\theta_i)/(1 + \delta)}{(\frac{\tan(\theta_i) + x\tan(\theta_i)}{1 + \delta})^2 + 1} dx \\ &\leq \delta \frac{(1 + \delta)\tan(\theta_i)}{\tan(\theta_i)^2 + (1 + \delta)^2} \\ &\leq \delta \frac{(1 + \delta)\sin \theta_i \cos \theta_i}{1 + (2\delta - \delta^2)\cos^2 \theta_i} \leq \delta \frac{\theta_i + \delta\theta_i \cos^2 \theta_i}{1 + \delta \cos^2 \theta_i} = \delta \theta_i. \end{aligned}$$

The final line is due to the fact that $\theta_i \geq \sin \theta_i \cos \theta_i$ on $[0, \pi/2]$, and also that $\theta_i \cos \theta_i \geq \sin \theta_i$ on $[0, \pi/2]$ (and that $2\delta - \delta^2 \geq \delta$).

Proof of Claim 3.

$$\hat{\mathbf{x}}_i(\mathbf{x}_i - \mathbf{x}_j) = \|\hat{\mathbf{x}}_i\| \|\mathbf{x}_i\| \cos(\theta_i - \hat{\theta}_i) - \|\hat{\mathbf{x}}_i\| \|\mathbf{x}_j\| \cos(\hat{\theta}_i - \theta_j)$$

$$\begin{aligned}
&= \|\hat{\mathbf{x}}_i\| \ell_N \cos(\theta_i - \hat{\theta}_i) - \|\hat{\mathbf{x}}_i\| \ell_N \cos(\hat{\theta}_i - \theta_j) \\
&\geq \ell_N^2 \left(\cos(\theta_i - \hat{\theta}_i) - \cos(\hat{\theta}_i - \theta_j) \right) && [\|\hat{\mathbf{x}}_i\| \geq \ell_N] \\
&= \ell_N^2 \left(2 \sin\left(\frac{|\theta_i - \theta_j|}{2}\right) \sin\left(\frac{|2\hat{\theta}_i - \theta_i - \theta_j|}{2}\right) \right) && [\cos x - \cos y = 2 \sin \frac{y+x}{2} \sin \frac{y-x}{2}, \quad \sin x = \sin |x|] \\
&\geq \frac{2}{\pi^2} \ell_N^2 (|\theta_i - \theta_j|) \left(|2\hat{\theta}_i - 2\theta_i + \theta_i - \theta_j| \right) && [\text{Jordan's inequality: } \sin x \geq \frac{2}{\pi}x] \\
&\geq \frac{2}{\pi^2} \ell_N^2 \left(|\theta_i - \theta_j| - 2|\hat{\theta}_i - \theta_i| \right)^2 \\
&\geq \frac{2}{\pi^2} \ell_n^2 (3\delta\theta_i - 2\delta\theta_i)^2. && [\text{Claims 2 and 1}]
\end{aligned}$$

□

C.2 Square-Shift Lower Bound for Two Items

Lemma 5. *There exists a bivariate distribution \mathcal{D} such that for its corresponding perturbed distribution $\hat{\mathcal{D}}$, $\text{REV}(\hat{\mathcal{D}}) \in \Omega\left(\left(\frac{1}{8}\right)^{1/7}\right) \text{BREV}(\hat{\mathcal{D}})$.*

Proof. First, we restate some known results that we will need.⁷

Let q^1, q^2, \dots be a finite or countably infinite sequence of points in $[0, 1]^k$. Define

$$\text{gap}^n = q^n \cdot q^n - \max_{j < n} q^j \cdot q^n.$$

Proposition 1 (Proposition 5.1 in [22]). *For every (finite or countably infinite) sequence q^1, q^2, \dots of points in $[0, 1]^k$ there exists a distribution F on \mathbb{R}_+^k such that $\text{BREV}(F) \leq 2k$, $\text{REV}(F) \geq \sum_n \text{gap}^n$ and $\mathbb{E}[v_1 + v_2] = \sum_n \|q^n\|_1$. Specifically, for each n , let $M^n = (2k)^n / (\prod_{j=1}^n \text{gap}^j)$. Then, F is the distribution that puts probability $1/M^n$ on the point $x^n = M^n q^n$, and puts the remaining probability on the point 0.*

Proposition 2 (Proposition 6.1 in [22]). *There exists an infinite sequence of points q^1, q^2, \dots of points in $[0, 1]^2$ with $\|q^k\|_2 \leq 1$, such that taking for all k , $\text{gap}^k \in \Omega(k^{-6/7})$. The construction puts points in “shells”. The N -th shell has $N^{3/4}$ points with length $\sum_{\ell=1}^N \ell^{-3/2} / (\sum_{\ell=1}^{\infty} \ell^{-3/2})$.*

Combining these two results we can get a distribution \mathcal{D} over n points, such that $\text{BREV}(\mathcal{D}) \leq 4$ and $\text{REV}(\mathcal{D}) \geq \sum_{k=1}^n \text{gap}^k = \sum_{k=1}^n k^{-6/7} \in \Omega(n^{1/7})$. The largest possible value for any of the two items is $M^n q^n \leq 2M^n \leq 2 \frac{4^n}{\prod_{j=1}^n \text{gap}^j} \leq 4^{n+1}$. The expected welfare of \mathcal{D} is $\sum_{k=1}^n \|q^k\|_1 \in \Theta(\sum_{k=1}^n \|q^k\|_2) = \Theta(\sum_{N=1}^{N_{max}} N^{3/4} \frac{\sum_{\ell=1}^N \ell^{-3/2}}{\sum_{\ell=1}^{\infty} \ell^{-3/2}}) = \Theta(N_{max}^{7/4})$, where N_{max} is the number of shells. The total number of points is n , and the N -th shell has $N^{3/4}$ points, therefore, $N_{max} \in \Theta(n^{4/7})$, and therefore the expected welfare is in $\Theta(n)$.

To complete the proof of the lemma, we find a δ that is large, but at the same time the gap between $\text{REV}(\hat{\mathcal{D}})$ and $\text{BREV}(\hat{\mathcal{D}})$ remains sufficiently large as well. We use the following theorem:

Theorem 1 ([32]). *Let D be a distribution, and D^+ be a distribution that stochastically dominates D . Let Δ denote the random function $v^+ - v$, when couples v^+ and v are sampled jointly from D^+ and D . Finally, let $\text{VAL}(\Delta)$ be the expected welfare of VCG with buyers whose types are distributed according to Δ . Then, for all such D, D^+, Δ , and for every mechanism M for D and every $\epsilon > 0$, there exists a mechanism M' for D^+ such that*

$$\text{REV}(M', D^+) \geq (1 - \epsilon) \cdot \left(\text{REV}(M, D) - \frac{\text{VAL}(\Delta)}{\epsilon} \right).$$

⁷These propositions appear in the arXiv version of [22] uploaded on 22 April, 2013.

We apply this theorem for $D = \mathcal{D}$, and $D^+ = \hat{\mathcal{D}}$, which stochastically dominates \mathcal{D} , and M the optimal mechanism for \mathcal{D} . Observe that $\text{VAL}(\Delta) = \mathbb{E}[\sum_{i=1}^2 (\hat{v}_i - v_i)] \leq 2\delta \mathbb{E}[\max_i v_i] \leq 2\delta \mathbb{E}[v_1 + v_2] \in \Theta(\delta n)$. We set $\epsilon = \frac{1}{2}$. Then:

$$\begin{aligned} \text{REV}(\hat{\mathcal{D}}) &\geq \text{REV}(M', \hat{\mathcal{D}}) \\ &\geq \frac{1}{2} \cdot (\text{REV}(\mathcal{D}) - \Theta(\delta n)) \\ &\geq \frac{n^{1/7}}{2} - \Theta(\delta n). \end{aligned}$$

Picking $n \in \Theta(\frac{1}{\delta})$ gives $\text{REV}(\hat{\mathcal{D}}) \in \Omega(n^{1/7})$, while $\text{BREV}(\hat{\mathcal{D}}) \leq \text{BREV}(\mathcal{D}) + \delta n \in O(1)$. \square

C.3 Square-Shift Lower Bound for Multiple Items

In this subsection we present a lower bound for the Square-Shift model with a single buyer interested in multiple items.

Theorem C.1. *For small enough δ , there exists a distribution over m items \mathcal{D} such that $\text{BREV}(\hat{\mathcal{D}}) \in O(m^2)$ but $\text{REV}(\hat{\mathcal{D}}) \in \Omega(m2^m)$.*

Proof. Enumerate all subsets of $[m]$ of size $\frac{m}{2}$ in some order: S_1, S_2, \dots, S_M , $M = \binom{m}{m/2}$, and let $C = \sum_{j=1}^M m^{-2j} = \frac{1}{m^2-1}$. For \mathcal{D} , the buyer will have type j with probability m^{-2j}/C . A buyer of type j is interested in each item in S_j equally, for a value of $2m^{2j}$ per item, and is not interested in items not in S_j .

The valuations in $\hat{\mathcal{D}}$ are not significantly different. The buyer with type j has a value of at most $\delta 2m^j$ for an item not in S_j and $(1 + \delta)2m^{2j}$ for an item in S_j . This means that for the buyer's value for the grand bundle is at most $\frac{m}{2}\delta 2m^{2j} + \frac{m}{2}(1 + \delta)2m^{2j} \leq m^{2j+1}(1 + 2\delta)$, which is smaller than $2m^{2(j+1)}$. Therefore, if the grand bundle is priced at $2m^{2j}$, only buyers with types $k \geq j$ will purchase. Therefore, the revenue of bundling is $\text{BREV}(\hat{\mathcal{D}}) \leq (1 + 2\delta)2m^{2j} \sum_{i \geq j}^k \frac{m^{-2i}}{C} \leq 4(1 + 2\delta)\frac{1}{C} \in O(m^2)$.

Consider the following mechanism. Offer the set S_j at a price m^{2j} . We show that a buyer with type j always prefers that tailored offer to any other subset. For $k > j$, the cost of S_k is at least m^{2j+2} . The value of j for S_j is at most $\frac{m}{2}2m^{2j}(1 + \delta) = m^{2j+1}(1 + \delta)$, which is less than the cost of S_k for $k > j$. For $k < j$, the utility of a type j for getting S_k for free is at most $V_k = \sum_{i=1}^{m-1} \hat{v}_i$, where \hat{v}_i is the value for the i -th valuable item after perturbing. The utility of j for S_j is at least $\sum_{i=1}^{m-1} \hat{v}_i + 2m^{2j} - m^{2j} = V_k + m^{2j}$. Therefore, type j purchases S_k for a price of m^{2j} . The total revenue is $\sum_{j=1}^k m^{2j} \frac{m^{-2j}}{C} = \frac{k}{C} \geq m2^m$. \square

D Upper Bounds

Polar Coordinates

It will turn out extremely useful for our analysis to look at valuations in polar coordinates as opposed to Cartesian coordinates. We use this space to remind the reader about some useful properties. Given a m -dimensional vector \mathbf{x} , its polar transformation is $(r, \boldsymbol{\theta})$, where r is a non-negative real number and $\boldsymbol{\theta}$ is an $m - 1$ dimensional vector of angles $\theta_i \in [0, \frac{\pi}{2}]$ such that $x_1 = r \cos \theta_1$, $x_2 = r \sin \theta_1 \cos \theta_2$, $x_3 = r \sin \theta_1 \sin \theta_2 \cos \theta_3$ and so on until $x_{m-1} = r \prod_{k=1}^{m-2} \sin \theta_k \cos \theta_{m-1}$, $x_m = r \prod_{k=1}^m \sin \theta_k$. Let $\text{trig}(\boldsymbol{\theta})$ be a vector such that $\text{trig}_j(\boldsymbol{\theta}) = \cos \theta_j \prod_{k=1}^{j-1} \sin \theta_k$. We sometimes write $(r \cdot \text{trig}(\boldsymbol{\theta}))$ or $(r, \text{trig}(\boldsymbol{\theta}))$ for the transformation from polar to cartesian coordinates.

Claim 12. $\max_{j=1,\dots,m} (\text{trig}_j(\boldsymbol{\theta})) \geq \frac{1}{\sqrt{m}}$.

Proof. $\sum_{j=1}^m \text{trig}(\boldsymbol{\theta})_j^2 = 1$ and therefore the largest entry (of the sum) must be greater than their average, $1/m$. Since the entries are non-negative we can take a square root; the Claim follows. \square

Let g_i and \hat{g}_i be the densities of \mathcal{D}_i and $\hat{\mathcal{D}}_i$ in polar coordinates. Claim 13 is a well-known consequence of multi-variable calculus.

Claim 13. Let $\mathbf{x} \in \mathbb{R}_m^+$ and let $(r, \boldsymbol{\theta})$ be its polar transform. For any probability density $f(\mathbf{x})$ defined on $\mathbf{x} \in \mathbb{R}_m^+$, the corresponding density in polar coordinates is given by $g(r, \boldsymbol{\theta}) = J_m f(\mathbf{x})$, where $J_m = r^{m-1} \prod_{i=1}^{m-2} \sin(\theta_i)^{m-1-j}$. In particular $J_2 = r$, $J_3 = r^2 \sin \theta_1$ and so on.

We use the following Lemma in all our upper bounds; it's a generalization of the fixed angle upper bound in the proof of Theorem 4.1.

Lemma 6. For any pricing p , for any angles $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_n$, where $\boldsymbol{\theta}_i = (\theta_{i,1}, \dots, \theta_{i,m-1})$,

$$\int_{\mathbf{r} \in [0, \infty)^n} \hat{g}(\mathbf{r} | \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_n) \sum_{i=1}^n p_i^*(\mathbf{r}, \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_n) d\mathbf{r} \leq \text{REV}(\hat{\mathcal{D}}_{\boldsymbol{\theta}_1}, \dots, \hat{\mathcal{D}}_{\boldsymbol{\theta}_n}),$$

where $\hat{\mathcal{D}}_{\boldsymbol{\theta}_i}$ is the single parameter distribution of buyer i (perhaps correlated with the distribution of another buyer j), where a length vector \mathbf{r} is drawn according to the density function $\hat{g}(\mathbf{r} | \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_n)$, and $\text{REV}(\hat{\mathcal{D}}_{\boldsymbol{\theta}})$ is the revenue of the optimal auction that is truthful in the same sense as p^* .

Proof. The LHS is the revenue of p conditioned on the angles being $\boldsymbol{\theta}$, i.e. the revenue of p for the distribution where we draw \mathbf{r} according to $\hat{g}(\mathbf{r} | \boldsymbol{\theta})$, and buyer i has value $r_i \text{trig}_j(\boldsymbol{\theta}_i)$ for item j . The RHS is the maximum revenue that a truthful auction can extract from the same distribution. \square

E Additive noise

We provide almost matching upper and lower bounds for the model where \mathcal{D} is supported on $[0, v_{\max}]^2$ and we perturb \mathbf{v} by adding a uniformly random vector from $[0, \delta v_{\max}]^2$.

Lemma 7. $\text{REV}(\hat{\mathcal{D}}) \in O(\ln(1/\delta)) \text{BREV}(\hat{\mathcal{D}})$

Proof. Observe that for any $\hat{\mathcal{D}}$, and any $x \geq \text{BREV}(\hat{\mathcal{D}})$, we have that $\Pr[\hat{v}_1 + \hat{v}_2 \geq x] \leq \frac{\text{BREV}(\hat{\mathcal{D}})}{x}$. Otherwise, if there exists $x^* \geq \text{BREV}(\hat{\mathcal{D}})$ such that $\Pr[\hat{v}_1 + \hat{v}_2 \geq x^*] > \frac{\text{BREV}(\hat{\mathcal{D}})}{x^*}$, setting a posted price of x^* for the grand bundle would yield strictly more revenue than $\text{BREV}(\hat{\mathcal{D}})$; a contradiction.

$$\begin{aligned} \mathbb{E}[\hat{v}_1 + \hat{v}_2] &= \int_{x=0}^2 \Pr[\hat{v}_1 + \hat{v}_2 \geq x] dx \\ &= \int_{x=0}^{\text{BREV}(\hat{\mathcal{D}})} \Pr[\hat{v}_1 + \hat{v}_2 \geq x] dx + \int_{x=\text{BREV}(\hat{\mathcal{D}})}^2 \Pr[\hat{v}_1 + \hat{v}_2 \geq x] dx \\ &\leq \text{BREV}(\hat{\mathcal{D}}) + \int_{x=\text{BREV}(\hat{\mathcal{D}})}^2 \frac{\text{BREV}(\hat{\mathcal{D}})}{x} dx \\ &= \text{BREV}(\hat{\mathcal{D}}) + \text{BREV}(\hat{\mathcal{D}}) \ln \left(\frac{2}{\text{BREV}(\hat{\mathcal{D}})} \right), \end{aligned}$$

which immediately implies that $\text{REV}(\hat{\mathcal{D}}) \leq \text{BREV}(\hat{\mathcal{D}}) + \text{BREV}(\hat{\mathcal{D}}) \ln \left(\frac{2}{\text{BREV}(\hat{\mathcal{D}})} \right)$. Furthermore, $\text{BREV}(\hat{\mathcal{D}}) \geq \delta$, by simply selling the grand bundle for a price of δ . The lemma follows. \square

Lemma 8. *There exists a bivariate distribution \mathcal{D} such that for its corresponding perturbed distribution $\hat{\mathcal{D}}$, $\text{REV}(\hat{\mathcal{D}}) \in \Omega(\ln(\frac{1}{3})^{1/7}) \text{BREV}(\hat{\mathcal{D}})$.*

Proof. First, we restate some known results that we will need.⁸

Let q^1, q^2, \dots be a finite or countably infinite sequence of points in $[0, 1]^k$. Define

$$\text{gap}^n = q^n \cdot q^n - \max_{j < n} q^j \cdot q^n.$$

Proposition 3 (Proposition 5.1 in [22]). *For every (finite or countably infinite) sequence q^1, q^2, \dots of points in $[0, 1]^k$ there exists a distribution F on \mathbb{R}_+^k such that $\text{BREV}(F) \leq 2k$, and $\text{REV}(F) \geq \sum_n \text{gap}^n$. Specifically, for each n , let $M^n = (2k)^n / (\prod_{j=1}^n \text{gap}^j)$. Then, F is the distribution that puts probability $1/M^n$ on the point $x^n = M^n q^n$, and puts the remaining probability on the point 0.*

Proposition 4 (Proposition 6.1 in [22]). *There exists an infinite sequence of points q^1, q^2, \dots of points in $[0, 1]^2$ with $\|q^k\|_2 \leq 1$, such that taking for all k , $\text{gap}^k \in \Omega(k^{-6/7})$.*

Combining these two results we can get a distribution D over n points in $[0, 1]^2$, such that $\text{BREV}(D) \leq 4$ and $\text{REV}(D) \geq \sum_{k=1}^n \text{gap}^k = \sum_{k=1}^n k^{-6/7} \in \Omega(n^{1/7})$. The largest possible value for any of the two items is $M^n q^n \leq 2M^n \leq 2 \frac{4^n}{\prod_{j=1}^n \text{gap}^j} \leq 4^{n+1}$.

Let \mathcal{D} be $D/4^{n+1}$. Then, $\text{BREV}(\mathcal{D}) \leq 4^{-n}$ and $\text{REV}(\mathcal{D}) \in \Omega(n^{1/7}/4^n)$. To complete the proof of the lemma, we find a δ that is large, but at the same time the gap between $\text{REV}(\hat{\mathcal{D}})$ and $\text{BREV}(\hat{\mathcal{D}})$ remains sufficiently large as well. We use the following theorem:

Theorem 2 ([32]). *Let D be a distribution, and D^+ be a distribution that stochastically dominates D . Let Δ denote the random function $v^+ - v$, when couples v^+ and v are sampled jointly from D^+ and D . Finally, let $\text{VAL}(\Delta)$ be the expected welfare of VCG with buyers whose types are distributed according to Δ . Then, for all such D, D^+, Δ , and for every mechanism M for D and every $\epsilon > 0$, there exists a mechanism M' for D^+ such that*

$$\text{REV}(M', D^+) \geq (1 - \epsilon) \cdot \left(\text{REV}(M, D) - \frac{\text{VAL}(\Delta)}{\epsilon} \right).$$

We apply this theorem for $D = \mathcal{D}$, and $D^+ = \hat{\mathcal{D}}$, which stochastically dominates \mathcal{D} , and M the optimal mechanism for \mathcal{D} . Our noise model adds at most δ to the value of each item, therefore $\text{VAL}(\Delta) \leq 2\delta$. We set $\epsilon = \frac{1}{2}$. Then:

$$\begin{aligned} \text{REV}(\hat{\mathcal{D}}) &\geq \text{REV}(M', \hat{\mathcal{D}}) \\ &\geq \frac{1}{2} \cdot (\text{REV}(\mathcal{D}) - 4\delta) \\ &\geq \frac{n^{1/7}}{4^{n+1}} - 2\delta. \end{aligned}$$

Finally, $\text{BREV}(\hat{\mathcal{D}}) \leq \text{BREV}(\mathcal{D}) + 2\delta \leq 4^{-n} + 2\delta$. Picking $n \in \Theta(\ln(\frac{1}{3}))$ completes the lemma. \square

F Unilluminating Alternative Models

In this section we briefly overview a few seemingly natural alternatives that unfortunately have an unilluminating analysis (similar to the additive model from Section 2). **It is certainly a**

⁸These propositions appear in the arXiv version of [22] uploaded on 22 April, 2013.

worthwhile direction for future research to identify different illuminating models, but the point of this section is just to provide further evidence that doing so requires care. This section only contains sketches and no formal proofs. Both alternatives below propose some restrictions on the class of distributions \mathcal{D} considered, and ask what guarantees can be made subject to these restrictions.

F.1 Lower-bounded Density

Consider the following restriction on distributions: for all \vec{v} in the support of \mathcal{D} , $f(\vec{v}) \in [\delta, 1/\delta]$. Let K denote the support of \mathcal{D} . There are two cases to consider:

First, maybe K is not convex. If K is not required to be convex, then we can take exactly the construction of [5, 22] and turn each discrete point into a tiny uniform ball where the density is in the range $[\delta, 1/\delta]$, all without affecting the properties at all.

Second, maybe K is required to be convex. If K is convex, then consider $L = \max_{\vec{v} \in K} \{v_1 + v_2\}$. First, observe that at least $1/4$ of the volume of K (not weighted by density in \mathcal{D}) lies above the line $v_1 + v_2 \geq L/2$. To see this, consider shrinking K by a factor of 2 in all dimensions, and moving it back to touch the line $v_1 + v_2 = L$. Then clearly this lies inside K , and has volume at least $\text{vol}(K)/4$. Now, observe that because the density lies in $[\delta, 1/\delta]$ everywhere, this means that the total *density* of \mathcal{D} above the line $v_1 + v_2 \geq L/2$ is at least $\delta^2/4$. Therefore, by simply setting price $L/2$ on the grand bundle, we can generate revenue $L\delta^2/8$. The optimum revenue is clearly at most L , so this guarantees an $O(\delta^2)$ approximation. Like with the initial silly model, we simply don't learn anything from this analysis. All we learn is that if the densities are sufficiently bounded, then you can get a decent approximation by picking a specific type and targeting them while ignoring the rest.

F.2 Bounded away from Axes

Consider the following restriction on distributions: with probability $1 - \delta$, $\vec{v} \leftarrow \mathcal{D}$ satisfies $v_1 \geq \delta v_2 \geq \delta^2 v_1$. In other words, \vec{v} is sufficiently bounded away from the axes with non-negligible probability. Unfortunately this buys us absolutely nothing. Consider the construction of [5, 22], and the distribution \mathcal{D}' which draws from their construction with probability δ , and otherwise draw from $U([1, 2] \times [1, 2])$. Then the optimal revenue of this distribution is still infinite, and the revenue of any mechanism with menu complexity C is at most $O(C)$, even though with high probability it is bounded away from the axes.

The above construction in fact works for any property. Consider restricting to distributions such that with probability $1 - \delta$, $\vec{v} \leftarrow \mathcal{D}$ has property X. Assuming that some \vec{v} exists with property X, let \mathcal{D}' draw from the [5, 22] construction with probability δ , and draw \vec{v} otherwise. Then the optimal revenue of this distribution is still infinite, yet the optimal revenue for any mechanism with menu complexity C is at most $O(C(v_1 + v_2))$.

References

- [1] Rashmi Adaval and Robert S Wyer. Conscious and nonconscious comparisons with price anchors: Effects on willingness to pay for related and unrelated products. *Journal of Marketing Research*, 48(2):355–365, 2011. doi: 10.1509/jmkr.48.2.355. URL <https://doi.org/10.1509/jmkr.48.2.355>.

- [2] 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.
- [3] MohammadHossein Bateni, Sina Dehghani, MohammadTaghi Hajiaghayi, and Saeed Seddighin. Revenue maximization for selling multiple correlated items. In *Algorithms-ESA 2015*, pages 95–105. Springer, 2015.
- [4] Luca Becchetti, Stefano Leonardi, Alberto Marchetti-Spaccamela, Guido Schäfer, and Tjark Vredeveld. Average-case and smoothed competitive analysis of the multilevel feedback algorithm. *Mathematics of Operations Research*, 31(1):85–108, 2006.
- [5] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S Matthew Weinberg. Pricing lotteries. *Journal of Economic Theory*, 156:144–174, 2015.
- [6] 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.
- [7] 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.
- [8] Gabriel Carroll. Robustness and separation in multidimensional screening. *Econometrica*, 85(2):453–488, 2017.
- [9] 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.
- [10] 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. ISBN 978-1-59593-653-0. doi: 10.1145/1250910.1250946. URL <http://doi.acm.org/10.1145/1250910.1250946>.
- [11] 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.
- [12] Shuchi Chawla, David Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015.
- [13] Ning Chen, Nick Gravin, and Pinyan Lu. Optimal competitive auctions. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 253–262. ACM, 2014.
- [14] Jacques Cremer and Richard P McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society*, pages 1247–1257, 1988.
- [15] Constantinos Daskalakis and Seth Matthew Weinberg. Symmetries and optimal multi-dimensional mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 370–387. ACM, 2012.

- [16] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. The complexity of optimal mechanism design. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1302–1318. SIAM, 2014.
- [17] Xiaotie Deng, Yansong Gao, and Jie Zhang. Smoothed and average-case approximation ratios of mechanisms: Beyond the worst-case analysis. *arXiv preprint arXiv:1705.07200*, 2017.
- [18] Hu Fu, Chris Liaw, Pinyan Lu, and Zhihao Gavin Tang. The value of information concealment. *arXiv preprint arXiv:1707.05875*, 2017.
- [19] Andrew V Goldberg and Jason D Hartline. Envy-free auctions for digital goods. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 29–35. ACM, 2003.
- [20] Andrew V Goldberg, Jason D Hartline, Anna R Karlin, Michael Saks, and Andrew Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.
- [21] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 656–656. ACM, 2012.
- [22] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. *arXiv preprint arXiv:1304.6116*, 2013.
- [23] Sergiu Hart and Philip J Reny. Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Theoretical Economics*, 10(3):893–922, 2015.
- [24] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, EC '09, pages 225–234, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-458-4. doi: 10.1145/1566374.1566407. URL <http://doi.acm.org/10.1145/1566374.1566407>.
- [25] 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.
- [26] Aradhna Krishna, Mary Wagner, Carolyn Yoon, and Rashmi Adaval. Effects of extreme-priced products on consumer reservation prices. *Journal of Consumer Psychology*, 16(2):176–190, 2006.
- [27] 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.
- [28] Nina Mazar, Botond Koszegi, and Dan Ariely. True context-dependent preferences? the causes of market-dependent valuations. *Journal of Behavioral Decision Making*, 27(3):200–208, 2014.
- [29] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [30] Joseph C. Nunes and Peter Boatwright. Incidental prices and their effect on willingness to pay. *Journal of Marketing Research*, 41(4):457–466, 2004. doi: 10.1509/jmkr.41.4.457.47014. URL <https://doi.org/10.1509/jmkr.41.4.457.47014>.
- [31] Amir Ronen. On approximating optimal auctions. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 11–17. ACM, 2001.

- [32] Aviad Rubinfeld 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.
- [33] 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.
- [34] Daniel A Spielman and Shang-Hua Teng. Smoothed analysis: an attempt to explain the behavior of algorithms in practice. *Communications of the ACM*, 52(10):76–84, 2009.
- [35] 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.