

A Simple Mechanism for a Budget-Constrained Buyer

YU CHENG, Duke University, USA

NICK GRAVIN, Shanghai University of Finance and Economics, China

KAMESH MUNAGALA, Duke University, USA

KANGNING WANG, Duke University, USA

We study a classic Bayesian mechanism design setting of monopoly problem for an additive buyer in the presence of budgets. In this setting a monopolist seller with m heterogeneous items faces a single buyer and seeks to maximize her revenue. The buyer has a budget and additive valuations drawn independently for each item from (non-identical) distributions. We show that when the buyer's budget is publicly known, the better of selling each item separately and selling the grand bundle extracts a constant fraction of the optimal revenue. When the budget is private, we consider a standard Bayesian setting where buyer's budget b is drawn from a known distribution B . We show that if b is independent of the valuations (which is necessary) and distribution B satisfies monotone hazard rate condition, then selling items separately or in a grand bundle is still approximately optimal.

CCS Concepts: • **Theory of computation** → **Algorithmic mechanism design; Computational pricing and auctions**; *Algorithmic game theory*.

Additional Key Words and Phrases: auctions, revenue maximization, budget, simple mechanisms

ACM Reference Format:

Yu Cheng, Nick Gravin, Kamesh Munagala, and Kangning Wang. 2020. A Simple Mechanism for a Budget-Constrained Buyer. *ACM Transactions on Economics and Computation* 1, 1 (August 2020), 24 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Revenue maximization is one of the fundamental problems in auction theory. The well-celebrated result of Myerson [Myerson 1981] characterizes the revenue-maximizing mechanism when there is only one item for sale. Specifically, in the single buyer case, the optimal solution is to post a take-it-or-leave-it price. Since Myerson's work, the optimal mechanism design problem has been studied extensively in computer science literature and much progress has been made [Alaei et al. 2012; Cai et al. 2012a,b, 2013a,b; Daskalakis 2015]. The problem of finding the optimal auction turned out to be so much more complex than the single-item case. Unlike the Myerson's single-item auction, the optimum can use randomized allocations and price bundles of items already for two items and a single buyer. It is also known that the gap between the revenue of the optimal randomized and optimal deterministic mechanism can be arbitrarily large [Briest et al. 2010; Hart and Nisan 2013], the optimal mechanism may require a menu with infinitely many options [Daskalakis et al. 2017;

Authors' addresses: Yu Cheng, Duke University, Durham, USA, yucheng@cs.duke.edu; Nick Gravin, Shanghai University of Finance and Economics, Shanghai, China, nikolai@mail.shufe.edu.cn; Kamesh Munagala, Duke University, Durham, USA, kamesh@cs.duke.edu; Kangning Wang, Duke University, Durham, USA, knwang@cs.duke.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

XXXX-XXXX/2020/8-ART \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Manelli and Vincent 2007], and the revenue of the optimal auction may decrease when the buyer's valuation distributions move upwards (in the stochastic dominance sense) [Hart and Reny 2015].

In light of these negative results for optimal auction design, many recent papers focused on the design of *simple* mechanisms that are *approximately* optimal. One such notable line of work initiated by Hart and Nisan [Hart and Nisan 2017] concerns a basic and natural setting of monopoly problem for the buyer with item values drawn independently from given distributions D_1, \dots, D_m and whose valuation for the sets of items is additive¹ (linear). A remarkable result by Babaioff et al. [Babaioff et al. 2014] showed that the better mechanism of either selling items separately, or selling the grand bundle extracts at least $(1/6)$ -fraction of the optimal revenue. It was also observed [Babaioff et al. 2014; Hart and Nisan 2013; Rubinstein and Weinberg 2015] that the independence assumption on the items is essentially necessary and without it no simple (any deterministic) mechanism cannot be approximately optimal.

Auction design with budget constraints is an even harder problem. Because the buyer's utility is no longer quasi-linear, many standard concepts do not carry over². For example, even for one buyer and one item, the optimal mechanism may require randomization when the budget is public [Chawla et al. 2011], and may need an exponential-size menu (exponential in the number of possible budgets) when the budget is private [Devanur and Weinberg 2017]. Despite many efforts [Abrams 2006; Bei et al. 2012; Bhattacharya et al. 2010a,b; Borgs et al. 2005; Chakrabarty and Goel 2010; Chawla et al. 2011; Che and Gale 2000; Chen et al. 2011; Daskalakis et al. 2015; Devanur et al. 2013; Devanur and Weinberg 2017; Dobzinski et al. 2012; Goel et al. 2015; Goldberg et al. 2001; Laffont and Robert 1996; Singer 2010], the theory of optimal auction design with budgets is still far behind the theory without budgets.

In this paper, we investigate the effectiveness of simple mechanisms in the presence of budgets. Our work is motivated by the following questions:

How powerful are simple mechanisms in the presence of budgets? In particular, is there a simple mechanism that is approximately optimal for a budget-constrained buyer with independent valuations?

To this end we consider one of the most basic and natural settings of extensively studied monopoly problem for an additive buyer. In this setting, a monopolistic seller sells m items to a single buyer. The buyer has additive valuations drawn independently for each item from an arbitrary (non-identical) distribution. We study two different budget settings: the *public budget* case where the buyer has a fixed budget known to the seller, and the *private budget* case where the buyer's budget is drawn from a distribution. The seller wishes to maximize her revenue by designing an auction subject to individual rationality, incentive compatibility, and budget constraints. We consider the Bayesian setting where the buyer knows his budget and his values for each item, but the seller only knows the prior distributions.

1.1 Our Results and Techniques

Our first result is that simple mechanisms remain approximately optimal when the buyer has a public budget.

THEOREM 1.1. *For an additive buyer with a known public budget and independent valuations, the better of selling each item separately and selling the grand bundle extracts a constant fraction of the optimal revenue.*

¹A buyer has additive valuations if his value for a set of items is equal to the sum of his values for the items in the set.

²E.g., the classic VCG mechanism may not be implementable and social efficiency may not be achievable in the budgeted-setting [Singer 2010].

Theorem 1.1 is among the few positive results in budget-constrained settings that hold for arbitrary distributions. Before our work, it is not clear that any mechanism extracting a constant fraction of the optimal revenue can be computed in polynomial time.

In Sections 3 and 4, we present two different approaches to prove Theorem 1.1. Both approaches truncate the valuation distribution V according to the budget b (in different ways) and then relate the revenues of the optimal/simple mechanisms on the truncated distribution to the revenues on the original valuations. The first approach uses the main result of [Babaioff et al. 2014] in a black-box way, and the second approach adapts the duality-based framework developed in [Cai et al. 2016].

It is worth pointing out that many of our structural lemmas hold for correlated valuations as well. Using these lemmas, we can generalize Theorem 1.1 to allow the buyer to have weakly correlated valuations. We call a distribution \widehat{V} *weakly correlated* if it is the result of conditioning an independent distribution V on the sum of $v \sim V$ being at most c : $\widehat{V} = V_{|\sum v_i \leq c}$ (See Definition 2.1 for the formal definition).

COROLLARY 1.2. *Let \widehat{V} be a weakly correlated distribution. For an additive buyer with a public budget and valuations drawn from \widehat{V} , the better of selling separately and selling the grand bundle extracts a constant fraction of the optimal revenue.*

In Section 5, we examine the private budget setting. The budget b is no longer fixed but is drawn from a distribution B . The seller only knows the prior distribution B but not the value of b . We first show that if the valuations can be correlated with the budget, the problem is at least as hard as budget-free mechanism design with correlated valuations, where simple mechanisms are known to be ineffective. In light of this negative result, we focus on the setting where the budget distribution B is independent of the valuations V . In this setting, we show that simple mechanisms are approximately optimal when the budget distribution satisfies the monotone hazard rate (MHR) condition.

THEOREM 1.3. *When the budget distribution B is MHR, the better mechanism of pricing items separately and selling a grand bundle achieves a constant fraction of the optimal revenue.*

We will show that it is sufficient to pretend the buyer has a public budget $b^* = \mathbb{E}_{b \sim B}[b]$. The proof of Theorem 1.3 uses the MHR condition, as well as the fact that for a public budget b , the (budget-constrained) optimal revenue is nondecreasing in b , but optimal revenue divided by b is nonincreasing in b .

1.2 Related Work

The most closely related to ours are the following two lines of work.

Simple Mechanisms. In a line of work initiated by Hart and Nisan [Babaioff et al. 2014; Cai and Zhao 2017; Chawla and Miller 2016; Hart and Nisan 2017; Li and Yao 2013], [Babaioff et al. 2014] first showed that for an additive buyer with independent valuations, either selling separately or selling the grand bundle extracts a constant fraction of the optimal revenue. This was later extended to multiple buyers [Yao 2015], as well as buyers with more general valuations (e.g., sub-additive [Rubinstein and Weinberg 2015], valuations with a common-value component [Batani et al. 2015], and valuations with complements [Eden et al. 2017]). Others have studied the trade-off between the complexity and approximation ratio of an auction, along with the design of small-menu mechanisms in various settings [Babaioff et al. 2017; Cheng et al. 2015; Dughmi et al. 2014; Hart and Nisan 2013; Tang and Wang 2017].

Auctions for Budget-Constrained Buyers. There has been a lot of work studying the impact of budget constraints on mechanism design. Most of the earlier work required additional assumptions on the valuations distributions, like regularity or monotone hazard rate ([Bhattacharya et al. 2010b; Che and Gale 2000; Laffont and Robert 1996; Pai and Vohra 2014]). We mention a few results that work for arbitrary distributions. For public budgets, [Chawla et al. 2011] designed approximately optimal mechanisms for several single-parameter settings and multi-parameter settings with unit-demand buyers (our paper deals with additive buyers). For private budgets, [Devanur and Weinberg 2017] characterized the structure of the optimal mechanism for one item and one buyer, and showed its menu complexity to be exponential in the number of possible budgets. [Daskalakis et al. 2015] gave a constant-factor approximation for additive bidders whose private budgets can be correlated with their values. However, they require the buyers' valuation distribution to be given explicitly, which is of exponential size in our setting. There are also approximation and hardness results in the prior-free setting [Abrams 2006; Borgs et al. 2005; Devanur et al. 2013], as well as designing Pareto optimal auctions [Dobzinski et al. 2012; Goel et al. 2015].

Other Related Work. Our work concerns revenue maximization for additive buyer. Another natural and basic scenario extensively studied in the literature concerns buyers with unit-demand preferences [Chawla et al. 2007, 2010, 2015]. Our work studies monopoly problem for additive budgeted buyer in the standard Bayesian approach. In this framework, the prior distribution is known to the seller and typically is assumed to be independent. Parallel to this framework, the (budgeted) additive monopoly problem has been studied in a new robust optimization framework [Carroll 2017; Gravin and Lu 2018]. Another group of papers on budget feasible mechanism design [Bei et al. 2012; Chen et al. 2011; Singer 2010; Singla and Krause 2013] studies different reverse auction settings and are concerned with value maximization.

2 PRELIMINARIES

2.1 Optimal Mechanism Design

We study the design of optimal auctions with one buyer, one seller, and m heterogeneous items labeled by $[m] = \{1, \dots, m\}$. There is exactly one copy of each item, and the items are indivisible. The buyer has additive valuation ($v(S) = \sum_{j \in S} v(\{j\})$ for any set $S \subseteq [m]$) and a publicly known budget b ³.

We use $v \in \mathbb{R}^m$ to denote the buyer's valuations, where v_j is the buyer's value for item j . We consider the Bayesian setting of the problem, in which the buyer's values are drawn from a discrete⁴ distribution V . Let $T = \text{supp}(V)$ be the set of all possible valuation profiles in V . We use $f(t)$ for any $t \in T$ to denote the probability mass function of V : $f(t) = \Pr_{v \sim V}[v = t]$. Let $T_j = \text{supp}(V_j)$. We say the valuation distribution V is independent across items if it can be expressed as $V = \times_j V_j$.

We assume the buyer is risk-neutral and has quasi-linear utility when the payment does not exceed his budget. Let $\pi : T \rightarrow [0, 1]^m$ and $p : T \rightarrow \mathbb{R}$ denote the allocation and payment rules of a mechanism respectively. That is, when the buyer reports type t , the probability that he will receive item j is $\pi_j(t)$, and his expected payment is $p(t)$ (over the randomness of the mechanism). Thus, if the buyer has type t , his (expected) value for reporting type t' is exactly $\pi(t')^\top t$,⁵ and his

³In this paper, we mostly focus on the public budget case. So we define notations and discuss backgrounds assuming the buyer has a public budget.

⁴Like previous work on simple and approximately optimal mechanisms, our results extend to continuous types as well, since we can discretize the types so that the revenues are arbitrarily close (see, e.g., [Cai et al. 2016] for a more detailed discussion).

⁵We use $x^\top y = \sum_{i=1}^m x_i y_i$ to denote the inner product of two vectors x and y .

(expected) utility for reporting type t' is

$$u(t, t') = \begin{cases} \pi(t')^\top t - p(t') & \text{if } p(t') \leq b, \text{ and} \\ -\infty & \text{otherwise.} \end{cases}$$

By the revelation principle, it is sufficient to consider mechanisms that are incentive compatible (i.e., “truthful”). A mechanism $M = (\pi, p)$ is (interim) incentive-compatible (IC) if the buyer is incentivized to tell the truth (over the randomness of mechanism), and (interim) individually rational (IR) if the buyer’s expected utility is non-negative whenever he reports truthfully. We use \emptyset for the option of not participating in the auction ($\pi(\emptyset) = 0, p(\emptyset) = 0$), and let $T^+ = T \cup \{\emptyset\}$. Then, the IC and IR constraints can be unified as follows:

$$u(t, t) \geq u(t, t') \quad \forall t \in T, t' \in T^+.$$

To summarize, when the seller faces a single buyer with budget b and valuation drawn from V , the optimal mechanism $M^* = (\pi^*, p^*)$ is the optimal solution to the following (exponential-size) linear program (LP):

$$\begin{aligned} & \text{maximize} && \sum_{t \in T} f(t)p(t) \\ & \text{subject to} && \pi(t')^\top t - p(t') \leq \pi(t)^\top t - p(t), \quad \forall t \in T, t' \in T^+. \\ & && 0 \leq \pi_j(t) \leq 1, \quad \forall t \in T, j \in [m]. \\ & && p(t) \leq b, \quad \forall t \in T. \\ & && \pi(\emptyset) = 0, p(\emptyset) = 0. \end{aligned} \tag{1}$$

A mechanism is called ex-post IC, ex-post IR, or ex-post budget-preserving respectively, if the corresponding constraints hold for all possible outcomes, without averaging over the randomness in the mechanism. We will show the better of pricing each item separately and pricing the grand bundle, which is a deterministic ex-post mechanism, can extract a constant fraction of the revenue of any interim mechanism.

2.2 Simple Mechanisms

For a buyer with valuation distribution V , we frequently use the following notations in our analysis:

- $\text{REV}(V)$: the revenue of the optimal truthful mechanism.
- $\text{SREV}(V)$: the maximum revenue obtainable by pricing each item separately.
- $\text{BREV}(V)$: the maximum revenue obtainable by pricing the grand bundle.
- $\text{REV}^b(V)$: the revenue of the optimal truthful mechanism, when the buyer has a budget b .
- $\text{SREV}^b(V)$: the maximum revenue that can be extracted by pricing each item separately, when the buyer has a public budget b .
- $\text{BREV}^b(V)$: the maximum revenue that can be extracted by pricing the grand bundle, when the buyer has a public budget b .

We know that $\text{SREV}(V)$ is obtained by running Myerson’s optimal auction separately for each item, and $\text{BREV}(V)$ is obtained by running Myerson’s auction viewing the grand bundle as one item. Similarly, $\text{BREV}^b(V)$ is a single-parameter problem as well, with the minor change that the posted price is at most b .

The case of $\text{SREV}^b(V)$ is more complicated. For example, when a budgeted buyer of type $t \in \mathbb{R}^m$ participates in an auction with posted price p_j for each item j , he will maximize his utility by solving a KNAPSACK problem. There exists a poly-time computable mechanism that extracts a constant fraction of $\text{SREV}^b(V)$ (e.g., [Bhattacharya et al. 2010a]). However, in this paper, we only need the mild assumption that the buyer will purchase a maximal subset of items – if the buyer values an item more than its price and he can still afford it, then he will buy it.

2.3 Weakly Correlated Distributions

We call a distribution like \widehat{V} *weakly correlated* if the only condition causing the correlation is a cap on its sum.

Definition 2.1. For an m -dimensional independent distribution V and a threshold $c > 0$, we remove the probability mass on any $t \in \text{supp}(V)$ with $\|t\|_1 > c$ and renormalize. Let $\widehat{V} := V_{(\|v\|_1 \leq c)}$ denote the resulting distribution. Formally,

$$\Pr_{\widehat{v} \sim \widehat{V}}[\widehat{v} = t] = \Pr_{v \sim V}[v = t \mid \|v\|_1 \leq c], \quad \forall t \in \text{supp}(V).$$

Weakly correlated distributions arise naturally in our analysis. We will show that if the buyer's valuations are weakly correlated, then the better of selling separately and selling the grand bundle is approximately optimal, and this holds with or without a (public) budget constraint.

2.4 First-Order Stochastic Dominance

Stochastic dominance is a partial order between random variables. A random variable X with $\text{supp}(X) \subseteq \mathbb{R}$ (*weakly*) *first-order stochastically dominates* another random variable Y with $\text{supp}(Y) \subseteq \mathbb{R}$ if and only if

$$\Pr[X \geq a] \geq \Pr[Y \geq a] \text{ for all } a \in \mathbb{R}.$$

This notion of stochastic dominance can be extended to multi-dimensional distributions. In this paper, we use the notion of coordinate-wise dominance.

Definition 2.2. Given two m -dimensional distributions D_1 and D_2 , we say D_1 *coordinate-wise stochastically dominates* D_2 ($D_1 \geq D_2$ or $D_2 \leq D_1$) if there exists a randomized mapping $f : \text{supp}(D_1) \rightarrow \text{supp}(D_2)$ such that $f(x) \sim D_2$ when $x \sim D_1$, and $f(x) \leq x$ coordinate-wise for all $x \in \text{supp}(D_1)$ with probability 1.

This notion helps us express the monotonicity of optimal revenues in some cases. For example, we can show that $\text{SREV}(V_1) \geq \text{SREV}(V_2)$ when $V_1 \geq V_2$. The mapping f allows us to couple the draws $v_1 \sim V_1$ and $v_2 \sim V_2$, so that for a set of fixed prices, if the buyer buys an item under v_2 , he will also buy it under v_1 .

3 PUBLIC BUDGET

In this section, we focus on the public budget case and prove our main result (Theorem 1.1). The buyer has a fixed budget b and valuations drawn from an independent distribution V .

Theorem 1.1. $\text{REV}^b(V) \leq 8\text{SREV}^b(V) + 23\text{BREV}^b(V)$.

It follows that the better of $\text{SREV}^b(V)$ and $\text{BREV}^b(V)$ is at least $\frac{\text{REV}^b(V)}{31}$.⁶

Overview of Our Approach. Instead of taking the Lagrangian dual of LP (1) to derive an upper bound on the optimal objective value $\text{REV}^b(V)$, we adopt a more combinatorial approach. Intuitively, we come up with a charging argument that splits $\text{REV}^b(V)$ and charges each part to either $\text{SREV}^b(V)$ or $\text{BREV}^b(V)$. Here is the proof outline:

First, we partition the buyer types $t \in \text{supp}(V)$ into two sets: *high-value* types where $\|t\|_\infty \geq b$ and *low-value* types where $\|t\|_\infty < b$. Note that we can already charge the revenue of high-value types to $\text{BREV}^b(V)$: If we sell the grand bundle at price b , all high-value types will exhaust their budgets.

⁶We do not optimize the constants in our proofs. In Section 4, we will give an alternative proof of Theorem 1.1 that shows $\text{REV}^b(V) \leq 5\text{SREV}^b(V) + 6\text{BREV}^b(V)$, thus improving this constant from 31 to 11.

We now examine the low-value types. Let V' denote the valuation distribution conditioned on the buyer having a low-value type. Observe that V' is independent because it is defined using ℓ_∞ -norm, and we can remove the budget to upper bound its revenue. For a budget-free additive buyer with independent valuations, we can apply the main result of [Babaioff et al. 2014], which states that either selling separately or grand bundling works for V' : $\text{REV}(V') = O(\text{SREV}(V') + \text{BREV}(V'))$.

Next, we will relate $\text{SREV}(V')$, $\text{BREV}(V')$ to $\text{SREV}^b(V')$, $\text{BREV}^b(V')$. We can assume the sum of $v' \sim V'$ is usually much smaller than b . Similar to standard tail bounds, if the sum $\|v'\|_1$ is often small and the random variables are independent and bounded (each v'_j is at most b), then $\|v'\|_1$ must have an exponentially decaying tail. Therefore, we can add back the budget, because the sum $\|v'\|_1$, which upper bounds the buyer's payment, is rarely very large.

Finally, we will show that $\text{SREV}^b(V') = O(\text{SREV}^b(V))$ and $\text{BREV}^b(V') \leq \text{BREV}^b(V)$. The BREV statement is easy to verify, but the SREV statement is more tricky. The monotonicity of $\text{SREV}(V)$ in the budget-free case (see Section 2.4) no longer holds when there is a budget. Fortunately, we can pay a factor of two and circumvent this non-monotonicity due to budget constraints.

We will now make our intuitions formal and present three key lemmas. Throughout the paper, we will always use $V' = V_{\|v\|_\infty \leq b}$ as defined below.

Definition 3.1. Fix an m -dimensional distribution $V = \times V_j$. Let V' be the independent distribution where every coordinate of V is capped at b . That is, $V' = \times_j V'_j$, and V'_j is given by $\Pr_{V'_j}[x] = \Pr_{v_j \sim V_j}[\min(v_j, b) = x]$.

LEMMA 3.2. $\text{REV}^b(V) \leq \text{REV}(V') + \text{BREV}^b(V)$.

LEMMA 3.3. Assume $\text{BREV}^b(V') < \frac{b}{10}$. Then, $\text{BREV}(V') \leq 3\text{BREV}^b(V')$ and $\text{SREV}(V') \leq \text{SREV}^b(V') + 4\text{BREV}^b(V')$.

LEMMA 3.4. $\text{BREV}^b(V') \leq \text{BREV}^b(V)$ and $\text{SREV}^b(V') \leq 2\text{SREV}^b(V)$.

We defer the proofs of these lemmas to Sections 3.1, 3.2, and 3.3, and first use them to prove Theorem 1.1.

PROOF OF THEOREM 1.1. If $\text{BREV}^b(V') \geq \frac{b}{10}$, then the theorem holds because the optimal revenue $\text{REV}^b(V)$ is at most the budget b . By Lemma 3.4, $\text{BREV}^b(V) \geq \text{BREV}^b(V') \geq \frac{b}{10} \geq \frac{\text{REV}^b(V)}{10}$.

We now assume $\text{BREV}^b(V') < \frac{b}{10}$. The theorem follows straightforwardly from Lemmas 3.2, 3.3, 3.4, and a black-box use of the main result of [Babaioff et al. 2014].

$$\begin{aligned}
\text{REV}^b(V) &\leq \text{REV}(V') + \text{BREV}^b(V) && \text{(Lemma 3.2)} \\
&\leq 4\text{SREV}(V') + 2\text{BREV}(V') + \text{BREV}^b(V) && \text{([Babaioff et al. 2014])} \\
&\leq 4\text{SREV}^b(V') + 22\text{BREV}^b(V') + \text{BREV}^b(V) && \text{(Lemma 3.3)} \\
&\leq 8\text{SREV}^b(V) + 22\text{BREV}^b(V) + \text{BREV}^b(V) && \text{(Lemma 3.4)} \\
&= 8\text{SREV}^b(V) + 23\text{BREV}^b(V). && \square
\end{aligned}$$

3.1 Proof of Lemma 3.2

We will prove the following lemma, which is a generalization of Lemma 3.2.

LEMMA 3.5. Fix $b > 0$ and $0 < c \leq b$. For any distribution \widehat{V} with $\text{supp}(\widehat{V}) \subseteq \text{supp}(V)$ and $\Pr_{\widehat{V}}[t] \geq \Pr_V[t]$ for any $\|t\|_1 \leq c$, we have $\text{REV}^b(V) \leq (b/c) \cdot \text{BREV}^b(V) + \text{REV}(\widehat{V})$.

Lemma 3.2 follows immediately from Lemma 3.5 by choosing $c = b$ and $\widehat{V} = V'$, because capping each coordinate at c does not create new support, and does not decrease probability mass on any type t whose sum is at most c .

Intuitively, Lemma 3.5 upper bounds the optimal revenue by splitting the buyer types t into two sets: when $\|t\|_1 > c$, we upper bound the seller's revenue by the budget b ; when $\|t\|_1 \leq c$, we run the optimal mechanism for $\text{REV}^b(V)$.

PROOF OF LEMMA 3.5. Let T and \widehat{T} , f and \widehat{f} denote the support and probability density function of V and \widehat{V} respectively. Let $M^* = (\pi^*, p^*)$ be the optimal mechanism that obtains $\text{REV}^b(V)$. Recall that π^* and p^* are the allocation and payment rules, and (π^*, p^*) is the optimal solution to LP (1) for f and T .

We split the optimal revenue into two parts:

$$\text{REV}^b(V) = \sum_{t \in T} f(t)p^*(t) = \sum_{\|t\|_1 > c} f(t)p^*(t) + \sum_{\|t\|_1 \leq c} f(t)p^*(t).$$

Since $p^*(t) \leq b$, the first term is at most $b \sum_{\|t\|_1 > c} f(t) = b \cdot \Pr[\|v\|_1 > c] \leq (b/c)\text{BREV}^b(V)$, because we can sell the grand bundle at price $p = c$.

The second term is at most $\text{REV}(\widehat{V})$, because M^* is a feasible solution to the LP for $\widehat{T} \subseteq T$. In other words, M^* satisfies the IC and IR constraints for \widehat{V} . The revenue of \widehat{V} is at least the revenue of M^* on \widehat{V} :

$$\text{REV}(\widehat{V}) \geq \sum_{t \in \widehat{T}} \widehat{f}(t)p^*(t) \geq \sum_{t \in T, \|t\|_1 \leq c} \widehat{f}(t)p^*(t) \geq \sum_{t \in T, \|t\|_1 \leq c} f(t)p^*(t).$$

Combining the upper bounds, we get $\text{REV}^b(V) \leq (b/c)\text{BREV}^b(V) + \text{REV}(\widehat{V})$. \square

3.2 Proof of Lemma 3.3

Lemma 3.3 states that when the sum of $v' \sim V'$ is often small, the budget does not matter too much for V' . Intuitively, because each coordinate of $v' \sim V'$ is independent and upper bounded by b , a concentration inequality implies that the sum has an exponentially decaying tail. Therefore, the budget constraint is less critical because it is very unlikely that the buyer's value for the grand bundle is much larger than the budget.

We formalize this intuition by proving the following lemma, which is similar to standard tail bounds. The main difference is that, instead of knowing the mean of $\|v'\|_1$ is small, we only know that $\text{BREV}^b(V')$ is small.

LEMMA 3.6. *If V' is independent and $\|v'\|_\infty \leq c$ for all $v' \sim V'$, then*

$$\Pr_{v' \sim V'} [\|v'\|_1 \geq x + y + c] \leq \Pr_{v' \sim V'} [\|v'\|_1 \geq x] \cdot \Pr_{v' \sim V'} [\|v'\|_1 \geq y] \quad \text{for all } x, y > 0.$$

In particular, if $\Pr_{v' \sim V'} [\|v'\|_1 \geq c] \leq q$, then for all integer $k \geq 0$,

$$\Pr_{v' \sim V'} [\|v'\|_1 \geq (2k + 1)c] \leq q^k \Pr_{v' \sim V'} [\|v'\|_1 \geq c].$$

We defer the proof of Lemma 3.6 to Appendix A, and first use this tail bound to prove Lemma 3.3.

PROOF OF LEMMA 3.3. Let $c = b$ and $q = \frac{1}{10}$. We know that $\Pr[\|v'\|_1 \geq c] \leq q$ from the assumption $\text{BREV}^b(V') \leq \frac{b}{10}$.

First observe that $\text{BREV}^b(V') = \max_{p \leq b} (p \cdot \Pr[\|v'\|_1 \geq p])$. If we price the grand bundle at price p where $(2k + 1)c < p \leq (2k + 3)c$ for some $k \geq 0$, by Lemma 3.6, the revenue is at most

$$p \cdot \Pr[\|v'\|_1 \geq (2k + 1)c] \leq (2k + 3)c \cdot q^k \Pr[\|v'\|_1 \geq c] \leq (2k + 3)q^k \text{BREV}^b(V').$$

As $(2k+3)q^k \leq 3$ for any integer k , we have $\text{BREV}(V') \leq 3\text{BREV}^b(V')$.

For $\text{SREV}(V')$, similar to Lemma 3.2, we can upper bound the revenue by allowing the seller to extract full revenue if $\|v'\|_1 > c$, and running the optimal budget-constrained mechanism when $\|v'\|_1 \leq c$:

$$\begin{aligned} \text{SREV}(V') &\leq \text{SREV}^c(V') + \mathbb{E}[\|v'\|_1 \mid \|v'\|_1 \geq c] \\ &\leq \text{SREV}^c(V') + \sum_{k=1}^{\infty} (2k+3)c \cdot \Pr[\|v'\|_1 \geq (2k+1)c] \\ &\leq \text{SREV}^c(V') + \sum_{k=0}^{\infty} (2k+3)q^k \text{BREV}^c(V') \\ &\leq \text{SREV}^c(V') + 4\text{BREV}^c(V'). \quad \square \end{aligned}$$

3.3 Proof of Lemma 3.4

Lemma 3.4 states that $\text{SREV}^b(V)$ and $\text{BREV}^b(V)$ are both (up to constant factors) monotone in V . We prove a more general version of the lemma that does not require V to be independent. Recall that $\widehat{V} \leq V$ means \widehat{V} is coordinate-wise stochastically dominated by V .

LEMMA 3.7. Fix $b > 0$ and $0 < c \leq b$. For any distribution $\widehat{V} \leq V$, $\text{BREV}^c(\widehat{V}) \leq \text{BREV}^b(V)$ and $\text{SREV}^c(\widehat{V}) \leq \max(1, \frac{2c}{b}) \text{SREV}^b(V)$.

Lemma 3.4 follows directly from Lemma 3.7, by choosing $c = b$ and $\widehat{V} = V'$.

Intuitively, we would like to prove that $\text{SREV}^b(\widehat{V}) \leq \text{SREV}^b(V)$ for any $\widehat{V} \leq V$. While this is true in the budget-free case (See Section 2.4), it is actually false in the presence of a budget. We give a counterexample in Appendix B. Fortunately, we can prove $\text{SREV}^b(V') \leq 2\text{SREV}^b(V)$. The intuition is that we can cap the price of each item at $b/2$, then the buyer either spends at least $b/2$, or will purchase everything he likes.

PROOF OF LEMMA 3.7. First consider BREV . Because $c \leq b$ and $\widehat{V} \leq V$,

$$\text{BREV}^c(\widehat{V}) = \max_{p \leq c} \left(p \cdot \Pr_{\widehat{v} \sim \widehat{V}}[\|\widehat{v}\|_1 \geq p] \right) \leq \max_{p \leq b} \left(p \cdot \Pr_{v \sim V}[\|v\|_1 \geq p] \right) = \text{BREV}^b(V).$$

For SREV , let \widehat{M} be the optimal mechanism that achieves $\text{SREV}^c(\widehat{V})$ by pricing each item separately. We construct a mechanism M to mimic \widehat{M} except the prices are capped at $b/2$. Consider applying M to a buyer with valuation drawn from V and a budget b . As $\widehat{V} \leq V$, we can couple the realizations $\widehat{v} \sim \widehat{V}$ and $v \sim V$ such that $\widehat{v} \leq v$ (coordinate-wise). For every (\widehat{v}, v) pair:

- If M gets a revenue of at least $\frac{b}{2}$ on v . This is at least $\frac{b}{2c}$ -fraction of the revenue \widehat{M} gets on \widehat{v} , because the latter is at most c .
- If M gets a revenue less than $\frac{b}{2}$ on v , then the buyer has enough budget left to buy any item. Therefore, the buyer can buy everything he wants. Because $\widehat{v} \leq v$, the revenue of M on v is at least that of \widehat{M} on \widehat{v} .

Thus, M can get at least $\min(1, \frac{b}{2c})$ -fraction of the revenue that \widehat{M} gets on \widehat{v} , which implies $\text{SREV}(\widehat{V}) \leq \max(1, \frac{2c}{b}) \text{SREV}^b(V)$. \square

4 PUBLIC BUDGET AND WEAKLY CORRELATED VALUATIONS

In this section, we present an alternative approach to prove our main result (Theorem 1.1). Recall that the buyer has a public budget b and valuations drawn from an independent distribution V .

Theorem 1.1 (stronger version). $\text{REV}^b(V) \leq 5\text{SREV}^b(V) + 6\text{BREV}^b(V)$.

Overview of Our Approach. We will truncate the input distribution V in a different way: instead of truncating $v \sim V$ in ℓ_∞ -norm (as in Section 3), we will truncate v in ℓ_1 -norm. This truncation produces a correlated distribution \widehat{V} . The upshot of truncating in ℓ_1 -norm is that we always have $\|\widehat{v}\|_1 \leq b$, so \widehat{V} can ignore the budget. In addition, as in Section 3, we can relate the optimal revenue to the revenue of \widehat{V} (Lemma 3.5), and we can relate the revenue of simple mechanisms on \widehat{V} back to revenue of simple mechanisms on V (Lemma 3.7).

We still need to argue that simple mechanisms work well for \widehat{V} . This is the main challenge in this approach. Because \widehat{V} is correlated, we cannot apply the result of [Babaioff et al. 2014] in a black-box way. Instead, we need to modify the analysis of previous work [Babaioff et al. 2014; Cai et al. 2016; Li and Yao 2013] and build on the key ideas like “core-tail” decomposition. More specifically, we generalize the duality-based framework developed in [Cai et al. 2016] to handle the specific type of correlation \widehat{V} has.

Weakly Correlated Valuations. It is worth mentioning that our structural lemmas (Lemmas 3.5 and 3.7) do not require the input distribution to be independent. This is why our techniques can be applied to more general settings. For example, in this section, we will generalize Theorem 1.1 to handle weakly correlated valuations (see Definition 2.1 for the formal definition).

Corollary 1.2. *Let \widehat{V} be a weakly correlated distribution (Definition 2.1). We have $\text{REV}^b(\widehat{V}) \leq 5\text{SREV}^b(\widehat{V}) + 6\text{BREV}^b(\widehat{V})$.*

Our main contribution in this section is Lemma 4.1. Lemma 4.1 shows that for any weakly correlated distribution \widehat{V} (see Definition 2.1), the better of $\text{SREV}(\widehat{V})$ and $\text{BREV}(\widehat{V})$ is a constant approximation to the optimal revenue $\text{REV}(\widehat{V})$.

LEMMA 4.1. *Fix $c > 0$. Let $\widehat{V} = V_{(\|\widehat{v}\|_1 \leq c)}$ for an independent distribution V . We have $\text{REV}(\widehat{V}) \leq 5\text{SREV}(\widehat{V}) + 4\text{BREV}(\widehat{V})$.*

We defer the proof of Lemma 4.1 to Appendix C. We first use these lemmas to prove Theorem 1.1 and Corollary 1.2.

PROOF OF THEOREM 1.1 AND COROLLARY 1.2. If $\min_{v \sim V} \|v\|_1 \geq b/2$, then the seller can price the grand bundle at $b/2$ and the buyer always buys it. In this case, the revenue is $b/2$ and $\text{REV}^b(V) \leq b \leq 2\text{BREV}^b(V)$. Thus, we focus on the more interesting case where $\Pr_{v \sim V}[\|v\|_1 \leq b/2] > 0$.⁷

Let $\widehat{V} := V_{(\|\widehat{v}\|_1 \leq c)}$ for $c = b/2$. We will reuse Lemmas 3.5 and 3.7 from Section 3. We can reuse both lemmas because they do not require V or \widehat{V} to be independent, \widehat{V} does not modify the small-sum part of V , and $\widehat{V} \leq V$ (which we will prove as Lemma D.2 in Appendix D).

$$\begin{aligned}
\text{REV}^b(V) &\leq (b/c) \cdot \text{BREV}^b(V) + \text{REV}(\widehat{V}) && \text{(Lemma 3.5)} \\
&\leq 2\text{BREV}^b(V) + 5\text{SREV}(\widehat{V}) + 4\text{BREV}(\widehat{V}) && \text{(Lemma 4.1)} \\
&= 2\text{BREV}^b(V) + 5\text{SREV}^c(\widehat{V}) + 4\text{BREV}^c(\widehat{V}) && (\|\widehat{v}\|_1 \leq c) \\
&\leq 2\text{BREV}^b(V) + 5\text{SREV}^b(V) + 4\text{BREV}^b(V). && \text{(Lemma 3.7)}
\end{aligned}$$

We now prove Corollary 1.2. Intuitively, Corollary 1.2 holds because simple mechanisms work well for weakly correlated valuations, and the weakly-correlated notion is closed under further capping the sum.

⁷Throughout the paper, when we consider the conditional distribution $\widehat{V} := V_{(\|\widehat{v}\|_1 \leq c)}$, we will always have $c > \min_{v \in \text{supp}(V)} \|v\|_1$, so that the event we condition on happens with non-zero probability.

Let $V = V_{(\|v\|_1 \leq c_2)}$ be the input distribution. If $c_2 \leq b$, then we can remove the budget constraint and apply Lemma 4.1 directly. If $c_2 > b$, then we can cap V at $c_1 = b/2$ to obtain a weakly correlated distribution \widehat{V} . One can verify that Lemmas 3.5 and 3.7 still hold for V and \widehat{V} , and Lemma 4.1 holds for \widehat{V} . The only difference is that we need to show $V_{(\|v\|_1 \leq c_1)} \leq V_{(\|v\|_1 \leq c_2)}$ for $c_1 \leq c_2$. We will prove this (Lemma D.1) in Appendix D. \square

5 PRIVATE BUDGET

In this section, we consider the case where the budget b is no longer fixed but instead drawn from a distribution B . We assume B is independent of V in light of the following observations:

If we allow B and V to be correlated, one natural model is that the buyer's budget b is first drawn from B , and then depending on the value of b , the buyer's valuations are drawn independently for each item. In this case, the problem is at least as hard as finding (approximately) optimal mechanisms for arbitrarily correlated valuations in the budget-free setting. Consider an instance in which all possible budgets are larger than $\max_{v \sim V} \|v\|_1$ so they are irrelevant. However, the budget can still be used as a signal (or a correlation device) to produce correlated valuations. It is known that for correlated distributions, the better of selling separately and bundling together [Hart and Nisan 2013], or even the best partition-based mechanism [Babaioff et al. 2014], does not offer a constant approximation.

Another natural model is we draw $v \sim V$ first where the V_i 's are independent, and then draw b depending on v . This is also at least as hard as the correlated-valuation case, since we can make $b = 0$ if we want v in the correlated distribution and $b > \|v\|_1$ if we do not.

These negative observations motivate us to study the private budget setting when the budget distribution B is independent of the valuation distributions V .

5.1 Monotone-Hazard-Rate Budgets

We focus on the case where the budget is independent of valuations, and it is drawn from a continuous⁸ monotone-hazard-rate (MHR) distribution. Let $g(\cdot)$ and $G(\cdot)$ be the probability density function and cumulative distribution function of B . The MHR condition says $\frac{g(b)}{1-G(b)}$ is non-decreasing in b .

LEMMA 5.1. *Let b^* be the expectation of an MHR distribution B . Let M^* be the optimal mechanism for a buyer with a public budget b^* . Then in expectation, M^* extracts at least $\frac{1}{2e}$ -fraction of the expected optimal revenue when the buyer has a private budget drawn from B .*

PROOF. Let $R(b, V)$ denote the expected revenue of M^* when the buyer has a public budget b and valuations drawn from V . Let $R(B, V) = \mathbb{E}_{b \sim B}[R(b, V)]$ denote the expected revenue of M^* when the buyer's budget is drawn from B .

$$\begin{aligned} R(B, V) &= \int_b g(b)R(b, V)db \geq \int_{b \geq b^*} g(b)R(b, V)db \\ &= \int_{b \geq b^*} g(b)R(b^*, V)db \geq e^{-1} \cdot R(b^*, V). \end{aligned}$$

The second last step uses $R(b, V) = R(b^*, V)$ when $b \geq b^*$, because M^* provides a menu of allocation/payment pairs for the buyer to choose from; A buyer with budget $b \geq b^*$ can afford any option on the menu so he will choose the same option as if he had budget b^* . The last inequality comes from the fact that for any MHR distribution B , $\Pr_{b \sim B}[b \geq b^*] \geq e^{-1}$ (see, e.g., [Barlow and Marshall 1965]).

⁸If the distribution is a discrete MHR distribution, similar results still hold. For discrete distributions we have $\Pr_{b \sim B}[b \geq \lfloor b^* \rfloor] \geq e^{-1}$ instead of $\Pr_{b \sim B}[b \geq b^*] \geq e^{-1}$.

Let $\text{REV}^B(V)$ denote the optimal revenue we can extract when the buyer has private budgets drawn from B .

$$\begin{aligned} \text{REV}^B(V) &\leq \int_{b < b^*} g(b) \text{REV}^b(V) db + \int_{b \geq b^*} g(b) \text{REV}^b(V) db \\ &\leq \int_{b < b^*} g(b) \text{REV}^{b^*}(V) db + \int_{b \geq b^*} g(b) \cdot \frac{b}{b^*} \cdot \text{REV}^{b^*}(V) db \\ &\leq \text{REV}^{b^*}(V) + \frac{\int_b g(b) b db}{b^*} \cdot \text{REV}^{b^*}(V) = 2\text{REV}^{b^*}(V). \end{aligned}$$

The first line is because the seller can only do better if she knows the buyer's budget b . The second line uses the fact that $\text{REV}^b(V) \leq \text{REV}^{b^*}(V)$ when $b < b^*$ and $\text{REV}^b(V) \leq \frac{b}{b^*} \text{REV}^{b^*}(V)$ when $b > b^*$. The third line is because $b^* = \mathbb{E}[b]$.

We have $\text{REV}^b(V) \leq \text{REV}^{b^*}(V)$ when $b < b^*$ because a buyer with budget b^* can afford all options from the menu that achieves $\text{REV}^b(V)$. When $b > b^*$, consider the menu that achieves $\text{REV}^b(V)$ and cap all prices at b^* . A buyer with budget $b > b^*$ either chooses the same option as if he had budget b^* , or chooses a different option whose price must be b^* , and therefore $\text{REV}^b(V) \leq \frac{b}{b^*} \text{REV}^{b^*}(V)$.

By definition $R(b^*, V) = \text{REV}^{b^*}(V)$. Therefore, $R(B, V) \geq \frac{1}{2e} \text{REV}^B(V)$. \square

Theorem 1.3. *When the budget distribution B is MHR, the better of pricing items separately and bundling them together achieves a constant fraction of the optimal revenue.*

PROOF. By pretending the budget is b^* ,

$$\text{SREV}^B(V) \geq \int_{b \geq b^*} g(b) \text{SREV}^{b^*}(V) db \geq \frac{1}{e} \text{SREV}^{b^*}(V).$$

Similarly, $\text{BREV}^B(V) \geq \frac{1}{e} \text{BREV}^{b^*}(V)$. Therefore, by Theorem 1.1 and Lemma 5.1, $\text{SREV}^B(V) + \text{BREV}^B(V) = \Omega(\text{SREV}^{b^*}(V) + \text{BREV}^{b^*}(V)) = \Omega(\text{REV}^{b^*}(V)) = \Omega(\text{REV}^B(V))$. \square

6 CONCLUSION AND FUTURE DIRECTIONS

In this paper, we investigated the effectiveness of simple mechanisms in the presence of budgets, and showed that for an additive buyer with independent valuations and a public budget, either selling separately or selling the grand bundle gives a constant approximation to optimal revenue.

The area of designing simple and approximately optimal auctions with budget constraints is still largely unexplored. Our work leaves many natural follow-up questions. We only considered selling to a single buyer. An immediate open question is whether our results can be extended to multiple bidders. A generalization to multiple bidders is known in the budget-free case [Cai et al. 2016; Yao 2015].

Question 1. *Is there a simple mechanism that is approximately optimal for multiple additive buyers, when each buyer has the same public budget b ?*

For private budgets where the budget is independent of the valuations, we showed that if the budget distribution satisfies monotone hazard rate, then we can extract a constant fraction of the revenue. The general case with arbitrary budget distributions appears to be nontrivial and is an interesting avenue for future work.

Question 2. *Is there a simple mechanism that is approximately optimal for an additive buyer with private budgets, when the budget distribution is independent of the valuations?*

ACKNOWLEDGMENTS

Yu Cheng is supported by NSF grants CCF-1527084, CCF-1535972, CCF-1637397, CCF-1704656, IIS-1447554, and NSF CAREER Award CCF-1750140. Kamesh Munagala is supported by NSF grants CCF-1408784, CCF-1637397, and IIS-1447554; and by an Adobe Data Science Research Award. Kangning Wang is supported by NSF grants CCF-1408784 and CCF-1637397.

REFERENCES

- Zoë Abrams. 2006. Revenue maximization when bidders have budgets. In *Proc. 17th ACM-SIAM Symp. on Discrete Algorithms*. 1074–1082.
- Saeed Alaei, Hu Fu, Nima Haghpanah, Jason D. Hartline, and Azarakhsh Malekian. 2012. Bayesian optimal auctions via multi- to single-agent reduction. In *Proc. 13th ACM Conf. on Electronic Commerce*. 17.
- Moshe Babaioff, Yannai A. Gonczarowski, and Noam Nisan. 2017. The menu-size complexity of revenue approximation. In *Proc. 49th ACM Symp. on Theory of Computing*. 869–877.
- Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. 2014. A simple and approximately optimal mechanism for an additive buyer. In *Proc. 55th IEEE Symp. on Foundations of Computer Science*. 21–30.
- Richard E Barlow and Albert W Marshall. 1965. Tables of bounds for distributions with monotone hazard rate. *J. Amer. Statist. Assoc.* 60, 311 (1965), 872–890.
- MohammadHossein Bateni, Sina Dehghani, MohammadTaghi Hajiaghayi, and Saeed Seddighin. 2015. Revenue maximization for selling multiple correlated items. In *Proc. 23rd European Symp. on Algorithms*. 95–105.
- Xiaohui Bei, Ning Chen, Nick Gravin, and Pinyan Lu. 2012. Budget feasible mechanism design: from prior-free to Bayesian. In *Proc. 44th ACM Symp. on Theory of Computing*. 449–458.
- Sayan Bhattacharya, Vincent Conitzer, Kamesh Munagala, and Lirong Xia. 2010a. Incentive compatible budget elicitation in multi-unit auctions. In *Proc. 21st ACM-SIAM Symp. on Discrete Algorithms*. 554–572.
- Sayan Bhattacharya, Gagan Goel, Sreenivas Gollapudi, and Kamesh Munagala. 2010b. Budget constrained auctions with heterogeneous items. In *Proc. 42nd ACM Symp. on Theory of Computing*. 379–388.
- Christian Borgs, Jennifer T. Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi. 2005. Multi-unit auctions with budget-constrained bidders. In *Proc. 6th ACM Conf. on Electronic Commerce*. 44–51.
- Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. 2010. Pricing randomized allocations. In *Proc. 21st ACM-SIAM Symp. on Discrete Algorithms*. 585–597.
- Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 2012a. An algorithmic characterization of multi-dimensional mechanisms. In *Proc. 44th ACM Symp. on Theory of Computing*. 459–478.
- Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 2012b. Optimal multi-dimensional mechanism design: reducing revenue to welfare maximization. In *Proc. 53rd IEEE Symp. on Foundations of Computer Science*. 130–139.
- Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 2013a. Reducing revenue to welfare maximization: Approximation algorithms and other generalizations. In *Proc. 24th ACM-SIAM Symp. on Discrete Algorithms*. 578–595.
- Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 2013b. Understanding incentives: Mechanism design becomes algorithm design. In *Proc. 54th IEEE Symp. on Foundations of Computer Science*. 618–627.
- Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. 2016. A duality based unified approach to Bayesian mechanism design. In *Proc. 48th ACM Symp. on Theory of Computing*. 926–939.
- Yang Cai and Mingfei Zhao. 2017. Simple mechanisms for subadditive buyers via duality. In *STOC17*. 170–183.
- Gabriel Carroll. 2017. Robustness and separation in multidimensional screening. *Econometrica* 85, 2 (2017), 453–488.
- Deeparnab Chakrabarty and Gagan Goel. 2010. On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and GAP. *SIAM J. Comput.* 39, 6 (2010), 2189–2211.
- Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. 2007. Algorithmic pricing via virtual valuations. In *Proc. 8th ACM Conf. on Electronic Commerce*. 243–251.
- Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proc. 42nd ACM Symp. on Theory of Computing*. 311–320.
- Shuchi Chawla, David L. Malec, and Azarakhsh Malekian. 2011. Bayesian mechanism design for budget-constrained agents. In *Proc. 12th ACM Conf. on Electronic Commerce*. 253–262.
- Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. 2015. The power of randomness in Bayesian optimal mechanism design. *Games and Economic Behavior* 91 (2015), 297–317.
- Shuchi Chawla and J Benjamin Miller. 2016. Mechanism design for subadditive agents via an ex ante relaxation. In *ECOM16*. 579–596.
- Yeon-Koo Che and Ian L. Gale. 2000. The optimal mechanism for selling to a budget-constrained buyer. *J. Economic Theory* 92, 2 (2000), 198–233.

- Ning Chen, Nick Gravin, and Pinyan Lu. 2011. On the approximability of budget feasible mechanisms. In *Proc. 22nd ACM-SIAM Symp. on Discrete Algorithms*. 685–699.
- Yu Cheng, Ho Yee Cheung, Shaddin Dughmi, Ehsan Emamjomeh-Zadeh, Li Han, and Shang-Hua Teng. 2015. Mixture selection, mechanism design, and signaling. In *Proc. 56th IEEE Symp. on Foundations of Computer Science*. 1426–1445.
- Constantinos Daskalakis. 2015. Multi-item auctions defying intuition? *SIGecom Exchanges* 14, 1 (2015), 41–75.
- Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2017. Strong Duality for a Multiple-Good Monopolist. *Econometrica* 85, 3 (2017), 735–767.
- Constantinos Daskalakis, Nikhil R. Devanur, and S. Matthew Weinberg. 2015. Revenue maximization and ex-post budget constraints. In *Proc. 16th ACM Conf. on Economics and Computation*. 433–447.
- Nikhil R. Devanur, Bach Q. Ha, and Jason D. Hartline. 2013. Prior-free auctions for budgeted agents. In *Proc. 14th ACM Conf. on Electronic Commerce*. 287–304.
- Nikhil R. Devanur and S. Matthew Weinberg. 2017. The optimal mechanism for selling to a budget constrained buyer: The general case. In *Proc. 18th ACM Conf. on Economics and Computation*. 39–40.
- Shahar Dobzinski, Ron Lavi, and Noam Nisan. 2012. Multi-unit auctions with budget limits. *Games and Economic Behavior* 74, 2 (2012), 486–503.
- Shaddin Dughmi, Li Han, and Noam Nisan. 2014. Sampling and representation complexity of revenue maximization. In *Proc. 10th Conference on Web and Internet Economics (WINE)*. 277–291.
- Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. 2017. A simple and approximately optimal mechanism for a buyer with complements. In *Proc. 18th ACM Conf. on Economics and Computation*. 323.
- Gagan Goel, Vahab S. Mirrokni, and Renato Paes Leme. 2015. Polyhedral clinching auctions and the adWords polytope. *J. ACM* 62, 3 (2015), 18:1–18:27.
- Andrew V. Goldberg, Jason D. Hartline, and Andrew Wright. 2001. Competitive auctions and digital goods. In *Proc. 12th ACM-SIAM Symp. on Discrete Algorithms*. 735–744.
- Nick Gravin and Pinyan Lu. 2018. Separation in Correlation-Robust Monopolist Problem with Budget. In *Proc. 29th ACM-SIAM Symp. on Discrete Algorithms*. 2069–2080.
- Sergiu Hart and Noam Nisan. 2013. The menu-size complexity of auctions. In *Proc. 14th ACM Conf. on Electronic Commerce*. 565–566.
- Sergiu Hart and Noam Nisan. 2017. Approximate revenue maximization with multiple items. *Journal of Economic Theory* 172 (2017), 313–347.
- Sergiu Hart and Philip J. Reny. 2015. Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Theoretical Economics* 10, 3 (2015), 893–922.
- Jean-Jacques Laffont and Jacques Robert. 1996. Optimal auction with financially constrained buyers. *Economics Letters* 52, 2 (1996), 181–186.
- Xinye Li and Andrew Chi-Chih Yao. 2013. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences* 110, 28 (2013), 11232–11237.
- Alejandro M. Manelli and Daniel R. Vincent. 2007. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic theory* 137, 1 (2007), 153–185.
- Roger B. Myerson. 1981. Optimal auction design. *Math. Oper. Res.* 6, 1 (1981), 58–73.
- Malles M. Pai and Rakesh Vohra. 2014. Optimal auctions with financially constrained buyers. *Journal of Economic Theory* 150 (2014), 383–425.
- Aviad Rubinfeld and S. Matthew Weinberg. 2015. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proc. 16th ACM Conf. on Economics and Computation*. 377–394.
- Yaron Singer. 2010. Budget Feasible Mechanisms. In *Proc. 51st IEEE Symp. on Foundations of Computer Science*. 765–774.
- Adish Singla and Andreas Krause. 2013. Truthful incentives in crowdsourcing tasks using regret minimization mechanisms. In *WWW*. 1167–1178.
- Pingzhong Tang and Ziheng Wang. 2017. Optimal mechanisms with simple menus. *Journal of Economic Theory* 69 (2017), 54–70.
- Andrew Chi-Chih Yao. 2015. An n -to-1 bidder reduction for multi-item auctions and its Applications. In *Proc. 26th ACM-SIAM Symp. on Discrete Algorithms*. 92–109.

A PROOF OF THE CONCENTRATION LEMMA IN SECTION 3.2

In this section, we prove Lemma 3.6. We first restate it for convenience.

Lemma 3.6. *If V' is independent and $\|v'\|_\infty \leq c$ for all $v' \sim V'$, then*

$$\Pr_{v' \sim V'} [\|v'\|_1 \geq x + y + c] \leq \Pr_{v' \sim V'} [\|v'\|_1 \geq x] \cdot \Pr_{v' \sim V'} [\|v'\|_1 \geq y] \quad \text{for all } x, y > 0.$$

In particular, if $\Pr_{v' \sim V'}[\|v'\|_1 \geq c] \leq q$, then for all integer $k \geq 0$,

$$\Pr_{v' \sim V'}[\|v'\|_1 \geq (2k+1)c] \leq q^k \Pr_{v' \sim V'}[\|v'\|_1 \geq c].$$

PROOF. Consider the probability of $\|v'\|_1 \geq x + y + c$ conditioned on $\|v'\|_1 \geq x$. We will show this probability is at most the probability of $\|v'\|_1 \geq y$.

For every v' with $\|v'\|_1 \geq x$, there is a unique $j \in [m]$ where $\sum_{i=1}^{j-1} v'_i < x$ but $\sum_{i=1}^j v'_i \geq x$. Now v'_j is at most c , so for the total sum to be at least $x + y + c$, the remaining sum $\sum_{i=j+1}^m v'_i$ must be at least y . Due to the independence of V' , this probability is the same conditioned on any values of (v'_1, \dots, v'_j) . Formally,

$$\begin{aligned} \Pr[\|v'\|_1 \geq x + y + c] &= \sum_{j=1}^m \Pr \left[\|v'\|_1 \geq x + y + c \wedge \sum_{i=1}^{j-1} v'_i < x \wedge \sum_{i=1}^j v'_i \geq x \right] \\ &\leq \sum_{j=1}^m \Pr \left[\sum_{i=j+1}^m v'_i \geq y \wedge \sum_{i=1}^{j-1} v'_i < x \wedge \sum_{i=1}^j v'_i \geq x \right] \\ &= \sum_{j=1}^m \Pr \left[\sum_{i=j+1}^m v'_i \geq y \right] \cdot \Pr \left[\sum_{i=1}^{j-1} v'_i < x \wedge \sum_{i=1}^j v'_i \geq x \right] \\ &= \Pr[\|v'\|_1 \geq y] \cdot \Pr[\|v'\|_1 \geq x]. \end{aligned}$$

The second statement can be proved inductively using the first statement. The inductive step chooses $x = c$ and $y = (2k - 1)c$. \square

B REVENUE NON-MONOTONICITY FOR SEPARATE SELLING TO A BUDGETED BUYER

We provide an example where $V_1 \leq V_2$ but $\text{SREV}^b(V_1) > \text{SREV}^b(V_2)$. Intuitively, a budget-constrained buyer solves a KNAPSACK problem when deciding which items to purchase, and the total volume (i.e, payment) of the optimal KNAPSACK solution is not monotone in the item values. Increasing the value of a cheap item might incentivize the buyer to purchase this item instead of a more expensive one, if the buyer does not have enough budget to buy both items.

Consider an auction with 3 items. T_1 and T_2 are two matrices defined as

$$T_1 = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 1 & 1 \\ 2 & 1 & 0 \end{bmatrix}, \quad T_2 = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 1 & 1 \\ 2 & 2 & 0 \end{bmatrix}.$$

The rows of T_1 are the support of V_1 , and similarly, the rows of T_2 are the support of V_2 . Each row is associated with a probability of $\frac{1}{3}$. Assume the budget $b = 2$.

One of the optimal mechanisms that obtains $\text{SREV}^b(V_1) = 2$ is to price the items at $(2, 1, 1)$. A buyer of type 1 and 3 will buy the first item, and a buyer of type 2 will buy the last two items. This is optimal because SREV^b cannot exceed the budget.

However, $\text{SREV}^b(V_2) < 2$. To prove it, we first notice $\text{SREV}^b(V_2) \leq 2$ because $b = 2$. It means we must get a revenue of 2 from all buyer types to make $\text{SREV}^b(V_2) = 2$. Thus, we must price the first item at 2, and each of the last two items at 1 to satisfy this constraint for the first two types. Nevertheless, with this pricing strategy, we only get revenue 1 for the last buyer type $(2, 2, 0)$, because he will only buy the second item. This shows $\text{SREV}^b(V_2) < 2 = \text{SREV}^b(V_1)$.

C SIMPLE MECHANISMS FOR WEAKLY CORRELATED VALUATIONS

This section is devoted to proving Lemma 4.1. Lemma 4.1 states that simple mechanisms (more specifically, the better of $\text{SREV}(\widehat{V})$ and $\text{BREV}(\widehat{V})$) are approximately optimal for any weakly correlated distribution $\widehat{V} = V_{(\|v\|_1 \leq c)}$. Note that there is no budget in this section, only a cap c on the ℓ_1 -norm of $\widehat{v} \sim \widehat{V}$.

Our approach builds on the ideas like “core-tail decomposition” from previous works that show $\text{REV}(V) = O(\text{SREV}(V) + \text{BREV}(V))$ for independent valuations V [Babaioff et al. 2014; Cai et al. 2016; Hart and Nisan 2017; Li and Yao 2013]. More specifically, we generalize the duality-based framework developed in [Cai et al. 2016] to handle weakly correlated distributions. The idea of [Cai et al. 2016] is to Lagrangify only the incentive constraints, then guess the Lagrangian multipliers to derive an upper bound on the maximum revenue.

We first highlight some of the difficulties in extending previous works to correlated distributions.

- (1) When the distribution is independent, one can upper bound the maximum revenue by Myerson’s *virtual value* of the bidder’s favorite item, plus the sum of the *values* of the remaining items. For correlated distributions, the virtual value of the favorite item depends on the other items.
- (2) In the core part of core-tail decomposition, we need the total value of the low-value items to concentrate around its expectation, so we can upper bound their values by BREV (by selling the grand bundle at a price slightly lower than that expectation). When the valuations are independent, we can show the variance is small, which may not be true for correlated distributions.

These difficulties are not surprising, because Lemma 4.1 cannot hold for arbitrary correlated distributions. As shown in [Briest et al. 2010; Hart and Nisan 2013], for correlated distributions, the gap between the best deterministic and randomized mechanisms can be unbounded. Hence, we have to take advantage of the special properties of \widehat{V} .

Notations. In this Section, because $\widehat{V} = V_{(\|v\|_1 \leq c)}$ is the distribution we focus on, we use T to denote the support of \widehat{V} . Given a (correlated) distribution D , we use D_j to denote D ’s marginal distribution on the j -th coordinate, and D_{-j} to denote D ’s marginal (joint) distribution on coordinates other than j . Let T_j and T_{-j} be the support of \widehat{V}_j and \widehat{V}_{-j} respectively. In addition, we will make use of the conditional distributions $\widehat{V}_j|_{\widehat{v}_{-j}=t_{-j}}$ and $\widehat{V}_{-j}|_{\widehat{v}_j=t_j}$; The former is the distribution of \widehat{v}_j for $\widehat{v} \sim \widehat{V}$ conditioned on $\widehat{v}_{-j} = t_{-j}$, and the latter is the distribution of \widehat{v}_{-j} conditioned on $v_j = t_j$. Abusing notation, we use $f(t)$, $f(t_j)$, $f(t_{-j})$, $f(t_j|t_{-j})$, and $f(t_{-j}|t_j)$ to denote the probability mass function of the correlated and conditional distributions we mentioned in this section. When the value of item j is drawn from $f(t_j|t_{-j})$, we use $\widetilde{\varphi}(t_j|t_{-j})$ to denote item j ’s (ironed) Myerson’s virtual value.

For a bidder type $t \in T \subseteq \mathbb{R}^m$, the favorite item of t is the one with the highest value (with ties broken lexicographically). We write $t \in R_j$ if and only if j is the favorite item of type t after tie-breaking. Formally,

$$R_j = \{t \in T \mid j \text{ is the smallest index with } t_j \geq t_k \text{ for all } k \in [m]\}.$$

Proof of Lemma 4.1. We extend the duality framework in [Cai et al. 2016] to handle correlated distributions. As we will see in Section C.2, we can upper bound the optimal revenue of \widehat{V} into three components. Recall that π is the allocation rule. For notational convenience, we write r for

SREV(\widehat{V}). Notice that in SINGLE we get $\widetilde{\varphi}(t_j|t_{-j})$ rather than $\widetilde{\varphi}(t_j)$.

$$\text{REV}(\widehat{V}) \leq \sum_{t \in T} f(t) \sum_{j \in [m]} \pi_j(t) \widetilde{\varphi}(t_j|t_{-j}) \cdot \mathbb{1}[t \in R_j] \quad (\text{SINGLE})$$

$$+ \sum_{t \in T} f(t) \sum_{j \in [m]} \pi_j(t) t_j \cdot \mathbb{1}[t \notin R_j] \cdot \mathbb{1}[t_j \leq 2r] \quad (\text{CORE})$$

$$+ \sum_{t \in T} f(t) \sum_{j \in [m]} \pi_j(t) t_j \cdot \mathbb{1}[t \notin R_j] \cdot \mathbb{1}[t_j > 2r]. \quad (\text{TAIL})$$

Lemma 4.1 follows directly from the statements of Lemmas C.1, C.2, and C.3.

LEMMA C.1. SINGLE \leq 2BREV(\widehat{V}) + SREV(\widehat{V}).

LEMMA C.2. TAIL \leq SREV(\widehat{V}).

LEMMA C.3. CORE \leq 2BREV(\widehat{V}) + 3SREV(\widehat{V}).

Organization. For completeness, we first recall the approach in [Cai et al. 2016] in Section C.1. In Section C.2, we show that there exists a choice of the Lagrangian multipliers that REV(\widehat{V}) can be upper bounded by SINGLE + CORE + TAIL. Lemmas C.1, C.2, and C.3 are proved in Section C.3.

C.1 The Duality Based Approach in [Cai et al. 2016].

The optimal mechanism $M^* = (\pi^*, p^*)$ for \widehat{V} is the optimal solution to the following LP:

$$\begin{aligned} & \text{maximize} && \sum_{t \in T} f(t) p(t) \\ & \text{subject to} && \pi(t)^\top t - p(t) \leq \pi(t')^\top t - p(t'), \quad \forall t \in T, t' \in T^+. \\ & && 0 \leq \pi_j(t) \leq 1, \quad \forall t \in T, j \in [m]. \\ & && \pi(\emptyset) = 0, p(\emptyset) = 0. \end{aligned} \quad (2)$$

We can upper bound the optimal primal value by Lagrangifying the incentive constraints.

$$\text{REV}(\widehat{V}) = \min_{\lambda \geq 0} \max_{\pi, p} L(\lambda, \pi, p),$$

where the Lagrangian dual of LP (2) is given by

$$\begin{aligned} L(\lambda, \pi, p) &= \sum_{t \in T} f(t) p(t) + \sum_{t \in T, t' \in T^+} \lambda(t, t') [(\pi(t) - \pi(t'))^\top t - (p(t) - p(t'))] \\ &= \sum_{t \in T} p(t) \left[f(t) - \sum_{t' \in T^+} \lambda(t, t') + \sum_{t' \in T} \lambda(t', t) \right] \\ &\quad + \sum_{t \in T} \pi(t)^\top \left[\sum_{t' \in T^+} \lambda(t, t') t - \sum_{t' \in T} \lambda(t', t) t' \right]. \end{aligned}$$

Because the $p(t)$'s are unconstrained variables, any dual solution with a finite value must have

$$f(t) - \sum_{t' \in T^+} \lambda(t, t') + \sum_{t' \in T} \lambda(t', t) = 0, \quad \forall t \in T. \quad (3)$$

From now on, we restrict our attention to only dual solution with finite values. We can then simplify $L(\lambda, \pi, p)$ by replacing $\sum_{t' \in T^+} \lambda(t, t')$ with $f(t) + \sum_{t' \in T} \lambda(t', t)$ to get rid of T^+ :

$$\begin{aligned} L(\lambda, \pi, p) &= \sum_{t \in T} \pi(t)^\top \left[f(t)t + \sum_{t' \in T} \lambda(t', t)t - \sum_{t' \in T} \lambda(t', t)t' \right] \\ &= \sum_{t \in T} f(t)\pi(t)^\top \left[t - \frac{1}{f(t)} \sum_{t' \in T} \lambda(t', t)(t' - t) \right]. \end{aligned}$$

We write $\Phi(t)$ as a shorthand for the term in the bracket: $\Phi(t) = t - \frac{1}{f(t)} \sum_{t' \in T} \lambda(t', t)(t' - t)$. We know that $L(\lambda, \pi, p) = \sum_{t \in T} f(t)\pi(t)^\top \Phi(t) \geq \text{REV}(\widehat{V})$ is an upper bound on the revenue of the optimal mechanism. We can rewrite Equation (3) as

$$f(t) + \sum_{t' \in T} \lambda(t', t) = \sum_{t' \in T^+} \lambda(t, t') = \sum_{t' \in T} \lambda(t, t') + \lambda(t, \emptyset), \quad \forall t \in T.$$

[Cai et al. 2016] interpreted these constraints as flow conservation constraints. Let $\lambda(t, t') \geq 0$ denote the amount of flow t sends to t' . The left-hand side is the total flow received by t , where every type t receives $f(t)$ units of flow from the source; and the right-hand side is the total flow send out from t , with all the excess flow sent to the sink (\emptyset).

They proposed a ‘‘canonical flow’’ which was shown to be a good guess for the Lagrangian multipliers. It turns out the same dual solution is sufficient to prove our results for correlated distributions. In the next section, we recall this canonical flow and use it to derive an upper bound on the optimal revenue.

C.2 Canonical Flow for Weakly Correlated Distributions

Recall that $t \in R_j$ if and only if j is the favorite item of type t . Formally, there exists $\lambda(t, t') \geq 0$ such that

- For every j , all flows entering R_j are from the source, and all flows leaving R_j are to \emptyset .
- For $t, t' \in R_j$, we can have $\lambda(t', t) > 0$ only if t and t' only differ on the j -th coordinate. When there is no ironing, $\lambda(t', t) > 0$ only if t'_j is the smallest value larger than t_j in T_j .

LEMMA C.4. *There exists a set of the Lagrangian multipliers λ that satisfies the flow conservation constraints, such that*

- (1) *If $t \notin R_j$, then $\Phi_j(t) = t_j$.*
- (2) *If $t \in R_j$, then $\Phi_j(t) \leq \widetilde{\varphi}_j(t_j | t_{-j})$, where $\widetilde{\varphi}_j(t_j | t_{-j})$ is item j 's (ironed) Myerson's virtual value conditioned on t_{-j} .*

PROOF. Recall that $\Phi(t) = t - \frac{1}{f(t)} \sum_{t' \in T} \lambda(t', t)(t' - t)$. For (1), assume that $t \in R_k$ for some $k \neq j$. If $\lambda(t', t) > 0$, it must be the case that t and t' are only different on the k -th coordinate, so $(t - t')_j = 0$ and $\Phi_j(t) = t_j$.

Now we prove (2). We first consider a canonical flow without ironing. Fix any j and $t \in R_j$. If t_j is the largest value in T_j , then $\lambda(t', t) = 0$ for all t' and $\Phi_j(t) = t_j$. If t_j is not the largest value in T_j , t receives flow from the source and exactly one other node t' where $t'_{-j} = t_{-j}$, and t'_j is the smallest value larger than t_j in T_j . The total flow from t' to t includes the flows from the source to all t^*

with $t_{-j}^* = t_{-j}$ and $t_j^* > t_j$:

$$\begin{aligned}\lambda(t', t) &= \sum_{t^* \in T} f(t^*) \cdot \mathbb{1}[t_{-j}^* = t_{-j} \wedge t_j^* > t_j] \\ &= f(t_{-j}) \sum_{t_j^* \in T_j, t_j^* > t_j} f(t_j^* | t_{-j}) = f(t_{-j}) (1 - F(t_j | t_{-j})).\end{aligned}$$

Substituting $\lambda(t', t)$ in the expression of $\Phi(t)$, this implies

$$\begin{aligned}\Phi_j(t) &= t_j - \frac{1}{f(t)} \sum_{t^* \in T} \lambda(t^*, t)(t_j^* - t_j) \\ &= t_j - \frac{1}{f(t)} \lambda(t', t)(t_j' - t_j) \\ &= t_j - \frac{f(t_{-j}) (1 - F(t_j | t_{-j}))}{f(t_{-j})f(t_j | t_{-j})} (t_j' - t_j) = \varphi(t_j | t_{-j}).\end{aligned}$$

Finally, we show that the flow can be modified to implement Myerson's ironing procedure. The analysis on modifying the flow to reflect ironing is given in [Cai et al. 2016], and we include it here for completeness. Suppose there exist two types $t, t' \in R_j$ such that $t_j < t_j'$ but $\Phi_j(t) > \Phi_j(t')$. We can add a cycle of w units of flow between t and t' , that is, we increase both $\lambda(t, t')$ and $\lambda(t', t)$ by w . Notice that the resulting flow is still valid, and $\Phi(t^*)$ for all $t^* \neq t, t'$ remain unchanged. Moreover, the change does not alter $\Phi_k(t)$ or $\Phi_k(t')$ for all $k \neq j$. The only effect of the change is to increase $\Phi_j(t)$ by $w(t_j' - t_j)/f(t)$, and decrease $\Phi_j(t')$ by $w(t_j' - t_j)/f(t')$. Therefore, we can choose w so that $\Phi_j(t) = \Phi_j(t')$ without changing any other virtual values.

Repeating this process allows us to simulate Myerson's ironing procedure. One technical issue is that we may cut off an ironing interval of $f(t_j | t_{-j})$ because it leaves the region R_j . However, we know that truncating an ironing interval I to $I' \subseteq I$ from below can only decrease the virtual value on I' . This is because the average virtual value on I' is smaller than the average virtual value on I , otherwise we would not iron the entire interval I in the first place. \square

C.3 Upper Bounds for SINGLE, CORE, and TAIL

We decompose the upper bound we had in the previous section into three components. Recall that $r = \text{SREV}(\widehat{V})$. By Lemma C.4, we know that

$$\begin{aligned}\text{REV}(\widehat{V}) &\leq \sum_{t \in T} f(t) \pi(t)^\top \Phi(t) \\ &= \sum_{t \in T} f(t) \sum_{j \in [m]} \pi_j(t) (\widetilde{\varphi}(t_j | t_{-j}) \cdot \mathbb{1}[t \in R_j] + t_j \cdot \mathbb{1}[t \notin R_j]) \\ &\leq \sum_{t \in T} f(t) \sum_{j \in [m]} \pi_j(t) \widetilde{\varphi}(t_j | t_{-j}) \cdot \mathbb{1}[t \in R_j] && \text{(SINGLE)} \\ &\quad + \sum_{t \in T} f(t) \sum_{j \in [m]} \pi_j(t) t_j \cdot \mathbb{1}[t \notin R_j] \cdot \mathbb{1}[t_j \leq 2r] && \text{(CORE)} \\ &\quad + \sum_{t \in T} f(t) \sum_{j \in [m]} \pi_j(t) t_j \cdot \mathbb{1}[t \notin R_j] \cdot \mathbb{1}[t_j > 2r]. && \text{(TAIL)}\end{aligned}$$

We now restate the lemmas that upper bounds each component.

Lemma C.1. $\text{SINGLE} \leq 2\text{BREX}(\widehat{V}) + \text{SREV}(\widehat{V})$.

PROOF. We first recall the expression of SINGLE.

$$\text{SINGLE} = \sum_{t \in T} f(t) \sum_{j \in [m]} \pi_j(t) \tilde{\varphi}(t_j | t_{-j}) \cdot \mathbb{1}[t \in R_j].$$

Because $\pi_j(t) \cdot \mathbb{1}[t \in R_j]$ is between 0 and 1, we can upper bound SINGLE by setting it to 1 whenever the ironed virtual value is positive, and 0 otherwise.

$$\begin{aligned} \text{SINGLE} &\leq \sum_{t \in T} f(t) \sum_{j \in [m]} \tilde{\varphi}(t_j | t_{-j}) \cdot \mathbb{1}[\tilde{\varphi}(t_j | t_{-j}) \geq 0] \\ &= \sum_{j \in [m]} \sum_{t \in T} f(t_{-j}) f(t_j | t_{-j}) \tilde{\varphi}(t_j | t_{-j}) \cdot \mathbb{1}[\tilde{\varphi}(t_j | t_{-j}) \geq 0] \\ &= \sum_{j \in [m]} \sum_{t_{-j} \in T_{-j}} f(t_{-j}) \cdot \text{REV}(t_j | t_{-j}). \end{aligned}$$

Intuitively, we need to show that knowing t_{-j} does not help us sell item j by too much. Observe that \widehat{V} is a correlated distribution obtained from capping an independent distribution. The proof of the lemma crucially relies on the following property of \widehat{V} : revealing t_{-j} gives the same amount of information as revealing only the sum of t_{-j} . For a fixed j , the revenue $\text{REV}(t_j | t_{-j})$ can be captured by the (non-disjoint) union of the following two cases:

- (1) $\|t_{-j}\|_1 < c/2$ and $t_j < c/2$ both hold. Conditioned on this event, the valuation of t_j is independent of t_{-j} . Hence, knowing t_{-j} does not provide additional information.
- (2) $\|t\|_1 \geq c/2$. In this case, buyer's value for the grand bundle is at least $c/2$, so we could charge this to $\text{BREV}(\widehat{V})$.

It is worth noting that t_j and t_{-j} are not independent when $\|t\|_1 < c/2$, so we have to condition on stricter events for them to become independent. Formally, if $\|t_{-j}\|_1 < c/2$,

$$f(t_j | t_{-j}, t_j < c/2) = f(t_j | t_j < c/2), \quad \forall j \in [m], t \in T.$$

We are now ready to bound SINGLE. For single-parameter distributions, the optimal auction simply sets a reserve price. Let p_j^* be the optimal reserve price for the distribution $f(t_j | t_j < c/2)$, and let $p_j(t_{-j})$ be the optimal reserve price for the distribution $f(t_j | t_{-j})$.

$$\begin{aligned} \text{SINGLE} &\leq \sum_{j \in [m]} \sum_{t_{-j} \in T_{-j}} f(t_{-j}) \cdot \text{REV}(t_j | t_{-j}) \\ &= \sum_{j \in [m]} \sum_{t \in T} f(t_{-j}) f(t_j | t_{-j}) p_j(t_{-j}) \cdot \mathbb{1}[t_j \geq p_j(t_{-j})] \\ &\leq \sum_j \sum_{\|t\|_1 \geq \frac{c}{2}} f(t) t_j + \sum_j \sum_{t_j < \frac{c}{2}, \|t_{-j}\|_1 < \frac{c}{2}} f(t) p_j(t_{-j}) \cdot \mathbb{1}[t_j \geq p_j(t_{-j})] \\ &\leq \sum_{\|t\|_1 \geq \frac{c}{2}} f(t) \sum_j t_j + \sum_j \sum_{t_j < \frac{c}{2}, \|t_{-j}\|_1 < \frac{c}{2}} f(t) p_j^* \cdot \mathbb{1}[t_j \geq p_j^*] \\ &\leq c \cdot \Pr_{t \sim \widehat{V}} \left[\|t\|_1 \geq \frac{c}{2} \right] + \sum_{t \in T} f(t) \sum_{j \in [m]} p_j^* \cdot \mathbb{1}[t_j \geq p_j^*] \\ &\leq 2\text{BREV}(\widehat{V}) + \text{SREV}(\widehat{V}). \end{aligned}$$

The last step uses the facts that (1) we can price the grand bundle at price $c/2$ and therefore $\text{BREV}(\widehat{V}) \geq (c/2) \cdot \Pr[\|t\|_1 \geq c/2]$; and (2) the second term is exactly the revenue we can obtain if we post each item j separately at price p_j^* . \square

Recall that $r = \text{SREV}(\widehat{V})$. We continue to upper bound TAIL and CORE.

Lemma C.2. $\text{TAIL} \leq \text{SREV}(\widehat{V})$.

PROOF. Recall that $R_j \subseteq T$ is the subset of buyer types whose favorite item is j .

$$\begin{aligned} \text{TAIL} &= \sum_{t \in T} f(t) \sum_{j \in [m]} \pi_j(t) t_j \cdot \mathbb{1}[t \notin R_j] \cdot \mathbb{1}[t_j > 2r] \\ &\leq \sum_{t \in T} f(t) \sum_{j \in [m]} t_j \cdot \mathbb{1}[t \notin R_j] \cdot \mathbb{1}[t_j > 2r] \\ &= \sum_{j \in [m]} \sum_{t_j > 2r} f(t_j) \sum_{t_{-j} \in T_{-j}} f(t_{-j}|t_j) t_j \cdot \mathbb{1}[t \notin R_j]. \end{aligned}$$

We first show that for any fixed $t_j^* > 2r$, knowing the value for item j to be t_j^* will not increase the probability of $t \notin R_j$ by more than a factor of 2. Intuitively, it is because $t_j > 2r$ is large enough. This is another place where we use the special property of the \widehat{V} we designed.

Recall that \widehat{V}_j is the marginal distribution of \widehat{V} on the j -th coordinate, and $f(t_j)$ is the probability mass function of \widehat{V}_j . The definition $r = \text{SREV}(\widehat{V})$ implies that $\Pr_{t_j \sim \widehat{V}_j}[t_j < 2r] \geq 1/2$, otherwise selling only item j at price $2r$ gives revenue more than r .

$$\begin{aligned} &\sum_{t_{-j} \in T_{-j}} f(t_{-j}|t_j^*) \cdot \mathbb{1}[(t_j^*, t_{-j}) \notin R_j] \\ &\leq 2 \Pr_{t_j \sim \widehat{V}_j}[t_j < 2r] \cdot \sum_{t_{-j} \in T_{-j}} f(t_{-j}|t_j^*) \cdot \mathbb{1}[(t_j^*, t_{-j}) \notin R_j] \\ &= 2 \sum_{t_j < 2r} f(t_j) \sum_{t_{-j} \in T_{-j}} f(t_{-j}|t_j^*) \cdot \mathbb{1}[(t_j^*, t_{-j}) \notin R_j] \\ &\leq 2 \sum_{t_j < 2r} f(t_j) \sum_{t_{-j} \in T_{-j}} f(t_{-j}|t_j) \cdot \mathbb{1}[(t_j, t_{-j}) \notin R_j] \\ &\leq 2 \sum_{t_{-j} \in T_{-j}} f(t_{-j}) \cdot \mathbb{1}[t \notin R_j]. \end{aligned}$$

The last step uses the fact that $f(t_{-j}) = \sum_{t_j} f(t_j) f(t_{-j}|t_j)$. The second last step states the event $t \notin R_j$ (that j is not the largest coordinate) is more likely to happen when the value of t_j is smaller. It uses the monotonicity of $f(t_{-j}|\cdot)$ and $t_j < 2r < t_j^*$. This fact is captured in Lemma D.1 and will be proved in Appendix D.

Lemma D.1. Let $\widehat{V} = V_{(\|v\|_1 \leq c_1)}$ and $\widetilde{V} = V_{(\|v\|_1 \leq c_2)}$ for any $c_1 \leq c_2$. We have $\widehat{V} \leq \widetilde{V}$.

Finally, we relate our upper bound on TAIL to r .

$$\begin{aligned} \text{TAIL} &\leq \sum_{j \in [m]} \sum_{t_j > 2r} f(t_j) \sum_{t_{-j} \in T_{-j}} f(t_{-j}|t_j) t_j \cdot \mathbb{1}[t \notin R_j] \\ &\leq 2 \sum_{j \in [m]} \sum_{t_j > 2r} f(t_j) \sum_{t_{-j} \in T_{-j}} f(t_{-j}) t_j \cdot \mathbb{1}[t \notin R_j] \\ &\leq 2 \sum_{j \in [m]} \sum_{t_j > 2r} f(t_j) \cdot r \leq \text{SREV}(\widehat{V}). \end{aligned}$$

The second last step is because for any j , selling each item separately at the same price t_j^* gives revenue at least $\sum_{t_{-j}} f(t_{-j}) t_j^* \cdot \mathbb{1}[t \notin R_j]$: If some item $k \neq j$ satisfies that $t_k \geq t_j^*$, then the buyer

would purchase at least one of such items. The last step holds because $\sum_{j \in [m]} \sum_{t_j \geq 2r} f(t_j) \cdot 2r$ is exactly the revenue of selling each item at $2r$, so it is at most $\text{SREV}(\widehat{V})$. \square

Lemma C.3. $\text{CORE} \leq 2\text{BREV}(\widehat{V}) + 3\text{SREV}(\widehat{V})$.

PROOF. Recall $r = \text{SREV}(\widehat{V})$. If we sell the grand bundle at price $\text{CORE} - 3r$, we show that the buyer will purchase it with probability at least $5/9$. This implies that $\text{BREV}(\widehat{V}) \geq \frac{5}{9}(\text{CORE} - 3r)$, or equivalently $\text{CORE} \leq \frac{9}{5}\text{BREV}(\widehat{V}) + 3r \leq 2\text{BREV}(\widehat{V}) + 3\text{SREV}(\widehat{V})$.

For the simplicity of presentation, we define a new random variable $s \in \mathbb{R}^m$ as follows: we first draw a sample \widehat{v} from \widehat{V} , and set $s_j = \min(\widehat{v}_j, 2r)$ for all $j \in [m]$.

$$\begin{aligned} \text{CORE} &= \sum_{t \in T} f(t) \sum_{j \in [m]} \pi_j(t) t_j \cdot \mathbb{1}[t \notin R_j] \cdot \mathbb{1}[t_j \leq 2r] \\ &\leq \sum_{t \in T} f(t) \sum_{j \in [m]} t_j \cdot \mathbb{1}[t_j \leq 2r] \\ &\leq \sum_{t \in T} f(t) \sum_{j \in [m]} (t_j \cdot \mathbb{1}[t_j \leq 2r] + 2r \cdot \mathbb{1}[t_j > 2r]) = \mathbb{E}[\|s\|_1]. \end{aligned}$$

Since we price the grand bundle at $\text{CORE} - 3r \leq \mathbb{E}[\|s\|_1] - 3r$, and the buyer's value for the grand bundle is $\|t\|_1$, it is sufficient to prove

$$\Pr_{t \sim \widehat{V}} [\|t\|_1 \geq \mathbb{E}[\|s\|_1] - 3r] \geq \frac{5}{9}.$$

Moreover, because \widehat{V} stochastically dominates s , it is sufficient to prove

$$\Pr[\|s\|_1 \geq \mathbb{E}[\|s\|_1] - 3r] \geq \frac{5}{9}.$$

Intuitively, the condition says that the ℓ_1 -norm of the random variable s concentrates around its expectation. This is the crucial reason why we design our \widehat{V} to be negatively correlated.

In the next claim (Lemma C.5), we are going to prove $\text{Var}[\|s\|_1] \leq 4r^2$. Assume this is true, we conclude the proof by applying the Chebyshev inequality,

$$\Pr[\|s\|_1 < \mathbb{E}[\|s\|_1] - 3r] \leq \frac{\text{Var}[\|s\|_1]}{9r^2} \leq \frac{4}{9}. \quad \square$$

LEMMA C.5. *Let $c \in \mathbb{R}^m$ be the random variable with $s_j = \min(\widehat{v}_j, 2r)$ for all j and $\widehat{v} \sim \widehat{V}$. We have $\text{Var}[\|s\|_1] \leq 4r^2$.*

PROOF. We first show that for any $i \neq j$, s_i and s_j are negatively correlated, i.e.,

$$\text{Cov}(s_i, s_j) = \mathbb{E}[s_i s_j] - \mathbb{E}[s_i] \cdot \mathbb{E}[s_j] \leq 0.$$

Observe that $\widehat{V}_j|_{\widehat{v}_i=x}$ stochastically dominates $\widehat{V}_j|_{\widehat{v}_i=y}$ for any two possible values $x < y$ of \widehat{v}_i . This is because, for any $a \in \mathbb{R}$,

$$\begin{aligned} & \Pr_{\widehat{v} \sim \widehat{V}} \left[\widehat{v}_j \leq a \mid \widehat{v}_i = x \right] \\ &= \Pr_{v \sim V} \left[v_j \leq a \mid v_i = x, \|v\|_1 \leq c \right] \\ &= \Pr_{v_j \sim V_j} \left[v_j \leq a \mid v_j \leq c - x - \sum_{k \neq i, j} v_k \right] \\ &\leq \Pr_{v_j \sim V_j} \left[v_j \leq a \mid v_j \leq c - y - \sum_{k \neq i, j} v_k \right] = \Pr_{\widehat{v} \sim \widehat{V}} \left[\widehat{v}_j \leq a \mid \widehat{v}_i = y \right]. \end{aligned}$$

Because $s_j = \min(\widehat{v}_j, 2r)$, we know that $s_j|_{v_i=x} \geq s_j|_{v_i=y}$ as well. It follows that $\mathbb{E}[s_j|v_i = x] \geq \mathbb{E}[s_j|v_i = y]$. In addition, since $\mathbb{E}[s_j|s_i = x] = \mathbb{E}[s_j|v_i = x]$ when $x < 2r$, and $\mathbb{E}[s_j|s_i = 2r] = \mathbb{E}[s_j|v_i \geq 2r]$, we can deduce that $\mathbb{E}[s_j|s_i]$ weakly decreases as s_i increases. Therefore, we get

$$\mathbb{E}[s_i s_j] = \sum_{s_i} \Pr[s_i] (s_i \cdot \mathbb{E}[s_j|s_i]) \leq \sum_{s_i} \Pr[s_i] (s_i \cdot \mathbb{E}[s_j]) = \mathbb{E}[s_i] \cdot \mathbb{E}[s_j],$$

by an application of (a generalization of) the rearrangement inequality.

Given the negative correlations between the s_i 's, we can upper bound the variance of their sum.

$$\text{Var} \left(\sum_{i \in [m]} s_i \right) = \sum_{i \in [m]} \text{Var}(s_i) + 2 \sum_{1 \leq i < j \leq m} \text{Cov}(s_i, s_j) \leq \sum_{i \in [m]} \text{Var}(s_i).$$

It remains to show $\sum_i \text{Var}(s_i) \leq 4r^2$. This part is identical to the analysis in earlier works [Babaioff et al. 2014; Cai et al. 2016; Li and Yao 2013], but we include it for completeness. Let $r_j \in \mathbb{R}$ denote the maximum revenue one can extract by selling item j alone. Note that $\text{SREV}(\widehat{V}) = r = \sum_j r_j$, so it is sufficient to show $\text{Var}(s_j) \leq 4rr_j$ for all j .

Fix some $j \in [m]$. We use $x = s_j$ as an alias for the random variable s_j . We know that

- (1) x is at most $2r$, and
- (2) the revenue of x is at most r_j : $(a \cdot \Pr[x \geq a]) \leq r_j$ for any $a \in \mathbb{R}$.

Combining these two facts gives the required bound on the variance of x . Let $0 < a_1 < \dots < a_\ell \leq 2r$ be the support of x , and let $a_0 = 0$.

$$\begin{aligned} \mathbb{E}[x^2] &= \sum_{k=1}^{\ell} \Pr[x = a_k] \cdot a_k^2 \\ &= \sum_{k=1}^{\ell} (a_k^2 - a_{k-1}^2) \cdot \Pr[x \geq a_k] \\ &< \sum_{k=1}^{\ell} 2(a_k - a_{k-1}) \left(a_k \cdot \Pr[x \geq a_k] \right) \\ &\leq r_j \sum_{k=1}^{\ell} 2(a_k - a_{k-1}) \\ &= 2r_j a_\ell \leq 4rr_j. \end{aligned} \quad \square$$

D PROOF OF LEMMA D.2 AND D.1

In this section we prove Lemma D.2 and D.1.

LEMMA D.1. $V_{|\|\cdot\|_1 \leq c_1} \leq V_{|\|\cdot\|_1 \leq c_2}$ for any $c_1 \leq c_2$.

PROOF. We are going to modify any \bar{v} dimension by dimension to reach \hat{v} . Define n random functions $\sigma_1, \sigma_2, \dots, \sigma_n$, where $\sigma_i(u)$ and u only differs in u_i . In other words, σ_i only modifies the i -th dimension of its input. Also define $\tau_i = \sigma_i \circ \sigma_{i-1} \circ \dots \circ \sigma_1$. Note that $\tau_n(v) \leq v$ as long as $\sigma_i(\tau_{i-1}(v)) \leq \tau_{i-1}(v)$ for all i .

In the first step, select a σ_1 such that

$$\Pr_{v \sim \bar{V}} [(\sigma_1(v))_1] = \Pr_{v \sim \bar{V}} \left[v_1 \left| \sum_{i=1}^n v_i \leq c_1 \right. \right].$$

Then we separately deal with v_2 's according to their $(\tau_1(v))_1$ value. Select a σ_2 such that

$$\Pr_{v \sim \bar{V}} [(\sigma_2(v))_2 \mid (\tau_1(v))_1] = \Pr_{v \sim \bar{V}} \left[v_2 \left| (\tau_1(v))_1, (\tau_1(v))_1 + \sum_{i=2}^n v_i \leq c_1 \right. \right].$$

We continue this procedure till we get σ_n . The k -th step will be setting

$$\begin{aligned} & \Pr_{v \sim \bar{V}} [(\sigma_k(v))_k \mid (\tau_1(v))_1, \dots, (\tau_{k-1}(v))_{k-1}] \\ &= \Pr_{v \sim \bar{V}} \left[v_k \left| (\tau_1(v))_1, \dots, (\tau_{k-1}(v))_{k-1}, \sum_{i=1}^{k-1} (\tau_i(v))_i + \sum_{i=k}^n v_i \leq c_1 \right. \right]. \end{aligned}$$

Next we show for all k , there exists σ_k satisfying $(\sigma_k(x))_k \leq x_k$ for all x . This is simply first-order stochastic dominance for one-dimensional distributions, and it is equivalent to the following condition:

$$\begin{aligned} & \Pr_{v \sim \bar{V}} [v_k \leq a \mid (\tau_1(v))_1, \dots, (\tau_{k-1}(v))_{k-1}] \\ & \leq \Pr_{v \sim \bar{V}} \left[v_k \leq a \left| (\tau_1(v))_1, \dots, (\tau_{k-1}(v))_{k-1}, \sum_{i=1}^{k-1} (\tau_i(v))_i + \sum_{i=k}^n v_i \leq c_1 \right. \right], \quad \forall a \in \mathbb{R}. \end{aligned}$$

Writing $p_k = (v_{k+1}, \dots, v_n)$, $q_k = ((\tau_1(v))_1, \dots, (\tau_{k-1}(v))_{k-1})$, and $r_k = (v_1, \dots, v_{k-1})$, the inequality above is true because

$$\begin{aligned} & \Pr_{v \sim \bar{V}} [v_k \leq a \mid q_k, v_k + \|p_k\|_1 + \|q_k\|_1 \leq c_1] \\ &= \sum_{p_k, r_k} \Pr_{v \sim \bar{V}} [p_k, r_k] \cdot \Pr_{v \sim \bar{V}} [v_k \leq a \mid q_k, p_k, r_k, v_k \leq \min(c_1 - \|p_k\|_1 - \|q_k\|_1, c_2 - \|p_k\|_1 - \|r_k\|_1)] \\ &\geq \sum_{p_k, r_k} \Pr_{v \sim \bar{V}} [p_k, r_k] \cdot \Pr_{v \sim \bar{V}} [v_k \leq a \mid q_k, p_k, r_k, v_k \leq \min(c_2 - \|p_k\|_1 - \|q_k\|_1, c_2 - \|p_k\|_1 - \|r_k\|_1)] \\ &= \Pr_{v \sim \bar{V}} [v_k \leq a \mid q_k]. \end{aligned}$$

Therefore τ_n is the random function that defines coordinate-wise stochastic dominance, as every σ_i satisfies $\sigma_i(x) \leq x$ for all x . \square

LEMMA D.2. $V_{|\|\cdot\|_1 \leq c} \leq V$ for any $c > 0$.

PROOF. It is implied by Lemma D.1 when $c_1 = c$ and $c_2 = \max_{v \in \text{supp}(V)} \|v\|_1$. \square