Simple Mechanisms for a Subadditive Buyer and Applications to Revenue Monotonicity

AVIAD RUBINSTEIN[1], UC Berkeley
S. MATTHEW WEINBERG[2], Princeton University

We study the revenue maximization problem of a seller with n heterogeneous items for sale to a single buyer whose valuation function for sets of items is unknown and drawn from some distribution D. We show that if D is a distribution over subadditive valuations with independent items, then the better of pricing each item separately or pricing only the grand bundle achieves a constant-factor approximation to the revenue of the optimal mechanism. This includes buyers who are k-demand, additive up to a matroid constraint, or additive up to constraints of any downwards-closed set system (and whose values for the individual items are sampled independently), as well as buyers who are fractionally subadditive with item multipliers drawn independently. Our proof makes use of the core-tail decomposition framework developed in prior work showing similar results for the significantly simpler class of additive buyers [Li and Yao 2013; Babaioff et al. 2014].

In the second part of the paper, we develop a connection between approximately optimal simple mechanisms and approximate revenue monotonicity with respect to buyers' valuations. Revenue non-monotonicity is the phenomenon that sometimes strictly *increasing* buyers' values for every set can strictly *decrease* the revenue of the optimal mechanism [Hart and Reny 2012]. Using our main result, we derive a bound on how bad this degradation can be (and dub such a bound a proof of *approximate* revenue monotonicity); we further show that better bounds on approximate monotonicity imply a better analysis of our simple mechanisms.

Categories and Subject Descriptors: Theory of computation [**Algorithmic game theory and mechanism design**]: Computational pricing and auctions

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Revenue optimization, combinatorial valuations, simple auctions, revenue monotonicity

1. INTRODUCTION

Consider a revenue-maximizing seller with n heterogeneous items for sale to a single buyer whose value for sets of items is unknown, but drawn from a known distribution D. When $n = 1$, seminal work of Myerson [Myerson 1981] and Riley and Zeckhauser [Riley and Zeckhauser 1983] shows that the optimal selling scheme simply sets the price $p^* = \arg\max\{p \cdot \Pr[v \geq p | v \leftarrow D]\}$. Thirty years later, understanding the structure of the optimal mechanism when $n > 1$ still remains a central open problem. Unfortunately, it is well-known that the optimal mechanism may require randomization, behave non-monotonically, and be computationally hard to find, even in very simple instances [Thanassoulis 2004; Pavlov 2011; Briest et al. 2010; Daskalakis et al. 2014; Chen et al. 2014; Hart and Nisan 2013; Hart and Reny 2012]. In light of this,

[1] Department of Electrical Engineering and Computer Science, UC Berkeley, aviad@eecs.berkeley.edu. This research was supported by NSF grants CCF0964033 and CCF1408635, and by Templeton Foundation grant 3966. This work was done in part at the Simons Institute for the Theory of Computing and while visiting Princeton University.
[2] Department of Computer Science, Princeton University, sethmw@cs.princeton.edu.

EC'15, June 15–19, 2015, Portland, OR, USA. Copyright © 2015 ACM 978-1-4503-3410-5/15/06 ...$15.00.
http://dx.doi.org/10.1145/http://dx.doi.org/10.1145/2764468.2764510

recent work began studying the performance of especially simple auctions through the lens of approximation. Remarkably, these works have shown that when the bidder's valuation is *additive*[1], and her value for each item is drawn independently, very simple mechanisms can achieve quite good approximation ratios. Specifically, techniques developed in this series of works proves that the better of setting Myerson's reserve on each item separately or setting Myerson's reserve on the grand bundle of all items together achieves a 6-approximation [Hart and Nisan 2012; Li and Yao 2013; Babaioff et al. 2014].

While this model of buyer values is certainly mathematically interesting and economically motivated, it is also perhaps too simplistic to have broad real-world applications. A central question left open by these works is whether or not simple mechanisms can still approximate optimal ones in more general settings. In this work we resolve this question in the affirmative, showing that the better of selling separately (we will henceforth use SREV to denote the revenue of the optimal such mechanism) or together (henceforth BREV) still obtains a constant-factor approximation to the optimal revenue (henceforth REV) when buyer values are combinatorial in nature but complement-free.

INFORMAL THEOREM 1. *Let D be any distribution over subadditive valuation functions with independent items. Then* $\max\{\text{SREV}, \text{BREV}\} \geq \Omega(1) \cdot \text{REV}$. *Furthermore, prices providing this guarantee can be found computationally efficiently.*

We postpone a formal definition of exactly what it means for D to have "independent items" to Section 2. We note here a few instantiations of our model in commonly studied settings (from least to most general):

— k-demand: The buyer has value v_i for item i, and the v_is are drawn independently. The buyer's value for a set S is $v(S) = \max_{T \subseteq S, |T| \leq k}\{\sum_{i \in T} v_i\}$.
— Additive up to constraints $\mathcal{I}$: $\mathcal{I}$ is some downwards-closed set system on $[n]$. The buyer has value v_i for item i, and the v_is are drawn independently. $v(S) = \max_{T \subseteq S, T \in \mathcal{I}}\{\sum_{i \in T} v_i\}$.
— Fractionally-subadditive: buyer has "possible values" $\{v_{ij}\}_j$ for item i, and the sets $\{v_{ij}\}_j$ are drawn independently across items (but may be correlated within an item). $v(S) = \max_j\{\sum_{i \in S} v_{ij}\}$.

A recent book of Hartline [Hartline 2011] provides a fantastic discussion of the role of approximation in mechanism design. Before proceeding, it is worth repeating some aspects of this discussion to view our result in the proper context. One should not interpret our main result as claiming that sellers should be satisfied with a constant fraction of the optimal obtainable revenue, but rather as studying the tradeoff between simplicity and optimality. Sometimes, the optimal mechanism simply isn't an option: perhaps it is prohibitively complex to implement, prohibitively frustrating for buyers to participate, or prohibitively difficult (computationally) to find. And even when the optimal mechanism is a feasible option, the desire for simplicity and transparency may outweigh the expected loss in revenue. Similarly, one should not interpret the ratios obtained in our main result (they are noticeably larger than 6) as ratios that one might expect to trade off in practice, as these are *provable* bounds for *worst-case* instances.

1.1. Challenges of Combinatorial Valuations

The design of simple, approximately optimal mechanisms for any non-trivial multi-item setting has been a large focus for much of the Algorithmic Game Theory com-

[1]A valuation function $v(\cdot)$ is additive if $v(S \cup T) = v(S) + v(T)$ for all $S \cap T = \emptyset$.

munity over the past decade. Even "simple" settings with additive or unit-demand valuations required significant breakthroughs. The key insight enabling these breakthroughs for additive buyers is that the buyer's valuation is separable across items. While the optimal mechanism can still be quite bizarre despite this realization [Hart and Reny 2012], this fact enables certain elementary decomposition theorems that are surprisingly powerful (e.g. the "Marginal Mechanism" [Cai and Huang 2013; Hart and Nisan 2012]). However, these theorems are extremely sensitive to being able to separate the marginal contribution of different items *exactly* (and not just via upper/lower bounds). This is due to the phenomenon that a slight miscalculation in estimating a buyer's value may cause her to change preferences entirely, resulting in a potentially unbounded loss of revenue. One of our main technical contributions is overcoming this obstacle by providing an *approximate* version of these decomposition theorems.

A further complication in applying these previous techniques is that they all make use of the fact that $\text{SREV}(D_1 \times \ldots \times D_n) = \sum_i \text{SREV}(D_i)$. This claim is not even approximately true for subadditive buyers, and the ratio between the two values could be as large as n (the right-hand side is always larger). To have any hope of applying these tools, we therefore need a proxy for SREV that at least approximately has this separability property.

For unit-demand buyers, the key insight behind the mechanisms designed in [Chawla et al. 2007, 2010a,b; Kleinberg and Weinberg 2012] is that every multi-dimensional problem instance has a related single-dimensional problem instance, and there is a correspondence between truthful mechanisms in the two instances. This realization means that one can instead design mechanisms for the single-dimensional setting, where optimal mechanisms are well understood due to Myerson's virtual values, and translate them in a black-box manner to mechanisms for the original instance. While these techniques have proven extremely fruitful in the design of mechanisms for multiple unit-demand buyers and sophisticated feasibility constraints, they have also proven to be limited in use to unit-demand settings. A special case of our results can be seen as providing an alternative proof of the single-buyer result of Chawla, Hartline, and Kleinberg [Chawla et al. 2007] (albeit with a significantly worse constant) that doesn't require virtual valuation machinery.

Aside from the difficulties in applying existing machinery to design optimal mechanisms for combinatorial valuations, formal barriers exist as well. For instance, it is a trivial procedure for an additive buyer to select his utility-maximizing set of items when facing an item-pricing, and finding the revenue-optimal item-pricing is also trivial (just find the optimal price for each item separately). Yet for a subadditive buyer, both tasks are quite non-trivial. Just *computing* the expected revenue obtained by a fixed item-pricing is NP-hard. Worse, the buyer's problem of just *selecting* her utility-maximizing set from a given item-pricing is also NP-hard! Therefore, buyers may behave quite unpredictably in the face of an item-pricing depending on how well they can optimize. Moreover, even if we are willing to assume that the buyer has the computational power to select her utility-maximizing set, it is known *still* that (without our independence assumption) finding an n^c-approximately optimal mechanism is NP-hard for all $c = O(1)$ [Cai et al. 2013]. We sidestep all these difficulties by not attempting to compute or approximate SREV at all, nor trying to predict bizarre buyer behavior. We instead perform our analysis on revenue contributions only of items purchased *when the buyer is not willing to purchase any others*. Buyer behavior in such instances is predictable and easily computable: simply purchase the unique item for which $v(\{i\}) > p_i$. It is surprising that such an analysis suffices, as it completely ignores any revenue contribution coming from the entirely plausible event that the buyer is willing to purchase multiple items.

1.2. Techniques

We prove our main theorem by making use of the core-tail decomposition framework introduced by Li and Yao [Li and Yao 2013]. There are three high-level steps to applying the framework. The first is proving a "core decomposition" lemma that separates the optimal revenue into contributions from items which the buyer values very highly (the "tail"), and items which the buyer values not so high (the "core"). The second is showing that the contribution from the tail can be approximated well by $\mathrm{SR}\textsc{ev}$. The third is showing that the contribution from the core can be approximated well by $\max\{\mathrm{SR}\textsc{ev}, \mathrm{BR}\textsc{ev}\}$.

The Core Decomposition Lemma. The proof of the original Core Decomposition Lemma in [Li and Yao 2013] was obtained by cleverly stringing together simple claims proved in [Hart and Nisan 2012]. As discussed above, these seemingly "obvious" claims may not extend beyond additive valuations over independent items, due to the fact that the buyer's value cannot be separated across items. Nevertheless, we are able to prove an approximate version of the core decomposition lemma for subadditive buyers (Lemma 3.6) by making use of ideas from reductions from ϵ-truthful mechanisms to fully truthful ones. Like in [Babaioff et al. 2014], our core decomposition lemma holds for many buyers. The proof for a single buyer, which is the focus of this paper can be found in Section 3.1. In the full version we also provide a more technically involved proof for many buyers which builds on heavier tools from [Bei and Huang 2011; Hartline et al. 2011; Daskalakis and Weinberg 2012].

Bounding the Tail's Contribution. Arguments for bounding the contribution from the tail in prior work (and ours) use the following reasoning. If the cutoff between core and tail is sufficiently high, then the probability that k items are simultaneously in the tail for a sampled valuation decays exponentially in k. If one can also show that the approximation guarantee of $\mathrm{SR}\textsc{ev}$ decays subexponentially in k, then we can bound the gap between $\mathrm{SR}\textsc{ev}$ and the tail's contribution by a constant factor. We show that indeed the approximation guarantee of $\mathrm{SR}\textsc{ev}$ decays only polynomially in k.

Bounding the Core's Contribution. Arguments for bounding the contribution from the core in prior work (and ours) use the following reasoning. The total expected value for items in the core is a subadditive function of independent random variables (bounded above by the core-tail cutoff). If the cutoff between core and tail is sufficiently low, then one of two things must happen. Either the expected contribution from the core is also small, in which case $\mathrm{SR}\textsc{ev}$ itself provides a good approximation, or the expected contribution is large, and therefore also large with respect to the cutoffs. In the latter case, a concentration bound implies that $\mathrm{BR}\textsc{ev}$ must provide a good approximation. In the additive case, the appropriate concentration bound is Chebyshev's inequality. In the subadditive case, we need heavier tools, and apply a concentration bound due to Schechtman [Schechtman 1999].

1.3. Connection to Approximate Revenue-Monotonicity

Consider designing revenue-optimal mechanisms for two different markets, and suppose that the valuations of the consumers in the first market first-order stochastically dominate[2] the valuations of the consumers in the second market. It then seems reasonable to expect that the optimal revenue achieved from the first market, $\mathrm{R}\textsc{ev}(D^+)$, should be at least as large as the revenue achieved from the dominated market,

[2]We say that a distribution D^+ over valuation functions v^+ first-order stochastically dominates distribution D over valuation functions v if the probability spaces can be coupled so that for every subset S, $v^+(S) \geq v(S)$.

$\text{REV}(D)$. When there is just a single item for sale, this is an easy corollary of the format for Myerson's optimal auction. Yet Hart and Reny provided an example where this intuition breaks even in a setting as simple as an *additive buyer with i.i.d. values for two items* [Hart and Reny 2012]. Surprisingly, their example shows that it is possible to make *strictly* more revenue in a market when buyers have *strictly* less value for your goods, and the market need not even be very complex for this phenomenon to occur.

A natural question to ask then, is how large this anomaly can be. For example, Hart and Reny's constructions exhibit a (multiplicative) gap of $33/32$ between $\text{REV}(D^+)$ and $\text{REV}(D)$ for an additive buyer with correlated values for two items, and $(1 + \frac{1}{7000000})$ for an additive buyer with i.i.d. values for two items. Interestingly, the simple mechanisms of [Hart and Nisan 2012; Li and Yao 2013; Babaioff et al. 2014] upper bound the possible gap of any instance where their results apply, since SREV and BREV are monotone for additive buyers (i.e. $\text{SREV}(D^+) \geq \text{SREV}(D)$ and $\text{BREV}(D^+) \geq \text{BREV}(D)$). Specifically, for an additive buyer the gap is at most $(1 + 1/e)$ for two i.i.d. items, 2 for two asymmetric independent items, and 6 for any number of independent items. In Section 4 we show that as a corollary of our results, the gap is also constant for a subadditive buyer with independent items. Interestingly, this connection between approximately optimal simple mechanisms and approximate revenue-monotonicity is also fruitful in the other direction: it turns out that improving the bound on approximate monotonicity for a subadditive buyer would also improve the constant in our main theorem. Finally, we show in Section 4.3 that for an additive buyer with correlated values for two items, the gap is potentially infinite. (This is the case for which Hart and Reny provide a gap of $33/32$.) The proof is by a black-box reduction from an example due to Hart and Nisan [Hart and Nisan 2013] that exhibits a similar gap between simple and optimal mechanisms, further demonstrating the connection between these two important research directions.

1.4. Discussion and Open Problems

Our work contributes to the recent growing literature on simple, approximately optimal mechanisms. We extend greatly beyond prior work, providing the first simple and approximately optimal mechanisms for buyers with combinatorial valuations. Prior to our work, virtually nothing was known about this setting (modulo the impossibility result of [Cai et al. 2013]). Our results also demonstrate the strength of the core-tail decomposition framework developed by Li and Yao to go beyond additive buyers. We suspect that this framework will continue to prove useful in other Bayesian mechanism design problems.

In our opinion, the most exciting open question in this area is extending these results to multiple buyers. A beautiful lookahead reduction was recently developed by Yao [Yao 2015] for additive buyers. Still, generalizing his tools beyond additive buyers seems quite challenging and is a very intriguing direction. Another important direction is extending our understanding of simple mechanisms to models of limited correlation over values for disjoint sets of items[3]. Recent independent work of Bateni et. al. [Bateni et al. 2015] addresses this direction, providing approximation guarantees on $\max\{\text{SREV}, \text{BREV}\}$ vs. REV for an additive buyer whose values for items are drawn from a common-base-value distribution and various extensions. Their results also make use of a core-tail decomposition, but the tools they develop beyond the decomposition are disjoint from ours. A natural question in this direction is whether our results extend to settings where buyer values are both combinatorial *and* exhibit lim-

[3]Note that as we mention in the previous section, for arbitrary correlated items the gap can be infinite [Hart and Nisan 2013; Briest et al. 2010].

ited correlation between disjoint sets of items, as the end goal is to have a model that encompasses as many real-world instances as possible.

2. PRELIMINARIES

We focus the body of the exposition on the single-buyer problem, and defer all details regarding auctions for multiple buyers to the full version. There is a single revenue-maximizing seller with n items for sale to a single buyer. The buyer has combinatorial valuations for the items (i.e. value $v(S)$ for receiving set S), and is quasi-linear and risk-neutral. That is, the buyer's utility for a randomized outcome that awards him set S with probability $A(S)$ while paying (expected) price p is $\sum_S A(S)\, v(S) - p$. The valuation $v(\cdot)$ is unknown to the seller, who has a prior D over possible buyer valuations that is *subadditive over independent items*, a term we describe below. By the taxation principle, the seller may restrict attention to only lottery systems. In other words, the seller provides a list of potential lotteries (distributions over sets) each with a price, and the buyer chooses the utility-maximizing option.

2.1. Subadditive valuations over independent items

We now carefully define what we mean by subadditive valuations over independent items. Intuitively, our model is such that the buyer has some private information x_i pertaining to item i,[4] and $D^{\vec{x}}$ is a product distribution over $\mathbb{R}^n$ representing the seller's prior over the private information possessed by the buyer. The buyer's valuation for set S is parametrized by the private information she has about items in that set, and can be written as $V(\langle x_i\rangle_{i\in S}, S)$. In economic terms, this models that the items not received by the buyer impose no externalities. We capture this formally in the definition below:

Definition 2.1. We say that a distribution D over valuation functions $v(\cdot) : \{0,1\}^n \to \mathbb{R}$ is *subadditive over independent items* if:

(1) All $v(\cdot)$ in the support of D exhibit *no externalities*.
Formally, let $\Omega_S = \bigtimes_{i\in S} \Omega_i$, where each Ω_i is a compact subset of a normed space. There exists a distribution $D^{\vec{x}}$ over $\Omega_{[n]}$ and functions $V_S : \Omega_S \to \mathbb{R}$ such that D is the distribution that first samples $\vec{x} \leftarrow D^{\vec{x}}$ and outputs the valuation function $v(\cdot)$ with $v(S) = V_S(\langle x_i\rangle_{i\in S})$ for all S.
(2) All $v(\cdot)$ in the support of D are *monotone*. That is, $v(S) \le v(S \cup T)$ for all S, T.
(3) All $v(\cdot)$ in the support of D are *subadditive*. That is, $v(S \cup T) \le v(S) + v(T)$.
(4) The private information is *independent across items*. That is, the $D^{\vec{x}}$ guaranteed in property 1 is a product distribution.

We describe now how to encode commonly studied valuation distributions in this model.

Example 2.2. The following types of distributions can be modeled as subadditive over independent items. (Recall that $\vec{x}$ is the vector of independently sampled attributes in the definition above.)

(1) Additive: Let $\Omega_i = [0,1]$ and interpret x_i as the buyer's value for item i. $V_S(\langle x_i\rangle_{i\in S}) = \sum_{i\in S} x_i$.
(2) k-demand: Let $\Omega_i = [0,1]$ and interpret x_i as the buyer's value for item i. $V_S(\langle x_i\rangle_{i\in S}) = \max_{T\subseteq S, |T|\le k}\{\sum_{i\in T} x_i\}$.
(3) Additive up to $\mathcal{I}$: Let $\Omega_i = [0,1]$ and interpret x_i as the buyer's value for item i. $V_S(\langle x_i\rangle_{i\in S}) = \max_{T\subseteq S, T\in\mathcal{I}}\{\sum_{i\in T} x_i\}$.

[4]Think of this information as "information about the buyer's preferences related to item."

(4) Fractionally subadditive: Let $\Omega_i = [0,1]^k$ for any finite k and interpret x_i as encoding the values $\{v_{ij}\}_{j\in[k]}$. $V_S(\langle x_i\rangle_{i\in S}) = \max_j\{\sum_{i\in S} v_{ij}\}$.

2.2. Notation

Definition 2.3. For any distribution D of buyer's valuation, we use the following notation, most of which is due to [Hart and Nisan 2012; Babaioff et al. 2014]:

— D_i: The distribution of $v(\{i\})$ when $v(\cdot) \leftarrow D$.
— t: the cutoff between core and tail. If $v(\{i\}) > t$, we say that item i is in the tail. Otherwise it is in the core.
— D_A: the distribution D, conditioned on A being exactly the set of items in the tail.
— D_A^T: the distribution D_A restricted just to items in the tail (i.e. A).
— D_A^C: the distribution D_A restricted just to items in the core (i.e. $\bar{A}$).
— p_i: the probability that element i is in the tail.
— p_A: the probability that A is exactly the set of items in the tail (note that $p_{\{i\}} \neq p_i$).
— $\text{REV}\,(D)$: The maximum revenue obtainable via a truthful mechanism from a buyer with valuation profile drawn from D.
— $\text{BREV}\,(D)$: The revenue obtainable from a buyer with valuation profile drawn from D by auctioning the grand bundle via Myerson's optimal auction.
— $\text{SREV}\,(D)$: The maximum revenue obtainable from a buyer with valuation profile drawn from D by pricing each item separately. Note that when the buyer is not additive, $\text{SREV}(D)$ behaves erratically and is NP-hard to find [Chen et al. 2014].
— $\text{REV}_q(D)$: For a one-dimensional distribution D, the optimal revenue obtained by a reserve price that sells with probability at most q.
— $\text{SREV}^*_{\vec{q}}(D) : \prod_i(1-q_i)\cdot\sum_i \text{REV}_{q_i}(D_i)$: a proxy for $\text{SREV}(D)$ that behaves nicer and is easy to compute.
— $\text{VAL}\,(D)$: the buyer's expected valuation for the grand bundle, $\mathsf{E}_{v\leftarrow D}\left[v\left([n]\right)\right]$.

When the distribution is clear from the context, we simply use REV, VAL, etc. Most of this notation is standard following [Hart and Nisan 2012], with the exception of REV_q and $\text{SREV}^*_{\vec{q}}$. We introduce $\text{SREV}^*_{\vec{q}}$ because it will serve as a proxy to SREV that behaves nicely and is easy to compute. Note that $\text{SREV}^*_{\vec{q}}$ is essentially computing the revenue of the best item pricing that sells item i with probability at most q_i, but only counting revenue from cases where the other values are too low to have possibly sold (and actually it undercounts this quantity).

Remark 2.4. In our definitions of $\text{REV}_q(D)$ and $\text{SREV}^*_{\vec{q}}(D)$ we assume without loss of generality that for every single-dimensional D and $q \in [0,1]$ it is possible to set a price that sells with probability exactly q. When D is a continuous distribution, this is true by the intermediate value theorem. When D has a point mass, this is no longer true per se. Fortunately, there are standard methods for reducing the study of arbitrary distributions to continuous ones with arbitrarily small loss. We briefly sketch one, a rounding scheme commonly attributed to Nisan (that appears also in [Chawla et al. 2007; Cai and Huang 2013]):

For any $\epsilon > 0$, D can be "smoothed" into a continuous distribution D^ϵ by multiplying samples from D by a random factor drawn uniformly from $[1, 1+\epsilon]$. For any smoothed D^ϵ, the desired prices exist by the intermediate value theorem. Using techniques similar in spirit to those of Section 3.1, it is easy (both computationally and conceptually) to convert mechanisms for D^ϵ to mechanisms for D, and vice versa, with negligible (dependent on ϵ) loss in revenue. Therefore, one may formally study D^ϵ for sufficiently small ϵ, and all results hold with respect to D as well with negligible loss (and taking

$\epsilon \to 0$ results in no loss at all). So for the remainder of the paper, we will assume w.l.o.g. that all distributions are continuous, and therefore the desired prices exist.

We conclude the preliminaries by stating a lemma of Hart and Nisan that we will use. We include the proof below for completeness, as well as to verify that it continues to hold when the valuations are not additive.

LEMMA 2.5 (SUB-DOMAIN STITCHING SPECIAL CASE [HART AND NISAN 2012]). $\textsc{Rev}(D) \leq \sum_A p_A \textsc{Rev}(D_A)$.

PROOF. Let M be an optimal mechanism for selling items with valuations sampled from D, and let $\textsc{Rev}_M(D) = \textsc{Rev}(D)$ denote its revenue. Then, $\textsc{Rev}_M(D) = \sum p_A \textsc{Rev}_M(D_A)$. Also, for each $A \subseteq [n]$, $\textsc{Rev}_M(D_A) \leq \textsc{Rev}(D_A)$. □

3. MAIN RESULT: CONSTANT-FACTOR APPROXIMATION FOR SUBADDITIVE BUYER

THEOREM 3.1. *When D is subadditive over independent items, there exists a probability vector $\vec{q}$ such that:*

$$\textsc{Rev}(D) \leq 314\textsc{SRev}^*_{\vec{q}}(D) + 24\textsc{BRev}(D)$$

Furthermore, $\vec{q}$ can be computed efficiently, as well as an induced item pricing that yields expected revenue at least $\textsc{SRev}^*_{\vec{q}}(D)$.

Proof outline. We follow the core-tail decomposition framework. First, we provide an approximate core decomposition lemma in Section 3.1. Then, we provide a bound on the contribution of the core with respect to $\max\{\textsc{BRev}, \textsc{SRev}^*\}$ in Section 3.2, and a bound on the contribution of the tail with respect to $\textsc{SRev}^*$ as a function of the cutoffs chosen in Section 3.3.

For ease of exposition, we simply set t so that the probability of having an empty tail is exactly half; i.e. $p_\emptyset = \prod (1 - p_i) = 1/2$. We also set $\vec{q} = \vec{p}$.

3.1. Approximate Core Decomposition

In this section, we prove our approximate core decomposition lemma. The key ingredient will be an approximate version of the "Marginal Mechanism" lemma from [Cai and Huang 2013; Hart and Nisan 2012] for subadditive buyers, stated below:

LEMMA 3.2. *("Approximate Marginal Mechanism")*
Let S, X be a partition of $[n]$, and let $D = (D^S, D^X)$ be the joint distribution for the valuations of items in S, X, respectively, for buyers with subadditive valuations. Then for any $0 < \epsilon < 1$,

$$\textsc{Rev}(D) \leq \left(\frac{1}{\epsilon} + \frac{1}{1-\epsilon}\right) \textsc{Val}\left(D^S\right) + \frac{1}{1-\epsilon} \mathsf{E}_{v^S \leftarrow D^S}\left[\textsc{Rev}\left(D^X \mid v^S\right)\right] \tag{1}$$

When D^S and D^X are independent, this simplifies to

$$\textsc{Rev}(D) \leq (\frac{1}{\epsilon} + \frac{1}{1-\epsilon})\textsc{Val}(D^S) + \frac{1}{1-\epsilon}\textsc{Rev}(D^X).$$

We outline the proof of Lemma 3.2 below. We first recall the original "Marginal Mechanism" lemma (that holds for an additive buyer without any multipliers). We provide a complete proof so that the reader can see where the argument fails for subadditive buyers.

LEMMA 3.3. ("Marginal Mechanism" [Cai and Huang 2013; Hart and Nisan 2012]) *Let $S \sqcup X$ be any partition of $[n]$, and let D^+ be any distribution over valuation functions such that $v^+(U) = v^+(U \cap S) + v^+(U \cap X)$ for all $U \subseteq n$, and v^+ in the support of D^+.*

Let also D^S denote D^+ restricted to items in S and D^X denote D^+ restricted to items in X. Then $\text{Rev}(D^+) \leq \text{Val}(D^S) + \mathsf{E}_{v^S \leftarrow D^S}[\text{Rev}(D^X|v^S)]$.

PROOF. We design a mechanism M^X (the "Marginal Mechanism") to sell items in X to consumers sampled from $D^X|v^S$ based on the optimal mechanism M for selling items in $S \sqcup X$ to consumers sampled from D^+. Define $A(v)$ to be the (possibly random) allocation of items awarded to type v in M, and $p(v)$ to be the price paid. Let M^X first sample a value $v^S \leftarrow D^S$. The buyer is then invited to report any type v^X, and M^X will award him the items in $A(v^S, v^X) \cap X$ and charge him price $p(v^S, v^X) - v^S(A(v^S, v^X) \cap S)$. In other words, the buyer will receive value from exactly the same items in M^X as he would have received in M, except he receives the actual items in X, whereas for items in S he is given a monetary rebate instead of his actual value.

We first claim that if M is truthful, then so is M^X. The utility of a buyer with type v^X for reporting w^X to M^X can be written as: $v^X(A(v^S, w^X) \cap X) + v^S(A(v^S, w^X) \cap S) - p(v^S, w^X) = (v^S, v^X)(A(v^S, w^X)) - p(v^S, w^X)$, which is exactly the utility of a buyer with type (v^S, v^X) for reporting (v^S, w^X) to M. As M was truthful, we know that a buyer with type (v^S, v^X) maximizes utility when reporting (v^S, v^X) over all possible (v^S, w^X). Therefore, a buyer with type v^X prefers to tell the truth as well.

Finally, we just have to compute the revenue of M^X. For each v^S, the marginal mechanism provides a concrete mechanism for the distribution $D^X|v^S$ that attains revenue at least $\text{Rev}(D^+|v^S) - v^S(S)$. So $\text{Rev}(D^X|v^S) \geq \text{Rev}(D^+|v^S) - v^S(S)$. Taking an expectation over all v^S (and an application of sub-domain stitching) yields the lemma. $\square$

Notice that it is *crucial* in the proof above that the buyer's value could be written as $v^S(\cdot) + v^X(\cdot)$. Otherwise the auctioneer does not know how much to "reimburse" the buyer, since the correct amount depends on the buyer's private information. The buyer can then manipulate his own report v^X to influence how much he gets reimbursed for the items in S.

A natural approach then, given any distribution D over subadditive valuations, is to define a new value distribution D^+ by redefining all $v(\cdot)$ to satisfy $v(U) = v(U \cap S) + v(U \cap X)$ (it is easy to see that all valuations in the support of D^+ are still subadditive). Unfortunately, even though D^+ (first-order stochastically) dominates D, due to non-monotonicity we could very well have $\text{Rev}(D^+) < \text{Rev}(D)$. Still, we bound the revenue lost as we move from D to D^+ by making use of tools for turning ϵ-truthful mechanisms into truly truthful ones. Lemma 3.4 and Corollary 3.5 below capture this formally.

LEMMA 3.4.
Consider two coupled distributions D and D^+, with $v(\cdot)$ and $v^+(\cdot)$ denoting a random sample from each. Define the random function $\delta(\cdot)$ so that $\delta(S) = v^+(S) - v(S)$ for all S. Suppose that $\delta(S) \geq 0$ for all S and that $\delta(\cdot)$ also satisfies $\mathsf{E}_D\left[\max_{S\subseteq[n]}\{\delta(S)\}\right] \leq \overline{\delta}$. Then for any $0 < \epsilon < 1$,

$$\text{Rev}(D^+) \geq (1-\epsilon)\left(\text{Rev}(D) - \overline{\delta}/\epsilon\right).$$

PROOF. Consider a mechanism M which achieves the optimal revenue $\text{Rev}(D)$. Let (ϕ_v, p_v) denote the lottery purchased by a buyer with type v in M, where ϕ_v is a (possibly randomized) allocation, and p is a price. Consider now the mechanism M^+ that offers the same menu as M, but with all prices discounted by a factor of $(1-\epsilon)$. Let (ϕ_v^+, p_v^+) denote the lottery that a buyer with type v^+ (coupled with v) chooses to purchase in M^+ (knowing that she would only pay $(1-\epsilon)p_v^+$ because of the discount). The following inequalities must hold (we will abuse notation and let $v(\psi) = \mathsf{E}_{S\leftarrow\psi}[v(S)]$):

$$v(\phi_v) - p_v \ \geq \ v(\phi_v^+) - p_v^+. \tag{2}$$
$$v^+(\phi_v^+) - (1-\epsilon)\, p_v^+ \ \geq \ v^+(\phi_v) - (1-\epsilon)\, p_v \tag{3}$$

Now, summing equations (2) and (3) (then making use of the definition of $\delta(\cdot)$ and the fact that it is non-negative), we have:

$$\begin{aligned} \epsilon p_v^+ + \delta(\phi_v^+) \ &\geq \ \epsilon p_v \\ \Rightarrow p_v^+ \ &\geq \ p_v - \delta(\phi_v^+)/\epsilon \end{aligned}$$

We can now bound the expected revenue by taking an expectation over all valuations:

$$\begin{aligned} \textsc{Rev}\left(D^+\right) \ &\geq \ \mathsf{E}_{v\leftarrow D}\left[(1-\epsilon)p_v^+\right] \\ &\geq \ (1-\epsilon)\,\mathsf{E}_{v\leftarrow D}\left[p_v - \delta(\phi_v^+)/\epsilon\right] \\ &\geq \ (1-\epsilon)\,\textsc{Rev}(D) - (1-\epsilon)\bar{\delta}/\epsilon \end{aligned}$$

□

COROLLARY 3.5.
For a given partition of $[n]$, $S \sqcup X$, and distribution D over subadditive valuations, define D_S to be D restricted to items in S, and D^+ to first sample $v \leftarrow D$, and output $v^+(\cdot)$ with $v^+(U) = v(U\cap S) + v(U \cap X)$. Then for all $\epsilon \in (0,1)$, $\textsc{Rev}(D) \leq \frac{\textsc{Rev}(D^+)}{1-\epsilon} + \frac{\textsc{Val}(D_S)}{\epsilon}$.

PROOF. By monotonicity, $v(U) \geq v(U\cap X)$ for all U, X. Therefore, $v^+(U) - v(U) \leq v(U \cap S) \leq v(S)$ for all U. Furthermore, by subadditivity, we have $v^+(U) \geq v(U)$ for all U. Together, this means that D and D^+ are coupled so that we can set $\delta(U) \leq v(S)$ for all U. Therefore, we may set $\bar{\delta} = \textsc{Val}(D_S)$ in the hypothesis of Lemma 3.4. The corollary follows by rearranging the inequality. □

The proof of Lemma 3.2 is now a combination of Corollary 3.5 and Lemma 3.3. We can now provide our approximate core decomposition by combining sub domain stitching (Lemma 2.5) and approximate marginal mechanism (Lemma 3.2).

LEMMA 3.6.
("Approximate Core Decomposition")
For any distribution D that is subadditive over independent items, and any $0 < \epsilon < 1$,

$$\textsc{Rev}\left(D\right) \leq \left(\frac{1}{\epsilon} + \frac{1}{1-\epsilon}\right)\textsc{Val}\left(D_\emptyset^C\right) + \frac{1}{1-\epsilon}\sum_{A\subseteq[n]} p_A \textsc{Rev}\left(D_A^T\right).$$

In particular, for $\epsilon = 1/2$, we have

$$\textsc{Rev}\left(D\right) \leq 4\textsc{Val}\left(D_\emptyset^C\right) + 2\sum_{A\subseteq[n]} p_A \textsc{Rev}\left(D_A^T\right)$$

PROOF. By the Approximate Marginal Mechanism Lemma (Lemma 3.2),

$$\textsc{Rev}\left(D_A\right) \ \leq \ \left(\frac{1}{\epsilon} + \frac{1}{1-\epsilon}\right)\textsc{Val}\left(D_A^C\right) + \frac{1}{1-\epsilon}\textsc{Rev}\left(D_A^T\right)$$

Also, for any $A \subseteq [n]$,

$$\textsc{Val}\left(D_A^C\right) \leq \textsc{Val}\left(D_\emptyset^C\right)$$

Finally, by sub-domain stitching (Lemma 2.5):

$$\begin{aligned}
\textsc{Rev}\,(D) &\leq \sum_{A\subseteq[n]} p_A \textsc{Rev}\,(D_A) \\
&\leq \sum_{A\subseteq[n]} p_A \left(\left(\frac{1}{\epsilon}+\frac{1}{1-\epsilon}\right)\textsc{Val}\left(D_A^C\right)+\frac{1}{1-\epsilon}\textsc{Rev}\left(D_A^T\right)\right) \\
&\leq \left(\frac{1}{\epsilon}+\frac{1}{1-\epsilon}\right)\textsc{Val}\left(D_\emptyset^C\right)+\frac{1}{1-\epsilon}\sum_{A\subseteq[n]} p_A \textsc{Rev}\left(D_A^T\right)
\end{aligned}$$

□

3.2. Core

Here, we show how to bound $\textsc{Val}(D_\emptyset^C)$ using $\max\{\textsc{SRev}^*_{\vec{q}}(D), \textsc{BRev}(D)\}$. We use a concentration result due to Schechtman [Schechtman 1999] that first requires a definition.

Definition 3.7. Let $D^{\vec{x}}$ denote a distribution of private information, V denote a valuation function $V(\vec{x},\cdot)$, and D denote the distribution that samples $\vec{x} \leftarrow D^{\vec{x}}$ and outputs the function $v(\cdot) = V(\vec{x},\cdot)$. Then D is *$c-$Lipschitz* if for all $\vec{x},\vec{y}$, and sets S and T we have:

$$|V(\vec{x},S)-V(\vec{y},T)| \leq c\cdot(|S\cup T|-|S\cap T|+|\{i\in S\cap T\colon x_i\neq y_i\}|)\,.$$

Before applying Schechtman's theorem, we show that $D_\emptyset^C$ is t-Lipschitz (recall that t is the cutoff between core and tail).

LEMMA 3.8.
Let D be any distribution that is subadditive over independent items where each $v(\{i\})\in[0,c]$ with probability 1. Then D is c-Lipschitz.

PROOF. For any $\vec{x},\vec{y},S,T$, let $U=\{i\in S\cap T|x_i=y_i\}$. Because of no externalities, we must have $V(\vec{x},U)=V(\vec{y},U)$, which we will denote by B. By monotonicity, we must have $V(\vec{x},S),V(\vec{y},T)\geq B$. By subadditivity and the fact that each $V(\vec{x},\{i\})\leq c$, we have $V(\vec{x},S)\leq c(|S|-|U|)+B$. Similarly, we have $V(\vec{y},T)\leq c(|T|-|U|)+B$. It's also clear that $|S|-|U|\leq|S\cup T|-|S\cap T|+|\{i\in S\cap T:x_i\neq y_i\}|$, and that $|T|-|U|\leq|S\cup T|-|S\cap T|+|\{i\in S\cap T:x_i\neq y_i\}|$. So we also must have $V(\vec{x},S),V(\vec{y},T)\leq B+c(|S\cup T|-|S\cap T|+|\{i\in S\cap T:x_i\neq y_i\}|)$. Therefore $V(\vec{x},S),V(\vec{y},S)\in[B,B+c(|S\cup T|-|S\cap T|+|\{i\in S\cap T:x_i\neq y_i\}|)]$, completing the proof. □

COROLLARY 3.9. *$D_\emptyset^C$ is t-Lipschitz.*

Now we state Schechtman's theorem and apply it to bound $\textsc{Val}(D_\emptyset^C)$.

THEOREM 3.10. *([Schechtman 1999]) Suppose that D is a distribution that is subadditive over independent items and c-Lipschitz. Then for any parameters $q,a,k>0$,*

$$\Pr_{v\leftarrow D}[v([n])\geq(q+1)\,a+k\cdot c]\leq\Pr[v([n])\leq a]^{-q}\,q^{-k}$$

In particular, if a is the median of $v([n])\,|_{v\leftarrow D}$ and $q=2$, we have

$$\Pr_{v\leftarrow D}[v([n])\geq 3a+k\cdot c]\leq 4\cdot 2^{-k}$$

COROLLARY 3.11. *Suppose that D is a distribution that is subadditive over independent items and c-Lipschitz. If a is the median of $v([n])\,|_{v\leftarrow D}$, then $\mathsf{E}_{v\leftarrow D}[v([n])]\leq 3a+4c/\ln 2$.*

PROOF. $\mathsf{E}[v([n])] = \int_0^\infty \Pr[v([n]) > y]\, dy$. We can upper bound this using the minimum of 1 and the bound provided in Theorem 3.10 to yield:

$$\mathsf{E}[v([n])] \leq 3a + \int_0^\infty 4 \cdot 2^{-y/c} dy = 3a + 4c/\ln 2$$

□

PROPOSITION 3.12. $\textsc{Val}(D_\emptyset^C) \leq 6\textsc{BRev} + 4t/\ln 2$.

PROOF. Since a is the median of $v([n])$, we can set price a on the grand bundle and extract revenue at least $a/2$. Therefore, $\textsc{BRev} \geq a/2$. The proposition then follows by combining Corollaries 3.9 and 3.11. □

Finally, if the cutoff t is not too large, we can recover a constant fraction of it by selling each item separately.

LEMMA 3.13. $\textsc{SRev}_{\vec{p}}^* \geq t \cdot p_\emptyset (1 - p_\emptyset)$. *In particular, if we choose* t *so that* $p_\emptyset = 1/2$*, then* $\textsc{SRev}_{\vec{p}}^* \geq t/4$.

PROOF. Clearly $\textsc{Rev}_{p_i}(D_i) \geq p_i t$, as we could set a price of t for item i. So $\textsc{SRev}_{\vec{p}}^* = p_\emptyset \sum_i \textsc{Rev}_{p_i}(D_i) \geq p_\emptyset t \sum_i p_i$ Finally, we observe that $\sum_i p_i$ is exactly the expected number of items in the tail, and $p_\emptyset$ is the probability that zero items are in the tail. So we clearly have $\sum_i p_i \geq 1 - p_\emptyset$. □

Combining Proposition 3.12 and Lemma 3.13 then yields:

PROPOSITION 3.14.

$$\textsc{Val}(D_\emptyset^C) \leq 6\textsc{BRev} + 24\textsc{SRev}_{\vec{p}}^*.$$

3.3. Tail

We now show that the revenue from the tail can be approximated by $\textsc{SRev}_{\vec{q}}^*$. We begin by proving a much weaker bound on the optimum revenue for any distribution of subadditive valuations over independent items:

LEMMA 3.15.

$$\textsc{Rev}(D) \leq 6n^{\log_2 6} \sum_i \textsc{Rev}(D_i).$$

PROOF. Babaioff et al. [Babaioff et al. 2014] prove that $\textsc{Rev} \leq n \sum_i \textsc{Rev}(D_i)$ for an additive buyer by recursively reducing the number of items by one at each step. Unfortunately, each step of the induction uses the Marginal Mechanism Lemma; when applying the approximate variant for subadditive valuations, we would incur an exponential factor.

Instead, we use a slightly more complicated argument along the lines of Hart and Nisan [Hart and Nisan 2012] that halves the number of items in each step. Let S and X be a partition of $[n]$ into subsets of size at most $\lceil n/2 \rceil$. Let $D_{S \geq X}$ be the distribution over valuations which is the same as D whenever $v(S) \geq v(T)$, and has valuation zero otherwise. Similarly, let $D_{S<X}$ be the distribution which is equal to D on $v(S) < (T)$. Then by sub-domain stitching (Lemma 2.5) we have,

$$\textsc{Rev}(D) \leq \textsc{Rev}(D_{S \geq X}) + \textsc{Rev}(D_{S<X}). \tag{4}$$

Now, by the Approximate Marginal Mechanism Lemma,

$$\textsc{Rev}\left(D_{S\geq X}\right) \leq \left(\frac{1}{\epsilon}+\frac{1}{1-\epsilon}\right)\textsf{Val}\left(D^X_{S\geq X}\right)+\frac{1}{1-\epsilon}\mathsf{E}_{v^X\leftarrow D^X_{S\geq X}}\left[\textsc{Rev}\left(D^S_{S\geq X}\mid v^X\right)\right] \tag{5}$$

One mechanism for selling items in S is to sample $v^X \leftarrow D^X_{S\geq X}$, and then use a mechanism that achieves $\textsc{Rev}\left(D^S_{S\geq X}\mid v^X\right)$. Thus we have,

$$\textsc{Rev}\left(D^S\right) \geq \mathsf{E}_{v^X\leftarrow D^X_{S\geq X}}\left[\textsc{Rev}\left(D^S_{S\geq X}\mid v^X\right)\right]. \tag{6}$$

Another way to sell items in S is to again sample $v^X \leftarrow D^X_{S\geq X}$, and offer the entire bundle for price $v^X(X)$. Therefore we also have,

$$\textsc{Rev}\left(D^S\right) \geq \mathsf{E}_{v\leftarrow D}\left[v(X)\mid \left(v(S)\geq v(X)\right)\right] = \mathsf{E}_{v\leftarrow D_{S\geq X}}\left[v(X)\right] = \textsf{Val}\left(D^X_{S\geq X}\right). \tag{7}$$

Combining equations (5)-(7), we have

$$\textsc{Rev}\left(D_{S\geq X}\right) \leq \left(\frac{1}{\epsilon}+\frac{2}{1-\epsilon}\right)\textsc{Rev}\left(D^S\right)$$

By symmetry, the same holds for $\textsc{Rev}\left(D_{S<X}\right)$ and $\textsc{Rev}\left(D^X\right)$. Therefore using (4),

$$\textsc{Rev}\left(D\right) \leq \left(\frac{1}{\epsilon}+\frac{2}{1-\epsilon}\right)\left(\textsc{Rev}\left(D^S\right)+\textsc{Rev}\left(D^X\right)\right)$$

Applying the recursion $\lceil \log n \rceil$ times, we have

$$\textsc{Rev}\left(D\right) \leq \left(\frac{1}{\epsilon}+\frac{2}{1-\epsilon}\right)^{\log_2 n+1}\sum_i \textsc{Rev}\left(D_i\right) = \left(\frac{1}{\epsilon}+\frac{2}{1-\epsilon}\right)n^{\log_2\left(\frac{1}{\epsilon}+\frac{2}{1-\epsilon}\right)}\sum_i \textsc{Rev}\left(D_i\right)$$

Choosing $\epsilon = 1/2$ yields $\left(\frac{1}{\epsilon}+\frac{2}{1-\epsilon}\right) = 6$. $\square$

Note that in Lemma 3.15, $\sum_i \textsc{Rev}(D_i)$ is *not* the same as $\textsc{SRev}(D)$ as the buyer is not necessarily additive. In fact, they can be off by a factor of n (in the case of a unit-demand buyer). Nonetheless, this weak bound suffices for our analysis of the tail, which is summarized in Proposition 3.16 below. Essentially, the proposition amplifies the bound in Lemma 3.15 greatly by making use of the fact that it is unlikely to see multiple items in the tail.

PROPOSITION 3.16. *Recall that* $p_i = \Pr[v(\{i\}) > t]$*, and* $p_\emptyset = \prod_i (1-p_i)$*. Then*

$$\sum_A p_A \textsc{Rev}(D^T_A) \leq \frac{6}{p_\emptyset}\left(1+7\ln(1/p_\emptyset)+6\ln(1/p_\emptyset)^2+\ln(1/p_\emptyset)^3\right)\cdot \textsc{SRev}^*_{\vec{p}}$$

In particular, if we choose t *so that* $p_\emptyset = 1/2$*, then* $\sum_A p_A \textsc{Rev}(D^T_A) \leq 109\cdot\textsc{SRev}^*_{\vec{p}}$

PROOF. Our proof builds on the intuition that the number of items in the tail is typically very small. By Lemma 3.15, we have that

$$\begin{aligned}\sum_{A\subseteq[n]} p_A \textsc{Rev}\left(D^T_A\right) &\leq \sum_{A\subseteq[n]} p_A 6\left|A\right|^{\log_2 6}\sum_{i\in A}\textsc{Rev}\left(D^T_i\right)\\ &= 6\sum_{i\in[n]} p_i \textsc{Rev}\left(D^T_i\right)\sum_{A\ni i}\left|A\right|^{\log_2 6} p_A/p_i \end{aligned} \tag{8}$$

For any i, the expression $\sum_{A\ni i}|A|\,p_A/p_i$ is also the expected number of items in the tail, conditioning on i being in the tail. Similarly, $\sum_{A\ni i}|A|^{\log_2 6}\,p_A/p_i$ is the expected value of (# items)$^{\log_2 6}$. Let b_j be the indicator random variable that is 1 whenever item j is in the tail. Noting that $\log_2 6 < 3$ and each b_j is 1 with probability exactly p_j and the b_j's are independent, we have:

$$\begin{aligned}
\sum_{A\ni i}|A|^{\log_2 6}\,p_A/p_i &\le \mathsf{E}_{b_j}\left[\left(1+\sum_{j\neq i}b_j\right)^3\right]\\
&= \mathsf{E}_{b_j}\left[1+3\left(\sum_{j\neq i}b_j\right)+3\left(\sum_{j\neq i}b_j\right)^2+\left(\sum_{j\neq i}b_j\right)^3\right]\\
&= 1+3\mathsf{E}\left[\sum_{j\neq i}b_j\right]+3\mathsf{E}\left[\sum_{j\neq i}b_j^2+\sum_{k\neq j\neq i}b_jb_k\right]\\
&\quad+\mathsf{E}\left[\sum_{j\neq i}b_j^3+3\sum_{k\neq j\neq i}b_j^2b_k+\sum_{l\neq k\neq j\neq i}b_jb_kb_l\right]\\
&= 1+7\mathsf{E}\left[\sum_{j\neq i}b_j\right]+6\mathsf{E}\left[\sum_{k\neq j\neq i}b_jb_k\right]+\mathsf{E}\left[\sum_{l\neq k\neq j\neq i}b_jb_kb_l\right] \qquad (9)\\
&\le 1+7\sum_{j\neq i}p_j+6\left(\sum_{j\neq i}p_j\right)^2+\left(\sum_{j\neq i}p_j\right)^3 \qquad (10)
\end{aligned}$$

(9) follows because $b_j\in\{0,1\}$. We continue to bound the last line as a function of just $p_\emptyset$. Note that $e^{-\sum_i p_i}\ge\prod_i(1-p_i)=p_\emptyset$, and therefore we have $\sum_i p_i\le\ln(1/p_\emptyset)$. Combining with (8) and (10), we have:

$$\sum_{A\subseteq[n]}p_A\text{Rev}\left(D_A^T\right)\le 6\left(1+7\ln(1/p_\emptyset)+6\ln(1/p_\emptyset)^2+\ln(1/p_\emptyset)^3\right)\cdot\sum_{i\in[n]}p_i\text{Rev}\left(D_i^T\right)$$

Now, we have to interpret $p_i\text{Rev}(D_i^T)$. We claim that in fact this is exactly $\text{Rev}_{p_i}(D_i)$. Why? It's clear that the optimal reserve for D_i^T is at least t, as D_i^T is not supported below t. It's also easy to see that for any reserve $r_i\ge t$, that the revenue obtained by selling to D_i^T is exactly $r_i\cdot\Pr[v(\{i\})>r_i]/p_i$, and therefore the same $r_i\ge t$ that is optimal for D_i is also optimal for D_i^T, and $p_i\text{Rev}(D_i^T)=\text{Rev}_{p_i}(D_i)$. Therefore,

$$\sum_{A\subseteq[n]}p_A\text{Rev}\left(D_A^T\right)\le 6\left(1+7\ln(1/p_\emptyset)+6\ln(1/p_\emptyset)^2+\ln(1/p_\emptyset)^3\right)\cdot\sum_{i\in[n]}\text{Rev}_{p_i}\left(D_i\right)$$

Plug in $\text{SRev}^*_{\vec{p}}=p_\emptyset\sum_{i\in[n]}\text{Rev}_{p_i}\left(D_i\right)$ to complete the proof. □

Note that Theorem 3.1 is now a corollary of Proposition 3.16, Proposition 3.14, and Lemma 3.6 (setting $\epsilon=1/2$). That the desired $\vec{q}$ can be computed efficiently is easy to see: simply do a binary search over cutoffs t until we find one that induces $p_\emptyset=1/2$. It is also easy to find an item pricing that guarantees revenue at least $\text{SRev}^*_{\vec{q}}$: for each item i, simply find the optimal reserve for D_i, subject to that reserve being at least t.

Finally, notice that the only bundle price we ever need to set to obtain our guarantee is the median of $v([n])$, when $v(\cdot) \leftarrow D^C_\emptyset$. It is also easy to see that our bounds degrade smoothly if we set a price that only approximates the median instead. For a discussion of exactly what access to D suffices in order for these prices/cutoffs to be truly easy to find, we refer the reader to [Babaioff et al. 2014]. We note here just that it should be clear that any reasonable, even minimal, access to D does indeed suffice.

4. SIMPLE AUCTIONS AND APPROXIMATE REVENUE MONOTONICITY

In this section we explore the rich connection between approximately optimal simple auctions, and approximate revenue monotonicity. By approximate revenue monotonicity, we formally mean the following:

Definition 4.1. We say that a class of distributions is *α-approximately revenue monotone* if for any two distributions D and D^+ in that class such that D^+ first-order stochastically dominates D,[5] $\alpha \cdot \textsc{Rev}(D^+) \geq \textsc{Rev}(D)$.

In the rest of the section we observe that subadditive valuations over independent items are α-approximately monotone for some constant factor (Subsection 4.1). We also note that a (significantly) tighter approximate monotonicity would yield a better factor of approximation in Theorem 3.1 (Subsection 4.2). Finally, for the class of (possibly correlated) additive valuations over n items, we prove a reduction from approximate revenue monotonicity to approximately optimal simple auctions (that loses a factor of n). Then we use an infinite gap between $\max\{\text{B}\textsc{Rev}, \text{S}\textsc{Rev}\}$ and $\textsc{Rev}$ for two correlated items due to Hart and Nisan [Hart and Nisan 2013] to prove an infinite lower bound on approximate revenue monotonicity (Subsection 4.3).

4.1. Approximately optimal simple auctions imply approximate monotonicity

As a corollary of our main theorem (Theorem 3.1) we deduce constant-factor approximate monotonicity for subadditive valuations over independent items:

COROLLARY 4.2. *The class of subadditive valuations over independent items is* 338-*approximately monotone.*

Similarly, from the 6-approximation of Babaioff et al. for additive yields.

COROLLARY 4.3. *The class of additive valuations over independent items is* 6-*approximately monotone.*

PROOF. For additive functions, B$\textsc{Rev}$ and S$\textsc{Rev}$ constitute of separate Myerson's auctions, and are therefore revenue monotone. Thus,

$$\begin{aligned} 6\textsc{Rev}\left(D^+\right) &\geq 6\max\left\{\text{B}\textsc{Rev}\left(D^+\right), \text{S}\textsc{Rev}\left(D^+\right)\right\} \\ &\geq 6\max\left\{\text{B}\textsc{Rev}\left(D\right), \text{S}\textsc{Rev}\left(D\right)\right\} \geq \textsc{Rev}\left(D\right) \end{aligned}$$

For subadditive functions, S$\textsc{Rev}(D)$ is no longer monotone, but S$\textsc{Rev}^*_{\vec{q}}(D)$ is. This is because S$\textsc{Rev}_q(D_i)$ is clearly monotone, and S$\textsc{Rev}^*_{\vec{q}}$ is just a (scaled) sum of S$\textsc{Rev}_{q_i}(D_i)$. So we get that there exists a $\vec{q}$ such that:

$$\begin{aligned} 338\textsc{Rev}\left(D^+\right) &\geq 338\max\left\{\text{B}\textsc{Rev}\left(D^+\right), \text{S}\textsc{Rev}^*_{\vec{q}}\left(D^+\right)\right\} \\ &\geq 338\max\left\{\text{B}\textsc{Rev}\left(D\right), \text{S}\textsc{Rev}^*_{\vec{q}}\left(D\right)\right\} \geq \textsc{Rev}\left(D\right) \end{aligned}$$

□

[5] Recall that we say D^+ first-order stochastically dominates D if they can be coupled so that when we sample v^+ from D^+ and v from D we have $v^+(S) \geq v(S)$ for all S.

4.2. Approximate monotonicity implies approximately optimal simple auctions

A closer look at the proof of our main theorem also yields the converse of the above corollaries, namely: a tighter approximate monotonicity for subadditive valuations would yield an improved factor of approximation by simple auctions, as well as a simpler proof.

COROLLARY 4.4. *If the class of subadditive valuations over independent items is α-approximately monotone, then*

$$\text{Rev} \leq \alpha\Big((37\alpha + 24)\, \text{SRev} + 6\text{BRev}\Big)$$

PROOF. (Sketch)
Recall that in the proof of the Approximate Marginal Mechanism Lemma (Lemma 3.2), we made use of Lemma 3.4 to bound the gap between $\text{Rev}(D^+)$ and $\text{Rev}(D)$, where D^+ first-order stochastically dominated D. Instead of the ϵ-truthful to truthful reduction, we could derive $\alpha\text{Rev}(D^+) \geq \text{Rev}\,(D)$ from approximate monotonicity. Then, we can directly apply Lemma 3.3 to get:

$$\text{Rev}\,(D) \leq \alpha \left(\text{Val}\left(D^S\right) + \mathsf{E}_{v^S \sim D^S}\left[\text{Rev}\left(D^X \mid \mathbf{v}^S\right)\right]\right)$$

Instead of

$$\text{Rev}\,(D) \leq 4\text{Val}\left(D^S\right) + 2\mathsf{E}_{v^S \sim D^S}\left[\text{Rev}\left(D^X \mid \mathbf{v}^S\right)\right]$$

If $\alpha \leq 2$, this indeed yields a tighter approximation. □

4.3. Correlated Distributions are not Approximately Monotone

So far we've shown that (for some valuation classes) approximately optimal simple mechanisms imply approximate revenue-monotonicity. Are all classes approximately revenue-monotone? In this subsection we provide a reduction from an instance where $\max\{\text{SRev}, \text{BRev}\}$ does not approximate Rev to show an infinite non-monotonicity for correlated items. We first prove a reduction from gaps between BRev and Rev to non-monotonicity.

PROPOSITION 4.5. *Let D be a distribution over subadditive valuations for n items for which $\text{Rev}(D) > c \cdot \text{BRev}(D)$. Then any class of distributions containing D and all single-dimensional distributions*[6] *is not (c/n)-approximately revenue monotone.*

PROOF. We define D^+ as follows. First, sample $v \leftarrow D$. Then let $i^* = \arg\max_i\{v(\{i\})\}$. Now, set $v^+(S) = v(\{i^*\}) \cdot |S|$ for all $|S|$. By subadditivity, it's clear that D^+ first order stochastically dominates D. Now, however, D^+ is a single-dimensional distribution, meaning that $\text{BRev}(D^+) = \text{Rev}(D^+)$ [Myerson 1981; Riley and Zeckhauser 1983]. Finally, we just need to compare $\text{BRev}(D)$ to $\text{BRev}(D^+)$.
Note that by monotonicity, we have $v^+([n]) \leq n \cdot v([n])$ for all v, v^+. Therefore, for any price p, if $v^+([n]) > p$, $v([n]) > p/n$. This immediately implies that $\text{BRev}(D) \geq \text{BRev}(D^+)/n$: let p be the optimal reserve for the grand bundle under D^+, then setting price p/n sells with at least the same probability under D. Putting both observations together, we see that: $\text{Rev}(D) > c\text{BRev}(D) \geq (c/n)\text{BRev}(D^+) = (c/n)\text{Rev}(D^+)$, meaning that any class containing D and D^+ is not (c/n)-approximately monotone. □

We apply Proposition 4.5 to a theorem of Hart and Nisan.

THEOREM 4.6. *(Hart and Nisan [Hart and Nisan 2013]) There exists a distribution D over correlated additive valuations for two items such that $\text{BRev}\,(D) \leq 1/2$, and $\text{Rev}\,(D) = \infty$.*

[6]A distribution is single-dimensional if for all v in its support, $v(S) = c|S|$ for some value c.

COROLLARY 4.7. *There exist distributions D and D^+ over over correlated additive valuations for two items such that D^+ first-order stochastically dominates D, $\text{REV}(D^+) = 1$, and yet $\text{REV}(D) = \infty$. Therefore, the class of correlated additive valuations for two items is not c-approximately revenue monotone for any finite c.*

ACKNOWLEDGMENT

We would like to thank Moshe Babaioff, Hu Fu, Nicole Immorlica, and Brendan Lucier for numerous suggestions and helpful discussions.

REFERENCES

BABAIOFF, M., IMMORLICA, N., LUCIER, B., AND WEINBERG, S. M. 2014. A simple and approximately optimal mechanism for an additive buyer. *FOCS*.

BATENI, M., DEHGHANI, S., HAJIAGHAYI, M., AND SEDDIGHIN, S. 2015. Revenue Maximization for Selling Multiple Correlated Items. *arxiv report*. http://arxiv.org/abs/1412.3187.

BEI, X. AND HUANG, Z. 2011. Bayesian Incentive Compatibility via Fractional Assignments. In *the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

BRIEST, P., CHAWLA, S., KLEINBERG, R., AND WEINBERG, S. M. 2010. Pricing randomized allocations. In *SODA*. 585–597.

CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2013. Understanding Incentives: Mechanism Design becomes Algorithm Design. In *the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. http://arxiv.org/pdf/1305.4002v1.pdf.

CAI, Y. AND HUANG, Z. 2013. Simple and Nearly Optimal Multi-Item Auctions. In *the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

CHAWLA, S., HARTLINE, J. D., AND KLEINBERG, R. 2007. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*. EC '07. 243–251.

CHAWLA, S., HARTLINE, J. D., MALEC, D. L., AND SIVAN, B. 2010a. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*.

CHAWLA, S., MALEC, D. L., AND SIVAN, B. 2010b. The Power of Randomness in Bayesian Optimal Mechanism Design. In *the 11th ACM Conference on Electronic Commerce (EC)*.

CHEN, X., DIAKONIKOLAS, I., PAPARAS, D., SUN, X., AND YANNAKAKIS, M. 2014. The complexity of optimal multidimensional pricing. In *the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

DASKALAKIS, C., DECKELBAUM, A., AND TZAMOS, C. 2014. The complexity of optimal mechanism design. In *the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

DASKALAKIS, C. AND WEINBERG, S. M. 2012. Symmetries and Optimal Multi-Dimensional Mechanism Design. In *the 13th ACM Conference on Electronic Commerce (EC)*.

HART, S. AND NISAN, N. 2012. Approximate revenue maximization with multiple items. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. EC '12. 656–656.

HART, S. AND NISAN, N. 2013. The menu-size complexity of auctions. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. EC '13. 565–566.

HART, S. AND RENY, P. J. 2012. Maximal Revenue with Multiple Goods: Nonmonotonicity and Other Observations. Discussion Paper Series dp630, The Center for the

Study of Rationality, Hebrew University, Jerusalem. Nov.

HARTLINE, J. D. 2011. *Approximation and Mechanism Design*.

HARTLINE, J. D., KLEINBERG, R., AND MALEKIAN, A. 2011. Bayesian Incentive Compatibility via Matchings. In *the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

KLEINBERG, R. AND WEINBERG, S. M. 2012. Matroid Prophet Inequalities. In *the 44th Annual ACM Symposium on Theory of Computing (STOC)*.

LI, X. AND YAO, A. C.-C. 2013. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences 110,* 28, 11232–11237.

MYERSON, R. B. 1981. Optimal Auction Design. *Mathematics of Operations Research 6,* 1, 58–73.

PAVLOV, G. 2011. A Property of Solutions to Linear Monopoly Problems. *B. E. Journal of Theoretical Economics 11,* 1. Article 4.

RILEY, J. AND ZECKHAUSER, R. 1983. Optimal Selling Strategies: When to Haggle, When to Hold Firm. *Quarterly J. Economics 98,* 2, 267–289.

SCHECHTMAN, G. 1999. Concentration, results and applications.

THANASSOULIS, J. 2004. Haggling Over Substitutes. *J. Economic Theory 117*, 217–245.

YAO, A. C.-C. 2015. An n-to-1 bidder reduction for multi-item auctions and its applications. In *To Appear in the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.