# A Simple and Approximately Optimal Mechanism for an Additive Buyer 

Moshe Babaioff*, Nicole Immorlica*, Brendan Lucier*, S. Matthew Weinberg ${ }^{\dagger}$<br>*Microsoft Research<br>${ }^{\dagger}$ EECS, MIT


#### Abstract

We consider a monopolist seller with $n$ heterogeneous items, facing a single buyer. The buyer has a value for each item drawn independently according to (non-identical) distributions, and his value for a set of items is additive. The seller aims to maximize his revenue. It is known that an optimal mechanism in this setting may be quite complex, requiring randomization [19] and menus of infinite size [15]. Hart and Nisan [17] have initiated a study of two very simple pricing schemes for this setting: item pricing, in which each item is priced at its monopoly reserve; and bundle pricing, in which the entire set of items is priced and sold as one bundle. Hart and Nisan [17] have shown that neither scheme can guarantee more than a vanishingly small fraction of the optimal revenue. In sharp contrast, we show that for any distributions, the better of item and bundle pricing is a constant-factor approximation to the optimal revenue. We further discuss extensions to multiple buyers and to valuations that are correlated across items.


## I. Introduction

A monopolist seller has a collection of $n$ items to sell. How should he sell the items to maximize revenue given that the buyers are strategic? When there is only a single item for sale, and a single buyer with value drawn from a distribution $F$, Myerson [23] shows that the optimal sale protocol is straightforward: the seller should post a fixed take-it-or-leave-it price $p$ chosen to maximize $p(1-F(p))$, the expected revenue. The optimality of this simple auction format extends to the case of multiple buyers, as well. ${ }^{1}$ Despite the simplicity of the single-item case, extending this solution to handle multiple items remains the primary open challenge in mechanism design. While recent work in the computer science literature has made progress on this front [2], [3], [5], [6], [8], [9], [10], [13], [15], [17], [21], it is still the case that very little is known about optimal multiitem auctions, and what is known lacks the simplicity of Myerson's single-item auction.

Consider even the simplest multi-item scenario [17]:

[^0]there is a single buyer ${ }^{2}$ with item values drawn independently from distributions $D_{1}, \ldots, D_{n}$, and whose value for a set of items is additive. Even when there are only two items for sale, it is known that the revenue-optimal mechanism may involve randomization [19], even to the extent of offering the buyer a choice among infinitely many lotteries [15], [18]. This is troubling not only from the perspective of analyzing optimal mechanisms, but also from the point of view of their usefulness. For an auction to be useful in practice, it should be simple to describe and transparent in its execution. Indeed, Myerson's single-item auction is exciting not only for its optimality, but also its practicality. ${ }^{3}$ The danger, then, is that revenue-optimal but complex mechanisms for multiple items may share the fate of other mathematically optimal designs, such as the Vickrey-Clarke-Groves mechanism, which are very rarely used in practice [4]. It is therefore crucial to pair the study of revenue optimization with an exploration of the power of simple auctions. In other words, what is the relative strength of simple versus complex mechanisms?

The above question was posed in general by Hartline and Roughgarden [20], and by Hart and Nisan specifically for the setting of a single additive buyer [17]. They proposed the following suggestion for a simple multi-item auction: sell each item separately, posting a fixed price on each one. The optimal price to set on item $i$ is then $\arg \max _{p} p\left(1-D_{i}(p)\right)$, mirroring the single-item scenario. At first glance, it appears that perhaps this simple approach should be optimal: the buyer's value for each item is sampled independently, and her value for item $i$ doesn't depend at all on what other items she receives due to additivity. There is absolutely no interaction between the items at all from the buyer's perspective, so why not sell the items separately? Somewhat counter-intuitively, it turns out that this mechanism need not achieve the optimal revenue. For example, suppose that there are $n$ items, and that the buyer's value for each item is distributed uniformly on $[0,1]$. Then the optimal price to set on a single item is $\arg \max _{p} p(1-p)=1 / 2$, with a per-item revenue of

[^1]$1 / 4$ and hence a total revenue of $n / 4$. However, there is a different and equally straightforward mechanism that performs much better: offer only the set of all items at a take-it-or-leave-it price of $n\left(\frac{1}{2}-\epsilon\right)$ for some small $\epsilon>0$. As $n$ grows large, the probability that the sum of item values exceeds this price approaches 1 , and hence the buyer is almost certain to buy. This leads to a revenue slightly less than $n / 2$, a significant improvement over $n / 4$. Hart and Nisan [17] show how to modify this example to exhibit a gap of $\Omega(\log (n))$ by replacing the uniform distribution with an EqualRevenue distribution. ${ }^{4}$

What is going on in this example? The inherent problem is that the buyer's value for the set of all items concentrates tightly around its expectation. This is potentially helpful for revenue generation, but the strategy of selling items separately cannot exploit this property. On the other hand, the mechanism designed to target such concentration (selling only the grand bundle at a fixed price) does very poorly in settings where concentration doesn't occur; Hart and Nisan show that this grand-bundle mechanism achieves only an $\Omega(n)$ approximation to the optimal revenue in general. We must conclude that neither of these two simple mechanisms approximate the optimal revenue to within a constant factor.

Our main result is that the maximum of the revenue generated by these two approaches - either selling all items separately or selling only the grand bundle - is a constant-factor approximation to the optimal revenue. In other words, for any distribution of buyer values, either selling items separately approximates the optimal revenue to within a constant factor, or else bundling all items together does. Since a good approximation to the expected revenue of each approach can be computed in polynomial time given an appropriate access to the distribution (see Appendix $G$ in the full version for a discussion of this claim), our results furthermore imply the first polytime constant-factor approximation mechanism for the case of an additive buyer with independently (and non-identically) distributed values, even without the restriction of simplicity. ${ }^{5}$ Furthermore, prior to our work, it was not even known if any de-

[^2]terministic mechanism could achieve a constant-factor approximation to the optimal mechanism, even without regard for simplicity or computational efficiency.

Main Result (Informal). In any market with a single additive buyer and arbitrary independent item value distributions, either selling every item separately or selling all items together as a grand bundle generates at least a constant fraction of the optimal revenue.

Our result nicely complements an active research area aimed at characterizing distributions and valuations in which simple mechanisms are precisely optimal [2], [17], [24], [25]. In contrast to that literature, we show that a maximum over simple mechanisms is approximately optimal, for arbitrary distributions and additive valuations. Our result also echoes a similar line of investigation for markets with unit-demand valuations in which a buyer's value for a set of items is his maximum value for an item in the set. In this setting, it is known [12], [13], [14] that selling items separately achieves a constant approximation to the optimal revenue. Our result illustrates that a similar approximation can be achieved for additive buyers, provided that we also consider selling all items together as a grand bundle.

To obtain some intuition into our result, recall the example above with $n$ items and uniformly-distributed values. This example illustrates that selling all items separately may be a poor choice when the value for the grand bundle concentrates around its expectation. What we show is that, in fact, this is the only scenario in which selling all items separately is a poor choice. We prove that if the total value for all items does not concentrate, then selling separately must generate a constant fraction of optimal revenue.

Our argument makes use of a core-tail decomposition technique introduced by Li and Yao [22] to study the revenue of selling items separately. Roughly speaking, the idea is to split the support of each item's value distribution into a "tail" (those values that are sufficiently large), and a "core" (the remainder). One then attributes the revenue of the optimal mechanism to the revenue extracted from values in the tail, plus the expected sum of values in the core. To bound the optimal revenue, it then suffices to bound each of these two quantities separately. Li and Yao define the tail of a distribution so that each value is in the tail with probability at most $1 / n$; they use this to prove that selling all items separately obtains a logarithmic approximation to the optimal revenue (which is tight).

We apply a similar approach, but we define the boundary between core and tail in a different way. We aim to strike a balance between two opposing goals:
we want the boundaries to be high enough that the probability of being in the tail is low, which will imply that the revenue from the tail is small relative to selling items separately. At the same time, we want values in the core to be small enough that, subject to their sum being large, the sum must necessarily concentrate around its expectation (which would imply that bundling all items together achieves good revenue). To meet these two goals, we design thresholds that are adapted to the revenue contributions of different items, which makes the core smaller (relative to non-adaptive thresholds) when the value distributions are highly asymmetric. This gives us the extra flexibility needed to derive a constant-factor approximation.

We apply the same methodology to prove that when there are many buyers (with valuations that are not necessarily samples from identical distributions), selling all items separately yields an $O(\log (n))$ approximation to the optimal mechanism. This bound is asymptotically tight, as Hart and Nisan have presented a lower bound that matches this for just a single buyer. Prior to our work, no non-trivial bounds were known on the revenue of selling separately to many buyers, or even on the revenue on any class of mechanisms. Furthermore, the observation that selling separately fails only under concentration has implications in this setting as well: we further show that unless the maximum attainable welfare (of all buyers together) concentrates, that selling items separately again obtains a constant-factor approximation. However, with many buyers the concentration of welfare does not imply that selling the grand bundle together obtains a constant-factor approximation. Indeed, unlike in the single-buyer case, one cannot improve the approximation ratio by using bundling: we prove that the $\Omega(\log (n))$ lower bound applies against the better of selling separately and together as well. This realization motivates our first open problem:

Open Problem 1. Is there a "simple," approximately optimal mechanism for many additive buyers with independent values?

In attempt to make progress on this problem, we turn to a subclass of deterministic mechanisms that we call "partition mechanisms." A partition mechanism first partitions the items into disjoint bundles, then sells each bundle separately. This natural class of mechanisms clearly generalizes both selling separately and selling together, so we study the performance of the optimal mechanism in this class relative to that of others. On this front, we show that unfortunately the revenue of the optimal mechanism for many independent buyers can still be an $\Omega(\log n)$ factor larger than that of the optimal partition mechanism, and further that revenue of the optimal partition mechanism can be an $\Omega(\log n)$ factor
larger than the better of selling separately and together. ${ }^{6}$ Following the initial submission of this paper, Yao developed a lookahead reduction for the case of many buyers [26]. He is then able to plug our single-buyer results into his reduction, arguably resolving Question 1 in the affirmative. We refer the reader to [26] for further details regarding the simplicity and approximation ratios of his auctions.

Finally, we study the performance of selling separately and together against partition mechanisms for a single buyer whose values for the items may be arbitrarily correlated. While neither class of mechanisms can guarantee any finite factor of the optimal revenue ([7], [17]), the question remains as to whether simple mechanisms can approximate more complex (though still suboptimal) mechanisms in the presence of correlation. To this end, we prove that selling items separately obtains an $O(\log n)$-approximation the optimal obtainable revenue by a partition mechanism, and that this is tight. In fact, we show a gap of $\Omega(\log n)$ between the better of selling separately and together versus the optimal partition mechanism. We include several tables in Appendix A of the full version displaying the relative power of the various classes of mechanisms studied in this paper, noting here that as of our work, all upper and lower bounds are (asymptotically) matching.

Our paper leaves several natural open problems for future work. The first was already stated and concerns extending our results to many buyers. A second problem concerns extending our results beyond additive valuations. As for both unit-demand and additive valuations a constant-factor approximation mechanism is now known, one could naturally ask if such a result is also achievable for valuations that generalize both unit-demand and additive. One potential instantiation is a buyer with a $k$-demand valuation; i.e., additive, but wants at most $k$ items. A significantly more challenging instantiation is the class of gross-substitute valuations.

Open Problem 2. Is there a "simple," approximately optimal mechanism for single buyer with a $k$-demand valuation? With a gross-substitute valuation?

Finally, a third problem concerns extending our results to settings with mild (but not aribtrary) correlation. This approach was fruitful in [14] for the "common base-value" model. ${ }^{7}$

Open Problem 3. Is there a "simple," approximately optimal mechanism for a single additive buyer whose

[^3]value for $n$ items is sampled from a common basevalue distribution? What about other models of limited correlation?

## II. Preliminaries

The setting we consider is that of a single monopolist seller with $n$ heterogeneous and indivisible items for sale to $m$ additive, risk-neutral, quasi-linear consumers (buyers). That is, each consumer $j$ has a value $v_{i j}$ for item $i$. While our main results are for the setting of a single buyer, we will define our setting more generally; this will be useful when discussing extensions. If a randomized outcome awards consumer $j$ item $i$ with probability $\pi_{i j}$ and charges him a price $q_{j}$ in expectation, then his utility for this outcome is $\sum_{i} v_{i j} \pi_{i j}-q_{j}$. Each value $v_{i j}$ is sampled independently from a known distribution $D_{i j}$. We make no assumptions on $D_{i j}$ whatsoever. We refer to $D$ as the joint $m n$-dimensional distribution over all consumers' values for all items, $D_{i}$ as the $m$ dimensional distribution over all consumers' values for item $i$. Furthermore, we denote by $\vec{v}$ a random sample from $D, \vec{v}_{i}$ a random sample from $D_{i}$. We also denote the maximum value for item $i$ as $v_{i}^{*}=\max _{j}\left\{v_{i j}\right\}$.

We are interested in analyzing mechanisms at BayesNash equilibrium of buyer behavior, with an eye toward maximizing revenue at equilibrium. By the revelation principle, we can restrict attention to mechanisms that are Bayesian Incentive Compatible (i.e., truthful). ${ }^{8}$ As usual, we also impose the individual rationality constraint, saying that every buyer's utility is non-negative when truthful.

We use the following terminology to discuss the revenue obtainable by various types of mechanisms, where the first three are taken from [17].

- $\operatorname{Rev}(D)$ : The optimal revenue obtained by any (possibly randomized) truthful mechanism when the consumer profile is drawn from $D$.
- $\operatorname{SREv}(D)$ : The optimal revenue obtained by auctioning items separately when the consumer profile is drawn from $D$. That is, the revenue obtained by running Myerson's optimal auction separately for each item.
- $\operatorname{BREv}(D)$ : The optimal revenue obtained by auctioning the grand bundle when the consumer profile is drawn from $D$. That is, the revenue obtained by running Myerson's optimal auction when treating the grand bundle as a single item.
- $\operatorname{PREv}(D)$ : The optimal revenue obtained by any partition mechanism when the consumer profile is drawn from $D$. That is, the maximal revenue obtained by first partitioning the items into disjoint bundles, and then running Myerson's optimal

[^4]auction separately for each bundle, treating each bundle as a single item.
Given a distribution $D$ over profiles, we will often consider the welfare $\sum_{i} v_{i}^{*}$ of a consumer profile $\vec{v}$ drawn from $D$. We will write $\operatorname{VaL}(D)$ for the expected welfare, so that $\operatorname{VAL}(D)=\mathbb{E}_{\vec{v} \sim D}\left[\sum_{i} v_{i}^{*}\right]$. We will also write $\operatorname{var}(D)=\operatorname{var}_{\vec{v} \sim D}\left(\sum_{i} v_{i}^{*}\right)$ for the variance of the welfare.

We will make use of some results from [17] that provide useful bounds on $\operatorname{REV}(D)$. We include proofs in Appendix B of the full version for completeness. Lemma 1 is stated and proved directly in [17]. Lemma 2 is not directly stated nor proved, but is similar to an implicit result from [17].

In the lemma below, we think of $D$ and $D^{\prime}$ as being distributions over values for disjoint sets of items, for the same set of $m$ consumers. The distribution $D \times D^{\prime}$ then draws values for those two sets of items, independently, from $D$ and $D^{\prime}$ respectively.

Lemma 1. ([17]) $\operatorname{REV}\left(D \times D^{\prime}\right) \leq \operatorname{VaL}(D)+$ $\operatorname{REv}\left(D^{\prime}\right)$.
Lemma 2. $\operatorname{Rev}(D) \leq n m \operatorname{SREv}(D)$.

## III. The Core Decomposition

We make use of an idea developed by Li and Yao [22] called the "core" of a value distribution for a single consumer. In order to obtain our stronger results for a single consumer and also extend to many consumers, we define the core differently but in the same spirit. The idea is to separate each $m$-dimensional value distribution for each item into the core and the tail, the tail being the part where some consumer has an unusually high value for the item. Then the core of the entire $n m$-dimensional distribution is the product of all the cores, and the tail is everything else.

## A. Defining the Core and Prior Results

Below we formalize the notion of the core. We introduce some notation that will be used throughout the paper. By the "null" distribution, we mean a distribution whose product with any other distribution is also a null distribution, and that outputs $\perp$ with probability 1 .

- $r_{i}$ : The optimal revenue obtainable by selling just item $i$ (using Myerson's optimal auction).
- $r: \sum_{i} r_{i}$. The same as $S \operatorname{Rev}(D)$ but cleaner to write in formulas.
- $t_{i}$ : A profile of parameters, one per item, to define the separation between the core and tail of distribution $D_{i}$. We will think of $t_{i}$ as a multiplier applied to $r_{i}$. The core for item $i$ will be supported on the interval $\left[0, t_{i} r_{i}\right]$, and the tail for item $i$ will be supported on $\left(t_{i} r_{i}, \infty\right)$. Different results
throughout the paper will specify different choices for $t_{i}$.
- $p_{i}: \operatorname{Pr}\left[v_{i}^{*}>t_{i} r_{i}\right]$, the probability that the highest value on item $i$ lies in the tail. Note that this may be 0 .
- $D_{i}^{C}$ : The core of $D_{i}$, the conditional distribution of $\vec{v}_{i}$ conditioned on $v_{i}^{*} \leq t_{i} r_{i}$. Note that this may be the null distribution if $p_{i}=1$.
- $D_{i}^{T}$ : The tail of $D_{i}$, the conditional distribution of $\vec{v}_{i}$ conditioned on $v_{i}^{*}>t_{i} r_{i}$. Note that this may be the null distribution if $p_{i}=0$.
- $A$ : Throughout our notation, we will use $A$ to represent a subset of items. We often think of $A$ as the items whose values lie in the tail of their respective distributions.
- $D_{A}^{T}: A$ is a subset of items, and $D_{A}^{T}$ is a product distribution equal to $\times_{i \in A} D_{i}^{T}$.
- $D_{A}^{C}: A$ is a subset of items, and $D_{A}^{C}$ is a product distribution equal to $\times{ }_{i \notin A} D_{i}^{C}$.
- $D_{A}: D_{A}^{C} \times D_{A}^{T}$. Note that this product is taken over the tail of items in $A$ and the core of items not in $A$. In other words, $D_{A}$ is the distribution $D$, conditioned on $v_{i}^{*}>t_{i} r_{i}$ if $i \in A$ and conditioned on $v_{i}^{*} \leq t_{i} r_{i}$ if $i \notin A$.
- $p_{A}: \operatorname{Pr}\left[\vec{v} \in \operatorname{support}\left(D_{A}\right)\right]$. This is equal to $\left(\prod_{i \in A} p_{i}\right)\left(\prod_{i \notin A}\left(1-p_{i}\right)\right)$.
Before stating our core decomposition lemma, we present some known results about the core. The lemmas below were either stated explicitly in [22] or [17], or use ideas from one of those papers. We put a citation in the statement of such lemmas, but include all proofs in Appendix C of the full version.

Lemma 3. ([22]) $p_{i} \leq 1 / t_{i}$ for all $i$.
Lemma 4. ([22]) $\operatorname{REv}\left(D_{i}^{C}\right) \leq r_{i}$ and $\operatorname{REv}\left(D_{i}^{T}\right) \leq$ $r_{i} / p_{i}$.
Lemma 5. ([17]) $\operatorname{REv}(D) \leq \sum_{A} p_{A} \operatorname{REv}\left(D_{A}\right)$.

## B. The Core Decomposition Lemma

In this section we state our Core Decomposition Lemma, which relates the optimal revenue from a distribution $D$ to the revenue and welfare that can be extracted from the tail and core of $D$. This result is similar in spirit to the core lemma of [22].

Our first result, Lemma 6, is our main decomposition lemma. The lemma states that the optimal revenue from distribution $D$ can be split into a contribution from the core of $D$ and a contribution from the tail of $D$. One might hope for a bound of the form "the optimal revenue from $D$ is at most the optimal revenue from the tail plus the optimal revenue from the core." Indeed, such a bound is attainable for a single buyer [22], but is problematic for many buyers, see Section 4.4 and

Appendix 3 in [17] for a discussion. We will therefore settle for a weaker bound: the optimal revenue from the tail plus the expected welfare from the core. We also note that the approach of Li and Yao eventually upper bounds the optimal revenue of the core with the expected welfare anyway.

Lemma 6 (Core Decomposition). $\operatorname{REv}(D) \leq$ $\operatorname{VAL}\left(D_{\emptyset}^{C}\right)+\sum_{A} p_{A} \operatorname{REv}\left(D_{A}^{T}\right)$

Proof: By Lemma 1,

$$
\operatorname{REv}\left(D_{A}\right) \leq \operatorname{VAL}\left(D_{A}^{C}\right)+\operatorname{REv}\left(D_{A}^{T}\right)
$$

for all $A$. Also, since $\operatorname{VAL}\left(D_{A}^{C}\right)$ is the expected sum of values for items not in $A$, we have

$$
\operatorname{VAL}\left(D_{A}^{C}\right) \leq \operatorname{VAL}\left(D_{\emptyset}^{C}\right) .
$$

By Lemma 5,

$$
\begin{aligned}
\operatorname{REv}(D) & \leq \sum_{A} p_{A} \operatorname{REV}\left(D_{A}\right) \\
& \leq \sum_{A} p_{A}\left(\operatorname{VAL}\left(D_{A}^{C}\right)+\operatorname{REV}\left(D_{A}^{T}\right)\right) \\
& \leq\left(\sum_{A} p_{A}\right) \operatorname{VAL}\left(D_{\emptyset}^{C}\right)+\sum_{A} p_{A} \operatorname{Rev}\left(D_{A}^{T}\right) .
\end{aligned}
$$

As $\sum_{A} p_{A}=1$ the desired result follows.

## IV. Revenue Bounds for a Single Buyer

In this section we focus on the case of a single buyer, $m=1$. We will work toward proving our main result, which is that max $\{\operatorname{SREv}(D), \operatorname{BREv}(D)\}$ is a constantfactor approximation to $\operatorname{REV}(D)$ in this setting. Our argument will make use of the core decomposition, described in the previous section. We will begin with a simpler result that illustrates our techniques: that $\operatorname{Rev}(D)$ is at most $(\ln n+3)$ times $\operatorname{SREv}(D)$. A logarithmic approximation was already established in [22]; we obtain a slightly tighter bound, but the primary purpose of presenting this result is as a warm-up to introduce our techniques and those of [22]. We will then show how this bound can be improved to a constant by considering the maximum of $\operatorname{SREV}(D)$ and $\operatorname{BREv}(D)$.
A. Warm-up: $(\ln n+3)$ SREV $\geq$ REV

Theorem 1. For a single buyer, and any $c>0$, $(2+1 / c+\ln c+\ln n) \operatorname{SREv}(D) \geq \operatorname{REv}(D)$. This is minimized at $c=1$, yielding $(\ln n+3) \operatorname{SREV}(D) \geq$ $\operatorname{REv}(D)$.

The idea of the proof is to consider the core decomposition of $D$, choosing $t_{i}=c n$ for each item $i$. By the Core Decomposition Lemma (Lemma 6), Theorem 1 follows if we can bound the optimal revenue from the tail and the expected welfare from the core, given this choice of $\left\{t_{i}\right\}_{i}$.

We begin with Proposition 1, which effectively shows that for constant $c$, the revenue from the tail is at most a constant times $\operatorname{SREV}(D)$. The intuition behind this result is that each item $i$ lies in the tail with probability $p_{i} \leq 1 / t_{i}=1 / c n$, and hence a large fraction of the time there will be at most a single item whose value lies in the tail. In this case, the revenue from the values in the tail is certainly no more than $\operatorname{SREv}(D)$, since the optimal mechanism can do no better than setting the optimal price for the single item present. To bound the revenue contribution when many values lie in the tail, the relatively weak bound in Lemma 2 will suffice.

Proposition 1. For a single buyer, and any $c>0$, if $t_{i}=$ cn for all $i$, then $\sum_{A} p_{A} \operatorname{REV}\left(D_{A}^{T}\right) \leq(1+$ $1 / c) \operatorname{SREv}(D)$.

Proof: By Lemma 2 and Lemma 4, $\operatorname{REv}\left(D_{A}^{T}\right) \leq$ $|A| \operatorname{SREv}\left(D_{A}^{T}\right) \leq \sum_{i \in A}|A| r_{i} / p_{i}$. Therefore, we may rewrite the sum by first summing over item $i$, and then sets $A$ containing $i$, obtaining:

$$
\sum_{A} p_{A} \operatorname{REv}\left(D_{A}^{T}\right) \leq \sum_{i} \sum_{A \ni i}|A| p_{A} r_{i} / p_{i} .
$$

We now wish to interpret the term $\sum_{A \ni i}|A| p_{A} / p_{i}$. Observe that $p_{A} / p_{i}$ is exactly the probability that the set $A$ of items are in the tail, conditioned on $i$ being in the tail, and $|A|$ is just the size of $A$. Summing over all $A \ni i$ therefore yields the expected size of the set of items in the tail, conditioned on $i$ being on the tail. ${ }^{9}$ Clearly this expectation is just $1+\sum_{j \neq i} p_{j}$, which is at most $1+1 / c$ by Lemma 3 .

As we have just observed that $\sum_{A \ni i}|A| p_{A} / p_{i} \leq$ $1+1 / c$, we have now shown that $\sum_{A} p_{A} \operatorname{REV}\left(D_{A}^{T}\right) \leq$ $\sum_{i}(1+1 / c) r_{i}$, which is exactly $(1+1 / c) \operatorname{SREv}(D)$.

Having established a bound on the revenue of the tail, we turn to the welfare of the core. For this, we use the definition of $r_{i}=\operatorname{SREv}\left(D_{i}\right)$ to directly bound $\operatorname{Pr}\left[v_{i}>x\right]$ for all $x$, and then take an expectation over the range of the core.

Proposition 2. For a single buyer, and any $c>0$, if $t_{i}=c n$ for all $i$, then $(1+\ln c+\ln n) \operatorname{SREv}(D) \geq$ $\operatorname{VAL}\left(D_{\emptyset}^{C}\right)$.

Proof: Note that $\operatorname{VAL}\left(D_{\emptyset}^{C}\right)=\sum_{i} \operatorname{VAL}\left(D_{i}^{C}\right) \leq$ $\sum_{i} \int_{0}^{c n r_{i}} \operatorname{Pr}\left[v_{i}>x\right] d x$. The last inequality would be equality if we replaced $v_{i}$ with a random variable drawn from $D_{i}^{C}$, but since $v_{i}$ stochastically dominates such a random variable, we get an inequality instead. As the optimal revenue of $D_{i}$ is $r_{i}$, this means that

[^5]$\operatorname{Pr}\left[v_{i}>x\right] \leq \min \left\{1, r_{i} / x\right\}$. So we have
\[

$$
\begin{aligned}
\operatorname{VAL}\left(D_{i}^{C}\right) & \leq \int_{0}^{r_{i}} d x+\int_{r_{i}}^{c n r_{i}} r_{i} / x d x \\
& \leq r_{i}+r_{i}\left(\ln \left(c n r_{i}\right)-\ln \left(r_{i}\right)\right) \\
& \leq r_{i}(1+\ln n+\ln c)
\end{aligned}
$$
\]

Summing this guarantee over all $i$ yields the proposition.
Combining Propositions 1 and 2 with Lemma 6 yields Theorem 1.

## B. Main Result: 6•甶ax\{SREv, BREv\} $\geq$ REv

In this section we prove our main result, showing that the best of selling items separately and bundling all of them together is a constant-factor approximation to the optimal mechanism. The proof will follow a similar skeleton to that of Section IV-A, by proving propositions similar to Propositions 1 and 2. The notable difference is that we will need to be more careful in defining the core, which makes proving the equivalent of Proposition 1 more technical.

When all $D_{i}$ are identical, the approach in Section IV-A (setting each $t_{i}=\mathrm{cn}$ ) can be leveraged to yield the bound $O(1) \cdot \operatorname{BREV} \geq \operatorname{REV}$ ([22]), but fails in the case that a small number $k$ of items contributes the majority of the optimal revenue. To see the problem, note that the definition of the core depends on the number of items $n$, but this can be made arbitrarily large by adding extra items of negligible value. The effect is that the core is potentially larger than necessary when value distributions are asymmetric. What we need instead is for $t_{i}$ to depend on the value distribution $D_{i}$. We let $t_{i}$ scale inverse proportionally to $r_{i}$, so that highrevenue items are more likely to occur in the tail. This allows us to capture scenarios in which revenue comes primarily from one heavy item (by analyzing the tail), as well as instances driven by the combined contribution of many light items (by analyzing the core). Indeed, note that if we set $t_{i}=c r / r_{i}$, then the boundary between core and tail becomes $t_{i} r_{i}=c r=c \operatorname{SREV}(D)$ for each item. This turns out to be precisely the threshold that we need to attain constant-factor approximation bounds for both the core and the tail, simultaneously.

Theorem 2. For a single buyer, $\operatorname{REv}(D) \leq$ $6 \max \{\operatorname{SREv}(D), \operatorname{BREv}(D)\}$.

As in Theorem 1, our approach will be to apply the Core Decomposition Lemma (Lemma 6) with an appropriate choice of values $t_{i}$, then bound separately the revenue from the tail and the welfare from the core.

Proposition 3. For a single buyer, when $t_{i}=r / r_{i}$ for each $i, \sum_{A} p_{A} \operatorname{REv}\left(D_{A}^{T}\right) \leq 2 \operatorname{SREv}(D)$.

Proof: We begin similarly to the proof of Proposition 1, using Lemma 2 and Lemma 4 to write $\operatorname{REv}\left(D_{A}^{T}\right) \leq|A| \operatorname{SREv}\left(D_{A}^{T}\right) \leq \sum_{i \in A}|A| r_{i} / p_{i}$. Again, summing this over all $A$ yields:

$$
\sum_{A} p_{A} \operatorname{REV}\left(D_{A}^{T}\right) \leq \sum_{i} \sum_{A \ni i}|A| p_{A} r_{i} / p_{i} .
$$

Just like in Proposition 1, $\sum_{A \ni i}|A| p_{A} / p_{i}$ is exactly the expected number of items in the tail, conditioned on $i$ being in the tail. It's again clear that this sum is exactly $1+\sum_{j \neq i} p_{j}$. By Lemma 3, this is at most $1+\sum_{j \neq i} 1 / t_{j}$. By our choice of $t_{i}$, the second term is upper bounded by 1 , as $t_{j}=r / r_{j}$ and $\sum_{j} r_{j}=r$. Therefore, $\sum_{A \ni i}|A| p_{A} / p_{i} \leq 2$, and $\sum_{A} p_{A} \operatorname{REv}\left(D_{A}^{T}\right) \leq 2 \operatorname{SREv}(D)$.

We now turn to bounding the welfare from the core. We will use the small range of the core to derive an upper bound on the variance of its welfare. This will allow us to conclude that the welfare is highly concentrated whenever it is sufficiently large relative to $\operatorname{SREv}(D)$. Thus, if the welfare is "small" compared to $\operatorname{SREV}(D)$, then selling separately extracts most of the welfare (within the core); otherwise the welfare concentrates and so bundling extracts most of the welfare (within the core). The following lemma of [22] will be helpful for this approach; its proof appears in Appendix D of the full version.

Lemma 7. ([22]) Let $F$ be a one-dimensional distribution with optimal revenue at most c supported on $[0, t c]$. Then $\operatorname{var}(F) \leq(2 t-1) c^{2}$.

Corollary 1. For a single buyer, and any choice of $t_{i}$, $\operatorname{var}\left(D_{i}^{C}\right) \leq 2 t_{i} r_{i}^{2}$.

Proof: $\operatorname{REV}\left(D_{i}^{C}\right) \leq r_{i}$, and the distribution $D_{i}^{C}$ is supported on $\left[0, t_{i} r_{i}\right]$. Therefore, plugging into Lemma 7 (and relaxing) yields the desired bound.

Proposition 4. For a single buyer, when all $t_{i}=r / r_{i}$, $\max \{\operatorname{SREV}(D), \operatorname{BREV}(D)\} \geq \frac{1}{4} \operatorname{VAL}\left(D_{\emptyset}^{C}\right)$.

Proof: There are two cases to consider. If $\operatorname{VAL}\left(D_{\emptyset}^{C}\right) \leq 4 r$, then we have that $\operatorname{SREV}(D)=r \geq$ $\frac{1}{4} \operatorname{VAL}\left(D_{\emptyset}^{C}\right)$ as required.

On the other hand, if $\operatorname{VAL}\left(D_{\emptyset}^{C}\right) \geq 4 r$, then Corollary 1 tells us that $\operatorname{var}\left(D_{i}^{C}\right) \leq 2 t_{i} r_{i}^{2}$. Summing over all $i$ and recalling that $t_{i}=r / r_{i}$ we get

$$
\operatorname{var}\left(D_{\emptyset}^{C}\right)=\sum_{i} \operatorname{var}\left(D_{i}^{C}\right) \leq 2 \sum_{i} t_{i} r_{i}^{2}=2 r^{2}
$$

So $\operatorname{var}\left(D_{\emptyset}^{C}\right) \leq 2 r^{2}$ and $\operatorname{VAL}\left(D_{\emptyset}^{C}\right) \geq 4 r$. By Cheby-
shev's inequality, we get

$$
\begin{aligned}
\operatorname{Pr}\left[\sum_{i} v_{i} \leq \frac{2}{5} \cdot \operatorname{VAL}\left(D_{\emptyset}^{C}\right)\right] & \leq \frac{2 r^{2}}{\left(1-\frac{2}{5}\right)^{2} \cdot \operatorname{VAL}\left(D_{\emptyset}^{C}\right)^{2}} \\
& \leq \frac{25 r^{2}}{72 r^{2}}=\frac{25}{72}
\end{aligned}
$$

Since $\operatorname{BREv}(D)$ is at least the revenue obtained by setting price $\frac{2}{5} \cdot \operatorname{VAL}\left(D_{\emptyset}^{C}\right)$ on the grand bundle, this implies $\operatorname{BREV}(D) \geq\left(\frac{2}{5} \cdot \operatorname{VAL}\left(D_{\emptyset}^{C}\right)\right) \cdot \frac{47}{72}=\frac{47}{180} \cdot \operatorname{VAL}\left(D_{\emptyset}^{C}\right)$. As $\frac{47}{180}>\frac{1}{4}, \operatorname{BREV}(D)>\frac{1}{4} \operatorname{VAL}\left(D_{\emptyset}^{C}\right)$ as required.

Combining Propositions 3 and 4 with Lemma 6 yields Theorem 2. To our knowledge, the best known lower bound on max\{SREV, BREV\} vs. REV is 1.05 , provided by an example in [15]. The example has two items with $D_{1}=U\{1,2\}$ and $D_{2}=U\{1,3\}$, with $\operatorname{REV}=2.625, \operatorname{SREV}=2.5$, and $\mathrm{BREV}=2.25$. It is an interesting open question to close the gap between 1.05 and 6 , either by tightening our analysis or providing better lower bounds.

## V. Revenue Bounds for Multiple Buyers

Here we extend our results to multiple buyers with valuations sampled independently (but not necessarily identically). We will refer to this as the independent setting, as the buyers' valuations are independent and furthermore each buyer's item values are also drawn independently. We first show in Theorem 3 that for the independent setting, selling items separately achieves a logarithmic (in $n$ ) approximation to the optimal revenue. We next show in Theorem 5 that like in the single buyer case, the only case in which selling items separately fails to achieve a good approximation, is the case that welfare is highly concentrated. Unfortunately, such concentration is no longer sufficient to achieve a constant approximation by selling all items together. This is so because even though the welfare is concentrated, the partition that provides such welfare can change dramatically between realizations. Indeed, in Proposition 8 we show not only that $\operatorname{BREv}(D)$ fails to provide a constant approximation to the optimal mechanism, but even $\operatorname{PREv}(D)$ fails, and this is so even when item values are sampled i.i.d. for all items and buyers. Finally, in Proposition 9 we show that in the independent setting, $\operatorname{PREv}(D)$ cannot be approximated well by $\max \{\operatorname{SREv}(D), \operatorname{BREv}(D)\}$.

## A. An Upper Bound: $(\ln n+6)$ SREV $\geq$ REV

We first show that selling items separately achieves a logarithmic (in $n$ ) approximation to the optimal revenue.

Theorem 3. For arbitrarily many buyers, in the independent setting, $\left(2+2 e^{1 / 4}+\ln 4+\ln n\right) \operatorname{SREV}(D) \geq$ $\operatorname{REV}(D)$. (Note that $2+2 e^{1 / 4}+\ln 4<6$.)

Our proof will proceed via amplification. We will begin with the (awful) bound on SREV vs. REv from Lemma 2, then show in Theorem 4 how to amplify any such bound into an improved bound. We will then iterate this amplification process over and over, until we reach the desired logarithmic approximation (which will be a fixed point of the amplification process). To prove the amplification theorem, we use an approach similar to the single-buyer analysis from Section IV-A. That is, we will apply the Core Decomposition Lemma (Lemma 6 ), then bound the revenue of the tail and the welfare of the core with respect to $\operatorname{SREV}(D)$.

Theorem 4 (Amplification). For arbitrarily many buyers in the independent setting, assume that for some $a>1$ it holds that $\operatorname{anSREv}(D) \geq \operatorname{REv}(D)$. Then, for any $c \geq 1,\left(2+2 e^{1 / c a} / c+\ln c+\ln a+\ln n\right) \operatorname{SREV}(D) \geq$ $\operatorname{REv}(D)$ as well. Setting $c=1$ yields $\left(2+2 e^{1 / a}+\ln a+\right.$ $\ln n) \operatorname{SREv}(D) \geq \operatorname{REv}(D)$.

To prove Theorem 4, we will apply the Core Decomposition Lemma (Lemma 6), using $t_{i}=c \cdot a \cdot n$ for each $i$. Theorem 4 will then follow from bounds on the revenue from the tail and the expected welfare from the core, which we establish in the following propositions. The proof of Proposition 5 appears in Appendix E of the full version, and is similar to that of Proposition 1.

Proposition 5. For arbitrarily many buyers in the independent setting, if $t_{i}=c \cdot a \cdot n$ for all $i$ and $\operatorname{anSREv}(D) \geq \operatorname{REv}(D)$, then $\sum_{A} p_{A} \operatorname{REv}\left(D_{A}^{T}\right) \leq$ $\left(1+2 e^{1 / c a} / c\right) \operatorname{SREv}(D)$.

The following bound on the welfare from the core follows in a manner similar to Proposition 2. We defer its proof to Appendix E of the full version.

Proposition 6. For arbitrarily many buyers in the independent setting, if $t_{i}=c \cdot a \cdot n$ for all $i$, then $(1+\ln c+\ln a+\ln n) \operatorname{SREV}(D) \geq \operatorname{VAL}\left(D_{\emptyset}^{C}\right)$.

Theorem 4 then follows from Propositions 5 and 6, together with Lemma 6. We now show how to prove Theorem 3 using Theorem 4.

Proof (of Theorem 3): By Lemma 2, we may apply Theorem 4 starting with $a=m$. This yields a bound of the form $a^{\prime} n \operatorname{SREV}(D) \geq \operatorname{REV}(D)$ for some new $a^{\prime}$. We can then apply Theorem 4 again, taking $a$ to be this new value $a^{\prime}$. We can iteratively apply Theorem 4 over and over until we either reach a fixed point (with respect to the value of $a$ ) or reach $a=1$. One can verify that, for all $n \geq 2$, no $a \geq 4$ is a fixed point and that the function $f(a)=\left(2+2 e^{1 / a}+\ln a+\ln n\right) / n$ is continuous. Therefore, we can always iterate until $a \leq 4$ and then apply Theorem 4 with $a=4$, yielding the desired bound.

## B. A Concentration Result

We next present a characterization of when $\operatorname{SREv}(D)$ is a constant-factor approximation to $\operatorname{REV}(D)$ for the independent setting with multiple buyers. We will show (in Theorem 5, below) that this occurs unless the welfare of $D$ is sufficiently well concentrated around its expectation.

We begin with a corollary of Theorem 3, which will be useful for our analysis.

Corollary 2. For arbitrarily many buyers in the independent setting, $4 n \operatorname{SREv}(D) \geq \operatorname{REV}(D)$.

Proof: This is a direct application of Theorem 3 and noting that $6+\ln n \leq 4 n$ for all $n \geq 2$.

We next prove an alternative bound on the revenue from the tail of the distribution $D$, using a familiar choice of $t_{i}$. The proof, which closely follows that of Proposition 3, appears in Appendix E of the full version.

Proposition 7. For arbitrarily many buyers in the independent setting, if we choose $t_{i}=4 r / r_{i}$ for all $i$, then $\sum_{A} p_{A} \operatorname{REv}\left(D_{A}^{T}\right) \leq 5 e^{1 / 4} \operatorname{SREv}(D)$.

We are now ready to establish the claimed bound between SREv and REv, subject to the welfare of $D$ not being too concentrated around its expectation.

Definition 1. We say that a one-dimensional distribution $F$ is d-concentrated if there exists a value $C$ such that $\operatorname{Pr}_{x \sim F}[|x-C| \leq C / 2] \geq d$.

Theorem 5. For arbitrarily many buyers in the independent setting, and any $c \geq 4 \sqrt{2}$, either $(c+$ $\left.5 e^{1 / 4}\right) \operatorname{SREV}(D) \geq \operatorname{REv}(D)$ or the welfare of $D$ (the random variable with expectation $\operatorname{VAL}(D)$ ) is $(3 / 4-$ $\frac{24}{c^{2}}$ )-concentrated.

Proof: Let all $t_{i}=4 r / r_{i}$. Then combining Proposition 7 and Lemma 6 yields

$$
5 e^{1 / 4} \operatorname{SREV}(D)+\operatorname{VAL}\left(D_{\emptyset}^{C}\right) \geq \operatorname{REv}(D)
$$

There are two cases to consider. Maybe $c \operatorname{SREv}(D) \geq$ $\operatorname{VAL}\left(D_{\emptyset}^{C}\right)$. In this case, we have $\left(c+5 e^{1 / 4}\right) \operatorname{SREV}(D) \geq$ $\operatorname{Rev}(D)$.

On the other hand, maybe $\operatorname{VAL}\left(D_{\emptyset}^{C}\right) \geq c \operatorname{SREV}(D)$. In this case, Corollary 1 tells us that $\operatorname{var}\left(D_{i}^{C}\right) \leq 2 t_{i} r_{i}^{2}$. Summing over all $i$ and recalling that $t_{i}=4 r / r_{i}$, we get

$$
\operatorname{var}\left(D_{\emptyset}^{C}\right) \leq 2 \sum_{i} t_{i} r_{i}^{2}=2 \sum_{i}(4 r) r_{i}=8 r^{2}
$$

So $\operatorname{var}\left(D_{\emptyset}^{C}\right) \leq 8 r^{2}$ and $\operatorname{VAL}\left(D_{\emptyset}^{C}\right) \geq c r$. By Cheby-
shev's inequality, we get

$$
\begin{aligned}
& \operatorname{Pr}\left[\left|\sum_{i} v_{i}^{*}-\operatorname{VAL}\left(D_{\emptyset}^{C}\right)\right| \geq \operatorname{VAL}\left(D_{\emptyset}^{C}\right) / 2\right] \\
& \leq \frac{8 r^{2}}{\operatorname{VAL}\left(D_{\emptyset}^{C}\right)^{2} / 4} \leq \frac{32 r^{2}}{c^{2} r^{2}}=\frac{32}{c^{2}}
\end{aligned}
$$

meaning that the welfare of $D_{\emptyset}^{C}$ is $\left(1-\frac{32}{c^{2}}\right)$-concentrated. The last step is observing that $\vec{v}$ is sampled in the support of $D_{\emptyset}^{C}$ with probability exactly $\prod_{i}\left(1-p_{i}\right)$. As $\sum_{i} p_{i} \leq 1 / 4$ and each $p_{i} \leq 1 / 4$, this is minimized when exactly one $p_{i}$ is $1 / 4$ and the rest are 0 , yielding $\prod_{i}\left(1-p_{i}\right)=3 / 4$. So with probability at least $3 / 4 \vec{v}$ is in the support of $D_{\emptyset}^{C}$. When this happens, the welfare is $\left(1-\frac{32}{c^{2}}\right)$ concentrated. So the welfare of $D$ is $\left(3 / 4-\frac{24}{c^{2}}\right)-$ concentrated.

## C. A Lower Bound: PREv $\leq \operatorname{REv} / \Omega(\log n)$ even for

 i.i.d. Item ValuesWe next show that there is a setting with many buyers with item valuations that are sampled i.i.d from the same distribution, for which $\operatorname{PREv}(D)$ (and thus also $\max \{\operatorname{SREV}(D), \operatorname{BREV}(D)\})$ provides a poor approximation to $\operatorname{Rev}(D)$.

Proposition 8. There exists a setting with $n$ items and many buyers, with item valuations that are sampled i.i.d from the same distribution, for which $\operatorname{PREv}(D) \leq$ $\operatorname{REv}(D) / \Omega(\log n)$.

Proof: Consider a setting with $n$ items and $\sqrt{n}$ buyers with the following value distributions. For every item $i$ and buyer $j$, the distribution $D_{i, j}$ such that the value is 0 with probability $1-1 / \sqrt{n}$, and with the remaining probability it is sampled from a distribution $F$ with CDF $F(x)=1-x^{-1}$ for $x \in\left[1, n^{1 / 8}\right]$ and $F(x)=1$ for $x>n^{1 / 8}$ (an Equal-Revenue distribution with all mass above $n^{1 / 8}$ moved to an atom at $n^{1 / 8}$ ). To prove the claim we show in Lemma 12 in Appendix E of the full version that $\operatorname{REv}(D) \in \Omega(n \log n)$ while $\operatorname{PREv}(D) \in O(n)$ (actually, since $\operatorname{SREV} \in \Omega(n)$ it holds that $\operatorname{PREV}(D) \in \Theta(n))$.

## D. A Lower Bound: max\{SREv, BREv\} $\leq$ $\operatorname{PREV} / \Omega(\log n)$

We next show that there is a setting with many buyers with item valuations that are sampled independently (but not identically), for which $\max \{\operatorname{SREv}(D), \operatorname{BREv}(D)\}$ provides a poor approximation to $\operatorname{PREv}(D)$.

Proposition 9. There exists a independent setting with $n$ items and many buyers for which $\max \{\operatorname{SREv}(D), \operatorname{BREv}(D)\} \leq \operatorname{PREv}(D) / \Omega(\log n)$.

Proof: Fix $n$ such that $\sqrt{n}$ is an integer. Consider a setting with $\sqrt{n}$ buyers, and a partition the items to
$\sqrt{n}$ disjoint sets of size $\sqrt{n}$. Buyer $k$ has value 0 for every item that is not in the $k$-th set of items, and for item in that set his value is sampled independently from an Equal-Revenue distribution.

Clearly, $\operatorname{SREv}(D)=n . \operatorname{BREv}(D)$ is the same as the revenue that $\operatorname{BREv}(D)$ can get in a setting with $\sqrt{n}$ buyers and only $\sqrt{n}$ items for which each item value is sampled i.i.d. from an Equal-Revenue distribution. That revenue is $O(\sqrt{n} \log \sqrt{n})$. We conclude that $\max \{\operatorname{SREv}(D), \operatorname{BREv}(D)\} \in O(n) . \operatorname{PREv}(D)$ on the other hand, can bundle each of the sets of size $\sqrt{n}$ separately and sell it to the interested buyer, getting a total revenue of $\sqrt{n} \cdot \Omega(\sqrt{n} \log \sqrt{n})=\Omega(n \log n)$.

## VI. One Buyer with Correlated Values

In this section, we study the relationship between $\operatorname{SREv}(\mathrm{D}), \operatorname{Max}\{\operatorname{SREv}(\mathrm{D}), \operatorname{BREv}(\mathrm{D})\}$, and $\operatorname{PREv}(\mathrm{D})$ for a single buyer with correlated values. The prior work of [17], [7] already shows that there is no hope of obtaining a finite bound between any of these quantities and $\operatorname{REV}(\mathrm{D})$ because they are all deterministic, even when there are only two items. But it is still important to understand the relationship between these mechanisms of varying complexity even if their revenue cannot compare to that of the optimal mechanism. We show in Theorem 6 that for any distribution $D$ for a single buyer, possibly even correlated, $\operatorname{SREV}(D)$ is a $O(\log n)$ approximation to $\operatorname{BREv}(D)$, and thus also to $\operatorname{Max}\{\operatorname{SREv}(\mathrm{D}), \operatorname{BREv}(\mathrm{D})\}$ and $\operatorname{PREv}(\mathrm{D}){ }^{10}$ We then show in Proposition 10 that this bound is tight, $\operatorname{Max}\{\operatorname{SREv}(\mathrm{D}), \operatorname{BREv}(\mathrm{D})\}$ $\leq \operatorname{PREv}(D) / \Omega(\log n)$. In other words, $\operatorname{SREV}(\mathrm{D})$ provides a logarithmic approximation to $\operatorname{PREV}(\mathrm{D})$, but taking max $\{\operatorname{SREV}(D), \operatorname{BREV}(D)\}$ can't guarantee anything better. The proof of Theorem 6 and Proposition 10 appear in Appendix $F$ of the full version.

Theorem 6. For any n-dimensional value distribution $D$ for a single buyer (possibly correlated across items), $\operatorname{BREv}(D) \leq 5 \ln (n) \operatorname{SREv}(D)$. Therefore, $\operatorname{PREV}(D) \leq$ $5 \ln (n) \operatorname{SREv}(D)$ as well.
Proposition 10. There exists a (correlated) distribution $D$ of the valuation of a single buyer over $n$ items for which $\max \{\operatorname{SREv}(D), \operatorname{BREv}(D)\} \leq$ $\operatorname{PREV}(D) / \Omega(\log n)$.

## VII. Acknowledgments

In an earlier version of this paper, we proved a factor of 7.5 in Theorem 2. This factor was later improved by Aviad Rubinstein to a factor of 6 . We thank Aviad for allowing us to include this improvement in our paper.

[^6]
## References

[1] Saeed Alaei. Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers. In the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2011.
[2] Saeed Alaei, Hu Fu, Nima Haghpanah, and Jason Hartline. The Simple Economics of Approximately Optimal Auctions. In the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2013.
[3] Saeed Alaei, Hu Fu, Nima Haghpanah, Jason Hartline, and Azarakhsh Malekian. Bayesian Optimal Auctions via Multi- to Single-agent Reduction. In the 13th ACM Conference on Electronic Commerce (EC), 2012.
[4] Lawrence M. Ausubel and Paul Milgrom. The lovely but lonely vickrey auction. In Combinatorial Auctions, chapter 1. MIT Press, 2006.
[5] Anand Bhalgat, Sreenivas Gollapudi, and Kamesh Munagala. Optimal Auctions via the Multiplicative Weight Method. Manuscript, http://www.seas.upenn.edu/ bhalgat/optimal_bic.pdf, 2013.
[6] Sayan Bhattacharya, Gagan Goel, Sreenivas Gollapudi, and Kamesh Munagala. Budget Constrained Auctions with Heterogeneous Items. In the 42nd ACM Symposium on Theory of Computing (STOC), 2010.
[7] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. Pricing randomized allocations. In SODA, pages 585-597, 2010.
[8] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. An Algorithmic Characterization of MultiDimensional Mechanisms. In the 44th Annual ACM Symposium on Theory of Computing (STOC), 2012.
[9] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization. In the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2012.
[10] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Reducing Revenue to Welfare Maximization: Approximation Algorithms and other Generalizations. In the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2013. http://arxiv.org/pdf/1305.4000v1.pdf
[11] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Understanding Incentives: Mechanism Design becomes Algorithm Design. In the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2013. http://arxiv.org/pdf/1305.4002v1.pdf.
[12] Shuchi Chawla, Jason D. Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In Proceedings of the 8th ACM Conference on Electronic Commerce, EC '07, pages 243-251, 2007.
[13] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In the 42nd ACM Symposium on Theory of Computing (STOC), 2010.
[14] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. The Power of Randomness in Bayesian Optimal Mechanism Design. In the 11th ACM Conference on Electronic Commerce (EC), 2010.
[15] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. The Complexity of Optimal Mechanism Design. In the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2014.
[16] Constantinos Daskalakis and S. Matthew Weinberg. Symmetries and Optimal Multi-Dimensional Mechanism Design. In the 13th ACM Conference on Electronic Commerce (EC), 2012.
[17] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. In Proceedings of the 13th ACM Conference on Electronic Commerce, EC '12, pages 656-656, 2012.
[18] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In Proceedings of the Fourteenth ACM Conference on Electronic Commerce, EC '13, pages 565-566, 2013.
[19] Sergiu Hart and Philip J. Reny. Maximal Revenue with Multiple Goods: Nonmonotonicity and Other Observations. Discussion Paper Series dp630, The Center for the Study of Rationality, Hebrew University, Jerusalem, November 2012.
[20] J.D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In EC'09, pages 225-234, 2009.
[21] Robert Kleinberg and S. Matthew Weinberg. Matroid Prophet Inequalities. In the 44th Annual ACM Symposium on Theory of Computing (STOC), 2012.
[22] Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. Proceedings of the National Academy of Sciences, 110(28):11232-11237, 2013.
[23] R. Myerson. Optimal auction design. Mathematics of Operations Research, 6(1):58-73, 1981.
[24] Gregory Pavlov. A Property of Solutions to Linear Monopoly Problems. B. E. Journal of Theoretical Economics, 11(1), 2011. Article 4.
[25] Pingzhong Tang and Zihe Wang. Optimal Mechanisms with Simple Menus. arxiv report, 2013. http://arxiv.org/abs/1311.5966.
[26] Andrew Chi-Chih Yao. An $n$-to-1 Bidder Reduction for Multi-Item Auctions and its Applications. arxiv report, 2014. http://arxiv.org/abs/1406.3278.


[^0]:    moshe@microsoft.com
    nicimm@ microsoft.com
    brlucier@microsoft.com
    smweinberg@csail.mit.edu. Supported by a Microsoft Graduate Research Fellowship and NSF CCF-6923736.
    ${ }^{1}$ This assumes regularity of the value distributions and that the buyers' values are drawn independently.

[^1]:    ${ }^{2}$ Note that if the seller has unlimited copies of each item for sale, then an auction for a single buyer directly extends to the case of multiple buyers.
    ${ }^{3}$ This simplicity again assumes regularity and independence.

[^2]:    ${ }^{4}$ The Equal-Revenue distribution has $\operatorname{CDF} F(x)=0$ for $x \leq 1$, and $F(x)=1-1 / x$ for $x \geq 1$.
    ${ }^{5}$ When the distributions are identical, and furthermore satisfy the Monotone Hazard Rate condition, [16] provides a PTAS. However, other recent results based on linear programming formulations ([1], [3], [2], [5], [8], [9], [10], [11]) all run in time polynomial in the support of $D$. In many correlated settings, this is the right runtime to shoot for, or the best one could hope for. But in our independent setting, this runtime will be exponential in $n$ when ideally we would like to run in time polynomial in $n$. We show that if we have meaningful access to the distributions in a way that allows us to compute the optimal per-item reserves efficiently, then our mechanism runs in polynomial time.

[^3]:    ${ }^{6}$ Clearly, no example can exhibit both gaps simultaneously as selling separately achieves an $O(\log n)$-approximation to the optimal revenue.
    ${ }^{7}$ In the common base-value model, the buyer has $n+1$ distributions $D_{0}, \ldots, D_{n}$, and samples $v_{i}$ from each $D_{i}$. Her value for item $i \in$ $\{1, \ldots, n\}$ is then $v_{0}+v_{i}$, and $v_{0}$ is called the "base-value."

[^4]:    ${ }^{8}$ As it turns out, all of the mechanisms we describe will also satisfy the stronger property of dominant strategy truthfulness.

[^5]:    ${ }^{9}$ This observation is due to Aviad Rubinstein, and we thank him for allowing us to include it. An earlier version of this paper presented a $(\ln n+5)$-approximation in Theorem 1 and a 7.5 -approximation in Theorem 2. This observation improved those factors to $(\ln n+3)$ and 6 , respectively.

[^6]:    ${ }^{10}$ As SREV approximates BREV for any set of items, it can do so for any part in the partition in PREV separately, and thus also approximate PREV.

