99% Revenue via Enhanced Competition

MICHAL FELDMAN, Computer Science, Tel-Aviv University, and Microsoft Research., Israel
OPHIR FRIEDLER, Computer Science, Tel-Aviv University., Israel
AVIAD RUBINSTEIN, Harvard University., USA

A sequence of recent studies show that even in the simple setting of a single seller and a single buyer with additive, independent valuations over m items, the revenue-maximizing mechanism is prohibitively complex. This problem has been addressed using two main approaches:

- *Approximation:* the best of two simple mechanisms (sell each item separately, or sell all the items as one bundle) gives 1/6 of the optimal revenue [1].
- *Enhanced competition:* running the simple VCG mechanism with additional m buyers extracts at least the optimal revenue in the original market [17].

Both approaches, however, suffer from severe drawbacks: On the one hand, losing 83% of the revenue is hardly acceptable in any application. On the other hand, attracting a linear number of new buyers may be prohibitive. We show that by combining the two approaches one can achieve the best of both worlds. Specifically, for any constant ϵ one can obtain a $(1 - \epsilon)$ fraction of the optimal revenue by running simple mechanisms — either selling each item separately or selling all items as a single bundle — with substantially fewer additional buyers: logarithmic, constant, or even none in some cases.

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

Additional Key Words and Phrases: Mechanism Design, Revenue, Simple Mechanisms, Pricing

1 INTRODUCTION

The scenario of a buyer with an additive, independent valuation over m items has become the paradigmatic setting for studying optimal (revenue-maximizing) mechanisms. In this setting, the buyer's valuation is drawn from a distribution D that is known to the seller, and the seller wishes to design a selling mechanism that extracts as much revenue as possible.

By now it is well known that the optimal mechanism requires randomization [22], infinite, uncountable menus [16], is non-monotone [23], and computationally intractable [15][1]. Thus it is mostly interesting as a theoretic benchmark to which one can compare more plausible mechanisms (much like the way an offline optimum serves as a benchmark in online settings). In recent years,

This work was partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122, Israel Science Foundation (grant number 317/17), and a Rabin Postdoctoral Fellowship.

[1]Similar undesired properties have been observed with respect to the optimal mechanism in additional related (multi-dimensional) models, e.g., unit-demand buyers [3, 13, 14, 31]

Authors' addresses: Michal Feldman, Computer Science, Tel-Aviv University, and Microsoft Research. Tel-Aviv, Israel, michal.feldman@cs.tau.ac.il; Ophir Friedler, Computer Science, Tel-Aviv University. Tel-Aviv, Israel, ophirfriedler@gmail.com; Aviad Rubinstein, Harvard University. Cambridge, MA, USA, aviad@seas.harvard.edu.

ACM EC'18, June 18–22, 2018, Ithaca, NY, USA. ACM ISBN 978-1-4503-5829-3/18/06…$15.00
https://doi.org/10.1145/3219166.3219202

two main approaches have been taken with respect to this challenge, both of which proposed simple mechanisms and measured their performance against the theoretic optimum:

The first line of work approaches this problem through the lens of *approximation* [1, 2, 6, 8–12, 18, 21, 26, 27, 29, 34–36]. In particular, the breakthrough result of [1] shows that the better of two simple mechanisms — selling each item separately or selling all items together in a grand bundle — obtains at least 1/6 (but at most 1/2 [33]) of the optimal revenue.

The second approach is to enhance the competition for the merchandise by increasing the population of potential buyers [4, 17, 28, 32]. This approach is often termed resource augmentatation in computer science. The state of the art for additive buyers is by Eden *et al.* [17], who showed that adding m additional buyers is sufficient to recover the original optimal revenue with a simple mechanism. It is also known that at least $\Omega(\log m)$ additional buyers are required to achieve this benchmark. This result from [17] generalizes the seminal work of Bulow and Klemperer [4] who showed that for a market with a single item, under a regularity assumption, running the second price auction with one additional buyer extracts at least as much revenue as the original optimal revenue. However, when m is large, adding m additional buyers may be prohibitive. We therefore believe that the most interesting question here is whether the linear dependence on m is necessary.

To summarize, the "optimal" benchmark is intractable. The approximation approach is stuck at a 1/6-approximation (forfeiting 83% of potential revenue). And if we wish to follow the enhanced competition approach, the best known bound on the number of additional buyers is linear in the number of items.

1.1 Our Contribution

We show that one can combine the two approaches (of approximation and enhanced competition) in a way that achieves the best of both worlds. We establish a host of results for various settings, but they all convey one theme: in order to obtain revenue that is very close to optimum, there is no need to recruit a linear number (in m) of additional buyers; that is:

Main take away (informal): *A seller can obtain 99% of the optimal revenue in a simple mechanism (selling each item separately or selling all items together in a singe bundle) with substantially fewer additional buyers — logarithmic (in m), constant, or even none in some cases.*

The constant 99% above is arbitrary of course. More generally, we show what we think of as the **mechanism-design analog of computational PTAS**: an efficient mechanism that obtains a $(1-\varepsilon)$-approximation, for any constant $\varepsilon > 0$ (where in our case *efficient* means "few additional buyers" instead of "polynomial time").

All of our results apply to the paradigmatic scenario of a seller who sells m items to a single buyer with additive, independent valuations over the items, with a known prior. Some of our results extend to the more general scenario of n i.i.d. buyers, namely where buyers' values are drawn independently according to the same distribution. The induced product distribution is then denoted by $\mathcal{D}$. Our first set of results consider the simple mechanism that sells each item separately.

Theorem 1.1. *For every constant $\varepsilon > 0$, selling each item separately to $O(\log m)$ i.i.d. buyers extracts at least $(1-\varepsilon)$-fraction of the optimal revenue achievable by a single buyer.*

This result improves upon the $O(m)$ bound shown in [17], at the loss of ε fraction in revenue. In fact, this result can be extended even to a setting with n i.i.d. buyers and m items. The following theorem fully characterizes, up to constant factors, the number of additional buyers necessary as a function of the number of items.

Theorem 1.2. *(implies Theorem 1.1) For a setting with n i.i.d. buyers and m items, for every constant $\varepsilon > 0$, selling each item separately achieves at least $(1-\varepsilon)$-fraction of the optimal revenue if we* *increase the number of buyers by a factor of $O(\log(2 + m/n))$, and this is tight up to constant factors. Moreover, if $m = o(n)$, then selling each item separately achieves $(1-\varepsilon)$-fraction of the optimal revenue even with no additional buyers.*

Theorem 1.2 essentially fully characterizes (up to constant factors) the number of additional buyers necessary for achieving $(1 - \varepsilon)$ of the optimal revenue. In particular, we consider three different regimes of m and n: For $m = \omega(n)$ we prove that increasing the number of buyers by a factor of $O(\log(m/n))$ is both necessary (Theorem 4.12) and sufficient (Theorem 4.1). For $m = \Theta(n)$, our new lower bound implies that the previous results of [17] (who showed that a linear number of additional buyer suffice) are essentially tight. Finally, for $m = o(n)$, we show that *no* additional buyers are necessary (Theorem 5.1). We note that our lower bound generalizes the special case of $\Omega(\log m)$-factor for the case of a single buyer in [17] [2].

Let us return to the single buyer setting. Theorem 1.2 implies that we can recover $(1-\varepsilon)$-fraction of the optimal revenue by adding $O(\log m)$ buyers. However, one may argue that for a large value of m, $O(\log m)$ is still too large. We address this issue by showing that the better of selling items separately and selling the grand bundle requires only a constant number of additional buyers[3].

THEOREM 1.3. *For every constant $\varepsilon > 0$, the better of selling each item separately and selling the grand bundle, to a constant number of i.i.d. buyers, extracts at least $(1 - \varepsilon)$-fraction of the optimal revenue achievable from one buyer.*

Up until now, we concentrated on *prior-dependent* mechanisms; namely, mechanisms that use the knowledge of the distribution of values. Our work, however, contributes also to the literature on *prior-independent* mechanisms. As in previous literature on prior-independent mechanisms, to achieve any meaningful result, we assume that the underlying single-dimensional distributions are regular (note that this does not generally imply regularity of the grand bundle's distribution). Since bidders are additive, the prior-independent VCG mechanism simply runs the second price auction for each item simultaneously. Thus, Theorem 1.2 combined with the original result of Bulow and Klemperer [4] immediately implies the following corollary:

COROLLARY 1.4. *If $\mathcal{D}$ is a product of regular distributions, then for every constant $\varepsilon > 0$, running the VCG mechanism with a factor $O(\log(2 + m/n))$ increase in the number of buyers (and with no additional buyers in the case of $m = o(n)$) extracts at least $(1 - \varepsilon)$-fraction of the original optimal revenue.*

It would be highly desirable to obtain such an analog to Theorem 1.3. An immediate barrier, however, is that the seller must know in advance whether to sell the items separately or sell the grand bundle. How can the seller determine which one of these mechanisms to run in the absence of a prior? This barrier is overcome by the surprising result that when $\mathcal{D}$ is a product of regular distributions, the seller *never* needs to sell items separately; selling the grand bundle is always the "correct" strategy:

THEOREM 1.5. *If $\mathcal{D}$ is a product of regular distributions, then for every constant $\varepsilon > 0$, selling the grand bundle in a second price auction to a constant number of i.i.d. additive buyers extracts at least $(1 - \varepsilon)$-fraction of the optimal revenue achievable from one buyer.*

[2]The lower bound in [17] was proven for mechanisms that target 100% of the optimal revenue, but it can be easily extended to mechanisms that target 99% or any constant fraction.

[3]There is no contradiction to the $\Omega(\log m)$ lower bound, which applies only for selling items separately.

1.2 The Case for 99% AGT

Constant factor approximations, and the heavy use of asymptotics in general, are one of the distinctive characteristics of Algorithmic Game Theory (in contrast to "standard" Economics), and we are by no means innocent. Note that in all our theorem statements hidden constants are lurking, and we do not even attempt to explicitly state them (obtaining reasonable constants remains an interesting open problem). We hope, as we always do when we make such asymptotic analysis, that our mathematically precise models tell an insightful story about how our guarantees degrade when we extend our single item / neighborhood grocery store intuition to auctioneers selling a huge number of items[4], especially online. (And this goes without mentioning our "standard" but absurd assumptions: rational bidders, independent valuation distributions, etc.)

That being said, we argue that some constants are more important than others, and constant factor approximation of the revenue are particularly problematic. Consider the 1/6-approximation by [1]. This is an extremely interesting result, but ultimately, losing 83% (or even the most optimistic 50%) of the revenue doesn't actually *solve* the auctioneer's mechanism problem[5]. On the other hand, a seller might be willing to compromise on optimality if guaranteed 99% of the optimal revenue. (For example, merchants around the world pay small fees to credit card companies in return for simple selling mechanisms). We therefore argue that in general the most interesting agenda, in 99% of mechanism design problems, should be obtaining 99% of the optimal revenue.

1.3 Relation to work on approximate revenue maximization

As already mentioned, our work is related to and inspired by a long line of research that aims at understanding what fraction of the optimal revenue can be guaranteed with simple mechanisms (without adding buyers), including [1, 20, 27] and references therein. It is interesting to note that obtaining a 0.99-approximation to the optimal revenue with k additional buyers has implications on approximation results without additional buyers. In particular, if one can obtain a 0.99-approximation to the optimal revenue using mechanism $\mathcal{M}$ with k additional buyers, then, by revenue submodularity, this immediately implies a $0.99/(k+1)$-approximation of the optimal revenue with a single buyer. Hence, our Theorem 1.1 implies the result of [27] that selling each item separately is an $\Omega(1/\log m)$-approximation of the optimal revenue. Similarly, our Theorem 1.3 implies the main result of [1] that the better of selling separately and selling the grand bundle yields a constant fraction of the optimal revenue. Of course, the reverse implication is not true, since the seller only has a single copy of each item to allocate to all buyers. Therefore, while our analysis builds on the techniques of [1, 27], it is significantly more challenging.

Recently, Goldner and Karlin [20] built on the result of Babaioff *et al.* [1], and showed that for regular distributions, the prior-independent mechanism which sells either each item separately or the grand bundle, and sets the price according to a single sample from the valuations distribution, obtains a constant fraction of the optimal revenue. As a corollary of our Theorem 1.5, we can obtain the following qualitative strengthening of [20, Corollary 2] for the case of a single buyer:

Corollary 1.6. *If $\mathcal{D}$ is a product of regular distributions, taking a single sample of the value of the grand bundle and selling the grand bundle for that price (to a single buyer) extracts at least a constant fraction of the optimal revenue.*

[4] For example, in 2012 Google was estimated to serve an average of 29,741,270,774 (≈ 30 billion) impressions per day[25].
[5] There are still many great reasons to study constant-factor approximations in mechanism design; see, e.g., [24] for an excellent discussion.

For the special case where the single-dimensional distributions satisfy the monotone hazard rate (MHR) condition[6], Cai and Huang [7] gave a PTAS for the revenue maximizing auction. En route to obtaining their computational result, they prove the following structural lemma, which holds even for heterogeneous buyers. 99% of the optimal revenue can be obtained by one of the following simple mechanisms: (i) selling every item separately; or (ii) selling all but a constant number of items via a VCG-like mechanism. In other words, they show that when the single-dimensional distributions satisfy the MHR condition, no additional buyers are necessary to obtain 99% of the optimal revenue with mechanisms that are simple (in the above sense). Moreover, inspired by [5], they show that for MHR distributions, when the number of buyers is larger than some constant, selling items separately obtains 99% of the *social welfare*. (In contrast, note that for regular distributions the ratio of social welfare to revenue may be unbounded, even when adding any number of buyers.)

1.4 Organization

In Section 2 we provide a succinct overview of the techniques for all our results. Model definitions and other preliminaries appear in Section 3. An abridged version of the proof of Theorem 1.2 appears in the following two sections: Section 4 deals with the regime where the number of items is greater than the number of buyers, whereas the regime of more buyers than items is considered in Section 5. The full version of this paper with detailed proofs of all theorems is available at https://arxiv.org/abs/1801.02908.

2 OVERVIEW OF PROOFS

Our techniques build upon the now-standard approach for approximately optimal revenue analysis: Separately reason about revenue contribution from rare events of extremely high value ("tail" events), and revenue contribution from lower values ("core" events). This approach is made formal via the core-tail decomposition framework of Li and Yao [27]. Separation between core and tail events is done by setting a cutoff for each item, and then proving a "core-tail decomposition" lemma that upper bounds the optimal revenue by the sum of total value from core events in which items are below their cutoffs (a.k.a. contribution from the core), and the total revenue from tail events in which items are above their cutoffs (a.k.a. contribution from the tail).

For approximation results, the next step would typically be to show that simple mechanisms approximate both the contribution from the core and the tail, hence the best simple mechanism approximates the optimal revenue. However, when targeting 99% of the optimal revenue, even if one shows that simple mechanisms fully recover the contribution from the core and the tail, it does not yet imply that simple mechanisms recover the *sum* of the contributions. Hence this approach alone cannot yield better than a 1/2-approximation. Instead, we carefully set the cutoffs so that the contribution from the core is almost fully recovered using simple mechanisms with either $O(1)$ buyers (Theorems 1.3 and 1.5) or $O(\log m)$ buyers (Theorem 1.2), and the contribution from the tail is only a tiny fraction of the revenue from simple mechanisms with $O(1)$ buyers[7].

Before delving into the specificities of each result, we provide a few general notes about the tail and the core.

Tail. Since tail events are rare, they are approximately "separable" across items and across buyers. First, since the probability of two or more items in the tail of any buyer is low, the revenue contribution from tail items is roughly "separable" across items, i.e., approximating revenue contribution of tail events by the revenue contribution from selling items separately loses only a moderate factor.

[6]Roughly speaking, MHR distribution (a special case of regular distributions) has a tail that is thinner than that of an exponential distribution.

[7]In other words, for heavy-tailed distributions, adding a few buyers *increases* the revenue by a lot.

Furthermore, each buyer is likely to have tail valuations for disjoint subsets of items ("separability across buyers"). Thus in an enhanced competition setting with more buyers we can simultaneously serve all of them, quickly increasing the revenue.

Core. Since the value of core items is bounded, a concentration bound typically suggests a near-optimal grand bundle price whenever the sum of item values is significantly larger than the revenue from selling items separately.

In what follows we elaborate on cases where interesting artifacts arise in the analysis. We use Rev, SRev, and BRev to denote the optimal revenue from any mechanism, the optimal revenue from selling each item separately, and the optimal revenue from selling the grand bundle, respectively. [8]

Theorem 1.3 ($\max\{SRev, BRev\}$). An optimistic approach for proving Theorem 1.3 would use the same cutoff for all items, and apply a concentration bound (Chebyshev's inequality) to suggest a bundle price that almost fully recovers the contribution from the core with constant probability, then improve this probability to almost 1 by considering additional buyers. Unfortunately, such strong concentration does not hold in general: a small number of items may still exhibit a large variance, even inside the core. To overcome this challenge, we separate the set of items to items that exceed their cutoffs with constant probability, which we call "high items", and the remaining we call "low items". We then use a concentration bound only for the sum of low items to get a good bundle price for the low items. We set the cutoff so that the number of high items is a constant, hence with probability close to 1, one of our additional $O(1)$ buyers simultaneously exceeds the cutoff of all the "high items". We conclude that BRev with $O(1)$ buyers recovers almost all of the contribution from the core with probability almost 1.

Theorem 1.5 (BRev, regular distributions). Single dimensional regular distributions are appealing since they have a "small tail" property. While the grand bundle distribution need not be regular even when the individual item distributions are, we show that the underlying regularity of the individual items still maintains some well behaved properties that we can exploit. Specifically, we can set cutoffs so that in the resulting core-tail decomposition, the contribution from the tail is significantly smaller than the contribution from the core. This is important since we are only allowed to sell the grand bundle, while the tail is typically covered by selling items separately.

It is now left to guarantee that the core is almost fully recovered. The challenge is that a small number of outlier items may have very large variance, ruining the naive concentration bound argument. (In Theorem 1.3 we sometimes sold the outliers separately, but here we are only allowed to sell the grand bundle.) We show that given a large (but still constant) number of additional buyers, two of them are likely to have high values simultaneously for all of the outlier items. We can then use a concentration bound to suggest a good price for the remaining items, so that VCG on the grand bundle gives high revenue.

Theorem 1.2 (SRev), $m = \omega(n)$ regime. Instead of a concentration bound for the grand bundle (which is irrelevant when considering SRev) we further separate the contribution from the core to the contributions from lower and higher values. We then show that the contributions of lower values to the core can be almost fully recovered with probability almost 1 by SRev with $O(n \cdot \log(m/n))$ buyers, while the contributions of higher values to the core form only a tiny fraction of the revenue that can be extracted by SRev with $O(n \cdot \log(m/n))$ buyers. The proof appears in Section 4.

[8] SRev and BRev essentially perform Myerson's mechanism either on each item separately or on the grand bundle respectively.

Theorem 1.2 (SRev), $m = o(n)$ regime. We show that selling each item separately achieves 99% of the optimal auction without adding *any buyers at all*. Proving this result requires yet new ideas. The intuition is simple: with so many buyers, we can afford to set the cutoff much higher – so high that *most buyers* have at most one item in the tail – while the contribution from the core can be almost fully recovered by selling each item to a tiny fraction of the buyers. Therefore, we can set aside a small subset of special buyers, and offer the items at high (tail) prices to all other buyers; we can then recover the rest of the revenue by auctioning the items not previously sold to the special buyers. (Note that this mechanism is for analysis purposes only — once we establish guarantees for any mechanism for selling items separately, we can use Myerson's optimal mechanism for selling items separately.) This part of the analysis also takes into account the probability that an item was sold in tail events.

Formalizing the above intuition is quite subtle. With high probability, there are still *some* buyers with multiple items in the tail – and each of those items has many other buyers whose valuations are in the tail, etc. Thus even after the core-tail decomposition, we have to reason about a multiple buyer, multiple item setting. To cope with this difficulty, we consider a bipartite graph of buyers and items, where we draw an edge $\{i, j\}$ whenever buyer i's value of item j is in the tail. We then argue that we can analyze the revenue from each connected component separately. In particular, while there are, w.h.p., large components in the graph (with $\omega(1)$ buyers and items), we use simple ideas from percolation theory to bound the expected size of each connected component. The proof appears in Section 5.

3 MODEL AND PRELIMINARIES

We consider a setting in which a monopolist seller sells a set $[m] = \{1, 2, \ldots, m\}$ of heterogeneous items to n additive buyers. Buyer i's value is additive if there exist values $v_1^i, \ldots, v_m^i$ such that buyer i's value for a set of items A is $\sum_{j \in A} v_j^i$. The seller does not know the buyers' values, but knows the distribution from which they are sampled. For every item j, the value of each agent i for item j is independently sampled from a single-dimensional distribution D_j.

A mechanism is given by a pair of an allocation function π, and a payment function $\mathbf{p}$. The mechanism receives a valuation profile $\mathbf{v} = \{v_j^i\}_{i,j}$ as input. Based on $\mathbf{v}$, the allocation function determines the (possibly random) allocation of items to buyers, and the payment function determines the payment $\mathbf{p}^i$ of every buyer i. Buyers are quasi-linear; namely, buyer i's utility from a mechanism that allocates her each item j with probability π_j^i and charges her a payment $\mathbf{p}_i$ is $\sum_j \pi_j^i \cdot v_j^i - \mathbf{p}^i$.

A mechanism is Bayesian Individually Rational (BIR) if every buyer's expected utility (over the randomness of the mechanism and other buyers' values, assuming they are drawn from $\times_j D_j$) is non-negative. A mechanism is Bayesian Incentive Compatible (BIC) if every buyer's expected utility is maximized when the buyer reports her valuation truthfully. Throughout the paper, a BIR-BIC mechanism is termed a *truthful* mechanism. The expected revenue of a truthful mechanism is $\mathbb{E}\left[\sum_i \mathbf{p}^i(\mathbf{v})\right]$, where for every buyer i and item j buyer i's value for item j is drawn independently from D_j. The *optimal revenue* is the optimal[9] expected revenue among all truthful mechanisms.

For any single dimensional distribution D and any $q \in [0, 1]$, we assume w.l.o.g. that there exists $v(q) \in \mathbb{R}$ such that $\Pr_{x \sim D}[x \geq v(q)] = q$. When the distribution is continuous this is true by the intermediate value theorem. When the distribution has point masses, we can smooth it with an infinitesimal perturbation; see e.g., [34] for a formal discussion. We note that when D is regular, the perturbed distribution may not be regular. However, it will not be hard to see that our proof in

[9]In general for distributions of infinite support, the optimum revenue may not be achieved by any mechanism (i.e. it is the supremum of all revenues achievable by truthful mechanisms). This is not so important for our purposes since we are only trying to approximate the optimum.

Section ?? will also work with distributions that are infinitesimally-close to regular distributions. For a set of items $A \subseteq [m]$, let $\bar{A} = [m] \setminus A$. We use $\mathbf{v}_A^i = \{\mathbf{v}_j^i\}_{j \in A}$, therefore $\mathbf{v}^i = (\mathbf{v}_A^i, \mathbf{v}_{\bar{A}}^i)$. Also, we use $\mathbf{v}_j = \{\mathbf{v}_j^i\}_{i \in [n]}$, and more generally, for a set $S \subseteq [n] \times [m]$, we use $\mathbf{v}_S = \{\mathbf{v}_j^i\}_{(i,j) \in S}$ and $\mathbf{v}_{-S} = \{\mathbf{v}_j^i\}_{(i,j) \notin S}$. Let $\mathcal{D} = \times_{j \in [m]} D_j$. Therefore, in a setting with n i.i.d. buyers, each drawn from $\mathcal{D}$, the prior distribution is $\mathcal{D}^n$[10].

Let $\mathcal{M}$ be a truthful mechanism with an allocation function π and a payment function $\mathbf{p}$, i.e., given the submitted bids $\mathbf{v}$, each buyer i is allocated item j with probability $\pi_j^i(\mathbf{v})$ and pays $\mathbf{p}^i(\mathbf{v})$. Let $\textsc{Rev}_{\mathcal{M}}(\cdot)$ be mechanism $\mathcal{M}$'s expected revenue, e.g., $\textsc{Rev}_{\mathcal{M}}(\mathcal{D}^h)$ is mechanism $\mathcal{M}$'s expected revenue from h buyers, where each buyer's valuation $\mathbf{v}^i$ is drawn i.i.d. from $\mathcal{D}$, i.e., $\textsc{Rev}_{\mathcal{M}}(\mathcal{D}^h) = \mathbb{E}_{\mathbf{v} \leftarrow D^h}\left[\sum_{i \in [h]} \mathbf{p}^i(\mathbf{v})\right]$. Let $\textsc{Rev}(\cdot)$ be the optimal expected revenue by any truthful mechanism, e.g., $\textsc{Rev}(\mathcal{D}^h)$ is the optimal expected revenue by any truthful mechanism for h i.i.d. buyers drawn from $\mathcal{D}$. Similarly, $\textsc{SRev}(\cdot)$ is the optimal expected revenue when selling the items separately, and $\textsc{BRev}(\cdot)$ is the optimal expected revenue when selling all items in a bundle. Let $\textsc{Val}(\cdot)$ be the expected value of the optimal allocation. Since we consider additive, independently drawn buyers, $\textsc{Val}(\mathcal{D}^h) = \sum_j \mathbb{E}\left[\max_{i \in [h]} \mathbf{v}_j^i\right]$.

For a single dimensional distribution D, and a number $p \in [0, 1]$, let $\textsc{Rev}_p(\cdot)$ denote the revenue of an optimal mechanism, among all truthful mechanisms that *sell with ex-ante probability of at most* p, i.e., mechanisms with allocation function π that satisfies $\mathbb{E}_{\mathbf{v}}\left[\sum_i \pi^i(\mathbf{v})\right] \leq p$.

Throughout the analysis, for h that is not an integer, we slightly abuse notation and use, e.g., D^h instead of $D^{\lceil h \rceil}$ to denote the product distribution $\times_{i \in \lceil h \rceil} D$. For a random variable X drawn from distribution D, we use $\textsc{Rev}(X)$ and $\textsc{Rev}(D)$ interchangeably (and similarly for $\textsc{SRev}, \textsc{BRev}$, etc.). Also, let $\mathbb{I}[\mathcal{E}]$ be the indicator random variable that equals 1 when event $\mathcal{E}$ occurs, and 0 otherwise.

Finally, we will use the following previously shown lemma.

Lemma 3.1. *[1] For any nm dimensional distribution $\mathcal{D}$, $\textsc{Rev}(\mathcal{D}) \leq n \cdot m \cdot \textsc{SRev}(\mathcal{D})$.*

3.1 Lemmas for single dimensional distributions.

We recall and develop tools for our analysis in the remaining sections. Let v be drawn from a single dimensional distribution D (i.e., $v \leftarrow D$), and consider some cutoff $\mathcal{T} \in \mathbb{R}$.

Lemma 3.2. *[21] $\textsc{Rev}(v | v > \mathcal{T}) = \frac{\textsc{Rev}(v \cdot \mathbb{I}[v > \mathcal{T}])}{\Pr[v > \mathcal{T}]}$.*

Lemma 3.3. *[27] Let $\gamma > 0$. Then $\Pr[v > \gamma \cdot \textsc{Rev}(D)] \leq \frac{1}{\gamma}$.*

Let $\mathrm{Var}(D)$ denote the variance of D.

Lemma 3.4. *[27] Let $\gamma > 0$, and supposed that both $\textsc{Rev}(D) \leq \gamma$ and that the support of D is in $[0, t\gamma]$. Then $\mathrm{Var}(D) \leq (2t - 1)\gamma^2$.*

The following lemma shows that for $\delta > 0$ that is not too small, the optimal revenue from $\lceil 1/\delta \rceil$ i.i.d. buyers drawn from D is at least a $1/(2\delta)$ factor larger than the optimal revenue that can be extracted from a buyer with value distributed according to $v \cdot \mathbb{I}[v > \mathcal{T}]$. This lemma will be used to bound the contribution (both to the core and tail) of higher values.

Lemma 3.5. *Suppose that $\delta \geq 2 \cdot \Pr[v > \mathcal{T}]$. Then $\textsc{Rev}(v \cdot \mathbb{I}[v > \mathcal{T}]) \leq 2\delta \cdot \textsc{Rev}(D^{1/\delta})$.*

The following lemma will be used to cover the contribution of values from the core.

Lemma 3.6. *Let $\mathcal{T}_\alpha$ satisfy $\Pr_{v \leftarrow D}[v > \mathcal{T}_\alpha] \geq \alpha$. For any $h \geq 1$: $\textsc{Rev}(D^h) \geq (1 - e^{-\alpha \cdot h}) \cdot \mathcal{T}_\alpha$.*

[10] By $\mathcal{D}^n$ we mean the product distribution $\times_{i \in [n]} \mathcal{D}$.

The following lemma is a specified analog to Lemma 3.6 that bounds the contribution from values in the core using the second price auction, and its proof is similar to the proof of Lemma 3.6. Let Rev_{VCG} denote the revenue from the second price auction.

Lemma 3.7. *Let* $\mathcal{T}_\alpha$ *satisfy* $\Pr_{v \leftarrow D}[v > \mathcal{T}_\alpha] \geq \alpha$. *For any* $h \geq 1$: $\text{Rev}_{VCG}\left(D^{2h}\right) \geq (1 - e^{-\alpha \cdot h})^2 \cdot \mathcal{T}_\alpha$

3.1.1 Regular distributions. For ease of exposition, we assume D has a density f and is strictly increasing. Nevertheless, our results extend to arbitrary regular distributions. A distribution D with density f is *regular* if $v - (1 - D(v))/f(v)$ is non-decreasing in v. Our analysis is done in quantile space, which is defined below.

Definition 3.8. (Quantiles, demand curve, and revenue curve.)

- Let $q(v) = \Pr_{x \leftarrow D}[x \leq v] = D(v)$ be the *quantile* of v.
- Let $V(q) = v$ for which $D(v) = q$, i.e., $V(q) = D^{-1}(q)$ is the *demand curve* of D.
- Let $\text{rev}(q) = (1 - q) \cdot V(q)$, i.e., the revenue from the posted price $V(q)$ that sells w.p. $1 - q$.

Observe that V is increasing in q. By change of variables, the expected value v can be computed as follows: [11] $\mathbb{E}[v] = \int_0^\infty v f(v) dv = \int_0^1 V(q) dq = \mathbb{E}_{q \leftarrow U[0,1]}[V(q)]$. The following lemma is a well known characterization of regular distributions.

Lemma 3.9. *[30] A distribution D is regular if and only if rev(q) is concave, i.e., for every* $\alpha, \beta, \gamma \in [0, 1]$, *it holds that* $\text{rev}(\gamma \cdot \alpha + (1 - \gamma)\beta) \geq \gamma \cdot \text{rev}(\alpha) + (1 - \gamma) \cdot \text{rev}(\beta)$.

The following corollary is the only way we use regularity (and is only used to prove Lemma 3.11).

Corollary 3.10. *If D is regular and* $0 \leq \beta \leq \alpha \leq 1$, *then* $\text{rev}(\beta) \geq \alpha \cdot \text{rev}(\beta) \geq \beta \cdot \text{rev}(\alpha)$.

The following lemma shows that if the probability of exceeding the cutoff $\mathcal{T}$ is ε, then $\mathcal{T}$ is a price that provides a $1 - \varepsilon$ approximation to the optimal revenue from the random variable $v \cdot \mathbb{I}[v > \mathcal{T}]$.

Lemma 3.11. *For a cutoff* $\mathcal{T} = V(1 - \varepsilon)$, *it holds that* $\text{rev}(1 - \varepsilon) \geq (1 - \varepsilon) \cdot \text{Rev}(v \cdot \mathbb{I}[v > \mathcal{T}])$.

The following lemma relates the contribution from the core to the contribution from the tail.

Lemma 3.12. *For a cutoff* $\mathcal{T} = V(1 - \varepsilon_1)$, *for every* $0 < \varepsilon_2 < 1$ *and* $k \in \mathbb{N}$ *such that* $\varepsilon_1 \cdot \varepsilon_2^{-k} < 1$ *it holds that:* $\mathbb{E}_{v \leftarrow D}[v \cdot \mathbb{I}[v \leq \mathcal{T}]] \geq k \cdot (1 - \varepsilon_2) \cdot (1 - \varepsilon_1 \cdot \varepsilon_2^{-k}) \text{Rev}(v \cdot \mathbb{I}[v > \mathcal{T}])$

3.1.2 Mechanisms with restricted probability of sale. We describe optimal single item mechanisms with restricted probability of sale. The lemmas established in this section are used in Section 5. In particular, we formally show that a naive generalization of Myerson's optimal auction for the unrestricted case is also optimal for the restricted case.

Fix a number $p \in [0, 1]$ and x i.i.d. buyers, each buyer i with value $\mathbf{v}^i$ drawn from a single dimensional distribution D. A mechanism with allocation function π that satisfies $\mathbb{E}_{\mathbf{v}}\left[\sum_{i \in [x]} \pi^i(\mathbf{v})\right] \leq p$, is said to *sell with ex-ante probability of at most p*.

Recall that by Myerson's theory, a mechanism with allocation rule π and payment $\mathbf{p}$ is truthful if and only if each buyer i's allocation rule π^i is monotone, and buyer i's payment function is fully determined by π^i via the equality $\mathbf{p}^i(\mathbf{v}) = \mathbf{v}^i \cdot \pi^i(\mathbf{v}) - \int_{z=0}^{\mathbf{v}^i} \pi^i(z, \mathbf{v}^{-i}) dz$. Furthermore, Myerson defines *virtual value* functions[12] $\varphi^D : \text{supp}(D) \to \mathbb{R}$ for single dimensional distributions D (we henceforth drop the superscript D). The virtual value function is particularly useful because one can re-amortize the expected payment to a term that uses the virtual value functions, and this term can be optimized point wise. This becomes apparent in Myerson's payment identity:

[11] By $v = V(q)$, we get $dv = V'(q)dq = [D^{-1}]'(q) = \frac{1}{f(D^{-1}(q))} = \frac{1}{f(V(q))}$, so $vf(v)dv = V(q)\frac{f(V(q))}{f(V(q))}dq$

[12] When D has a density f, its virtual value function is defined by $\varphi(v) = v - \frac{1 - D(v)}{f(v)}$. See e.g., [6, 19] for the general case.

Lemma 3.13. *[30] (Payment identity) In every truthful mechanism with allocation function π and payment function $\mathbf{p}$, for every buyer i it holds that* $\mathbb{E}\left[\varphi(\mathbf{v}^i)\cdot\pi^i(\mathbf{v})\right] = \mathbb{E}\left[\mathbf{p}^i(\mathbf{v})\right]$

A regular distribution is a distribution whose corresponding virtual value function φ is monotone. For irregular distributions, Myerson defines yet another transformation that transforms the virtual value function to an *ironed* virtual value function $\bar{\varphi}$ (the details are not so important for our application; for more details see e.g., [24, 30]). The key point is that ironed virtual value functions are always monotone, and furthermore, maintain the following property:

Theorem 3.14. *([24, Theorem 3.12]) For every truthful mechanism with allocation function π, for every buyer i,* $\mathbb{E}\left[\varphi(\mathbf{v}^i)\cdot\pi^i(\mathbf{v})\right] \leq \mathbb{E}\left[\bar{\varphi}(\mathbf{v}^i)\cdot\pi^i(\mathbf{v})\right]$, *with equality if $\pi^i(\cdot,\mathbf{v}^{-i})$ is constant on ironed intervals.*

In order to construct the optimal mechanism with restricted probability of sale, we will need the following definitions. Let $\mathbf{v}^* = \max_{i\in[x]}\mathbf{v}^i$, and let $S = \{y \geq 0 : \Pr[\bar{\varphi}(\mathbf{v}^*) \geq y] \leq p\}$ and $\phi = \inf S$.

Note that if $\phi \in S$, then it must be that $\phi = 0$ or that $\Pr[\bar{\varphi}(\mathbf{v}^*) \geq \phi] = p$. [13] Otherwise, it must be that $\Pr[\bar{\varphi}(\mathbf{v}^*) > \phi] \leq p$, [14] but $\Pr[\bar{\varphi}(\mathbf{v}^*) \geq \phi] > p$, i.e, $\Pr[\bar{\varphi}(\mathbf{v}^*) = \phi] = \Pr[\bar{\varphi}(\mathbf{v}^*) \geq \phi] - \Pr[\bar{\varphi}(\mathbf{v}^*) > \phi] > 0$.

Consider the mechanism $\mathcal{M}$ that sells a single item to x buyers as follows. If $\phi \in S$, the mechanism allocates the item to a random buyer i among those with $\bar{\varphi}(\mathbf{v}^i) = \bar{\varphi}(\mathbf{v}^*)$ whenever $\bar{\varphi}(\mathbf{v}^*) \geq \phi$, and otherwise does not allocate.

If $\phi \notin S$, then the mechanism allocates the item to a random buyer i among those with $\bar{\varphi}(\mathbf{v}^i) = \bar{\varphi}(\mathbf{v}^*)$ whenever $\bar{\varphi}(\mathbf{v}^*) > \phi$, and whenever $\bar{\varphi}(\mathbf{v}^*) = \phi$ the mechanism draws a Bernoulli random variable that equals 1 w.p. $\frac{p-\Pr[\bar{\varphi}(\mathbf{v}^*)>\phi]}{\Pr[\bar{\varphi}(\mathbf{v}^*)=\phi]}$, and if the variable equals 1 the mechanism allocates to a random buyer i among those with $\bar{\varphi}(\mathbf{v}^i) = \phi$.

Lemma 3.15. *For x i.i.d. buyers drawn from a single dimensional distribution D, and $p, \phi, \mathbf{v}^*$ as above, the mechanism $\mathcal{M}$ described above sells with ex-ante probability of at most p, and extracts expected revenue of:* $\mathbb{E}\left[\bar{\varphi}(\mathbf{v}^*)\cdot\mathbb{I}[\bar{\varphi}(\mathbf{v}^*) > \phi]\right] + \phi\cdot(p - \Pr[\bar{\varphi}(\mathbf{v}^*) > \phi])$.

Moreover, every buyer contributes the same amount of expected revenue.

We will now prove the optimality of $\mathcal{M}$ among mechanisms with restricted probability of sale.

Lemma 3.16. *For D, x, p as above, the mechanism $\mathcal{M}$ above is a revenue maximizing mechanism among all truthful mechanisms that sell to x i.i.d. buyers drawn from D with ex-ante probability of at most p, i.e.,* $\textsc{Rev}_p(D^x) = \textsc{Rev}_{\mathcal{M}}(D^x)$.

Corollary 3.17. *Let $0 \leq \delta < 1$. Then* $\textsc{Rev}_p(D^x) \leq \frac{1}{1-\delta}\textsc{Rev}_p(D^{(1-\delta)\cdot x})$

Lemma 3.18. $\textsc{Rev}_p(D^k)$ *is non-decreasing in k.*

Lemma 3.19. *Fix a cutoff $\mathcal{T}$. Let $\mathbf{v} \leftarrow D^n$. Let $\hat{D}$ denote the distribution of $\mathbf{v}^i|(\mathbf{v}^i > \mathcal{T})$ for every buyer i. Then for every $k \leq n$ and $p \in [0,1]$:* $\textsc{Rev}_p\left(\hat{D}^k\right) \leq \frac{\textsc{Rev}_p(D^n)}{\Pr[|\{i:\mathbf{v}^i>\mathcal{T}\}|\geq k]}$

3.2 Partition of the support.

Recall that for every buyer i and item j, $\mathbf{v}^i_j \leftarrow D_j$, and that $\mathcal{D} = \times_{j\in[m]}D_j$. Every item j is assigned a cutoff $\mathcal{T}_j$. If the value of an item is greater than its cutoff (i.e., $\mathbf{v}^i_j > \mathcal{T}_j$) then it is said to be in the tail w.r.t. $\mathbf{v}^i$, and otherwise it is in the core w.r.t $\mathbf{v}^i$.

[13] Otherwise $\Pr[\bar{\varphi}(\mathbf{v}^*) \geq \phi] = p - \epsilon$ but for every $0 < a < \phi$ we have $\Pr[\phi > \bar{\varphi}(\mathbf{v}^*) \geq a] > \epsilon$.

[14] Otherwise $\Pr[\bar{\varphi}(\mathbf{v}^*) > \phi] = p + \epsilon$ but for every $a > \phi$ we have $\Pr[a > \bar{\varphi}(\mathbf{v}^*) > \phi] \geq \epsilon$.

For a set of items A, let $\mathcal{D}^A$ be the distribution $\mathcal{D}$ conditioned on A being exactly the set of items in the tail. Let p_A be the probability of A being exactly the set of items in the tail, i.e., $p_A = \prod_{j \in A} \Pr[\mathbf{v}^i_j > \mathcal{T}_j] \cdot \prod_{j \notin A} \Pr[\mathbf{v}^i_j \leq \mathcal{T}_j]$. Let $\mathcal{D}^A_C$ (resp., $\mathcal{D}^A_T$) be the distribution $\mathcal{D}^A$ restricted to just items in the core (resp., tail).

For every buyer i consider some $\mathcal{A}^i \subseteq [m]$, and let $\mathcal{A} = \{\mathcal{A}^i\}_{i \in [n]}$. Let $\mathcal{A}$ represent the event that for every i it holds that $\mathcal{A}^i$ is exactly the set of items that are in the tail with respect to buyer i's valuation, i.e., for every item $j \in \mathcal{A}^i$ it holds that $\mathbf{v}^i_j > \mathcal{T}_j$, and for every $j \notin \mathcal{A}^i$ it holds that $\mathbf{v}^i_j \leq \mathcal{T}_j$. Denote by $p_{\mathcal{A}}$ the probability of event $\mathcal{A}$, i.e., $p_{\mathcal{A}} = \prod_i p_{\mathcal{A}^i}$, and let $\mathcal{D}^{\mathcal{A}}$ be the distribution $\mathcal{D}^n$ conditioned on the event $\mathcal{A}$, i.e., $\mathcal{D}^{\mathcal{A}} = \times_{i \in [n]} \mathcal{D}^{\mathcal{A}^i}$.

The following lemma appeared previously and is included for completeness.

Lemma 3.20. *[21] (Sub-domain stitching) Let $\mathcal{L} = \{\mathcal{A} : \forall i, \mathcal{A}^i \subseteq [m]\}$ be all possible $\mathcal{A}$. Then $\textsc{Rev}(\mathcal{D}^n) \leq \sum_{\mathcal{A} \in \mathcal{L}} p_{\mathcal{A}} \cdot \textsc{Rev}(\mathcal{D}^{\mathcal{A}})$*

4 MANY ITEMS: SRev WITH $O(n \cdot \log \frac{m}{n})$ BUYERS

In this section we prove that for $m \gg n$, increasing the number of buyers by a factor of $O(\log(m/n))$ suffices to recover 99% of the optimal revenue by selling the items separately.

Theorem 4.1 (Theorem 1.2, case $m \gg n$). *For any constant $\varepsilon > 0$ there exists a constant $\delta = \delta(\varepsilon) > 0$ such that whenever $m \geq 2n/\delta$,*

$$(1-\varepsilon)\textsc{Rev}(\mathcal{D}^n) \leq \textsc{SRev}(\mathcal{D}^{\frac{n \cdot \log \frac{m}{n}}{\delta}})$$

This implies the following corollary for the special case of a single buyer:

Corollary 4.2 (Theorem 1.1). *For any constant $\varepsilon > 0$ there exists a constant $\delta(\varepsilon) > 0$ such that whenever $m \geq 2/\delta$, $(1-\varepsilon)\textsc{Rev}(\mathcal{D}) \leq \textsc{SRev}(\mathcal{D}^{\frac{\log m}{\delta}})$*

4.1 Proof outline

We first partition the domain of the valuations distribution into sub-domains (using the Subdomain Stiching Lemma (Lemma 3.20)). We henceforth condition on the event (or sub-domain) $\mathcal{A}$ that describes which items are in the tail or core for which buyer (we still don't know their values within the tail/core).

We now describe a "marginal mechanism lemma" (Lemma 4.3) that shows that the optimal revenue from $\mathcal{D}^{\mathcal{A}}$, is bounded by the sum of expected item values in the core in event $\mathcal{A}$, plus the revenue from selling to each buyer the items that are in her tail in event $\mathcal{A}$ (conditioned on them being in the tail). This implies a core-tail decomposition lemma (Lemma 4.4) that bounds the optimal revenue using two terms: the core and tail.

We bound the tail in Section 4.3. In Lemma 4.5 we show that the tail is bounded by the revenue from selling items separately and using only prices that are higher than the cutoffs. Since the probability that two buyers are interested in the same item at a high price is extremely low, the supply constraint is hardly restricting. Therefore when we increase the number of buyers, the revenue increases almost linearly. This is used in Lemma 4.6 that shows that the tail is a tiny fraction of the revenue after multiplying the number of buyers by a (large) constant.

We bound the core in Section 4.4. We separate the core into regions of lower values and higher values. Lower values are upper bounded by a maximum of two cutoffs, and are handled in two lemmata (4.7 and 4.8).

In Lemma 4.10 we bound the contribution from "higher" values (still in the core). First, we observe that they lie in a bounded region: bounded from above since they are in the core, and bounded from below since we handle low values separately. We use this property to show that their expected value is at most $O(\log(m/n))$-factor larger than the revenue from selling each item separately at "higher" prices. Now, similarly to the tail, since at high prices buyers are likely to want disjoint sets of items, the above revenue scales almost linearly when increasing the number of buyers. Therefore, multiplying the number of buyers by $O(\log(m/n))$ suffices to extract revenue much larger than the contribution from high values.

Finally, we complete the proof by combining the core-tail decomposition with the above arguments.

4.2 Core-tail decomposition.

The following lemma may be seen as a generalization of Lemma 24 in [21] for $n > 1$ buyers. In [21], the authors note that the proof of their Lemma 24 implies a generalization for $n > 1$, by considering for each item the maximal value among all buyers' values. This method still allows to partition the support of $\mathcal{D}^n$ to only 2^m subsets (according to the set of items that have a maximal value above the cutoff). We show a more fine-grained generalization that partitions the support of $\mathcal{D}^n$ to $2^{m \cdot n}$ subsets, thereby leading to a tighter analysis.

Recall that for some $\mathcal{A}^i$, $\mathcal{D}_T^{\mathcal{A}^i}$ is the distribution $\mathcal{D}$ conditioned on $\mathcal{A}^i$ being exactly the set of items in the tail, restricted only to the items in the tail ($\mathcal{A}^i$).

Lemma 4.3. *(Marginal Mechanism on Sub-Domain) For $\mathcal{A}$ as above,*

$$\textsc{Rev}\left(\mathcal{D}^{\mathcal{A}}\right) \leq \sum_{j \in [m]} \mathop{\mathbb{E}}_{\mathbf{v} \leftarrow \mathcal{D}^n}\left[\max_{i : j \notin \mathcal{A}^i} \mathbf{v}_j^i \middle| \mathcal{A}\right] + \sum_{i \in [n]} \textsc{Rev}\left(\mathcal{D}_T^{\mathcal{A}^i}\right)$$

The following core-tail decomposition upper bounds the optimal revenue from n buyers. In contrast to the decomposition in [1], our decomposition provides a tail that relates to single buyer settings.

To simplify notation, once cutoffs $\mathcal{T}_j$ are fixed for every j, we use $\mathcal{D}_{\textsc{Tail}}$ to denote the product distribution of the random variables $\{v_j \cdot \mathbb{I}[v_j > \mathcal{T}_j]\}_{j \in [m]}$, and $\mathcal{D}_{\textsc{Core}}$ to denote the product distribution of the random variables $\{v_j \cdot \mathbb{I}[v_j \leq \mathcal{T}_j]\}_{j \in [m]}$, where $v_j \leftarrow D_j$ for every j.

Lemma 4.4. *(Core-tail decomposition) $\textsc{Rev}(\mathcal{D}^n) \leq \textsc{Val}(\mathcal{D}^n_{\textsc{Core}}) + n \cdot \sum_{A \subseteq [m]} p_A \cdot \textsc{Rev}(\mathcal{D}_T^A)$*

Cutoff setting. For every item j fix $\mathcal{T}_j$ to be so that $\Pr_{\mathbf{v}_j \leftarrow D_j}[\mathbf{v}_j > \mathcal{T}_j] = m^{-1}$.

4.3 Tail

The main lemma of this section is Lemma 4.6 which states that the contribution from the tail is a tiny fraction of the revenue from SRev with $n \cdot C$ buyers where C is some constant. Lemma 4.5 shows (in the spirit of Proposition 1 in [1]) that we can approximately recover the tail of one buyer ($\sum_{A \subseteq [m]} p_A \cdot \textsc{Rev}(\mathcal{D}_T^A)$) by selling items separately, only using prices that are higher than the cutoffs. Recall that $\mathcal{D}_{\textsc{Tail}}$ is the product distribution of the random variables $\{\mathbf{v}_j \cdot \mathbb{I}[\mathbf{v}_j > \mathcal{T}_j]\}_{j \in [m]}$, and therefore $\textsc{SRev}(\mathcal{D}_{\textsc{Tail}}) = \sum_{j \in [m]} \textsc{Rev}(\mathbf{v}_j \cdot \mathbb{I}[\mathbf{v}_j > \mathcal{T}_j])$.

Lemma 4.5. *$\sum_{A \subseteq [m]} p_A \cdot \textsc{Rev}(\mathcal{D}_T^A) \leq 2 \cdot \textsc{SRev}(\mathcal{D}_{\textsc{Tail}})$*

Combining the above lemma with Lemma 3.5 shows that SRev with $n \cdot C$ buyers (where C is a constant) extracts *much more* revenue than the contribution from the tail.

Lemma 4.6. *Let $\delta \geq 2n/m$. Then $n \cdot \sum_{A \subseteq [m]} p_A \textsc{Rev}(\mathcal{D}_T^A) \leq 4 \cdot \delta \cdot \textsc{SRev}(\mathcal{D}^{n/\delta})$*

4.4 Core

For ease of exposition, let $\mathbf{v}_j^*$ denote the random variable $\mathbf{v}_j^* = \max_{i\in[n]}\{\mathbf{v}_j^i \cdot \mathbb{I}[\mathbf{v}_j^i \leq \mathcal{T}_j]\}$ where $\mathbf{v}_j^i \leftarrow D_j$ for each $i \in [n]$. In this section we upper bound the core $\text{VAL}(\mathcal{D}^n_{\text{CORE}}) = \sum_{j\in[m]} \mathbb{E}[\mathbf{v}_j^*]$.

Cutoff setting in the core. Fix some constant $\varepsilon_1 \leq 1$. For each item j, set $\mathcal{T}_j{}'$ so that $\Pr_{v\leftarrow D_j}[v > \mathcal{T}_j{}'] = \frac{\varepsilon_1}{n\cdot\log\frac{m}{n}}$, and set $\mathcal{T}_j{}'' = \max\{\frac{n}{m}\cdot\mathcal{T}_j, \mathcal{T}_j{}'\}$. In order to bound the contribution from the core, for each item we separate to "tiny values" (at most $\frac{n}{m}\mathcal{T}_j$), "low values" (at most $\mathcal{T}_j{}'$), and "higher values" (above $\max\{\frac{n}{m}\cdot\mathcal{T}_j, \mathcal{T}_j{}'\}$).

In the following we bound the contribution from values that are at most $\mathcal{T}_j{}''$. In Lemma 4.7 we show that "tiny values" contribute a tiny fraction of the revenue that is achievable by increasing the number of buyers by a constant factor, and in Lemma 4.8 we show that the contribution of "lower values" to the core can be almost completely covered by the revenue that is achievable by increasing the number of buyers by a factor of $O(\log\frac{m}{n})$.

LEMMA 4.7. *(Tiny values) Let $\delta \geq 2n/m$. Then for each item j $\frac{n}{m}\mathcal{T}_j \leq 2\delta \cdot \text{REV}(D_j^{n/\delta})$*

LEMMA 4.8. *(Lower values) Let $\delta < \varepsilon_1$. Then for each item j, $\mathcal{T}_j{}' \leq (1-e^{-\frac{\varepsilon_1}{\delta}})^{-1}\cdot \text{REV}(D_j^{n\cdot\log(\frac{m}{n})/\delta})$*

Putting Lemmas 4.7 and 4.8 gives the following corollary

COROLLARY 4.9. *Let $\frac{2n}{m} \leq \delta \leq \min\{\frac{1}{2}, \varepsilon_1\}$. Then for each item j, $\mathcal{T}_j{}'' \leq (1-e^{-\frac{\varepsilon_1}{\delta}})^{-1}\cdot\text{REV}(D_j^{n\cdot\log(\frac{m}{n})/\delta})$*

Lemma 4.10 shows that "higher values" in the core contribute a tiny fraction of the revenue that is achievable by increasing the number of buyers by a factor of $O(\log\frac{m}{n})$.

LEMMA 4.10. *Let $\delta < 2\varepsilon_1$. Then for each j, $\mathbb{E}[\mathbf{v}_j^*\cdot\mathbb{I}[\mathbf{v}_j^* > \mathcal{T}_j{}'']] \leq 8\cdot\varepsilon_1\cdot\text{REV}(D_j^{\frac{n\cdot\log\frac{m}{n}}{\delta}})$.*

We are now ready to bound the contribution from the core. Let ε_2 be so that $1+\varepsilon_2 = (1-e^{-\frac{\varepsilon_1}{\delta}})^{-1}$.

LEMMA 4.11. *Let $2n/m \leq \delta \leq \min\{1/2, \varepsilon_1\}$. Then, $\text{VAL}(\mathcal{D}^n_{\text{CORE}}) \leq (1+\varepsilon_2+8\cdot\varepsilon_1)\cdot\text{SREV}(\mathcal{D}^{\frac{n\cdot\log\frac{m}{n}}{\delta}})$.*

4.5 Matching Lower Bound for Many Items

THEOREM 4.12. *Let $m = \Omega(n)$. There exists a distribution $\mathcal{D}$ so that for every $\varepsilon > 0$ there exists some $\delta > 0$ so that*

$$\text{SREV}(\mathcal{D}^{\delta\cdot n\cdot\ln\frac{m}{n}}) \leq \varepsilon\cdot\text{REV}(\mathcal{D}^n).$$

Proof sketch: Each item is drawn from the equal revenue (ER) distribution on the support $[1, m]$. Let $\mathcal{ER} = \times_{j\in[m]}\text{ER}$. One can easily show that $\text{SREV}(\mathcal{ER}^{\delta\cdot n\cdot\ln(\frac{m}{n})}) \leq \delta\cdot m\cdot n\cdot\ln\left(\frac{m}{n}\right)$. Therefore, it suffices to show a mechanism that achieves revenue of $\Omega(m\cdot n\cdot\ln\left(\frac{m}{n}\right))$. Consider the mechanism that orders the buyers (arbitrarily), and offers each buyer in her turn to purchase her favorite bundle of size $\frac{m}{4\cdot n}$ at price $\frac{m}{8}\cdot\ln\frac{m}{n}$. As $\frac{m}{n}$ grows to a large enough constant, one can show that for every buyer w.h.p. there are at least $\frac{m}{4n}$ items valued above n. Conditioned on that event, the expected value of the sum of these items is at least $\frac{m}{4}\cdot\ln\frac{m}{n}$. Also, w.h.p. every buyer will purchase a bundle (observe that there are always enough items), thus extracting revenue of $\frac{n\cdot m}{8}\cdot\ln\frac{m}{n}$, w.h.p., as required.

5 MANY BUYERS: SREV WITH NO ADDITIONAL BUYERS

In this section we show that adding more buyers is not required for a number of buyers n that is sufficiently larger than the number of items m. Specifically, we prove the following theorem:

Theorem 5.1 (Theorem 1.2, case $n \gg m$). *For any constant $\varepsilon > 0$ there exists a constant $\delta(\varepsilon) > 0$ such that whenever $n \geq \frac{m}{\delta}$:*

$$(1-\varepsilon)\textsc{Rev}(\mathcal{D}^n) \leq \textsc{SRev}(\mathcal{D}^n)$$

5.1 Proof outline (a tale of three tails)

We first partition the domain of the valuation distributions into sub-domains, using the Subdomain Stiching Lemma (Lemma 3.20). In our analysis, we condition on the event (or sub-domain) $\mathcal{A}$ that describes which items are in the tail or core for which buyer (we still don't know their values within the tail/core). For each event we define a bipartite graph with buyer-nodes $[n]$ and item-nodes $[m]$, where an edge $\{i, j\}$ between buyer i and item j exists if and only if item j is in buyer i's tail. Crucially, the revenue is separable across connected components (because in different connected components, disjoint sets of buyers are interested in disjoint sets of items).

We now establish a "marginal mechanism lemma" (Lemma 5.2) that shows that the optimal revenue from $\mathcal{D}^{\mathcal{A}}$ is bounded from above by the sum of expected item values in the core in event $\mathcal{A}$, plus the revenue from selling items in each connected component to the buyers in that connected component (conditioned on them being in the tail).

For a connected component that contains only one item (i.e., the buyers in this connected component have only this item in their tails), the lemma provides an even tighter bound. Namely, we can restrict attention to mechanisms that sell the item to buyers in the connected component with probability at most p, where p is the upper bound on the ex-ante probability that the optimal mechanism (for event $\mathcal{A}$) sells this item. Then, with probability $1-p$ we can resell the item to other buyers. Essentially, this means that we can use most buyers and some items to (almost) recover the contribution from the tail, and then use the rest of the buyers and items to (almost) recover the contribution from the core. Note that while we partition the buyers in advance, the partitioning of the items is only done after seeing which items are bought by the first subset of buyers. Note also that the mechanism of partitioning the buyers is for analysis purpose only: once we establish that it achieves a good revenue by selling each item separately, we can only improve the revenue by running Myerson's optimal mechanism.

Since we consider a setting with many buyers, the expected number of buyers that have a specific item in their tail (henceforth termed *interested buyers*) is concentrated. This implies a core-tail decomposition lemma (Lemma 5.3) that bounds the optimal revenue using four terms: one core, and three tails. The three tails are (1) the contribution from selling item j whenever the number of interested buyers does not exceed the expected number of interested buyers by much, (2) the contribution from selling item j whenever the number of interested buyers *does* exceed the expected number of interested buyers by a significant constant, and (3) the contribution from selling items in connected components with at least two items.

We bound the three tails in Section 5.4. In Lemma 5.5 we show that the first tail is almost completely recovered by the revenue from selling items separately to all but a small fraction of the buyers, and restricting the ex-ante probability of sale as described above. Note that the tail is not fully recovered, but this is a sufficient guarantee due to the probability of sale restriction (see above). In particular, the items can be resold to other buyers with sufficiently high probability.

In Lemma 5.6 we show that the second tail is a tiny fraction of the revenue from selling items separately. This is done by arguing that the event where the number of interested buyers significantly exceeds the expectation is extremely rare.

In Lemma 5.9 we show that the third tail is also a tiny fraction of the revenue from selling items separately. This part is more involved, since we need to bound the expected size of a connected component, as well as some of the higher moments. Here we use some simple ideas from percolation theory. See Section 5.4.1 for details.

Finally, we conclude the proof in Section 5.5 by integrating the above arguments into the core-tail decomposition.

5.2 The item-buyer bipartite graph

We adopt all the notation from Section 3.2, and introduce some more notation. For a single dimensional distribution $v \leftarrow D$ and cutoff $\mathcal{T}$, let $\hat{D}$ denote the distribution of $v|v > \mathcal{T}$. Fix an event $\mathcal{A}$ as described in Section 3.2. Consider the bipartite graph $H_{\mathcal{A}} = ([n], [m], E[\mathcal{A}])$ over buyers and items. The set of edges is defined to be $E[\mathcal{A}] = \{\{i, j\} : j \in \mathcal{A}^i\}$, i.e., for each item j in the tail of i, ($\mathbf{v}^i_j > \mathcal{T}_j$) there exists an edge $\{i, j\}$. For each node k, let $\mathcal{A}[k]$ denote k's neighbors in $H_{\mathcal{A}}$, i.e., for an item $j \in [m]$, $\mathcal{A}[j] = \{i : j \in \mathcal{A}^i\}$, and for a buyer $i \in [n]$, $\mathcal{A}[i] = \mathcal{A}^i$. Note that $H_{\mathcal{A}}$ is a random graph where each edge $\{i, j\}$ exists independently w.p. $\Pr[\mathbf{v}^i_j > \mathcal{T}_j]$. Let $P^{\mathcal{A}}$ be the partition of the graph $H_{\mathcal{A}}$ to connected components. For a connected component $X \in P^{\mathcal{A}}$, let $X_n = X \cap [n]$ be the nodes in X associated with buyers, and similarly $X_m = X \cap [m]$ are the nodes in X associated with items. For a connected component $X \in P^{\mathcal{A}}$, let $\mathcal{D}^X_T$ be the product distribution over the buyers X_n, where each buyer i has only values $\mathbf{v}^i_j$ for items j in $\mathcal{A}[i]$, and each such $\mathbf{v}^i_j$ is drawn from the conditional distribution $\hat{D}_j$, i.e., $D^X_T = \times_{i \in X_n} \mathcal{D}^{\mathcal{A}[i]}_T$.

5.3 Core-tail decomposition.

We introduce a new "marginal mechanism" lemma (Lemma 5.2) that is suitable for the case $n \gg m$. Fix some $\mathcal{A}$ as in Section 3.2, and let $\mathcal{M}$ be an optimal mechanism w.r.t. $\mathcal{D}^{\mathcal{A}}$, i.e., $\textsc{Rev}_{\mathcal{M}}(\mathcal{D}^{\mathcal{A}}) = \textsc{Rev}(\mathcal{D}^{\mathcal{A}})$. Let π be the mechanism's allocation function, and let $\bar{\pi}_j = \sum_{i \in \mathcal{A}[j]} \mathbb{E}_{\mathbf{v} \leftarrow \mathcal{D}^{\mathcal{A}}}[\pi^i_j(\mathbf{v})]$. Recall that for a connected component $X \in P^{\mathcal{A}}$, $X_n = X \cap [n]$ and similarly $X_m = X \cap [m]$. Also, recall that given $v \leftarrow D_j$, we use $\hat{D}_j$ to denote the distribution of $v|v > \mathcal{T}_j$.

Lemma 5.2. *(Marginal Mechanism on Sub-Domain) Fix $\mathcal{A}$, and let $\bar{\pi}$ be as above. Let $P^{\mathcal{A}}_2$ be all the connected components X in $P^{\mathcal{A}}$ so that $|X_m| \geq 2$. Then*

$$\textsc{Rev}(\mathcal{D}^{\mathcal{A}}) \leq \sum_{j \in [m]} \left(\textsc{Rev}_{\bar{\pi}_j}\left(\hat{D}_j^{|\mathcal{A}[j]|}\right) + \mathcal{T}_j \cdot (1 - \bar{\pi}_j) \right) + \sum_{X \in P^{\mathcal{A}}_2} \textsc{Rev}(\mathcal{D}^X_T)$$

I.e., the revenue from $\mathcal{D}^{\mathcal{A}}$ is bounded by the revenue from selling items separately, each item j to $|\mathcal{A}[j]|$ buyers in the tail at $\bar{\pi}_j$ ex-ante probability of sale, plus the contribution to the core ($\mathcal{T}_j \cdot (1 - \bar{\pi}_j)$), plus the revenue from connected components with more than one item.

One can show that for fixed item j, cutoff $\mathcal{T}_j$, and $\delta = 2 \cdot \Pr_{v \leftarrow D_j}[v > \mathcal{T}_j]$, for any $k \leq 1/\delta$: $\textsc{Rev}(\hat{D}^k_j) \leq 4 \cdot k \cdot \textsc{Rev}(D^{1/\delta}_j)$, (see full version for details). This in turns allows to show the following core-tail decomposition,

Lemma 5.3. *(Core-tail decomposition) For cutoffs $\mathcal{T}_j$ so that $\Pr_{v \leftarrow D_j}\left[v > \mathcal{T}_j\right] \geq \frac{1}{2n}$, and a cutoff $\mathcal{T}^*$*

$$\begin{aligned}
\textsc{Rev}(\mathcal{D}^n) \leq & \sum_{j \in [m]} \mathcal{T}_j \cdot \mathbb{E}\left[1 - \bar{\pi}_j\right] && \textit{Core} \\
& + \sum_{j \in [m]} \mathbb{E}\left[\textsc{Rev}_{\bar{\pi}_j}\left(\hat{D}_j^{|\mathcal{A}[j]|}\right) \cdot \mathbb{I}\left[|\mathcal{A}[j]| \leq \mathcal{T}^*\right]\right] && \textit{Tail 1} \\
& + 4 \cdot \sum_{j \in [m]} \textsc{Rev}\left(\mathcal{D}^n_j\right) \cdot \mathbb{E}\left[|\mathcal{A}[j]| \cdot \mathbb{I}\left[|\mathcal{A}[j]| > \mathcal{T}^*\right]\right] && \textit{Tail 2} \\
& + 4 \cdot \sum_{j \in [m]} \textsc{Rev}\left(\mathcal{D}^n_j\right) \cdot \mathbb{E}\left[\left|X^j_m\right| \cdot \left|X^j_n\right|^2 \cdot \mathbb{I}\left[\left|X^j_m\right| \geq 2\right]\right] && \textit{Tail 3}
\end{aligned}$$

Cutoff setting. Fix some $0 < \varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4 < 1$. Throughout this section, for every item j fix $\mathcal{T}_j$ to be the value for which $\Pr_{v \leftarrow D_j}\left[v > \mathcal{T}_j\right] = \frac{1}{n \cdot \varepsilon_1}$. Clearly the expected number of buyers that have item j in their tail is $\frac{1}{\varepsilon_1}$. We set $\mathcal{T}^* = \frac{(1+\varepsilon_2)}{\varepsilon_1}$, i.e., we use ε_2 as a deviation from this expected size.

We will first sell items separately at high prices to a $1 - \varepsilon_4$ fraction of the buyers. In this case the expected number of buyers that have item j in their tail is $\frac{1-\varepsilon_4}{\varepsilon_1}$. We will use ε_3 as a deviation from this expected size.

Whenever an item j is not sold to one of the $1 - \varepsilon_4$ fraction of the buyers, we will sell the item to one of the remaining ε_4 buyers at price $\mathcal{T}_j$. This will suffice to to almost fully recover the contribution from the core.

Lemma 5.4. *(Core.) For every item j, $\mathcal{T}_j \cdot \mathbb{E}\left[(1 - \bar{\pi}_j)\right] \leq \mathbb{E}\left[(1 - \bar{\pi}_j)\right](1 - e^{-\varepsilon_4/\varepsilon_1})^{-1} \textsc{Rev}(D_j^{\varepsilon_4 \cdot n})$*

5.4 Tail.

The following lemma essentially combines Lemma 3.19 with a concentration bound, to show that Tail 1 can be almost fully recovered by selling items separately to a $1 - \varepsilon_4$ fraction of the buyers. Let $\varepsilon_5 = \varepsilon_5(\varepsilon_1, \varepsilon_2, \varepsilon_4, \varepsilon_3)$ be so that $1 + \varepsilon_5 = \frac{1+\varepsilon_2}{(1-\varepsilon_4)\cdot(1-\varepsilon_3)} \cdot \left(1 - \exp\left(-\frac{(1-\varepsilon_4)\cdot\varepsilon_3^2}{2\cdot\varepsilon_1}\right)\right)^{-1}$. It is only important to note that $\varepsilon_5 \to 0$ as $\varepsilon_1 \to 0$.

Lemma 5.5. *(Tail 1). Fix an item j. For $k \leq \frac{(1+\varepsilon_2)}{\varepsilon_1}$, $\textsc{Rev}_p\left(\hat{D}_j^k\right) \leq (1 + \varepsilon_5) \cdot \textsc{Rev}_p(D_j^{(1-\varepsilon_4)\cdot n})$*

Lemma 5.6 shows that the coefficient of $\textsc{SRev}$ in Tail 2 tends to 0 with ε_1. Let $\varepsilon_6 = \frac{4}{\varepsilon_1} \cdot \exp\left(-\frac{\varepsilon_2^2}{8\cdot\varepsilon_1}\right)$.

Lemma 5.6. *(Tail 2) Fix item j. $\mathbb{E}\left[|\mathcal{A}[j]| \cdot \mathbb{I}\left[|\mathcal{A}[j]| > \mathcal{T}^*\right]\right] \leq \varepsilon_6$*

5.4.1 *The 3rd tail (a non-trivial connected component).*

The following lemmas ultimately show that the coefficient of $\textsc{SRev}(\mathcal{D}^n)$ in Tail 3 tends to 0 as as ε_1 tends to 0. This will be proved in Lemma 5.9. We will develop the proof in steps. Recall that for a connected component $X \in P^{\mathcal{A}}$, $X_n = X \cap [n]$ and $X_m = X \cap [m]$. Lemma 5.7 upper bounds the term we wish to bound into two terms. Note that in both terms, at a loss of a constant factor, we free ourself from the dependence on the number of buyers that have at most one item in their tail (and are in j's connected component).

Lemma 5.7. *For every $\mathcal{A}$, let $S \subseteq [n]$ denote the buyers i with at least two items in their tail, i.e., $\mathcal{A}[i] \geq 2$. For every item j, let $\hat{\mathbb{I}} = \mathbb{I}[\left|X_m^j\right| \geq 2]$, then*

$$\mathbb{E}[|X_m^j| \cdot |X_n^j|^2 \cdot \hat{\mathbb{I}}] \leq 3 \cdot \varepsilon_1^{-1} \cdot \mathbb{E}[(|X_m^j| + |X_n^j \cap S|)^3 \cdot \hat{\mathbb{I}}] + \varepsilon_1^{-2} \cdot \mathbb{E}[|X_m^j| \cdot \hat{\mathbb{I}}]$$

Clearly, showing that $\mathbb{E}\left[\left(\left|X_m^j\right| + \left|X_n^j \cap S\right|\right)^3\right]$ is small implies that $\mathbb{E}\left[\left|X_m^j\right|\right]$ is small. For every j it holds that $\mathbb{E}\left[\left|X_m^j\right|\right] \leq \frac{1}{1-\frac{m}{n\cdot\varepsilon_1^2}}$ (this follows from simple ideas from percolation theory, see full version for details). Similar techniques can be used to prove Lemma 5.8.

Lemma 5.8. *For every $\mathcal{A}$, let $S \subseteq [n]$ denote the buyers i with at least two items in their tail. There exists a (universal) constant c so that for every item j, $\mathbb{E}\left[\left(\left|X_m^j\right| + \left|X_n^j \cap S\right|\right)^3\right] \leq \frac{1+c\cdot\frac{m}{n\cdot\varepsilon_1^2}}{1-c\cdot\frac{m}{n\cdot\varepsilon_1^2}}$*

Lemma 5.8 is then used to prove that Tail 3 contributes only a tiny fraction. Let $\varepsilon_7 = \varepsilon_7(m, n, \varepsilon_1)$ be so that for the universal constant c from Lemma 5.8, $\varepsilon_7 = \frac{4}{n\cdot\varepsilon_1^4} + \frac{m}{n} \cdot \frac{6\cdot c}{\varepsilon_1^3\cdot\left(1-c\cdot\frac{m}{n\cdot\varepsilon_1^2}\right)} + \frac{m}{n} \cdot \frac{1}{\varepsilon_1^4\cdot\left(1-\frac{m}{n\cdot\varepsilon_1^2}\right)}$.

Lemma 5.9. *(Tail 3) For every j: $\mathbb{E}\left[\left|X_m^j\right| \cdot \left|X_n^j\right|^2 \cdot \mathbb{I}\left[\left|X_m^j\right| \geq 2\right]\right] \leq \varepsilon_7$*

5.5 Proof of Theorem 5.1

Combine the core-tail decomposition (Lemma 5.3) with Lemma 5.4, Lemma 5.5 with $p = \bar{\pi}_j$, Lemma 5.6, and Lemma 5.9 to get:

$$\textsc{Rev}\left(\mathcal{D}^n\right) \leq \left(1 - e^{-\varepsilon_4/\varepsilon_1}\right)^{-1} \cdot \sum_{j\in[m]} \mathbb{E}\left[(1-\bar{\pi}_j)\right] \cdot \textsc{Rev}(D_j^{\varepsilon_4 \cdot n}) + (1+\varepsilon_5) \cdot \sum_{j\in[m]} \mathbb{E}\left[\textsc{Rev}_{\bar{\pi}_j}(D_j^{(1-\varepsilon_4)\cdot n})\right]$$
$$+ 4 \cdot \varepsilon_6 \cdot \sum_{j\in[m]} \textsc{Rev}\left(\mathcal{D}_j^n\right) + 4 \cdot \varepsilon_7 \cdot \sum_{j\in[m]} \textsc{Rev}\left(\mathcal{D}_j^n\right)$$

Consider a seller that sells the items separately. Each item j is first auctioned to $(1-\varepsilon_4)\cdot n$ buyers using an optimal mechanism with restricted ex-ante probability of sale of $\bar{\pi}_j$. Whenever the mechanism does not allocate the item, the seller proceeds and sells the item to the remaining $\varepsilon_4 \cdot n$ buyers optimally. Clearly, the revenue of the suggested mechanism is at least[15] $\textsc{Rev}_{\bar{\pi}_j}(D_j^{(1-\varepsilon_4)\cdot n}) + \textsc{Rev}(D_j^{\varepsilon_4 \cdot n}) \cdot (1-\bar{\pi}_j)$. The above mechanism, by definition, does not make more than $\textsc{Rev}(D_j^n)$ revenue, hence, $\textsc{Rev}\left(\mathcal{D}^n\right) \leq \left(\max\{(1-e^{-\varepsilon_4/\varepsilon_1})^{-1}, (1+\varepsilon_5)\} + 4\cdot(\varepsilon_6+\varepsilon_7)\right) \cdot \sum_{j\in[m]} \textsc{Rev}(D_j^n)$ As required.

REFERENCES

[1] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. 2014. A Simple and Approximately Optimal Mechanism for an Additive Buyer. In *Proceedings of the 2014 IEEE 55th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, 21–30.

[2] MohammadHossein Bateni, Sina Dehghani, MohammadTaghi Hajiaghayi, and Saeed Seddighin. 2015. Revenue Maximization for Selling Multiple Correlated Items. In *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*. 95–105. DOI: http://dx.doi.org/10.1007/978-3-662-48350-3_9

[3] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S Matthew Weinberg. 2010. Pricing randomized allocations. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 585–597.

[4] Jeremy Bulow and Paul Klemperer. 1996. Auctions Versus Negotiations. *The American Economic Review* (1996), 180–194.

[5] Yang Cai and Constantinos Daskalakis. 2011. On Minmax Theorems for Multiplayer Games. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*. 217–234. DOI: http://dx.doi.org/10.1137/1.9781611973082.20

[6] Yang Cai, Nikhil R Devanur, and S Matthew Weinberg. 2016. A duality based unified approach to Bayesian mechanism design. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 926–939.

[7] Yang Cai and Zhiyi Huang. 2013. Simple and Nearly Optimal Multi-Item Auctions. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

[8] Yang Cai and Mingfei Zhao. 2017. Simple Mechanisms for Subadditive Buyers via Duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 170–183.

[9] Shuchi Chawla, Jason D Hartline, and Robert Kleinberg. 2007. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM conference on Electronic commerce*. ACM, 243–251.

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

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

[12] Shuchi Chawla and J. Benjamin Miller. 2016. Mechanism Design for Subadditive Agents via an Ex Ante Relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC '16)*. ACM, New York, NY, USA, 579–596. DOI: http://dx.doi.org/10.1145/2940716.2940756

[13] Xi Chen, Ilias Diakonikolas, Anthi Orfanou, Dimitris Paparas, Xiaorui Sun, and Mihalis Yannakakis. 2015. On the Complexity of Optimal Lottery Pricing and Randomized Mechanisms. In *IEEE 56th Annual Symposium on Foundations*

[15] The mechanism does not extract exactly this quantity, because it might sell to the $(1-\varepsilon_4)\cdot n$ buyers with ex-ante probability that is *lower* than $\bar{\pi}_j$, in which case it will make even more than the above.

of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015. 1464–1479. DOI:http://dx.doi.org/10.1109/FOCS.2015.93

[14] Xi Chen, Ilias Diakonikolas, Dimitris Paparas, Xiaorui Sun, and Mihalis Yannakakis. 2014. The Complexity of Optimal Multidimensional Pricing. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*. 1319–1328. DOI:http://dx.doi.org/10.1137/1.9781611973402.97

[15] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2014. The complexity of optimal mechanism design. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 1302–1318.

[16] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2016. Strong Duality for a Multiple-Good Monopolist. *Econometrica* (2016).

[17] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. 2017. The Competition Complexity of Auctions: A Bulow-Klemperer Result for Multi-Dimensional Bidders. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*. 343. DOI:http://dx.doi.org/10.1145/3033274.3085115

[18] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. 2017. A Simple and Approximately Optimal Mechanism for a Buyer with Complements: Abstract. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*. 323. DOI:http://dx.doi.org/10.1145/3033274.3085116

[19] Edith Elkind. 2007. Designing and learning optimal finite support auctions. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 736–745.

[20] Kira Goldner and Anna Karlin. 2016. A Prior-Independent Revenue-Maximizing Auction for Multiple Additive Bidders. In *The 12th Conference on Web and Internet Economics (WINE)*.

[21] Sergiu Hart and Noam Nisan. 2012. Approximate revenue maximization with multiple items. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 656–656.

[22] Sergiu Hart and Noam Nisan. 2013. The menu-size complexity of auctions. In *Proceedings of the 14th ACM Conference on Economics and Computation*. 565–566.

[23] Sergiu Hart and Philip J. Reny. 2012. Maximal Revenue with Multiple Goods: Nonmonotonicity and Other Observations. *Discussion Paper Series dp630, The Center for the Study of Rationality, Hebrew University, Jerusalem* (2012).

[24] Jason D Hartline. 2013. Mechanism design and approximation. (2013).

[25] Larry Kim. 2012. How Many Ads Does Google Serve In A Day? Read more at https://www.business2community.com/online-marketing/how-many-ads-does-google-serve-in-a-day-0322253. *Bussines2Community* (Nov 2012). https://www.business2community.com/online-marketing/how-many-ads-does-google-serve-in-a-day-0322253

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

[27] Xinye Li and Andrew Chi-Chih Yao. 2013. On Revenue Maximization for Selling Multiple Independently Distributed Items. *Proceedings of the National Academy of Sciences* 110, 28 (2013), 11232–11237.

[28] Siqi Liu and Christos-Alexandros Psomas. 2018. On the Competition Complexity of Dynamic Mechanism Design. In *SODA*. To appear.

[29] Will Ma and David Simchi-Levi. 2015. Reaping the Benefits of Bundling under High Production Costs. *CoRR* abs/1512.02300 (2015). http://arxiv.org/abs/1512.02300

[30] Roger B. Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.

[31] Jean-Charles Rochet and Philippe Chone. 1998. Ironing, Sweeping, and Multidimensional Screening. *Econometrica* 66, 4 (July 1998), 783–826.

[32] Tim Roughgarden, Inbal Talgam-Cohen, and Qiqi Yan. 2015. Robust Auctions for Revenue via Enhanced Competition. (2015). Working paper; a preliminary version appeared in EC'12.

[33] Aviad Rubinstein. 2016. On the computational complexity of optimal simple mechanisms. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*. ACM, 21–28.

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

[35] Andrew Chi-Chih Yao. 2017. On Revenue Monotonicity in Combinatorial Auctions. *CoRR* abs/1709.03223 (2017). arXiv:1709.03223 http://arxiv.org/abs/1709.03223

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