The Menu-Size Complexity of Auctions (Working Paper)

Sergiu Hart * Noam Nisan †

August 24, 2012

Abstract

We consider the *menu size* of auctions as a measure of auction complexity and study how it affects revenue. Our setting has a single revenue-maximizing seller selling $k \geq 2$ heterogenous items to a single buyer whose private values for the items are drawn from a (possibly correlated) known distribution, and whose valuation for sets of items ("bundles") is additive over the items. We show that the revenue may increase arbitrarily with menu size and that a bounded menu size can *not* ensure any positive fraction of the optimal revenue. (For bounded valuations, we show that a finite menu size can ensure an arbitrarily good *additive* approximation of revenue.)

The menu size turns out to "nail down" the revenue properties of *deterministic* auctions: their menu size may be at most exponential in the number of items and indeed their revenue may be larger than that achievable by the simplest types of auctions by a factor that is exponential in the number of items but no larger. In particular our results imply an infinite separation between the revenues achievable by deterministic and general randomized auctions even when selling two items, answering a question left open in Briest et al. [2010].

1 Introduction

Are complex auctions better than simple ones? Myerson's classic result (Myerson [1981]) shows that if you are aiming to maximize revenue when selling a single item, then the

*Center for the Study of Rationality, Institute of Mathematics, and Department of Economics, Hebrew University of Jerusalem. Research partially supported by a European Research Council Advanced Investigator grant.

†Microsoft Research and Hebrew University of Jerusalem. Research partially supported by a grant from the Israeli Science Foundation and by a Google grant on Electronic Markets and Auctions.

answer is "no." The optimal auction is very simple, allocating to the highest bidder (using either 1st or 2nd price) as long as he bids above a single deterministically chosen reserve price.

However, when selling multiple items the situation turns out to be more complex. There has been significant work both in the economics literature and in the computer science literature showing that, when selling multiple items, simple auctions are no longer optimal. Specifically, it is known that randomized auctions may yield more revenue than deterministic ones[1] and that bundling the items may yield higher (or lower) revenue than selling each of them separately. This is true even in the very simple setting where there is a single bidder (buyer) and where his valuation is additive over the different items.

In this paper we consider such a simple setting: a single seller, who aims to maximize his expected revenue, is selling k heterogenous items to a single buyer whose private values for the items are drawn from an arbitrary (possibly correlated) but known prior distribution, and whose value for bundles is additive over the items in the bundle. Since we are considering only a single buyer, this work may alternatively be interpreted as dealing with monopolist pricing for multiple items.[2]

In a previous paper (Hart and Nisan [2012]) we considered the case where the buyer's values for the different items are *independent*, in which case we showed that simple auctions are *approximately* optimal: Selling each item separately (deterministically) for its Myerson price extracts a constant fraction of the optimal revenue. Specifically, for the case of two items, it gives at least a half of the optimal revenue, and when selling k items, at least a $\Omega(1/\log^2 k)$ fraction of the optimal revenue. In this paper we start by showing that the picture changes completely when the item valuations are correlated.

For a prior distribution $\mathcal{F}$ on $\Re_+^k$, let us denote by $\textsc{Rev}(\mathcal{F})$ the optimal revenue achievable by an auction when selling k items to a single buyer whose values for the items are jointly distributed according to the distribution $\mathcal{F}$. We would like to compare this revenue to the one achievable by "simple" auctions. For start let us take "deterministic" as an "upper bound" of modeling "simple"; denote by $\text{DR}\textsc{ev}(\mathcal{F})$ the optimal revenue achievable by any *deterministic* auction. We have:

Theorem A *For every $k \geq 2$ there exists a k-item distribution $\mathcal{F}$ on $\Re_+^k$ such that*

$$\textsc{Rev}(\mathcal{F}) = \infty \quad \text{and} \quad \text{DR}\textsc{ev}(\mathcal{F}) = 1.$$

Notice the contrast with the independent case for which, as mentioned above, for

[1] See Hart and Reny [2012] for a most simple and transparent such example, together with a discussion of why this phenomenon can occur only when there are two or more items.

[2] Our main separation results obtained in the single-buyer setting apply directly also to multiple-buyer auctions. In Appendix 4 we discuss the multiple-buyer case further.

every $k \geq 2$ there exists a fixed c_k such that $\textsc{Rev}(\mathcal{F}) \leq c_k \cdot \textsc{DRev}(\mathcal{F})$ (for two items we have $c_2 \leq 2$).

In Briest et al. [2010] a similar question was considered in a combinatorially simpler model where the buyer has a unit demand. They gave an infinite separation for the case $k \geq 3$ and left the case $k = 2$ as an open problem, providing some partial results indicating that for this case the gap may be bounded. While the models are different, we show in Appendix 1 that the gap between our models is at most exponential in the number of items and thus our infinite separation carries over to their model as well, answering the open problem.

Of course, to get infinite revenue the support of $\mathcal{F}$ must be unbounded. However, by truncating and then rescaling the distribution one immediately gets a similar result for distributions with bounded support: deterministic auctions may get only an arbitrarily small fraction of the optimal revenue.

Corollary 1.1 *For every $k \geq 2$ and every $\varepsilon > 0$ there exists a k-item distribution $\mathcal{F}$ on $[0,1]^k$ such that*

$$\textsc{DRev}(\mathcal{F}) < \varepsilon \cdot \textsc{Rev}(\mathcal{F}).$$

Looking at the proof of this separation as well as the earlier one in Briest et al. [2010], one observes that having a large menu size seems to be the crucial attribute of the high-revenue auctions: it enables the sophisticated screening between different buyer types required for high revenue extraction. So, in the rest of this paper we focus on the ***menu size*** as a ***complexity measure of auctions*** and study the revenue extraction capabilities of auctions that are limited in their menu size. The menu size of an auction is simply the number of possible outcomes of the auction, where an outcome is a vector $(q_1, \ldots, q_k)$ in which q_i specifies the probability of allocating item[3] i. It is well known ("the taxation principle") that in our setting any auction can be put into the normal form of offering a fixed menu with a fixed price for every menu item and letting the buyer choose among these options. Notice that while deterministic auctions can have a menu size of at most $2^k - 1$ (since each q_i must be 0 or 1, and we are not counting $(0, \ldots, 0)$), randomized auctions may have either finite or infinite menu size.

For a prior distribution $\mathcal{F}$ on $\Re_+^k$, we will use $[m]$-$\textsc{Rev}(\mathcal{F})$ to denote the optimal revenue achievable by an auction whose menu size is at most m (when selling k items to a single buyer whose values for the items are distributed according to the distribution $\mathcal{F}$.) For a single item, $k = 1$, Myerson's result implies that $[1]$-$\textsc{Rev}(\mathcal{F}) = \textsc{Rev}(\mathcal{F})$, but this is not true for more than a single item: the revenue may increase as we allow the menu size to increase.

[3]Our menu-size measure does not count the "empty" allocation $(0, \ldots, 0)$ that we assume without loss of generality is always available, by individual rationality.

We start by looking at the simplest auctions according to this complexity measure, those with a single menu item. It is not difficult to verify that the optimal single-menu-item auction is always a *bundling* auction that sells the whole bundle (i.e., $q_i = 1$ for all i) for the optimal (Myerson) bundle-price: $[1]\text{-}\textsc{Rev}(\mathcal{F}) = \textsc{BRev}(\mathcal{F})$, where BRev denotes the optimal bundling revenue. Next, one easily observes that the revenue may increase at most linearly in the menu size; thus

$$[m]\text{-}\textsc{Rev}(\mathcal{F}) \leq m \cdot [1]\text{-}\textsc{Rev}(\mathcal{F}) = m \cdot \textsc{BRev}(\mathcal{F}). \tag{1}$$

Our proof of Theorem A implies that the revenue may indeed strictly increase with the menu size. A more precise estimate is the following:[4]

Theorem B *For every $k \geq 2$ there exists a k-item distribution $\mathcal{F}$ on $\Re_+^k$ with $0 < [1]\text{-}\textsc{Rev}(\mathcal{F}) < \infty$ such that for all $m \geq 1$,*

$$[m]\text{-}\textsc{Rev}(\mathcal{F}) \geq \Omega(m^{1/7}) \cdot [1]\text{-}\textsc{Rev}(\mathcal{F}).$$

The proof constructs the distribution together with the menu of allocations: for every menu entry it puts a mass point on a valuation that chooses that menu entry. The trick is to have a *large* number of menu entries, and to extract a significant amount of revenue from each one—while not letting the simple bundling auction, which uses a *single* price, do the same.

As mentioned above, deterministic auctions on k items can have menu size of at most $2^k - 1$ and thus we have $\textsc{DRev}(\mathcal{F}) \leq (2^k - 1) \cdot [1]\text{-}\textsc{Rev}(\mathcal{F})$ while $\textsc{Rev}(\mathcal{F}) \geq \lim_{m\to\infty}[m]\text{-}\textsc{Rev}(\mathcal{F}) = \infty$ so this theorem actually implies Theorem A. We continue by showing that the limited menu size—rather than not using lotteries—is the main bottleneck of deterministic auctions. Specifically we show that the exponential-in-k upper bound on deterministic revenue implied by menu size (see (1)),

$$\textsc{DRev}(\mathcal{F}) \leq [2^k - 1]\text{-}\textsc{Rev}(\mathcal{F}) \leq (2^k - 1) \cdot [1]\text{-}\textsc{Rev}(\mathcal{F}),$$

is essentially tight:

Theorem C *For every $k \geq 2$ there exists a k-item distribution $\mathcal{F}$ on $[0,1]^k$ with $0 < [1]\text{-}\textsc{Rev}(\mathcal{F}) < \infty$ such that*

$$\textsc{DRev}(\mathcal{F}) \geq \Omega\left(\frac{2^k}{k}\right) \cdot [1]\text{-}\textsc{Rev}(\mathcal{F}).$$

[4]The increase is at a polynomial rate in m. For menu sizes that are at most exponential in the number of items, the proof of Theorem C below actually shows that the growth can be almost linear.

This again is in contrast to the independent case studied in Hart and Nisan [2012] for which[5] $\text{DREV}(\mathcal{F}) \leq \text{O}(k) \cdot [1]\text{-REV}(\mathcal{F})$.

Extending the results of Hart and Nisan [2012] to the case of correlated valuations, we compare the revenue of such single-menu-item auctions to that of the other type of "simple" auctions, those that sell the items *separately* in completely separate auctions, whose revenue is denoted SREV. As opposed to the previous theorems, the bounds that we get here are the same as in the independent case and as Hart and Nisan [2012] shows are tight. Specifically, single-menu-item auctions (i.e. bundling auctions) may raise at most $\text{O}(\log k)$ times more revenue than selling the items separately and at least a $1/k$ fraction of that: $\Omega(1/\log k) \cdot [1]\text{-REV}(\mathcal{F}) \leq \text{SREV}(\mathcal{F}) \leq k \cdot [1]\text{-REV}(\mathcal{F})$. The linear upper bound on the gap holds despite the fact that, technically, the separate auction may have exponential menu size since any subset of the items may be acquired. This suggests a more refined complexity measure, presented in section 8, which we term "additive-menu-size," a measure to which all previous results extend as well.

While we have seen that auctions with bounded menu size cannot guarantee any fixed fraction of revenue, it turns out that they can provide a good *additive approximation* of revenue. For an additive approximation to make sense we need of course to normalize the range of values, so we consider without loss of generality distributions $\mathcal{F}$ on $[0,1]^k$. As in Corollary 1.1, this bounded domain does not matter for multiplicative approximation as one can easily modify the distribution of Theorem B (by truncating and then rescaling) to get, for every $m \geq 1$ and $\varepsilon > 0$, a distribution $\mathcal{F}_{m,\varepsilon}$ on $[0,1]^2$ such that

$$[m]\text{-REV}(\mathcal{F}_{m,\varepsilon}) < \varepsilon \cdot \text{REV}(\mathcal{F}_{m,\varepsilon}).$$

Yet, for additive approximation we show:

Theorem D *For every $k \geq 2$ and $\delta > 0$ there is $m_0 = (k/\delta)^{\text{O}(k)}$ such that for every k-item distribution $\mathcal{F}$ on $[0,1]^k$ and every $m \geq m_0$,*

$$[m]\text{-REV}(\mathcal{F}) \geq \text{REV}(\mathcal{F}) - \delta.$$

One may get a multiplicative version of this theorem if $\mathcal{F}$'s support is bounded both from below and from above, that is, if the ratio between the highest and the lowest values is bounded. In particular if $\mathcal{F}$ is a distribution on, say, $[1,H]^k$ then this theorem implies that $[m]\text{-REV}(\mathcal{F}) \geq (1-\delta)\text{REV}(\mathcal{F})$ whenever[6] $m > m_0$ for $m_0 = (kH/\delta)^{\text{O}(k)}$.

[5] If one looks at the gap between $\text{DREV}(\mathcal{F})$ and the revenue obtained by selling the items separately $\text{SREV}(\mathcal{F})$, then the same exponential gap holds, a rare *doubly-exponential* contrast with the independent case considered in Hart and Nisan [2012] where $\text{DREV}(\mathcal{F}) \leq \text{O}(\log^2 k) \cdot \text{SREV}(\mathcal{F})$.

[6] In Appendix 2, we tighten the dependence on H to be poly-logarithmic, for the case of two items.

Organization of the Paper

Next, in Section 2, we briefly go over related previous work; and then in Section 3 we define our model and notations. We start in Section 4 by listing the basic properties of the menu-size complexity measure. Next, in Section 5 we describe our basic reduction from questions of auction revenue to combinatorial constructions and as a demonstration we also provide a construction that implies Theorem C. The heart of the paper, which is the combinatorial construction implying Theorems A and B, appears in Section 6. Section 7 provides positive results on approximation by simple auctions and proves Theorem D. Finally, in section, 8, we prove the relations between the menu-size notion of simple auctions and the separate auction, and introduce the more refined "additive menu size" complexity measure. Several additional results are postponed to appendices. In Appendix 1, we prove the close relationship between our model and that of Briest et al. [2010] showing why our results answer their question as well. Appendix 2 gives the tighter multiplicative approximation when the items' values are bounded both from below and from above. Appendix 3 is devoted to precise comparisons between the bundling auction and other classes of auctions, and Appendix 4 deals with multiple buyers.

2 Existing Work

Within the economics literature the issue of simple vs. complex auctions is mostly implicit in the study of randomized vs. deterministic auctions. This was first studied in McAfee and McMillan [1988] who identified sufficient conditions for the optimal mechanism to be deterministic. However, Thanassoulis [2004] and Manelli and Vincent [2006] found a technical error in the paper and exhibited counter-examples. Other examples where randomization helps, including in the case where the values of the two items are independent and identically distributed, are provided in Pavlov [2011] and Hart and Reny [2012] (the latter is essentially the simplest possible such example). In general, it is still not clear what optimal mechanisms for selling two items look like and, in particular, it is not clear when they are deterministic. Good surveys of the work within economic theory on such questions appear in Thanassoulis [2004] and Manelli and Vincent [2006], with more recent work in Fang and Norman [2006], Jehiel et al. [2007], Hart and Reny [2010], Lev [2011], Hart and Reny [2012].

This question was studied more explicitly in a line of work in computer science (Chawla et al. [2007], Chawla et al. [2010a], Chawla et al. [2010b], Daskalakis and Weinberg [2011]) that considered approximating the optimal revenue by simple mechanisms and quantified the amount of the loss incurred by simple mechanisms relative to the optimal. This was done for various settings, especially unit-demand settings and some generalizations. In

particular, in unit-demand settings, Briest et al. [2010] show that the gap between the revenue of randomized and deterministic mechanisms may be unbounded for three or more items, and leave the case of two items open, while showing that for a restricted class of randomized auctions, the gap is only constant.

3 Notation and Preliminaries

3.1 Mechanisms

A mechanism for selling k items specifies a (possibly randomized) protocol for interaction between a seller (who has no private information and commits to the mechanism) and a buyer who has a private valuation for the items. The outcome of the mechanism is an allocation specifying the probability of getting each of the k items and an (expected)[7] payment that the buyer gives to the seller. We will use the following notations:

- **Buyer valuation:** $x = (x_1, \ldots, x_k)$ where $x_i \geq 0$ denotes the value of the buyer for getting item i.
- **Allocation:** $q = (q_1, \ldots, q_k) \in [0,1]^k$, where $q_i = q_i(x)$ denotes the probability that item i is allocated to the buyer when his valuation is x (alternatively, one may interpret q_i as the fractional quantity of item i that the buyer gets).
- **Seller revenue:** $s = s(x)$ denotes the expected payment[8] that the seller receives from the buyer when the buyer's valuation is x.
- **Buyer utility:** $b = b(x)$ denotes the utility of the buyer when his valuation is x, i.e., $b(x) = \sum_i x_i q_i(x) - s(x) = x \cdot q(x) - s(x)$.

We will be discussing mechanisms that are:

- **IR – (Ex-post) Individually Rational**: $b(x) \geq 0$ for all x.
- **IC – Incentive Compatible**: For all x, x': $\sum_i x_i q_i(x) - s(x) \geq \sum_i x_i q_i(x') - s(x')$.

The IC requirement simply captures the notion that the buyer acts strategically in the mechanism. Since we are discussing a single buyer, this is in a simple decision-theoretic sense and in particular there is no distinction between the dominant strategy and the Bayes-Nash implementation notions.

[7] We only consider risk-neutral agents.

[8] In the literature this is also called transfer, cost, price, revenue, and denoted by p, t, c, etc. We hope that using the mnemonic s for the ***S***eller's final payoff and b for the ***B***uyer's final payoff will avoid confusion.

The following lemma gives well known and easily proven equivalent conditions for incentive compatibility. A short proof using our notations can be found in Hart and Nisan [2012] (see also Hart and Reny [2012] for a slightly tighter characterization).

Lemma 3.1 *The following three definitions are equivalent for a mechanism with* $b(x) = x \cdot q(x) - s(x) = \sum_i x_i q_i(x) - s(x)$*:*

1. *The mechanism is IC.*

2. *The allocation* q *is weakly monotone, in the sense that for all* x, x' *we have* $(x - x') \cdot (q(x) - q(x')) \geq 0$*, and the payment to the seller satisfies* $x' \cdot (q(x) - q(x')) \leq s(x) - s(x') \leq x \cdot (q(x) - q(x'))$ *for all* x, x'*.*

3. *The buyer's utility* b *is a convex function of* x *and for all* x *the allocation* $q(x)$ *is a subgradient of* b *at* x*, i.e., for all* x' *we have* $b(x') - b(x) \geq q(x) \cdot (x' - x)$*. In particular* b *is differentiable almost everywhere and there* $q_i(x) = \partial b(x)/\partial x_i$*.*

Note that this in particular implies that any convex function b with $0 \leq \partial b(x)/\partial x_i \leq 1$ for all i defines an incentive compatible mechanism by setting $q_i(x) = \partial b(x)/\partial x_i$ (at non-differentiability points take q to be an arbitrary subgradient of b) and $s(x) = x \cdot q(x) - b(x)$.

When $x_1, \ldots, x_k$ are distributed according to the joint cumulative distribution function $\mathcal{F}$ on[9] $\Re_+^k$, the expected revenue of the mechanism given by b is

$$R(b; \mathcal{F}) = \mathbb{E}_{x \sim \mathcal{F}}(s(x)) = \int \cdots \int \left(\sum_{i=1}^{k} x_i \frac{\partial b(x)}{\partial x_i} - b(x) \right) \mathrm{d}\mathcal{F}(x_1, \ldots, x_k).$$

3.2 Revenue

For a cumulative distribution $\mathcal{F}$ on $\Re_+^k$ (for $k \geq 1$), we consider the optimal revenue obtainable from selling k items to a (single, additive) buyer whose valuation for the k items is jointly distributed according to $\mathcal{F}$:

- $\textsc{Rev}(\mathcal{F})$ is the maximal revenue obtainable by any incentive compatible and individually rational mechanism. Formally, $\textsc{Rev}(\mathcal{F}) = sup_b R(b; \mathcal{F})$ where b ranges over all convex functions with $0 \leq \partial b(x)/\partial x_i \leq 1$ for all i and $b(0) = 0$.

- $\textsc{DRev}(\mathcal{F})$ is the maximum revenue obtainable by a *deterministic* auction, i.e., in our notation when $q_i(x) \in \{0, 1\}$ for all i and x.

[9] We write this as $x = (x_1, \ldots, x_k) \sim \mathcal{F}$.

- SREV($\mathcal{F}$) is the maximal revenue obtainable by selling each item *separately*. Formally, by Myerson's characterization, this means considering those deterministic auctions where there are single-item prices $p_1, \ldots, p_k$ such that $s(x) = \sum_i q_i(x)p_i$ for all x.

- BREV($\mathcal{F}$) is the maximal revenue obtainable by *bundling* all items together. By Myerson's characterization, this means that for each x, $q(x)$ is either $(1, \ldots, 1)$ or $(0, \ldots, 0)$.

- $[m]$-REV($\mathcal{F}$) is the maximal revenue obtainable by an auction whose menu size is at most m; that is, by auctions that have at most m possible outcomes, i.e., such that[10] $|\{q(x) : x \in \Re_+^k\} \backslash \{(0, \ldots, 0)\}| \leq m$.

4 Basic Results for Menu Size

In this section we list the basic relations related to menu-size complexity. All the statements here are immediate.

We start by showing that the optimal auction with menu of size 1 is a bundling one. Recall the notation $[m]$-REV($\mathcal{F}$) for the revenue achievable by an auction whose menu size is at most m.

Proposition 4.1 $[1]$-REV($\mathcal{F}$) = BREV($\mathcal{F}$).

Proof. Consider the optimal single-item-menu auction selling some fractions of items for some price p, and now consider selling the whole bundle for the same price p. The buyer will buy whenever he did in the original auction, so the revenue can only increase. ∎

For higher menu complexities, the revenue can increase at most linearly.

Proposition 4.2 $[m]$-REV($\mathcal{F}$) $\leq m \cdot$ BREV($\mathcal{F}$).

Proof. Just consider the menu entry that brings in the largest fraction of revenue. ∎

This is essentially tight e.g. for the case of a distribution that puts value 2^i on item i with probability 2^{-i} (and zero otherwise), where these events are disjoint. The bundling auction can get revenue of at most 2, while for each $1 \leq m \leq k$ selling each of the items $i = 1, \ldots, m$ for a price of 2^i gets revenue m. We will later see that the increase with m need not be bounded by k, and in fact even for 2 items, increasing m can increase revenue by a factor of at least $m^{1/7}$.

[10] In our counting m does not include the pay-nothing-get-nothing outcome that we always assume without loss of generality is available, by individual rationality.

This directly implies the basic relation between deterministic auctions and simple ones.

Proposition 4.3 $\text{DREV}(\mathcal{F}) \leq (2^k - 1) \cdot \text{BREV}(\mathcal{F})$.

Proof. A deterministic auction has at most $2^k - 1$ menu items, one for each non-empty set of items. ∎

For the special case of $k = 2$ we will show in Appendix 3 the tight bound $\text{DREV}(\mathcal{F}) \leq 5/2 \cdot \text{BREV}(\mathcal{F})$.

5 Combinatorial Constructions for the Support of Distributions

Our main results are gaps between the revenue achievable by different types of auctions. To find distributions for which this gap materializes we will use various combinatorial constructions for the support of the distribution.

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

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

($q' \cdot q''$ is the scalar product $\sum_{i=1}^k q'_i q''_i$).

Proposition 5.1 *For every (finite or countably infinite) sequence $q^1, q^2, \ldots$ of points in $[0,1]^k$ there exists a distribution $\mathcal{F}$ on $\Re_+^k$ such that $\text{SREV}(\mathcal{F}) \leq 2k$, $\text{BREV}(\mathcal{F}) \leq 2k$, and $\text{REV}(\mathcal{F}) \geq \sum_n \text{gap}^n$. Moreover, in the auction yielding revenue of $\sum_n \text{gap}^n$, for each n there is a menu item that gives allocation q^n and extracts expected revenue of gap^n from buyer types that choose it.*

In Appendix 3 we will state and prove a somewhat tighter relation between $\text{BREV}(\mathcal{F})$ and $\text{REV}(\mathcal{F})$, again in terms of $\sum_n \text{gap}^n$.

Proof. Without loss of generality assume that $\text{gap}^n > 0$ for all n (remove all points with non-positive gap; this can only increase the gaps for the remaining points). For each n, let $M^n = (2k)^n / (\Pi_{j=1}^n \text{gap}^j)$ (so $M^n \geq 2M^{n-1}$ since $\text{gap}^n \leq k$). For every n, let $\mathcal{F}$ be the distribution that for each n puts probability $1/M^n$ on the point $x^n := M^n q^n \in \Re_+^k$ (these probabilities sum up to at most 1 since $M^n \geq 2^n$), and puts the remaining probability on the point 0.

First notice that the revenue obtainable from any single item x_i is at most 2. This is a direct consequence of the fact that for every n the probability that x_i is at least M^n is at most $\sum_{j \geq n} 1/M^j$ (since $x_i^j = M^j q_i^j \leq M^j < M^n$ for $j < n$), and this sum is at most $2/M^n$ because the M^j grow by a factor of at least 2. A similar reasoning shows that $\text{BREV}(\mathcal{F}) \leq 2k$, as the probability that $\sum_i x_i$ is at least kM^n is at most $2/M^n$.

Now consider the auction whose menu contains, for all n in the sequence, the allocation q^n with payment $s^n := M^n \operatorname{gap}^n$. The utility that a buyer with valuation x^n gets from the j-th menu entry (q^j, s^j) is $M^n(q^j \cdot q^n) - M^j \operatorname{gap}^j$. First, for $j < n$ this utility is bounded from above by $M^n(q^j \cdot q^n)$, which by the definition of gap^n is bounded from above by $M^n(q^n \cdot q^n) - M^n \operatorname{gap}^n$, which is precisely the utility that the buyer gets from the n-th menu item (q^n, s^n). Second, for $j > n$ we have $M^n(q^j \cdot q^n) - M^j \operatorname{gap}^j \leq M^n k - M^{j-1} 2k < 0$. Thus the buyer with valuation x^n will choose the n-th menu item and pay for it $M^n \operatorname{gap}^n$, and since this happens with probability $1/M^n$, the expected revenue from this type of buyer is exactly gap^n. ∎

To demonstrate the use of this result let us show that the separation between the simple auctions of bundling or selling separately and an arbitrary deterministic auction may be exponential in the number of items; a statement that includes that of Theorem C.

Lemma 5.2 *There exists a k-item distribution $\mathcal{F}$ such that* $\text{SREV}(\mathcal{F}) \leq 2k$, $\text{BREV}(\mathcal{F}) \leq 2k$, *and* $\text{DREV}(\mathcal{F}) \geq 2^k - 1$.

Proof. Let us enumerate the $2^k - 1$ non-empty subsets L^n of $\{1, \ldots, k\}$ such that the size of L^n is weakly increasing in n, and let q^n be the indicator vector of L^n (i.e., $q_i^n = 1$ for $i \in L^n$ and $q_i^n = 0$ for $i \notin L^n$). For $j < n$ we have $q^j \cdot q^n = |L^j \cap L^n| < |L^n| = q^n \cdot q^n$, and thus $\operatorname{gap}^n \geq 1$. Use the "moreover" statement of Proposition 5.1. ∎

6 Complex Auctions May Be Infinitely Better

We will construct the sequence of points q^n for Proposition 5.1, in order to prove an infinite separation between the revenue of an arbitrary auction for $k = 2$ items and that of selling the items separately (for larger k one may add items with zero values).

Proposition 6.1 *There exists an infinite sequence $q^1, q^2, \ldots$ of points in $[0,1]^2$ with $||q^n|| \leq 1$ such that for all n,* $\operatorname{gap}^n = \Omega(n^{-6/7})$.

Proof. The sequence of points that we will build is composed of a sequence of "shells", each containing multiple points. The shells will be getting closer and closer to each other,

approaching, as the shell, N, goes to infinity, the unit sphere: All the points q^j in the N-th shell will be of length $||q^j|| = \sum_{\ell=1}^{N} \ell^{-3/2}/\alpha$, where $\alpha = \sum_{\ell=1}^{\infty} \ell^{-3/2}$ (which indeed converges). Each shell N will contain $N^{3/4}$ different points in it so the angle between any two of them is at least $\Omega(N^{-3/4})$.

Now we estimate $q^j \cdot q^{j'} = ||q^j|| \cdot ||q^{j'}|| \cdot \cos(\theta)$ where θ is the angle between q^j and $q^{j'}$. Let N denote j's shell. For $j' < j$ there are two possibilities, either $q^{j'}$ is in the same shell as q^j or it is in a smaller shell. In the first case we have $\theta \geq \Omega(N^{-3/4})$ and thus $\cos(\theta) \leq 1 - \Omega(N^{-3/2})$ (since $\cos(x) = 1 - x^2/2 + x^4/24 - \ldots$) and since $||q^j|| = \Theta(1)$ we have $q^j \cdot q^j - q^{j'} \cdot q^j \geq \Omega(N^{-3/2})$. In the second case, $||q^j|| - ||q^{j'}|| \geq N^{-3/2}$, so again since $||q^j|| = \Theta(1)$ we have $(q^j \cdot q^j - q^{j'} \cdot q^j \geq \Omega(N^{-3/2})$. Thus for any point q^j in the N-th shell we have $\text{gap}^j = \Omega(N^{-3/2})$. Since the first N shells contain together $\sum_{\ell=1}^{N} \ell^{3/4} = \Theta(N^{7/4})$ points, we have $j = \Theta(N^{7/4})$ and thus $\text{gap}^j = \Omega(N^{-3/2}) = \Omega(j^{-6/7})$. ∎

This directly implies Theorems A and B.

Proof of Theorems A and B and Corollary 1.1. Take the countably infinite sequence of points q^n constructed in Proposition 6.1, and apply Proposition 5.1 to get the distribution $\mathcal{F}$. An auction that has only the first m menu items $q^1, \ldots, q^m$ will extract $\sum_{j=1}^{m} \text{gap}^j = \Omega(\sum_{j=1}^{m} j^{-6/7}) = \Omega(m^{1/7})$ from $\mathcal{F}$ (by the "moreover" statement of Proposition 5.1), and thus $[m]$-$\textsc{Rev}(\mathcal{F}) \geq \Omega(m^{1/7})$. However for this $\mathcal{F}$, $\textsc{SRev}(\mathcal{F}) \leq 4$ and $\textsc{BRev}(\mathcal{F}) \leq 4$, so this proves Theorem B. As $\textsc{DRev}(\mathcal{F}) \leq 3 \cdot \textsc{BRev}(\mathcal{F}) \leq 12$ by Proposition 8.5, scaling $\mathcal{F}$ by the constant factor 12 yields Theorem A. Finally, for Corollary 1.1 we take the distribution $\mathcal{F}_n$ generated by Proposition 5.1 for the finite sequence $q^1, \ldots, q^n$ with large enough n. ∎

7 Additive Approximation

Theorem D is a direct corollary of the following lemma.

Lemma 7.1 *Let n be a positive integer and $m = (n+1)^k - 1$. Then for every distribution $\mathcal{F}$ on $[0,1]^k$ we have $[m]$-$\textsc{Rev}(\mathcal{F}) \geq Rev(\mathcal{F}) - (2k)/\sqrt{n}$.*

Proof. Let $\mathcal{F}$ have support in $[0,1]^k$, and let b be a mechanism with corresponding (q, s).

Define a new mechanism $\tilde{b}$ as follows: for each $x \in [0,1]^n$, let $\tilde{q}(x)$ be the rounding up of $q(x)$ to the $1/n$-grid on[11] $[0,1]^k$, and let $\tilde{s}(x) := (1 - 1/\sqrt{n})s(x)$. Since $\tilde{q}$ can take at most $(n+1)^k$ different values, the menu size of $\tilde{b}$ is at most $(n+1)^k - 1 = m$.

If $\tilde{q}(x) \cdot x - \tilde{s}(x) \leq \tilde{q}(y) \cdot x - \tilde{s}(y)$ then (recall that $q(x) \cdot x - s(x) \geq q(y) \cdot x - s(y)$) we must have $(1/n) \sum_{i=1}^{k} x_i \geq (1/\sqrt{n})(s(x) - s(y))$, hence $s(y) \geq s(x) - k/\sqrt{n}$ (since $\sum_i x_i \leq k$),

[11] I.e., $\tilde{q}_i(x) = \lceil nq_i(x) \rceil / n$ for each i and x.

which implies that the seller's revenue at x from $\tilde{b}$ must be $\geq (1-1/\sqrt{n})(s(x)-k/\sqrt{n})$. Therefore $R(\tilde{b};\mathcal{F}) \geq (1-1/\sqrt{n})R(b;\mathcal{F}) - k/\sqrt{n} \geq R(b;\mathcal{F}) - 2k/\sqrt{n}$ (since $R(b;\mathcal{F}) \leq \sum_i x_i \leq k$). ∎

One may get a relative approximation result from this one by bounding the values both from above and from below. Specifically by letting $\mathcal{F}$ be a distribution on $[1,H]^k$. In this case, the following variant of Theorem D may be stated:

Theorem 7.2 *For every $k \geq 2$, $\delta > 0$, and $H > 0$ there is $m_0 \geq (kH/\delta)^{O(k)}$ such that for every k-item distribution $\mathcal{F}$ on $[1,H]^k$ and every $m \geq m_0$,*

$$[m]\text{-}\textsc{Rev}(\mathcal{F}) \geq (1-\delta)\cdot \textsc{Rev}(\mathcal{F}).$$

Proof. We first scale $[1,H]$ to $[1/H,1]$, which for multiplicative approximation is the same. We then design an auction which gives an additive approximation to within ε/H, which using Theorem D requires a menu size m as stated. Now, since $\mathcal{F}$ is bounded from below by $1/H$, its revenue is similarly bounded $\textsc{Rev}(\mathcal{F}) \geq 1/H$, and thus an ε/H additive approximation is also a $(1-\varepsilon)$ multiplicative approximation as required. ∎

We do not know whether the dependence on k needs to be exponential, however we do show, in appendix 2, that, at least for $k=2$ items, the dependence on H may in fact be poly-logarithmic rather than polynomial.

8 The Separate Auction and Additive Menu Size

In this section we study the separate auction and its relationship to the bundling auction.

Proposition 8.1 $\textsc{BRev}(\mathcal{F}) \leq k \cdot \textsc{SRev}(\mathcal{F})$.

Proof. Let BRev be achieved with price p. If the separate auction offers price p/k for each item then whenever $\sum_i X_i \geq p$ we have $X_i \geq p/k$ for some i, and so one of the k items will be acquired in the separate auction. ∎

This is tight for $k=2$ and correlated values. Consider the distribution of $(x,1-x) \in \Re_+^2$ where x is chosen uniformly in $[0,1]$. The bundling auction can get all the value by offering price 1 and selling always. Each item is distributed uniformly in $[0,1]$ so the best obtainable revenue for each is $1/4$, obtained at price $1/2$.

For larger values of k, we can get a stronger result.

Proposition 8.2 $\textsc{BRev}(\mathcal{F}) \leq O(\log k)\cdot \textsc{SRev}(\mathcal{F})$.

Proof. Let $\text{BREV}(\mathcal{F})$ be achieved with price p. We can first assume without loss of generality that for each x in the support of $\mathcal{F}$, either $\sum_i x_i = p$ or $\sum_i x_i = 0$. (This is without loss of generality since for this new "truncated distribution" $\mathcal{F}'$, $\text{BREV}(\mathcal{F}') = \text{BREV}(\mathcal{F})$ while $\text{SREV}(\mathcal{F}') \leq \text{SREV}(\mathcal{F})$ using Myerson's characterization.) We can now make another assumption without loss of generality, that in fact we always have that $\sum_i x_i = p$ in the support of $\mathcal{F}$ (so in fact $\text{BREV}(\mathcal{F}) = p$). (This is without loss of generality as if we look at the conditional distribution on those x's with $\sum_i x_i = p$, both BREV and SREV increase by a factor of exactly $Pr[\sum_i x_i = p]$.)

At this point there are two different ways to proceed; we present both, as they may lead to different extensions.

Proof 1: Denote $e_i = \mathbb{E}_{X\sim\mathcal{F}}(X_i)$ be the expected value of the i'th item, so now (using our assumptions) $\sum_i e_i = p$. The claim is that the i'th item can be sold in a separate auction yielding a revenue of at least $(e_i - p/(2k))/(2(1 + log_2 k))$. The lemma is implied by summing over all i. Consider the distribution of X_i and split the range of values of X_i into $(2 + \log_2 k)$ sub-ranges: A "low" range for which $X_i \leq p/(2k)$ and for each $j = 0 \ldots \log_2 k$ a subrange where: $p/(2^{j+1}) < X_i \leq p/(2^j)$ (notice that since $X_i \leq p$ we have covered the whole support of X_i). The low subrange contributes at most $p/(2k)$ to the expectation of X_i, and thus one of the other $1 + \log_2 k$ sub-ranges contributes at least $((e_i - p/(2k))/(1 + \log_2 k)$ to this expectation. The lower bound of this sub-range, $p/(2^{j+1})$, is smaller by a factor of at most 2 than any value in the subrange, so asking it as the item price will yield revenue which is at least half of the contribution of this sub-range to the expectation, i.e. at least $((e_i - p/(2k))/(2(1 + \log_2 k))$.

Proof 2: Let $r_i := \text{REV}(X_i) = \sup_t t(1 - F_i(t))$ (where F_i is the i-th marginal of $\mathcal{F}$), then $1 - F_i(t) \leq \min\{r_i/t, 1\}$ and so

$$\mathbb{E}[X_i] = \int_0^\infty (1 - F_i(t))dt \leq \int_0^{r_i} 1dt + \int_{r_i}^{p} \frac{r_i}{t}dt = r_i(1 + \log p - \log r_i).$$

The function $r(1 + \log p - \log r)$ is concave in r, hence

$$\frac{p}{k} = \frac{1}{k}\sum_{i=1}^{k} \mathbb{E}[X_i] \leq \bar{r}\,(1 + \log p - \log \bar{r}),$$

where $\bar{r} := (1/k)\sum_i r_i$. Thus $p/(k\bar{r}) \leq 1 + \log k + \log(p/(k\bar{r}))$, from which it follows that[12] $\text{BREV}(\mathcal{F})/\text{SREV}(\mathcal{F}) = p/(k\bar{r}) < 4\log k$. ∎

The relation in the opposite direction is simple and given by:

Proposition 8.3 $\text{SREV}(\mathcal{F}) \leq k \cdot \text{BREV}(\mathcal{F})$.

[12]The function $f(x) = x - \log x - 1 - \log k$ is increasing, and $f(4\log k) > 0$ for every $k \geq 2$.

Proof. Let F_i denote the marginal distribution on the i'th item, so by definition $\text{SREV}(\mathcal{F}) = \sum_i \text{REV}(F_i)$. But the optimal auction for item i separately has, by Myerson, a single menu entry so is at most $[1]\text{-REV}(\mathcal{F}) = \text{BREV}(\mathcal{F})$. ∎

This is tight for every k, even for independent items; see Hart and Nisan [2012].

Combining with Proposition 4.3 we get the following corollary, which by Theorem C is essentially tight.

Corollary 8.4 $\text{DREV}(\mathcal{F}) \leq \text{O}(2^k \log k) \cdot \text{SREV}(\mathcal{F})$.

For the special case of $k = 2$ items, we have the somewhat tighter

Proposition 8.5 *For $k = 2$,* $\text{DREV}(\mathcal{F}) \leq 3 \cdot \text{SREV}(\mathcal{F})$.

Proof. Take the optimal deterministic auction which has three menu items: either selling just one of the items, or selling the bundle. Its revenue is the sum of that obtained for a single item and that obtained for the bundle, thus $\text{DREV}(\mathcal{F}) \leq \text{SREV}(\mathcal{F}) + \text{BREV}(\mathcal{F})$. The proof is completed using Proposition 8.1. ∎

Notice that the separate auction has menu size $2^k - 1$ (the buyer can buy any set $L \subset \{1, \ldots, k\}$ for $\sum_{i \in L} p_i$), and yet its revenue is bounded by only k times that of the bundling revenue rather than $2^k - 1$ (as is the case for general deterministic auctions; recall Proposition 4.3). Intuitively, this seems connected to the fact that the separate auction only has k "degrees of freedom" or "parameters" (the k prices). This may be formalized by defining a more refined "additive-menu-size" complexity:

- An auction has *additive menu size* m if it offers at most m "basic" menu entries to the buyer, each with its own price, but here the buyer may buy any *combination* of basic menu entries, provided that no item is allocated more than once in total, for the sum of the entry prices.[13]
- $[m]^*\text{-REV}(\mathcal{F})$ is the maximal revenue obtainable by an auction whose additive menu size is at most m.

Clearly $[m]^*\text{-REV}(\mathcal{F}) \geq [m]\text{-REV}(\mathcal{F})$, and $[1]^*\text{-REV}(\mathcal{F}) = [1]\text{-REV}(\mathcal{F}) = \text{BREV}(\mathcal{F})$ (cf. Proposition 4.1). Interestingly, the linear bound of Proposition 4.2 holds here too:

Proposition 8.6 $[m]^*\text{-REV}(\mathcal{F}) \leq m \cdot \text{BREV}(\mathcal{F})$.

[13] Formally, (q, s) is a *combination* of $(q^1, s^1), \ldots, (q^n, s^n)$ if $q = q^1 + \cdots + q^n$, $s = s^1 + \cdots + s^n$, and $q \in [0, 1]^k$ (i.e., (q, s) is a feasible choice).

Proof. For each $n = 1, \ldots, m$, let s^n be the seller's payment from the n-th basic menu item, and let α^n be the *total* probability that the buyer takes this menu item (thus $\sum_n \alpha^n$ may be as high as k). The total revenue is $\sum_n s^n \alpha^n$, and, as in the proofs of Propositions 4.1 and 4.2, each term is bounded from above by $\text{BREV}(\mathcal{F})$. ∎

This essentially implies that the results in this paper apply also to this more refined complexity measure as well. This complexity measure really captures the power of selling the items separately as the additive menu complexity of the separate auction is at most k (cf. Propositions 8.3 and 8.6).

Acknowledgments

The authors would like to thank Motty Perry and Phil Reny for introducing us to the subject and for many useful discussions.

Appendix 1: The Unit Demand Model

In this section we shortly compare our model to the unit-demand model considered in several papers, in particular in Briest et al. [2010]. There are k items for sale and a single buyer. There are two basic differences between our model and the unit-demand one. First, in the unit-demand model, the buyers are modelled as having unit-demand valuations. Additionally, the unit-demand model was defined to only allow the auction to offer single items rather than bundles of items as in our model. This second restriction does not turn out to matter.

More formally, in the unit demand model there is a single buyer with a unit demand valuation, i.e., for a set $L \subset \{1, \ldots, k\}$ of items, its value is $\max_{i \in S} x_i$ (rather than $\sum_{i \in L} x_i$). A deterministic auction is this setting would offer a price p_i for each item i. For unit-demand buyers this is equivalent to a completely general deterministic auction as there is no need to offer prices for bundles since the buyer is not interested in them. Thus for example an auction offering price p_1 for item 1, p_2 for item 2 and p_{12} for both items, would be the same as asking price $\min(p_1, p_{12})$ for item 1 and price $\min(p_1, p_{12})$ for item 2.

A randomized auction in this model is allowed to offer a set of lotteries, each with its own price, where a lottery is a vector of probabilities $\alpha_1, \ldots, \alpha_k$ of getting the items, with $\sum_i \alpha_i \leq 1$ (in contrast to our additive buyer, where $q_i \leq 1$ for each i). Again, for unit-demand buyers this is equivalent to general randomized auctions that are also allowed to offer lotteries for bundles of items. For example a menu item offering the lottery "item 1

with probability 2/9; item 2 with probability 3/9; and both items with probability 4/9" for a certain price can be replaced by the two menu items "item 1 with probability 6/9; item 2 with probability 3/9" and "item 1 with probability 2/9; item two with probability 7/9", each for the same price as the original menu item.

Let us use the notation $\textsc{Rev}^{\text{UD}}(\mathcal{F})$ to denote the revenue obtainable from a unit demand buyer whose valuation for the k single items is distributed according to $\mathcal{F}$. Similarly $\textsc{DRev}^{\text{UD}}(\mathcal{F})$ will denote the revenue achievable by deterministic auctions. We can compare these revenues to those achievable in our model from an additive buyer whose valuation for the k single items is distributed according to the same $\mathcal{F}$.

Proposition 9.1 *For every $k \geq 2$ and every k-item distribution $\mathcal{F}$ on $\Re_+^k$,*

- $\textsc{Rev}^{\text{UD}}(\mathcal{F}) \leq \textsc{Rev}(\mathcal{F}) \leq k \cdot \textsc{Rev}^{\text{UD}}(\mathcal{F})$.
- $\textsc{DRev}^{\text{UD}}(\mathcal{F}) \leq \textsc{DRev}(\mathcal{F}) \leq k2^k \cdot \textsc{DRev}^{\text{UD}}(\mathcal{F})$.

Proof. The lower bounds in both cases follow since any auction in the unit-demand model offers only unit-demand menu entries, and on these the unit-demand buyer and the additive buyer have the same preferences; thus offering the same menu in our setting would give exactly the same revenue as it does in the unit-demand setting.

For the upper bound for randomized algorithms notice that if we replace each menu entry of an auction in our model that asks price s for allocation $(q_1, \ldots, q_k)$ (where $0 \leq q_i \leq 1$ for each i) by a menu entry asking price s/k for the allocation $(q_1/k, \ldots, q_k/k)$, then we do not change the preferences of the buyer between the different menu items, and thus the revenue drops by a factor of exactly k. However, the new auction only gives unit-demand allocations, and moreover the unit-demand buyer and the additive buyer behave the same.

For the deterministic upper bound we start with a deterministic auction in our model. Since it has at most $2^k - 1$ menu entries, a fraction of at least 2^{-k} of the revenue must come when allocating one of them to, say, the set L of items. An auction that only offers to sell the set L of items for the same price s as the original one did, will thus make at least a 2^{-k} fraction of the revenue of the original one. Now consider the unit-demand auction that offers the price $s/|L|$ for each of the items in L: whenever the additive buyer in the additive auction bought L we are guaranteed that his value for at least one of the items in L was at least $s/|L|$, in which case the unit-demand buyer will also acquire an item. ∎

The interesting gap in the above lemma is the exponential one for deterministic auctions, and indeed we can show that this is tight.

Lemma 9.2 *There exists a k-item distribution $\mathcal{F}$ such that* $\textsc{DRev}(\mathcal{F}) \geq \Omega(2^k/k) \cdot \textsc{DRev}^{\text{UD}}(\mathcal{F})$.

Proof. Take the distribution $\mathcal{F}$ constructed in the proof of Lemma 5.2 that satisfies $\textsc{SRev}(\mathcal{F}) \leq 2k$ and $\textsc{DRev}(\mathcal{F}) \geq 2^k - 1$. Now $\textsc{DRev}^{\text{UD}}(\mathcal{F}) \leq \textsc{SRev}(\mathcal{F})$ as the item prices used in any deterministic auction in the unit-demand model can only yield more revenue in our additive model where the buyer may buy more than a single item. ∎

Despite the exponential separation, for finite k it is still constant, and so a super-constant separation between randomized and deterministic auctions in our setting is equivalent to the same separation in the unit-demand setting.

Appendix 2: Multiplicative Approximation with Bounded Domains

In this appendix we give a stronger version of Theorem 7.2, with a poly-logarithmic dependence on H, for the case of two items.

Theorem 10.1 *For every $\delta > 0$ and $H > 0$ there is $m_0 \geq \Omega(\delta^{-5} \log^2 H)$ such that for every two-item distribution $\mathcal{F}$ on $[1, H]^2$ and every $m \geq m_0$,*

$$[m]\text{-}\textsc{Rev}(\mathcal{F}) \geq (1 - \delta) \cdot \textsc{Rev}(\mathcal{F}).$$

This is a direct corollary of the following lemma.

Lemma 10.2 *For every two-item mechanism (q, s), every $H > 1$ and every $\delta > 0$, there exists a mechanism $(\tilde{q}, \tilde{s})$ with menu size at most $\mathrm{O}(\delta^{-5} \log^2 H)$, such that for every x with $1 \leq s(x) \leq H$ we have $\tilde{s}(x) \geq (1 - \delta)s(x)$.*

Proof. We will discretize the menu of the the original mechanism. Our first step will be to discretize the payments s, and the second to discretize the allocations $q = (q_1, q_2)$.

We start by splitting the range $[1, H]$ into K sub-ranges, each of them with ratio at most $H^{1/K}$ between its end-points, where K is chosen so that $H^{1/K} \leq \delta^2$, i.e. $K = \mathrm{O}(\delta^{-2} \log H)$. We define a real function $\phi(s)$ by rounding s up to the top of its range and then multiplying by $1 - \delta$. So we have $(1 - \delta)s < \phi(s) < (1 - \delta)(1 + \delta^2)s$. Then for any $s' \leq s(1 - \delta)$ we have $\phi(s) - \phi(s') \leq (1 - \delta)(s - s' + s\delta^2) \leq (s - s') - s\delta^2 + s\delta^2 = s - s'$.

Now we take every menu entry (q, s) of the original mechanism and replace s with $\phi(s)$. The previous property of ϕ ensures that any buyer who previously preferred (q, s) to some other menu item (q', s') with $s' < (1 - \delta)s$ will still prefer $(q, \phi(s))$ in the new menu; thus

in the new menu he will pay $\phi(s')$ for some $s' \geq (1-\delta)s$, and $\phi(s') > (1-\delta)s' \geq (1-\delta)^2 s$; his payment in the new menu is therefore at least $(1-\delta)^2$ times his payment in the original menu.

We now have a menu with only K distinct price levels $s^1 < \cdots < s^K$. Before we continue, we scale it down by a factor of $(1-\delta)$, i.e. multiply both the q's and the s's by $(1-\delta)$. This does not change the menu choice of any buyer; reduces the payments by a factor of exactly $1-\delta$; and ensures that $q_1, q_2 \leq 1-\delta$. We now round down each q_1 and each q_2 to an integer multiple of δ/K, and then add $\delta i/K$ to each menu entry whose price is s^i. Notice that rounding down reduced each q by at most δ/K, and since higher paying menu entries got a boost which is larger by at least δ/K than any lower paying menu entry, any buyer that previously chose an entry paying s, can now only choose an entry paying some $s' \geq s$.

All in all we have reached a new mechanism whose payment is at least $(1-\delta)^3 \geq 1-3\delta$ times that of the original one (so redefine the δ in the proof to be 1/3 of the δ in the statement). There are $K = \mathrm{O}(\delta^{-2} \log H)$ price levels and $\delta^{-1}K = \mathrm{O}(\delta^{-3} \log H)$ different allocation levels for both q_1 and for q_2. However notice that for a fixed price level s and a fixed q_1 there can only be a single value of q_2 that is actually used in the menu (as lower ones will be dominated), and so the total number of possible allocations is $\mathrm{O}(\delta^{-5} \log^2 H)$.

∎

Notice that the poly-logarithmic dependence of m on H is "about right" since the distribution $\mathcal{F}$ induced by the first m_0 points in the construction used for proving theorem A has $H = m_0^{\mathrm{O}(m_0)}$, and the $m_0^{1/7}$ gap between $\textsc{Rev}(\mathcal{F})$ and [1]-$\textsc{Rev}(\mathcal{F})$ implies that for $m = \mathrm{O}((\log H)^{1/8})$ we have [m]-$\textsc{Rev}(\mathcal{F}) = o(\textsc{Rev}(\mathcal{F}))$.

Appendix 3: Exact Comparisons to Bundling

In this section we will provide bounds on the bundling revenue: on the one hand, it can be arbitrarily low relative to the optimal revenue, and on the other hand it guarantees a certain fixed fraction (such as 2/3, or 1/2, or 5/2) of the optimal revenue over certain interesting subclasses of mechanisms (such as deterministic, or separate prices, mechanisms).

The following proposition provides a precise tool for all these comparisons.

Proposition 11.1 *Let* $b(x) = \max_{0 \leq n \leq m}\{q^n \cdot x - s^n\}$ *be a finite* k*-dimensional mechanism, where* $q^n \in [0,1]^k$, $s^n \geq 0$, $q^0 = (0, \ldots, 0)$, $s^0 = 0$, *and the indices are such that*

$s^{n-1} \leq s^n$ *for all* $n \geq 1$. *For each* $n \geq 1$ *let* $t^n := \inf\{\sum_{i=1}^k x_i : s(x) \geq s^n\}$, *and define*

$$T(b) := \sum_{n=1}^{m} \frac{s^n - s^{n-1}}{t^n}$$

(where we take $(0-0)/0$ *as* 0, *and the* inf *of an empty set as*[14] ∞*). Then:*

(i) *For any* k*-item distribution* $\mathcal{F}$ *on* $\Re_+^k$,

$$R(b; \mathcal{F}) \leq T(b) \cdot \text{BREV}(\mathcal{F}). \tag{2}$$

(ii) *There exists a* k*-item distribution* $\mathcal{F}$ *on* $[0,1]^k$ *such that*

$$R(b; \mathcal{F}) = T(b) \cdot \text{BREV}(\mathcal{F}). \tag{3}$$

Moreover, if b *is symmetric (i.e.,* $b(x) = b(\pi x)$ *for every* $x \in \Re_+^k$ *and every permutation of the indices*[15] π*, then* $\mathcal{F}$ *can be taken to be a symmetric distribution*[16] *(i.e.,* $\mathcal{F}(x) = \mathcal{F}(\pi x)$ *for every* $x \in \Re_+^k$ *and every permutation of the indices* π*).*

Before proving the proposition it is instructive to see the computation of $T(b)$ in a few examples with $k = 2$ items.

- For $b(x) = \max\{0, x_1 - 1, x_2 - 2, x_1 + x_2 - 4\}$ we have $(s^0, s^1, s^2, s^3) = (0, 1, 2, 4)$ and $(t^1, t^2, t^3) = (1, 2, 5)$ (attained at the points $(1, 0), (0, 2)$, and $(2, 3)$, respectively), and so $T(b) = (1-0)/1 + (2-1)/2 + (4-2)/5 = 12/5$.
- For $b(x) = \max\{0, x_1 - 2, x_2 - 2, x_1 + x_2 - 4\}$ we have $(s^0, s^1, s^2, s^3) = (0, 2, 2, 4)$ and $(t^1, t^2, t^3) = (2, 2, 4)$ ($t^1 = t^2 = 2$ is attained at both $(2, 0)$ and $(0, 2)$, and $t^3 = 4$ at $(2, 2)$), and so $T(b) = (2-0)/2 + (2-2)/2 + (4-2)/4 = 3/2$.
- For $b(x) = \max\{0, x_1 - 5, x_2 - 2, x_1 + x_2 - 4\}$ we have $(s^0, s^1, s^2, s^3) = (0, 2, 4, 5)$ and $(t^1, t^2, t^3) = (2, 4, \infty)$, and ($t^1 = 2$ is attained at $(0, 2)$ and $t^2 = 4$ at $(x_1, 4 - x_1)$ with $2 \leq x_1 \leq 4$), and so $T(b) = (2-0)/2 + (4-2)/4 + (5-4)/\infty = 3/2$.

In general, the t^n are computed as follows: First, for each $n \geq 1$ let τ^n be the minimum of $x_1 + \cdots + x_k$ over the set $\{x \geq 0 : q^n \cdot x - s^n \geq q^j \cdot x - s^j \text{ for all } j\}$; this is a linear programming problem. Then, for each $n \geq 1$ put $t^n = \min\{\tau^j : s^j \geq s^n\}$.

[14] $t^n = 0$ if and only if $s^n = 0$ (and then $s^{n-1} = 0$). The sum in $T(b)$ can thus be started with $n = n_0$, the first index with $s^{n_0} > 0$, and ending at $n = m_0$, the first index with $t^{m_0+1} = \infty$, where we put $t^{m+1} := \infty$.

[15] I.e., π is a one-to-one mapping from $\{1, \ldots, k\}$ onto itself, and $\pi(x_1, \ldots, x_k) = (x_{\pi(1)}, \ldots, x_{\pi(k)})$.

[16] Also known as an "exchangeable" distribution.

Proof of Proposition 11.1. If $s^n = 0$ for all n then $R(b;F) = 0$ for all F and $T(b) = 0$ and there is nothing to prove. Otherwise $T(b) > 0$, and, as in footnote 14, we only consider those n between n_0 and m_0 for which $0 < t^n < \infty$; for simplicity assume that these are all n, i.e., $1 \le n \le m$. Let X be a k-dimensional random variable with distribution $\mathcal{F}$. By definition of t^n we have $[\sum_i X_i \ge t^n] \supset [s(X) \ge s^n]$, and so

$$\textsc{BRev}(\mathcal{F}) \ge t^n \cdot \mathbb{P}\left[\sum_{i=1}^{k} X_i \ge t^n\right] \ge t^n \cdot \mathbb{P}\left[s(X) \ge s^n\right] = t^n \cdot \sum_{j=n}^{m} \sigma^j,$$

where $\sigma^j := \mathbb{P}\left[s^j \le s(X) < s^{j+1}\right]$ (thus $\sigma^j = 0$ when $s^j = s^{j+1}$, and $\sigma^j = \mathbb{P}\left[s(X) = s^j\right]$ when $s^j < s^{j+1}$). Therefore

$$\begin{aligned}\sum_{i=1}^{m} \frac{s^n - s^{n-1}}{t^n} \textsc{BRev}(\mathcal{F}) &\ge \sum_{i=1}^{m} (s^n - s^{n-1}) \sum_{j=i}^{m} \sigma_j = \sum_{j=1}^{m} \sigma_j \sum_{i=1}^{j} (s^n - s^{n-1}) \\ &= \sum_{j=1}^{m} \sigma_j s_j = \mathbb{E}\left[s(X)\right] = R(b;\mathcal{F}).\end{aligned}$$

(recall that $s_0 = 0$). The left-hand side is precisely $T(b) \cdot \textsc{BRev}(\mathcal{F})$, and we have proved (2).

To construct a distribution for which we have equality, for each n let $x^n \in \Re^k_+$ be a point where the minimum t^n is attained; thus $\sum_i x^n_i = t^n$ and $s(x^n) \ge s^n$. Let $\mathcal{F}_b$ put probability $\alpha^n := t^1/t^n - t^1/t^{n+1} \ge 0$ on the point x^n; note that $\sum_{j \ge n} \alpha^j = t^1/t^n$, (recall that $t^{m+1} = \infty$) and in particular $\sum_{j=1}^{\infty} \alpha^j = 1$. The sequence t^n is nondecreasing (since the sequence s^n is nondecreasing), and so the revenue of the bundled mechanism $b^n(x) = \max\{0, \sum_i x_i - t^n\}$ is $R(b^n;\mathcal{F}_b) = t^n \sum_{j \ge n} \alpha^j = t^n(t^1/t^n) = t^1$. Therefore $\textsc{BRev}(\mathcal{F}_b) = t^1$ (prices different from some t^n can only yield lower revenues). For the original mechanism b, its revenue at x^n is $s(x^n) \ge s^n$, and so

$$\begin{aligned}R(b;\mathcal{F}_b) &= \sum_{n=1}^{m} s(x^n)\ \alpha^n \ge \sum_{n=1}^{m} s^n \left(\frac{t^1}{t^n} - \frac{t^1}{t^{n+1}}\right) = t^1 \sum_{n=1}^{m} \left(\frac{s^n}{t^n} - \frac{s^{n-1}}{t^n}\right) \\ &= T(b)\ t^1 = T(b)\ \textsc{BRev}(\mathcal{F}_b) \ge R(b;\mathcal{F}_b)\end{aligned}$$

(recall that $s^0/t^1 = s^m/t^{m+1} = 0$), where the last inequality is (2). We thus have equality throughout.

Now rescaling all s^n by $\lambda > 0$ rescales all the t^n, and thus the support of F_b, by λ, while it does not affect T. Specifically, let $\hat{b}$ be the mechanism $\{(\hat{q}^n, \hat{s}^n)\}_{n=0}^m$ where $\hat{q}^n = q^n$ and $\hat{s}^n = \lambda s^n$; then $t^n = \lambda t^n$, and so $T(\hat{b}) = T(b)$. Thus the distribution $F_{\hat{b}}$ constructed above for $\hat{b}$ also satisfies (3) (with the same $T(b)$), and its support is included

in $\{\sum_i x_i \le \hat{t}^m = \lambda t^m\}$; take $\lambda = 1/t^m$.

Finally, the revenue of any symmetric mechanism b remains the same when we "symmetrize" $\mathcal{F}$; i.e., $R(b;\tilde{\mathcal{F}}) = R(b;\mathcal{F})$, where $\tilde{\mathcal{F}}(x) := (k!)^{-1}\sum_\pi \mathcal{F}(\pi x)$ for all x. Bundled mechanisms are symmetric, thus $\text{BREV}(\tilde{\mathcal{F}}) = \text{BREV}(\mathcal{F})$, and so one can use the symmetric distribution $\tilde{\mathcal{F}}_b$ instead of $\mathcal{F}_b$ in (3). ∎

As we will see below (Proposition 11.3), the bound $T(b)$ can be arbitrarily large. However, there are various interesting classes of mechanisms for which $T(b)$ is uniformly bounded, and then Proposition 11.1 yields uniform bounds.

For a class C of k-dimensional mechanisms (for instance: all mechanisms, or all deterministic mechanisms), we denote by $\text{C-REV}(\mathcal{F})$ the maximal revenue obtainable over all mechanisms in C; i.e., $\text{C-REV}(\mathcal{F}) := \sup_{b\in\text{C}} R(b;\mathcal{F})$.

Corollary 11.2 *Let* C *be a class of finite k-dimensional mechanisms, and define*

$$T(\text{C}) := \sup_{b\in\text{C}} T(b).$$

Then for every k-item distribution $\mathcal{F}$ on $\Re_+^k$

$$\text{C-REV}(\mathcal{F}) \le T(\text{C})\cdot\text{BREV}(\mathcal{F}); \tag{4}$$

moreover, the bound $T(\text{C})$ is tight.

Proof. By (2) and the definition of $T(\text{C})$, we have $R(b;\mathcal{F})/\text{BREV}(\mathcal{F}) \le T(b) \le T(\text{C})$ for every mechanism $b \in \text{C}$; taking sup over $b \in \text{C}$ gives (4).

To see that the bound $T(\text{C})$ is tight, take any $T' < T(\text{C})$, then there exists $b \in \text{C}$ with $T(b) > T'$. Let $\mathcal{F}_b$ be a distribution satisfying (3), then $\text{C-REV}(\mathcal{F}_b)/\text{BREV}(\mathcal{F}_b) \ge R(b;\mathcal{F}_b)/\text{BREV}(\mathcal{F}_b) = T(b) > T'$; thus no $T' < T(\text{C})$ can serve as an upper bound in (4). ∎

Proposition 11.3 *Let $[m]$-*FIN *be the class of all k-dimensional mechanisms with menu of size at most m, then*

$$T([m]\text{-FIN}) = \Omega(m^{1/7}).$$

Proof. Let $q^1, q^2, \ldots, q^m \in [0,1]^k$ be given by Proposition 6.1 so $\text{gap}^n = \Omega(n^{-6/7})$. Let b be given by the menu of size m consisting of (q^n, s^n) for $1 \le n \le m$, where $s^n := 2^n$.

For each $n \ge 1$ let $\mu^n := s^n/\text{gap}^n$ and $x^n := \mu^n q^n$. We have

$$q^n\cdot x^n - s^n = \mu^n(q^n\cdot q^n - \text{gap}^n) = \mu^n \max_{j<n} q^j\cdot q^n = \max_{j<n} q^j\cdot x^n \ge \max_{j<n}(q^j\cdot x^n - s^j),$$

and so a buyer with valuation x^n will choose only among those menu items (q^j, s^j) with[17] $j \geq n$, and all these s^j are $\geq s^n$; thus $s(x^n) \geq s^n$.

Therefore $t^n = \inf\{x \cdot e : s(x) \geq s^n\} \leq x^n \cdot e = \mu^n \; q^n \cdot e \leq \mu^n k = k2^n/\text{gap}^n$, and we get

$$T(b) = \sum_{n=1}^{m} \frac{s^n - s^{n-1}}{t^n} \geq \sum_{n=1}^{m} \frac{2^n - 2^{n-1}}{k2^n/\text{gap}^n} = \frac{1}{2k} \sum_{n=1}^{m} \text{gap}^n = \Omega(m^{1/7}).$$

∎

Remark. One can easily modify the construction of the mechanism above as to obtain a *symmetric* mechanism b with $T(b) = \Omega(m^{1/7})$. Indeed, the menu will include each (q^n, s^n) together with all its permutations $(\pi q^n, s^n)$, which increases its size by a factor of at most 2^k and does not affect $T(b)$ (since all permutations of q^n come with the same s^n).

Corollary 11.4

$$\sup_{\mathcal{F}} \frac{[m]\text{-REV}(\mathcal{F})}{\text{BREV}(\mathcal{F})} \geq \Omega\left(m^{1/7}\right) \qquad \text{and} \qquad \sup_{\mathcal{F}} \frac{\text{REV}(\mathcal{F})}{\text{BREV}(\mathcal{F})} = \infty,$$

where the supremum is over all k-item distributions $\mathcal{F}$ on $\Re_+^k$, or over all symmetric such distributions.

We will now provide precise bounds on how good the bundling auction is relative to deterministic auctions and to separate-items auctions when there are two items.

Proposition 11.5 (i) *For every two-item distribution $\mathcal{F}$ on $\Re_+^2$,*

$$\text{BREV}(\mathcal{F}) \geq \frac{2}{5} \; \text{DREV}(\mathcal{F}) \qquad \text{and} \qquad \text{BREV}(\mathcal{F}) \geq \frac{1}{2} \; \text{SREV}(\mathcal{F}).$$

(ii) *For every two-item symmetric distribution $\mathcal{F}$ on $\Re_+^2$,*

$$\text{BREV}(\mathcal{F}) \geq \frac{2}{3} \; \text{DREV}(\mathcal{F}) \qquad \text{and} \qquad \text{BREV}(\mathcal{F}) \geq \frac{2}{3} \; \text{SREV}(\mathcal{F}).$$

Moreover, all four bounds are tight.

Proof. We will compute $T(\text{C})$ for the four appropriate classes of mechanisms: (2) and (4) give (i), and (1) and (3) give (ii) (when $\mathcal{F}$ is symmetric we can limit ourselves to symmetric mechanisms).

[17] When indifferent, the buyer chooses the menu item with highest payment to the seller s; this can be assumed without loss of generality, see Hart and Reny [2012].

(1) A *symmetric deterministic* mechanism is of the form $b(x_1, x_2) = \max\{0, x_1 - p_1, x_2 - p_1, x_1 + x_2 - p_2\}$ with $p_i \geq 0$.

If $p_1 \leq p_2$ then $s_1 = p_1$ and $s_2 = p_2$, and we get $t_1 = p_1$ and $t_2 \leq 2(p_2 - p_1)$ (since $x_1 + x_2 - p_2 \geq x_1 - p_1$ implies $x_2 \geq p_2 - p_1$, and $x_1 + x_2 - p_2 \geq x_2 - p_1$ implies $x_1 \geq p_2 - p_1$, and so $x_1 + x_2 \geq 2(p_2 - p_1)$); therefore $T(b) \leq (p_1 - 0)/p_1 + (p_2 - p_1)/(2(p_2 - p_1)) \leq 3/2$ (when $p_1 = 0$ the first term is 0 and then $T(b) \leq 1/2$).

If $p_1 > p_2$ then $b(x_1, x_2) = \max\{0, x_1 + x_2 - p_2\}$ (which is bundled) and then $T(b) = (p_2 - 0)/p_2 \leq 1$.

Thus $T(b) \leq 3/2$. A mechanism b with $T(b) = 3/2$ is, for example, $b(x_1, x_2) = \max\{0, x_1 - 1, x_2 - 1, x_1 + x_2 - 3\}$.

(2) A *deterministic* mechanism is of the form $b(x_1, x_2) = \max\{0, x_1 - p_1, x_2 - p_2, x_1 + x_2 - p_3\}$; without loss of generality assume that $p_1 \leq p_2$.

If $0 \leq p_1 \leq p_2 \leq p_3$ then $s_i = p_i$ and we have: $t_1 \geq p_1$ (since $x_1 - p_1 \geq 0$ implies $x_1 + x_2 \geq p_1$); next, $t_2 \geq p_2 \geq p_2 - p_1$ (since $x_2 - p_2 \geq 0$ implies $x_1 + x_2 \geq p_2$); and finally $t_3 \geq 2(p_3 - p_2)$ (since $x_1 + x_2 - p_3 \geq x_1 - p_1$ implies $x_2 \geq p_3 - p_1 \geq p_3 - p_2$, and $x_1 + x_2 - p_3 \geq x_2 - p_2$ implies $x_1 \geq p_3 - p_2$, and thus $x_1 + x_2 \geq 2(p_3 - p_2)$). Therefore $T(b) \leq 1 + 1 + 1/2 = 5/2$.

If $0 \leq p_1 \leq p_3 \leq p_2$ then $b(x_1, x_2) = \max\{x_1 - p_1, x_1 + x_2 - p_3\}$ and then $s_1 = p_1, s_2 = p_3, t_1 = p_1$, and $t_2 \geq p_3 - p_1$ (since $x_1 + x_2 - p_3 \geq x_1 - p_1$ implies $x_2 \geq p_3 - p_1$), thus $T(b) \leq 1 + 1 = 2$.

If $0 \leq p_3 \leq p_1 \leq p_2$ then $b(x_1, x_2) = \max\{0, x_1 + x_2 - p_3\}$ and $T(b) = 1$.

Thus $T(b) \leq 5/2$. The computations above show that the bound $5/2$ cannot be attained; to approach it, for large M take $b(x_1, x_2) = \max\{0, x_1 - 1, x_2 - M, x_1 + x_2 - M^2\}$, then $T(b) > 5/2 - 2/M$.

(3) A *symmetric separate-price* mechanism is of the form $b(x_1, x_2) = \max\{0, x_1 - p, x_2 - p, x_1 + x_2 - 2p\}$.We have $s_1 = p_1 = t_1$ (attained at $(p, 0)$ and $(0, p)$) and $s_2 = 2p = t_2$ (attained at (p, p); therefore $T(b) = p/p + (2p - p)/2p = 3/2$.

(4) A *separate-price* mechanism is of the form $b(x_1, x_2) = \max\{0, x_1 - p_1, x_2 - p_2, x_1 + x_2 - p_1 - p_2\}$; without loss of generality assume that $p_1 \leq p_2$. As above, we get: $s_1 = p_1 = t_1$ (attained at $(p_1, 0)$); $s_2 = p_2 = t_2$ (attained at $(0, p_2)$); and $s_3 = p_1 + p_2 = t_3$ (attained at (p_1, p_2)). Therefore $T(b) = p_1/p_1 + (p_2 - p_1)/p_2 + p_1/(p_1 + p_2) = 2 - p_1^2/((p_2(p_1 + p_2)) < 2$, and the sup is 2 (take $p_2/p_1 \to \infty$). ∎

Examples of distributions for which these bounds are tight are easily obtained from the above proof together with the construction in Proposition 11.1 (ii) (we have slightly simplified them):

- For DREV, take the distribution $\mathcal{F}_M$ that puts probabilities $c, c/M$, and $c/(2M^2)$,

on the points $(1,0)$, $(0,M)$, and (M^2,M^2), respectively, where $c = (1+1/M+1/(2M^2))^{-1}$ (note that $c \to 1$ as $M \to \infty$). Then $\text{BREV}(\mathcal{F}_M) = \max\{1\cdot 1,\ M\cdot(c/M + c/(2M^2)),\ 2M^2\cdot c/2M^2\} \to 1$, and the deterministic mechanism $b(x_1,x_2) = \max\{0, x_1-1, x_2-M, x_1+x_2-M^2\}$ yields $R(b;\mathcal{F}_M) = 1\cdot c + M\cdot c/M + M^2\cdot c/(2M^2) = 5/2c \to 5/2$.

- For SREV, take X_1 and X_2 to be independent, $\mathbb{P}[X_1=0] = \mathbb{P}[X_1=2] = 1/2$, $\mathbb{P}[X_2=0] = 1-1/M$, $\mathbb{P}[X_2=M] = 1/M$ (where $M>2$), then $\text{REV}(X_1) = 2\cdot 1/2 = 1$, $\text{REV}(X_2) = M\cdot 1/M = 1$, and so $\text{SREV}(X) = 2$; and $\text{BREV}(X) = \max\{0\cdot 1, 2\cdot(1/2+1/(2M)), M\cdot 1/M, (M+1)\cdot 1/(2M)\} = 1+1/M \to 1$.

- For the symmetric case, take the distribution $\mathcal{F}$ that puts probability $1/4$ on each of $(1,0)$ and $(0,1)$ and probability $1/2$ on $(1,1)$. We have $\text{BREV}(\mathcal{F}) = \max\{1\cdot 1, 2\cdot 1/2\} = 1$, and the mechanism $b(x_1,x_2) = \max\{0, x_1-1, x_2-1, x_1+x_2-2\}$ yields $\text{DREV}(\mathcal{F}) = \text{SREV}(\mathcal{F}) = R(b;\mathcal{F}) = 3/2$.

Appendix 4: Multiple Buyers

This paper has concentrated on a single-buyer scenario that may also be interpreted to be a monopolist pricing setting. One may naturally ask the same questions for more general auction settings involving multiple buyers. An immediate observation is that since our main results (Theorems A, B, and C) are separations, they apply directly also to multiple-buyer settings, simply by considering a single "significant" buyer together with multiple "negligible" (in the extreme, 0-value for all items) buyers. The issue of extending the results to the multiple-buyer settings is thus relevant for the upper bounds in the paper, both the significant ones (Theorem D and Proposition 8.2) and the simple ones (Propositions 4.1, 4.2, and 4.3). In this appendix we discuss why these can all be extended to the multi-buyer scenario, at least if we are willing to incur a *loss that is linear in the number of buyers*. It is not completely clear where and how this loss may be avoided.

In the case of multiple buyers, we must first choose our notion of implementation: Bayes-Nash or Dominant-strategy. Also, we need to specify whether we are assuming independence between buyers' valuations or allow them to be correlated. The discussion here will be coarse enough as to apply to all these variants at the same time, with differences noted explicitly.

The next issue is how should we define the menu size for the case of multiple buyers. In the single-buyer case we defined it to be simply the size of the outcome space: $|\{q(x)|x \in \Re_+^k\}\backslash\{(0,\ldots,0)\}|$ and interpreted it as the number of options from which the buyer may

choose. In the case of multiple buyers, these are two separate notions. For example consider deterministic auctions of k items among n buyers. There are a total of $(n+1)^k$ different outcomes (each item may go to any buyer or to no one), but each buyer only considers 2^k possibilities (whether *he* gets each item or not). Moreover, the total set of outcomes cannot be interpreted as a menu from which the buyers may choose, since each buyer can only choose among the possibilities that he sees in front of him (and these choices need not be overall feasible). It takes the combined actions of all the buyers together in order for the auction outcome to be determined. For this reason we prefer to define the menu size of a multi-buyer auction by considering its menu size from the point of view of the different buyers. Since the menu that a buyer sees is a function of the bids of the others, we will take the maximum.

- An n-buyer auction has *menu size* at most m if for every buyer $j = 1, \ldots, n$ and every (direct) bids of the other buyers[18] $x^{-j} \in (\Re_+^k)^{n-1}$, the number of non-zero choices that buyer j faces is at most m, i.e., $|\{q^j(x^j, x^{-j}) : x^j \in \Re_+^k\} \backslash \{(0, \ldots, 0)\}| \leq m$.

Note: If the original mechanism was incentive compatible in dominant strategies then the mechanism induced by x^{-j} on player j is also incentive compatible. If the original mechanism was incentive compatible in the Bayesian sense then this need not be the case, but we will still have individual rationality[19] of the induced mechanism which suffices for what comes next.

Let us first analyze the simplest auctions, those with a single non-trivial menu item for each buyer. Clearly bundling auctions satisfy this property, however not every auction that has a single non-trivial menu-item for each buyer can be converted to a bundling auction. Let us also be careful with the meaning of bundling auction. Clearly in the case of correlated buyer valuations, the optimal auction for selling even a single item (the whole bundle in our case) is not necessarily to sell it to the highest bidder, but rather use the bids of the others as to set the reserve price to each bidder. (Consider for example the case of two buyers with a common value, where the bid of one of them should be used as the asking price from the other.) Thus in the rest of the discussion below we used BRev to denote the best auction that always sells the bundle as a whole – not necessarily to the highest bidder or with a uniform reserve price. For the case of independent buyer values, the simpler version that sells it to the highest bidder with a fixed reserve price will do too.

[18] Superscripts are used here for the buyers.

[19] This assumes that the original mechanism was ex-post individually rational, which one may verify does not really loose generality relative to ex-ante individual rationality.

What can be easily observed is that by focusing solely on the buyer that yields the largest fraction of revenue we can reduce back to the single-buyer case and extract at least $1/n$ fraction of revenue by selling the bundle to that single buyer. A full bundling auction can only make more, so this gives us the analog to Proposition 4.1 for the case of n buyers:[20] $\text{BREV}^{[n]}(\mathcal{F}) \leq [1]\text{-REV}^{[n]}(\mathcal{F}) \leq n \cdot \text{BREV}^{[n]}(\mathcal{F})$. The loss of the factor of n can be seen to be justified when considering independent buyer values and the restricted definition of bundling auctions even when there is one good (i.e., $k = 1$) by considering the distribution where each buyer j gets value M^j with probability M^{-j}, and zero otherwise (all independently), for a large enough but fixed M.

A similar argument that focuses on the single buyer that provides the largest fraction of revenue provides the generalization of Proposition 4.2: $[m]\text{-REV}^{[n]}(\mathcal{F}) \leq nm \cdot \text{BREV}^{[n]}(\mathcal{F})$, and of Proposition 4.3: $\text{DREV}^{[n]}(\mathcal{F}) \leq n \cdot (2^k - 1) \cdot \text{BREV}^{[n]}(\mathcal{F})$. It turns out that the linear loss in n is required here too, again for independent buyer values and the restricted interpretation of bundling auctions: take the construction of Theorem C on each of the n different buyers and combine it with the argument above. That is, whenever the construction has a valuation x with probability p, let buyer j have valuation $M^j x$ with probability $M^{-j} p$ (independently over the buyers).

Versions of Theorem D and Proposition 8.2 that incur a linear loss in n are also easily implied, but do not seem to be interesting. It would seem that in both cases one might get sharper results that avoid the additional loss due to the number of buyers.

References

Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. Pricing randomized allocations. In *SODA*, pages 585–597, 2010.

Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic pricing via virtual valuations. In *ACM Conference on Electronic Commerce*, pages 243–251, 2007.

Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *STOC*, pages 311–320, 2010a.

Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. In *Proceedings of the 11th ACM conference on Electronic commerce*, 2010b.

[20] The superscript $[n]$ on the various revenues denotes the number of buyers.

Constantinos Daskalakis and S. Matthew Weinberg. On optimal multi-dimensional mechanism design. *Electronic Colloquium on Computational Complexity (ECCC)*, 18:170, 2011.

Hanming Fang and Peter Norman. To bundle or not to bundle. *RAND Journal of Economics*, 37(4):946–963, 2006.

Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. In *ACM Conference on Electronic Commerce*, 2012.

Sergiu Hart and Philip J. Reny. Revenue maximization in two dimensions, manuscript, 2010.

Sergiu Hart and Philip J. Reny. Maximizing revenue with multiple goods: Nonmonotonicity and other observations, manuscript, 2012.

Philippe Jehiel, Moritz Meyer ter Vehn, and Benny Moldovanu. Mixed bundling auctions. *Journal of Economic Theory*, 134(1):494 – 512, 2007.

Omer Lev. A two-dimensional problem of revenue maximization. *Journal of Mathematical Economics*, 47:718 – 727, 2011.

Alejandro M. Manelli and Daniel R. Vincent. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory*, 127(1):1 – 35, 2006. ISSN 0022-0531. doi: 10.1016/j.jet.2005.08.007.

R. Preston McAfee and John McMillan. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory*, 46(2):335 – 354, 1988. ISSN 0022-0531. doi: 10.1016/0022-0531(88)90135-4.

R.oger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1): 58–73, 1981.

Gregory Pavlov. Optimal mechanism for selling two goods. *B.E. Journal of Theoretical Economics: Advances*, 11(1):Article 3, 2011.

John Thanassoulis. Haggling over substitutes. *Journal of Economic Theory*, 117(2):217 – 245, 2004. ISSN 0022-0531. doi: 10.1016/j.jet.2003.09.002.