The Menu-Size Complexity of Auctions*

Sergiu Hart† Noam Nisan‡

November 6, 2017

Abstract

We consider the *menu size* of auctions and mechanisms in general as a measure of their complexity, and study how it affects revenue. Our setting has a single revenue-maximizing seller selling two or more heterogeneous goods to a single buyer whose private values for the goods are drawn from a (possibly correlated) known distribution, and whose valuation is additive over the goods. We show that the revenue may increase arbitrarily with menu size and that a bounded menu size *cannot* ensure any positive fraction of the optimal revenue. The menu size turns out also to "pin down" the revenue properties of deterministic mechanisms: their menu size is at most exponential in the number of goods, and indeed their revenue may be larger than that achievable by the simplest types of mechanisms by a similar factor.

*Previous versions: August 2012; April 2013 (The Hebrew University of Jerusalem, Center for Rationality DP-637 and http://arxiv.org/abs/1304.6116). Presented at the 2013 ACM Conference on Electronic Commerce. Research partially supported by a European Research Council Advanced Investigator grant (Hart), by an Israel Science Foundation grant (Nisan), and by a Google grant (Nisan). We thank Motty Perry and Phil Reny for introducing us to the subject and for many useful discussions. A presentation that covers some of this work is available at `http://www.ma.huji.ac.il/hart/abs/2good-p.html`

†The Hebrew University of Jerusalem (Federmann Center for the Study of Rationality, Department of Economics, and Institute of Mathematics). *E-mail*: `hart@huji.ac.il` *Web site*: `http://www.ma.huji.ac.il/hart`

‡The Hebrew University of Jerusalem (Federmann Center for the Study of Rationality, and School of Computer Science and Engineering), and Microsoft Research. *E-mail*: `noam@cs.huji.ac.il` *Web site*: `http://www.cs.huji.ac.il/~noam`

Our model is related to a previously studied "unit-demand" model and our results also answer an open problem in that model.

Contents

1 Introduction

Are complex auctions better than simple ones? Myerson's classic result (Myerson 1981; see also Riley and Samuelson 1981 and Riley and Zeckhauser 1983) shows that if one is aiming to maximize revenue when selling a single good, then the answer is "no." The optimal auction is very simple, allocating the good to the highest bidder (using either first or second price) as long as he bids above a single deterministically chosen reserve price.

However, when selling multiple goods the situation turns out to be more complex. There has been significant work both in economics and in computer science[1] showing that, for selling multiple goods, simple auctions are no longer optimal. Specifically, it is known that randomized auctions may yield more revenue than deterministic ones, and that bundling the goods 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).

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

In our previous paper, Hart and Nisan (2017),[3] we considered the setup where the buyer's values for the different goods are *independent*, in which case we showed that simple auctions are *approximately* optimal: selling each good separately (deterministically) for its optimal price extracts a constant fraction of the optimal revenue.

In this paper we start by showing that the picture changes completely when the *valuations of the goods are correlated*, in which case "complex" auctions can become arbitrarily better than "simple" ones. For a start, let

[1] See Section 2 for the relevant literature.

[2] See Appendix A.4 for the extension of our results from the single-buyer to the multiple-buyer setting.

[3] Originally circulated in 2012.

us take "deterministic" as an "upper bound" of modeling "simple."

Consider $k \geq 2$ goods, whose valuation to the single buyer is given by a random variable $X = (X_1, X_2, ..., X_k)$ with values in $\mathbb{R}^k_+$; we emphasize that we allow for arbitrary dependence between the coordinates of X. We denote by $\text{Rev}(X)$ the optimal revenue achievable by any mechanism (or auction) when selling k goods to a single additive buyer with a valuation given by X, and denote by $\text{DRev}(X)$ the optimal revenue achievable by any *deterministic* mechanism. We have:

Theorem A *For every $k \geq 2$:*

(i) *For every $\varepsilon > 0$ there exists a k-good random valuation X with values in $[0, 1]^k$ such that*

$$\text{DRev}(X) < \varepsilon \cdot \text{Rev}(X).$$

(ii) *There exists a k-good random valuation X such that*

$$\text{DRev}(X) = 1 \quad \text{and} \quad \text{Rev}(X) = \infty.$$

What (ii) says is that if we allow the valuation to be unbounded, then we can get the fraction ε in (i) to go all the way down to 0. It is easy to see that (ii) implies (i) (just truncate X beyond a high enough value); the construction that yields (i) however is simpler and explicit. Part (i) is proved in Section 6 and part (ii) in Section 7.

In the independent case, Hart and Nisan (2017) showed that selling the goods separately—which is a simple deterministic mechanism—is guaranteed to yield at least[4] $1/2$ of the optimal revenue for $k = 2$ goods, and at least a fraction of the order of[5] $1/\log^2 k$ of the optimal revenue for $k > 2$ goods. By contrast, in the general correlated case, Theorem A says that neither separate selling, nor bundling, nor any deterministic mechanism can guarantee any positive fraction of the optimal revenue, for two or more goods. In terms of the "Guaranteed Fraction of Optimal Revenue" (GFOR) of Hart and Nisan

[4]Improved to 62% by Hart and Reny (2016).

[5]Improved to $c/\log k$ by Li and Yao (2013).

(2017),[6] we can state this as

$$\begin{aligned}\text{GFOR}(\textsc{deterministic};\ k\ \textit{goods}) &= \text{GFOR}(\textsc{separate};\ k\ \textit{goods})\\ &= \text{GFOR}(\textsc{bundled};\ k\ \textit{goods})\\ &= 0 \qquad (1)\end{aligned}$$

for any number $k \geq 2$ of goods.

As mentioned, the gap between the revenues of deterministic and general mechanisms has received significant attention and, in particular, in a model where the buyer wants to get *only one* of the k goods (the so-called "unit demand" setup), Briest, Chawla, Kleinberg, and Weinberg (2015) proved that there is an infinite gap between these revenues when there are three or more goods.[7] The case of two goods was left open, with some partial results indicating that for this case the gap may be bounded. While the unit demand model and our model are different, we show in Appendix A.3 that they are closely related: the various revenues in the two models are within constant factors of one another (these constants being at most exponential in the number of goods). On the one hand, this implies that Theorem A for $k \geq 3$ goods follows from the result of Briest *et al.* (2015). On the other hand, Theorem A for $k = 2$ goods solves their open problem: there is an infinite gap between the deterministic revenue and the optimal revenue in the unit-demand model, already for two goods.

The proofs of Briest *et al.* (2015) and of our Theorem A construct a random valuation X together with a set of possible outcomes (an outcome specifies the probabilities of getting each one of the k goods together with the payment that the buyer pays to the seller), such that for every outcome in the set there is a valuation that chooses that outcome (i.e., a buyer type for which this is the best outcome). The trick is to have a *large* number of

[6]This is the reciprocal of the so-called "competitive ratio" used in the computer science literature. While the two notions are clearly equivalent, using the optimal revenue as the benchmark (i.e., 100%) and measuring everything relative to this basis—as GFOR does—seems to come more naturally.

[7]In the case of independent goods, Chawla *et al.* (2007) show a constant upper bound on the gap between deterministic and randomized mechanisms in the unit-demand model.

possible outcomes, and to have a *significant* payment in each, so that the total revenue extracted is high—while, at the same time, not letting any simple mechanism do the same. As Briest *et al.* (2015) showed, this is significantly more difficult to do for two goods than for larger numbers of goods.

Looking at these proofs one observes that having a large set of possible outcomes—a large "menu" from which the buyer chooses, according to his valuation (or type)—seems to be the crucial attribute of the high-revenue mechanisms: it enables the sophisticated screening between different buyer types that is required for high-revenue extraction. Thus, we focus on ***menu size*** as a ***complexity measure of mechanisms and auctions*** and study the revenue extraction capabilities of mechanisms that are limited in their menu size.

Formally, we define the *menu size* of a mechanism to be the number of possible outcomes of the mechanism, where an outcome (or "menu entry") is, as stated above, an element $(q_1, q_2, \ldots, q_k; s)$ in $[0,1]^k \times \mathbb{R}$ in which q_i is the probability of allocating good i and s is the payment (our menu-size measure does not count the "null" outcome $(0, 0, \ldots, 0; 0)$ of getting nothing and paying nothing, which is always available, as it corresponds to the participation or individual rationality constraint). It is well known (by the so-called "taxation principle") that in our setting any mechanism can be put into the normal form of offering a fixed menu and letting the buyer choose among these menu entries. Notice that while deterministic mechanisms for k goods can have a menu size of at most $2^k - 1$ (since each q_i must be 0 or 1), randomized mechanisms can have an arbitrarily large, even infinite, menu size.

For a k-good random valuation X, we will use $\textsc{Rev}_{[m]}(X)$ to denote the optimal revenue achievable by any mechanism whose menu size is at most m, when selling k goods to a single additive buyer whose values for the goods are given by X. For a single good, $k = 1$, the characterization of optimal mechanisms of Myerson (1981) implies that $\textsc{Rev}_{[1]}(X)$ is already the same as the optimal revenue $\textsc{Rev}(X)$, but this is no longer true for more than a single good: the revenue may strictly increase as we allow the menu size to increase.

We start by looking at the simplest mechanisms according to this complexity measure, those with a single menu entry. It is not difficult to verify that the optimal single-menu-entry mechanism is always a *bundling* auction that sells the whole bundle (i.e., $q_i = 1$ for all i) for its optimal bundle price. In the same way, no menu entry can extract more than the bundling revenue. We thus have

Theorem B *For every $k \geq 2$:*

(i) *For every k-good random valuation X and every $m \geq 1$,*

$$\textsc{Rev}_{[m]}(X) \leq m \cdot \textsc{Rev}_{[1]}(X) = m \cdot \textsc{BRev}(X). \tag{2}$$

(ii) *There exists a k-good random valuation X with $0 < \textsc{Rev}_{[1]}(X) < \infty$ and a constant $c > 0$ such that for all $m \geq 1$,*

$$\textsc{Rev}_{[m]}(X) \geq cm^{1/7} \cdot \textsc{Rev}_{[1]}(X) = cm^{1/7} \cdot \textsc{BRev}(X).$$

This can be restated as

$$cm^{1/7} \leq \sup_X \frac{\textsc{Rev}_{[m]}(X)}{\textsc{Rev}_{[1]}(X)} \leq m$$

where the supremum is over all k-good random valuations[8] X. As discussed above, (i), whose simple proof is in Section 4, says that the revenue may grow at most linearly in the menu size; as for (ii), which is obtained from our construction in the proof of Theorem A(ii) in Section 7, it says that the revenue may well grow at least polynomially in the menu size.[9] The fact that for every m the revenue $\textsc{Rev}_{[m]}(X)$ from mechanisms with menu size at most m is bounded from above by a constant (namely, m) multiple of $\textsc{BRev}(X)$,

[8]Throughout this paper, we consider ratios of revenues only for those random valuations for which these ratios are well defined (i.e., neither $0/0$ nor ∞/∞).

[9]The increase is at a polynomial rate in m, and we do not think that the constant of $1/7$ we obtain is tight. For larger values of k the construction in Briest *et al.* (2015) implies a somewhat better polynomial dependence on m. For m that is at most exponential in k, Theorem C below shows that the growth can be almost linear in m.

and thus of $\text{DR}\textsc{ev}(X)$ (which can be only higher), implies that the result of Theorem A holds also with $\textsc{Rev}_{[m]}(X)$ instead $\text{DR}\textsc{ev}(X)$. In terms of the Guaranteed Fraction of Optimal Revenue (cf. (1)), we thus have

$$\text{GFOR}(\textsc{menu size} \leq m\ ;\ \ k\ \textit{goods}) = 0$$

for every finite m and $k \geq 2$ goods. It thus does not matter exactly how "simple" mechanisms are defined; so long as their menu size is bounded (which is natural, as unbounded menu size can hardly be considered simple), the (perhaps surprising) result is:

> ***Simple mechanisms cannot guarantee any positive fraction of the optimal revenue.***

Returning to deterministic mechanisms, whose menu size is, as seen above, at most $2^k - 1$, we have the following.

Theorem C *For every $k \geq 2$:*

(i) *For every k-good random valuation X,*

$$\text{DR}\textsc{ev}(X) \leq (2^k - 1) \cdot \textsc{Rev}_{[1]}(X).$$

(ii) *There exists a k-good random valuation X with values in $[0,1]^k$ and $0 < \textsc{Rev}_{[1]}(X) < \infty$ such that*[10]

$$\text{DR}\textsc{ev}(X) \geq \left(\frac{2^k - 1}{k}\right) \cdot \textsc{Rev}_{[1]}(X).$$

The upper bound in (i) follows immediately from Theorem B(i); as for (ii), which is proved using the techniques of Theorem A(ii) in Section 7, it shows that this bound is essentially tight (the factor k is much smaller than $2^k - 1$). Note again the contrast to the independent case, for which the bound

[10] We obtain in fact a slightly better bound than $(2^k - 1)/k$; for large k, it is close to twice that. See Proposition 10 and the Remarks following it in Section 6.

is linear in k, rather than exponential[11]: Lemma 28 in Hart and Nisan (2017) shows that for k independent goods $\text{DR}\textsc{ev}(X) \le \textsc{Rev}(X) \le ck\cdot\textsc{Rev}_{[1]}(X)$ for some $c > 0$.

The result of Theorem C can be written as

$$\frac{2^k-1}{k} \le \sup_X \frac{\text{DR}\textsc{ev}(X)}{\textsc{Rev}_{[1]}(X)} \le \sup_X \frac{\textsc{Rev}_{[2^k-1]}(X)}{\textsc{Rev}_{[1]}(X)} \le 2^k - 1. \tag{3}$$

The deterministic revenue is thus essentially the same as the revenue obtainable from mechanisms with a menu of size $2^k - 1$. This suggests that the reason that deterministic mechanisms yield low revenue (cf. Theorem A) is *not* that they are deterministic, but rather that being deterministic limits their menu size (to $2^k - 1$); any mechanism with that menu size will do just as badly.

Next, we compare the maximal revenue obtainable by selling each good separately (at its one-good optimal price)—we denote this separate revenue by SRev—to the single-menu-entry (i.e., bundling) revenue. We have

Theorem D *There exists a constant $c > 0$ such that for every $k \ge 2$ and every k-good random valuation X,*

$$\frac{c}{\log k} \cdot \textsc{Rev}_{[1]}(X) \le \text{SR}\textsc{ev}(X) \le k \cdot \textsc{Rev}_{[1]}(X).$$

This theorem is proved in Section 8. Unlike our previous results, the bounds here are the same as those we obtained for independently distributed goods, and they are tight already in that case; see Propositions 12(iv), 13(ii), and 14(i), and Example 27 in Hart and Nisan (2017).

Now the mechanism that sells the k goods separately has menu size $2^k - 1$ (since the buyer may acquire any subset of the goods, and so there are $2^k - 1$ possible outcomes), but its revenue is only k times, rather than $2^k - 1$ times, the bundling revenue. Moreover, selling separately seems intuitively to be

[11]Proposition 26 in Appendix A.2 shows that the same exponential-in-k gap exists between deterministic mechanisms and separate selling: there is X such that $\text{DR}\textsc{ev}(X) \ge (2^k - 1)/k\cdot\text{SR}\textsc{ev}(X)$. This provides a rare *doubly exponential* contrast with the independent case in which $\text{DR}\textsc{ev}(X) \le c\log^2 k\cdot\text{SR}\textsc{ev}(X)$ for some constant c (by Theorem C in Hart and Nisan 2017).

much simpler than this exponential-in-k complexity measure suggests: one just needs to determine the k prices. All this leads us to define a stronger notion of mechanism complexity, one that assigns to separate selling its more natural complexity k. This measure allows "additive menus" in which the buyer may choose not just a single menu entry but any set of menu entries (so long as their combination is feasible, namely, all goods are allocated with probabilities at most 1). We present this *additive menu size* complexity measure in Section 8, and show that all our results apply to this stronger complexity measure as well.

Finally, looking at the constructions used in the proofs of Theorems A and B, one sees that the range of valuations (i.e., the support of X) is exponential in the gap obtained; more precisely, if we restrict the values of each good to being in a range that is bounded (from above as well as from below, i.e., away from[12] 0), say, in the range $[1, H]$, then the gap becomes bounded by some constant power of $\log H$; see Section 9, where we show that this exponential blowup in the range is indeed needed. If the range is bounded, however, arbitrarily good approximations of the optimal revenue can be obtained by a menu size that is only poly-logarithmic in the range size.

Theorem E *Let $k = 2$. There exists a constant $c < \infty$ such that for every $1 < H < \infty$ and $\varepsilon > 0$,*

$$\textsc{Rev}_{[m]}(X) \geq (1 - \varepsilon) \cdot \textsc{Rev}(X)$$

for every two-good random valuation X with values in $[1, H]^2$ and menu size

$$m \geq \frac{c}{\varepsilon^5} \log^2 \left(\frac{H}{\varepsilon} \right).$$

This theorem is proved in Section 9. Again, contrast this result with the unbounded range case: when the upper bound on valuations is infinite there is X with $\textsc{Rev}(X) = \infty$ while $\textsc{Rev}_{[m]}(X) \leq m$ for every finite m (by Theorems A(ii) and B(i)), and when the lower bound on valuations is zero,

[12]Both bounds are needed, as rescaling X rescales all revenues and so does not affect the ratios between revenues.

for every m there is X with $\textsc{Rev}_{[m]}(X)/\textsc{Rev}(X) < m\varepsilon$ (by Theorems A(i) and B(i)).

For a larger number of goods $k > 2$ we are only able to prove upper bounds on the menu size of the form $(H/\varepsilon)^{ck}$ for some constant[13] c, which are essentially equivalent to known bounds (cf. Daskalakis and Weinberg 2012).

1.1 Organization of the Paper

In Section 2 immediately below we briefly go over the related literature. Section 3 defines formally our model and notations. We then start in Section 4 with the basic properties of the menu-size complexity measure. Next, in Section 5 we provide a sharp comparison tool between the revenue of an arbitrary mechanism and the bundling revenue (Theorem 6). This is then used in Section 6 to construct the two-good valuation that proves Theorem A(i), and in Section 7 to provide a general construction that proves Theorems A(ii) and B(ii), as well as Theorem C(ii). Section 8 studies separate selling, proves Theorem D, and introduces the more refined "additive menu size" complexity measure. We conclude in Section 9 with positive results on approximation for the case where the valuations are in a bounded domain, thus proving Theorem E for two goods and its weaker variants for more goods. Additional results are provided in the appendices. Appendix A.1 gives the precise comparison between deterministic and bundling revenues for two goods, and Appendix A.2 carries out the comparison of the revenue of an arbitrary mechanism to the separate revenue. In Appendix A.3 we prove the close relationships between our setup and the unit-demand setup. Finally, in Appendix A.4 we discuss the multiple-buyer case.

2 Literature

We briefly describe some of the existing work on these issues.

[13]Later improved by Dughmi, Han, and Nisan (2014); see footnote 33 in Section 9.

The realization that maximizing revenue with multiple goods is complex has a long history in economic theory and more recently in the computer science literature as well. McAfee and McMillan (1988) identified cases where the optimal mechanism is deterministic. However, Thanassoulis (2004) and Manelli and Vincent (2006) found a technical error in the paper and presented counterexamples.[14] These papers contain good surveys of the work within economic theory, with more recent studies by Fang and Norman (2006), Pycia (2006), Manelli and Vincent (2007, 2012), Jehiel, Meyer-ter-Vehn, and Moldovanu (2007), Lev (2011), Pavlov (2011), and Hart and Reny (2015). In the past few years algorithmic work on these types of topics was carried out. One line of work (e.g., Briest, Chawla, Kleinberg, and Weinberg 2015; Cai, Daskalakis, and Weinberg 2012a; Alaei, Fu, Haghpanah, Hartline, and Malekian 2012) showed that for discrete distributions the optimal mechanism can be found by linear programming in rather general settings. Another line of work in computer science directly related to this work (Chawla, Hartline, and Kleinberg 2007; Chawla, Hartline, Malec, and Sivan 2010; Chawla, Malec, and Sivan 2010; Alaei, Fu, Haghpanah, Hartline, and Malekian 2012; Cai, Daskalakis, and Weinberg 2012b) attempted to *approximate* the optimal revenue by simple mechanisms in various settings, where simplicity was defined qualitatively. In this line of research, Hart and Nisan (2017) considered mechanisms that sell the goods either separately or as a single bundle to be simple mechanisms, and showed that when the values of the goods are independently distributed then a non-trivial fraction of the optimal revenue can be ensured by simple mechanisms. On the other hand, Briest *et al.* (2015) considered deterministic mechanisms to be simple, and, in the unit-demand setting with at least 3 goods,[15] proved that deterministic mechanisms cannot ensure any positive fraction of the revenue of general mechanisms.

Since the circulation in 2012 and 2013 of early versions of our papers—Hart and Nisan (2017) and the present paper—there has been a flurry of work on optimal mechanisms for multiple goods in various settings: Daskalakis,

[14]See Hart and Reny (2015) for a simple and transparent such example, together with a discussion of why this phenomenon can occur only when there are two or more goods.

[15]The unit-demand setting and its relation to our additive setting are discussed in Appendix A.3.

Deckelbaum, and Tzamos (2013, 2017), Giannakopoulos (2014), Giannakopoulos and Koutsoupias (2014), Menicucci, Hurkens, and Jeon (2015), and Tang and Wang (2017). In addition, various improved approximation results have been obtained for wider classes of setups with *independently distributed* goods values: Li and Yao (2013), Babaioff, Immorlica, Lucier, and Weinberg (2014), Yao (2014), and Rubinstein and Weinberg (2015).

The menu size was further studied in Dughmi, Han, and Nisan (2014) which showed that exponential menu size is required in order to get any approximation ratio that is better than logarithmic in the range of valuations, by Babaioff, Gonczarowski, and Nisan (2017) who obtained lower and upper bounds on the menu size for independently distributed valuations, and by Gonczarowski (2017) who proved a polynomial (in the approximation rate) lower bound on the menu size even for the case of two goods with independent and identically distributed valuations. Further measures of the complexity of mechanisms were also studied, in particular related to the difficulty of learning (see Morgenstern and Roughgarden 2016 and the references there), and to Kolmogorov complexity (Dughmi, Han, and Nisan 2014).

3 Preliminaries

3.1 The Model

The basic model is standard, and the notations follow our previous paper Hart and Nisan (2017), which the reader may consult for further details (see also Hart and Reny 2015).

One seller (or "monopolist") is selling a number $k \geq 1$ of goods (or "items," "objects," etc.) to one buyer.

The goods have no value or cost to the seller. Let $x_1, x_2, ..., x_k \geq 0$ be the buyer's values for the goods. The value for getting a set of goods is *additive*: getting the subset $I \subseteq \{1, 2, ..., k\}$ of goods is worth $\sum_{i \in I} x_i$ to the buyer (and so, in particular, the buyer's demand is *not* restricted to one good only). The valuation of the goods is given by a random variable $X = (X_1, X_2, ..., X_k)$ that takes values in $\mathbb{R}_+^k$ (we thus assume that valuations are always nonnegative); we will refer to X as a *k-good random valuation.* The realization $x = (x_1, x_2, ..., x_k) \in \mathbb{R}_+^k$ of X is known to the buyer, but not to the seller, who knows only the distribution F of X (which may be viewed as the seller's belief); we refer to a buyer with valaution x also as a buyer of *type* x. The buyer and the seller are assumed to be risk neutral and to have quasi-linear utilities.

The objective is to *maximize* the seller's (expected) *revenue.*

As was well established by the so-called "Revelation Principle" (starting with Myerson 1981; see for instance the book of Krishna 2010), we can restrict ourselves to "direct mechanisms" and "truthful equilibria." A direct *mechanism* μ consists of a pair of functions (q, s), where $q = (q_1, q_2, ..., q_k) : \mathbb{R}_+^k \to [0, 1]^k$ and $s : \mathbb{R}_+^k \to \mathbb{R}$, which prescribe the *allocation* of goods and the *payment*, respectively. Specifically, if the buyer reports a valuation vector $x \in \mathbb{R}_+^k$, then $q_i(x) \in [0, 1]$ is the probability that the buyer receives good[16] i (for $i = 1, 2, ..., k$), and $s(x)$ is the payment that the seller receives from the buyer; we refer to $(q(x), s(x))$ as an *outcome.* When the buyer reports his value x truthfully, his payoff is[17] $b(x) = \sum_{i=1}^k q_i(x)x_i - s(x) = q(x) \cdot x - s(x)$, and the seller's payoff is $s(x)$.

The mechanism $\mu = (q, s)$ satisfies *individual rationality* (**IR**) if $b(x) \geq 0$ for every $x \in \mathbb{R}_+^k$; it satisfies *incentive compatibility* (**IC**) if $b(x) \geq q(\tilde{x}) \cdot x - s(\tilde{x})$ for every alternative report $\tilde{x} \in \mathbb{R}_+^k$ of the buyer when his value is x, for every $x \in \mathbb{R}_+^k$; and it satisfies the *no positive transfer* (**NPT**) property if $s(x) \geq 0$ for every $x \in \mathbb{R}_+^k$.

The (expected) revenue of a mechanism $\mu = (q, s)$ from a buyer with random valuation X, which we denote by $R(\mu; X)$, is the expectation of the payment received by the seller; i.e., $R(\mu; X) = \mathbb{E}[s(X)]$. We now define:

- $\textsc{Rev}(X)$, the *optimal revenue*, is the maximal revenue that can be obtained: $\textsc{Rev}(X) = \sup_\mu R(\mu; X)$, where the supremum is taken over all IC and IR mechanisms μ.

[16]When the goods are infinitely divisible and the valuations are linear in quantities, q_i may be alternatively viewed as the *quantity* of good i that the buyer gets.

[17]The scalar product of two n-dimensional vectors $y = (y_1, ..., y_n)$ and $z = (z_1, ..., z_n)$ is $y \cdot z = \sum_{i=1}^n y_i z_i$.

As seen in Hart and Nisan (2017), when maximizing revenue we can limit ourselves without loss of generality to IR and IC mechanisms that are in addition NPT mechanisms (for which $s(0,0,...,0) = b(0,0,...,0) = 0$).

From now on we will assume that all mechanisms μ are given in direct form, i.e., $\mu = (q,s)$, and that they satisfy IR, IC, and NPT.

When there is only one good, i.e., when $k = 1$, Myerson's (1981) result is that

$$\textsc{Rev}(X) = \sup_{p\geq 0} p \cdot \mathbb{P}\left[X \geq p\right] = \sup_{p\geq 0} p \cdot \mathbb{P}\left[X > p\right] = \sup_{p\geq 0} p \cdot (1 - F(p)), \quad (4)$$

where F is the cumulative distribution function of X. Optimal mechanisms correspond to the seller "posting" a price p and the buyer buying the good for the price p whenever his value is at least p; in other words, the seller makes the buyer a "take-it-or-leave-it" offer to buy the good at price p.

Besides the maximal revenue $\textsc{Rev}(X)$, we are also interested in what can be obtained from certain classes of mechanisms.

- $\textsc{SRev}(X)$, the *separate revenue*, is the maximal revenue that can be obtained by selling each good separately. Thus

$$\textsc{SRev}(X) = \textsc{Rev}(X_1) + \textsc{Rev}(X_2) + ... + \textsc{Rev}(X_k).$$

- $\textsc{BRev}(X)$, the *bundling revenue*, is the maximal revenue that can be obtained by selling all goods together in one "bundle." Thus

$$\textsc{BRev}(X) = \textsc{Rev}(X_1 + X_2 + ... + X_k).$$

- $\textsc{DRev}(X)$, the *deterministic revenue*, is the maximal revenue that can be obtained by deterministic mechanisms; these are the mechanisms in which every good $i = 1, 2, ..., k$ is either fully allocated or not at all: $q_i(x) \in \{0, 1\}$ for all valuations $x \in \mathbb{R}^k_+$ (rather than $q_i(x) \in [0, 1]$).

While the separate and the bundling revenues are obtained by solving one-dimensional problems (using (4)), for each good in the former, and for the bundle in the latter, the deterministic revenue is a true multidimensional problem.

3.2 Menu and Menu Size

Given a k-good mechanism $\mu = (q, s)$, we define its *menu* as the range of its nonzero outcomes, i.e.,

$$\text{MENU}(\mu) := \{(q(x), s(x)) : x \in \mathbb{R}_+^k\} \backslash \{(0, 0, ..., 0), 0)\} \subset [0, 1]^k \times \mathbb{R}_+$$

(we thus ignore the zero outcome, $((0, 0, ..., 0), 0)$, which is always included without loss of generality[18]). We will refer to each outcome in the menu as a *menu entry*. Conversely, any set of outcomes $M \subset [0, 1]^k \times \mathbb{R}_+$ generates a mechanism $\mu = (q, s)$ with $(q(x), s(x)) \in \arg\max_{(g,t)}(g \cdot x - t)$ where (g, t) ranges over $M \cup \{((0, 0, ..., 0), 0)\}$ (the mechanism is well defined up to tie-breaking; see Hart and Reny 2015 for more details). The menu of μ is included in M (as some outcomes in M may never be chosen; it will be convenient at times to ignore this and refer to such a μ as a mechanism with menu M).

The *menu size* of a mechanism μ is defined as the cardinality of its menu, i.e., the number of elements of $\text{MENU}(\mu)$, which may well be infinite:

$$\text{MENU-SIZE}(\mu) := |\text{MENU}(\mu)|.$$

Since a menu cannot contain two entries (g, t) and (g, t') with the same allocation $g \in [0, 1]^k$ but with different payments t, t' (if, say, $t' > t$ then (g, t') will never be chosen, as (g, t) is better for any buyer type), the menu size is identical to the cardinality of the set of nonzero allocations; i.e.,

$$\text{MENU-SIZE}(\mu) = |\{q(x) : x \in \mathbb{R}_+^k \text{ and } q(x) \neq (0, 0, ..., 0)\}|.$$

The corresponding revenue is:

[18]We thus slightly depart from Hart and Reny (2015), where the menu includes the zero outcome as well. This yields simple relations (such as Theorem B(i)) between menu size and revenue, for which the zero outcome does not count.

- $\textsc{Rev}_{[m]}(X)$, the *menu size-m revenue*, is the maximal revenue that can be obtained by mechanisms whose menu size is at most m.

In Section 8 we will define the *additive menu size*, a useful refinement of menu size, with its corresponding revenue notion $\textsc{Rev}_{[m]*}(X)$.

3.3 Guaranteed Fraction of Optimal Revenue (GFOR)

Given a class $\mathbb{X}$ of random valuations (e.g., k goods, two independent goods, and so on), and a class $\mathcal{N}$ of mechanisms (e.g., separate mechanisms, deterministic mechanisms, and so on), the corresponding *Guaranteed Fraction of Optimal Revenue* (GFOR) is defined in Hart and Nisan (2017) as the maximal fraction α such that, for any random valuation X in $\mathbb{X}$, there are mechanisms in the class $\mathcal{N}$ that yield at least the fraction α of the optimal revenue; that is,

$$\text{GFOR}(\mathcal{N};\mathbb{X}) := \inf_{X\in\mathbb{X}} \frac{\mathcal{N}\text{-}\textsc{Rev}(X)}{\textsc{Rev}(X)},$$

where $\mathcal{N}\text{-}\textsc{Rev}(X) := \sup_{\mu\in\mathcal{N}} R(\mu;X)$ denotes the maximal revenue that can be obtained by any mechanism in the class $\mathcal{N}$.

4 Basic Results on Menu Size

We start with a few simple and immediate relations concerning menu-size complexity, which yield Theorems B(i) and C(i). Recall the notation $\textsc{Rev}_{[m]}(X)$ for the revenue achievable by mechanisms whose menu size is at most m.

First, the optimal mechanism with a menu of size 1 is bundled selling.

Proposition 1 *For every $k \geq 2$ and every k-good random valuation X,*

$$\textsc{Rev}_{[1]}(X) = \textsc{BRev}(X).$$

Proof. Let μ be any mechanism with a single menu entry, say (g,t). If the seller were to offer instead to sell the whole bundle at the same price t, the buyer would surely buy whenever he did so in μ, and the revenue could only increase. Thus $R(\mu; X) \leq$BREV(X). Conversely, BREV(X) is achieved by a single menu entry by Myerson's result (4). ■

Next, the revenue can increase at most linearly in the menu size.

Proposition 2 *For every $k \geq 2$, every k-good random valuation X, and every $m \geq 1$,*

$$\text{REV}_{[m]}(X) \leq m \cdot \text{BREV}(X).$$

Proof. Let (g, t) be a menu entry. If (g, t) is chosen when the valuation is x then $t \leq g \cdot x \leq \sum_{i=1}^{k} x_i$ (the first inequality by IR, the second because $g_i \in [0, 1]$ and $x_i \geq 0$), and so the probability that (g, t) is chosen, which we denote by α, is at most $\mathbb{P}\left[\sum_{i=1}^{k} X_i \geq t\right]$. Therefore $t \cdot \alpha$, the portion of the expected revenue that comes from (g, t), is at most the revenue obtained from setting the bundle price at t, that is, $t \cdot \alpha \leq$ BREV(X). Sum over the m menu entries. ■

For small menu size m this bound is tight, as the example below shows for m not exceeding the number of goods k. It remains essentially tight for m up to $2^k - 1$ by Theorem C; as for large m, Theorem B(ii) shows that REV$_{[m]}(X)$ can be as large as[19] $\Omega(m^{1/7})\cdot$BREV(X).

Example 3 Let $1 \leq m \leq k$. Take a large[20] $M > 0$, and consider the following random valuation X. For each $i = 1, ..., m$, with a probability α_i that is proportional to $1/M^{i-1}$, good i is valued at M^{i-1} and all the other goods are valued at 0; thus $\alpha_i = c/M^{i-1}$ where $c := 1/(1 + 1/M + ... + 1/M^{m-1})$. Bundling yields a revenue of 1 (because setting the bundle price at M^{i-1} yields a revenue of $c(1/M^{i-1} + ... + 1/M^{m-1})$, which is maximal at $i = 1$, and the revenue there is 1). Selling each good $i = 1, ..., m$ at the price M^{i-1} yields a revenue of c from each good; this is obtained at distinct valuations,

[19]It is convenient to use the standard O and Ω notations. For two expressions F and G that depend on certain variables, we write $F = O(G)$ if $\sup F/G < \infty$, and $F = \Omega(G)$ if $\inf F/G > 0$; i.e., there is a constant $0 < c < \infty$ such that $F \leq cG$, respectively $F \geq cG$, for any values of the variables in the relevant range.

[20]Theorem 6 in Section 5 below provides the right tool to easily generate such examples.

and so the mechanism consisting of these m menu entries yields a revenue of mc, which is close to m for large M.

An immediate consequence of Proposition 2 provides a comparison between deterministic and bundling revenues.

Corollary 4 *For every $k \geq 2$ and every k-good random valuation X,*

$$\text{DRev}(X) \leq (2^k - 1) \cdot \text{BRev}(X).$$

Proof. A deterministic mechanism has at most $2^k - 1$ menu entries, one for each non-empty set of goods. ■

For $k = 2$ goods the tight bound turns out to be 5/2; we show this in Proposition 23 in Appendix A.1 (using, again, Theorem 6). For large k the bound $2^k - 1$ is tight within a factor of the order of k, which is small relative to $2^k - 1$ (see Proposition 10 in Section 7). The comparison of the deterministic revenue to the separate revenue is provided in Corollary 16 in Section 8; see also Appendix A.2.

5 Comparing to the Bundling Revenue

Having seen in the previous section the basic role of the bundling revenue, we now provide a precise tool that measures how much better a mechanism can be relative to bundling. It will then be used in the next sections to construct random valuations together with corresponding mechanisms that yield revenues that are arbitrarily higher than the bundling revenue, and thus than any other simple revenue as well.

Let $\mu = (q, s)$ be a k-good mechanism. For every $t > 0$ we define $v(t)$ as the minimum, over all valuations whose payment in μ is at least t, of the value of the bundle of all goods. Formally,

$$v(t) := \inf\{||x||_1 : x \in \mathbb{R}^k_+ \text{ and } s(x) \geq t\},$$

where $||x||_1 = \sum_{i=1}^k |x_i|$ is the 1-norm in $\mathbb{R}^k$; for x in $\mathbb{R}^k_+$, this is $\sum_i x_i$, the value to the buyer of the bundle of all goods.[21] As usual, the infimum of an empty set is taken to be ∞, and so $v(t) = \infty$ when t is higher than any possible payment $s(x)$. It is immediate that the function v is nondecreasing and satisfies $v(t) \geq t$ for every $t > 0$ (because $\sum_i x_i \geq q(x) \cdot x \geq s(x)$ by IR). Finally, we define $\beta(\mu)$ by

$$\beta(\mu) := \int_0^\infty \frac{1}{v(t)} \, \mathrm{d}t. \tag{5}$$

The function $1/v(t)$ is nonnegative, nonincreasing, and vanishes beyond the maximal possible payment (i.e., for $t > \sup_x s(x)$). Its integral may well be zero or infinite: $0 \leq \beta(\mu) \leq \infty$ (with $\beta(\mu) = 0$ only when $v(t) = \infty$ for every $t > 0$, which is the case only for the null mechanism with $s(x) = 0$ for all x). When μ has a finite menu, say $\{(g_n, t_n)\}_{n=1}^m$, ordered so that the sequence t_n is nondecreasing, we have $v(t) = v(t_n)$ for every $t_{n-1} < t \leq t_n$ (some of these intervals may well be empty[22]), and so

$$\beta(\mu) = \sum_{n=1}^m \frac{t_n - t_{n-1}}{v(t_n)} \tag{6}$$

(computing the numbers $v(t_n)$ amounts to solving m linear programming problems).

It may be instructive to compute $\beta(\mu)$ in a few examples with $k = 2$ goods.

Example 5 Let μ be given by the menu[23] $\{x_1 - p_1, x_2 - 2, x_1 + x_2 - 4\}$, where we allow p_1 to vary.

When $p_1 = 1$ we have $(t_1, t_2, t_3) = (1, 2, 4)$ and $(v(t_1), v(t_2), v(t_3)) = (1, 2, 5)$ (attained, respectively, at the points $(1, 0)$, $(0, 2)$, and $(2, 3)$); there-

[21] We use the 1-norm notation (rather than the sum of coordinates) because the parallel analysis for the separate revenue turns out to be very similar, except for using the ∞-norm instead; see Appendix A.2.

[22] If $v(t_n) = v(t_{n+1})$ then we may eliminate t_n altogether from the sum: $(t_n - t_{n-1})/v(t_n) + (t_{n+1} - t_n)v(t_{n+1}) = (t_{n+1} - t_{n-1})/v(t_{n+1})$.

[23] We write a menu entry (g, t) here as $g \cdot x - t$ (the payoff of the buyer with valuation x is thus $b(x) = \max\{0, x_1 - p_1, x_2 - 2, x_1 + x_2 - 4\}$).

fore $\beta(\mu) = (1-0)/1 + (2-1)/2 + (4-2)/5 = 19/10$.

When $p_1 = 2$ we have $(t_1, t_2, t_3) = (2, 2, 4)$ and $(v(t_1), v(t_2), v(t_3)) = (2, 2, 4)$ (with $v(2)$ attained at $(2, 0)$ and also at $(0, 2)$, and $v(4)$ at $(2, 2)$); therefore $\beta(\mu) = (2-0)/2 + (2-2)/2 + (4-2)/4 = 3/2$.

When $p_1 = 5$ we have $(t_1, t_2, t_3) = (2, 4, 5)$ and $(v(t_1), v(t_2), v(t_3)) = (2, 4, \infty)$ (with the first two attained at $(0, 2)$ and $(2, 2)$, and $v(5)$ infinite since $x_1 - 5$ is never chosen by the buyer, as it is always strictly worse than $x_1 + x_2 - 4$); therefore $\beta(\mu) = (2-0)/2 + (4-2)/4 + (5-4)/\infty = 3/2$.

We state now the basic result showing that $\beta(\mu)$ provides the comparison between the revenue from μ and the bundling revenue.

Theorem 6 *For every $k \geq 2$ and every k-good mechanism $\mu = (q, s)$,*

$$\sup_X \frac{R(\mu; X)}{\text{BRev}(X)} = \beta(\mu),$$

where the supremum is taken over all k-good random valuations X with[24] $0 < \text{BRev}(X) < \infty$.

Proof. (i) First, we show that

$$\frac{R(\mu; X)}{\text{BRev}(X)} \leq \beta(\mu)$$

for every k-good random valuation X. Indeed,

$$\begin{aligned} R(\mu; X) &= \mathbb{E}[s(X)] = \int_0^\infty \mathbb{P}[s(X) \geq t] \, \mathrm{d}t \leq \int_0^\infty \mathbb{P}[||X||_1 \geq v(t)] \, \mathrm{d}t \\ &\leq \int_0^\infty \frac{\text{BRev}(X)}{v(t)} \, \mathrm{d}t = \beta(\mu) \cdot \text{BRev}(X), \end{aligned}$$

where we have used: $s(X) \geq 0$ by NPT; $s(X) \geq t$ implies $||X||_1 \geq v(t)$ by the definition of $v(t)$; and $u \cdot \mathbb{P}[||X||_1 \geq u] \leq \text{BRev}(X)$ for every $u > 0$.

[24]The remarks immediately following the proof show that one may restrict X to have finite support or to be bounded and take values in, say, $[0, 1]^k$.

(ii) Second, we show that for every $\beta' < \beta(\mu)$ (which means any arbitrarily large β' when $\beta(\mu)$ is infinite) there exists a k-good random valuation X such that $0 <\text{BRev}(X) < \infty$ and

$$\frac{R(\mu; X)}{\text{BRev}(X)} > \beta'. \tag{7}$$

Indeed, the function $1/v(t)$ is nonincreasing and nonnegative, and so there exist $0 = t_0 < t_1 < ... < t_M < t_{M+1} = \infty$ with $0 = v(t_0) < v(t_1) < v(t_2) < ... < v(t_M) < v(t_{M+1}) = \infty$ such that

$$\beta'' := \sum_{n=1}^{M} \frac{t_n - t_{n-1}}{v(t_n)} > \beta'.$$

Let $\varepsilon > 0$ be small enough so that $\beta'' > (1+\varepsilon)\beta'$ and $v(t_{n+1}) > (1+\varepsilon)v(t_n)$ for all $1 \leq n \leq M$. By the definition of v we can choose for every $1 \leq n \leq M$ a point[25] $x_n \in \mathbb{R}_+^k$ such that $s(x_n) \geq t_n$ and $v(t_n) \leq ||x_n||_1 < (1+\varepsilon)v(t_n)$; then

$$\sum_{n=1}^{M} \frac{t_n - t_{n-1}}{||x_n||_1} > \sum_{n=1}^{M} \frac{t_n - t_{n-1}}{v(t_n)(1+\varepsilon)} = \frac{\beta''}{1+\varepsilon} > \beta'. \tag{8}$$

Put $\xi_n := ||x_n||_1$; the sequence ξ_n is strictly increasing (because $(1+\varepsilon)v(t_n) < v(t_{n+1})$) and $\xi_1 > 0$ (because $v(t_1) > 0$). Let X be a random variable with support $\{x_1, ..., x_M\}$ and distribution $\mathbb{P}[X = x_n] = \xi_1/\xi_n - \xi_1/\xi_{n+1}$ for every $1 \leq n \leq M$, where we put $\xi_{M+1} := \infty$; thus $\mathbb{P}[X \in \{x_n, ..., x_M\}] = \xi_1/\xi_n$ for every $n \geq 1$.

To compute $\text{BRev}(X)$, we only need to consider the bundle prices ξ_n for $1 \leq n \leq M$ (these are the possible values of $\sum_i X_i = ||X||_1$), for which we have

$$\xi_n \cdot \mathbb{P}[||X||_1 \geq \xi_n] = \xi_n \cdot \mathbb{P}[X \in \{x_n, ..., x_M\}] = \xi_n \frac{\xi_1}{\xi_n} = \xi_1,$$

and so

$$\text{BRev}(X) = \xi_1. \tag{9}$$

[25]Subscripts n, m, and j are used for sequences, whereas i is used exclusively for coordinates; thus x_n is a vector in $\mathbb{R}_+^k$, and x_i is the i-th coordinate of x.

Finally, the revenue $R(\mu;X)$ that μ extracts from X is

$$\begin{aligned} R(\mu;X) &\geq \sum_{n=1}^{M} s(x_n)\mathbb{P}\left[X=x_n\right] \geq \sum_{n=1}^{M} t_n \left(\frac{\xi_1}{\xi_n} - \frac{\xi_1}{\xi_{n+1}}\right) \\ &= \sum_{n=1}^{M}(t_n - t_{n-1})\frac{\xi_1}{\xi_n} > \xi_1\beta' = \beta' \cdot \mathrm{BRev}(X) \end{aligned} \tag{10}$$

(use $\xi_{M+1} = \infty$,(8), and (9)). ■

Remarks. *(a)* For any $m < M$, let μ_m denote a mechanism obtained by restricting the menu of μ to the entries chosen by[26] $x_1, ..., x_m$. The computation of $R(\mu_m;X)$ is the same as in (10), but with the sum going only up to m instead of M, and thus with a final term of $t_m(\xi_1/\xi_{m+1})$ that needs to be subtracted; this gives

$$R(\mu_m;X) > \left(\sum_{n=1}^{m} \frac{t_n - t_{n-1}}{\xi_n} - \frac{t_m}{\xi_{m+1}}\right) \cdot \mathrm{BRev}(X) \tag{11}$$

(recall (9)), a result that will be used in Proposition 9 in Section 7.

(b) The random valuation X that we have constructed in part (ii) of the proof has finite support, and is thus bounded from above; one may thus rescale it (which does not affect the ratio Rev/BRev) so that it takes values in, say, $[0,1]^k$.

(c) If the mechanism μ has a finite menu of size m, then $v(t)$ can take at most m distinct values, and so $M \leq m$ and the support of the resulting X is of size at most m.

(d) If the mechanism μ has a finite menu of size m, then $\beta(\mu) \leq m$ (because $v(t) \geq t$ implies that each term in the sum (6) is ≤ 1). This implies the linear-in-menu-size bound of Theorem B(i) and Proposition 2 (Example 3 in the previous section is obtained by making each term close to 1). The function β may thus be viewed as a more refined complexity measure than the menu size: higher values of β go with more complex mechanisms (and

[26] Formally, $\mu_m = (q_m, s_m)$ satisfies $(q_m(x), s_m(x)) = (q(x), s(x))$ for $x = x_1, ..., x_m$, and $(q_m(x), s_m(x)) \in \arg\max_{1 \leq n \leq m}(q(x_n) \cdot x - s(x_n))$ otherwise (the ties are broken arbitrarily).

this goes beyond just counting the number of terms in (6), which is the menu size), and, indeed, yield higher revenues by Theorem 6.

In Appendix A.2 we will provide a similar analysis relative to the separate instead of the bundling revenue; it uses the ∞-norm instead of the 1-norm.

6 Complex Mechanisms are Arbitrarily Better

Based on the result of the previous section we can now construct mechanisms whose revenues may be arbitrarily higher than the bundling revenue.

Proposition 7 *Let $k = 2$. For every finite $M \geq 1$ there exists a two-good mechanism μ with a menu of size M such that*

$$\beta(\mu) > \frac{1}{2} \ln M - 1.$$

Proof. Let $M = (N+1)^2 - 1$ where $N \geq 2$ is an integer. Let $g_0, g_1, ..., g_M$ be the $M+1 = (N+1)^2$ points of the $1/N$-grid of $[0,1]^2$ arranged in lexicographic order; i.e., in order of increasing first coordinate, and for equal first coordinate in order of increasing second coordinate (thus $g_0 = (0,0)$ and $g_M = (1,1)$).

For each $n \geq 1$, let $g_n = (i_1/N, i_2/N)$ with $i_1 \equiv i_1^n$ and $i_2 \equiv i_2^n$ integers between 0 and N, and define $y_n := (N+1-i_2, 1)$. We claim that for every $0 \leq j < n$ we have

$$(g_n - g_j) \cdot y_n \geq \frac{1}{N}. \tag{12}$$

Indeed, let $g_j = (\ell_1/N, \ell_2/N)$. Now $j < n$ implies either $i_1 = \ell_1$ and $i_2 \geq \ell_2 + 1$, in which case $(g_n - g_j) \cdot y_n = i_2/N - \ell_2/N \geq 1/N$, or $i_1 \geq \ell_1 + 1$, in which case $(g_n - g_j) \cdot y_n = (i_1/N - \ell_1/N)(N+1-i_2) + (i_2/N - \ell_2/N) \geq 1/N$ because $i_1 - \ell_1 \geq 1$ and $i_2 - \ell_2 \geq 0 - N = -N$.

Let $t_n := N^{n-1}$ and $x_n := N^n y_n$, and consider the mechanism $\mu = (q, s)$ with menu $\{(g_n, t_n)\}_{n=1}^M$ that is "seller-favorable"; i.e., when indifferent, the buyer chooses the outcome with the highest payment (that is, he breaks ties in favor of the seller; see Hart and Reny 2015). For every $0 \leq j < n$ we have

$$g_n \cdot x_n - g_j \cdot x_n = N^n (g_n - g_j) \cdot y_n \geq N^{n-1} = t_n \geq t_n - t_j,$$

and so $g_n \cdot x_n - t_n \geq g_j \cdot x_n - t_j$. Therefore a buyer of type x_n will not choose any menu entry (g_j, t_j) with $j < n$ (by seller-favorability when there is indifference, because $t_j < t_n$), and so $s(x_n)$ is one of $\{t_n, t_{n+1}, ..., t_M\}$, which implies that $s(x_n) \geq t_n$. Thus $v(t_n) \leq ||x_n||_1 = N^n(N + 2 - i_2^n)$, and so

$$\begin{aligned} \beta(\mu) &= \sum_{n=1}^{M} \frac{t_n - t_{n-1}}{v(t_n)} \geq \sum_{n=1}^{M} \frac{N^{n-1} - N^{n-2}}{N^n(N + 2 - i_2^n)} \\ &\geq \sum_{i_1=1}^{N} \sum_{i_2=1}^{N} \frac{1}{N} \frac{1}{N + 2 - i_2} = \sum_{\ell=2}^{N+1} \frac{1}{\ell} \\ &> \ln(N + 2) - 1 > \frac{1}{2} \ln M - 1 \end{aligned}$$

(in the second line we have dropped the terms with $i_1 = 0$ or $i_2 = 0$). ■

Corollary 8 *Let $k = 2$. For every finite $M \geq 1$ there exist a two-good random valuation X with support of size M and values in $[0, 1]^2$, and a mechanism μ with a menu of size M such that*

$$\frac{\text{Rev}(X)}{\text{BRev}(X)} \geq \frac{R(\mu; X)}{\text{BRev}(X)} > \frac{1}{2} \ln M - 1. \tag{13}$$

Proof. The mechanism μ is given by Proposition 7, and the valuation X is then constructed by Theorem 6; see the remarks following its proof concerning the support of X. ■

Remark. An explicit random valuation X that satisfies (13) is easily obtained from the proof of Proposition 7. Take $x_n = N^n y_n = N^n(N+1-i_2^n, 1)$, put $\xi_n := ||x_n||_1$, and let X have support $\{x_1, ..., x_M\}$ and distribution $\mathbb{P}[X = x_n] = \xi_1/\xi_n - \xi_1/\xi_{n+1}$ for every $1 \leq n \leq M$. Then $\text{BRev}(X) = \xi_1$ and $R(\mu; X) > \xi_1((1/2)\ln M - 1)$ (cf. the proof of Theorem 6). To get the valuations in $[0, 1]^2$ one just needs to rescale: divide everything by N^M.

Theorem A(i) immediately follows.

Proof of Theorem A(i). For $k = 2$ goods, take M large enough so that $(1/2)\ln M - 1 > 1/(3\varepsilon)$; then the random valuation X given by Corollary 8 satisfies, by Corollary 4, $\text{DR}\textsc{ev}(X) \leq 3 \cdot \text{BR}\textsc{ev}(X) < \varepsilon \cdot \text{R}\textsc{ev}(X)$. For $k > 2$ goods, just append to X another $k - 2$ goods with constant valuation 0; this does not affect any of the revenues. ■

7 A General Construction

We now generalize the above construction, yielding a mechanism μ with infinite $\beta(\mu)$, together with a corresponding random valuation X that has infinite optimal revenue, whereas all its simple revenues—bundled, separate, deterministic, finite menu—are bounded. The construction will also be useful for comparing the deterministic revenue to the bundling revenue.

Proposition 9 *Let $(g_n)_{n=0}^M$ be a finite or countably infinite sequence in $[0,1]^k$ starting with $g_0 = (0, ..., 0)$, and let $(y_n)_{n=1}^M$ be a sequence of points in $\mathbb{R}_+^k$ such that*

$$\text{gap}_n := \min_{0 \leq j < n} (g_n - g_j) \cdot y_n > 0$$

for all $n \geq 1$. Then for every $\varepsilon > 0$ there exists a k-good random valuation X with $0 < \text{BR}\textsc{ev}(X) < \infty$ and a mechanism μ with menu $\{(g_n, t_n)\}_{n=1}^M$ (for an appropriate sequence $(t_n)_{n=1}^M$) such that

$$\frac{\text{R}\textsc{ev}(X)}{\text{BR}\textsc{ev}(X)} \geq \frac{R(\mu; X)}{\text{BR}\textsc{ev}(X)} \geq (1 - \varepsilon) \sum_{n=1}^{M} \frac{\text{gap}_n}{||y_n||_1}. \qquad (14)$$

Moreover, for every finite $1 \leq m < M$ let μ_m denote the mechanism obtained by restricting μ to its first m menu entries $\{(g_n, t_n)\}_{n=1}^m$; then

$$\frac{\text{R}\textsc{ev}_{[m]}(X)}{\text{BR}\textsc{ev}(X)} \geq \frac{R(\mu_m; X)}{\text{BR}\textsc{ev}(X)} \geq (1 - \varepsilon) \sum_{n=1}^{m} \frac{\text{gap}_n}{||y_n||_1} - \varepsilon. \qquad (15)$$

Proof. Let $x_n := (t_n/\text{gap}_n) y_n$ where the sequence of positive numbers $(t_n)_{n \geq 1}$ increases fast enough so that the sequence $\xi_n := ||x_n||_1 = t_n ||y_n||_1/\text{gap}_n$

is increasing and $t_{n+1}/t_n > 1/\varepsilon$ for all $n \geq 1$. We have $\xi_n \geq t_n$ (because $\text{gap}_n \leq g_n \cdot y_n \leq ||y_n||_1$) and thus when M is infinite $(t_n)_n$ and $(\xi_n)_n$ both increase to infinity; when M is finite, we put $t_{M+1} = \xi_{M+1} = \infty$. For every $0 \leq j < n$,

$$g_n \cdot x_n - g_j \cdot x_n = \frac{t_n}{\text{gap}_n}(g_n - g_j) \cdot y_n \geq t_n \geq t_n - t_j$$

(for $j = 0$ put as usual $t_0 = 0$). Therefore, in the seller-favorable mechanism $\mu = (q, s)$ with menu $\{(g_n, t_n)\}_{n=1}^M$, the buyer of type x_n prefers the menu entry (g_n, t_n) to any entry (g_j, t_j) with $0 \leq j < n$. Therefore $s(x_n) \geq t_n$, and so $v(t_n) \leq ||x_n||_1$, and we get

$$\begin{aligned}\beta(\mu) &= \sum_{n=1}^{M} \frac{t_n - t_{n-1}}{v(t_n)} \geq \sum_{n=1}^{M} \frac{t_n - t_{n-1}}{||x_n||_1} \\ &= \sum_{n=1}^{M} \frac{t_n - t_{n-1}}{t_n} \frac{\text{gap}_n}{||y_n||_1} > (1-\varepsilon) \sum_{n=1}^{M} \frac{\text{gap}_n}{||y_n||_1}.\end{aligned}$$

Theorem 6 implies that there is a random valuation X satisfying (14); to get (15) for a finite $m < M$, use (11) and $\xi_{m+1} \geq t_{m+1} > t_m/\varepsilon$. ■

Before showing how to obtain the infinite separation of Theorem A(ii), we use this result for deterministic mechanisms, proving Theorem C(ii).

Proposition 10 *For every $k \geq 2$,*

$$\sup_X \frac{\text{DRev}(X)}{\text{BRev}(X)} \geq \sum_{\ell=1}^{k} \frac{1}{\ell} \binom{k}{\ell} > \frac{2^k - 1}{k}. \tag{16}$$

Proof. Let $I_0, I_1, I_2, ..., I_{2^k-1}$ be the 2^k subsets of $\{1, \ldots, k\}$ ordered in nondecreasing size (i.e., $|I_n|$ is nondecreasing in n), and let g_n be the indicator vector of I_n (i.e., $g_{n,i} = 1$ for $i \in I_n$ and $g_{n,i} = 0$ for $i \notin I_n$). Take $y_n = g_n$ (thus $||y_n||_1 = |I_n|$); then for $0 \leq j < n$ we have $g_j \cdot g_n = |I_j \cap I_n| < |I_n| = g_n \cdot g_n$ (the strict inequality is because otherwise I_n would be a subset of I_j, contradicting $|I_j| \leq |I_n|$ and $j \neq n$), and thus $\text{gap}_n \geq 1$ (in fact, $\text{gap}_n = 1$: take I_j

to be a subset of I_n with one less element). Thus

$$\sum_{n=1}^{2^k-1} \frac{\text{gap}_n}{||y_n||_1} \geq \sum_{n=1}^{2^k-1} \frac{1}{|I_n|} = \sum_{\ell=1}^{k} \frac{1}{\ell}\binom{k}{\ell},$$

and we use Proposition 9. Replacing each $1/\ell$ with the lower $1/k$ yields the final inequality. ■

Remarks. Let d_k denote the binomial sum in (16).

(a) A better lower bound on d_k, easily obtained by replacing each $1/\ell$ with the lower $1/(\ell+1)$, is[27] $d_k \geq (2^{k+1} - k - 2)/(k+1) \sim 2 \cdot (2^k - 1)/k$.

(b) For large k most of the mass of the binomial coefficients, whose sum is $2^k - 1$, is at those ℓ that are close to $k/2$, and so $d_k \sim 1/(k/2) \cdot (2^k - 1) = 2 \cdot (2^k - 1)/k$ (formally, use a standard large deviations inequality; in (a) above we got this only as a lower bound on d_k).

(c) For $k = 2$ goods we have $d_2 = \binom{2}{1}/1 + \binom{2}{2}/2 = 5/2$, which turns out to be tight: see Proposition 23 in Appendix A.1.

(d) Proposition 26 in Appendix A.2 shows that the same lower bound of $(2^k - 1)/k$ also holds relative to the separate (instead of the bundling) revenue, and even relative to the maximum of the two.

We now construct an infinite sequence of points for which the appropriate sum of gaps in Proposition 9 is infinite.

Proposition 11 *There exists an infinite sequence of points* $(g_n)_{n=1}^{\infty}$ *in* $[0,1]^2$ *with* $||g_n||_2 \leq 1$ *such that taking* $y_n = g_n$ *for all* n *we have* $\text{gap}_n = \Omega(n^{-6/7})$.

Proof. The sequence of points that we build is composed of a sequence of "shells," each containing multiple points. The shells get closer and closer to each other, approaching the unit sphere as the shell, N, goes to infinity: all the points g_n in the N-th shell are of length $||g_n||_2 = \sum_{\ell=1}^{N} \ell^{-3/2}/\alpha$, where $\alpha = \sum_{\ell=1}^{\infty} \ell^{-3/2}$ (which indeed converges; thus $||g_n||_2$ approaches 1 as n increases), and each shell N contains $N^{3/4}$ different points in it so that the angle between any two of them is at least $\Omega(N^{-3/4})$.

[27] The standard notation $f(k) \sim g(k)$ means that $f(k)/g(k) \to 1$ as $k \to \infty$.

We now estimate $g_n \cdot g_j = ||g_n||_2 \cdot ||g_j||_2 \cdot \cos(\theta)$ where θ denotes the angle between g_n and g_j. Let N be g_n's shell. For $j < n$ there are two possibilities: either g_j is in the same shell, N, as g_n or it is in a smaller shell $N' < N$. 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 $||g_n||_2 = \Theta(1)$ we have $g_n \cdot g_n - g_n \cdot g_j \geq \Omega(N^{-3/2})$. In the second case, $||g_n||_2 - ||g_j||_2 = \sum_{\ell=N'+1}^{N} \ell^{-3/2}/\alpha \geq N^{-3/2}/\alpha$, and so again since $||g_n||_2 = \Theta(1)$ we have $g_n \cdot g_n - g_n \cdot g_j \geq \Omega(N^{-3/2})$. Thus for any point g_n in the N-th shell we have $\text{gap}_n = \Omega(N^{-3/2})$. Since the first N shells together contain $\sum_{\ell=1}^{N} \ell^{3/4} = \Theta(N^{7/4})$ points, we have $n = \Theta(N^{7/4})$ and thus $\text{gap}_n = \Omega(N^{-3/2}) = \Omega(n^{-6/7})$. ■

This directly implies Theorem A(ii)—the infinite separation between the optimal revenue and the deterministic revenue—and also B(ii)—the revenue may increase polynomially in the menu size.

Proof of Theorems A(ii) and B(ii). For $k = 2$ the infinite sequence of points $(g_n)_{n=1}^{\infty}$ constructed in Proposition 11, together with $y_n = g_n$ for all n, satisfies $\sum_{n=1}^{m} \text{gap}_n/||g_n||_1 \geq \sum_{n=1}^{m} \text{gap}_n/\sqrt{2} \geq \Omega(\sum_{n=1}^{m} n^{-6/7})$ (recall that $||g_n||_2 \leq 1$ and so $||g_n||_1 \leq \sqrt{2}$). When $m = \infty$ this sum is infinite, and when m is finite it is $\Omega(m^{1/7})$. Applying Proposition 9 then gives a two-good random valuation X that satisfies $0 < \text{BRev}(X) < \infty$ (and thus $0 < \text{DRev}(X) < \infty$ as well), $\text{Rev}(X) = \infty$, and $\text{Rev}_{[m]}(X) = \Omega(m^{1/7})$ for every finite m. This proves the two results (for Theorem A(ii) just rescale X to make DRev equal to 1). ■

8 Separate Selling and Additive Menu Size

In this section we study the separate selling mechanism, and then refine our menu size measure.

We start with a simple comparison between the bundling and separate revenues.

Proposition 12 *For every $k \geq 2$ and every k-good random valuation X,*

$$\text{BRev}(X) \leq k \cdot \text{SRev}(X).$$

Proof. Let BREV(X) be achieved for a bundle price of p. If the separate auction offers each good for the price of p/k then whenever $\sum_i x_i \geq p$ we have $x_i \geq p/k$ for some i, and so one of the k goods will be acquired in the separate auction. ■

This is tight for $k = 2$ and correlated values.

Example 13 Let X_1 be distributed uniformly on $[0, 1]$, and consider the two-good random valuation $X = (X_1, 1 - X_1)$. The bundling revenue is 1, since the bundle is always worth 1 to the buyer. Each good is distributed uniformly on $[0, 1]$ and so the optimal revenue from each good is, by (4), 1/4 (obtained at price 1/2).

For larger values of k, we can get a stronger result, which gives the first inequality of Theorem D.

Proposition 14 *There exists a constant $c < \infty$ such that for every $k \geq 2$ and every k-good random valuation X,*

$$\text{BRev}(X) \leq c \log k \cdot \text{SRev}(X).$$

Proof. Let BREV(X) be achieved for bundle price p. We first assume without loss of generality that the support of X contains only points x with $\sum_i x_i = p$ or $\sum_i x_i = 0$. (This is without loss of generality, since the random variable X' defined by $X' := 0$ when $\sum_i X_i < p$ and $X := (p/\sum_i X_i)X$ satisfies BREV(X') =BREV(X), while SREV(X') ≤SREV(X) because $X' \leq X$ everywhere.[28]) We now make another assumption without loss of generality, namely, that $\sum_i X_i = p$ (and so BREV(X) $= p$). (This is without loss of generality because if we replace X with its conditional on $\sum_i x_i = p$, then all revenues are just rescaled by a factor of $1/\mathbb{P}\left[\sum_i X_i = p\right]$.)

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

[28] Use the monotonicity of the one-good revenue (Hart and Reny 2015 or Hart and Nisan 2017), or Myerson's (1981) characterization (4).

Proof 1: Let $e_i := \mathbb{E}[X_i]$ be the expected value of good i; then (using our assumptions) $\sum_i e_i = p$. The claim is that good i can be sold in a separate auction yielding a revenue of at least $(e_i - p/(2k))/(2(1+\log_2 k))$. The result is then implied by summing over all i.

Indeed, split the range of values of X_i into $(2+\log_2 k)$ subranges: a "low" subrange for which $X_i \le p/(2k)$, and, for each $j = 0, \ldots, \log_2 k$, a subrange where $p/(2^{j+1}) < X_i \le p/(2^j)$ (notice that since $X_i \le p$ 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 $1 + \log_2 k$ subranges contributes at least $((e_i - p/(2k))/(1 + \log_2 k)$ to this expectation. The lower bound of this subrange, $p/(2^{j+1})$, is smaller by a factor of at most 2 than any value in the subrange, and so setting it as the price for good i yields a revenue that is at least half of the contribution of this subrange to the expectation, i.e., at least $((e_i - p/(2k))/(2(1 + \log_2 k))$.

Proof 2: Let $r_i :=$Rev$(X_i) = \sup_{t>0} t \cdot (1 - F_i(t))$ (where F_i denotes the cumulative distribution function of X_i); then $1 - F_i(t) \le r_i/t$ and so (recall that $X_i \le p$ because $\sum_i X_i = p$)

$$\mathbb{E}[X_i] = \int_0^\infty (1 - F_i(t))\mathrm{d}t \le \int_0^{r_i} 1\mathrm{d}t + \int_{r_i}^{p} \frac{r_i}{t}\mathrm{d}t = r_i(1 + \ln p - \ln r_i).$$

Averaging over i and using the concavity in r of the function $r(1+\ln p - \ln r)$ yields

$$\frac{p}{k} = \frac{1}{k}\sum_{i=1}^{k} \mathbb{E}[X_i] \le \frac{s}{k}\left(1 + \ln p - \ln\frac{s}{k}\right)$$

where $s := \sum_i r_i$. Thus $p/s \le 1 + \ln(p/s) + \ln k$, from which it follows that[29] BRev$(X)/$SRev$(X) = p/s < 4 \ln k$. ■

To complete the proof of Theorem D we only need the simple relation in the opposite direction:

[29]The function $x - \ln x - 1 - \ln k$ is increasing in x, and is positive at $x = 4 \ln k$ (because $k \ge 2$ implies $k^3/\ln k > 4e$).

Proposition 15 *For every $k \geq 2$ and every k-good random valuation X,*

$$\text{SRev}(X) \leq k \cdot \text{BRev}(X).$$

Proof. For each good i we have $X_i \leq \sum_\ell X_\ell$, which implies that[30] $\text{Rev}(X_i) \leq \text{Rev}(\sum_\ell X_\ell) = \text{BRev}(X)$. Sum over i. ■

This is tight for every k, even for independent goods; see Example 27 in Hart and Nisan (2017).

Together with Corollary 4 we get the following bound, which by Theorem C is essentially tight.

Corollary 16 *There exists a constant $c < \infty$ such that for every $k \geq 2$ and every k-good random valuation X,*

$$\text{DRev}(X) \leq c 2^k \log k \cdot \text{SRev}(X).$$

For the special case of $k = 2$ goods, we have a somewhat tighter bound.

Proposition 17 *Let $k = 2$. For every two-good random valuation X,*

$$\text{DRev}(X) \leq 3 \cdot \text{SRev}(X).$$

Proof. A deterministic mechanism has at most three menu entries: either selling just one of the goods, or selling the bundle. The portion of the revenue that comes from those types that buy only good i cannot exceed $\text{Rev}(X_i)$, and the portion that comes from those that buy the bundle cannot exceed $\text{BRev}(X)$; in total, $\text{DRev}(X) \leq \text{SRev}(X) + \text{BRev}(X)$. The proof is completed using Proposition 12. ■

Optimal separate mechanisms sell each good i at a price p_i, and so have menu size $2^k - 1$ (the buyer can buy any set of goods $I \subset \{1, \ldots, k\}$ for the price $\sum_{i \in I} p_i$), and yet Proposition 15 shows that the separate revenue is at most k times the bundling revenue, rather than $2^k - 1$ times that (as is the

[30] See footnote 28 above.

case for menu size $2^k - 1$, and in particular for deterministic mechanisms; see Theorems B and C). Intuitively, this seems related to the fact that separate mechanisms have only k "degrees of freedom" or "parameters" (the k prices). To formalize this we introduce a more refined "additive menu size" complexity measure, as follows.

A mechanism has *additive menu size* m if it offers m *basic* menu entries to the buyer, who may now choose any *subset* of basic menu entries for which no good is allocated more than once in total. Thus, basic menu entries $(g_1, t_1), (g_2, t_2), ..., (g_m, t_m) \in [0, 1]^k \times \mathbb{R}$ translate to actual menu entries $(\sum_{n \in N} g_n, \sum_{n \in N} t_n)$ with N ranging over all nonempty subsets of $\{1, 2, ..., m\}$ for which $\sum_{n \in N} g_n \in [0, 1]^k$.

The corresponding revenue is:

- $\textsc{Rev}_{[m]*}(X)$, the *additive menu size-m revenue*, is the maximal revenue that can be obtained by mechanisms whose additive menu size is at most m.

The additive menu size of a mechanism can thus be only lower than its menu size; for separate selling, it is indeed at most k, rather than $2^k - 1$. Interestingly, the basic properties of the menu size of Propositions 1 and 2 of Section 4, namely, menu size 1 yields the bundling revenue, and the increase in revenue is at most linear in the menu size, hold for the additive menu size as well.

Proposition 18 *For every $k \geq 2$ and every k-good random valuation X,*

(i) $\textsc{Rev}_{[1]*}(X) = \textsc{Rev}_{[1]}(X) = \textsc{BRev}(X)$, *and*

(ii) $\textsc{Rev}_{[m]}(X) \leq \textsc{Rev}_{[m]*}(X) \leq m \cdot \textsc{BRev}(X)$ *for every $m \geq 1$.*

Proof. The only claim that is not immediate is the last inequality. Let (g, t) be a basic menu entry. If (g, t) is chosen (i.e., is part of the chosen subset) when the valuation is x then $g \cdot x - t \geq 0$ (otherwise, dropping it from the chosen subset would increase the buyer's payoff at x); therefore $\sum_i x_i \geq g \cdot x \geq t$, and we continue exactly as in the proof of Proposition 2:

the probability[31] α that (g,t) is chosen is at most $\mathbb{P}\left[\sum_{i=1}^{k} X_i \geq t\right]$, and thus the portion of the expected revenue, $t \cdot \alpha$, that comes from (g,t), is at most BREV(X). ■

This essentially implies that the results in this paper apply also to this more refined complexity measure. Moreover, by Propositions 15 and 18, it captures well the complexity of selling the goods separately, whose additive menu size is at most k.

9 Approximation for Bounded Domains

In this section we prove Theorem E: for two goods, when the valuations are bounded, say, in the range $[1,H]^2$, then mechanisms need not have more than a polylogarithmic in H menu size in order to obtain arbitrarily good approximations. It is a direct corollary of the following lemma that shows how to incur, with an appropriate bounded menu size, only a small loss of payment for every valuation x.

Lemma 19 *Let $k = 2$. For every $H > 1$ and $\varepsilon > 0$ there exists $m = \mathrm{O}(\varepsilon^{-5} \log^2 H)$ such that for every two-good mechanism $\mu = (q,s)$ whose nonzero payments lie in the range $[1,H]$ (i.e., $s(x) = 0$ or $s(x) \in [1,H]$ for all x) there exists a mechanism $\tilde{\mu} = (\tilde{q},\tilde{s})$ with menu size at most m that satisfies $\tilde{s}(x) \geq (1-\varepsilon)s(x)$ for all x.*

Proof. We will discretize the menu of the given mechanism μ. 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 subranges, each with a ratio of at most $H^{1/K}$ between its endpoints, where K is chosen so that $H^{1/K} \leq \varepsilon^2$, i.e., $K = \mathrm{O}(\varepsilon^{-2} \log H)$. We define a real function $\phi(s)$ by rounding s up to the top of its range and then multiplying by $1-\varepsilon$. Hence we have

[31] The sum of these probabilities over all basic menu entries may be as high as m, as these events are no longer disjoint.

$(1-\varepsilon)s < \phi(s) < (1-\varepsilon)(1+\varepsilon^2)s$. Then for any $s' < s(1-\varepsilon)$ we have $\phi(s) - \phi(s') < (1-\varepsilon)(1+\varepsilon^2)s - (1-\varepsilon)s' < s - s'$.

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 entry (q', s') with $s' < (1-\varepsilon)s$ still prefers $(q, \phi(s))$ in the new menu; thus in the new menu he pays $\phi(s')$ for some $s' \geq (1-\varepsilon)s$, and $\phi(s') > (1-\varepsilon)s' \geq (1-\varepsilon)^2 s$; his payment in the new menu is therefore at least $(1-\varepsilon)^2$ 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 $(1-\varepsilon)$, i.e., multiply both the q's and the s's by $(1-\varepsilon)$. This does not change the menu choice of any buyer, reduces the payments by a factor of exactly $1-\varepsilon$, and ensures that $q_1, q_2 \leq 1-\varepsilon$. We now round down each q_1 and each q_2 to an integer multiple of ε/K, and then add $\varepsilon i/K$ to each menu entry whose price is s^i. Notice that rounding down reduces each q by at most ε/K, and since higher-paying menu entries got a boost that is larger by at least ε/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 obtained a new mechanism whose payment is at least $(1-\varepsilon)^3 \geq 1-3\varepsilon$ times that of the original one (and so we redefine the ε in the proof to be 1/3 of the ε in the statement). There are $K = \mathrm{O}(\varepsilon^{-2} \log H)$ price levels and $\varepsilon^{-1}K = \mathrm{O}(\varepsilon^{-3} \log H)$ 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 $\mathrm{O}(\varepsilon^{-5} \log^2 H)$. ■

Proof of Theorem E. Let X be a two-good random valuation with values in $[1, H]^2$, and let $\mu = (q, s)$ be a two-good mechanism. We have $s(x) \leq q(x) \cdot x \leq 2H$ for every $x \in [1, H]^2$; and, because the revenue from X is at least 2 (obtained, for instance, by selling each good at the price 1), we can assume without loss of generality that $R(\mu; X) \geq 2$. First, we eliminate from the menu of μ all entries whose payment is less than[32] 2ε; any type x with $s(x) < 2\varepsilon$ then either pays 0, or some $s(y) \geq 2\varepsilon$. The loss in revenue, if any, is thus at most $2\varepsilon \cdot \mathbb{P}[s(X) < 2\varepsilon] \leq 2\varepsilon$. Let $\mu' = (q', s')$ denote the resulting mechanism; then the range of its nonzero payments is $[2\varepsilon, 2H]$. Applying Lemma 19 to μ' yields a new mechanism $\tilde{\mu} = (\tilde{q}, \tilde{s})$ with a menu of size $\mathrm{O}(\varepsilon^{-5} \log^2(H/\varepsilon))$, such that $\tilde{s}(x) \geq (1-\varepsilon)s(x)$ for all x, and thus

$$
\begin{aligned}
R(\tilde{\mu}; X) &\geq (1-\varepsilon)R(\mu'; X) \geq (1-\varepsilon)(R(\mu; X) - 2\varepsilon) \\
&\geq (1-2\varepsilon)R(\mu; X)
\end{aligned}
$$

(recall that $R(\mu; X) \geq 2$). ■

Notice that the polylogarithmic dependence of m on H is "about right" since the valuation X induced by the first m points in the construction of Proposition 11 (used for proving Theorem A(ii)) has $H = m^{\mathrm{O}(m)}$, and the $\Omega(m^{1/7})$ gap between $\mathrm{Rev}(X)$ and $\mathrm{Rev}_{[1]}(X)$ implies that for, say, $m = \mathrm{O}((\log H)^{1/8})$, we get $\mathrm{Rev}_{[m]}(X) = \mathrm{o}(\mathrm{Rev}(X))$.

For more than two goods, i.e., $k > 2$, we have a somewhat weaker result, stating that the menu size need only be polynomial in[33] H.

Proposition 20 *For every $k \geq 2$ and $\varepsilon > 0$ there is $m_0 = (k/\varepsilon)^{\mathrm{O}(k)}$ such that for every k-good random valuation X with values in $[0,1]^k$ and every $m \geq m_0$,*

$$\mathrm{Rev}_{[m]}(X) \geq \mathrm{Rev}(X) - \varepsilon.$$

This result is directly implied by the following lemma.

Lemma 21 *Let $m = (n+1)^k - 1$ where $n \geq 1$ is an integer. Then for every k-good random valuation X with values in $[0,1]^k$,*

$$\mathrm{Rev}_{[m]}(X) \geq \mathrm{Rev}(X) - \frac{2k}{\sqrt{n}}.$$

[32] Formally, for every x with $s(x) < 2\varepsilon$ we take $(q'(x), s'(x))$ to be a maximizer of $q(y) \cdot x - s(y)$ over all y with $s(y) \geq 2\varepsilon$ or $s(y) = 0$.

[33] This has been improved to $(\log H/\varepsilon)^{\mathrm{O}(k)}$ by Dughmi, Han and Nisan (2014), who also showed that the exponential dependence on k is essential.

Proof. Let X have values in $[0,1]^k$, and let $\mu = (q, s)$ be a mechanism.

Define a new mechanism $\tilde{\mu} = (\tilde{q}, \tilde{s})$ 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 $[0,1]^k$, 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 $\tilde{\mu}$ 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 $\sum_i x_i \leq k$), which implies that the seller's revenue at x from $\tilde{\mu}$ must be $\geq (1 - 1/\sqrt{n})(s(x) - k/\sqrt{n})$. Therefore $R(\tilde{\mu}; X) \geq (1-1/\sqrt{n})R(\mu; X) - k/\sqrt{n} \geq R(\mu; X) - 2k/\sqrt{n}$ (since $R(\mu; X) \leq \sum_i x_i \leq k$). ■

From Proposition 20 we can derive an essentially equivalent multiplicative approximation result.

Proposition 22 *For every $k \geq 2$, $\varepsilon > 0$, and $H > 1$, there is $m_0 = (H/\varepsilon)^{\mathrm{O}(k)}$ such that for every k-good random valuation X with values in $[1, H]^k$ and every $m \geq m_0$,*

$$\textsc{Rev}_{[m]}(X) \geq (1 - \varepsilon) \cdot \textsc{Rev}(X).$$

Proof. We first rescale $[1, H]$ to $[1/H, 1]$, which for multiplicative approximations is the same. We then design a mechanism that gives an additive approximation to within $\varepsilon k/H$, which using Proposition 20 requires a menu size m as stated. Now, since each X_i is bounded from below by $1/H$, the revenue of X is at least k/H (each good is sold for sure at the price $1/H$), and thus an $\varepsilon k/H$-additive approximation is also a $(1 - \varepsilon)$-multiplicative approximation, as required. ■

A Appendix

A.1 Comparing to the Bundling Revenue for Two Goods

Using the function β of Theorem 6, we compute here the precise bound on the ratio of the deterministic and separate revenues to the bundling revenue for two goods. In particular, it follows that the $d_2 = 5/2$ bound of Proposition 10 (see Remark (c) there) is tight.

Proposition 23 *For $k = 2$ goods,*

$$\sup_X \frac{\text{DRev}(X)}{\text{BRev}(X)} = \frac{5}{2} \quad \text{and} \quad \sup_X \frac{\text{SRev}(X)}{\text{BRev}(X)} = 2$$

where the supremum is over all two-good random valuations X.

Proof. We will compute the supremum of $\beta(\mu)$ over deterministic and separate mechanisms μ for two goods.

(i) A deterministic mechanism μ is given by nonnegative prices p_1, p_2, p_{12} for good 1, good 2, and the bundle (i.e., the buyer of type x chooses the maximum in $\{0, x_1 - p_1, x_2 - p_2, x_1 + x_2 - p_{12}\}$). Without loss of generality we assume that $p_1 \leq p_2 \leq p_{12}$; the first inequality because we can interchange the two coordinates, and the second because if $p_i > p_{12}$ then the menu entry $x_i - p_i$ is never chosen, and so replacing p_i with $p_i' := p_{12}$ does not affect the revenue.

If $p_1 > 0$ then $\beta(\mu) = (p_1 - 0)/v(p_1) + (p_2 - p_1)/v(p_2) + (p_{12} - p_2)/v(p_{12})$. Now $v(p_1) = p_1$ (attained at $x = (p_1, 0)$) and $v(p_2) = p_2$ (attained at $x = (0, p_2)$); as for $v(p_{12})$, if $s(x) = p_{12}$ then $x_1 + x_2 - p_{12} \geq x_i - p_i$ for $i = 1, 2$, which implies $x_{3-i} \geq p_{12} - p_i \geq p_{12} - p_2$, and so $x_1 + x_2 \geq 2(p_{12} - p_2)$. Therefore $\beta(\mu) \leq 1 + 1 + 1/2 = 5/2$.

If $p_1 = 0 < p_2$ then $\beta(\mu) = (p_2 - 0)/v(p_2) + (p_{12} - p_2)/v(p_{12}) \leq 1 + 1/2 = 3/2$.

If $p_1 = p_2 = 0 < p_{12}$ then $\beta(\mu) \leq 1$.

Finally, if $p_1 = p_2 = p_{12} = 0$ then $\beta(\mu) = 0$.

Thus $\beta(\mu) \leq 5/2$ in all cases; taking, say, $p_1 = 1$, $p_2 = M$, and $p_{12} = M^2$ for large[34] M shows that $\sup_\mu \beta(\mu)$ over deterministic mechanisms is indeed 5/2.

(ii) A separate selling mechanism μ is a deterministic mechanism with $p_{12} = p_1 + p_2$. In this case $v(p_1 + p_2) = p_1 + p_2$ (attained at $x = (p_1, p_2)$), and

[34] Alternatively, use the bound of Proposition 10 (see Remark (b) there).

so $\beta(\mu) = 1 + 1 - p_1/p_2 + p_1/(p_1 + p_2)$, which is less than 2, but can be made arbitrarily close to 2 by taking, say, $p_1 = 1$ and $p_2 = M$ for large M. ∎

Remark. For *symmetric* deterministic mechanisms, where $p_1 = p_2$, we get $\beta(\mu) \leq p_1/p_1 + (p_{12} - p_1)/(2(p_{12} - p_2)) = 3/2$, with equality for, say, $p_1 = p_2 = 1$ and $p_{12} = 2$ (which is in fact a symmetric separate selling mechanism). Thus $\sup_\mu \beta(\mu)$ over all symmetric deterministic mechanisms μ, and $\sup_\mu \beta(\mu)$ over all symmetric separate selling mechanisms μ, are both equal to 3/2.

A.2 Comparing to the Separate Revenue

In this appendix we compare the revenue of an arbitrary mechanism to the revenue obtainable from selling the goods separately. The analysis is parallel to the one carried out with respect to the bundling revenue in Sections 5–7, but we use now the maximum norm $||x||_\infty = \max_i |x_i|$ instead of the 1-norm.

Let $\mu = (q, s)$ be a k-good mechanism. For every $t > 0$ we define

$$w(t) := \inf\{||x||_\infty : x \in \mathbb{R}^k_+ \text{ and } s(x) \geq t\},$$

and then

$$\gamma(\mu) := \int_0^\infty \frac{1}{w(t)} \, \mathrm{d}t.$$

Clearly $\beta(\mu) \leq \gamma(\mu) \leq k\beta(\mu)$ (because $||x||_1 \geq ||x||_\infty \geq ||x||_1/k$).

Let

$$\text{S-B-R}\textsc{ev}(X) := \max\{\text{SR}\textsc{ev}(X), \text{BR}\textsc{ev}(X)\}$$

stand for the maximal revenue that can be obtained by selling either separately or bundled.

Theorem 24 *Let $\mu = (q, s)$ be a k-good mechanism. Then*

$$\frac{1}{k}\gamma(\mu) \leq \sup_X \frac{R(\mu; X)}{\text{S-B-R}\textsc{ev}(X)} \leq \sup_X \frac{R(\mu; X)}{\text{SR}\textsc{ev}(X)} \leq \gamma(\mu),$$

where the supremum is over all k-good random valuations X.

Proof. First, for every $t > 0$ we have

$$\begin{aligned}\mathbb{P}[s(X) \geq t] &\leq \mathbb{P}[||X||_\infty \geq w(t)] = \mathbb{P}[\cup_i \{X_i \geq w(t)\}] \leq \sum_i \mathbb{P}[X_i \geq w(t)] \\ &\leq \sum_i \frac{\text{Rev}(X_i)}{w(t)} = \frac{\text{SRev}(X)}{w(t)}.\end{aligned}$$

Integrating over t yields $R(\mu, X) \leq \gamma(\mu) \cdot \text{SRev}(X)$.

Second, we show that for every $\gamma' < \gamma(\mu)$ there exists a k-good random valuation X such that $0 <$BRev$(X) < \infty$ and $R(\mu; X)/\text{S-B-Rev}(X) > \gamma'/k$. Let $0 = t_0 < t_1 < ... < t_M < t_{M+1} = \infty$ with $0 = w(t_0) < w(t_1) < w(t_2) < ... < w(t_M) < w(t_{M+1}) = \infty$ such that

$$\gamma'' := \sum_{n=1}^{M} \frac{t_n - t_{n-1}}{w(t_n)} > \gamma'.$$

Take $\varepsilon > 0$ be small enough so that $\gamma'' > (1+\varepsilon)\gamma'$ and $w(t_{n+1}) > (1+\varepsilon)w(t_n)$ for all $1 \leq n \leq M$, and choose for each $1 \leq n \leq M$ a point $x_n \in \mathbb{R}^k_+$ such that $s(x_n) \geq t_n$ and $w(t_n) \leq ||x_n||_\infty < (1+\varepsilon)w(t_n)$; then

$$\sum_{n=1}^{M} \frac{t_n - t_{n-1}}{||x_n||_1} > \sum_{n=1}^{M} \frac{t_n - t_{n-1}}{w(t_n)(1+\varepsilon)} = \frac{\gamma''}{1+\varepsilon} > \gamma'. \tag{17}$$

Let X be a random variable with support $\{x_1, ..., x_M\}$ and distribution $\mathbb{P}[X = x_n] = \xi_1/\xi_n - \xi_1/\xi_{n+1}$ for every $1 \leq n \leq M$, where $\xi_n := ||x_n||_\infty$ and we put $\xi_{M+1} := \infty$; thus $\mathbb{P}[X \in \{x_n, ..., x_M\}] = \xi_1/\xi_n$ for every $n \geq 1$.

Consider good i. For every $u \in (\xi_{n-1}, \xi_n]$ (with $1 \leq n \leq M$) we have

$$u \cdot \mathbb{P}[X_i \geq u] \leq u \cdot \mathbb{P}[X \in \{x_n, ..., x_M\}] = u\frac{\xi_1}{\xi_n} \leq \xi_1$$

(because $X = x_j$ for some $j \leq n-1$ implies $X_i \leq ||x_j||_\infty \leq ||x_{n-1}||_\infty = \xi_{n-1} < u$). Therefore $\text{Rev}(X_i) = \sup_{u>0} u \cdot \mathbb{P}[X_i \geq u] \leq \xi_1$ for every good i, and so $\text{SRev}(X) \leq k\xi_1$.

Next, for every $u \in (k\xi_{n-1}, k\xi_n]$ we have

$$u \cdot \mathbb{P}\left[\sum_{i=1}^{k} X_i \geq u\right] \leq u \cdot \mathbb{P}\left[X \in \{x_n, ..., x_M\}\right] = u\frac{\xi_1}{\xi_n} \leq k\xi_1$$

(because $X = x_j$ for some $j \leq n-1$ implies $\sum_i X_i \leq k||x_j||_\infty \leq k||x_{n-1}||_\infty = k\xi_{n-1} < u$). Therefore $\text{BRev}(X) = \sup_{u>0} u \cdot \mathbb{P}\left[X_i \geq u\right] \leq k\xi_1$, and $\text{BRev}(X) > 0$ because X does not vanish. Altogether,

$$0 < \text{S-B-Rev}(X) \leq k\xi_1.$$

The computation of $R(\mu; X)$ that μ gets from X is

$$\begin{aligned} R(\mu; X) &\geq \sum_{n=1}^{M} s(x_n)\mathbb{P}\left[X = x_n\right] \geq \sum_{n=1}^{M} t_n \left(\frac{\xi_1}{\xi_n} - \frac{\xi_1}{\xi_{n+1}}\right) \\ &= \xi_1 \sum_{n=1}^{M} \frac{t_n - t_{n-1}}{\xi_n} > \xi_1\gamma' \geq \frac{\gamma'}{k} \cdot \text{S-B-Rev}(X) \end{aligned}$$

(recall (17)). ■

Remarks. *(a)* As in Theorem 6 (see the remarks following its proof), X may be taken so that its values are in $[0, 1]^k$, and its support is at most the size of the menu of μ.

(b) A better upper bound for $R(\mu; X)/\text{S-B-Rev}(X)$ is obtained from Theorem 6 (because $\beta(\mu) \leq \gamma(\mu)$); thus

$$\frac{1}{k}\gamma(\mu) \leq \sup_X \frac{R(\mu; X)}{\text{S-B-Rev}(X)} \leq \beta(\mu).$$

(c) The gap of k in Theorem 24 is correct. Take two goods. For μ that sells the bundle for the price of 1 we have $w(1) = 1/2$ (attained at $x = (1/2, 1/2)$) and so $\gamma(\mu) = 2$; the two-good random valuation X of Example 13 in Section 8 has $R(\mu; X)/\text{SRev}(X) = 1/(1/2) = \gamma(\mu)$. For μ that sells each good separately for the price of $1/2$ we have $w(1/2) = 1/2$

and $w(1)=1$, and so[35] $\gamma(\mu)=2$, but $R(\mu;X)/\text{SRev}(X)\le 1=\gamma(\mu)/k$ for any X (with equality for, say, the constant valuation $(1/2,1/2)$).

We come now to the constructions of Section 7. The parallel result for $\text{S-B-Rev}(X)=\max\{\text{SRev}(X),\text{BRev}(X)\}$ is:

Proposition 25 *Let $(g_n)_{n=0}^M$ be a finite or countably infinite sequence in $[0,1]^k$ starting with $g_0=(0,...,0)$, and let $(y_n)_{n=1}^M$ be a sequence of vectors in $\mathbb{R}_+^k$ such that*

$$\text{gap}_n := \min_{0\le j<n}(g_n-g_j)\cdot y_n > 0$$

for all $n\ge 1$. Then for every $\varepsilon>0$ there exists a k-good random valuation X with $0<\text{BRev}(X)<\infty$ and a mechanism μ with menu $\{(g_n,t_n)\}_{n=1}^M$ (for an appropriate sequence $(t_n)_{n=1}^M$) such that

$$\frac{\text{Rev}(X)}{\text{S-B-Rev}(X)} \ge \frac{R(\mu;X)}{\text{S-B-Rev}(X)} \ge (1-\varepsilon)\frac{1}{k}\sum_{n=1}^{M}\frac{\text{gap}_n}{||y_n||_\infty}.$$

Moreover, for every finite $1\le m<M$ let μ_m denote the mechanism obtained by restricting μ to its first m menu entries $\{(g_n,t_n)\}_{n=1}^m$; then

$$\frac{\text{Rev}_{[m]}(X)}{\text{S-B-Rev}(X)} \ge \frac{R(\mu_m;X)}{\text{S-B-Rev}(X)} \ge (1-\varepsilon)\frac{1}{k}\sum_{n=1}^{m}\frac{\text{gap}_n}{||y_n||_\infty}-\varepsilon.$$

The proof is omitted, as it is identical to that of Proposition 9, but with the ∞-norm instead of the 1-norm, and the construction is as in Theorem 24 instead of Theorem 6.

As a consequence, for deterministic mechanisms we get (see Corollary 16 in Section 8 for the opposite inequality):

Proposition 26 *For every $k\ge 2$,*

$$\sup_X \frac{\text{DRev}(X)}{\text{SRev}(X)} \ge \sup_X \frac{\text{DRev}(X)}{\text{S-B-Rev}(X)} \ge \frac{2^k-1}{k}.$$

[35]It is easy to see that $\gamma(\mu)=k$ for every k-good mechanism μ that sells the goods separately for positive prices.

Proof. We proceed exactly as in the proof of Propositon 10, but now we have $||y_n||_\infty = 1$, and so

$$\frac{1}{k}\sum_{n=1}^{2^k-1}\frac{\text{gap}_n}{||y_n||_\infty} = \frac{1}{k}\sum_{\ell=1}^{k}\binom{k}{\ell} = \frac{2^k-1}{k}.$$

■

Remark. For $k = 2$ goods, the supremum of $\gamma(\mu)$ over all deterministic mechanisms equals 3 (attained in the limit as $M \to \infty$ by prices $p_1 = 1, p_2 = M, p_{12} = M^2$; cf. the proof of Proposition 23). Thus,

$$\begin{aligned}\frac{3}{2} &\leq \sup_X \frac{\text{DRev}(X)}{\text{SRev}(X)} \leq 3 \\ \frac{3}{2} &\leq \sup_X \frac{\text{DRev}(X)}{\text{S-B-Rev}(X)} \leq \frac{5}{2}\end{aligned}$$

(cf. Proposition 23: the supremum of DRev/BRev is exactly 5/2).

A.3 The Unit Demand Model

In this section we briefly compare our model to the unit-demand model considered in several papers, in particular in Briest *et al.* (2015). There are k goods 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 modeled as having unit-demand valuations. Additionally, the unit-demand model was defined to allow the mechanism to offer only single goods, rather than bundles of goods 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 $I \subset \{1, \ldots, k\}$ of goods, its value is $\max_{i \in I} x_i$ (rather than $\sum_{i \in I} x_i$). A deterministic mechanism in this setting would offer a price p_i for each good i. For unit-demand buyers this is equivalent to a completely general deterministic mechanism as there is no need to offer prices for bundles since the buyer is not interested in them. Thus, for example, a mechanism asking price p_1 for good 1, price p_2 for good 2, and price p_{12} for both goods would be the same as asking price $\min(p_1, p_{12})$ for good 1 and price $\min(p_1, p_{12})$ for good 2.

A randomized mechanism 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 goods, 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 mechanisms that are also allowed to offer lotteries for bundles of goods. For example, a menu entry offering the lottery "good 1 with probability 2/9; good 2 with probability 3/9; and both goods with probability 4/9" for a certain price can be replaced by the two menu entries "good 1 with probability 6/9; good 2 with probability 3/9" and "good 1 with probability 2/9; good 2 with probability 7/9," each for the same price as the original menu entry.

Let us use the notation $\text{Rev}^{UD}(X)$ to denote the revenue obtainable from a unit-demand buyer with a k-good random valuation X. Similarly $\text{DRev}^{UD}(X)$ denotes the revenue achievable by deterministic mechanisms. We can compare these revenues to those achievable in our model from an additive buyer whose valuation for the k goods is given by the same X.

Proposition 27 *For every $k \geq 2$ and every k-good random valuation X,*

(i) $\text{Rev}^{UD}(X) \leq \text{Rev}(X) \leq k \cdot \text{Rev}^{UD}(X)$*, and*

(ii) $\text{DRev}^{UD}(X) \leq \text{DRev}(X) \leq k2^k \cdot \text{DRev}^{UD}(X)$*.*

Proof. The lower bounds in both cases are obtained by noting that any mechanism in the unit-demand model offers only unit-demand menu entries, and on these both the unit-demand buyer and the additive buyer have the same preferences; thus offering the same menu in our setting gives exactly the same revenue as it does in the unit-demand setting.

For the upper bound for randomized mechanisms in (i), notice that if we replace each menu entry $(g_1, ..., g_k; t)$ in our model (where $0 \leq g_i \leq 1$ for each i) by the menu entry $(g_1/k, \ldots, g_k/k; t/k)$, then we do not change the preferences of the buyer between the different menu entries, and thus the revenue drops by a factor of exactly k. However, the new mechanism gives only unit-demand allocations (because $g_1/k + ... + g_k/k \leq 1$), where the unit-demand buyer and the additive buyer behave the same.

For the upper bound for deterministic mechanisms in (ii), consider a deterministic mechanism 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 from one of them, which allocates, say, a set I of goods. A mechanism that offers to sell only this set I of goods for the same price t as the original mechanism did will thus make at least a 2^{-k} fraction of the revenue of the original one. Now consider the unit-demand mechanism that offers each one of the goods in I for the price $t/|I|$; whenever the additive buyer in the additive mechanism buys I we are guaranteed that his value for at least one of the goods in I is at least $t/|I|$, in which case the unit-demand buyer will also acquire that good at $t/|I|$ in the unit-demand mechanism. ■

The interesting gap in the above proposition is the exponential one for deterministic mechanisms in (ii), and indeed we can show that this is essentially tight.

Proposition 28 *For every $k \geq 2$,*

$$\sup_X \frac{\text{DRev}(X)}{\text{DRev}^{UD}(X)} \geq \frac{2^k - 1}{k}.$$

Proof. For every X we have $\text{DRev}^{UD}(X) \leq \text{SRev}(X)$ because the good prices used in any deterministic mechanism in the unit-demand model can only yield more revenue in our additive model where the buyer may buy more than a single good. Use Proposition 26 in Appendix A.2. ■

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

A.4 Multiple Buyers

This paper has concentrated on a single-buyer scenario that may also be interpreted to be a monopolistic price 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, with 0-value for all goods) buyers. The issue of extending the results to the multiple-buyer settings is thus relevant to the upper bounds in the paper, both the significant ones (Theorem 20 and Proposition 14) and the simple ones (Propositions 1, 2, and Corollary 4). 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: dominant strategy or Bayesian Nash. Also, we need to specify whether we assume independence between buyers' valuations or allow them to be correlated. The discussion here will be coarse enough 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 in the case of multiple buyers. In the single-buyer case we defined it as the number of options from which the buyer may choose, which is the same as the number of allocations $|\{q(x) : x \in \mathbb{R}^k_+\}\backslash\{(0, \ldots, 0)\}|$. In the case of multiple buyers, these are two separate notions. For example, consider deterministic auctions of k goods among n buyers. There are a total of $(n+1)^k$ different allocations (each good may go to any buyer or to no one), but each buyer considers only 2^k possibilities (whether *he* gets each good or not). Moreover, the set of allocations cannot be interpreted as a menu from which the buyers may choose, since each buyer can choose only from the possibilities offered to him (and these choices need not be feasible overall). It takes the combined actions of all the buyers together in order for the mechanism outcome to be determined. For this reason we prefer to define the menu size of a multiple-buyer mechanism 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 take the maximum. We thus define:

- An n-buyer mechanism has *menu size* at most m if for every buyer $j = 1, \ldots, n$ and every $(n-1)$-tuple of (direct) bids of the other buyers[36] $x^{-j} \in (\mathbb{R}_+^k)^{n-1}$, the number of nonzero choices that buyer j faces is at most m, i.e., $|\{q^j(x^j, x^{-j}) : x^j \in \mathbb{R}_+^k\} \backslash \{(0, \ldots, 0)\}| \leq m$.

Note: If the original mechanism was incentive compatible in dominant strategies then the mechanism induced on player j by x^{-j} is also incentive compatible. If the original mechanism was incentive compatible in the Bayesian Nash sense then this need not be the case, but we still have individual rationality[37] of the induced mechanism, which suffices for what comes next.

Let us first analyze the simplest mechanisms, those with a single non-trivial menu entry for each buyer. Clearly, bundling auctions satisfy this property; however, not every mechanism that has a single non-trivial menu entry for each buyer can be converted to a bundling mechanism. We also need to be careful with the meaning of a bundling mechanism. Clearly, in the case of correlated buyer valuations, the optimal mechanism for selling even a single good (the whole bundle in our case) is not necessarily to sell it to the highest bidder, but rather to use the bids of the others to set the reserve price for 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 for the other.) Thus, in the rest of the discussion below we use BRev to denote the optimal revenue from mechanisms that sell the bundle only as a whole—not necessarily to the highest bidder or at a uniform reserve price. For the case of independent buyer values, the simpler version that sells it to the highest bidder at a fixed reserve price will suffice as well.

What can be easily observed is that by focusing solely on the buyer that pays the largest fraction of the revenue, we can reduce the problem to the

[36]Superscripts are used here for the buyers.

[37]This assumes that the original mechanism was ex-post individually rational, which one may verify is without loss of generality relative to ex-ante individual rationality.

single-buyer case and extract at least a $1/n$ fraction of revenue by selling the bundle to that single buyer. A full bundling mechanism can only do better, which gives us the analog to Proposition 1 for the case of n buyers:[38]

$$\mathrm{BR{\small EV}}^n(X) \leq \mathrm{R{\small EV}}^n_{[1]}(X) \leq n \cdot \mathrm{BR{\small EV}}^n(X).$$

The loss of the factor of n can be seen to be justified by considering independent buyer values and the restricted definition of bundling mechanisms already in the case of one good (i.e., $k = 1$): take the distribution where each buyer $j = 1, ..., n$ values the single good at M^j with probability M^{-j}, and zero otherwise (independently over buyers), for a large enough but fixed M.

A similar argument that focuses on the single buyer that provides the largest fraction of revenue yields the generalization of Proposition 2:

$$\mathrm{R{\small EV}}^n_{[m]}(X) \leq n \cdot m \cdot \mathrm{BR{\small EV}}^n(X),$$

and of Corollary 4:

$$\mathrm{DR{\small EV}}^n(X) \leq n \cdot (2^k - 1) \cdot \mathrm{BR{\small EV}}^n(X).$$

It turns out that the linear loss in n is required here too, again for independent buyer values and the restricted interpretation of bundling mechanisms: take the construction of Theorem C for 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 Propositions 20 and 14 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 sharper results, in which the additional loss due to the number of buyers is avoided, might be obtained.

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

References

Alaei, S., H. Fu, N. Haghpanah, J. Hartline, and A. Malekian (2012), “Bayesian Optimal Auctions via Multi- to Single-Agent Reduction,” *EC 2012: Proceedings of the 13th ACM Conference on Electronic Commerce.*

Babaioff, M., Y. Gonczarowski, and N. Nisan (2017), “The Menu-Size Complexity of Revenue Approximation,” *STOC 2017: Proceedings of the 49th ACM Symposium on Theory of Computing*, forthcoming.

Babaioff, M., N. Immorlica, B. Lucier, and S. M. Weinberg (2014), “A Simple and Approximately Optimal Mechanism for an Additive Buyer,” *FOCS 2014: Proceedings of the 55th Annual Symposium on Foundations of Computer Science*, 21–30.

Briest, P., S. Chawla, R. Kleinberg, and M. Weinberg (2015), “Pricing Randomized Allocations,” *Journal of Economic Theory* 156, 144–174.

Cai, Y., C. Daskalakis, and S. M. Weinberg (2012a), “An Algorithmic Characterization of Multi-Dimensional Mechanisms,” *STOC 2012: Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, 459–478.

Cai, Y., C. Daskalakis, and S. M. Weinberg (2012b), “Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization,” *FOCS 2012: Proceedings of the 55th Annual Symposium on Foundations of Computer Science.*

Chawla, S., J. D. Hartline, and R. D. Kleinberg (2007), “Algorithmic Pricing via Virtual Valuations,” *EC 2007: Proceedings of the 8th ACM Conference on Electronic Commerce*, 243–251.

Chawla, S., J. D. Hartline, D. L. Malec, and B. Sivan (2010), “Multi-Parameter Mechanism Design and Sequential Posted Pricing,” *STOC 2010: Proceedings of the 42nd ACM Symposium on Theory of Computing*, 311–320.

Chawla, S., D. L. Malec, and B. Sivan (2010), “The Power of Randomness in Bayesian Optimal Mechanism Design,” *EC 2010: Proceedings of the 11th ACM Conference on Electronic Commerce*, 149–158.

Daskalakis, C., A. Deckelbaum, and C. Tzamos (2013), “Mechanism Design via Optimal Transport,” *EC 2013: Proceedings of the 14th ACM Conference on Electronic Commerce*, 269–286.

Daskalakis, C., A. Deckelbaum, and C. Tzamos (2014), "The Complexity of Optimal Mechanism Design," *SODA 2014: Proceedings of the 25th Annual ACM–SIAM Symposium on Discrete Algorithms*, 1302–1318.

Daskalakis, C., A. Deckelbaum, and C. Tzamos (2017), "Strong Duality for a Multiple-Good Monopolist," *Econometrica* 85, 735–767.

Daskalakis, C., and S. M. Weinberg (2012), "Symmetries and Optimal Multi-dimensional Mechanism Design," *EC 2012: Proceedings of the 13th ACM Conference on Electronic Commerce*.

Dughmi, S., L. Han, and N. Nisan (2014), "Sampling and Representation Complexity of Revenue Maximization," *Lecture Notes in Computer Science* 8877, 277–291.

Fang, H. and P. Norman (2006), "To Bundle or Not to Bundle," *RAND Journal of Economics* 37, 946–963.

Giannakopoulos, Y. (2014), "Bounding Optimal Revenue in Multiple-Items Auctions," *arXiv* 1402.2382.

Giannakopoulos, Y. and E. Koutsoupias (2014), "Duality and Optimality of Auctions for Uniform Distributions," *EC 2014: Proceedings of the 15th ACM Conference on Electronic Commerce*, 259–276.

Gonczarowski, Y. A. (2017), "Bounding the Menu-Size of Approximately Optimal Auctions via Optimal-Transport Duality," *arXiv* 1708.08907.

Hart, S. and N. Nisan (2017), "Approximate Revenue Maximization with Multiple Items," *Journal of Economic Theory* 172, 313–347.

Hart, S. and P. J. Reny (2015), "Maximal Revenue with Multiple Goods: Nonmonotonicity and Other Observations," *Theoretical Economics* 10, 893–922.

Hart, S. and P. J. Reny (2016), "The Better Half of Selling Separately," mimeo.

Jehiel, P., M. Meyer-ter-Vehn, and B. Moldovanu (2007), "Mixed Bundling Auctions," *Journal of Economic Theory* 134, 494–512.

Krishna, V. (2010), *Auction Theory*, Second Edition, Academic Press.

Lev, O. (2011), "A Two-Dimensional Problem of Revenue Maximization," *Journal of Mathematical Economics* 47, 718–727.

Li, X. and A. C.-C. Yao (2013), “On Revenue Maximization for Selling Multiple Independently Distributed Items,” *Proceedings of the National Academy of Sciences* 110, 11232–11237.

Manelli, A. M. and D. R. Vincent (2006), “Bundling as an Optimal Selling Mechanism for a Multiple-Good Monopolist,” *Journal of Economic Theory* 127, 1–35.

Manelli, A. M. and D. R. Vincent (2007), “Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly,” *Journal of Economic Theory* 137, 153–185.

Manelli, A. M. and D. R. Vincent (2012), “Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly. A Corrigendum,” *Journal of Economic Theory* 147, 2492–2493.

McAfee, R. P. and J. McMillan (1988), “Multidimensional Incentive Compatibility and Mechanism Design,” *Journal of Economic Theory* 46, 335–354.

Menicucci, D., S. Hurkens, and D.-S. Jeon (2015), “On the Optimality of Pure Bundling for a Monopolist,” *Journal of Mathematical Economics* 60, 33–42.

Morgenstern, J. and T. Roughgarden (2016), “Learning Simple Auctions,” *JLMR: Workshop and Conference Proceedings* 49, 1–21.

Myerson, R. B. (1981), “Optimal Auction Design,” *Mathematics of Operations Research* 6, 58–73.

Pavlov, G. (2011), “Optimal Mechanism for Selling Two Goods,” *The B.E. Journal of Theoretical Economics: Advances* 11, Issue 1, Article 3.

Pycia, M. (2006), “Stochastic vs Deterministic Mechanisms in Multidimensional Screening,” MIT (mimeo).

Riley, J. G. and W. F. Samuelson (1981), “Optimal Auctions,” *American Economic Review* 71, 381–392.

Riley, J. and R. Zeckhauser (1983), “Optimal Selling Strategies: When to Haggle, When to Hold Firm,” *Quarterly Journal of Economics* 98, 267–289.

Rubinstein, A. and S. M. Weinberg (2015), "Simple Mechanisms for a Subadditive Buyer and Applications to Revenue Monotonicity," *EC 2015: Proceedings of the Sixteenth ACM Conference on Economics and Computation*, 377–394.

Tang, P. and Z. Wang (2017), "Optimal Mechanisms with Simple Menus," *Journal of Mathematical Economics* 69, 54–70.

Thanassoulis, J. (2004), "Haggling over Substitutes," *Journal of Economic Theory* 117, 217–245.

Yao, A. C.-C. (2014), "An n-to-1 Bidder Reduction for Multi-Item Auctions and its Applications," *Proceedings of the Twenty-Sixth Annual ACM–SIAM Symposium on Discrete Algorithms*, 92–109.