Selling Multiple Correlated Goods: Revenue Maximization and Menu-Size Complexity*

Sergiu Hart† Noam Nisan‡

March 14, 2019

Abstract

We consider the well known, and notoriously difficult, problem of 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 when there are two (or more) goods, *simple mechanisms*—such as selling the goods separately or as a bundle—*may yield only a negligible fraction of the optimal revenue.* This resolves the open problem of Briest, Chawla, Kleinberg, and Weinberg (*JET* 2015) who prove the result for at least three goods in the related setup of a unit-demand buyer. We also introduce the *menu*

*Previous versions, titled "The Menu-Size Complexity of Auctions": August 2012; April 2013 (The Hebrew University of Jerusalem, Center for Rationality DP-637 and http://arxiv.org/abs/1304.6116); November 2017. 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 insightful discussions, and the referees and editors of this journal for their many useful comments. 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`

size as a simple measure of the complexity of mechanisms, and show that the revenue may increase polynomially with menu size and that no bounded menu size can ensure any positive fraction of the optimal revenue. The menu size also turns out to "pin down" the revenue properties of deterministic mechanisms.

Contents

1 Introduction

Are complex auctions better than simple ones? Myerson's (1981) classic result (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 much 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 the goods separately. This is true even in the very simple setting where there is a single buyer.

In this paper we consider such a simple setting: a single seller, who aims 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, originally circulated in 2012), we considered the setup where the buyer's values for the different

[1] See Section 2 for a literature survey.

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

goods are *independent*, in which case we showed that simple mechanisms are *approximately* optimal: selling each good separately (deterministically) for its optimal price extracts a constant fraction of the optimal revenue. In this paper we show that the picture changes completely when the valuations of the goods are *correlated*, in which case "complex" mechanisms can become arbitrarily better than "simple" ones.

The setup is that of k 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. The buyer's valuation for a bundle of goods is additive over the goods; thus, for example, getting the first two goods is worth X_1+X_2 to the buyer. We denote by $\textsc{Rev}(X)$ the optimal revenue achievable by any mechanism for selling k goods to an additive buyer with a random valuation X.

Consider first the case of just two goods, i.e., $k = 2$. When the valuations of the two goods are independent (i.e., X_1 and X_2 are independent random variables), Hart and Nisan (2017) showed that selling the goods separately—each one at its optimal one-good price—is guaranteed to yield at least 50% of the optimal revenue, a bound that was later improved to 62% by Hart and Reny (2017). This can be stated in terms of the "Guaranteed Fraction of Optimal Revenue" (GFOR)[3] as

$$\text{GFOR}(\textsc{separate};\ 2\ \textit{independent goods}) \geq \frac{\sqrt{e}}{\sqrt{e}+1} \approx 0.62.$$

How does this fraction change when the two goods need not be independent? Our first result is that it drops all the way down to zero:

$$\text{GFOR}(\textsc{separate};\ 2\ \textit{goods}) = 0.$$

Indeed, we show that[4]

[3]See Hart and Nisan (2017): given a class of mechanisms $\mathcal{N}$ and a class of valuations $\mathbb{X}$, we define $\text{GFOR}(\mathcal{N};\mathbb{X})$ as the maximal fraction of the optimal revenue that can be achieved by mechanisms in $\mathcal{N}$ for any valuation in $\mathbb{X}$ (cf. Section 3.4).

[4]A stronger result is in fact proved; see Section 4 for precise statements of the main results.

For every $\varepsilon > 0$ ***there exists a two-good random valuation*** X ***with values in*** $[0,1]^2$ ***such that***

$$\text{SRev}(X) < \varepsilon \cdot \text{Rev}(X),$$

where SRev stands for the "separate revenue" achievable by selling the goods separately. Thus, for correlated goods, selling separately may yield only an arbitrarily small fraction of the optimal revenue. We emphasize that, while we provide specific such random valuations X, none of the constructions in this paper are knife-edge or pathological (see Remark 6.2).

This suggests considering the other one-dimensional mechanism, namely, that of selling the two goods as a bundle. That does not help: the guaranteed fraction of optimal revenue is still zero; i.e.,

$$\text{GFOR}(\text{bundled};\ 2\ \textit{goods}) = 0.$$

In fact, even the larger class of all "deterministic" mechanisms—in which the seller sets a price for each good separately as well as a price for the bundle—does not fare any better:

$$\text{GFOR}(\text{deterministic};\ 2\ \textit{goods}) = 0. \tag{1}$$

This immediately extends to any number of goods $k \geq 2$ (just add $k-2$ goods with zero valuation):

$$\text{GFOR}(\text{deterministic};\ k \geq 2\ \textit{goods}) = 0. \tag{2}$$

While these results are new in the case of $k = 2$ goods, they have already been established for $k \geq 3$ goods in the related model of a *unit-demand* (instead of additive) buyer—i.e., a buyer who wants to get *only one* of the k goods—by Briest, Chawla, Kleinberg, and Weinberg (2015, originally circulated in 2010); the case of two goods was left open, with some partial results indicating that GFOR may be bounded away from zero for $k = 2$. While the unit-demand model and our additive model are different, they are closely related: the various revenues in the two models are within constant factors of one another (see Appendix A.3 for precise statements). On the one hand, this implies that our result (2) for $k \geq 3$ goods follows from the above-mentioned result of Briest et al. (2015); on the other hand, our result (1) solves their open problem for $k = 2$: there is an infinite gap between the deterministic revenue and the optimal revenue in the unit-demand model, already for two goods.

What these results say is that allowing for probabilistic outcomes, where the buyer gets some goods with probabilities that are strictly between 0 and 1, makes a huge difference in terms of revenue. But is it really the probabilistic vs. deterministic distinction that matters here? A deterministic mechanism for k goods consists of setting prices for nonempty subsets of goods and thus provides to the buyer at most $2^k - 1$ nonzero outcomes to choose from. Suppose we were to limit the seller to provide the same number, i.e., $2^k - 1$, of outcomes, but allow these outcomes to be probabilistic; would that significantly increase the revenue? The answer is that it would not! As we will see, the guaranteed fraction of optimal revenue remains zero for *any* fixed bound on the number of outcomes.

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") specifies for each good i the probability q_i that it is allocated to the buyer, together with the payment s that the buyer pays to the seller;[5] it turns out to be convenient not to count the "zero" outcome of getting nothing and paying nothing (this outcome is always available, as it corresponds to the individual rationality or participation constraint). It is easy to see, and well known, 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. Let $\text{REV}_{[m]}(X)$ denote the optimal revenue achievable by mechanisms whose menu size is at most

[5]See Dobzinski (2011) for an earlier use of menu size in the context of combinatorial auctions.

m. For a single good, $k = 1$, the characterization of optimal mechanisms of Myerson (1981) implies that $\text{Rev}_{[1]}(X)$ is already the same as the optimal $\text{Rev}(X)$, 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.

Our general result is that for any fixed m, mechanisms that have at most m menu entries cannot guarantee any positive fraction of the optimal revenue:

$$\text{GFOR}(\text{MENU SIZE} \leq m;\ k\ \textit{goods}) = 0 \tag{3}$$

for any number of goods $k \geq 2$ and any menu size $m \geq 1$. Thus, 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. As stated above, taking $m = 2^k - 1$ yields result (2), which suggests that (2) is not driven by the mechanisms being deterministic, but rather by their being limited in the number of outcomes that they can offer.

Result (3) says that it does not matter exactly how "simple" mechanisms are defined; as long as their menu size is bounded (which is natural, as unbounded menu size can hardly be considered simple[6]), we have

> ***For multiple goods, simple mechanisms cannot guarantee any positive fraction of the optimal revenue.***

The fact that all simple mechanisms look equally "bad" when compared to the optimal revenue-maximizing mechanism does not however preclude some mechanisms from being better than others in terms of their revenues. This leads us to compare mechanisms by taking as a benchmark the simplest *basic* revenue (rather than the optimal revenue), which we take to be $\text{Rev}_{[1]}(X)$, the revenue that is achievable from *a single* take-it-or-leave it offer (i.e., a single menu entry); as we will see in Section 3.3, this basic revenue turns out to be nothing other than the revenue from selling the bundle of all goods at

[6]See the discussion in Section 3.2 on other complexity measures that do not use the "normal form" menu representation.

its optimal price, which we denote by BRev(X).[7] Thus, given a mechanism μ we define the "Multiple of Basic revenue" of μ, or MoB(μ) for short, to be the maximum, over all (relevant) valuations X, of the ratio of the revenue that μ extracts from X to the basic revenue $\text{Rev}_{[1]}(X)$ from X (the definition is then extended to classes of mechanisms by taking the maximum over the mechanisms in the class). Thus MoB(μ) measures how many times better the revenue from μ can be relative to the basic revenue.

The MoB measure turns out to be a useful tool for the analysis: first, it is given by a simple explicit formula (see Theorem 5.1); second, finding a sequence of mechanisms whose MoB goes to infinity is equivalent to proving that GFOR(bounded) = 0 (see Lemma 3.4(i)); and third, for any class $\mathcal{N}$ of mechanisms with a bounded MoB—such as deterministic mechanisms, or mechanisms with bounded menu size—the result that GFOR$(\mathcal{N}) = 0$ follows immediately from GFOR(bounded) = 0 (see Lemma 3.4(ii)).

We also show that the relation between MoB and menu size is polynomial (see Theorem C); that MoB of deterministic mechanisms is exponential in the number of goods (specifically, for many goods, i.e., large k, MoB of deterministic mechanisms is essentially the same as MoB of mechanisms with the same menu size, i.e., $2^k - 1$; see Theorem D); and, finally, that MoB of separate-selling mechanisms is linear in the number of goods (specifically, it equals the number of goods k; see Theorem E).

Our results thus show that the menu size, although it is just a simple and crude measure of the complexity of mechanisms, is nevertheless strongly related to revenue-extraction capabilities.

1.1 Organization of the Paper

In Section 2 immediately below we briefly go over some of the related literature. Section 3 presents our model, defines the menu-size complexity measure and the revenue comparison tools GFOR and MoB, and provides some preliminary results. The main results are then stated in Section 4. Section 5 deals with the MoB measure, which is then used in Sections 6 and 7 to

[7] The "B" in BRev may thus stand also for "Basic."

construct valuations that prove our results; see Section 4.3 for a detailed guide to the proofs. Section 8 studies separate selling, and introduces a more refined "additive menu size" complexity measure. We conclude in Section 9 with positive approximation results for the case where the valuations are in a bounded domain. Additional results are relegated to the appendices: the computation of MoB for two-good deterministic mechanisms (Appendix A.1); the use of the separate-selling revenue, instead of the bundled revenue, as the "basic" revenue (Appendix A.2); the relations between our setup and the unit-demand setup (Appendix A.3); and the multiple-buyer case (Appendix A.4).

2 Literature

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

The realization that maximizing revenue with multiple goods is a complex problem has had a long history in economic theory and more recently in the computer science literature as well. McAfee and McMillan (1988) identify 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.[8] These papers contain good surveys of the related 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), Hart and Reny (2015). In the past few years algorithmic work on these types of topics has been carried out. One line of work shows that for discrete distributions the optimal mechanism can be found by linear programming in rather general settings: Briest, Chawla, Kleinberg, and Weinberg (2010/2015),[9] Cai, Daskalakis, and Weinberg (2012a), Alaei, Fu, Haghpanah, Hartline, and Malekian (2012). Another line of work deals with optimal mecha-

[8]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 is more than one good.

[9]By "2010/2015" we mean "conference publication in 2010 and journal publication in 2015."

nisms for multiple goods in various settings: Daskalakis, Deckelbaum, and Tzamos (2013, 2017), Giannakopoulos (2014), Giannakopoulos and Koutsoupias (2014), Menicucci, Hurkens, and Jeon (2015), Tang and Wang (2017). Yet another line of work attempts to approximate the optimal revenue by simple mechanisms in various settings, where simplicity is defined qualitatively: 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). In this line of research, Hart and Nisan (2012/2017) consider mechanisms that sell the goods either separately or as a single bundle to be simple mechanisms, and show that when the values of the goods are *independently* distributed then a nontrivial fraction of the optimal revenue can be ensured by simple mechanisms. This was followed by various improved approximation results for independently distributed goods: Li and Yao (2013), Babaioff, Immorlica, Lucier, and Weinberg (2014), Yao (2014), Rubinstein and Weinberg (2015), Babaioff, Nisan, and Rubinstein (2018). By contrast, Briest, Chawla, Kleinberg, and Weinberg (2010/2015) consider deterministic mechanisms to be simple, and, in the unit-demand setting with at least 3 *correlated* goods, prove that deterministic mechanisms cannot ensure any positive fraction of the revenue of general mechanisms.

Approaches to quantifying the complexity of mechanisms are studied by Balcan, Blum, Hartline, and Mansour (2008), Dughmi, Han, and Nisan (2014), Morgenstern and Roughgarden (2015); we discuss these in Section 3.2. Since the circulation in 2013 of early versions of the present paper there has been additional work on menu-size complexity; see Babaioff, Gonczarowski, and Nisan (2017), Gonczarowski (2017), the tutorial of Goldner and Gonczarowski (2018), and the references there.

When the valuations are bounded, the approximation of auctions and mechanisms by various discretizations is studied by Hartline and Koltun (2005), Balcan, Blum, Hartline, and Mansour (2008) (where the construction is attributed to Nisan), Briest, Chawla, Kleinberg, and Weinberg (2010/2015), Daskalakis and Weinberg (2012), Dughmi, Han, and Nisan (2014); see the discussion following the statement of Theorem B in Section 4.

3 Preliminaries

3.1 The Model

The basic model is standard, and the notation follows 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 values of the goods to the buyer. The value of 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 valuation 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[10] (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[11] i (for $i = 1, 2, ..., k$), and $s(x)$ is the payment that the seller receives from the

[10] All functions in this paper are assumed to be Borel-measurable (cf. Hart and Reny 2015, footnotes 10 and 48).

[11] 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.

buyer; we refer to $(q(x), s(x))$ as an *outcome.* When the buyer reports his value x truthfully, his payoff is[12] $b(x) = \sum_{i=1}^{k} q_i(x)x_i - s(x) = q(x) \cdot x - s(x)$, 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 $x \in \mathbb{R}_+^k$.

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., $R(\mu; X) = \mathbb{E}\left[s(X)\right]$. We now define

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

As seen in Hart and Nisan (2017), when maximizing revenue we can limit ourselves without loss of generality to IR and IC mechanisms that satisfy in addition the *no positive transfer* (**NPT**) property: $s(x) \geq 0$ for every $x \in \mathbb{R}_+^k$ (and so $s(0, 0, ..., 0) = b(0, 0, ..., 0) = 0$).

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

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

where F is the cumulative distribution function of X. Thus, there are optimal mechanisms where the seller "posts" a price p and the buyer buys 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 $\text{Rev}(X)$, we are also interested in what can be obtained from certain classes of mechanisms.

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

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

$$\text{SR}\textsc{ev}(X) = \text{R}\textsc{ev}(X_1) + \text{R}\textsc{ev}(X_2) + ... + \text{R}\textsc{ev}(X_k).$$

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

$$\text{BR}\textsc{ev}(X) = \text{R}\textsc{ev}(X_1 + X_2 + ... + X_k).$$

- $\text{DR}\textsc{ev}(X)$, 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 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 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.,

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

(we ignore the zero outcome, $((0, 0, ..., 0), 0)$, which is always included without loss of generality as it corresponds to the IR constraint[13]). 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 $M \cup \{((0, 0, ..., 0), 0)\}$, whose

[13]We thus slightly depart from Hart and Reny (2015) (where the menu includes the zero outcome as well); this yields simple relations (such as Proposition 3.1) between menu size and revenue.

menu is included in[14] M (the mechanism is well defined up to tie-breaking; see Hart and Reny 2015 for more details).

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

$$\textsc{menu size}(\mu) := |\textsc{menu}(\mu)|.$$

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 and t' (if, say, $t' > t$ then (g, t') will never be chosen, as (g, t) is strictly preferred to it by every buyer type), the menu size is identical to the cardinality of the set of nonzero allocations; i.e.,

$$\textsc{menu size}(\mu) = |\{q(x) : x \in \mathbb{R}^k_+ \text{ and } q(x) \neq (0, 0, ..., 0)\}|.$$

The corresponding revenue is:

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

We will refer to the menu-size-1 revenue $\textsc{Rev}_{[1]}$ as the *basic revenue*: it is the revenue achievable from a single take-it-or-leave-it offer.

Interestingly, Babaioff, Gonczarowski, and Nisan (2017) have recently shown that the *communication complexity* of a mechanism is precisely the base 2 logarithm of its menu size.

Menu size is clearly a very crude measure of the complexity of a mechanism. In particular, it is based on the "normal" form of the mechanism (namely, a menu), and so it ignores the fact that a large menu may well be representable in a very succinct manner. Such an approach, namely, a *Kolmogorov complexity* notion, is used by Dughmi, Han and Nisan (2014). The *additive menu size*, a refinement of menu size that we introduce in Section 8, is also a step in this direction. Another approach is based on *learning-like* notions of "dimension": see Balcan, Blum, Hartline, and Mansour (2008) and

[14] 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.

Morgenstern and Roughgarden (2015). The advantage of our menu-size measure is that it is simple, it is defined for each mechanism separately (rather than for classes of mechanisms), and, as we will see below,[15] it provides useful connections to revenue-extraction capabilities.

3.3 Basic Results on Menu Size

We provide here a few simple and immediate relations concerning menu-size complexity and revenue.

Proposition 3.1 *For every $k \geq 2$ and every k-good random valuation X we have*

(i)

$$\textsc{Rev}_{[1]}(X) = \textsc{BRev}(X); \tag{5}$$

(ii) for any integers $m_1, m_2 \geq 1$

$$\textsc{Rev}_{[m_1+m_2]}(X) \leq \textsc{Rev}_{[m_1]}(X) + \textsc{Rev}_{[m_2]}(X);$$

(iii) the sequence $\frac{1}{m}\textsc{Rev}_{[m]}(X)$ is weakly decreasing in m, and thus, in particular, for every integer $m \geq 1$

$$\textsc{Rev}_{[m]}(X) \leq m \cdot \textsc{Rev}_{[1]}(X); \tag{6}$$

(iv)

$$\textsc{DRev}(X) \leq (2^k - 1) \cdot \textsc{Rev}_{[1]}(X) = (2^k - 1) \cdot \textsc{BRev}(X).$$

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

[15] See the tutorial of Goldner and Gonczarowski (2018) and the references there for additional such results.

(ii) Let $\mu = (q, s)$ be any mechanism with menu $(g_n, t_n)_{n=1}^{m_1+m_2}$. For each menu entry (g_n, t_n) let π_n be the probability that it is chosen (when the valuation is X), then the revenue from μ is $\sum_{n=1}^{m_1+m_2} \pi_n t_n$. Let μ_1 and μ_2 be mechanisms with menus $(g_n, t_n)_{n=1}^{m_1}$ and $(g_n, t_n)_{n=m_1+1}^{m_1+m_2}$, respectively. The probability that (g_n, t_n) for $n \le m_1$ is chosen is at least as large in μ_1 as it is in μ (since every valuation $x \in \mathbb{R}_+^k$ that prefers this menu item in μ continues to prefer it in μ_1, and all t_n are ≥ 0 by NPT), which implies that the revenue from μ_1 is at least $\sum_{n=1}^{m_1} \pi_n t_n$. A similar argument shows that the revenue from μ_2 is at least $\sum_{n=m_1+1}^{m_1+m_2} \pi_n t_n$.

(iii) Let $\mu = (q, s)$ be any mechanism with menu $(g_n, t_n)_{n=1}^{m}$; for each menu entry (g_n, t_n) let π_n be the probability that it is chosen. Without loss of generality order the menu entries so that the sequence $\pi_n t_n$ is weakly decreasing. Let $m' < m$; the mechanism μ' with menu $(g_n, t_n)_{n=1}^{m'}$ yields as revenue at least $\sum_{n=1}^{m'} \pi_n t_n$, which is at least $(m'/m) \sum_{n=1}^{m} \pi_n t_n$ (because $\pi_n t_n$ is weakly decreasing). Thus $R(\mu'; X) \ge (m'/m) R(\mu; X)$.

(iv) A deterministic mechanism has menu size at most $2^k - 1$. ■

For small menu size m the inequalites in (ii) and (iii) are tight, as the example below shows that $\text{Rev}_{[m]}(X) = m \cdot \text{Rev}_{[1]}(X)$ for m not exceeding the number of goods k. They remain essentially tight for m up to $2^k - 1$ (by Theorem D below); as for large m, we will see that $\text{Rev}_{[m]}(X)$ can be as large as[16] $\Omega(m^{1/7}) \cdot \text{Rev}_{[1]}(X)$ (see Theorem C).

Example 3.2 Let $1 \le m \le k$. Take a large[17] $H > 0$, and consider the following random valuation X. For each $i = 1, ..., m$, with a probability α_i that is proportional to $1/H^{i-1}$, good i is valued at H^{i-1} and all the other goods are valued at 0; thus $\alpha_i = c/H^{i-1}$, where $c := 1/(1 + 1/H + ... + 1/H^{m-1})$. Bundling yields a revenue of 1 (because setting the bundle price at H^{i-1} yields a revenue of $c(1/H^{i-1} + ... + 1/H^{m-1})$, which is maximal at $i = 1$, and the revenue there is 1). Selling each good $i = 1, ..., m$ at price H^{i-1}

[16] It is convenient to use the standard O and Ω notations. For two expressions F and G that depend on certain variables, we write $F = \text{O}(G)$ 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 $F \le cG$, respectively $F \ge cG$, for any values of the variables in the relevant range.

[17] Theorem 5.1 below provides the tool to easily generate such examples.

yields a revenue of c from each good; this is obtained at distinct valuations, and so the mechanism consisting of these m menu entries yields a revenue of mc, which is close to m for large H.

3.4 Revenue Comparisons: GFOR and MoB

To evaluate how good mechanisms are we compare the revenue that they can extract to two benchmarks, a "high" one and a "low" one. The high benchmark is the *optimal* revenue $\textsc{Rev}$, and the low benchmark is the *basic* revenue $\textsc{Rev}_{[1]} = \text{B}\textsc{Rev}$.[18] As discussed in the Introduction, when the valuations of the goods are correlated the former yields infinite gaps in most cases of interest, and so the latter is needed to provide useful comparisons.

Formally, let $\mathbb{X}$ be a class of random valuations (e.g., k goods, two independent goods, and so on), and let $\mathcal{N}$ be a class of mechanisms (e.g., separate mechanisms, deterministic mechanisms, and so on).

We define

- GFOR($\mathcal{N}$; $\mathbb{X}$), the *Guaranteed Fraction of Optimal Revenue* (Hart and Nisan 2017), as the maximal fraction α such that, for any random valuation X in $\mathbb{X}$, mechanisms in the class $\mathcal{N}$ yield a revenue that is at least the fraction α of the optimal revenue; that is,[19]

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

  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[20] $\mathcal{N}$.

[18] See Appendix A.2 for a similar, but slightly less sharp, approach where the basic revenue is taken to be the separate-selling revenue $\text{S}\textsc{Rev}$.

[19] When taking the infimum we ignore the cases $0/0$ and ∞/∞ (because the inequality $\mathcal{N}\text{-}\textsc{Rev}(X) \geq \alpha\ \textsc{Rev}(X)$ holds for any α in these cases). The same applies when taking the supremum and, more generally, when dealing with any ratio of revenues throughout the paper.

[20] GFOR 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. See the remarks in Section 2.2 of Hart and Nisan (2017), which, in particular, explain why ratios are used.

- $\textsc{MoB}(\mathcal{N};\mathbb{X})$, the *Multiple of Basic revenue*, as the minimal multiple β such that, for any random valuation X in $\mathbb{X}$, mechanisms in the class $\mathcal{N}$ achieve a revenue that is at most the multiple β of the basic revenue; that is,

$$\textsc{MoB}(\mathcal{N};\mathbb{X}) := \sup_{X\in\mathbb{X}} \frac{\mathcal{N}\text{-}\textsc{Rev}(X)}{\textsc{Rev}_{[1]}(X)} = \sup_{X\in\mathbb{X}} \frac{\mathcal{N}\text{-}\textsc{Rev}(X)}{\textsc{BRev}(X)};$$

when $\mathcal{N}$ consists of a single k-good mechanism μ and $\mathbb{X}$ is the class of all k-good random valuations we write $\textsc{MoB}(\mu)$ for short.

Thus, $\textsc{MoB}(\mathcal{N};\mathbb{X})$ is the highest multiple of the basic revenue that may be achieved by the mechanisms in $\mathcal{N}$ for valuations in $\mathbb{X}$.

Putting $\alpha =$GFOR$(\mathcal{N};\mathbb{X})$ and $\beta =\textsc{MoB}(\mathcal{N};\mathbb{X})$, we then have

$$\alpha \cdot \textsc{Rev}(X) \leq \mathcal{N}\text{-}\textsc{Rev}(X) \leq \beta \cdot \textsc{Rev}_{[1]}(X)$$

for every random valuation X in $\mathbb{X}$, and both bounds are tight: i.e., for every $\alpha' > \alpha$ there is X in $\mathbb{X}$ with $\alpha'{\cdot}\textsc{Rev}(X) > \mathcal{N}\text{-}\textsc{Rev}(X)$, and for every $\beta' < \beta$ there is X in $\mathbb{X}$ with $\mathcal{N}\text{-}\textsc{Rev}(X) > \beta'{\cdot}\textsc{Rev}_{[1]}(X)$.

Remark 3.3 The results of Proposition 3.1(iii)–(iv) can be thus restated as

$$\begin{aligned} \textsc{MoB}(\textsc{menu size} \leq m;\ k\ \textit{goods}) &\leq m \quad \text{and} \\ \textsc{MoB}(\textsc{deterministic};\ k\ \textit{goods}) &\leq 2^k - 1. \end{aligned}$$

The following lemma provides simple but useful connections between GFOR and MoB.

Lemma 3.4 *Let $\mathcal{M}$ be the class of all (IC and IR) mechanisms, let $\mathcal{N} \subset \mathcal{M}$, and let $\mathbb{X}$ be a class of valuations. Then:*

(i)

$$\text{GFOR}(\textsc{bundled};\mathbb{X}) = \frac{1}{\textsc{MoB}(\mathcal{M};\mathbb{X})}; \quad \textit{and}$$

(ii)

$$\begin{aligned}\mathrm{GFOR}(\mathcal{N};\mathbb{X}) &\leq \mathrm{M{\scriptsize OB}}(\mathcal{N};\mathbb{X})\cdot \mathrm{GFOR}(\textsc{bundled};\mathbb{X}) \\ &= \frac{\mathrm{M{\scriptsize OB}}(\mathcal{N};\mathbb{X})}{\mathrm{M{\scriptsize OB}}(\mathcal{M};\mathbb{X})}.\end{aligned}$$

Proof. (i) GFOR(BUNDLED) $= \inf_X \mathrm{BR{\scriptsize EV}}(X)/\mathrm{R{\scriptsize EV}}(X)$ and $\mathrm{M{\scriptsize OB}}(\mathcal{M}) = \sup_X \mathrm{R{\scriptsize EV}}(X)/\mathrm{BR{\scriptsize EV}}(X)$.

(ii) $\mathcal{N}$-REV/REV$=$ ($\mathcal{N}$-REV/BREV) $\cdot$ (BREV/REV). ■

Thus, showing that there are mechanisms μ with arbitrarily large MOB proves that GFOR(BUNDLED) $= 0$ (by (i)), which then implies that GFOR(MENU SIZE $\leq m$) $= 0$ for any fixed m, and, in particular, GFOR(DETERMINISTIC) $= 0$ (by (ii) and Remark 3.3 above).

4 Main Results

We now state formally the main results, first for the Guaranteed Fraction of Optimal Revenue (GFOR), and then for the Multiple of Basic revenue (MOB), followed by an outline of the way in which these results are proved.

4.1 Results for GFOR

The results here are, first, that GFOR equals 0 for simple mechanisms, including those with bounded menu size, and, second, that in the case of bounded valuations GFOR becomes close to 1 for an appropriately large enough menu size.

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

(i)

$$\begin{aligned}\mathrm{GFOR}(\textsc{bundled};\ k\ \textit{goods}) &= 0;\\ \mathrm{GFOR}(\textsc{separate};\ k\ \textit{goods}) &= 0;\\ \mathrm{GFOR}(\textsc{deterministic};\ k\ \textit{goods}) &= 0;\\ \mathrm{GFOR}(\textsc{menu size} \leq m;\ k\ \textit{goods}) &= 0\end{aligned}$$

for every finite menu size $m \geq 1$.

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

$$\mathrm{DR{\small EV}}(X) < \varepsilon \cdot \mathrm{R{\small EV}}(X).$$

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

$$\mathrm{DR{\small EV}}(X) = 1 \quad \text{and} \quad \mathrm{R{\small EV}}(X) = \infty.$$

As discussed in the Introduction, our contribution lies in the result for the case of $k = 2$ goods, as for $k \geq 3$ it follows from Briest et al. (2014). In part (i), once we have the result that GFOR for BUNDLED is 0 all the other results immediately follow, because $\mathrm{R{\small EV}}_{[m]} \leq m \cdot \mathrm{BR{\small EV}}$ (by Proposition 3.1 and Lemma 3.4(ii) above). Clearly, (ii) and (iii) each yield (i). If we allow the valuations to be unbounded then we can get the fraction ε in (ii) to go all the way down to 0, which is (iii). Clearly (iii) implies (ii) (just truncate X beyond a high enough value); the construction that yields (ii) however is simpler and explicit. Part (ii) (and thus (i)) is proved in Section 6 and part (iii) in Section 7.[21]

Looking at the constructions used in the proof of Theorem A, 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[22] 0), say, in the range $[L, H]$, then the gap becomes bounded by some constant power of $\log(H/L)$; see Section 9, where we show that this exponential blowup in the range is indeed needed. Our result is:

Theorem B *Let $k = 2$. There exists a constant $c < \infty$ such that for every*

[21] Given the marginal distributions of the valuations of the two goods—which determine the separate revenue—we obtain joint distributions for which the revenue becomes arbitrarily *large*; by contrast, Carroll (2017) looks at the *smallest* joint revenue for given marginals.

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

$0 < L < H < \infty$ *and* $\varepsilon > 0$,

$$\text{GFOR}(\text{MENU SIZE} \leq m;\ 2 \textit{ goods with values in } [L,H]^2) \geq 1 - \varepsilon$$

holds for every menu size m *that satisfies*

$$m \geq \frac{c}{\varepsilon^5} \log^2 \left(\frac{H}{L} \frac{1}{\varepsilon} \right).$$

This theorem is proved in Section 9. Again, contrast this result with the unbounded range case: when the upper bound H is infinite (and $L > 0$) there is a valuation X with $\text{Rev}(X) = \infty$ while $\text{Rev}_{[m]}(X) \leq m$ for every finite m (by[23] Theorem A(iii) and (6)), and when the lower bound L is zero (and H is finite) for every finite m there is a valuation X with $\text{Rev}_{[m]}(X)/\text{Rev}(X) < m\varepsilon$ (by Theorem A(ii) and (6)).

Thus arbitrarily good approximations of the optimal revenue can be obtained, for two goods, by a menu size m that is only *polylogarithmic* in the range size H/L. This improves results obtainable by known techniques (Hartline and Koluim 2005, Balcan et al. 2008, Briest et al. 2015, and our Proposition 9.2 below), which yield a *polynomial* dependence on H/L (i.e., $m \geq (H/L\varepsilon)^{ck}$). Recently Dughmi, Han, and Nisan (2014) have extended the polylogarithmic result to all k (i.e., $m \geq (\log(H/(L\varepsilon))/\varepsilon)^{ck}$), and shown that the exponential dependence on k is necessary.

4.2 Results for MoB

The results here show the relations between MoB and menu size (polynomial), and, for deterministic and separate-selling mechanisms, between MoB and the number of goods (exponential for the former and linear for the latter).

Theorem C *There exists a constant* $c > 0$ *such that for every* $k \geq 2$ *and* $m \geq 1$

$$cm^{1/7} \leq \text{MoB}(\text{MENU SIZE} \leq m;\ k \textit{ goods}) \leq m.$$

[23] For the boundedness away from 0, see Remark 7.2.

As discussed above, the right-hand side inequality, whose simple proof is in Proposition 3.1, says that the revenue may grow at most linearly in the menu size; as for the left-hand side, which is obtained from our construction in the proof of Theorem A(iii) in Section 7, it says that the revenue may grow at least polynomially in menu size.[24]

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

Theorem D *For every*[25] $k \geq 2$*:*

$$\begin{aligned}\frac{2^k-1}{k} &\leq \text{MoB}(\text{DETERMINISTIC};\ k\ \textit{goods}) && (7)\\ &\leq \text{MoB}(\text{MENU SIZE} \leq 2^k-1;\ k\ \textit{goods}) \leq 2^k-1. && (8)\end{aligned}$$

The upper bound (8) is given, again, by Proposition 3.1; as for the lower bound (7), which is proved using the techniques of the proof of Theorem A(iii) in Section 7, it shows that the exponential-in-k bound is essentially tight (the factor k being much smaller than 2^k-1 for large k). Note again the contrast to the independent case, for which the bound is linear, rather than exponential,[26] in k: Lemma 28 in Hart and Nisan (2017) implies that for k independent goods $\text{DRev}(X) \leq \text{Rev}(X) \leq ck \cdot \text{Rev}_{[1]}(X)$ for some $c > 0$, and thus MoB(DETERMINISTIC; k *independent goods*) $\leq ck$.

The two inequalities in Theorem D say that the revenue that can be extracted by deterministic mechanisms is, for large k, of the same order of magnitude as for arbitrary mechanisms with a menu of size 2^k-1. This suggests that the reason that deterministic mechanisms yield low revenue

[24]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 D below shows that the growth can be almost linear in m.

[25]We obtain in fact a lower bound that is somewhat better than $(2^k-1)/k$; for large k, it is close to twice as much. See Proposition 7.3 and Remark 7.4.

[26]Proposition A.10 in Appendix A.2 below shows that the same exponential-in-k gap exists between deterministic mechanisms and separate selling: there is X such that $\text{DRev}(X) \geq (2^k-1)/k \cdot \text{SRev}(X)$. This provides a rare *doubly exponential* contrast with the independent case in which $\text{DRev}(X) \leq c \log^2 k \cdot \text{SRev}(X)$ for some constant c (by Theorem C in Hart and Nisan 2017).

(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.

Finally, we consider the maximal revenue SRev obtainable by selling each good separately (at its one-good optimal price). We have

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

$$\textsc{MoB}(\textsc{separate};\ k \text{ goods}) = k.$$

This theorem is proved in Section 8. Unlike in our previous results, the bound here is the same as the one we have obtained for independently distributed goods, and it is tight already in that case; see Proposition 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 may be at most k times, rather than $2^k - 1$ times, the bundling revenue. Moreover, selling separately seems intuitively to be much simpler than this exponential-in-k menu-size measure suggests: one needs to determine only k prices. All this leads us to define a stronger notion of mechanism complexity, one that assigns to separate selling its more natural complexity, namely, k. This new measure allows "additive menus" in which the buyer may choose not just single menu entries but also sets of menu entries. We present this *additive menu size* complexity measure in Section 8, and show that in fact our results hold with respect to this stronger complexity measure as well.

4.3 Outline of the Proofs

We present now a short but hopefully useful outline of the proofs in the following sections.

- In Section 5 we provide an explicit formula for MoB of a mechanism, and construct random valuations where MoB is (almost) attained (Theorem 5.1).

- In Section 6 we construct mechanisms with an arbitrarily large MoB, which shows that $\mathrm{MOB}(\mathcal{M}) = \infty$ and so $\mathrm{GFOR}(\textsc{bundled}) = 0$, thus proving Theorem A(i) and (ii).
- In Section 7 we construct a random valuation, for Theoreom A(iii), with an infinite gap between the revenue from simple mechanisms and the optimal revenue; we also prove the lower bound of Theorem D for deterministic mechanisms.
- In Section 8 we prove Theorem E for the separate revenue, and then introduce and analyze the more refined "additive-menu-size" measure.
- In Section 9 we deal with valuations in bounded domains and prove Theorem B.

5 The Multiple of Basic Revenue (MoB)

We start by providing 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. Recall that for a single k-good mechanism μ we write $\mathrm{MoB}(\mu)$ for short for $\mathrm{MoB}(\{\mu\}; k \text{ goods})$.

Theorem 5.1 *Let $\mu = (q, s)$ be a k-good mechanism. Then*

$$\mathrm{MoB}(\mu) = \int_0^\infty \frac{1}{v(t)}\, \mathrm{d}t,$$

where for every $t > 0$ we define[27]

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

[27]The 1-norm $||x||_1 = \sum_{i=1}^k |x_i|$ on $\mathbb{R}^k$ gives, for nonnegative x, the value $\sum_{i=1}^k x_i$ of the bundle of all goods to the buyer of type x. The infimum of an empty set is taken to be ∞, and so $v(t) = \infty$ when t is higher than any possible payment $s(x)$.

Thus $v(t)$ is the minimal value of the bundle, $x_1 + ... + x_k$, among all the valuations x where the payment to the seller is at least t. Geometrically, this says that the supporting hyperplane with normal $(1, ..., 1)$ to the set $\{x \in \mathbb{R}^k_+ : s(x) \geq t\}$ is $x_1 + ... + x_k = v(t)$. The function v is weakly increasing and satisfies $v(t) \geq t$ for every $t > 0$ (because $\sum_i x_i \geq q(x) \cdot x \geq s(x)$ for every x by IR); the function $1/v$ is nonnegative, weakly decreasing, and vanishes beyond the maximal possible payment (i.e., for $t > \sup_x s(x)$). Its integral may well be zero or infinite, i.e., $0 \leq \textsc{MoB}(\mu) \leq \infty$ (with $\textsc{MoB}(\mu) = 0$ 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 weakly increasing, we have $v(t) = v(t_n)$ for every $t_{n-1} < t \leq t_n$ (some of these intervals may well be empty[28]), and so

$$\textsc{MoB}(\mu) = \sum_{n=1}^{m} \frac{t_n - t_{n-1}}{v(t_n)} \tag{9}$$

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

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

Example 5.2 Let μ be given by the menu[29] $\{x_1 - p_1, x_2 - 2, x_1 + x_2 - 4\}$, and allow p_1 to vary.

(i) 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 $(2, 3)$; see Figure 1). Therefore $\textsc{MoB}(\mu) = (1 - 0)/1 + (2 - 1)/2 + (4 - 2)/5 = 19/10$. As we will see in the proof of Theorem 5.1 below, $\textsc{MoB}(\mu)$ is attained for the random valuation X that takes the values $(1, 0)$, $(0, 2)$, and $(2, 3)$ with probabilities $1/v(1) - 1/v(2) = 1/2$, $1/v(2) - 1/v(4) = 3/10$, and $1/v(4) = 1/5$, respectively; indeed, $\textsc{BRev}(X) = \max\{(1 + 0) \cdot 1, (0 + 2) \cdot (1/2), (2 +$

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

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

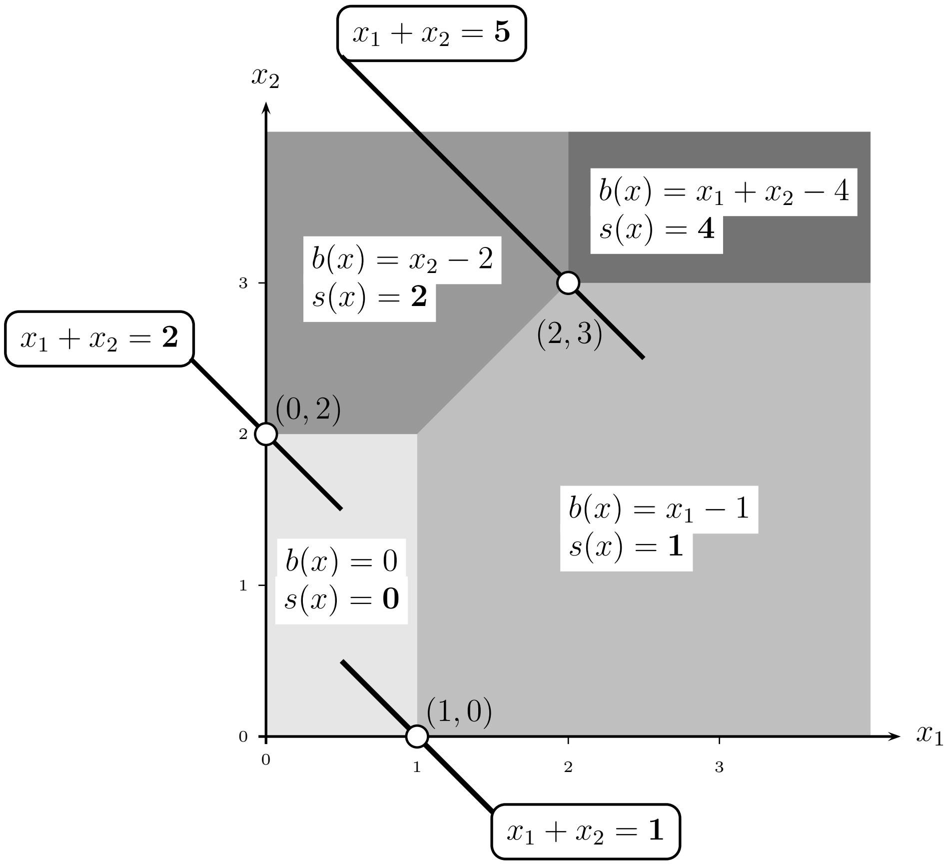

Figure 1: The function v in Example 5.2(i): $v(1) = ||(1,0)||_1 = 1$, $v(2) = ||(0,2)||_1 = 2$, $v(4) = ||(2,3)||_1 = 5$

$3) \cdot (1/5)\} = 1$ and[30] $R(\mu; X) = 1 \cdot (1/2) + 2 \cdot (3/10) + 4 \cdot (1/5) = 19/10$.

(ii) 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 $\text{MoB}(\mu) = (2-0)/2 + (2-2)/2 + (4-2)/4 = 3/2$.

(iii) 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 $\text{MoB}(\mu) = (2-0)/2 + (4-2)/4 + (5-4)/\infty = 3/2$.

Proof of Theorem 5.1. Put $\beta := \int_0^\infty 1/v(t)\, dt$.

[30] Assume without loss of generality that the buyer breaks ties in favor of the seller (i.e., the mechanism μ is "seller-favorable"); see Hart and Reny (2015).

(i) First, we show that

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

for every k-good random valuation X. Indeed,

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

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 $u \cdot \mathbb{P}[||X||_1 \geq u] \leq$ BRev(X) for every $u > 0$.

(ii) Second, we show that for every $\beta' < \beta$ (which, when β is infinite, is taken to mean any arbitrarily large β'), there exists a k-good random valuation X with $0 <$BRev$(X) < \infty$ and

$$\frac{R(\mu; X)}{\text{BRev}(X)} > \beta'. \tag{10}$$

Indeed, the function $1/v(t)$ is weakly decreasing and nonnegative, and its integral is β, and so there exist $0 = t_0 < t_1 < ... < t_N < t_{N+1} = \infty$ with $0 = v(t_0) < v(t_1) < v(t_2) < ... < v(t_N) < v(t_{N+1}) = \infty$ such that

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

Let $\varepsilon > 0$ be small enough so that $\beta'' > (1+\varepsilon)\beta'$ and $v(t_{n+1}) > (1+\varepsilon)v(t_n)$ for all $1 \leq n \leq N$. By the definition of v we can choose for every $1 \leq n \leq N$ a point[31] $x_n \in \mathbb{R}^k_+$ such that $s(x_n) \geq t_n$ and $v(t_n) \leq ||x_n||_1 < (1+\varepsilon)v(t_n)$; then

$$\sum_{n=1}^{N} \frac{t_n - t_{n-1}}{||x_n||_1} > \sum_{n=1}^{N} \frac{t_n - t_{n-1}}{v(t_n)(1+\varepsilon)} = \frac{\beta''}{1+\varepsilon} > \beta'. \tag{11}$$

[31] 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.

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 $\{x_1, ..., x_N\}$ and distribution $\mathbb{P}[X = x_n] = \xi_1/\xi_n - \xi_1/\xi_{n+1}$ for every $1 \leq n \leq N$, where we put $\xi_{N+1} := \infty$; thus $\mathbb{P}[X \in \{x_n, ..., x_N\}] = \xi_1/\xi_n$ for every[32] $n \geq 1$.

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

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

and so

$$\text{BR}\textsc{ev}(X) = \xi_1. \tag{12}$$

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

$$\begin{aligned} R(\mu; X) &\geq \sum_{n=1}^{N} s(x_n)\mathbb{P}[X = x_n] \geq \sum_{n=1}^{N} t_n \left(\frac{\xi_1}{\xi_n} - \frac{\xi_1}{\xi_{n+1}} \right) \\ &= \sum_{n=1}^{N} (t_n - t_{n-1}) \frac{\xi_1}{\xi_n} > \xi_1 \beta' = \beta' \cdot \text{BR}\textsc{ev}(X) \end{aligned} \tag{13}$$

(use $\xi_{N+1} = \infty$, (11), and (12)). ■

Remark 5.3 *(a)* In the proof of part (ii) above: for any $m < N$ let μ_m be obtained by restricting the menu of μ to the entries chosen by $x_1, ..., x_m$ in μ (with ties broken the same way as in μ for $x_1, ..., x_m$, and arbitrarily otherwise).[33] The computation of $R(\mu_m; X)$ is the same as in (13), but the sum is now going up only to m instead of N, and thus there is a final term

[32]Since the payment $s(x_n)$ increases with n, we want to put as much probability as possible on points x_n with high n, subject to the constraint that the bundled revenue is kept fixed, specifically, equal to $\xi_1 = ||x_1||_1$; for illustration see the random valuation X in Example 5.2(i) above.

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

of $t_m(\xi_1/\xi_{m+1})$ that needs to be subtracted; this gives

$$R(\mu_m; X) > \left(\sum_{n=1}^{m} \frac{t_n - t_{n-1}}{\xi_n} - \frac{t_m}{\xi_{m+1}} \right) \cdot \mathrm{BR{\small EV}}(X) \tag{14}$$

(recall (12)). This result will be used in Proposition 7.1 below.

(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 therefore rescale it (which does not affect the ratio of revenues) 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 $N \leq m$ 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 $\mathrm{M{\small O}B}(\mu) \leq m$ (because $v(t) \geq t$ implies that each term in the sum (9) is ≤ 1). This is the linear-in-menu-size bound of Proposition 3.1(iii); Example 3.2 in Section 3.3 above is obtained by making each term close to 1.

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

6 The Guaranteed Fraction of Optimal Revenue (GFOR)

Based on the result of the previous section we can now construct mechanisms whose revenues may be arbitrarily higher than the bundling revenue, which yields the GFOR$=0$ result.

Proposition 6.1 *Let $k = 2$. For every finite $m \geq 1$ there exists a two-good mechanism μ with a menu of size m such that*

$$\mathrm{MOB}(\mu) > \frac{1}{2} \ln m - 1.$$

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 the 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 $n \geq 1$, by writing the vector g_n as $g_n = (i_1/N, i_2/N)$ with $i_1 \equiv i_1^{(n)}$ and $i_2 \equiv i_2^{(n)}$ integers between 0 and N, we define $y_n := (N+1-i_2, 1)$. We claim that for every $0 \leq j < n$ we have

$$(g_n - g_j) \cdot y_n \geq \frac{1}{N}. \tag{15}$$

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

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, ties are broken 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

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

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

Thus $\text{MoB}(\text{menu size} \leq m)$ is at least of the order of $\log m$; in the next section we will improve this lower bound and show that it is polynomial in m. From Proposition 6.1 we immediately get parts (i) and (ii) of Theorem A.

Proof of Theorem A(i) and (ii). We prove this for $k = 2$ goods; for $k > 2$ goods we take the two-good random valuation and append $k-2$ goods with constant valuation 0, which does not affect any of the revenues.

(i) We have $\text{MoB}(\mathcal{M}; 2\ \textit{goods}) = \sup_\mu \text{MoB}(\mu) = \infty$ by Proposition 6.1, and so $\text{GFOR}(\text{bundled}; 2\ \textit{goods}) = 1/\text{MoB}(\mathcal{M}; 2\ \textit{goods}) = 0$ (see Lemma 3.4(ii) in Section 3.4).

(ii) For every finite $m \geq 1$, let μ be the mechanism given by Proposition 6.1, and then let X be a random valuation in $[0,1]^2$ with support of size m, as constructed by Theorem 5.1 (see Remark 5.3(b) and (c)), that satisfies

$$\frac{\text{Rev}(X)}{\text{BRev}(X)} \geq \frac{R(\mu; X)}{\text{BRev}(X)} > \frac{1}{2} \ln m - 1. \tag{16}$$

An explicit random valuation X that satisfies (16) is easily obtained from the proof of Proposition 6.1. 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 $\text{Rev}(X) \geq R(\mu; X) > \xi_1((1/2)\ln m - 1)$ (cf. the proof of Theorem 5.1). To get the valuations in $[0,1]^2$ one just needs to rescale: divide everything by N^m. Taking m large enough so that $(1/2)\ln m - 1 > 1/(3\varepsilon)$ then yields (use Proposition 3.1(iv)) $\text{DRev}(X) \leq 3{\cdot}\text{BRev}(X) < \varepsilon{\cdot}\text{Rev}(X)$. ■

Remark 6.2 *Any random valuation X' that is close to the above random valuation X will yield a similar gap between the optimal revenue and the simple revenues.*[34] The same applies to all our constructions, and so none of our results is knife-edge or pathological.

[34]For formal revenue continuity results, see Hart and Reny (2017, Appendix A).

7 A General Construction

We now generalize the construction of the previous section, and obtain a mechanism μ with infinite MoB, together with a corresponding random valuation X for which the optimal revenue is infinite, whereas all its simple revenues—bundled, separate, deterministic, finite-menu—are bounded; this proves Theorem A(iii). Proposition 7.1 below will turn out to be useful also for evaluating MoB of deterministic mechanisms, thereby proving Theorem D.

Proposition 7.1 *Let $(g_n)_{n=0}^N$ be a finite or countably infinite sequence in $[0,1]^k$ starting with $g_0 = (0,...,0)$, and let $(y_n)_{n=1}^N$ be a sequence of points in $\mathbb{R}_+^k$ such that*

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

for all $n \ge 1$. Then for every $\varepsilon > 0$ there exist a sequence $(t_n)_{n=1}^N$ of positive real numbers, a k-good mechanism μ with menu $\{(g_n, t_n)\}_{n=1}^N$, and a k-good random valuation X with $0 <\mathrm{BRev}(X) < \infty$ such that

$$\begin{aligned} \mathrm{MoB}(\mu) \quad &\ge \quad (1-\varepsilon)\sum_{n=1}^{N} \frac{\mathrm{gap}_n}{||y_n||_1}, \\ \frac{\mathrm{Rev}(X)}{\mathrm{BRev}(X)} \ge \frac{R(\mu;X)}{\mathrm{BRev}(X)} \quad &\ge \quad (1-\varepsilon)\sum_{n=1}^{N} \frac{\mathrm{gap}_n}{||y_n||_1}, \quad \textit{and} \qquad (17) \\ \frac{\mathrm{Rev}_{[m]}(X)}{\mathrm{BRev}(X)} \ge \frac{R(\mu_m;X)}{\mathrm{BRev}(X)} \quad &\ge \quad (1-\varepsilon)\sum_{n=1}^{m} \frac{\mathrm{gap}_n}{||y_n||_1} - \varepsilon \qquad (18) \end{aligned}$$

for every finite $1 \le m < N$, where μ_m denotes the mechanism obtained by restricting μ to its first m menu entries $\{(g_n, t_n)\}_{n=1}^m$.

Proof. Let $x_n := (t_n/\mathrm{gap}_n)y_n$ where the sequence of positive numbers $(t_n)_{n\ge1}$ increases fast enough so that the sequence $\xi_n := ||x_n||_1 = t_n||y_n||_1/\mathrm{gap}_n$ is increasing and $t_{n+1}/t_n \ge 1/\varepsilon$ for all $n \ge 1$. We have $\xi_n \ge t_n$ (because $\mathrm{gap}_n \le g_n \cdot y_n \le ||y_n||_1$) and thus, when N is infinite, $(t_n)_n$ and $(\xi_n)_n$ both increase to infinity; when N is finite, we put $t_{N+1} = \xi_{N+1} = \infty$. 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 $t_0 = 0$). Thus, in the seller-favorable mechanism $\mu = (q, s)$ with menu $\{(g_n, t_n)\}_{n=1}^N$, 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 $v(t_n) \leq ||x_n||_1 = \xi_n$, and we get

$$\begin{aligned} \text{MoB}(\mu) &= \sum_{n=1}^{N} \frac{t_n - t_{n-1}}{v(t_n)} \geq \sum_{n=1}^{N} \frac{t_n - t_{n-1}}{\xi_n} \\ &= \sum_{n=1}^{N} \frac{t_n - t_{n-1}}{t_n} \frac{\text{gap}_n}{||y_n||_1} \geq (1-\varepsilon) \sum_{n=1}^{N} \frac{\text{gap}_n}{||y_n||_1} \end{aligned} \tag{19}$$

(the final inequality follows from $t_{n-1}/t_n \leq \varepsilon$). As in the proof of Theorem 5.1, let X take the value x_n with probability $\xi_1/\xi_n - \xi_1/\xi_{n+1}$, then $R(\mu; X) \geq \xi_1 \cdot \sum_{n=1}^{N}(t_n - t_{n-1})/\xi_n$ and $\text{BRev}(X) \leq \xi_1$, which implies (17) (use (19)); to get (18) for a finite $m < N$, use (14) and $\xi_{m+1} \geq t_{m+1} \geq t_m/\varepsilon$. ■

Remark 7.2 The random valuation X that we have constructed in Proposition 7.1 is bounded away from zero: $||X||_1 \geq ||x_1||_1 = \xi_1 > 0$.

Before showing how to obtain the infinite separation of Theorem A(iii), we use Proposition 7.1 for deterministic mechanisms, proving the lower bound on MoB(DETERMINISTIC) of Theorem D (recall that the upper bound of $2^k - 1$ is immediate; see Proposition 3.1(iv)).

Proposition 7.3 *For every $k \geq 2$,*

$$\text{MoB}(\text{DETERMINISTIC};\ k \text{ goods}) \geq \sum_{\ell=1}^{k} \frac{1}{\ell}\binom{k}{\ell} > \frac{2^k - 1}{k}. \tag{20}$$

Proof. Let $I_0, I_1, I_2, ..., I_{2^k-1}$ be the 2^k subsets of $\{1, \ldots, k\}$ ordered in weakly increasing size (i.e., $|I_n| \geq |I_{n-1}|$ for all n), and let g_n be the indicator vector of I_n (i.e., the i-th coordinate of g_n is 1 for $i \in I_n$ and 0 for

$i \notin I_n$). Take $y_n = g_n$ (thus $||y_n||_1 = |I_n|$); then for $0 \le j < n$ we have $g_j \cdot g_n = |I_j \cap I_n| < |I_n| = g_n \cdot g_n$ (the strict inequality holds because otherwise I_n would be a subset of I_j, contradicting $|I_j| \le |I_n|$ and $j \neq n$), and thus $\text{gap}_n \ge 1$ (in fact, $\text{gap}_n = 1$: take I_j to be a subset of I_n with one less element). Thus

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

and we use Proposition 7.1. Replacing each $1/\ell$ with the lower $1/k$ yields the final inequality. ■

Remark 7.4 Let d_k denote the binomial sum in (20).

(a) A better lower bound on d_k, easily obtained by replacing each $1/\ell$ with the lower $1/(\ell+1)$, is[35] $d_k \ge (2^{k+1} - k - 2)/(k+1) \sim 2 \cdot (2^k - 1)/k$.

(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 deviation inequality; in (a) above we got this estimate 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 the exact value of MoB; see Proposition A.1 in Appendix A.1 (proved by using, again, Theorem 5.1).

(d) Proposition A.10 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 revenues.

We now construct, already for two goods, an infinite sequence of points for which the appropriate sum of gaps in Proposition 7.1 is infinite.

Proposition 7.5 *There exists an infinite sequence of points* $(g_n)_{n=1}^{\infty}$ *in* $[0,1]^2$ *with* $||g_n||_2 \le 1$ *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:

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

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 $\Omega(N^{-3/4})$.

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 $\cos(\theta) \leq 1 - \Omega(N^{-3/2})$ (because $\cos(x) = 1 - x^2/2 + x^4/24 - \ldots$) 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(iii), i.e., the infinite separation between the optimal revenue and the deterministic revenue, and also the lower bound in Theorem C, i.e., the revenue may increase polynomially in the menu size.

Proof of Theorems A(iii) and C. For $k = 2$ the infinite sequence of points $(g_n)_{n=1}^{\infty}$ constructed in Proposition 7.5, 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 7.1 gives a two-good random valuation X that satisfies $0 < \text{DRev}(X) < \infty$ (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. For $k > 2$, again, add $k-2$ goods with constant valuation 0. This proves the two results (for Theorem A(iii) just rescale X to make DRev equal to 1; and the upper bound in Theorem C is by Proposition 3.1). ■

8 Additive Menu Size

We start by proving Theorem E, which says that MoB of selling separately k goods equals the number of goods k.

Proof of Theorem E. For each good i we have $X_i \leq \sum_\ell X_\ell$, which implies that[36] $\text{Rev}(X_i) \leq \text{Rev}(\sum_\ell X_\ell) = \text{BRev}(X)$. Summing over i yields $\text{SRev}(X) \leq k \cdot \text{BRev}(X)$.

Example 27 in Hart and Nisan (2017) shows that this bound is tight for every k, even for independent goods. ■

Now optimal separate mechanisms sell each good i at a price p_i, and so have a menu size of at most $2^k - 1$ (the buyer can buy any set of goods $I \subseteq \{1, \ldots, k\}$ for the price $\sum_{i \in I} p_i$), and yet Theorem E shows that the separate revenue is at most k times the bundling revenue, rather than $2^k - 1$ times that (as is the case for menu size $2^k - 1$, and in particular for deterministic mechanisms; see Theorem D). Intuitively, this seems related to the fact that separate-selling mechanisms have only k "degrees of freedom" or "parameters" (the k prices).[37] To formalize this we introduce a more refined "additive menu size" complexity measure, as follows.

Let μ be a k-good mechanism with menu $M \subseteq [0,1]^k \times \mathbb{R}_+$. An *additive representation of* M is a subset $M_0 = \{(g_1, t_1), (g_2, t_2), ..., (g_m, t_m)\} \subseteq M$ of menu entries, which we will refer to as *basic menu entries*, such that every menu entry (g, t) in M can be represented as a sum of basic menu entries in M_0, i.e., $(g, t) = \sum_{n \in N} (g_n, t_n)$ for some $N \subseteq M_0$, and moreover every partial sum $\sum_{n \in N'} (g_n, t_n)$ with $N' \subset N$ is also a menu entry in[38] M. The *additive menu size* of a mechanism μ is defined as the minimal size $|M_0|$ of an additive representation of its menu M. Since taking M_0 equal to M trivially yields an additive representation, the additive menu size can thus only be lower than its menu size. For separate selling of k goods, the additive menu size is at most k, rather than $2^k - 1$: the basic menu entries consist of selling each

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

[37] This is related to the fact that menu size is defined using the "normal" form of a mechanism—its menu—rather than its other, possibly simpler, descriptions.

[38] Our definition is just one of several possible definitions. Indeed, basic entries may be combined in other ways—such as taking the allocation probabilites to be independent (as in Briest et al. 2015), or adding them and then capping the sum at 1. What matters (see the proof of Proposition 8.1 below) is that any chosen basic entry should yield a nonnegative payoff (i.e., if (g, t) is part of the set chosen by type x then $g \cdot x - t \geq 0$); the variants mentioned above satisfy this.

good by itself at its price.[39]

The corresponding revenue is

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

Interestingly, the basic properties of the menu size, namely, that menu size 1 yields the bundling revenue, and that the increase in revenue is at most linear in the menu size (Proposition 3.1 in Section 3.2), hold for the additive menu size as well.

Proposition 8.1 *For every $k \geq 2$ and every k-good random valuation X,*

(i) $\text{REV}_{[1]*}(X) = \text{REV}_{[1]}(X) = \text{BREV}(X)$, *and*

(ii) $\text{REV}_{[m]}(X) \leq \text{REV}_{[m]*}(X) \leq m \cdot \text{BREV}(X)$ *for every $m \geq 1$.*

Proof. The only claim that is not immediate is the last inequality. Let M_0 with $|M_0| = m$ be a minimal additive representation of the menu. Let $(g, t) \in M_0$ be a basic menu entry. If the buyer with valuation x chooses (g, t) (i.e., (g, t) is part of the chosen subset $N \subseteq M_0$), then $g \cdot x - t \geq 0$ (otherwise, dropping it from the chosen subset—i.e., switching to $N \backslash \{(g, t)\}$, which yields an available menu entry—would strictly increase the buyer's payoff at x); hence $\sum_i x_i \geq g \cdot x \geq t$ (the first inequality is due to $(1, ..., 1) \geq g$ and $x \geq 0$). Therefore the total probability[40] π that (g, t) is chosen is at most $\mathbb{P}\left[\sum_{i=1}^k X_i \geq t\right]$, and so that part of the expected revenue that comes from (g, t), namely $t \cdot \pi$, is at most $t \cdot \mathbb{P}\left[\sum_{i=1}^k X_i \geq t\right] \leq \text{BREV}(X)$. This holds for each one of the m basic menu entries in M_0. ■

Proposition 8.1 thus implies that the results in this paper hold also for this more refined complexity measure; specifically, in each one of Theorems A–D one may replace MENU SIZE with MENU SIZE*. Moreover, by Theorem E, this measure captures well the complexity of selling the goods separately: its *additive* menu size is at most k.

[39] More precisely, it is the number of goods whose price is positive.

[40] The sum of these probabilities over all basic menu entries may be as high as m, as these events need not be disjoint (in contrast to standard menu items, where they are disjoint).

9 Bounded Valuations

In this section we deal with valuations in bounded domains, i.e., $[L, H]^k$ for $0 < L < H < \infty$. Since rescaling valuations by a constant factor of $1/L$ changes the range from $[L, H]^k$ to $[1, H/L]^k$ without affecting ratios of revenues, we take without loss of generality $L = 1$ and the range $[1, H]^k$. We first prove Theorem B: for two goods with valuations in $[1, H]^2$, 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 9.1 *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., for each x either $s(x) = 0$ or $s(x) \in [1, H]$) 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 $(1 - \varepsilon)s < \phi(s) < (1 - \varepsilon)(1 + \varepsilon^2)s$. 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 $\varepsilon j/K$ to each menu entry whose price is s^j. Notice that rounding down reduces each q by at most ε/K, and since higher-paying menu entries got a boost that is at least ε/K greater than any lower-paying menu entry, any buyer that previously chose an entry that pays s can now choose only an entry that pays 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 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 B. 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 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[41] 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 $2\varepsilon \cdot \mathbb{P}\left[s(X) < 2\varepsilon\right] \leq 2\varepsilon$. Let $\mu' = (q', s')$ denote the resulting mechanism; then the range of its nonzero payments is $[2\varepsilon, 2H]$. Applying Lemma 9.1 to μ' yields a new mechanism $\tilde{\mu} = (\tilde{q}, \tilde{s})$ with a menu of size

[41]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 such that either $s(y) = 0$ or $s(y) \geq 2\varepsilon$.

$O(\varepsilon^{-5}\log^2(H/\varepsilon))$, 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 7.5 (used for proving Theorem A(ii)) has $H = m^{O(m)}$, and the $\Omega(m^{1/7})$ gap between $\text{REV}(X)$ and $\text{REV}_{[1]}(X)$ implies that for, say, $m = O((\log H)^{1/8})$, we get $\text{REV}_{[m]}(X) = o(\text{REV}(X))$.

For more than two goods, i.e., $k > 2$, we obtain the somewhat weaker result that the menu size need only be polynomial in H.

Proposition 9.2 *For every $k \geq 2$ and $\varepsilon > 0$ there is $m_0 = (k/\varepsilon)^{O(k)}$ such that for every k-good random valuation X with values in $[0,1]^k$ and every $m \geq m_0$,*

$$\text{REV}_{[m]}(X) \geq \text{REV}(X) - \varepsilon.$$

This result is directly implied by the following lemma.

Lemma 9.3 *Let $m = (n+1)^k - 1$, where $n \geq 1$ is an integer. Then for every k-good random valuation X with values in $[0,1]^k$,*

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

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 9.2 we can derive an essentially equivalent multiplicative approximation result.

Proposition 9.4 *For every $k \geq 2$, $\varepsilon > 0$, and $H > 1$, there is $m_0 = (H/\varepsilon)^{O(k)}$ 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, by Proposition 9.2, 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 Two-Good Deterministic Mechanisms

Using the formula of Theorem 5.1 we can show that the Multiple of Basic revenue for two-good deterministic mechanisms equals precisely the $d_2 = 5/2$ bound of Proposition 7.3 (see Remark 7.4(c)).

Proposition A.1 *For $k = 2$ goods,*

$$\textsc{MoB}(\textsc{deterministic};\ 2\ \textit{goods}) = \frac{5}{2}.$$

Proof. We compute the supremum of $\textsc{MoB}(\mu)$, as given by Theorem 5.1, over all deterministic mechanisms μ. Such a mechanism is given by nonnegative prices p_1, p_2, and p_{12} for good 1, good 2, and the bundle, respectively (thus $b(x) = \max\{0, x_1 - p_1, x_2 - p_2, x_1 + x_2 - p_{12}\}$). Without loss of generality we assume that $p_1 \le p_2 \le p_{12}$; 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. We have four cases:

- If $p_1 > 0$ then $\mathrm{MoB}(\mu) = (p_1 - 0)/v(p_1) + (p_2 - p_1)/v(p_2) + (p_{12} - p_2)/v(p_{12})$. 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 $x_1 + x_2 - p_{12} \ge x_i - p_i$ for $i = 1, 2$, which implies $x_{3-i} \ge p_{12} - p_i \ge p_{12} - p_2$, and so $x_1 + x_2 \ge 2(p_{12} - p_2)$. Therefore $\mathrm{MoB}(\mu) \le 1 + 1 + 1/2 = 5/2$.

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

- If $p_1 = p_2 = 0 < p_{12}$ then $\mathrm{MoB}(\mu) \le 1$.

- If $p_1 = p_2 = p_{12} = 0$ then $\mathrm{MoB}(\mu) = 0$.

Thus $\mathrm{MoB}(\mu) \le 5/2$ in all cases; taking, say, $p_1 = 1$, $p_2 = H$, and $p_{12} = H^2$ for large[42] H shows that $\sup_\mu \mathrm{MoB}(\mu)$ over all deterministic mechanisms μ is indeed $5/2$. ■

For *separate-selling* mechanisms we have in addition $p_{12} = p_1 + p_2$, and then $v(p_1 + p_2) = p_1 + p_2$ (attained at $x = (p_1, p_2)$), and so $\mathrm{MoB}(\mu) = 1 + 1 - p_1/p_2 + p_1/(p_1 + p_2)$, which is less than 2, but can be made arbitrarily close to 2 by taking, say, $p_1 = 1$ and $p_2 = H$ for large H. This shows that MoB(SEPARATE; 2 goods) = 2; cf. Theorem E. For *symmetric* deterministic mechanisms we have $p_1 = p_2$, and so $\mathrm{MoB}(\mu) \le p_1/p_1 + (p_{12} - p_1)/(2(p_{12} - p_2)) = 3/2$, with equality for, say, $p_1 = p_2 = 1$ and $p_{12} = 2$ (which is in fact a symmetric separate-selling mechanism). Thus MoB(SYMMETRIC DETERMINISTIC; 2 goods) = MoB(SYMMETRIC SEPARATE; 2 goods) = 3/2.

[42]Alternatively, use the bound of Proposition 7.3 (see Remark 7.4(b)).

A.2 The Multiple of Separate Revenue (MoS)

Our MoB measure takes as basic revenue the bundling revenue, obtained by menu-size-1. We now consider using the separate revenue instead:

$$\text{MoS}(\mathcal{N};\mathbb{X}) := \sup_{x\in\mathbb{X}} \frac{\mathcal{N}\text{-Rev}(X)}{\text{SRev}(X)}$$

(MoS stands for "Multiple of Separate revenue").

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

Proposition A.2 *For every $k \geq 2$,*

$$\text{MoS}(\text{bundled};\ k \text{ goods}) \leq k.$$

Proof. Let $\text{BRev}(X)$ be achieved for a bundle price of p. If the separate auction offers each good at a 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; thus $\text{BRev}(X) \leq k{\cdot}\text{SRev}(X)$. ■

This is tight for $k = 2$.

Example A.3 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.

Proposition A.4 *There exists a constant $c < \infty$ such that for every $k \geq 2$ and every k-good random valuation X,*

$$\text{MoS}(\text{bundled};\ k \text{ goods}) \leq c \log k.$$

Proof. Let $\textsc{BRev}(X)$ 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 $\textsc{BRev}(X') = \textsc{BRev}(X)$, while $\textsc{SRev}(X') \leq \textsc{SRev}(X)$ because $X' \leq X$ everywhere.[43]) We now make another assumption without loss of generality, namely, that $\sum_i X_i = p$ (and so $\textsc{BRev}(X) = p$). (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.

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 $X_i \leq p/(2k)$, and, for each $j = 0, \ldots, \log_2 k$, a subrange where $p/(2^{j+1}) < X_i \leq p/(2^j)$ (notice that since $X_i \leq p$ we have covered the whole support of X_i). The low subrange contributes at most $p/(2k)$ to the expectation of X_i, and thus one of the other $1 + \log_2 k$ 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 $r_i := \textsc{Rev}(X_i) = \sup_{t>0} t \cdot (1 - F_i(t))$ (where F_i denotes the cumulative distribution function of X_i); then $1 - F_i(t) \leq r_i/t$ and so (recall that $X_i \leq p$ because $\sum_i X_i = p$)

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

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

[43] See footnote 36 above.

yields

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

where $s := \sum_i r_i$. Thus $p/s \leq 1 + \ln(p/s) + \ln k$, from which it follows that[44] $\text{BRev}(X)/\text{SRev}(X) = p/s < 4 \ln k$. ■

Corollary A.5 *There exists a constant $c < \infty$ such that for every $k \geq 2$,*

$$\text{MoS}(\text{deterministic}; \ k \ \textit{goods}) \leq c 2^k \log k.$$

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

Proposition A.6 *Let $k = 2$. Then*

$$\text{MoS}(\text{deterministic}; \ 2 \ \textit{goods}) \leq 3.$$

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 A.2. ■

We now study MoS; the analysis is analogous to the one carried out with respect to the bundling revenue in Sections 5–7, but we now use the maximum norm $||x||_\infty = \max_i |x_i|$ instead of the 1-norm.

We have

Theorem A.7 *Let $\mu = (q, s)$ be a k-good mechanism. Then*

$$\frac{1}{k}\int_0^\infty \frac{1}{w(t)} \, \mathrm{d}t \leq \text{MoS}(\mu) \leq \int_0^\infty \frac{1}{w(t)} \, \mathrm{d}t,$$

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

where for every $t > 0$ we define

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

Unlike Theorem 5.1, here we do not get a sharp formula for MoS, but only an integral that is within a factor of k from it (see Remark A.8(b) below).

Proof. Let $\gamma := \int_0^\infty 1/w(t)\ \mathrm{d}t$.

First, for every $t > 0$ we have

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

Integrating over t yields $R(\mu, X) \leq \gamma \cdot \mathrm{SRev}(X)$, proving that $\mathrm{MoS}(\mu) \leq \gamma$.

Second, we show that for every $\gamma' < \gamma$ there exists a k-good random valuation X such that $0 < \mathrm{SRev}(X) < \infty$ and $R(\mu; X)/\mathrm{SRev}(X) > \gamma'/k$. Let $0 = t_0 < t_1 < ... < t_N < t_{N+1} = \infty$ with $0 = w(t_0) < w(t_1) < w(t_2) < ... < w(t_N) < w(t_{N+1}) = \infty$ be such that

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

Let $\varepsilon > 0$ be small enough so that $\gamma'' > (1+\varepsilon)\gamma'$ and $w(t_{n+1}) > (1+\varepsilon)w(t_n)$ for all $1 \leq n \leq N$, and choose for each $1 \leq n \leq N$ 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

$$\sum_{n=1}^{N} \frac{t_n - t_{n-1}}{||x_n||_1} > \sum_{n=1}^{N} \frac{t_n - t_{n-1}}{w(t_n)(1+\varepsilon)} = \frac{\gamma''}{1+\varepsilon} > \gamma'. \tag{21}$$

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

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

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

(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 $\text{SRev}(X) \leq k\xi_1$ (which is finite; also $\text{SRev}(X) > 0$ because X does not vanish).

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

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

(recall (21)). ■

Remark A.8 *(a)* As in Theorem 5.1 (see Remark 5.3 following its proof), the random valuation X in the second part of the proof 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) The gap of k in Theorem A.7 is correct. Take two goods. For μ that sells the bundle at 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 A.3 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[45] $\gamma(\mu) = 2$, but $R(\mu; X)/\text{SRev}(X) \leq 1 = \gamma(\mu)/k$ for any X (with equality for, say, the constant valuation $(1/2, 1/2)$).

(c) Recalling the definition of $v(t)$ in Theorem 5.1, we have $1/v(t) \leq 1/w(t) \leq k/v(t)$ for every t (because $||x||_1 \geq ||x||_\infty \geq ||x||_1/k$), and so for every mechanism μ we have $\text{MoB}(\mu) \leq \int 1/w(t)\, dt \leq k \cdot \text{MoB}(\mu)$.

(d) We can take as benchmark the maximum of the two one-dimensional

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

mechanisms, bundled and separate (cf. Babaioff et al. 2014). Thus, putting

$$\text{MoBS}(\mu) := \sup_X \frac{R(\mu; X)}{\max\{\text{BRev}(X), \text{SRev}(X)\}},$$

we have

$$\frac{1}{k}\int_0^\infty \frac{1}{w(t)}\,\mathrm{d}t \le \text{MoBS}(\mu) \le \int_0^\infty \frac{1}{v(t)}\,\mathrm{d}t. \tag{22}$$

Indeed, in the second part of the proof of Theorem A.7 above, for every $u \in (k\xi_{n-1}, k\xi_n]$,

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

(because $X = x_j$ for some $j \le n-1$ implies $\sum_i X_i \le k||x_j||_\infty \le k||x_{n-1}||_\infty = k\xi_{n-1} < u$), and so $\text{BRev}(X) = \sup_{u>0} u \cdot \mathbb{P}[X_i \ge u] \le k\xi_1$ as well, which yields the first inequality in (22). For the second inequality we use $\text{MoBS}(\mu) \le \text{MoB}(\mu) \le \int 1/v$ (which, by (c) above, yields a better inequality than $\text{MoBS}(\mu) \le \text{MoS}(\mu) \le \int 1/w$).

The analogous result to the construction of Section 7 is

Proposition A.9 *Let $(g_n)_{n=0}^N$ be a finite or countably infinite sequence in $[0,1]^k$ starting with $g_0 = (0, ..., 0)$, and let $(y_n)_{n=1}^N$ 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 exist a sequence $(t_n)_{n=1}^N$ of positive real numbers, a k-good mechanism μ with menu $\{(g_n, t_n)\}_{n=1}^N$, and a k-good random valuation X with $0 < \text{BRev}(X) < \infty$, such that

$$\text{MoS}(X) \ge \text{MoBS}(X) > (1-\varepsilon)\frac{1}{k}\sum_{n=1}^N \frac{\text{gap}_n}{||y_n||_\infty}.$$

The proof is omitted, as it is identical to that of Proposition 7.1, except that it uses throughout the ∞-norm instead of the 1-norm (and the construction of the appropriate random valuation is as in Theorem A.7 and Remark A.8(d) above instead of Theorem 5.1).

As a consequence, for deterministic mechanisms we get (see Corollary A.5 for the opposite inequality):

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

$$\textsc{MoS}(\textsc{deterministic};\ k\ \textit{goods}) \geq \textsc{MoSB}(\textsc{deterministic};\ k\ \textit{goods}) \geq \frac{2^k - 1}{k}.$$

Proof. We proceed exactly as in the proof of Proposition 7.3, but now we have $||y_n||_\infty = 1$, and so we get

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

■

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

$$\begin{aligned} \frac{3}{2} &\leq \textsc{MoS}(\textsc{deterministic};\ 2\ \textit{goods}) \leq 3 \\ \frac{3}{2} &\leq \textsc{MoSB}(\textsc{deterministic};\ 2\ \textit{goods}) \leq \frac{5}{2} \end{aligned}$$

(cf. Proposition A.1, which shows that MoB is exactly 5/2).

A.3 The Unit-Demand Model

In this section we briefly compare our model to the unit-demand model that is considered in many papers. 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 requires 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., the valuation of a set $I \subseteq \{1, \ldots, k\}$ of goods 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" at 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 at the same price as in the original menu entry.

Let us use the notation $\textsc{Rev}^{UD}(X)$ to denote the revenue obtainable from a unit-demand buyer with a k-good random valuation X. Similarly $\textsc{DRev}^{UD}(X)$ 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 A.11 *For every $k \geq 2$ and every k-good random valuation X,*

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

(ii) $\textsc{DRev}^{UD}(X) \leq \textsc{DRev}(X) \leq k 2^k \cdot \textsc{DRev}^{UD}(X)$.

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 for 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 $0 \le g_i \le 1$ 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 $g_1/k + ... + g_k/k \le 1$), and for these 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 at 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 at 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 A.12 *For every $k \ge 2$,*

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

Proof. For every X we have $\mathrm{DRev}^{UD}(X) \le \mathrm{SRev}(X)$ 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 A.10 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 More Than One Buyer

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 in more general settings involving multiple buyers. An immediate observation is that since our main results (Theorems A, C, and D) 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 multiple-buyer settings is thus relevant to the upper bounds in the paper, both the significant ones (Propositions 9.2 and A.4) and the simple ones (Proposition 3.1). In this appendix we discuss why these can all be extended to the multiple-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 $|\{q(x) : x \in \mathbb{R}^k_+\}\setminus\{(0,\ldots,0)\}|$. 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 a *menu size* of at most m if for every buyer $j = 1, \ldots, n$ and every $(n-1)$-tuple of (direct) bids of the other buyers[46] $x^{-j} \in (\mathbb{R}^k_+)^{n-1}$, 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 that if the original mechanism was incentive compatible in dominant strategies then the mechanism induced on player j by x^{-j} is also incentive compatible. However, 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[47] 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 mechanisms 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

[46] Superscripts are used here for the buyers.

[47] 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.

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 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 3.1(i) for the case of n buyers:[48]

$$\mathrm{BR\textsc{ev}}^n(X) \le \mathrm{R\textsc{ev}}^n_{[1]}(X) \le n \cdot \mathrm{BR\textsc{ev}}^n(X).$$

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 H^j with probability H^{-j}, and zero otherwise (independently over buyers), for a large enough but fixed H.

A similar argument that focuses on the single buyer that provides the largest fraction of revenue yields the generalization of Proposition 3.1(iii) and (iv):

$$\begin{aligned} \mathrm{R\textsc{ev}}^n_{[m]}(X) &\le n \cdot m \cdot \mathrm{BR\textsc{ev}}^n(X) \text{ and} \\ \mathrm{DR\textsc{ev}}^n(X) &\le n \cdot (2^k - 1) \cdot \mathrm{BR\textsc{ev}}^n(X). \end{aligned}$$

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 D 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 $H^j x$ with probability $H^{-j} p$ (independently over the buyers).

Versions of Propositions 9.2 and A.4 that incur a linear loss in n are also easily implied, but do not seem to be interesting. It would seem that in

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

both cases sharper results, in which the additional loss due to the number of buyers is avoided, might be obtained.

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*, 869–877.

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.

Babaioff, M., N. Nisan, and A. Rubinstein (2018), "Optimal Deterministic Mechanisms for an Additive Buyer," *EC 2018: Proceedings of the 19th ACM Conference on Economics and Computation*, 429.

Balcan, M. F., A. Blum, J. D. Hartline, and Y. Mansour (2008), "Reducing Mechanism Design to Algorithm Design via Machine Learning," *Journal of Computer and System Sciences* 74, 1245–1270.

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*.

Carroll, G. (2017), "Robustness and Separation in Multidimensional Screening," *Econometrica* 85, 453–488.

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.*

Dobsinski, S. (2011), "An Impossibility Result for Truthful Combinatorial Auctions with Submodular Valuations," *STOC 2011: Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, 139–148.

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.

Goldner, K. and Y. A. Gonczarowski (2018), "The Menu Size of Precise and Approximate Revenue-Maximizing Auctions: A Tutorial at EC 2018."

Hart, S. and N. Nisan (2013), "The Menu-Size Complexity of Aunctions," *arXiv* 1304.6116; *EC 2013: Proceedings of the 14th ACM Conference on Electronic Commerce,* 565–566 (early version of the present paper).

Hart, S. and N. Nisan (2017), "Approximate Revenue Maximization with Multiple Items," *Journal of Economic Theory* 172, 313–347 (early version: *arXiv* 1204.1846; *EC 2012: Proceedings of the 13th ACM Conference on Electronic Commerce,* 656).

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 (2017), "The Better Half of Selling Separately," Hebrew University, Center for Rationality DP-712 (December 2017); *arXiv* 1712.08973.

Hartline, J. and V. Koltun (2005), "Near Optimal Pricing in Near-Optimal Time," *WADS 2005: Workshop on Algorithms and Data Structures*, 422–431.

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.