The Menu-Size Complexity of Revenue Approximation

Moshe Babaioff[*] Yannai A. Gonczarowski[†] Noam Nisan[‡]

March 27, 2016

Abstract

We consider a monopolist that is selling n items to a single additive buyer, where the buyer's values for the items are drawn according to independent distributions $F_1, F_2, \ldots, F_n$ that possibly have unbounded support. It is well known that — unlike in the single item case — the revenue-optimal auction (a pricing scheme) may be complex, sometimes requiring a continuum of menu entries. It is also known that simple auctions can extract a constant fraction of the optimal revenue. Nonetheless, the question of the possibility to extract an arbitrarily high fraction of the optimal revenue via a finite menu size remained open.

In this paper, we give an affirmative answer to this open question, showing that for every n and for every $\varepsilon > 0$, there exists a complexity bound $C = C(n, \varepsilon)$ such that auctions of menu size at most C suffice for obtaining a $(1 - \varepsilon)$ fraction of the optimal revenue. We prove upper and lower bounds on the revenue approximation complexity $C(n, \varepsilon)$.

1 Introduction

As familiar economic institutions are moving to computerized platforms, they are are gaining unprecedented sizes and complexities. These new complexities often become the defining feature of the computerized economic scenario, like in the case of spectrum auctions or of ad auctions. The use of the word "complexity" here was intentionally vague, and can refer to a wide variety of computational, informational, or descriptive measures of complexity. A high-level goal of the field of "Economics and Computation" is analyzing such complexities and understanding the degree to which they are indeed a bottleneck to achieving desired economic properties.

This paper studies exactly such a question in the recently well-studied scenario of pricing multiple items. The scenario is that of a monopolist seller that is selling n items to a single additive buyer. The buyer has a private value v_i for each item i, where each v_i is distributed according to a commonly known prior distribution F_i, independently of the values of the other items. The valuation of the buyer is assumed to be additive, so that her value for a subset S of the items is simply $\sum_{i \in S} v_i$, and the seller's goal is to design an "auction" (really just a a pricing scheme) that maximizes her revenue. The classical economic analysis (Myerson, 1981) shows that for a single item, the optimal auction is to simply sell the item at some fixed price. On the other hand, when there is more than a single item, it is known that the optimal auction may be surprisingly complex, randomized, and unintuitive (McAfee and McMillan, 1988; Thanassoulis, 2004; Manelli

[*]Microsoft Research, *E-mail*: moshe@microsoft.com.

[†]Einstein Institute of Mathematics, Rachel & Selim Benin School of Computer Science & Engineering and Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, Israel; and Microsoft Research, *E-mail*: yannai@gonch.name.

[‡]Rachel & Selim Benin School of Computer Science & Engineering and Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, Israel; and Microsoft Research, *E-mail*: noam.nisan@gmail.com.

and Vincent, 2006; Giannakopoulos and Koutsoupias, 2014, 2015; Hart and Reny, 2015; Daskalakis et al., 2013).

A significant amount of recent work has studied whether "simple" auctions may at least yield approximately-optimal revenue. Following a sequence of results (Chawla et al., 2010; Hart and Nisan, 2012; Li and Yao, 2013), it was shown by Babaioff et al. (2014) that one of the following two "simple" auctions always yields at least a constant fraction ($1/6$) of the optimal revenue: either sell all items as a single take-it-or-leave-it bundle (for some carefully chosen price) or sell each item separately for its Myerson price. This was further extended (with different constants) to the case of multiple buyers (Yao, 2015) and to buyers with sub-additive valuations (Rubinstein and Weinberg, 2015), but is in contrast to the case where the item values come from a joint (correlated) distribution, a case for which no finite approximation is possible by finite auctions (Briest et al., 2010; Hart and Nisan, 2013).

In this work, we study the tradeoff between the complexity of an auction and the extent to which it can approximate the optimal revenue. One may choose various measures of auction complexity (Hart and Nisan, 2013; Dughmi et al., 2014; Morgenstern and Roughgarden, 2015), and we will focus on the simplest one, the *menu size* suggested in Hart and Nisan (2013). The menu size of an auction (for a single buyer, an auction is just a pricing scheme) is defined to be the number of possible different outcomes of the auction. More specifically, every single-buyer auction is equivalent to one that offers a *menu* of options to the buyer, where each option — *entry* $(\vec{x};p)$ — in the menu specifies a probability x_i of acquiring each item i as well as a price p to be paid for the combination $\vec{x}$, and where the buyer chooses an entry that maximizes her own utility $\sum_i x_i \cdot v_i - p$. The number of entries in the menu is defined to be the *menu size* complexity of the auction. It is known that for some distributions, the optimal auction has infinite menu size (Daskalakis et al., 2013), and that a constant fraction of the optimal revenue may be extracted by a finite-complexity auction (Babaioff et al., 2014).

Our first and main result shows that, in fact, finite complexity suffices to get *arbitrarily close* to the optimal revenue.

Definition 1.1 ($\mathcal{R}ev_C$; $\mathcal{R}ev$)**.** For a distribution F on n items, we denote by $\mathcal{R}ev_C(F)$ the maximal (formally, the supremum) revenue obtainable by an individually rational incentive-compatible auction that has at most C menu entries that sells the n items to a single additive buyer whose values for the items are distributed according to F. We also denote $\mathcal{R}ev(F) = \mathcal{R}ev_\infty(F)$ to be the maximal revenue obtainable without any complexity restrictions on the auction.

Formally, our result shows that $\lim_{C\to\infty} \frac{\mathcal{R}ev_C(F)}{\mathcal{R}ev(F)} = 1$, *uniformly* across all product distributions F. In other words:

Theorem 1.1. *For every number of items n and every $\varepsilon > 0$, there exists $C = C(n, \varepsilon) \in \mathbb{R}_+$ such that for every $F_1, F_2, \ldots, F_n \in \Delta(\mathbb{R}_+)$, we have that $\mathcal{R}ev_C(F_1\times\cdots\times F_n) \geq (1-\varepsilon)\cdot\mathcal{R}ev(F_1\times\cdots\times F_n)$.*

Theorem 1.1 gives a positive answer to Open Problem 6 from Hart and Nisan (2014)[1], which asks precisely whether the statement of Theorem 1.1 holds. It is natural to ask what is the rate of the uniform convergence of the sequence $\frac{\mathcal{R}ev_C(F)}{\mathcal{R}ev(F)}$. In other words, *how complex* must a revenue-approximating auction be?

Definition 1.2 (Revenue Approximation Complexity)**.** For every number of items n and every $\varepsilon > 0$, we define the *revenue approximation complexity* $C(n, \varepsilon) \in \mathbb{R}_+$ to be the smallest value $C \in \mathbb{R}_+$ such that $\mathcal{R}ev_C(F_1 \times \cdots \times F_n) \geq (1-\varepsilon) \cdot \mathcal{R}ev(F_1 \times \cdots \times F_n)$ for every $F_1, F_2, \ldots, F_n \in \Delta(\mathbb{R}_+)$.

[1] Hart and Nisan (2014) is a manuscript combining Hart and Nisan (2012) and Hart and Nisan (2013).

The construction used in the proof of Theorem 1.1 gives an upper bound on $C(n, \varepsilon)$ (i.e., a lower bound on the rate of uniform convergence of $\frac{\mathcal{R}ev_C(F)}{\mathcal{R}ev(F)}$).

Theorem 1.1 (Quantitative Version[2]). $C(n, \varepsilon) \leq (n/\varepsilon)^{O(n)}$.

This bound is exponential in n, and so the next natural question is whether polynomial size suffices. At first sight the answer seems to be obviously "no": the menu-size complexity measure is quite weak, and even the auction that sells each item separately has exponential menu size (since a menu entry is needed for each possible subset of the items). This, however, is premature; in fact, we show that polynomial menu size turns out to suffice for approximating the revenue obtainable from selling items separately. Let us denote the revenue obtainable by selling each item separately for its optimal price by $\mathcal{S}Rev(F_1 \times \cdots \times F_n)$.

Theorem 1.2. *For every $\varepsilon > 0$, there exists $d = d(\varepsilon)$ such that for every number of items n and $F_1, F_2, \ldots, F_n \in \Delta(\mathbb{R}_+)$, we have for $C = n^d$ that $\mathcal{R}ev_C(F_1 \times \cdots \times F_n) \geq (1-\varepsilon) \cdot \mathcal{S}Rev(F_1 \times \cdots \times F_n)$.*

The same bound applies also to the revenue obtainable by selling arbitrary prepartitioned collection bundles. Using the result of Babaioff et al. (2014), this immediately implies that polynomial menu size suffices for extracting a constant fraction of the *optimal* revenue

Corollary 1.3. *There exist fixed constants d and $\alpha > 0$ such that $C(n, 1-\alpha) \leq O(n^d)$.*

The above reasoning shows that Corollary 1.3 holds for every α arbitrarily close to $1/6$, which is the constant fraction of the optimal revenue shown by Babaioff et al. (2014) to be obtainable by the better of bundled selling and separate selling.

Does polynomial menu size suffice for extracting revenue *arbitrarily close* to the optimal revenue? We prove that this is not the case, at least for ε that is polynomially small in n.

Theorem 1.4. $C(n, 1/n) \geq 2^{\Omega(n)}$.

This proof shows, in fact, that polynomial dependence on ε is impossible even for approximating the revenue from selling the items separately. At this point, we leave two main open problems. The first one asks whether for every *fixed* $\varepsilon > 0$, polynomial menu size suffices for approximating the optimal revenue. We conjecture a positive answer, as is the case for approximating the revenue from selling item separately.

Conjecture 1.5. *For every $\varepsilon > 0$, there exists $d = d(\varepsilon)$ such that $C(n, \varepsilon) \leq O(n^d)$.*

The second question asks whether stronger notions of auction complexity may allow for better revenue in polynomial complexity. This may be asked for any complexity measure, but for concreteness let us focus on the *additive menu size* complexity defined by Hart and Nisan (2013), which allows presenting menus from which the buyer is allowed to take *any combination* of menu entries for the sum of their prices. The lower bounds shown by Hart and Nisan (2013) and Dughmi et al. (2014), which are stated for the case of correlated item values, apply also to additive menu size. We are not able to prove any lower bounds in our case for this additive menu-size complexity, nor are we able to obtain any improvement in our upper bounds; doing one of these is our second open problem.

[2]Some more careful analysis can in fact show that $C(n, \varepsilon) \leq (\log n/\varepsilon)^{O(n)}$. See Section 2.7 for more details.

2 Upper Bound on Revenue-Approximation Complexity

In this section, we prove Theorem 1.1, which states that $C(n, \varepsilon)$ is finite for every number of items $n \in \mathbb{N}$ and $\varepsilon > 0$, and moreover, that $C(n, \varepsilon) \leq (n/\varepsilon)^{O(n)}$. The proof proceeds in four steps. Section 2.1 provides a rough overview of the proof strategy, Sections 2.3 through 2.6 provide the full details of each of the four steps of the proof, and Section 2.7 connects the dots by providing a proof of the theorem by combining the four steps. Finally, Section 2.8 concludes with a short discussion of the application of the proof steps to obtain uniform approximation results for correlated distributions over a restricted valuation space that generalize bounded distributions.

2.1 Proof Overview

Let $F_1, F_2, \ldots, F_n$ be the respective distributions of the values of the n items. We will construct an auction with (finite) menu size $(n/\varepsilon)^{O(n)}$ that guarantees a $(1 - \varepsilon)$ multiplicative approximation to the optimal revenue.

One possible approach, used by Li and Yao (2013), Babaioff et al. (2014) and Rubinstein and Weinberg (2015) to guarantee a constant fraction of the optimal revenue, is using a core/tail decomposition, and bounding the revenue from (or the welfare of) the core and the revenue from the tail. Unfortunately, such a decomposition inherently entails a nonnegligible loss, either due to bounding the welfare of the core instead of the revenue from it, or due to estimating the total revenue using the revenues obtained by selling to the core and to the tail separately. Therefore, we take a more subtle approach, handling core and tail regions together: We first show that one does not lose much revenue by disregarding what can be described as "second-order" tails — valuations where two or more of the item prices lie in the tail, and then, starting with an optimal auction (which may be arbitrarily complex, even infinite in size) for the valuation space consisting of the core plus all first-order tails, we show how to gradually simplify this auction while losing only a tiny fraction of the revenue in each step, where for every modification that we perform to the menu, we must simultaneously check that we do not significantly hurt the revenue from neither core nor (first-order) tail buyers.

Our first step, taken in Section 2.3, simplifies the valuation space by showing that since item prices are independent, finding an approximately optimal auction under the assumption that at most one item has price higher than H (i.e., has price that lies in the H-tail), for some $H = \mathrm{poly}(n, 1/\varepsilon)$, entails a very small loss compared to doing so without this assumption. This is possible, very roughly speaking, because the probability of two item prices lying in the tail can be thought of as being of order ε^2, while the revenue conditioned upon being in this "second-order" tail (i.e., conditioned upon the prices of both of these items lying in the H-tail) is of order $1/\varepsilon$. Therefore, it is enough to construct our finite approximation for the distribution conditioned upon being in the valuation space comprised of the core and first-order tail, i.e., the valuation space where at most one item price lies in the tail; we call distributions over this valuation space *exclusively unbounded* distributions. We note that this is the only step in which the independence of the item prices is used; indeed, combining the remaining steps shows that the revenue from all exclusive unbounded distributions (even highly correlated distributions not originating from a product distribution over $\mathbb{R}_+^n$) can be uniformly approximated using finite-size menus (see Proposition 2.6 in Section 2.8).

The second step, taken in Section 2.4, starts with an optimal (possibly arbitrarily complex) revenue-maximizing auction for some exclusively unbounded distribution. In this step, we simplify the "expensive" part of the menu, i.e., the part of the menu consisting of all menu entries that cost more than E, for some $E = \mathrm{poly}(n, 1/\varepsilon)$, so that each expensive menu entry can only allocate a single item with nonzero probability. This means that while in the "cheap" part of the menu we can allocate arbitrary combinations of items, once the price increases beyond E, our auction must act like a unit-demand one and never allocate more than a single item. We call such an auction E-*exclusive*. This is possible since, roughly speaking, due to the assumption of exclusive unboundedness, most of the value from an expensive menu entry chosen by some buyer type comes only from the unique item whose price lies in the tail for the valuation of that buyer type. So, instead of offering that (nonexclusive) menu entry, we offer an (exclusive) entry with only the corresponding winning probability of that item, for a slightly discounted price. While this step in fact in most natural cases increases the size of the expensive part of the menu (as each expensive menu entry possibly becomes n exclusive menu entries, each allocating a distinct item with nonzero probability), this simplification allows the next step to significantly reduce the size of this part of the menu.

The third step, taken in Section 2.5, reduces the size of the expensive part of the menu to at most $2n$ menu entries. This is the most technically elaborate of all steps. Since E-exclusivity means that the expensive auction entries "look like" separate auctions for each of the n items, we show that we are able to carefully use the analysis of Myerson (1981) to replace each of these separate expensive auctions with a simple "almost-deterministic" one. In contrast with Myerson's analysis, we require two menu entries (and not only one) for each item: one (deterministic) analogous to Myerson's "optimal price" entry, and the other (randomized) — to the "opt out" zero entry in Myerson's auction. The function of the latter entry is to make sure that buyers are not incentivized to "jump" from the expensive part to the cheap part of the menu following the reduction of the size of the former.

The fourth and final step, taken in Section 2.6, is that of simplifying the cheap part of the menu by "rounding" the menu entries into a discrete set. This is not totally trivial, as slight changes in allocation probabilities may result in large revenue changes (indeed, the valuation space is not bounded but only exclusively unbounded), but the "nudge and round" techniques that allowed this rounding to be done with only negligible loss of revenue for bounded valuations (Daskalakis and Weinberg, 2012; Hart and Nisan, 2013; Dughmi et al., 2014) can be carefully extended to be used here as well: Roughly speaking, we construct n discretizations of each cheap menu entry, where each discretization rounds the price and all but one allocation; rounding in the right direction guarantees that at least one of these discretizations is still a leading candidate for any buyer type that previously chose the corresponding original (nondiscretized) menu entry. As all but one coordinate of each of the discretized menu entries lie on a grid, only finitely many of them are in fact chosen by any buyer type.

While the second and third (and first) steps each entail a slight multiplicative revenue drop, the fourth step entails also a slight additive revenue drop. To altogether obtain only a slight multiplicative drop (with no additional additive drop), when connecting all steps, in Section 2.7, we assume w.l.o.g. that $\mathrm{Max}_i\, \mathcal{R}ev(F_i)$ is normalized[3] (by scaling the currency) to a suitable value such that the additive drop in the fourth step can be quantified to be less than a slight multiplicative drop. Clearly, as the overall cumulative revenue drop that we show is purely multiplicative, the proof also implies the same multiplicative drop for all (even nonnormalized) distributions.

2.2 Preliminaries

Definition 2.1 (Notation)**.**

[3]The cases in which $\mathrm{Max}_i\, \mathcal{R}ev(F_i)$ cannot be normalized, i.e., when it is 0 or infinite, are easy to handle separately; the former as there is nothing to show, and the latter as it implies that $\mathcal{R}ev(F_i) = \infty$ for some $i \in [n]$, and so by the theorem of Myerson (1981), arbitrarily high revenue can be extracted using a take-it-or-leave-it offer for item i.

- (Naturals). We denote the strictly positive natural numbers by $\mathbb{N} \triangleq \{1,2,3,\ldots\}$.
- ([n]). For every $n \in \mathbb{N}$, we define $[n] \triangleq \{1,2,\ldots,n\}$.
- (Nonnegative Reals). We denote the nonnegative reals by $\mathbb{R}_+ \triangleq \{r \in \mathbb{R} \mid r \geq 0\}$.

Definition 2.2 (Outcome; Type; Utility). Let $n \in \mathbb{N}$ be a number of items.

a. An *outcome* is an $(n+1)$-tuple $(\vec{x};p) = (x_1,x_2,\ldots,x_n;p) \in [0,1]^n \times \mathbb{R}_+$, denoting an allocation (to the buyer) of every item $i \in [n]$ with probability x_i, for a total price (paid by the buyer) of p.

b. We denote the (expected) *utility* of a (risk-neutral additive) buyer with *type* (respective item valuations) $v = (v_1,v_2,\ldots,v_n) \in \mathbb{R}_+^n$ from an outcome $e = (\vec{x};p) \in [0,1]^n \times \mathbb{R}_+$ by

$$u_e(v) \triangleq \sum_{i=1}^{n} x_i \cdot v_i - p.$$

Definition 2.3 (IC Auction as Menu). Let $n \in \mathbb{N}$ be a number of items. By the taxation principle, we identify any *incentive-compatible (IC) n-item auction* with a (possibly infinitely large) menu of outcomes (the *entries* in the menu are all the possible outcomes of the auction), where by IC the buyer chooses an entry that maximizes her utility.[4] If the auction is *individually rational (IR)*, then we assume w.l.o.g. that the menu includes the entry $(\vec{0};0)$ that allocates no item and costs nothing. (Conversely, if the menu includes the entry $(\vec{0};0)$, then the auction is IR.) Following Hart and Nisan (2013), we define the *menu size* of an IC and IR auction as the number of entries, other than $(\vec{0};0)$, in the menu of that auction.

Definition 2.4 ($\mathcal{R}ev_{\mathcal{M}}$; $\mathcal{R}ev_C$; $\mathcal{R}ev$). Let $n \in \mathbb{N}$ be a number of items and let $F \in \Delta(\mathbb{R}_+^n)$ be a distribution over $\mathbb{R}_+^n$.

a. Given an IC and IR n-item auction $\mathcal{M}$, we denote the (expected) revenue obtainable by $\mathcal{M}$ from (a single risk-neutral additive buyer with type distributed according to) F, by

$$\mathcal{R}ev_{\mathcal{M}}(F) \triangleq \mathbb{E}_{v\sim F}\big[p(v)\big],$$

where $p(v)$ is the price of the entry from $\mathcal{M}$ that maximizes the utility of v, with ties broken in favor of higher prices.[5,6]

b. Given $C \in \mathbb{N}$, we denote the highest revenue (more accurately, the supremum of the revenues) obtainable from F by an IC and IR n-item auction with at most C menu entries by

$$\mathcal{R}ev_C(F) \triangleq \sup_{\substack{\mathcal{M}\subseteq[0,1]^n\times\mathbb{R}_+: \\ |\mathcal{M}|\leq C}} \mathcal{R}ev_{\mathcal{M}}(F).$$

[4] If the menu is infinite, then the fact that it corresponds to an IC auction guarantees that some menu entry maximizes the utility of each buyer type. See Appendix A.1.1 for more details.

[5] The results of this paper hold regardless of the tie-breaking rule chosen. See Appendix A.2 for more details.

[6] If the menu is infinite, then the fact that a utility-maximizing menu entry exists for every buyer type does not guarantee that a utility-maximizing entry *with maximal price* (among all utility-maximizing entries) exists for every buyer type. (I.e., it is not guaranteed that the supremum price over all utility-maximizing entries is attained as a maximum.) Indeed, to be completely general, a more subtle definition of the revenue obtainable by an IC auction would have been needed. Nonetheless, for the auctions considered in this paper, this subtle definition is not required as we make sure that they all possess, for each buyer type, a utility-maximizing entry with maximal price. See Appendix A.1.2 for more details.

c. We denote the highest revenue (more accurately, the supremum of the revenues) obtainable from F by an IC and IR n-item auction by

$$\mathcal{R}ev(F) \triangleq \sup_{\mathcal{M}\subseteq[0,1]^n\times\mathbb{R}_+} \mathcal{R}ev_{\mathcal{M}}(F).$$

Theorem 2.1 (Hart and Nisan (2012))**.** *$\mathcal{R}ev(F\times G) \leq 2\cdot\big(\mathcal{R}ev(F)+\mathcal{R}ev(G)\big)$, for every $m,n \in \mathbb{N}$, $F \in \Delta(\mathbb{R}_+^m)$, and $G \in \Delta(\mathbb{R}_+^n)$.*

2.3 At Most One High Price

As outlined above, our first step toward proving Theorem 1.1, which we take in this section, simplifies the valuation space by showing that since item prices are independent, any auction that extracts most of the revenue under the assumption that the valuation space is restricted to some H-*exclusively unbounded* valuation space, i.e., to a valuation space where for each buyer type at most one item has price higher than some $H = \text{poly}(n, 1/\varepsilon)$, also extracts most of the revenue without this assumption. This step is formalized by Lemma 2.2.

Definition 2.5 ($\mathcal{EU}_H^n$; Exclusively Unbounded Type Distribution)**.** Let $H \in \mathbb{R}_+$ and $n \in \mathbb{N}$.

a. We denote the subset of $\mathbb{R}_+^n$ where at most one coordinate is greater than H by

$$\mathcal{EU}_H^n \triangleq \big\{(v_1, v_2, \ldots, v_n) \in \mathbb{R}_+^n \;\big|\; |\{i \in [n] \mid v_i > H\}| \leq 1\big\}.$$

b. We say that a type distribution $F \in \Delta(\mathbb{R}_+^n)$ is *(H-)exclusively unbounded* if $\text{supp}(F) \subseteq \mathcal{EU}_H^n$.

Definition 2.6 ($F|_A$)**.** For a set A and a distribution F defined over some superset of A s.t. $F(A) > 0$, we denote the conditional distribution of $v \sim F$ conditioned upon $v \in A$ by $F|_A$. Formally, for every measurable set $B \subseteq A$, we define $F|_A(B) \triangleq \frac{F(B)}{F(A)}$.

Lemma 2.2. *Let $n \in \mathbb{N}$ s.t. $n \geq 2$, let $R \in \mathbb{R}_+$, let $\varepsilon \in (0,1)$, and let $H \geq \frac{2\cdot n\cdot(n-1)\cdot R}{\varepsilon}$. For every $F = F_1 \times F_2 \times \cdots \times F_n \in \Delta(\mathbb{R}_+)^n$ s.t. $\text{Max}_{i\in[n]} \mathcal{R}ev(F_i) \leq R$, all of the following hold.*

a. $F(\mathcal{EU}_H^n) > 0$ (hence the exclusively unbounded conditioned distribution $F|_{\mathcal{EU}_H^n}$ is well defined).

b. $\mathcal{R}ev(F|_{\mathcal{EU}_H^n}) \geq (1-\varepsilon)\cdot\mathcal{R}ev(F)$.

c. For every $a \in (0,1]$ and for every IC and IR n-item auction $\mathcal{M}$, if $\mathcal{R}ev_{\mathcal{M}}(F|_{\mathcal{EU}_H^n}) \geq a\cdot\mathcal{R}ev(F|_{\mathcal{EU}_H^n})$, then $\mathcal{R}ev_{\mathcal{M}}(F) \geq (1-\varepsilon)\cdot a\cdot\mathcal{R}ev(F)$.

Proof. For every $i \in [n]$, we denote the probability of $v \sim F_i$ being greater than H by $p_i \triangleq F_i\big((H,\infty)\big)$. We first note that

$$p_i \leq \frac{\varepsilon}{2\cdot n\cdot(n-1)} \tag{1}$$

for every $i \in [n]$. Indeed, the revenue from F_i of the auction selling (item i) for a take-it-or-leave-it price of H is at least $H\cdot p_i$, and by definition of $\mathcal{R}ev$ we therefore have $H\cdot p_i \leq \mathcal{R}ev(F) \leq R$ and so $p_i \leq R/H \leq \frac{\varepsilon}{2\cdot n\cdot(n-1)}$, as claimed. In particular, since $\varepsilon < 1$, we note that this implies that $F_i([0,H]) = 1 - p_i > 0$ and so, since $\mathcal{EU}_H^n \supset [0,H]^n$, we obtain that $F(\mathcal{EU}_H^n) \geq (1-p_i)^n > 0$, proving Part a. (Thus, $F|_{\mathcal{EU}_H^n}$ is well defined.) For the proof of Parts b and c, we will need the following sublemma.

Sublemma 2.2.1. *Let $n \in \mathbb{N}$ and let $F \in \Delta(\mathbb{R}_+^n)$.*

a. $F(B) \cdot \mathcal{R}ev(F|_B) \le F(A) \cdot \mathcal{R}ev(F|_A)$ for every $B \subseteq A \subseteq \mathbb{R}_+^n$.[7]

b. For every $m \in \mathbb{N}$ and every $A_1, A_2, \ldots, A_m \subseteq \mathbb{R}_+^n$ s.t. $\bigcup_{i=1}^m A_m = \mathbb{R}_+^n$, we have that $\mathcal{R}ev(F) \le \sum_{i=1}^m F(A_i) \cdot \mathcal{R}ev(F|_{A_i})$.

The proof of Sublemma 2.2.1 is given after the proof of Lemma 2.2. We now proceed to prove Parts b and c of Lemma 2.2. We note that for every $i \in [n]$, if $p_i > 0$, then

$$\mathcal{R}ev(F_i|_{(H,\infty)}) \le {}^{\mathcal{R}ev(F_i)}/_{p_i}. \tag{2}$$

Indeed, by Sublemma 2.2.1(a), we have that $p_i \cdot \mathcal{R}ev(F_i|_{(H,\infty)}) \le \mathcal{R}ev(F_i)$, and so $\mathcal{R}ev(F_i|_{(H,\infty)}) \le {}^{\mathcal{R}ev(F_i)}/_{p_i}$, as required.

For every $1 \le i < j \le n$, we let $B_{i,j} \triangleq \big\{(v_1, v_2, \ldots, v_n) \in \mathbb{R}_+^n \bigm| v_i > H \;\;\&\;\; v_j > H\big\}$ (the "double-tail" w.r.t. i and j) and $p_{i,j} \triangleq F(B_{i,j}) = p_i \cdot p_j$. We claim that

$$p_{i,j} \cdot \mathcal{R}ev(F|_{B_{i,j}}) \le \frac{\varepsilon}{\binom{n}{2}} \cdot \mathcal{R}ev(F), \tag{3}$$

for every $1 \le i < j \le n$. Since the claim trivially holds when $p_{i,j} = 0$, we need only prove it when $p_i > 0$ and $p_j > 0$. In this case, by Theorem 2.1 (applied twice) and by Eqs. (1) and (2), we have:

$$\begin{aligned}
p_{i,j} \cdot \mathcal{R}ev(F|_{B_{i,j}}) &= p_{i,j} \cdot \mathcal{R}ev\Big(F_i|_{(H,\infty)} \times F_j|_{(H,\infty)} \times \bigtimes_{k \in [n]\setminus\{i,j\}} F_k\Big) \le \\
&\le 4 \cdot p_{i,j} \cdot \Big(\mathcal{R}ev\big(F_i|_{(H,\infty)}\big) + \mathcal{R}ev\big(F_j|_{(H,\infty)}\big) + \mathcal{R}ev\big(\bigtimes_{k \in [n]\setminus\{i,j\}} F_k\big)\Big) \le \\
&\le 4 \cdot p_{i,j} \cdot \Big(\frac{\mathcal{R}ev(F_i)}{p_i} + \frac{\mathcal{R}ev(F_j)}{p_j} + \mathcal{R}ev\big(\bigtimes_{k \in [n]\setminus\{i,j\}} F_k\big)\Big) = \\
&= 4 \cdot \Big(p_j \cdot \mathcal{R}ev\big(F_i\big) + p_i \cdot \mathcal{R}ev\big(F_j\big) + p_i \cdot p_j \cdot \mathcal{R}ev\big(\bigtimes_{k \in [n]\setminus\{i,j\}} F_k\big)\Big) \le \\
&\le 4 \cdot \frac{\varepsilon}{2 \cdot n \cdot (n-1)} \Big(\mathcal{R}ev\big(F_i\big) + \mathcal{R}ev\big(F_j\big) + \mathcal{R}ev\big(\bigtimes_{k \in [n]\setminus\{i,j\}} F_k\big)\Big) \le \\
&\le 2 \cdot \frac{\varepsilon}{n \cdot (n-1)} \cdot \mathcal{R}ev(F) = \frac{\varepsilon}{\binom{n}{2}} \cdot \mathcal{R}ev(F),
\end{aligned}$$

as claimed.

We define $p_{\mathcal{EU}} \triangleq F(\mathcal{EU}_H^n)$. As $\mathbb{R}_+^n \setminus \mathcal{EU}_H^n = \bigcup_{1 \le i < j \le n} B_{i,j}$, by Sublemma 2.2.1(b) we have that $\mathcal{R}ev(F) \le p_{\mathcal{EU}} \cdot \mathcal{R}ev(F|_{\mathcal{EU}_H^n}) + \sum_{1 \le i < j \le n} p_{i,j} \cdot \mathcal{R}ev(F|_{B_{i,j}})$. Therefore, by Eq. (3), we have that

$$p_{\mathcal{EU}} \cdot \mathcal{R}ev(F|_{\mathcal{EU}_H^n}) \ge \mathcal{R}ev(F) - \sum_{1 \le i < j \le n} p_{i,j} \cdot \mathcal{R}ev(F|_{B_{i,j}}) \ge \mathcal{R}ev(F) - \binom{n}{2} \cdot \frac{\varepsilon}{\binom{n}{2}} \cdot \mathcal{R}ev(F) = (1-\varepsilon) \cdot \mathcal{R}ev(F). \tag{4}$$

In particular, $\mathcal{R}ev(F|_{\mathcal{EU}_H^n}) \ge p_{\mathcal{EU}} \cdot \mathcal{R}ev(F|_{\mathcal{EU}_H^n}) \ge (1-\varepsilon) \cdot \mathcal{R}ev(F)$, proving Part b.

Let $a \in (0,1]$ and let $\mathcal{M}$ be an IC and IR n-item auction with $\mathcal{R}ev_{\mathcal{M}}(F|_{\mathcal{EU}_H^n}) \ge a \cdot \mathcal{R}ev(F|_{\mathcal{EU}_H^n})$. By definition of $\mathcal{R}ev$ and by Eq. (4), we have that

$$\begin{aligned}
\mathcal{R}ev_{\mathcal{M}}(F) = p_{\mathcal{EU}} \cdot \mathcal{R}ev_{\mathcal{M}}(F|_{\mathcal{EU}_H^n}) + (1 - p_{\mathcal{EU}}) \cdot \mathcal{R}ev_{\mathcal{M}}(F|_{\mathbb{R}_+^n \setminus \mathcal{EU}_H^n}) \ge \\
\ge p_{\mathcal{EU}} \cdot \mathcal{R}ev_{\mathcal{M}}(F|_{\mathcal{EU}_H^n}) \ge p_{\mathcal{EU}} \cdot a \cdot \mathcal{R}ev(F|_{\mathcal{EU}_H^n}) \ge a \cdot (1-\varepsilon) \cdot \mathcal{R}ev(F),
\end{aligned}$$

proving Part c. □

[7] If $F(C) = 0$ for some $C \subseteq \mathbb{R}_+^n$, then while $F|_C$ is not defined, we henceforth define $F(C) \cdot \mathcal{R}ev(F|_C)$ to equal 0.

Proof of Sublemma 2.2.1. For Part a, if $F(B) = 0$, then there is nothing to prove, so we assume henceforth that $F(B) > 0$ and hence also $F(A) \geq F(B) > 0$ and thus $F|_A$ and $F|_A|_B = F|_B$ are well defined. We begin by noting that $\mathcal{R}ev(F|_A) \geq F|_A(B)\cdot\mathcal{R}ev(F|_B)$. Indeed, this inequality holds since for any auction $\mathcal{M}$ (in particular an auction obtaining revenue close to optimal from $F|_B$), we have $\mathcal{R}ev_{\mathcal{M}}(F|_A) \geq F|_A(B) \cdot \mathcal{R}ev_{\mathcal{M}}(F|_B)$; by definition of $\mathcal{R}ev$, the inequality follows. Therefore, we have that $F(A) \cdot \mathcal{R}ev(F|_A) \geq F(A) \cdot F|_A(B) \cdot \mathcal{R}ev(F|_B) = F(B) \cdot \mathcal{R}ev(F|_B)$, as required.

For Part b, we start by defining $B_i \triangleq A_i \setminus \bigcup_{j=1}^{i-1} A_j$. By definition, $(B_i)_{i=1}^m$ is a partition of $\mathbb{R}_+^n$. We first claim that $\mathcal{R}ev(F) \leq \sum_{i=1}^m F(B_i) \cdot \mathcal{R}ev(F|_{B_i})$. Indeed, for any auction $\mathcal{M}$ (in particular an auction obtaining revenue close to optimal from F), we have that $\mathcal{R}ev_{\mathcal{M}}(F) = \sum_{i=1}^m F(B_i) \cdot \mathcal{R}ev_{\mathcal{M}}(F|_{B_i})$;[8] by definition of $\mathcal{R}ev$, the inequality follows. By Part a, we have that $F(B_i) \cdot \mathcal{R}ev(F|_{B_i}) \leq F(A_i) \cdot \mathcal{R}ev(F|_{A_i})$ for every $i \in [m]$. Combining both of these, we obtain that $\mathcal{R}ev(F) \leq \sum_{i=1}^m F(B_i) \cdot \mathcal{R}ev(F|_{B_i}) \leq \sum_{i=1}^m F(A_i) \cdot \mathcal{R}ev(F|_{A_i})$, as required. ☐

2.4 Exclusivity at Expensive Menu Entries

Now that we have proven Lemma 2.2, from this point onward we phrase and prove the next steps for arbitrary exclusively unbounded distributions, i.e., not necessarily product distributions conditioned upon $\mathcal{EU}_H^n$. As outlined above, our second step toward proving Theorem 1.1, which we take in this section, shows that in any auction over some exclusively unbounded distribution, the "expensive" part of the menu, i.e., the part of the menu consisting of all menu entries that cost more than some $E = \text{poly}(n, 1/\varepsilon)$, can be simplified without significant loss in revenue to make the auction E-*exclusive*, i.e., to make each expensive menu entry only allocate a single item with nonzero probability. This step is formalized by Lemma 2.3.

Definition 2.7 (Exclusive Auction)**.** Let $n \in \mathbb{N}$ and let $E \in \mathbb{R}_+$. We say that an n-item auction is E-*exclusive* if it allocates (with positive probability) at most one item whenever it charges strictly more than E.[9]

Lemma 2.3. *Let $n \in \mathbb{N}$ s.t. $n \geq 2$, let $H \in \mathbb{R}_+$, let $\varepsilon \in (0,1)$, and set $E \triangleq \frac{4\cdot(n-1)\cdot H}{\varepsilon^2}$. For every $F \in \Delta(\mathcal{EU}_H^n)$ and for every IC and IR n-item auction $\mathcal{M}$, there exists an E-exclusive IC and IR n-item auction $\mathcal{M}'$ such that $\mathcal{R}ev_{\mathcal{M}'}(F) \geq (1-\varepsilon) \cdot \mathcal{R}ev_{\mathcal{M}}(F)$.*

Proof. Set $\tilde{\varepsilon} \triangleq \varepsilon/2$. We construct a new IC and IR auction $\mathcal{M}'$ as follows:

- For every menu entry $e = (\vec{x}; p) \in \mathcal{M}$ with $p \leq E$, we add the menu entry e to $\mathcal{M}'$, unmodified.
- For every menu entry $e = (\vec{x}; p) \in \mathcal{M}$ with $p > E$, we add the following n menu entries to $\mathcal{M}'$: $\big(x_1, \vec{0}_{-1}; (1-\tilde{\varepsilon}) \cdot p\big)$, $\big(x_2, \vec{0}_{-2}; (1-\tilde{\varepsilon}) \cdot p\big)$,$\cdots$,$\big(x_n, \vec{0}_{-n}; (1-\tilde{\varepsilon}) \cdot p\big)$;[10] each of these menu entries is a modified version of e that completely "unallocates" all but one of the items, while giving a slight multiplicative price discount of $(1-\tilde{\varepsilon})$.

Finally, we define $\mathcal{M}'$ to be the closure of the set of menu entries that we have just added to it.[11]

By definition, $\mathcal{M}'$ is E-exclusive. We note that since $\mathcal{M}$ is IR, it contains the menu entry $(\vec{0}; 0)$, and therefore so does $\mathcal{M}'$, and hence $\mathcal{M}'$ is IR as well. It remains to show that $\mathcal{M}'$ obtains revenue

[8] Similarly, if $F(B_i) = 0$ for some $i \in [n]$, then we define $F(B_i) \cdot \mathcal{R}ev_{\mathcal{M}}(F|_{B_i})$ to equal 0.

[9] While this also implies that the allocated item is quite expensive, and therefore some may say exclusive, the exclusivity discussed in the definition is that of only that item being sold.

[10] As is customary in the literature, we use the notation $(y', \vec{y}_{-i}; q)$, for $\vec{y} \in [0,1]^n$, $i \in [n]$, $y' \in [0,1]$, and $q \in \mathbb{R}_+$, to denote the outcome $(y_1, \ldots, y_{i-1}, y', y_{i+1}, \ldots, y_n; q)$, i.e., an outcome that is identical to $(\vec{y}; q)$ in price and all winning probabilities, except for that of item i, which is set to y'.

[11] Taking the closure ensures that a utility-maximizing entry with maximal price exists for every buyer type. See Appendix A.1 for more details.

at least $(1-\varepsilon)\cdot \mathcal{R}ev_{\mathcal{M}}(F)$ from F. Let us compare the payments that $\mathcal{M}'$ and $\mathcal{M}$ extract from a buyer of each type $v=(v_1,v_2,\ldots,v_n)\in \mathcal{EU}^n_H$. We reason by cases, according to the menu entry of choice of buyer type v from $\mathcal{M}$, which we denote by $e=(\vec{x};p)$,[12] showing that in either case, the payment extracted from a buyer of type v decreases by at most a multiplicative factor of $(1-2\tilde{\varepsilon})$ in $\mathcal{M}'$ compared to $\mathcal{M}$.

- If $p\le E$, then by definition $e\in\mathcal{M}'$. We claim that v weakly prefers e over all menu entries $f'=(\vec{y}';q')\in\mathcal{M}'$ s.t. $q'<(1-\tilde{\varepsilon})\cdot p$.[13] Indeed, for every such menu entry f', we have that $q'<(1-\tilde{\varepsilon})\cdot E$ and so by definition we have that $f'\in\mathcal{M}$ and so, by definition of e, we have that v weakly prefers e over f'. Therefore, the price of the menu entry chosen by v from $\mathcal{M}'$ is at least $(1-\tilde{\varepsilon})\cdot p$, and so the payment extracted from a buyer of type v decreases by at most a multiplicative factor of $(1-\tilde{\varepsilon})$ in $\mathcal{M}'$ compared to $\mathcal{M}$.
- Otherwise, i.e., if $p>E$, then since $E>n\cdot H$, by IR there must exist $i\in[n]$ s.t. $v_i>H$; as $v\in\mathcal{EU}^n_H$, we have that $v_k\le H$ for every $k\in[n]\setminus\{i\}$. Let $e'=(\vec{x}';p')\triangleq\big(x_i,\vec{0}_{-i};(1-\tilde{\varepsilon})\cdot p\big)$ be the menu entry in $\mathcal{M}'$ corresponding to e that unallocates all items but item i. We claim that v weakly prefers e' over all menu entries $f'=(\vec{y}';q')\in\mathcal{M}'$ s.t. $q'<p'-\frac{(1-\tilde{\varepsilon})\cdot(n-1)\cdot H}{\tilde{\varepsilon}}$.[14] Let f' be such a menu entry and denote the menu entry corresponding to f' in $\mathcal{M}$ by $f=(\vec{y};q)$ (where either $f'=f$ or $f'=(x_j,\vec{0}_{-j},(1-\tilde{\varepsilon})\cdot p)$ for some $j\in[n]$). Noting that $q\le\frac{q'}{1-\tilde{\varepsilon}}<p-\frac{(n-1)\cdot H}{\tilde{\varepsilon}}$, we indeed have that

$$u_{e'}(v)=x'_i\cdot v_i-p'=x_i\cdot v_i-p+\tilde{\varepsilon}\cdot p\ge-(n-1)\cdot H+\sum_{k=1}^{n}x_k\cdot v_k-p+\tilde{\varepsilon}\cdot p=$$

$$=-(n-1)\cdot H+u_e(v)+\tilde{\varepsilon}\cdot p\ge-(n-1)\cdot H+u_f(v)+\tilde{\varepsilon}\cdot p=-(n-1)\cdot H+\sum_{k=1}^{n}y_k\cdot v_k-q+\tilde{\varepsilon}\cdot p>$$

$$>-(n-1)\cdot H+\sum_{k=1}^{n}y_k\cdot v_k-q+\tilde{\varepsilon}\cdot\Big(q+\tfrac{(n-1)\cdot H}{\tilde{\varepsilon}}\Big)=-(n-1)\cdot H+\sum_{k=1}^{n}y_k\cdot v_k-q\cdot(1-\tilde{\varepsilon})+(n-1)\cdot H\ge$$

$$\ge\sum_{k=1}^{n}y_k\cdot v_k-q'\ge\sum_{k=1}^{n}y'_k\cdot v_k-q'=u_{f'}(v).$$

  Therefore, the payment that $\mathcal{M}'$ extracts from a buyer of type v is at least

$$p'-\tfrac{(1-\tilde{\varepsilon})\cdot(n-1)\cdot H}{\tilde{\varepsilon}}=p'-(1-\tilde{\varepsilon})\cdot\tilde{\varepsilon}\cdot E>p'-(1-\tilde{\varepsilon})\cdot\tilde{\varepsilon}\cdot p\ge$$
$$\ge p'-\tilde{\varepsilon}\cdot p'=(1-\tilde{\varepsilon})\cdot p'=(1-\tilde{\varepsilon})^2\cdot p>(1-2\tilde{\varepsilon})\cdot p,$$

  and so the payment extracted from a buyer of type v decreases by at most a multiplicative factor of $(1-2\tilde{\varepsilon})$ in $\mathcal{M}'$ compared to $\mathcal{M}$.

To summarize, the revenue from each buyer type $v\in\mathcal{EU}^n_H$ decreases by at most a multiplicative factor of $(1-2\tilde{\varepsilon})=(1-\varepsilon)$ in $\mathcal{M}'$ compared to $\mathcal{M}$, and so the (overall) revenue that $\mathcal{M}'$ obtains from $\mathcal{EU}^n_H$ is at least a $(1-\varepsilon)$ fraction of the revenue that $\mathcal{M}$ obtains from $\mathcal{EU}^n_H$, as required. □

[12] If more than one utility-maximizing menu entry with maximal price exists, then here and in similar situations henceforth we choose one such entry arbitrarily.

[13] As we show weak preference, and as q' is defined via a strict inequality, by continuity of the utility function correctness of the claim for all f' before taking the closure of $\mathcal{M}'$ implies correctness for all f' in the closure as well.

[14] See Footnote 13.

2.5 Trimming the Expensive Part of the Menu

As outlined above, our third step toward proving Theorem 1.1, which we take in this section, shows that in any exclusive auction over some exclusively unbounded distribution, the expensive part of the menu can be simplified without significant loss in revenue, so that it contains at most $2n$ menu entries. This step, which is the most technically elaborate of all steps, is formalized by Lemma 2.4.

Lemma 2.4. *Let $n \in \mathbb{N}$ s.t. $n \geq 2$, let $H \in \mathbb{R}_+$, let $\varepsilon \in (0,1)$, and let $E \geq \mathrm{Max}\big\{n \cdot H + 2, \frac{(n-1)\cdot H+1}{\varepsilon}\big\}$. For every $F \in \Delta(\mathcal{EU}_H^n)$ and for every E-exclusive IC and IR n-item auction $\mathcal{M}$, there exists an E-exclusive IC and IR n-item auction $\mathcal{M}'$, such that both of the following hold.*

- $\mathcal{R}ev_{\mathcal{M}'}(F) \geq (1-\varepsilon) \cdot \mathcal{R}ev_{\mathcal{M}}(F)$.
- *(The set of menu entries in) $\mathcal{M}'$ coincides with a subset of the set of menu entries in $\mathcal{M}$ that cost at most E, with the addition of at most $2n$ menu entries.*

Proof. Throughout the proof, we assume w.l.o.g. that each menu entry in $\mathcal{M}$ (except possibly $(\vec{0};0)$) is chosen by at least one buyer type $v \in \mathcal{EU}_H^n$.[15] For every $i \in [n]$, we define

$$\mathcal{M}_i \triangleq \big\{(\vec{x};p) \in \mathcal{M} \;\big|\; p > E \;\;\&\;\; x_i > 0\big\},$$

and

$$W_i \triangleq \{v \in \mathcal{EU}_H^n \mid v\text{'s menu entry of choice from } \mathcal{M} \text{ is in } \mathcal{M}_i\}.$$

As $\mathcal{M}$ is E-exclusive, we have that $x_j = 0$ for every $j \in [n] \setminus \{i\}$ and $(\vec{x};p) \in \mathcal{M}_i$. Thus, W_i is the set of buyer types that choose to pay more than E, and receive in return a positive probability for winning item i and zero probability for winning any other item. Our goal is to apply the single-dimensional analysis of Myerson (1981) in order to replace the plethora of menu entries in each $\mathcal{M}_i$ with a small constant number of menu entries. **Until further notice, fix $\boldsymbol{i \in [n]}$ s.t. $\boldsymbol{\mathcal{M}_i \neq \emptyset}$.**

We define $b_i \triangleq \inf\big\{p \;\big|\; (\vec{x};p) \in \mathcal{M}_i\big\}$ and $s_i \triangleq \inf\big\{x_i \;\big|\; (\vec{x};p) \in \mathcal{M}_i\big\}$, and set $z_i \triangleq (s_i, \vec{0}_{-i}; b_i)$. One may intuitively think about z_i (charging b_i for an s_i probability of winning item i) as the cheapest entry, allocating the least probability, in $\mathcal{M}_i$, although formally (as $\mathcal{M}$ is not closed, and also due to the way in which $\mathcal{M}_i$ is defined) z_i need not necessarily be in $\mathcal{M}_i$.[16] Our strategy is to show that, in a precise sense, $\mathcal{M}_i$ behaves as follows on W_i: First allocate the buyer z_i (i.e., provide a "starting winning probability" of s_i for item i, and charge a "base price" of b_i), and then hold a "continuation auction" for possibly allocating some or all of the remaining $(1-s_i)$ probability of winning item i; in this "continuation auction," we allow the buyer to swap z_i for a different entry from $\mathcal{M}_i$, paying the difference in costs and increasing the probability of getting item i accordingly. We now make this statement precise. **Until further notice, fix $\boldsymbol{i \in [n]}$ s.t. $\boldsymbol{F(W_i) > 0}$ (and thus, by definition, $\boldsymbol{\mathcal{M}_i \neq \emptyset}$) and $\boldsymbol{s_i < 1}$.** We note that as we will see below, we require that $s_i < 1$ in order for this "continuation auction" to be well defined (and in fact make sense, otherwise there is no "remaining probability" to sell) and that $F(W_i) > 0$ for the valuation distribution of that auction to be well defined. As we will show, if $s_i = 1$ then $|\mathcal{M}_i| = \big|\{z_i\}\big| = 1$ (see Sublemma 2.4.4(c) below), and so there will be no need to reduce the number of menu entries in $\mathcal{M}_i$ in this case (and if $F(W_i) = 0$, then the revenue of the original auction $\mathcal{M}$ from W_i is zero, and so we will be able to simply delete $\mathcal{M}_i$ from the original auction, without replacing it with anything).

[15]While $\mathcal{M}$, after the removal of all menu entries that are chosen by no buyer type, is not necessarily closed, it does possess a utility-maximizing entry with maximal price (see Appendix A.1 for a discussion) for every buyer type.

[16]Nonetheless, due to the assumption that every menu entry is chosen by some buyer type, we are able to show that z_i is in the closure of $\mathcal{M}_i$; see Sublemma 2.4.4(d) below. This property is heavily used throughout our proof.

For every n-dimensional buyer type $v \in W_i$, we define the single-dimensional valuation for the "remaining probability" of winning item i by

$$\alpha(v) \triangleq v_i \cdot (1 - s_i).$$

We define the corresponding single-dimensional buyer type space of our "continuation auction" as

$$W_i^1 \triangleq \{\alpha(v) \mid v \in W_i\} \in \mathbb{R}_+,$$

and define a distribution $F_i^1 \in \Delta(W_i^1) \subseteq \Delta(\mathbb{R}_+)$ over it by

$$F_i^1(V) \triangleq F|_{W_i}\big(\alpha^{-1}(V)\big)$$

for every measurable set $V \subseteq W_i^1$. (Recall that $F(W_i) > 0$; therefore, $F|_{W_i}$ is well defined.) Very roughly speaking, F_i^1 is defined such that its density at every $v^1 \in W_i^1$ can be informally thought of as the sum of the densities of $F|_{W_i}$ at all $v \in W_i$ s.t. $\alpha(v) = v^1$.

Having defined the single-dimensional buyer type space of the "continuation auction," we now turn to defining the (menu of the) auction itself. For an outcome (not necessarily a menu entry in $\mathcal{M}$) of the form $e = (x_i, \vec{0}_{-i}; p)$, we define

$$\beta(e) = \beta(x_i, \vec{0}_{-i}; p) \triangleq \left(\tfrac{x_i - s_i}{1 - s_i}; p - b_i\right),$$

i.e., an entry selling a fraction of the remaining probability $(1 - s_i)$ so that (in addition to the starting winning probability s_i) the overall winning probability is x_i, in exchange for "upping" the price (from the base price b_i) to p. (Recall that $s_i < 1$; therefore, β is well defined.) We define the "continuation auction" by[17]

$$\mathcal{M}_i^1 \triangleq \big\{\beta(e) \mid e \in \mathcal{M}_i \cup \{z_i\}\big\}.$$

We note that $\mathcal{M}_i^1 \in [0,1] \times \mathbb{R}_+$ by definition of s_i and b_i, and furthermore that $(0; 0) = \beta(z_i) \in \mathcal{M}_i^1$; therefore, $\mathcal{M}_i^1$ is an IC and IR 1-item auction. We now make precise the above claim regarding $\mathcal{M}_i$ acting as if the buyer is allocated z_i and then bids in the single-dimension "continuation auction" $\mathcal{M}_i^1$, thus relating the revenue of the single-dimensional "continuation auction" $\mathcal{M}_i^1$ with that of the original auction $\mathcal{M}$ from W_i.

Sublemma 2.4.1 (Relation between $\mathcal{M}_i^1$ and $\mathcal{M}$)**.**

a. *For every $v \in W_i$, we have that $\beta(e)$ is the menu entry of choice of $\alpha(v)$ from $\mathcal{M}_i^1$, where $e \in \mathcal{M}_i$ is the menu entry of choice of v from $\mathcal{M}$.*

b. $\mathcal{R}ev_{\mathcal{M}_i^1}(F_i^1) = \mathcal{R}ev_{\mathcal{M}}(F|_{W_i}) - b_i$.

The proof of Sublemma 2.4.1 is given after the proof of Lemma 2.4. We note that the only delicate part of the proof (apart from the definitions, of course), is showing that even if z_i is not an entry of $\mathcal{M}_i$, then in the notation of the sublemma no buyer type chooses z_i over e (equivalently, $(0; 0)$ over $\beta(e)$); see Sublemma 2.4.6(a) below.

Now that we have formally constructed a suitable single-dimensional auction, we apply the celebrated theorem of Myerson (1981), to obtain a price $c_i \in \mathbb{R}_+$ s.t. the IC and IR 1-item auction $\mathcal{M}_i^* \triangleq \big\{(0; 0), (1; c_i)\big\}$ maximizes the revenue from F_i^1, i.e., $\mathcal{R}ev_{\mathcal{M}_i^*}(F_i^1) = \mathcal{R}ev(F_i^1)$. Ideally, we would have now wanted to replace all entries from $\mathcal{M}_i$ in $\mathcal{M}$ by the two entries $z_i = \beta^{-1}\big((0; 0)\big)$

[17] We emphasize that we do not take the closure (see Appendix A.1 for a discussion) of $\mathcal{M}_i^1$, as there is no "need" to do so. Indeed, as shown by Sublemma 2.4.1(a) below, $\mathcal{M}_i^1$ already possesses a utility-maximizing entry with maximal price for each buyer type.

and $o_i \triangleq (1, \vec{0}_{-i}; c_i + b_i) = \beta^{-1}\big((1; c_i)\big)$, however it is not clear that following such a replacement, none of the buyer types in W_i whose utility from this replacement has dropped would not actually now choose some "cheap" entry instead, which might cause a sharp decrease in revenue. Therefore, in order to make sure that no buyer type has an incentive to switch to a very cheap entry, instead of replacing $\mathcal{M}_i$ with $\{z_i, o_i\}$, we replace it with $\{z'_i, o'_i\}$, where z'_i and o'_i are slightly discounted versions of z_i and o_i, respectively. Formally, we define $z'_i \triangleq \big(s_i, \vec{0}_{-i}; b_i - (n-1) \cdot H - 1\big)$ and $o'_i \triangleq \big(1, \vec{0}_{-i}; c_i + b_i - (n-1) \cdot H - 1\big)$. The menu entries z'_i and o'_i sell the same winning probabilities as z_i and o_i (i.e., s_i and 1), respectively, but with a slight (when compared to E) additive price discount of $(n-1) \cdot H + 1$.

Having defined the new menu entries for every i separately, we define our "trimmed" auction:[18]

$$\mathcal{M}' \triangleq \Big(\mathcal{M} \cap \big([0,1]^n \times [0, E]\big)\Big) \cup \bigcup_{\substack{i \in [n]: \\ F(W_i) > 0 \ \& \\ s_i = 1}} \{z_i\} \cup \bigcup_{\substack{i \in [n]: \\ F(W_i) > 0 \ \& \\ s_i < 1}} \{z'_i, o'_i\}.$$

We note that since $\mathcal{M}$ is IR, it contains the menu entry $(\vec{0}; 0)$, and therefore so does $\mathcal{M}'$ and hence $\mathcal{M}'$ is IR as well. By definition, $\mathcal{M}'$ contains at most 2 *expensive* (i.e., costing more than E) items for each $i \in [n]$, for a total of at most $2n$ expensive items, as required. It remains to reason about the revenue of $\mathcal{M}'$. To do so, we now make precise what can be thought of as a "converse" for Sublemma 2.4.1, thus relating the revenue of the new trimmed auction $\mathcal{M}'$ from W_i with that of the optimal single-dimensional Myerson auction $\mathcal{M}^*_i$.

Sublemma 2.4.2 (Relation between $\mathcal{M}'$ and $\mathcal{M}^*_i$)**.** *Let $i \in [n]$ s.t. $F(W_i) > 0$ and $s_i < 1$.*

a. For every $v \in W_i$, the menu entry of choice of v from $\mathcal{M}'$ is either o'_i or z'_i:

- *o'_i is the menu entry of choice of v from $\mathcal{M}'$ iff $(1; c_i)$ is the menu entry of choice of $\alpha(v)$ from $\mathcal{M}^*_i$.*
- *z'_i is the menu entry of choice of v from $\mathcal{M}'$ iff $(0; 0)$ is the menu entry of choice of $\alpha(v)$ from $\mathcal{M}^*_i$.*

b. $\mathcal{R}ev_{\mathcal{M}'}(F|_{W_i}) = \mathcal{R}ev_{\mathcal{M}^*_i}(F^1_i) + b_i - (n-1) \cdot H - 1$.

The proof of Sublemma 2.4.2 is given after the proof of Lemma 2.4. We note that, as implied above, the delicate part of the proof is showing that no buyer type suddenly chooses a very cheap menu entry (or, hypothetically, an expensive menu entry from a different "expensive part" of $\mathcal{M}'$) due to the move from $\mathcal{M}$ with $\mathcal{M}'$. Showing this strongly depends on the assumption that every menu entry in $\mathcal{M}$ is chosen by some buyer type; see Sublemma 2.4.6(b,c) below.

We now have all the pieces of the puzzle needed to show that our trimmed auction loses no significant revenue from W_i for every $i \in [n]$ s.t. $s_i < 1$. In fact, the following sublemma shows not only that, but also that from no other part of the buyer type space, "wealthy" or "poor", does our trimmed auction lose significant revenue compared to the original untrimmed auction $\mathcal{M}$.

Sublemma 2.4.3 (Comparison between the Revenues of $\mathcal{M}'$ and $\mathcal{M}$ from Restrictions of F)**.**

---

[18] As before, we emphasize that we do not take the closure (see Appendix A.1 for a discussion) of $\mathcal{M}'$, as there is no "need" to do so. Indeed, $\mathcal{M}'$ already possess a utility-maximizing entry with maximal price for each buyer type except perhaps for a set of buyer types of measure zero: For $v \in \mathcal{EU}^n_H \setminus \bigcup_{i \in [n]: F(W_i) = 0 \text{ or } s_i < 1} W_i$ since its entry of choice from $\mathcal{M}$ is in $\mathcal{M}'$ (as mentioned above, when $s_i = 1$ we have that $\mathcal{M}_i = \{z_i\}$; see Sublemma 2.4.4(c) below) and only finitely many additional menu entries were added, and for $v \in \bigcup_{i \in [n]: F(W_i) > 0 \ \& \ s_i < 1} W_i$ as shown by Sublemma 2.4.2(a) below.

a. For every $i \in [n]$ *s.t.* $F(W_i) > 0$ & $s_i < 1$*, we have that* $\mathcal{R}ev_{\mathcal{M}'}(F|_{W_i}) \geq (1-\varepsilon)\cdot\mathcal{R}ev_{\mathcal{M}}(F|_{W_i})$*.*

b. For every $i \in [n]$ *s.t.* $F(W_i) > 0$ & $s_i = 1$*, we have that* $\mathcal{R}ev_{\mathcal{M}'}(F|_{W_i}) = \mathcal{R}ev_{\mathcal{M}}(F|_{W_i})$*.*

c. Letting $P \triangleq \mathcal{EU}_H^n \setminus \bigcup_{i\in[n]} W_i$*, if* $F(P) > 0$*, then* $\mathcal{R}ev_{\mathcal{M}'}(F|_P) \geq (1-\varepsilon)\cdot\mathcal{R}ev_{(\mathcal{M}}(F|_P)$*.*

The proof of Sublemma 2.4.3 is given after the proof of Lemma 2.4. We conclude the proof of Lemma 2.4, as by Sublemma 2.4.3 we have that

$$\mathcal{R}ev_{\mathcal{M}'}(F) = F(P)\cdot\mathcal{R}ev_{\mathcal{M}'}(F|_P) + \sum_{i=1}^{n} F(W_i)\cdot\mathcal{R}ev_{\mathcal{M}'}(F|_{W_i}) \geq$$

$$\geq (1-\varepsilon)\cdot\Big(F(P)\cdot\mathcal{R}ev_{\mathcal{M}}(F|_P) + \sum_{i=1}^{n} F(W_i)\cdot\mathcal{R}ev_{\mathcal{M}}(F|_{W_i})\Big) = (1-\varepsilon)\cdot\mathcal{R}ev_{\mathcal{M}}(F),$$

as required. □

Proof of Sublemmas 2.4.1 through 2.4.3 Before proving Sublemmas 2.4.1 through 2.4.3, we first phrase and prove three auxiliary sublemmas, which in fact encompass most of the delicate details and technical complexity in the proof of Sublemmas 2.4.1 through 2.4.3.

Sublemma 2.4.4 (Properties of $\mathcal{M}_i$)**.** *Let* $i \in [n]$ *s.t.* $\mathcal{M}_i \neq \emptyset$*.*

a. For every $e = (x_i, \vec{0}_{-i}; p) \in \mathcal{M}_i$ *and* $f = (y_i, \vec{0}_{-i}; q) \in \mathcal{M}_i$*, either i)* $e = f$*, or ii)* $x_i < y_i$ *and* $p < q$*, or iii)* $x_i > y_i$ *and* $p > q$*.*

b. If there exists $e = (x_i, \vec{0}_{-i}; p) \in \mathcal{M}_i$ *s.t.* $x_i = s_i$ *(i.e., if* s_i *is attained as a minimum rather than merely an infimum), then* $p = b_i$*, i.e.,* $z_i = e$ *and so* $z_i \in \mathcal{M}_i$*.*

c. If $s_i = 1$*, then* $\mathcal{M}_i = \{z_i\}$*.*

d. For every $i \in [n]$*, there exists a sequence* $\big(z_i^m = (s_i^m, \vec{0}_{-i}; b_i^m)\big)_{m=1}^{\infty} \subseteq \mathcal{M}_i$ *s.t. i)* $(s_i^m)_{m=1}^{\infty}$ *is weakly decreasing, ii)* $(b_i^m)_{m=1}^{\infty}$ *is weakly decreasing, and iii)* $\lim_{m\to\infty} z_i^m = z_i = (s_i, \vec{0}_{-i}, b_i)$*.*

Proof. We start with Part a. If $p = q$, then we claim that $x_i = y_i$; indeed, assuming w.l.o.g. that $x_i \geq y_i$ we note that if we had $x_i > y_i$, then all buyer types v with $u_f(v) \geq 0$ (and hence $v_i > 0$) would have strictly preferred e over f, as the former sells (for the same price!) a strictly larger probability for winning item i, contradicting the assumption that each menu entry in $\mathcal{M}$ is chosen by at least one buyer type. If $p < q$, then we claim that $x_i < y_i$; indeed, otherwise all buyer types would strictly prefer e over f since e would sell a weakly larger probability for winning item i for a strictly lower price, contradicting the assumption that f is chosen by at least one buyer type. Similarly, if $p > q$, then $x_i > y_i$.

We move on to Part b. If there exists $e = (x_i, \vec{0}_{-i}; p) \in \mathcal{M}_i$ s.t. $x_i = s_i$, then for every $s = (y_i, \vec{0}_{-i}; q) \in \mathcal{M}_i$ we have by definition that $y_i \geq s_i = x_i$ and so by Part a that $q \geq p$, and so by definition $b_i \geq p$. Since $e \in \mathcal{M}_i$, we also have by definition that $b_i \leq p$, and hence $b_i = p$. Therefore, $z_i = e$, as required.

We move on to Part c. Assume that $s_i = 1$ and let $e = (x_i, \vec{0}_{-i}; p) \in \mathcal{M}_i$. By definition $s_i \leq x_i$, however we also have that $s_i = 1 \geq x_i$, and so $s_i = x_i$. By Part b, we therefore have that $e = z_i$. As $\mathcal{M}_i \neq \emptyset$ by assumption, the proof of Part c is complete.

We move on to Part d. If s_i is attained as a minimum, then by Part b, $z_i \in \mathcal{M}_i$ and we may set $z_i^m \triangleq z$ for all $m \in \mathbb{N}$, and so $(z_i^m)_{m=1}^{\infty}$ trivially meets all of the desired requirements. It remains to analyze the case in which s_i is not attained as a minimum but rather only as an infimum; assume henceforth, therefore, that this is the case.

By definition of s_i, there exists a sequence $\big(z_i^m = (s_i^m, \vec{0}_{-i}; b_i^m)\big)_{m=1}^{\infty} \subseteq \mathcal{M}_i$ s.t. $(s_i^m)_{m=1}^{\infty}$ is weakly decreasing and s.t. $\lim_{m\to\infty} s_i^m = s_i$. By Part a, $(b_i^m)_{m=1}^{\infty}$ is weakly decreasing as well. It is enough, therefore, to show that $\lim_{m\to\infty} b_i^m = b_i$. As $(b_i^m)_{m=1}^{\infty}$ is weakly decreasing and bounded from below by b_i (by definition of b_i), we have that it indeed converges, and moreover, that $\lim_{m\to\infty} b_i^m \geq b_i$; assume for contradiction that $\lim_{m\to\infty} b_i^m > b_i$. Therefore, by definition of b_i there exists $e = (x_i, \vec{0}_{-i}; p) \in \mathcal{M}_i$ s.t. $\lim_{m\to\infty} b_i^m > p$ and so $b_i^m > p$ for all $m \in \mathbb{N}$. By Part a, we therefore have that $s_i^m > x_i$ for all $m \in \mathbb{N}$, and so $s_i \geq x_i$; as $s_i \leq x_i$ by definition, we obtain that $s_i = x_i$, contradicting the assumption that s_i is not attained as a minimum. ☐

Sublemma 2.4.5. *Let $i \in [n]$ and let $e = (x_i, \vec{0}_{-i}; p)$ and $f = (y_i, \vec{0}_{-i}; q)$ be two outcomes (not necessarily menu entries in $\mathcal{M}$) assigning zero probability for winning any item other than i, s.t. $x_i > y_i$. For every $v \in \mathbb{R}_+^n$, we have that both of the following hold.*

- $u_e(v) \geq u_f(v)$ *iff* $v_i \geq \frac{p-q}{x_i-y_i}$.
- $u_f(v) \geq u_e(v)$ *iff* $\frac{p-q}{x_i-y_i} \geq v_i$.

Proof. $u_e(v) \geq u_f(v) \Leftrightarrow v_i \cdot x_i - p \geq v_i \cdot y_i - q \Leftrightarrow v_i \cdot (x_i - y_i) \geq p - q \Leftrightarrow v_i \geq \frac{p-q}{x_i-y_i}$, where in the last equivalence we used the fact that $x_i - y_i > 0$. The proof of the second statement is identical, with all inequalities flipped. ☐

Sublemma 2.4.6 (Preferences of W_i)**.** *Let $i \in [n]$ s.t. $\mathcal{M}_i \neq \emptyset$, let $v \in W_i$, and let $e \in \mathcal{M}_i$ be the menu entry of choice of v from $\mathcal{M}$.*

a. $u_e(v) \geq u_{z_i}(v)$.

b. $u_{z_i'}(v) > u_f(v)$, *for every* $f \in \mathcal{M} \setminus \mathcal{M}_i$.

c. $\min\big\{u_e(v), u_{z_i'}(v)\big\} > u_f(v)$, *for every outcome f of the form $(y_j, \vec{0}_{-j}; q)$ s.t. $j \in [n] \setminus \{i\}$ and $q \geq E - (n-1) \cdot H - 1$ (where f need not necessarily be a menu entry in $\mathcal{M}$).*

Proof. Let $(z_i^m)_{m=1}^{\infty}$ be as in Sublemma 2.4.4(d). We start with Part a. By definition of e, we have that $u_e(v) \geq u_{z_i^m}(v)$ for every $m \in \mathbb{N}$. By continuity of u, we therefore also have that $u_e(v) \geq u_{\lim_{m\to\infty} z_i^m}(v) = u_{z_i}(v)$, as required.

We move on to Part b. Denote $e = (x_i, \vec{0}_{-i}; p) \in \mathcal{M}_i$ (e is still the menu entry of choice of v from $\mathcal{M}$). If $x_i = s_i$, then by Sublemma 2.4.4(b), $e = z_i$, and the proof of Part b is complete, as v (like any other buyer type) strictly prefers z_i' over $z_i = e$, which is the menu entry of choice of v from $\mathcal{M}$, and is therefore weakly preferred by v over f. Otherwise, $x_i > s_i$, and so there exists $N \in \mathbb{N}$ s.t. $x_i > s_i^m$ for every $m \geq N$; assume w.l.o.g. that $N = 1$.

Denote $f = (\vec{y}; q)$ and let $f' \triangleq (y_i, \vec{0}_{-i}; q)$ (we emphasize that f' need not necessarily be a menu entry in $\mathcal{M}$). We claim that $u_{z_i^m}(v) \geq u_{f'}(v)$ for every $m \in \mathbb{N}$. Indeed, let $m \in \mathbb{N}$ and let $w \in W_i$ be a buyer type that chooses z_i^m from $\mathcal{M}$. As $x_i > s_i^m$, by Sublemma 2.4.5 (applied twice) we have that $v_i \geq \frac{p-b_i^m}{x_i-s_i^m} \geq w_i$. We note that $s_i^m > y_i$. Indeed, if $q > E$, then as $f \in \mathcal{M} \setminus \mathcal{M}_i$, we have that $f \in \mathcal{M}_j$ for some $j \in [n] \setminus \{i\}$, and so (since $\mathcal{M}$ is E-exclusive) $y_i = 0 < s_i^m$; otherwise, i.e., if $q \leq E < b_i^m$, then if $y_i \geq s_i^m$, then we would have $u_f(w) = w_i \cdot y_i - q > w_i \cdot s_i^m - b_i^m = u_{z_i^m}(w)$, contradicting the definition of w. By definition, w weakly prefers z_i^m over f, and by definition of f', we have that w (as well as all other buyer types) weakly prefers f over f'. Therefore, w weakly prefers z_i^m over f', and so, since $v_i \geq w_i$ and by Sublemma 2.4.5, we obtain that $v_i \geq w_i \geq \frac{b_i^m - q}{s_i^m - y_i}$.

Therefore, using Sublemma 2.4.5 once more, we have that v weakly prefers z_i^m over f', as claimed. By continuity of u, we therefore have that $u_{z_i}(v) = u_{\lim_{m\to\infty} z_i^m}(v) \geq u_{f'}(v)$ as well.

As $\mathcal{M}$ is IR, since $e \in \mathcal{M}_i$ we have that $v_i > E > H$, and so since $v \in \mathcal{EU}_H^n$, we have that $v_j \leq H$ for every $j \in [n] \setminus \{i\}$. Therefore, $u_{z'_i}(v) = u_{z_i}(v) + (n-1)\cdot H + 1 \geq u_{f'}(v) + (n-1)\cdot H + 1 = u_f(v) - \sum_{j\in[n]\setminus\{i\}} v_j \cdot y_j + (n-1)\cdot H + 1 > u_f(v)$, completing the proof of Part b.

We conclude by proving Part c. Recall that $\mathcal{M}$ is IR; therefore, we have by definition of v that $u_e(v) \geq 0$, and since $(\vec{0};0) \in \mathcal{M} \setminus \mathcal{M}_i$ we also have by Part b that $u_{z'_i}(v) \geq u_{(\vec{0};0)}(v) = 0$. Therefore, to prove Part c it suffices to show that $u_f(v) < 0$. Indeed, recalling that $v_j \leq H$, we have that $u_f(v) = v_j \cdot y_j - q \leq H - q \leq H - E + (n-1)\cdot H + 1 = n \cdot H + 1 - E < 0$. □

Proof of Sublemma 2.4.1. We start with Part a. Recall that $i \in [n]$ s.t. $F(W_i) > 0$ and $s_i < 1$. We start with Part a. We first claim that $u_{\beta(f)}\big(\alpha(v)\big) = u_f(v) - u_{z_i}(v)$ for every $f \in \mathcal{M}_i \cup \{z_i\}$. Indeed, denoting $f = (y_i, \vec{0}_{-i}; q)$, we have

$$u_{\beta(f)}\big(\alpha(v)\big) = v_i\cdot(1-s_i)\cdot\tfrac{y_i-s_i}{1-s_i} - q + b_i = v_i\cdot(y_i - s_i) - q + b_i = (v_i\cdot y_i - q) - (v_i \cdot s_i - b_i) = u_f(v) - u_{z_i}(v).$$

For every $\beta(f), \beta(g) \in \mathcal{M}_i^1$ (where $f, g \in \mathcal{M}_i \cup \{z_i\}$), we therefore have that

$$u_{\beta(f)}\big(\alpha(v)\big) \geq u_{\beta(g)}\big(\alpha(v)\big) \Leftrightarrow u_f(v) - u_{z_i}(v) \geq u_g(v) - u_{z_i}(v) \Leftrightarrow u_f(v) \geq u_g(v),$$

Therefore, denoting the set of utility-maximizing entries for v in $\mathcal{M}_i \cup \{z_i\}$ by M, we have that the set of utility-maximizing entries for v in $\mathcal{M}_i^1$ is $M^1 \triangleq \big\{\beta(f) \mid f \in M\big\}$. As β is strictly monotone in the price coordinate, and since e is an entry in M with maximal price (indeed, e is not surpassed by z_i in utility for v by Sublemma 2.4.6(a), and in price as $p \geq b_i$ by definition), we have that $\beta(e)$ is an entry in M^1 with maximal price, as required.

We proceed to Part b. For every $v \in W_i$, denoting the menu entry of choice of v from $\mathcal{M}$ by $e = (x_i, \vec{0}_{-i}; p) \in \mathcal{M}_i$, we have by Part a that the payment that $\mathcal{M}_i^1$ extracts from a buyer of type $\alpha(v)$ is the price of $\beta(e)$, which equals $p - b_i$, i.e., precisely the payment that $\mathcal{M}$ extracts from a buyer of type v, minus b_i. By definition of F_i^1, the proof is therefore complete. □

Proof of Sublemma 2.4.2. We start with Part a. By Sublemma 2.4.6(b), v strictly prefers z'_i over every menu entry in $\mathcal{M} \setminus \mathcal{M}_i$ (which, by Sublemma 2.4.4(c), also contains $\mathcal{M}_j$ for every $j \in [n]$ s.t. $F(W_j) > 0$ and $s_j = 1$); by Sublemma 2.4.6(c), v strictly prefers z'_i over z'_j and over o'_j for every $j \in [n] \setminus \{i\}$ s.t. $F(W_j) > 0$ and $s_j < 1$. Therefore, by definition of $\mathcal{M}'$, the menu entry of choice of v from $\mathcal{M}'$ is either z'_i or o'_i. (By definition, the menu entry of choice of $\alpha(v)$ from $\mathcal{M}_i^*$ is either $(0;0)$ or $(1;c_i)$.) We note that

$$u_{o'_i}(v) = v_i - c_i - b_i + (n-1)\cdot H + 1 = (v_i\cdot s_i - b_i + (n-1)\cdot H + 1) + (v_i\cdot(1-s_i) - c_i) = u_{z'_i}(v) + u_{(1;c_i)}\big(\alpha(v)\big),$$

and so v weakly (resp. strictly) prefers o'_i over z'_i iff $\alpha(v)$ weakly (resp. strictly) prefers $(1;c_i)$ over $(0;0)$, completing the proof of Part a as β is strictly monotone in the price coordinate.

We proceed to Part b. For every $v \in W_i$, we have by Part a that the payment that $\mathcal{M}'$ extracts from a buyer of type v equals precisely the payment that $\mathcal{M}_i^*$ extracts from a buyer of type $\alpha(v)$, plus $b_i - (n-1)\cdot H - 1$. By definition of F_i^1, the proof is therefore complete. □

Proof of Sublemma 2.4.3. We start with Part a. By Sublemma 2.4.2(b), the definition of $\mathcal{M}_i^*$ as the revenue-maximizing auction for F_i^1, and Sublemma 2.4.1(b), we have

$$\begin{aligned}\mathcal{R}ev_{\mathcal{M}'}(F|_{W_i}) &= \mathcal{R}ev_{\mathcal{M}_i^*}(F_i^1) + b_i - (n-1)\cdot H - 1 \geq \mathcal{R}ev_{\mathcal{M}_i^1}(F_i^1) + b_i - (n-1)\cdot H - 1 = \\ &= \mathcal{R}ev_{\mathcal{M}}(F|_{W_i}) - (n-1)\cdot H - 1 \geq \mathcal{R}ev_{\mathcal{M}}(F|_{W_i}) - \varepsilon\cdot E \geq (1-\varepsilon)\cdot\mathcal{R}ev_{\mathcal{M}}(F|_{W_i}),\end{aligned}$$

where the last inequality is since by definition of W_i we have that $\mathcal{R}ev_{\mathcal{M}}(F|_{W_i}) \geq E$.

We move on to Part b. Since $s_i = 1$, by Sublemma 2.4.4(c) we have that $\mathcal{M}_i = \{z_i\}$, and so by definition, z_i is the menu entry of choice of every $v \in W_i$ from $\mathcal{M}$ (which, by Sublemma 2.4.4(c), also contains $\mathcal{M}_j$ for every $j \in [n] \setminus \{i\}$ s.t. $F(W_j) > 0$ and $s_j = 1$). As by definition $z_i \in \mathcal{M}'$, it is enough to show that v strictly prefers z_i over all menu entries in $\mathcal{M}' \setminus \mathcal{M}$. Indeed, by Sublemma 2.4.6(c), v strictly prefers z_i over both z'_j and o'_j, for every $j \in [n]$ s.t. $F(W_j) > 0$ and $s_j < 1$ (note that $j \neq i$ since $s_j < 1 = s_i$), i.e., over every menu entry in $\mathcal{M}' \setminus \mathcal{M}$, and so the proof of Part b is complete.

We conclude by proving Part c. Let $v \in P$ and let $e = (\vec{x}; p) \in \mathcal{M}$ be the menu entry of choice of v from $\mathcal{M}$. By definition of P and since $\mathcal{M}$ is E-exclusive, we have that $p \leq E$. Since by definition we have that $e \in \mathcal{M} \cap \big([0,1]^n \times [0,E]\big) \subseteq \mathcal{M}'$, and since the price of every menu entry in $\mathcal{M}' \setminus \mathcal{M}$ is at least $\min_{i \in [n]: F(W_i) > 0 \ \& \ s_i < 1} b_i - (n-1) \cdot H - 1 \geq E - (n-1) \cdot H - 1$, we have that the payment that $\mathcal{M}'$ extracts from a buyer of type v is either $p > (1-\varepsilon) \cdot p$, or at least $E - (n-1) \cdot H - 1 \geq (1-\varepsilon) \cdot E \geq (1-\varepsilon) \cdot p$, and the proof of Part c is complete. ☐

2.6 Discretizing the Cheap Part of the Menu

As outlined above, our fourth and final step toward proving Theorem 1.1, which we take in this section, shows that in any auction over some exclusively unbounded distribution, the "cheap" part of the menu can be simplified without significant loss in revenue and without increasing the number of menu entries in the expensive part, so that (if both the parameter H defining the exclusive unboundedness of the distribution and the parameter E defining the cheap part of the menu are polynomial in n and $1/\varepsilon$) it contains at most $(n/\varepsilon)^{O(n)}$ menu entries. This step is formalized by Lemma 2.5.

Lemma 2.5. *Let $n \in \mathbb{N}$, let $H \in \mathbb{R}_+$, let $E \in \mathbb{R}_+$, and let $\varepsilon \in (0,1)$. For every $F \in \Delta(\mathcal{EU}^n_H)$ and for every IC and IR n-item auction $\mathcal{M}$, there exists an IC and IR n-item auction $\mathcal{M}'$ such that all of the following hold.*

- *$\mathcal{R}ev_{\mathcal{M}'}(F) \geq (1-\varepsilon) \cdot \mathcal{R}ev_{\mathcal{M}}(F) - \varepsilon$.*
- *The menu entries that cost more than $(1-\varepsilon) \cdot E$ in $\mathcal{M}'$ are precisely the menu entries that cost more than E in $\mathcal{M}$, each given a multiplicative price discount of $(1-\varepsilon)$. In particular, there are as many menu entries that cost more than $(1-\varepsilon) \cdot E$ in $\mathcal{M}'$ as there are that cost more than E in $\mathcal{M}$.*
- *There are less than $n \cdot \left\lceil \frac{n \cdot H}{\varepsilon^2} + 1 \right\rceil^{n-1} \cdot \left\lceil \frac{n \cdot (1-\varepsilon) \cdot E}{\varepsilon^2} + 1 \right\rceil$ entries in $\mathcal{M}'$ that cost at most $(1-\varepsilon) \cdot E$.*

Proof. We start by defining an "interim" auction $\mathcal{M}''$ that will help us define the required auction $\mathcal{M}'$. While $\mathcal{M}''$ may contain infinitely many menu entries, we will show that w.l.o.g. only finitely many of these menu entries are in fact chosen by any buyer type; consequently, we will derive the required auction $\mathcal{M}'$ by defining it to be the auction offering precisely these menu entries.

Let $X \triangleq \left\lceil \frac{n \cdot H}{\varepsilon^2} \right\rceil$ and $P \triangleq \left\lceil \frac{n \cdot (1-\varepsilon) \cdot E}{\varepsilon^2} \right\rceil$. Let $\chi \triangleq 1/X \leq \frac{\varepsilon^2}{n \cdot H}$ and $\psi \triangleq \frac{(1-\varepsilon) \cdot E}{P} \leq \frac{\varepsilon^2}{n}$. For every $\delta > 0$ and for every real number $r \in \mathbb{R}_+$, we denote the rounding-down of r to the "δ-grid" by

$$[r]_\delta \triangleq \delta \cdot \lfloor r/\delta \rfloor.$$

We furthermore denote the coordinate-wise rounding-down of every vector $\vec{r} = (r_1, \ldots, r_n) \in \mathbb{R}^n_+$ to the δ-grid by $[\vec{r}]_\delta \triangleq \big([r_1]_\delta, [r_2]_\delta, \ldots, [r_n]_\delta\big)$. We construct a new IC and IR auction $\mathcal{M}''$ as follows:

- For every menu entry $e = (\vec{x}; p) \in \mathcal{M}$ with $p \leq E$, we define $p' \triangleq (1-\varepsilon) \cdot p$ (the price of e, after a slight multiplicative discount of $(1-\varepsilon)$), and add the following n menu entries to $\mathcal{M}''$:

$\big(x_1, ([\vec{x}]_\chi)_{-1}, [p']_\psi\big), \big(x_2, ([\vec{x}]_\chi)_{-2}, [p']_\psi\big), \cdots, \big(x_n, ([\vec{x}]_\chi)_{-n}; [p']_\psi\big)$; each of these menu entries is a modified version of a discounted e (i.e., e with the price modified to p') that rounds-down all but one of the allocation probabilities to the χ-grid, and rounds-down the (discounted) price to the ψ-grid.

- For every menu entry $e = (\vec{x}; p) \in \mathcal{M}$ with $p > E$, we add the menu entry $(\vec{x}; p')$ to $\mathcal{M}''$, where once again $p' \triangleq (1-\varepsilon) \cdot p$; this menu entry is a discounted version of e, by the same slight multiplicative discount of $(1-\varepsilon)$ as above (but without any rounding).

By slight abuse of notation, we write $[r]_0 \triangleq r$ for every $r \in \mathbb{R}_+$, and $[\vec{r}]_0 \triangleq r$ for every $\vec{r} \in \mathbb{R}^n_+$; using this notation, we note that every menu entry that we have added to $\mathcal{M}''$ is of the form $\big(x_i, ([\vec{x}]_{\chi'})_{-i}; [(1-\varepsilon) \cdot p]_{\psi'}\big)$, for some $(\vec{x}; p) \in \mathcal{M}$, $i \in [n]$, and $(\chi', \psi') \in \big\{(\chi, \psi), (0,0)\big\}$. Finally, we define $\mathcal{M}''$ to be the closure of the set of menu entries that we have just added to it.[19] We note that since $\mathcal{M}$ is IR, it contains the menu entry $(\vec{0}; 0)$, and therefore (as $[0]_\delta = 0$ for every δ) so does $\mathcal{M}''$, and hence $\mathcal{M}''$ is IR as well.

We claim that $\mathcal{R}ev_{\mathcal{M}''}(F) \geq (1-\varepsilon) \cdot \mathcal{R}ev_{\mathcal{M}}(F) - \varepsilon$. Indeed, let us compare the payments that $\mathcal{M}''$ and $\mathcal{M}$ extract from a buyer of each type $v = (v_1, \ldots, v_n) \in \mathcal{EU}^n_H$. Let $e = (\vec{x}; p) \in \mathcal{M}$ be the menu entry of choice of v from $\mathcal{M}$. As $v \in \mathcal{EU}^n_H$, there exists $i \in [n]$ s.t. $v_k \leq H$ for every $k \in [n] \setminus \{i\}$. Let $e' \triangleq \big(x_i, ([\vec{x}]_{\chi'})_{-i}; [(1-\varepsilon) \cdot p]_{\psi'}\big) \in \mathcal{M}''$ be the menu entry in $\mathcal{M}''$ corresponding to e that does not round the allocation probability of item i (if $p \leq E$, then $(\chi', \psi') = (\chi, \psi)$; otherwise, $(\chi', \psi') = (0,0)$). We claim that v weakly prefers e' over all menu entries $f' = \big(y_j, ([\vec{y}]_{\chi''})_{-j}; [(1-\varepsilon) \cdot q]_{\psi''}\big) \in \mathcal{M}''$ with $q < p - \varepsilon$ (where $j \in [n]$, $f = (\vec{y}; q) \in \mathcal{M}$, and $(\chi'', \psi'') \in \big\{(\chi, \psi), (0,0)\big\}$).[20] Indeed, we have that:

$$u_{e'}(v) = v_i \cdot x_i + \sum_{k \in [n] \setminus \{i\}} v_k \cdot [x_k]_{\chi'} - [(1-\varepsilon) \cdot p]_{\psi'} \geq \sum_{k=1}^{n} v_k \cdot x_k - (n-1) \cdot H \cdot \chi' - (1-\varepsilon) \cdot p =$$

$$= u_e(v) - (n-1) \cdot H \cdot \chi' + \varepsilon \cdot p \geq u_f(v) - (n-1) \cdot H \cdot \chi' + \varepsilon \cdot p = \sum_{k=1}^{n} v_k \cdot y_k - q - (n-1) \cdot H \cdot \chi' + \varepsilon \cdot p \geq$$

$$\geq v_j \cdot y_j + \sum_{k \in [n] \setminus \{j\}} v_k \cdot [y_k]_{\chi''} - [(1-\varepsilon) \cdot q]_{\psi''} - \varepsilon \cdot q - \psi'' - (n-1) \cdot H \cdot \chi' + \varepsilon \cdot p =$$

$$= u_{f'}(v) - \varepsilon \cdot q - \psi'' - (n-1) \cdot H \cdot \chi' + \varepsilon \cdot p = u_{f'}(v) + \varepsilon \cdot (p-q) - \psi'' - (n-1) \cdot H \cdot \chi' \geq$$

$$\geq u_{f'}(v) + \varepsilon \cdot (p-q) - \tfrac{\varepsilon^2}{n} - (n-1) \cdot \tfrac{\varepsilon^2}{n} = u_{f'}(v) + \varepsilon \cdot (p-q) - \varepsilon^2 > u_{f'}(v) + \varepsilon^2 - \varepsilon^2 = u_{f'}(v).$$

Therefore, the price of the menu entry chosen by v from $\mathcal{M}''$ is at least $(1-\varepsilon)(p-\varepsilon) - \psi \geq (1-\varepsilon) \cdot (p-\varepsilon) - \varepsilon^2 = (1-\varepsilon) \cdot p - \varepsilon$, and so the payment extracted from a buyer of type v in $\mathcal{M}''$ compared to $\mathcal{M}$ decreases by at most a multiplicative factor of $(1-\varepsilon)$ followed by an additive decrease of at most ε, so overall we obtain $\mathcal{R}ev_{\mathcal{M}''}(F) \geq (1-\varepsilon) \cdot \mathcal{R}ev_{\mathcal{M}}(F) - \varepsilon$, as claimed.

As mentioned above, while $\mathcal{M}''$ may contain infinitely many menu entries, we now show that w.l.o.g. only finitely many of these menu entries are in fact chosen by any buyer type, and define the auction $\mathcal{M}'$ to be the auction offering precisely these menu entries. Let $L \triangleq \bigcup_{i \in [n]} \big(\{i\} \times \{0, \chi, 2 \cdot \chi, \ldots, X \cdot \chi\}^{[n] \setminus \{i\}}\big) \times \big\{0, \psi, 2 \cdot \psi, \ldots, P \cdot \psi\big\} \subset \bigcup_{i=1}^{n} \big(\{i\} \times [0,1]^{[n] \setminus \{i\}}\big) \times [0, E]$. By definition, every menu entry in $\mathcal{M}''$ that costs at most $(1-\varepsilon) \cdot E$ is of the form $(x_i, \vec{x}_{-i}; p)$ for some $(i, \vec{x}_{-i}, p) \in L$

[19] As before, taking the closure ensures that a utility-maximizing entry with maximal price exists for every buyer type. See Appendix A.1 for a more details.

[20] As we show weak preference, and as q is defined via a strict inequality, by continuity of the utility function correctness of the claim for all f' before taking the closure of $\mathcal{M}''$ implies correctness for all f' in the closure as well.

and $x_i \in [0,1]$. For every $(i,\vec{x}_{-i},p) \in L$, we set $S_{(i,\vec{x}_{-i},p)} \triangleq \{x_i \mid (x_i,\vec{x}_{-i};p) \in \mathcal{M}''\} \subseteq [0,1]$, and if $S_{(i,\vec{x}_{-i},p)} \neq \emptyset$, we also set $s_{(i,\vec{x}_{-i},p)} \triangleq \operatorname{Max} S_{(i,\vec{x}_{-i},p)}$.[21] We define a new IC and IR auction $\mathcal{M}'$ as follows:

$$\mathcal{M}' \triangleq \Big\{\big(s_{(i,\vec{x}_{-i},p)},\vec{x}_{-i};p\big) \;\Big|\; (i,\vec{x}_{-i},p) \in L \;\;\&\;\; S_{(i,\vec{x}_{-i},p)} \neq \emptyset\Big\} \;\cup\; \Big\{\big(\vec{x};(1-\varepsilon)\cdot p\big) \;\Big|\; (\vec{x};p) \in \mathcal{M} \;\;\&\;\; p > E\Big\}.$$

By definition, $\mathcal{M}'$ is a subset of $\mathcal{M}''$, obtained by removing from $\mathcal{M}''$ only menu entries that w.l.o.g. no buyer type chooses. (Indeed, a buyer type that chooses from $\mathcal{M}''$ some entry $(x;p) \in \mathcal{M}''$ weakly prefers $(s_{(i,\vec{x}_{-i},p)},\vec{x}_{-i};p)$ over $(x_i,\vec{x}_{-i};p) = (x;p)$, since by definition $x_i \in S_{(i,\vec{x}_{-i},p)}$. As both of these entries have the same price, we can assume w.l.o.g. that this buyer type in fact chooses $\big(s_{(i,\vec{x}_{-i},p)},\vec{x}_{-i};p\big)$ from $\mathcal{M}''$.) Therefore, every buyer type chooses from $\mathcal{M}'$ the same menu entry (or, at least, a menu entry with the same price) as from $\mathcal{M}''$, and so $\mathcal{R}ev_{\mathcal{M}'}(F) = \mathcal{R}ev_{\mathcal{M}''}(F) \geq (1-\varepsilon)\cdot\mathcal{R}ev_{\mathcal{M}}(F) - \varepsilon$.

We conclude the proof by noting that indeed, there are less than $|L| = n\cdot(X+1)^{n-1}\cdot(P+1) = n\cdot\left\lceil\frac{n\cdot H}{\varepsilon^2}+1\right\rceil^{n-1}\cdot\left\lceil\frac{n\cdot(1-\varepsilon)\cdot E}{\varepsilon^2}+1\right\rceil$ menu entries (not including $(\vec{0};0)$) that cost at most $(1-\varepsilon)\cdot E$ in $\mathcal{M}'$, and that the menu entries that cost more than $(1-\varepsilon)\cdot E$ in $\mathcal{M}'$ are precisely the menu entries that cost more than E in $\mathcal{M}$, each given a multiplicative price discount of $(1-\varepsilon)$. □

2.7 Connecting the Dots

We are now ready to "connect the dots" and use Lemmas 2.2 through 2.5 to prove Theorem 1.1, which states that $C(n,\varepsilon)$ is finite for every number of items $n \in \mathbb{N}$ and $\varepsilon > 0$, and moreover, that $C(n,\varepsilon) \leq ({n}/{\varepsilon})^{O(n)}$.

Proof of Theorem 1.1. Let $n \in \mathbb{N}$, let $\varepsilon \in (0,1)$, and let $F = F_1 \times \cdots \times F_n \in \Delta(\mathbb{R}_+)^n$. If $n = 1$, then by the theorem of Myerson (1981), we are done. If $\operatorname{Max}_{i\in[n]} \mathcal{R}ev(F_i) = \infty$, i.e., if there exists $i \in [n]$ s.t. $\mathcal{R}ev(F_i) = \infty$, then by the same theorem of Myerson (1981), we are done in this case as well, as a single menu entry suffices to get arbitrarily high revenue;[22] otherwise, by Theorem 2.1, $\mathcal{R}ev(F) < \infty$. If $\operatorname{Max}_{i\in[n]} \mathcal{R}ev(F_i) = 0$, then the valuation of each item is 0 with probability 1 and so $\mathcal{R}ev(F) = 0$ and we are done as well. Assume, therefore, that $\operatorname{Max}_{i\in[n]} \mathcal{R}ev(F_i) \in (0,\infty)$ and set $\tilde{\varepsilon} \triangleq {\varepsilon}/{6}$. By scaling the currency we assume w.l.o.g. that $R \triangleq \operatorname{Max}_{i\in[n]} \mathcal{R}ev(F_i) = (1-\tilde{\varepsilon})^{-5} < ({6}/{5})^5 < {5}/{2}$. (This indeed is w.l.o.g. as our goal is to prove a multiplicative approximation.)

Let $H \triangleq \frac{2\cdot n\cdot(n-1)\cdot R}{\tilde{\varepsilon}}$, and let $E \triangleq \frac{4\cdot(n-1)\cdot H}{\tilde{\varepsilon}^2}$. By Lemma 2.2(a), $F(\mathcal{EU}_H^n) > 0$ and so $F|_{\mathcal{EU}_H^n}$ is well defined. Furthermore (e.g., by Sublemma 2.2.1(a)), $F(\mathcal{EU}_H^n)\cdot\mathcal{R}ev(F|_{\mathcal{EU}_H^n}) \leq \mathcal{R}ev(F) < \infty$, and so, as $F(\mathcal{EU}_H^n) > 0$, we have that $\mathcal{R}ev(F|_{\mathcal{EU}_H^n}) < \infty$. Therefore, by definition of $\mathcal{R}ev$, there exists an IC and IR n-item auction $\mathcal{M}$ s.t. $\mathcal{R}ev_{\mathcal{M}}(F|_{\mathcal{EU}_H^n}) \geq (1-\tilde{\varepsilon})\cdot\mathcal{R}ev(F|_{\mathcal{EU}_H^n})$. By Lemma 2.3, there exists an E-exclusive IC and IR n-item auction $\mathcal{M}'$ s.t. $\mathcal{R}ev_{\mathcal{M}'}(F|_{\mathcal{EU}_H^n}) \geq (1-\tilde{\varepsilon})\cdot\mathcal{R}ev_{\mathcal{M}}(F|_{\mathcal{EU}_H^n})$. By Lemma 2.4, there exists an IC and IR n-item auction $\mathcal{M}''$, with at most $2n$ menu entries that cost more than E, s.t. $\mathcal{R}ev_{\mathcal{M}''}(F|_{\mathcal{EU}_H^n}) \geq (1-\tilde{\varepsilon})\cdot\mathcal{R}ev_{\mathcal{M}'}(F|_{\mathcal{EU}_H^n})$. By Lemma 2.5, there exists an IC and IR n-item auction $\mathcal{M}'''$, with at most $n\cdot\left\lceil\frac{n\cdot H}{\tilde{\varepsilon}^2}+1\right\rceil^{n-1}\cdot\left\lceil\frac{n\cdot(1-\tilde{\varepsilon})\cdot E}{\tilde{\varepsilon}^2}+1\right\rceil$ menu entries that cost at most $(1-\tilde{\varepsilon})\cdot E$ and (by definition of $\mathcal{M}''$) at most $2n$ menu entries that cost more than $(1-\tilde{\varepsilon})\cdot E$, s.t. $\mathcal{R}ev_{\mathcal{M}'''}(F|_{\mathcal{EU}_H^n}) \geq (1-\tilde{\varepsilon})\cdot\mathcal{R}ev_{\mathcal{M}''}(F|_{\mathcal{EU}_H^n}) - \tilde{\varepsilon}$. By definitions of $\mathcal{M}'''$, $\mathcal{M}''$, $\mathcal{M}'$, and $\mathcal{M}$, we have that

[21] Since $\mathcal{M}''$ is closed, $S_{(i,\vec{x}_{-i},p)}$ is compact.

[22] By convention, $(1-\varepsilon)\cdot\infty = \infty$.

$$\mathcal{R}ev_{\mathcal{M}'''}(F|_{\mathcal{EU}^n_H}) \geq (1-\tilde{\varepsilon})\cdot\mathcal{R}ev_{\mathcal{M}''}(F|_{\mathcal{EU}^n_H}) - \tilde{\varepsilon} \geq (1-\tilde{\varepsilon})^2\cdot\mathcal{R}ev_{\mathcal{M}'}(F|_{\mathcal{EU}^n_H}) - \tilde{\varepsilon} \geq$$
$$\geq (1-\tilde{\varepsilon})^3\cdot\mathcal{R}ev_{\mathcal{M}}(F|_{\mathcal{EU}^n_H}) - \tilde{\varepsilon} \geq (1-\tilde{\varepsilon})^4\cdot\mathcal{R}ev(F|_{\mathcal{EU}^n_H}) - \tilde{\varepsilon} \geq (1-\tilde{\varepsilon})^5\cdot\mathcal{R}ev(F|_{\mathcal{EU}^n_H}),$$

where the last inequality is since by Lemma 2.2(b), $\mathcal{R}ev(F|_{\mathcal{EU}^n_H}) \geq (1-\tilde{\varepsilon})\cdot\mathcal{R}ev(F) \geq (1-\tilde{\varepsilon})\cdot R = (1-\tilde{\varepsilon})^{-4}$. Therefore, by Lemma 2.2(c), we have that

$$\mathcal{R}ev_{\mathcal{M}'''}(F) \geq (1-\tilde{\varepsilon})^6\cdot\mathcal{R}ev(F) > (1-6\cdot\tilde{\varepsilon})\cdot\mathcal{R}ev(F) = (1-\varepsilon)\cdot\mathcal{R}ev(F).$$

We conclude the proof by noting that the number of menu entries (not including $(\vec{0};0)$ in $\mathcal{M}'''$ is less than

$$n\cdot\left\lceil\frac{n\cdot H}{\tilde{\varepsilon}^2}+1\right\rceil^{n-1}\cdot\left\lceil\frac{n\cdot(1-\tilde{\varepsilon})\cdot E}{\tilde{\varepsilon}^2}+1\right\rceil+2n =$$
$$n\cdot\left\lceil\frac{2\cdot n^2\cdot(n-1)\cdot R}{\tilde{\varepsilon}^3}+1\right\rceil^{n-1}\cdot\left\lceil\frac{8\cdot n^2\cdot(n-1)^2\cdot(1-\tilde{\varepsilon})\cdot R}{\tilde{\varepsilon}^5}+1\right\rceil+2n \leq$$
$$n\cdot\left\lceil\frac{5\cdot n^2\cdot(n-1)}{\tilde{\varepsilon}^3}+1\right\rceil^{n-1}\cdot\left\lceil\frac{20\cdot n^2\cdot(n-1)^2\cdot(1-\tilde{\varepsilon})}{\tilde{\varepsilon}^5}+1\right\rceil+2n \leq ({}^{n}/_{\varepsilon})^{O(n)}. \qquad \square$$

We note that, as mentioned in Section 1, it is possible to strengthen the upper bound on $C(n,\varepsilon)$ from $({}^{n}/_{\varepsilon})^{O(n)}$ to $({}^{\log n}/_{\varepsilon})^{O(n)}$, but at a cost of a messier presentation. As this stronger bound is still exponential in n, we have decided not to present the more involved proof for this bound. Nonetheless, we comment that the only change required to prove this stronger bound is to Lemma 2.5, where by following along the same lines as the proof above (i.e., discretizing each cheap menu entry n times, each time rounding-down all but one allocation probability, etc.), but utilizing discretization techniques of Hart and Nisan (2013) and Dughmi et al. (2014) to round to a more carefully chosen grid, one may show an improved upper bound for the number of menu entries into which it is possible to discretize the cheap part of the menu in Lemma 2.5. (Note that other than for the cheap part of the menu, the remaining lemmas show that we only need at most $2n$ more entries for the expensive part of the menu, and therefore no modification is required to these lemmas, which either way contain most of the conceptual and technical "beef" of the proof, in order to show the stronger upper bound.)

2.8 Arbitrary Exclusively Unbounded Type Distributions

We emphasize once more that apart from the first step (Lemma 2.2) of the proof of Theorem 1.1, the remaining three steps (i.e., Lemmas 2.3 through 2.5) hold for any exclusively unbounded type distribution, and not merely one obtained by conditioning product distributions upon $\mathcal{EU}^n_H$. Therefore, these steps allow us to derive results also for such distributions. Indeed, a proof similar to that of Theorem 1.1 (only without using Lemma 2.2) yields:

Proposition 2.6. *For every $n\in\mathbb{N}$, $H\in\mathbb{R}_+$, and $\varepsilon>0$, there exists $C=C(n,H,\varepsilon)\leq\left(\frac{n\cdot H}{\varepsilon}\right)^{O(n)}$ such that for every (possibly even highly correlated) $F\in\Delta(\mathcal{EU}^n_H)$, we have that $\mathcal{R}ev_C(F)\geq(1-\varepsilon)\cdot\mathcal{R}ev(F)-\varepsilon$.*

Proposition 2.6 states that a menu size of $\left(\frac{n\cdot H}{\varepsilon}\right)^{O(n)}$ guarantees a multiplicative approximation of at least $(1-\varepsilon)$ for the revenue from any exclusively unbounded distribution (e.g., in particular for any distribution over $\mathbb{R}_+\times[1,H]^{n-1}$), thereby generalizing a similar result by Hart and Nisan (2013) that shows the same upper bound of $\left(\frac{n\cdot H}{\varepsilon}\right)^{O(n)}$, however only for distributions where the valuations of all items are bounded (i.e., distributions over $[1,H]^n$). We note that Dughmi et al. (2014) (following Hart and Nisan, 2013, who showed this for two items) improve the upper bound that Hart and Nisan (2013) show for distributions over $[1,H]^n$ to $\left(\frac{\log n+\log H}{\varepsilon}\right)^{O(n)}$. Similarly to the discussion concluding Section 2.7, the techniques of Hart and Nisan (2013) and Dughmi et al. (2014) can be built upon in precisely the same way to improve the upper bound in Proposition 2.6 to the same upper bound of $\left(\frac{\log n+\log H}{\varepsilon}\right)^{O(n)}$, therefore generalizing the result of Dughmi et al. (2014) from bounded to exclusively unbounded valuations as well.

Finally, we note that an exclusively unbounded valuation space is essentially the maximal valuation space for which a result along the lines of Proposition 2.6 can be shown, since Hart and Nisan (2013) show that for two items with unbounded valuations (and as a result, for any number of items, where the valuations of at least two items are unbounded), no finite menu size can guarantee any fixed fraction of the optimal revenue. Proposition 2.6 in fact shows that the use of distributions where the valuations of the two items grow arbitrarily large together (like the distribution used in the proof of Hart and Nisan, 2013) cannot be avoided in any proof of this impossibility result.

3 A Small Menu for Item Pricing

In this section, we prove Theorem 1.2, which states that for every $\varepsilon > 0$, there exists $d = d(\varepsilon)$ such that for every number of items $n \in \mathbb{N}$ and $F_1, F_2, \ldots, F_n \in \Delta(\mathbb{R}_+)$, we have for $C = n^d$ that $\mathcal{R}ev_C(F_1 \times \cdots \times F_n) \geq (1-\varepsilon)\cdot \mathcal{S}Rev(F_1 \times \cdots \times F_n)$. Concretely, we show that the theorem holds for $d \leq O(\varepsilon^{-5})$.

Definition 3.1 ($\mathcal{S}Rev$)**.** Let $n \in \mathbb{N}$ be a number of items and let $F = F_1 \times F_2 \times \cdots \times F_n \in \Delta(\mathbb{R}_+)^n$ be a product distribution over $\mathbb{R}_+^n$. We denote the (expected) revenue obtainable from F by selling each of the n items separately via a revenue-maximizing auction by

$$\mathcal{S}Rev(F) = \mathcal{S}Rev(F_1 \times F_2 \times \cdots \times F_n) \triangleq \sum_{i=1}^{n} \mathcal{R}ev(F_i).$$

Remark 3.1. By the theorem of Myerson (1981), the menu size (not including $(\vec{0};0)$) of the auction obtaining revenue $\mathcal{S}Rev(F)$ by selling each of the n items separately via a revenue-maximizing auction is at most $2^n - 1$.

The main idea underlying the proof of Theorem 1.2, which we give in full detail below, is that instead of holding n separate auctions (one for each item), which may result in an exponential-size menu, we hold exponentially fewer separate auctions, which result in a polynomial-size menu. Recall that for every item i, by the theorem of Myerson (1981), the optimal separate-selling revenue $\mathcal{R}ev(F_i)$ can be obtained via a take-it-or-leave-it offer for selling item i (with probability 1) for a certain price that we henceforth denote by c_i. Let p_i be the probability item i is sold (if it is offered for the price c_i). The separate auctions that we perform are as follows:

- For each item for which c_i is very small (say, at most an ε/n fraction) compared to $\mathcal{S}Rev(F)$, the optimal (separate-selling) revenue $\mathcal{R}ev(F_i)$ is very small compared to $\mathcal{S}Rev(F)$ as well, and therefore the sum of the optimal revenues from all of these items is small compared to $\mathcal{S}Rev(F)$. Therefore, we allow ourselves to not sell any of these items at all (or alternatively, give them to the buyer for free).

- For each item for which c_i is very large (say, by at least a factor of n/ε) compared to $\mathcal{R}ev(F_i)$, the probability that the item is sold (when selling items separately) is small. Hence, the probability that two or more of these items are sold is small. Therefore, we can afford to allow the buyer to only buy at most one of these items (for the same price c_i as when selling each item separately) without incurring a significant loss in revenue. We therefore offer the buyer at most $n+1$ choices for these items.

- We partition the remaining items (those with "nonextreme" values for c_i) into $O(\log n)$ many bundles, each to be offered (via a separate auction) to the buyer for a take-it-or-leave-it price described below, where the ratio between the (optimal separate-selling) prices of any two items in a single bundle is small. Moreover, we show that this can be done s.t. each such bundle has $\sum p_i$ either very large or very small.
  - For each bundle with $\sum p_i$ very large, we show that the buyer's valuation of the bundle is tightly concentrated, allowing us to extract almost all of this valuation by offering the bundle for a price slightly below the expectation of this valuation.
  - For each bundle with $\sum p_i$ very small, somewhat similarly to the case of high-costing items above, we show that the probability for two or more of the items in the bundle to be sold (when selling separately) is small. Since the prices of all of these items are similar, instead of allowing the buyer to buy at most one of these items (for its separate-selling price c_i) as in the case of the high-costing items above (thereby offering as many choices to the buyer for these items as there are items in the bundle), we simply offer the entire bundle for the cheapest (optimal separate-selling) price of any of the items in the bundle. We can afford to do so without incurring a significant loss in revenue since all of the items in the bundle have similar prices.

  As each of these $O(\log n)$ many bundles is offered via a take-it-or-leave-it price, we therefore offer the buyer altogether at most poly(n) choices for (buying any subset of) these bundles.

Proof of Theorem 1.2. Let $\varepsilon \in (0,1)$. We prove the theorem for $d \triangleq 8194/\varepsilon^5$,[23] and note that the auction that we construct is deterministic.

Let $n \in \mathbb{N}$, and let $F = F_1 \times F_2 \times \cdots \times F_n \in \Delta(\mathbb{R}_+)^n$. For every $i \in [n]$, denote by $r_i \triangleq \mathcal{R}ev(F_i)$ the maximal revenue obtainable from F_i. By definition, $\mathcal{S}Rev(F) = \mathcal{S}Rev(F_1 \times \cdots \times F_n) = \sum_{i=1}^n r_i$. If $n=1$, then by Remark 3.1 we are done, as selling the single item "separately" (and obtaining revenue $\mathcal{S}Rev(F)$) requires at most $1 = n^d$ menu entry (not including $(\vec{0};0)$). If $2 \leq n < 4/\varepsilon$, then we are similarly done, as selling each item separately (and obtaining revenue $\mathcal{S}Rev(F)$) requires at most $2^n - 1 < n^{4/\varepsilon} < n^d$ menu entries. Assume henceforth, therefore, that $n \geq 4/\varepsilon$. If $\mathrm{Max}_{i\in[n]}\, r_i = \infty$, i.e., if there exists $i \in [n]$ s.t. $r_i = \infty$, then we are done as well, as by the theorem of Myerson (1981) a single menu entry suffices to get arbitrarily high revenue.[24] If $\mathrm{Max}_{i\in[n]}\, r_i = 0$, then there is nothing to prove. Assume, therefore, that $\mathrm{Max}_{i\in[n]}\, r_i \in (0,\infty)$. By scaling the currency we assume w.l.o.g. that $\mathrm{Max}_{i\in[n]}\, r_i = 1$. (This indeed is w.l.o.g. as our goal is to prove a multiplicative approximation.) Therefore, $\mathcal{S}Rev(F) \geq 1$ (and also $\mathcal{S}Rev(F) \leq n$).

For every $i \in [n]$, by the theorem of Myerson (1981), the revenue r_i can be obtained from F_i via a take-it-or-leave-it offer for selling item i (with probability 1) for a certain price that we henceforth denote by c_i. Let $p_i \triangleq \mathbb{P}_{v_i \sim F_i}\left[v_i \geq c_i\right]$ be the probability that c_i is accepted. We note that $r_i = p_i \cdot c_i$. As outlined above, instead of holding n separate auctions (one for each item), which may result in an exponential-size menu, we hold exponentially fewer separate auctions, which results in a polynomial-size menu. To describe these separate auctions, we first partition the n items into

[23] This constant multiplier is used for ease of presentation, and is far from tight.

[24] As already noted above, by convention, $(1-\varepsilon)\cdot\infty = \infty$.

buckets based on the optimal price c_i for each item $i \in [n]$, and then describe the separate auction that we hold for the items in each bucket. Set $\tilde{\varepsilon} \triangleq \varepsilon/4$.

- The *low* bucket L includes all items $i \in [n]$ for which $c_i < \tilde{\varepsilon}/n$.
- The *high* bucket H includes all items $i \in [n]$ for which $c_i \geq n/\tilde{\varepsilon}$.
- Let $m \triangleq \lceil \log_{1+\tilde{\varepsilon}} n/\tilde{\varepsilon} \rceil$. We partition the remaining items (i.e., the items that are not already in the low or the high bucket) into $2m$ *regular* buckets, where the ratio of any two prices in each bucket is less than $1+\tilde{\varepsilon}$. Specifically, for each integer $-m \leq b < m$, regular bucket B_b includes all items $i \in [n] \setminus (L \cup H)$ such that $(1+\tilde{\varepsilon})^b \leq c_i < (1+\tilde{\varepsilon})^{b+1}$.

We now construct our auction by describing the separate auction that we hold for the items in each bucket: (As the buyer's valuation is additive, holding a compound auction comprised of a number of separate IC and IR auctions for pairwise-disjoint sets of items, is itself IC and IR.)

- We bundle all of the items in the low bucket L together, and give this bundle to the buyer for free. (Alternatively, we could give any predefined subset of the low bucket to the buyer for free.)
- For the high bucket H, we offer the buyer the option of purchasing at most one of the items $i \in H$ (to be chosen by the buyer), for the price c_i of that item.
- For a regular bucket B_b, we define $\mu_b \triangleq \sum_{i \in B_b} p_B$. We say that B_b is *dense* if $\mu_b > \tilde{\varepsilon}^{-3}$.
  - If B_b is dense, then we bundle all of the items in the bucket B_b together, and offer this bundle to the buyer for a take-it-or-leave-it price of $(1-\tilde{\varepsilon}) \cdot \mu_b \cdot (1+\tilde{\varepsilon})^b$.
  - Otherwise, i.e., if B_b is not dense, we partition the bucket B_b into bundles $B_b^1, B_b^2, \ldots$, s.t. for each such bundle B_b^j and for each $i \in B_b^j$, we have that $\sum_{k \in B_b^j \setminus \{i\}} p_k \leq \tilde{\varepsilon}$. We offer each such bundle, separately, to the buyer for a take-it-or-leave-it price of $(1+\tilde{\varepsilon})^b$. (So the auction for the items in the bucket B_b is itself comprised of a number of separate auctions, one for each bundle.) By Lemma 3.1 (see below, after this proof; the lemma is applied after scaling by $\tilde{\varepsilon}$), we have that no more than $\lceil \tilde{\varepsilon}^{-4} \rceil$ such bundles are needed when partitioning the bucket B_b.

Altogether, the number of bundles of items from regular buckets that are offered to the buyer in (separate) take-it-or-leave-it auctions is at most

$$\left|\left\{b \in [-m, m-1] \cap \mathbb{Z} \mid \mu_b > \tilde{\varepsilon}^{-3}\right\}\right| + \lceil \tilde{\varepsilon}^{-4} \rceil \cdot \left|\left\{b \in [-m, m-1] \cap \mathbb{Z} \mid \mu_b \leq \tilde{\varepsilon}^{-3}\right\}\right| \leq$$
$$\leq \lceil \tilde{\varepsilon}^{-4} \rceil \cdot 2m = 2 \cdot \lceil \tilde{\varepsilon}^{-4} \rceil \cdot \lceil \log_{1+\tilde{\varepsilon}} n/\tilde{\varepsilon} \rceil = 2 \cdot \lceil \tilde{\varepsilon}^{-4} \rceil \cdot \left\lceil \frac{\log_2 n + \log_2(1/\tilde{\varepsilon})}{\log_2(1+\tilde{\varepsilon})} \right\rceil \leq$$
$$\leq 2 \cdot \left(\tilde{\varepsilon}^{-4} + 1\right) \cdot \left(\frac{2 \cdot \log_2 n}{\tilde{\varepsilon}} + 1\right) \leq 4 \cdot \tilde{\varepsilon}^{-4} \cdot \frac{2 \cdot \log_2 n}{\tilde{\varepsilon}} = \log_2 n \cdot 8/\tilde{\varepsilon}^5.$$

Therefore, recalling that $\tilde{\varepsilon} = \varepsilon/4$, the total number of menu entries in the compound auction (comprised of the separate auctions held for each low, high, or regular bucket as described above) is less than

$$1 \cdot \left(|H|+1\right) \cdot 2^{\log_2 n \cdot 8/\tilde{\varepsilon}^5} = \left(|H|+1\right) \cdot n^{8/\tilde{\varepsilon}^5} \leq (n+1) \cdot n^{8/\tilde{\varepsilon}^5} < n^{8/\tilde{\varepsilon}^5+2} = n^{8192/\varepsilon^5+2} < n^d.$$

It remains to analyze the revenue from the compound auction, which is, by linearity of expectation, the sum of the revenues from the separate auctions held for each of the buckets.

- The revenue from the auction held for the low bucket L is obviously 0. (Note that selling each item from L separately, the revenue would have been $\sum_{i\in L} r_i \leq \sum_{i\in L} c_i < \sum_{i\in L} \tilde{\varepsilon}/n \leq \tilde{\varepsilon}$, and so the loss in revenue is at most an additive $\tilde{\varepsilon}$.)

- We claim that the revenue from the auction held for the high bucket H is at least $(1-\tilde{\varepsilon})\cdot\sum_{i\in H} r_i$. To show this, we note that a specific item $i \in H$ in this bucket is sold whenever the buyer values it by at least c_i (which happens with probability p_i), perhaps unless the buyer values one of the other items $j \in H\setminus\{i\}$ in this bucket by at least c_j. As for every $j \in H\setminus\{i\}$, we have that $r_j \leq 1$ by normalization, and that $c_j \geq n/\tilde{\varepsilon}$ by definition of H, we obtain that $p_j = r_j/c_j \geq \tilde{\varepsilon}/n$, and therefore we have that item i is sold with probability at least
$$p_i \cdot \bigl(1 - \sum_{j\in H\setminus\{i\}} p_j\bigr) \geq p_i \cdot (1-(n-1)\cdot\tilde{\varepsilon}/n) > (1-\tilde{\varepsilon})\cdot p_i.$$
Therefore, the (expected) revenue from the auction held for the high bucket is
$$\sum_{i\in H} \mathbb{P}_{v\sim F}\bigl[\text{item } i \text{ is sold}\bigr]\cdot c_i > \sum_{i\in H}(1-\tilde{\varepsilon})\cdot p_i\cdot c_i = (1-\tilde{\varepsilon})\cdot\sum_{i\in H} r_i.$$
(Note that selling each item from H separately, the revenue would have been $\sum_{i\in H} r_i$, and so the loss in revenue is at most a multiplicative $(1-\tilde{\varepsilon})$ factor.)

- We claim that the revenue from the auction held for a dense regular bucket B_b is at least $(1-\tilde{\varepsilon})^3\cdot\sum_{i\in B_b} r_i$. As the price for which we offer the bundle of all of items the in this bucket is $(1-\tilde{\varepsilon})\cdot\mu_b\cdot(1+\tilde{\varepsilon})^b$, it suffices to show that this bundle is sold with probability at least $(1-\tilde{\varepsilon})$, as this implies that the revenue from this auction is at least
$$(1-\tilde{\varepsilon})\cdot(1-\tilde{\varepsilon})\cdot\mu_b\cdot(1+\tilde{\varepsilon})^b = (1-\tilde{\varepsilon})^2\cdot\sum_{i\in B_b} p_i\cdot(1+\tilde{\varepsilon})^b > (1-\tilde{\varepsilon})^3\cdot\sum_{i\in B_b} p_i\cdot c_i = (1-\tilde{\varepsilon})^3\cdot\sum_{i\in B_b} r_i.$$
For every $i \in B_b$, let X_i be an indicator random variable for the buyer valuing item i by at least c_i; thus, X_i is a Bernoulli variable taking value 1 with probability p_i. Let $X \triangleq \sum_{i\in b} X_i$ and note the expectation of X is $\mu_b > \tilde{\varepsilon}^{-3}$. Clearly, if $X \geq (1-\tilde{\varepsilon})\cdot\mu_b$, then the buyer values the bundle by at least $(1-\tilde{\varepsilon})\cdot\mu_b\cdot\min_{i\in B_b} c_i \geq (1-\tilde{\varepsilon})\cdot\mu_b\cdot(1+\tilde{\varepsilon})^b$, and the bundle is sold. Therefore, to show that the bundle is sold with probability at least $(1-\tilde{\varepsilon})$, it suffices to show that $\mathbb{P}_{v\sim F}\bigl[X<(1-\tilde{\varepsilon})\cdot\mu_b\bigr] \leq \tilde{\varepsilon}$. Denoting the standard deviation of X by σ_b, since $(X_i)_{i\in B_b}$ are independent we have that $\sigma_b = \sqrt{\sum_{i\in B_b} p_i\cdot(1-p_i)} \leq \sqrt{\mu_b}$. By Chebyshev's inequality, we therefore have that
$$\mathbb{P}_{v\sim F}\bigl[X<(1-\tilde{\varepsilon})\cdot\mu_b\bigr] \leq \mathbb{P}_{v\sim F}\bigl[X<\mu_b-(\tilde{\varepsilon}\sqrt{\mu_b})\cdot\sigma_b\bigr] \leq \frac{1}{\tilde{\varepsilon}^2\cdot\mu_b} < \tilde{\varepsilon}^{-2}\cdot\tilde{\varepsilon}^3 = \tilde{\varepsilon},$$
as required. (Note that selling each item from B_b separately, the revenue would have been $\sum_{i\in B_b} r_i$, and so the loss in revenue is at most a multiplicative $(1-\tilde{\varepsilon})^3$ factor, where the triple loss of $(1-\tilde{\varepsilon})$ is once for lower-bounding c_i by $(1+\tilde{\varepsilon})^b$, a second time for multiplying the requested price by $(1-\tilde{\varepsilon})$ in order to guarantee a sale with high probability, and a third time for the $\tilde{\varepsilon}$ probability of not selling the bundle.)

- Finally, we claim that the revenue from the auction held for a bundle B_b^j from a nondense regular bucket B_b is at least $(1-\tilde{\varepsilon})^2\cdot\sum_{i\in B_b^j} r_i$. As the price that we charge for this bundle is no more than the value of c_i for any item $i \in B_b^j$ in the bundle, we note that in particular, the bundle is sold if the buyer values some item $i \in B_b^j$ by at least c_i. Therefore, the probability that this bundle is sold is at least the probability that the buyer values *exactly* one of the items $i \in B_b^j$ by at least c_i, which, by definition of B_b^j, is at least $\sum_{i \in B_b^j} p_i \cdot \big(1 - \sum_{k \in B_b^j \setminus \{i\}} p_k\big) \geq (1-\tilde{\varepsilon}) \cdot \sum_{i \in B_b^j} p_i$. Therefore, the revenue from this auction is at least

$$(1-\tilde{\varepsilon}) \cdot \big(\sum_{i \in B_i^j} p_i\big) \cdot (1+\tilde{\varepsilon})^b > (1-\tilde{\varepsilon})^2 \cdot \sum_{i \in B_b^j} p_i \cdot c_i = (1-\tilde{\varepsilon})^2 \cdot \sum_{i \in B_b^j} r_i.$$

(Note that selling each item from B_b^j separately, the revenue would have been $\sum_{i \in B_b^j} r_i$, and so the loss in revenue is at most a multiplicative $(1-\tilde{\varepsilon})^2$ factor, where the double loss of $(1-\tilde{\varepsilon})$ is once for lower-bounding c_i by $(1+\tilde{\varepsilon})^b$, and a second time for the $\tilde{\varepsilon}$ probability of not selling the bundle.)

The total revenue is, therefore, at least

$$(1-\tilde{\varepsilon}) \cdot \big(\sum_{i \in H} r_i\big) + (1-\tilde{\varepsilon})^3 \cdot \big(\sum_{\substack{i \in B_b: \\ \mu_b > \tilde{\varepsilon}^{-3}}} r_i\big) + (1-\tilde{\varepsilon})^2 \cdot \big(\sum_{\substack{i \in B_b: \\ \mu_b \leq \tilde{\varepsilon}^{-3}}} r_i\big) \geq (1-\tilde{\varepsilon})^3 \cdot \big(\sum_{i \in [n] \setminus L} r_i\big) \geq$$

$$\geq (1-\tilde{\varepsilon})^3 \cdot \big(\mathcal{S}Rev(F) - \tilde{\varepsilon}\big) \geq (1-\tilde{\varepsilon})^4 \cdot \mathcal{S}Rev(F) > (1 - 4 \cdot \tilde{\varepsilon}) \cdot \mathcal{S}Rev(F) = (1-\varepsilon) \cdot \mathcal{S}Rev(F),$$

as required. □

Lemma 3.1. *Let $m \in \mathbb{R}_+$, let $p_1, p_2, \ldots, p_m \in \mathbb{R}$ be strictly positive numbers, and let $s \triangleq \sum_{i=1}^m p_i$. There exists a partition of $[m]$ into at most $\lceil s \rceil$ sets B^j, s.t. for every set B^j in the partition, and for every index $i \in B^j$, we have that $\sum_{k \in B^j \setminus \{i\}} p_k \leq 1$.*

Proof. We prove the claim by induction over $\lceil s \rceil$. We note that the constraint on each set B^j is equivalent to demanding that $\sum_{k \in B^j \setminus \{i\}} p_k \leq 1$ for $i \in \arg\min_{k \in B^j} p_k$.

Assume w.l.o.g. that $p_1 \geq p_2 \geq \cdots \geq p_m$ and let $\ell \triangleq \mathrm{Max}\big\{\ell \in \{0, 1, \ldots, m-1\} \bigm| \sum_{k=1}^{\ell} p_k \leq 1\big\} \in \{0, 1, \ldots, m-1\}$. We define $B^1 \triangleq [\ell+1] \subseteq [m]$. By monotonicity of $(p_k)_{k=1}^m$, we have that $i \triangleq \ell+1 \in \arg\min_{k \in B^1} p_k$, and by definition of ℓ, we have that $\sum_{k \in B^1 \setminus \{i\}} p_k \leq 1$. If $\ell = m-1$ (and so $i = m$), then we are done. Otherwise, by definition, of ℓ we have that $\sum_{k \in B^1} p_k = \sum_{k=1}^{\ell+1} p_k > 1$ (and hence $\lceil s \rceil > 1$), and so $\big\lceil \sum_{k \in [m] \setminus B^1} p_k \big\rceil \leq \lceil s \rceil - 1$. Therefore, by the induction hypothesis, the proof is complete as $[m] \setminus B^1$ can be partitioned into at most $\lceil s \rceil - 1$ sets B^j for $j \geq 2$ s.t. for every set B^j in the partition and for every $i \in B^j$, we have that $\sum_{k \in B^j \setminus \{i\}} p_k \leq 1$. □

We comment that in Lemma 3.1, the bound of $\lceil s \rceil$ on the number of sets in the partition is in fact tight. Indeed, this can be demonstrated by setting $m = 2 \cdot \lceil s \rceil - 1$, and setting $p_1 \triangleq p_2 \triangleq \cdots \triangleq p_{m-1} \triangleq 1/2 + \varepsilon$ and $p_m \triangleq s - \sum_{k=1}^{m-1} p_k$, for $\varepsilon > 0$ sufficiently small s.t. $p_m > 0$. Indeed, no set B^j that contains two of the indices $1, \ldots, m-1$ may contain any other index without violating the constraint of the lemma.

4 Lower Bound on Revenue-Approximation Complexity

In this section, we prove Theorem 1.4, which states that $C(n, 1/n) \geq 2^{\Omega(n)}$. We consider the simple distribution where the value of each item is (independently) either 0 or 1, each with probability 50%. Thus, the buyer is interested in (i.e., has value > 0 for the items in) a random subset of the items (where each subset is chosen with probability exactly 2^{-n}), and has value 1 for each item in this subset. A revenue-maximizing auction for this setting offers each item, separately, for a take-it-or-leave-it price of 1; viewed as a menu, this auction has a menu entry for each nonempty bundle of items $\emptyset \neq S \subseteq [n]$, which sells this bundle S for a price of $|S|$. This auction clearly maximizes the revenue, as it extracts the full social welfare (i.e., expected sum of item valuations), which is $n/2$, as revenue. An alternative auction, which has only a single menu entry, sells the whole bundle for a price that is slightly less than $n/2$, so that with high probability the size of the set of desired items is at least this price (in which case the bundle is sold). Choosing the price to be $n/2 - \omega(\sqrt{n})$ is required in order for the probability of selling the bundle to be subconstant. Therefore, this simple auction loses (about) a $\frac{1}{\sqrt{n}}$ fraction of the revenue. In the proof below, we show that such a $1/n^c$ loss of revenue is necessary in every simple-enough auction, i.e., in every auction having subexponentially many menu entries.

Before proceeding to the proof, we roughly sketch the intuition underlying it. Consider an auction that loses less than, say, a $\frac{1}{10n}$ fraction of the revenue. This means that for most of the 2^n possible bundles of desired items, the payment extracted from a buyer interested in (precisely the items in) that bundle must be higher than the bundle size minus $1/2$. Let us say that a buyer type that is interested in (precisely the items in) a bundle S *pays full price* if the auction extracts a payment higher than $|S| - 1/2$ from a buyer of this type. The basic idea is that, roughly speaking, if a buyer type that is interested in a bundle S pays full price, then this buyer type "should" have a menu entry "for itself", which allocates S for a price close to $|S|$. The reason is that two buyer types interested in bundles of different sizes cannot both choose the same entry and both pay full price (as the buyer type interested in the smaller-sized bundle would not want to buy at the "full price" of the other buyer). We note that it is true that buyer types interested in identically sized bundles $T \neq S$ can still "share" the same menu entry and both pay full price; indeed, this can happen if the entry that they both choose allocates the union $S \cup T$ for a price close to $|S| = |T|$. This, however, would imply that a buyer type interested in $S \cup T$ does not pay full price, since this menu entry already offers a discount for the bundle $S \cup T$. Our proof formalizes this intuition: If "too many" buyer types pay full price and choose the same menu entry (which must be the case with a small menu that extracts almost all of the revenue), then we show that buyer types interested in many subsets or their union cannot pay full price, leading to a significant loss in revenue.

Proof of Theorem 1.4. We will prove that $C\big(n, \frac{1}{10n}\big) > 2^{n/10}$. We note that this implies the theorem, as $C\big(10n, \frac{1}{10n}\big) \geq C\big(n, \frac{1}{10n}\big)$, which can be seen by fixing the value of $9n$ of the $10n$ items to 0 with probability 1. Let $n \in \mathbb{N}$ (throughout the proof, we will add requirements for n to be large enough as to accommodate for certain conditions). As stated above, we focus on the product distribution $F \in \Delta(\mathbb{R}_+)^n$, where the value of each of the n items is either 0 or 1, each independently with probability 50%.[25] Assume for contradiction that there exists an IC and IR auction that has at most $2^{n/10}$ menu entries, but that however obtains from F revenue of at least $\big(1 - \frac{1}{10n}\big) \cdot \mathcal{R}ev(F) = \big(1 - \frac{1}{10n}\big) \cdot n/2$.

For each set of items $S \subseteq [n]$, we identify S with the buyer type valuing each item in S by 1 and each item not in S by 0. We say that a buyer type $S \subseteq [n]$ *pays full price* if the payment extracted from a buyer of this type is strictly higher than $|S| - 1/2$. Our strategy is to define a certain distribution over pairs of buyer types (where the set of items associated with the first buyer type is contained in the one associated with the second, i.e., the first buyer type is interested in a subset of the items that interest the second), and reach a contradiction by showing that the

[25] So, following the above remark, we in fact prove the theorem by showing that $C\big(10n, \frac{1}{10n}\big) > 2^{(10n)/100}$ by focusing on a product distribution over $10n$ items, where $9n$ of the items have value 0 with probability 1, and each of the remaining n items has value either 0 or 1, each independently with probability 50%.

probability that both of these buyer types pay full price should on one hand be large due to the assumption of a slight revenue loss, but on the other hand be small due to the the assumption of a small menu size (recall the informal discussion preceding the proof, roughly explaining why two buyer types that are interested in two sets ordered by inclusion and that are both paying full price must "usually" choose distinct menu entries). We begin by defining this distribution over pairs of buyer types.

Construction We choose two buyer types (sets of items) S_0, S_1 at random as follows: we first choose a set of items uniformly at random among all 2^n sets, and then choose a random item $i \in [n]$ and take the symmetric difference of the first set with $\{i\}$ to obtain a second set. Thus, exactly one of the two chosen sets contains the element i, while the two sets are identical in their containment of all other items. Let us denote the set that does not contain i by S_0, and the set that contains i by S_1. We note that while the distributions of each of S_0 and S_1 are not uniform (e.g., the expected size of S_0 is $\frac{n-1}{2}$ and of S_1 is $\frac{n+1}{2}$, while the expected size of a uniformly chosen set is $n/2$), it is nonetheless the case that the probability that any of these distributions gives to any event is at most twice the probability that the uniform distribution gives the same event (since the average of the probabilities given to this event by the distributions of S_0 and of S_1 is exactly that given by the uniform distribution).

Lower Bound We lower-bound the probability that both buyer types S_0 and S_1 pay full price by using the union bound. We first note that at most $1/10$ of the buyer types $S \subseteq [n]$ do *not* pay full price; otherwise, the revenue loss due to all such buyer types would have been higher than an additive $\frac{1}{20}$, i.e., higher than a $\frac{1}{10n}$ fraction of the maximal revenue, contradicting the assumption that the given auction loses at most a $\frac{1}{10n}$ fraction of the maximal revenue. Next, notice that before the labeling, each of S_0 and S_1 was uniformly distributed among all possible 2^n buyer types. Therefore, for each one of the (unlabeled) buyer types, the probability that it does not pay full price is at most $1/10$. Hence, taking the union bound, the probability that at least one of the (unlabeled) buyer types does not pay full price is at most $1/5$, and so the probability that both pay full price is at least $4/5$. Following the strategy outlined above, we will reach a contradiction by showing that this probability need in fact be less than $4/5$ when taking into account the assumption that the given auction has a small menu size.

Upper Bound The derivation of an upper bound on the probability that both buyer types S_0 and S_1 pay full price is more subtle. Consider the following way of obtaining the exact same distribution for (S_0, S_1): first, S_0 is chosen according to the marginal distribution on it, and then S_1 is obtained by adding a random item $i \notin S_0$ (uniformly chosen from all items not in S_0) to S_0.

In our derivation of the upper bound, we condition upon whether the menu entry chosen by S_0 is chosen by "many" buyer types that pay full price. Formally, let us say that a menu entry from the given auction is *tiny* if at most $2^{0.8n}$ buyer types that pay full price choose this entry. Our upper bound on the probability that both buyer types S_0 and S_1 pay full price will be the sum of the probability that S_0 pays full price and chooses a tiny menu entry (which is obviously not smaller than the probability that S_0 and S_1 both pay full price, and S_0 chooses a tiny menu entry) and the probability that S_1 pays full price, conditioned upon S_0 paying full price and not choosing a tiny menu entry (which is obviously not smaller than the probability that S_0 and S_1 both pay full price, and S_0 does not choose a tiny menu entry). We now estimate each of these two probabilities.

Since there are at most $2^{n/10}$ menu entries in the given auction, the total number of buyer types that pay full price and choose a tiny menu entry is at most $2^{0.9n}$. Thus, the probability that S_0

pays full price and chooses a tiny menu entry is at most $2 \cdot \frac{2^{0.9n}}{2^n}$ (by the above observation that the probability of any event for S_0 is at most twice the probability of that event for a uniform set), i.e., this probability decreases exponentially with n.

Following the strategy outlined above, we move on to upper-bound the probability that S_1 pays full price, conditioned upon S_0 paying full price and not choosing a tiny menu entry. Assume, therefore, that S_0 pays full price and does not choose a tiny menu entry, and let us henceforth denote the (nontiny) menu entry chosen by S_0 by e. We consider the probabilities with which e allocates each of the items. By IR, every item that is contained in every set that (as a buyer type) pays full price and chooses e, must be allocated by e with probability strictly greater than $1/2$ (since if for some buyer type $S \subseteq [n]$, even a single item $i \in S$ is allocated by e with probability at most $1/2$, then S would receive negative utility from paying strictly more than $|S| - 1/2$ for e). Let U be the set of items that are allocated with probability strictly greater than $1/2$ by e. Since e is not tiny, there are more than $2^{0.8n}$ buyer types S that pay full price and choose it, and so $|U| > 0.8n$ (since the total number of subsets of U is at most $2^{|U|}$.)

We now come to the crux of our argument. Consider a set T that is contained in U, but is a proper *superset* of one of the sets (buyer types) $S \subseteq [n]$ that pay full price and choose e. By IC, we claim that T cannot be paying full price in the given auction. Indeed, e already offers buyer type T utility strictly greater than $1/2$ (since the utility of T from e is greater by strictly more than $1/2$ than the utility of buyer type S from e, which by IR is nonnegative), and a buyer that pays full price has, by definition, utility strictly less than $1/2$.

Since S_1 is a strict superset of S_0, and since S_0 chooses e and pays full price, we therefore have that if $S_1 \subseteq U$, then S_1 does not pay full price. Therefore, the probability that S_1 is not contained in U is an upper bound on the probability that S_1 pays full price (when both are conditioned upon S_0 paying full price and not choosing a tiny menu entry). Recall that in our "second way" of obtaining the distribution for (S_0, S_1), the set S_1 is obtained by adding a random item $i \notin S_0$ (uniformly chosen from all items not in S_0) to S_0. For any S_0, the probability that $i \notin U$ is thus at most $\frac{n-|U|}{n-|S_0|} < \frac{0.2n}{n-|S_0|}$, and since the probability that $|S_0| > 0.6n$ (i.e., that $\frac{0.2n}{n-|S_0|} > 1/2$) decreases exponentially with n (even after conditioning upon S_0 paying full price and not choosing a tiny menu entry[26]), we get that conditioned upon S_0 paying full price and not choosing a tiny menu entry, the probability that S_1 pays full price is, for large enough n, at most (say) $3/5$.

Contradiction Following the strategy outlined above, we can now upper-bound the probability that both S_0 and S_1 pay full price, by summing the probability that S_0 pays full price and chooses a tiny menu entry (which decreases exponentially with n) and the probability that S_1 pays full price conditioned upon S_0 paying full price and not choosing a tiny menu entry (which we upper-bounded by $3/5$ for large enough n). Together, for large enough n, we get strictly less than $4/5$, which is the lower bound computed above for the probability of the same event — a contradiction. □

Acknowledgments

Yannai Gonczarowski is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities. The work of Yannai Gonczarowski was supported in part by the European

[26]Indeed, recall that at least $9/10$ of the buyer types pay full price; therefore, the probability that S_0 pays full price is at least $8/10$ (by the above observation that the probability of any event for S_0 is at most twice the probability of that event for a uniform set). Recall also that the probability that S_0 pays full price and chooses a tiny menu entry decreases exponentially with n; therefore, for large enough n, the probability that S_0 pays full price and does not choose a tiny menu entry is at least $7/10$.

Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no. [249159]. The work of Noam Nisan was supported by ISF grant 1435/14 administered by the Israeli Academy of Sciences and by Israel-USA Bi-national Science Foundation (BSF) grant number 2014389.

References

M. Babaioff, N. Immorlica, B. Lucier, and S. M. Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *Proceedings of the IEEE 55th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 21–30, 2014.

P. Briest, S. Chawla, R. Kleinberg, and S. M. Weinberg. Pricing randomized allocations. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 585–597, 2010.

S. Chawla, D. L. Malec, and B. Sivan. The power of randomness in Bayesian optimal mechanism design. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*, pages 149–158, 2010.

C. Daskalakis and S. M. Weinberg. Symmetries and optimal multi-dimensional mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pages 370–387, 2012.

C. Daskalakis, A. Deckelbaum, and C. Tzamos. Mechanism design via optimal transport. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce (EC)*, pages 269–286, 2013.

S. Dughmi, L. Han, and N. Nisan. Sampling and representation complexity of revenue maximization. In *Proceedings of the 10th Conference on Web and Internet Economics (WINE)*, pages 277–291, 2014.

Y. Giannakopoulos and E. Koutsoupias. Duality and optimality of auctions for uniform distributions. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation (EC)*, pages 259–276, 2014.

Y. Giannakopoulos and E. Koutsoupias. Selling two goods optimally. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 650–662, 2015.

S. Hart and N. Nisan. Approximate revenue maximization with multiple items. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, page 656, 2012.

S. Hart and N. Nisan. The menu-size complexity of auctions. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce (EC)*, page 565, 2013.

S. Hart and N. Nisan. How good are simple mechanisms for selling multiple goods? Discussion Paper 666, Center for the Study of Rationality, Hebrew University of Jerusalem, 2014.

S. Hart and P. J. Reny. Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Theoretical Economics*, 10(3):893–922, 2015.

X. Li and A. C.-C. Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences (PNAS)*, 110(28):11232–11237, 2013.

A. M. Manelli and D. R. Vincent. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory*, 127(1):1–35, 2006.

R. P. McAfee and J. McMillan. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory*, 46(2):335–354, 1988.

J. H. Morgenstern and T. Roughgarden. On the pseudo-dimension of nearly optimal auctions. In *Proceedings of Advances in Neural Information Processing Systems 28 (NIPS)*, pages 136–144, 2015.

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

A. Rubinstein and S. M. Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation (EC)*, pages 377–394, 2015.

J. Thanassoulis. Haggling over substitutes. *Journal of Economic Theory*, 117(2):217–245, 2004.

A. C.-C. Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 92–109, 2015.

A Technical Notes

A.1 Infinite Menus

A.1.1 Utility-Maximizing Entries

As is well known in the literature, by the taxation principle every (single-buyer) IC auction with finitely many possible outcomes can be identified with a finite menu of possible choices for the buyer (where by IC the buyer chooses an entry that maximizes her utility), and *vice versa*. For IC auctions with infinitely many possible outcomes, while it is still true that each such auction can be identified with an (infinite) menu of possible choices for the buyer, it is no longer the case that every such menu defines some IC auction. Indeed, in a general infinite menu, a utility-maximizing entry for each buyer type does not necessarily exist, yet in menus corresponding to IC auctions, such an entry always exists. We note that one way to make sure that a (possibly infinite) menu that we construct indeed defines an IC auction is to make sure that this menu is closed (as a subset of $[0,1]^n \times \mathbb{R}_+$). (This is the technique employed in our proofs of Lemmas 2.3 and 2.5.[27]) Indeed, for a buyer type $v = (v_1, \ldots, v_n) \in \mathbb{R}_+^n$, all nonnegative-utility entries from a closed menu $\mathcal{M}$ lie in the compact set $\mathcal{M} \cap \Big([0,1]^n \times \big[0, \sum_{i=1}^n v_i\big]\Big)$, and by continuity of the utility function, this function attains a maximum value in this compact set.

A.1.2 Tie Breaking by Prices

Definition 2.4(a), which is the standard definition of the revenue obtainable by an IC auction, specifies that ties (in utility) between menu entries are broken in favor of higher prices. We note that even if, in some menu, a utility-maximizing menu entry exists for some buyer type (or for all buyer types), then it is not guaranteed that a utility-maximizing entry *with maximal price* (among all

[27]The reasoning in our proof of Lemma 2.4 is more delicate and is detailed in that proof.

utility-maximizing entries) exists for this buyer type. (I.e., it is not guaranteed that the supremum price over all utility-maximizing entries for this buyer type is attained as a maximum.) Indeed, to be completely general, a more subtle definition of the revenue obtainable by an IC auction would have been needed (taking, roughly speaking, the supremum revenue over all tie-breaking rules), and even under such a definition, the menu entry of choice of a specific buyer type would not have been well defined (or would have become a limit of menu entries), making reasoning about such auctions quite cumbersome. Nonetheless, similarly to above, if an IC auction has a closed menu, then it possesses, for each buyer type, a utility-maximizing entry with maximal price, and so the revenue is well defined by Definition 2.4(a), without the need for a more subtle definition. Indeed, the set of utility-maximizing entries from a given closed menu for a given buyer type is (by the reasoning given in Appendix A.1.1) a compact set; therefore, it contains an entry with maximum price. The only point in the proofs in this paper where we do not explicitly construct an auction by specifying its menu is in the proof of Theorem 1.1, where we start with an IC (and IR) auction (possibly of infinite size) that obtains revenue close to $\mathcal{R}ev(F)$; to justify the fact that we can assume w.l.o.g. that the menu of such an auction is closed, we note that if some menu entry e is weakly preferred by some buyer type v over all menu entries in some menu $\mathcal{M}$, then by continuity of the utility function, e is weakly preferred by v also over all menu entries in the $\overline{\mathcal{M}}$ — the *closure* of $\mathcal{M}$ (in $[0,1]^n \times \mathbb{R}_+$). Therefore, given an IC auction $\mathcal{M}$ that obtains revenue R from a distribution F under some tie-breaking rule, we have that $\overline{\mathcal{M}}$, under price-maximization tie breaking (which is well defined since $\overline{\mathcal{M}}$ is closed), obtains revenue at least R from F.

A.2 Arbitrary Tie Breaking

As already discussed above, the (standard) definition that we use for the revenue obtainable by a given IC auction (Definition 2.4(a)) depends on tie breaking being performed in favor of higher prices. Nonetheless, we emphasize that the definition of the revenue obtainable from a given distribution F, whether constrained by the menu size (i.e., $\mathcal{R}ev_C$) or unconstrained (i.e., $\mathcal{R}ev$), does not depend on the tie-breaking rule. Indeed, if an auction $\mathcal{M}$ obtains revenue R w.r.t. tie breaking in favor of high prices, then for arbitrarily small $\varepsilon > 0$, multiplying the price of each menu entry in $\mathcal{M}$ by $(1-\varepsilon)$ and taking the closure of the resulting menu (equivalently, multiplying the price of each menu entry in $\overline{\mathcal{M}}$ by $(1-\varepsilon)$), gives an auction with the same menu size (since a finite menu is always closed) that gives revenue at least $(1-\varepsilon)\cdot R$ w.r.t. *any* tie-breaking rule (since multiplying each price by $(1-\varepsilon)$ breaks ties in favor of higher-priced menu entries, and can only cause a buyer type to "jump" to even higher-costing menu entries). As the definitions of $\mathcal{R}ev_C$ and $\mathcal{R}ev$ therefore do not depend on the tie-breaking rule, our results hold for any tie-breaking rules, e.g., even tie-breaking in favor of low prices in the definition of $\mathcal{R}ev_C$ and in favor of high prices in the definition of $\mathcal{R}ev$.