



The menu-size complexity of revenue approximation[☆]

Moshe Babaioff¹, Yannai A. Gonczarowski^{*,2}, Noam Nisan³



ARTICLE INFO

Article history:

Received 11 April 2020

Available online 8 April 2021

Keywords:

Mechanism design

Revenue maximization

Approximate revenue maximization

Menu size

Auction

Communication complexity

ABSTRACT

Consider a monopolist selling n items to an additive buyer whose item values are drawn from independent distributions F_1, F_2, \dots, F_n possibly having unbounded support. Unlike in the single-item case, it is well known that the revenue-optimal selling mechanism (a pricing scheme) may be complex, sometimes requiring a continuum of menu entries. Also known is that simple mechanisms with a bounded number of menu entries can extract a constant fraction of the optimal revenue. Nonetheless, whether an arbitrarily high fraction of the optimal revenue can be extracted via a bounded menu size remained open.

We give an affirmative answer: for every n and $\varepsilon > 0$, there exists $C = C(n, \varepsilon)$ s.t. mechanisms of menu size at most C suffice for obtaining $(1 - \varepsilon)$ of the optimal revenue from any F_1, \dots, F_n . We prove upper and lower bounds on the revenue-approximation complexity $C(n, \varepsilon)$ and on the deterministic communication complexity required to run a mechanism achieving such an approximation.

© 2021 Elsevier Inc. All rights reserved.

1. Introduction

As familiar economic institutions move to computerized platforms, they are reaching unprecedented sizes and levels of complexity. These new levels of complexity often become the defining feature of the computerized economic scenario, as in the cases of spectrum auctions and ad auctions. The use of the word “complexity” here is 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 to analyze such measures of complexity and understand 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 who 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” or “mechanism” (really just a pricing scheme) that maximizes her revenue. The classical economic analysis (Myerson, 1981) shows that for a single item, the optimal mechanism is simply to 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

[☆] An extended abstract of this work appeared in STOC 2017 (Babaioff et al., 2017).

^{*} Corresponding author.

E-mail addresses: moshe@microsoft.com (M. Babaioff), yannai@gonch.name (Y.A. Gonczarowski), noam.nisan@gmail.com (N. Nisan).

¹ Microsoft Research.

² Microsoft Research. Research conducted while also affiliated with the 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.

³ Rachel & Selim Benin School of Computer Science & Engineering and Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, Israel. Research conducted while also affiliated with Microsoft Research.

mechanism may be surprisingly complex, randomized, and unintuitive (McAfee and McMillan, 1988; Thanassoulis, 2004; Manelli 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” mechanisms may yield at least 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” mechanisms 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; Cai and Zhao, 2017), but is in contrast to the case where the item values are distributed according to a joint (correlated) distribution, a case for which no finite approximation is possible by simple mechanisms (Briest et al., 2010; Hart and Nisan, 2013).

In this work, we study the trade-off between the complexity of a mechanism and the extent to which it can approximate the optimal revenue. One may choose various measures of mechanism 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 a mechanism (for a single buyer, an auction or mechanism is just a pricing scheme) is defined to be the number of different possible outcomes of the mechanism. More specifically, every single-buyer mechanism is equivalent to one that offers a *menu* of options to the buyer, where each option — *entry* $(\bar{x}; p)$ — 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 \bar{x} , and the buyer chooses an entry that maximizes her own expected 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 mechanism. As we observe, the menu size is tightly related to the *deterministic communication complexity* of the mechanism: if the mechanism is considered common knowledge, then the amount of information⁴ that the buyer must send to the seller so that the outcome (allocation probabilities and price) of the mechanism for this buyer can be determined by the seller,⁵ is exactly equal to the logarithm of the menu size of the mechanism. It is known that for some distributions, the optimal mechanism has infinite menu size (Daskalakis et al., 2013) — or equivalently by the above, requires infinite communication complexity — but a constant fraction of the optimal revenue may be extracted by a finite-complexity mechanism (Babaioff et al., 2014). Is it possible to extract an arbitrarily high fraction of the optimal revenue via a finite menu size (equivalently, via mechanisms with finite communication complexity)?

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

Definition 1.1 (REV_C ; REV). For a distribution F on n items, we denote by $\text{REV}_C(F)$ the maximal (formally, the supremum) revenue obtainable by an individually rational incentive-compatible mechanism that has at most C menu entries and sells the n items to a single additive buyer whose values for the items are distributed according to F . We denote by $\text{REV}(F) = \text{REV}_\infty(F)$ the maximal revenue obtainable without any complexity restrictions on the mechanism.

Formally, our result shows that $\lim_{C \rightarrow \infty} \frac{\text{REV}_C(F)}{\text{REV}(F)} = 1$ uniformly across all product distributions F . In other words:

Theorem 1.1 (Qualitative version). For every number of items n and every $\varepsilon > 0$, there exists a finite menu size $C = C(n, \varepsilon)$ such that for every $F_1, F_2, \dots, F_n \in \Delta(\mathbb{R}_+)$, we have that $\text{REV}_C(F_1 \times \dots \times F_n) \geq (1 - \varepsilon) \cdot \text{REV}(F_1 \times \dots \times F_n)$.

Theorem 1.1 gives a positive answer to Open Problem 6 from Hart and Nisan (2014),⁶ which asks precisely whether the statement of Theorem 1.1 holds. From Theorem 1.1, we immediately get that finite communication complexity uniformly suffices for arbitrarily approximating the optimal revenue:

Corollary 1.2 (Qualitative version). For every number of items n and every $\varepsilon > 0$, there exists a finite communication complexity $D = D(n, \varepsilon)$ such that for every $F_1, F_2, \dots, F_n \in \Delta(\mathbb{R}_+)$, there exists an n -item mechanism with a deterministic communication complexity of only D that approximates the optimal revenue from $F_1 \times \dots \times F_n$ up to a multiplicative ε loss.

It is natural to ask what is the rate of the uniform convergence of the sequence $\frac{\text{REV}_C(F)}{\text{REV}(F)}$. In other words, *how complex* must a revenue-approximating mechanism 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 $\text{REV}_C(F_1 \times \dots \times F_n) \geq (1 - \varepsilon) \cdot \text{REV}(F_1 \times \dots \times F_n)$ for every $F_1, F_2, \dots, F_n \in \Delta(\mathbb{R}_+)$.

⁴ See Appendix A for a formal definition.

⁵ Since only the buyer has private information (her values) in our setting, an arbitrary interactive protocol between the buyer and the seller must use at least as many bits as this one-way communication. Thus, this notion captures general two-way deterministic complexity as well. See Appendix A for more details, as well as a discussion regarding randomized communication complexity.

⁶ 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 $\frac{\text{REV}_C(F)}{\text{REV}(F)}$).

Theorem 1.1 (Quantitative version). $C(n, \varepsilon) \leq \left(\frac{\log n}{\varepsilon}\right)^{O(n)}$.

Corollary 1.2 (Quantitative version). For every number of items n , every $\varepsilon > 0$, and every $F_1, F_2, \dots, F_n \in \Delta(\mathbb{R}_+)$, there exists an n -item mechanism with a deterministic communication complexity of only $D(n, \varepsilon) = \log C(n, \varepsilon) = O\left(n \log\left(\frac{\log n}{\varepsilon}\right)\right)$ that approximates the optimal revenue from $F_1 \times \dots \times F_n$ up to a multiplicative ε loss.

This bound on the menu size is exponential in n , and so the next natural question is whether polynomial menu size suffices. At first glance the answer seems to be “obviously not”: the menu-size complexity measure is quite weak, and even the mechanism that sells each item separately has exponential menu size (since, when presented as a menu, a menu entry is needed for each possible subset of the items), which has been the source of one of the main criticisms of the menu-size complexity measures. This answer, however, is premature; in fact, we show that polynomial menu size turns out to suffice for approximating the revenue obtainable from selling items separately, thereby also appeasing this criticism.⁷ Let us denote by $\text{SREV}(F_1 \times \dots \times F_n)$ the revenue obtainable by selling each item separately for its optimal price.

Theorem 1.3. For every $\varepsilon > 0$, there exists $d(\varepsilon)$ such that for every number of items n and $F_1, F_2, \dots, F_n \in \Delta(\mathbb{R}_+)$, we have for $C = n^{d(\varepsilon)}$ that $\text{REV}_C(F_1 \times \dots \times F_n) \geq (1 - \varepsilon) \cdot \text{SREV}(F_1 \times \dots \times F_n)$.

The same bound applies also to the revenue obtainable by selling the items after arbitrarily repartitioning them into 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 1.4. There exist a fixed constant number d and a fixed constant fraction $\alpha > 0$ such that $C(n, 1 - \alpha) \leq O(n^d)$.

The above reasoning shows that Corollary 1.4 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 1.5. $C(n, 1/n) \geq 2^{\Omega(n)}$.

The proof of Theorem 1.5 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 problems open. The first one is whether for every fixed $\varepsilon > 0$, polynomial (or at least quasi-polynomial) menu size suffices for approximating the optimal revenue to within a multiplicative ε . In terms of communication complexity, this translates to whether logarithmic or polylogarithmic deterministic communication suffices⁸ for every fixed value of ε .

Open Problem 1.6. Is it true that for every $\varepsilon > 0$, there exists $d(\varepsilon)$ such that $C(n, \varepsilon) \leq O(n^{d(\varepsilon)})$? How about $C(n, \varepsilon) \leq O(2^{\log^{d(\varepsilon)} n})$?

The second open problem (or rather, class of open problems) is whether stronger notions of mechanism description complexity may allow for better revenue in polynomial complexity. This may be asked with respect to any complexity measure, and it is not clear which specific natural choice to consider, so an identification of such a measure is part of what is left open.⁹

⁷ That is, by our next result, any approximation impossibility that can be shown for polynomial menu-size mechanisms would immediately apply also to separate selling as well.

⁸ The authors have opposing conjectures regarding the answer to this open problem.

⁹ One may be tempted to consider the *additive menu size* complexity defined by Hart and Nisan (2013), which allows the seller to present menus from which the buyer is allowed to take any combination of menu entries for the sum of their prices. (So, for example, selling each item separately has linear additive menu size.) There are several possible interpretations here regarding whether buying lotteries translates into winning at the sum of probabilities (capped at one) or translates into not winning at the product of the probabilities of not winning. Similarly, the mechanism could either allow or not allow the buyer to adaptively decide whether and which additional menu entry to buy based on the realization of the lottery in the first, already purchased, menu entry. Follow-up work by Babaioff et al. (2018) showed that regardless of which interpretation is adopted, additive menus *cannot* achieve the optimal revenue. In other words, restricting to additive menus in fact entails for some distributions a constant multiplicative loss on the revenue, and so additive menus (regardless of how large their menu size is) cannot be used to obtain arbitrarily small losses compared to the optimal revenue. There are several more general definitions that are possible, but we have not found a truly satisfactory one.

1.1. Further and follow-up literature

In the short time since the first appearance of this paper, there has been a rise in the research attention dedicated to menu sizes, in two major directions. Most related to this work, Gonczarowski (2018) analyzes the required menu size when holding the number of items n fixed, as a function of the allowed loss ε (in a sense symmetrically to our Open Problem 1.6 above, which asks about the menu size for loss ε held fixed, as a function of the number of items n), and shows that already for two items with bounded i.i.d. valuations, a polynomial dependence on $1/\varepsilon$ as in our bound from Theorem 1.1 is required. As Gonczarowski (2018) shows, this lower bound, together with our upper bound (Theorem 1.1) and with our observation that the deterministic communication complexity of running a mechanism is the logarithm of its menu size (Proposition A.1 in Appendix A), implies the following: for any fixed number of items n , there exists a tight $\Theta(\log 1/\varepsilon)$ bound on the minimal communication complexity that is guaranteed for any product distribution to suffice for running a mechanism that maximizes revenue up to a multiplicative ε .

Our arguments readily extend beyond additive valuations, to subadditive valuations. Following up on our work, Kothari et al. (2019) consider a more permissive notion of menu size that they define, *symmetric menu size*, which allows for a single menu entry to represent many menu entries that can be obtained from the original menu entry by permuting the identities of the items. They extend and tighten our analysis for the case of unit-demand valuations (a different subclass of subadditive valuations) and for their menu-size notion (their most notable change is to the last part of our analysis – the discretization of the “cheap part” of the menu), and prove that a quasipolynomial symmetric menu size suffices for obtaining a $(1 - \varepsilon)$ fraction of the optimal revenue. That is, they give a positive answer to a unit-demand/symmetric-menu-size variant of the second part of our Open Problem 1.6. They still explicitly leave open the question of obtaining a similar result for additive buyers (as in our Open Problem 1.6), even in their more permissive symmetric menu size model.

Slightly farther apart is the line of study of the menu size of mechanisms that lie, in a sense, between single- and multi-dimensional mechanisms – mechanisms whose study was originated by Fiat et al. (2016) with their study of the “FedEx auction.” Necessary and sufficient menu sizes for precise revenue maximization in the FedEx setting and variants thereof have been studied by Fiat et al. (2016), Devanur and Weinberg (2017), Saxena et al. (2018), and Devanur et al. (2020), while the menu size for approximate revenue maximization up to a multiplicative loss of $\varepsilon = 1/n^2$ has been studied by Saxena et al. (2018). This latter analysis asks a somewhat similar question to our Theorem 1.5 (where ε also shrinks as n grows, though with $\varepsilon = 1/n$), however in the multi-item model that we study the corresponding menu size is exponential while in the “one-and-a-half dimensional” FedEx model that they study it is polynomial. For more details on all of these results, we refer the interested reader to Goldner and Gonczarowski (2018).

2. Upper bound on revenue-approximation complexity

In this section, we prove our main result, 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 $C(n, \varepsilon) \leq \left(\frac{\log n}{\varepsilon}\right)^{O(n)}$. To somewhat ease presentation and to emphasize which elements in the proof are used for the qualitative result and which only for the precise quantitative one, we will first prove that $C(n, \varepsilon) \leq (n/\varepsilon)^{O(n)}$, which makes the same qualitative statement and still requires virtually all of the novel technical “beef” of the proof, but saves some clutter at the end. 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 formal details of each of the four steps of the proof, and Section 2.7 connects the dots by combining the four steps. Section 2.8 gives an overview of the modifications required to prove the quantitatively stronger upper bound of $\left(\frac{\log n}{\varepsilon}\right)^{O(n)}$, with the details relegated to the appendix. Finally, Section 2.9 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, which generalize bounded distributions.

2.1. Proof overview

Let F_1, F_2, \dots, F_n be the respective distributions of the values of the n items. We will construct a mechanism with (finite) menu size $(n/\varepsilon)^{O(n)}$ that guarantees a $(1 - \varepsilon)$ multiplicative approximation to the optimal revenue.¹⁰

Limitations of existing techniques One possible approach to approximate revenue maximization, taken by Li and Yao (2013), Babaioff et al. (2014), Rubinstein and Weinberg (2015), and more recently Cai and Zhao (2017), is to use a core/tail decomposition and bound the revenue from (or the welfare of) the core and the revenue from the tail. Unfortunately, such a decomposition inherently entails a nonnegligible revenue loss (it only guarantees a constant fraction of the optimal revenue), 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, while this approach makes no assumptions regarding the valuation space (beyond independence), this technique, as used in the literature so far, is unsuitable for guaranteeing negligible loss in revenue as in the result that we seek.

¹⁰ The overview of the proof of the stronger upper bound of $\left(\frac{\log n}{\varepsilon}\right)^{O(n)}$ is identical up to the last step, which is also quite similar but more intricate for the stronger bound; See Section 2.8 for the details.

Another possible approach, taken by Daskalakis and Weinberg (2012), Hart and Nisan (2013), and Dughmi et al. (2014), is to round all possible menu entries onto a discrete grid via “nudge and round” operations. Unfortunately, for the grid (and thus the menu size) to be finite and for the revenue loss to indeed be negligible, the above papers all require that the valuation space be bounded. Therefore, while this approach can guarantee negligible loss in revenue, this technique, as used in the literature so far, is unsuitable for the analysis of unbounded valuation spaces as in our setting.

To overcome the above-described limitations of the core/tail decomposition technique, we take a more subtle approach, by analyzing core and tail regions together. We first show (in Step 1 below) that one does not lose much revenue by disregarding what can be described as “second-order” tails, i.e., valuations where two or more of the item values lie in the tail. Then, we show how to gradually simplify an optimal mechanism (which may be arbitrarily complex, even infinite in size) for the valuation space consisting of the core plus all first-order tails while losing only a tiny fraction of the revenue in each step. For every modification that we perform to the menu, we must “simultaneously” check that we do not significantly hurt the revenue from either core or (first-order) tail buyers. To gradually simplify the menu, we first carefully modify the menu so that only a small number of menu entries have a high price (this is the most technically elaborate part of the proof, performed in Steps 2 and 3 below), and then (in Step 4 below) round the menu entries with low prices to a finite grid using “nudge and round” operations. At this point, the use of “nudge and round” onto a finite grid is possible without significant revenue loss since the price of the menu entries that we round is bounded. Nonetheless, care still has to be taken beyond previous “nudge and round” uses, to ensure that this rounding does not incentivize buyers in the (first-order) tails to switch to buying a lower-priced rounded entry. Before moving on to the definitions and formal statements and proof, we first give a somewhat more detailed, yet still high-level, overview of each of the four steps of the proof.

Step 1: move to an “almost bounded” valuation space This step, taken in Section 2.3, simplifies the valuation space by showing that since item values are independent, finding an approximately optimal mechanism under the assumption that at most one item has a value higher than H (i.e., has a value that lies in the H -tail), for some $H = \text{poly}(n, 1/\varepsilon)$, entails a very small loss compared to doing so without this assumption. This is possible, very roughly speaking, because the probability of two item values lying in the tail, for H as above, 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 values 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 the first-order tail, i.e., the valuation space where at most one item value 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 values 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 2.9).

Step 2: modify expensive menu items to behave “almost like” single-item mechanisms This step, taken in Section 2.4, starts with an optimal (possibly arbitrarily complex) revenue-maximizing mechanism 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 $E = \text{poly}(n, 1/\varepsilon)$, so that each expensive menu entry allocates only a single item with non-zero 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 mechanism must act like a unit-demand one and never allocate more than a single item. We call such a mechanism 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 value lies in the tail for the valuation of that buyer type. Thus, 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 in most natural cases, this step in fact 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 non-zero probability), this simplification allows the next step to significantly reduce the size of this part of the menu.

Step 3: apply Myerson’s result to obtain “almost one” expensive entry per item This 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 step. Since E -exclusivity means that the expensive menu entries “look like” separate single-dimensional mechanisms 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 mechanisms with a simple “almost deterministic” one. In contrast to Myerson’s single (non-zero) menu entry, we require two menu entries for each item: a deterministic one analogous to Myerson’s “optimal price” entry, and an additional randomized one analogous to the “opt out” zero entry in Myerson’s mechanism. 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.

Step 4: discretize cheap menu entries “almost to a grid” This final step, taken in Section 2.6, simplifies the cheap part of the menu by “rounding” the menu entries into a discrete set. We note that even at this point in the proof, the “nudge and round” techniques that allowed this rounding to be done with only negligible loss of revenue for bounded valuations in previous

papers (Daskalakis and Weinberg, 2012; Hart and Nisan, 2013; Dughmi et al., 2014) cannot just be used “out of the box” in this step. Indeed, slight changes in allocation probabilities may result in large revenue changes, since the valuation space is not bounded but only exclusively unbounded. Nonetheless, these techniques 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, we show that only finitely many of the menu entries 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. Recall, however, that we aim to achieve only a slight multiplicative drop (with no additional additive drop) in overall revenue. To obtain this result, when combining all of the above steps in Section 2.7 we assume w.l.o.g. that $\text{Max}_i \text{Rev}(F_i)$ is normalized¹¹ (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 obtained bound on the overall cumulative revenue drop for normalized mechanisms is purely multiplicative, the proof also implies the same multiplicative bound for all (even nonnormalized) distributions.

2.2. Preliminaries

Definition 2.1 (Notation).

- (Naturals). We denote the strictly positive natural numbers by $\mathbb{N} \triangleq \{1, 2, 3, \dots\}$.
- ($[n]$). For every $n \in \mathbb{N}$, we define $[n] \triangleq \{1, 2, \dots, n\}$.
- (Nonnegative Reals). We denote the nonnegative reals by $\mathbb{R}_+ \triangleq \{r \in \mathbb{R} \mid r \geq 0\}$.
- ($\Delta(\cdot)$). For a set A , we denote by $\Delta(A)$ the set of probability distributions over A .

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

- An *outcome* is an $(n+1)$ -tuple $(\vec{x}; p) = (x_1, x_2, \dots, x_n; p) \in [0, 1]^n \times \mathbb{R}_+$, 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 .
- We denote the (expected) *utility* of a (risk-neutral additive) buyer with *type* (respective item valuations) $v = (v_1, v_2, \dots, v_n) \in \mathbb{R}_+^n$ 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 mechanism as menu). Let $n \in \mathbb{N}$ be a number of items. By the taxation principle, we identify any *incentive-compatible* (IC) n -item *mechanism* with a (possibly infinitely large) menu of outcomes (the *entries* in the menu are all the possible outcomes of the mechanism), where by IC the buyer chooses an entry that maximizes her utility.¹² If the mechanism 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 mechanism is IR.) Following Hart and Nisan (2013), we define the *menu size* of an IC and IR mechanism as the number of entries, except $(\vec{0}; 0)$, in the menu of that mechanism.

Definition 2.4 ($\text{REV}_{\mathcal{M}}$; REV_C ; REV). 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 .

- Given an IC and IR n -item mechanism \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

$$\text{REV}_{\mathcal{M}}(F) \triangleq \mathbb{E}_{v \sim F}[p(v)],$$

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.^{13,14}

¹¹ The cases in which $\text{Max}_i \text{Rev}(F_i)$ cannot be normalized, i.e., when it is 0 or infinite, are easy to handle separately. In the former case, there is nothing to show. In the latter case, $\text{Rev}(F_i) = \infty$ for some $i \in [n]$, and so by the theorem of Myerson (1981), an arbitrarily high revenue can be extracted using a take-it-or-leave-it offer for item i .

¹² If the menu is infinite, then the fact that it corresponds to an IC mechanism guarantees that some menu entry maximizes the utility of each buyer type. See Appendix B.1 for more details.

¹³ The results of this paper hold regardless of the tie-breaking rule chosen. See Appendix B.2 for more details.

¹⁴ 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-

- 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 mechanism with at most C menu entries by

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

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

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

Theorem 2.1 (Hart and Nisan, 2012). $\text{REV}(F \times G) \leq 2 \cdot (\text{REV}(F) + \text{REV}(G))$, 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 value

As outlined above, our first step toward proving Theorem 1.1, which we take in this section, simplifies the valuation space. It does so by showing that since item values are independent, any mechanism 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 value higher than some $H = \text{poly}(n, 1/\varepsilon)$, also extracts most of the revenue without this assumption. This subsection is dedicated to the statement and proof of Lemma 2.2, which formalizes this step.

Definition 2.5 ($\mathbb{E}\mathbb{U}_H^n$; 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 strictly greater than H by

$$\mathbb{E}\mathbb{U}_H^n \triangleq \{(v_1, v_2, \dots, v_n) \in \mathbb{R}_+^n \mid |\{i \in [n] \mid v_i > H\}| \leq 1\}.$$

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

Definition 2.6 ($F|_A$). For a set A and a distribution F defined over some superset of A s.t. A is measurable and $F(A) > 0$, 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 $F = F_1 \times F_2 \times \dots \times F_n \in \Delta(\mathbb{R}_+^n)$ s.t. $\max_{i \in [n]} \text{REV}(F_i) \leq R$, all of the following hold.

- $F(\mathbb{E}\mathbb{U}_H^n) > 0$ (hence the exclusively unbounded conditioned distribution $F|_{\mathbb{E}\mathbb{U}_H^n}$ is well defined).
- $\text{REV}(F|_{\mathbb{E}\mathbb{U}_H^n}) \geq (1 - \varepsilon) \cdot \text{REV}(F)$.
- For every $a \in (0, 1]$ and for every IC and IR n -item mechanism \mathcal{M} , if $\text{REV}_{\mathcal{M}}(F|_{\mathbb{E}\mathbb{U}_H^n}) \geq a \cdot \text{REV}(F|_{\mathbb{E}\mathbb{U}_H^n})$, then $\text{REV}_{\mathcal{M}}(F) \geq (1 - \varepsilon) \cdot a \cdot \text{REV}(F)$.

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

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

for every $i \in [n]$. Indeed, the revenue from F_i of the mechanism selling (item i) for a take-it-or-leave-it price of H is at least $H \cdot p_i$, and by definition of REV we therefore have $H \cdot p_i \leq \text{REV}(F_i) \leq R$ 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 $\mathbb{E}\mathbb{U}_H^n \supseteq [0, H]^n$, we obtain that $F(\mathbb{E}\mathbb{U}_H^n) \geq \prod_{i=1}^n (1 - p_i) > 0$, proving Part a. (Thus, $F|_{\mathbb{E}\mathbb{U}_H^n}$ is well defined.)

For the proof of Parts b and c, we will need the following sublemma (the second part of this sublemma is a slightly generalized version of the “subdomain stitching” lemma of Babaioff et al. (2014); we give a full proof below for completeness).

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

maximizing entries is attained as a maximum.) Indeed, to be completely general, a more subtle definition of the revenue obtainable by an IC mechanism would have been needed. Nonetheless, for the mechanisms 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 B.1.2 for more details.

- a. $F(B) \cdot \text{REV}(F|_B) \leq F(A) \cdot \text{REV}(F|_A)$ for every¹⁵ $B \subseteq A \subseteq \mathbb{R}_+^n$.
b. For every $m \in \mathbb{N}$ and every $A_1, A_2, \dots, A_m \subseteq \mathbb{R}_+^n$ s.t. $\bigcup_{i=1}^m A_i = \mathbb{R}_+^n$, we have that $\text{REV}(F) \leq \sum_{i=1}^m F(A_i) \cdot \text{REV}(F|_{A_i})$.

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

$$\text{REV}(F_i|_{(H,\infty)}) \leq \text{REV}(F_i)/p_i. \quad (2)$$

Indeed, by Sublemma 2.2.1(a), we have that $p_i \cdot \text{REV}(F_i|_{(H,\infty)}) \leq \text{REV}(F_i)$, and so $\text{REV}(F_i|_{(H,\infty)}) \leq \text{REV}(F_i)/p_i$, as required.

For every $1 \leq i < j \leq n$, we let $B_{i,j} \triangleq \{(v_1, v_2, \dots, v_n) \in \mathbb{R}_+^n \mid v_i > H \text{ \& \& } v_j > H\}$ (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

$$p_{i,j} \cdot \text{REV}(F|_{B_{i,j}}) \leq \frac{\varepsilon}{\binom{n}{2}} \cdot \text{REV}(F), \quad (3)$$

for every $1 \leq i < j \leq n$. 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

$$\begin{aligned} p_{i,j} \cdot \text{REV}(F|_{B_{i,j}}) &= p_{i,j} \cdot \text{REV}\left(F_i|_{(H,\infty)} \times F_j|_{(H,\infty)} \times \bigtimes_{k \in [n] \setminus \{i,j\}} F_k\right) \leq \\ &\leq 4 \cdot p_{i,j} \cdot \left(\text{REV}(F_i|_{(H,\infty)}) + \text{REV}(F_j|_{(H,\infty)}) + \text{REV}\left(\bigtimes_{k \in [n] \setminus \{i,j\}} F_k\right)\right) \leq \\ &\leq 4 \cdot p_{i,j} \cdot \left(\frac{\text{REV}(F_i)}{p_i} + \frac{\text{REV}(F_j)}{p_j} + \text{REV}\left(\bigtimes_{k \in [n] \setminus \{i,j\}} F_k\right)\right) = \\ &= 4 \cdot \left(p_j \cdot \text{REV}(F_i) + p_i \cdot \text{REV}(F_j) + p_i \cdot p_j \cdot \text{REV}\left(\bigtimes_{k \in [n] \setminus \{i,j\}} F_k\right)\right) \leq \\ &\leq 4 \cdot \frac{\varepsilon}{2 \cdot n \cdot (n-1)} \left(\text{REV}(F_i) + \text{REV}(F_j) + \text{REV}\left(\bigtimes_{k \in [n] \setminus \{i,j\}} F_k\right)\right) \leq \\ &\leq 2 \cdot \frac{\varepsilon}{n \cdot (n-1)} \cdot \text{REV}(F) = \frac{\varepsilon}{\binom{n}{2}} \cdot \text{REV}(F), \end{aligned}$$

as claimed.

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

$$\begin{aligned} p_{\mathcal{E}\mathcal{U}} \cdot \text{REV}(F|_{\mathbb{E}\mathcal{U}_H^n}) &\geq \text{REV}(F) - \sum_{1 \leq i < j \leq n} p_{i,j} \cdot \text{REV}(F|_{B_{i,j}}) \geq \\ &\geq \text{REV}(F) - \binom{n}{2} \cdot \frac{\varepsilon}{\binom{n}{2}} \cdot \text{REV}(F) = (1 - \varepsilon) \cdot \text{REV}(F). \end{aligned} \quad (4)$$

In particular, $\text{REV}(F|_{\mathbb{E}\mathcal{U}_H^n}) \geq p_{\mathcal{E}\mathcal{U}} \cdot \text{REV}(F|_{\mathbb{E}\mathcal{U}_H^n}) \geq (1 - \varepsilon) \cdot \text{REV}(F)$, proving Part b.

Let $a \in (0, 1]$ and let \mathcal{M} be an IC and IR n -item mechanism with $\text{REV}_{\mathcal{M}}(F|_{\mathbb{E}\mathcal{U}_H^n}) \geq a \cdot \text{REV}(F|_{\mathbb{E}\mathcal{U}_H^n})$. By definition of REV and by Eq. (4), we have that

$$\begin{aligned} \text{REV}_{\mathcal{M}}(F) &= p_{\mathcal{E}\mathcal{U}} \cdot \text{REV}_{\mathcal{M}}(F|_{\mathbb{E}\mathcal{U}_H^n}) + (1 - p_{\mathcal{E}\mathcal{U}}) \cdot \text{REV}_{\mathcal{M}}(F|_{\mathbb{R}_+^n \setminus \mathbb{E}\mathcal{U}_H^n}) \geq \\ &\geq p_{\mathcal{E}\mathcal{U}} \cdot \text{REV}_{\mathcal{M}}(F|_{\mathbb{E}\mathcal{U}_H^n}) \geq p_{\mathcal{E}\mathcal{U}} \cdot a \cdot \text{REV}(F|_{\mathbb{E}\mathcal{U}_H^n}) \geq a \cdot (1 - \varepsilon) \cdot \text{REV}(F), \end{aligned}$$

proving Part c. \square

Proof of Sublemma 2.2.1. For Part a, if $F(B) = 0$ then there is nothing to prove, so we assume henceforth that $F(B) > 0$. Therefore, 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 $\text{REV}(F|_A) \geq F|_{A(B)} \cdot \text{REV}(F|_B)$. Indeed, this inequality holds since for any mechanism \mathcal{M} (in particular, any mechanism obtaining close to optimal revenue from $F|_B$), we have $\text{REV}_{\mathcal{M}}(F|_A) \geq F|_{A(B)} \cdot \text{REV}_{\mathcal{M}}(F|_B)$; by definition of REV , the inequality follows. Therefore, we have that $F(A) \cdot \text{REV}(F|_A) \geq F(A) \cdot F|_{A(B)} \cdot \text{REV}(F|_B) = F(B) \cdot \text{REV}(F|_B)$, as required.

¹⁵ If $F(C) = 0$ for some $C \subseteq \mathbb{R}_+^n$, then even though $F|_C$ is not defined, we henceforth define $F(C) \cdot \text{REV}(F|_C)$ to equal 0.

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 $\text{REV}(F) \leq \sum_{i=1}^m F(B_i) \cdot \text{REV}(F|_{B_i})$. This is the “subdomain stitching” lemma of Babaioff et al. (2014); for completeness, we will repeat the one-sentence proof: for any mechanism \mathcal{M} (in particular, any mechanism obtaining close to optimal revenue from F), we have that $\text{REV}_{\mathcal{M}}(F) = \sum_{i=1}^m F(B_i) \cdot \text{REV}_{\mathcal{M}}(F|_{B_i})$; ¹⁶ by definition of REV , the inequality follows. Now, by Part a, we have that $F(B_i) \cdot \text{REV}(F|_{B_i}) \leq F(A_i) \cdot \text{REV}(F|_{A_i})$ for every $i \in [m]$. Combining both of these, we obtain that $\text{REV}(F) \leq \sum_{i=1}^m F(B_i) \cdot \text{REV}(F|_{B_i}) \leq \sum_{i=1}^m F(A_i) \cdot \text{REV}(F|_{A_i})$, as required. \square

2.4. Exclusivity at expensive menu entries

Having proven Lemma 2.2, we phrase and prove the next steps for arbitrary exclusively unbounded distributions, i.e., not necessarily product distributions conditioned upon $\mathbb{E}\mathbb{U}_H^n$. As outlined above, our second step toward proving Theorem 1.1, which we take in this section, shows that in any mechanism 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 mechanism E -exclusive, i.e., to make each expensive menu entry only allocate a single item with non-zero probability. This subsection is dedicated to the statement and proof of Lemma 2.3, which formalizes this step.

Definition 2.7 (*Exclusive mechanism*). Let $n \in \mathbb{N}$ and let $E \in \mathbb{R}_+$. We say that an n -item mechanism is E -exclusive if it allocates (with positive probability) at most one item whenever it charges strictly more than E .¹⁷

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 $F \in \Delta(\mathbb{E}\mathbb{U}_H^n)$ and for every IC and IR n -item mechanism \mathcal{M} , there exists an E -exclusive IC and IR n -item mechanism \mathcal{M}' such that $\text{REV}_{\mathcal{M}'}(F) \geq (1 - \varepsilon) \cdot \text{REV}_{\mathcal{M}}(F)$.

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

- For every menu entry $e = (\tilde{x}; p) \in \mathcal{M}$ with $p \leq E$, we add the menu entry e , unmodified, to \mathcal{M}' .
- For every menu entry $e = (\tilde{x}; p) \in \mathcal{M}$ with $p > E$, we add the following¹⁸ n menu entries to \mathcal{M}' : $(x_1, \tilde{0}_{-1}; (1 - \tilde{\varepsilon}) \cdot p)$, $(x_2, \tilde{0}_{-2}; (1 - \tilde{\varepsilon}) \cdot p)$, \dots , $(x_n, \tilde{0}_{-n}; (1 - \tilde{\varepsilon}) \cdot p)$. 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 (topological) closure of the set of menu entries added above to \mathcal{M}' .¹⁹

By definition, \mathcal{M}' is E -exclusive. We note that since \mathcal{M} is IR, it contains the menu entry $(\tilde{0}; 0)$. Therefore, \mathcal{M}' also contains this menu entry, and hence \mathcal{M}' is IR as well. It remains to show that \mathcal{M}' obtains a revenue of at least $(1 - \varepsilon) \cdot \text{REV}_{\mathcal{M}}(F)$ from F . Let us compare the payments that \mathcal{M}' and \mathcal{M} extract from a buyer of each type $v = (v_1, v_2, \dots, v_n) \in \mathbb{E}\mathbb{U}_H^n$. We reason by cases according to the menu entry of choice²⁰ of buyer type v from \mathcal{M} , which we denote by $e = (\tilde{x}; p)$, and show 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 $p \leq E$, then by definition $e \in \mathcal{M}'$. We claim that v weakly prefers e to all menu entries $f' = (\tilde{y}'; q') \in \mathcal{M}'$ with²¹ $q' < (1 - \tilde{\varepsilon}) \cdot p$. 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 to 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 $E > n \cdot H$ (since $n \geq 2$ and $\varepsilon < 1$), by IR there must exist $i \in [n]$ s.t. $v_i > H$. Since $v \in \mathbb{E}\mathbb{U}_H^n$, we have that $v_k \leq H$ for every $k \in [n] \setminus \{i\}$. Let $e' = (\tilde{x}'; p') \triangleq (x_i, \tilde{0}_{-i}; (1 - \tilde{\varepsilon}) \cdot p)$ be the menu entry in \mathcal{M}' corresponding to e that unallocates all items except item i . We claim that v weakly prefers e' to all menu entries $f' = (\tilde{y}'; q') \in \mathcal{M}'$ with²² $q' < p' - \frac{(1 - \tilde{\varepsilon}) \cdot (n-1) \cdot H}{\varepsilon}$. Let f' be such a menu entry and denote the menu entry corresponding

¹⁶ Similarly, if $F(B_i) = 0$ for some $i \in [n]$, then we define $F(B_i) \cdot \text{REV}_{\mathcal{M}}(F|_{B_i})$ to equal 0.

¹⁷ 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 solely this specific item being sold.

¹⁸ Following standard notation, we use $(y', \tilde{y}_{-i}; q)$, for $\tilde{y} \in [0, 1]^n$, $i \in [n]$, $y' \in [0, 1]$, and $q \in \mathbb{R}_+$, to denote the outcome $(y_1, \dots, y_{i-1}, y', y_{i+1}, \dots, y_n; q)$, i.e., an outcome that is identical to $(\tilde{y}; q)$ in price and all winning probabilities, except the winning probability of item i , which is set to y' .

¹⁹ Taking the closure ensures that a utility-maximizing entry with maximal price exists for every buyer type. See Appendix B.1 for more details.

²⁰ If more than one utility-maximizing menu entry with maximal price exists, then here and whenever we henceforth refer to the “menu entry of choice” of some buyer type, we choose one such entry arbitrarily.

²¹ As we show weak preference and as q' is defined via a strict inequality, by continuity of the utility function the correctness of the claim for all f' before taking the closure of \mathcal{M}' implies its correctness for all f' in the closure as well.

²² See Footnote 21.

to f' in \mathcal{M} by $f = (\vec{y}; q)$ (where either $f' = f$ or $f' = (y_j, \vec{0}_{-j}, (1 - \tilde{\varepsilon}) \cdot q)$ for some $j \in [n]$). Recall that $u_e(v) \geq u_f(v)$ by definition of e . Noting that $q \leq \frac{q'}{1 - \tilde{\varepsilon}} < p - \frac{(n-1) \cdot H}{\tilde{\varepsilon}}$, we indeed get

$$\begin{aligned} u_{e'}(v) &= x'_i \cdot v_i - p' = x_i \cdot v_i - p + \tilde{\varepsilon} \cdot p \geq -(n-1) \cdot H + \sum_{k=1}^n x_k \cdot v_k - p + \tilde{\varepsilon} \cdot p = \\ &= -(n-1) \cdot H + u_e(v) + \tilde{\varepsilon} \cdot p \geq -(n-1) \cdot H + u_f(v) + \tilde{\varepsilon} \cdot p = \\ &= -(n-1) \cdot H + \sum_{k=1}^n y_k \cdot v_k - q + \tilde{\varepsilon} \cdot p > -(n-1) \cdot H + \sum_{k=1}^n y_k \cdot v_k - q + \tilde{\varepsilon} \cdot \left(q + \frac{(n-1) \cdot H}{\tilde{\varepsilon}} \right) = \\ &= -(n-1) \cdot H + \sum_{k=1}^n y_k \cdot v_k - q \cdot (1 - \tilde{\varepsilon}) + (n-1) \cdot H \geq \sum_{k=1}^n y_k \cdot v_k - q' \geq \\ &\geq \sum_{k=1}^n y'_k \cdot v_k - q' = u_{f'}(v). \end{aligned}$$

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

$$\begin{aligned} p' - \frac{(1-\tilde{\varepsilon}) \cdot (n-1) \cdot H}{\tilde{\varepsilon}} &= p' - (1 - \tilde{\varepsilon}) \cdot \tilde{\varepsilon} \cdot E > p' - (1 - \tilde{\varepsilon}) \cdot \tilde{\varepsilon} \cdot p \geq \\ &\geq p' - \tilde{\varepsilon} \cdot p' = (1 - \tilde{\varepsilon}) \cdot p' = (1 - \tilde{\varepsilon})^2 \cdot p > (1 - 2\tilde{\varepsilon}) \cdot p, \end{aligned}$$

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 $v \in \mathbb{EU}_H^n$ 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 \mathbb{EU}_H^n is at least a $(1 - \varepsilon)$ fraction of the revenue that \mathcal{M} obtains from \mathbb{EU}_H^n , as required. \square

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 mechanism 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 subsection is dedicated to the statement and proof of Lemma 2.4, which formalizes this step, which is the most technically elaborate of all steps of our proof of Theorem 1.1.

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 $E \geq \text{Max}\{n \cdot H + 2, \frac{(n-1) \cdot H + 1}{\varepsilon}\}$. For every $F \in \Delta(\mathbb{EU}_H^n)$ and for every E -exclusive IC and IR n -item mechanism \mathcal{M} , there exists an E -exclusive IC and IR n -item mechanism \mathcal{M}' , such that both of the following hold.*

- $\text{REV}_{\mathcal{M}'}(F) \geq (1 - \varepsilon) \cdot \text{REV}_{\mathcal{M}}(F)$.
- (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} (with the possible exception of $(\vec{0}; 0)$) is chosen by at least one buyer type $v \in \mathbb{EU}_H^n$.²³ For every $i \in [n]$, we define

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

and

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

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 noted otherwise, fix $i \in [n]$ s.t. $\mathcal{M}_i \neq \emptyset$.**

²³ While \mathcal{M} is not necessarily closed after the removal of all menu entries that are not chosen by any buyer type, it does possess a utility-maximizing entry with maximal price for every buyer type. See Appendix B.1 for a discussion.

We define $b_i \triangleq \inf\{p \mid (\vec{x}; p) \in \mathcal{M}_i\}$ and $s_i \triangleq \inf\{x_i \mid (\vec{x}; p) \in \mathcal{M}_i\}$, and set $z_i \triangleq (s_i, \vec{0}_{-i}; b_i)$. One may intuitively think of z_i (charging b_i for an s_i probability of winning item i) as the cheapest entry, which also allocates the least probability, in \mathcal{M}_i , although formally (since \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 .²⁴ Our strategy is to show that, in a precise sense, \mathcal{M}_i behaves on W_i as follows: 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 mechanism” for possibly allocating some or all of the remaining $(1 - s_i)$ probability of winning item i . In this “continuation mechanism,” 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 noted otherwise, fix $i \in [n]$ s.t. $F(W_i) > 0$ (and so, by definition, $\mathcal{M}_i \neq \emptyset$) and $s_i < 1$.** As we will see below, we require that $s_i < 1$ in order for this “continuation mechanism” 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 mechanism to be well defined. As we will show, if $s_i = 1$ then $|\mathcal{M}_i| = |\{z_i\}| = 1$ (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 mechanism \mathcal{M} from W_i is zero, and so we will be able to simply delete \mathcal{M}_i from the original mechanism, without replacing it with anything).

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 mechanism” as

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

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

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

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 mechanism,” we now turn to defining the (menu of the) mechanism itself. For an outcome (not necessarily a menu entry in \mathcal{M}) of the form $e = (x_i, \vec{0}_{-i}; p)$, we define

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

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 mechanism” by²⁵

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

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 mechanism. We now make precise the above claim regarding \mathcal{M}_i acting as if the buyer were allocated z_i and then bid in the single-dimensional “continuation mechanism” \mathcal{M}_i^1 , thus relating the revenue of the single-dimensional “continuation mechanism” \mathcal{M}_i^1 with that of the original mechanism \mathcal{M} from W_i .

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

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

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 mechanism, 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 mechanism $\mathcal{M}_i^* \triangleq \{(0; 0), (1; c_i)\}$ maximizes the

²⁴ 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.

²⁵ We emphasize that we do not take the closure (see Appendix B.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.

revenue from F_i^1 , i.e., $\text{REV}_{\mathcal{M}_i^*}(F_i^1) = \text{REV}(F_i^1)$. Ideally, we would have now wanted to replace all entries from \mathcal{M}_i in \mathcal{M} by the two entries $z_i = \beta^{-1}((0; 0))$ and $o_i \triangleq (1, \vec{0}_{-i}; c_i + b_i) = \beta^{-1}((1; c_i))$, however it is not clear that none of the buyer types in W_i whose utility has dropped as a result of this replacement 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 $z'_i \triangleq (s_i, \vec{0}_{-i}; b_i - (n-1) \cdot H - 1)$ and $o'_i \triangleq (1, \vec{0}_{-i}; c_i + b_i - (n-1) \cdot H - 1)$. 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” mechanism:²⁶

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

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) menu entries for each $i \in [n]$, for a total of at most $2n$ expensive menu entries, as required. It remains to reason about the revenue of \mathcal{M}' . To do so, we now formulate what can be thought of as a “converse” of Sublemma 2.4.1, thus relating the revenue of the new trimmed mechanism \mathcal{M}' from W_i with the revenue of the optimal single-dimensional Myerson mechanism \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. $\text{REV}_{\mathcal{M}'}(F|_{W_i}) = \text{REV}_{\mathcal{M}_i^*}(F_i^1) + b_i - (n-1) \cdot H - 1$.

The proof of Sublemma 2.4.2 is given after the proof of Lemma 2.4. As mentioned 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 mechanism loses no significant revenue from W_i for every $i \in [n]$ s.t. $s_i < 1$. 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 mechanism lose significant revenue compared to the original untrimmed mechanism \mathcal{M} .

Sublemma 2.4.3 (Comparison between the revenues of \mathcal{M}' and \mathcal{M} from restrictions of F).

- a. *For every $i \in [n]$ s.t. $F(W_i) > 0$ and $s_i < 1$, we have that $\text{REV}_{\mathcal{M}'}(F|_{W_i}) \geq (1 - \varepsilon) \cdot \text{REV}_{\mathcal{M}}(F|_{W_i})$.*
- b. *For every $i \in [n]$ s.t. $F(W_i) > 0$ and $s_i = 1$, we have that $\text{REV}_{\mathcal{M}'}(F|_{W_i}) = \text{REV}_{\mathcal{M}}(F|_{W_i})$.*
- c. *Letting $P \triangleq \mathbb{E} U_H^n \setminus \bigcup_{i \in [n]} W_i$, if $F(P) > 0$, then $\text{REV}_{\mathcal{M}'}(F|_P) \geq (1 - \varepsilon) \cdot \text{REV}_{\mathcal{M}}(F|_P)$.*

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

$$\begin{aligned} \text{REV}_{\mathcal{M}'}(F) &= F(P) \cdot \text{REV}_{\mathcal{M}'}(F|_P) + \sum_{i=1}^n F(W_i) \cdot \text{REV}_{\mathcal{M}'}(F|_{W_i}) \geq \\ &\geq (1 - \varepsilon) \cdot \left(F(P) \cdot \text{REV}_{\mathcal{M}}(F|_P) + \sum_{i=1}^n F(W_i) \cdot \text{REV}_{\mathcal{M}}(F|_{W_i}) \right) = (1 - \varepsilon) \cdot \text{REV}_{\mathcal{M}}(F), \end{aligned}$$

as required. \square

²⁶ As before, we emphasize that we do not take the closure (see Appendix B.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 buyer types $v \in \mathbb{E} U_H^n \setminus \bigcup_{i \in [n]: F(W_i)=0 \text{ or } s_i < 1} W_i$ such an entry exists since the entry of choice of v 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. For buyer types $v \in \bigcup_{i \in [n]: F(W_i) > 0 \text{ \& } s_i < 1} W_i$ we show in Sublemma 2.4.2(a) below that such an entry exists.

2.5.1. Proof of Sublemmas 2.4.1 through 2.4.3

In this subsection, we prove Sublemmas 2.4.1 through 2.4.3. Before doing so, 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. $\mathcal{M}_i \neq \emptyset$.

- For every $e = (x_i, \bar{0}_{-i}; p) \in \mathcal{M}_i$ and $f = (y_i, \bar{0}_{-i}; q) \in \mathcal{M}_i$, either i) $e = f$, or ii) $x_i < y_i$ and $p < q$, or iii) $x_i > y_i$ and $p > q$.
- If there exists $e = (x_i, \bar{0}_{-i}; p) \in \mathcal{M}_i$ 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 $z_i \in \mathcal{M}_i$.
- If $s_i = 1$, then $\mathcal{M}_i = \{z_i\}$.
- For every $i \in [n]$, there exists a sequence $(z_i^m = (s_i^m, \bar{0}_{-i}; b_i^m))_{m=1}^\infty \subseteq \mathcal{M}_i$ s.t. i) $(s_i^m)_{m=1}^\infty$ is weakly decreasing, ii) $(b_i^m)_{m=1}^\infty$ is weakly decreasing, and iii) $\lim_{m \rightarrow \infty} z_i^m = z_i = (s_i, \bar{0}_{-i}; b_i)$.

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 to 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 to 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 proceed to Part b. If there exists $e = (x_i, \bar{0}_{-i}; p) \in \mathcal{M}_i$ s.t. $x_i = s_i$, then for every $s = (y_i, \bar{0}_{-i}; q) \in \mathcal{M}_i$ 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 proceed to Part c. Assume that $s_i = 1$ and let $e = (x_i, \bar{0}_{-i}; p) \in \mathcal{M}_i$. 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 $z_i^m \triangleq z_i$ 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 $(z_i^m = (s_i^m, \bar{0}_{-i}; b_i^m))_{m=1}^\infty \subseteq \mathcal{M}_i$ s.t. $(s_i^m)_{m=1}^\infty$ is weakly decreasing and s.t. $\lim_{m \rightarrow \infty} s_i^m = s_i$. By Part a, $(b_i^m)_{m=1}^\infty$ is weakly decreasing as well. It is enough, therefore, to show that $\lim_{m \rightarrow \infty} b_i^m = b_i$. 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 $\lim_{m \rightarrow \infty} b_i^m \geq b_i$; assume for contradiction that $\lim_{m \rightarrow \infty} b_i^m > b_i$. Therefore, by definition of b_i there exists $e = (x_i, \bar{0}_{-i}; p) \in \mathcal{M}_i$ s.t. $\lim_{m \rightarrow \infty} b_i^m > p$ 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 $s_i > x_i$. As by definition also $s_i \leq x_i$, we obtain that $s_i = x_i$, contradicting the assumption that s_i is not attained as a minimum. \square

Sublemma 2.4.5. Let $i \in [n]$ and let $e = (x_i, \bar{0}_{-i}; p)$ and $f = (y_i, \bar{0}_{-i}; q)$ 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. \square

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} .

- $u_e(v) \geq u_{z_i}(v)$.
- $u_{z_i'}(v) > u_f(v)$, for every $f \in \mathcal{M} \setminus \mathcal{M}_i$.
- $\min\{u_e(v), u_{z_i'}(v)\} > u_f(v)$, for every outcome f of the form $(y_j, \bar{0}_{-j}; q)$ 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 $u_e(v) \geq u_{\lim_{m \rightarrow \infty} z_i^m}(v) = u_{z_i}(v)$, as required.

We turn to Part b. Denote $e = (x_i, \bar{0}_{-i}; p) \in \mathcal{M}_i$ (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' to $z_i = e$,

which is the menu entry of choice of v from \mathcal{M} and, therefore, is weakly preferred by v to 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 $f = (\bar{y}; q)$ and let $f' \triangleq (y_i, \bar{0}_{-i}; q)$ (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 $q \leq E < b_i^m$), 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 to f , and by definition of f' , we have that w (as well as all other buyer types) weakly prefers f to f' . Therefore, w weakly prefers z_i^m to 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 to f' , as claimed. By continuity of u , we therefore have that $u_{z_i}(v) = u_{\lim_{m \rightarrow \infty} z_i^m}(v) \geq u_{f'}(v)$ as well.

As \mathcal{M} is IR and $e \in \mathcal{M}_i$, we have that $v_i > E > H$. Hence, since $v \in \mathbb{E}U_H^n$ 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 $(\bar{0}; 0) \in \mathcal{M} \setminus \mathcal{M}_i$ we also have by Part b that $u_{z_i'}(v) \geq u_{(\bar{0}; 0)}(v) = 0$. Therefore, to prove Part c it suffices to show that $u_f(v) < 0$. Indeed, recalling that $v_j \leq H$, we have that $u_f(v) = v_j \cdot y_j - q \leq H - q \leq H - E + (n-1) \cdot H + 1 = n \cdot H + 1 - E < 0$. \square

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

$$\begin{aligned} u_{\beta(f)}(\alpha(v)) &= v_i \cdot (1 - s_i) \cdot \frac{y_i - s_i}{1 - s_i} - q + b_i = v_i \cdot (y_i - s_i) - q + b_i = \\ &= (v_i \cdot y_i - q) - (v_i \cdot s_i - b_i) = u_f(v) - u_{z_i}(v). \end{aligned}$$

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

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

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 $M^1 \triangleq \{\beta(f) \mid f \in M\}$. As β is strictly monotone in the price coordinate, and since e is an entry in M with maximal price (indeed, e is surpassed by z_i neither in utility for v (by Sublemma 2.4.6(a)) nor 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 $e = (x_i, \bar{0}_{-i}; p) \in \mathcal{M}_i$, 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. \square

Proof of Sublemma 2.4.2. We start with Part a. By Sublemma 2.4.6(b), v strictly prefers z_i' to 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' to z_j' and 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

$$\begin{aligned} u_{o_i'}(v) &= v_i - c_i - b_i + (n-1) \cdot H + 1 = \\ &= (v_i \cdot s_i - b_i + (n-1) \cdot H + 1) + (v_i \cdot (1 - s_i) - c_i) = u_{z_i'}(v) + u_{(1; c_i)}(\alpha(v)), \end{aligned}$$

and so v weakly (resp. strictly) prefers o_i' to z_i' iff $\alpha(v)$ weakly (resp. strictly) prefers $(1; c_i)$ to $(0; 0)$, completing the proof of Part a as the price of o_i' is higher than that of z_i' and the price of $(1; c_i)$ is higher than that of $(0; 0)$ (and so in the case of indifference, v chooses o_i' and $\alpha(v)$ chooses $(1; c_i)$).

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. \square

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 mechanism for F_i^1 , and Sublemma 2.4.1(b), we have

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

where the last inequality is since $\text{REV}_{\mathcal{M}}(F|_{W_i}) \geq E$ by definition of W_i .

We proceed 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 to all menu entries in $\mathcal{M}' \setminus \mathcal{M}$. Indeed, by Sublemma 2.4.6(c), v strictly prefers z_i to 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., v strictly prefers z_i to 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 $e = (\bar{x}; p) \in \mathcal{M}$ 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 $e \in \mathcal{M} \cap ([0, 1]^n \times [0, E]) \subseteq \mathcal{M}'$, and since the price of every menu entry in $\mathcal{M}' \setminus \mathcal{M}$ is at least $\min_{i \in [n]: F(W_i) > 0 \text{ \& } s_i < 1} b_i - (n-1) \cdot H - 1 \geq E - (n-1) \cdot H - 1$, 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 so the proof of Part c is complete. \square

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 mechanism 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 subsection is dedicated to the statement and proof of Lemma 2.5, which formalizes this step.²⁷

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 $F \in \Delta(\mathbb{U}_H^n)$ and for every IC and IR n -item mechanism \mathcal{M} , there exists an IC and IR n -item mechanism \mathcal{M}' such that all of the following hold.*

- $\text{REV}_{\mathcal{M}'}(F) \geq (1-\varepsilon) \cdot \text{REV}_{\mathcal{M}}(F) - \varepsilon$.
- 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 fewer 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” mechanism \mathcal{M}'' that will help us define the required mechanism \mathcal{M}' . While \mathcal{M}'' may contain infinitely many menu entries, we show below that w.l.o.g. only finitely many of them are in fact chosen by any buyer type; consequently, we will derive the required mechanism \mathcal{M}' by defining it to be the mechanism 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

$$\lfloor r \rfloor_\delta \triangleq \delta \cdot \lfloor r/\delta \rfloor.$$

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

- For every menu entry $e = (\bar{x}; p) \in \mathcal{M}$ 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}'' : $(x_1, (\lfloor \bar{x} \rfloor_\chi)_{-1}; \lfloor p' \rfloor_\psi)$, $(x_2, (\lfloor \bar{x} \rfloor_\chi)_{-2}; \lfloor p' \rfloor_\psi)$, \dots , $(x_n, (\lfloor \bar{x} \rfloor_\chi)_{-n}; \lfloor p' \rfloor_\psi)$. 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 $e = (\bar{x}; p) \in \mathcal{M}$ with $p > E$, we add the menu entry $(\bar{x}; p')$ 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 $\lfloor r \rfloor_0 \triangleq r$ for every $r \in \mathbb{R}_+$, and $\lfloor \vec{r} \rfloor_0 \triangleq \vec{r}$ 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 $(x_i, (\lfloor \bar{x} \rfloor_{\chi'})_{-i}; \lfloor (1-\varepsilon) \cdot p \rfloor_{\psi'})$, for some $(\bar{x}; p) \in \mathcal{M}$, $i \in [n]$,

²⁷ See Appendix C for a stronger version of this lemma, strengthening the upper bound of $(n/\varepsilon)^{O(n)}$ to $(\frac{\log n}{\varepsilon})^{O(n)}$.

and $(\chi', \psi') \in \{(\chi, \psi), (0, 0)\}$. Finally, we define \mathcal{M}'' to be the closure of the set of menu entries added above to \mathcal{M} .²⁸ We note that since \mathcal{M} is IR, it contains the menu entry $(\vec{0}; 0)$. Therefore (as $[0]_\delta = 0$ for every δ), \mathcal{M}'' also contains this menu entry, and hence \mathcal{M}'' is IR as well.

We claim that $\text{REV}_{\mathcal{M}''}(F) \geq (1 - \varepsilon) \cdot \text{REV}_{\mathcal{M}}(F) - \varepsilon$. Indeed, let us compare the payments that \mathcal{M}'' and \mathcal{M} extract from a buyer of each type $v = (v_1, \dots, v_n) \in \mathbb{E}\mathbb{U}_H^n$. Let $e = (\vec{x}; p) \in \mathcal{M}$ be the menu entry of choice of v from \mathcal{M} . As $v \in \mathbb{E}\mathbb{U}_H^n$, there exists $i \in [n]$ s.t. $v_k \leq H$ for every $k \in [n] \setminus \{i\}$. Let $e' \triangleq (x_i, (\lfloor \vec{x} \rfloor_{\chi'} - i; \lfloor (1 - \varepsilon) \cdot p \rfloor_{\psi'})) \in \mathcal{M}''$ 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' to all menu entries $f' = (y_j, (\lfloor \vec{y} \rfloor_{\chi''} - j; \lfloor (1 - \varepsilon) \cdot q \rfloor_{\psi''})) \in \mathcal{M}''$ with²⁹ $q < p - \varepsilon$ (where $j \in [n]$, $f = (\vec{y}; q) \in \mathcal{M}$, and $(\chi'', \psi'') \in \{(\chi, \psi), (0, 0)\}$). Indeed, we have that

$$\begin{aligned} u_{e'}(v) &= v_i \cdot x_i + \sum_{k \in [n] \setminus \{i\}} v_k \cdot \lfloor x_k \rfloor_{\chi'} - \lfloor (1 - \varepsilon) \cdot p \rfloor_{\psi'} \geq \sum_{k=1}^n v_k \cdot x_k - (n-1) \cdot H \cdot \chi' - (1 - \varepsilon) \cdot p = \\ &= u_e(v) - (n-1) \cdot H \cdot \chi' + \varepsilon \cdot p \geq u_f(v) - (n-1) \cdot H \cdot \chi' + \varepsilon \cdot p = \sum_{k=1}^n v_k \cdot y_k - q - (n-1) \cdot H \cdot \chi' + \varepsilon \cdot p \geq \\ &\geq v_j \cdot y_j + \sum_{k \in [n] \setminus \{j\}} v_k \cdot \lfloor y_k \rfloor_{\chi''} - \lfloor (1 - \varepsilon) \cdot q \rfloor_{\psi''} - \varepsilon \cdot q - \psi'' - (n-1) \cdot H \cdot \chi' + \varepsilon \cdot p = \\ &= u_{f'}(v) - \varepsilon \cdot q - \psi'' - (n-1) \cdot H \cdot \chi' + \varepsilon \cdot p = u_{f'}(v) + \varepsilon \cdot (p - q) - \psi'' - (n-1) \cdot H \cdot \chi' \geq \\ &\geq u_{f'}(v) + \varepsilon \cdot (p - q) - \frac{\varepsilon^2}{n} - (n-1) \cdot \frac{\varepsilon^2}{n} = u_{f'}(v) + \varepsilon \cdot (p - q) - \varepsilon^2 > u_{f'}(v) + \varepsilon^2 - \varepsilon^2 = u_{f'}(v). \end{aligned}$$

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 . Overall, we therefore obtain that $\text{REV}_{\mathcal{M}''}(F) \geq (1 - \varepsilon) \cdot \text{REV}_{\mathcal{M}}(F) - \varepsilon$, 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 mechanism \mathcal{M}' to be the mechanism offering precisely these menu entries. Let $L \triangleq \bigcup_{i \in [n]} (\{i\} \times \{0, \chi, 2 \cdot \chi, \dots, X \cdot \chi\}^{[n] \setminus \{i\}}) \times \{0, \psi, 2 \cdot \psi, \dots, P \cdot \psi\} \subset \bigcup_{i=1}^n (\{i\} \times [0, 1]^{[n] \setminus \{i\}}) \times [0, E]$. 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 $(i, \vec{x}_{-i}, p) \in L$ 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 set³⁰ $s_{(i, \vec{x}_{-i}, p)} \triangleq \text{Max } S_{(i, \vec{x}_{-i}, p)}$. We define a new IC and IR mechanism \mathcal{M}' as follows:

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

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 $(s_{(i, \vec{x}_{-i}, p)}, \vec{x}_{-i}; p)$ to $(x_i, \vec{x}_{-i}; p) = (x; p)$, 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 $(s_{(i, \vec{x}_{-i}, p)}, \vec{x}_{-i}; p)$ 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 $\text{REV}_{\mathcal{M}'}(F) = \text{REV}_{\mathcal{M}''}(F) \geq (1 - \varepsilon) \cdot \text{REV}_{\mathcal{M}}(F) - \varepsilon$.

We conclude the proof by noting that, as required, there are fewer than $|L| = n \cdot (X + 1)^{n-1} \cdot (P + 1) = n \cdot \lceil \frac{n \cdot H}{\varepsilon^2} + 1 \rceil^{n-1} \cdot \lceil \frac{n \cdot (1 - \varepsilon) \cdot E}{\varepsilon^2} + 1 \rceil$ 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 $(1 - \varepsilon)$. \square

2.7. Connecting the dots

We are now ready to “connect the dots” and use Lemmas 2.2 through 2.5 to prove the qualitative version of 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, to prove a weaker version of the quantitative version of that theorem, i.e., that $C(n, \varepsilon) \leq (n/\varepsilon)^{O(n)}$.

²⁸ As before, taking the closure ensures that a utility-maximizing entry with maximal price exists for every buyer type. See Appendix B.1 for a more details.

²⁹ As we show weak preference and as q is defined via a strict inequality, by continuity of the utility function the correctness of the claim for all f' before taking the closure of \mathcal{M}'' implies its correctness for all f' in the closure as well.

³⁰ Since \mathcal{M}'' is closed, $S_{(i, \vec{x}_{-i}, p)}$ is compact.

Proof. 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 $\max_{i \in [n]} \text{REV}(F_i) = \infty$ (i.e., if there exists $i \in [n]$ s.t. $\text{REV}(F_i) = \infty$), 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;³¹ otherwise, by Theorem 2.1, $\text{REV}(F) < \infty$. If $\max_{i \in [n]} \text{REV}(F_i) = 0$, then the valuation of each item is 0 with probability 1 and so $\text{REV}(F) = 0$ and we are done as well. Assume, therefore, that $\max_{i \in [n]} \text{REV}(F_i) \in (0, \infty)$ and set $\tilde{\varepsilon} \triangleq \varepsilon/6$. By scaling the currency we assume w.l.o.g. that $R \triangleq \max_{i \in [n]} \text{REV}(F_i) = (1 - \tilde{\varepsilon})^{-5} < (6/5)^5 < 5/2$. (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), $F(\mathbb{E}U_H^n) > 0$ and so $F|_{\mathbb{E}U_H^n}$ is well defined. Furthermore (e.g., by Sublemma 2.2.1(a)), $F(\mathbb{E}U_H^n) \cdot \text{REV}(F|_{\mathbb{E}U_H^n}) \leq \text{REV}(F) < \infty$, and so, as $F(\mathbb{E}U_H^n) > 0$, we have that $\text{REV}(F|_{\mathbb{E}U_H^n}) < \infty$. Therefore, by definition of REV , there exists an IC and IR n -item mechanism \mathcal{M} s.t. $\text{REV}_{\mathcal{M}}(F|_{\mathbb{E}U_H^n}) \geq (1 - \tilde{\varepsilon}) \cdot \text{REV}(F|_{\mathbb{E}U_H^n})$. By Lemma 2.3, there exists an E -exclusive IC and IR n -item mechanism \mathcal{M}' s.t. $\text{REV}_{\mathcal{M}'}(F|_{\mathbb{E}U_H^n}) \geq (1 - \tilde{\varepsilon}) \cdot \text{REV}_{\mathcal{M}}(F|_{\mathbb{E}U_H^n})$. By Lemma 2.4, there exists an IC and IR n -item mechanism \mathcal{M}'' with at most $2n$ menu entries that cost more than E , s.t. $\text{REV}_{\mathcal{M}''}(F|_{\mathbb{E}U_H^n}) \geq (1 - \tilde{\varepsilon}) \cdot \text{REV}_{\mathcal{M}'}(F|_{\mathbb{E}U_H^n})$. By Lemma 2.5, there exists an IC and IR n -item mechanism \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. $\text{REV}_{\mathcal{M}'''}(F|_{\mathbb{E}U_H^n}) \geq (1 - \tilde{\varepsilon}) \cdot \text{REV}_{\mathcal{M}''}(F|_{\mathbb{E}U_H^n}) - \tilde{\varepsilon}$. By definition of \mathcal{M}''' , \mathcal{M}' , \mathcal{M} , and \mathcal{M} , we have that

$$\begin{aligned} \text{REV}_{\mathcal{M}'''}(F|_{\mathbb{E}U_H^n}) &\geq (1 - \tilde{\varepsilon}) \cdot \text{REV}_{\mathcal{M}''}(F|_{\mathbb{E}U_H^n}) - \tilde{\varepsilon} \geq (1 - \tilde{\varepsilon})^2 \cdot \text{REV}_{\mathcal{M}'}(F|_{\mathbb{E}U_H^n}) - \tilde{\varepsilon} \geq \\ &\geq (1 - \tilde{\varepsilon})^3 \cdot \text{REV}_{\mathcal{M}}(F|_{\mathbb{E}U_H^n}) - \tilde{\varepsilon} \geq (1 - \tilde{\varepsilon})^4 \cdot \text{REV}(F|_{\mathbb{E}U_H^n}) - \tilde{\varepsilon} \geq (1 - \tilde{\varepsilon})^5 \cdot \text{REV}(F|_{\mathbb{E}U_H^n}), \end{aligned}$$

where the last inequality is since $\text{REV}(F|_{\mathbb{E}U_H^n}) \geq (1 - \tilde{\varepsilon}) \cdot \text{REV}(F) \geq (1 - \tilde{\varepsilon}) \cdot R = (1 - \tilde{\varepsilon})^{-4}$ by Lemma 2.2(b). Therefore, by Lemma 2.2(c), we have that

$$\text{REV}_{\mathcal{M}'''}(F) \geq (1 - \tilde{\varepsilon})^6 \cdot \text{REV}(F) > (1 - 6 \cdot \tilde{\varepsilon}) \cdot \text{REV}(F) = (1 - \varepsilon) \cdot \text{REV}(F).$$

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

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

2.8. A stronger upper bound

The proof of Theorem 1.1 (in its full strength) strengthens the upper bound on $C(n, \varepsilon)$ from $(n/\varepsilon)^{O(n)}$ to $\left(\frac{\log n}{\varepsilon}\right)^{O(n)}$, but at a cost of a more conceptually elaborate presentation. As this stronger bound carries the same qualitative message and is still exponential in n , we have decided, as noted above, to first present the proof of the weaker upper bound of $(n/\varepsilon)^{O(n)}$. We note 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, but utilizing discretization techniques of Hart and Nisan (2013) (see also Dughmi et al., 2014) to discretize to a more carefully chosen logarithmic-sized grid — which is achievable by not only rounding but also “compensating” between coordinates — one may show an improved upper bound of $\left(\frac{\log n}{\varepsilon}\right)^{O(n)}$ on the number of menu entries into which the cheap part of the menu can be discretized 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.) The statement of this stronger version of Lemma 2.5 is given as Lemma C.1 in Appendix C, along with its full proof.

2.9. Arbitrary exclusively unbounded type distributions

We emphasize once more that in contrast to 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, as well as the stronger version of the latter, Lemma C.1) hold for any exclusively unbounded type distribution, and not merely for one obtained by conditioning product distributions upon $\mathbb{E}U_H^n$. Therefore,

³¹ By convention, $(1 - \varepsilon) \cdot \infty = \infty$.

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 $C = C(n, H, \varepsilon) \leq \left(\frac{\log n + \log H}{\varepsilon}\right)^{O(n)}$ such that for every (possibly even highly correlated) $F \in \Delta(\mathbb{E}\mathbb{U}_H^n)$, we have that $\text{REV}_C(F) \geq (1 - \varepsilon) \cdot \text{REV}(F) - \varepsilon$.*

Proposition 2.6 states that a menu size of $\left(\frac{\log n + \log H}{\varepsilon}\right)^{O(n)}$ guarantees³² a multiplicative-and-then-additive (with at most an ε loss in each) approximation to the optimal revenue from any exclusively unbounded type distribution. If the L^1 -norm of the support of this distribution is furthermore bounded away from zero, then so is the optimal revenue (since the optimal revenue is at least the lower bound on that norm, as the seller could price the grand bundle of all items at that lower bound and sell with probability 1), and so in this case the additive approximation from Proposition 2.6 is absorbed into the multiplicative approximation, guaranteeing a purely multiplicative approximation of at least $(1 - \varepsilon)$ for the revenue. (For example, this applies to any distribution over $\mathbb{R}_+ \times [1, H]^{n-1}$.) Proposition 2.6 therefore generalizes a similar result by Dughmi et al. (2014) (which followed Hart and Nisan, 2013, who showed it for two items and gave a weaker result for more items) that shows the same upper bound of $\left(\frac{\log n + \log H}{\varepsilon}\right)^{O(n)}$ for guaranteeing a multiplicative approximation of at least $(1 - \varepsilon)$, but only for distributions where the valuations of *all* items are bounded, i.e., distributions over $[1, H]^n$.

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 therefore 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 1.3, which states that for every $\varepsilon > 0$, there exists $d = d(\varepsilon)$ such that for every number of items $n \in \mathbb{N}$ and $F_1, F_2, \dots, F_n \in \Delta(\mathbb{R}_+)$, we have for $C = n^d$ that $\text{REV}_C(F_1 \times \dots \times F_n) \geq (1 - \varepsilon) \cdot \text{SREV}(F_1 \times \dots \times F_n)$. Concretely, we show that the theorem holds for $d \leq O(\varepsilon^{-5})$.

Definition 3.1 (SREV). Let $n \in \mathbb{N}$ be a number of items and let $F = F_1 \times F_2 \times \dots \times F_n \in \Delta(\mathbb{R}_+)^n$ 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 mechanism by

$$\text{SREV}(F) = \text{SREV}(F_1 \times F_2 \times \dots \times F_n) \triangleq \sum_{i=1}^n \text{REV}(F_i).$$

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

The main idea underlying the proof of Theorem 1.3, which we give in full detail below, is that instead of using n separate mechanisms (one for each item), which may result in an exponential-size menu, we use exponentially fewer separate mechanisms, which result in a polynomial-size menu. Recall that for every item i , by the theorem of Myerson (1981), the optimal separate-selling revenue $\text{REV}(F_i)$ can be obtained via a take-it-or-leave-it offer for selling item i (with probability 1) for a certain price c_i . Let p_i be the probability item i is sold (if it is offered for the price c_i). The separate mechanisms that we use are as follows:

- For each item for which c_i is very small (say, at most an ε/n fraction) compared to $\text{SREV}(F)$, the optimal (separate-selling) revenue $\text{REV}(F_i)$ is very small compared to $\text{SREV}(F)$ as well, and therefore the sum of the optimal revenues from all of these items is small compared to $\text{SREV}(F)$. 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 $\text{REV}(F_i)$, 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 buy at most only 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.

³² Once again, following the proof above would give a weaker upper bound of $\left(\frac{n \cdot H}{\varepsilon}\right)^{O(n)}$, while replacing Lemma 2.5 with Lemma C.1 from Appendix C gives the stronger upper bound.

- We partition the remaining items (those with “nonextreme” values of c_i) into $O(\log n)$ many bundles, each to be offered (via a separate mechanism) to the buyer at 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, not unlike the case of the high-costing items above, we show that the probability of two or more of the items in the bundle being 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 offer the buyer altogether at most $\text{poly}(n)$ choices for (buying any subset of) these bundles.

Proof of Theorem 1.3. Let $\varepsilon \in (0, 1)$. We prove the theorem for³³ $d \triangleq 8194/\varepsilon^5$, and note that the mechanism 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 $r_i \triangleq \text{REV}(F_i)$ the maximal revenue obtainable from F_i . By definition, $\text{SREV}(F) = \text{SREV}(F_1 \times \cdots \times F_n) = \sum_{i=1}^n r_i$. If $n = 1$, then by Remark 3.1 we are done, as selling the single item “separately” (and obtaining revenue $\text{SREV}(F)$) requires at most $1 = n^d$ menu entry (not including $(0; 0)$). If $2 \leq n < 4/\varepsilon$, then we are similarly done, as selling each item separately (and obtaining revenue $\text{SREV}(F)$) requires at most $2^n - 1 < n^{4/\varepsilon} < n^d$ menu entries. Assume henceforth, therefore, that $n \geq 4/\varepsilon$. If $\text{Max}_{i \in [n]} r_i = \infty$, 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.³⁴ If $\text{Max}_{i \in [n]} r_i = 0$, then there is nothing to prove. Assume, therefore, that $\text{Max}_{i \in [n]} r_i \in (0, \infty)$. By scaling the currency we assume w.l.o.g. that $\text{Max}_{i \in [n]} r_i = 1$. (This indeed is w.l.o.g. as our goal is to prove a multiplicative approximation.) Therefore, $\text{SREV}(F) \geq 1$ (and also $\text{SREV}(F) \leq n$).

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 c_i . Let $p_i \triangleq \mathbb{P}_{v_i \sim F_i}[v_i \geq c_i]$ be the probability that c_i is accepted. We note that $r_i = p_i \cdot c_i$. As outlined above, instead of using n separate mechanisms (one for each item), which may result in an exponential-size menu, we use exponentially fewer separate mechanisms, which results in a polynomial-size menu. To describe these separate mechanisms, we first partition the n items into *buckets* based on the optimal price c_i for each item $i \in [n]$, and then describe the separate mechanism that we use 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 mechanism by describing the separate mechanism that we hold for the items in each bucket. (As the buyer’s valuation is additive, holding a compound mechanism comprised of a number of separate IC and IR mechanisms 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, \dots , 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 $(1 + \tilde{\varepsilon})^b$. (Thus, the mechanism for the items in the bucket B_b is itself comprised of a number

³³ The following constant multiplier is used for ease of presentation, and is far from tight.

³⁴ As already noted above, by convention, $(1 - \varepsilon) \cdot \infty = \infty$.

of separate mechanisms, 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 mechanisms is at most

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

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

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

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

- The revenue from the mechanism used for the low bucket L is obviously 0. (Note that had each item from L been sold 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 mechanism used for the high bucket H is at least $(1 - \tilde{\varepsilon}) \cdot \sum_{i \in H} r_i$. To show this, we note that if the buyer values *exactly* one item $i \in H$ by at least c_i , then this item is sold. 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 $p_j = r_j/c_j \geq \tilde{\varepsilon}/n$. Thus, item i is sold with probability at least

$$p_i \cdot \left(1 - \sum_{j \in H \setminus \{i\}} p_j\right) \geq p_i \cdot (1 - (n-1) \cdot \tilde{\varepsilon}/n) > (1 - \tilde{\varepsilon}) \cdot p_i.$$

Therefore, the (expected) revenue from the mechanism used for the high bucket is

$$\sum_{i \in H} \mathbb{P}_{v \sim F}[\text{item } i \text{ is sold}] \cdot c_i > \sum_{i \in H} (1 - \tilde{\varepsilon}) \cdot p_i \cdot c_i = (1 - \tilde{\varepsilon}) \cdot \sum_{i \in H} r_i.$$

(Note that had each item from H been sold 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 mechanism used 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 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 mechanism 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 , and so X_i is a Bernoulli variable taking value 1 with probability p_i . Let $X \triangleq \sum_{i \in B_b} X_i$ and note that 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. To show that the bundle is sold with probability at least $(1 - \tilde{\varepsilon})$, it therefore suffices to show that $\mathbb{P}_{v \sim F}[X < (1 - \tilde{\varepsilon}) \cdot \mu_b] \leq \tilde{\varepsilon}$. 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

$$\mathbb{P}_{v \sim F}[X < (1 - \tilde{\varepsilon}) \cdot \mu_b] \leq \mathbb{P}_{v \sim F}[X < \mu_b - (\tilde{\varepsilon} \sqrt{\mu_b}) \cdot \sigma_b] \leq \frac{1}{\tilde{\varepsilon}^2 \cdot \mu_b} < \tilde{\varepsilon}^{-2} \cdot \tilde{\varepsilon}^3 = \tilde{\varepsilon},$$

as required. (Note that had each item from B_b been sold 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 due once to lower-bounding c_i by $(1 + \tilde{\varepsilon})^b$, a second time to multiplying the requested price by $(1 - \tilde{\varepsilon})$ in order to guarantee a sale with high probability, and a third time to the $\tilde{\varepsilon}$ probability of not selling the bundle.)

- Finally, we claim that the revenue from the mechanism used 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 $\sum_{i \in B_b^j} p_i \cdot (1 - \sum_{k \in B_b^j \setminus \{i\}} p_k) \geq (1 - \tilde{\varepsilon}) \cdot \sum_{i \in B_b^j} p_i$. Therefore, the revenue from this mechanism is at least

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

(Note that had each item from B_b^j been sold 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 due once to lower-bounding c_i by $(1 + \tilde{\varepsilon})^b$, and a second time to the $\tilde{\varepsilon}$ probability of not selling the bundle.)

The total revenue is, therefore, at least

$$\begin{aligned} (1 - \tilde{\varepsilon}) \cdot \left(\sum_{i \in H} r_i \right) + (1 - \tilde{\varepsilon})^3 \cdot \left(\sum_{\substack{i \in B_b: \\ \mu_b > \tilde{\varepsilon}^{-3}}} r_i \right) + (1 - \tilde{\varepsilon})^2 \cdot \left(\sum_{\substack{i \in B_b: \\ \mu_b \leq \tilde{\varepsilon}^{-3}}} r_i \right) &\geq (1 - \tilde{\varepsilon})^3 \cdot \left(\sum_{i \in [n] \setminus L} r_i \right) \geq \\ &\geq (1 - \tilde{\varepsilon})^3 \cdot (\text{SREV}(F) - \tilde{\varepsilon}) \geq (1 - \tilde{\varepsilon})^4 \cdot \text{SREV}(F) > (1 - 4 \cdot \tilde{\varepsilon}) \cdot \text{SREV}(F) = (1 - \varepsilon) \cdot \text{SREV}(F), \end{aligned}$$

as required. \square

Lemma 3.1. *Let $m \in \mathbb{R}_+$, let $p_1, p_2, \dots, p_m \in \mathbb{R}$ 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 $p_1 \geq p_2 \geq \dots \geq p_m$ and let $\ell \triangleq \text{Max}\{\ell \in \{0, 1, \dots, m-1\} \mid \sum_{k=1}^{\ell} p_k \leq 1\} \in \{0, 1, \dots, m-1\}$. 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 $\lceil \sum_{k \in [m] \setminus B^1} p_k \rceil \leq \lceil s \rceil - 1$. 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 $\sum_{k \in B^j \setminus \{i\}} p_k \leq 1$. \square

We conclude by noting that in Lemma 3.1, the bound of $\lceil s \rceil$ on the number of sets in the partition is in fact tight. This can be demonstrated by setting $m = 2 \cdot \lceil s \rceil - 1$, and setting $p_1 \triangleq p_2 \triangleq \dots \triangleq p_{m-1} \triangleq 1/2 + \varepsilon$ 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 $1, \dots, m-1$ 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 1.5, 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 mechanism for this setting offers each item, separately, for a take-it-or-leave-it price of 1; viewed as a menu, this mechanism 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 mechanism 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 mechanism, 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 mechanism 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 mechanism, i.e., in every mechanism having subexponentially many menu entries.

Before proceeding to the proof, we roughly sketch the intuition underlying it. Consider a mechanism that loses less than, say, a $\frac{1}{10n}$ fraction of the revenue.³⁵ 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 mechanism 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” that 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 1.5. We will prove that $C(n, \frac{1}{10n}) > 2^{n/10}$. We note that this implies the theorem, as $C(10n, \frac{1}{10n}) \geq C(n, \frac{1}{10n})$, 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 to accommodate 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%.³⁶ Assume for contradiction that there exists an IC and IR mechanism that has at most $2^{n/10}$ menu entries, but that obtains from F a revenue of at least $(1 - \frac{1}{10n}) \cdot \text{Rev}(F) = (1 - \frac{1}{10n}) \cdot n/2$.

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 probability that both of these buyer types pay full price should on the one hand be large due to the assumption of a slight revenue loss, but on the other hand be small due to 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 to 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 mechanism 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 mechanism 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 .

³⁵ Note that since the optimal revenue is $n/2$, a $\frac{1}{10n}$ fraction of the optimal revenue is $1/20$, so we will in fact also show that when all item distributions are supported on $[0, 1]$, to guarantee an additive loss of at most a fixed $1/20$, a menu size that is exponential in n is required.

³⁶ So, we in fact prove the theorem by showing that $C(10n, \frac{1}{10n}) > 2^{(10n)/100}$ 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%.

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 mechanism is *tiny* if at most $2^{0.9n}$ 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 mechanism, 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 proceed 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 but chooses some nontiny menu entry 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 mechanism. 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 entry³⁷), 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). Summing these two probabilities, 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. \square

Acknowledgments

Yannai Gonczarowski was supported in part 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 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, by Israel-USA Bi-national Science Foundation (BSF) grant number 2014389, and by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 740282). We thank the anonymous referees for helpful feedback.

Appendix A. Menu size and communication complexity

One may view the menu size of a mechanism as exactly characterizing its (deterministic) communication complexity. To see this easy correspondence, we should first carefully define the communication complexity. In the communication complexity scenario, we have a commonly known mechanism (in the intended application it may be determined by the

³⁷ 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 $7/10$.

commonly known distributions F_1, \dots, F_n on item values). Our single buyer also has a privately known valuation, and the seller has no private knowledge. The mechanism is just a function that maps the buyer's valuation to the allocation and payment for this valuation.

Definition A.1 (*Deterministic communication complexity*). The deterministic communication complexity of a given mechanism is the minimal number of bits that must be exchanged, in the worst case, by the buyer and seller in order for both of them to compute the allocation and payment for the buyer's valuation.

We can immediately notice that since the seller has no private information, then the buyer may simulate him completely, and thus there is never any need for the seller to transmit anything. It follows that without loss of generality, the communication is one-sided: the buyer transmits some information about his valuation, which suffices to determine the outcome of the mechanism (allocation and payment). Thus, if some (prefix free) communication protocol for the mechanism uses at most c bits of communication, then there are at most 2^c possible outcomes, i.e., the menu size of the mechanism is at most 2^c . Conversely, if the mechanism has a menu size of C , then clearly the buyer need only transmit her chosen menu entry, which requires $\lceil \log_2 C \rceil$ bits of information. We have thus observed the following.

Proposition A.1. *The deterministic communication complexity of a mechanism is exactly the logarithm (base 2, rounded up) of the menu size of the mechanism.*

Note also that once the mechanism is incentive compatible, this characterization of the communication complexity holds whether or not the communication protocol is required to be incentive compatible: the lower bound on the communication complexity applies even without any incentive requirements, while the upper bound protocol is clearly incentive compatible.

The definition of randomized communication complexity may be more subtle: one definition would just require a randomized protocol that computes, with high probability, the required outcome. A more natural definition, however, would just require the distribution of the outcomes of the protocol to be that specified by the mechanism.

Definition A.2 (*Randomized communication complexity*). The randomized communication complexity of a mechanism is the expected number of bits of communication required, for the worst case input, in a randomized protocol whose probability of allocating each item and whose expected payment are as specified by the mechanism.

As it turns out, this definition of randomized communication complexity is a much easier benchmark to satisfy, at least as long as we allow “public coins” (i.e., as long as we assume that the buyer and seller share a randomly generated string, whose length is not counted toward the communication). For instance, one may in fact implement any single-item mechanism with a single bit of communication under a public-coin model. By a theorem of Myerson (1981), an IC single-item mechanism is completely specified by 1) the probability of allocation $x(v)$ for every value v , where $x(\cdot)$ is a nondecreasing function, and 2) the payment p_0 of a buyer with valuation 0. Suppose that the buyer and seller choose (jointly, without any communication cost) a price p at random so that $x(p)$ is exactly uniform on $[0, 1]$. To get the correct allocation, it suffices for the buyer to get the item if $v \geq p$ (requiring one bit of communication), which indeed happens with probability exactly $x(v)$. If we charge the buyer exactly $p + p_0$ when she gets the item and p_0 when she does not get the item, then the expected payment turns out to be as required.

It is not completely clear to us how large the public-coin randomized complexity can be in the case of multiple items, nor is it clear how large the private-coin randomized complexity can be.

Appendix B. Technical notes

B.1. Infinite menus

B.1.1. Utility-maximizing entries

As is well known in the literature, by the taxation principle every (single-buyer) IC mechanism 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 mechanisms with infinitely many possible outcomes, while it is still true that each such mechanism 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 mechanism. Indeed, in a general infinite menu, a utility-maximizing entry for each buyer type does not necessarily exist, yet in menus corresponding to IC mechanisms, 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 mechanism 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 2.5.³⁸) Indeed, for a buyer type $v = (v_1, \dots, v_n) \in \mathbb{R}_+^n$, all nonnegative-utility entries from a closed menu

³⁸ The reasoning in our proof of Lemma 2.4 is more delicate and is detailed in that proof.

\mathcal{M} lie in the compact set $\mathcal{M} \cap ([0, 1]^n \times [0, \sum_{i=1}^n v_i])$, and by continuity of the utility function, this function attains a maximum value in this compact set.

B.1.2. Tie-breaking by prices

Definition 2.4(a), which is the standard definition of the revenue obtainable by an IC mechanism, 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 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 mechanism 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 mechanisms quite cumbersome. Nonetheless, similarly to above, if an IC mechanism 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 B.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 a mechanism by specifying its menu is in the proof of Theorem 1.1, where we start with an IC (and IR) mechanism (possibly of infinite size) that obtains revenue close to $\text{REV}(F)$; to justify the fact that we can assume w.l.o.g. that the menu of such a mechanism is closed, we note that if some menu entry e is weakly preferred by some buyer type v to all menu entries in some menu \mathcal{M} , then by continuity of the utility function, e is weakly preferred by v also to all menu entries in $\overline{\mathcal{M}}$ – the *closure* of \mathcal{M} (in $[0, 1]^n \times \mathbb{R}_+$). Therefore, given an IC mechanism \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 .

B.2. Arbitrary tie-breaking

As discussed above, the (standard) definition that we use for the revenue obtainable by a given IC mechanism (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., REV_C) or unconstrained (i.e., REV), does not depend on the tie-breaking rule. Indeed, if a mechanism \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 $(1 - \varepsilon)$), yields a mechanism with the same menu size (since a finite menu is always closed) that obtains 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 REV_C and REV 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 REV_C and in favor of high prices in the definition of REV .

Appendix C. A tighter version of Lemma 2.5

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

- $\text{REV}_{\mathcal{M}'}(F) > (1 - \varepsilon) \cdot \text{REV}_{\mathcal{M}}(F) - \varepsilon$.
- There are as many menu entries that cost more than $(1 - \varepsilon/3)^2 \cdot E$ in \mathcal{M}' as there are that cost more than E in \mathcal{M} .
- There are fewer than $(\log E/\varepsilon)^{O(n)}$ entries in \mathcal{M}' that cost at most $(1 - \varepsilon/3)^2 \cdot E$.

Proof. The proof uses discretization techniques from Hart and Nisan (2013) (see also Dughmi et al., 2014). Similarly to the proof of Lemma 2.5, we will gradually define the mechanism \mathcal{M}' through interim mechanisms, however we will have more than one such interim mechanism. We will first define an interim mechanism $\mathcal{M}^{(2)}$ that will have prices (of cheap menu entries) discretized, and we will then use it to sequentially define two more interim mechanisms $\mathcal{M}^{(3)}$ and $\mathcal{M}^{(4)}$ that will also have allocations discretized. Finally, as in the proof of Lemma 2.5, we will derive the required mechanism \mathcal{M}' by defining it to be the mechanism offering only those menu entries of $\mathcal{M}^{(4)}$ that are in fact chosen by any buyer type. Set $\tilde{\varepsilon} \triangleq \varepsilon/3$. Let $K \triangleq \left\lceil \frac{\log E/\varepsilon}{\log(1+\tilde{\varepsilon}^2)} \right\rceil = O(1/\tilde{\varepsilon}^2 \cdot \log E/\varepsilon)$ and let $\pi = \sqrt[K]{E/\varepsilon}$. Note that K is the smallest natural number such that $\pi \leq (1 + \tilde{\varepsilon}^2)$.

We start by defining $\mathcal{M}^{(2)}$. Each of the cheap menu entries of this mechanism, except for $(\vec{0}; 0)$, will only have one of the $K + 1$ prices $p_\ell \triangleq (1 - \varepsilon) \cdot \pi^\ell \cdot \varepsilon$ for $\ell \in \{0, \dots, K\}$. (Note that $p_K = (1 - \varepsilon) \cdot E$.) We construct $\mathcal{M}^{(2)}$ as follows:

- For every menu entry $e = (\bar{x}; p) \in \mathcal{M}$ with $p \in [\varepsilon, E]$, we define p' to be $(1 - \tilde{\varepsilon}) \cdot p$ rounded up to the nearest price in the set $\{p_0, \dots, p_K\}$,³⁹ and add the menu entry $e' = (\bar{x}; p')$ to $\mathcal{M}^{(2)}$. (Note that $p' \leq p_K = (1 - \tilde{\varepsilon}) \cdot E$.)
- For every menu entry $e = (\bar{x}; p) \in \mathcal{M}$ with $p > E$, as well as for the menu entry $e = (\bar{0}; 0) \in \mathcal{M}$, we define $p' = (1 - \tilde{\varepsilon}) \cdot p$ (but without any rounding), and add the menu entry $e' = (\bar{x}; p')$ to $\mathcal{M}^{(2)}$. (In particular, $\mathcal{M}^{(2)}$ is IR.)
- We ignore any other menu entries of \mathcal{M} (i.e., any menu entries other than $(\bar{0}; 0)$, whose price is less than ε).

We note that for each entry $e' = (\bar{x}; p')$ that is added to $\mathcal{M}^{(2)}$, the price of its counterpart $e = (\bar{x}; p)$ from \mathcal{M} satisfies $(1 - \tilde{\varepsilon}) \cdot p \leq p' < \pi \cdot (1 - \tilde{\varepsilon}) \cdot p \leq (1 + \tilde{\varepsilon}^2) \cdot (1 - \tilde{\varepsilon}) \cdot p$. Finally, as in the proof of Lemma 2.5, we define $\mathcal{M}^{(2)}$ to be the closure of the set of menu entries added above to $\mathcal{M}^{(2)}$.

We note that $\text{REV}_{\mathcal{M}^{(2)}}(F) \geq (1 - \tilde{\varepsilon})^2 \cdot (\text{REV}_{\mathcal{M}}(F) - \varepsilon)$. The ε additive loss is due to possibly losing all revenue from buyers who pay less than ε in \mathcal{M} . The multiplicative bound on the loss is ensured by the fact that for any two menu entries $e = (\bar{x}; p)$ and $f = (\bar{y}; q)$ from \mathcal{M} , if $q < (1 - \tilde{\varepsilon}) \cdot p$, then $p' - q' < (1 + \tilde{\varepsilon}^2) \cdot (1 - \tilde{\varepsilon}) \cdot p - (1 - \tilde{\varepsilon}) \cdot q < p - q - \tilde{\varepsilon}^3 \cdot p < p - q$ (and recall that the allocations of f and e are unchanged), so any buyer who chooses e from \mathcal{M} , prefers e' over any menu entry from $\mathcal{M}^{(2)}$ that costs less than $(1 - \tilde{\varepsilon}) \cdot p'$, and so pays at least $(1 - \tilde{\varepsilon}) \cdot p' \geq (1 - \tilde{\varepsilon})^2 \cdot p$ in $\mathcal{M}^{(2)}$.

We now define $\mathcal{M}^{(3)}$ by taking $\mathcal{M}^{(2)}$ and multiplying every allocation probability and every price by $(1 - \tilde{\varepsilon})$. Therefore, $\mathcal{M}^{(3)}$ is IR, has revenue $\text{REV}_{\mathcal{M}^{(3)}}(F) = (1 - \tilde{\varepsilon}) \cdot \text{REV}_{\mathcal{M}^{(2)}}(F) \geq (1 - \tilde{\varepsilon})^3 \cdot (\text{REV}_{\mathcal{M}}(F) - \varepsilon) > (1 - \varepsilon) \cdot \text{REV}_{\mathcal{M}}(F) - \varepsilon$, its menu entries that cost more than $(1 - \tilde{\varepsilon})^2 \cdot E$ are in one-to-one correspondence with the entries of \mathcal{M} that costs more than E , and its other menu entries apart from $(\bar{0}; 0)$ only have prices in $\{(1 - \tilde{\varepsilon}) \cdot p_1, (1 - \tilde{\varepsilon}) \cdot p_2, \dots, (1 - \tilde{\varepsilon}) \cdot p_K\}$. Also, each allocation probability of $\mathcal{M}^{(3)}$ is at most $1 - \tilde{\varepsilon}$.

Finally, we define $\mathcal{M}^{(4)}$. Set $\chi \triangleq \frac{\tilde{\varepsilon}}{K+1}$. We construct $\mathcal{M}^{(4)}$ as follows:

- We add the menu entry $e = (\bar{0}; 0) \in \mathcal{M}^{(3)}$ to $\mathcal{M}^{(4)}$, unchanged.
- For every other menu entry $e = (\bar{x}; p) \in \mathcal{M}^{(3)}$ with $p \leq (1 - \tilde{\varepsilon})^2 \cdot E$, recall that $p = (1 - \tilde{\varepsilon}) \cdot p_\ell$ for some $\ell \in \{0, 1, \dots, K\}$. For each $i \in [n]$, we define $x'_i = \lfloor x_i \rfloor_\chi + \ell \cdot \chi$, i.e., x_i rounded down to the nearest multiple of χ , plus a “bonus” probability that increases with p_ℓ (note that $x'_i \leq 1$ since $x_i \leq 1 - \tilde{\varepsilon}$), and add the menu entry $(\bar{x}'; p)$ to $\mathcal{M}^{(4)}$.
- For every menu entry $e = (\bar{x}; p) \in \mathcal{M}^{(3)}$ with $p > (1 - \tilde{\varepsilon})^2 \cdot E$, we add that menu entry to $\mathcal{M}^{(4)}$, with each allocation probability increased by $\tilde{\varepsilon}$. (These are a well defined probabilities since in $\mathcal{M}^{(3)}$ all allocation probabilities are at most $1 - \tilde{\varepsilon}$).

Finally, we define $\mathcal{M}^{(4)}$ to be the closure of the set of menu entries added above to $\mathcal{M}^{(4)}$. We note that all allocation probabilities have been weakly increased from $\mathcal{M}^{(3)}$ to $\mathcal{M}^{(4)}$, and for every two menu entries $e, f \in \mathcal{M}^{(3)}$, if the price of e is higher than that of f , then each allocation probability in e was weakly increased (even when taking into account also the rounding) by no less than the corresponding allocation probability in f , and therefore the revenue from no buyer type decreases, and so $\text{REV}_{\mathcal{M}^{(4)}}(F) \geq \text{REV}_{\mathcal{M}^{(3)}}(F) > (1 - \varepsilon) \cdot \text{REV}_{\mathcal{M}}(F) - \varepsilon$.

Finally, we will define \mathcal{M}' to be the mechanism offering only the menu entries from $\mathcal{M}^{(4)}$ that are in fact chosen by any buyer type. We conclude the proof by noting that, as required, the number of menu entries (not including $(\bar{0}; 0)$) that cost at most $(1 - \tilde{\varepsilon})^2 \cdot E$ in \mathcal{M}' is at most $(K + 1) \cdot (1/\chi + 1)^{n-1} = (K + 1) \cdot (\frac{K+1}{\tilde{\varepsilon}} + 1)^{n-1} = (\log E/\varepsilon)^{O(n)}$, and the number of menu entries that cost more than $(1 - \tilde{\varepsilon})^2 \cdot E$ in \mathcal{M}' is the number of menu entries that cost more than E in \mathcal{M} . \square

References

- Babaioff, M., Immorlica, N., Lucier, B., Weinberg, S.M., 2014. A simple and approximately optimal mechanism for an additive buyer. In: *Proceedings of the IEEE 55th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 21–30.
- Babaioff, M., Gonczarowski, Y.A., Nisan, N., 2017. The menu-size complexity of revenue approximation. In: *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 869–877.
- Babaioff, M., Nisan, N., Rubinstein, A., 2018. Optimal deterministic mechanisms for an additive buyer. In: *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, p. 429.
- Briest, P., Chawla, S., Kleinberg, R., Weinberg, S.M., 2010. Pricing randomized allocations. In: *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 585–597.
- Cai, Y., Zhao, M., 2017. Simple mechanisms for subadditive buyers via duality. In: *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 170–183.
- Chawla, S., Malec, D.L., Sivan, B., 2010. The power of randomness in Bayesian optimal mechanism design. In: *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*, pp. 149–158.
- Daskalakis, C., Weinberg, S.M., 2012. Symmetries and optimal multi-dimensional mechanism design. In: *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pp. 370–387.
- Daskalakis, C., Deckelbaum, A., Tzamos, C., 2013. Mechanism design via optimal transport. In: *Proceedings of the Fourteenth ACM Conference on Electronic Commerce (EC)*, pp. 269–286.
- Devanur, N.R., Weinberg, S.M., 2017. The optimal mechanism for selling to a budget constrained buyer: the general case. In: *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pp. 39–40.
- Devanur, N.R., Goldner, K., Saxena, R.R., Schwartzman, A., Weinberg, S.M., 2020. Optimal mechanism design for single-minded agents. Mimeo.

³⁹ Note that indeed $(1 - \tilde{\varepsilon}) \cdot p \leq (1 - \tilde{\varepsilon}) \cdot E = p_K$.

- Dughmi, S., Han, L., Nisan, N., 2014. Sampling and representation complexity of revenue maximization. In: Proceedings of the 10th Conference on Web and Internet Economics (WINE), pp. 277–291.
- Fiat, A., Goldner, K., Karlin, A.R., Koutsoupias, E., 2016. The FedEx problem. In: Proceedings of the 17th ACM Conference on Economics and Computation (EC), pp. 21–22.
- Giannakopoulos, Y., Koutsoupias, E., 2014. Duality and optimality of auctions for uniform distributions. In: Proceedings of the Fifteenth ACM Conference on Economics and Computation (EC), pp. 259–276.
- Giannakopoulos, Y., Koutsoupias, E., 2015. Selling two goods optimally. In: Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP), pp. 650–662.
- Goldner, K., Gonczarowski, Y.A., 2018. The menu size of precise and approximate revenue-maximizing auctions. In: Tutorial, the 19th ACM Conference on Economics and Computation (EC). <http://yannai.gonch.name/scientific/ec18-menu-size-tutorial/>.
- Gonczarowski, Y.A., 2018. Bounding the menu-size of approximately optimal auctions via optimal-transport duality. In: Proceedings of the 50th Annual ACM Symposium on Theory of Computing (STOC), pp. 123–131.
- Hart, S., Nisan, N., 2012. Approximate revenue maximization with multiple items. In: Proceedings of the 13th ACM Conference on Electronic Commerce (EC), p. 656.
- Hart, S., Nisan, N., 2013. The menu-size complexity of auctions. In: Proceedings of the Fourteenth ACM Conference on Electronic Commerce (EC), p. 565.
- Hart, S., Nisan, N., 2014. How good are simple mechanisms for selling multiple goods? Discussion Paper 666. Center for the Study of Rationality, Hebrew University of Jerusalem.
- Hart, S., Reny, P.J., 2015. Maximal revenue with multiple goods: nonmonotonicity and other observations. *Theor. Econ.* 10 (3), 893–922.
- Kothari, P., Mohan, D., Schwartzman, A., Singla, S., Weinberg, S.M., 2019. Approximation schemes for a unit-demand buyer with independent items via symmetries. In: Proceedings of the IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), pp. 220–232.
- Li, X., Yao, A.C.-C., 2013. On revenue maximization for selling multiple independently distributed items. *Proc. Natl. Acad. Sci. (PNAS)* 110 (28), 11232–11237.
- Manelli, A.M., Vincent, D.R., 2006. Bundling as an optimal selling mechanism for a multiple-good monopolist. *J. Econ. Theory* 127 (1), 1–35.
- McAfee, R.P., McMillan, J., 1988. Multidimensional incentive compatibility and mechanism design. *J. Econ. Theory* 46 (2), 335–354.
- Morgenstern, J., Roughgarden, T., 2015. On the pseudo-dimension of nearly optimal auctions. In: Proceedings of Advances in Neural Information Processing Systems 28 (NIPS), pp. 136–144.
- Myerson, R., 1981. Optimal auction design. *Math. Oper. Res.* 6 (1), 58–73.
- Rubinstein, A., Weinberg, S.M., 2015. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In: Proceedings of the Sixteenth ACM Conference on Economics and Computation (EC), pp. 377–394.
- Saxena, R.R., Schwartzman, A., Weinberg, S.M., 2018. The menu complexity of “one-and-a-half-dimensional” mechanism design. In: Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 2026–2035.
- Thanassoulis, J., 2004. Haggling over substitutes. *J. Econ. Theory* 117 (2), 217–245.
- Yao, A.C.-C., 2015. An n -to-1 bidder reduction for multi-item auctions and its applications. In: Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 92–109.