2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)

Approximation Schemes for a Unit-Demand Buyer with Independent Items via Symmetries

Pravesh Kothari
Carnegie Mellon University
kotpravesh@gmail.com

Divyarthi Mohan
Princeton University
dm23@cs.princeton.edu

Ariel Schvartzman
Princeton University
acohenca@cs.princeton.edu

Sahil Singla
Princeton University
Institute for Advanced Study
singla@cs.princeton.edu

S. Matthew Weinberg
Princeton University
smweinberg@princeton.edu

***Abstract*—We consider a revenue-maximizing seller with n items facing a single buyer. We introduce the notion of *symmetric menu complexity* of a mechanism, which counts the number of distinct options the buyer may purchase, up to permutations of the items. Our main result is that a mechanism of quasi-polynomial symmetric menu complexity suffices to guarantee a $(1-\varepsilon)$-approximation when the buyer is unit-demand over independent items, even when the value distribution is unbounded, and that this mechanism can be found in quasi-polynomial time.**

Our key technical result is a polynomial-time, (symmetric) menu-complexity-preserving black-box reduction from achieving a $(1-\varepsilon)$-approximation for unbounded valuations that are subadditive over independent items to achieving a $(1-O(\varepsilon))$-approximation when the values are bounded (and still subadditive over independent items). We further apply this reduction to deduce approximation schemes for a suite of valuation classes beyond our main result.

Finally, we show that selling separately (which has exponential menu complexity) can be approximated up to a $(1-\varepsilon)$ factor with a menu of *efficient-linear* ($f(\varepsilon)\cdot n$) symmetric menu complexity.

I. Introduction

Multi-item mechanism design has been at the forefront of Mathematical Economics since Myerson's seminal work resolved the single-item case [Mye81]. Once it became clear that optimal multi-item mechanisms were prohibitively complex, even with just a single buyer (e.g., [RC98], [Tha04], [MV07]), the problem also entered the Theory of Computation through the lens of approximation. As a result, there is now a long line of work developing auctions which are simple, computationally-efficient, and *approximately* optimal [CHK07], [CHMS10], [CMS15], [KW12], [HN17], [LY13], [BILW14], [Yao15], [RW15], [CM16], [CDW16], [CZ17].

These works take a binary view on simplicity, and aim to discover the best approximation guarantees achievable by simple mechanisms. Only recently have works begun to explore the tradeoff between simplicity and optimality, aiming instead to discover how complex a mechanism must be, as a function of ε, in order to guarantee a $(1-\varepsilon)$-approximation to the optimum. This question is studied formally through the lens of *computational complexity* (how much computation is required to find a mechanism guaranteeing a $(1-\varepsilon)$-approximation on a given instance?) and *menu complexity* (how many distinct outcomes need a mechanism induce in order to guarantee a $(1-\varepsilon)$-approximation?).[1] Prior to our work, neither subexponential upper bounds nor superpolynomial lower bounds were known in any multi-dimensional setting for either measure. Our main results provide the first subexponential upper bound through both lenses.

A. Main Result Part 1: Quasi-Polynomial Computational Complexity

Our main results concern a single *unit-demand buyer* with independently drawn values for n items, the same setting considered in seminal work of Chawla, Hartline and Kleinberg which introduced this domain to TCS [CHK07]. Specifically, there is a single seller with n heterogeneous items facing a single buyer with value v_i for each item i and value $\max_{i\in S}\{v_i\}$ for any set S. The seller has independent Bayesian priors D_i over each v_i (so we say that the buyer is drawn from $D := \prod_i D_i$). The seller presents the buyer with a

[1]See Section II for a formal definition of menu complexity, and formal statement of the computational problem.

2575-8454/19/$31.00 ©2019 IEEE
DOI 10.1109/FOCS.2019.00023

IEEE computer society

menu of (randomized allocation, price) pairs (S, p), and the buyer purchases whichever option maximizes her expected utility $(v(S) - p)$.[2] The seller's goal is to find, over all menus, the one which optimizes her expected revenue.

Our first main result considers the computational complexity of this problem, which is known to be computationally hard to solve exactly (unless $P^{NP} = P^{\#P}$), even when each D_i has support three [CDO+15]. On the other hand, works of [CHK07], [CHMS10], [CMS15] establish that a 1/4-approximation can be found in polynomial time. It was previously unknown whether a $(1-\varepsilon)$-approximation (or even a $(1/4+\varepsilon)$-approximation) could be achieved in subexponential time. Part one of our main result provides the first subexponential-time approximation scheme.

Main Result 1 (Informal, see Theorem IV.5)**.** For all $\varepsilon > 0$, a $(1-\varepsilon)$-approximation to the optimal revenue for a unit-demand buyer with independent values for n items can be found in time quasi-polynomial in n.

B. Main Result Part 2: Quasi-Polynomial Symmetric Menu Complexity

Part two of our main result considers the same problem through the lens of menu complexity. Menu complexity, first defined in [HN13], is widely regarded as an insightful yet imperfect measure. Imagine for example the menu which allows the buyer to purchase any desired set S for a price of $|S|$ ("selling separately"). This is ubiquitously accepted as a fairly simple menu (perhaps it has "intrinsic complexity" n, since there are n non-trivial "kinds" of possible outcomes), yet technically it has menu complexity 2^n (because there are 2^n different sets the buyer can purchase).

Prior work addresses this concern in two ways. The first simply proposes alternative definitions, such as additive menu complexity [HN13].[3] This particular definition, however, was later shown to be ill-defined (in the sense that there exist optimal menus for which the additive menu complexity is undefined) [BNR18], and there are no prior alternative proposals. The second approach is to argue that while selling separately may technically have menu complexity 2^n, it is always well-approximated by a menu of polynomial size [BGN17]. The implication is that while the definition is still imperfect, it is at least impossible for a distribution for which selling separately is optimal to witness a super-polynomial lower bound on the menu complexity required for a $(1-\varepsilon)$-approximation.

We propose the first alternative which is well-defined for all menus, the *symmetric menu complexity.* Informally, a menu respects a permutation group Σ if whenever (S, p) is an option on the menu, $(\sigma(S), p)$ is an option as well for all $\sigma \in \Sigma$. A menu has symmetric menu complexity C if there exists a Σ such that the menu is symmetric with respect to Σ and contains at most C distinct equivalence classes under Σ.[4]

Observe that in our motivating example where a buyer could pick S at price $|S|$ has symmetric menu complexity n (the menu is invariant under all permutations), and that importantly the symmetric menu complexity is well-defined for all menus (take Σ to contain only the identity permutation). Of course, the definition is still imperfect, as selling separately at n distinct prices still has symmetric menu complexity 2^n, but the improvement over standard menu complexity is significant (discussed in Section I-G). Part two of our main result establishes that quasi-polynomial symmetric menu complexity suffices for a $(1-\varepsilon)$-approximation. Importantly, note that both parts of our main result hold even for unbounded distributions.

Main Result 2 (Informal, see Theorem IV.5)**.** For all $\varepsilon > 0$, a $(1-\varepsilon)$-approximation to the optimal revenue for a unit-demand buyer with independent values for n items exists with symmetric menu complexity quasi-polynomial in n.

C. Main Result 3: A Reduction from Unbounded to Almost-Bounded

Our proof of the above main results is cleanly broken down into two steps, the first of which we now overview. We provide a black-box reduction from proving computational/menu/symmetric menu complexity bounds for unbounded distributions to proving the same bounds for *almost-bounded* distributions. Roughly, a distribution D is almost-bounded if for each i, distribution D_i is supported on $[0,1] \cup \{W\}$ (think of W as some large number $\gg n$). That is, D_i has at most one value in its support exceeding 1.

This step in our proof applies quite generally, in fact to any distribution which is *subadditive over independent items* (see Section II for definition). This constitutes a key result in its own right due to

[2]Throughout this paper, we abuse notation and write $v(S) := \mathbb{E}[v(S)]$ when S is a set-valued random variable and v is fixed.

[3]A formal definition is not relevant to this discussion, but essentially the definition is designed to address the specific concern raised: the menu contains a list of (randomized allocation, price) pairs, and the buyer may (adaptively or non-adaptively) select any subset of options to purchase. So "selling separately" has additive menu complexity n.

[4]For ease of exposition, this definition is slightly imprecise, see Section II-D for a formal definition.

significant gaps in tractability between unbounded and almost-bounded instances. For example, it was only recently shown that *some* $f(n,\varepsilon) < \infty$ menu complexity suffices for a $(1-\varepsilon)$-approximation on all unbounded distributions which are additive[5] over n independent items [BGN17] (and the proof is quite involved), whereas the analogous result follows for almost-bounded distributions by a folklore discretization argument.

Main Result 3 (Informal, see Theorem III.1)**.** There is a polynomial-time reduction from a multiplicative $(1-\varepsilon)$-approximation for unbounded distributions which are subadditive over independent items to an additive $O(\varepsilon^5)$-approximation for almost-bounded distributions which are subadditive over independent items. If the $O(\varepsilon^5)$-approximation produced on the almost-bounded instance has (symmetric) menu complexity C, the $(1-\varepsilon)$-approximation for the unbounded instance has (symmetric) menu complexity $\leq nC+n$.

Readers familiar with [BGN17] may notice a relationship to their main result, and a detailed comparison is warranted. The main result of [BGN17] asserts that a $(1-\varepsilon)$-approximation for an additive buyer over independent items can be achieved with bounded $((\ln(n)/\varepsilon)^{O(n)})$ menu complexity, which can now alternatively be deduced from Theorem III.1 plus the aforementioned folklore discretization argument (also called a "nudge-and-round"). In comparison to [BGN17], the main qualitative improvement in Theorem III.1 is that we provide a true reduction from unbounded to almost-bounded distributions.[6] The main quantitative improvements are an extension to subadditive over independent items (versus additive) and that the approximation required on the almost-bounded distribution is independent of n (versus $O(\varepsilon^3/n^3)$). It is worth noting that this quantitative improvement is necessary for our previous quasi-polynomial results (see discussion following Theorem III.1), so the removal of dependence on n is significant. It is also worth noting that our proof of Theorem III.1 indeed makes use of several ideas developed in [BGN17], and we identify the connections where appropriate.

D. Approximating Almost-Bounded Distributions via Symmetries

Theorem III.1 takes care of reducing unbounded distributions to almost-bounded ones, but we still need to figure out how to get an additive $O(\varepsilon^5)$-approximation on almost-bounded distributions that are unit-demand over independent items. The folklore discretization argument (roughly: round all values down to the nearest multiple of $O(\varepsilon^{10})$) establishes only that an exponential $1/\varepsilon^{O(n)}$ computational/menu complexity suffices. Perhaps shockingly, no better bounds were previously known, so our remaining task is to improve this.

The unique special case where progress was previously made is if D is heavily symmetric (that is, D is i.i.d., or there are only $o(n)$ distinct marginals of D) [DW12]. In this case, [DW12] establish that an additive ε-approximation with symmetric menu complexity $n^{O(s/\varepsilon^2)}$ can be found in time $n^{O(s/\varepsilon^2)}$ for any distribution D that is unit-demand over independent items with at most s distinct marginals. Of course, our given D may have n distinct marginals, rendering a direct application of their theorem useless. So our key argument here is to show that every D which is almost-bounded and unit-demand over independent items is "close" in a precise metric to some D' which is almost-bounded and unit-demand over independent items with at most $\ln(n)^{1/\varepsilon^{O(1)}}$ distinct marginals.

E. Extensions

Beyond our main results, Theorem III.1 also allows us to conclude the following corollaries:

- For all $\varepsilon > 0$ and $n \in \mathbb{N}$, there exists a finite $f(n,\varepsilon)$ s.t. for all subadditive D over n independent items, a $(1-\varepsilon)$-approximation can be found in $f(n,\varepsilon)$ time which has menu complexity $f(n,\varepsilon)$ (Theorem IV.1).
- For all $\varepsilon > 0$ and all unit-demand D over *i.i.d.* items, a $(1-\varepsilon)$-approximation can be found in polynomial time which has polynomial symmetric menu complexity (Theorem IV.4).
- For all $\varepsilon > 0$ and all additive D over *i.i.d.* items, a $(1-\varepsilon)$-approximation can be found in quasi-polynomial time which has quasi-polynomial symmetric menu complexity (Theorem IV.4).
- For all $\varepsilon > 0$ and all additive D over independent items *where for all* i, $|support(D_i)| = O(1)$, a $(1-\varepsilon)$-approximation can be found in quasi-polynomial time which has quasi-polynomial symmetric menu complexity (Theorem IV.6). Note that the supports of each D_i may be distinct.

The proofs of the first three bullets follow by first applying our reduction (Theorem III.1), and then applying standard (albeit somewhat subtle) nudge-and-round arguments on the resulting almost-bounded distribution.

[5] A valuation function is additive if $v(S) = \sum_{i \in S} v(\{i\})$.

[6] In contrast, [BGN17] wraps up their additive approximation on a bounded distribution via a specific nudge-and-round tailored to the rest of their proof, and are explicit that care is required in this step.

Note that exploiting symmetries for i.i.d. distributions is considerably simpler than for non-i.i.d. distributions as the input is already symmetric (so [DW12] can be applied almost immediately). The final bullet considers a non-i.i.d. setting, but again establishes that all distributions whose marginals have constant support are "close" to symmetric ones. Note that exact solutions for this setting are computationally intractable [DDT14], [CDO$^+$15], and that no subexponential-time approximation schemes were previously known (and posed as an open problem in [DDT14], even for marginals with support two).

Additionally, like [BGN17], we also consider the (symmetric) menu complexity necessary to approximate the revenue achieved by selling separately. For the standard menu complexity, [BGN17] establishes that there is always a menu of size $n^{1/\varepsilon^{O(1)}}$ which guarantees a $(1-\varepsilon)$-fraction of the best revenue achievable by selling separately. We establish an even stronger claim for symmetric menu complexity: there is always a menu of *efficient-linear* size which guarantees a $(1-\varepsilon)$-fraction of selling separately.

Informal Theorem (See Theorem V.1)**.** For all menus M which sell separately, and all D which are additive over independent items, there exists a menu with symmetric menu complexity $f(\varepsilon) \cdot n$ which achieves a $(1-\varepsilon)$-approximation to the revenue of M when buyers are drawn from D (here $f(\varepsilon) = 2^{O(1/\varepsilon^3)}$).

Again, recall that symmetric menu complexity is still an imperfect definition which assigns complexity 2^n to a menu which sells separately at n distinct prices (whereas the "intrinsic complexity" of such a menu is n). But Theorem V.1 asserts that for all ε, there is a menu of linear symmetric menu complexity which achieves the same (up to $(1-\varepsilon)$) revenue guarantees.

Finally, in Section VI we analyze a barrier example for extending our quasi-polynomial bounds for a unit-demand buyer to an additive buyer. Essentially, the distribution is almost-bounded, but (provably) no previous approaches, nor our approach for almost-bounded unit-demand distributions can guarantee better than an 8/9-approximation. We believe that resolving this example will be a fruitful direction for future work to circumvent current barriers.

F. Related Work

The most related work to ours is [BGN17], whose main result establishes that finite menu complexity suffices to guarantee a $(1-\varepsilon)$-approximation for an additive buyer over independent items. As discussed above, our black-box reduction provides both qualitative and quantitative improvements on their work, and makes use of tools they develop. There are numerous other works which study the menu complexity of optimal and approximately optimal auctions, but there is not much technical overlap [BCKW15], [HN13], [FGKK16], [SSW18], [Gon18].

(Quasi-)Polynomial Time Approximation Schemes for a single buyer have been considered in prior works from a few different perspectives. For example, [CD11] develops a PTAS for the optimal *deterministic* item pricing for a unit-demand buyer over independent MHR items, and a QPTAS for a unit-demand buyer over independent regular items.[7] [Rub16] develops a PTAS for the optimal "partition mechanism" for an additive buyer over independent items.[8] The simplest comparison to these works is that we are searching for a good approximation to the optimal (possibly randomized) mechanism, versus a restricted class of mechanisms. [DW12] develop a PTAS for a bounded unit-demand buyer over i.i.d. items by exploiting symmetries. As noted previously, our work provides the first approximation schemes towards the true optimum in unrestricted settings (and also the first application of [DW12] in asymmetric settings).

A series of works also considers the multi-bidder, multi-item case. Works such as [DW12], [CH13] consider special cases (such as i.i.d., MHR, etc.), and are able to exploit symmetries or concentration to prove that simple auctions can approach optimal guarantees. In the general case, [CDW12a], [CDW12b], [CDW13] develop fully-polynomial randomized approximation schemes. These works achieve polynomial dependence on the number of bidders, and the size of a single bidder's support (so with independent items, this would be exponential in n), and bear no technical similarity. Indeed, one of the open questions left by these works is whether it is possible to improve the dependence on n when the items are independent, and our work resolves this affirmatively in the case of a single unit-demand buyer.

Finally, note that the interesting questions indeed surround a $(1-\varepsilon)$-approximation, and not exact solutions. For example, [MV07], [DDT17] establish that optimal mechanisms may have *uncountable* menu complexity. Moreover, even in the case where the marginal of each item has constant support (two, for additive [DDT14], three for unit-demand [CDO$^+$15]),

[7] That is, each $f_i(v)/(1-F_i(v))$ is monotone non-decreasing (MHR) or $v-(1-F_i(v))/f_i(v)$ monotone non-decreasing (regular).

[8] A partition mechanism partitions the items into disjoint bundles and allows the buyer to purchase any subset of bundles.

exact solutions are computationally intractable and subexponential-time approximation schemes were unknown prior to our work (Theorem IV.6).

G. Discussion and Open Problems

We introduce the notion of symmetric menu complexity, and provide the first subexponential time approximation schemes and subexponential bounds on the symmetric menu complexity of $(1-\varepsilon)$-approximately optimal auctions for an unbounded unit-demand buyer over independent items. Our main technical innovations are: (a) a black-box reduction from computational/menu/symmetric menu complexity bounds on unbounded distributions to almost-bounded ones (Theorem III.1), and (b) establishing that a wide class of (asymmetric) almost-bounded distributions are "close to" symmetric distributions in a formal sense (including unit-demand over independent items, or additive over independent items of constant support). We also conclude approximation-schemes for a suite of additional classes of valuation functions.

The notion of symmetric menu complexity itself will likely be of independent interest for future work. Symmetric menu complexity is well defined for any menu, and is always at most the menu complexity. Additionally, an additive or unit-demand buyer can always find their favorite option on a menu of symmetric menu complexity C in poly(n, C) value queries (see the full version of this paper for a short proof). Moreover, Theorem V.1 establishes that selling separately can be approximated arbitrarily well by a menu of linear symmetric menu complexity. These arguments suggest that symmetric menu complexity is a convincing simplicity measure for additive/unit-demand buyers, and the following open questions are directly relevant:

Open Question 1. Does there exist a polynomial-time (respectively, polynomial symmetric menu complexity) approximation scheme for a single unit-demand buyer over independent items?

Does there exist a subexponential-time (respectively, subexponential symmetric menu complexity) approximation scheme for a single additive (or subadditive) buyer over independent items?

Open Question 2. Does there exist a subexponential-time approximation scheme for *multiple* buyers who are (unit-demand/additive/subadditive) over independent items?

Still, symmetric menu complexity by no means "dominates" the traditional menu complexity (for example, a subadditive buyer can find her favorite option on a menu of menu complexity C in poly(n, C) value queries, but the same is not necessarily true for a menu of symmetric menu complexity C, or additive menu complexity C). It is therefore also an important open question (Open Problem 1.6 in [BGN17]) to understand the standard menu complexity required to achieve a $(1-\varepsilon)$ approximation in any of the settings considered in this paper. In this direction, note importantly that our Theorem III.1 allows future work to restrict attention only to almost-bounded distributions.

H. Roadmap

Preliminaries are split into two sections: Section II contains the minimal notation necessary to formally state and overview our results. Extended preliminaries can be found in the full version of this paper. Section III contains a formal statement and brief overview of our reduction from unbounded to almost-bounded instances. Section IV overviews our use of symmetries to derive approximation schemes for asymetric distributions. Section V overviews the connection between selling separately and symmetric menu complexity. Section VI overviews a barrier example for an additive buyer. Missing details and full proofs can be found in the full version of this paper.

II. Preliminaries

In the interest of brevity, we first provide the minimal notation necessary to understand our precise statements and proof overviews. Additional notation for detailed proofs is provided in the full version of this paper.

A. Classes of Distributions

This paper considers instances with a single buyer and n items. The buyer's valuation function $v(\cdot)$ for the items is drawn from some distribution D (written as $v \leftarrow D$), which will always have *independent items*: $D := \prod_i D_i$. We will consider the following classes of valuations:

- **k-demand over independent items.** Each D_i is a single-dimensional distribution. The buyer's value v_i for item i is drawn independently from D_i, and her value for a set S is $\max_{U \subseteq S, |U| \leq k}\{\sum_{i \in U} v_i\}$. When $k = 1$ we say the distribution is **unit-demand** and when $k = n$ we say it is **additive**.
- **Subadditive over independent items.** Each D_i is an arbitrary distribution.[9] Denote by X_i a random variable with distribution D_i. There exists

[9] [RW15] defines D_i to be a distribution over a compact subset of a normed space, but this is not necessary.

a function $V(\cdot,\cdot)$: support$(D) \times 2^{[n]} \to \mathbb{R}_+$ for which a buyer with type $\vec{X}$ has valuation function $v_{\vec{X}}(\cdot)$ satisfying $v_{\vec{X}}(S) := V(\vec{X}, S)$. Moreover, for all $\vec{X} \in$ support(D), function $V(\vec{X}, \cdot)$ is monotone and subadditive.[10] We will often abuse notation and think of the valuation function $v(\cdot)$ as being drawn directly from D.[11]

We will use $D_S := \prod_{i \in S} D_i$ to refer to the distribution D restricted to items in S. Note that we will often abuse notation and use $v(S)$ to refer to $\mathbb{E}_S[v(S)]$ when S is a randomized allocation.

B. Revenue Benchmarks

We will also be interested in the following quantities. If the parameter D is clear from context, we may drop it (but sometimes it will not be clear, and we will make sure to include it).

- $\mathsf{Rev}(D)$: the optimal revenue achievable by any mechanism for a single buyer drawn from D (formally, the supremum of achievable revenues). We will always assume that $\mathsf{Rev}(D)$ is finite.
- $\mathsf{Rev}_M(D)$: for mechanism M, the expected revenue of M for distribution D.
- $\mathsf{Val}(D)$: the expected value of $v([n])$, when $v(\cdot)$ is drawn from D (not necessarily finite).

C. Bounded and Truncated Distributions

In order to formally state our results, we will be interested in the following restrictions on D.

- A distribution D is *unbounded* if $\mathsf{Rev}(D) < \infty$ (but maybe $\mathsf{Val}(D) = \infty$, no other constraints).
- A distribution D is *c-bounded* if $\Pr_{v \leftarrow D}\left[v(\{i\}) \leq c \cdot \mathsf{Rev}(D)\right] = 1$ for all i.
- D is *almost c-bounded* if $\exists X \in \mathbb{R}$ so that for all i, $\Pr_{v \leftarrow D}\left[v(\{i\}) \in [0, c \cdot \mathsf{Rev}(D)] \cup \{X\}\right] = 1$.

Importantly, observe that whenever D is almost c-bounded, we can normalize so that all $v(\{i\}) \in [0,1] \cup \{X/(c \cdot \mathsf{Rev}(D))\}$ with probability 1 (by dividing all values by $c \cdot \mathsf{Rev}(D)$). Now, an additive ε/c-approximation to $\mathsf{Rev}(D)$ immediately implies a multiplicative $(1-\varepsilon)$-approximation to $\mathsf{Rev}(D)$.

Our main results will also involve truncating unbounded distributions into ones which are nearly-bounded. Below we define these truncations formally. The definition below is parameterized by a value $T > 0$ and a vector $\vec{p}$. Intuitively, the truncation operation first replaces all item values $> T$ with exactly T, and then for each item i independently sets a huge value $n^2 \cdot (\max\{1, T\})^3$.

Definition II.1 (Canonical truncations)**.** Let D be subadditive over independent items. Let $T \in \mathbb{R}_+$, and let $\vec{p} \in \mathbb{R}^n$ be a vector of probabilities. Denote by $D(T, \vec{p})$ the truncation of D with respect to $T, \vec{p}$. To sample from the distribution $D(T, \vec{p})$:

1) Draw $v \leftarrow D$. For each item i such that $v(\{i\}) > T$, add i to S. These items will have their value truncated at T.
2) For each item i, independently add i to W with probability $\min\left\{\frac{p_i}{n^2 \cdot (\max\{1,T\})^3}, 1\right\}$. Update $S := S \setminus W$. These items will have their value set at $n^2 \cdot (\max\{1, T\})^3$.
3) Set $v'(\{i\}) := T$ for all $i \in S$, $v'(\{i\}) = n^2 \cdot (\max\{1, T\})^3$ for all $i \in W$.
4) *(Additive truncation)* Output $v'(\cdot)$ with $v'(U) := v(U \cap \bar{S} \cap \bar{W}) + \sum_{i \in U \cap (S \cup W)} v'(\{i\})$.
5) *(Max truncation)* Output $v'(\cdot)$ with $v'(U) := \max\left\{v(U \cap \bar{S} \cap \bar{W}), \max_{i \in U \cap (S \cup W)}\{v'(\{i\})\}\right\}$.

We also use the notation $D(T) := D(T, \vec{0})$.

Our reduction from unbounded to almost-bounded requires truncating the original distribution, and holds for either the additive or max truncation (or many others), so we will not emphasize which is used. We quickly parse what is going on in the definition. Both truncations first initialize $v'(\{i\}) := \min\{v(\{i\}), T\}$. For each i, both truncations then independently select each i with tiny probability[12] and update $v'(\{i\}) := n^2 \cdot \max\{1, T\}^3$. Afterwards, in order to output a complete set function, $v'(\cdot)$ must be defined on all sets (not just the singletons), and the two truncations extend differently.

Observe that when D is additive over independent items and $\Pr\left[v(\{i\}) \leq T\right] = 1$, then $D(T) = D$ under the additive truncation. The same holds for unit-demand and the max truncation. If all we know is that D is subadditive (and $\Pr\left[v(\{i\}) \leq T\right] = 1$), then $D(T)$ does not necessarily equal D under either truncation (but this is fine from the perspective of our results). More importantly, observe that if D is subadditive (resp. XOS, submodular) over independent items, then $D(T, \vec{p})$ is also subadditive (resp. XOS, submodular) over independent items under both truncations. If D is additive (resp. gross substitutes) over independent items, then $D(T, \vec{p})$ is additive (resp. gross substitutes) over independent items under the additive truncation. If D is unit-demand over independent items, then

[10] That is, for all S, T: $v(S) \leq v(S \cup T) \leq v(S) + v(T)$.

[11] We refer the reader to [RW15] for some examples of natural distributions satisfying this definition.

[12] In all applications of this definition, we will have $p_i/T \ll \varepsilon$.

$D(T, \vec{p})$ is unit-demand over independent items under the max truncation. So all of these classes are "closed" under (at least) one of the canonical truncations.

D. Menu Complexity

We will consider two menu complexity measures in this paper. Recall that the Taxation Principle [HDSM79], [GO81] asserts that any mechanism for a single buyer can be represented as a menu of (randomized allocation, non-negative price) pairs, where the buyer selects their favorite pair from the menu (that is, the pair which maximizes the buyer's expected value for the randomized allocation minus the price paid). We will therefore directly refer to a mechanism M as a menu/list of such pairs (which implicitly includes the pair $(\emptyset, 0)$). The first notion we consider is the standard menu complexity from [HN13].

Definition II.2 (Menu Complexity [HN13]). The *menu complexity* of a menu M is simply the size of the list $|M|$. We denote by $\mathrm{MC}(M)$ the menu complexity of M.

The following two definitions introduce our notion of symmetric menu complexity.

Definition II.3 (Symmetries in a Menu). Let S be a randomized allocation, p be a price, and Σ be a subgroup of permutations of $[n]$. Then we denote by (S, p, Σ) the set of (randomized allocation, price) pairs $\bigcup_{\sigma \in \Sigma}\{(\sigma(S), p)\}$. That is, the set (S, p, Σ) contains, for all $\sigma \in \Sigma$, the option to receive for price p the randomized allocation which instantiates the random set S, and then permutes the items according to σ.

Definition II.4 (Weak/Strong Symmetric Menu Complexity). We say that a mechanism M has *strong symmetric menu complexity* equal to the smallest c such that there exists an index set $\mathcal{I}$ of size c, collection of (randomized allocation, price) pairs $\{(S_i, p_i)\}_{i \in \mathcal{I}}$, and subgroup Σ of item permutations such that M can be written as $\bigcup_{i \in \mathcal{I}}\{(S_i, p_i, \Sigma)\}$. We refer to the strong symmetric menu complexity of M as $\mathrm{SSMC}(M)$.

We say that M has *weak symmetric menu complexity* equal to the smallest c such that there exists an index set $\mathcal{I}$ and menus $\{M_i\}_{i \in \mathcal{I}}$ such that each $\mathrm{SSMC}(M_i) = d_i$ for all i, menu $M = \bigcup_i M_i$, and $\sum_i d_i = c$. We will refer to the weak symmetric menu complexity of M as $\mathrm{WSMC}(M)$.

Above, the idea is that the mechanism designer can present any mechanism M to the buyer with a description of Σ via its generating set, together with a list of $\mathrm{SSMC}(M)$ (randomized allocation, price) pairs. Similarly, the designer can present any mechanism M to the buyer with a set of such lists, totaling $\mathrm{WSMC}(M)$ (randomized allocation, price) pairs (again representing each Σ_i via its generating set).

In principle, one might find some subgroups Σ to be simpler than others (e.g., the subgroup of all permutations, or all permutations on even elements, etc.), but Jerrum's filter establishes that all subgroups have a generating set of size at most n [Jer82]. So while some subgroups may indeed be conceptually simpler than others, from the point of view of how much space is needed to define Σ, the space is always $n^2 \ln(n)$ (this sanity checks, for instance, that it is not the case that all menus have low symmetric menu complexity simply because they can be cleverly partitioned into few heavily-symmetric parts. See further discussion in the full version of this paper).

Note also that the weak/strong symmetric menu complexity is well-defined for any menu M (by taking Σ to be the trivial subgroup), and that for all M, we have $\mathrm{WSMC}(M) \leq \mathrm{SSMC}(M) \leq \mathrm{MC}(M)$. This is in contrast to previously posed notions such as "additive menu complexity" [HN13], as some menus may simply not admit an additive description (and therefore their additive menu complexity is undefined) [BNR18].

To simplify presentation, we formally define what it means for a class of distributions to have a low $(1-\varepsilon)$-approximation menu complexity.

Definition II.5 (ε-Menu Complexity of a Class of Distributions). Let $\mathcal{D}$ be a class of distributions. Define the ε-Menu Complexity $\mathrm{MC}(\mathcal{D}, \varepsilon)$ of $\mathcal{D}$ to be the minimum c such that for all $D \in \mathcal{D}$ there exists a menu M with $\mathrm{MC}(M) \leq c$ and $\mathrm{Rev}_M(D) \geq (1-\varepsilon)\mathrm{Rev}(D)$. We also define $\mathrm{WSMC}(\mathcal{D}, \varepsilon)$ and $\mathrm{SSMC}(\mathcal{D}, \varepsilon)$ similarly.

E. Computational Problems

Finally, we define the computational problem we consider for our PTAS/QPTAS. Below, when we describe a distribution D as being input, we do not explicitly specify how the input is given, other than (a) it is possible to sample from D in time $\mathrm{poly}(n)$ and (b) for any $T \in \mathbb{R}$, $\varepsilon > 0$, and all items i, it is possible to find $\sup_{p \geq T}\left\{p \cdot \Pr[v(\{i\}) \geq p]\right\}$, along with an r satisfying $r \cdot \Pr[v(\{i\}) \geq r] \geq (1-\varepsilon) \cdot \sup_{p \geq T}\left\{p \cdot \Pr[v(\{i\}) \geq p]\right\}$ in time $\mathrm{poly}_{\varepsilon}(n)$.[13] Observe that if the support of each D_i is explicitly listed and of size $\mathrm{poly}(n)$, then both these properties are satisfied (even though the support of D is exponential in n).

[13]Observe that this supremum is always finite when $\mathrm{Rev}(D)$ is finite.

Definition II.6 (Implicit description of a menu [DDT14])**.** An implicit description of a menu M is a Turing machine which takes as input a valuation $v(\cdot)$ and outputs $\arg\max_{(S,p)\in M\cup\{(\emptyset,0)\}}\{v(S)-p\}$. The description has overhead c if on input $v(\cdot)$ described using b bits, the Turing machine terminates in time $\text{poly}(c,b)$.

Definition II.7 (Computational Revenue Maximization)**.** A $(1-\varepsilon)$ approximation for the problem $\textsc{RevMax}_{\mathcal{D}}$ takes as input $D\in\mathcal{D}$ and outputs an implicit description of a menu M such that $\text{Rev}_M(D)\geq(1-\varepsilon)\text{Rev}(D)$. Whenever we say that an algorithm for $\textsc{RevMax}_{\mathcal{D}}$ runs in time c, we mean both that the implicit description is found in time c, and that the implicit description itself has overhead c.

III. Overview: Reduction from Unbounded to Bounded

In this section, we overview our polynomial-time (symmetric) menu-complexity preserving reduction from unbounded distributions to almost-bounded distributions. Theorem III.1 is the main result of this section.

Theorem III.1. *For any $n\in\mathbb{N}$ and $\varepsilon>0$, let $\mathcal{D}$ be a class of distributions that is closed under one of the canonical truncations, such that every $D\in\mathcal{D}$ is subadditive over n independent items. Let $\mathcal{D}_B$ denote the subset of $\mathcal{D}$ that is also almost $1/\varepsilon^4$-bounded. Then there is a $\text{poly}(n,1/\varepsilon)$-time reduction from achieving a $(1-\varepsilon)$-approximation to $\textsc{RevMax}_{\mathcal{D}}$ to achieving a $(1-O(\varepsilon))$-approximation to $\textsc{RevMax}_{\mathcal{D}_B}$. Moreover:*

$$MC(\mathcal{D},\varepsilon)\leq n+nMC(\mathcal{D}_B,O(\varepsilon))\quad \text{and}$$
$$WSMC(\mathcal{D},\varepsilon)\leq n+nWSMC(\mathcal{D}_B,O(\varepsilon)).$$

The proof of Theorem III.1 is broken down into two main parts. The first half, captured in Proposition III.2,[14] asserts that for any distribution which is subadditive over independent items, there exists a $(1-\varepsilon)$-approximate menu of a particular form. Readers familiar with [BGN17] will notice simliarity to their Lemma 2.4; we discuss the differences shortly after.

Proposition III.2. *Let $E\geq\text{Rev}(D)/\varepsilon^3$. Then for all D that are subadditive over independent items, there exists a menu M such that $\text{Rev}_M(D)\geq(1-O(\varepsilon))\text{Rev}(D)$ and:*

- *For all $(S,p)\in M$, either $p\leq E$, or there exists at most one i such that $\Pr[i\in S]>0$.*
- *For each item i, there exist at most two distinct $(S,p)\in M$ such that $p>E$ and $\Pr[i\in S]>0$.*

[14]See full paper for a more precise version of Proposition III.2.

The structure of the promised M is identical to Lemma 2.4 of [BGN17]. The key difference is that we take $E\geq\text{Rev}(D)/\varepsilon^3$, versus their $E\geq n^3\text{Rev}(D)/\varepsilon^2$. This quantitative improvement is crucial for Theorem III.1: without it, instead of reducing to $1/\varepsilon^4$-bounded distributions, we would only reduce to $\text{poly}(n/\varepsilon)$-bounded distributions. Our positive results for c-bounded distributions require runtime/symmetric menu complexity exponential in c, so the quantitative difference is significant. The second difference is the extension to distributions which are subadditive over independent items (their Lemma 2.4 holds for additive).

The second half of Theorem III.1 is Proposition III.5 below.[15] We first need some definitions. Definition III.3 describes an operation which appears in [BGN17], which takes a menu M and replaces all "expensive" options in M with options which award at most a single item with non-zero probability. Definition III.4 defines a new operation, which takes a menu M and concatenates it with n new options which each offer a single item deterministically.

Definition III.3 (Making a menu E-exclusive)**.** For a given menu M, let the menu $M|_E$ ("M made exclusive above E") denote the menu constructed from M as follows:

- For any $(S,p)\in M$ with $p\leq E$, add $(S,(1-\varepsilon)p)$ to $M|_E$.
- For any $(S,p)\in M$ with $p>E$, and all items i, let $X_i(S)$ denote the (randomized) set that is $\{i\}$ with probability $\Pr[i\in S]$ and $\emptyset$ otherwise. Add $(X_i(S),(1-\varepsilon)p)$ to $M|_E$.

Definition III.4 (Concatenating a menu with exclusive options)**.** Let $\vec{r}\in\mathbb{R}^{[n]}$ be a vector of reserve prices, and $T\in\mathbb{R}$. Let also (S_i,q_i) be the option in M that would be purchased by a buyer with value $v(\cdot)$ satisfying $v(S)=T\cdot\mathbb{I}(i\in S)$.[16] Then $M^{T,\vec{r}}$ ("M concatenated with exclusive options $\vec{r}$") is the menu M with the n additional options $\bigcup_i\left\{(\{i\}\ ,\ q_i+r_i\cdot(1-\Pr[i\in S_i]))\right\}$, and then multiplying all prices by $(1-\varepsilon)$.

That is, concatenating a menu with exclusive options $\vec{r}$ adds, for all i, an option to purchase item i deterministically. The price-per-additional-probability of getting item i beyond what is already allocated by S_i is r_i (and then all prices are multiplied by $(1-\varepsilon)$).

[15]Again, see full paper for a more precise version of Proposition III.5.

[16]Morally, one should think of (S_i,q_i) as the option in M which awards i with highest probability. However, if M has infinite menu complexity, then this option need not be well-defined. The definition is given as such to avoid overly cumbersome notation with supremums.

Proposition III.5. *Let M be the mechanism promised by Proposition III.2, and let $T \geq E/\varepsilon \geq \text{Rev}(D)/\varepsilon^4$. Also, let $p_i := \sup_{r \geq T} \left\{ r \cdot \Pr[v(\{i\}) \geq r] \right\}$ and let $r_i \geq T$ be such that $r_i \cdot \Pr[v(\{i\}) \geq r_i] \geq (1-\varepsilon)p_i$. Then:*

- $(1 - O(\varepsilon)) \cdot \text{Rev}_M(D) \leq \text{Rev}(D(T, \vec{p}))$.
- *For any M', we have* $\text{Rev}_{(M'|_E)^{T,\vec{r}}}(D) \geq \text{Rev}_{M'}(D(T, \vec{p})) - O(\varepsilon) \cdot \text{Rev}(D)$.

Proposition III.5 establishes both that the optimal revenue for D and $D(T, \vec{p})$ is close, and also that any mechanism for $D(T, \vec{p})$ can be efficiently transformed into one which achieves similar guarantees for D. Proposition III.5 has no real analogue in [BGN17], but replaces their Lemma 2.5. Their Lemma 2.5 specifies a particular discretization of the "cheap part" of the menu promised in Lemma 2.4 (priced $\leq E$) which is compatible with the remaining $2n$ expensive options (priced $> E$). The entire challenge of this process is ensuring that a buyer with value $v(\cdot)$ who chooses to purchase one of the $2n$ expensive options from the promised M does not all of a sudden wish to purchase a cheap option instead after discretization. As a result, one cannot simply view the cheap part and expensive part as separate subproblems. The main insight in [BGN17]'s Lemma 2.5 is that a particular discretization of the cheap part does not interfere with the $2n$ expensive options. The main insight in our Proposition III.5 is more general: the *only* property of the cheap part which interacts with the expensive part is the maximum probability with which item i is ever allocated (and this is captured in a somewhat roundabout way by the inserted point-masses at $n^2(\max\{1, T\})^3$). So in comparison to Lemma 2.5 in [BGN17], the key contribution of Proposition III.5 is that it provides a true reduction from unbounded to almost-bounded distributions. Future work (and the remainder of the present work) can simply focus on almost-bounded distributions, rather than separately ensuring that the resulting menu is compatible with Proposition III.2.

Observe that Theorem III.1 now follows from Propositions III.2 and III.5. For any $D \in \mathcal{D}$, $D(T, \vec{p})$ is almost $1/\varepsilon^4$-bounded. So if we can find a $(1 - O(\varepsilon))$-approximation for $D(T, \vec{p})$, we can efficiently make it E-exclusive and concatenate the expensive options, and these changes increase the (symmetric) menu complexity by at most a factor of n, plus an additional n. Chaining the inequalities in Propositions III.2 and III.5 establishes that the resulting menu is a $(1 - O(\varepsilon))$-approximation. A more detailed outline and complete proofs can be found in the full version of this paper.

IV. Overview: Symmetries and Optimal Mechanisms

With Theorem III.1 in hand, our task is now to design good menus for almost-bounded distributions. Unfortunately, there is not much prior work in this direction (even for bounded distributions). It is only known that, via an application of nudge-and-round arguments, one can discretize all values into multiples of ε/n while losing at most an ε fraction of the optimal revenue, at which point an exponentially large linear program can output an explicit description of the optimal mechanism (for the discretized distribution). Note that the linear program has size polynomial in the support of D, which would remain exponential in n even if D is unit-demand/additive and each marginal has support size 2 (as D is a product distribution). We omit the overview of this linear program and defer it to the full version of this paper.

Despite this simple result for bounded distributions, prior work of [BGN17] was the first to establish even that *some* bounded menu complexity suffices for unbounded additive distributions over independent items, and Theorem III.1 now allows us to do the same for unbounded distributions which are subadditive over independent items. A proof of the following theorem appears in the full version of this paper, which formalizes the above paragraph and mostly follows from Theorem III.1.

Theorem IV.1. *Let $\mathcal{D}$ be the class of distributions which are subadditive over n independent items. Then for all ε, there exists a finite number $C(n, \varepsilon) < \infty$ such that $MC(\mathcal{D}, \varepsilon) \leq C(n, \varepsilon)$.*

One special case where progress was made is if the underlying distribution D is *symmetric* [DW12]. Specifically, if D is invariant under all permutations in Σ, [DW12] shows that the canonical LP referenced above can be simplified to have size only $|\text{support}(D)/\Sigma|$ (that is, the LP needs to only consider one representative from each equivalence class in the support of D under Σ), which allowed them to conclude a PTAS when D was bounded and unit-demand over i.i.d. items. We recap (a slight generalization of) their main result below, and include a proof in the full version of this paper for completeness.

Definition IV.2 (Invariant under item permutations)**.** We say that Σ is an *item permutation group* if there exists a partition $T_1 \sqcup \ldots \sqcup T_s$ of $[n]$ such that Σ is the subgroup generated by $\bigcup_i \{(x, y)\}_{x, y \in T_i}$. That is, Σ is generated by all swaps of pairs of elements in the same

T_i.[17] We further let s (the number of parts necessary for the partition) denote the *partition size* of Σ.

We say that D is *symmetric with respect to* Σ if for all $\sigma \in \Sigma$, the distribution which first draws $v(\cdot) \leftarrow D$ and then outputs $v'(\cdot)$ with $v'(S) := v(\sigma(S))$ is itself the distribution D.

For example, if D is k-demand over i.i.d. items, then D is symmetric with respect to the group of all permutations on n items, and this permutation group has partition size one. If D is k-demand over independent items and all values for even items are drawn i.i.d., and odd items are drawn i.i.d. (but from a different distribution than the even items), D is symmetric with respect to the item permutation group Σ with T_1 := even items and T_2 := odd items, and therefore Σ has partition size two.

Theorem IV.3 ([DW12])**.** *For any item permutation Σ of partition size s, let $\mathcal{D}$ be the class of distributions D which are k-demand over independent items, with each $|support(D_i)| \leq c$, and symmetric with respect to Σ. Then an optimal solution to* RevMax$_\mathcal{D}$ *can be found in time* poly(n^{cs}). *Moreover, the mechanism output M has $SSMC(M) \leq n^{cs}$.*

Our applications of Theorem IV.3 will first discretize a distribution into one which is symmetric with respect to a Σ of low partition size, and also where each marginal is supported on not many values. Theorem IV.4 considers i.i.d. items, where the input is already heavily symmetric and we just need to bound the loss from discretizing values. We treat this case as a warmup, but even here our results make quantitative improvements on [DW12], and also now extend their work to the unbounded case due to Theorem III.1. In particular, observe that Theorem IV.4 concludes a PTAS for distributions which are unit-demand over i.i.d. items, and a QPTAS for distributions which are additive over i.i.d. items.

Theorem IV.4. *Let $\mathcal{D}$ be the class of valuations which are k-demand over i.i.d. items. Then $WSMC(\mathcal{D}, \varepsilon) \leq n^{O(\ln(k)/\varepsilon^{12})}$, and there exists a $(1-\varepsilon)$-approximation to* RevMax$_\mathcal{D}$ *which runs in time $n^{O(\ln(k)/\varepsilon^{12})}$.*

Our main application of Theorem IV.3 will be on *arbitrary* distributions which are unit-demand over independent items. That is, the initial distribution might have no symmetries whatsoever. Still, we show that it is possible to discretize the distribution in a way which creates symmetries. Note that discretizing only the values clearly no longer suffices, as there are fully asymmetric distributions even when each marginal has support size two. So we additionally need to discretize the probabilities. The main challenge here is that unless our discretization is excessively fine (that is, too fine to improve the runtime/menu complexity), there will almost certainly be *some* item for which the original and the discretized values are quite different. So we need to carefully dive into the details of an advanced nudge-and-round argument to figure out exactly which item values contribute to lost revenue. This careful dive is possible for unit-demand valuations because optimal menus award at most one item without loss of generality. For additive valuations, there is a barrier to this approach, which we expound in Section VI.

Theorem IV.5. *Let $\mathcal{D}$ be the class of valuations which are unit-demand over independent items. Then there exists a $(1-\varepsilon)$-approximation to* RevMax$_\mathcal{D}$ *which runs in time $n^{O(\ln(n/\varepsilon))^{1/\varepsilon^7}}$ and $WSMC(\mathcal{D}, \varepsilon) \leq n^{O(\ln(n/\varepsilon))^{1/\varepsilon^7}}$.*

Finally, Theorem IV.6 below establishes that a similarly careful nudge-and-round yields a quasi-polynomial approximation scheme for distributions where each marginal has support at most c (even if that support is distinct for each marginal). Recall that even when $c = 2$, no subexponential approximation schemes are previously known, and this is left open by [DDT14].

Theorem IV.6. *Let $\mathcal{D}$ be the class of valuations which are k-demand over independent items and satisfy $|support(D_i)| \leq c$ for all i. Then there exists a $(1-\varepsilon)$-approximation to* RevMax$_\mathcal{D}$ *which runs in time $n^{O(\ln(n/\varepsilon))^c}$ and $WSMC(\mathcal{D}, \varepsilon) \leq n^{O(\ln(n/\varepsilon))^c}$.*

V. Overview: Selling Separately with Low Symmetric Menu Complexity

One justified critique of menu complexity is that it assigns menu complexity 2^n to the "selling separately" mechanism, which places price p_i on each item and allows the buyer to purchase any set S for price $\sum_{i \in S} p_i$. Symmetric menu complexity is still imperfect in this regard: if all p_i are distinct, the menu will still have strong/weak symmetric menu complexity 2^n.

[BGN17] provide a nice response to this critique, by proving that while technically selling separately is deemed to have 2^n menu complexity, for every M which sells separately and for all D which are additive over independent items, there exists another M' for which $\mathrm{Rev}_{M'}(D) \geq (1-\varepsilon)\mathrm{Rev}_M(D)$ and $\mathrm{MC}(M') \leq$

[17] Put another way, Σ contains exactly permutations which separately permute items in T_i, for all i.

$n^{1/\varepsilon^{O(1)}}$.[18] So while the definition of menu complexity is certainly still imperfect, we at least now know that if (for instance) a distribution D admits no good mechanisms of polynomial menu complexity, it is not because selling separately is close to optimal.

We provide an even stronger response in the case of symmetric menu complexity: when D is additive over independent items, and M sells separately, there exists another M' for which $\mathsf{Rev}_{M'}(D) \geq (1-\varepsilon)\mathsf{Rev}_M(D)$ and $\mathrm{SSMC}(M') \leq f(\varepsilon)n$. That is, the blow up from the "intrinsic complexity" (or description complexity) of selling separately n items to the strong symmetric menu complexity of a menu that is almost as good is just a multiplicative factor independent of n. The proof can be found in the full version of this paper.

Theorem V.1. *Let D be additive over independent items and M be a mechanism which sells separately. Then there exists a mechanism M' with $\mathsf{Rev}_{M'}(D) \geq (1-\varepsilon)\mathsf{Rev}_M(D)$ and $SSMC(M') \leq f(\varepsilon)n$, for $f(\varepsilon) = 2^{O(1/\varepsilon^3)}$.*

VI. A Barrier Example for an Additive Buyer

In this section we highlight an example of a $(1/\varepsilon)$-bounded distribution which is additive over independent items but serves as a barrier to proving good (symmetric) menu complexity bounds. In a formal sense, the known approaches for bounding the menu complexity of $(1-\varepsilon)$-approximately optimal mechanisms for a bounded distribution are:

- Argue that $\mathsf{SRev}(D) \geq (1-\varepsilon)\mathsf{Val}(D)$, perhaps because each D_i is nearly a point-mass (recall that while selling separately does not itself have low menu complexity, this suffices by Theorem V.1).
- Argue that $\mathsf{BRev}(D) \geq (1-\varepsilon)\mathsf{Val}(D)$, perhaps because for all i, $v(\{i\}) \leq \varepsilon^2\mathsf{Val}(D)$ with probability one, and therefore $v([n])$ concentrates tightly around its expectation.
- (New, from Section IV) Argue that D is "close" to a highly symmetric distribution D'. Then use Theorem IV.3 to argue that D' has a near-optimal mechanism of low symmetric menu complexity, followed by a claim to argue that this menu (with discounts) also suffices for D.

We provide an example for which all three of these approaches fail, highlighting the main challenge for future work. We overview the construction in Example VI.1 below, and highlight its main features in Proposition VI.2, deferring a proof of Proposition VI.2 to the full version of this paper.

Example VI.1. For even n, an $\varepsilon < 1$ (one interesting choice discussed below is $\varepsilon = 1/9$), and $k = \Theta(\frac{\ln(n)}{\varepsilon})$, consider n vectors in $\vec{r}_1, \ldots, \vec{r}_n \in \{0,1\}^k$ with $|\vec{r}_i|_1 = k/2$ for all i (i.e., each has exactly $k/2$ ones and $k/2$ zeroes). Let it also be the case that for all $i, j \in [n]$, we have $|\{\ell : r_{i\ell} \neq r_{j\ell}\}| \geq k/6$.[19] Define D_i so that:

$$\Pr[v_i = x] = \begin{cases} \frac{\varepsilon}{n} \cdot \mathbb{I}(i \text{ is even}) & x = 1 \\ \frac{2\varepsilon}{n} \cdot \mathbb{I}(i \text{ is odd}) & x = \frac{1}{2} \\ \frac{(\ln n)(1-\varepsilon)^{-\ell}}{nk} \cdot \mathbb{I}(r_{i\ell} = 1) & x = \frac{\varepsilon(1-\varepsilon)^{\ell}}{\ln n}, \\ & \ell = 0, \ldots, k-1 \\ 1 - \Pr[v_i > x] & x = 0 \end{cases}$$

Proposition VI.2. *Any distribution D satisfying the definition in Example VI.1 has the following properties:*

- $\mathsf{Val}(D) = 3\varepsilon/2$.
- $\mathsf{SRev}(D) = \varepsilon$.
- $\mathsf{BRev}(D) \leq \varepsilon$.
- *$v_i \leq 1$ for all i with probability 1. Therefore, D is $1/\varepsilon$-bounded.*
- *For all Σ with partition size s, and all D' such that D' is symmetric with respect to Σ, the coupling distance $\delta(D, D') \geq \varepsilon^2 \cdot \frac{n-s}{6n}$. In particular, if $s = o(n)$, then $\delta(D, D') \geq \varepsilon^2/6 - o(\varepsilon^2)$* [20].

Observe that Proposition VI.2 rules out any of the known approaches achieving better than a 8/9-approximation. Indeed, in order to prove that either of $\mathsf{SRev}(D)$ or $\mathsf{BRev}(D)$ beats a 2/3-approximation, a better bound than $\mathsf{Val}(D)$ on the optimal revenue would be necessary. This may indeed be the right approach, but because D is $1/\varepsilon$-bounded, it is already well inside the range where techniques like those of our Section III or [BGN17] can yield traction. This rules out the first two approaches. In addition, the partition size of Σ appears in the *exponent* of the symmetric menu complexity for optimal mechanisms on distributions that are symmetric with respect to Σ, so the final bullet asserts that a direct application of the approach in Section IV cannot beat a $(1-\sqrt{\varepsilon}/8)$-approximation with subexponential symmetric menu complexity. In particular, the construction is valid for any $\varepsilon < 1$, so we can take ε small enough to have $\sqrt{\varepsilon}/8 < 1/9$, which would rule out an 8/9-approximation via any of the three known approaches.

Focusing on arguments for this example (and slight generalizations) should be illuminating for future

[18] Note that "selling separately" is not obviously simple when D is not additive over independent items, so it is not clear that one should expect/demand such an M' unless D is additive over independent items.

[19] We prove (in the full version) that such vectors exist, which follows by the probabilistic method.

[20] See the full version of this paper for the definition of coupling distance $\delta(D, D')$ between two distributions D, D'.

progress. It seems that the missing ingredient is a near-optimal bound on the optimal revenue without relying on coupling with a symmetric distribution. The interesting feature of D is that it is highly asymmetric, but the values in the support of the marginals which contribute to the asymmetry are small (much smaller than $\mathsf{Val}(D)$). Normally, this would imply that the expected value for the grand bundle concentrates, but the point masses at 1 (or 1/2) ruin such a concentration.

References

[BCKW15] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. Pricing lotteries. *J. Economic Theory*, 156:144–174, 2015.

[BGN17] Moshe Babaioff, Yannai A. Gonczarowski, and Noam Nisan. The menu-size complexity of revenue approximation. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 869–877, 2017.

[BILW14] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2014.

[BNR18] Moshe Babaioff, Noam Nisan, and Aviad Rubinstein. Optimal deterministic mechanisms for an additive buyer. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, page 429, 2018.

[CD11] Yang Cai and Constantinos Daskalakis. Extreme-value theorems for optimal multidimensional pricing. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 522–531, 2011.

[CDO+15] Xi Chen, Ilias Diakonikolas, Anthi Orfanou, Dimitris Paparas, Xiaorui Sun, and Mihalis Yannakakis. On the complexity of optimal lottery pricing and randomized mechanisms. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 1464–1479, 2015.

[CDW12a] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. An Algorithmic Characterization of Multi-Dimensional Mechanisms. In *the 44th Annual ACM Symposium on Theory of Computing (STOC)*, 2012.

[CDW12b] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization. In *the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2012.

[CDW13] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Understanding Incentives: Mechanism Design becomes Algorithm Design. In *the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2013.

[CDW16] Yang Cai, Nikhil Devanur, and S. Matthew Weinberg. A duality based unified approach to bayesian mechanism design. In *Proceedings of the 48th ACM Conference on Theory of Computation(STOC)*, 2016.

[CH13] Yang Cai and Zhiyi Huang. Simple and Nearly Optimal Multi-Item Auctions. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2013.

[CHK07] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic Pricing via Virtual Valuations. In *the 8th ACM Conference on Electronic Commerce (EC)*, 2007.

[CHMS10] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*, 2010.

[CM16] Shuchi Chawla and J. Benjamin Miller. Mechanism design for subadditive agents via an ex ante relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 579–596, 2016.

[CMS15] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015.

[CZ17] Yang Cai and Mingfei Zhao. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 170–183, 2017.

[DDT14] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. The Complexity of Optimal Mechanism Design. In *the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2014.

[DDT17] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Strong duality for a multiple-good monopolist. *Econometrica*, 85(3):735–767, 2017.

[DW12] Constantinos Daskalakis and S. Matthew Weinberg. Symmetries and Optimal Multi-Dimensional Mechanism Design. In *the 13th ACM Conference on Electronic Commerce (EC)*, 2012.

[FGKK16] Amos Fiat, Kira Goldner, Anna R. Karlin, and Elias Koutsoupias. The fedex problem. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 21–22, 2016.

[GO81] Roger Guesnerie and Claude Oddou. Second best taxation as a game. *Journal of Economic Theory*, 25(1):67–91, 1981.

[Gon18] Yannai A. Gonczarowski. Bounding the menu-size of approximately optimal auctions via optimal-transport duality. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 123–131, 2018.

[HDSM79] Peter Hammond, Partha Dasgupta, and Eric S Maskin. The implementation of social choice rules: Some general results on incentive compatibility. *Review of Economic Studies*, 46:185–216, 04 1979.

[HN13] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In *the 14th ACM Conference on Electronic Commerce (EC)*, 2013.

[HN17] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. *J. Economic Theory*, 172:313–347, 2017.

[Jer82] Mark Jerrum. A compact representation for permutation groups. In *23rd Annual Symposium on Foundations of Computer Science, Chicago, Illinois, USA, 3-5 November 1982*, pages 126–133, 1982.

[KW12] Robert Kleinberg and S. Matthew Weinberg. Matroid Prophet Inequalities. In *the 44th Annual ACM Symposium on Theory of Computing (STOC)*, 2012.

[LY13] Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28):11232–11237, 2013.

[MV07] A. M. Manelli and D. R. Vincent. Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly. *Journal of Economic Theory*, 137(1):153–185, 2007.

[Mye81] Roger B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[RC98] Jean-Charles Rochet and Philippe Chone. Ironing, sweeping, and multidimensional screening. *Econometrica*, 66(4):783–826, 1998.

[Rub16] Aviad Rubinstein. Settling the complexity of computing approximate two-player nash equillibria. In *FOCS*, 2016.

[RW15] Aviad Rubinstein and S. Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the 16th ACM Conference on Electronic Commerce*, 2015.

[SSW18] Raghuvansh R. Saxena, Ariel Schvartzman, and S. Matthew Weinberg. The menu complexity of "one-and-a-half-dimensional" mechanism design. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 2026–2035, 2018.

[Tha04] John Thanassoulis. Haggling over substitutes. *Journal of Economic Theory*, 117:217–245, 2004.

[Yao15] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2015.