

# Simplicity in Auctions Revisited: The Primitive Complexity

MOSHE BABAIOFF<sup>\*</sup>, Hebrew University of Jerusalem, Israel SHAHAR DOBZINSKI<sup>\*</sup>, Weizmann Institute of Science, Israel RON KUPFER<sup>\*</sup>, Harvard University, USA

In this paper we revisit the notion of simplicity in mechanisms. We consider a seller of m heterogeneous items, facing a single buyer with valuation v. We observe that previous attempts to define complexity measures often fail to classify mechanisms that are intuitively considered simple (e.g., the "selling separately" mechanism) as such. We suggest to view a menu as simple if a bundle that maximizes the buyer's profit can be found by conducting a few primitive operations that are considered simple. The *primitive complexity of a menu* is the number of primitive operations needed to (adaptively) find a profit-maximizing entry in the menu. In this paper, the primitive operation that we study is essentially computing the outcome of the "selling separately" mechanism.

Does the primitive complexity capture the simplicity of other auctions that are intuitively simple? We consider *bundle-size pricing*, a common pricing method in which the price of a bundle depends only on its size. Our main technical contribution is determining the primitive complexity of bundle-size pricing menus in various settings. First, we connect the notion of primitive complexity to the vast literature on query complexity. We then show that for any distribution  $\mathcal{D}$  over weighted matroid rank valuations, even distributions with arbitrary correlation among their values, there is always a bundle-size pricing menu. As part of this proof we provide a randomized algorithm that for any weighted matroid rank valuation v and integer k, finds the most valuable set of size k with only a poly-logarithmic number of demand and value queries. We show that this result is essentially tight in several aspects. For example, if the valuation v is submodular, then finding the most valuable set of size k requires exponentially many queries (this solves an open question of Badanidiyuru et al. [EC' 12]). We also show that any deterministic algorithm that finds the most valuable set of size k requires  $\Omega(\sqrt{m})$  demand and value queries, even for additive valuations.

# CCS Concepts: • Theory of computation $\rightarrow$ Algorithmic game theory.

#### **ACM Reference Format:**

Moshe Babaioff, Shahar Dobzinski, and Ron Kupfer. 2023. Simplicity in Auctions Revisited: The Primitive Complexity. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC '23), July 9–12, 2023, London, United Kingdom.* ACM, New York, NY, USA, 30 pages. https://doi.org/10.1145/3580507.3597695

# **1 INTRODUCTION**

The search for simple mechanisms is a central theme in the Mechanism Design literature (e.g., [Cai and Zhao 2017; Dütting et al. 2011; Hartline and Roughgarden 2009; Ronen 2001; Rubinstein and Weinberg 2018]). While complex mechanisms might be theoretically superior in terms of, e.g., extracting more revenue from the participants, they are often inferior in practice. For example,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

EC '23, July 9–12, 2023, London, United Kingdom © 2023 Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 979-8-4007-0104-7/23/07... https://doi.org/10.1145/3580507.3597695

<sup>\*</sup>Part of the work done while the authors were affiliated with Microsoft Research. The second author is supported by ISF grant 2185/19 and BSF-NSF grant (BSF number: 2021655, NSF number: 2127781).

complicated rules might be harder for the designer to implement and for the bidders to understand and to interact with, thus making them less attractive.

Of course, simplicity is a vague and elusive notion, and, unfortunately, there is little hope of finding a formal mathematical definition that sharply separates simple mechanisms from complex ones. Furthermore, a mechanism that is simple in one setting or for one group of participants might not be considered simple in other situations. To a large extent, simplicity is in the eye of the beholder. But to an even larger extent, the "I know it when I see it" test perfectly applies here.

Yet, a mathematical treatment of simplicity in auctions must be based on some formal definition, as imperfect as may be. Broadly speaking, many papers identify simplicity with particular forms of mechanisms (e.g., second price auctions [Moldovanu and Tietzel 1998; Thompson and Leyton-Brown 2013], ascending auctions [Cramton 1998; Gul and Stacchetti 2000; Mishra and Parkes 2007], posted prices auctions [Chawla et al. 2010; Hartline and Roughgarden 2009]) and analyze these classes of mechanisms.

In some settings, e.g., selling a single item, identifying simplicity with a specific auction format is an extremely useful idea. The canonical example is Myerson's optimal auction characterization that shows that second price auctions with reserve are optimal when values are drawn i.i.d. from a regular distribution; Other papers show that such auctions are approximately optimal in some other settings [Alaei 2014; Alaei et al. 2019; Chawla et al. 2010; Hartline and Roughgarden 2009]. Yet, the more complex the setting is, the less likely it is that a rigid list of permissible "simple" auction formats will provide optimal or approximately optimal results. Note also that a binary classification of auctions as either "simple" or not, does not allow the ranking and quantification of different auctions: which auction is simpler, a "selling separately" auction where each item has a different price, or the auction that sells any bundle of 10 items at price 1? And by how much?

This calls for using simplicity *measures* as an additional tool for analyzing auctions – focusing on quantitative approaches to simplicity. One of the most influential measures, the *menu complexity*, was suggested by Hart and Nisan [2019]. Consider a pricing problem in which a monopolist wants to sell a set M of m heterogeneous items. A deterministic *menu* M is a set of pairs  $(S, p_S)$ , in which every such pair specifies the price of the bundle  $S \subseteq M$ . The buyer has a valuation  $v : 2^M \to \mathbb{R}_{\geq 0}$  that specifies her value for every possible subset of the items. The valuation is drawn from some known distribution  $\mathcal{D}$ . Given deterministic menu  $\mathcal{M}$ , the buyer is assigned a bundle O that maximizes her profit, that is, bundle  $O \in \arg \max_{S \subseteq M} v(S) - p_S$ . The revenue of the mechanism is the expected payment of the buyer. The *menu complexity* of the mechanism is the minimal number of pairs (of a bundle and its price) needed to describe the mechanism.

In many cases the notion of menu complexity captures the simplicity of auctions very well. Menus with few entries tend to be "simpler" than menus with many entries, whatever the precise meaning of simplicity is. Indeed, in recent years we have seen the notion of menu complexity grows in popularity and being extended to more settings [Babaioff et al. 2021; Chawla et al. 2020; Dughmi et al. 2014; Gonczarowski 2018; Saxena et al. 2018] and to richer classes of valuations [Eden et al. 2021; Rubinstein and Weinberg 2018]. In general, the literature tends to draw the simplicity/complexity borderline by treating menu complexity poly(m) as a proxy for simplicity.

However, intuitive simplicity and menu complexity do not always go hand by hand. Consider an additive valuation over items and the mechanism that sells item separately, with item j sold at price  $p_j$ . The menu complexity of this mechanism is exponential<sup>1</sup>, as for each set S out of the  $2^m - 1$  non-empty sets, it needs to list it with price  $\sum_{j \in S} p_j$ . Yet, this menu is intuitively very simple. Moreover, a simple variant of this mechanism was shown to have very attractive properties. In fact,

<sup>&</sup>lt;sup>1</sup>The complexity of the "selling separately" menu is exponential also for the symmetric menu complexity [Kothari et al. 2019], which is a generalization of the menu complexity.

Babaioff et al. [2020] show that for additive valuations, when the value of every item j is drawn independently from a known distribution  $\mathcal{D}_j$ , then one of the following mechanisms extracts a constant fraction of the optimal revenue: sell the bundle of all items at the monopolist price (with respect to the distribution of the bundle of all items), or separately sell each item j at the monopolist price of the distribution  $\mathcal{D}_j$ . Most would agree that the Babaioff et al. mechanism is simple, but unfortunately its menu complexity is huge due to the "selling separately" component.

Another example of the limits of the notion of menu complexity can be found in the popular and practically-used "bundle-size" pricing (see, e.g., [Abdallah et al. 2021; Chu et al. 2011] and their followups) which prices all bundles of the same size at the same price. Although this menu is intuitively simple, its menu complexity is high. Indeed, observe that when the price of every bundle of size  $\frac{m}{2}$  is 1, the menu complexity of this mechanism is exponential, since each of the exponentially many bundles of size  $\frac{m}{2}$  requires an entry in the menu.

The Primitive Complexity of Auctions. We have exhibited several examples of mechanisms that pass the "I know it when I see it" test for simplicity, yet have high menu complexity. This calls for a more nuanced approach toward measuring simplicity. Before presenting our approach, we would like to stress again that an "ultimate" mathematical definition of simplicity is unlikely to exist. In all likeliness, inevitably, as any other simplicity notion, our new notion will fail for some mechanisms that "should" be considered simple and will include mechanisms that "should" be classified as complex. Yet, we believe that our approach would better capture the simplicity of many mechanisms. For other mechanisms, different approaches, possibly tailored to the specific application, might be useful.

The basic intuition that leads our work is that mechanisms are often considered simple if they can be implemented by applying only a small number of primitive operations that are considered "simple". In our case, the primitive operation is computing the outcome of the "selling separately" auction, perhaps the canonical example for a simple auction that is not captured as such by the notion of menu complexity. That is, given prices per item  $p_1, \ldots, p_m$ , return a bundle *S* that maximizes the buyer's profit (return  $S \in \arg \max_T v(T) - \sum_{j \in T} p_j$ ), as well as the value v(S) of the bundle *S*. The *primitive complexity of a menu* is the number of times the primitive operation has to be (adaptively) applied to find a bundle that maximizes the buyer's profit for any given valuation. The fewer times the primitive operation has to be invoked, the simpler the menu is.

The primitive complexity was defined here for deterministic algorithms using worst case approach on valuations, yet the definition naturally extends to randomized algorithms and to valuations sampled from a Bayesian prior.<sup>2</sup> Similarly, the definition can be extended by restricting the valuations to belong to a specific class (e.g., only additive or submodular valuations). Note that in principle, only the number of queries that the algorithm makes is restricted, not the running time, though all algorithms that we develop in this paper are computationally efficient. The definition of primitive complexity can also be naturally extended to randomized menus, i.e., menus that allow lotteries over bundles.

Our focus is in understanding whether the primitive complexity of the intuitively-simple class of bundle-size pricing menus is indeed low (mostly when the primitive operation is computing the outcome of the "selling separately" mechanism). Towards this end, we rely on (and advance) the literature on query complexity and valuation functions. In fact, since we will observe that the primitive complexity is essentially equivalent to a query complexity of finding a buyer's profit-maximizing bundle (as discussed below), our work also suggests that simple mechanisms are those for which a profit-maximizing bundle can be "easily" found.

 $<sup>^{2}</sup>$ A related notion is the *randomized* communication complexity of finding the profit-maximizing bundle in the menu. See [Rubinstein and Zhao 2021].

*Connection to Query Complexity.* We now discuss the connection of primitive complexity to query complexity. Recall that as usual in algorithmic game theory, the size of a naive description of the valuation v is exponential in the number of items. Thus it is common to assume that v is given as a black box that can only answer a limited number of types of queries. The two standard queries are *value queries* (given *S*, what is v(S)?) and *demand queries* (given item prices  $p_1, \ldots, p_m$ , return a bundle *S* that maximizes the profit of the buyer. That is, find  $S \in \arg \max_T v(T) - \sum_{j \in T} p_j$ ). It is not hard to see that the outcome of a "selling seperately" operation can be simulated by a demand query. Also note that a value query to a bundle *S* can be simulated by considering the outcome of one "selling seperately" operation returns the value of the demanded set). Thus, the primitive complexity and the query complexity are related up to a constant multiplicative factor.

Value queries are extensively used in various optimization problems [Calinescu et al. 2011; Nemhauser et al. 1978; Vondrák 2008]. Demand queries are standard in the algorithmic game theory literature and appear naturally in various posted prices auctions [Assadi and Singla 2019; Feldman et al. 2014], as the separation oracle needed to solve the natural LP relaxation for combinatorial auctions [Nisan and Segal 2006], and in various (not necessarily incentive compatible) approximation algorithms [Dobzinski et al. 2006; Feige and Vondrak 2006; Feige and Vondrak 2010].<sup>3</sup>

With this interpretation of primitive complexity in mind, the primitive complexity of a deterministic menu is at most its menu complexity: every menu with menu complexity c can be implemented by making c value queries to query v(S) for each bundle S which has an entry in the menu. The converse is far from being true: the primitive complexity of the mechanism that separately sells each item j at price  $p_j$  is just 1, while its menu complexity is exponential.

*Our Results.* In this paper we analyze the primitive complexity of the extensively studied class of bundle-size pricing menus for various classes of valuations. Recall that bundle-size pricing menu gives a price of  $p_r$  for every number of items  $r \in [m]$ . We start by considering the family of additive valuations ( $v(S) = \sum_{j \in S} v(\{j\})$ ) for every bundle S). We then move on to more complicated families of valuations. For additive valuations, the primitive complexity of every menu is at most m: querying the value  $v(\{j\})$  of every item j gives the entire valuation and thus suffices to compute a profit-maximizing bundle. Thus, in the context of additive valuations, simplicity will be captured by sub-linear primitive complexity, ideally achieving complexity that is poly-logarithmic in m, or even a constant. We prove that the primitive complexity of (approximately) maximizing the revenue is much better than linear in m, showing that it is only poly-logarithmic.

**Theorem I:** Let  $\mathcal{D}$  be some distribution over additive valuations. Then, for any  $\varepsilon > 0$  there is a bundle-size pricing menu  $\mathcal{M}$  with primitive complexity  $poly(\log m, \frac{1}{\varepsilon})$  such that the revenue of  $\mathcal{M}$  is in expectation at least  $(1 - \varepsilon)$  of the revenue of any other bundle-size pricing menu on  $\mathcal{D}$ .

Note that the distribution  $\mathcal{D}$  can be arbitrary. In particular, we do not assume that the values of the items are drawn from independent distirbutions<sup>4</sup>. To prove the theorem we show that for any bundle-size pricing menu  $\mathcal{M}'$ , there is a low complexity bundle-size pricing menu  $\mathcal{M}$  with revenue comparable to that of  $\mathcal{M}'$ . Specifically, we show that although  $\mathcal{M}$  might contain many distinct prices and bundles sizes, only  $poly(\log m, \frac{1}{\varepsilon})$  bundle sizes need to be considered (in expectation over the

<sup>&</sup>lt;sup>3</sup>Mathematically speaking, if *b* is the known number of bits used to represent numbers then a value query can be computed with *b* demand queries, whereas computing a demand query might require exp(m) value queries [Blumrosen and Nisan 2010]. To some extent, some would argue that in practice it is common to solve a demand query (what would you buy in the grocery store?) where as value queries are harder (what is your value for 12 eggs, bread, and a bottle of orange juice?)

<sup>&</sup>lt;sup>4</sup>For additive valuations, when item values are sampled independently, Babaioff et al. [2021] prove that the auction that separately sells each item can be well approximated by a menu with a polynomial size. In contrast, our result does not assume independence, it holds for any menu, and moreover, as we will see later, generalizes to much richer classes of valuations.

distribution  $\mathcal{D}$ ) to find a profit-maximizing bundle. Thus, the profit-maximization problem was reduced to the problem of finding a set of highest value for a given set size (maximization under a cardinality constraint), using "selling seperately" operations or, almost equivalently, value and demand queries: given a valuation v and bundle size k, find a bundle  $S_k$  that maximizes v(S) subject to  $|S_k| = k$  (in this case we say that  $S_k$  is a *k*-optimal set). Let  $p_k$  be the price of bundles of size k in  $\mathcal{M}$ . A profit maximizing bundle of  $\mathcal{M}$  is in arg max{ $v(S_k) - p_k$ } (or the empty set, if this maximum profit is negative). Thus, to complete the proof we prove the following algorithmic result:

**Theorem II:** Fix some bundle size k. There exists a randomized algorithm that given an additive valuation v finds a bundle  $S_k \in \arg \max_{S:|S|=k} v(S)$  by making, in expectation,  $poly \log(m)$  value and demand queries (the expectation is over the randomness of the algorithm).

Note that our algorithm for finding a k-optimal set does not assume that the valuations are drawn from some distribution: the guarantee is in the worst case, for every possible additive valuation. Furthermore, our algorithm *always* finds a value-maximizing bundle of size k, randomization is only use to accelerate the running time. Moreover, a significant challenge in developing our algorithms is that we want the algorithms to work with every implementation of the demand oracle. That is, the demand query is required to return a profit-maximizing bundle, but if there are several such bundles we want our algorithms to work with any implementation of the tie-breaking rule, even an adversarial one.

We show that our algorithm for additive valuations is qualitatively optimal in multiple respects:

- Deterministic algorithms with value queries: We show that any deterministic algorithm that uses only value queries must make at least m 1 queries in order to find the item with the highest value (k = 1).
- Randomized algorithms with value queries: We provide two different proofs that show with no demand queries,  $\Omega(\frac{m}{\log m})$  value queries are needed to find the *k*-optimal bundle, even if the algorithm is randomized. One proof assumes  $k = \frac{m}{2}$  and is based on a simple counting argument. The second proof is based on a slightly more involved communication complexity argument and shows this impossibility even for the simple case of k = 1 (i.e., finding the most valuable item).
- Randomization is required: We prove that even if demand queries are allowed but the algorithm must be deterministic, then  $\Omega(\sqrt{m})$  queries are required.

The proofs of the first two impossibilities are easier than the proof of the last result, which is more subtle and involved. We then move on to consider richer valuation classes, starting with weighted matroid rank functions:

**Theorem III:** Fix a bundle size k. There is a randomized algorithm that finds a bundle  $S_k \in \arg \max_{S:|S|=k} v(S)$  for every weighted matroid rank valuation v and makes in expectation  $poly \log(m)$  value and demand queries.

Using Theorem III we are able to extend Theorem I to hold for weighted matroid rank valuations, not just additive ones. For weighted matroid rank valuations, the greedy algorithm finds a k-optimal set with  $\Theta(m \cdot k)$  value queries. We thus see that with demand queries it is possible to find a k-optimal set exponentially faster. Furthermore, even if all weights of the items are either 0 and 1, the number of matroid rank functions is doubly exponential [Knuth 1974]. Nevertheless, we find a k-optimal set in only  $poly \log(m)$  queries.

This paper is not the first to consider maximization subject to cardinality constraint with value and demand queries. The first was [Badanidiyuru et al. 2012], and it considered richer classes: submodular, XOS, and subadditive valuations. However, while the current paper considers exact

optimization algorithms, the focus of [Badanidiyuru et al. 2012] was in approximation algorithms. Approximation algorithms are not useful for finding a profit maximizing bundle in a menu, which must be done exactly, otherwise the incentive constraints are likely to be violated. One of the main results of [Badanidiyuru et al. 2012] is a  $\frac{9}{8}$ -approximation algorithm for maximizing a submodular function subject to a cardinality constraint using value and demand queries (recall that with value queries only, the greedy algorithm provides an approximation ratio of  $\frac{e}{e-1}$  [Nemhauser et al. 1978]). However, the paper [Badanidiyuru et al. 2012] proves no impossibility at all for this setting, asking whether an *exact* solution can be found with polynomially many value and demand queries. We solve this open question:

**Theorem IV:** Fix a randomized algorithm A for maximizing a submodular function subject to cardinality constraint that succeeds with constant probability. Then, A makes at least exp(m) value and demand queries.

This result highlights the importance of the valuations class in measuring the primitive complexity: the primitive complexity of the menu that corresponds to maximization subject to a cardinality constraint is poly logarithmic for the class of additive valuations, but exponential for the richer class of submodular valuations.

Connections to other Problems. The problem of maximization subject to cardinality constraint with demand queries has interesting connections to some well-studied problems. One such problem is unordered partial sorting [Chambers 1971]: we are given an array that contains n numbers and the goal is to find a set with k highest numbers (in any order). Note that unordered partial sorting is equivalent to finding a k-optimal set in an additive valuation.

Unordered partial sorting can be solved in O(m) time by QuickSelect [Hoare 1961], a variant of QuickSort. Note that the pivot procedure of QuickSort and QuickSelect – divide the array into two, one includes all numbers bigger than some p, and the other contains all number smaller than p – is essentially a demand query at price p per item (with some tie-breaking rule that depends on the implementation). Indeed, QuickSelect makes in expectation  $O(\log m)$  pivot calls. Our lower bound for deterministic algorithms implies that randomization is essential to QuickSelect and other pivot-based algorithms in the sense that any deterministic algorithm requires  $\Omega(\sqrt{m})$  pivot and value queries. Note that algorithms like IntroSelect [Musser 1997] use more advanced methods to deterministically ensure a good selection of the pivot for QuickSelect. However, these algorithms use many value queries, and our results show that this is unavoidable.

There are also connections to various coin weighting problems. For example, consider the following problem studied in [Bshouty 2009; Djackov 1975; Du et al. 2000; Lindstrom 1975]: we are given n coins. We know that d of them are counterfeit. The weight of each real coin is  $w_1$  and the weight of each counterfeit coin is  $w_2$ . We are also given a spring weight. How many weightings are needed to find all counterfeit coins? Note that every use of the spring weight is equivalent to a value query. If d is big then a simple counting argument – similar to the one that prove that even randomized algorithms must make many value queries – shows that many weightings are needed. However, our communication-complexity based proof shows that if there are three types of coins with  $w_3 > w_2 > w_1$  then finding just one coin with weight  $w_3$  requires almost linear number of weightings, a result that was not known before, to the best of our knowledge.

*Future Directions*. In this paper we introduced a new measure of complexity for auctions: the primitive complexity. We have examined this notion in the context of a pricing problem, and obtained some algorithms with low primitive complexity as well as some impossibilities. Obviously, studying the primitive complexity of other problems is an exciting future direction.

We propose a number of open questions. We have shown that bundle-size pricing menus have poly-logarithmic primitive complexity if the valuation belongs to the class of weighted matroid rank functions, and that if the valuation belongs to the class of submodular valuations the primitive complexity is exponential. The class of gross substitutes valuations contains all weighted matroid rank functions and is contained in the class of submodular valuations. We know that the greedy algorithm finds a *k*-optimal bundle with polynomially many value queries, but can we find a *k*-optimal bundle with polynomially mand queries?

In addition, it will be extremely interesting to understand whether bundle-size pricing can well approximate the revenue that can be obtained by any deterministic mechanism. Of course, for this question to make sense we have to consider some kind of symmetry in the distribution. For additive valuations, if the values of the items are sampled i.i.d. then the mechanisms of Li and Yao [2013] and Babaioff et al. [2020] already imply that bundle-size pricing can provide a constant fraction of the optimal revenue. But what if the joint distribution of item values is symmetric, yet item values are not sampled i.i.d.? Can bundle-size pricing provide a constant fraction of the optimal revenue that can be achieved by a deterministic mechanism then? See also [Babaioff et al. 2018] for some related work.

Finally, a fascinating direction is to obtain mechanisms with polylogarithmic primitive complexity that obtain  $(1 - \epsilon)$  fraction of the optimal revenue. We do not know how to obtain such a mechanism even for independent distributions and additive valuations.

# 2 MODEL AND PRELIMINARIES

*Valuations.* Given a set  $M = [m] = \{1, 2, ..., m\}$  of *m* indivisible *items*, a *valuation function*  $v : 2^M \to \mathbb{R}_+$  determines a non-negative *value* v(S) for each bundle  $S \subseteq M$ . We make the standard assumptions that any valuation function *v* is normalized  $(v(\emptyset) = 0)$  and monotone (for  $S \subseteq T$  it holds that  $v(S) \leq v(T)$ ). With a slight abuse of notation, for valuation *v* and an item  $a \in M$  we use v(a) to denote  $v(\{a\})$ . We consider several standard classes of valuations (each of the classes is strictly contained in the class that follows it):

- A valuation v is *additive* if for all  $S \subseteq M$  we have that  $v(S) = \sum_{a \in S} v(\{a\})$ . In this case we may represent the function as a vector  $v = (v_1, ..., v_m) \in \mathbb{R}^m_+$ , where  $v_i \ge 0$  is the value of the *i*'th item.
- A valuation function v is called *weighted matroid-rank valuation* if there exists a matroid<sup>5</sup> over the set of elements M and a weight function<sup>6</sup>  $\omega : M \to \mathbb{R}_+$  such that  $v(S) = \omega(S)$  for every  $S \subseteq M$ .
- A valuation v is called *submodular* if it exhibits the diminishing returns property, i.e.,  $v(S \cup \{a\}) v(S) \ge v(T \cup \{a\}) v(T)$  for all  $S \subseteq T \subseteq M$  and  $a \in M$ .

*Maximization Subject to a Cardinality Constraint.* Given a valuation v over set M, the *cardinality maximization problem* with parameter k is the problem of finding a maximum value set of size k. I.e., finding a set S that is k-optimal:  $S \in \arg \max_{S:|S|=k} v(S)$ .

<sup>&</sup>lt;sup>5</sup>A *matroid* is a pair (M, I), with M being a finite set of elements and  $I \subseteq 2^M$  is a non-empty family of subsets of M such that: (1) If  $B \subseteq A \subseteq M$  and  $A \in I$ , then  $B \in I$  and (2) If  $A, B \in I$  and |B| > |A|, then there exists an item  $b \in B \setminus A$  such that  $A \cup \{b\} \in I$ . A set  $A \in I$  is called an *independent set*. An independent set  $A \in I$  is called a *base* of the matroid, if it is not contained in any larger independent set. The *rank* of the matroid is the size of any base (all have the same size).

<sup>&</sup>lt;sup>6</sup>A weight function  $\omega : M \to \mathbb{R}_+$  assigns a weight to each element. The weight function is extended to sets as follows: the weight of a set  $S \subseteq M$  is defined to be  $\omega(S) = \max_{A \subseteq S, A \in I} \sum_{a \in A} \omega(a)$ . An independent set  $A \in I$  is called a *maximal weight independent set* if there is no independent set of larger weight.

*Mechanisms.* We consider a setting with one seller holding a set M of m items, that faces a single buyer with a valuation v. A *deterministic menu* M is a set of pairs  $\{S, p_S\}$  of bundles and prices.<sup>7</sup> We assume that any menu includes the option of getting no item and paying 0. Given a menu M, a set S is a *demanded set* (or a *most profitable set*) of a buyer with valuation v if  $S \in \arg \max_{S' \subseteq M} v(S') - p_{S'}$ . The family of sets that are demanded are called the *demand* of the buyer. We assume that a buyer with valuation v that faces M selects a set that she demands, but make no assumption about how she picks between different demanded sets. In a Bayesian setting, the valuation v is drawn from a known distribution  $\mathcal{F}$ , and the revenue of the menu M is measured in expectation over  $\mathcal{F}$ :

- REV<sub>M</sub>(v): the expected revenue of the seller from menu  $\mathcal{M}$  when the buyer's valuation is v, i.e., if the buyer picks a demanded set S with probability  $g_S$  then the revenue is  $\sum_{S \subseteq \mathcal{M}} p_S \cdot g_S$ .
- REV(M, F): the expected revenue where the expectation is over the buyer's valuation v ~ F, that is REV(M, F) = 𝔅<sub>v∼F</sub> [REV<sub>M</sub>(v)].

A specific class of mechanisms that is considered in this paper is *bundle-size pricing*: the price of each bundle of size r is  $p_r$ . Since the valuations are monotone, we assume that for any two bundle sizes  $q_i > q_j$  we have that  $p_i > p_j$  (all inequalities are strict).

*Queries.* In this paper we consider "selling seperately" operations: given a price  $p_i$  for each item *i*, find some bundle in the demand, a bundle in  $\arg \max_{S \subseteq M} (v(S) - \sum_{i \in S} p_i))$ , and return this bundle and its value.

The literature on multi-item auctions has extensively studied two types of queries as means of accessing a valuation v (which might have a large representation): value queries and demand queries. A *value query* is given a set S and simply returns v(S), the value of the bundle S. A *demand query* asks for a bundle of maximum profit at some given item prices, i.e., a most profitable set for the given prices (a demanded set). Formally, the query is given an item-price vector  $p = (p_1, \dots, p_m) \in \mathbb{R}^m_+$  and returns an arbitrary set  $D \in \arg \max_{S \subseteq M} (v(S) - p(S))$  in the demand, where  $p(S) = \sum_{i \in S} p_i$ . We assume that the demand query also returns the value of the set S (this can always be done at the cost of an additional value query). We make no assumption about the way ties are broken between sets in the demand. Ties might be broken adversarially and this issue creates significant challenges which we need to address. When all the coordinates of p have the same value  $t \in \mathbb{R}$ , we refer to the corresponding query as a *uniform demand query* for price t.

Obviously, a "selling seperately" operation can simulate a demand query. It can also simulate a value query for a bundle S: set the price of each item in S to 0, and the price of every other item to  $\infty$ . It is also straightforward to see that any "selling seperately" operation can be simulated by one demand query followed by one value query. Hence, any algorithm that uses only t "selling seperately" operations can be simulated with 2t value and demand queries, and every algorithm that uses t value and demand queries can be implemented with t "selling seperately" operations. Thus, we will freely switch between these two similar points of view.

The Primitive Complexity. In this paper we suggest to consider the primitive complexity of menus. Fix a class of valuations  $\mathcal{V}$ . The primitive complexity of an algorithm with respect to  $\mathcal{V}$  is the maximal number of "selling seperately" operations that it (adaptively) makes on any  $v \in \mathcal{V}$ . The primitive complexity of a menu is the minimal primitive complexity of any algorithm that for any valuation  $v \in \mathcal{V}$  computes a most profitable bundle in this menu.<sup>8</sup> If the algorithm is randomized, the randomized primitive complexity is the *expected* number of queries that the algorithm makes, where

<sup>&</sup>lt;sup>7</sup>We assume that valuations are monotone non-decreasing (free disposal). Under this assumption, it is wlog to assume that for any  $S, T \subseteq M$  such that  $S \subseteq T$  we have that  $p_S \leq p_T$  (otherwise S is never being sold and the menu entry can be removed). <sup>8</sup>As stated, this definition considers deterministic menus (as the menu entries are bundles). This definition naturally extends to randomized menus in which an entry might be a lottery over bundles.

expectation is over the internal random coins of the algorithm. Similarly, if the valuations are drawn from some distribution, the distributional primitive complexity is the expected number of queries that the algorithm makes, where expectation is over valuations sampled from the prior distribution. We note that our randomized algorithms will always return a correct solution, not just with high probability. Randomization is only used to reduce the (expected) number of queries. In contrast, all of our lower bounds hold even for algorithms that only succeed with constant probability.

# **3 ALGORITHMS FOR MAXIMIZATION SUBJECT TO A CARDINALITY CONSTRAINT**

In this section we first present an algorithm that finds a k-optimal set for any additive valuation, and then present an algorithm that does the same for any weighted matroid rank valuation. Our algorithms are randomized and make only  $poly \log(m)$  queries in expectation. Note that this is an exponential improvement in the number of queries over the "obvious" algorithms: the trivial algorithm for additive valuations finds a k-optimal set with m queries (one for each item). For weighted matroid rank valuations, a k-optimal set can be found by running the greedy algorithm (poly(m) value queries).

We start with the case of additive valuations. We then solve the more general case of weighted matroid-rank valuations by first finding a maximal weight independent set R using  $poly \log(m)$  queries, and then applying the algorithm for additive valuations as a sub-procedure.

THEOREM 3.1. There exists a randomized algorithm that for any additive valuation v over m items, finds a k-optimal set using value and demand queries, and in expectation makes  $O(\log^3 m)$  queries.

We further extend the result to weighted matroid-rank valuations.

THEOREM 3.2. There exists a randomized algorithm that for any weighted matroid-rank valuation v over m items, finds a k-optimal set using value and demand queries, and in expectation makes  $O(\log^3 m)$  queries.

Proofs of these theorems can be found in Appendix A.1 and Appendix A.2. We now provide some intuition for the case of additive valuations. We start with presenting an algorithm for the special case in which all items have distinct values. That is, for each two items  $i \neq j$  we have that  $v_i \neq v_j$ . In this case, we can easily find the set of k highest value items with  $O(\log m)$  queries: select an item uniformly at random and denote the value of this item by v. Make a uniform demand query with a price of v per item. Let D be the returned demanded set and denote m' = |D|. If  $k \ge m'$  we know that all the m' items of D are among the k highest values, so we pick them and remove them from the set, update the number of items we still need to pick to k' = k - m', and continue recursively on the remaining items that are not in D to select an additional k = k' items. If k < m' we similarly remove the items not in D, and continue recursively, aiming to pick k out of the items in D. The expected number of iterations is  $O(\log m)$  since in each iteration, in expectation, half of the items are removed from consideration (either picked or discarded).

The problem is more challenging if the values are not distinct. That is, if there are items of equal value  $(v_i = v_j \text{ for items } i, j \neq i)$  and we make a demand query with price  $p = v_i = v_j$  per item. The challenge is that, unless assuming a specific tie-breaking rule (which we do not), the demanded set might contain all items of value p, none of them, or some arbitrary subset of them. We present an algorithm that works for any implementation of the demand query, and even if the tie breaking is adversarial.

### 4 THE PRIMITIVE COMPLEXITY OF BUNDLE-SIZE PRICING

Recall that a *bundle-size pricing menu* is a set of offers  $\{(q_i, p_i)\}_i$ , each of the form "pay  $p_i$  and choose any set of size  $q_i$  of items to receive". We call the number of different quantities that are offered the *size* of the bundle-size pricing menu. Given a bundle-size pricing menu, a buyer that wants to find a profit maximizing set can do so by finding a  $q_i$ -optimal set for each of quantity  $q_i$  specified by the menu, and picking the one of highest profit among the candidates sets. As with *m* items there can be *m* different quantities specified, this approach will result in polynomial (in *m*) number of *k*-optimal set problems that need to be solved, and thus require at least polynomial number of demand queries. We aim for sub-linear number of queries.

We first observe, using a variant of a result of Hart and Nisan [2019], that any bundle-size pricing menu can be transformed to another bundle-size pricing, losing only  $\epsilon$ -fraction of the revenue, but with the size of the new menu only depending on the revenue loss parameter  $\epsilon$  and on the largest ratio of prices in the original bundle-size pricing menu, but not on *m*. Specifically, the new bundle-size pricing menu size will only be polynomial in  $1/\epsilon$  and in the logarithm of the maximal ratio of prices in the menu, but independent of the number of items *m*. For valuation classes for which the *k*-optimal set problem is solvable in poly-logarithmic number of demand queries (as additive and weighted matroid-rank valuations), we can thus derive a bound on the number of demand queries needed to pick a profit-maximizing bundle from this smaller bundle-size pricing menu. Yet, the number of demand queries in above result depends on the maximum price-ratio not being too large, and will be polynomial in *m* when this ratio is exponential.

Our main result in this section is that we can get rid of this limitation when optimizing the *expected* revenue for a given distribution  $\mathcal{D}$  over valuations (rather than ex-post, for any given valuation). We consider the expected revenue of the menu for the given distribution  $\mathcal{D}$  and further shrink the menu size by removing some of the priced bundles sizes. Specifically, we remove each bundle size that contributes at most  $\frac{\epsilon}{m}$  fraction of the expected revenue (over  $\mathcal{D}$ ), losing another  $\epsilon$  fraction of the total revenue. Yet, even after this additional processing step that reduces the size of the menu, the menu size might still be large (not poly-logarithmic). Thus, it might well still be the case that the size of this bundle-size menu is not small enough to get a poly-logarithmic number of queries by always finding a profit-maximizing bundle for each size and taking the best one. Nevertheless, we show that either the same revenue can be obtained by replacing the menu with a menu that only sells the grand bundle, or that for the same menu a poly-logarithmic number of queries is sufficient *in expectation*. That is, although the menu might contain many distinct prices and bundles sizes, only  $poly(\log m, \frac{1}{\epsilon})$  bundle sizes need to be considered (in expectation over the distribution  $\mathcal{D}$ ) to find a profit-maximizing bundle for a valuation v sampled from  $\mathcal{D}$ . The proof is in Appendix B.

THEOREM 4.1. Given a distribution  $\mathcal{D}$  over weighted matroid-rank valuations over a set M of m items, a bundle-size pricing menu  $\mathcal{M}_1$ , and  $\epsilon > 0$ , there exists a bundle-size pricing menu  $\mathcal{M}_2$  such that  $\text{Rev}(\mathcal{M}_2, \mathcal{D}) \ge (1 - \epsilon) \text{Rev}(\mathcal{M}_1, \mathcal{D})$  and such that a profit-maximizing bundle for  $\mathcal{M}_2$  can be found in  $\mathcal{O}(\epsilon^{-3}\log^4 m)$  value and demand queries in expectation, where the expectation is over the internal random coins of the algorithm and the distribution  $\mathcal{D}$ .

#### 5 HARDNESS OF MAXIMIZATION SUBJECT TO A CARDINALITY CONSTRAINT

We have presented a randomized algorithm that for weighted matroid-rank valuations finds a *k*-optimal set using poly-logarithmic number of demand and value queries (Theorem 3.2). In this section we present lower bounds for several related problems, showing that our results can not be strengthened in multiple ways. Namely, we show that:

- For **submodular** valuations, any randomized algorithm that for every valuation succeeds with constant probability to find a *k*-optimal set, must use, in expectation, an exponential number of value and demand queries (Section 5.1).
- Any **deterministic** algorithm that finds a 1-optimal set (an item with the highest value) requires  $\Omega(\sqrt{m})$  value and demand queries, even when the valuation is additive (Section 5.2).
- For additive valuations, we show, using two proof techniques, that any algorithm that given k finds a k-optimal set and succeeds with constant probability requires  $\Omega(m/\log m)$  value queries, even if randomization is allowed. Furthermore, for deterministic algorithms we show that m 1 value queries are needed. These results are presented in Section 5.3.

# 5.1 An Exponential Lower Bound for Submodular Valuations

We next consider submodular valuations and show that any randomized algorithm that with value and demand queries finds (with a constant probability) a *k*-optimal set makes in expectation exponentially many queries:

THEOREM 5.1. Let A be a randomized algorithm that given a submodular valuation and k finds a k-optimal set by making value and demand queries. For m that is large enough, if A succeeds with probability at least  $\frac{1}{2}$  then A makes at least  $1.3^m$  queries.

This solves an open question of [Badanidiyuru et al. 2012] that provided a  $\frac{9}{8}$ -approximation for this problem, but did not even rule out the possibility that a *k*-optimal set can be found with a polynomial number of value and demand queries.

The rest of this subsection is devoted to outlining the proof of the theorem. To prove a bound for randomized algorithms, it is enough to provide a distribution  $\mathcal{D}$  over valuations such that the probability that a deterministic algorithm that makes subexponentially many value and demand queries finds a k-optimal set in a valuation that is sampled from  $\mathcal{D}$  is small, by Yao's principle.

We will prove the theorem for an even m and  $k = \frac{m}{2}$ . Let  $\mathcal{D}$  be the following distribution over submodular valuations over a set of items M: each valuation v is defined by a family  $\mathcal{B}_v$  of sets of size k + 1 and a set  $G_v$  of size k ( $G_v$  will be the k-optimal bundle). Each set of size k + 1 is included in  $\mathcal{B}_v$  with probability  $\frac{1}{m^2}$ , independently at random. Out of the sets of size k that are not contained in any of the sets in  $\mathcal{B}_v$ , we choose one random set and denote it by  $G_v$ .<sup>9</sup> It will also be convenient to define  $\mathcal{R}_v$  to be the family of all sets of size k - 1 that are not contained in any of the sets in  $\mathcal{B}_v$ . The valuation v is then defined as follows:

$$v(S) = \begin{cases} k & |S| > k+1 \\ k & |S| = k+1, \ S \in \mathcal{B}_v \\ k-3/11 & |S| = k+1, \ S \notin \mathcal{B}_v \\ k-6/11 & |S| = k, \ S = G_v \\ k-7/11 & |S| = k, \ S \neq G_v \\ k-1 & |S| = k-1, \ S \in \mathcal{R}_v \\ k-14/11 & |S| = k-1, \ S \notin \mathcal{R}_v \\ |S| & |S| < k-1 \end{cases}$$

<sup>&</sup>lt;sup>9</sup>There is an exponentially small probability that every set of size k is contained in some set of  $\mathcal{B}_v$ . In this case  $G_v$  is not defined and all sets of size k have the same value. We thus condition our analysis on having that this event does not happen and that  $G_v$  is defined.



Fig. 1. The relation between bundles in valuations in the support of  $\mathcal{D}$ . Blue squares denote bundles that are in  $\mathcal{B}_v$ , the green square denote the bundle  $G_v$ , and red square denote the bundles in  $\mathcal{R}_v$ . A line between squares denotes a possible containment relationship between the bundles. Note that all subsets of size k - 1 of a set in  $\mathcal{B}_v$  have value k - 14/11.

The set  $G_v$  is the *k*-optimal set. Roughly speaking, sets from the families  $\mathcal{B}_v$  and  $\mathcal{R}_v$  guarantee that it is very unlikely that any information about the identity of  $G_v$  will be provided by any demand or value query. See Figure 1 for an illustration of the relations between the sets.

We prove our exponential lower bound in two steps. We onsider any algorithm that uses only value queries and on valuation that is samples from  $\mathcal{D}$  finds a *k*-optimal set with non-negligible probability. We show that any such algorithm makes, in expectation, an exponential number of value queries (Lemma 5.2). We complete the proof by showing that, with high probability over  $\mathcal{D}$ , all demand queries on a valuation *v* sampled from  $\mathcal{D}$  can be simulated by value queries with only a polynomial blowup in the number of queries (Lemma 5.3).

LEMMA 5.2. Fix some deterministic algorithm A that makes only value queries and the set of those queries is in a canonical form<sup>10</sup>. Suppose that A makes  $t < 1.9^m$  value queries on valuations that are sampled from  $\mathcal{D}$ . Then, for a large enough m, the probability (over  $\mathcal{D}$ ) that A finds a k-optimal set is at most  $\frac{t}{1.9^m}$ .

LEMMA 5.3. Fix a deterministic algorithm A that uses t demand and value queries for valuations sampled from  $\mathcal{D}$ . For any  $\alpha > 1$ , with probability  $1 - \frac{1}{\alpha}$ , A can be implemented using at most  $2m^5 \cdot t^2 \cdot \alpha$  value queries.

The proof of Lemma 5.2 is presented in Appendix C.1.2. Before proving Lemma 5.3 in Section 5.1.2, we present some definitions and auxiliary claims.

5.1.1 Definitions and Auxiliary Claims. We first show that every valuation v in the support of  $\mathcal{D}$  is indeed submodular. We then present some definitions and prove several claims that will be helpful in the proof of the theorem. All proofs in this section are deferred to Appendix C.1.1.

LEMMA 5.4. Every valuation v in the support of D is submodular.

<sup>&</sup>lt;sup>10</sup>Later (Appendix C.1), we formally define what it means for a set of value queries to be in a canonical form. We comment now that any set value queries of size t can be converted to a canonical form by making poly(t,m) additional value queries.

A valuation in the support of  $\mathcal{D}$  is completely defined by the values of all sets of size k and k + 1. We say that a set of value queries Q is in a *canonical form* if all queries in Q are for sets of size k or k + 1, and for every query of size k all of its supersets of size k + 1 are also in Q. Essentially, all information that a set of value queries conveys about a valuation can also be conveyed by some set of queries that is in a canonical form and is not much larger. The next proposition shows that we can assume that the query set is in a canonical form at a cost of a polynomial blow-up in the number of queries:

PROPOSITION 5.5. Let A' be an algorithm that makes t value queries on a valuation in the support of  $\mathcal{D}$ . Then, there is an algorithm A that simulates A' while making  $m^2 \cdot t$  value queries on a valuation in the support of  $\mathcal{D}$ . Moreover, the set of queries that A makes has a canonical form.

We next present several useful definitions and notations. Fix some deterministic algorithm A that makes only value queries and runs on valuations from  $\mathcal{D}$ . Fix any valuation v from the support of  $\mathcal{D}$ , and let  $Q_v$  denote the list of t bundles that A queried together with their values. Let  $\mathcal{D}_{Q_v}$  denote the distribution over valuations that is obtained by sampling according to  $\mathcal{D}$  a valuation that is consistent with the queries in  $Q_v$ . Let  $\mathcal{B}_{Q_v}^y$  be the family of sets that includes every set S such that  $\Pr_{v'\sim \mathcal{D}_{Q_v}}[S \in \mathcal{B}_{v'}] = 1$ . Similarly, let  $\mathcal{B}_{Q_v}^n$  be the family of sets that includes every set S such that  $\Pr_{v'\sim \mathcal{D}_{Q_v}}[S \in \mathcal{B}_{v'}] = 0$ . Let  $\mathcal{K}_{Q_v}$  be the family of sets of size k that were queried in  $Q_v$ .

We now claim that assuming queries are in a canonical form, the conditional distribution for sets not queried is essentially identical to the prior.

LEMMA 5.6. Fix any valuation v sampled from  $\mathcal{D}$  and assume  $Q_v$  is in a canonical form. It holds that:

- For any set S of size k + 1 it holds that  $\Pr_{v' \sim \mathcal{D}_{Q_v}}[S \in \mathcal{B}_{v'}] \in \{0, 1, \frac{1}{m^2}\}$ .
- The conditional probabilities are independent: for any family  $\mathcal{F}$  of sets of size k + 1 it holds that  $\Pr_{v' \sim \mathcal{D}_{Q_v}}[\forall S \in \mathcal{F}, S \notin \mathcal{B}_{v'}] = \prod_{S \in \mathcal{F}} \Pr_{v' \sim \mathcal{D}_{Q_v}}[S \notin \mathcal{B}_{v'}].$

5.1.2 Proof of Lemma 5.3: Simulating Demand Queries by Value Queries. In this section we prove Lemma 5.3, showing that for valuations drawn from  $\mathcal{D}$ , demand queries can be simulated by value queries, and with high probability polynomial number of queries is sufficient for the simulation.

PROOF OF LEMMA 5.3. First, we prove by induction on t that a set of value queries in a canonical form that is followed by a demand query, can be simulated by a set of value queries that has a canonical form, and the expected size of that set is at most  $t + 2 \cdot m^5$ . The claim trivially holds for t = 0 and Claim 5.7 proves the induction step. Second, given the claim, we use Markov's inequality, to argue that for any  $\alpha > 0$ , with probability at most  $\frac{1}{t \cdot \alpha}$ , more than  $2m^5 \cdot t \cdot \alpha$  value queries are needed for the implementation of a query. Hence, using the union bound, with probability  $1 - \frac{t}{t \cdot \alpha} = 1 - \frac{1}{\alpha}$  all t demand queries can be implemented using at most  $2m^5 \cdot t^2 \cdot \alpha$  value queries.

CLAIM 5.7. Fix a deterministic algorithm A that runs on valuations sampled from  $\mathcal{D}$  and up to some point has used a set of t value queries, and the set has a canonical form. Fix any demand query for price vector p. Then, it is possible to simulate all these queries (including the demand query) by a set of value queries that has a canonical form, and the expected size of that set is at most  $t + 2 \cdot m^5$ .

PROOF. For  $m \le 3$ , the total number of subsets of  $2^m$  is smaller than  $2 \cdot m^5$  and the claim trivially holds. We next assume that m is even and  $m \ge 4$ .

We first present some intuition for the proof. Consider a demand query with price vector p for valuation v. For a set S, denote by U(S, p) = v(S) - p(S) the profit from buying set S at price  $p(S) = \sum_{i \in S} p_i$ . A demand query returns a set S that has maximal profit under price vector p. Clearly, if we can find a most profitable bundle of size d for every  $d \in [m]$  then we can return a most profitable

set (with a set from these *m* bundles that is most profitable). While finding a most profitable bundle of every size is clearly sufficient, it turns out it is not necessary, and we show that a most profitable set can be found with polynomially many value queries, without always knowing a most profitable bundle of size *k*. We first show that for each  $d \neq k$ , a most profitable bundle of size *d* can indeed be found by value queries to a family  $\mathcal{F}_d^p$  of sets of size *d* that we can specify. Second, we show that if every set that is a most profitable set overall is of size *k*, then such a set can also be found by value queries to a family  $\mathcal{F}_k^p$  of sets of size *k* that we can specify. Finally, we show that the set of all queries (the *t* value queries as well as value queries to new sets that are in these families of sets) is only polynomially larger than *t*. We next present the formal claim and its proof.

For each  $d \in [m]$  we define a family  $\mathcal{F}_d^p$  of sets of size d, such that:

- If  $d \neq k$  then some set of size d that has the highest profit among all sets of size d in the family  $\mathcal{F}_d^p$ .
- If every demanded set is of size d = k, then a demanded set of size k belongs to the family  $\mathcal{F}_k^p$ .
- For even  $m \ge 4$  it holds that  $\mathbb{E}_{v \sim D} \left[ \sum_{d \in [m]} |\mathcal{F}_d^p \setminus Q_v| \right] \le 2m^3$ .

Assume algorithm A is running on valuation v sampled from  $\mathcal{D}$ , and the algorithm was using the set  $Q_v$  of value queries that is a canonical form. For each size d we consider the list of size d from cheapest to most expensive (breaking ties arbitrarily). Let  $S_d^p$  denote a cheapest set of size d.

- For d < k 1 or d > k + 1, all bundles of size d have the same value, thus a profit maximizing set of size d is simply some cheapest set of size d, so we define \$\mathcal{F}\_d^p = {S\_d^p}\$.
  For d = k + 1, there are two possible values for a bundle of size k + 1, depending on whether the
- For d = k + 1, there are two possible values for a bundle of size k + 1, depending on whether the bundle is in B<sub>v</sub> or not. Let S<sub>B<sub>v</sub></sub> be a cheapest bundle in B<sub>v</sub> that is the first in order of set prices. A most profitable set of size d = k + 1 is then either S<sub>B<sub>v</sub></sub> or S<sup>p</sup><sub>k+1</sub>. We add to F<sup>p</sup><sub>k+1</sub> the cheapest sets of size d = k + 1 in increasing order of price, till we find the first set that belongs to B<sub>v</sub>. Note that S<sup>p</sup><sub>k+1</sub> is the first added set and is always in F<sup>p</sup><sub>k+1</sub>. Since Q<sub>v</sub> is in a canonical form, by Lemma 5.6, it holds that either S ∈ Q<sub>v</sub> or that the probability that S ∈ B<sub>v'</sub> conditional on v' being sampled according to D<sub>Q<sub>v</sub></sub> is <sup>1</sup>/<sub>m<sup>2</sup></sub>. Hence the expected number of cheapest bundles of size k + 1 that are not in Q<sub>v</sub> till a set in B<sub>v</sub> is found is at most m<sup>2</sup>, that is, E [|F<sup>p</sup><sub>k+1</sub> \ Q<sub>v</sub>|] ≤ m<sup>2</sup>.
- For d = k 1, there are two possible values for a bundle of size k 1, depending on whether the bundle is in  $\mathcal{R}_v$  or not. Let  $S_{\mathcal{R}_v}$  be a cheapest bundle in  $\mathcal{R}_v$  that is the first in order of set prices. A most profitable set of size d = k - 1 is then either  $S_{\mathcal{R}_v}$  or  $S_{k-1}^p$ . We add to  $\mathcal{F}_{k-1}^p$  the cheapest sets of size d = k - 1 in increasing order of price, till we find the first set that belongs to  $\mathcal{R}_v$ . Note that  $S_{k-1}^p$  is the first added set and is always in  $\mathcal{F}_{k-1}^p$ .

Consider some set S of size k - 1, we need bound the probability that S is in  $\mathcal{R}_{v'}$  given that v' is sampled according to  $\mathcal{D}_{Q_v}$ . For given  $v', S \in \mathcal{R}_{v'}$  if none of its supersets are in  $\mathcal{B}_{v'}$ . Hence, by Lemma 5.6, since  $Q_v$  is in a canonical form, for any set S of size k - 1 we have that either  $\Pr_{v' \sim \mathcal{D}_{Q_v}}[S \in \mathcal{R}_{v'}] = 0$  or  $\Pr_{v' \sim \mathcal{D}_{Q_v}}[S \in \mathcal{R}_{v'}] \ge (1 - \frac{1}{m^2})^{(m-k)(m-k-1)/2} > (1 - \frac{1}{m^2})^{\frac{m^2-1}{2}} > e^{-0.5} > 0.5$ . Hence the expected number of cheapest bundles of size k - 1 that are not in  $Q_v$  till a set in  $\mathcal{R}_{v'}$  is found (for  $v' \sim \mathcal{D}_{Q_v}$ ) is at most 2, that is,  $\mathbb{E}\left[|\mathcal{F}_{k-1}^{\mathcal{P}} \setminus Q_v|\right] \le 2$ . • For d = k, there are two possible values for a bundle of size k, depending on whether the

• For d = k, there are two possible values for a bundle of size k, depending on whether the bundle is  $G_v$  or not. A most profitable set of size d = k is either  $G_v$  or  $S_k^p$ . In Claim 5.8 we show that when it is not  $S_k^p$  then it must be either a set of size k that is a subset of a set in  $\mathcal{F}_{k+1}^p$ , or a set of size k that is a superset of a set in  $\mathcal{F}_{k-1}^p$ . Thus we define  $\mathcal{F}_k^p$  to include  $S_k^p$ , all sets of size k that are a subset of a set in  $\mathcal{F}_{k+1}^p$ , and all sets of size k that are a superset of a set in  $\mathcal{F}_{k+1}^p$ , and all sets of size k that are a superset of a set in  $\mathcal{F}_k^p$  satisfies  $\mathbb{E}_{v\sim D}\left[|\mathcal{F}_k^p \setminus Q_v| \cdot (m-k)\right] + \mathbb{E}_{v\sim D}\left[|\mathcal{F}_{k+1}^p \setminus Q_v| \cdot (k+1)\right] + 1$ .

Simplicity in Auctions Revisited: The Primitive Complexity

In total,

$$\mathbb{E}_{v \sim D}\left[\sum_{d \in [m]} |\mathcal{F}_d^p \setminus Q_v|\right] \leq (m-3) + \mathbb{E}_{v \sim D}\left[|\mathcal{F}_{k-1}^p \setminus Q_v| \cdot (m-k+1)\right] + \mathbb{E}_{v \sim D}\left[|\mathcal{F}_{k+1}^p \setminus Q_v| \cdot (k+2)\right] + 1$$
$$\leq m+2(m-k+1) + m^2(k+2) = m^3/2 + 2m^2 + 2m + 2 \leq 2m^3$$

as  $k = \frac{m}{2}$  and  $m \ge 4$ .

Thus, each demand query can implemented by  $2m^3$  value queries in expectation. As we mentioned earlier, moving to a canonical form requires replacing each value query with at most  $m^2$  value queries. Therefore, the demand query can be implemented while remaining in a canonical form using  $2m^5$  value queries in expectation.

We now complete the proof by showing that for d = k, the family  $\mathcal{F}_d^p$  contains a most profitable bundle whenever the most profitable bundle is of size k.

CLAIM 5.8. For any price vector p, if for valuation v every demanded set is of size k, then any most profitable set (of size k) is either  $S_k^p$  (cheapest set of size k), a set of size k that is a subset of a set in  $\mathcal{F}_{k+1}^p$ , or a set of size k that is a superset of a set in  $\mathcal{F}_{k-1}^p$ .

PROOF. Assume that every demanded set is of size k. Fix any set that most profitable set of size k and denote it by  $U_k^p$ . If  $S_k^p = G_v$  then it must be that  $U_k^p = G_v$ , that is,  $G_v$  must be the unique most profitable set of size k (as  $G_v$  has higher value than any other set of size k). So we can assume that  $S_k^p \neq G_v$ . If  $U_k^p \neq G_v$  then  $U_k^p$  must be  $S_k^p$ . We thus assume that  $U_k^p = G_v$  (and  $S_k^p \neq G_v$ ), and the value of  $U_k^p$  is thus k - 6/11.

The family  $\mathcal{F}_{k-1}^p$  includes a set R from  $\mathcal{R}_v$  with value k-1. As  $U_k^p = G_v$  is more profitable than R, it holds that  $v(G_v) - p(G_v) = k - 6/11 - p(G_v) > k - 1 - p(R)$  and thus  $p(G_v) - 5/11 < p(R)$ . If there is an item  $i \in G_v$  of price at least 5/11 then the set  $G_v \setminus \{i\}$  has price smaller than  $p(G_v \setminus \{i\}) \le p(G_v) - 5/11 < p(R)$ , and thus the set  $G_v \setminus \{i\}$  is in  $\mathcal{F}_{k-1}^p$  which implies that  $G_v \in \mathcal{F}_k^p$  as needed.

The family  $\mathcal{F}_{k+1}^p$  includes a set *B* from  $\mathcal{B}_v$  with value *k*. As  $U_k^p = G_v$  is more profitable than *B*, it holds that  $v(G_v) - p(G_v) = k - 6/11 - p(G_v) > k - p(B)$  and thus  $p(B) - p(G_v) > 6/11$ . If there is an item  $i \notin G_v$  of price lower than 6/11 then the set  $G_v \cup \{i\}$  has price smaller than  $p(G_v) + 6/11 < p(B)$ , and thus the set  $G_v \cup \{i\}$  is in  $\mathcal{F}_{k+1}^p$  which implies that  $G_v \in \mathcal{F}_k^p$  as needed.

Otherwise, the price of every item in  $U_k^p$  is less than  $\frac{5}{11}$ , and the price of every item not in  $U_k^p$  is more than  $\frac{6}{11}$ , and thus  $U_k^p = S_k^p$  is the unique cheapest bundle of size k, a contradiction to  $G_v = U_k^p \neq S_k^p$ .

This completes the proof of Claim 5.7.

This completes the proof of Lemma 5.3.

We now conclude the proof of Theorem 5.1. By Lemma 5.3, we have that with probability  $1 - \frac{1}{\alpha}$  over  $\mathcal{D}$ , a deterministic algorithm that makes t demand and value queries can be implemented using  $2m^5 \cdot t^2 \cdot \alpha$  values queries in a canonical form. Let  $t' = 2m^5 \cdot t^2 \cdot \alpha$ . By lemma 5.2, implementation that uses a set of t' value queries that is in a canonical form, for  $t' < 1.9^m$  and large enough m, has a probability of at most  $\frac{t'}{1.9^m}$  for finding  $G_v$ . Hence, the original algorithm fails with probability at least  $1 - \alpha^{-1} - \frac{t'}{1.9^m}$ . Taking  $t = 1.3^m$ ,  $\alpha = 3$ , and m large enough, we have that a deterministic algorithm that makes at most  $1.3^m$  queries, fails with probability at least  $1 - \frac{6 \cdot m^5 \cdot 1.3^{2m}}{1.9^m} - \frac{1}{3} > \frac{1}{2}$  over  $\mathcal{D}$ .

#### 5.2 An $\Omega(\sqrt{m})$ Deterministic Lower Bound for Additive Valuations

In this section we show that randomization is inherently required for maximizing an additive valuation subject to a cardinality constraint. We show that every deterministic algorithm that always finds an item with the smallest value (equivalently, finds an (m - 1)-optimal set) must make  $\sqrt{m} - 1$  value and demand queries. We prove the following theorem:

THEOREM 5.9. Let A be a deterministic algorithm that for any additive valuation finds an (m-1)-optimal set using value and demand queries. Then, A makes at least  $\Omega(\sqrt{m})$  queries.

Since we discuss only additive valuations in this section, we abuse notation and sometimes refer to additive valuations as vectors in  $\mathbb{R}^m$ , with each element representing the value of the corresponding item. We start with several lemmas regarding linear constraints that will be useful in the proof. First, recall that *basic feasible solutions (BFS)* are non-negative solutions to a linear system with minimal support (see, e.g., [Eisenbrand and Shmonin 2006; Matousek and Gärtner 2007]).

LEMMA 5.10 ([MATOUSEK AND GÄRTNER 2007]). For  $A \in \mathbb{R}^{\ell \times m}$  and  $b \in \mathbb{R}^{\ell}$ , if the system  $(A \cdot v = b, v \ge 0)$  has a solution, it has a solution with support of size at most  $\ell$ .

The following simple observation will be useful later in the proof.

CLAIM 5.11. Let  $\alpha, \beta \in \mathbb{R}$  such that  $\alpha + \beta = 1$ . If  $\vec{x}, \vec{y}$  are solutions of the linear system  $A \cdot v = b$ , then so is  $\vec{z} = \alpha \cdot \vec{x} + \beta \cdot \vec{y}$ .

**PROOF.** Since both  $\vec{x}$  and  $\vec{y}$  are solutions of the linear system  $A \cdot v = b$ , we have that

$$A \cdot \vec{z} = A \cdot (\alpha \cdot \vec{x} + \beta \cdot \vec{y}) = \alpha \cdot A \cdot \vec{x} + \beta \cdot A \cdot \vec{y} = (\alpha + \beta) \cdot b = b$$

LEMMA 5.12. Let  $A \in \mathbb{R}^{\ell \times m}$  be a matrix of rank at most m - 1. The linear system  $A \cdot v = A \cdot \vec{I}$  with the constraint  $v \ge 0$  has two different solutions which disagree on the identity of the item with the smallest value.

PROOF. The vector  $\vec{\mathbf{1}}$  is a solution. In addition, since *A*'s rank is smaller than *m*, there exists some solution  $\vec{w} \in \mathbb{R}^m$  such that  $\vec{w}$  and  $\vec{\mathbf{1}}$  are independent. Let  $\vec{x} = (1 - \epsilon)\vec{\mathbf{1}} + \epsilon \cdot \vec{w}$  and  $\vec{y} = (1 + \epsilon)\vec{\mathbf{1}} - \epsilon \cdot \vec{w}$  for a small enough  $\epsilon > 0$  to ensure that all entries in  $\vec{x}$  and  $\vec{y}$  are positive. Note that by Claim 5.11 both  $\vec{x}$  and  $\vec{y}$  are solutions of the linear system. The values in  $\vec{x}$  and  $\vec{y}$  are ordered according to  $\vec{w}$  where a minimal item in  $\vec{x}$  is maximal and  $\vec{y}$  and vice versa. Since  $\vec{w}$  is not the all zero vector, these two solutions disagree on the identity of the items with the smallest value.

We are now ready to prove Theorem 5.9. Throughout the proof, we assume that in all demand queries that *A* makes the prices are strictly positive. This assumption only doubles the number of queries: if there is a demand query that gives 0 prices for some set of items *S*, we can replace all 0 prices with  $\infty$ , run the new demand query and return the union of *S* and the answer *T* of the new demand query. Since the demand query should also return the value of  $S \cup T$  we can make one value query v(S) and return v(S) + v(T) (recall that a demand query also returns the value of the most demanded set). In addition, we assume that the first query is a value query for the entire set *M*, increasing the total queries made by the algorithm by at most one. As the algorithm must work with any implementation of the demand query, it must work with the one we specify here.

PROOF OF THEOREM 5.9. We show that for every deterministic algorithm A there is an adversary that can answer all queries in a way such that as long as the algorithm did not make many queries, there are two different valuations v, v', both consistent with the queries asked, each has a unique

item with minimal value, but the items with the minimal value in v and in v' are different. Thus, the algorithm does not distinguish between v and v' and does not find an (m - 1)-optimal bundle.

We will describe a set of valuations and inductively show that every valuation in this set is consistent with the queries asked so far. Thus, every item that is a minimal item of a valuation in the set is a possible solution.

The next claim is the heart of the proof. To give some intuition, let us examine some of the possible answers of the adversary. For every bundle *S* whose value is queried, the algorithm will return the value |S|. Suppose that a demand query, all with positive prices, is made. The adversary now "defines" a set of items *L* whose value we "set" to be very big. The items in *L* are the items that the demand query returns. We consider the values of items that are not in *L* to be very low, although the adversary does not commit on their specific values (so any of them might be the minimal item). In the next value queries we will treat every  $S \subseteq M - L$  as having value  $v(S) = \varepsilon' \cdot |S|$ , for  $\varepsilon \gg \varepsilon'$  ( $\varepsilon'$  will be smaller than the smallest price in the demand query). Thus, the demand query will only return the items in *L*.

There are several challenges in achieving this. The first is making sure that after the first demand query we are consistent with the value queries that were done so far. The second is to be able to answer not just the first demand query but also the following ones. The third is to make sure that the set of items L is small, otherwise the adversary has to commit on the values of all items too quickly. The next claim handles all these challenges.

CLAIM 5.13. Consider an execution of the algorithm after a set of Q queries. Let  $t_r$  denote the number of value queries made before the r'th demand query in Q. Suppose that Q contains i demand queries and x value queries. Let x' denote the number of value queries made after the last (i'th) demand query. There exists an adversary, a non-empty set of valuations  $\mathcal{V}^Q$ , a set of items  $L_Q \subseteq M$  where  $|L_Q| = \sum_{r \leq i} t_r$ , a matrix  $A_Q \in \{0, 1\}^{(t_i+x') \times (m-|L_Q|)}$  and  $\varepsilon_Q > 0$  such that:

- All valuations in  $v \in \mathcal{V}^Q$  are consistent with all the queries made so far.
- For every  $v, v' \in \mathcal{V}^Q$  and  $j \in L_Q$ ,  $v(\{j\}) = v'(\{j\})$ .
- By renaming, assume without loss of generality that the items that are not in  $L_Q$  are indexed  $1, \ldots, m |L_Q|$ . For each  $v \in \mathcal{V}^Q$  it holds that:

$$A_Q \cdot v = \varepsilon_Q \cdot A_Q \cdot \mathbf{1}$$

PROOF. We prove the claim by induction over the number of queries. The base case is when Q includes a single query, a value query for M that returns a value of m. Let  $L_Q = \emptyset$ ,  $\varepsilon_Q = 1$ , and  $A_Q$  which is a single all-one line. The claim trivially holds for the case the valuation  $\vec{\mathbf{1}}$  satisfies all three requirements.

We now assume the induction hypothesis for Q', and prove for Q, where Q is Q' with an additional query. We describe how the adversary answers the additional query and how to obtain  $A_Q$ ,  $L_Q$ ,  $\varepsilon_Q$ , for which the claim holds after each query. For each item  $j \in L_{Q'}$ , let  $L_{Q'}(j)$  be the value of item j for all valuations in  $\mathcal{V}^{Q'}$ .

We first consider the case where the additional query is a value query, and then the case where the additional query is a demand query. If the algorithm makes a value query for some set *S*, the adversary answers that the value of *S* is  $\varepsilon_{Q'} \cdot |S \setminus L_{Q'}| + \sum_{j \in S \cap L_{Q'}} L_{Q'}(j)$ . This answer is consistent with the previous queries since there exists a valuation  $v \in \mathcal{V}^{Q'}$  that obeys the conditions in the statement of the lemma: e.g.,  $v(\{j\}) = L_{Q'}(j)$  for every  $j \in L_{Q'}$  and  $v(\{j\}) = \varepsilon_{Q'}$  for every  $j \notin L_{Q'}$ . Define  $L_Q = L_{Q'}, \varepsilon_Q = \varepsilon_{Q'}$  and the set of linear equations  $A_Q$  to be  $A_{Q'}$  with the additional constraint that  $v(S \setminus L_{Q'}) = \varepsilon_{Q'} \cdot |S \setminus L_{Q'}|$ .

If the additional query is a demand query, the adversary answers it as follows. Find a solution for the system  $A_{Q'} \cdot v = \varepsilon_{Q'} \cdot A_{Q'} \cdot \vec{\mathbf{1}}$  and  $v \ge 0$ , with the smallest support size. Since the vector

 $\varepsilon_{Q'} \cdot \vec{1}$  is a solution, by Lemma 5.10 there exists a solution  $\vec{r} \in \mathbb{R}^{m-|L_{Q'}|}_+$  with at most  $t_{i+1} = t_i + x'$  non-zero coordinates (recall that *i* is the number of demand queries in Q'). According to Claim 5.11, the vector  $\vec{r'} = (1 - \frac{\varepsilon_Q}{\varepsilon_{Q'}})\vec{r} + \varepsilon_Q \cdot \vec{1}$  is also a valid solution.

We now show that the induction hypothesis holds for Q with any  $\varepsilon_Q < \varepsilon_{Q'}$ . Let  $L_Q$  be  $L_{Q'}$  with all items that their value is non-zero in  $\vec{r}$ . We set the common value  $L_Q(j)$  of any such newly added item  $j \in L_Q \setminus L_{Q'}$  to be its value in the solution  $\vec{r'}$ . We define  $\mathcal{V}^Q$  to be all valuations in  $\mathcal{V}^{Q'}$  that have  $v(j) = L_Q(j)$  for every  $j \in L_Q \setminus L_{Q'}$ .

The set of linear equations  $A_Q$  is obtained from  $A_{Q'}$  by updating every constraint  $\sum_{a \in S \setminus L_{Q'}} v(a) = b$ in  $A_{Q'}$  to  $\sum_{a \in S \setminus L_Q} v(a) = b - \sum_{a \in S \cap L_Q} v(a) = \varepsilon_Q \cdot |S \setminus L_Q|$ .

Note that  $\mathcal{V}^Q$  is not empty as it contains the valuation defined by  $\vec{r'}$  and the values of the items in  $L_{Q'}$ . Furthermore, all valuations in  $\mathcal{V}^Q$  are consistent with all queries made so far, and in particular with the last demand query. We now set  $\varepsilon_Q$  to be small enough to make sure that only items in  $L_Q$  could be in the demand set. If  $p_{min}$  is the minimal price of an item in the (i + 1)'th (last) demand query,  $\varepsilon_Q$  is chosen to be strictly smaller than  $\frac{p_{min}}{m}$ . Since the first line in  $A_Q$  has that  $\sum_{a \in M \setminus L_Q} v(a) = \varepsilon_Q \cdot |M \setminus L_Q|$ , this guarantees that for any item  $j \notin L_Q$ , the maximal possible value for j is smaller than  $m \cdot \varepsilon_Q < p_{min}$  and thus j is not in any most profitable set for the i + 1'th demand query. Hence, committing on the values of items in  $L_Q$  is sufficient for the implementation of the demand query. The adversary then answers the demand query according to the valuation defined by the solution  $\vec{r'}$ .

Using Claim 5.13 we are now ready to complete the proof of the theorem. Let  $q_v$  be the total number of value queries and  $q_d$  be the total number of demand queries made by the algorithm. Let Q be the list of queries. Since the *i*'th demand query adds at most  $t_i$  new items to the set  $L_Q$  and  $t_i$  is at most  $q_v$ , the total number of items in  $L_Q$  is  $q_v \cdot q_d$ . Thus, the linear system has at most  $q_v$  constraints. Furthermore, each item  $j \in L_Q$  corresponds to a linear constraint of the form v(j) = t for some known value *t*. That is, the set  $\mathcal{V}^Q$  is non-empty and defined by at most  $q_v \cdot q_d + q_v$  linear constraints. When  $q_d, q_v \leq \sqrt{m} - 1$ , we have that the total number of constraints is at most  $q_v \cdot q_d + q_v < m - 1$ . By Lemma 5.12 the algorithm fails for some valuation v.

#### 5.3 Impossibilities for Additive Valuations using Value Queries

For deterministic algorithms, we show that m - 1 value queries are needed to find an (m - 1)-optimal set. For randomized algorithms, we show two different proofs that  $\Omega\left(\frac{m}{\log m}\right)$  queries are needed to find a *k*-optimal set (with each proof using a different *k*) with non-negligible success probability.

PROPOSITION 5.14. Let A be a deterministic algorithm that given an additive valuation v makes q value queries and returns an (m - 1)-optimal set (equivalently, finds the item with the smallest value). Then, the algorithm makes at least m - 1 value queries.

PROOF. We construct a hard valuation v "on the fly" by observing the queries that A makes and (partially) defining v appropriately (A is deterministic so the queries that it makes are only a function of the values that were returned in the previous queries). We will see that for every set of queries that A makes, after m - 2 queries there are at least two ways to complete the definition of v, and each completion has a different minimal-value item. Thus, A cannot determine the minimal item after m - 2 queries.

Specifically, in each of the first m - 2 queries that A makes, we partially define v as follows: whenever the algorithm queries v(S), set v(S) = |S|. Note that this construction is indeed valid in the sense that it is easy to extend it to a fully defined additive valuation (e.g., by setting v(S) = |S| for every bundle S).

The queries and answers naturally define a set of m - 2 linear equations with m variables, each equation has the form  $\sum_{j \in S} x_j = |S|$ , for some S. If there are less than m - 2 independent rows in the corresponding matrix, we add some independent rows of the same form. Note that the newly added rows correspond to additional queries that the algorithm makes.

Since there are *m* equations and m - 2 variables, the matrix has infinitely many solutions. In particular, there are m - 2 variables such that if we set each of those variable to 1, there are still two free variables, denoted *w*, *z*. Note that *w*, *z* are constrained by at most one linear equation, and since  $x_i = 1$  for every *i* is a valid solution, this equation must take the form w + z = 2. Thus, both  $w = \frac{1}{2}$ ,  $z = 2 - \frac{1}{2}$ , and  $w = 2 - \frac{1}{2}$ ,  $z = \frac{1}{2}$  complete setting every other variable to 1 to two different valid solutions to the set of equations. Moreover, each of these solutions defines an additive valuation with a different minimal-value item. Therefore, *A* cannot distinguish between these two valuations and find the minimal-value item if it makes less than m - 1 queries.

PROPOSITION 5.15. Let A be a randomized algorithm that given an additive valuation v makes q value queries, and with a constant positive probability returns a (m/2)-optimal set. Then, the algorithm makes in expectation at least  $\Omega\left(\frac{m}{\log m}\right)$  value queries.

PROOF. Consider the following distribution over additive valuations: choose at random set of m/2 items and set the value of each item in the set to 1. Set the value of the rest of the items to 0. By Yao's principle, it is enough to show that every deterministic algorithm that succeeds with probability  $\frac{2}{3}$  makes at least  $\Omega\left(\frac{m}{\log m}\right)$  value queries in expectation.

Fixing such a deterministic algorithm, there is a family of  $\frac{2}{3} \cdot \left(\frac{m}{2}\right)$  additive valuations on which the algorithm succeeds. We associate each such valuation with a transcript of the run of the algorithm on this valuation. Such transcript consists of q answers to value queries, each takes  $\log \frac{m}{2}$  bits to write down since the possible values are integers between 0 to  $\frac{m}{2}$ . Since the algorithm is deterministic we have that the identity of the bundle S that is queried is only a function of the answers to the previous queries. Thus representing each transcript takes  $q \cdot \log \frac{m}{2}$  bits. Note that each such transcript must be different for every valuation that the algorithm succeeds on, and thus we have that  $\frac{2}{3} \cdot \log {\binom{m}{m/2}} \le q \cdot \log(m/2)$ , and since  ${\binom{m}{m/2}} > \frac{2^m}{m+1}$  we have that  $q \ge \frac{2}{3} \cdot \frac{m-\log(m+1)}{\log(m/2)} = \Omega\left(\frac{m}{\log m}\right)$  as needed.

PROPOSITION 5.16. Let A be a randomized algorithm that given an additive valuation v makes q value queries and returns a maximal value item (1-optimal set) with a constant positive probability. Then, the algorithm makes in expectation at least  $\Omega\left(\frac{m}{\log m}\right)$  value queries.

PROOF. We prove our lower bound using a reduction to the Set Disjointness problem. In this problem, Alice and Bob, are given two input vectors  $x = (x_1, ..., x_m)$ ,  $y = (y_1, ..., y_m) \in \{0, 1\}^m$  respectively, and they wish to determine whether there exists an index *i* such that  $x_i = y_i = 1$ . The randomized communication complexity of DISJ is  $\Omega(m)$ , see, e.g., [Razborov 1990].

Given an input  $x, y \in \{0, 1\}^m$  for the disjointness problem, define an additive valuation by setting the value of each item  $w_i$  to be  $v(\{w_i\}) = x_i + y_i$ . We follow the run of *A* and simulate the value queries that *A* makes: when *A* queries the value of the bundle *S*, Alice uses log *m* bits to send  $\sum_{j \in S} x_j$ and Bob uses log *m* bits to send  $\sum_{j \in S} y_j$ . The value v(S) is simply the sum of these two numbers. Thus we have that if *A* makes *q* queries then it can be simulated for this family of instances by a communication protocol that takes  $2 \cdot q \cdot \log m$  bits.

Observe that the maximal value of a single item in v equals to 2 if and only if there exists an index *i* for which the *i*-th variable has value 1 for both Alice and Bob. Therefore, the communication protocol must use  $\Omega(m)$  bits, which immediately implies that  $q = \Omega\left(\frac{m}{\log m}\right)$ .

#### REFERENCES

- Tarek Abdallah, Arash Asadpour, and Josh Reed. 2021. Large-Scale Bundle-Size Pricing: A Theoretical Analysis. Operations Research (2021).
- Saeed Alaei. 2014. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM J. Comput.* 43, 2 (2014), 930–972.
- Saeed Alaei, Jason Hartline, Rad Niazadeh, Emmanouil Pountourakis, and Yang Yuan. 2019. Optimal auctions vs. anonymous pricing. Games and Economic Behavior 118 (2019), 494–510.
- Sepehr Assadi and Sahil Singla. 2019. Improved truthful mechanisms for combinatorial auctions with submodular bidders. In 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS). IEEE, 233–248.
- Moshe Babaioff, Yannai A Gonczarowski, and Noam Nisan. 2021. The menu-size complexity of revenue approximation. Games and Economic Behavior (2021).
- Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. 2020. A simple and approximately optimal mechanism for an additive buyer. *Journal of the ACM (JACM)* 67, 4 (2020), 1–40.
- Moshe Babaioff, Noam Nisan, and Aviad Rubinstein. 2018. Optimal Deterministic Mechanisms for an Additive Buyer. In ACM Conference on Economics and Computation (ACM-EC).
- Ashwinkumar Badanidiyuru, Shahar Dobzinski, and Sigal Oren. 2012. Optimization with demand oracles. In Proceedings of the 13th ACM conference on electronic commerce. 110–127.
- Liad Blumrosen and Noam Nisan. 2010. On the computational power of demand queries. *SIAM J. Comput.* 39, 4 (2010), 1372–1391.
- Nader H Bshouty. 2009. Optimal Algorithms for the Coin Weighing Problem with a Spring Scale. In COLT, Vol. 2009. 82.
- Yang Cai and Mingfei Zhao. 2017. Simple mechanisms for subadditive buyers via duality. In Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing. 170–183.
- Gruia Calinescu, Chandra Chekuri, Martin Pal, and Jan Vondrák. 2011. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM J. Comput.* 40, 6 (2011), 1740–1766.
- JM Chambers. 1971. Algorithm 410: partial sorting. Commun. ACM 14, 5 (1971), 357-358.
- Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*. 311–320.
- Shuchi Chawla, Yifeng Teng, and Christos Tzamos. 2020. Menu-size complexity and revenue continuity of buy-many mechanisms. In *Proceedings of the 21st ACM Conference on Economics and Computation*. 475–476.
- Chenghuan Sean Chu, Phillip Leslie, and Alan Sorensen. 2011. Bundle-size pricing as an approximation to mixed bundling. *The American Economic Review* (2011), 263–303.
- Peter Cramton. 1998. Ascending auctions. European Economic Review 42, 3-5 (1998), 745–756.
- AG Djackov. 1975. On a search model of false coins. In *Topics in Information Theory (Colloquia Mathematica Societatis Janos Bolyai 16). Budapest, Hungary: Hungarian Acad. Sci.* 163–170.
- Shahar Dobzinski, Noam Nisan, and Michael Schapira. 2006. Truthful randomized mechanisms for combinatorial auctions. In Proceedings of the thirty-eighth annual ACM symposium on Theory of computing. 644–652.
- Dingzhu Du, Frank K Hwang, and Frank Hwang. 2000. Combinatorial group testing and its applications. Vol. 12. World Scientific.
- Shaddin Dughmi, Li Han, and Noam Nisan. 2014. Sampling and representation complexity of revenue maximization. In International Conference on Web and Internet Economics. Springer, 277–291.
- Paul Dütting, Felix Fischer, and David C Parkes. 2011. Simplicity-expressiveness tradeoffs in mechanism design. In Proceedings of the 12th ACM conference on Electronic commerce. 341–350.
- Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S Matthew Weinberg. 2021. A simple and approximately optimal mechanism for a buyer with complements. *Operations Research* 69, 1 (2021), 188–206.
- Friedrich Eisenbrand and Gennady Shmonin. 2006. Carathéodory bounds for integer cones. *Operations Research Letters* 34, 5 (2006), 564–568.
- Uriel Feige and Jan Vondrak. 2006. Approximation algorithms for allocation problems: Improving the factor of 1-1/e. In 2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06). IEEE, 667–676.
- Uriel Feige and Jan Vondrák. 2010. The submodular welfare problem with demand queries. *Theory of Computing* 6, 1 (2010), 247–290.
- Michal Feldman, Nick Gravin, and Brendan Lucier. 2014. Combinatorial auctions via posted prices. In Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms. SIAM, 123–135.
- Yannai A Gonczarowski. 2018. 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. 123–131.
- Faruk Gul and Ennio Stacchetti. 2000. The English auction with differentiated commodities. *Journal of Economic theory* 92, 1 (2000), 66–95.

- Sergiu Hart and Noam Nisan. 2019. Selling multiple correlated goods: Revenue maximization and menu-size complexity. Journal of Economic Theory 183 (2019), 991–1029.
- Jason D Hartline and Tim Roughgarden. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*. 225–234.

C. A. R. Hoare. 1961. Algorithm 65: Find. Commun. ACM 4, 7 (1961), 321-322.

- Donald E Knuth. 1974. The asymptotic number of geometries. *Journal of Combinatorial Theory, Series A* 16, 3 (1974), 398–400.
- Pravesh Kothari, Sahil Singla, Divyarthi Mohan, Ariel Schvartzman, and S Matthew Weinberg. 2019. Approximation schemes for a unit-demand buyer with independent items via symmetries. In 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS). IEEE, 220–232.
- Xinye Li and Andrew Chi-Chih Yao. 2013. On revenue maximization for selling multiple independently distributed items. Proceedings of the National Academy of Sciences 110, 28 (2013), 11232–11237.
- Bernt Lindstrom. 1975. Determining subsets by unramified experiments. A Survey of Statistical Design and Linear Models (1975).
- Jiri Matousek and Bernd Gärtner. 2007. Understanding and using linear programming. Springer Science & Business Media.
- Debasis Mishra and David C Parkes. 2007. Ascending price Vickrey auctions for general valuations. *Journal of Economic Theory* 132, 1 (2007), 335–366.
- Benny Moldovanu and Manfred Tietzel. 1998. Goethe's second-price auction. *Journal of Political Economy* 106, 4 (1998), 854–859.
- David R Musser. 1997. Introspective sorting and selection algorithms. *Software: Practice and Experience* 27, 8 (1997), 983–993.
- George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. 1978. An analysis of approximations for maximizing submodular set functions—I. *Mathematical programming* 14, 1 (1978), 265–294.
- Noam Nisan and Ilya Segal. 2006. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory* 129, 1 (2006), 192–224.
- Alexander A Razborov. 1990. On the distributional complexity of disjointness. In International Colloquium on Automata, Languages, and Programming. Springer, 249–253.
- Amir Ronen. 2001. On approximating optimal auctions. In Proceedings 3rd ACM Conference on Electronic Commerce (EC-2001), Tampa, Florida, USA, October 14-17, 2001, Michael P. Wellman and Yoav Shoham (Eds.). ACM, 11–17. https://doi.org/10.1145/501158.501160
- Aviad Rubinstein and S Matthew Weinberg. 2018. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. ACM Transactions on Economics and Computation (TEAC) 6, 3-4 (2018), 1–25.
- Aviad Rubinstein and Junyao Zhao. 2021. The randomized communication complexity of randomized auctions. In Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing. 882–895.
- Raghuvansh R Saxena, Ariel Schvartzman, and S Matthew Weinberg. 2018. The menu complexity of "one-and-a-halfdimensional" mechanism design. In Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM, 2026–2035.
- David RM Thompson and Kevin Leyton-Brown. 2013. Revenue optimization in the generalized second-price auction. In Proceedings of the fourteenth ACM conference on Electronic commerce. 837–852.
- Jan Vondrák. 2008. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*. 67–74.

# A MISSING PROOFS FROM SECTION 3

## A.1 Proof of Theorem 3.1

The following two lemmas will be useful for constructing an algorithm to find *k*-optimal set for additive valuations with small number of value and demand queries.

LEMMA A.1. For a subadditive valuation v over a set M of m items and a threshold t. Let L be the set of elements of value (as a singleton) larger than t, that is  $L = \{k \in M | v(k) > t\}$ .

There exists an algorithm such that for any v and t, if L is not empty returns an element from L picked uniformly at random. The algorithm makes  $O(\log m)$  value and demand queries.

PROOF. Denote the set returned by a demand query with uniform price t by D. If  $t \cdot |D| = v(D)$  we have that the profit from D is zero (v(D) - p(D) = 0) which implies that there is no item  $d \in M$ 

such that v(d) > t, otherwise the profit from buying item d alone is positive, contradicting that D is a set in demand. Thus, if  $t \cdot |D| = v(D)$  then L is empty and the algorithm terminates.

Otherwise,  $v(D) > t \cdot |D|$ . By subadditivity, there exists  $d \in D$  such that v(d) > t and thus L is non-empty. Given t, the algorithm works as follows: It keeps a set N containing all items who might be in L, initiated as N = M. It then goes iteratively:

- If |N| = 1, the single element of N is in L, and we return it.
- Otherwise, split the set N into two random sets,  $N_0$  and  $N_1$ , of sizes as equal as possible (equal up to one item, that is,  $||N_0| |N_1|| \le 1$ ). For each one of the two sets  $N_0$  and  $N_1$ , the algorithm makes a uniform price demand query at price t. Denote the returned sets by  $D_0$  and  $D_1$ . If the returned set  $D_i$  is empty or satisfies  $t \cdot |D_i| = v(D_i)$ , as before, the set contains no element from L, and is discarded. As  $v(D) > t \cdot |D|$ , by subadditivity,  $v(D_i) > t \cdot |D_i|$  for at least one  $i \in \{0, 1\}$ , so at least one of these non-discarded sets at random. Denote it corresponding set  $N_i$  by N, and recursively run the algorithm on this new set N and threshold t.

First, observe that the algorithm terminates in  $O(\log m)$  demand queries, as the size of D shrinks by a constant factor (about 2) at every iteration.

Finally, we observe that if *L* is non-empty then every element in *L* has the same probability of being picked. First observe that only elements in *L* are ever returned. Secondly, at each iteration, one set  $N_i$  is discarded. It can be since  $v(D_i) = t \cdot |D_i|$ , in that case no item from *L* is discarded. In the other case, both  $N_0$  and  $N_1$  contains items from *L* and the probability any item from *L* to be discarded in this round is equal. As this claim is true for every  $j \in L$ , all elements in *L* have the same probability of being selected.

If the valuation is additive, we can use Lemma A.1 to sample a random element *smaller* than a given threshold in  $O(\log m)$  demand queries.

LEMMA A.2. For an additive valuation v over a set M of m items and threshold t, let Z be the set of elements of value (as a singleton) smaller than t, that is  $Z = \{k \in M | v(k) < t\}$ .

There exists a randomized algorithm such that for any v and t, if Z is not empty returns an element from Z picked uniformly at random, using  $O(\log m)$  value and demand queries in expectation.

PROOF. Let W be a large enough constant, say W = v(M) + 1. Consider the additive valuation v' in which for every item  $j \in M$  the value of j is  $v'_j = W - v_j$ . Picking a random item from  $Z = \{k \in M | v(k) < t\}$  is equivalent to picking a random item from  $L' = \{k \in M | v'(k) > W - t\}$ . We now use Lemma A.1 with valuation v' and threshold W - t to sample a random item from L' (or equivalently, a random element form Z). To use the lemma we observe that both value and demand queries for v' can be simulated by value and demand queries on v. For value query on a set S the value v'(S) is simply  $v'(S) = |S| \cdot W - v(S)$ , and a demand query on v' with price  $p_j$  for item j is replaced by a demand query on v with price  $W - p_j$ .

We next use the above lemma to find, for any fixed value t, all items of value exactly t.

LEMMA A.3. There exists a randomized algorithm that for any additive valuation v over a set M of size m and a value  $t \in \mathbb{R}$ , using value and demand queries finds all items in M with value exactly t. The algorithm makes in expectation  $O(\log^2 m)$  queries.

PROOF. We start with a uniform demand query where the price of each item is t, splitting the items into a demand set D and its complement  $C = M \setminus D$ . Any item of value exactly t can belong to any of the two sets. We show how to find all items of value exactly t in D, the algorithm for C is the same up to trivial adjustments.

First, observe that either the set *D* does not contain items of value *t* or that *t* is the minimal value in the set. A uniform demand query for a price that is strictly larger than *t* but strictly smaller than any other item in *D* will return all items of value larger than *t* in *D*, and only them. In order to find such a price, we run the following iterative algorithm: We maintain a set *S*, initiated to have all items in *D*. The set will contain any item in *D* that might have value *t*. If the average value of an item in *S* is  $t (v(S) = t \cdot |S|)$ , we are done (all elements in *S* must have value *t*). Otherwise, there is at least one element of value higher than *t*.

We now use Lemma A.1 on the set S as the set of all items and the threshold t to pick an element from  $L = \{k \in S | v(k) > t\}$  uniformly at random, in  $O(\log m)$  demand queries. Let q be the value of the element picked.

We have that in expectation, q is smaller or equal to at least half of the items in L. Let  $p = \frac{q+t}{2}$  and note that q > p > t. A uniform demand query on S at price p returns a set of items, each of value larger than t, and that set is of size at least half the size of L, in expectation (as it includes all items of value at least q). We remove all these demanded items from S, and iterate. As the problem size shrinks by factor of two,  $O(\log m)$  rounds, each with  $O(\log m)$  demand queries, suffice in expectation in order of identifying all the set L and remove it from S completely, leaving in S exactly the set of items from D that have value exactly t.

The set *C* is handled in a similar way with some minor changes. First, when using Lemma A.2 we do so to pick an element from  $Z = \{k \in S | v(k) < t\}$  uniformly at random (instead of picking from *L*) and denote its value by *v*. Unlike for *D*, to keep items of value *t* at *S* we now update *S* to include items that are **in** the demand for uniform price  $p = \frac{v+t}{2}$  which now satisfies t > p > q.

We can now conclude the proof of Theorem 3.1.

The algorithm gradually builds a k-optimal set K by maintaining a set S that contains items that are still candidates for inclusion in the set K. It runs in iterations, in each iteration items are either moved from S to K (if we determine that they are among the k items with the highest value), or removed from S (if we know for sure they are not). In expectation, S will shrink by a constant factor at each round, so the expected number of rounds is  $O(\log m)$  rounds. The algorithm continues until K is an k-optimal set. We move to present the algorithm more formally.

Given the parameter k, the algorithm works as follows. It first initializes S = M and  $K = \emptyset$ . Then it runs the following iterative procedure:

- Sample a random item  $t \in S$ . Make a uniform demand query with its value as the price for each item. In addition, using Lemma A.3, find all items of value t, i.e., find a partition of S into three sets:  $S_H = \{j \in S \mid v_j > t\}$ ,  $S_M = \{j \in S \mid v_j = t\}$  and  $S_L = \{j \in S \mid v_j < t\}$ .
  - If  $k |K| < |S_H|$ , we update S to be the set  $S_H$  (no new items are added to K). Reiterate with the updated S.
  - If  $|S_H| \le k |K| \le |S_H \cup S_M|$ , the algorithm adds  $S_H$  to K and additionally adds arbitrary items from  $S_M$  to K to complete filling it with k items altogether. It then returns K and terminate.
  - If  $|S_H \cup S_M| < k |K|$  then it adds  $S_H \cup S_M$  to K, update S to be the set  $S_L$ . Reiterate with the updated S.

We first argue that the algorithm indeed returns a k-optimal set. Since the valuation function is additive, for any r < k, a r-optimal set can be extended to a k-optimal set by adding items or largest value that are not in the set. The algorithm does this till k items are added.

Each iteration of the algorithm makes  $O(\log^2 m)$  demand queries in expectation. In expectation, at each round, at least half of the items in *S* are classified and the algorithm terminates after  $O(\log m)$  rounds in expectation.

### A.2 Proof of Theorem 3.2

We first find a maximum weight independent set *R*, and denote its rank by *r*. When  $k \le r$ , finding a *k*-optimal set reduces to finding a *k*-optimal set in *R*. Since the restriction of the valuation *v* to the set of items *R* is additive (since *R* is a base), we can use the algorithm described in Theorem 3.1 to find a *k*-optimal set. If k > r then *k*-optimal set can be constructed by adding k - r arbitrary items to *R*.

We next show that a maximum weight independent set *R* can be found using  $O(\log^2 m)$  demand queries in expectation. As by Theorem 3.1 we can find a *k*-optimal set in *R* using  $O(\log^3 m)$  demand queries in expectation, the bound on the number of queries follows.

In the next proof we use the notation  $v(\cdot|R)$  to denote the marginal valuation given a set *R*. I.e., for a set *S* we have that  $v(S|R) = v(S \cup R) - v(R)$ . Given query oracle for *v*, both value and demand queries can be easily implemented for  $v(\cdot|R)$ : value queries by querying  $v(S \cup R)$  and v(R), and demand queries by setting zero prices for all items in *R*, which by monotonicity guarantees that all items in *R* are in the demand.

LEMMA A.4. There exists a randomized algorithm that for any weighted matroid-rank valuation v over m items, finds a maximum weight independent set R using value and demand queries. The algorithm makes in expectation  $O(\log^2 m)$  queries.

PROOF. The algorithm is iterative. First, initialize *R* to be the empty set. At each round we consider the valuation  $v(\cdot|R)$  and observe that it is subadditive. Thus, we can use Lemma A.1 to pick a uniform random item *r* from the set of items satisfying v(r|R) > 0, using  $O(\log m)$  demand queries in expectation. We then make a demand query with price p = v(r|R)/2 for any item in  $M \setminus R$ , and zero for all items in *R*. The returned set *D* contains *R* and is a subset of a maximum weight independent set. We update *R* to be *D* (which is a superset of the prior *R*), and if for the updated *R* it holds that v(R) < v(M) we reiterate, again picking a random *r* and so on. The process ends when *R* satisfies v(R) = v(M), and thus is a maximum weight independent set.

Clearly, if the algorithm terminates with *R* satisfying v(R) = v(M) then *R* is indeed a maximum weight independent set - in every iteration *R* is an independent set, and as v(R) = v(M) it has maximal weight. The algorithm must terminate with *R* satisfying v(R) = v(M) as *R* size monotonically increases and is bounded, and unless v(R) = v(M) there is always an item not in *R* with positive marginal, so some item can be picked at the next iteration.

Finally, we claim that the expected number of demand queries the algorithm makes is  $O(\log^2 m)$ . Since *r* is selected randomly from all items with positive marginal value relative to *R*, we have that, in expectation, at least half of the items in  $M \setminus R$  have marginal of at least *p* relative to *R*. Each such item is either in the returned set *D* (and hence in the updated *R*) or has a negative marginal utility when added to *D*. That is, at each round, in expectation, at least half of the items are either in *R* or have a marginal 0 when added to it. Hence, after  $O(\log m)$  rounds in expectation, all items are either in *R* or have a marginal 0 relative to *R*. When that is the situation, *R* is a maximum weight independent set. Since each round requires  $O(\log m)$  queries, the total number of queries is still  $O(\log^2 m)$ .

# **B PROOF OF THEOREM 4.1**

Given a menu  $\mathcal{M}$ , we normalize<sup>11</sup> the minimal positive price to 1 and denote the highest price in the menu by  $H_{\mathcal{M}}$ .

We next use Lemma 9.1 of [Hart and Nisan 2019] to show that for any bundle-size pricing menu  $\mathcal{M}$  there exists a bundle-size pricing menu with bundle-size pricing menu size  $O(\epsilon^{-2} \ln(H_{\mathcal{M}}))$  that

<sup>&</sup>lt;sup>11</sup>It is trivial to adapt the queries to this normalization, e.g., divide the price of any demand query with this normalization factor.

obtains at least  $1 - \epsilon$  fraction of  $\mathcal{M}$ 's revenue. Our lemma makes simple observations regarding the result of [Hart and Nisan 2019] when applied to a bundle-size pricing menu.

LEMMA B.1. Given a bundle-size pricing menu  $\mathcal{M}_1$ , for any  $\epsilon > 0$  there exists a bundle-size pricing menu  $\mathcal{M}_2$  that offers at most  $2 + \frac{5}{\epsilon^2} \cdot \ln H_{\mathcal{M}_1}$  different bundles' sizes, such that for any monotone valuation v,  $\operatorname{REV}_{\mathcal{M}_2}(v) \geq (1 - \epsilon)\operatorname{REV}_{\mathcal{M}_1}(v)$ . Moreover, if  $H_{\mathcal{M}_1} > 2$ , the ratio of two different prices in  $\mathcal{M}_2$  is at least  $e^{\frac{\epsilon^2}{7}}$ .

PROOF. For  $\epsilon \ge 1$  the claim is trivially true with  $\mathcal{M}_2$  being the menu that only offers nothing for zero payment. We now assume that  $\epsilon < 1$ . Denote  $H = H_{\mathcal{M}_1}$ . We start by splitting the range [1, H] into K subranges, each with a ratio of  $H^{1/K}$  between its endpoints where K is the smallest integer such that  $H^{1/K} \le 1 + \frac{\epsilon^2}{4}$ , i.e.,  $K \le 1 + \ln^{-1}(1 + \frac{\epsilon^2}{4}) \cdot \log H$ , which is at most  $1 + \frac{5}{\epsilon^2} \cdot \ln H$  for  $\epsilon < 1$ . All prices in the same subranges are rounded to the same single price in  $\mathcal{M}_2$ . By the monotonicity assumption, if two different bundle sizes are offered for the same price, a larger size bundle will always be selected. Hence, for all price in  $\mathcal{M}_2$  we keep only entries offering the largest bundle sizes for that price and  $\mathcal{M}_2$  will have at most K + 1 different price levels.

We now explain the rounding schema. For any price *s* in  $\mathcal{M}_1$ , we apply the transform  $\phi(s)$  by rounding *s* up to the next multiple of  $H^{1/K}$  and then multiplying it by  $1 - \frac{\epsilon}{2}$ . Hence we have that,  $(1 - \frac{\epsilon}{2})s \leq \phi(s)$  and  $\phi(s) < s \cdot (1 - \frac{\epsilon}{2}) \cdot H^{1/K} < (1 - \frac{\epsilon}{2})(1 + \frac{\epsilon^2}{4}) \cdot s$  for any price *s*. We then have that for any two prices *s*, *s*':

$$\phi(s) - \phi(s') \le \left(1 - \frac{\epsilon}{2}\right) \left(1 + \frac{\epsilon^2}{4}\right) s - \left(1 - \frac{\epsilon}{2}\right) s' = s - s' - \left(\frac{\epsilon}{2}\right) \left(\left(1 - \frac{\epsilon}{2} + \frac{\epsilon^2}{4}\right) s - s'\right) < s - s', \quad (1)$$

When the last inequality holds whenever  $\left(1 - \frac{\epsilon}{2} + \frac{\epsilon^2}{4}\right)s - s' > 0$ .

To complete the proof we show that for any valuation v, if the profit-maximizing set in  $\mathcal{M}_1$  was A and has generated revenue of s, then the profit-maximizing set A' in  $\mathcal{M}_2$  generates revenue of at least  $s \cdot (1 - \epsilon)$ . Denote by s' the price of A' in  $\mathcal{M}_1$ . Since A is selected in  $\mathcal{M}_1$ , we have that  $v(A) - s \ge v(A') - s'$ . A necessary condition for a buyer to select A' in  $\mathcal{M}_2$  is that  $v(A) - \phi(s) \le v(A') - \phi(s')$ . Combining the two we get that a necessary condition to select A' in  $\mathcal{M}_2$  is that  $\phi(s) - \phi(s') \ge s - s'$ . For this not to contradict Equation (1) it must holds that  $\left(1 - \frac{\epsilon}{2} + \frac{\epsilon^2}{4}\right)s \le s'$ . From this inequality we derive that for A' that is picked in  $\mathcal{M}_2$  the payment is  $\phi(s') > (1 - \frac{\epsilon}{2})s' > \left(1 - \frac{\epsilon}{2} + \frac{\epsilon^2}{4}\right)(1 - \frac{\epsilon}{2})s > (1 - \epsilon)s$ , and we conclude that  $\text{REV}_{\mathcal{M}_2}(v) \ge (1 - \epsilon)\text{REV}_{\mathcal{M}_1}(v)$ . In addition, we have that for H > 2 and  $0 < \epsilon < 1$ , the parameter K is smaller than  $\frac{7}{\epsilon^2} \ln H$  which implies that  $H^{1/K} > H^{(\frac{7}{\epsilon^2} \ln H)^{-1}} = e^{\frac{\epsilon^2}{7}}$  and this is the minimal ratio between two different prices in  $\mathcal{M}_2$ .

Recall that by Theorem 3.2, for any weighted matroid-rank valuation v over a set of size m, there exists a randomized algorithm that finds a k-optimal set and has in expectation makes  $O(\log^3 m)$  demand queries. Combining this with the lemma above we get as an immediate corollary that for any bundle-size pricing menu  $\mathcal{M}_1$ , there exists a bundle-size pricing menu  $\mathcal{M}_2$  such that for any weighted matroid-rank valuation v it holds that  $\operatorname{REV}_{\mathcal{M}_2}(v) \ge (1 - \epsilon)\operatorname{REV}_{\mathcal{M}_1}(v)$ , and such that the buyer can find a profit-maximizing set with  $O(\log^3 m \cdot \epsilon^{-2} \log H_{\mathcal{M}})$  demand queries in expectation.

The number of demand queries in above result depends on the price-ratio  $H_M$  being not too large. Our main result in this section is that we can get rid of the dependence on  $H_M$  when optimizing the *expected* revenue for a given distribution over valuations (rather than ex-post, for any given valuation).

Before proving our main result, we prove a lemma showing that there are cases in which many entries of a menu can be removed without harming the *expected* revenue by much.

LEMMA B.2. Given a bundle-size pricing menu  $\mathcal{M}_1$  over m items, a distribution  $\mathcal{D}$  over valuations, and  $\epsilon > 0$ , let d be smallest bundle size in  $\mathcal{M}_1$  such that the expected revenue from selling bundles of size d is at least  $\frac{\epsilon}{m} \text{REV}(\mathcal{M}_1, \mathcal{D})$ . Let  $\mathcal{M}_2$  be the menu obtained from  $\mathcal{M}_1$  by removing all entries of bundles smaller then d. Then,  $\text{REV}(\mathcal{M}_2, \mathcal{D}) \ge (1 - \epsilon) \text{REV}(\mathcal{M}_1, \mathcal{D})$ .

PROOF. For any valuation in which a bundle of size  $d' \ge d$  was selected in  $\mathcal{M}_1$ , it was picked although the buyer had the option to buy a smaller and cheaper bundle. Removing an option that was not picked will not change the selection and hence, even if every cheaper bundle is removed, bundle d' will also selected in  $\mathcal{M}_2$ . That is, removing bundles of size k smaller than d from the menu will result in revenue loss that is bounded by the revenue contribution of valuations that selected those bundles of size k in  $\mathcal{M}_1$ , and that loss is at most  $\frac{\epsilon}{m} \operatorname{Rev}(\mathcal{M}_1, \mathcal{D})$ . Since at most m bundles are removed,  $\operatorname{Rev}(\mathcal{M}_2, \mathcal{D}) \ge (1 - \epsilon) \operatorname{Rev}(\mathcal{M}_1, \mathcal{D})$ .

We now prove our main result in this section by using the above two lemmas and applying our main upper bound of Theorem 3.2 (which shows that for weighted matroid-rank valuations we can find a k-optimal set in poly-logarithmic number of demand queries).

We are now ready to prove Theorem 4.1.

OF THEOREM 4.1. Denote  $H = H_{\mathcal{M}_1}$ . Given  $\mathcal{M}_1$  we use Lemma B.1 to construct a bundle-size pricing menu  $\mathcal{M}'_2$  with  $O(\epsilon^{-2} \log H)$  price levels such that  $\operatorname{Rev}(\mathcal{M}'_2, \mathcal{D}) \ge (1 - \epsilon)\operatorname{Rev}(\mathcal{M}_1, \mathcal{D})$ . For  $H \le m$ , we have that the number of price levels in  $\mathcal{M}'_2$  is  $O(\epsilon^{-2} \log m)$ , and by solving for each price level separately, we can find a profit-maximizing set with  $O(\log^4 m \cdot \epsilon^{-2})$  demand queries in expectation.

Else, we have that H > m and we now modify  $\mathcal{M}'_2$  using the information that valuations are drawn from  $\mathcal{D}$ . We use Lemma B.2 to construct a bundle-size pricing menu  $\mathcal{M}''_2$ , by removing entries from  $\mathcal{M}'_2$  that have low expected revenue contribution, and get a menu such that  $\operatorname{Rev}(\mathcal{M}''_2, \mathcal{D}) \ge$  $(1 - \epsilon)\operatorname{Rev}(\mathcal{M}'_2, \mathcal{D})$ . Note that since  $\mathcal{M}''_2$  is obtained by deleting entries, the ratio between prices is still at least  $e^{7/\epsilon^2}$ . Let *p* the price of the cheapest bundle in  $\mathcal{M}''_2$ . By the definition of  $\mathcal{M}''_2$  we have that  $p \ge p \cdot \Pr[$ cheapest bundle selected by the buyer $] \ge \frac{\epsilon}{m} \operatorname{Rev}(\mathcal{M}''_2, \mathcal{D})$ .

Let  $t = \epsilon^{-1}m^3p$ . If  $\Pr_{v \sim \mathcal{D}}[v(M) \ge t] > m^{-2}$ , we define  $\mathcal{M}_2$  to be the menu that only sells the grand bundle for price t, having an expected revenue which is at least  $t \cdot m^{-2} = mp/\epsilon \ge \operatorname{REV}(\mathcal{M}''_2, \mathcal{D}) > (1 - 2\epsilon)\operatorname{REV}(\mathcal{M}_1, \mathcal{D})$ . This menu is implementable using a single value query.

Else  $\Pr_{v \sim \mathcal{D}} [v(M) \ge t] \le m^{-2}$ . In this case we define  $\mathcal{M}_2$  to be  $\mathcal{M}_2''$  and we implement it using the following algorithm: We start by querying the value of the grand bundle v(M). We then consider two cases.

In the first case  $v(M) \ge t$ . In this case we use the fact that for matroid-rank valuations the greedy algorithm finds a *k*-optimal set for every  $k \in [m]$ : at each point selecting the item with the highest marginal value with respect to the set selected so far. The greedy algorithm requires  $O(m^2)$  value queries. Since the probability of running this algorithm is at most  $m^{-2}$ , and when running the number of queries is  $O(m^2)$ , this case contributes only a constant to the expected query complexity, where the expectation is taken over the distribution  $\mathcal{D}$ .

Else, we are in the second case in which  $v(M) \le t = e^{-1}m^3p$ . By monotonicity no bundle is sold for a price higher than t in this realization of v. Since p is the cheapest price in the menu and  $t = e^{-1}m^3p \ge v(M)$ , the ratio between the highest and lowest prices of sold bundles is at most  $e^{-1}m^3$ . As H > 2 (which is the case since H > m), by Lemma B.1 the ratio between any two different prices in the menu is at least  $e^{e^2/7}$ . Thus, there are at most  $21e^{-3} \cdot \log m$  bundle sizes in the range between p and  $e^{-1}m^3p$  that are needed to be considered. For any weighted matroid-rank valuations we can find a profit-maximizing set for each bundle size separately using  $O(\log^3 m)$  demand queries in expectation, by applying our main upper bound presented in Theorem 3.2. Among those bundles, we

return the one with the highest profit when paying its price. In total, the expected number of queries is then  $O(\epsilon^{-3}\log^4 m)$  demand queries, where the expectation is taken over the internal random coins of the algorithm and the distribution  $\mathcal{D}$ .

In both of the possible menus,  $\mathcal{M}_2''$  and the grand bundle, the expected revenue is at least (1 - $2\epsilon$ )REV( $\mathcal{M}_1, \mathcal{D}$ ). Setting  $\epsilon' = \epsilon/2$  in the theorem statement completes the proof. 

As an immediate corollary of the theorem we get that it is possible to take the optimal bundle-size pricing revenue on  $\mathcal{D}$ , and while losing at most  $\epsilon$  fraction of the revenue, convert it to another bundle-size pricing menu that has a primitive complexity of  $poly(\log m, \frac{1}{c})$ .

#### LOWER BOUNDS С

#### C.1 Submodular Valuation (Proof of Theorem 5.1)

In this section we prove Theorem 5.1. Some definitions and auxiliary claims are presented in Subsection C.1.1. We then prove Lemma 5.2 in Section C.1.2.

C.1.1 Definitions and Auxiliary Claims. We first show that every valuation v in the support of  $\mathcal{D}$ is indeed submodular. We then present some definitions and prove several claims that will be helpful in the proof of the theorem.

LEMMA C.1. Every valuation v in the support of  $\mathcal{D}$  is submodular.

**PROOF.** It is sufficient to show that for any  $S \subseteq M$  and  $a, b \in M \setminus S$  it holds that  $v(S \cup \{a\}) + v(S \cup \{a\})$  $\{b\} \ge v(S \cup \{a, b\}) + v(S)$ . We split into cases and verify the inequality holds for each case:

- |S| > k + 1:  $v(S \cup \{a\}) + v(S \cup \{b\}) = 2k = v(S \cup \{a, b\}) + v(S)$ .
- $|S| = k + 1, S \in \mathcal{B}_{v}: v(S \cup \{a\}) + v(S \cup \{b\}) = 2k = v(S \cup \{a, b\}) + v(S).$
- $|S| = k + 1, S \notin \mathcal{B}_v: v(S \cup \{a\}) + v(S \cup \{b\}) = 2k > 2k \frac{3}{11} = v(S \cup \{a, b\}) + v(S).$
- $|S| = k, S = G_v: v(S \cup \{a\}) + v(S \cup \{b\}) = 2k \frac{6}{11} = v(S \cup \{a, b\}) + v(S).$   $|S| = k, S \neq G_v: v(S \cup \{a\}) + v(S \cup \{b\}) \ge 2k \frac{6}{11} > 2k \frac{7}{11} = v(S \cup \{a, b\}) + v(S).$   $|S| = k 1, S \notin \mathcal{R}_v: v(S \cup \{a\}) + v(S \cup \{b\}) \ge 2k \frac{14}{11} \ge v(S \cup \{a, b\}) + v(S).$   $|S| = k 1, S \notin \mathcal{R}_v: v(S \cup \{a\}) + v(S \cup \{b\}) \ge 2k \frac{14}{11} \ge v(S \cup \{a, b\}) + v(S).$

- |S| < k 1:  $v(S \cup \{a\}) + v(S \cup \{b\}) \ge 2k \frac{28}{14} \ge v(S \cup \{a, b\}) + v(S)$ .

A valuation in the support of  $\mathcal{D}$  is completely defined by the values of all sets of size k and k + 1. We say that a set of value queries Q is in a *canonical form* if all queries in Q are for sets of size k or k + 1, and for every query of size k all of its supersets of size k + 1 are also in Q. Essentially, all information that a set of value queries conveys about a valuation can also be conveyed by some set of queries that is in a canonical form and is not much larger. The next proposition shows that we can assume that the query set is in a canonical form at a cost of a polynomial blow-up in the number of queries:

PROPOSITION C.2. Let A' be an algorithm that makes t value queries on a valuation in the support of  $\mathcal{D}$ . Then, there is an algorithm A that simulates A' while making  $m^2 \cdot t$  value queries on a valuation in the support of  $\mathcal{D}$ . Moreover, the set of queries that A makes has a canonical form.

**PROOF.** We show how instead of querying directly a set S we can compute its value by  $m^2$  value queries to bundles of size k and k + 1. We split into cases:

- |S| < k 1 or |S| > k + 1: the value of S is known and the query is discarded.
- |S| = k + 1: the query remains the same.

- |S| = k 1: the query is replaced with value queries to all  $\frac{(m-k)(m-k-1)}{2}$  sets of size k + 1 that contain S. If the value of all these sets of size k + 1 is k then  $S \in \mathcal{R}_v$  and its value is k 1, otherwise  $S \notin \mathcal{R}_v$  and its value is  $k \frac{14}{11}$ .
- |S| = k: the query remains the same and we additionally make value queries to all (m k) sets of size k + 1 containing S.

We next present several useful definitions and notations. Fix some deterministic algorithm A that makes only value queries and runs on valuations from  $\mathcal{D}$ . Fix any valuation v from the support of  $\mathcal{D}$ , and let  $Q_v$  denote the list of t bundles that A queried together with their values. Let  $\mathcal{D}_{Q_v}$  denote the distribution over valuations that is obtained by sampling according to  $\mathcal{D}$  a valuation that is consistent with the queries in  $Q_v$ . Let  $\mathcal{B}_{Q_v}^y$  be the family of sets that includes every set S such that  $\Pr_{v'\sim \mathcal{D}_{Q_v}}[S \in \mathcal{B}_{v'}] = 1$ . Similarly, let  $\mathcal{B}_{Q_v}^n$  be the family of sets that includes every set S such that  $\Pr_{v'\sim \mathcal{D}_{Q_v}}[S \in \mathcal{B}_{v'}] = 0$ . Let  $\mathcal{K}_{Q_v}$  be the family of sets of size k that were queried in  $Q_v$ .

We now claim that assuming queries are in a canonical form, the conditional distribution for sets not queried is essentially identical to the prior.

LEMMA C.3. Fix any valuation v sampled from D and assume  $Q_v$  is in a canonical form. It holds that:

- For any set S of size k + 1 it holds that  $\Pr_{v' \sim \mathcal{D}_{Q_v}}[S \in \mathcal{B}_{v'}] \in \{0, 1, \frac{1}{m^2}\}$ .
- The conditional probabilities are independent: for any family  $\mathcal{F}$  of sets of size k + 1 it holds that  $\Pr_{v' \sim \mathcal{D}_{Q_n}}[\forall S \in \mathcal{F}, S \notin \mathcal{B}_{v'}] = \prod_{S \in \mathcal{F}} \Pr_{v' \sim \mathcal{D}_{Q_n}}[S \notin \mathcal{B}_{v'}].$

PROOF. We start with proving the first part. First, if  $S \in \mathcal{B}_{Q_v}^n$  or  $S \in \mathcal{B}_{Q_v}^y$  then by definition  $S \in \mathcal{B}_{v'}$  with probability 0 or 1, respectively. Thus, the first part of the claim holds for all sets of size k + 1 that were queried in  $Q_v$ . Next, suppose that S was not queried in  $Q_v$ . Since  $Q_v$  is in a canonical form, if T is a set of size k that was queried, all of its supersets of size k + 1 were queried. Thus, as S was not queried, none of its subsets was queried.

Recall that the sampling process that defines  $\mathcal{B}_v$  picks every set of size k + 1 to belong to  $\mathcal{B}_v$  independently with probability  $\frac{1}{m^2}$ . The posterior probability given  $Q_v$  is still  $\frac{1}{m^2}$ , as none of sets of size k that were queried is a subset of S, and conditional on this event the value of S is sampled independently with probability  $\frac{1}{m^2}$ . Thus, by the principle of deferred decisions we can think of the membership of S in  $\mathcal{B}_v$  as determined after the values  $Q_v$  are given, and thus:

$$\Pr_{v'\sim\mathcal{D}_{Q_v}}[S\in\mathcal{B}_{v'}] = \Pr_{v'\sim\mathcal{D}}[S\in\mathcal{B}_{v'}|Q_v=Q_{v'}] = \frac{\Pr_{v'\sim\mathcal{D}}[S\in\mathcal{B}_{v'}\&Q_v=Q_{v'}]}{\Pr_{v'\sim\mathcal{D}}[Q_v=Q_{v'}]}$$

The claim follows since  $Q_v$  is in canonical form and as  $S \notin Q_v$  it holds that:

$$\Pr_{v'\sim\mathcal{D}}[S\in\mathcal{B}_{v'}\&Q_v=Q_{v'}]=\frac{1}{m^2}\cdot\Pr_{v'\sim\mathcal{D}}[Q_v=Q_{v'}]$$

For the second bullet, if  $\mathcal{F} \cap \mathcal{B}_{Q_v}^y$  is not empty, we have that  $\Pr_{v' \sim \mathcal{D}_{Q_v}}[\forall S \in \mathcal{F}, S \notin \mathcal{B}_{v'}] = 0 = \prod_{S \in \mathcal{F}} \Pr_{v' \sim \mathcal{D}_{Q_v}}[S \notin \mathcal{B}_{v'}]$ . Otherwise, since the membership of each set *S* in  $\mathcal{B}_{v'}$  is independent in  $\mathcal{D}$  for sets of size k + 1, and none of the sets of size k in  $Q_v$  is contained in any set in  $\mathcal{F} \setminus \mathcal{B}_{Q_v}^n$ , we again can similarly apply Bayes' rule and the priciple of deferred decisions and get that:

$$\Pr_{v'\sim\mathcal{D}_{Q_v}}[\forall S\in\mathcal{F},S\notin\mathcal{B}_{v'}] = \Pr_{v'\sim\mathcal{D}}[\forall S\notin\mathcal{F}\setminus\mathcal{B}_{Q_v}^n,S\in\mathcal{B}_{v'}] = (1-\frac{1}{m^2})^{|\mathcal{F}\setminus\mathcal{B}_{Q_v}^n|} = \prod_{S\in\mathcal{F}}\Pr_{v'\sim\mathcal{D}_{Q_v}}[S\notin\mathcal{B}_{v'}]$$

*C.1.2* An Impossibility for Value Queries. We next prove Lemma 5.2, our lower bound for value queries.

PROOF OF LEMMA 5.2. Fix any valuation v in the support of  $\mathcal{D}$ , and let  $Q_v$  denote the list of t bundles that A queried together with their values and assume that  $Q_v$  has a canonical form.

We now analyze the conditional distribution (given  $Q_v$ ) that a specific set S (of size k) is k-optimal, showing that this probability is exponentially small even after the algorithm makes its t queries, as long as t is not huge. Thus with high probability the algorithm cannot determine which bundle is the good set. We start with showing that by "ignoring" queries to bundles of size k. Then, we will show that the effect of such queries is small.

Let  $\mathcal{D}^{\mathcal{B}^y_{Q_o}, \mathcal{B}^y_{Q_o}}$  be the distribution over valuations obtained by sampling from  $\mathcal{D}$  a valuation that agrees on  $\mathcal{B}^y_{Q_o}$  and  $\mathcal{B}^n_{Q_o}$ .

CLAIM C.4. For any set S, 
$$|S| = k$$
,  $\Pr_{v' \sim \mathcal{D}^{\mathcal{B}^y_{Q_v}, \mathcal{B}^n_{Q_v}}}[S = G_{v'}] < \frac{2}{\binom{m}{k} - mt}$ .

PROOF. Observe that a S of size k that is a subset of some set  $S' \in \mathcal{B}_{Q_v}^y$  cannot be  $G_{v'}$  for any  $v' \in supp(\mathcal{D}^{\mathcal{B}_{Q_v}^y, \mathcal{B}_{Q_v}^n})$ . Each set of size k + 1 contains k + 1 sets of size k. Thus there are at most  $(k + 1) \cdot t \leq m \cdot t$  sets of size k for which  $\mathcal{B}_{Q_v}^y$  dictates that they cannot be  $G_{v'}$  for any  $v' \in supp(\mathcal{D}^{\mathcal{B}_{Q_v}^y, \mathcal{B}_{Q_v}^n})$ .

Each of the remaining sets is not contained in any set from  $\mathcal{B}_{Q_v}^y$  and is contained in m - k other sets of size k + 1. There are at least  $\binom{m}{k} - m \cdot t$  such sets. Next, for every two sets  $S_1$  and  $S_2$  that are not contained in any set from  $\mathcal{B}_{Q_v}^y$ , we compare the probabilities of the events  $S_1 = G_{v'}$  and of  $S_2 = G_{v'}$  in a valuation v' sampled from  $\mathcal{D}_{Q_v}^{\mathcal{B}_{Q_v}^y, \mathcal{B}_{Q_v}^n}$ .

For a valuation v' let  $\mathcal{F}_{v'}$  be the family of sets of size k who are not contained in any of the sets from  $\mathcal{B}_{v'}$ . Recall that all sets of size k that are in  $\mathcal{F}_{v'}$  have the same ex-ante probability of being  $G_{v'}$ .

Consider some set *S* of size *k*. The probability that  $S = G_{v'}$  when the valuation v' is sampled from  $\mathcal{D}^{\mathcal{B}_{Q_v}^y, \mathcal{B}_{Q_v}^n}$  is 0 whenever  $\mathcal{B}_{Q_v}^y$  contains a superset of *S*. Else, none of the supersets of *S* is in  $\mathcal{B}_{Q_v}^y$ , and by Lemma 5.6, if the number of its supersets that are in  $\mathcal{B}_{Q_v}^n$  is  $\ell$  then  $\Pr_{v' \sim \mathcal{D}^{\mathcal{B}_{Q_v}^y, \mathcal{B}_{Q_v}^n}}[S = G_{v'}] =$ 

$$\Pr_{\substack{v' \sim \mathcal{D}^{\mathcal{B}_{Q_v}^{y}, \mathcal{B}_{Q_v}^{n}}} \left[S \in \mathcal{F}_{v'}\right] \cdot \frac{1}{|\mathcal{F}_{v'}|} = \left(1 - \frac{1}{m^2}\right)^{m-k-\ell} \cdot \frac{1}{|\mathcal{F}_{v'}|}$$

For two sets  $S_1$  and  $S_2$  that are not contained in any set from  $\mathcal{B}_{Q_v}^y$ , let  $\ell_i$  be the number of sets in  $\mathcal{B}_{Q_v}^n$  that contain  $S_i$ . We have that:

$$\frac{\Pr_{\substack{v' \sim \mathcal{D}}} \mathcal{B}_{Q_{v}}^{y} \mathcal{B}_{Q_{v}}^{n}}{\Pr_{v' \sim \mathcal{D}} \mathcal{B}_{Q_{v}}^{y} \mathcal{B}_{Q_{v}}^{n}} \left[S_{1} = G_{v'}\right]}{\left(1 - \frac{1}{m^{2}}\right)^{m-k-\ell_{1}}} = \frac{1}{\left(1 - \frac{1}{m^{2}}\right)^{(\ell_{1}-\ell_{2})}} \leq 2$$

where the last inequality uses the fact that since  $0 \le \ell_1, \ell_2 \le m - k \le m$ , we have that  $\ell_1 - \ell_2 < m$  and thus the ratio is at most  $\frac{1}{\left(1 - \frac{1}{m^2}\right)^m}$ , which approaches to 1 as *m* goes to infinity.

Assume by contradiction that for some *S* we have that  $\Pr_{v' \sim \mathcal{D}^{\mathcal{B}_{Qv}^{y}, \mathcal{B}_{Qv}^{n}}}[S = G_{v'}] > \frac{2}{\binom{m}{k} - m \cdot t}$ , as the ratio of probabilities for any two sets  $S_1, S_2$  to be the *k*-optimal set is bounded by 2, the minimal probability for each of the sets of size *k* (of positive probability) to be the *k*-optimal set is at least  $\frac{1}{\binom{m}{k} - m \cdot t}$ . As there are at least  $\binom{m}{k} - m \cdot t$  sets with positive probability, summing the probabilities over all of them exceeds 1, a contradiction.

We have shown that when conditioning on  $\mathcal{B}_{Q_v}^y$  and  $\mathcal{B}_{Q_v}^n$  the probability of a set being the good set is low. To fully condition on  $Q_v$  we also need to condition on  $\mathcal{K}_{Q_v}$ . We show that even after conditioning on  $\mathcal{K}_{Q_v}$  the probability for finding  $G_v$  is still small.

CLAIM C.5. For  $|\mathcal{K}_{Q_v}| < 1.9^m$  and a large enough m, if for any set S',  $\Pr_{v' \sim \mathcal{D}^{\mathcal{B}_{Q_v}^y, \mathcal{B}_{Q_v}^n}}[S' = G_{v'}] < \frac{1}{1.95^m}$ , then for any set S, either  $S \in Q_v$  or  $\Pr_{v' \sim \mathcal{D}_{Q_v}}[S = G_{v'}] < \frac{2}{1.95^m}$ . Moreover,  $\Pr_{v' \sim \mathcal{D}^{\mathcal{B}_{Q_v}^y, \mathcal{B}_{Q_v}^n}}[S \in \mathcal{K}_{Q_v}] \le \frac{2|\mathcal{K}_{Q_v}|}{1.95^m}$ .

PROOF. We prove the claim by induction on  $r = |\mathcal{K}_{Q_v}|$ . For r = 0 the claim trivially holds. Assume the claim holds for any r' < r, we prove that it holds for r.

For a set  $S \in \mathcal{K}_{Q_v}$ , let  $Q^S = Q_v \setminus \{S\}$ . By the induction hypothesis, the probability that a set  $S \in Q_v$  is  $G_{v'}$  for v' sampled from  $\mathcal{D}$  conditioned on the values of sets in  $Q^S$ , is at most  $\frac{2}{1.95^m}$ .

By the induction hypothesis, since  $|Q^S| = r - 1$ , with probability at least  $1 - \frac{2(r-1)}{1.95^m}$ , none of the sets in  $Q^S$  is  $G_v$ . Together, with probability at least  $1 - \frac{2r}{1.95^m}$ , none of the sets in  $Q_v$  is  $G_v$ .

Fix any set  $T \notin \mathcal{K}_{Q_v}$  of size k. Let C be the event that all sets in  $\mathcal{K}_{Q_v}$  are not k-optimal (which also determines their exact values). By Bayes' theorem, conditioned on C we have that

$$\Pr_{v' \sim \mathcal{D}_{Q_v}}[T = G_{v'}|C] = \Pr_{v' \sim \mathcal{D}^{\mathcal{B}_{Q_v}^y, \mathcal{B}_{Q_v}^n}}[T = G_{v'}|C] = \frac{\Pr_{v' \sim \mathcal{D}^{\mathcal{B}_{Q_v}^y, \mathcal{B}_{Q_v}^n}}[T = G_{v'}]}{\Pr_{v' \sim \mathcal{D}^{\mathcal{B}_{Q_v}^y, \mathcal{B}_{Q_v}^n}}[C]} < \frac{\frac{1}{1.95^m}}{1 - \frac{2r}{1.95^m}} < \frac{2}{1.95^m}$$

where the last inequality holds for  $r < 1.9^m$  and a large enough *m*.

For  $t < 1.9^m$ ,  $k = \frac{m}{2}$  and a large enough *m*, by Claim C.4, we have that  $\Pr_{v' \sim \mathcal{D}_{Q_v}}[S = G_{v'}] < \frac{2}{\binom{m}{k} - mt} < \frac{2}{1.95^m}$  and the conditions of Claim C.5 hold.

Considering the set *S* returned by the algorithm as a t + 1'th query, the probability that either  $G_v \in Q_v$  or  $S = G_v$  is at most  $\frac{2(t+1)}{1.95^m}$  with probability over  $\mathcal{D}$ . For large enough *m*, we have that  $\frac{2(t+1)}{1.95^m} < \frac{t}{1.9^m}$ .