

Approximate Modularity Revisited*

Uriel Feige[†], Michal Feldman[‡], Inbal Talgam-Cohen[§]

Abstract

Set functions with convenient properties (such as submodularity) appear in application areas of current interest, such as algorithmic game theory, and allow for improved optimization algorithms. It is natural to ask (e.g., in the context of data driven optimization) how robust such properties are, and whether small deviations from them can be tolerated. We consider two such questions in the important special case of linear set functions.

One question that we address is whether any set function that approximately satisfies the modularity equation (linear functions satisfy the modularity equation exactly) is close to a linear function. The answer to this is positive (in a precise formal sense) as shown by Kalton and Roberts [1983] (and further improved by Bondarenko, Prymak, and Radchenko [2013]). We revisit their proof idea that is based on expander graphs, and provide significantly stronger upper bounds by combining it with new techniques. Furthermore, we provide improved lower bounds for this problem.

Another question that we address is that of how to learn a linear function h that is close to an approximately linear function f , while querying the value of f on only a small number of sets. We present a deterministic algorithm that makes only linearly many (in the number of items) nonadaptive queries, by this improving over a previous algorithm of Chierichetti, Das, Dasgupta and Kumar [2015] that is randomized and makes more than a quadratic number of queries. Our learning algorithm is based on the Hadamard transform.

1 Introduction

A set function f over a universe U of n items assigns a real value $f(S)$ to every subset $S \subseteq U$ (including the empty set \emptyset). Equivalently, it is a function whose domain is the Boolean n -dimensional cube $\{0, 1\}^n$, where each coordinate corresponds to an item, and a vector in $\{0, 1\}^n$ corresponds to the indicator vector of a set. Set functions appear in numerous applications, some of which are briefly mentioned in Section 1.2. Though set functions are defined over domains of size 2^n , one is often interested in optimizing over them in time polynomial in n . This offers several challenges, not least of which is the issue of representing f . An explicit representation of the truth table of f is of exponential size, and hence other representations are sought.

Some classes of set functions have convenient structure that leads to a polynomial size representation, from which the value of every set can easily be computed. A prime example for this is the class of linear functions. Formally, a set function f is *linear* if there exist constants c_0, c_1, \dots, c_n

*A conference version appeared in STOC'17. Part of this work was done at Microsoft Research, Herzliya. Part of the work of U. Feige was done while visiting Princeton University. This work has received funding from the Israel Science Foundation (grants No. 1388/16 and No. 336/18), the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement No. 337122, and the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 708935.

[†]Weizmann Institute of Science, Rehovot, Israel, uriel.feige@weizmann.ac.il.

[‡]Tel-Aviv University, Tel-Aviv, and Microsoft Research, Herzliya, Israel, michal.feldman@cs.tau.ac.il.

[§]Hebrew University, Jerusalem, Israel, inbaltalgam@gmail.com.

such that for every set S , $f(S) = c_0 + \sum_{i \in S} c_i$. The constants (c_0, c_1, \dots, c_n) may serve as a polynomial size representation of f . More generally, for set functions that arise naturally, one typically assumes that there is a so called *value oracle*, such that for every set S one can query the oracle on S and receive $f(S)$ in reply (either in unit time, or in polynomial time, depending on the context). The value oracle serves as an abstraction of either having some explicit polynomial time representation (e.g., a Boolean circuit) from which the value of f on any given input can be computed, or (in cases in which set functions model some physical reality) having a physical process of evaluating f on the set S (e.g., by making a measurement).

Optimizing over general set functions (e.g., finding the maximum, the minimum, maximizing subject to size constraints, etc.) is a difficult task, requiring exponentially many value queries if only a value oracle is given, and NP-hard if an explicit representation is given. However, for some special classes of set functions various optimization problems can be solved in polynomial time and with only polynomially many value queries. Notable nontrivial examples are minimization of submodular set functions [28, 16], and welfare maximization when the valuation function of each agent satisfies the *gross substitutes* property (see, e.g., [23]). For the class of linear set functions, many optimization problems of interest can be solved in polynomial time, often by trivial algorithms.

A major concern regarding the positive algorithmic results for some nice classes of set functions is their *stability*. Namely, if f is not a member of that nice class, but rather is only *close* to being a member (under some natural notion of *closeness*), is optimizing over f still easy? Can one obtain solutions that are “close” to optimal? Or are the algorithmic results “unstable” in the sense that a small divergence from the nice class leads to a dramatic deterioration in the performance of the associated algorithms?

A complicating factor is that there is more than one way of defining *closeness*. For example, when considering two functions, one may consider the variational distance between them, the mean square distance, the Hamming distance, and more. The situation becomes even more complicated when one wishes to define how close f is to a given *class* C of nice functions (rather than to a particular function). One natural definition is in terms of a distance to the function $g \in C$ closest to f . But other definitions make sense as well, especially if the class C is defined in terms of properties that functions in C have. For example, the distance from being submodular can be measured also by the extent to which the submodularity condition $f(S) + f(T) \geq f(S \cup T) + f(S \cap T)$ might be violated by f (as in [22, 20]), or even by the so called *supermodular degree* [12].

Following the lead of Chierichetti, Das, Dasgupta, and Kumar [2015], the goal of this work is to study questions such as the above in a setting that is relatively simple, yet important, namely, that of linear set functions. As illustrated by the work of [8], even this relatively simple setting is challenging.

1.1 Our results and techniques

Kalton constants One question that we address is the relation between two natural notions of being close to a linear function. The first notion is that of being close point-wise in an additive sense. We say that f is Δ -*linear* if there is a linear set function g such that $|f(S) - g(S)| \leq \Delta$ for every set S . Under this notion, the smaller Δ is, the closer we consider f to being linear. The other notion of closeness concerns a different but equivalent definition of linear functions, namely, as those functions that satisfy the *modular equation* $f(S) + f(T) = f(S \cup T) + f(S \cap T)$ for every two sets S and T . This form of defining linear functions is the key to generalizing linear functions to other classes of functions of interest, and specifically to submodular functions that satisfy $f(S) + f(T) \geq f(S \cup T) + f(S \cap T)$, and are considered to be the discrete analog of convex functions. Formally, we say that f is

- ϵ -modular if $|f(S) + f(T) - f(S \cup T) - f(S \cap T)| \leq \epsilon$ for every two sets S and T ;
- *weakly* ϵ -modular if the inequality holds for every two *disjoint* sets S and T .

The smaller ϵ is, the closer we consider f to being linear (or equivalently, to being modular).

It can easily be shown that every Δ -linear function is ϵ -modular for $\epsilon \leq 4\Delta$. Establishing a reverse implication is more difficult. In [8] it was shown that every ϵ -modular function is Δ -linear for $\Delta = O(\epsilon \log n)$. However, the authors of [8] were not aware of the earlier work of [18], which already showed that $\Delta \leq O(\epsilon)$ (this work was brought to our attention by Assaf Naor). We shall use K_s to denote the smallest constant such that every ϵ -modular function is $K_s\epsilon$ -linear, and refer to K_s as the *strong Kalton constant*. The bound provided in [18] was $K_s \leq \frac{89}{2}$, and this was subsequently improved by Bondarenko, Prymak, and Radchenko [6] to $K_s \leq 35.8$.

The approach initiated by Kalton and Roberts (and used almost as a blackbox in [6]) makes essential use of expander graphs in deriving upper bounds on K_s . We revisit this approach, simplify it and add to it new ingredients. Technically, the advantage that we get by our new ingredients is that we can use bipartite graphs in which only small sets of vertices expand, whereas previous work needed to use bipartite graphs in which large sets expand. This allows us to use expanders with much better parameters, leading to substantially improved upper bounds on K_s . For concreteness, we prove in this paper the following upper bound.

Theorem 1.1 (Upper bound, strong Kalton). *Every ϵ -modular function is Δ -linear for $\Delta < 12.62\epsilon$. Hence $K_s < 12.62$.*

We remark that our technique for upper bounding K_s adds a lot of versatility to the expander approach, which is not exploited to its limit in the current version of the paper: the upper bound that we report strikes a balance between simplicity of the proof and quality of the upper bound. Directions for further improvements are mentioned in Section 3.7.

Obtaining good lower bounds on K_s is also not easy. Part of the difficulty is that even if one comes up with a function f that is a candidate for a lower bound, verifying that it is ϵ -modular involves checking roughly 2^{2n} approximate modularity equations (one equation for every pair S and T of sets), making a computer assisted search for good lower bounds impractical. For set functions over 4 items, we could verify that the worst possible constant is $1/2$. Checking ϵ -modularity is much easier for symmetric functions (for which the value of a set depends only on its size), but we show that for such functions $K_s = \frac{1}{2}$, and this is also the case for ϵ -modular functions that are submodular (see Section 2.3). The only lower bound on K_s that we could find in previous work is $K_s \geq \frac{3}{4}$, implicit in [25]. We consider a class of functions that enjoys many symmetries (we call such functions (k, M) -*symmetric*), and for some function in this class we provide the following lower bound.

Theorem 1.2 (Lower bound, strong Kalton). *There is an integer valued set function over 70 items that is 2-modular and tightly 2-linear. Hence $K_s \geq 1$.*

In addition, we shall use K_w to denote the smallest constant such that every *weakly* ϵ -modular function is $K_w\epsilon$ -linear, and refer to K_w as the *weak Kalton constant*. For the weak Kalton constant we prove the following theorem.

Theorem 1.3 (Upper and lower bounds, weak Kalton). *Every weakly ϵ -modular function is Δ -linear for $\Delta < 24\epsilon$, so $K_w < 24$. Moreover, there is an integer valued set function over 20 items that is weakly 2-modular and 3-linear, so $K_w \geq 3/2$.*

Learning algorithms Another part of our work concerns the following setting. Suppose that one is given access to a value oracle for a function f that is Δ -linear. In such a setting, it is desirable to obtain an explicit representation for some linear function h that is close to f , because using such an h one can approximately solve optimization problems on f by solving them exactly on h . The process of learning h involves two complexity parameters, namely the number of queries made to the value oracle for f , and the computation time of the learning procedure, and its quality can be measured by the distance $d(h, f) = \max_S \{|h(S) - f(S)|\}$ of h from f . It was shown in [8] that for every learning algorithm that makes only polynomially many value queries, there will be cases in which $d(h, f) \geq \Omega(\Delta\sqrt{n}/\log n)$. This might sound discouraging but it need not be. One reason is that in many cases $\Delta < \log n/\sqrt{n}$. For example, this may be the case when f itself is actually linear, but is modeled as Δ -linear due to noise in measurements of values of f . By investing more in the measurements, Δ can be decreased. Another reason is that the lower bounds are for worst case f , and it is still of interest to design natural learning algorithms that might work well in practice. Indeed, [8] designed a randomized learning algorithm that makes $O(n^2 \log n)$ nonadaptive queries and learns h within distance $O(\Delta\sqrt{n})$ from f . We improve upon this result in two respects, one being reducing the number of queries, and the other being removing the need for randomization.

Theorem 1.4 (Learning algorithm). *There is a deterministic polynomial time learning algorithm that given value oracle access to a Δ -linear function on n items, makes $O(n)$ nonadaptive value queries and outputs a linear function h , such that $h(S)$ is $O(\Delta(1 + \sqrt{\min\{|S|, n - |S|\}}))$ -close to $f(S)$ for every set S .*

The learning algorithm is based on the Hadamard basis. A similar technique is applied in [9] for a different purpose (as brought to our attention by Moni Naor) – attacking the privacy of an n -sized binary database, and recovering a binary vector whose Hamming distance from the true database is $o(n)$.

The Hadamard basis is an orthogonal basis of \mathbb{R}^n consisting of vectors with ± 1 entries, which can be constructed for every n that is a power of 2. It has the following property: for every two linear functions that are $O(\Delta)$ -close on the basis vectors (where each basis vector can be interpreted as a difference between two sets, one corresponding to its $+1$ indices, the other to its -1 indices), the linear functions are $O(\Delta(1 + \sqrt{\min\{|S|, n - |S|\}}))$ -close on *every* set S . (This property follows from the vectors having large norms, and thus a large normalization factor by which the distance of Δ is divided.) Given $O(n)$ value queries to a Δ -linear function f , an algorithm can learn the values of the linear function g that is Δ -close to f for the n Hadamard basis vectors, up to an additive error of $O(\Delta)$. This is enough information to construct a linear function h that is $O(\Delta(1 + \sqrt{\min\{|S|, n - |S|\}}))$ -close to $g(S)$, and thus to $f(S)$, for every set S .

1.2 Additional Related Work

The stability of linearity and the connection between approximate modularity and approximate linearity, which we study on the discrete hypercube, have been extensively studied in mathematics on continuous domains (e.g., Banach spaces). An early example is [15], and this result together with its extensions to other classes of functions are known as the Hyers-Ulam-Rassias theory [17].

Our work is related to the literature on *data-driven optimization* (see, e.g., [4, 29, 14, 2]). This literature studies scenarios in which one wishes to optimize some objective function, whose parameters are derived from real-world data and so can be only approximately evaluated. Such scenarios arise in applications like machine learning [21], sublinear algorithms and property testing [5], and algorithmic game theory [22]. This motivates the study of optimization of functions that satisfy some properties only approximately (e.g., submodularity [20], gross substitutes [27], or

convexity [3]). Our work is related to this strand of works in that we are also interested in functions that satisfy some property (modularity in our case) approximately, but unlike the aforementioned works, we are interested in the characterization and learning of functions, not in their optimization.

Learning of (exactly) submodular functions was studied by [1] and [13], in the PMAC model and general query model, respectively. We shall discuss learning of approximately modular functions in a value query model.

Organization In Section 2 we present preliminaries. Section 3 shows improved upper bounds on the weak and strong Kalton constants, and Section 4 shows improved lower bounds on the weak and strong Kalton constants. Section 5 describes how to learn an approximately linear function from an approximately modular one.

2 Preliminaries

2.1 Approximate Modularity and Linearity

Let $U = \{1, \dots, n\}$ be a ground set of $n \geq 2$ items (also called elements). For every set $S \subseteq U$ of items, let \bar{S} denote its complement $U \setminus S$. A *collection* \mathcal{G} is a multiset of sets, and its complement $\bar{\mathcal{G}}$ is the collection of complements of the sets in \mathcal{G} . Given a set function $f : 2^U \rightarrow \mathbb{R}$, the *value* of set S is $f(S)$. Throughout we focus on additive closeness, and say that two values $x, y \in \mathbb{R}$ are Δ -close (or equivalently, that x is Δ -close to y) if $|x - y| \leq \Delta$.

A set function f is ϵ -modular if for every two sets $S, T \subseteq U$,

$$|f(S) + f(T) - f(S \cup T) - f(S \cap T)| \leq \epsilon. \quad (1)$$

If $\epsilon = 0$ then f is *modular*. A set function f is *weakly ϵ -modular* if Condition (1) holds for every two *disjoint* sets $S, T \subseteq U$. The following proposition shows the relation:

Proposition 2.1. *Every weakly ϵ -modular set function is a 2ϵ -modular set function.*

Proof. Let f be a weakly ϵ -modular set function, we show it must be 2ϵ -modular. For every two sets S and T , since f is weakly ϵ -modular we have that $f(T) = f(S \cap T) + f(T \setminus S) - f(\emptyset) \pm \epsilon$, and that $f(S \cup T) \pm \epsilon = f(S) + f(T \setminus S) - f(\emptyset)$. Therefore, $f(S) + f(T) = f(S) + f(S \cap T) + f(T \setminus S) - f(\emptyset) \pm \epsilon = f(S \cup T) + f(S \cap T) \pm 2\epsilon$, completing the proof. ■

Observation 2.2 (Approximate modularity for ≥ 2 sets). *Let f be a weakly ϵ -modular set function, and let $S \subseteq U$ be a set of items with a partition (S_1, \dots, S_s) (i.e., the disjoint union $S_1 \uplus \dots \uplus S_s$ is equal to S). By iterative applications of weak ϵ -modularity,*

$$f(S) + (s - 1)(f(\emptyset) - \epsilon) \leq f(S_1) + \dots + f(S_s) \leq f(S) + (s - 1)(f(\emptyset) + \epsilon).$$

A set function f is *linear* if there exist constants c_0, c_1, \dots, c_n such that for every set S , $f(S) = c_0 + \sum_{i \in S} c_i$. A linear set function is *additive* if $c_0 = 0$. The *zero function* has $c_i = 0$ for every $0 \leq i \leq n$. A set function f is Δ -linear if there exists a linear set function g that is Δ -close to f , i.e., $f(S)$ and $g(S)$ are Δ -close for every set S . We say that a set function f is *tightly Δ -linear* if it is Δ -linear and for every $\Delta' < \Delta$, f is not Δ' -linear. A *closest linear function* to f is a linear set function g that is Δ -close to f where f is tightly Δ -linear.

Proposition 2.3 (Modularity = linearity). *A set function is modular if and only if it is linear.*

Proof. Let f be a modular function, we show it must be linear: Let $c_0 = f(\emptyset)$, and let $c_i = f(\{i\}) - c_0$. For every set S , by modularity, $f(S) = f(S_1) + f(S_2) - c_0$ for any S_1, S_2 whose disjoint union is S . Applying this equality recursively gives $f(S) = \sum_{i \in S} c_i + |S|c_0 - |S-1|c_0 = \sum_{i \in S} c_i + c_0$, as required.

For the other direction, let f be a linear function such that $f(S) = c_0 + \sum_{i \in S} c_i$. For any two sets S, T , it holds that $f(S) + f(T) = 2c_0 + \sum_{i \in S} c_i + \sum_{i \in T} c_i$, and $f(S \cup T) + f(S \cap T) = 2c_0 + \sum_{i \in S \cup T} c_i + \sum_{i \in S \cap T} c_i$. The desired equality follows by $\sum_{i \in S} c_i + \sum_{i \in T} c_i = \sum_{i \in S \cup T} c_i + \sum_{i \in S \cap T} c_i$. ■

Proposition 2.4. *Every Δ -linear set function is 4Δ -modular. This result is tight, even for symmetric functions.*

Proof. Let f be a Δ -linear function, and let g be a linear function such that $g(S) = c_0 + \sum_{i \in S} c_i$ and $|f(S) - g(S)| \leq \Delta$ for every set S . For every two sets S and T , $f(S) + f(T) - f(S \cup T) - f(S \cap T) \leq (c_0 + \sum_{i \in S} c_i + \Delta) + (c_0 + \sum_{i \in T} c_i + \Delta) - (c_0 + \sum_{i \in S \cup T} c_i - \Delta) - (c_0 + \sum_{i \in S \cap T} c_i - \Delta) = 4\Delta$. Similarly, $f(S) + f(T) - f(S \cup T) - f(S \cap T) \geq (c_0 + \sum_{i \in S} c_i - \Delta) + (c_0 + \sum_{i \in T} c_i - \Delta) - (c_0 + \sum_{i \in S \cup T} c_i + \Delta) - (c_0 + \sum_{i \in S \cap T} c_i + \Delta) = -4\Delta$. This establishes that f is 4Δ -modular.

The proposition is tight, even for symmetric functions: Consider the 1-linear function on 4 items in which sets of size 0 and 4 are worth 0, sets of size 1 and 3 are worth -1 , and sets of size 2 are worth $+1$. If S and T are two different sets of size 2 that intersect, the modularity equation is violated by 4. ■

We shall often refer to set functions whose closest linear function is the zero function.

Observation 2.5 ([8]). *For every ϵ -modular (resp., weakly ϵ -modular) set function f that is tightly Δ -linear, there is an ϵ -modular (resp., weakly ϵ -modular) set function f' that is tightly Δ -linear and whose closest linear function is the zero function.*

The function f' can be defined as follows: if g is a closest linear function to f then $f'(S) = f(S) - g(S)$ for every set S .

2.2 Kalton Constants

Let K_w denote the smallest constant such that every weakly ϵ -modular set function is $K_w\epsilon$ -linear. Let K_s denote the smallest constant such that every ϵ -modular set function is $K_s\epsilon$ -linear. Notice that $K_s \leq K_w$. We refer to K_w as the weak Kalton constant (the possibility that there is such a constant K_w independent of n was advocated in the work of Nigel Kalton [19]), and to K_s as the strong Kalton constant. Formally:

Definition 2.6. $K_s \in \mathbb{R}_{\geq 0}$ (resp., $K_w \in \mathbb{R}_{\geq 0}$) is the strong (weak) Kalton constant if:

- for every $n \in \mathbb{N}_{\geq 0}, \epsilon \in \mathbb{R}_{\geq 0}$, every (weakly) ϵ -modular set function over $[n]$ is $K_s\epsilon$ -linear ($K_w\epsilon$ -linear); and
- for every $\kappa < K_s$ ($\kappa < K_w$) and for every $\epsilon \in \mathbb{R}_{\geq 0}$, there exists a sufficiently large n and a (weakly) ϵ -modular set function f over $[n]$ such that f is not $\kappa\epsilon$ -linear (it is sufficient that there exist a (weakly) ϵ -modular set function that is tightly $K_s\epsilon$ -linear ($K_w\epsilon$ -linear)).

The propositions in Section 2.1 imply the following corollaries on Kalton constants.

Corollary 2.7 (of Proposition 2.1). $K_w/2 \leq K_s \leq K_w$.

Corollary 2.8 (of Observation 2.5). *If K_s is the strong Kalton constant for set functions whose closest linear function is the zero function, then K_s is the strong Kalton constant (for general set functions). Similarly, if K_w is the weak Kalton constant for set functions whose closest linear function is the zero function, then K_w is the weak Kalton constant (for general set functions).*

2.3 Kalton Constants for Special Cases

The analysis of Kalton constants becomes much easier in the following special cases.

2.3.1 Kalton Constants for Symmetric Set Functions

Here we give tight bounds on K_s for symmetric set functions. A set function f is *symmetric* if for every two sets S and T , $|S| = |T| \implies f(S) = f(T)$. A symmetric set function over $[n]$ can be represented as a function over integers $f : \{0, \dots, n\} \rightarrow \mathbb{R}$. The following example shows that $K_s \geq 1/2$ for symmetric set functions, and Proposition 2.11 shows this is tight.

Example 2.9. *For every $n \in \mathbb{N}_{\geq 0}$, consider the symmetric set function f_n over $[n]$ where $f_n([n]) = -\epsilon$ and $f_n(S) = 0$ for every other set S . Let g_δ be a symmetric linear set function over $[n]$ such that $g_\delta(\{j\}) = -\frac{\epsilon j}{n} + \delta$ for every $0 \leq j \leq n$. Observe that g_δ is $\max\{\delta, \epsilon - \frac{\epsilon}{n} - \delta\}$ -close to f_n . When $\delta = \frac{\epsilon}{2} - \frac{\epsilon}{2n}$, the distance is minimized and g_δ is a closest linear function to f_n . Thus f_n is tightly δ -linear. This shows that for every $\kappa < 1/2$, there exists n such that f_n is tightly δ -linear for $\delta = \frac{\epsilon}{2} - \frac{\epsilon}{2n} > \kappa\epsilon$.*

The next observation will be used in the proof of Proposition 2.11, which establishes tightness.

Observation 2.10. *For every symmetric ϵ -modular set function f that is tightly Δ -linear, there is a symmetric ϵ -modular set function f' that is tightly Δ -linear and whose closest linear function is the zero function.*

Proof. Let g be a linear set function Δ -close to f . We show a symmetric set function g' that is Δ -close to f : for every $k \in [n]$, let $g'(k)$ be the average value of k items according to g . The proof follows as in Observation 2.5. ■

Proposition 2.11. *The strong Kalton constant for symmetric set functions is $K_s = \frac{1}{2}$.*

Proof. Let f be a symmetric ϵ -modular set function, we argue that f must be $\frac{1}{2}\epsilon$ -linear. By Observation 2.10 we can assume without loss of generality that f 's closest linear function is the zero function. Let M be the maximum absolute value of f , then there exist $k_1 < k < k_2$ such that either $f(k_1) = f(k_2) = -M$ and $f(k) = M$, or $f(k_1) = f(k_2) = M$ and $f(k) = -M$. Without loss of generality assume the former. Suppose that $k \geq n/2$. Using ϵ -modularity we get that $2f(k) \leq f(k_2) + f(2k - k_2) + \epsilon \leq -M + M + \epsilon$. On the other hand, $2f(k) = 2M$. We get $2M \leq -M + M + \epsilon$, implying that $M \leq \epsilon/2$, as desired. If $k < n/2$, an analogous argument is invoked using k_1 . ■

2.3.2 Kalton Constants for Submodular Set Functions

We now give tight bounds on K_s for submodular set functions. A set function f is *submodular* if for every two sets S, T it holds that $f(S) + f(T) \geq f(S \cup T) + f(S \cap T)$. Example 2.9 shows that $K_s \geq 1/2$ not only for symmetric set functions but also for submodular ones.

Proposition 2.12. *The strong Kalton constant for submodular set functions is $K_s = \frac{1}{2}$.*

Proof. Without loss of generality we normalize f such that $f(\emptyset) = 0$. Since f is submodular and normalized, it belongs to the class of XOS functions,¹ and so there exists an additive set function g such that for every set S , $f(S) \geq g(S)$, and for the set $S = U$, $f(U) = g(U)$ (see, e.g., [10]). It remains to show that for every S , $f(S) \leq g(S) + \epsilon$. Assume for contradiction that $f(S) > g(S) + \epsilon$. Then $f(S) + f(\bar{S}) > g(S) + \epsilon + g(\bar{S}) = g(U) + \epsilon = f(U) + \epsilon$, where we use that $g(S) + g(\bar{S}) = g(U)$ by additivity. We have shown a contradiction to ϵ -modularity, completing the proof. ■

2.4 Results from the Literature

2.4.1 Upper and Lower Bounds

Kalton and Roberts [18] prove that $K_w \leq \frac{89}{2}$. This upper bound was subsequently improved to $K_w \leq 38.8$ by [6], who also show that $K_s \leq 35.8$. Let us provide more details on how the known upper bounds on K_w are achieved.

Definition 2.13 (Expander). *For $k \in \mathbb{N}_{\geq 0}$ and $\alpha, r, \theta \in \mathbb{R}_{\geq 0}$ such that $\alpha, \theta < 1$ and $r > 2$, we say that a bipartite graph $G_k(V, W; E)$ is an (α, r, θ) -expander if $|V| = 2k$, $|W| = 2\theta k$, $|E| = 2kr$, and every set $S \subset V$ of at most $2k\alpha$ vertices has at least $|S|$ neighbors in W (and hence a perfect matching into W).*

We say that (α, r, θ) -expanders exist if there is some $k' \in \mathbb{N}_{\geq 0}$ such that for every integer multiple k of k' , there exists an (α, r, θ) -expander G_k .

The following theorem (rephrased from [18]) is the key to the known upper bounds on K_w .

Theorem 2.14 ([18]). *Suppose that for fixed r and θ and all sufficiently large k there are $(\frac{1}{2}, r, \theta)$ -expanders $\{G_k\}$. Then an upper bound on the weak Kalton constant is:*

$$K_w \leq \frac{7 + 4r - 2\theta}{2(1 - \theta)}.$$

Pippenger [26] shows that $(\frac{1}{2}, r, \theta)$ -expanders exist with $r = 6$ and $\theta = \frac{2}{3}$, if k is sufficiently large ($k \geq 3$). This together with Theorem 2.14 implies that $K_w \leq \frac{89}{2}$. In [6] it is shown that r can be reduced to 5.05, thus leading to the improved bound of $K_w \leq 38.8$.

As for lower bounds, [25] gives a high-level sketch of a construction that shows the following for the weak Kalton constant. His construction implies also a lower bound for the strong Kalton constant, as follows (a detailed proof appears for completeness in Appendix A).

Theorem 2.15 (Lower bound [25]). *Lower bounds on the Kalton constants are $K_w \geq \frac{3}{2} - \Theta(1/n)$ and $K_s \geq \frac{3}{4} - \Theta(1/n)$.*

2.4.2 A Characterization of Chierichetti et al. [8]

For a set function f , let $M = \max_S \{|f(S)|\}$ be the maximum absolute value of f (also called f 's extreme value). We say that a set S has value M if $f(S) = M$, and value $-M$ if $f(S) = -M$. Given a distribution (p_1, \dots, p_κ) over sets S_1, \dots, S_κ , the *marginal probability* of item i according to this distribution is the probability that i appears in a set randomly selected according to the distribution, i.e., $\sum_{j|i \in S_j} p_j$.

¹A set function f belongs to the class XOS if there exists a collection of additive set functions $\{g_S\}$ such that for every set S , the corresponding additive function g_S satisfies $f(T) \geq g_S(T)$ for every T , and $f(S) = g_S(S)$.

Lemma 2.16 ([7, 8]). *The closest linear function to a set function f is the zero function if and only if there exist probability distributions P^+ and P^- with rational probabilities over sets with value M and sets with value $-M$, respectively, such that for every item i , the marginal probabilities of i according to P^+ and P^- are the same.*

Lemma 2.16 appears as Lemma 10 in [8]. It also appears as Proposition 3.2 in [7], but stated using different terminology. The rationality of the probability distributions follows since they are obtained as solutions to a linear program.

Definition 2.17 (Positive and Negative Supports). *For a set function f whose closest linear set function is the zero function, let P^+, P^- be the distributions guaranteed by Lemma 2.16. Then the positive support $\text{PS} = (P_1, \dots, P_\kappa)$ of f is the support of P^+ (sets assigned a non-zero probability by P^+), and the negative support $\text{NS} = (N_1, \dots, N_\nu)$ of f is the support of P^- .*

We emphasize that all sets in PS have value M and all sets in NS have value $-M$. A geometric interpretation of Lemma 2.16 is that the convex hull of PS , treated as a collection of vertices of the hypercube, and the convex hull of collection NS have a non-empty intersection.

3 Improved Upper Bounds

Our main goal in this section is to provide improved upper bounds for the strong Kalton constant K_s , which was previously known to be ≤ 35.8 [6]. Along the way we also provide an improved upper bound for the weak Kalton constant K_w , which was previously known to be ≤ 38.8 [6].

The section is organized as follows. In Section 3.1 we give an overview of the techniques. In Sections 3.2 and 3.3 we provide preliminaries and our main lemmas. Sections 3.4-3.5 are dedicated to K_w , and Sections 3.6-3.7 focus on K_s .

3.1 Roadmap

Let us first review the known upper bounds on K_w, K_s and how they are derived. For K_w , the basic approach of [18] is outlined in Theorem 2.14. There the value of K_w is related to parameters of $(\frac{1}{2}, r, \theta)$ -expanders. Rearranging the bound from Theorem 2.14, denoting the denominator by D and the numerator by $N_r + N_1$ where N_r depends on r and N_1 does not, it is shown that:

$$K_w \leq \frac{N_r + N_1}{D},$$

where $D = 1 - \theta$, $N_r = 2r$ and $N_1 = \frac{7}{2} - \theta$. Kalton and Roberts [18] use a previously known expander construction of [26] to get an upper bound of 44.5. The improved upper bound of 38.8 of [6] comes from constructions of expanders with a smaller value of r . As they explain, their techniques do not allow the value of r to be substantially reduced further (without changing θ), and so the approach of constructing better expanders is unlikely to significantly further reduce K_w . As for the strong Kalton constant K_s , the additional restrictions on f (ϵ -modularity instead of only weak ϵ -modularity) were used in [6] to reduce N_1 from $\frac{7}{2} - \theta$ to $\frac{5}{2} - \theta$, achieving an upper bound of 35.8.

3.1.1 Our approach to improving K_w and the role of expanders

In Section 3.4 we improve upon the upper bound of 38.8 on K_w by reducing N_1 from $\frac{7}{2} - \theta$ to $-\frac{1}{2} - \theta$, giving an upper bound of 26.8 using the expander construction of [6]. The path to this

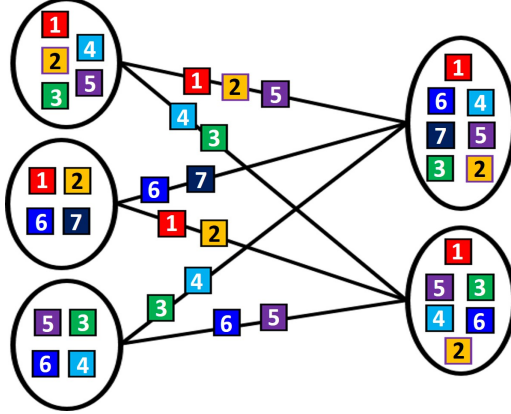


Figure 1: Split and merge rearrangement

improvement lies in extracting the main idea from the proof of [18], cleaning away redundancies and using Lemma 3.8 (which establishes the existence of complementary collections with values approximately equal to the function’s extreme value or its negation). It seems that the value of N_1 cannot be substantially reduced further within this framework of ideas, and hence new ideas appear to be needed if one wishes to obtain significant improvements in the upper bound on K_w . In Section 3.5 we upper-bound K_w using the minimum between two expressions rather than a bound of the form $(N_r + N_1)/D$, and show that $K_w \leq 23.82$.

We now provide a high-level view of the proof in Section 3.4, emphasizing in particular where expanders fit in. Let f be a weakly ϵ -modular set function whose closest linear set function is the zero function, and let M be the maximum absolute value of f . We need to show that M is at most 26.8ϵ . The proof is divided into 3 main steps:

1. Show that the existence of a certain “split and merge” rearrangement of a collection of approximately M -valued sets is sufficient to prove the theorem.
2. Establish that such a rearrangement is guaranteed for collections of sets in which items are not too frequent, given appropriate families of bipartite expander graphs.
3. Show existence of a collection of approximately M -valued sets in which items are not too frequent.

In what follows we expand upon the first two steps (shown formally in Lemmas 3.11 and 3.12). For the third step, see Lemma 3.8, which uses a characterization of set function whose closest linear set function is the zero function.

In Figure 1, we depict the split and merge rearrangement of Step 1. In the figure, $k = 3$ “source” sets (on the left) are split into rk intermediate sets (middle), where $r = 2$; then disjoint intermediate sets are merged back into θk “target” sets (on the right), where $\theta = \frac{2}{3}$. The smaller the parameters r and θ are, the more challenging it becomes to get existence of a split and merge rearrangement. For example, since the target sets must be the result of merging disjoint intermediate sets, if an item appears in the source sets more than θr times, then some target set will necessarily have two copies and so such a rearrangement does not exist. If a split and merge rearrangement does exist, then the intermediate sets can be utilized to relate the values of the source and target sets. The property of ϵ -modularity we use here is that if source set S is a disjoint union of r intermediate sets I_1, \dots, I_r , then $f(S) = \sum_{j=1}^r f(I_j) \pm (r - 1)$ (and the same holds for target set T). By applying this to source sets with values of approximately M we can get an upper bound on M .




Figure 2: Split and merge of sparse sets via expanders

In Figure 2 we illustrate Step 2 and its connection to expanders. As mentioned above, for a split and merge rearrangement to exist, no item should appear in too many source sets. Indeed, the depicted items 1, 4 and 6 are all only α -frequent for $\alpha = \frac{2}{3}$ (i.e., they appear in only $\frac{2}{3}$ of the source sets). Let's assume there exist bipartite expanders with α -expansion, k vertices on the left-hand side, rk edges and θk vertices on the right-hand side. Then every set of αk vertices on the left-hand side has a matching to the right-hand side. As depicted in Figure 2, the items can be mapped to target sets along these matchings, which assures that no target set will have multiple copies of any item – thus we have a valid split and merge rearrangement.

3.1.2 Our approach to improving K_s

Our approach to improving K_s makes more extensive use of ϵ -modularity. Rather than considering sets from a collection with approximately the extreme value M , we shall consider *intersections* of these sets. Using ϵ -modularity we shall be able to show that the function values of intersections are also close to M . In fact, using some averaging arguments we shall obtain even stronger estimates on how close these values are to M (a point that is relevant to controlling the value of N_1). Thereafter, we will no longer be restricted to using (α, r, θ) -expanders with $\alpha = \frac{1}{2}$. We will be able to use $\alpha = \frac{1}{4}$ instead (for intersections of two sets), or even $\alpha = \frac{1}{8}$ (for intersections of three sets), and so on. The advantage of reducing α is that for smaller values of α , expanders with lower values of r and θ exist, leading to better upper bounds on K_s . However, we cannot reduce α to arbitrarily small values because each reduction of α by a factor of two is accompanied by an increase in N_1 , and one needs to balance between these two factors.

3.2 Upper Bounds Preliminaries

We present definitions and preliminary results used to establish our upper bounds.

3.2.1 Families of expanders

As discussed when describing our approach (Section 3.1), we utilize the existence of expanders with a range of parameters. Our existence argument (see Appendix B) uses the probabilistic method as in [26], and so results in expanders that are biregular (all vertices on the same side of the bipartite graph have the same degree). More complicated expander constructions that are not biregular may achieve even better parameters, as in [6].

Lemma 3.1. *The following families of (biregular) expanders exist:*

1. $(\frac{1}{2}, 5, \frac{5}{7})$ -expanders.
2. $(\frac{3}{10}, 4, \frac{4}{7})$ -expanders.
3. $(\frac{1}{4}, 5, \frac{1}{2})$ -expanders.
4. $(\frac{1}{16}, 4, \frac{4}{15})$ -expanders.
5. $(\frac{1}{64}, 3, \frac{3}{11})$ -expanders.
6. $(\frac{1}{256}, 3, \frac{3}{19})$ -expanders.

3.2.2 Weakly ϵ -modular set functions

Throughout this section, let f be a weakly ϵ -modular set function whose closest linear set function is the zero function. Let M be the maximum absolute value of f .

Observation 3.2. $-\epsilon \leq f(\emptyset) + f(U) \leq \epsilon$.

Proof. Denote $\beta = f(\emptyset) + f(U)$. If $M = 0$ the claim follows trivially, otherwise by Lemma 2.16 there exist sets P, N with values $M, -M$, respectively. Consider the values of \bar{P}, \bar{N} . By weak ϵ -modularity and the definition of M , we have $-M \leq f(\bar{P}) \leq -f(P) + \beta + \epsilon = -M + \beta + \epsilon$, and $M + \beta - \epsilon = -f(N) + \beta - \epsilon \leq f(\bar{N}) \leq M$. We conclude that $0 \leq \beta + \epsilon$ and $\beta - \epsilon \leq 0$, completing the proof. \blacksquare

The following is a direct corollary of Observation 3.2 and weak ϵ -modularity:

Corollary 3.3 (Value of complement set). *Let $\beta = f(\emptyset) + f(U)$. Then for every set S , $-f(S) - \epsilon + \beta \leq f(\bar{S}) \leq -f(S) + \epsilon + \beta$.*

Definition 3.4. *A set S has deficit $d \geq 0$ if $f(S) = M - d$, and has surplus $s \geq 0$ if $f(S) = -M + s$. A collection has average deficit d (resp., surplus s) if the expected deficit (resp. surplus) of its sets with respect to the uniform distribution is d (resp. s).*

The next observation follows directly from Corollary 3.3.

Observation 3.5 (Average deficit/surplus of complement collection). *Let $\beta = f(\emptyset) + f(U)$. Let \mathcal{G} be a collection with average deficit d and surplus s . Then the average surplus of its complement $\bar{\mathcal{G}}$ is at most $d + \epsilon + \beta$, and the average deficit of $\bar{\mathcal{G}}$ is at most $s + \epsilon - \beta$.*

Definition 3.6. *An item is α -frequent in a collection if it appears in exactly an α -fraction of the sets. A collection has α -frequent items if every item is α -frequent in it.*

To give geometric intuition, we mention that a $1/2$ -frequent collection of sets corresponds to a collection of hypercube vertices (weighted if the same set appears in the collection multiple times), whose convex combination is simply the center of the cube.

Observation 3.7. *The complement of a collection with α -frequent items has $(1 - \alpha)$ -frequent items.*

Lemma 3.8 (Complementary collections). *There exists $k' \in \mathbb{N}_{\geq 0}$ such that for every k which is an integer multiple of k' , f has a collection PS^* of $2k$ sets (the same set might appear multiple times in the collection and we treat these appearances as distinct) with $1/2$ -frequent items and average deficit d , whose complement collection NS^* has $1/2$ -frequent items and average surplus s , and $d + s \leq \epsilon$.*

Proof. Consider the positive and negative supports PS, NS of f , where $\text{PS} = \{P_1, \dots, P_\kappa\}$ and $\text{NS} = \{N_1, \dots, N_\nu\}$, as defined in Definition 2.17. By Lemma 2.16, there exist distributions $P^+ = (p_1^+, \dots, p_\kappa^+)$ and $P^- = (p_1^-, \dots, p_\nu^-)$ with rational probabilities over PS and NS, whose marginals are equal for all items. Consider the complement collections $\overline{\text{PS}}$ and $\overline{\text{NS}}$; let d' be the average deficit of $\overline{\text{NS}}$ and let s' be the average surplus of $\overline{\text{PS}}$. Since the average deficit of PS and the average surplus of NS are 0, then by Observation 3.5, $d' \leq \epsilon - \beta$ and $s' \leq \epsilon + \beta$ and so $d' + s' \leq 2\epsilon$.

Towards constructing the collections PS^* and NS^* , consider the collections $\text{PS} \cup \overline{\text{NS}}$ and $\overline{\text{PS}} \cup \text{NS}$ of $\kappa + \nu$ sets each. We define a distribution Q over $\kappa + \nu$ sets, which can be associated with both $\text{PS} \cup \overline{\text{NS}}$ and $\overline{\text{PS}} \cup \text{NS}$, to be $Q = (\frac{1}{2}p_1^+, \dots, \frac{1}{2}p_\kappa^+, \frac{1}{2}p_1^-, \dots, \frac{1}{2}p_\nu^-)$. This distribution has the following properties:

- First, if sets are randomly drawn from $\text{PS} \cup \overline{\text{NS}}$ according to this distribution, the probability of selecting a set with value M is at least $1/2$, since the total weight on sets in PS is exactly $1/2$. Similarly, the probability of selecting a set from $\overline{\text{PS}} \cup \text{NS}$ with value $-M$ when sampling according to Q is at least $1/2$.
- Second, for $\text{PS} \cup \overline{\text{NS}}$ and for every item i , the probability of selecting a set with item i (i.e., the marginal of i) is exactly $1/2$, since it is equal to $\sum_{j|i \in P_j} \frac{1}{2}p_j^+ + \sum_{j|i \in \overline{N}_j} \frac{1}{2}p_j^- = \sum_{j|i \in P_j} \frac{1}{2}p_j^+ + \frac{1}{2} - \sum_{j|i \notin \overline{N}_j} \frac{1}{2}p_j^-$, and from the equality of the marginals of P^+, P^- we have that $\sum_{j|i \in P_j} p_j^+ = \sum_{j|i \in N_j} p_j^- = \sum_{j|i \notin \overline{N}_j} p_j^-$. The same holds for item marginals when the distribution is taken over $\overline{\text{PS}} \cup \text{NS}$.
- Third, the probabilities of the distribution Q are all rational and strictly positive.

By the third property, we can duplicate sets in $\text{PS} \cup \overline{\text{NS}}$ and in $\overline{\text{PS}} \cup \text{NS}$ to construct the collections PS^* and NS^* , such that sampling a set uniformly at random from PS^* is equivalent to sampling a set according to Q from $\text{PS} \cup \overline{\text{NS}}$, and similarly for NS^* and $\overline{\text{PS}} \cup \text{NS}$. By the first property, the average deficit of PS^* is $d \leq d'/2$ and the average surplus of NS^* is $s \leq s'/2$, and $d + s \leq \epsilon$. By the second property, every item must appear in exactly half the sets in PS^* and half the sets in NS^* , meaning that the number of sets in PS^* and NS^* is even, and we can denote it by $2k'$ for some integer $k' > 0$. We have thus shown the existence of a collection PS^* and its complement NS^* with k' sets each whose average surplus and deficit guarantee $d + s \leq \epsilon$. Observe that existence of such collections of size k holds for every integer multiple $k = ck'$, since taking c copies of PS^* and c copies of NS^* satisfies all the conditions of the lemma. ■

3.3 Main Lemmas for Upper Bounds

We present the two key lemmas used to establish our upper bounds. We begin with Lemma 3.11, which is a version of an argument used in Lemma 3.1 of [18]. The next two definitions are used in the lemma statement.

Definition 3.9 (Refined partition). *A refined partition \mathcal{G}' of a collection of sets \mathcal{G} is a collection of sets obtained from \mathcal{G} by partitioning each set $S \in \mathcal{G}$ into subsets $S_1, \dots, S_p \in \mathcal{G}'$ (where p can be 1).*

Definition 3.10 (Recombining by disjoint unions). *A collection of sets \mathcal{G} can be recombined by disjoint unions into another collection \mathcal{G}' if \mathcal{G}' is a refined partition of \mathcal{G} .*

Lemma 3.11 (Using expanders for set recombination). *Let $k \in \mathbb{N}_{\geq 0}$ and $\alpha, r, \theta \in \mathbb{R}_{\geq 0}$ be such that there exists an (α, r, θ) -expander G_k . Consider a collection \mathcal{G} of $2k$ sets with α -frequent items (referred to as the source sets). Then \mathcal{G} has a refined partition into a total of $2kr$ subsets (referred to as the intermediate subsets), which can be recombined by disjoint unions into a collection of $2k\theta$ sets with α/θ -frequent items (referred to as the target sets).*

Proof. Let $G_k = G_k(V, W; E)$. Align the $2k$ source sets with the $2k$ vertices of V in the (α, r, θ) -expander G_k . Because every item appears in $2k\alpha$ sets, then for every item i there are $2k\alpha$ vertices in V corresponding to i (i.e., aligned with the source sets that contain i). By the expansion property of G_k , for every item i there exists a perfect matching M_i between the vertices in V that correspond to i , and some $2k\alpha$ vertices in W .

We now use these matchings to label the edges: For every item i , add i to the labels of the matched edges in M_i . Every edge in E is now labeled by a subset of items (some labels may be the empty set). Let these labels be the intermediate subsets. Their total number is $|E| = 2kr$, as desired. For every vertex $v \in V$, the items of the source set S_v corresponding to v are partitioned among the edges leaving v , and hence the intermediate subsets indeed reflect a refined partitioning of the source sets.

Observe that the edges entering a vertex $w \in W$ are labeled by disjoint intermediate subsets (since for every item i , the edges labeled by subsets containing i form a matching). Let the set S_w corresponding to w be the disjoint union of the subsets labeling the edges adjacent to w . The sets corresponding to the vertices in W can thus be the target sets.

Notice that by construction, every item i appears in the same number of source and target sets, so if the source sets have α -frequent items, the target sets have α/θ -frequent items, completing the proof. \blacksquare

Let f be a weakly ϵ -modular set function whose closest linear set function is the zero function. Let M be the maximum absolute value of f . The next two lemmas upper-bound M , where the first is simpler and the second is more nuanced. The expert reader may skip directly to Lemma 3.13.

Lemma 3.12. *Let $k \in \mathbb{N}_{\geq 0}$ and $\alpha, r, \theta \in \mathbb{R}_{\geq 0}$ be such that there exists an (α, r, θ) -expander G_k . Let \mathcal{G} and \mathcal{G}' be collections of $2k$ sets each, both with α -frequent items, such that the average deficit of \mathcal{G} is at most d and the average surplus of \mathcal{G}' is at most s . Then*

$$M \leq \frac{\frac{1}{2}(d + s) + 2\epsilon(r - 1)}{1 - \theta} + \epsilon.$$

Proof. We first apply Lemma 3.11 to partition and disjointly recombine the sets of collection \mathcal{G} using the expander $G_k = G_k(V, W; E)$. We use the following notation: For a vertex $v \in V$ (resp., $w \in W$) let S_v (resp., S_w) be the source (resp., target) set corresponding to v (resp., w). Denote the neighboring vertices of a vertex v by $N(v)$ and its degree by $\deg(v)$. Let $S_{v,w}$ be the intermediate subset that labels (corresponds to) edge $(v, w) \in E$.

By Lemma 3.11, for every $v \in V$, the intermediate subsets labeling the edges adjacent to v are disjoint, and the same holds for every $w \in W$. We can thus apply Observation 2.2 to get

$$\sum_{w \in N(v)} f(S_{v,w}) \geq f(S_v) + (\deg(v) - 1)(f(\emptyset) - \epsilon) \quad \forall v \in V. \quad (2)$$

$$\sum_{v \in N(w)} f(S_{v,w}) \leq f(S_w) + (\deg(w) - 1)(f(\emptyset) + \epsilon) \quad \forall w \in W. \quad (3)$$

Since collection \mathcal{G} has average deficit of at most d , by summing over vertices $v \in V$ (where $|V| = 2k$) we get $\sum_{v \in V} f(S_v) \geq 2k(M - d)$. Since the target sets of \mathcal{G} have value at most M , by summing over vertices $w \in W$ (where $|W| = 2\theta k$) we get $\sum_{w \in W} f(S_w) \leq 2\theta kM$. Clearly in the bipartite graph

G_k , $\sum_{v \in V} \deg(v) = \sum_{w \in W} \deg(w)$, and by the parameters of G_k both are equal to $2kr$. Therefore, by Equations (2) and (3), summing over $v \in V$ and $w \in W$ we get

$$\begin{aligned} 2kM - 2kd + (2kr - 2k)(f(\emptyset) - \epsilon) &\leq \sum_{(v,w) \in E} f(S_{v,w}) \\ &\leq 2\theta kM + (2kr - 2\theta k)(f(\emptyset) + \epsilon). \end{aligned}$$

Dividing the resulting inequality by $2k$ and rearranging gives

$$\begin{aligned} (1 - \theta)M &\leq d + (r - 1)(\epsilon - f(\emptyset)) + (r - \theta)(\epsilon + f(\emptyset)) \\ &= d + 2\epsilon(r - 1) + (1 - \theta)(\epsilon + f(\emptyset)). \end{aligned} \quad (4)$$

Similarly, using that the average deficit of collection \mathcal{G}' is at most s and the value of its target sets is at most M , $\sum_{v \in V} f(S_v) \leq 2k(-M + s)$, and $\sum_{w \in W} f(S_w) \geq -2\theta kM$. Therefore, by Equations (2) and (3)

$$\begin{aligned} -2kM + 2ks + (2kr - 2k)(f(\emptyset) + \epsilon) &\geq \sum_{(v,w) \in E} f(S_{v,w}) \\ &\geq -2\theta kM + (2kr - 2\theta k)(f(\emptyset) - \epsilon). \end{aligned}$$

Dividing the resulting inequality by $2k$ and rearranging gives

$$\begin{aligned} (1 - \theta)M &\leq s + (r - 1)(\epsilon + f(\emptyset)) + (r - \theta)(\epsilon - f(\emptyset)) \\ &= s + 2\epsilon(r - 1) + (1 - \theta)(\epsilon - f(\emptyset)). \end{aligned} \quad (5)$$

Rearranging Inequalities (4) and (5) as well as averaging the resulting inequalities implies the theorem. \blacksquare

Lemma 3.13 (Strong Version of Lemma 3.12). *Let $k \in \mathbb{N}_{\geq 0}$ and $\alpha, r, \theta \in \mathbb{R}_{\geq 0}$ be such that there exists an (α, r, θ) -expander G_k . Let \mathcal{G} and \mathcal{G}' (possibly $\mathcal{G} = \mathcal{G}'$) be collections of $2k$ sets each, both with α -frequent items, such that the average deficit of \mathcal{G} is $\leq d$, the average surplus of \mathcal{G}' is $\leq s$, and their target sets have average deficit $\geq d'$ and average surplus $\geq s'$, respectively. Then*

$$M \leq \frac{\frac{1}{2}(d + s - \theta(d' + s')) + 2\epsilon(r - 1)}{1 - \theta} + \epsilon.$$

Proof. We first apply Lemma 3.11 to partition and disjointly recombine the sets of collection \mathcal{G} via the expander $G_k = G_k(V, W; E)$. We use the following notation: For a vertex $v \in V$ (resp., $w \in W$) let S_v (resp., S_w) be the source (resp., target) set corresponding to v (resp., w). Denote the neighboring vertices of a vertex v by $N(v)$ and its degree by $\deg(v)$. Let $S_{v,w}$ be the intermediate subset that labels (corresponds to) edge $(v, w) \in E$.

By Lemma 3.11, for every $v \in V$, the intermediate subsets labeling the edges adjacent to v are disjoint, and the same holds for every $w \in W$. We can thus apply Observation 2.2 to get

$$\sum_{w \in N(v)} f(S_{v,w}) \geq f(S_v) + (\deg(v) - 1)(f(\emptyset) - \epsilon) \quad \forall v \in V. \quad (6)$$

$$\sum_{v \in N(w)} f(S_{v,w}) \leq f(S_w) + (\deg(w) - 1)(f(\emptyset) + \epsilon) \quad \forall w \in W. \quad (7)$$

Denote the maximum absolute value of f by M . Since collection \mathcal{G} has average deficit $\leq d$, by summing over vertices $v \in V$ (where $|V| = 2k$) we get $\sum_{v \in V} f(S_v) \geq 2k(M - d)$. Since the target

sets of \mathcal{G} have average deficit $\geq d'$, by summing over vertices $w \in W$ (where $|W| = 2\theta k$) we get $\sum_{w \in W} f(S_w) \leq 2\theta k(M - d')$. Clearly in the bipartite graph G_k , $\sum_{v \in V} \deg(v) = \sum_{w \in W} \deg(w)$, and by the parameters of G_k both are equal to $2kr$. Therefore, by Equations (6) and (7), summing over $v \in V$ and $w \in W$ we get

$$\begin{aligned} 2kM - 2kd + (2kr - 2k)(f(\emptyset) - \epsilon) &\leq \sum_{(v,w) \in E} f(S_{v,w}) \\ &\leq 2\theta kM - 2\theta kd' + (2kr - 2\theta k)(f(\emptyset) + \epsilon). \end{aligned}$$

Dividing the resulting inequality by $2k$ and rearranging gives

$$\begin{aligned} (1 - \theta)M &\leq d - \theta d' + (r - 1)(\epsilon - f(\emptyset)) + (r - \theta)(\epsilon + f(\emptyset)) \\ &= d - \theta d' + 2\epsilon(r - 1) + (1 - \theta)(\epsilon + f(\emptyset)). \end{aligned} \quad (8)$$

Similarly, using that the average deficit of collection \mathcal{G}' is $\leq s$ and the average deficit of its target sets is $\geq s'$, $\sum_{v \in V} f(S_v) \leq 2k(-M + s)$, and $\sum_{w \in W} f(S_w) \geq -2\theta k(M - s')$. Therefore, by Equations (6) and (7)

$$\begin{aligned} -2kM + 2ks + (2kr - 2k)(f(\emptyset) + \epsilon) &\geq \sum_{(v,w) \in E} f(S_{v,w}) \\ &\geq -2\theta kM + 2\theta ks' + (2kr - 2\theta k)(f(\emptyset) - \epsilon). \end{aligned}$$

Dividing the resulting inequality by $2k$ and rearranging gives

$$\begin{aligned} (1 - \theta)M &\leq s - \theta s' + (r - 1)(\epsilon + f(\emptyset)) + (r - \theta)(\epsilon - f(\emptyset)) \\ &= s - \theta s' + 2\epsilon(r - 1) + (1 - \theta)(\epsilon - f(\emptyset)). \end{aligned} \quad (9)$$

Rearranging Inequalities (8) and (9) as well as averaging the resulting inequalities implies the theorem. \blacksquare

3.4 Upper-Bounding the Weak Kalton Constant K_w

To upper bound K_w in the next two sections, we focus without loss of generality on a weakly 1-modular set function f , whose closest linear set function is the zero function (Corollary 2.8). In this section we show the following upper bound:

Lemma 3.14 (Upper bound on K_w). *Suppose that for fixed $r, \theta \in \mathbb{R}_{\geq 0}$ there exist $(\frac{1}{2}, r, \theta)$ -expanders. Then the weak Kalton constant satisfies: $K_w \leq \frac{2r - \frac{1}{2} - \theta}{1 - \theta}$.*

Proof. By Lemma 3.8, there exists k' such that for every k that is a member of the arithmetic progression $k', 2k', 3k', \dots$, the function f has collections PS^* and NS^* of $2k$ sets with $\frac{1}{2}$ -frequent items, and whose average deficit d and average surplus s , respectively, satisfy $d + s \leq 1$. The assumption that $(\frac{1}{2}, r, \theta)$ -expanders exist implies that there is another arithmetic progression $k'', 2k'', 3k'', \dots$ such that for every k in this sequence an expander G_k with the above parameters exist. As the two arithmetic progressions must meet at $k'k''$, there is a common k in both progressions. The upper bound follows from applying Lemma 3.12 to collections PS^*, NS^* to get $K_w \leq \frac{\frac{1}{2} + 2r - 2}{1 - \theta} + 1$. \blacksquare

Theorem 3.15. *The weak Kalton constant satisfies $K_w \leq 26.8$.*

Proof. In Lemma 3.14, we can get $K_w \leq 26.8$ if we substitute $r = 5.05$ and $\theta = \frac{2}{3}$ by using the expanders from [6].² \blacksquare

²Using the expanders from [26] (with $r = 6$ instead of $r = 5.05$) would result in $K_w \leq 32.5$.

3.5 Strengthening the Upper-Bound on K_w

In this section we use Lemma 3.13 (the strengthened version of Lemma 3.12) to further improve our upper bound on the weak Kalton constant. Recall the collections PS^* and NS^* guaranteed to exist by Lemma 3.8, and denote their average deficit and surplus by d and s , respectively, and the average deficit and surplus of their target sets by d' and s' , respectively.

Lemma 3.16 (Strengthened upper bound on K_w). *Suppose that for fixed $r, r', \theta, \theta' \in \mathbb{R}_{\geq 0}$ there exist $(\frac{1}{2}, r, \theta)$ -expanders and $(1 - \frac{1}{2\theta}, r', \theta')$ -expanders. Then the weak Kalton constant satisfies:*

$$K_w \leq \min \left\{ \frac{2r - \frac{1}{2} - \theta - \theta \frac{d'+s'}{2}}{1 - \theta}, \frac{2r' - \theta' + \frac{d'+s'}{2}}{1 - \theta'} \right\}.$$

Proof. As in the proof of Lemma 3.14, observe that there exists k such that f has collections PS^* and NS^* of $2k$ sets, and expanders G_k and G'_k with the above parameters. The first upper bound is found in Lemma 3.14. Now consider the collection \mathcal{G} that is the *complement of the target sets* of PS^* . Since PS^* has $\frac{1}{2}$ -frequent items, then the target sets have $\frac{1}{2\theta}$ -frequent items (Lemma 3.11), and \mathcal{G} has $(1 - \frac{1}{2\theta})$ -frequent items (Observation 3.7). By Observation 3.5 with $\epsilon = 1$, the average surplus of \mathcal{G} is $\leq d' + 1 + \delta$, and clearly the average surplus of its target sets is ≥ 0 . Similarly, let \mathcal{G}' be the complement of the target sets of NS^* , which has $(1 - \frac{1}{2\theta'})$ -frequent items and average deficit $\leq s' + 1 - \delta$. Clearly the average deficit of the target sets of \mathcal{G}' is ≥ 0 . Applying Lemma 3.13 thus gives an upper bound of

$$K_w \leq \frac{\frac{1}{2}(d' + s' + 2) + 2r' - 2}{1 - \theta'} + 1 = \frac{2r' - 1 + \frac{d'+s'}{2}}{1 - \theta'} + 1. \quad (10)$$

Taking the minimum of the bound in Lemma 3.14 and the bound in (10) completes the proof. \blacksquare

We can now prove the upper bound in Theorem 1.3, by which the weak Kalton constant satisfies $K_w \leq 23.811$.

Proof. [Proof of First Part of Theorem 1.3] By using the expanders with $\alpha = \frac{1}{2}$ from Lemma 3.1 we can substitute $r = 5$ and $\theta = \frac{5}{7}$. Since $1 - \frac{1}{2\theta} = \frac{3}{10}$, we can complement these parameters by using the expanders with $\alpha' = \frac{3}{10}$ from Lemma 3.1 for which $r' = 4$ and $\theta' = \frac{4}{7}$. This implies that $K_w \leq \min\{30.75 - 1.25(d' + s'), \frac{52}{3} + \frac{7}{6}(d' + s')\} \leq 23.811$. \blacksquare

3.6 Upper-Bounding the Strong Kalton Constant K_s

As in the previous section, we focus here without loss of generality on a 1-modular set function f whose closest linear set function is the zero function, and on its collections PS^* and NS^* as defined in Lemma 3.8. Recall that d and s are the average deficit and surplus of PS^* and NS^* , respectively, and $d + s \leq 1$. Denote the average deficit of an intersection of ℓ sets in PS^* by d_ℓ and the average surplus of an intersection of ℓ sets in NS^* by s_ℓ .

Lemma 3.17. $d_\ell + s_\ell \leq \lceil \frac{5\ell}{2} \rceil - 2$.

Proof. We first prove the following claim: $d_2 + s_2 \leq 3$.

To prove the claim, we begin with a simple observation: Recall that PS^* and NS^* are complements. For every $P_1, P_2 \in \text{PS}^*$ whose complements are $N_1, N_2 \in \text{NS}^*$ it holds that $f(P_1 \cup P_2) =$

$f(\overline{P_1 \cap P_2}) = f(\overline{N_1 \cap N_2}) \leq -f(N_1 \cap N_2) + \epsilon + \beta$ (Corollary 3.3). Similarly, $f(N_1 \cup N_2) = f(\overline{N_1 \cap N_2}) = f(\overline{P_1 \cap P_2}) \geq -f(P_1 \cap P_2) - \epsilon + \beta$. By 1-modularity we thus have $f(P_1 \cap P_2) \geq f(P_1) + f(P_2) - f(P_1 \cup P_2) - 1 \geq f(P_1) + f(P_2) + f(N_1 \cap N_2) - \beta - 2$, and $f(N_1 \cap N_2) \leq f(N_1) + f(N_2) - f(N_1 \cup N_2) + 1 \leq f(N_1) + f(N_2) + f(P_1 \cap P_2) - \beta + 2$.

Recall that M is the maximum absolute value of f . We now take the average over P_1, P_2 . This gives $M - d_2 \geq 2M - 2d + (-M + s_2) - \beta - 2$. Similarly, the average over N_1, N_2 gives $-M + s_2 \leq -2M + 2s + (M - d_2) - \beta + 2$. Subtracting the second inequality from the first we get $2M - d_2 - s_2 \geq 4M - 2d - 2s - 2M + d_2 + s_2 - 4$, so $d_2 + s_2 \leq d + s + 2 \leq 3$, completing the proof of the claim.

We can now prove Lemma 3.17 by induction: It holds for $\ell = 1$ and $\ell = 2$ using the above claim. For every ℓ sets $P_1, \dots, P_\ell \in \text{PS}^*$, by 1-modularity the value of their intersection satisfies $f((P_1 \cap \dots \cap P_{\ell_1}) \cap (P_{\ell_1+1} \cap \dots \cap P_\ell)) \geq f(P_1 \cap \dots \cap P_{\ell_1}) + f(P_{\ell_1+1} \cap \dots \cap P_\ell) - M - 1$. Taking the average over the ℓ sets we get $M - d_\ell \geq M - d_{\ell_1} - d_{\ell - \ell_1} - 1$. Similarly, $-M + s_\ell \leq -M + s_{\ell_1} + s_{\ell - \ell_1} + 1$. So $d_\ell + s_\ell \leq d_{\ell_1} + d_{\ell - \ell_1} + s_{\ell_1} + s_{\ell - \ell_1} + 2$.

If ℓ is even we can take $\ell_1 = 2$ and using the induction hypothesis get $d_\ell + s_\ell \leq 3 + \frac{5(\ell-2)}{2} - 2 + 2 = \frac{5\ell}{2} - 2$. If ℓ is odd we can take $\ell_1 = 1$ and using the induction hypothesis get $d_\ell + s_\ell \leq 1 + \frac{5(\ell-1)}{2} - 2 + 2 = \frac{5(\ell-1)}{2} + 1$. This completes the proof. \blacksquare

Observation 3.18. *Every item is contained in $1/2^\ell$ of the intersections of ℓ sets in PS^* or NS^* .*

Recall that the case of $\ell = 1$ was established in Lemma 3.8. Observation 3.18 allows us to use the main lemma for upper-bounding the Kalton constants (Lemma 3.11) with $\alpha = 1/2^\ell$. Together with Lemma 3.17 that establishes the average deficit and surplus, this enables us to obtain the following result.

Theorem 3.19. *Suppose that for fixed $r, \theta \in \mathbb{R}_{\geq 0}$ there exist (α, r, θ) -expanders. Then the strong Kalton constant satisfies $K_s \leq \frac{2r + N_2 - \theta}{1 - \theta}$, where if $\alpha = \frac{1}{2}$ then $N_2 = -\frac{1}{2}$; if $\alpha = \frac{1}{4}$ then $N_2 = \frac{1}{2}$; if $\alpha = \frac{1}{8}$ then $N_2 = 2$; if $\alpha = \frac{1}{16}$ then $N_2 = 3$; and if $\alpha = \frac{1}{32}$ then $N_2 = 4.5$.*

Proof. For every ℓ , by taking the intersections of ℓ sets in PS^* and NS^* we get collections \mathcal{G} and \mathcal{G}' , both with $\frac{1}{2^\ell}$ -frequent items (Observation 3.18), whose average deficit and surplus are d_ℓ and s_ℓ , respectively. Moreover, there exists some k such that \mathcal{G} and \mathcal{G}' have $2k$ sets each, and an expander G_k with the above parameters exists. We now apply Lemma 3.12, which gives an upper bound of $K_s \leq \frac{\frac{1}{2}(d_\ell + s_\ell) + 2r - 2}{1 - \theta} + 1$. By Lemma 3.17, $d_\ell + s_\ell$ is at most 3, 6, 8, 11 for $\ell = 2, 3, 4, 5$, respectively. This implies the theorem. \blacksquare

To use Theorem 3.19 one needs to substitute in parameters of expanders. Theorem 3.19 combined with the $(\frac{1}{16}, 4, \frac{4}{15})$ -expanders from Lemma 3.1 implies that the value of the strong Kalton constant satisfies $K_s \leq \frac{2 \cdot 4 + 3 - 0.26}{1 - 0.26} \leq 14.637$.

3.7 Strengthening the Upper Bound on K_s

Using our techniques, the upper bound on K_s can be improved even further. The sources of these improvements are twofold. First, all of our results are derived using bi-regular expanders; better bounds can be obtained using more sophisticated expanders. Second, improvements can be obtained by using additional properties of ϵ -modular functions to improve the bounds in Lemma 3.17. Improvements of the second type are demonstrated in this section, where the main result is showing that $K_s < 12.62$ (establishing Theorem 1.1).

Definition 3.20. Given $0 \leq \alpha \leq \frac{1}{2}$ and $\mu \geq 0$, an (α, μ) -collection-pair is a pair of collections D, S , such that in both collections items are α -frequent, the average value of sets in D is at least $M - d$, the average value of sets in S is at most $-M + s$, and $\frac{1}{2}(d + s) = \mu$.

For a given $\mu \geq 0$, let $\alpha(\mu)$ denote the smallest $0 \leq \alpha \leq \frac{1}{2}$ such that there is some $\mu' \leq \mu$ for which an (α, μ') -collection-pair exists. By Lemma 3.17, we may assume for our given 1-modular function f that $\alpha(\frac{1}{2}) \leq \frac{1}{2}$, $\alpha(\frac{3}{2}) \leq \frac{1}{4}$, and $\alpha(4) \leq \frac{1}{16}$.

Fix some $\frac{3}{2} < \delta < 4$ whose value will be optimized later. (Intuitively, we are aiming at δ satisfying $\alpha(\delta) \simeq \frac{1}{8}$.) Now we consider two cases, each addressed in its own lemma.

Lemma 3.21. Suppose that $\alpha(\delta) \leq \frac{\alpha(\frac{3}{2})}{2}$. Then $\alpha(2\delta - \frac{1}{2}) \leq \frac{1}{64}$.

Proof. Consider the $\mu' \leq \delta$ for which an $(\alpha(\delta), \mu')$ -collection-pair (D, S) with deficit d' and surplus s' exists, and let m denote the number of sets in D and in S (sets appearing more than once are counted more than once). Consider now the two collection-pairs (D_\cap, S_\cap) and (D_\cup, S_\cup) obtained by taking all m^2 pairwise intersections or unions (respectively) of sets from D, S . (A pair is generated by picking one set and then another set, with repetitions.) Let α_\cap and α_\cup be the item frequencies associated with these two collection-pairs, respectively.³ Let d_\cup, d_\cap be the deficits of D_\cup, D_\cap , respectively, and let s_\cup, s_\cap be the surpluses of S_\cup, S_\cap , respectively. Let μ_\cup (resp., μ_\cap) be the average of d_\cup, s_\cup (resp., d_\cap, s_\cap). Then:

$$\alpha_\cap = (\alpha(\delta))^2 \leq \left(\frac{\alpha(\frac{3}{2})}{2} \right)^2 \leq \frac{1}{64} \quad (11)$$

and

$$\alpha_\cup = 1 - (1 - \alpha(\delta))^2 < 2\alpha(\delta) \leq \alpha(\frac{3}{2}) \quad (12)$$

It follows from Equation (12) that $\mu_\cup \geq \frac{3}{2}$. The 1-modularity condition implies that $2d' - d_\cup + 1 \geq d_\cap$ and $2s' - s_\cup + 1 \geq s_\cap$, and so by averaging $2\mu' - \mu_\cup + 1 \geq \mu_\cap$. Substitute $\mu_\cup \geq 3/2$ and $\mu' \leq \delta$ to get $2\delta - \frac{1}{2} \geq \mu_\cap$. By Equation (11), it follows that $\alpha(2\delta - \frac{1}{2}) \leq \frac{1}{64}$, as required. ■

Lemma 3.22. Suppose that $\alpha(\delta) > \frac{\alpha(\frac{3}{2})}{2}$. Then $\alpha(9 - \delta) \leq \frac{1}{256}$.

Proof. Consider the $\mu' \leq \frac{3}{2}$ for which an $(\alpha(\frac{3}{2}), \mu')$ -collection-pair (D, S) with deficit d' and surplus s' exists. For the collection-pair (D_\cap, S_\cap) (that we shall denote by (\tilde{D}, \tilde{S})) the item frequencies are $(\alpha(\frac{3}{2}))^2$, and the average of the deficit and surplus $\tilde{\mu}$ is at most $2\mu' + 1 \leq 4$ (using 1-modularity).⁴ Consider now $(\tilde{D}_\cap, \tilde{S}_\cap)$ and $(\tilde{D}_\cup, \tilde{S}_\cup)$. Let $\tilde{\alpha}_\cap$ and $\tilde{\alpha}_\cup$ be the item frequencies associated with these two collections respectively. Then

$$\tilde{\alpha}_\cap = \left(\alpha(\frac{3}{2}) \right)^4 \leq \frac{1}{256} \quad (13)$$

³Notice that both D_\cap and S_\cap have the same item frequency, since an item's frequency in D_\cap (resp., S_\cap) is the probability of taking two random sets from collection D (resp., S) and having the item appear in both – where the item's frequencies in D and S are equal. A similar argument holds for D_\cup and S_\cup .

⁴For completeness we give the details: For every 2 sets $U, V \in \tilde{D}$, by 1-modularity $f(U \cap V) \geq f(U) + f(V) - f(U \cup V) - 1 \geq f(U) + f(V) - M - 1$. On average this is at least $2M - 2d' - M - 1$. Similarly, the average surplus of \tilde{S} is $2s' + 1$. Taking the average completes the argument.

and

$$\tilde{\alpha}_U < 2 \left(\alpha \left(\frac{3}{2} \right) \right)^2 < \alpha(\delta), \quad (14)$$

where the last inequality uses the fact that $\alpha(\frac{3}{2}) \leq \frac{1}{4}$ and the premise of the lemma.

Similarly to the proof of Lemma 3.21, it follows from Equation (14) that $\tilde{\mu}_U \geq \delta$. The 1-modularity condition then implies that $9 - \delta \geq 2 \cdot 4 - \delta + 1 \geq 2\tilde{\mu} - \tilde{\mu}_U + 1 \geq \tilde{\mu}_\Gamma$. This completes the proof of the lemma. \blacksquare

Corollary 3.23. *For every $\frac{3}{2} < \delta < 4$, every 1-modular function f as above has an (α, μ) -collection-pair with either $\alpha \leq \frac{1}{64}$ and $\mu \leq 2\delta - \frac{1}{2}$, or $\alpha \leq \frac{1}{256}$ and $\mu \leq 9 - \delta$.*

Improved upper bound no. 1: $K_s < 13.25$. We now apply Lemma 3.12, which gives an upper bound of $K_s \leq \frac{\mu+2r-2}{1-\theta} + 1$ when there are (α, r, θ) -expanders. By Corollary 3.23 we get

$$K_s \leq \max \left\{ \frac{2\delta + 2r_1 - 2.5}{1 - \theta_1}, \frac{-\delta + 2r_2 + 7}{1 - \theta_2} \right\} + 1,$$

where r_1, θ_1 are the best expander parameters for $\alpha = \frac{1}{64}$ (recall from Lemma 3.1 we can use 3 and $\frac{3}{11}$, respectively), and r_2, θ_2 are the best expander parameters for $\alpha = \frac{1}{256}$ (recall from Lemma 3.1 we can use 3 and $\frac{3}{19}$, respectively). Plugging in $\delta = 2.7$, and using Lemma 3.1 we get that $K_s \leq 13.24$.

Improved upper bound no. 2: $K_s < 12.62$. The 13.25 bound can be further improved by using Lemma 3.13 (which is an improved version of Lemma 3.12). To make use of Lemma 3.13 we need $d' + s'$ (the bounds on the average deficit and surplus of the target sets) to be sufficiently large. We therefore distinguish between two cases as shown next.

Suppose f has an (α, μ) -collection-pair with $\alpha \leq \frac{1}{64}$ and $\mu \leq 2\delta - 1/2$ (the first case in Corollary 3.23), and consider an $(\alpha, 3, \frac{3}{11})$ -expander (which exists by Lemma 3.1). Consider the target sets of the collection pair guaranteed by Lemma 3.11; their frequency is $\frac{\alpha}{3/11} \leq \frac{1/64}{3/11} < 1/16$. If $d' + s' \leq 5.04$, then we apply Lemma 3.12 with the target sets above serving as the source sets, and using the existence of $(\frac{1}{16}, 4, \frac{4}{15})$ -expanders (Lemma 3.1), to get $K_s < 12.62$ (by substituting $d' + s' \leq 5.04$, $r = 4$, $\theta = 4/15$). Similarly, if f has an (α, μ) -collection-pair with $\alpha \leq \frac{1}{256}$ and $\mu \leq 9 - \delta$ (the second case in Corollary 3.23), repeat an analogous analysis using an $(\alpha, 3, \frac{3}{19})$ -expander (which exists by Lemma 3.1) to get $K_s < 12.62$, as above (this is valid since the frequency of the target sets is $\frac{\alpha}{3/19} \leq \frac{1/256}{3/19} < 1/16$).

So now we apply Lemma 3.13 in the case where $d' + s' > 5.04$. By Corollary 3.23, we get that for every $\frac{3}{2} < \delta < 4$ it holds that K_s is upper-bounded by

$$\max \left\{ \frac{2\delta - \frac{1}{2} - \frac{5.04\theta_1}{2} + 2(r_1 - 1)}{1 - \theta_1}, \frac{9 - \delta - \frac{5.04\theta_2}{2} + 2(r_2 - 1)}{1 - \theta_2} \right\} + 1,$$

where r_1, θ_1 are the best expander parameters for $\alpha = \frac{1}{64}$ (recall from Lemma 3.1 we can use 3 and $\frac{3}{11}$, respectively), and r_2, θ_2 are the best expander parameters for $\alpha = \frac{1}{256}$ (recall from Lemma 3.1 we can use 3 and $\frac{3}{19}$, respectively). Plugging in $\delta = 2.818$, and using Lemma 3.1 we get that $K_s < 12.62$. Together, we get that $K_s < 12.62$.

Remark The bounds shown in this section illustrate the techniques we use. Clearly, these techniques can be extended to give even better bounds by, e.g., establishing better expanders and applying our ideas recursively. We save these extensions for future work.

4 Improved Lower Bounds

In this section we prove the lower bound stated in Theorem 1.2 on the strong Kalton constant; i.e., we prove that $K_s \geq 1$. The lower bound stated in Theorem 1.3 on K_w is proved in Appendix C.

We begin by explicitly constructing a set function f that establishes the lower bound. The idea underlying f 's construction is to achieve symmetry properties which facilitate its analysis. In general, checking whether a set function over n items is approximately modular involves verifying that roughly 2^{2n} modular equations (one equation for every pair of sets) approximately hold. Verifying approximate modularity for f becomes an easier task thanks to its symmetries. This idea is developed in Section 4.1 where we explicitly state the properties of the class to which f belongs. The analysis appears in Section 4.2.

Proof. [Proof of Theorem 1.2] The following set function with $n = 70$ is 2-modular and tightly 2-linear, thus showing $K_s \geq 1$. Let $M = 2$. The positive support PS and the negative support NS of f (Definition 2.17) are as follows: There are 8 sets in PS, each with 35 out of the 70 items, such that every item appears in exactly 4 of the 8 sets (notice that $\binom{8}{4} = 70$). (We can fix any such collection of sets as PS without loss of generality.) The negative support NS is the complement collection of PS, and it does not overlap with PS (notice that the intersection of any two sets from PS has exactly 20 elements, so they cannot complement each other). Notice that PS, NS support uniform distributions P^+, P^- with equal item marginals, as required. The values of sets under f are determined by the following rules, where the first applicable rule applies (the last statement necessarily makes f well defined):

1. Each positive support set $S \in \text{PS}$ has value $f(S) = M = 2$.
2. For every two distinct sets S_1 and S_2 in PS and every set R , we have $f(S_1 \cup (S_2 \cap R)) = 1$, and likewise $f(S_1 \cap (S_2 \cup R)) = 1$. (In particular, for every two sets $S_1, S_2 \in \text{PS}$, $f(S_1 \cup S_2) = f(S_1 \cap S_2) = 1$.)
3. We impose $-f(S) = f(\bar{S})$ and derive from this sets with negative value.
4. All other sets have value 0.

Claim 4.1. f is tightly 2-linear.

Claim 4.2. f is 2-modular.

The proof of the latter claim is based on an analysis of a subset of cases, which are sufficient to establish ϵ -modularity for f due to its symmetry properties (see Lemma 4.7). The proofs of the claims appear in Section 4.2, thus completing the proof of Theorem 1.2. ■

4.1 (k, M) -Symmetric Set Functions

Definition 4.3. A collection $\mathcal{G} = \{S_1, S_2, \dots, S_g\}$ of g subsets of U is generating if the following conditions hold: (1) It is covering, i.e., $\bigcup_{S_j \in \mathcal{G}} S_j = U$ (every item is contained in at least one set). (2) It is item-differentiating, i.e., $\bigcap_{S_j \in \mathcal{G} | i \in S_j} S_j = \{i\}$ for every item i .⁵

⁵Equivalently, for every pair of items, there is a set containing one but not the other. This implies in particular that $\bigcap_{S_j \in \mathcal{G}} S_j = \emptyset$ (no item is contained in all sets).

If \mathcal{G} is generating then every subset of U can be generated from sets in \mathcal{G} by a sequence of intersections and unions; and the complement collection $\bar{\mathcal{G}}$ is also generating (De Morgan).

Definition 4.4. A generating collection \mathcal{G} is canonical if the number g of generating sets is even (we denote $g = 2k$ for some positive integer k), every item is contained in exactly k sets from \mathcal{G} , every set in \mathcal{G} contains exactly $n/2$ items, and $n = \binom{2k}{k}$.

Given a canonical generating collection \mathcal{G} where $\mathcal{G} = \{S_1, S_2, \dots, S_{2k}\}$, items can be thought of as balanced vectors in $\{\pm 1\}^{2k}$, where coordinate j of item i is $+1$ if $i \in S_j$, and -1 if $i \in \bar{S}_j$. Observe also that if \mathcal{G} is a canonical generating collection, then so is its complement $\bar{\mathcal{G}}$.

Definition 4.5. For integers $k \geq 2$ and $M \geq 1$, we say that an integral set function f over a set U of $n = \binom{2k}{k}$ items is (k, M) -symmetric if it has the following properties:

1. Antisymmetry: for every set S and its complement \bar{S} it holds that $f(S) = -f(\bar{S})$.
2. Canonical generating sets: positive support PS contains $2k$ sets P_1, \dots, P_{2k} (of value M), negative support NS contains $2k$ sets $N_1 = \bar{P}_1, \dots, N_{2k} = \bar{P}_{2k}$ (of value $-M$, these are the complements of the sets in PS), and PS (and likewise NS) is a canonical generating collection.

Proposition 4.6. Every (k, M) -symmetric function f is tightly M -linear, and the 0 function is a linear function closest to f .

Proof. Consider Property 2 (canonical generating sets) in Definition 4.5. By the virtue of PS being a canonical generating set, the uniform distribution over PS has marginal $1/2$ for every item in U , and likewise for NS. Hence Lemma 2.16 implies that the 0 function is a linear function closest to f . As the maximum value of f is M , it then follows that f is tightly M -linear. ■

We shall design certain (k, M) -symmetric functions and would like to verify that they are ϵ -modular, typically for $\epsilon = 2$. To do so we can restrict our verification as follows.

Lemma 4.7. To verify that a (k, M) -symmetric function f is ϵ -modular it suffices to verify the approximate modularity equation $|f(S) + f(T) - f(S \cap T) - f(S \cup T)| \leq \epsilon$ for either S and T or \bar{S} and \bar{T} . In particular, we can assume $f(S) \geq |f(T)|$.

Proof. Given S and T consider the complement sets \bar{S} and \bar{T} . By antisymmetry they satisfy $f(S) = -f(\bar{S})$, $f(T) = -f(\bar{T})$, $f(S \cap T) = -f(\bar{S} \cup \bar{T})$, and $f(S \cup T) = -f(\bar{S} \cap \bar{T})$. Hence the approximate modularity condition for \bar{S} and \bar{T} implies it for S and T .

Assuming $|f(S)| \geq |f(T)|$ is clearly without loss of generality. Even if $f(S) < |f(T)|$ so that S and T do not satisfy the condition $f(S) \geq |f(T)|$, then \bar{S} and \bar{T} do satisfy the condition, and approximate modularity holds for \bar{S} and \bar{T} if and only if it holds for S and T . Thus assuming $f(S) \geq |f(T)|$ is also without loss of generality for the purpose of verifying approximate modularity. ■

4.2 Proofs of Claims 4.1 and 4.2

Proof. [Proof of Claim 4.1] It is not hard to verify that the function f defined above is (k, M) -symmetric according to Definition 4.5 for $k = 4$ and $M = 2$: the main thing to check is that there are no two positive sets that are complements of each other (and consequently antisymmetry is enforced by rule 3), and this is implied by the proof of Claim 4.8 below. Proposition 4.6 implies that f is tightly 2-linear. ■

We say that S is a *positive set* if $f(S) > 0$, a *negative set* if $f(S) < 0$, and a *zero set* if $f(S) = 0$.

Claim 4.8. *There is no pair of sets S and T such that one of them is positive and the other is negative and $S \subset T$.*

Proof. We show that no positive set is contained in a negative set. The opposite direction can be shown analogously. It is sufficient to show that no minimal positive set is contained in a maximal negative set. View each item as a balanced vector in $\{\pm 1\}^8$. Let S be a minimal positive set. S is the intersection of two sets in PS, it thus has two coordinates (out of eight), say 1 and 2, fixed to +1 and contains all items that agree with both. Let T be a maximal negative set. T is the union of two sets in NS, it thus has two coordinates, say 3 and 4, fixed to -1 and contains all vectors that agree with at least one of the coordinates. The item corresponding to the vector that has +1 on all these four coordinates (here is where we used the fact that $k \geq 4$) is in S but not in T . Similarly, in the case where the set of coordinates for S coincides with the set of coordinates for T or overlaps it, there is an item that belongs to S but not to T . This completes the proof. ■

Proof. [Proof of Claim 4.2] We show that f is 2-modular. Consider two sets S and T . By Lemma 4.7, the cases in the following case analysis suffices in order to establish 2-modularity.

1. $f(S) = 2$ (namely, $S \in \text{PS}$). In this case both $f(S \cup T) \geq 0$ and $f(S \cap T) \geq 0$, by Claim 4.8, a fact that will be implicitly used in the subcases below.
 - (a) $f(T) = 2$ (namely, $T \in \text{PS}$). In this case $f(S \cup T) = f(S \cap T) = 1$, by definition of f (rule 2). Thus, $f(S) + f(T) - f(S \cup T) - f(S \cap T) = 2$.
 - (b) $f(T) = 1$. In this case $f(S) + f(T) = 3$. So it suffices to show that either $f(S \cup T) \geq 1$ or $f(S \cap T) \geq 1$. By the definition of f , there are two possibilities:
 - i. T is contained in some PS set S_2 , and then $f(S \cup T) \geq 1$ by the condition $f(S_1 \cup (S_2 \cap R)) = 1$ (with S_1 serving as S).
 - ii. T contains some PS set S_2 , and then $f(S \cap T) \geq 1$ by the condition $f(S_1 \cap (S_2 \cup R)) = 1$.
 - (c) $f(T) = 0$. In this case $f(S) + f(T) = 2$ and $0 \leq f(S \cup T) + f(S \cap T) \leq 4$, satisfying the 2-modularity condition.
 - (d) $f(T) < 0$. Given that $f(S) > 0$ and $f(T) < 0$ we have that $f(S \cup T) = f(S \cap T) = 0$, by Claim 4.8. Hence the stronger 1-modularity condition holds.
2. $f(S) = 1$. Both $f(S \cup T) \geq 0$ and $f(S \cap T) \geq 0$, by Claim 4.8.
 - (a) $f(T) = 1$. The 2-modularity condition then holds since $f(S) + f(T) = 2$ and $0 \leq f(S \cup T) + f(S \cap T) \leq 4$.
 - (b) $f(T) = 0$. Observe that it cannot be that both $f(S \cup T) = 2$ and $f(S \cap T) = 2$ since then both are in PS meaning that $S = T$. Hence $f(S) + f(T) = 1$ and $0 \leq f(S \cup T) + f(S \cap T) \leq 3$, satisfying the 2-modularity condition.
 - (c) $f(T) = -1$. Given that $f(S) > 0$ and $f(T) < 0$ we have that $f(S \cup T) = f(S \cap T) = 0$, by Claim 4.8. Hence $f(S) + f(T) = f(S \cup T) + f(S \cap T)$ in this case.
3. $f(S) = f(T) = 0$. We need to show that $-2 \leq f(S \cup T) + f(S \cap T) \leq 2$. This might be violated only if $\max\{|f(S \cup T)|, |f(S \cap T)|\} = 2$. Lemma 4.7 implies that it suffices to consider the cases that $f(S \cup T) = 2$ or $f(S \cap T) = 2$ (by taking the complements of S and T if necessary and using that $f(\bar{S}) = f(S) = 0$ and similarly for $f(\bar{T})$).

- (a) $S \cup T \in \text{PS}$. Then necessarily $|S \cap T| < 35$ and $f(S \cap T) < 2$. We show that $f(S \cap T) \neq 1$, and then indeed $-2 \leq f(S \cup T) + f(S \cap T) \leq 2$. Given that $|S \cap T| < 35$, the only rule that might cause $f(S \cap T) = 1$ is that $S \cap T = S_1 \cap (S_2 \cup R)$, where $S_1, S_2 \in \text{PS}$ and $R \neq \emptyset$. But given that $k \geq 4$ and that $S \cup T$ is also in PS, it follows that $S_i = S \cup T$ for $i = 1$ or $i = 2$. (No set in PS contains the intersection of two other sets from PS.) So $S_1 \cap S_2 \subset S \cap T \subset S \subset S \cup T = S_i$. We conclude $f(S) = 1$, contradicting the assumption.
- (b) $S \cap T \in \text{PS}$. Similarly to Case (a), it is sufficient to show that $f(S \cup T) \neq 1$. The only relevant rule is $S \cup T = S_1 \cup (S_2 \cap R)$, but then $S_i = S \cap T$ for $i = 1$ or $i = 2$. (No union of sets in PS contains another set from PS.) So $S_i = S \cap T \subset S \subset S \cup T \subset S_1 \cup S_2$. We conclude $f(S) = 1$, contradicting the assumption.

Hence we established that f is 2-modular. ■

5 Learning Δ -Linear and ϵ -Modular Functions

Consider the following natural question: We have value-query access to a Δ -linear set function f . Our goal is to learn in polynomial-time a “hypothesis” linear set function h that is δ -close to f . How small can δ be as a function of Δ and n ? We say that an algorithm δ -learns a set function f if given value-query access to f it returns in polynomial-time a linear set function that is δ -close to f .

The work of Chierichetti et al. [8] (Theorem 4) presents an algorithm that $O(\epsilon\sqrt{n})$ -learns ϵ -modular set functions. Since every Δ -linear set function is 4Δ -modular (Proposition 2.4), their algorithm also $O(\Delta\sqrt{n})$ -learns Δ -linear set functions. The algorithm of [8] is randomized, and after making $O(n^2 \log n)$ non-adaptive queries to the function it is learning, returns a δ -close linear set function with probability $1 - o(1)$.

In this section we present an alternative algorithm for $O(\Delta\sqrt{n})$ -learning Δ -linear set functions (and since every ϵ -modular set function is $K_s\epsilon$ -linear for a constant K_s , also for $O(\epsilon\sqrt{n})$ -learning ϵ -modular set functions). Our algorithm (Algorithm 1) is simple, *deterministic*, and makes a *linear* number of non-adaptive queries to the function it is learning. A similar technique is applied in [9].

5.1 Learning with the Hadamard Basis

A *Hadamard basis* is an orthogonal basis $\{v_1, \dots, v_n\}$ of \mathbb{R}^n which consists of vectors in $\{\pm 1\}^n$. Let $\{y_1, \dots, y_n\} = \{v_1/\sqrt{n}, \dots, v_n/\sqrt{n}\}$, then $\{y_1, \dots, y_n\}$ is an orthonormal basis of \mathbb{R}^n . Easy recursive constructions of a Hadamard basis are known whenever n is a power of 2 [30] (and also for some other values of n [24]). For all values of n for which a Hadamard basis exists, and for every choice of one particular vector $v \in \{\pm 1\}^n$, we may assume that v is a member of the Hadamard basis. This can be enforced by taking the first vector v_1 in an arbitrary Hadamard basis of \mathbb{R}^n , and flipping $-$ in all vectors of the basis $-$ those coordinates in which v and v_1 do not agree.

For every set of items $S \subseteq [n]$, let $v_S \in \{0, 1\}^n$ be the indicator vector of S . Vector v_S can be written as a linear combination $\sum_{i=1}^n \lambda_i y_i$ of the basis vectors $\{y_1, \dots, y_n\}$ where $\lambda_i = v_S \cdot y_i$. We shall make use of the following property of orthonormal bases: Express v_S as $\sum_{i=1}^n \lambda_i y_i$; since $\{y_1, \dots, y_n\}$ is orthonormal, then by Parseval’s identity,

$$\sum_{i=1}^n (\lambda_i)^2 = |S|. \tag{15}$$

Given a basis vector v_i we denote by S_i the set of indices in which v_i has entries $+1$, and then \bar{S}_i is the set of indices in which v_i has entries -1 .

Throughout this section, $i \in [n]$ specifies the index of a basis vector (as in v_i or y_i), whereas to index other objects (such as coordinates of vectors, or items) we shall use j rather than i . When a vector is indexed by a set (such as v_S , or $v_{\{j\}}$) then the vector is the indicator vector (in $\{0, 1\}^n$) for the set.

Claim 5.1 (Expressing a linear set function using S_i, \bar{S}_i). *Consider a linear set function g where $g(S) = \sum_{j \in S} g_j + g(\emptyset)$ for every set $S \subseteq [n]$. Then for every v_S it holds that $g(S) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \lambda_i (g(S_i) - g(\bar{S}_i)) + g(\emptyset)$, where here $\lambda_1, \dots, \lambda_n$ are the unique coefficients for which $v_S = \sum_{i=1}^n \lambda_i y_i$. In particular $g_j = g(\{j\})$ can be obtained by substituting in the unique $\lambda_1, \dots, \lambda_n$ for which $v_{\{j\}} = \sum_{i=1}^n \lambda_i y_i$.*

Proof. Let g' be the additive function defined as $g'(S) = g(S) - g(\emptyset)$. Extend the domain of g' from $\{0, 1\}^n$ to \mathbb{R}^n , giving the additive function \tilde{g} over \mathbb{R}^n defined as follows: $\tilde{g}(x) = \sum_{j \in S} g_j x_j$ for every $x \in \mathbb{R}^n$. Observe that for every Hadamard basis vector v_i , $\tilde{g}(v_i) = g(S_i) - g(\bar{S}_i)$ (by definition of the sets S_i, \bar{S}_i). By additivity of \tilde{g} , for every set S whose indicator vector v_S can be expressed as $\sum_{i=1}^n \lambda_i y_i$, we have that $\tilde{g}(v_S) = \sum_{i=1}^n \lambda_i \tilde{g}(y_i) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \lambda_i \tilde{g}(v_i) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \lambda_i (g(S_i) - g(\bar{S}_i))$. Finally, $g(S) - g(\emptyset) = g(S) - g_0 = \tilde{g}(v_S)$ for every set S . ■

Lemma 5.2. *Let h and g be two linear set functions such that for every $i \in [n]$, $h(S_i) - h(\bar{S}_i)$ is $O(\Delta)$ -close to $g(S_i) - g(\bar{S}_i)$, and for every $S' \in \{\emptyset, U\}$, $h(S')$ is $O(\Delta)$ -close to $g(S')$. Then for every set $S \subseteq [n]$, $h(S)$ is $O(\Delta(1 + \sqrt{\min\{|S|, n - |S|\}}))$ -close to $g(S)$.*

Proof. The lemma clearly holds for $S \in \{\emptyset, U\}$, since in this case $O(\Delta(1 + \sqrt{\min\{|S|, n - |S|\}})) = O(\Delta \cdot 1)$. For the remainder of the proof assume $0 < |S| < n$. By Claim 5.1 applied to h and g , $h(S) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \lambda_i (h(S_i) - h(\bar{S}_i)) + h(\emptyset)$, and $g(S) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \lambda_i (g(S_i) - g(\bar{S}_i)) + g(\emptyset)$. For every $i \in [n]$, $h(S_i) - h(\bar{S}_i)$ estimates $g(S_i) - g(\bar{S}_i)$ up to an additive error of $O(\Delta)$, and $h(\emptyset)$ estimates $g(\emptyset)$ up to an additive error of $O(\Delta)$. Thus we get that $h(S)$ estimates $g(S)$ up to an additive error of $\frac{O(\Delta)}{\sqrt{n}} \sum_{i=1}^n |\lambda_i| + O(\Delta)$. By the Cauchy-Schwartz inequality we have $\sum_{i=1}^n |\lambda_i| \leq \sqrt{n \sum_{i=1}^n (\lambda_i)^2} = \sqrt{n|S|}$, where the last equality is by invoking (15). We conclude that an upper bound on the additive estimation error is

$$\frac{O(\Delta)}{\sqrt{n}} \sqrt{n|S|} + O(\Delta) = O(\Delta) \sqrt{|S|} + O(\Delta) = O(\Delta \sqrt{|S|}). \quad (16)$$

Since $h(S) = h(\emptyset) + h(U) - h(\bar{S})$ and $g(S) = g(\emptyset) + g(U) - g(\bar{S})$, and since we know from (16) that $h(\bar{S})$ estimates $g(\bar{S})$ up to an additive error of $O(\Delta \sqrt{|\bar{S}|}) = O(\Delta \sqrt{n - |S|})$, we also get that the maximum additive estimation error is

$$O(\Delta) + O(\Delta \sqrt{n - |S|}) = O(\Delta \sqrt{n - |S|}). \quad (17)$$

Taking the minimum among (16) and (17) completes the proof. ■

5.2 The Algorithm

For simplicity, we state the following theorem for values of n that are a power of 2. In Remark 5.5 we explain how to extend it beyond powers of 2. In the following Theorem, when referring to Algorithm 1, we mean Algorithm 1 run with the first Hadamard basis vector v_1 being the all-ones vector.

ALGORITHM 1: Using the Hadamard basis to learn a linear set function.

Input: Value-query access to a set function f .

Output: A linear set function h where $h(S) = h_0 + \sum_{j \in S} h_j$.

% Query linearly many values of f

query $f(\emptyset)$

for every $i \in [n]$ **do** % Consider the i th Hadamard basis vector v_i

 query $f(S_i)$ and $f(\bar{S}_i)$ % S_i, \bar{S}_i are sets (possibly empty) of the $+1, -1$ entries of v_i

end for

% Compute h

set $h_0 = f(\emptyset)$

for every $j \in [n]$ **do**

 express $v_{\{j\}}$ as $\sum_{i=1}^n \lambda_i y_i$ % y_i is the Hadamard basis vector v_i after normalization

 set $h_j = \frac{1}{\sqrt{n}} \sum_{i=1}^n \lambda_i (f(S_i) - f(\bar{S}_i))$

end for

Theorem 5.3. *Let n be a power of 2. Given value-query access to a Δ -linear set function f over n items, Algorithm 1 returns in polynomial time a linear set function h such that for every set $S \subseteq [n]$, $h(S)$ is $O(\Delta(1 + \sqrt{\min\{|S|, n - |S|\}}))$ -close to $f(S)$. Algorithm 1 thus $O(\Delta\sqrt{n})$ -learns Δ -linear set functions over n items.*

Proof. First note that Algorithm 1 clearly runs in polynomial time. By the construction of h in Algorithm 1 and by Claim 5.1, h is the function that assigns to every set $S = \sum_{i=1}^n \lambda_i y_i$ the value $h(S) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \lambda_i (f(S_i) - f(\bar{S}_i)) + f(\emptyset)$. So for every $i \in [n]$, $h(S_i) - h(\bar{S}_i) = f(S_i) - f(\bar{S}_i)$, and $h(\emptyset) = f(\emptyset)$. Let g be a linear set function Δ -close to f , then for every $i \in [n]$, $h(S_i) - h(\bar{S}_i)$ is 2Δ -close to $g(S_i) - g(\bar{S}_i)$, and $h(\emptyset)$ is Δ -close to $g(\emptyset)$. Recall that v_1 is the all-ones vector, which is the indicator vector of U . So $h(U)$ is also 2Δ -close to $g(U)$. Invoking Lemma 5.2, we get that for every set S , $h(S)$ is $O(\Delta(1 + \sqrt{\min\{|S|, n - |S|\}}))$ -close to $g(S)$. Since g is Δ -close to f , this completes the proof. ■

Theorem 5.3 is essentially the best possible, in the following strong sense: Corollary 23 of [8] shows that no algorithm (deterministic or randomized) that performs polynomially-many value queries can find a $o(\Delta\sqrt{n/\log n})$ -close linear set function, even if the queries are allowed to be adaptive. We include a proof sketch of this tightness result in Appendix D for completeness. Our tightness proof holds even for learning *monotone* Δ -linear set functions (where a set function f is monotone if for every $S \subseteq T$, $f(S) \leq f(T)$).

Remark 5.4 (LP-based approach). *We describe an alternative way to derive the linear function h : After querying $f(S)$ for $S = \emptyset$ and for $S = S_i, \bar{S}_i \forall i \in [n]$, one solves a linear program (LP). The $n + 1$ variables x_j of the LP are intended to have the value g_j for every j . The constraints are $f(S) - \Delta \leq \sum_{j \in S} x_j \leq f(S) + \Delta$ for every set S that was queried (thus one needs to know the value of Δ , which is a disadvantage of the LP-based approach). Let h be the linear set function defined by $c_j = x_j^*$ for every j where x^* is a feasible solution of the LP. We claim that h is $O(\Delta(1 + \sqrt{\min\{|S|, n - |S|\}}))$ -close to f : On each of the queried sets S , the values of h and g differ by at most 2Δ . Since we may assume that v_1 is the all-ones vector, the set $S = U$ is one of the queried sets. Invoking Lemma 5.2 and using that g is Δ -close to f shows the claim. The advantage of this LP-based approach is that additional constraints can easily be incorporated once more data of f is collected, potentially leading to better accuracy. Likewise, one can easily enforce desirable properties such as non-negativity on h (if g is nonnegative).*

Remark 5.5 (Beyond powers of 2). *Using either Algorithm 1 or the LP-based algorithm described in Remark 5.4, we can $O(\Delta\sqrt{n})$ -learn Δ -linear set functions over n items even when n is not a power of 2, as follows. Extend f to f' over $n' \geq n$ items U' , where n' is a power of 2 and $U \subseteq U'$, by setting $f'(S) = f(S \cap U)$ for every set $S \subseteq U'$. Extend g to g' over U' in the same way. Notice that the extended versions f', g' are still Δ -close.*

Given h' over U returned by Algorithm 1, we define h over U by setting $h(S) = h'(S)$ for every set $S \subseteq U$. The proof of Theorem 5.3 holds verbatim with the single following modification: the Hadamard basis vector $v_1 \in \mathbb{R}^{n'}$ is no longer the all-ones vector, but rather the vector that is $+1$ on the first n coordinates and -1 on the $n' - n$ auxiliary variables. This ensures that we can apply Lemma 5.2 to h instead of h' .

For the LP-based algorithm, we can formulate the LP with $n' + 1$ variables, but since we know there is a feasible solution in which $x_j = 0$ for every $j > n$ (namely $x = g$), we can add these constraints so that the resulting linear function h is over n items.

A Missing Proof from Section 2 (Preliminaries)

Proof. [Proof of Theorem 2.15] Let $\epsilon = 1$, and let f be the function given in Pawlik’s paper “Approximately Additive Set Functions” [25] over the items $X \uplus Y$. We present the function f and the proof for completeness (and due to typos and brevity in [25]). Let $f(S \cup T)$ denote the value of the union of a set $S \subseteq X$ and a set $T \subseteq Y$. Let X' (resp., Y') denote a non-empty proper subset of X (resp., Y). The function f is defined as follows: $f(\emptyset \cup \emptyset) = 0$, $f(X' \cup \emptyset) = 1$, $f(X \cup \emptyset) = 3$, $f(\emptyset \cup Y') = -1$, $f(X' \cup Y') = 0$, $f(X \cup Y') = 1$, $f(\emptyset \cup Y) = -3$, $f(X' \cup Y) = -1$, $f(X \cup Y) = 0$.

As claimed in [25], f is weakly 1-modular. We observe that f is 2-modular, but no better than 2 modular (consider two non-disjoint proper subsets of X whose union is X).

Let μ be the closest linear function to f . Consider a set $C = X - x + y$ ($x \in X, y \in Y$ will be specified below). By definition $f(X) = 3$ and $f(C) = 0$. We argue that f is no closer than $3/2$ -close to μ .

Assume for contradiction that f is $(3/2 - \delta)$ -close to μ for $\delta > 0$. We show there is large enough k for which this leads to contradiction (k is the number of items in X , which equals the number of items in Y).

First notice that $\mu(X) \geq 3/2 + \delta$ (otherwise f and μ are more than $3/2 - \delta$ apart on X). If it were to hold that $\mu(x) \leq \mu(y)$, we’d get that $\mu(C) = \mu(X) - \mu(x) + \mu(y) \geq 3/2 + \delta$ (i.e., f and μ are more than $3/2 - \delta$ apart on C). So $\mu(x) > \mu(y)$ for every x, y .

We know that $\mu(X)$ can’t be too big (since it can’t exceed $f(X) = 3$ by too much). For the same reason and since $f(Y) = -3$ by definition, $\mu(Y)$ is negative but can’t be too small. By letting k grow large, it becomes apparent that $\mu(x)$ cannot be bounded away from 0 from above, similarly $\mu(y)$ cannot be bounded away from 0 from below. Using $\mu(x) > \mu(y)$ we conclude that $\mu(x) \geq 0$ for every x and $\mu(y) \leq 0$ for every y , and that for most items $\mu(x) \rightarrow 0$ and $\mu(y) \rightarrow 0$.

Now choose x and y that minimize $|\mu(x)| + |\mu(y)|$. For large enough k we get that $\mu(x) - \mu(y) < \delta$, so $\mu(C) = \mu(X) - \mu(x) + \mu(y) > 3/2 + \delta - \delta > f(C) + 3/2 - \delta$, contradiction. The bounds asserted in the theorem follow. ■

B Missing Proof from Section 3 (Upper Bounds)

In this appendix we give a proof sketch for Lemma 3.1, establishing the existence of expanders with a range of parameters. Recall Stirling’s approximation, by which for every positive integer n

it holds that

$$\sqrt{2\pi n}n^{n+1/2}e^{-n} \leq n! \leq en^{n+1/2}e^{-n}.$$

In Claim B.1 we state a direct consequence of Stirling's approximation (a proof appears for completeness in Appendix B.2 of the online version; see [11]), which we then use to prove Lemma 3.1.

Claim B.1 (Stirling). *For every $c, d \in \mathbb{R}$ where $c > d > 0$, for every sufficiently large integer m such that cm and dm are integers,*

$$\binom{cm}{dm} = \left(\frac{c^c}{d^d(c-d)^{c-d}} \right)^m \Theta \left(\frac{1}{\sqrt{m}} \right).$$

Proof. [Proof of Lemma 3.1 – sketch based on an approach of [26]] We give the proof for $(\frac{1}{4}, 5, \frac{1}{2})$ -expanders. Existence for the other families in Lemma 3.1 follows from the same argument. Let $k = 2m$. Suppose that there are $4m$ left-hand side vertices and consequently $2m$ right-hand side vertices (because $\theta = \frac{1}{2}$). The edges are determined by making $r = 5$ copies of each left-hand side vertex, taking a random permutation over all $2kr = 20m$ copies, and connecting to the right-hand side in a round-robin fashion. The probability that there is a set of $2k\alpha = m$ left-hand side vertices with fewer than m neighbors is at most:

$$\begin{aligned} & \binom{2k}{2k\alpha} \binom{2k\theta}{2k\alpha} \frac{\binom{2kr\frac{\alpha}{\theta}}{2kr\alpha}}{\binom{2kr}{2kr\alpha}} = \binom{4m}{m} \binom{2m}{m} \frac{\binom{10m}{5m}}{\binom{20m}{5m}} \\ & = \left(\frac{4^4}{3^3 \cdot 1^1} \cdot \frac{2^2}{1^1 \cdot 1^1} \cdot \frac{10^{10}}{5^5 \cdot 5^5} \cdot \frac{5^5 \cdot 15^{15}}{20^{20}} \right)^m \Theta \left(\frac{1}{m} \right) \\ & = \left(\frac{27}{32} \right)^{4m} \Theta \left(\frac{1}{m} \right) < 1, \end{aligned}$$

where the second equality follows from Claim B.1, and the inequality holds by taking sufficiently large m to overcome the constants in $\Theta(1/m)$. Hence such expanders exist for every integer multiple k of sufficiently large k' . \blacksquare

C Lower Bound on K_w

We prove the lower bound stated in Theorem 1.3 on K_w .

Proof. [Proof of Second Part of Theorem 1.3] We show there exists a weakly 2-modular set function with $n = 20$ that is tightly 3-linear. Thus, $K_w \geq \frac{3}{2}$. Consider the following (k, M) -symmetric function f with $k = 3$ (hence with $n = \binom{6}{3} = 20$ items) and $M = 3$. The positive support sets (PS) form a canonical generating collection. The values of sets under f is determined by the first applicable rule (the last statement necessarily makes f well defined):

1. Each positive support set $S \in \text{PS}$ has value $f(S) = M = 3$.
2. If there is some set $P \in \text{PS}$ such that $S \subset P$ and for every set $N \in \text{NS}$ (where NS is the collection of negative support sets) it holds that $S \not\subset N$, then $f(S) = 1$. Likewise, if there is some set $P \in \text{PS}$ such that $P \subset S$ and for every set $N \in \text{NS}$ it holds that $N \not\subset S$, then $f(S) = 1$.
3. Enforcing the antisymmetry property of (k, M) -symmetric functions, we impose $-f(S) = f(\bar{S})$ and derive from this sets with negative value.

4. All other sets have value 0.

One can verify (proof omitted) that the function f defined above is indeed (k, M) -symmetric according to Definition 4.5. The reason why we choose $k = 3$ (and not smaller) is because we shall use the following claim.

Claim C.1. *For every two sets $P_1, P_2 \in \text{PS}$ and two sets $N_3, N_4 \in \text{NS}$ all the following hold:*

1. $P_1 \cap P_2 \not\subset N_3$.
2. $N_3 \cap N_4 \not\subset P_1$.
3. $P_1 \not\subset N_3 \cup N_4$.
4. $N_3 \not\subset P_1 \cup P_2$.

Proof. We prove only item 1, as the proofs of the remaining items are similar. Suppose without loss of generality that P_i (for $i \in \{1, 2\}$) contains those items whose vector representation in $\{\pm 1\}^{2k}$ has bit i set to 1, and that N_3 contains those items whose vector representation in $\{\pm 1\}^{2k}$ has bit 3 set to -1 . (If in N_3 it was bit 1 that was set to -1 , the proof would be immediate.) Then because $k \geq 3$ the set $P_1 \cap P_2$ contains a vector for which all three bits 1, 2 and 3 are set to 1, and hence is not in N_3 . ■

Proposition 4.6 implies that f is tightly 3-linear. It remains to show that f is weakly 2-modular. Consider two disjoint nonempty sets S and T . As $S \cap T = \emptyset$, we have that $f(S \cap T) = 0$. Hence to prove weak 2-modularity, one needs to show that $|f(S) + f(T) - f(S \cup T)| \leq 2$. Similar to the proof of Lemma 4.7 it can be shown that one can assume without loss of generality that $f(S) \geq |f(T)|$. We proceed by a case analysis.

1. $f(S) = 3$ (namely, $S \in \text{PS}$). In this case, because T is disjoint from S , we have that $T \subset \bar{S}$ where $\bar{S} \in \text{NS}$. The definition of f then implies that $f(T) \leq 0$. We consider three cases.
 - (a) $f(T) = 0$. Given that $T \subset \bar{S}$ and $\bar{S} \in \text{NS}$, for $f(T) = 0$ to hold there must be some set $P \in \text{PS}$ such that $T \subset P$. We claim that $f(S \cup T) = 1$ (which establishes weak 2-modularity, because $f(S) + f(T) = 3$). Suppose for the sake of contradiction that $f(S \cup T) \neq 1$. As $S \subset S \cup T$ and $S \in \text{PS}$, this can happen only if there is a set $N \in \text{NS}$ such that $N \subset S \cup T$. But then $N \cap \bar{S} \subset T \subset P$. Hence we found two sets in NS whose intersection lies in a set in PS, contradicting item 2 of Claim C.1.
 - (b) $f(T) = -1$. In this case $f(S) + f(T) = 2$. We also have that $f(S \cup T) \geq 0$ because $S \subset S \cup T$ and $S \in \text{PS}$. Hence regardless of the value of $f(S \cup T)$, weak 2-modularity holds.
 - (c) $f(T) = -3$. In this case $T = \bar{S}$ and $S \cup T = U$, implying that $f(S \cup T) = 0 = f(S) + f(T)$.
2. $f(S) = 1$. Observe that the only way by which it may happen that $f(S \cup T) < 0$ is if there is some $P \in \text{PS}$ such that $S \subset P$ (leading to $f(S) = 1$), and there is some $N \in \text{NS}$ such that $N \subset S \cup T$, and $f(S \cup T) = -1$. (If $S \cup T \subset N$ then also $S \subset N$ and it cannot be that $f(S) = 1$.)

- (a) $f(T) = 1$. We claim that $f(S \cup T) \geq 0$ (which implies that weak 2-modularity holds). Assume for the sake of contradiction that $f(S \cup T) = -1$. As noted above this implies that there is $P \in \text{PS}$ with $S \subset P$, another $P' \in \text{PS}$ with $T \subset P'$ (because we can interchange S and T), and $N \in \text{NS}$ such that $N \subset S \cup T$. But then $N \subset P \cup P'$, contradicting item 4 of Claim C.1.
 - (b) $f(T) = 0$. In this case $f(S) + f(T) = 1$, and regardless of the value of $f(S \cup T)$ (which lies in the range $[-1, 3]$), weak 2-modularity holds.
 - (c) $f(T) = -1$. We need to show that $|f(S \cup T)| \neq 3$. This follows from the fact that S is not contained in any set in NS, and T is not contained in any set in PS.
3. $f(S) = f(T) = 0$. We need to show that $|f(S \cup T)| \neq 3$. We show that $S \cup T \notin \text{PS}$ (and the proof that $S \cup T \notin \text{NS}$ is similar). Suppose for the sake of contradiction that $S \cup T = P$ with $P \in \text{PS}$. Then $S \subset P$, and the fact that $f(S) = 0$ implies that there is a set $N_1 \in \text{NS}$ such that $S \subset N_1$. Likewise, $T \subset N_2$ for some $N_2 \in \text{NS}$. Hence $P \subset N_1 \cup N_2$, contradicting item 3 of Claim C.1.

Hence we established that f is weakly 2-modular, completing the proof. ■

D Tightness of Learning Algorithm

The following theorem establishes tightness of the learning algorithm. Results similar to Theorem D.1 appear in the literature – see for example [29]. See also [8] (Corollary 23).

Theorem D.1. *For $\Delta \leq \sqrt{\log n/n}$ and $\delta = o(\sqrt{n/\log n})$, no δ -learning algorithm exists for Δ -linear set functions f , even if f is monotone, and even if one allows for unbounded computation time (but only polynomially-many value queries).*

Proof. [Proof (Sketch)] Consider a random linear function g , for which there is a random set T of cardinality $\frac{n}{2}$, in which each item has value $q = \Delta/\sqrt{n \log n}$, and the other items have value 0. Then $g(T) = (\Delta\sqrt{n})/(2\sqrt{\log n})$ (which is at most 1), whereas $g(\bar{T}) = 0$. Consider now a function f constructed as follows. Call a set S balanced if $||S \cap T| - \frac{|S|}{2}| \leq \sqrt{n \log n}$, positively unbalanced if $|S \cap T| - \frac{|S|}{2} \geq \sqrt{n \log n}$, and negatively unbalanced if $\frac{|S|}{2} - |S \cap T| \geq \sqrt{n \log n}$. Define $f(S) = \frac{|S|}{2}q$ for balanced sets, $f(S) = g(S) - \sqrt{n \log n}$ for positively unbalanced sets, and $f(S) = g(S) + \sqrt{n \log n}$ for negatively unbalanced sets. Observe that g approximates f because $q\sqrt{n \log n} = \Delta$. If there are only polynomially many queries then w.h.p., for every query S , the underlying set is balanced. Hence these queries are not informative in exposing T , and $\delta \geq \frac{1}{2}(f(T) - f(\bar{T})) = \frac{\Delta\sqrt{n}}{4\sqrt{\log n}} - \Delta$. ■

Acknowledgements

We thank Assaf Naor for helpful discussions and for directing us to the paper of [18]. We thank Moni Naor for pointing us to the result of [9]. We are grateful for the anonymous referees' numerous helpful comments that have greatly improved the paper.

References

- [1] M.-F. BALCAN AND N. J. A. HARVEY, *Learning submodular functions*, in Proceedings of the 43rd Annual ACM Symposium on Theory of Computing, 2011, pp. 793–802.
- [2] E. BALKANSKI, A. RUBINSTEIN, AND Y. SINGER, *The limitations of optimization from samples*, in Proceedings of the 48th Annual ACM Symposium on Theory of Computing, 2017, pp. 1016–1027.
- [3] A. BELLONI, T. LIANG, H. NARAYANAN, AND A. RAKHLIN, *Escaping the local minima via simulated annealing: Optimization of approximately convex functions*, in Proceedings of the 28th Conference on Learning Theory, 2015, pp. 240–265.
- [4] D. BERTSIMAS AND A. THIELE, *Robust and data-driven optimization: Modern decision making under uncertainty*, INFORMS PubsOnline, 2014, ch. 5, pp. 95–122. TutORials in Operations Research.
- [5] M. BLUM, M. LUBY, AND R. RUBINFELD, *Self-testing/correcting with applications to numerical problems*, J. Comput. Syst. Sci., 47 (1993), pp. 549–595.
- [6] A. V. BONDARENKO, A. PRYMAK, AND D. RADCHENKO, *On concentrators and related approximation constants*, J. Math. Anal. Appl., 402 (2013), pp. 234–241.
- [7] Y. A. BRUDNYI AND N. J. KALTON, *Polynomial approximation on convex subsets of R^n* , Constr. Approx., 16 (2000), pp. 161–199.
- [8] F. CHERICHETTI, A. DAS, A. DASGUPTA, AND R. KUMAR, *Approximate modularity*, in Proceedings of the 56th Symposium on Foundations of Computer Science, 2015, pp. 1143–1162.
- [9] C. DWORK AND S. YEKHANIN, *New efficient attacks on statistical disclosure control mechanisms*, in Proceedings of the 28th Annual International Cryptology Conference, 2008, pp. 469–480.
- [10] U. FEIGE, *On maximizing welfare when utility functions are subadditive*, SIAM J. Comput., 39 (2009), pp. 122–142.
- [11] U. FEIGE, M. FELDMAN, AND I. TALGAM-COHEN, *Approximate modularity revisited*. Supplemented version, available as version “v1” at <https://arxiv.org/abs/1612.02034v1>, 2016.
- [12] U. FEIGE AND R. IZSAK, *Welfare maximization and the supermodular degree*, in Proceedings of the 4th Innovations in Theoretical Computer Science, 2013, pp. 247–256.
- [13] M. X. GOEMANS, N. J. A. HARVEY, S. IWATA, AND V. S. MIRROKNI, *Approximating submodular functions everywhere*, in Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms, 2009, pp. 535–544.
- [14] A. HASSIDIM AND Y. SINGER, *Submodular optimization under noise*, in Proceedings of the 34th International Conference on Machine Learning, 2017, pp. 1069–1122.
- [15] D. H. HYERS, *On the stability of the linear functional equation*, PNAS, 27 (1941), pp. 222–224.

- [16] S. IWATA, L. FLEISCHER, AND S. FUJISHIGE, *A combinatorial strongly polynomial algorithm for minimizing submodular functions*, J. ACM, 48 (2001), pp. 761–777, <https://doi.org/10.1145/502090.502096>, <http://doi.acm.org/10.1145/502090.502096>.
- [17] S. M. JUNG, *Hyers-Ulam-Rassias stability of functional equations in nonlinear analysis*, Springer, 2011.
- [18] N. KALTON AND J. W. ROBERTS, *Uniformly exhaustive submeasures and nearly additive set functions*, Transactions of the American Mathematical Society, 278 (1983), pp. 803–816.
- [19] N. J. KALTON, *The three space problem for locally bounded F -spaces*, Compositio Math., 37 (1978), pp. 243–276.
- [20] A. KRAUSE AND V. CEVHER, *Submodular dictionary selection for sparse representation*, in Proceedings of the 27th International Conference on Machine Learning, 2010, pp. 567–574.
- [21] A. KRAUSE, A. P. SINGH, AND C. GUESTRIN, *Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies*, Journal of Machine Learning Research, 9 (2008), pp. 235–284.
- [22] B. LEHMANN, D. LEHMANN, AND N. NISAN, *Combinatorial auctions with decreasing marginal utilities*, Games and Economic Behavior, 55 (2006), pp. 270–296.
- [23] R. PAES LEME, *Gross substitutability: An algorithmic survey*. To appear in Games and Economic Behavior, 2017.
- [24] R. E. A. C. PALEY, *On orthogonal matrices*, Journal of Mathematics and Physics, 12 (1933), pp. 311–320.
- [25] B. PAWLIK, *Approximately additive set functions*, Colloquium Mathematicae, 54 (1987), pp. 163–164.
- [26] N. PIPPENGER, *Superconcentrators*, SIAM J. Comput., 6 (1977), pp. 298–304.
- [27] T. ROUGHGARDEN, I. TALGAM-COHEN, AND J. VONDRÁK, *When are welfare guarantees robust?*, in Proceedings of the 20th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, 2017, pp. 22:1–22:23.
- [28] A. SCHRIJVER, *A combinatorial algorithm minimizing submodular functions in strongly polynomial time*, Journal of Combinatorial Theory, 80 (2000), pp. 346–355.
- [29] Y. SINGER AND J. VONDRÁK, *Information-theoretic lower bounds for convex optimization with erroneous oracles*, in Proceedings of the 28th Annual Conference on Neural Information Processing Systems, 2015, pp. 3204–3212.
- [30] J. J. SYLVESTER, *Thoughts on inverse orthogonal matrices, simultaneous sign successions, and tessellated pavements in two or more colours, with applications to Newton’s rule, ornamental tile-work, and the theory of numbers*, Philosophical Magazine, 34 (1867), pp. 461–475.