

Gradual Small-Bias Sample Spaces

Avraham Ben-Aroya*

Gil Cohen†

Abstract

A (k, ε) -biased sample space is a distribution over $\{0, 1\}^n$ that ε -fools every nonempty linear test of size at most k . Since they were introduced by Naor and Naor [NN93], these sample spaces have become a central notion in theoretical computer science with a variety of applications.

When constructing such spaces, one usually attempts to minimize the seed length as a function of n, k and ε . Given such a construction, if we reverse the roles and consider a fixed seed length, then the smaller we pick k , the better the bound on the bias ε becomes. However, once the space is constructed we have a *single* bound on the bias of all tests of size at most k .

In this work we consider the problem of getting “more mileage” out of (k, ε) -biased sample spaces. Namely, we study a generalization of these sample spaces which we call *gradual* (k, ε) -biased sample spaces. Roughly speaking, these are sample spaces that ε -fool linear tests of size *exactly* k and moreover, the bound on the bias of linear tests of size $i \leq k$ decays as i gets smaller.

We show how to construct gradual (k, ε) -biased sample spaces of size comparable to the (non-gradual) spaces constructed by Alon et al. [AGHP92]. Our construction is based on the lossless expanders of Guruswami et al. [GUV09], combined with the Quadratic Character Construction of [AGHP92].

*Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot 76100, Israel. Email: avraham.ben-aroya@weizmann.ac.il.

†Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot 76100, Israel. Email: gil.cohen@weizmann.ac.il. Research supported by Israel Science Foundation (ISF) grant.

1 Introduction

An ε -biased sample space S over $\{0, 1\}^n$ is a sample space with the following property: for every nonempty $T \subseteq [n]$, the random variable $s_T \triangleq \bigoplus_{i \in T} s_i$, where s is sampled from S , has bias at most ε . In other words, a sample space is ε -biased if it ε -fools every nontrivial linear test. When it is not desired or not important to specify ε , one usually refers to such a sample space as a *small-bias sample space*.

The notion of a small-bias sample space was introduced in the seminal paper of Naor and Naor [NN93] and has become a fundamental notion in theoretical computer science, with a variety of applications [Nao92, BNS92, NN93, HPS93, AR94, MW02, BSSVW03, VW08, Vio09].

Several explicit constructions of small-bias sample spaces that attempt to minimize the sample space size in terms of n and ε are known [AGHP92, ABN⁺92, NN93, BT09]. These constructions give incomparable sizes. Unfortunately, all known constructions fall short from achieving sample spaces of size $O(n/\varepsilon^2)$, which are guaranteed to exist by a simple probabilistic argument. Another research direction, which this work falls into, studies variations and generalizations of small-bias sample spaces [AIK⁺90, RSW93, EGL⁺92, AM95, MST06, Shp06].

A relaxation of the notion of a small-bias sample space requires only that *small* linear tests will be fooled. Formally, a (k, ε) -biased sample space is a sample space S over $\{0, 1\}^n$ such that for every nonempty $T \subseteq [n]$ of size at most k , the random variable s_T has bias at most ε , where again s is sampled from S . The advantage of this relaxed notion is that fooling only small tests, rather than every nontrivial test, can be achieved by much smaller sample spaces. The original motivation for studying (k, ε) -biased sample spaces was to obtain almost k -wise independent random variables. However, (k, ε) -biased sample spaces had proved to be useful in their own right, and found several applications [SZ94, Raz05, CRS12].

Naor and Naor [NN93] gave a general method for constructing (k, ε) -biased sample spaces from ε -biased sample spaces. Their method yields (k, ε) -biased sample spaces that are exponentially smaller in terms of n than what is possible for ε -biased sample spaces. In terms of seed-length, they showed that a seed of length $O(\log k + \log \log n + \log \varepsilon^{-1})$ is sufficient in order to fool tests of size k , while it is known that a seed of length $\Omega(\log n + \log \varepsilon^{-1})$ is necessary in order to fool every nontrivial linear test (see, e.g. [AGHP92, Alo09]).

Gradual small-bias sample spaces. Consider two pairs (k_1, ε_1) and (k_2, ε_2) such that

$$s = \log k_1 + \log \varepsilon_1^{-1} = \log k_2 + \log \varepsilon_2^{-1}.$$

Potentially, one could hope that a seed of length $O(s + \log \log n)$ would be sufficient to ε_1 -fool tests of size k_1 and *simultaneously* to ε_2 -fool tests of size k_2 . In other words, we are considering a (k, ε) -biased sample space that has the following property: for tests of size $t < k$, the “spare” $\log k - \log t$ bits of the seed are utilized to reduce the bias. In this paper we initiate the study of such sample spaces, which have a better bound on the bias for smaller tests.

Definition 1.1. A sample space S over $\{0, 1\}^n$ is called gradual (k, ε) -biased if for every nonempty $T \subseteq [n]$ of size at most k ,

$$\left| \mathbb{E}_{s \sim S} \left[(-1)^{\sum_{i \in T} s_i} \right] \right| \leq \varepsilon \cdot \frac{|T|}{k}.$$

A few words about the definition are in order. First, note that when T is of size exactly k , the bound on the bias is ε , i.e., a gradual (k, ε) -biased sample space is, in particular, (k, ε) -biased. One may consider a more general definition, which allows an arbitrary decaying function as the bound on the bias (say, $\varepsilon \cdot$

$(|T|/k)^c$ for some constant c). We choose this function to be $\varepsilon \cdot |T|/k$ in the definition, though this choice is not very restricting. Indeed, in Section 4 we give two simple methods to amplify the decaying exponent, that is, methods that transform a gradual (k, ε) -biased sample space to a sample space with a stronger decaying function. Moreover, the two methods require only black-box access to the sample space.

1.1 Motivation

Why should we care about gradual small-bias sample spaces? For one, we believe that the notion is simply a natural strengthening of a (k, ε) -biased sample space, and as such, is interesting in its own right. Moreover, we believe that gradual small-bias sample spaces provide an example of a more general phenomenon, which we now explain. The entropy in the seed of a gradual small-bias sample space is utilized to the fullest. Namely, if we have prepared our sample space with a seed long enough to fool large linear tests, and in practice a small test is used, the extra entropy in the seed is not wasted, but is rather channeled towards reducing the bias of the test. Another example of this general phenomenon arises in the setting of randomness extraction. Roughly speaking, a (k, ε) extractor E is a randomized function that when applied on a distribution with min-entropy at least k , results in a distribution which is ε -close to uniform. When an extractor is fed with a distribution of much higher min-entropy, this extra entropy could potentially go to waste. However, there are extractors which siphon this entropy to reduce the error ε . The extractor that is based on a random walk on an expander is one such example.

Finally we observe that the Fourier spectrum of a gradual small-bias sample space has the following nice structure. The bound on the Fourier coefficients is stronger for coefficients in the lower levels. Although this observation is trivial, we feel that it provides another neat perspective on gradual small-bias sample spaces.

1.2 Main Result

The following theorem is our main result:

Theorem 1.2. *For any integers n and $k \leq n$, for any $\varepsilon > 0$, and for any constant $\delta > 0$ ¹ there exists an explicit construction of a gradual (k, ε) sample space of size*

$$m = O_\delta \left(\left(\frac{k}{\varepsilon} \right)^{2+\delta} + \left(\frac{\log n}{\log k} \right)^{2+4/\delta} k^{1+\delta} \right),$$

where the O_δ hides a multiplicative constant that depends only on δ .

Obviously, one can find a value for δ that minimizes m as a function of n, k and ε . However, when no assumptions are made on the relations between n, k and ε , the expression one would get is cumbersome and non-informative. Moreover, when conducting such minimization one can no longer ignore the multiplicative dependency in δ that is hidden under the big O_δ notation. We therefore choose to specify our bound in the more readable way presented above.

In the following table we consider three natural ranges for k in terms of n , and for those we give the minimum value of m with respect to δ . We also make a comparison with the size of the (non-gradual) (k, ε) -biased sample space from [AGHP92], which equals to $m_{\text{AGHP}} = (k\varepsilon^{-1} \log n)^2$. The comparison is meant to show that using our construction for gradual small-bias sample space, one does not pay much more in the sample space size for having a decaying bound on the bias.

¹In fact, the construction works without assuming δ is constant, and this assumption appears only to simplify the presentation of the theorem. See Theorem 3.1 for a more general statement.

Range of k	Sample space size	
	[AGHP92]	Theorem 1.2
$k = n^\gamma$, for any constant $\gamma \leq 1$	$(k/\varepsilon)^{2+o(1)}$	$O_\delta((k/\varepsilon)^{2+\delta})$ for any constant $\delta > 0$
$k = \log^c n$, for any constant $c \geq 6$, and $\varepsilon = \Omega(1)$	$(\log n)^{2(c+1)}$	$(\log n)^{2(c+1)+3}$
$k = O(1)$ and $\varepsilon = \log^{-c} n$, for any constant $c \geq 1/2$	$(\log n)^{2(c+1)}$	$O_c((\log n)^{2(c+1)})$

1.3 Informal Description of the Construction

Our high-level strategy is composed of two steps:

1. Obtaining a gradual (n, ε) -biased sample space, over $\{0, 1\}^n$.
2. Transforming it into a gradual (k, ε) -biased sample space with a shorter seed.

A similar approach was used by [NN93] to construct (non-gradual) (k, ε) -biased sample spaces. We now elaborate on each of the steps.

1.3.1 Gradual small-bias sample spaces from quadratic characters

For the first step, we use the Quadratic Character Construction of small-bias sample spaces of [AGHP92], which we now describe.

Let q be an odd prime power. Denote by \mathbb{F}_q the finite field with q elements. The quadratic character $\chi : \mathbb{F}_q \rightarrow \{-1, 0, 1\}$ is defined as

$$\chi(x) = \begin{cases} 0, & x = 0; \\ 1, & \exists y \in \mathbb{F}_q \setminus \{0\} \text{ such that } x = y^2; \\ -1, & \text{otherwise.} \end{cases}$$

The sample space in this construction consists of q strings, in correspondence with the elements of \mathbb{F}_q . A string in the sample space is composed of n bits, which are indexed by elements from some arbitrarily chosen set $I \subseteq \mathbb{F}_q$ of size $|I| = n$. For $i \in I$ and $x \in \mathbb{F}_q$, the i^{th} bit of the x^{th} string is given by $\chi(x + i)$.²

The bias of this construction for linear tests of size k is essentially the expectation of χ over the image of some degree k polynomial. Weil's Theorem (see Theorem 2.1) bounds precisely expectations of this form. Moreover, the bound this theorem provides is linear in k , the degree of the aforementioned polynomial. This implies a better bound for smaller tests. It follows that this space is indeed a gradual (n, ε) -biased sample space.

²Formally, the support of the sample space should be $\{0, 1\}^n$. This minor technicality is resolved by mapping ± 1 to $\{0, 1\}$, and 0 arbitrarily.

1.3.2 Shortening the seed length

In order to obtain a (k, ε) -biased sample space clearly it suffices to construct an (n, ε) -biased sample space (since being (n, ε) -biased implies being (k, ε) -biased). We now describe a cleverer way to transform (n, ε) -biased sample spaces into (k, ε) -biased ones. This transformation is due to [NN93].

Let U_n denote the uniform distribution over $\{0, 1\}^n$. We say that a linear transformation $T : \{0, 1\}^n \rightarrow \{0, 1\}^N$ generates a k -wise independent space if the N random variables $(Z_i)_{i=1}^N$ defined by

$$Z_i = T(U_n)_i$$

are k -wise independent. Suppose S is an (n, ε) -biased sample space over $\{0, 1\}^n$ and suppose that $T : \{0, 1\}^n \rightarrow \{0, 1\}^N$ generates a k -wise independent space. Then, in [NN93] it is shown that the sample space over $\{0, 1\}^N$ defined by $T(S)$ is (k, ε) -biased. The advantage of this transformation is that it allows N to be significantly larger than n , thus shortening the seed length as a function of the output length.

Similar to this approach, we also suggest a general way to transform a gradual (n, ε) -biased sample space into a gradual (k, ε) -biased one. The idea is to use a linear transformation $T : \{0, 1\}^n \rightarrow \{0, 1\}^N$, which generates a k -wise independent space, but which is also *sparse*, in the sense that each output bit depends only on a small number of input bits. We claim that in this case, provided that S is a gradual (n, ε) -biased sample space, $T(S)$ is a gradual (k, ε) -biased sample space.

Let us sketch the proof idea. Suppose T is such a transformation, in which each output bit is a sum of at most ℓ input bits, and suppose S is a gradual (n, ε) -biased sample space. To sample from the new sample space, we first sample s from S and then output $T(s)$. Consider a linear test $A(\cdot)$ of size $r \leq k$, applied to $T(s)$. By the sparsity of T , it follows that $A(T(s))$ is a sum of at most $\ell \cdot r$ bits of s . Moreover, since T generates a k -wise independent space, this sum is not empty. Thus, the bias of $A(T(s))$ is the bias of some test of weight $\ell \cdot r$ in the sample space S , and the claim follows. It might be useful to note that this transformation works even if S is a gradual $(\ell \cdot r, \varepsilon)$ -biased sample space (as apposed to S being a gradual (n, ε) -biased sample space).

We present an explicit construction of such a transformation T , based on expanders (see Section 2.2). The construction that we use is essentially the parity-check matrix of the codes of Sipser and Spielman [SS96] when combined with the unbalanced expanders of [GUV09].

1.4 Organization

In Section 2 we state some preliminary definitions and results that we need. In Section 3 we present a construction of a gradual small-bias sample space and prove Theorem 1.2. In Section 4 we address the problem of achieving a stronger decay in the bound on the bias. Section 5 contains concluding remarks and some open problems.

2 Preliminaries

All logarithms in this paper are in base 2. For a natural number n we define $[n] = \{1, 2, \dots, n\}$.

2.1 Quadratic Characters

We denote by χ_q the quadratic character over \mathbb{F}_q . When the field is understood from the context, we omit the subscript and simply denote this character by χ . We use a special case of Weil's Theorem regarding character sums (see e.g., [Sch76]).

Theorem 2.1 (Weil's Theorem). *Let q be an odd prime power. Let $f \in \mathbb{F}_q[x]$ be a degree d polynomial. Assume that $f(x) \neq c \cdot g(x)^2$ for any $c \in \mathbb{F}_q, g \in \mathbb{F}_q[x]$. Then,*

$$\left| \sum_{x \in \mathbb{F}_q} \chi(f(x)) \right| \leq (d-1)\sqrt{q}.$$

2.2 Expanders and Codes

We associate a bipartite graph $G = (L, R, E)$ with $\ell = |L|$ left-vertices and $r = |R|$ right-vertices and left-degree d with the adjacency function $G : L \times [d] \rightarrow R$, where $G(x, i) = y$ if and only if y is the i th neighbor of x . For a set of left-vertices $A \subseteq L$ we denote by $G(A)$ the set of neighbors of A .

Definition 2.2. *A bipartite graph $G : L \times [d] \rightarrow R$ is a k -unique-neighbor expander if for any nonempty subset $A \subseteq L$ of size at most k , there exists some $y \in R$ that is adjacent to exactly one vertex in A .*

Definition 2.3. *A bipartite graph $G : L \times [d] \rightarrow R$ is a $(\leq k, \alpha)$ expander if for any subset $A \subseteq L$ of size at most k ,*

$$|G(A)| \geq \alpha \cdot |A|.$$

We will need the well known fact that a graph whose expansion is greater than half of the degree is also a unique-neighbor expander.

Fact 2.4. *If $G : L \times [d] \rightarrow R$ is a $(\leq k, \alpha)$ expander for $\alpha > d/2$ then G is a k -unique-neighbor expander.*

Proof: Consider a nonempty set of left-vertices $A \subseteq L$ of size at most k . The number of outgoing edges from A is $|A| \cdot d$. Suppose that each vertex in $G(A)$ is adjacent to at least two vertices in A . This implies, $|A| \cdot d \geq 2 \cdot |G(A)|$, which contradicts the fact that G is a $(\leq k, \alpha)$ expander for $\alpha > d/2$. \square

We will make use of the following expanders, constructed by [GUV09].

Theorem 2.5 ([GUV09, Theorem 3.2]). *Let q be a prime power.³ For every integers $\ell, r, h \geq 1$ there exists an explicit construction of a graph $G : [q^\ell] \times [q] \rightarrow [q^{r+1}]$ which is an $(\leq h^r, q - (\ell - 1)(h - 1)r)$ expander. In particular, G is an h^r -unique-neighbor expander when $q > (\ell - 1)(h - 1)r/2$.*

In essence our construction uses the error-correcting code whose parity check matrix is defined by the above graph (as in [SS96]). In comparison, the (non-gradual) (k, ε) -biased sample spaces of [NN93, AGHP92] use the BCH code. Unlike the BCH code, the expander code is a low-density parity-check code, and this property plays a crucial role in our construction.

3 The Construction

In this section we describe our construction of a gradual (k, ε) -biased sample space, and prove Theorem 1.2. For simplicity we combine the two conceptual steps that appear in the informal description of the construction (Section 1.3).

³For this construction to be explicit, the characteristic of \mathbb{F}_q should be small characteristic. In our construction we can take it to be 3.

Let $r \geq 2$ be an integer.⁴ Let q be an odd prime power to be determined later. Set $\ell = \lceil \frac{\log n}{\log q} \rceil$. For the construction, we assume that we have a bipartite graph $G = (L, R, E)$ which is a k -unique-neighbor expander with $|L| = q^\ell$, $|R| = q^{r+1}$, and left-degree q . By our choice of ℓ we have $|L| \geq n$. Fix an arbitrary subset L' of L such that $|L'| = n$. Set $m = q^{r+1}$ and identify R with the finite field \mathbb{F}_m . For every vertex $v \in L'$ define the polynomial $p_v(x) \in \mathbb{F}_m[x]$ by

$$p_v(x) = \prod_{w : (v,w) \in E} (x - w).$$

We now describe the sample space S over $\{0, 1\}^n$.⁵ Each element in S corresponds to a field element in \mathbb{F}_m , that is, $S = \{s_x : x \in \mathbb{F}_m\}$. The string s_x is indexed by elements from the set L' . In particular, for every $x \in \mathbb{F}_m$ and $v \in L'$, we define

$$(s_x)_v = \begin{cases} \frac{1 - \chi_m(p_v(x))}{2}, & p_v(x) \neq 0; \\ 0, & \text{otherwise.} \end{cases} \quad (3.1)$$

The following theorem readily implies Theorem 1.2 by setting $\delta = 4/(r-1)$.

Theorem 3.1. *For every integers n, k, r such that $n \geq k$ and $r \geq 2$, and for any $\varepsilon > 0$, there is a way to choose q such that the construction defined above is an explicit gradual (k, ε) sample space over $\{0, 1\}^n$ with size*

$$m \leq \max \left\{ (10r^2)^{r+1} \left(\frac{\log n}{\log k} \right)^{r+1} k^{1+1/r}, 3^{r+1} \left(\frac{2k}{\varepsilon} \right)^{2+4/(r-1)} \right\}.$$

To prove Theorem 3.1 we prove the following two claims.

Claim 3.2. *If*

$$q \geq \left(\frac{2k}{\varepsilon} \right)^{2/(r-1)}$$

then the sample space defined above is gradual (k, ε) -biased.

Claim 3.3. *If*

$$q \geq 3.3 \cdot \frac{\log n}{\log k} \cdot k^{1/r} r^2,$$

then we have an explicit construction of the k -unique-neighbor expander graph $G = (L, R, E)$ required by the above construction.

Before proving the two claims we derive Theorem 3.1 from them. By choosing

$$q \geq \max \left\{ 3.3 \cdot \frac{\log n}{\log k} \cdot k^{1/r} r^2, \left(\frac{2k}{\varepsilon} \right)^{2/(r-1)} \right\}, \quad (3.2)$$

Claim 3.3 assures us that we can obtain the graph G that we need in the construction. Having this graph, Claim 3.2 guarantees that the above sample space is gradual (k, ε) -biased. Certainly one can efficiently find

⁴The parameter r is related to the parameter δ that appears in Theorem 1.2. In particular $r = 1 + 4/\delta$.

⁵In fact, we define S as a multi-set. The sample space is induced in the natural way, namely, to sample from the sample space, one sample an element $s \in S$ with probability proportional to the multiplicity of s in S .

a choice for $q = 3^z$ which is at most three times the right hand side of Equation (3.2).⁶ As $m = q^{r+1}$ we get the following upper bound on m , the sample space size

$$m \leq \max \left\{ (10r^2)^{r+1} \left(\frac{\log n}{\log k} \right)^{r+1} k^{1+1/r}, 3^{r+1} \left(\frac{2k}{\varepsilon} \right)^{2+4/(r-1)} \right\},$$

hence Theorem 3.1 follows.

Proof of Claim 3.2: Let $T \subseteq L'$ be a non-empty set of size at most k . Define

$$p_T(x) = \prod_{v \in T} p_v(x).$$

Since $p_T(x)$ is defined as a product of $|T|$ polynomials, each of degree at most q , we have that $\deg(p_T(x)) \leq q \cdot |T|$. Moreover, we claim that $p_T(x)$ has a simple root. Indeed, T is a nonempty set of size at most k of $L' \subseteq L$. By our assumption, G is a k -unique-neighbor expander, and so there exists a vertex $w \in R$ with exactly one neighbor, v , in T . This implies that w is a simple root of $p_v(x)$, while for every $u \in T \setminus \{v\}$, $p_u(w) \neq 0$. Hence, by the definition of $p_T(x)$ we have that w is a simple root of $p_T(x)$. Now, the bias of the linear test defined by T is

$$\sum_{x \in \mathbb{F}_m} (-1)^{\sum_{v \in T} (s_x)_v} = \sum_{x \in \mathbb{F}_m} \prod_{v \in T} (-1)^{(s_x)_v}. \quad (3.3)$$

Suppose x is not a root of $p_T(x)$. Then the value such an x contributes to the sum in Equation (3.3) is

$$\prod_{v \in T} (-1)^{(s_x)_v} = \prod_{v \in T} \chi_m(p_v(x)) = \chi_m \left(\prod_{v \in T} p_v(x) \right) = \chi_m(p_T(x)),$$

where the middle equality follows from the fact that χ is a multiplicative homomorphism. As $p_T(x)$ has at most $\deg(p_T) \leq q \cdot |T|$ roots, we have that

$$\left| \sum_{x \in \mathbb{F}_m} (-1)^{\sum_{v \in T} (s_x)_v} \right| \leq \left| \sum_{x \in \mathbb{F}_m} \chi_m(p_T(x)) \right| + q \cdot |T|$$

Since $p_T(x)$ has a simple root, $p_T(x)$ is not of the form $c \cdot g(x)^2$ for any $c \in \mathbb{F}_m$ and $g \in \mathbb{F}_m[x]$. Therefore, we can apply Weil's Theorem (Theorem 2.1) to get

$$\left| \sum_{x \in \mathbb{F}_m} \chi_m(p_T(x)) \right| < q \cdot |T| \cdot \sqrt{m}.$$

Hence,

$$\frac{1}{m} \left| \sum_{x \in \mathbb{F}_m} (-1)^{\sum_{v \in T} (s_x)_v} \right| \leq \frac{2q \cdot |T|}{\sqrt{m}} = 2|T| \cdot q^{(1-r)/2}.$$

⁶Observe also that this solves the minor issue regarding the need for small characteristic for the explicitness requirements of Theorem 2.5.

To get a bound of at most ε on the bias for tests of size exactly k , we require that

$$2k \cdot q^{(1-r)/2} \leq \varepsilon,$$

or

$$q \geq \left(\frac{2k}{\varepsilon}\right)^{2/(r-1)}. \quad (3.4)$$

□

Proof of Claim 3.3: We use the expanders from Theorem 2.5 with $h = \lceil k^{1/r} \rceil$. If

$$q - (\ell - 1)(h - 1)r \geq 0.51q$$

then G is a k -unique-neighbor expander. By the definition of ℓ , for the above equation to hold, it is enough to require

$$q \log q \geq 2.05 \cdot \log n \cdot k^{1/r} r. \quad (3.5)$$

We use the following simple claim that can be easily verified.

Claim 3.4. For every $x, y > 1$, if

$$x \geq 1.6 \cdot \frac{y}{\log y}$$

then $x \log x \geq y$.

By the Claim 3.4, for equation (3.5) to hold, it is enough to require that

$$q \geq 1.6 \cdot \frac{2.05 \cdot \log n \cdot k^{1/r} r}{\log(k^{1/r})} = 3.28 \cdot \frac{\log n}{\log k} \cdot k^{1/r} r^2,$$

which concludes the proof. □

4 Amplifying the Decay Exponent

The definition of a gradual (k, ε) -biased sample space requires a bound of the form $\varepsilon \cdot |T|/k$ on the bias for any nonempty set T of size at most k . The construction we suggest in this paper indeed has such a linear decay. However, an easy application of the probabilistic method shows that a random sample space S over $\{0, 1\}^n$ of size $m = O(k\varepsilon^{-2} \log n)$ has the following property with high probability: for every nonempty $T \subseteq [n]$ of size at most k ,

$$\left| \mathbb{E}_{s \sim S} \left[(-1)^{\sum_{i \in T} s_i} \right] \right| \leq \varepsilon \cdot \sqrt{\frac{|T|}{k}}.$$

We note that the bound has a non-linear dependency in $|T|/k$. This might suggest that Definition 1.1 should be generalized in order to model the spaces we want to study – there is nothing special about a linear decay rule, besides the fact that our construction obeys it.

In this section we present two simple methods for converting a gradual small-bias sample space S into a sample space with a larger decaying exponent. The two methods need only a black-box access to S . This suggests that the linear decay in the definition is not restricting in the sense that such a sample space can be converted in a black-box manner to one that has a stronger decay function.

Of course, one must pay in the sample space size in order to get a stronger decay. The two methods we suggest give incomparable sample space sizes. One is better than the other depending on how the size of S depends on n, k and ε . Starting with our construction, the two methods give roughly the same sample space size.

For simplicity, we show how to amplify the decay exponent from 1 to 2, that is, we describe two methods for transforming a gradual small-bias sample space S to a sample space that has a bound of the form $\varepsilon \cdot (|T|/k)^2$ on the bias, for nonempty sets of size at most k . It is straightforward to generalize these methods to exponents other than 2.

The first method is based on the following trivial observation: if S is a gradual $(k, \varepsilon/k)$ -biased sample space, then

$$\left| \mathbb{E}_{s \sim S} \left[(-1)^{\sum_{i \in T} s_i} \right] \right| \leq \varepsilon \cdot \frac{|T|}{k^2} \leq \varepsilon \cdot \left(\frac{|T|}{k} \right)^2.$$

That is, choosing a smaller error to begin with, will result in a better decaying function, and in fact, even better than what we are aiming for (the $|T|$ is not squared). For constructions where the dependency in ε^{-1} is small, and this is the case in our construction, the above method is quite effective. However, for a construction that suffers a worse dependency on ε^{-1} , the following method would be preferred: given a gradual small-bias sample space S , use the sample space $S + S$.⁷ More formally,

Lemma 4.1. *Let S be a gradual $(k, \sqrt{\varepsilon})$ -biased sample space. Then,*

$$\left| \mathbb{E}_{s \sim S+S} \left[(-1)^{\sum_{i \in T} s_i} \right] \right| \leq \varepsilon \cdot \left(\frac{|T|}{k} \right)^2.$$

Proof: For a sample space X , define a function $p_X : \{0, 1\}^n \rightarrow \mathbb{R}$ by

$$p_X(x) = \Pr[X = x].$$

Then, for every $T \subseteq [n]$,

$$\widehat{p}_S(T) = 2^{-n} \cdot \mathbb{E}_{s \sim S} \left[(-1)^{\sum_{i \in T} s_i} \right].$$

By basic Fourier analysis [O'D]

$$p_{S+S} = 2^n \cdot p_S * p_S,$$

and so

$$\widehat{p}_{S+S}(T) = 2^n \cdot \widehat{p}_S * \widehat{p}_S(T) = 2^n \cdot \widehat{p}_S(T)^2.$$

Hence,

$$\left| \mathbb{E}_{s \sim S+S} \left[(-1)^{\sum_{i \in T} s_i} \right] \right| = \left(\mathbb{E}_{s \sim S} \left[(-1)^{\sum_{i \in T} s_i} \right] \right)^2 \leq \left(\sqrt{\varepsilon} \cdot \frac{|T|}{k} \right)^2 = \varepsilon \cdot \left(\frac{|T|}{k} \right)^2.$$

□

⁷The sample space $S + S$ is defined by sampling s_1 and s_2 , independently, from S and then outputting $s_1 \oplus s_2$.

5 Concluding Remarks and Open Problems

Our method for transforming a gradual ε -biased sample space into a gradual (k, ε) -biased sample space uses, as a black box, the unbalanced expanders of [GUV09]. Hence, improved constructions of unbalanced expanders, or low-density parity-check codes in general, may lead to improved constructions of gradual (k, ε) -biased sample spaces. Indeed, our general method has the potential to generate very good gradual (k, ε) -biased sample spaces from the Quadratic Characters Construction given better constructions of unbalanced expanders. Specifically, using the unbalanced expanders given by the probabilistic construction (see, e.g., [GUV09]), our method yields a gradual (k, ε) -biased sample space of size $O((k\varepsilon^{-1} \log n)^2)$. This is as good as the non-gradual (k, ε) -biased sample space of [AGHP92]. It would therefore be interesting to construct a gradual small-bias sample space that matches the parameters of the non-gradual sample space of [AGHP92].

For non-gradual small-bias sample spaces there are a few explicit constructions with incomparable size [AGHP92, ABN⁺92, NN93, BT09]. Finding an explicit construction of a gradual small-bias sample space with better (or incomparable) size to ours is therefore a natural research goal. One possible route towards this goal is to construct a gradual small-bias sample space that has an incomparable size with that of the Quadratic Character Construction. We are not aware of such a construction in the literature.

The original motivation for studying (non-gradual) (k, ε) -biased sample spaces was to construct a sample space S that is almost k -wise independent. Using a gradual (k, ε) -biased sample space instead of the non-gradual one improves the size of S by a mere multiplicative constant factor. Nevertheless, we hope that applications that exploit the gradual bound on the bias would be found.

Acknowledgements

The second author would like to thank his advisor Ran Raz for his continuous support and encouragement.

References

- [ABN⁺92] N. Alon, J. Bruck, J. Naor, M. Naor, and R. Roth. Construction of asymptotically good low-rate error-correcting codes through pseudo-random graphs. *IEEE Transactions on Information Theory*, 38:509–516, 1992.
- [AGHP92] N. Alon, O. Goldreich, J. Håstad, and R. Peralta. Simple construction of almost k -wise independent random variables. *Random Structures and Algorithms*, 3(3):289–304, 1992.
- [AIK⁺90] M. Ajtai, H. Iwaniec, J. Komlos, J. Pintz, and E. Szemerédi. Construction of a thin set with small fourier coefficients. *Bulletin of the London Mathematical Society*, 22:583–590, 1990.
- [Alo09] N. Alon. Perturbed identity matrices have high rank: proof and applications. *Combinatorics, Probability and Computing*, 18(1-2):3–15, 2009.
- [AM95] N. Alon and Y. Mansour. epsilon-discrepancy sets and their application for interpolation of sparse polynomials. *Information Processing Letters*, 54(6):337–342, 1995.
- [AR94] N. Alon and Y. Roichman. Random cayley graphs and expanders. *Random Structures and Algorithms*, 5(2):271–285, 1994.

- [BNS92] Laslo Babai, Naom Nisan, and Mario Szegedy. Multipart protocols, pseudo-random generators for logspace, and time-space trade-offs. *J. of Computer and System Sciences*, 45(2):204–232, 1992.
- [BSSVW03] E. Ben-Sasson, M. Sudan, S. Vadhan, and A. Wigderson. Randomness-efficient low degree tests and short PCPs via epsilon-biased sets. In *Proceedings of the 35th STOC*, pages 612–621, 2003.
- [BT09] A. Ben-Aroya and A. Ta-Shma. Constructing small-bias sets from algebraic-geometric codes. In *Proceedings of the 50th annual IEEE symposium on foundations of computer science (FOCS)*, 2009.
- [CRS12] G. Cohen, R. Raz, and G. Segev. Non-malleable extractors with short seeds and applications to privacy amplification. In *Proceedings of the 27rd Annual CCC*, 2012.
- [EGL⁺92] G. Even, O. Goldreich, M. Luby, N. Nisan, and B. Velickovic. Approximations of general independent distributions. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, pages 10–16, 1992.
- [GUV09] V. Guruswami, C. Umans, and S. Vadhan. Unbalanced expanders and randomness extractors from Parvaresh–Vardy codes. *J. ACM*, 56(4):1–34, 2009.
- [HPS93] J. Håstad, S. Phillips, and S. Safra. A well-characterized approximation problem. *Information Processing Letters*, 47(6):301–305, 1993.
- [MST06] E. Mossel, A. Shpilka, and L. Trevisan. On epsilon-biased generators in NC^0 . *Random Structures and Algorithms*, 29(1):56–81, 2006.
- [MW02] R. Meshulam and A. Wigderson. Expanders from symmetric codes. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 669–677, 2002.
- [Nao92] M. Naor. Constructing Ramsey graphs from small probability spaces. Technical report, IBM Research Report RJ 8810, 1992.
- [NN93] J. Naor and M. Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM J. on Computing*, 22(4):838–856, 1993.
- [O’D] R. O’Donnell. Analysis of boolean functions. <http://analysisofbooleanfunctions.org/>.
- [Raz05] R. Raz. Extractors with weak random seeds. In *Proceedings of the 37th Annual STOC*, pages 11–20, 2005.
- [RSW93] A. A. Razborov, E. Szemerédi, and A. Wigderson. Constructing small sets that are uniform in arithmetic progressions. *Combinatorics, Probability & Computing*, 2:513–518, 1993.
- [Sch76] W. M. Schmidt. *Equations over Finite Fields: An elementary approach*. Springer-Verlag, 1976.
- [Shp06] A. Shpilka. Constructions of low-degree and error-correcting in-biased generators. In *21st Annual IEEE Conference on Computational Complexity*, pages 33–45, 2006.

- [SS96] M. Sipser and D. Spielman. Expander codes. *Information Theory, IEEE Transactions on*, 42(6):1710–1722, 1996.
- [SZ94] A. Srinivasan and D. Zuckerman. Computing with very weak random sources. In *35th Annual Symposium on Foundations of Computer Science*, pages 264–275. IEEE, 1994.
- [Vio09] E. Viola. The sum of small-bias generators fools polynomials of degree d . *Computational Complexity*, 18(2):209–217, 2009.
- [VW08] E. Viola and A. Wigderson. Norms, xor lemmas, and lower bounds for $GF(2)$ polynomials and multiparty protocols. *Theory of Computing*, (4):137–168, 2008.