

On Approximability of Satisfiable k -CSPs: II

Amey Bhangale*

Subhash Khot[†]Dor Minzer[‡]

Abstract

Let Σ be an alphabet and μ be a distribution on Σ^k for some $k \geq 2$. Let $\alpha > 0$ be the minimum probability of a tuple in the support of μ (denoted by $\text{supp}(\mu)$).¹ We treat the parameters Σ, k, μ, α as fixed and constant.

We say that the distribution μ has a linear embedding if there exist an Abelian group G (with the identity element 0_G) and mappings $\sigma_i : \Sigma \rightarrow G$, $1 \leq i \leq k$, such that at least one of the mappings is non-constant and for every $(a_1, a_2, \dots, a_k) \in \text{supp}(\mu)$, $\sum_{i=1}^k \sigma_i(a_i) = 0_G$. In [5], the authors asked the following analytical question.

Let $f_i : \Sigma^n \rightarrow [-1, 1]$ be bounded functions, such that at least one of the functions f_i essentially has degree at least d , meaning that the Fourier mass of f_i on terms of degree less than d is negligible, say at most δ .² If μ has no linear embedding (over any Abelian group), then is it necessarily the case that

$$\left| \mathbb{E}_{(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k) \sim \mu^{\otimes n}} [f_1(\mathbf{x}_1) f_2(\mathbf{x}_2) \cdots f_k(\mathbf{x}_k)] \right| = o_{d, \delta}(1),$$

where the right hand side $\rightarrow 0$ as the degree $d \rightarrow \infty$ and $\delta \rightarrow 0$?

In this paper, we answer this analytical question fully and in the affirmative for $k = 3$. We also show the following two applications of the result.

1. The first application is related to hardness of approximation. Using the reduction from [5], we show that for every 3-ary predicate $P : \Sigma^3 \rightarrow \{0, 1\}$ such that P has no linear embedding, an *SDP integrality gap instance* of a P -CSP instance with gap $(1, s)$ can be translated into a dictatorship test with completeness 1 and soundness $s + o(1)$, under certain additional conditions on the instance.
2. The second application is related to additive combinatorics. We show that if the distribution μ on Σ^3 has no linear embedding, marginals of μ are uniform on Σ , and $(a, a, a) \in \text{supp}(\mu)$ for every $a \in \Sigma$, then every large enough subset of Σ^n contains a triple $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$ from $\mu^{\otimes n}$ (and in fact a significant density of such triples).

*Department of Computer Science and Engineering, University of California, Riverside. Email: ameyb@ucr.edu.

[†]Department of Computer Science, Courant Institute of Mathematical Sciences, New York University. Supported by the NSF Award CCF-1422159, NSF Award CCF-2130816, and the Simons Investigator Award. Email: khot@cs.nyu.edu.

[‡]Department of Mathematics, Massachusetts Institute of Technology. Supported by a Sloan Research Fellowship, NSF CCF award 2227876 and NSF CAREER award 2239160. Email: dminzer@mit.edu.

¹The support of μ is the set of all tuples in Σ^k that have a positive probability mass under μ .

²In particular, $|\mathbb{E}[f_i]| \leq \delta$. The Fourier representation is w.r.t. the marginal of μ on the i^{th} co-ordinate, denoted (Σ, μ_i) .

1 Introduction

The motivation for this paper is to study the following quantity associated with the product of functions $f_1, f_2, \dots, f_k : \Sigma^n \rightarrow \mathbb{R}$,

$$\mathbb{E}_{(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k) \sim \mu^{\otimes n}} [f_1(\mathbf{x}_1) f_2(\mathbf{x}_2) \cdots f_k(\mathbf{x}_k)], \quad (1)$$

where each coordinate of $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k)$ is distributed independently, according to the same distribution μ on Σ^k . We assume that all the functions are bounded, i.e., $\|f_i\|_\infty \leq 1$. This expression appears naturally in many areas including additive combinatorics, social choice, pseudorandomness and hardness of approximation. Here are a few examples.

1. Example 1: For $1 \leq i \leq 3$, let $f_i : \mathbb{Z}_p^n \rightarrow \{0, 1\}$ be the indicator functions of the sets $A_i \subseteq \mathbb{Z}_p^n$. Let μ be the uniform distribution on the three-term arithmetic progressions $(x, x + y, x + 2y)$ in \mathbb{Z}_p . Then the quantity $\mathbb{E}_{(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3) \sim \mu^{\otimes n}} [f_1(\mathbf{x}_1) f_2(\mathbf{x}_2) f_3(\mathbf{x}_3)]$, up to a normalization factor, precisely counts the number of arithmetic progressions $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$ from \mathbb{Z}_p^n such that $\mathbf{x}_i \in A_i$ for every $i \in [3]$.
2. Example 2: Consider a Boolean function $f : \{-1, +1\}^n \rightarrow \{-1, +1\}$. For a given $\rho \in [-1, 1]$, the stability of f , $\text{Stab}_\rho(f)$, is defined as $\mathbb{E}[f(\mathbf{x})f(\mathbf{y})]$ where for each $i \in [n]$, \mathbf{x}_i and \mathbf{y}_i are uniformly distributed, and $\mathbb{E}[\mathbf{x}_i \mathbf{y}_i] = \rho$. The Majority is Stablest Theorem [18], which is instrumental in the area of hardness of approximation and the theory of social choice, is about estimating $\text{Stab}_\rho(f)$ for the class of so-called low-influence functions.
3. Example 3: Fix a predicate $P : \Sigma^k \rightarrow \{0, 1\}$ and a distribution μ on Σ^k . Dictatorship tests corresponding to a predicate P and a distribution μ are extensively studied in hardness of approximation. Here, one is given a function $f : \Sigma^n \rightarrow \Sigma$ and the acceptance probability of the test is precisely

$$\Pr_{(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k) \sim \mu^{\otimes n}} [(f(\mathbf{x}_1), f(\mathbf{x}_2), \dots, f(\mathbf{x}_k)) \in P^{-1}(1)].$$

One is interested in estimating this probability for the class of low influence functions. Using the multilinear expansions of P and f , the above expectation can be expressed as a linear combination of expectations of the form (1).

Let $c = \Pr_{(a_1, a_2, \dots, a_k) \sim \mu} [(a_1, a_2, \dots, a_k) \in P^{-1}(1)]$. It is seen that the test accepts any Dictatorship function, namely $f(\mathbf{x}) = \mathbf{x}_{i_0}$ for a fixed co-ordinate $i_0 \in [n]$, with probability c . While tests with imperfect completeness, namely with $c < 1$, are interesting and well-studied in hardness of approximation,³ in the current paper, we exclusively focus on tests with perfect completeness, namely with $c = 1$. That is, we assume that $\text{supp}(\mu) \subseteq P^{-1}(1)$. In fact, we will generally assume that μ has full support, i.e. $\text{supp}(\mu) = P^{-1}(1)$ and then talk interchangeably in terms of either the predicate P or the distribution μ . In terms of hardness of approximation, this amounts to studying approximability of Constraint Satisfaction Problems (CSPs) on (fully) satisfiable instances, and this indeed has been the main motivation for authors' work in [5], continuing in the current paper.

One way to analyze the expectation from (1) is to write each function f_i as the sum of two functions $g_i + h_i$, where g_i is the structured part of f_i and h_i is the remaining unstructured part (resembling noise). The idea is that whenever the term h_i appears in the product of functions, then the expectation is negligible. Therefore, the expectation can be estimated by replacing each f_i by its structured part g_i . For instance,

³Indeed, Example 2 corresponds to the hardness of approximation result for the Max-Cut problem. Here the predicate is $x \neq y$ over a binary alphabet, μ is the ρ -correlated distribution on $\{-1, 1\}^2$ as mentioned, completeness $c = \frac{1-\rho}{2}$, and $-1 < \rho < 0$.

in Example 1, Roth's Theorem [23] estimates the desired density of arithmetic progressions; therein, the structured part is taken as all the heavy-weight Fourier terms of f_i . It is shown that the contribution of the unstructured part is negligible; formally, if we let \hat{f}_i be the Fourier terms of f_i , then we have

$$\left| \mathbb{E}_{(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3) \sim \mu^{\otimes n}} [f_1(\mathbf{x}_1) f_2(\mathbf{x}_2) f_3(\mathbf{x}_3)] \right| \leq \min_{1 \leq i \leq 3} \|\hat{f}_i\|_\infty.$$

On the other hand, it is often useful (especially in hardness of approximation) to take the structured part as the low-degree part of f_i . In this case, after replacing the functions f_i by their low degree parts g_i , provided that g_i are low influence functions, it is possible to estimate the expectation well using invariance principles. Here, one replaces the discrete inputs from Σ^n by Gaussian inputs and then the expectation is estimated using bounds in the Gaussian space. Still, the question remains as to when one can argue that the expectation is negligible for the unstructured, i.e. the high-degree, part of the functions. Specifically, one is naturally led to the following analytic question.

Question 1.1. (Informal) *Find the necessary and sufficient condition on the distribution μ on Σ^k , such that*

$$\left| \mathbb{E}_{(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k) \sim \mu^{\otimes n}} [f_1(\mathbf{x}_1) f_2(\mathbf{x}_2) \cdots f_k(\mathbf{x}_k)] \right| \rightarrow 0 \quad \text{as } d \rightarrow \infty, \quad (2)$$

where the functions are bounded in $[-1, 1]$ and at least one function (essentially) has degree at least d .

Mossel [17] showed a sufficient condition: if the distribution μ is *connected*, then Conclusion (2) as above holds. The connectedness condition is defined as follows: for every pair of tuples $(a_1, a_2, \dots, a_k) \in \text{supp}(\mu)$ and $(a'_1, a'_2, \dots, a'_k) \in \text{supp}(\mu)$, there is a way to convert the first tuple to the second by replacing only one coordinate at a time such that every intermediate tuple remains in $\text{supp}(\mu)$.

The connectedness condition however is not necessary. An example is noted implicitly in [4]. Let G be a non-Abelian group with no dimension one representation. Consider the group-equation predicate $P : G^3 \rightarrow \{0, 1\}$, $P^{-1}(1) = \{(x, y, z) \mid x \cdot y \cdot z = 1_G\}$, along with the distribution μ that is uniform on $P^{-1}(1)$. The distribution μ is (clearly) not connected and Conclusion (2) still holds as can be shown using basic representation theory.

A certain necessary condition was observed in [5] (for Conclusion (2) to hold), namely that the distribution μ has no linear embedding as defined below. To illustrate that this condition is necessary, one considers the contra-positive: if the distribution μ does have a linear embedding (in particular, it is not connected), then there do exist high-degree, bounded functions that make the expectation in (2) non-negligible.

Definition 1.2. *We say that a distribution μ on Σ^k has a linear embedding (or that μ satisfies a linear equation or simply that μ is linear) if there exists an Abelian group G and mappings $\sigma_i : \Sigma \rightarrow G$, $1 \leq i \leq k$, such that (i) at least one of the maps σ_i is non-constant and (ii) for every $(a_1, a_2, \dots, a_k) \in \text{supp}(\mu)$, $\sum_{i=1}^k \sigma_i(a_i) = 0_G$.*

The illustration is as follows. Suppose μ does have a linear embedding as in the definition. We show that it is possible to achieve non-negligible expectation in (2). To see this, let χ be any non-trivial character of the Abelian group G , namely a non-trivial group homomorphism $\chi : G \rightarrow \mathbb{C}$, and define $f_i(\mathbf{x}_i) =$

$\prod_{j=1}^n \chi(\sigma_i((\mathbf{x}_i)_j))$. Now,

$$\begin{aligned}
f_1(\mathbf{x}_1) f_2(\mathbf{x}_2) \cdots f_k(\mathbf{x}_k) &= \prod_{i=1}^k \prod_{j=1}^n \chi(\sigma_i((\mathbf{x}_i)_j)) \\
&= \prod_{j=1}^n \prod_{i=1}^k \chi(\sigma_i((\mathbf{x}_i)_j)) \\
&= \prod_{j=1}^n \chi\left(\sum_{i=1}^k \sigma_i((\mathbf{x}_i)_j)\right) \\
&= \prod_{j=1}^n \chi(0_G) \\
&= 1.
\end{aligned}$$

Here one uses the multiplicativity of the character χ and that $\chi(0_G) = 1$. For every $1 \leq j \leq n$, we have $\sum_{i=1}^k \sigma_i((\mathbf{x}_i)_j) = 0_G$ noting that the tuple $((\mathbf{x}_1)_j, \dots, (\mathbf{x}_k)_j) \in \text{supp}(\mu)$ and using the definition of the linear embedding. Moreover, for large n , whenever σ_i is non-constant, the corresponding f_i is a (essentially) high-degree function.⁴

Motivated by these examples and certain long-term applications to approximability of constraint satisfaction problems (CSPs) on satisfiable instances, authors of [5] hypothesized that the non-linearity is indeed the necessary and sufficient condition. We state the hypothesis below.

Hypothesis 1.3. (Informal): *The necessary and sufficient condition on a distribution μ on Σ^k so that the Conclusion (2) holds is that μ has no linear embedding over any Abelian group.*

In [5], the authors were able to prove the hypothesis for a sub-class of 3-ary predicates referred to therein as *semi-rich* predicates. A predicate $P: \Sigma^3 \rightarrow \{0, 1\}$ is called semi-rich if for each $(x, y) \in \Sigma \times \Sigma$, there exists a $z \in \Sigma$ such that $(x, y, z) \in P^{-1}(1)$ and also, for every $(x, z) \in \Sigma \times \Sigma$, there exists a $y \in \Sigma$ such that $(x, y, z) \in P^{-1}(1)$. We recall that while considering predicates, we always have an underlying distribution μ (in this case on Σ^3) such that $\text{supp}(\mu) = P^{-1}(1)$ and we may interchangeably talk in terms of either the predicate P or the distribution μ .

In this paper, we prove the hypothesis for **all** 3-ary predicates. The result, referred to as the Main Lemma in the rest of the paper, is stated below. It is more convenient (and general) to work with distributions μ on $\Sigma \times \Gamma \times \Phi$, allowing a different alphabet for each co-ordinate. In this case, a linear embedding consists of maps into an Abelian group G , $\sigma: \Sigma \rightarrow G$, $\gamma: \Gamma \rightarrow G$, $\phi: \Phi \rightarrow G$, not all constant, such that $\sigma(x) + \gamma(y) + \phi(z) = 0_G$ for all $(x, y, z) \in \text{supp}(\mu)$. We assume, unless stated otherwise, that the marginals of μ have full support on Σ, Γ, Φ respectively. In the following, m denotes the maximum size of Σ, Γ, Φ and $\alpha > 0$ denotes the minimum probability of a tuple in $\text{supp}(\mu)$. We always treat μ as fixed and m, α as fixed constants.

Lemma 1.4 (Main Analytical Lemma). *Suppose $|\Sigma|, |\Gamma|, |\Phi| \leq m$ and μ is a distribution over $\Sigma \times \Gamma \times \Phi$ such that*

⁴The functions here are complex valued with absolute value 1; one can take their real part if one insists on having real valued functions.

- The support of μ cannot be linearly embedded.
- $\mu(x, y, z) \geq \alpha$ for some $\alpha > 0$ and all $(x, y, z) \in \text{supp}(\mu)$.
- Marginals of μ (denoted as μ_x, μ_y, μ_z resp.) have full support on Σ, Γ, Φ respectively.

Considering m and α as fixed, for all $\varepsilon > 0$, there are $\xi, \delta > 0$ such that the following holds. If $f: \Sigma^n \rightarrow [-1, 1]$, $g: \Gamma^n \rightarrow [-1, 1]$, $h: \Phi^n \rightarrow [-1, 1]$ and $\text{Stab}_{1-\xi}(h; \mu_z) \leq \delta$, then we have that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right| \leq \varepsilon.$$

We clarify the condition that $\text{Stab}_{1-\xi}(h) \leq \delta$. Note that we have dropped μ_z from the notation for convenience. The parameter $\text{Stab}_{1-\xi}(h)$ denotes the stability of h under the noise parameter ξ . It is defined as $\langle h, T_{1-\xi}h \rangle$ where $T_{1-\xi}$ is the standard Beckner (noise) operator. We refer to Section 2 for all analytic definitions and basic tools.

The condition that $\text{Stab}_{1-\xi}(h) \leq \delta$ serves as a proxy for the condition that the function h is essentially of high degree. Indeed, if $\text{Stab}_{1-\xi}(h) \leq \delta$, it implies that the Fourier mass of h on terms of degree less than $\frac{1}{\xi}$ is at most $O(\delta)$. Conversely, if the Fourier mass on terms of degree less than $O(\frac{1}{\xi} \log(\frac{1}{\delta}))$ is at most $\frac{\delta}{2}$, then $\text{Stab}_{1-\xi}(h) \leq \delta$. Hence the low-stability condition is a proxy for the high-degree condition and turns out to be more convenient to work with.

One may wonder when a function h is bounded in $[-1, 1]$ as well as essentially of high degree. A natural example is when $h' : \Phi^n \rightarrow [-1, 1]$ is an arbitrary function and $h = h' - T_{1-\xi}h'$. In this case, since h' is bounded and $T_{1-\xi}$ is an averaging operator, h is also bounded. In addition, the operator $T_{1-\xi}$, roughly speaking, retains only the low-degree part of h' , and hence $h = h' - T_{1-\xi}h'$, roughly speaking, corresponds to the high-degree part of h' . More precisely, the Fourier mass of h on terms of degree less than $\frac{\delta}{\xi}$ is at most δ .⁵ In applications, it is almost always the case that the lemma is applied with $h = h' - T_{1-\xi}h'$ for some bounded function h' . One refers to h as a soft-truncation of h' , as opposed to a hard-truncation that would simply drop terms of degree less than a certain degree threshold. The advantage of using soft-truncation is that it preserves boundedness of functions whereas the hard-truncation in general does not.

1.1 Applications

In this section, we state a couple of applications of our main analytical lemma.

Hardness of approximation: Our first application is new results on dictatorship tests from integrality gap instances of constraint satisfaction problems (CSPs). Given a predicate $P : \Sigma^k \rightarrow \{0, 1\}$, for some alphabet Σ , a P -CSP instance consists of a set of variables x_1, x_2, \dots, x_n and a collection of *local* constraints C_1, C_2, \dots, C_m . Each constraint is of the type $P(x_{i_1}, x_{i_2}, \dots, x_{i_k})$. The constraints might involve *literals* instead of just the variables. An algorithmic task is to decide if there exists an assignment to the variables that satisfies all the constraints. In a related problem, called the Max- P -CSP problem, the task is to find an assignment to the variables that satisfies the maximum fraction of the constraints. An α -approximation algorithm is a polynomial-time algorithm which always returns an assignment that satisfies at least $\alpha \cdot \text{OPT}$ fraction of the constraints, where OPT is the value of the optimum assignment.

⁵Given the connection between stability and degree before, h also has low stability, albeit with somewhat different parameters.

Assuming the Unique Games Conjecture [15], Raghavendra [22] gave optimal hardness of approximation result for every Max- P -CSP. His work can be succinctly described as a two-step scheme:

SDP integrality gap \implies A dictatorship test \implies A hardness of approximation result.

However in his work, one necessarily loses perfect completeness and the hardness result does not hold on CSP instances that are (fully) satisfiable.

In order to prove hardness results on satisfiable instances, one would need a similar scheme that preserves perfect completeness in both the steps. Towards this goal, the Rich 2-to-1 Games Conjecture was introduced in [7] and further explored in [6]. Under this conjecture, [7, 6] showed how to convert, in certain specific cases, dictatorship test with completeness 1 and soundness s to a hardness result on satisfiable CSP instances with hardness threshold $s + \varepsilon$, for every constant $\varepsilon > 0$. This result can be interpreted as fulfilling the second step in the scheme above (albeit only morally speaking, since the implication is not entirely seamless and general yet).

It thus remains to fulfill the first step in the scheme while preserving perfect completeness. The authors [5] made progress on this question, showing that a $(1, s)$ integrality gap instance for certain CSPs can be converted into a dictatorship test with completeness 1 and soundness $s + \varepsilon$. Their result however was limited to (non-linear) 3-ary predicates satisfying the aforementioned semi-richness condition, and this was because in [5], the authors were able to prove analytic Lemma 1.4 only under the additional semi-richness condition. Since we are now able to prove the lemma for all (non-linear) 3-ary predicates, we now get the integrality gap to dictatorship test implication for all such predicates. The formal statement of our result appears below (one wishes that the condition (2b) therein could be dropped; if so, we would have a full-proof implication). For definitions and a more detailed discussion, we refer to Section 8 and the introductory section of [5].

Theorem 1.5. *Let $P: \Sigma^3 \rightarrow \{0, 1\}$ be any predicate that satisfies the following conditions: (1) P has no linear embedding, (2a) there exists an instance of Max- P -CSP that has a $(1, s)$ -integrality gap for the basic SDP relaxation, (2b) on every constraint, the local distribution in the SDP solution is not linearly embeddable. Then for every $\varepsilon > 0$, there is a dictatorship test for P -CSP that has perfect completeness and soundness $s + \varepsilon$.*

Counting Progressions: In additive combinatorics, finding a certain fixed progression (i.e. a pattern) in a subset of a given group is a cornerstone question. Such questions have had huge implications in understanding the pseudo-random properties of subsets of a group. Below we list a few of these results answering this question in different settings.

Fix a finite Abelian group $(G, +)$. A subset $A \subseteq G$ is said to be three term arithmetic progression (3-AP) free if there is no arithmetic progression of size 3 in A . In other words, there are no elements $x, y, z \in A$ such that $x + z = 2y$. The famous Roth's Theorem [23] shows that any 3-AP free subset of \mathbb{Z}_N must be of size $o(N)$. In the contrapositive, any constant density subset of \mathbb{Z}_N contains a 3-term AP. Szemerédi [24] generalized Roth's Theorem to any k -term AP. In these and similar results quoted next, one actually shows that a density δ subset of the group contains an ε fraction of all the progressions; the precise dependence of ε as a function of δ is also interesting, but for the sake of conciseness, we skip quantitative statements to that effect.

Now let (G, \cdot) be a finite group that is not necessarily Abelian. A subset of G is called product free if it does not contain three elements x, y, z with $x \cdot y = z$. If G is any Abelian group, then it is easy to come up with product-free sets of constant density. Gowers [12] showed that this is not true for a class of non-Abelian

groups called *quasirandom groups*.⁶ That is, every constant density subset of a quasirandom group contains the progression (x, y, xy) . Tao [25] extended Gowers' result to other progressions of the form (x, xg, xg^2) and (x, xg, xg^2, xg^3) for some very specific quasirandom groups. Bergelson and Tao [2] established it for progressions (x, xg, gx) and (g, x, xg, gx) for every quasirandom group. Recently, following the work by Peluse [20], Bhangale, Harsha and Roy [3] established it for the progression (x, xg, xg^2) for every quasirandom group. In a high-dimensional setting, finding the largest size of the 3-AP free set in \mathbb{F}_3^n has received considerable attention [8, 16, 1]. Ellenberg and Gijswijt [11], building on a beautiful work by Croot, Lev, Pach [10], obtained a substantial quantitative improvement over Roth's Theorem (applied to \mathbb{F}_3^n).

We now state our general theorem that establishes a similar result in high-dimensional setting for arbitrary 3-ary progression provided that the progression has no linear embedding (along with a couple of other conditions).

Theorem 1.6. *Suppose μ is a distribution over Σ^3 such that (1) the marginal distributions μ_x, μ_y, μ_z are uniform on Σ , (2) $\{(x, x, x) \mid x \in \Sigma\} \subseteq \text{supp}(\mu)$, and (3) $\text{supp}(\mu)$ cannot be linearly embedded. Then for all $\delta > 0$, there exists $\varepsilon > 0$ such that for $S \subseteq \Sigma^n$ with $|S| \geq \delta|\Sigma|^n$,*

$$\Pr_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [\mathbf{x} \in S, \mathbf{y} \in S, \mathbf{z} \in S] \geq \varepsilon.$$

Note that the condition (2) is necessary for such a conclusion to hold. This can be seen by the following example. Consider $\Sigma = \{0, 1, 2\}$ and μ be uniform on $\Sigma^3 \setminus \{(0, 0, 0)\}$. It is easy to check that μ is not linearly embeddable. Now, if we take $S \subseteq \Sigma^n$ to be $S = \{\mathbf{x} \in \Sigma^n \mid x_1 = 0\}$, then clearly the conclusion does not hold. Our theorem is comparable to the result by Hązła, Holenstein and Mossel [13] with the same conclusion under the additional condition that the distribution μ is *connected*. As there are distributions that are not linearly embeddable as well as not connected, Theorem 1.6 extends their result.

1.2 Techniques

In this section, we elaborate on the ideas involved in proving Lemma 1.4. We focus only on a few high-level ideas here. Since we will skip many technical (and even conceptual) details, there might be some discrepancies between the high-level exposition here and formal proofs appearing later.

Let μ be a distribution on $\Sigma \times \Gamma \times \Phi$ such that $\text{supp}(\mu)$ is not linearly embeddable. We wish to show that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right| \approx 0, \tag{3}$$

where $f : \Sigma^n \rightarrow [-1, 1]$, $g : \Gamma^n \rightarrow [-1, 1]$, $h : \Phi^n \rightarrow [-1, 1]$, are ℓ_∞ -bounded and at least one of the functions essentially has high degree. We begin by sketching Mossel's proof [17] that works in the 2-ary case, i.e. for a (non-linear) distribution μ on $\Sigma \times \Gamma$. This will help us understand various hurdles and new ideas needed to overcome these hurdles in our proof of the 3-ary case as above.

1.2.1 The 2-ary Case: Sketch of Mossel's Proof

Let μ be a distribution on $\Sigma \times \Gamma$ such that $\text{supp}(\mu)$ is not linearly embeddable. It is easily seen that the non-linearity condition, in this special 2-ary case, is same as saying that $\text{supp}(\mu)$, viewed as a bipartite

⁶A group (or rather a family of groups) is quasirandom if the minimum dimension of any non-trivial group representation grows with the size of the group.

graph G_μ on the vertex set $\Sigma \cup \Gamma$, is connected. Indeed, if this graph were disconnected, with components $C_0 \cup D_0, \dots, C_{r-1} \cup D_{r-1}$, then an embedding $\sigma : C_j \rightarrow j, \gamma : D_j \rightarrow -j$ is an embedding of Σ and Γ respectively into \mathbb{Z}_r and for all $(x, y) \in \text{supp}(\mu)$ (i.e. the edges of the graph G_μ), we have $\sigma(x) + \gamma(y) = 0$ in \mathbb{Z}_r .

We intend to show that if $f : \Sigma^n \rightarrow \mathbb{R}, g : \Gamma^n \rightarrow \mathbb{R}$ are n -dimensional ℓ_∞ -bounded functions where g has high degree, then $|\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})]|$ is small. For simplicity of exposition, we assume that g in fact has full degree n .⁷ In this case, we are able to show that $|\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})]| \leq (1 - \tau)^n \|f\|_2 \|g\|_2$ for some constant $\tau = \tau(\mu) > 0$. We emphasize here that one gets an upper bound in terms of the ℓ_2 -norm of the functions. This of course implies an upper bound in terms of the ℓ_∞ -norms. Thus we really do not need the n -dimensional functions to be ℓ_∞ -bounded in the 2-ary case. This is one aspect (among many) in which the 3-ary case is fundamentally different, where one does need the n -dimensional functions to be ℓ_∞ -bounded (as we will soon demonstrate via an example).

Continuing the consideration of the 2-ary case, the proof proceeds in two steps: first establishing a base case inequality (for $n = 1$) and then observing that the inequality tensorizes, leading to an inductive proof and the desired bound for the general case of n -dimensional functions. The base case inequality is necessarily an ℓ_2 -inequality and this fact is essential for the inductive proof (and the same holds in the 3-ary case).

Towards stating the base case inequality, let $f : \Sigma \rightarrow \mathbb{R}, g : \Gamma \rightarrow \mathbb{R}$ be functions. By Cauchy-Schwarz,

$$\left| \mathbb{E}_{(x,y) \sim \mu} [f(x)g(y)] \right| \leq \|f\|_2 \|g\|_2.$$

We refer to this essentially trivial inequality as the (base case) sanity check inequality. The inequality that is actually needed is that when $\mathbb{E}[f] = \mathbb{E}[g] = 0$, we in fact have the improvement

$$\left| \mathbb{E}_{(x,y) \sim \mu} [f(x)g(y)] \right| \leq (1 - \tau) \|f\|_2 \|g\|_2, \quad \mathbb{E}[f] = \mathbb{E}[g] = 0, \quad (4)$$

for some constant $\tau = \tau(\mu) > 0$. It is not difficult to see that this follows from the connectedness of the distribution μ (or equivalently the graph G_μ), but we skip the proof. An equivalent way to express the inequality is that the operator $T : \tilde{L}_2(\Gamma; \mu_y) \rightarrow \tilde{L}_2(\Sigma; \mu_x)$ defined as $Tg(x) = \mathbb{E}_{(x',y) \sim \mu} [g(y) | x' = x]$ has operator norm at most $1 - \tau$. Here $\tilde{L}_2(\Gamma; \mu_y)$ denotes the subspace of $L_2(\Gamma; \mu_y)$ consisting of those functions g for which $\mathbb{E}[g] = 0$ (and similarly for $\tilde{L}_2(\Sigma; \mu_x)$). The operator norm of T , denoted $\|T\| = \max_{g: \mathbb{E}[g]=0} \|Tg\|_2 / \|g\|_2$, is at most $1 - \tau$ according to the equivalent interpretation of the inequality (4), which can then be derived as:

$$\left| \mathbb{E}_{(x,y) \sim \mu} [f(x)g(y)] \right| = |\langle f, Tg \rangle| \leq \|f\|_2 \|Tg\|_2 \leq \|f\|_2 \|T\| \|g\|_2 \leq (1 - \tau) \|f\|_2 \|g\|_2.$$

Now we consider the n -dimensional case. Let $f : \Sigma^n \rightarrow \mathbb{R}, g : \Gamma^n \rightarrow \mathbb{R}$ be n -dimensional functions. As mentioned before, we assume that g has full degree, which amounts to saying that $g \in \tilde{L}_2(\Gamma; \mu_y)^{\otimes n}$. In this case, it follows directly that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})] \right| \leq (1 - \tau)^n \|f\|_2 \|g\|_2,$$

⁷This amounts to saying that after restricting any $n - 1$ co-ordinates, the expectation of g over the remaining co-ordinate is zero.

using the well-known fact that the operator norm is multiplicative (i.e. it tensorizes), namely that $\|T^{\otimes n}\| = \|T\|^n \leq (1 - \tau)^n$. Using this fact, one immediately concludes that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})] \right| = |\langle f, T^{\otimes n} g \rangle| \leq \|f\|_2 \|T^{\otimes n} g\|_2 \leq \|f\|_2 \|T^{\otimes n}\| \|g\|_2 \leq (1 - \tau)^n \|f\|_2 \|g\|_2,$$

as desired. If one wishes, one can prove the multiplicativity of operator norm by induction and view the overall proof as an inductive proof, using the base case inequality (4) and “gaining” a factor $1 - \tau$ in each step of the induction. While we don’t demonstrate it here, we mention it because the proof for the 3-ary case proceeds along similar lines, albeit with many conceptual and technical hurdles. Therein, it is rather challenging even to formulate the “correct” base case inequality.

1.2.2 Towards 3-ary Base Case: Restoring Sanity First

Moving onto the 3-ary case, let μ be a distribution on $\Sigma \times \Gamma \times \Phi$ such that $\text{supp}(\mu)$ is not linearly embeddable. One hopes to write down a suitable base case inequality and use it towards an inductive proof. However, it turns out that even the sanity check inequality fails in general! That is, for $f : \Sigma \rightarrow \mathbb{R}, g : \Gamma \rightarrow \mathbb{R}, h : \Phi \rightarrow \mathbb{R}$, while we desire a base case inequality (say when $\mathbb{E}[f] = 0$) of the form

$$\left| \mathbb{E}_{(x, y, z) \sim \mu} [f(x)g(y)h(z)] \right| \leq (1 - \tau) \|f\|_2 \|g\|_2 \|h\|_2, \quad (5)$$

it may actually happen that

$$\left| \mathbb{E}_{(x, y, z) \sim \mu} [f(x)g(y)h(z)] \right| > \|f\|_2 \|g\|_2 \|h\|_2.$$

In other words, we may not even have the upper bound of $\|f\|_2 \|g\|_2 \|h\|_2$ in the 3-ary case whereas the corresponding upper bound in the 2-ary case is the essentially trivial application of Cauchy-Schwarz! Here is an example.

Suppose that $\Sigma = \Gamma = \Phi$, $|\Sigma| = m \geq 54$, and μ has a probability mass of $1 - \varepsilon$ uniformly spread on the triples $\{(x, x, x) | x \in \Sigma\}$ and the remaining probability mass of ε uniformly spread on all the remaining triples in Σ^3 . Clearly, $\text{supp}(\mu) = \Sigma^3$ and hence μ is not linearly embeddable. The marginals of μ are uniform on Σ . We can certainly construct a function $f : \Sigma \rightarrow \mathbb{R}$ such that $\mathbb{E}[f(x)] = 0$ and $\mathbb{E}[f(x)^3] > \|f\|_2^3$. For instance, f could take the values $2m, -m, -m$ at three distinct points in Σ and zero at the remaining points in Σ . In this case, $\mathbb{E}[f(x)] = 0$, $\mathbb{E}[f(x)^2] = 6m$, and $\mathbb{E}[f(x)^3] = 6m^2$, and thus $\mathbb{E}[f(x)^3] \geq \sqrt{m/6} \cdot \|f\|_2^3 \geq 3 \|f\|_2^3$. Letting $f = g = h$ and recalling that the triples (x, x, x) receive $1 - \varepsilon$ of the probability mass, it follows that

$$\mathbb{E}_{(x, y, z) \sim \mu} [f(x)g(y)h(z)] \geq (1 - \varepsilon) \mathbb{E}[f(x)^3] - \varepsilon \cdot O_m(1) \geq 2 \cdot \|f\|_2^3 = 2 \cdot \|f\|_2 \|g\|_2 \|h\|_2,$$

by making ε sufficiently small. This example also shows that in order to claim the desired bound for n -dimensional functions as in Equation (3), we must use the fact that the functions are ℓ_∞ -bounded (i.e. in $[-1, 1]$)! Indeed, consider the same example here and let n -dimensional functions $\tilde{f} = \tilde{g} = \tilde{h} : \Sigma^n \rightarrow \mathbb{R}$ be all equal to $f^{\otimes n} / \|f^{\otimes n}\|_2$. Then these have all ℓ_2 -norm 1, whereas

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [\tilde{f}(\mathbf{x})\tilde{g}(\mathbf{y})\tilde{h}(\mathbf{z})] = \mathbb{E}_{(x, y, z) \sim \mu} [f(x)g(y)h(z)]^n \cdot \frac{1}{\|f\|_2^{3n}} \geq 2^n.$$

We thus face a seemingly intractable hurdle and a contradictory set of constraints: (i) we do need the ℓ_∞ -boundedness of the n -dimensional functions, (ii) an inductive proof is some form of tensorization argument and hence inherently an ℓ_2 -proof; consequently, the intermediate functions arising during the induction can only be assumed to have ℓ_2 norm at most 1, (iii) the inductive argument requires a base case ℓ_2 -inequality such as (5) which actually happens to fail miserably!

We now show how to overcome this hurdle step-by-step. This is achieved in a round-about manner, by carefully transforming the distribution and the alphabet $(\Sigma \times \Gamma \times \Phi, \mu)$ to another distribution and alphabet $(\tilde{\Sigma} \times \tilde{\Gamma} \times \tilde{\Phi}, \tilde{\mu})$. Formally, we show that

- If μ was not linearly embeddable to begin with, then $\tilde{\mu}$ isn't either.
- If Lemma 1.4 (i.e. our Main Lemma/Result) holds for $\tilde{\mu}$, then it also holds for μ .

In this sense, we are able to reduce our task of proving the lemma for the original distribution μ to proving the same lemma for the new distribution $\tilde{\mu}$. In fact, there will be a series of such transformations. The (first) transformation will ensure that the marginal of $\tilde{\mu}$ on $\tilde{\Gamma} \times \tilde{\Phi}$ is a uniform, product distribution. Once we have this additional property, we at least have the (base case) sanity check inequality as demonstrated next. For the sake of notational convenience, we rename the new distribution and the alphabet again as $(\Sigma \times \Gamma \times \Phi, \mu)$ and assume that the marginal of μ on $\Gamma \times \Phi$ is a uniform, product distribution. If so, it is easily seen that we get the (base case) sanity check inequality, namely that for $f : \Sigma \rightarrow \mathbb{R}, g : \Gamma \rightarrow \mathbb{R}, h : \Phi \rightarrow \mathbb{R}$, we have

$$\left| \mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h(z)] \right| \leq \|f\|_2 \|g\|_2 \|h\|_2.$$

Indeed, by Cauchy-Schwarz,

$$\begin{aligned} \mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h(z)]^2 &\leq \mathbb{E}_{x \sim \mu_x} [f(x)^2] \mathbb{E}_{(y,z) \sim \mu_{y,z}} [g(y)^2 h(z)^2] \\ &= \mathbb{E}_{x \sim \mu_x} [f(x)^2] \mathbb{E}_{y \sim \mu_y} [g(y)^2] \mathbb{E}_{z \sim \mu_z} [h(z)^2] \\ &= \|f\|_2^2 \|g\|_2^2 \|h\|_2^2, \end{aligned} \tag{6}$$

where in the second step, we used the property that (y, z) are uniform and independent! It is also possible to ensure (after the transformation) another property of μ that is quite convenient: for all pairs $(y, z) \in \Gamma \times \Phi$, there is a unique $x \in \Sigma$ such that $(x, y, z) \in \text{supp}(\mu)$ (we then say that (y, z) determine x). The details of this transformation and related proofs appear in Section 7; some of its ingredients are borrowed from authors' earlier work [5].

1.2.3 The 3-ary Relaxed Base Case: Overcoming the Horn-SAT Obstruction

We will henceforth assume that the distribution μ on $\Sigma \times \Gamma \times \Phi$ has no linear embedding and has uniform marginal on $\Gamma \times \Phi$. Now that we at least have the sanity check inequality, we ask ourselves whether we can claim the desired base case inequality as below:

Question 1.7. (Desired, Hypothetical Base Case Inequality:) *If μ has no linear embedding and has uniform marginal on $\Gamma \times \Phi$, is it necessarily the case that for $f : \Sigma \rightarrow \mathbb{R}, g : \Gamma \rightarrow \mathbb{R}, h : \Phi \rightarrow \mathbb{R}$,*

$$\left| \mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h(z)] \right| \leq (1 - \tau(\theta)) \|f\|_2 \|g\|_2 \|h\|_2, \quad |\mathbb{E}[f]| \leq (1 - \theta) \|f\|_2. \tag{7}$$

To avoid the trivial case when f, g, h are all constant functions, we added here the condition that f is non-constant and has some variance, the condition captured by the requirement $|\mathbb{E}[f]| \leq (1 - \theta)\|f\|_2$.

We note that such a base case inequality seems necessary towards an inductive proof since one hopes to “gain” a factor of $1 - \tau$ in each step of the induction. However it turns out that such an inequality need not necessarily hold and there could be an obstruction that we refer to as the Horn-SAT obstruction (and this is the only possible obstruction).

Definition 1.8. Assume that a distribution μ on $\Sigma \times \Gamma \times \Phi$ has no linear embedding and its marginal on $\Gamma \times \Phi$ is uniform. We say that μ has a Horn-SAT embedding if there are Boolean functions $f : \Sigma \rightarrow \{0, 1\}$, $g : \Gamma \rightarrow \{0, 1\}$, $h : \Phi \rightarrow \{0, 1\}$, such that

- For all $(x, y, z) \in \text{supp}(\mu)$, we have $f(x) = g(y)h(z)$.
- f is non-constant (and in that case so must be g and h).

The condition $f(x) = g(y)h(z)$ for Boolean functions is equivalent to the conjunction of clauses $\overline{f(x)} \vee g(y)$, $\overline{f(x)} \vee h(z)$, $f(x) \vee \overline{g(y)}$, $f(x) \vee \overline{h(z)}$. These are all Horn-SAT clauses (i.e. having at most one positive literal), explaining the term *Horn-SAT embedding*. We now make several remarks towards understanding how a Horn-SAT embedding is an obstruction towards the desired inequality (7) and how it is the only possible obstruction.

- Firstly, we note that having a Horn-SAT embedding violates inequality (7). Indeed, since $f(x) = g(y)h(z)$ in $\text{supp}(\mu)$ and (y, z) are uniform and independent, we have $\|f\|_2 = \|g\|_2\|h\|_2$ and then

$$\mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h(z)] = \mathbb{E}_{(y,z) \sim \mu_{y,z}} [g(y)^2h(z)^2] = \|g\|_2^2\|h\|_2^2 = \|f\|_2\|g\|_2\|h\|_2.$$

One also notes that since f is Boolean and non-constant, it does have constant variance.

- Secondly, we note that if the inequality (7) is not possible, then there is necessarily a Horn-SAT embedding. A sketch of the proof is as follows. For a fixed θ , suppose that there are functions that violate the inequality for all $\tau \rightarrow 0$. Then by standard compactness argument, there are functions $f : \Sigma \rightarrow \mathbb{R}$, $g : \Gamma \rightarrow \mathbb{R}$, $h : \Phi \rightarrow \mathbb{R}$, such that

$$\mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h(z)] = \|f\|_2\|g\|_2\|h\|_2,$$

i.e. achieving an exact equality. This means that the application of Cauchy-Schwarz in Equation (6) must be tight and therefore $f(x) = g(y)h(z)$ in $\text{supp}(\mu)$ (as equality of real numbers). If $f(x)$ is always non-zero, then so are $g(y)$ and $h(z)$. In this case, if at least one of f, g, h has a non-constant sign (i.e. positive and negative), then turning $f(x), g(y), h(z)$ into their $\{+1, -1\}$ - signs, $\text{sign}(f(x)) = \text{sign}(g(y))\text{sign}(h(z))$, yields a linear embedding of μ , a contradiction. On the other hand, if f, g, h all have constant sign, then w.l.o.g. this sign is positive, and then $\log f(x) = \log g(y) + \log h(z)$ gives a linear embedding, again a contradiction. Here f has some variance, so $\log f(x)$ is non-constant and the embedding is non-trivial. The embedding is not into a finite Abelian group, but this is not difficult to fix. One concludes therefore that $f(x)$ takes the zero value for some $x \in \Sigma$ and of course also takes a non-zero value for some $x' \in \Sigma$. We can now define the Horn-SAT embedding by turning $f(x), g(y), h(z)$ into Boolean 1 if the value is non-zero and Boolean 0 if the value is zero!

- In the definition, if f is non-constant, then so must be g and h . Let's suppose on the contrary that g is constant (the same proof applies for h). If $g \equiv 0$, then the condition $f(x) = g(y)h(z)$ implies that $f \equiv 0$, reaching a contradiction. If $g \equiv 1$, then one concludes that $f(x) = h(z)$ for all $(x, z) \in \text{supp}(\mu_{x,z})$. Since μ is not linearly embeddable, its marginals are not linearly embeddable either.⁸ In particular, $\mu_{x,z}$ has no linear embedding and hence is connected, implying that both f and h are constant, again a contradiction.

Considering these remarks, if μ does not have a Horn-SAT embedding, then we do have the base case inequality (7) and we can hope to carry out the induction. However, if μ does have a Horn-SAT embedding as in Definition 1.8, then the embedding serves as a violation of the inequality and we are stuck with a similar hurdle as before. The Horn-SAT embedding leads to n -dimensional functions $\tilde{f} = f^{\otimes n} / \|f^{\otimes n}\|$, $\tilde{g} = g^{\otimes n} / \|g^{\otimes n}\|$, $\tilde{h} = h^{\otimes n} / \|h^{\otimes n}\|$, with ℓ_2 -norm 1, and

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [\tilde{f}(\mathbf{x})\tilde{g}(\mathbf{y})\tilde{h}(\mathbf{z})] = 1.$$

As before, this hinders the possibility of proving the n -dimensional inequality (3) by induction: there is no base case inequality and there is a counter-example if one allows functions to have ℓ_2 norm 1 instead of ℓ_∞ norm 1.

We overcome this hurdle in a similar manner as before, albeit with even more subtleness. We carefully transform the distribution and the alphabet $(\Sigma \times \Gamma \times \Phi, \mu)$ to another distribution and alphabet $(\tilde{\Sigma} \times \tilde{\Gamma} \times \tilde{\Phi}, \tilde{\mu})$. Formally, we show that

- If Lemma 1.4 (i.e. our Main Lemma/Result) holds for $\tilde{\mu}$, then it also holds for μ .
- All the key properties of μ are retained by $\tilde{\mu}$ which has further additional properties.

In this sense, we are able to reduce our task of proving the lemma for the original distribution μ to proving the same lemma for the new distribution $\tilde{\mu}$. Now we state what additional properties $\tilde{\mu}$ has. For the sake of notational convenience, we rename the new distribution and the alphabet as $(\Sigma \times \Gamma \times \Phi, \mu)$ again. The key additional property is stated below, referred to as the relaxed base case inequality.

Definition 1.9. (*Relaxed Base Case Inequality*) Suppose a distribution μ on $\Sigma \times \Gamma \times \Phi$ has no linear embedding and has uniform support on $\Gamma \times \Phi$. We say that μ satisfies the relaxed base case inequality if:

- There is some $\Sigma' \subseteq \Sigma$, $|\Sigma'| \geq 2$, and constants $C > 0$ and $0 < c < 1$ such that the following holds. For all $\tau > 0$, let functions $f: \Sigma \rightarrow \mathbb{R}$, $g: \Gamma \rightarrow \mathbb{R}$ and $h: \Phi \rightarrow \mathbb{R}$ be such that f has variance at least $\tau \|f\|_2^2$ on Σ' , that is

$$\mathbb{E}_{x, x' \in \Sigma'} [(f(x) - f(x'))^2] \geq \tau \|f\|_2^2.$$

Then

$$\left| \mathbb{E}_{(x, y, z) \sim \mu} [f(x)g(y)h(z)] \right| \leq \max(1 - \tau^C, c) \|f\|_2 \|g\|_2 \|h\|_2.$$

- Furthermore, the distribution on $\Sigma' \times \Gamma \times \Phi$, derived as $(x, y, z) \sim \mu$ conditioned on $x \in \Sigma'$, cannot be linearly embedded.

⁸This is seen easily from the definition of a linear embedding, Definition 1.2. If marginal of μ on a subset of co-ordinates has a linear embedding, then so does μ by letting the embedding on other co-ordinates to be 0_G .

We remark that if μ did not have a Horn-SAT embedding, no transformation is needed, and one can simply take $\Sigma' = \Sigma$ in the above definition. However in general there might be a Horn-SAT embedding and the transformation would be needed. The transformation is rather subtle and while we do consider it to be one of the key ideas, we skip the discussion here and refer to Section 7.3 for details. To summarize, we reduce the task of proving our Main Lemma 1.4 to the same task with the additional property that μ satisfies the relaxed base case inequality, i.e. to the task of proving the lemma stated below. In the following lemma, properties numbered 1 and 2 are as before, 3 and 4 can be assumed from the authors' earlier work as discussed in Section 1.2.2, and that numbered 5 is the key relaxed base case inequality.

Lemma 1.10. (Main Analytical Lemma under Relaxed Base Case Inequality) *Suppose $|\Sigma|, |\Gamma|, |\Phi| \leq m$ and μ is a distribution over $\Sigma \times \Gamma \times \Phi$ such that:*

1. $\mu(x, y, z) \geq \alpha$ for some $\alpha > 0$ and all $(x, y, z) \in \text{supp}(\mu)$.
2. $\text{supp}(\mu)$ cannot be linearly embedded.
3. The marginal $\mu_{y,z}$ is uniform and independent over $\Gamma \times \Phi$.
4. For all $(y, z) \in \Gamma \times \Phi$, there is a unique $x \in \Sigma$ such that $(x, y, z) \in \text{supp}(\mu)$ (i.e. y, z determine x).
5. μ satisfies the relaxed base case inequality as in Definition 1.9.

Then for all $\varepsilon > 0$, there are $\xi, \delta > 0$ such that the following holds. If $f: \Sigma^n \rightarrow [-1, 1]$, $g: \Gamma^n \rightarrow [-1, 1]$ and $h: \Phi^n \rightarrow [-1, 1]$ satisfy that either $\text{Stab}_{1-\xi}(g) \leq \delta$ or $\text{Stab}_{1-\xi}(h) \leq \delta$, then we have that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right| \leq \varepsilon.$$

1.2.4 The Inductive Argument (without the Horn-SAT Obstruction)

Armed with the ‘‘correct’’ relaxed base case inequality, we now give an overview of the inductive proof (of Lemma 1.10). It is instructive and less cumbersome to first consider the special case when there is no Horn-SAT embedding and we already have the base case inequality as in (7). We will indicate how to incorporate the relaxed base case inequality later. Formal proofs appear in Sections 3, 4, 5, and 6.

So let us focus on this special case and assume the base case inequality (7) holds. The inductive proof proceeds in several steps. We emphasize again that an inductive proof must necessarily work with ℓ_2 norms of functions that arise as intermediate functions during the induction and we have no control over their ℓ_∞ norms.⁹ We are given that either g or h has essentially high degree, so let’s say this holds for g , formalized in terms of its low-stability. The first step towards the inductive proof is to note that it is sufficient (and necessary as far as our proof goes) to focus on the case when f, g, h are homogenous functions. We will skip details regarding how this is sufficient towards the general case. Therefore let’s assume that f, g, h are homogenous and define the parameter

$$\beta_{n, d_1, d_2, d_3} = \sup_{f, g, h} \frac{|\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})]|}{\|f\|_2 \|g\|_2 \|h\|_2},$$

⁹When there is no Horn-SAT embedding, we do not need an ℓ_∞ bound on the original functions either. This is indeed the special case we are considering here. When there is a Horn-SAT embedding, as noted before, we must somehow use the fact that the original functions f, g, h do have ℓ_∞ norm at most 1. We still have no control however over the ℓ_∞ norm of the intermediate functions. This issue is addressed later.

where the maximum is taken over all $f: \Sigma^n \rightarrow \mathbb{R}$, $g: \Gamma^n \rightarrow \mathbb{R}$, $h: \Phi^n \rightarrow \mathbb{R}$ homogenous of degrees d_1, d_2, d_3 respectively. Since we assumed that g had high degree, we think of d_2 as (roughly) the largest among the degrees. Indeed, it is sufficient to consider the case when $d_1, d_3 \leq 10d_2$, and we make this assumption skipping the details. We will be able to show an exponential decay, namely

$$\beta_{n,d_1,d_2,d_3} \leq (1 - \Omega_{\alpha,m}(1))^{d_2},$$

completing the proof. We now describe how this exponential decay is proved. First, we reduce the dimension n so that $n \leq O(d_2)$. Then comes the core inductive argument, where we “gain” a factor $1 - \Omega_{\alpha,m}(1)$ in each step of the induction, reducing the degree d_2 by one, until we have reduced it to say $\frac{d_2}{2}$.

Reducing Dimension: We show here that it is sufficient to consider the case when $n \leq O(d_2)$ (and we already assume that $d_1, d_3 \leq 10d_2$). The idea is as follows. As long as $n \gg d_2$, we can find a coordinate $i \in [n]$ which has very small “influence” on f, g and h ; assume without loss of generality that this coordinate is $i = n$. If the influence was zero, then f, g and h would only be functions of the first $n - 1$ co-ordinates, and hence we would conclude that $\beta_{n,d_1,d_2,d_3} \leq \beta_{n-1,d_1,d_2,d_3}$, making “progress” in reducing n . However in general, that influence may be very small but still non-zero. In that case one may write the decompositions

$$f = f_1 + f', \quad g = g_1 + g', \quad h = h_1 + h',$$

where f_1, g_1, h_1 depend only on the first $n - 1$ co-ordinates, and f', g', h' do depend on the n^{th} co-ordinate but have very small ℓ_2 -norm (which is precisely what influence is). Since f', g', h' have very small norm, one doesn't expect them to contribute much, and one still hopes to deduce that $\beta_{n,d_1,d_2,d_3} \leq \beta_{n-1,d_1,d_2,d_3}$. Alas, this doesn't quite work. While their contribution is very small, it is still non-zero, and a naive application of this idea would only give $\beta_{n,d_1,d_2,d_3} \leq \beta_{n-1,d_1,d_2,d_3} + o(1)$, and the $o(1)$ error terms will keep accumulating in successive inductive steps. To overcome this difficulty, we perform a more detailed analysis, and need more refined decompositions of f, g and h . For the sake of simplicity, we consider only a specialized scenario that allows us to write

$$f = f_1 + f_2 f'_2, \quad g = g_1 + g_2 g'_2, \quad h = h_1 + h_2 h'_2,$$

where f_1, g_1, h_1 depend only on the first $n - 1$ coordinates and have the same degrees as f, g, h , the functions f_2, g_2, h_2 also depend only on the first $n - 1$ coordinates but have degrees one less than f, g, h respectively, and f'_2, g'_2, h'_2 are functions that only depend on the last coordinate and have very small ℓ_2 -norm. Using this decomposition, we can write

$$\begin{aligned} \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] &= \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n-1}} [f_1(\mathbf{x})g_1(\mathbf{y})h_1(\mathbf{z})] \\ &+ \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n-1}} [f_2(\mathbf{x})g_2(\mathbf{y})h_1(\mathbf{z})] \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu} [f'_2(\mathbf{x})g'_2(\mathbf{y})] \\ &+ \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n-1}} [f_2(\mathbf{x})g_1(\mathbf{y})h_2(\mathbf{z})] \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu} [f'_2(\mathbf{x})h'_2(\mathbf{z})] \\ &+ \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n-1}} [f_2(\mathbf{x})g_2(\mathbf{y})h_2(\mathbf{z})] \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu} [f'_2(\mathbf{x})g'_2(\mathbf{y})h'_2(\mathbf{z})] \\ &+ \text{Other terms.} \end{aligned}$$

The other terms are zero thanks to the fact that $\mu_{y,z}$ is uniform and independent. The first term is the dominant term, the second and the third terms constitute as error terms, and the fourth term can be ignored

when compared to the second and third terms. Roughly speaking, the reason is that if ε denotes the small norm of f'_2, g'_2, h'_2 , then the corresponding expectations are of the order ε^2 in the second and third terms, and of the order ε^3 in the fourth term.

The second and third terms are error terms, which however cannot be ignored altogether (as said before) and require care. Skipping many details, it turns out that the key is to bound the expectation

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu} [f'_2(\mathbf{x})(g'_2(\mathbf{y}) + h'_2(\mathbf{z}))].$$

This can be upper bounded by $(1 - \Omega(1))\|f'_2\|_2 \sqrt{\|g'_2\|_2^2 + \|h'_2\|_2^2}$. We emphasize here that this is an inequality on functions of a single co-ordinate. It is referred to as the additive base case inequality (see Lemma 2.17). Using this bound, one can obtain an effective enough bound on the second and third terms above, somehow recover the loss from these error terms and get that $\beta_{n,d_1,d_2,d_3} \leq \beta_{n-1,d_1,d_2,d_3}$ as desired.

The Core Induction: We now show the core inductive step giving the exponential decay, namely that $\beta_{n,d_1,d_2,d_3} \leq (1 - \Omega_{\alpha,m}(1))^{d_2}$. We assume that $n \leq O(d_2)$ as discussed and that $d_1, d_3 \leq 10d_2$. Skipping details, it is sufficient to assume further that $d_1 \geq \Omega(d_2)$ as well. It follows from these assumptions that average influence of a coordinate on f is $\frac{d_1}{n} \geq \Omega(1)$. Let us assume that the coordinate n has influence $\Omega(1)$ on f . For the sake of simplicity, consider furthermore only a specialized scenario that allows us to write f, g and h as

$$f = f_1 f'_1, \quad g = g_1 g'_1, \quad h = h_1 h'_1,$$

where f_1, g_1, h_1 depend only on the first $n - 1$ co-ordinates and have degrees one less than f, g, h , and the functions f'_1, g'_1, h'_1 depend only on the single coordinate n , and f'_1 has constant norm (which amounts to the said influence). In this case, we would have that

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] = \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n-1}} [f_1(\mathbf{x})g_1(\mathbf{y})h_1(\mathbf{z})] \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu} [f'_1(\mathbf{x})g'_1(\mathbf{y})h'_1(\mathbf{z})].$$

By the inductive hypothesis, the first term is at most $\beta_{n-1,d_1-1,d_2-1,d_3-1}$ and by the base case inequality, $|\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu} [f'_1(\mathbf{x})g'_1(\mathbf{y})h'_1(\mathbf{z})]| \leq \lambda = 1 - \Omega(1)$. Hence we get that

$$\beta_{n,d_1,d_2,d_3} \leq \lambda \beta_{n-1,d_1-1,d_2-1,d_3-1},$$

as desired, and iterating this gives an exponential decay.

In general, the main complication is that f, g and h need not take the specialized form as above, and instead one has to decompose them in a more complicated manner (amounting to decomposing a tensor into a sum of mutually orthogonal rank one tensors). Using a more complicated argument (but vaguely similar in spirit) one can still recover that $\beta_{n,d_1,d_2,d_3} \leq \lambda \beta_{n-1,d_1-1,d_2-1,d_3-1}$.

1.2.5 The Inductive Argument (incorporating the Relaxed Base Case Inequality)

As discussed before, in general the base case inequality (7) does not hold and we are able to use only the relaxed base case inequality in Definition 1.9. We now indicate the main modification necessary in the inductive proof, skipping most other details from this overview.

Let $\Sigma' \subseteq \Sigma$ be the subset that exhibits the relaxed base case inequality in Definition 1.9. We consider the *effective influence* and *effective degree* of the function $f : \Sigma^n \rightarrow \mathbb{R}$. We recall that the standard influence of the i^{th} co-ordinate is

$$\mathbb{E}_{\substack{\mathbf{x}_{-i}, \\ x_i, x'_i \in \Sigma}} [(f(\mathbf{x}_{-i}, x_i) - f(\mathbf{x}_{-i}, x'_i))^2].$$

That is, the influence is the variance of the function on the i^{th} co-ordinate after randomly restricting the rest of the co-ordinates. We define the effective influence as

$$\mathbb{E}_{\substack{\mathbf{x}_{-i}, \\ x_i, x'_i \in \Sigma'}} [(f(\mathbf{x}_{-i}, x_i) - f(\mathbf{x}_{-i}, x'_i))^2],$$

which is similar, except that the variance is considered only over the subset Σ' .

We also indicate the related notion of the effective degree of f . We set up a suitable orthonormal basis \mathbf{B} of *characters* for (single co-ordinate) functions in $L_2(\Sigma; \mu_x)$. We ensure that $\mathbf{B} = \mathbf{B}_1 \cup \mathbf{B}_2$ so that characters in \mathbf{B}_1 span all functions that are constant on Σ' (including the All-1 function), and characters in \mathbf{B}_2 are zero outside Σ' . The effective degree of a monomial is then the degree when only the characters in \mathbf{B}_2 are counted towards the degree. The inductive proof is now carried out assuming that f not only has high degree, but also has high effective degree.

We do mention a crucial detail here. We do need to argue that starting with the original function $f : \Sigma^n \rightarrow [-1, 1]$ that has essentially high degree, we can “reduce” to the case where it has high effective degree as well. This argument does need that the original functions f, g, h are ℓ_∞ -bounded.¹⁰ As noted before, Lemmas 1.4, 1.10 could simply be false (for certain distributions μ) if only ℓ_2 -norm of the functions is assumed to be 1.

1.3 Organization

The rest of the paper is organized as follows. We start with preliminaries in Section 2. We set up the necessary machinery in Sections 3 and 4 that are needed to formulate the inductive statement towards proving the main analytical lemma under relaxed base case inequality. The proof of this lemma, which is divided into two parts, spans Sections 5 and 6. Finally, in Section 7 we derive our main analytical lemma, Lemma 1.4, from Lemma 1.10. In this section, we show how to get around the issue of the Horn-SAT embedding.

Section 8 is devoted to proving applications of our main analytical lemma.

2 Preliminaries

In this section, we record some basic definitions and tools from analysis of Boolean functions that will be used throughout the paper (see O’Donnell’s book [19] for reference). We begin with some notations.

We denote $A \lesssim B$ to refer to the fact that $A \leq C \cdot B$ for some absolute constant $C > 0$; we denote $A \gtrsim B$ to refer to the fact that $A \geq c \cdot B$ for some absolute constant $c > 0$. If this constant depends on some parameter, say m , we denote this fact by $A \lesssim_m B$. We use the normal letters x, y, z to denote elements from the domain Σ, Γ, Φ , respectively, and the bold face letters $\mathbf{x}, \mathbf{y}, \mathbf{z}$ to denote strings from $\Sigma^n, \Gamma^n, \Phi^n$, respectively.

2.1 Degrees and Homogenous Functions

We start with the definitions of a monomial and a degree- d monomial.

Definition 2.1. *Let Γ be a finite set, and let ν be some probability measure over Γ . A monomial over Γ is a function $\chi : (\Gamma, \nu) \rightarrow \mathbb{R}$ whose expectation according to ν is 0.*

¹⁰One needs ℓ_∞ -boundedness also while transforming the original distribution μ to achieve additional properties.

Definition 2.2. Let Γ be a finite set, $n \geq 1$, and let ν be some probability measure over Γ . A function $\chi: (\Gamma^n, \nu^{\otimes n}) \rightarrow \mathbb{R}$ is a degree d monomial if there are distinct indices i_1, \dots, i_d and monomials $\chi_{i_1}, \dots, \chi_{i_d}: \Gamma \rightarrow \mathbb{R}$ with respect to ν , such that

$$\chi(\mathbf{y}) = \prod_{j=1}^d \chi_{i_j}(y_{i_j}).$$

Based on these definitions, we now define homogeneous functions of degree d .

Definition 2.3. Let Γ be a finite set, $n \geq 1$, and let ν be some probability measure over Γ . A function $g: (\Gamma^n, \nu^{\otimes n}) \rightarrow \mathbb{R}$ is a homogenous degree d function if it can be written as a linear combination of monomials of degree d with respect to ν .

2.2 Efron-Stein Decomposition

For a product space $(\Gamma^n, \nu^{\otimes n})$, we will use the standard Efron-Stein decomposition. Given a function $g: (\Gamma^n, \nu^{\otimes n}) \rightarrow \mathbb{R}$, one may write (in a unique manner)

$$g(\mathbf{y}) = \sum_{i=0}^n g^{=i}(\mathbf{y}),$$

where $g^{=i}$ is a homogenous function of degree i . We denote by $V^{=i}(\Gamma^n, \nu^{\otimes n})$ the space of homogenous functions of degree i , and often omit the domain and the measure if these are clear from the context. Hence $g^{=i} \in V^{=i}$. The Efron-Stein decomposition is a refinement of the above. For each $i = 0, \dots, n$, one may write

$$g^{=i}(\mathbf{y}) = \sum_{S \subseteq [n], |S|=i} g^{=S}(\mathbf{y}),$$

where $g^{=S}(\mathbf{y})$ is a homogenous function of degree i whose value on \mathbf{y} only depends on \mathbf{y}_S , namely the co-ordinates in the set S . We denote by $V^{=S}(\Gamma^n, \nu^{\otimes n})$ the space of degree $|S|$ homogenous functions depending only on co-ordinates in S . So, $g^{=S} \in V^{=S}$. Furthermore, this decomposition is unique, and satisfies that $g^{=S}$ and $g^{=T}$ are orthogonal for any $S \neq T$, i.e. $\langle g^{=S}, g^{=T} \rangle = \mathbb{E}_{\mathbf{y} \sim \nu^{\otimes n}} [g^{=S}(\mathbf{y})g^{=T}(\mathbf{y})] = 0$. We define

$$g^{\geq T}(\mathbf{y}) = \sum_{S \supseteq T} g^{=S}(\mathbf{y}), \quad g^{\leq d}(\mathbf{y}) = \sum_{i=0}^d g^{=i}(\mathbf{y}), \quad g^{> d}(\mathbf{y}) = \sum_{i=d+1}^n g^{=i}(\mathbf{y}).$$

Next, we define the standard notion of the influence of a variable and the total influence of a function as follows:

Definition 2.4. The influence of a variable i on $g: (\Gamma^n, \mu^{\otimes n}) \rightarrow \mathbb{R}$ is defined as

$$I_i[g] = \mathbb{E}_{\mathbf{y}_{-i} \sim \mu^{n-1}, a, b \sim \mu} \left[(g(\mathbf{y}_{-i}, y_i = a) - g(\mathbf{y}_{-i}, y_i = b))^2 \right].$$

Definition 2.5. The total influence of g is defined as $I[g] = \sum_{i=1}^n I_i[g]$.

The following fact relates the total influence with the Efron-Stein decomposition of a function.

Fact 2.6. $I[g] = 2 \sum_{S \subseteq [n]} |S| \|g^{\neg S}\|_2^2 = 2 \sum_{i=1}^n i \|g^{\neg i}\|_2^2.$

Definition 2.7. The variance of $g: (\Gamma^n, \mu^{\otimes n}) \rightarrow \mathbb{R}$ is defined as

$$\text{var}(g) = \frac{1}{2} \mathbb{E}_{\mathbf{y}, \mathbf{y}' \sim \mu^{\otimes n}} [(g(\mathbf{y}) - g(\mathbf{y}'))^2] = \mathbb{E}_{\mathbf{y}} [g(\mathbf{y})^2] - \mathbb{E}[g(\mathbf{y})]^2.$$

2.3 A Markov Chain Lemma

We mention below a bound on the second largest eigenvalue of a Markov chain. We include the proof since we could not find this specific bound in literature (we need dependence on ξ to be linear and not quadratic).

Lemma 2.8. Let Σ be a finite alphabet of size at most m and $\mu(x, y)$ be a symmetric, connected distribution over $\Sigma \times \Sigma$. Suppose that $\mu(x) := \sum_{y \in \Sigma} \mu(x, y) \geq \alpha$ for all $x \in \Sigma$ and whenever $\mu(x, y) > 0$, we have $\mu(x, y) \geq \xi$. Let T be the associated Markov chain on Σ .¹¹ Then the second largest eigenvalue $\lambda_2(T) \leq 1 - \Omega_{\alpha, m}(\xi)$.

Proof. Note that μ is the stationary distribution of T . Let $f: \Sigma \rightarrow \mathbb{R}$ be the eigenvector of T corresponding to $\lambda_2(T)$. We have,

$$\mathbb{E}_{\mathbf{x} \sim \mu} [f(\mathbf{x})] = 0, \quad \mathbb{E}_{\mathbf{x} \sim \mu} [f(\mathbf{x})^2] = 1, \quad \lambda_2(T) = \langle f, Tf \rangle.$$

It follows that there is $x \in \Sigma$ such that $|f(x)| \geq 1$, and without loss of generality $f(x) \geq 1$. Since $\mathbb{E}[f] = 0$, there is y such that $f(y) \leq 0$. Since T is connected, there is a path $x = x^0 \rightarrow x^1 \rightarrow \dots \rightarrow x^\ell = y$, where $\ell \leq |\Sigma| \leq m$, so that $\mu(x^i, x^{i+1}) > 0$ for all $i \leq \ell - 1$. We note that

$$1 \leq |f(x) - f(y)| \leq \sum_{i=0}^{\ell-1} |f(x^{i+1}) - f(x^i)|,$$

so it follows that there is i such that $|f(x^{i+1}) - f(x^i)| \geq \frac{1}{\ell} \geq \frac{1}{m}$. Therefore,

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mu} [(f(\mathbf{x}) - f(\mathbf{y}))^2] \geq \mu(x, y) \frac{1}{m^2} \geq \frac{\xi}{m^2}.$$

On the other hand, the left hand side is $2\|f\|_2^2 - 2\langle f, Tf \rangle$, so we get that $\lambda_2(T) = \langle f, Tf \rangle \leq 1 - \frac{\xi}{2m^2}$, and the proof is concluded. \square

Lemma 2.9. In the setting of Lemma 2.8, if $f: (\Sigma^n, \mu^{\otimes n}) \rightarrow \mathbb{R}$ is any function, then

$$\|T^{\otimes n} f^{\neg S}\|_2 \leq (1 - \Omega_{\alpha, m}(\xi))^{|S|} \|f^{\neg S}\|_2.$$

Proof. Assume without loss of generality that $S = \{1, \dots, t\}$.

$$\|T^{\otimes n} f^{\neg S}\|_2 = \|T_1((T_2 \circ \dots \circ T_t) f^{\neg S})\|_2 \leq (1 - \Omega_{\alpha, m}(\xi)) \|(T_2 \circ \dots \circ T_t) f^{\neg S}\|_2,$$

where we used Lemma 2.8 and the fact that $(T_2 \circ \dots \circ T_t) f^{\neg S}$ has expectation 0 over x_1 for any setting of the remaining co-ordinates. The proof is concluded by iterating the above inequality over x_2, \dots, x_t . \square

¹¹Alternately, T is the adjacency matrix of a random walk on Σ where one leaves a vertex x with the edge probabilities $\mu(x, y)/\mu(x)$.

2.4 Effective Degrees and Effective Influences

Fix a distribution μ on $\Sigma \times \Gamma \times \Phi$ as in Lemma 1.10, and fix $\Sigma' \subseteq \Sigma$ to be the subset evidencing the fact that it satisfies the relaxed base case, Definition 1.9. We may choose an orthonormal basis of $L_2(\Sigma; \mu_x)$ as $B = B_1 \cup B_2$ where B_1 consists of functions that are constant on Σ' , and B_2 consists of functions only supported on Σ' and orthogonal to B_1 . Thus, a basis for $L_2(\Sigma^n; \mu_x^n)$ is given by

$$B^{\otimes n} = \left\{ \chi(\mathbf{x}) = \prod_i \chi_i(x_i) \mid \chi_i \in B_1 \cup B_2 \forall i \in [n] \right\},$$

and a given function $f: \Sigma^n \rightarrow \mathbb{R}$ can be uniquely written as

$$f(\mathbf{x}) = \sum_{\chi \in B^{\otimes n}} \hat{f}(\chi) \chi(\mathbf{x}), \quad \text{where } \hat{f}(\chi) = \langle f, \chi \rangle.$$

We now define the effective degree of a character χ , which is the number of components from $\{\chi_i\}$ that are not constant on Σ' .

Definition 2.10. Given $\chi \in B^{\otimes n}$, we define the effective degree of χ as

$$\text{effdeg}(\chi) = |\{i \in [n] \mid \chi_i \in B_2\}|.$$

To be compatible with the standard notion of degree, we will introduce a special notation for the trivial character in B_1 which is the constant 1 function on Σ , and we denote it by χ_{const} . Thus, we note that degree of a character χ is $|\{i \mid \chi_i \neq \chi_{\text{const}}\}|$, and as χ_{const} is in B_1 , one has that $\text{effdeg}(\chi) \leq \text{deg}(\chi)$ for all $\chi \in B^{\otimes n}$.

We also define the effective influence of a variable analogous to Definition 2.4, except we only resample the assignment to the variable if it belongs to Σ' .

Definition 2.11. For a function $f: \Sigma^n \rightarrow \mathbb{R}$ and $i \in [n]$, we define

$$I_{i,\text{effective}}[f] = \mathbb{E}_{(\mathbf{x}, \mathbf{y})} [(f(\mathbf{x}) - f(\mathbf{y}))^2],$$

where we sample $\mathbf{x} \sim \mu_x^{\otimes n}$, take $y_j = x_j$ for all $j \neq i$, and for the i^{th} co-ordinate, if $x_i \in \Sigma \setminus \Sigma'$ we take $y_i = x_i$, and if $x_i \in \Sigma'$, we sample $y_i \in \Sigma'$ independently.

Definition 2.12. For a function $f: \Sigma^n \rightarrow \mathbb{R}$, the total effective influence is defined as

$$I_{\text{effective}}[f] = \sum_{i=1}^n I_{i,\text{effective}}[f].$$

Based on how we defined the effective degree and the effective influences, we have the following fact analogous to Fact 2.6.

Fact 2.13. For a function $f: \Sigma^n \rightarrow \mathbb{R}$ and $i \in [n]$, we have

$$I_{i,\text{effective}}[f] = 2 \sum_{\chi: \chi_i \in B_2} \hat{f}(\chi)^2, \quad I_{\text{effective}}[f] = 2 \sum_{\chi} \text{effdeg}(\chi) \hat{f}(\chi)^2.$$

2.5 High Degree is preserved under Random Restrictions

We will often consider random restrictions of functions and would need to argue that if the original function has high degree, then so is the restricted function w.h.p. We recall that having high degree is formalized in terms of having low-stability. In the lemma below, a random restriction of a function $g : (\Sigma^n, \mu^n) \rightarrow \mathbb{R}$ includes every co-ordinate in the set I with probability s and then samples each co-ordinate in I according to ν_1 . For the restricted function $g_{I \rightarrow \mathbf{y}}$ on $[n] \setminus I$, each co-ordinate has marginal ν_2 . It is easily checked that to get the marginals "correct", we need $\mu = s\nu_1 + (1-s)\nu_2$.

Lemma 2.14. *Let ν_1, ν_2 be distributions over Σ whose support is full and the probability of each atom is at least α and $|\Sigma| \leq m$. Let $\mu = s\nu_1 + (1-s)\nu_2$. Then for some $c = c(\alpha) > 0$ we have*

$$\mathbb{E}_{\substack{|I| \sim sn \\ \mathbf{y} \sim \nu_1^I}} [\text{Stab}_{1-\xi}(g_{I \rightarrow \mathbf{y}}; \nu_2)] \leq \text{Stab}_{1-c(1-s)\xi}(g; \mu).$$

Proof. The left hand side is

$$\mathbb{E}_{(\mathbf{x}, \mathbf{x}')} [g(\mathbf{x})g(\mathbf{x}')],$$

where $(\mathbf{x}, \mathbf{x}')$ are sampled by taking, for each $i \in [n]$ independently, $x_i = x'_i \sim \nu_1$ with probability s ; otherwise with probability $1 - \xi$ we take $x_i = x'_i \sim \nu_2$, and with probability ξ we take $x_i, x'_i \sim \nu_2$ independently. We denote this Markov chain by \mathbb{T} , and note that the stationary distribution of it is μ , and also that as the support of ν_2 is full, \mathbb{T} is connected. Additionally, for every $a, b \in \Sigma$ we have that $\mathbb{T}(a, b) \geq (1-s)\xi\alpha^2$. We may now write the above expectation as

$$\langle g, \mathbb{T}^{\otimes n} g \rangle = \sum_{S, T} \langle g^{\neg S}, \mathbb{T}^{\otimes n} g^{\neg T} \rangle = \sum_S \langle g^{\neg S}, \mathbb{T}^{\otimes n} g^{\neg S} \rangle \leq \sum_S \|g^{\neg S}\|_2 \|\mathbb{T}^{\otimes n} g^{\neg S}\|_2.$$

Using Lemma 2.9, we get that

$$\langle g, \mathbb{T}^{\otimes n} g \rangle \leq \sum_S (1 - \Omega_{\alpha, m}((1-s)\xi))^{|S|} \|g^{\neg S}\|_2^2 = \text{Stab}_{1-c(1-s)\xi}(g),$$

where $c = c(m, \alpha) > 0$. □

2.6 Embedding into the Infinite Cyclic Group

In our proofs, towards arriving at a contradiction, we will often get an embedding of $\text{supp}(\mu)$ into an infinite cyclic group. Here, we argue that if $\text{supp}(\mu)$ cannot be embedded into a finite Abelian group, then it also cannot be embedded into an infinite cyclic group (say, $[0, 1)$ with addition mod 1). Therefore, the finiteness of the Abelian group is not essential to the definition of linear embeddability.

Claim 2.15. *Suppose that a finite set $S \subseteq \Sigma \times \Gamma \times \Phi$ cannot be linearly embedded. Then S cannot be embedded non-trivially into the infinite cyclic group.*

Proof. Suppose towards contradiction it can be. Then there are $\sigma : \Sigma \rightarrow [0, 1)$, $\phi : \Phi \rightarrow [0, 1)$ and $\gamma : \Gamma \rightarrow [0, 1)$ not all constant such that $\sigma(x) + \phi(y) + \gamma(z) = 0 \pmod{1}$. Without loss of generality, σ is non-constant. Consider the set of numbers $S = \text{Image}(\sigma) \cup \text{Image}(\phi) \cup \text{Image}(\gamma)$, let $r = |\Sigma| + |\Phi| + |\Gamma|$ and let $N = N(r) \in \mathbb{N}$ to be determined. Then $|S| \leq r$, so by Dirichlet's approximation theorem we may find integers p_i, q such that for each $s_i \in S$ we have that $\left| s_i - \frac{p_i}{q} \right| \leq \frac{1}{qN^{1/r}}$.

Let

$$\alpha = \min_{x, x' \sigma(x) \neq \sigma(x')} \min_{z \in \mathbb{Z}} |z + \sigma(x) - \sigma(x')|.$$

We choose $N = \left(\frac{3}{\alpha}\right)^r$, define σ' by $\sigma'(x) = \frac{p_i}{q} \pmod{1}$ if $\sigma(x) = s_i$, and similarly define ϕ', γ' .

1. First, we show that σ', ϕ', γ' is an embedding. Fix $(x, y, z) \in \text{supp}(\mu)$; then we have

$$\sigma'(x) + \phi'(y) + \gamma'(z) = \sigma(x) + \phi(y) + \gamma(z) + \Delta,$$

where $|\Delta| \leq \frac{3}{qN^{1/r}}$. Noting that $\sigma(x) + \phi(y) + \gamma(z)$ is an integer (as it is 0 mod 1), it follows that $\sigma'(x) + \phi'(y) + \gamma'(z)$ is very close to an integer, up to $\frac{3}{qN^{1/r}} < 1/q$. On the other hand, by definition of σ', ϕ', γ' it is a number of the form P/q for some integer P , hence it can either be an integer or at least $1/q$ far from one. It follows that it is an integer, so $\sigma'(x) + \phi'(y) + \gamma'(z) = 0 \pmod{1}$.

2. Second, we show that at least one of them is not constant. Indeed, take $x, x' \in \Sigma$ on which σ differs, and suppose towards contradiction that $\sigma'(x) - \sigma'(x') = 0$. Let i, j be such that $\sigma'(x) = p_i/q$, $\sigma'(x') = p_j/q$, and $\sigma(x) = s_i$, $\sigma(x') = s_j$. Then we get that $(p_i - p_j)/q$ is an integer, and as $s_i - s_j$ is Δ -close to it for $|\Delta| \leq \frac{2}{qN^{1/r}}$, we get that $s_i - s_j$ is $\frac{\alpha}{3}$ close to an integer. This contradicts the definition of α . \square

2.7 The Additive Base Case

In this section, we deduce a certain auxiliary (base case) inequality that follows solely under the assumption that the distribution μ has no linear embedding. We emphasize that it holds irrespective of whether or not μ has a Horn-SAT embedding. The inequality is used while reducing the dimension and making it comparable to the degree during our inductive proof (as in the overview Section 1.2.4).

Claim 2.16. *Let μ be a distribution on $\Sigma \times \Gamma \times \Phi$ that has no linear embedding. Then there exists $c_1 = c_1(\mu) > 0$, such that for $f: \Sigma \rightarrow \mathbb{R}$, $g: \Phi \rightarrow \mathbb{R}$ and $h: \Gamma \rightarrow \mathbb{R}$ that each have average equal to 0 we have that*

$$\left| \mathbb{E}_{(x,y,z) \sim \mu} [f(x)(g(y) + h(z))] \right| \leq (1 - c_1) \|f\|_2 \|g + h\|_2.$$

Proof. Assume this is not the case, so that we may find a sequence of functions (f_m, g_m, h_m) such that $\|f_m\|_2 = \|g_m + h_m\|_2 = 1$, $\mathbb{E}[f_m] = \mathbb{E}[g_m] = \mathbb{E}[h_m] = 0$ and $|\mathbb{E}_{(x,y,z) \sim \mu} [f_m(x)(g_m(y) + h_m(z))]| \geq 1 - \frac{1}{m}$. Passing to subsequences, we may assume that f_m converges to a function f , g_m converges to a function g and h_m converges to a function h , so that we get $\mathbb{E}[f] = \mathbb{E}[g] = \mathbb{E}[h] = 0$, $\|f\|_2 = \|g + h\|_2 = 1$ and $|\mathbb{E}_{(x,y,z) \sim \mu} [f(x)(g(y) + h(z))]| \geq 1$. By Cauchy-Schwarz we have that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu} [f(x)(g(y) + h(z))] \right| \leq \sqrt{\mathbb{E}_{(x,y,z) \sim \mu} [f(x)^2]} \sqrt{\mathbb{E}_{(x,y,z) \sim \mu} [(g(y) + h(z))^2]} = \|f\|_2 \|g + h\|_2 = 1,$$

hence we get that Cauchy-Schwarz is tight and so $f(x) = \theta(g(y) + h(z))$ for some $\theta \in \{-1, 1\}$ for all $(x, y, z) \in \text{supp}(\mu)$. As the 2-norm of f is 1 and its average is 0, the function f is not constant, and so after dividing them by a sufficiently large constant and adding a constant (so that their image is contained in $[0, 1]$), we get that $f, -\theta g, -\theta h$ form a non-trivial embedding of μ . Together with Claim 2.15, this contradicts the assumption that μ has no linear embedding. \square

Lemma 2.17. *Let μ be a distribution on $\Sigma \times \Gamma \times \Phi$ that has no linear embedding and $\mu_{y,z}$ is uniform. Then there exists $c_1 = c_1(\mu) > 0$, such that for $f: \Sigma \rightarrow \mathbb{R}$, $g: \Phi \rightarrow \mathbb{R}$ and $h: \Gamma \rightarrow \mathbb{R}$ we have that*

$$\left| \mathbb{E}_{(x,y,z) \sim \mu} [f(x)(g(y) + h(z))] \right| \leq \|f\|_2 \sqrt{(\mathbb{E}[g] + \mathbb{E}[h])^2 + (1 - c_1)(W_{=1}[g] + W_{=1}[h])}.$$

Here $W_{=1}[g]$ denotes the variance $\mathbb{E}[g^2] - \mathbb{E}[g]^2$ (and similarly for $W_{=1}[h]$).

Proof. Write $f = \mathbb{E}[f] + f^{\perp 1}$, $g = \mathbb{E}[g] + g^{\perp 1}$ and $h = \mathbb{E}[h] + h^{\perp 1}$ so that

$$\mathbb{E}_{(x,y,z) \sim \mu} [f(x)(g(y) + h(z))] = \mathbb{E}[f](\mathbb{E}[g] + \mathbb{E}[h]) + \mathbb{E}_{(x,y,z) \sim \mu} [f^{\perp 1}(x)(g^{\perp 1}(y) + h^{\perp 1}(z))].$$

By Claim 2.16 we get that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu} [f^{\perp 1}(x)(g^{\perp 1}(y) + h^{\perp 1}(z))] \right| \leq (1 - c_1) \|f^{\perp 1}\|_2 \|g^{\perp 1} + h^{\perp 1}\|_2,$$

and so

$$\begin{aligned} \left| \mathbb{E}_{(x,y,z) \sim \mu} [f(x)(g(y) + h(z))] \right| &\leq |\mathbb{E}[f]| |\mathbb{E}[g] + \mathbb{E}[h]| + (1 - c_1) \|f^{\perp 1}\|_2 \|g^{\perp 1} + h^{\perp 1}\|_2 \\ &\leq \sqrt{\mathbb{E}[f]^2 + \|f^{\perp 1}\|_2^2} \sqrt{(\mathbb{E}[g] + \mathbb{E}[h])^2 + (1 - c_1)^2 \|g^{\perp 1} + h^{\perp 1}\|_2^2} \\ &= \|f\|_2 \sqrt{(\mathbb{E}[g] + \mathbb{E}[h])^2 + (1 - c_1)^2 \|g^{\perp 1} + h^{\perp 1}\|_2^2} \end{aligned}$$

where we used Cauchy-Schwarz. As $\mu_{y,z}$ is uniform we get that $\|g^{\perp 1} + h^{\perp 1}\|_2^2 = \|g^{\perp 1}\|_2^2 + \|h^{\perp 1}\|_2^2 = W_{=1}[g] + W_{=1}[h]$ and the proof is concluded \square

3 The main homogenous statement

In this section, we reduce Lemma 1.10 to a similar statement about homogenous functions. Namely, we state Lemma 3.1, and show that it implies Lemma 1.10. There are two key differences between the two lemmas. The first of which is that whereas Lemma 1.10 is only concerned with bounded functions, Lemma 3.1 applies to general real-valued functions. It can be seen, however, that if the base case fails, i.e. if there are non-trivial functions for which (7) fails with $\tau = 1$, then Lemma 1.10 would not hold for unbounded functions. Therefore, in exchange for relaxing ℓ_∞ -boundedness to ℓ_2 -boundedness, we get to assume that the functions f, g, h are all homogenous, and the effective degree of f is significant. A precise statement follows.

Lemma 3.1. *There are $C(m, \alpha) > 0$ and $D(m, \alpha) \in \mathbb{N}$ such that the following holds for all $d \geq D$. Suppose μ is a distribution over $\Sigma \times \Gamma \times \Phi$ satisfying the conditions of Lemma 1.10, and $f: \Sigma^n \rightarrow \mathbb{R}$, $g: \Gamma^n \rightarrow \mathbb{R}$, $h: \Phi^n \rightarrow \mathbb{R}$ are functions such that*

1. *f is homogenous of degree at most $d \log^{10} d$, g is homogenous of degree at most $d \log^{10} d$ and at least $\frac{d}{\log^{10} d}$, and h is homogenous of degree at most $2d \log^{10} d$.*

2. The effective degree of f is at least $\frac{d}{\log^{200} d}$. Namely, for all $\chi \in B^n$ such that $\widehat{f}(\chi) \neq 0$, we have that $\text{eff-degree}(\chi) \geq \frac{d}{\log^{200} d}$.

Then

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right| \leq (1 - \delta)^{d/\log^C(d)} \|f\|_2 \|g\|_2 \|h\|_2.$$

The rest of this section is devoted to showing that Lemma 3.1 implies Lemma 1.10. We prove the implication using the following sequence.

$$\text{Lemma 3.1} \implies \text{Lemma 3.3} \implies \text{Lemma 3.2} \implies \text{Lemma 1.10.} \quad (8)$$

Then, in Sections 4, 5 and 6 we prove Lemma 3.1.

3.1 Soft truncations

In this section, we define a couple of noise operators that will be used in stating the intermediate lemmas from (8).

The noise operators. Let $\xi \in (0, 1]$ be some parameter.

1. We define the operator $T_{1-\xi}$ acting on $L_2(\Sigma)$ as follows. Consider the Markov chain on Σ that on $x \in \Sigma$ generates $x' \sim T_{1-\xi}x$ by: with probability $1 - \xi$ we take $x' = x$, and otherwise we re-samples $x' \sim \mu_x$. We let $T_{1-\xi}$ be the corresponding averaging operator on $L_2(\Sigma)$, i.e. $T_{1-\xi}f(x) = \mathbb{E}_{x' \sim T_{1-\xi}x} [f(x')]$.
2. We define analogs of the operator $T_{1-\xi}$ on $L_2(\Gamma)$, $L_2(\Phi)$ in the same way. For notational convenience, we will use the same notation for them, i.e. $T_{1-\xi}$, and it will be clear from the context which operator is applied.
3. We define the operators $E_{1-\xi}$. Let $\Sigma' \subseteq \Sigma$ be evidencing the fact that μ satisfies the relaxed base case. Consider the Markov chain on Σ that on $x \in \Sigma$ generates $x' \sim E_{1-\xi}x$ by: if $x \in \Sigma \setminus \Sigma'$, we take $x' = x$. Otherwise, with probability $1 - \xi$ we take $x' = x$, and with probability ξ we re-sample $x' \sim \mu_x | \Sigma'$. We let $E_{1-\xi}$ be the corresponding averaging operator on $L_2(\Sigma)$, i.e. $E_{1-\xi}f(x) = \mathbb{E}_{x' \sim E_{1-\xi}x} [f(x')]$.

The noise operator $T_{1-\xi}$ when applied on a function f dampens (in the ℓ_2 measure) the high-degree terms from f . Analogously, as we will see, the operator $E_{1-\xi}$ when applied on a function dampens the high *effective-degree* terms from the function.

3.2 Intermediate Lemmas

With these operators, we can now state the relaxed analogs of Lemma 3.1 wherein the degree conditions are replaced with analogous soft truncations (but in return, we still retain the boundedness of the functions).

3.2.1 Softly truncating g from both sides, and f from above

In the first relaxation, we softly truncate the degree of the function g from both sides and the degree of the function f from above using the noise operators $T_{1-\xi}$.

Lemma 3.2. *For all $\alpha > 0$, $m \in \mathbb{N}$, and $M > 0$, there is $\xi_0 > 0$ such that the following holds for all $0 < \xi \leq \xi_0$. Suppose μ is a distribution over $\Sigma \times \Gamma \times \Phi$ satisfying the conditions of Lemma 1.10, and $f: \Sigma^n \rightarrow [-1, 1]$, $g: \Gamma^n \rightarrow [-1, 1]$, $h: \Phi^n \rightarrow [-1, 1]$ are functions. Then*

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} \left[T_{1-M\xi/\log(1/\xi)^3} f(\mathbf{x}) (T_{1-\xi/2} - T_{1-\xi}) g(\mathbf{y}) h(\mathbf{z}) \right] \right| \lesssim \frac{1}{\log^6(1/\xi)}.$$

3.2.2 Softly truncating the effective degree of f from below

In the next relaxation, we further softly truncate the *effective degree* of f from below. We do this using the operator $E_{1-\xi}$.

Lemma 3.3. *For all $\alpha > 0$, $m \in \mathbb{N}$, and $M > 0$, there is $\xi_0 > 0$ such that the following holds for all $0 < \xi \leq \xi_0$. Suppose μ is a distribution over $\Sigma \times \Gamma \times \Phi$ satisfying the conditions of Lemma 1.10, and $f: \Sigma^n \rightarrow [-1, 1]$, $g: \Gamma^n \rightarrow [-1, 1]$, $h: \Phi^n \rightarrow [-1, 1]$ are functions. Then*

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} \left[(I - E_{1-M\xi \log(1/\xi)^{100}}) T_{1-M\xi/\log(1/\xi)^3} f(\mathbf{x}) (T_{1-\xi/2} - T_{1-\xi}) g(\mathbf{y}) h(\mathbf{z}) \right] \right| \lesssim \frac{1}{\log^6(1/\xi)}.$$

We defer the proofs of the implications Lemma 3.1 \implies Lemma 3.3 \implies Lemma 3.2 \implies Lemma 1.10 to the Appendix. All these proofs follow from the standard arguments that use the fact that the noise operator essentially gets rid of the high degree part of the functions.

4 The main inductive statement, and set up

In this section, we re-phrase and state a sharper version of Lemma 3.1, which will be more convenient for us to work with.

4.1 The parameter β_{n,d_1,d_2,d_3}

We begin by defining the parameter β'_{n,d_1,d_2,d_3} for all $d_1, d_2, d_3, n \in \mathbb{N}$ (which is a close variant of β_{n,d_1,d_2,d_3} which will be shortly defined):

$$\beta'_{n,d_1,d_2,d_3} = \max_{\substack{f: \Sigma^n \rightarrow \mathbb{R} \text{ degree } d_1 \text{ homogenous,} \\ \text{eff-deg}(f) \geq d_1 / \log^{20}(d_1) \\ g: \Gamma^n \rightarrow \mathbb{R} \text{ degree } d_2 \text{ homogenous,} \\ h: \Phi^n \rightarrow \mathbb{R} \text{ degree at most } d_3 \text{ homogenous}}} \frac{|\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})]|}{\|f\|_2 \|g\|_2 \|h\|_2}.$$

Using Cauchy-Schwarz, it is clear that $\beta_{n,d_1,d_2,d_3} \leq 1$ always. The following lemma asserts that if d_1, d_2, d_3 are roughly the same up to poly-logs, then β_{n,d_1,d_2,d_3} is actually almost exponentially decaying in d .

Lemma 4.1. *For all $K, m \in \mathbb{N}$ and $\alpha > 0$, there are $C, d_0 \in \mathbb{N}$ such that the following holds. Let μ be a distribution over $\Sigma \times \Gamma \times \Phi$ as in Lemma 3.1, and let $d_1, d_2, d_3 \leq n$ be such that $d_i \leq d_j \log^K(d_j)$ for all $i, j = 1, 2, 3$, and $d_i \geq d_0$ for all i . Then*

$$\beta'_{n,d_1,d_2,d_3} \leq 2^{-\frac{d_1}{\log C(d_1)}}.$$

We note that Lemma 4.1 immediately implies Lemma 3.1, hence we will focus henceforth on proving that 4.1 is true. We will actually need to adjust the parameter β'_{n,d_1,d_2,d_3} a bit and state a statement similar to Lemma 4.1 but stronger. The main difference will be that instead of working with the function f over \mathbf{x} , we will work with the function $F(\mathbf{y}, \mathbf{z})$. Given a function f , as in μ we have that \mathbf{y}, \mathbf{z} implies \mathbf{x} , we may define $F: (\Gamma^n \times \Phi^n, \mu_{\mathbf{y},\mathbf{z}}^{\otimes n}) \rightarrow \mathbb{R}$ by

$$F(\mathbf{y}, \mathbf{z}) = f(\mathbf{x}),$$

where \mathbf{x} is the unique point in Σ^n such that $(x_i, y_i, z_i) \in \text{supp}(\mu)$ for all $i \in [n]$. We denote this as an operator

$$W: L_2(\Sigma^n, \mu_x^{\otimes n}) \rightarrow L_2(\Gamma^n \times \Phi^n, \mu_{\mathbf{y},\mathbf{z}}^{\otimes n}), \text{ so that } F = Wf.$$

It will be more convenient for us to work with F since the distribution over y, z is uniform (whereas the distribution over x may not be), but to facilitate that we need to translate the information that f is homogenous and has high effective degree.

4.1.1 Setting up the basis to define the effective degree of F

Consider any two basis elements $\chi, \chi' \in B_1 \cup B_2$, and note that

$$\langle W\chi, W\chi' \rangle = \langle \chi, \chi' \rangle = 1_{\chi=\chi'},$$

so $\{W\chi\}_{B_1 \cup B_2}$ form a partial orthonormal basis for $L_2(\Gamma^n \times \Phi^n, \mu_{\mathbf{y},\mathbf{z}}^{\otimes n})$. We may complete it to an orthonormal basis using some set $\{\tilde{\chi}\}_{\tilde{\chi} \in C}$. Thus, any function $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{R}$ may be written as

$$F(\mathbf{y}, \mathbf{z}) = \sum_{\tilde{\chi} \in (WB_1 \cup WB_2 \cup C)^{\otimes n}} \widehat{F}(\tilde{\chi}) \tilde{\chi}(\mathbf{y}, \mathbf{z}),$$

where

$$\tilde{\chi}(\mathbf{y}, \mathbf{z}) = \prod_{i=1}^n \tilde{\chi}_i(y_i, z_i), \quad \text{and } \widehat{F}(\tilde{\chi}) = \langle F, \tilde{\chi} \rangle.$$

Noting that the function $W\chi_{\text{const}}$ is the all 1 function, we define the following notion of degree for monomials over \mathbf{y}, \mathbf{z} .

Definition 4.2. *The degree of $\tilde{\chi} \in (WB_1 \cup WB_2 \cup C)^{\otimes n}$ is defined to be the number of coordinates i such that $\tilde{\chi}_i \neq W\chi_{\text{const}}$.*

We also define the effective degree.

Definition 4.3. *The effective degree of $\tilde{\chi} \in (WB_1 \cup WB_2 \cup C)^{\otimes n}$ is defined to be the number of coordinates i such that $\tilde{\chi}_i \in WB_2$.*

Finally, we define a property of functions $F(\mathbf{y}, \mathbf{z})$ which is equivalent to it being in the image of the operator W .

Definition 4.4. Consider the graph $G = (V, E)$ whose vertices are $\Gamma \times \Phi$ and (y, z) and (y', z') are adjacent if there is some $x \in \Sigma$ such that $(x, y, z), (x, y', z') \in \text{supp}(\mu)$. Consider the graph $G^{\otimes n} = (V^{\otimes n}, E')$ where $E' = \{((\mathbf{y}, \mathbf{z}), (\mathbf{y}', \mathbf{z}')) \mid ((y_i, z_i), (y'_i, z'_i)) \in E \forall i\}$.

We say a function $F: \Gamma^{\otimes n} \times \Phi^{\otimes n} \rightarrow \mathbb{R}$ is constant on connected components if F is constant on all of the connected components of $G^{\otimes n}$.

We have the following claim.

Claim 4.5. For a function $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{R}$, the following are equivalent:

1. F is constant on connected components;
2. There is $f: \Sigma^n \rightarrow \mathbb{R}$ such that $F = Wf$.
3. $\widehat{F}(\chi) = 0$ for $\chi \notin (WB_1 \cup WB_2)^{\otimes n}$.

Proof. It is clear that the third item implies the second item, and that the second item implies the first item. We next show that the first item implies the third item. Towards this end, assume that F is constant on connected component, and let $\chi \notin (WB_1 \cup WB_2)^{\otimes n}$. Then $\chi_i \in C$ for some $i = 1, \dots, n$, without loss of generality assume that $i = 1$. Then

$$\widehat{F}(\chi) = \mathbb{E}_{(\mathbf{y}_{-1}, \mathbf{z}_{-1}) \sim \mu_{y,z}^{\otimes n-1}} \left[\prod_{i=2}^n \chi_i(y_i, z_i) \mathbb{E}_{(y_1, z_1) \sim \mu_{y,z}} [F(\mathbf{y}, \mathbf{z}) \chi_1(y_1, z_1)] \right].$$

Fix $\mathbf{y}_{-1} = y_{-1}$ and $\mathbf{z}_{-1} = z_{-1}$; we show that the inner expectation is 0. To see this, first note that since $F_{\{2, \dots, n\} \rightarrow (y_{-1}, z_{-1})}(y_1, z_1)$, as a function of y_1, z_1 , is constant on connected components, it suffices to show that each basis element in C is perpendicular to functions that are constant on connected components. To see that, it suffices to show that for each connected component C of G , the indicator of it $1_C(y_1, z_1)$ is in $\text{span}(WB_1 \cup WB_2)$, and we next show this is true.

Indeed, let C be some connected component. Define

$$f(x) = 1_C(y, z)$$

where (y, z) are chosen so that $(x, y, z) \in \text{supp}(\mu)$. We note that this is well defined, since if we have two (y, z) and (y', z') such that (x, y, z) and (x, y', z') are both in $\text{supp}(\mu)$, then they lie in the same connected component of H and hence $1_C(y, z) = 1_C(y', z')$. It follows that $1_C = Wf$, hence $1_C \in \text{span}(WB_1 \cup WB_2)$. \square

The following claim encapsulates the properties we need about $F = Wf$.

Claim 4.6. Suppose that $f: \Sigma^n \rightarrow \mathbb{R}$ is a degree d homogenous function, and for all χ such that $\widehat{f}(\chi) \neq 0$ we have $\text{eff} - \deg(\chi) \geq d'$. Then $F = Wf$ is a degree d homogenous function, it is constant on connected components and for all $\tilde{\chi}$ such that $\widehat{F}(\tilde{\chi}) \neq 0$ we have that $\text{eff} - \deg(\chi) \geq d'$.

Proof. Note that writing $f(\mathbf{x}) = \sum_{\chi \in (B_1 \cup B_2)^{\otimes n}} \widehat{f}(\chi) \chi(\mathbf{x})$, we get that

$$F(\mathbf{y}, \mathbf{z}) = Wf(\mathbf{y}, \mathbf{z}) = \sum_{\chi \in (B_1 \cup B_2)^{\otimes n}} \widehat{f}(\chi) (W\chi)(\mathbf{y}, \mathbf{z}),$$

so we get that $\widehat{F}(\tilde{\chi})$ is non-zero only if $\tilde{\chi} \in (WB_1 \cup WB_2)^{\otimes n}$, and it is equal to $\widehat{f}(\chi)$ where $\tilde{\chi} = W\chi$. From this and Claim 4.5, the assertions of the claim immediately follow. \square

4.1.2 The main inductive statement

We are now ready to define β_{n,d_1,d_2,d_3} and state a stronger form of Lemma 4.1. We define

$$\beta_{n,d_1,d_2,d_3} = \max_{\substack{F: \Gamma^n \times \Phi^n \rightarrow \mathbb{R} \text{ degree } d_1 \text{ homogenous,} \\ F \text{ is constant on connected components,} \\ \text{eff-deg}(F) \geq d_1 / \log^{20}(d_1), \\ g: \Gamma^n \rightarrow \mathbb{R} \text{ degree } d_2 \text{ homogenous,} \\ h: \Phi^n \rightarrow \mathbb{R} \text{ degree at most } d_3 \text{ homogenous.}}} \frac{|\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [F(\mathbf{y}, \mathbf{z})g(\mathbf{y})h(\mathbf{z})]|}{\|F\|_2 \|g\|_2 \|h\|_2}.$$

Lemma 4.7. *For all $K, m \in \mathbb{N}$ and $\alpha > 0$, there are $C, d_0 \in \mathbb{N}$ such that the following holds. Let μ be a distribution over $\Sigma \times \Gamma \times \Phi$ as in Lemma 3.1, and let $d_1, d_2, d_3 \leq n$ be such that $d_i \leq d_j \log^K(d_j)$ for all $i, j = 1, 2, 3$, and $d_i \geq d_0$ for all i . Then*

$$\beta_{n,d_1,d_2,d_3} \leq 2^{-\frac{d_1}{\log^C(d_1)}}.$$

Lemma 4.1 follows from Lemma 4.7 by the virtue of Claim 4.6.

The proof of Lemma 4.7 spans Sections 5, 6, and in the rest of the current section we set up some machinery and explain the high level overview of the argument. Our overall approach to the proof of Lemma 4.7 is inductive, and so we will need to be able to write down a function over n variables as linear combination of products of a function of $n - 1$ variables and a function over a single variable. Towards this end, we employ the singular-value decomposition as presented in the next section.

Throughout, we will have a partition the set of coordinates $[n]$ into $I \cup J$ where $|I| = n - 1$ and $|J| = 1$, which for now will be arbitrary (we will later explain how to choose it so that it satisfies several additional properties that we need).

4.2 SVD decompositions

In this section, we state claims asserting an SVD decomposition for homogenous and non-homogenous functions. The proofs of these claims are deferred to the Appendix. Throughout, I, J is a partition of $[n]$ wherein $|J| = 1$ and $|I| = n - 1$.

4.2.1 The SVD decompositions for homogenous functions

The following decomposition claim is phrased in terms of the function g . However, it applies to functions over z , as well as to functions over y, z . We will use it for both g and h , and may use it also for F . However, for F we need a few additional properties, which we establish in Claim 4.10.

Claim 4.8. *If $g: \Gamma^n \rightarrow \mathbb{R}$ is a homogenous function of degree d and $\|g\|_2 = 1$, then we may write*

$$g(\mathbf{y}) = \sum_{r=1}^m \lambda_r g_r(\mathbf{y}_I) g'_r(\mathbf{y}_J),$$

and $R = \{i \mid \lambda_i \neq 0\}$ where

1. For $r \in R$, $g_r: \Gamma^I \rightarrow \mathbb{R}$ is an orthonormal set of functions.
2. If $1 \in R$, then g_1 is homogenous of degree d and for $r \geq 2$ in R the function g_r is homogenous of degree $d - 1$.

3. For $r \in R$, $g'_r: \Gamma^J \rightarrow \mathbb{R}$ is an orthonormal set of functions.

4. $g'_1 \equiv 1$.

5. $\sum_{r=1}^m \lambda_r^2 = 1$.

Proof. We defer this proof to the Appendix. □

A natural question is what can be said about the coefficients λ_i in the above SVD decomposition, and indeed this will help us in choosing an appropriate partition. We have

Claim 4.9. *Let $g: \Gamma^n \rightarrow \mathbb{R}$ is a homogenous function of degree d_1 and $\|g\|_2 = 1$, and write*

$$g(\mathbf{y}) = \sum_{r=1}^m \lambda_r g_r(\mathbf{y}_I) g'_r(\mathbf{y}_J),$$

as in Claim 4.8. Then if j is the unique variable in the set J in the partition $[n] = I \cup J$, then

$$\lambda_1^2 = 1 - \frac{1}{2} I_j[g].$$

Proof. Consider $I_j[g]$, and note that

$$\begin{aligned} I_j[g] &= \mathbb{E}_{\mathbf{y} \sim \mu^{n-1}, a, b} [(g(\mathbf{y}_I, a) - g(\mathbf{y}_I, b))^2] = \mathbb{E}_{\mathbf{y} \sim \mu^{n-1}, a, b} \left[\left(\sum_r \lambda_r g_r(\mathbf{y}_I) (g'_r(a) - g'_r(b)) \right)^2 \right] \\ &= \sum_{r_1, r_2} \lambda_{r_1} \lambda_{r_2} \langle g_{r_1}, g_{r_2} \rangle \mathbb{E}_{a, b} [(g'_{r_1}(a) - g'_{r_1}(b))(g'_{r_2}(a) - g'_{r_2}(b))]. \end{aligned}$$

For $r_1 \neq r_2$ we have $\langle g_{r_1}, g_{r_2} \rangle = 0$, so the last sum is equal to

$$\sum_r \lambda_r^2 \mathbb{E}_{a, b} [(g'_r(a) - g'_r(b))^2] = 2 \sum_r \lambda_r^2 \text{var}(g'_r) = 2 \sum_{r \neq 1} \lambda_r^2,$$

as the variance of g'_1 is 0, and the variance of any other g'_r is 1. Hence

$$1 - \frac{I_j[g]}{2} = 1 - \sum_{r \neq 1} \lambda_r^2 = \lambda_1^2.$$

□

We next state an SVD decomposition statement that addresses the function F .

Claim 4.10. *Suppose $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{R}$ is a homogenous function of degree d which is constant on connected components, $\|F\|_2 = 1$, and the effective degree of each monomial in F is at least d' . Then we may write*

$$F(\mathbf{y}, \mathbf{z}) = \sum_{t=1}^m \gamma_t F_t(\mathbf{y}_I, \mathbf{z}_I) F'_t(\mathbf{y}_J, \mathbf{z}_J),$$

and $T = \{t \mid \gamma_t \neq 0\}$ where

1. For $t \in T$, $F_t: \Gamma^I \times \Phi^I \rightarrow \mathbb{R}$ is an orthonormal set of functions.
2. If $1 \in T$, then F_1 is homogenous of degree d and for $t \geq 2$ in T the function F_t is homogenous of degree $d - 1$.
3. If $1 \in T$, then in F_1 the effective degree of monomial is at least d' . and for $T \geq 2$ in T the effective degree of each monomial in F_t is at least $d' - 1$.
4. The functions F_t and F'_t are constant on connected components for all t .
5. For $t \in T$, $F'_t: \Gamma^J \times \Phi^J \rightarrow \mathbb{R}$ is an orthonormal set of functions.
6. $F'_1 \equiv 1$.
7. $\sum_{t=1}^m \gamma_t^2 = 1$.

Proof. The proof of this claim is also deferred to the Appendix. □

4.2.2 The SVD decompositions for non-homogenous functions

In this section, we state the decomposition for non-homogenous functions. For such functions, we do not get the guarantee that one of the functions in the decomposition is the constant 1 function (and hence we do not have one of the functions in the decomposition to have *full degree*).

Claim 4.11. *If $g: \Gamma^n \rightarrow \mathbb{R}$ satisfies that $g^{\leq d} \equiv 0$ and $\|g\|_2 = 1$, then we may write*

$$g(\mathbf{y}) = \sum_{r=1}^m \lambda_r g_r(\mathbf{y}_I) g'_r(\mathbf{y}_J),$$

and $R = \{i \mid \lambda_i \neq 0\}$ where

1. For $r \in R$, $g_r: \Gamma^I \rightarrow \mathbb{R}$ is an orthonormal set of functions.
2. For $r \in R$ we have that $(g_r)^{\leq d-1} \equiv 0$.
3. For $r \in R$, $g'_r: \Gamma^J \rightarrow \mathbb{R}$ is an orthonormal set of functions.
4. $\sum_{r=1}^m \lambda_i^2 = 1$.

Proof. The proof is similar to the proof of Claim 4.8 and is deferred to the Appendix. □

Claim 4.12. *Suppose $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{R}$ satisfies $F^{\leq d} \equiv 0$, is constant on connected components, $\|F\|_2 = 1$, and the effective degree of each monomial in F is at least d' . Then we may write*

$$F(\mathbf{y}, \mathbf{z}) = \sum_{t=1}^m \gamma_t F_t(\mathbf{y}_I, \mathbf{z}_I) F'_t(\mathbf{y}_J, \mathbf{z}_J),$$

and $T = \{t \mid \gamma_t \neq 0\}$ where

1. For $t \in T$, $F_t: \Gamma^I \times \Phi^I \rightarrow \mathbb{R}$ is an orthonormal set of functions.

2. For $t \in T$, $(F_t)^{\leq d-1} \equiv 0$, and each monomial in F_t has effective degree at least $d' - 1$.
3. The functions F_t and F_t' are constant on connected components for all t .
4. For $t \in T$, $F_t': \Gamma^J \times \Phi^J \rightarrow \mathbb{R}$ is an orthonormal set of functions.
5. $\sum_{t=1}^m \gamma_t^2 = 1$.

Proof. The proof is similar to the proof of Claim 4.10, as is omitted. □

4.3 Roadmap of the proof of Lemma 4.7

The proof of Lemma 4.7 comprises two steps. In this section, we give an overview of these two steps. In Section 5, we formally show how to perform the first step of reducing to the case of nearly-linear degree and in Section 6, we show how to prove Lemma 4.7 when the degrees of the functions are nearly-linear in the number of variables.

The first step: reducing to the case of nearly-linear degree. In the first step, we show an inductive step that manages to show (roughly speaking) if n is much larger than $\max(d_1, d_2, d_3)$, then one has

$$\beta_{n,d_1,d_2,d_3} \leq \min(\beta_{n-1,d_1,d_2,d_3}, (1-c)\beta_{n-1,d_1-1,d_2-1,d_3-1}),$$

where $c = c(m, \alpha) > 0$. Iterating this bound, we either manage to prove that β_{n,d_1,d_2,d_3} is at most $(1-c)^{d_1/2}$, or else we reduce n to be small enough so that the inductive step can no longer be made. In the first case we are done, and in the second case we have shown that

$$\beta_{n,d_1,d_2,d_3} \leq \beta_{n',d'_1,d'_2,d'_3}$$

where $d'_i \geq d_i/2$ for $i = 1, 2, 3$ and $n' \lesssim \max(d'_1, d'_2, d'_3)$. Since d_1, d_2, d_3 are originally of the same order up to poly-logs, we are reduced to proving a variant of Lemma 4.7 in the case that n', d'_1, d'_2, d'_3 are of the same order up to poly-logs; we refer to this case as the “nearly-linear degree case”.

The overarching idea of our argument is that since n is much larger than d_1, d_2, d_3 , we may choose the partition $I \cup J$ so that the variable in J has small influence in all of F, g, h , and hence when we use the SVD decompositions from Claims 4.8 and 4.10 (and using Claim 4.9), we get that most of the mass of the functions lies on F_1, g_1 and h_1 which are functions on one variable less that are homogenous of the same degree, and the lower bound on the effective degree of F_1 still holds. If all of the mass lied only on these functions, we would immediately get that $\beta_{n,d_1,d_2,d_3} \leq \beta_{n-1,d_1,d_2,d_3}$. However in general there may be some small mass outside these. Intuitively, since this is a very small mass, it shouldn't matter that much; indeed our arguments show that something along these lines is true.

The second step: the case of nearly-linear degree. Thus, it remains to bound β_{n,d_1,d_2,d_3} when n, d_1, d_2, d_3 are all the same, up to poly-logs. In this case, when we take F, g, h that achieve β_{n,d_1,d_2,d_3} , it follows that each monomial in F has effective degree at least $\tilde{\Omega}(d_1)$, and from this it follows that there is a variable (actually, many of them) i such that coordinate i has significant effective influence on F (more precisely, of the order $\Omega\left(\frac{1}{\log^c d_1}\right)$). We show that this implies that when we write F according to its SVD decomposition

as in Claim 4.10 according to the partition $[n] = I \cup J$ where J has size 1 and it contains a variable with significant effective influence, we have

$$F(\mathbf{y}, \mathbf{z}) = \sum_t \gamma_t F_t(\mathbf{y}_I, \mathbf{z}_I) F_t'(\mathbf{y}_J, \mathbf{z}_J),$$

and we show that there is some t such that $\gamma_t F_t'$ has significant variance in $\Sigma' \subseteq \Sigma$ (here Σ' is chosen so as to satisfy the conditions of the relaxed base case). This means that when bounding the expectation of $F_t' g_r' h_s'$ over $\mathbf{y}_J, \mathbf{z}_J$, we may appeal to the relaxed base case to obtain a stronger bound than we originally had (the trivial bound is 1). Indeed, our argument proceeds in a similar way (though not exactly in this way due to technical reasons), and we (morally) prove that

$$\beta_{n,d_1,d_2,d_3} \leq \left(1 - \Omega\left(\frac{1}{\log^C d_1}\right)\right) \beta_{n-1,d_1-1,d_2-1,d_3-1}. \quad (9)$$

Iterating this argument for $d_1/2$ times, we get that

$$\beta_{n,d_1,d_2,d_3} \leq \left(1 - \Omega\left(\frac{1}{\log^C d_1}\right)\right)^{d_1/2} \beta_{n-d_1/2,d_1/2,d_2-d_1/2,d_3-d_1/2} \leq 2^{-\frac{d_1}{\log^C(d_1)}},$$

concluding the proof.

Our formal argument follows the same spirit as the moral argument above, with two distinctions. First, we do not know how to establish inequality (9). Our inductive argument necessitates appealing to an inductive assumption on functions over $n - 1$ variables that are non-homogenous functions (even if the original functions F , g and h were homogenous). Hence we cannot really work with the parameter β_{n,d_1,d_2,d_3} , and we define a similar parameter to it, $\gamma_{n,d}$, where the conditions that the function F is homogenous of degree d is replaced with the condition that $F^{\leq d} \equiv 0$ (i.e. that F only contains monomials of degree d and more), and the conditions about g and h are dropped altogether.¹² The upside of moving to the parameter $\gamma_{n,d}$ is that it facilitates an inductive argument as described above. The downside of moving to the parameter $\gamma_{n,d}$ is that we may no longer use SVD decompositions as given in Claims 4.8, 4.10 due to the lack of homogeneity. Nevertheless, one still has a similar SVD decomposition as given in Claims 4.11, 4.12, where the main difference is that we can only say that the degree of each F_t is at least $d - 1$. Thus, we lose 1 in the degree parameter for each iteration. Such tradeoff would not be beneficial in the context of the previous step, since there we do not necessarily manage to show that the parameter β_{n,d_1,d_2,d_3} actually decreases when we increase d_1 . In the context of this step however, we manage to gain a factor of $\left(1 - \Omega\left(\frac{1}{\log^C d_1}\right)\right)$ from each iteration, and since (as we show) we can perform $\Omega(d)$ such iterations, the move to $\gamma_{n,d}$ is affordable and leads to an exponentially decaying bound.

5 The first step in the proof of Lemma 4.7: reducing to $n \lesssim \max(d_1, d_2, d_3)$

In this section, we prove the following lemma. It states that either we reduce the number of variables by one without reducing any of the degrees and without gaining any factor, *or* we reduce at least one of the degrees by one along with the reduction in the number of variables by one and gain an additional multiplicative factor of $(1 - \varepsilon)$ for some constant $\varepsilon > 0$.

¹²It is also possible to replace the condition on g with a condition of the form $g^{\leq d_2} \equiv 0$, but this is not necessary for us.

Lemma 5.1. *For all $\alpha > 0$ and $m \in \mathbb{N}$, there exist $\varepsilon > 0$ and $L \in \mathbb{N}$ such that the following holds. Suppose $n, d_1, d_2, d_3 \in \mathbb{N}$ are parameters, and $n \geq L \cdot \max(d_1, d_2, d_3)$. Then letting*

$$\beta' = \max \left(\beta_{n-1, d_1-1, d_2, d_3}, \beta_{n-1, d_1, d_2-1, d_3}, \beta_{n-1, d_1, d_2, d_3-1}, \beta_{n-1, d_1-1, d_2-1, d_3}, \right. \\ \left. \beta_{n-1, d_1, d_2-1, d_3-1}, \beta_{n-1, d_1-1, d_2, d_3-1}, \beta_{n-1, d_1-1, d_2-1, d_3-1}, \beta_{n-1, d_1, d_2, d_3} \right)$$

We have that

$$\beta_{n, d_1, d_2, d_3} \leq \max(\beta_{n-1, d_1, d_2, d_3}, (1 - \varepsilon)\beta').$$

As can be observed easily, if we iterate the above lemma, then either we gain the $(1 - \varepsilon)$ factor enough number of times to get the conclusion in Lemma 4.7, or we arrive at a situation when the degrees are nearly-linear in the number of variables and thereby finishing the first step in the proof of Lemma 4.7. This is formally shown in Corollary 5.6 later.

The rest of this section is devoted to the proof of Lemma 5.1. The proof proceeds by an inductive argument over n . Below, we consider the functions g and h of 2-norm 1 that achieve that value β_{n, d_1, d_2, d_3} , and partition the set of coordinates $[n]$ into $I \cup J$ where $|I| = n - 1$ and $|J| = 1$.

Fix F, g, h achieving β_{n, d_1, d_2, d_3} .

5.1 Warm-up

To motivate the argument, we begin with considering a simplistic case in which F can be written as $F'(\mathbf{y}_I, \mathbf{z}_I)F''(\mathbf{y}_J, \mathbf{z}_J)$, g can be written as $g'(\mathbf{y}_I)g''(\mathbf{y}_J)$ and h could be written as well as $h'(\mathbf{z}_I)h''(\mathbf{z}_J)$ (Note that using the SVD decompositions, F, g and h may be written as sum of such terms). The inductive step would be very easy. Indeed, we then have that

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [F(\mathbf{y}, \mathbf{z})g(\mathbf{y})h(\mathbf{z})] = \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^I} [F'(\mathbf{y}, \mathbf{z})g'(\mathbf{y})h'(\mathbf{z})] \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^J} [F''(\mathbf{y}, \mathbf{z})g''(\mathbf{y})h''(\mathbf{z})].$$

It is always clear, by Cauchy-Schwarz, that

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^J} [F''(\mathbf{y}, \mathbf{z})g''(\mathbf{y})h''(\mathbf{z})] \leq \|F''\|_2 \|g''\|_2 \|h''\|_2.$$

1. If either g'' or h'' are constant, we show using our additive base cases, e.g. Lemma 2.17, that this bound may be improved to $(1 - \varepsilon)\|F''\|_2 \|g''\|_2 \|h''\|_2$; combining this with

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^I} [F'(\mathbf{y}, \mathbf{z})g'(\mathbf{y})h'(\mathbf{z})] \leq \beta'$$

that follows by the inductive step, gives that $\beta_{n, d_1, d_2, d_3} \leq (1 - \varepsilon)\beta'$.

2. If neither g'' nor h'' are constant, it follows that if F'' is constant then

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^J} [F''(\mathbf{y}, \mathbf{z})g''(\mathbf{y})h''(\mathbf{z})] = 0,$$

using the fact that (y, z) are independent and there is nothing to prove.

Thus, the only case left to consider is the case that F'' , g'' and h'' are all non-constant. We do not know how to effectively handle this case (since the degrees of F' , g' and h' have reduced but we do not know how to gain a $(1 - \varepsilon)$ factor), and hence we will try to avoid having to give effective bounds on such terms. Indeed, such terms would have small weight (this is the way we will choose the partition $I \cup J$ and will hence be negligible when compared to terms from the item (1) above.

5.2 The actual proof

We proceed with the general case in which F , g and h need not take this special form as above. Using the singular-value decomposition we may write F , g and h as sum of at most m such functions satisfying some orthogonality properties (see Claim 4.8 and Claim 4.10 for precise statements):

$$F(\mathbf{y}, \mathbf{z}) = \sum_{t \in T} \gamma_t F_t(\mathbf{y}_I, \mathbf{z}_I) F_t'(\mathbf{y}_J, \mathbf{z}_J), \quad g(\mathbf{y}) = \sum_{r \in R} \lambda_r g_r(\mathbf{y}_I) g_r'(\mathbf{y}_J), \quad h(\mathbf{z}) = \sum_{s \in S} \mu_s h_s(\mathbf{y}_I) h_s'(\mathbf{y}_J), \quad (10)$$

where each one of the sets $\{F_t\}_{t \in T}$, $\{F_t'\}_{t \in T}$, $\{g_r\}_r$, $\{g_r'\}_r$, $\{h_s\}_s$, $\{h_s'\}_s$ is orthonormal, and $F_1', g_1', h_1' \equiv 1$, and $\sum_t \gamma_t^2 = \sum_r \lambda_r^2 = \sum_s \mu_s^2 = 1$. Thus, we have that

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [F(\mathbf{y}, \mathbf{z}) g(\mathbf{y}) h(\mathbf{z})] = \sum_{r, s, t} \gamma_t \lambda_r \mu_s \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [F_t(\mathbf{y}, \mathbf{z}) g_r(\mathbf{y}) h_s(\mathbf{z})] \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [F_t'(\mathbf{y}, \mathbf{z}) g_r'(\mathbf{y}) h_s'(\mathbf{z})],$$

which appears as a sort of weighted sum over the cases considered in the above simplistic case. Indeed, we identify certain parts of this sum which are 0, certain parts which are negligible, and use additive base cases to bound the rest; the main point of the argument is to have the gain from the additive base case overcomes the error terms, and we manage to achieve that.

It will be convenient for us to denote $\widehat{F}_t(r, s) = \langle F_t, g_r h_s \rangle$ as well as $\widehat{F}_t'(r, s) = \langle F_t', g_r' h_s' \rangle$. This is justified because $(g_r h_s)_{r \in R, s \in S}$ form an orthonormal set in $L_2(y_I, z_I; \mu_{y, z}^{\otimes I})$, and it can be completed to an orthonormal basis, in which case the coefficient $\widehat{F}_t(r, s)$ appear in front of $g_r h_s$ in the representation of F_t . In particular, by Parseval we get that

$$1 = \|F_t\|_2^2 \geq \sum_{r \in R, s \in S} \widehat{F}_t(r, s)^2.$$

Similar reasoning applies to the notation $\widehat{F}_t'(r, s)$. Thus, we get that

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [F(\mathbf{y}, \mathbf{z}) g(\mathbf{y}) h(\mathbf{z})] = \sum_{r, s, t} \gamma_t \lambda_r \mu_s \widehat{F}_t(r, s) \widehat{F}_t'(r, s), \quad (11)$$

Throughout the proof, we will assume that

$$\beta_{n, d_1, d_2, d_3} \geq \max(\beta_{n-1, d_1, d_2, d_3}, (1 - \varepsilon)\beta').$$

since otherwise we are done. We will also assume that $\gamma_1, \lambda_1, \mu_1 \geq 1 - \eta$, and next show that as long as n is much larger than $\max(d_1, d_2, d_3)$, we may choose a partition $[n] = I \cup J$ so that this occurs.

5.2.1 The parameters and choosing the partition

We will use several parameters throughout this section, obeying the following relations:

$$0 \ll \eta \ll c \ll m^{-1}, \alpha \leq 1. \quad (12)$$

We show that if $n \geq \frac{10}{\eta} \max(d_1, d_2, d_3)$, then we may find a partition $[n] = I \cup J$ with $|J| = 1$ so that in the SVD decomposition in (10), we have $\gamma_1, \lambda_1, \mu_1 \geq 1 - \eta$. To see that, choose the partition randomly, and note that by Claim 4.9 we have that $\mathbb{E}_{I, J} [1 - \gamma_1^2] = \frac{1}{2} \mathbb{E}_{J=\{j\}} [I_j[f]] \leq \frac{d_1}{n} \leq \frac{\eta}{10}$ and similarly

$$\mathbb{E}_{I, J} [1 - \lambda_1^2] = \frac{1}{2} \mathbb{E}_{J=\{j\}} [I_j[g]] \leq \frac{d_2}{n} \leq \frac{\eta}{10}, \quad \mathbb{E}_{I, J} [1 - \mu_1^2] = \frac{1}{2} \mathbb{E}_{J=\{j\}} [I_j[h]] \leq \frac{d_3}{n} \leq \frac{\eta}{10}.$$

By Markov's inequality, we get that

$$\Pr_{I,J} [1 - \gamma_1^2 \geq \eta \vee 1 - \lambda_1^2 \geq \eta \vee 1 - \mu_1^2 \geq \eta] \leq \frac{3}{10} < 1,$$

so we may find a partition $[n] = I \cup J$ with $|J|$ such that $\gamma_1^2, \lambda_1^2, \mu_1^2 \geq 1 - \eta$. We fix this partition henceforth.

5.2.2 The main inductive argument

We wish to use the fact that $\gamma_1^2, \lambda_1^2, \mu_1^2 \geq 1 - \eta$ in order to effectively give an upper bound on β_{n,d_1,d_2,d_3} . Towards this, we split the expression from (11) as follows.

$$\beta_{n,d_1,d_2,d_3} = \sum_t \gamma_t A_t, \quad (13)$$

where A_t is defined as

$$A_t = \lambda_1 \mu_1 \widehat{F}_t(1, 1) \widehat{F}'_t(1, 1) + \sum_{r \neq 1} \lambda_r \mu_1 \widehat{F}_t(r, 1) \widehat{F}'_t(r, 1) + \sum_{s \neq 1} \lambda_1 \mu_s \widehat{F}_t(1, s) \widehat{F}'_t(1, s) + B_t,$$

$$B_t = \sum_{r,s \neq 1} \lambda_r \mu_s \widehat{F}_t(r, s) \widehat{F}'_t(r, s).$$

To simplify notation, we will omit the $r \neq 1$ and $s \neq 1$ from the sums below, as whenever we sum over r or s in this subsection, the summation does not include 1. We make a couple of preliminary observations.

1. Note that for $t \neq 1$, as F'_t is orthogonal to $F'_1 \equiv 1$, we get that its average is 0 and so $\widehat{F}'_t(1, 1) = 0$ for $t \neq 1$.
2. As F'_1 is constant, and y, z are independent, we get that $\widehat{F}'_1(r, s) = 0$ for all $(r, s) \neq (1, 1)$, and $\widehat{F}'_1(1, 1) = 1$.

The next two claims give bounds on A_i s and B_i s. This following claim will let us ignore the contribution from B_t in (13).

Claim 5.2. (Bounding B_i s) *We have:*

1. $B_1 = 0$.
2. For $t \neq 1$, $|B_t| \leq \sqrt{(1 - \mu_1^2)(1 - \lambda_1^2)} \beta'$.

The next claim handles A_i s. In this claim, we use the additive base case inequality, Lemma 2.17, in order to gain enough so that the contributions from B_t in (13) becomes negligible.

Claim 5.3. (Bounding A_i s) *There exists a constant $c = c(\mu) > 0$ such that the following holds.*

1. $|A_1| \leq \lambda_1 \mu_1 \beta_{n-1,d_1,d_2,d_3}$
2. For all $t \neq 1$

$$|A_t| \leq \sqrt{(1 - c) \mu_1^2 \sum_r \lambda_r^2 \widehat{F}_t(r, 1)^2 + (1 - c) \lambda_1^2 \sum_s \mu_s^2 \widehat{F}_t(1, s)^2} + E_t,$$

where $E_t \lesssim_m \sqrt{|1 - \mu_1| |1 - \lambda_1|} \beta'$.

Before we see the proofs of the above two claims, let us see why these are enough to prove Lemma 5.1.

Proof of Lemma 5.1 Let $T_1 = \{1\}$, $T_2 = T \setminus T_1$. Then

$$\sum_t \gamma_t A_t = \underbrace{\sum_{t \in T_1} \gamma_t A_t}_{(I)} + \underbrace{\sum_{t \in T_2} \gamma_t A_t}_{(II)}.$$

By Claim 5.3 we have that $|(I)| \leq \gamma_1 \lambda_1 \mu_1 \beta_{n-1, d_1, d_2, d_3}$. For (II), we use Claim 5.3 and bound it as:

$$\sum_{t \in T_2} \gamma_t \sqrt{(1-c) \left(\mu_1^2 \sum_r \lambda_r^2 \widehat{F}_t(r, 1)^2 + \lambda_1^2 \sum_s \mu_s^2 \widehat{F}_t(1, s)^2 \right)} + \gamma_t E_t$$

where $E_t \lesssim_m \sqrt{(1-\lambda_1^2)(1-\mu_1^2)} \beta'$. Thus, by Cauchy–Schwarz, (II) is at most

$$\begin{aligned} & \sqrt{\sum_{t \in T_2} \gamma_t^2} \sqrt{(1-c) \left(\mu_1^2 \sum_r \lambda_r^2 \sum_{t \in T_2} \widehat{F}_t(r, 1)^2 + \lambda_1^2 \sum_s \mu_s^2 \sum_{t \in T_2} \widehat{F}_t(1, s)^2 \right)} \\ & + O_m \left(\sqrt{(1-\gamma_1^2)(1-\lambda_1^2)(1-\mu_1^2)} \beta' \right). \end{aligned}$$

Note that for all r ,

$$\sum_{t \in T_2} \widehat{F}_t(r, 1)^2 = \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^I} \left[\tilde{F} g_r h_1 \right]^2,$$

where $\tilde{F} = \frac{\sum_{t \in T_2} \widehat{F}_t(r, 1) F_t}{\sqrt{\sum_{t \in T_2} \widehat{F}_t(r, 1)^2}}$, and so by definition $\sum_{t \in T_2} \widehat{F}_t(r, 1)^2 \leq \beta_{n-1, d_1-1, d_2-1, d_3}^2 \leq \beta'^2$. Similarly, $\sum_{t \in T_2} \widehat{F}_t(1, s)^2 \leq \beta'^2$, so we get that

$$(II) \leq \sqrt{1-\gamma_1^2} \sqrt{(1-c) \beta'^2 \left(\mu_1^2 \sum_r \lambda_r^2 + \lambda_1^2 \sum_s \mu_s^2 \right)} + O_m \left(\sqrt{(1-\gamma_1^2)(1-\lambda_1^2)(1-\mu_1^2)} \beta' \right).$$

so

$$(II) \leq \sqrt{1-\gamma_1^2} \sqrt{(1-c) \beta'^2 (\mu_1^2 (1-\lambda_1^2) + \lambda_1^2 (1-\mu_1^2))} + O_m \left(\sqrt{(1-\gamma_1^2)(1-\lambda_1^2)(1-\mu_1^2)} \beta' \right).$$

Set $\beta'' = \max(\beta_{n-1, d_1, d_2, d_3}, (1-c/8) \beta')$. Combining the bounds on (I) and (II) and using Cauchy–Schwarz, we get

$$\begin{aligned} (I) + (II) & \leq \sqrt{1-\gamma_1^2 + \gamma_1^2} \sqrt{\lambda_1^2 \mu_1^2 \beta_{n-1, d_1, d_2, d_3}^2 + (1-\frac{c}{2}) \beta'^2 (\mu_1^2 (1-\lambda_1^2) + \lambda_1^2 (1-\mu_1^2))} \\ & + O_m \left(\sqrt{(1-\gamma_1^2)(1-\lambda_1^2)(1-\mu_1^2)} \beta' \right) \\ & \leq \beta'' \sqrt{1-\frac{c}{4} (\mu_1^2 (1-\lambda_1^2) + \lambda_1^2 (1-\mu_1^2))} + O_m \left(\sqrt{(1-\gamma_1^2)(1-\lambda_1^2)(1-\mu_1^2)} \beta'' \right) \\ & \leq \beta'' \left(1 - \frac{c}{8} (\mu_1^2 (1-\lambda_1^2) + \lambda_1^2 (1-\mu_1^2)) \right) + O_m \left(\sqrt{(1-\gamma_1^2)(1-\lambda_1^2)(1-\mu_1^2)} \beta'' \right) \\ & \leq \beta'', \end{aligned}$$

where the last inequality is because that

$$\sqrt{(1 - \gamma_1^2)(1 - \lambda_1^2)(1 - \mu_1^2)} \leq \sqrt{\eta} \left((1 - \lambda_1^2) + (1 - \mu_1^2) \right) \lesssim \sqrt{\eta} \left((1 - \lambda_1^2)\mu_1^2 + (1 - \mu_1^2)\lambda_1^2 \right),$$

and $\eta \ll c$. Plugging this into (13) finishes the proof of Lemma 5.1. \square

Bounding A_i s and B_i s. We now prove the bounds on A_i s and B_i s.

Claim 5.4 (Restatement of Claim 5.2, Bounding B_i s). *We have:*

1. $B_1 = 0$.
2. For $t \neq 1$, $|B_t| \leq \sqrt{(1 - \mu_1^2)(1 - \lambda_1^2)}\beta'$.

Proof. The first item follows since $\widehat{F}'_1(r, s) = 0$ for all $(r, s) \neq (1, 1)$. For the second item, we have

$$\begin{aligned} |B_t| &\leq \sum_{r,s} \lambda_r \mu_s \left| \widehat{F}_t(r, s) \right| \left| \widehat{F}'_t(r, s) \right| \leq \sqrt{\sum_{r,s} \lambda_r^2 \mu_s^2 \left| \widehat{F}_t(r, s) \right|^2} \sqrt{\sum_{r,s} \left| \widehat{F}'_t(r, s) \right|^2} \\ &\leq \sqrt{\sum_{r,s} \lambda_r^2 \mu_s^2 \beta'^2 \sqrt{1}} \\ &\leq \sqrt{(1 - \mu_1^2)(1 - \lambda_1^2)}\beta'. \end{aligned}$$

The first transition is by Cuachy-Schwarz. The second transition is because

$$\widehat{F}_t(r, s) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes I}} [F_t(\mathbf{y}, \mathbf{z}) g_r(\mathbf{y}) h_s(\mathbf{z})],$$

and as F_t, g_r, h_s are all homogenous of degrees $d_1 - 1, d_2 - 1, d_3 - 1$ we get by definition that

$$\left| \widehat{F}_t(r, s) \right| \leq \beta_{n-1, d_1-1, d_2-1, d_3-1},$$

which is at most β' . \square

Claim 5.5. (Restatement of Claim 5.3, Bounding A_i s) *There exists a constant $c = c(\mu) > 0$ such that the following holds.*

1. $|A_1| \leq \lambda_1 \mu_1 \beta_{n-1, d_1, d_2, d_3}$
2. For all $t \neq 1$

$$|A_t| \leq \sqrt{(1 - c)\mu_1^2 \sum_r \lambda_r^2 \widehat{F}_t(r, 1)^2 + (1 - c)\lambda_1^2 \sum_s \mu_s^2 \widehat{F}_t(1, s)^2} + E_t,$$

where $E_t \lesssim_m \sqrt{|1 - \mu_1| |1 - \lambda_1|} \beta'$.

Proof. Note that $A_1 = \lambda_1 \mu_1 \widehat{F}_1(1, 1)$. The result follows since by definition we have that

$$\left| \widehat{F}_1(1, 1) \right| = \left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^t} [F_1(\mathbf{y}, \mathbf{z}) g_1(\mathbf{y}) h_1(\mathbf{z})] \right| \leq \beta_{n-1, d_1, d_2, d_3},$$

as F_1 is homogenous, constant on connected component, has the same degree and effective degree as F (see Claim 4.10), and g_1, h_1 are homogenous of degrees d_2, d_3 respectively.

For the second item, consider

$$\tilde{g}(y) = \sum_r \lambda_r \widehat{F}_t(r, 1) g'_r, \quad \tilde{h}(y) = \sum_s \mu_s \widehat{F}_t(1, s) h'_s.$$

First, note that

$$\langle \tilde{g} + \tilde{h}, F'_t \rangle = \sum_r \lambda_r \widehat{F}_t(r, 1) \widehat{F}'_t(r, 1) + \sum_s \mu_s \widehat{F}_t(1, s) \widehat{F}'_t(1, s).$$

Therefore,

$$\begin{aligned} \langle \tilde{g} + \tilde{h}, F'_t \rangle &= \sum_r \mu_1 \lambda_r \widehat{F}_t(r, 1) \widehat{F}'_t(r, 1) + \sum_s \lambda_1 \mu_s \widehat{F}_t(1, s) \widehat{F}'_t(1, s) \\ &\quad + O\left(\left((1 - \lambda_1) \sqrt{1 - \mu_1} + (1 - \mu_1) \sqrt{1 - \lambda_1}\right) \beta'\right). \end{aligned} \quad (14)$$

In the last transition, we used the fact that

$$\left| \sum_r (1 - \mu_1) \lambda_r \widehat{F}_t(r, 1) \widehat{F}'_t(r, 1) \right| \leq (1 - \mu_1) \sqrt{\sum_r \lambda_r^2} \sqrt{\sum_r \widehat{F}'_t(r, 1)^2} \max_r \left| \widehat{F}_t(r, 1) \right|,$$

and $\sum_r \lambda_r^2 = 1 - \lambda_1^2 \lesssim 1 - \lambda_1$, $\max_r \left| \widehat{F}_t(r, 1) \right| \leq \beta'$ and $\sum_r \widehat{F}'_t(r, 1)^2 \leq \|F'_t\|_2^2 = 1$. Similarly, we also have that

$$\left| \sum_s (1 - \lambda_1) \mu_s \widehat{F}_t(1, s) \widehat{F}'_t(1, s) \right| \lesssim (1 - \lambda_1) \sqrt{1 - \mu_1} \beta'.$$

Combining (14) with Claim 5.2 gives $|A_t| = \left| \langle \tilde{g} + \tilde{h}, F'_t \rangle \right| + O\left(\sqrt{(1 - \lambda_1)(1 - \mu_1)} \beta'\right)$.

By Lemma 2.17,

$$\left| \langle \tilde{g} + \tilde{h}, F'_t \rangle \right| \leq \sqrt{(1 - c) \sum_r \lambda_r^2 \widehat{F}_t(r, 1)^2 + (1 - c) \sum_s \mu_s^2 \widehat{F}_t(1, s)^2},$$

implying that

$$\left| \langle \tilde{g} + \tilde{h}, F'_t \rangle \right| \leq \sqrt{(1 - c) \mu_1^2 \sum_r \lambda_r^2 \widehat{F}_t(r, 1)^2 + (1 - c) \lambda_1^2 \sum_s \mu_s^2 \widehat{F}_t(1, s)^2 + E},$$

where $E \lesssim |1 - \mu_1^2| |1 - \lambda_1^2| \beta'^2$. □

5.3 Concluding this section

To conclude this section, we use Lemma 5.1 in order to show that for given parameters n, d_1, d_2, d_3 , one either has the conclusion of the lemma or else can assume that n is at most $O(\max(d_1, d_2, d_3))$. Formally:

Corollary 5.6. *For all $\alpha > 0$ and $m \in \mathbb{N}$, there are $C > 0$ and $\varepsilon' > 0$ such that the following holds. For parameters $n, d_1, d_2, d_3 \in \mathbb{N}$, we either have that*

1. $\beta_{n,d_1,d_2,d_3} \leq (1 - \varepsilon')^{\min(d_1,d_2,d_3)}$, or else
2. $\beta_{n,d_1,d_2,d_3} \leq \beta_{n',d'_1,d'_2,d'_3}$ where $n' \leq C \min(d'_1, d'_2, d'_3)$ and $d'_1 \geq d_1/2$, $d'_2 \geq d_2/2$ and $d'_3 \geq d_3/2$.

Proof. We iterate Lemma 5.1, and divide iterations into two types: those in which we gain a factor of $(1 - \varepsilon)$, and those that we do not (in which case we leave d_1, d_2, d_3 as is and decrease n). Eventually, we stop at n', d'_1, d'_2, d'_3 where $n' \geq C \max(d'_1, d'_2, d'_3)$. Hence, there must have been at least $\max(d_1 - d'_1, d_2 - d'_2, d_3 - d'_3)$ iterations in which we gained a factor of $1 - \varepsilon$. Thus, if $d'_1 \leq d_1/2$, $d'_2 \leq d_2/2$ or $d'_3 \leq d_3/3$, then we gained $(1 - \varepsilon)^{\min(d_1/2, d_2/2, d_3/2)} \leq (1 - \varepsilon')^{\min(d_1, d_2, d_3)}$, and we are done. \square

Looking at Lemma 4.7, we see that in the first case of Corollary 5.6 we get that

$$\beta_{n,d_1,d_2,d_3} \leq (1 - \varepsilon)^{\min(d_1,d_2,d_3)} \leq (1 - \varepsilon)^{d_1/\log^K(d_1)} \leq 2^{-d_1/\log^{C'}(d_1)},$$

and the conclusion of the Lemma 4.7 holds. Thus, we turn our attention to handling the second case of Corollary 5.6, in which case we get that $\beta_{n,d_1,d_2,d_3} \leq \beta_{n',d'_1,d'_2,d'_3}$ where $d'_i \leq d'_j \log^K(d'_j)$ for all i, j and $n' \leq C \max(d'_1, d'_2, d'_3)$.

6 The second step in the proof of Lemma 4.7: the near-linear degree case

In our argument so far, we have not used the relaxed based case inequality. In this section, we crucially use the inequality in the second step of the inductive proof.

Now that we are in a case where the number of variable is of the same magnitude as the degrees of the functions, we can work with a simpler parameter that is defined next.

6.1 The parameter $\gamma_{n,d}$

Define the parameter

$$\gamma_{n,d} = \max_{\substack{F: \Gamma^n \times \Phi^n \rightarrow \mathbb{R}, F < d \equiv 0 \\ F \text{ is constant on connected components,} \\ \text{eff-deg}(F) \geq d/\log^{20}(d) \\ g: \Gamma^n \rightarrow \mathbb{R} \\ h: \Phi^n \rightarrow \mathbb{R}}} \frac{|\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [F(\mathbf{y}, \mathbf{z})g(\mathbf{y})h(\mathbf{z})]|}{\|F\|_2 \|g\|_2 \|h\|_2}.$$

Clearly $\beta_{n',d'_1,d'_2,d'_3}$ from the previous section is at most γ_{n',d'_1} . Our goal in this section is to prove the following result.

Lemma 6.1. *For all $\alpha > 0$, $m \in \mathbb{N}$ and $C > 0$ there exists $R > 0$ and d_0 such that if $n, d \in \mathbb{N}$ are parameters such that $n \leq d \log^C d$ and $d \geq d_0$, then*

$$\gamma_{n,d} \leq \left(1 - \frac{1}{\log^R(d)}\right) \gamma_{n-1,d-1}.$$

Before proving Lemma 6.1, we show the quick derivation of Lemma 4.7 from it (using the previous section).

Proof of Lemma 4.7. From Corollary 5.6 (see the discussion after it), we are either done by Corollary 5.6 or else we find d'_1, d'_2, d'_3 and n' as there such that $\beta_{n,d_1,d_2,d_3} \leq \beta_{n',d'_1,d'_2,d'_3} \leq \gamma_{n',d'_1}$. Note that $n' \lesssim \max(d'_1, d'_2, d'_3)$, so $n' \leq \frac{d}{2} \log^{C'}(\frac{d}{2})$, for $d = d'_1$. We now apply Lemma 6.1 $d/2$ times and get that

$$\gamma_{n',d} \leq \left(1 - \frac{1}{\log^R(d)}\right) \gamma_{n'-1,d-1} \leq \dots \leq \left(1 - \frac{1}{\log^R(d)}\right)^{d/2} \gamma_{n'-d/2,d/2} \leq \left(1 - \frac{1}{\log^R(d)}\right)^{d/2},$$

which is at most $2^{-d/\log^R(d)}$. \square

The rest of this section is devoted to the proof of Lemma 6.1. We assume henceforth that

$$\gamma_{n,d} \geq \frac{\gamma_{n,d-1}}{2},$$

since otherwise we are done.

6.2 Set up for the proof of Lemma 6.1

In this section, we present the set up for the proof of Lemma 6.1. We fix n and d as there. We will use the following hierarchy of parameters:

$$\begin{aligned} 0 \ll R_4^{-1} \ll R_3^{-1} \ll R_2^{-1} \ll R_1^{-1} \ll R_0^{-1} \ll R^{-1} \ll c \ll m^{-1}, \alpha, C^{-1} \leq 1, \\ 0 < \delta = \frac{1}{\log^{R_4} d} \leq \tau = \frac{1}{\log^{R_3} d} \leq \zeta = \frac{1}{\log^{R_2} d} \leq v = \frac{1}{\log^{R_1} d} \leq u = \frac{1}{\log^{R_0} d} \leq w = \frac{1}{\log^R d} \leq 1. \end{aligned} \quad (15)$$

By definition, we may find $F: \Sigma^n \rightarrow \mathbb{R}$ which is constant on connected components, has 2-norm 1, and satisfies that $F^{<d} \equiv 0$ and all monomials in F have effective degree at least $d' = d/\log^{20} d$, as well as $g: \Gamma^n \rightarrow \mathbb{R}$, $h: \Phi^n \rightarrow \mathbb{R}$ of 2-norm equal to 1 such that

$$\gamma_{n,d} = \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^n} [F(\mathbf{y}, \mathbf{z})g(\mathbf{y})h(\mathbf{z})].$$

We employ an SVD decomposition according to Claims 4.11 and 4.12; if we apply it using a partition $[n] = I \cup J$ where $|J| = 1$, we get

$$F(\mathbf{y}, \mathbf{z}) = \sum_{t \in T} \gamma_t F_t(\mathbf{y}_I, \mathbf{z}_I) F'_t(\mathbf{y}_J, \mathbf{z}_J), \quad g(\mathbf{y}) = \sum_{r \in R} \lambda_r g_r(\mathbf{y}_I) g'_r(\mathbf{y}_J), \quad h(\mathbf{z}) = \sum_{s \in S} \mu_s h_s(\mathbf{y}_I) h'_s(\mathbf{y}_J), \quad (16)$$

where each one of the sets $\{F_t\}_{t \in T}$, $\{F'_t\}_{t \in T}$, $\{g_r\}_r$, $\{g'_r\}_r$, $\{h_s\}_s$, $\{h'_s\}_s$ is orthonormal, and for all $t \in T$ the function F_t has $(F_t)^{<d-1} \equiv 0$ and all monomials in it have effective degree at least $d' - 1$, and $\sum_t \gamma_t^2 = \sum_r \lambda_r^2 = \sum_s \mu_s^2 = 1$. Then we have (following the Fourier coefficients notation from the last section) that

$$\gamma_{n,d} = \sum_{r,s,t} \gamma_t \lambda_r \mu_s \widehat{F}_t(r, s) \widehat{F}'_t(r, s). \quad (17)$$

We will want our partition I, J to satisfy a certain property that will be helpful for us. Namely, if $J = \{j\}$, we will want the variable j to have significant *effective influence* on F . Formally, we have defined effective influences for functions over \mathbf{x} , but as F is constant on connected components we can view it, using Claim 4.5, as $F = Wf$ for some $f: \Sigma^n \rightarrow \mathbb{R}$ of the same 2-norm, effective degrees etc. We define $I_{j,\text{effective}}[F] = I_{j,\text{effective}}[f]$. Note that

$$\frac{1}{n} \sum_{j=1}^n I_{j,\text{effective}}[F] = \frac{1}{n} I_{\text{effective}}[f] = \frac{2}{n} \sum_{\chi} \text{eff-deg}(\chi) \widehat{f}(\chi)^2 \geq \frac{2d'}{n} \|f\|_2^2 = \frac{2d'}{n}$$

where we used Fact 2.13. Since $n \leq d \log^C d$ and $d \leq d' \log^{20} d$ it follows that $\frac{1}{n} \sum_{j=1}^n I_{j,\text{effective}}[F] \geq \frac{1}{\log^{20+C} d}$, and therefore we may find a variable j such that $I_{j,\text{effective}}[F] \geq \frac{1}{\log^{20+C} d}$. We denote $v = \frac{1}{\log^{20+C} d}$, and choose the partition $J = \{j\}$ and $I = [n] \setminus J$.

Since F'_t are all constant on connected components, there are unique $f'_t: \Sigma \rightarrow \mathbb{R}$ of 2-norm 1 such that $F'_t = Wf'_t$. We fix such f'_t ; similarly, we have $F_t = Wf_t$, hence

$$F = W \sum_t \gamma_t f_t f'_t.$$

Note that

$$\langle f_{t_1}, f_{t_2} \rangle_{\mu_x} = \langle Wf_{t_1}, Wf_{t_2} \rangle_{\mu_{y,z}} = \langle F_{t_1}, F_{t_2} \rangle_{\mu_{y,z}} = \mathbf{1}_{t_1=t_2},$$

and similarly for f'_{t_1} and f'_{t_2} .

Next, we note that since F has effective influence at least v , the variance of f'_t over Σ' (this is $\Sigma' \subseteq \Sigma$ for which the relaxed base case holds) is significant. For notational convenience, we denote

$$\text{var}_{\Sigma'}(f'_t) = \mathbb{E}_{x,x' \in \Sigma'} [(f'_t(x) - f'_t(x'))^2].$$

Claim 6.2. $\sum_t \gamma_t^2 \text{var}_{\Sigma'}(f'_t) \geq w$.

Proof. Consider $I_{j,\text{effective}}[F]$. On the one hand, it is at least v by choice. On the other hand, consider the distribution over (a, b) where $a \sim \mu_x$, and $b = a$ if $a \in \Sigma \setminus \Sigma'$ and otherwise $b \sim \mu_x \mid b \in \Sigma'$. Then $I_{j,\text{effective}}[F]$ is equal to

$$\begin{aligned} I_{j,\text{effective}}[f] &= \mathbb{E}_{\mathbf{x} \sim \mu_x^{\otimes n-1}, a, b} [(f(\mathbf{x}_I, a) - f(\mathbf{x}_I, b))^2] = \mathbb{E}_{\mathbf{x} \sim \mu_x^{\otimes n-1}, a, b} \left[\left(\sum_t \gamma_t f_t(\mathbf{x}_I) (f'_t(a) - f'_t(b)) \right)^2 \right] \\ &= \sum_{t_1, t_2} \gamma_{t_1} \gamma_{t_2} \langle f_{t_1}, f_{t_2} \rangle \mathbb{E}_{a, b} [(f'_{t_1}(a) - f'_{t_1}(b))(f'_{t_2}(a) - f'_{t_2}(b))]. \end{aligned}$$

For $t_1 \neq t_2$ we have $\langle f_{t_1}, f_{t_2} \rangle = 0$, so the last sum is equal to

$$\sum_t \gamma_t^2 \mathbb{E}_{a, b} [(f'_t(a) - f'_t(b))^2] \lesssim_{\alpha, m} \sum_t \gamma_t^2 \text{var}_{\Sigma'}(f'_t). \quad \square$$

Getting a gap between values close to 0 and bounded away from 0. Consider the values of λ_r, μ_s and γ_t . We claim that there are $0 \ll \delta \ll \zeta$ (as in (15)) such that none of these values fall in the interval $[\delta, \zeta)$. Indeed, start from sufficiently small $\zeta_1 \ll \tau$ and take $\zeta_2 \ll \zeta_1$, if one of these values falls in the interval $[\zeta_2, \zeta_1)$, then we proceed with taking $\zeta_3 \ll \zeta_2$ and considering the interval $[\zeta_3, \zeta_2)$, and continue with the choice iteratively. Since these intervals are disjoint, and there are at most $3m$ distinct values of $\lambda_r, \mu_s, \gamma_t$, we will reach an interval not containing any of them after at most $2m + m^2 + 1$ steps as required.

In other words, for all r , we either have $\lambda_r \geq \zeta$ or $\lambda_r \leq \delta$ and similarly for all s . Define

$$R' = \{r \in R \mid \lambda_r \geq \zeta\}, \quad S' = \{s \in S \mid \mu_s \geq \zeta\}, \quad T' = \{t \in T \mid \gamma_t \geq \zeta\}.$$

Intuitively, one should think of r outside R' as having its associated masses λ_r as being 0; we cannot quite say that, but for the argument to go through it suffices to have a sufficiently large gap between ζ and δ . The same goes for μ_s, γ_t .

Getting rid of the error term. We would like to replace the right hand side of (17) by restricting the sum to go only over R', S', T' , and argue that this only incurs a small loss. Consider for example the sum where $r \in R \setminus R'$ and $s \in S, t \in T$. Then we can bound

$$\left| \sum_{R \in R \setminus R'} \lambda_r \mu_s \gamma_t \widehat{F}_t(r, s) \widehat{F}'_t(r, s) \right| \leq \delta \max_{r, s, t} |\widehat{F}_t(r, s)| \lesssim_m \delta \gamma_{n-1, d-1}.$$

Similarly, the other sums are also bounded by $\lesssim_m \delta \gamma_{n-1, d-1}$, hence we get that (17) implies that

$$\gamma_{n, d} \leq \sum_{r \in R', s \in S', t \in T'} \lambda_r \mu_s \gamma_t \widehat{F}_t(r, s) \widehat{F}'_t(r, s) + O_m(\delta \gamma_{n, d}). \quad (18)$$

6.3 Proof of Lemma 6.1

We begin by observing that $\gamma_{n, d} \leq \gamma_{n-1, d-1} + O_m(\delta \gamma_{n, d})$. Indeed, by Cauchy–Schwarz

$$\begin{aligned} \sum_{r \in R', s \in S', t \in T'} \lambda_r \mu_s \gamma_t \widehat{F}_t(r, s) \widehat{F}'_t(r, s) &\leq \sqrt{\sum_{r \in R', s \in S', t \in T'} \lambda_r^2 \mu_s^2 \widehat{F}_t(r, s)^2} \sqrt{\sum_{r \in R', s \in S', t \in T'} \gamma_t^2 \widehat{F}'_t(r, s)^2} \\ &\leq \sqrt{\sum_{r \in R', s \in S'} \lambda_r^2 \mu_s^2 \sum_{t \in T'} \widehat{F}_t(r, s)^2} \sqrt{\sum_{t \in T'} \gamma_t^2 \sum_{r \in R', s \in S'} \widehat{F}'_t(r, s)^2} \\ &\leq \sqrt{\sum_{r \in R', s \in S'} \lambda_r^2 \mu_s^2 \gamma_{n-1, d-1}^2} \sqrt{\sum_{t \in T'} \gamma_t^2} \\ &\leq \gamma_{n-1, d-1}. \end{aligned} \quad (19)$$

In a high level, our proof inspects the potential equality cases in the above chain of inequalities, and demonstrates it cannot happen. Indeed, we have three inequalities that potentially could be improved. The first one is the Cauchy–Schwarz inequality, and if we knew that an improved version of it held, i.e. that

$$\sum_{r, s, t} \lambda_r \mu_s \cdot \widehat{F}_t(r, s) \gamma_t \widehat{F}'_t(r, s) \leq (1 - \tau) \sqrt{\sum_{r, s, t} \lambda_r^2 \mu_s^2 \widehat{F}_t(r, s)^2} \sqrt{\sum_{r, s, t} \gamma_t^2 \widehat{F}'_t(r, s)^2},$$

we would be done. Indeed, if so we would then get that

$$\sum_{r \in R', s \in S', t \in T'} \lambda_r \mu_s \gamma_t \widehat{F}_t(r, s) \widehat{F}'_t(r, s) \leq (1 - \tau) \gamma_{n-1, d-1},$$

so by (18) we get that

$$\gamma_{n, d} \leq (1 - \tau) \gamma_{n-1, d-1} + O_m(\delta \gamma_{n, d}) \leq (1 - \varepsilon) \gamma_{n-1, d-1},$$

as $\delta \ll \tau$ and $\varepsilon \ll \tau$. We therefore assume henceforth that

$$\sum_{r \in R', s \in S', t \in T'} \lambda_r \mu_s \cdot \widehat{F}_t(r, s) \gamma_t \widehat{F}'_t(r, s) \geq (1 - \tau) \sqrt{\sum_{r \in R', s \in S', t \in T'} \lambda_r^2 \mu_s^2 \widehat{F}_t(r, s)^2} \sqrt{\sum_{r \in R', s \in S', t \in T'} \gamma_t^2 \widehat{F}'_t(r, s)^2}. \quad (20)$$

The second inequality we have is that $\sum_{r \in R', s \in S'} \widehat{F}'_t(r, s)^2 \leq 1$ for all $t \in T'$, and if we were able to improve on it for $t \in T'$ that have sufficient weight, i.e. if we had that

$$\sum_{t \in T'} \gamma_t^2 \sum_{r \in R', s \in S'} \widehat{F}'_t(r, s)^2 \leq (1 - \tau) \sum_t \gamma_t^2 = 1 - \tau,$$

then we would also be done again in a similar manner. We henceforth assume that

$$\sum_{t \in T'} \gamma_t^2 \sum_{r \in R', s \in S'} \widehat{F}'_t(r, s)^2 \geq 1 - \tau. \quad (21)$$

Finally, in a similar way we may assume that

$$\sum_{r \in R', s \in S'} \lambda_r^2 \mu_s^2 \sum_{t \in T'} \widehat{F}_t(r, s)^2 \geq (1 - \tau) \gamma_{n-1, d-1}^2. \quad (22)$$

High level description of the argument. Taking inequalities (20), (21) and (22) to the extreme, we conclude that the case the argument is tight, i.e., that $\lambda_r \mu_s \widehat{F}_t(r, s)$ is proportional to $\gamma_t \widehat{F}'_t(r, s)$, $\sum_{t \in T'} \widehat{F}_t(r, s)^2 =$

$\gamma_{n-1, d-1}^2$ and $\sum_{r \in R', s \in S'} \widehat{F}'_t(r, s)^2 = 1$, and use these to reach a contradiction. We show that consider-

ing the $n - 1$ variate functions $\tilde{g} = \sum_{r \in R'} \lambda_r \psi_r g_r$, $\tilde{h} = \sum_{s \in S'} \mu_s \tilde{\psi}_s h_s$ and $\tilde{F} = \frac{\sum_t \langle F_t, \tilde{g} \tilde{h} \rangle F_t}{\sqrt{\sum_t \langle F_t, \tilde{g} \tilde{h} \rangle^2}}$ where ψ_r and

$\tilde{\psi}_s$ are random variables that preserve the 2-norm of the function (which can be thought of as random signs for now, though that doesn't quite work), one may make an appropriate choice of the signs so that $\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^t} [\tilde{F}(\mathbf{y}, \mathbf{z}) \tilde{g}(\mathbf{y}) \tilde{h}(\mathbf{z})]$ exceeds $\gamma_{n-1, d-1}$.

The formal argument. Denote $M_1 = \sum_{r \in R'} \lambda_r^2$ and $M_2 = \sum_{s \in S'} \mu_s^2$. We shall consider collections of real numbers $(\psi_r)_{r \in R'}$ and $(\tilde{\psi}_s)_{s \in S'}$ such that

$$\sum_{r \in R'} \lambda_r^2 \psi_r^2 = M_1, \quad \sum_{s \in S'} \mu_s^2 \tilde{\psi}_s^2 = M_2. \quad (23)$$

Later on, we will consider the uniform distribution over $(\psi_r)_{r \in R'}$ and $(\tilde{\psi}_s)_{s \in S'}$ satisfying these equations, but for now they can be thought of as fixed. The point of these equations is that defining the functions $\tilde{g}: \Gamma^I \rightarrow \mathbb{R}$ and $\tilde{h}: \Phi^I \rightarrow \mathbb{R}$ as above we have that $\|\tilde{g}\|_2^2 = M_1 \leq 1$, $\|\tilde{h}\|_2^2 = M_2 \leq 1$. Thus, $\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^I} \left[\tilde{F}(\mathbf{y}, \mathbf{z}) \tilde{g}(\mathbf{y}) \tilde{h}(\mathbf{z}) \right] \right| \leq \gamma_{n-1, d-1}$ for all choices of ψ_r and $\tilde{\psi}_s$ satisfying (23). Consider $\tilde{F} = \frac{\sum_t \langle F_t, \tilde{g} \tilde{h} \rangle F_t}{\sqrt{\sum_t \langle F_t, \tilde{g} \tilde{h} \rangle^2}}$, and note that $\|\tilde{F}\|_2^2 = 1$. Thus, we may define $p((\psi_r)_{r \in R'}, (\tilde{\psi}_s)_{s \in S'})$ as

$$p((\psi_r)_{r \in R'}, (\tilde{\psi}_s)_{s \in S'}) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^I} \left[\tilde{F}(\mathbf{y}, \mathbf{z}) \tilde{g}(\mathbf{y}) \tilde{h}(\mathbf{z}) \right]^2,$$

and get that

$$p((\psi_r)_{r \in R'}, (\tilde{\psi}_s)_{s \in S'}) \leq \gamma_{n-1, d-1}^2 \quad (24)$$

for every input satisfying (23). On the other hand, we note that

$$\begin{aligned} p((\psi_r)_{r \in R'}, (\tilde{\psi}_s)_{s \in S'}) &= \sum_t \langle F_t, \tilde{g} \tilde{h} \rangle^2 \\ &= \sum_t \left(\sum_{r, s} \lambda_r \psi_r \mu_s \tilde{\psi}_s \hat{F}_t(r, s) \right)^2 \\ &= \sum_{\substack{r, r' \in R' \\ s, s' \in S'}} \lambda_r \lambda_{r'} \mu_s \mu_{s'} \psi_r \psi_{r'} \tilde{\psi}_s \tilde{\psi}_{s'} \sum_{t \in T'} \hat{F}_t(r, s) \hat{F}_t(r', s'), \end{aligned}$$

so $p((\psi_r)_{r \in R'}, (\tilde{\psi}_s)_{s \in S'})$ is a polynomial in its input variables. We summarize this discussion with the following claim.

Claim 6.3. $p((\psi_r)_{r \in R'}, (\tilde{\psi}_s)_{s \in S'})$ is a polynomial equal to

$$\sum_{\substack{r, r' \in R' \\ s, s' \in S'}} \lambda_r \lambda_{r'} \mu_s \mu_{s'} \psi_r \psi_{r'} \tilde{\psi}_s \tilde{\psi}_{s'} \sum_{t \in T'} \hat{F}_t(r, s) \hat{F}_t(r', s')$$

that satisfies that $\left| p((\psi_r)_{r \in R'}, (\tilde{\psi}_s)_{s \in S'}) \right| \leq \gamma_{n-1, d-1}^2$ for every input satisfying (23).

Roughly speaking, in the rest of the proof we will show that the expected value of p over a uniform choice of input satisfying (23) is very close to $\gamma_{n-1, d-1}^2$, hence p is close to being a constant function. On the other hand, we reach a contradiction by directly arguing that the variance of p is large, hence concluding the proof.

For $r \in R'$ and $s \in S'$, define the vector $V_{r, s} \in \mathbb{R}^{T'}$ by $V_{r, s}(t) = \hat{F}_t(r, s)$. Then we may write p as

$$p((\psi_r)_{r \in R'}, (\tilde{\psi}_s)_{s \in S'}) = \sum_{\substack{r, r' \in R' \\ s, s' \in S'}} \lambda_r \lambda_{r'} \mu_s \mu_{s'} \psi_r \psi_{r'} \tilde{\psi}_s \tilde{\psi}_{s'} \langle V_{r, s}, V_{r', s'} \rangle. \quad (25)$$

6.3.1 Analyzing the expectation of p and upper bounding the variance of p

Claim 6.4. For all $r \in R'$ and $s \in S'$ we have that $\left| \|V_{r,s}\|_2^2 - \gamma_{n-1,d-1}^2 \right| \lesssim \frac{\tau}{\zeta^4} \gamma_{n-1,d-1}^2$.

Proof. It is clear that $\|V_{r,s}\|_2^2 \leq \gamma_{n-1,d-1}^2$ by definition of $\gamma_{n-1,d-1}$, and in the rest of the argument we show the lower bound. From (22) we get that

$$\sum_{r,s} \lambda_r^2 \mu_s^2 (\gamma_{n-1,d-1}^2 - \|V_{r,s}\|_2^2) \leq \tau \gamma_{n-1,d-1}^2,$$

hence for $r \in R'$ and $s \in S'$ we have that $\left| \gamma_{n-1,d-1}^2 - \|V_{r,s}\|_2^2 \right| \leq \frac{\tau}{\zeta^4} \gamma_{n-1,d-1}^2$. \square

As a corollary, we get:

Claim 6.5. $\mathbb{E}_{\psi, \tilde{\psi}} [p(\psi, \tilde{\psi})] \geq (1 - \sqrt{\tau}) \gamma_{n-1,d-1}^2$.

Proof. Note that for $r \neq r'$, the expectation of $\psi_r \psi_{r'}$ is 0, since the distribution of (ψ_r) is invariant under changing a sign of any ψ_r . Thus,

$$\begin{aligned} \mathbb{E}_{\psi, \tilde{\psi}} [p(\psi, \tilde{\psi})] &= \mathbb{E} \left[\sum_{r \in R', s \in S'} \lambda_r^2 \psi_r^2 \mu_s^2 \tilde{\psi}_s^2 \|V_{r,s}\|_2^2 \right] \geq (1 - O_\zeta(\tau)) \gamma_{n-1,d-1}^2 \mathbb{E} \left[\sum_{r \in R', s \in S'} \lambda_r^2 \psi_r^2 \mu_s^2 \tilde{\psi}_s^2 \right] \\ &= (1 - O_\zeta(\tau)) \gamma_{n-1,d-1}^2 M_1 M_2. \end{aligned}$$

As $M_1, M_2 \geq 1 - O_m(\delta)$, we get that

$$\mathbb{E}_{\chi, \tilde{\psi}} [p(\chi, \tilde{\psi})] \geq (1 - O_\zeta(\tau))(1 - O_m(\delta)) \gamma_{n-1,d-1}^2 \geq (1 - \sqrt{\tau}) \gamma_{n-1,d-1}^2. \quad \square$$

We can now upper bound the variance of p as:

Claim 6.6. $\text{var}(p) \leq 2\sqrt{\tau} \gamma_{n-1,d-1}^4$

Proof. By definition,

$$\text{var}(p) = \mathbb{E}[p^2] - \mathbb{E}[p]^2.$$

Note that by (24), $\mathbb{E}[p^2] \leq \gamma_{n-1,d-1}^4$, whereas by Claim 6.5 $\mathbb{E}[p] \geq (1 - \sqrt{\tau}) \gamma_{n-1,d-1}^2$, and the result follows. \square

6.3.2 Lower bounding the variance of p

Inspecting (25) it becomes apparent that to lower bound the variance of p , we must show that the vectors $V_{r,s}$ cannot be mutually orthogonal. We prove the following lemma in the next section. The proof proceeds by showing that if (21) holds and if the vectors are mutually (almost) orthogonal, then this contradicts Claim 6.2. This step crucially uses **the relaxed base case** condition.

Lemma 6.7. There are $r, r' \in R'$ and $s, s' \in S'$ such that $(r, s) \neq (r', s')$ and $|\langle V_{r,s}, V_{r',s'} \rangle| \gtrsim c \gamma_{n-1,d-1}^2$.

We are now ready to prove a lower bound, assuming the above lemma, on the variance of p . Towards this end, we write $p = p_1 + p_2 + p_3 + p_4$ so that

$$\begin{aligned} p_1 &= \sum_{r \in R', s \neq s' \in S'} \lambda_r^2 \mu_s \mu_{s'} \psi_r^2 \tilde{\psi}_s \tilde{\psi}_{s'} \langle V_{r,s}, V_{r,s'} \rangle, & p_2 &= \sum_{r \neq r' \in R', s \in S'} \lambda_r \lambda_{r'} \mu_s^2 \psi_r \psi_{r'} \tilde{\psi}_s^2 \langle V_{r,s}, V_{r',s} \rangle, \\ p_3 &= \sum_{r \neq r' \in R', s \neq s' \in S'} \lambda_r \lambda_{r'} \mu_s \mu_{s'} \psi_r \psi_{r'} \tilde{\psi}_s \tilde{\psi}_{s'} \langle V_{r,s}, V_{r',s'} \rangle, & p_4 &= \sum_{r \in R', s \in S'} \lambda_r^2 \mu_s^2 \psi_r^2 \tilde{\psi}_s^2 \langle V_{r,s}, V_{r,s} \rangle. \end{aligned}$$

Claim 6.8. $\text{var}(p) \geq \frac{1}{2} (\mathbb{E}[p_1^2] + \mathbb{E}[p_2^2] + \mathbb{E}[p_3^2]) - \tau^{1/8} \gamma_{n-1,d-1}^4$.

Proof. Note that by Claim 6.4

$$\begin{aligned} |p_4 - M_1 M_2 \gamma_{n-1,d-1}^2| &\leq \sum_{r \in R', s \in S'} \lambda_r^2 \mu_s^2 \psi_r^2 \tilde{\psi}_s^2 \left| \|V_{r,s}\|_2^2 - \gamma_{n-1,d-1}^2 \right| \\ &\lesssim \zeta \tau \gamma_{n-1,d-1}^2, \end{aligned}$$

and by Claim 6.5 and Claim 6.3

$$\left| \mathbb{E}[p] - M_1 M_2 \gamma_{n-1,d-1}^2 \right| \lesssim \sqrt{\tau} \gamma_{n-1,d-1}^2,$$

so together we get that $|\mathbb{E}[p] - p_4| \lesssim \sqrt{\tau} \gamma_{n-1,d-1}^2$. Thus, we get that

$$\begin{aligned} \text{var}(p) &= \mathbb{E}[(p_1 + p_2 + p_3 + p_4 - \mathbb{E}[p])^2] \\ &= \mathbb{E}[(p_1 + p_2 + p_3)^2] + 2\mathbb{E}[(p_1 + p_2 + p_3)(p_4 - \mathbb{E}[p])] + \mathbb{E}[(p_4 - \mathbb{E}[p])^2]. \end{aligned}$$

We bound

$$\begin{aligned} \left| \mathbb{E}[(p_1 + p_2 + p_3)(p_4 - \mathbb{E}[p])] \right| &\leq \sqrt{\mathbb{E}[(p_1 + p_2 + p_3)^2] \mathbb{E}[(p_4 - \mathbb{E}[p])^2]} \\ &\leq \tau^{1/8} \mathbb{E}[(p_1 + p_2 + p_3)^2] + \tau^{-1/8} \mathbb{E}[(p_4 - \mathbb{E}[p])^2], \end{aligned}$$

so

$$\begin{aligned} \text{var}(p) &\geq (1 - \tau^{1/8}) \mathbb{E}[(p_1 + p_2 + p_3)^2] - \tau^{-1/8} \mathbb{E}[(p_4 - \mathbb{E}[p])^2] \\ &\geq (1 - \tau^{1/8}) \mathbb{E}[(p_1 + p_2 + p_3)^2] - \tau^{1/8} \gamma_{n-1,d-1}^4. \end{aligned}$$

Finally, we note that $\mathbb{E}[p_1 p_2] = \mathbb{E}[p_1 p_3] = \mathbb{E}[p_2 p_3] = 0$ by the invariance of $\psi_r, \tilde{\psi}_s$ to sign change. The result follows. \square

Claim 6.9. $\text{var}(p) \gtrsim_m \zeta^8 \gamma_{n-1,d-1}^4$.

Proof. Let $r_1, r_2 \in R'$ and $s_1, s_2 \in S'$ be from Lemma 6.7 so that $(r_1, s_1) \neq (r_2, s_2)$ and $|\langle V_{r_1, s_1}, V_{r_2, s_2} \rangle| \geq c \gamma_{n-1,d-1}^2$. There are several cases, depending on if $r_1 = r_2, s_1 = s_2$ or none of them occur.

The case that $r_1 \neq r_2$ and $s_1 \neq s_2$ We get that

$$\mathbb{E} [p_3^2] = \mathbb{E} \left[\sum_{r \neq r' \in R', s \neq s' \in S'} \lambda_r^2 \lambda_{r'}^2 \mu_s^2 \mu_{s'}^2 \psi_r^2 \psi_{r'}^2 \tilde{\psi}_s^2 \tilde{\psi}_{s'}^2 \langle V_{r,s}, V_{r',s'} \rangle^2 \right]$$

as the other terms vanish by invariance under changing signs. This is at least

$$\gtrsim_m \gamma_{n-1,d-1}^4 \mathbb{E} \left[\lambda_{r_1}^2 \lambda_{r_2}^2 \mu_{s_1}^2 \mu_{s_2}^2 \psi_{r_1}^2 \psi_{r_2}^2 \tilde{\psi}_{s_1}^2 \tilde{\psi}_{s_2}^2 \right],$$

which is at least

$$\gtrsim_m \zeta^8 \gamma_{n-1,d-1}^4 \mathbb{E} \left[\psi_{r_1}^2 \psi_{r_2}^2 \tilde{\psi}_{s_1}^2 \tilde{\psi}_{s_2}^2 \right] = \zeta^8 \gamma_{n-1,d-1}^4 \mathbb{E} [\psi_{r_1}^2 \psi_{r_2}^2] \mathbb{E} [\tilde{\psi}_{s_1}^2 \tilde{\psi}_{s_2}^2].$$

By Claim 6.13 we get that this is at least $\gtrsim_m \zeta^8 \gamma_{n-1,d-1}^4$, so overall we get that $\mathbb{E} [p_3^2] \gtrsim_m \zeta^8 \gamma_{n-1,d-1}^4$. Hence by Claim (6.8) and (15) we get that $\text{var}(p) \gtrsim_m \zeta^8 \gamma_{n-1,d-1}^4$.

The case that $r_1 = r_2$ and $s_1 \neq s_2$. In this case,

$$\mathbb{E} [p_1^2] \geq \mathbb{E} \left[\sum_{r,r' \in R', s \neq s' \in S'} \lambda_r^2 \lambda_{r'}^2 \mu_s^2 \mu_{s'}^2 \psi_r^2 \psi_{r'}^2 \tilde{\psi}_s^2 \tilde{\psi}_{s'}^2 \langle V_{r,s}, V_{r,s'} \rangle^2 \right],$$

and the rest of the proof is identical.

The case that $r_1 \neq r_2$ and $s_1 = s_2$. Analogous to the previous case. □

6.3.3 Finishing the proof

Claims 6.6 and 6.9 directly contradict each other by (15). That means that our initial hypothesis is false, i.e. not all of (20), (21) and (22) can hold, and therefore as explained in the beginning of Section 6.3, the conclusion of Lemma 6.1 follows.

6.3.4 Showing $\{V_{r,s}\}_{r \in R', s \in S'}$ cannot be roughly orthogonal

This section is devoted to the proof of Lemma 6.7. We assume towards contradiction that this is false, i.e. that $|\langle V_{r,s}, V_{r',s'} \rangle| < c\gamma_{n-1,d-1}^2$ for all $(r,s) \neq (r',s')$.

Claim 6.10. $\sum_{r \in R', s \in S'} \widehat{F}'_t(r,s)^2 \geq 1 - \sqrt{\tau}$ for all $t \in T'$.

Proof. Note that

$$\sum_{t \in T'} \gamma_t^2 \geq 1 - \sum_{t \in T \setminus T'} \gamma_t^2 \geq 1 - m^2 \delta^2.$$

Thus, by (21) we get that

$$\sum_{t \in T'} \gamma_t^2 \left(1 - \sum_{r \in R', s \in S'} \widehat{F}'_t(r,s)^2 \right) \lesssim_m \delta^2 + \tau \lesssim_m \tau.$$

In particular, for all $t \in T'$ we have

$$\gamma_t^2 \left(1 - \sum_{r \in R', s \in S'} \widehat{F}'_t(r, s)^2 \right) \lesssim_m \tau,$$

so

$$1 - \sum_{r \in R', s \in S'} \widehat{F}'_t(r, s)^2 \lesssim_m \frac{\tau}{\zeta^2} \leq \sqrt{\tau},$$

establishing the claim. \square

Claim 6.11. For all r, s we have $\sum_{t \in T'} \widehat{F}'_t(r, s)^2 > 1 - v$.

Proof. Assume this is not the case, and that there are r, s such that $\sum_{t \in T'} \widehat{F}'_t(r, s)^2 \leq 1 - v$.

Claim 6.12. $|T'| \leq |R'| |S'| - 1$.

Proof. Summing Claim 6.10 over $t \in T'$ yields

$$|T'| \leq \frac{1}{1 - \sqrt{\tau}} \sum_{t \in T'} \sum_{r \in R', s \in S'} \widehat{F}'_t(r, s)^2 = \sum_{r \in R', s \in S'} \sum_{t \in T'} \widehat{F}'_t(r, s)^2.$$

For all $r \in R'$ and $s \in S'$, we have that $\sum_{t \in T'} \widehat{F}'_t(r, s)^2 \leq 1$, and by assumption for some r, s we have

$\sum_{t \in T'} \widehat{F}'_t(r, s)^2 \leq 1 - v$. Hence,

$$\begin{aligned} |T'| &\leq \frac{1}{1 - \sqrt{\tau}} \sum_{r \in R', s \in S'} \sum_{t \in T'} \widehat{F}'_t(r, s)^2 \leq \frac{(|R'| |S'| - 1) + 1 - v}{1 - \sqrt{\tau}} = |R'| |S'| - \frac{\sqrt{\tau} |R'| |S'| - v}{1 - \sqrt{\tau}} \\ &< |R'| |S'|, \end{aligned}$$

as $\sqrt{\tau} |R'| |S'| \leq m^2 \sqrt{\tau} < v$. As $|T'|$ is an integer, the statement of Claim 6.12 follows. \square

Thus, applying Lemma 6.14 we get that

$$\sum_{(r, s) \neq (r', s')} \left\langle \frac{V_{r, s}}{\|V_{r, s}\|_2}, \frac{V_{r', s'}}{\|V_{r', s'}\|_2} \right\rangle^2 \geq \frac{1}{|T'|} \geq \frac{1}{m^2},$$

and therefore there are $(r, s) \neq (r', s')$ such that

$$|\langle V_{r, s}, V_{r', s'} \rangle| \geq \frac{1}{m^3} \|V_{r, s}\|_2 \|V_{r', s'}\|_2 > c\gamma_{n-1, d-1}^2,$$

where we used Claim 6.4. This is a contradiction, hence proving the assertion of Claim 6.11. \square

Proof of Lemma 6.7. Let $V'_{r,s} = (\widehat{F}'_t(r,s))_{t \in T'}$, and define

$$\tilde{F}_{r,s} = \frac{\sum_{t \in T'} \widehat{F}'_t(r,s) F'_t}{\sqrt{\sum_{t \in T'} \widehat{F}'_t(r,s)^2}}.$$

Note that as $\tilde{F}_{r,s}$ is constant on connected components, we may write $\tilde{F}_{r,s} = W \tilde{f}_{r,s}$. From Claim 6.11, it follows that for all r, s

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^J} \left[\tilde{f}_{r,s}(\mathbf{x}) g'_r(\mathbf{y}) h'_s(\mathbf{z}) \right]^2 = \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^J} \left[\tilde{F}_{r,s}(\mathbf{y}, \mathbf{z}) g'_r(\mathbf{y}) h'_s(\mathbf{z}) \right]^2 \geq 1 - v,$$

hence by **the relaxed base case** we get that $\text{var}_{\Sigma'}(\tilde{f}_{r,s}) \leq u$. We now show that this implies that $\text{var}_{\Sigma'}(f'_t)$ is small contradicting Claim 6.2.

We have that

$$\|\tilde{F}_{r,s} - g'_r h'_s\|_2^2 = \|\tilde{F}_{r,s}\|_2^2 + \|g'_r h'_s\|_2^2 - 2 \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^J} \left[\tilde{F}_{r,s}(\mathbf{y}, \mathbf{z}) g'_r(\mathbf{y}) h'_s(\mathbf{z}) \right] \leq 2v.$$

Therefore, for all $t \in T'$ we have

$$\|F'_t - \sum_{r,s} \widehat{F}'_t(r,s) \tilde{F}_{r,s}\|_2 \leq \|F'_t - \sum_{r,s} \widehat{F}'_t(r,s) g'_r h'_s\|_2 + \|\sum_{r,s} \widehat{F}'_t(r,s) (g'_r h'_s - \tilde{F}_{r,s})\|_2. \quad (26)$$

Note that

$$\|F'_t - \sum_{r,s} \widehat{F}'_t(r,s) g'_r h'_s\|_2^2 = \|F'_t\|_2^2 - \sum_{r,s} \widehat{F}'_t(r,s)^2 = 1 - \sum_{r,s} \widehat{F}'_t(r,s)^2 \leq \sqrt{\tau},$$

where in the last inequality we used Claim 6.10. Also,

$$\|\sum_{r,s} \widehat{F}'_t(r,s) (g'_r h'_s - \tilde{F}_{r,s})\|_2 \leq \sum_{r,s} \widehat{F}'_t(r,s) \|g'_r h'_s - \tilde{F}_{r,s}\|_2 \leq \sum_{r,s} \widehat{F}'_t(r,s) 2v \leq 2m^2 v.$$

Plugging these bounds into (26) yields $\|F'_t - \sum_{r,s} \widehat{F}'_t(r,s) \tilde{F}_{r,s}\|_2 \leq \tau^{1/4} + 2m^2 v \leq \sqrt{v}$. Thus, as

$$\text{var}_{\Sigma'}\left(\sum_{r,s} \widehat{F}'_t(r,s) \tilde{f}_{r,s}\right) \lesssim_m u,$$

we get that

$$\text{var}_{\Sigma'}(W^{-1} F'_t) \lesssim_m u + \sqrt{v}$$

for all $t \in T'$, which contradicts Claim 6.2 as $u, v \lesssim w$ due to (15). \square

6.4 Auxiliary statements

Claim 6.13. $\mathbb{E} [\psi_{r_1}^2 \psi_{r_2}^2] \gtrsim_m 1$

Proof. Note that the random vector $(\lambda_r \psi_r)$ is distributed uniformly over the unit sphere in $\mathbb{R}^{B'}$, hence

$$\mathbb{E} [(\lambda_{r_1} \psi_{r_1})^2 (\lambda_{r_2} \psi_{r_2})^2] \gtrsim_m 1,$$

and $\mathbb{E} [\psi_{r_1}^2 \psi_{r_2}^2]$ is at least as large. \square

Lemma 6.14. *Suppose we have a collection of unit vectors $v_1, \dots, v_q \in \mathbb{R}^T$ such that $q > T$. Then*

$$\sum_{i \neq j} |\langle v_i, v_j \rangle|^2 \geq \frac{q(q-T)}{T}.$$

Proof. Consider the matrix M whose (i, j) entry is $\langle v_i, v_j \rangle$. Then M is symmetric and the rank of M is at most T . Note that

$$\sum_{i,j} |\langle v_i, v_j \rangle|^2 = \text{Tr}(M^2) = \sum_{\ell=1}^T \lambda_\ell(M)^2 \geq \frac{1}{T} \left(\sum_{\ell=1}^T \lambda_\ell(M) \right)^2 = \frac{1}{T} \text{Tr}(M)^2.$$

Therefore,

$$\sum_{i \neq j} |\langle v_i, v_j \rangle|^2 \geq \frac{1}{T} \text{Tr}(M)^2 - q = \frac{1}{T} q^2 - q = \frac{q(q-T)}{T}. \quad \square$$

7 Proof of reductions

In this section, we prove Lemma 1.4 assuming the correctness of Lemma 1.10. As discussed in the introduction, there are two main differences between these lemma statements. The first difference is that we assume the marginal distribution on (y, z) is uniform and independent. The second difference is in the use of the relaxed base case. In Section 7.1, we show that the marginal distribution on (y, z) can be assumed to be uniform and independent. In Section 7.2 we show that the relaxed base case inequality is sufficient towards proving the main analytical lemma.

We use the following lemma as an intermediate lemma in showing this. Compared to Lemma 1.4, in the following lemma, we get to assume that the distribution $\mu_{y,z}$ is uniform over $\Gamma \times \Phi$.

Lemma 7.1. *For all $\alpha > 0$, $m \in \mathbb{N}$ there exists $\xi > 0$ such that for all $\varepsilon > 0$ there is $\delta > 0$ such that the following holds.*

Suppose $|\Sigma|, |\Gamma|, |\Phi| \leq m$ and μ is a distribution over $\Sigma \times \Gamma \times \Phi$ such that (a) the support of μ cannot be linearly embedded, (b) $\mu(x, y, z) \geq \alpha$ for all $(x, y, z) \in \text{supp}(\mu)$, (c) the distribution $\mu_{y,z}$ is uniform over $\Gamma \times \Phi$. Then for all $f: \Sigma^n \rightarrow [-1, 1]$, $g: \Gamma^n \rightarrow [-1, 1]$ and $h: \Phi^n \rightarrow [-1, 1]$ satisfying that $\text{Stab}_{1-\xi}(g; \mu_y) \leq \delta$ we have that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right| \leq \varepsilon.$$

7.1 Lemma 7.1 implies Lemma 1.4

Let f, g, h be functions as in the statement of Lemma 1.4, and suppose without loss of generality that $\text{Stab}_{1-\xi}(g) \leq \delta$.

7.1.1 The path trick

The argument herein is virtually identical to the arguments in [5, Sections 3.1, 3.2, 3.3], but we give it for sake of completeness. We let

$$0 \ll \delta \ll \varepsilon, \xi \ll \alpha, m^{-1} \leq 1,$$

and fix f, g and h as in the statement of the lemma. Without loss of generality we assume that $\text{Stab}_{1-\xi}(g) \leq \delta$.

Our first goal will be to reduce to the case the distribution μ has full support over $\Gamma \times \Phi$. Define

$$\tilde{h}(z) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z})} [f(\mathbf{x})h(\mathbf{z}) \mid \mathbf{z} = z].$$

We note that $\tilde{h}: \Phi^n \rightarrow [-1, 1]$ as the range of f, g is $[-1, 1]$, and also that by Cauchy–Schwarz

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right|^2 = \left| \mathbb{E}_{\mathbf{z} \sim \mu_z} [h(\mathbf{z})\tilde{h}(\mathbf{z})] \right|^2 \leq \|h\|_2^2 \|\tilde{h}\|_2^2 \leq \|\tilde{h}\|_2^2 = \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu} [f(\mathbf{x})g(\mathbf{y})\tilde{h}(\mathbf{z})],$$

so it suffices to prove the statement for f, g, \tilde{h} . To simplify notations, we henceforth assume that $h = \tilde{h}$ to begin with.

For each $r \geq 1$, consider the distribution \mathcal{D}_r over

$$(x^1, x^{1'}, x^2, x^{2'}, \dots, x^{2^{r-1}}, x^{2^{r-1}'}, y^1, y^2, y^3, \dots, y^{2^{r-1}+1}, z^1, z^2, \dots, z^{2^{r-1}})$$

defined as follows:

1. Sample $y^1 \sim \mu_y$;
2. Sample $(x^1, z^1) \sim \mu|y^1$;
3. Sample $(x^{1'}, y^2) \sim \mu|z^1$;
4. Iteratively, for $j \leq 2^{r-1}$, after sampling y^j we sample $(x^j, z^j) \sim \mu|y^j$;
5. Iteratively, for $j \leq 2^{r-1}$, after sampling z^j we sample $(x^{j'}, y^{j+1}) \sim \mu|z^j$.

This distribution can be viewed as a labeled random walk of length 2^r between the y and z side on the bipartite graph on $H = (\Gamma \cup \Phi, E)$ whose directed edges are $E = \text{supp}(\mu_{y,z})$; each directed edge is labeled by an x showing that the 3-tuple is in the support of μ . We note that reversing the order of the random walk, i.e. viewing this as a random walk from $y^{2^{r-1}}$ to y^1 , yields the same distribution.

Claim 7.2. *For all $r \geq 1$ it holds that*

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu} [\tilde{f}(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right|^{2^r} \leq \mathbb{E}_{(\vec{x}, \vec{y}, \vec{z}) \sim \mathcal{D}_r} [F(\vec{x})g(y^1)g(y^{2^{r-1}+1})],$$

where $F(\vec{x}) = \prod_{i=1}^{2^{r-1}} f(x^i)f(x^{i'})$.

Proof. This is repeated application of Cauchy–Schwarz, and is done by induction on r . For $r = 1$ this is true as

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right|^2 = \left| \mathbb{E}_{\mathbf{z} \sim \mu_z} [h(\mathbf{z}) \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \bar{\mathbf{z}}) \sim \mu} [f(\mathbf{x})g(\mathbf{y}) \mid \bar{\mathbf{z}} = \mathbf{z}]] \right|^2,$$

and using Cauchy–Schwarz over z yields this is at most

$$\mathbb{E}_{\mathbf{z} \sim \mu_z} \left[\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \bar{\mathbf{z}}) \sim \mu} [f(\mathbf{x})g(\mathbf{y}) \mid \bar{\mathbf{z}} = \mathbf{z}] \right|^2 \right] = \mathbb{E}_{(\bar{\mathbf{x}}, \bar{\mathbf{y}}, \bar{\mathbf{z}}) \sim \mathcal{D}_1} \left[f(\mathbf{x}^1)f(\mathbf{x}^{1'})g(\mathbf{y}^1)g(\mathbf{y}^2) \right].$$

Suppose the statement is true for r , and prove for $r + 1$. Then by induction hypothesis

$$\begin{aligned} \left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right|^{2^{r+1}} &\leq \mathbb{E}_{(\bar{\mathbf{x}}, \bar{\mathbf{y}}, \bar{\mathbf{z}}) \sim \mathcal{D}_r} \left[F_r(\bar{\mathbf{x}})g(\mathbf{y}^1)g(\mathbf{y}^{2^{r-1}+1}) \right]^2 \\ &= \mathbb{E}_{y^{2^{r-1}+1}} \left[g(y^{2^{r-1}+1}) \mathbb{E}_{(\bar{\mathbf{x}}, \bar{\mathbf{y}}, \bar{\mathbf{z}}) \sim \mathcal{D}_r} \left[F_r(\bar{\mathbf{x}})g(\mathbf{y}^1) \mid y^{2^{r-1}+1} \right] \right]^2, \end{aligned}$$

where $F_r(\bar{\mathbf{x}}) = \prod_{i=1}^{2^{r-1}} f(x^i)f(x^{i'})$. Hence by Cauchy–Schwarz this is bounded by

$$\begin{aligned} &\|g\|_2^2 \mathbb{E}_{y^{2^{r-1}+1}} \left[\mathbb{E}_{(\bar{\mathbf{x}}, \bar{\mathbf{y}}, \bar{\mathbf{z}}) \sim \mathcal{D}_r} \left[F_r(\bar{\mathbf{x}})g(\mathbf{y}^1) \mid y^{2^{r-1}+1} \right]^2 \right] \\ &\leq \mathbb{E}_{y^{2^{r-1}+1}} \left[\mathbb{E}_{\substack{(\bar{\mathbf{x}}, \bar{\mathbf{y}}, \bar{\mathbf{z}}) \sim \mathcal{D}_r \\ (\bar{\mathbf{x}}'', \bar{\mathbf{y}}'', \bar{\mathbf{z}}'') \sim \mathcal{D}_r}} \left[F_r(\bar{\mathbf{x}})F_r(\bar{\mathbf{x}}'')g(\mathbf{y}^1)g(\mathbf{y}''^1) \mid y^{2^{r-1}+1} = \mathbf{y}''^{2^{r-1}+1} = y^{2^{r-1}+1} \right] \right]. \end{aligned}$$

We may view $(\bar{\mathbf{x}}, \bar{\mathbf{y}}, \bar{\mathbf{z}})$ as a random walk starting at $y^{2^{r-1}+1}$, and $(\bar{\mathbf{x}}'', \bar{\mathbf{y}}'', \bar{\mathbf{z}}'')$ as an independently chosen random walk starting at $y^{2^{r-1}+1}$, therefore $(\bar{\mathbf{x}}, \text{reverse}(\bar{\mathbf{x}}''), \bar{\mathbf{y}}, \text{reverse}(\bar{\mathbf{y}}''), \bar{\mathbf{z}}, \text{reverse}(\bar{\mathbf{z}}''))$ describes a random walk of length $2 \cdot 2^{r-1} = 2^r$. We note that $\bar{\mathbf{y}}$ and $\text{rev}(\bar{\mathbf{y}}'')$ overlap in their starting point, and after removing this overlap we get that this random walk matches the distribution \mathcal{D}_{r+1} , so the inductive proof is complete. \square

Claim 7.3. For all $r \geq 1$ it holds that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right|^{2^r} \leq \mathbb{E}_{(\bar{\mathbf{x}}, \bar{\mathbf{y}}, \bar{\mathbf{z}}) \sim \mathcal{D}_r} \left[F(\bar{\mathbf{x}})g(\mathbf{y}^1)h(\mathbf{z}^{2^{r-1}}) \right],$$

where $F(\bar{\mathbf{x}}) = \prod_{i=1}^{2^{r-1}} f(\mathbf{x}^i) \cdot \prod_{i=1}^{2^{r-1}-1} f(\mathbf{x}^{i'})$.

Proof. The statement follows from applying the previous claim and noting that fixing $z^{2^{r-1}}$, the distribution of $(x^{2^{r-1}'}, y^{2^{r-1}+1})$ is $\mu_{x,y} \mid z = z^{2^{r-1}}$, hence

$$\mathbb{E}_{x^{2^{r-1}'}, y^{2^{r-1}+1}} \left[f(x^{2^{r-1}'})g(y^{2^{r-1}+1}) \mid z^{2^{r-1}} \right] = h(z^{2^{r-1}}).$$

\square

Claim 7.4. *The graph H is connected.*

Proof. Otherwise, we could write $\Gamma = \Gamma_0 \cup \Gamma_1$, $\Phi = \Phi_0 \cup \Phi_1$ non-trivial partitions and have that there are no edges between Γ_0, Φ_1 as well as between Γ_1, Φ_0 . We could then define the embedding over \mathbb{F}_2 defined as $\gamma(y) = i$ if $y \in \Gamma_i$, $\phi(z) = i$ if $z \in \Phi_i$ and have that $\gamma(y) + \phi(z) = 0$ in the support of μ , hence μ is linearly embeddable, and contradiction. \square

Thus, it follows that taking $r = \lceil 2 + \log m \rceil$, the distribution of (y^1, z^{2^r-1}) in \mathcal{D}_r has full support over $\Gamma \times \Phi$, and each element in $\Gamma \times \Phi$ has probability at least α^{2^r} in \mathcal{D}_r . We define the distribution of

$$\nu = ((x^1, x^{1'}, x^2, x^{2'} \dots, x^{2^r-1}, x^{2^r-1'}, x^{2^r-1}), y^1, z^{2^r-1})$$

as distribution over $\Sigma^{2^r-1} \times \Gamma \times \Phi$.

Claim 7.5. *The distribution ν is not linearly embeddable.*

Proof. Otherwise, we would have a non-trivial embedding $\sigma: \Sigma^{2^r-1} \rightarrow G$, $\gamma: \Gamma \rightarrow G$ and $\phi: \Phi \rightarrow G$. Hence, at least two of these functions are not constant. Note that for $(x, y, z) \in \text{supp}(\mu)$ we have that $(\vec{x}, y, z) \in \text{supp}(\nu)$ for $\vec{x} = (x, x, \dots, x)$, so $\sigma': \Sigma \rightarrow G$, defined by $\sigma'(x) = \sigma(\vec{x})$ forms an embedding, together with γ and ϕ , of $\text{supp}(\mu)$ to G , and contradiction. \square

Thus, moving from μ and f, g, h to ν and F, g, h , we get that in order to prove Lemma 1.4 it suffices to prove the following lemma.

Lemma 7.6. *For all $\alpha > 0$, $m \in \mathbb{N}$ there exists $\xi > 0$ such that for all $\varepsilon > 0$ there is $\delta > 0$ such that the following holds.*

Suppose $|\Sigma|, |\Gamma|, |\Phi| \leq m$ and μ is a distribution over $\Sigma \times \Gamma \times \Phi$ such that (a) the support of μ cannot be linearly embedded, (b) $\mu(x, y, z) \geq \alpha$ for all $(x, y, z) \in \text{supp}(\mu)$, (c) the support of $\mu_{y,z}$ is $\Gamma \times \Phi$. Then for all $f: \Sigma^n \rightarrow [-1, 1]$, $g: \Gamma^n \rightarrow [-1, 1]$ and $h: \Phi^n \rightarrow [-1, 1]$ satisfying that $\text{Stab}_{1-\xi}(g; \mu_y) \leq \delta$ we have that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right| \leq \varepsilon.$$

Next, we reduce proving the above lemma to proving Lemma 7.1 in which we get to assume that the distribution $\mu_{y,z}$ is uniform over $\Gamma \times \Phi$.

Lemma 7.1 implies Lemma 7.6. Let f, g, h and μ be as in Lemma 7.6. We may write $\mu = (1-s)\nu_1 + s\nu_2$, where $s \gtrsim_1 \alpha$ and ν_1, ν_2 have the same support as μ , the probability of each atom in them is at least $\alpha'(\alpha) > 0$ and $(\nu_2)_{y,z}$ is uniform over $\Gamma \times \Phi$. Thus, we choose a subset $I \subseteq [n]$ by including each element in it with probability $1-s$, and sampling $(\mathbf{x}_I, \mathbf{y}_I, \mathbf{z}_I) \sim \nu_1^I$, $(\mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}) \sim \nu_2^{\bar{I}}$ yields a sample of $\mu^{\otimes n}$. We may thus write

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] = \mathbb{E}_{(\mathbf{x}_I, \mathbf{y}_I, \mathbf{z}_I) \sim \nu_1^{\otimes n}} \left[\mathbb{E}_{(\mathbf{x}', \mathbf{y}', \mathbf{z}') \sim \nu_2^{\bar{I}}} [f'(\mathbf{x}')g'(\mathbf{y}')h'(\mathbf{z}')] \right], \quad (27)$$

where

$$f' = f_{I \rightarrow \mathbf{x}_I}, \quad g' = g_{I \rightarrow \mathbf{y}_I}, \quad h' = h_{I \rightarrow \mathbf{z}_I}.$$

We note that by Lemma 2.14

$$\mathbb{E}_{\mathbf{y}_I} \left[\text{Stab}_{1-\xi'}(g'; \nu_1^{\bar{I}}) \right] \leq \text{Stab}_{1-\alpha\xi'}(g)$$

for $c = c(s, m, \alpha) > 0$. Hence, taking $\xi' = \xi/c$ we get that $\mathbb{E}_{y_I} \left[\text{Stab}_{1-\xi'}(g'; \nu_1^{\bar{I}}) \right] \leq \delta$, and by Markov we get that

$$\Pr_{y_I} \left[\text{Stab}_{1-\xi'}(g'; \nu_1^{\bar{I}}) \geq \sqrt{\delta} \right] \leq \sqrt{\delta}.$$

Using that in (27), along with the trivial bound $\left| \mathbb{E}_{(\mathbf{x}', \mathbf{y}', \mathbf{z}') \sim \nu_2^{\bar{I}}} [f'(\mathbf{x}')g'(\mathbf{y}')h'(\mathbf{z}')] \right| \leq 1$ for all x_I, y_I, z_I yields that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right| \leq \sqrt{\delta} + \mathbb{E}_{(\mathbf{x}_I, \mathbf{y}_I, \mathbf{z}_I) \sim \nu_1^{\otimes n}} \left[\left| \mathbb{E}_{(\mathbf{x}', \mathbf{y}', \mathbf{z}') \sim \nu_2^{\bar{I}}} [f'(\mathbf{x}')g'(\mathbf{y}')h'(\mathbf{z}')] 1_{\text{Stab}_{1-\xi'}(g'; \nu_1^{\bar{I}}) \leq \sqrt{\delta}} \right| \right]$$

Using Lemma 7.1, for sufficiently small δ the inner expectation is at most $\varepsilon/2$, hence restricting to $\delta \leq \varepsilon^2/4$ gives $\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right| \leq \varepsilon$, and the proof is complete. \square

7.2 Lemma 1.10 implies Lemma 7.1

The remainder of this section is devoted to showing that Lemma 1.10 implies Lemma 7.1, and we begin with a remark. Note that in the statement of Lemma 7.1, the roles of x and y are exchangeable. Namely, to prove it, it is enough to prove that the conclusion of it holds for a distribution μ that satisfies (a), (b), and that $\mu_{x,y}$ is uniform; we assume this henceforth. We fix functions f, g, h as in the statement of Lemma 7.1, and we prove that the conclusion holds.

7.2.1 Defining the distribution $\tilde{\nu}$

We take r a large enough constant, and consider the distribution \mathcal{D}_r resulting from μ by performing the path-trick, so that by Claim 7.3 we have that for some bounded functions F, \tilde{h} it holds that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right|^{2^r} \leq \left| \mathbb{E}_{(\tilde{\mathbf{x}}, \mathbf{y}, \mathbf{z}) \sim \mathcal{D}_r} [F(\tilde{\mathbf{x}})g(\mathbf{y})\tilde{h}(\mathbf{z})] \right|; \quad (28)$$

here, we think of \mathcal{D}_r as only recording the sequence of x 's used, as well as y^1, z^{2^r-1} . Thus, as observed above, provided that r is large enough with respect to m , we get that in \mathcal{D}_r the support of (y, z) is full; we fix such r .

Denote $\tilde{\Sigma} = \Sigma^{2^r}$, and define a graph X on $\tilde{\Sigma}$ where \tilde{x} and \tilde{x}' are adjacent if there is $(y, z) \in \Gamma \times \Phi$ such that (\tilde{x}, y, z) and (\tilde{x}', y, z) are both in the support of \mathcal{D}_r . We look at the connected components of X , and from each one of them we pick a representative arbitrarily; denote by $\text{rep}(\tilde{x})$ the chosen representative form the connected component of \tilde{x} , and let

$$\Sigma_{\text{fin}} = \left\{ \text{rep}(\tilde{x}) \mid \tilde{x} \in \tilde{\Sigma} \right\} \subseteq \tilde{\Sigma}$$

denote the set of these representatives.

Next, we define a distribution $\tilde{\nu}$. The distribution $\tilde{\nu}$ first samples $(y, z) \in \Gamma \times \Phi$ uniformly, and then takes the unique $\tilde{x} \in \Sigma_{\text{fin}}$ such that $(\tilde{x}, y, z) \in \text{supp}(\mathcal{D}_r)$, and outputs (\tilde{x}, y, z) . It is clear that the probability of each atom is at least $\alpha'(\alpha, m) > 0$, that $\tilde{\nu}_{y,z}$ is uniform and that in $\tilde{\nu}$, y, z implies x . It is also easy to observe that as by Claim 7.5 the support of ν cannot be linearly embedded, the support of $\tilde{\nu}$ can also not be linearly embedded.

In the remainder of the proof, we will show that $\tilde{\nu}$ satisfies the relaxed base case, and therefore we may apply Lemma 1.10 to it. We then use a random-restriction argument to deduce Lemma 7.1.

Lemma 7.7. *The distribution $\tilde{\nu}$ satisfies the relaxed base case.*

Proof. Since $\text{supp}(\tilde{\nu}) = \text{supp}(\nu)$, Claim 7.5 implies that $\tilde{\nu}$ cannot be linearly embedded. The rest of the proof is deferred to Sections 7.3, 7.4. \square

Together from Claim 7.13 and Lemma 7.7 we get that $\tilde{\nu}$ satisfies the relaxed base case. Therefore, we can apply Lemma 1.10 on it. In the next section, we show how to do so and establish Lemma 7.1.

7.2.2 Proving Lemma 7.1: the merging and the random restriction arguments

Let

$$0 \ll \delta \ll \delta' \ll \varepsilon' \ll \gamma \ll \varepsilon \ll \xi \ll \xi' \ll \alpha, m^{-1} \leq 1;$$

we show that $\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right| \leq \varepsilon$. Indeed, towards contradiction to Lemma 7.1, assume that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right| > \varepsilon. \quad (29)$$

Thus, by (28) we get that

$$\left| \mathbb{E}_{(\vec{\mathbf{x}}, \mathbf{y}, \mathbf{z}) \sim \mathcal{D}_r^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right| \geq \varepsilon^{2^r}.$$

Consider the distribution \mathcal{D}'_r over $\Sigma_{\text{fin}} \times \Gamma \times \Phi$, defined by taking a sample $(\vec{x}, y, z) \sim \mathcal{D}_r$ and outputting $(\text{rep}(\vec{x}), y, z)$.

Claim 7.8. *There is $n' \geq \gamma n$ and functions $f': \Sigma_{\text{fin}}^{n'} \rightarrow [-1, 1]$, $g': \Gamma^{n'} \rightarrow [-1, 1]$, $h': \Phi^{n'} \rightarrow [-1, 1]$ such that:*

1. $\left| \mathbb{E}_{(\vec{\mathbf{x}}, \mathbf{y}, \mathbf{z}) \sim \mathcal{D}'_r \otimes n'} [f'(\vec{\mathbf{x}})g(\mathbf{y})h(\mathbf{z})] \right| \geq \varepsilon'$.
2. $\text{Stab}_{1-\xi'}(g') \leq \delta'$.

Proof. This is a direct consequence of [5, Lemma 3.10]. In the notation therein, Σ' represents the set of connected components of the graph X we defined, and noting the resulting distribution there (μ' in the notation therein) is exactly \mathcal{D}'_r . \square

We fix f', g', h' from the above claim henceforth.

Note that for some $s \gg_{\alpha, m} 1$ we may write $\mathcal{D}'_r = (1-s)\mathcal{D} + s\tilde{\nu}$, where \mathcal{D} is a distribution and $\tilde{\nu}$ is the distribution we have defined earlier; this is true since the probability of each atom in \mathcal{D}'_r is at least $\alpha'(\alpha, m) > 0$ and $\text{supp}(\tilde{\nu}) = \text{supp}(\mathcal{D}'_r)$, so $\mathcal{D}'_r - s\tilde{\nu}$ is non-negative for $s = \alpha'/2$, and re-normalizing it we get the distribution \mathcal{D} .

Consider a random restriction on f', g', h' as follows: sample $I \subseteq [n]$ by including each element with probability $(1-s)$, sample $(x, y, z) \sim \mathcal{D}^I$, and define $f'': \Sigma^{[n] \setminus I} \rightarrow [-1, 1]$, $g'': \Gamma^{[n] \setminus I} \rightarrow [-1, 1]$, $h'': \Phi^{[n] \setminus I} \rightarrow [-1, 1]$ as

$$f''(\mathbf{x}') = f'_{I \rightarrow \mathbf{x}}(\mathbf{x}'), \quad g''(\mathbf{y}') = g'_{I \rightarrow \mathbf{y}}(\mathbf{y}'), \quad h''(\mathbf{z}') = h'_{I \rightarrow \mathbf{z}}(\mathbf{z}').$$

Consider the following events:

1. $E_1: \text{Stab}_{1-\xi'/cs}(g''; \tilde{\nu}_y) \leq \sqrt{\delta}$ for some $c = c(m, \alpha) > 0$;

2. $E_2: \left| \mathbb{E}_{(\mathbf{x}', \mathbf{y}', \mathbf{z}') \sim \tilde{\nu}^{[n] \setminus I}} [f''(\mathbf{x}')g''(\mathbf{y}')h''(\mathbf{z}')] \right| \geq \frac{\varepsilon}{2};$
3. $E_3: |[n] \setminus I| \geq \frac{\varepsilon}{2}n.$

Claim 7.9. $\Pr [E_1] \geq 1 - \sqrt{\delta}.$

Proof. Noting that by Lemma 2.14 we have that $\mathbb{E}_y [\text{Stab}_{1-\xi'/sc}(g'')] \leq \text{Stab}_{1-\xi'}(g') \leq \delta$, the result follows by Markov's inequality. \square

Claim 7.10. $\Pr [E_2] \geq \frac{\varepsilon}{2}.$

Proof. Setting $T = \mathbb{E}_{(\mathbf{x}', \mathbf{y}', \mathbf{z}') \sim \tilde{\nu}^{[n] \setminus I}} [f''(\mathbf{x}')g''(\mathbf{y}')h''(\mathbf{z}')]$, the expected value of T is at least ε , and $|T| \leq 1$ always, hence we get that $T \geq \varepsilon/2$ with probability at least $\varepsilon/2$. \square

Claim 7.11. $\Pr [E_3] \geq 1 - o_n(1) \geq 1 - \frac{\varepsilon}{4}.$

Proof. This is an immediate consequence of Chernoff's bound. \square

Thus, $\Pr [E_1 \cap E_2 \cap E_3] > 0$, so we may fix x, y, z that satisfy these events. Then

$$\left| \mathbb{E}_{(\mathbf{x}', \mathbf{y}', \mathbf{z}') \sim \tilde{\nu}^{[n] \setminus I}} [f''(\mathbf{x}')g''(\mathbf{y}')h''(\mathbf{z}')] \right| \geq \frac{\varepsilon}{2},$$

however by Lemma 1.10 the last expression tends to 0 as δ' goes to 0, and as $\delta' \ll \varepsilon$ we get a contradiction to the assumption 29. This shows that Lemma 7.1 follows from Lemma 1.10. \square

7.3 Proof of Lemma 7.7: set up and the compactness argument

Consider $\mathcal{F} \subseteq P(\Sigma_{\text{fin}})$ defined as follows: we have $F \in \mathcal{F}$ if there are functions $f: \Sigma_{\text{fin}} \rightarrow \mathbb{R}$, $g: \Gamma \rightarrow \mathbb{R}$, $h: \Phi \rightarrow \mathbb{R}$ not all constant such that $f(\vec{x}) = g(y)h(z)$ in the support of $\tilde{\nu}$, and $\text{supp}(f) = F$. Note that the collection \mathcal{F} is closed under intersection, since if we have $F_1 \in \mathcal{F}$ using the functions f_1, g_1, h_1 and $F_2 \in \mathcal{F}$ using the functions f_2, g_2, h_2 , then we have $F_1 \cap F_2 \in \mathcal{F}$ using the functions $f_1 f_2, g_1 g_2, h_1 h_2$.

Thus, we may consider minimal sets in \mathcal{F} , namely sets $\emptyset \neq F \in \mathcal{F}$ that do not strictly contain any set from \mathcal{F} . Let F_1, \dots, F_s be all minimal sets in \mathcal{F} . For $x \in \Sigma$, define $a(x) = (x, x, \dots, x) \in \tilde{\Sigma}$, and let $b(x) \in \Sigma_{\text{fin}}$ be the representative symbol from the connected component of $a(x)$, i.e. $\text{rep}(a(x))$.

Claim 7.12. For all $x \in \Sigma$, $b(x) \notin F_1 \cup \dots \cup F_s$.

Proof. Suppose towards contradiction that $b(x^*) \in F_1 \cup \dots \cup F_s$ for some $x^* \in \Sigma$, so that we may find functions f, g, h such that $f(\vec{x}) = g(y)h(z)$ on $\text{supp}(\tilde{\nu})$, and additionally $f(b(x^*)) \neq 0$. Note that since f, g, h cannot be all constant, it follows that either g or h must be non-constant.

Define $f': \Sigma \rightarrow \mathbb{R}$ by $f'(x) = f(b(x))$, and note that it follows that $f'(x) = g(y)h(z)$ on $\text{supp}(\mu)$. Also, $f'(x^*) \neq 0$, so f' is not identically 0. Note that since the support of (x, y) in $\text{supp}(\mu)$ is full, it follows that the function g can never vanish on Γ (since for all y , there is some z such that $(x^*, y, z) \in \text{supp}(\mu)$). This implies that $f'(x) \neq 0$ iff $h(z) \neq 0$ for every $(x, y, z) \in \text{supp}(\mu)$. We now consider two cases:

1. If the function f' vanishes sometimes, i.e. $\text{supp}(f') \subsetneq \Sigma$, then we define $\sigma(x) = 1_{f'(x)=0}$, $\phi(z) = 1_{h(z)=0}$, and note that in $\text{supp}(\mu)$ we have that $\sigma(x) + \phi(z) = 0 \pmod{2}$ for all $(x, y, z) \in \text{supp}(\mu)$, and that moreover this is a non-trivial embedding in $(\mathbb{F}_2, +)$. This is a contradiction to the fact $\text{supp}(\mu)$ is not linearly embeddable.

2. Else, $\text{supp}(f') = \Sigma$, and it follows that g and h also never vanish. If the sign of any of f' , g or h is not constant, then we get that $\sigma(x) = \text{sign}(f(x))$, $\gamma(y) = \text{sign}(g(y))$, $\phi(z) = \text{sign}(h(z))$ form a linear embedding of $\text{supp}(\mu)$ into $(\{-1, 1\}, \cdot)$ (which is isomorphic to $(\mathbb{F}_2, +)$). Else, we may assume that the signs of each one of f' , g and h is always the same, and by multiplying them by -1 if necessary we may assume that they are all positive. By multiplying them by a large enough constant, we may assume that their range is $(1, 2^M)$ for some absolute constant M , hence

$$\sigma(x) = \log(f(x)), \quad \gamma(y) = \log(g(y)), \quad \phi(z) = \log(h(z))$$

forms a non-trivial embedding of $\text{supp}(\mu)$ in $((0, M), (\text{mod } M))$, and by Claim 2.15 it follows that $\text{supp}(\mu)$ can be linearly embedded, in contradiction. \square

We take $\Sigma' = \{b(x) \mid x \in \Sigma\}$, which is disjoint from $F_1 \cup \dots \cup F_s$, and observe that if f, g, h satisfy that $f(x) = g(y)h(z)$, then $f|_{\Sigma'} \equiv 0$; indeed, otherwise the collection \mathcal{F} would contain $\text{supp}(f)$ which intersects Σ' , hence there would be some minimal set in \mathcal{F} intersecting Σ' , and contradiction to the choice of Σ' .

Claim 7.13. *The set $S = \{(\vec{x}, y, z) \in \text{supp}(\tilde{\nu}) \mid \vec{x} \in \Sigma'\}$ cannot be linearly embedded.*

Proof. Otherwise there would be an Abelian group $(G, +)$ and embeddings $\sigma: \Sigma' \rightarrow G$, $\gamma: \Gamma \rightarrow G$ and $\phi: \Phi \rightarrow G$ not all constant such that $\sigma(\vec{x}) + \gamma(y) + \phi(z) = 0$ in S . Note that γ, ϕ cannot both be constant, and that defining $\sigma'(x) = \sigma(b(x))$ we get that σ', γ, ϕ form a non-trivial embedding of $\text{supp}(\mu)$ to G , and contradiction. \square

We move on to prove the heart of the relaxed base case, asserting that if $\mathbb{E}_{x, x' \in R\Sigma'} [(f(x) - f(x'))^2] \geq \tau$, and f, g, h have 2-norm equal to 1, then $|\mathbb{E}_{(x, y, z) \sim \tilde{\nu}} [f(x)g(y)h(z)]|$ is bounded away from 1; moreover, the gap from 1 is at least polynomial in τ . We begin the proof of this assertion with the following claim, that employs a compactness argument and handles the case that $\tau > c$, for c which is an absolute constant depending on the alphabet sizes. After that, we will handle the case that τ is small by a closer inspection of the compactness argument.

Claim 7.14. *For all $\tau > 0$, there exists $\lambda \in (0, 1)$ such that if $\mathbb{E}_{x, x' \in R\Sigma'} [(f(x) - f(x'))^2] \geq \tau \|f\|_2^2$, then*

$$\left| \mathbb{E}_{(x, y, z) \sim \tilde{\nu}} [f(x)g(y)h(z)] \right| \leq \lambda \|f\|_2 \|g\|_2 \|h\|_2.$$

Proof. Assume that the statement is false. Thus, we may find sequences of functions f_ℓ, g_ℓ, h_ℓ of 2-norm 1 such that $\mathbb{E}_{x, x' \in R\Sigma'} [(f_\ell(x) - f_\ell(x'))^2] \geq \tau$ for all ℓ , and

$$\left| \mathbb{E}_{(x, y, z) \sim \tilde{\nu}} [f_\ell(x)g_\ell(y)h_\ell(z)] \right| \geq 1 - \frac{1}{\ell}.$$

Passing to limit, we find f, g, h of 2-norm 1 such that

$$\left| \mathbb{E}_{(x, y, z) \sim \tilde{\nu}} [f(x)g(y)h(z)] \right| \geq 1$$

and $\mathbb{E}_{x,x' \in R\Sigma'} [(f(x) - f(x'))^2] \geq \tau$, and we next show that $f(x) = g(y)h(z)$ (or $-f(x) = g(y)h(z)$) in $\text{supp}(\tilde{\nu})$, thereby reaching a contradiction (as $f|_{\Sigma'} \neq 0$). Indeed, by Cauchy-Schwarz we have

$$\left| \mathbb{E}_{(x,y,z) \sim \tilde{\nu}} [f(x)g(y)h(z)] \right|^2 \leq \mathbb{E}_{(x,y,z) \sim \tilde{\nu}} [f(x)^2] \mathbb{E}_{(x,y,z) \sim \tilde{\nu}} [g(y)^2h(z)^2] = \|f\|_2^2 \|g\|_2^2 \|h\|_2^2 = 1$$

where we used the fact that y, z are independent in $\tilde{\nu}$. Thus, we get that the above Cauchy-Schwarz inequality is tight, and so f is proportional to gh . Considering the 2-norms of these functions, we get that either $f(x) = g(y)h(z)$ or $-f(x) = g(y)h(z)$, say without loss of generality $f(x) = g(y)h(z)$. Note that as $\mathbb{E}_{x,x' \in R\Sigma'} [(f(x) - f(x'))^2] \geq \tau$, f is not constant, so by definition we get that $\text{supp}(f) \in \mathcal{F}$, and hence $f|_{\Sigma'} \equiv 0$, and contradiction to $\mathbb{E}_{x,x' \in R\Sigma'} [(f(x) - f(x'))^2] \geq \tau$. \square

7.4 Proof of Lemma 7.7: unraveling compactness

We now move on to handle the case τ is sufficiently small. We prove:

Lemma 7.15. *There is $c = c(m, \alpha) > 0$, such that for all $0 < \tau \leq c$, if $\mathbb{E}_{x,x' \in R\Sigma'} [(f(x) - f(x'))^2] \geq \tau \|f\|_2^2$, then*

$$\left| \mathbb{E}_{(x,y,z) \sim \tilde{\nu}} [f(x)g(y)h(z)] \right| \leq (1 - \tau^{50m}) \|f\|_2 \|g\|_2 \|h\|_2.$$

The rest of this section is devoted to the proof of Lemma 7.15.

Fix f, g and h as in the statement of the claim, and assume without loss of generality that their 2-norms are all 1. Let ν' be the distribution of $(x, y, z) \sim \tilde{\nu}$ conditioned on $x \in \Sigma'$. We write

$$(I) = \mathbb{E}_{(x,y,z) \sim \tilde{\nu}} [f(x)g(y)h(z)],$$

and assume towards contradiction that $(I) \geq 1 - \tau^{50m}$.

Claim 7.16. $\mathbb{E}_{x \in R\Sigma'} [f(x)^2] \gtrsim_m \tau$.

Proof. Follows as $\mathbb{E}_{x,x' \in R\Sigma'} [(f(x) - f(x'))^2] \geq \tau$. \square

Claim 7.17. *We have that $|f(x) - g(y)h(z)| \leq \tau^{20m}$ for all $(x, y, z) \in \text{supp}(\nu)$.*

Proof. Note that

$$(I) = 1 - \frac{1}{2} \mathbb{E}_{(x,y,z) \sim \tilde{\nu}} [(f(x) - g(y)h(z))^2].$$

Thus, if the conclusion of the claim fails, then $|(I)| \leq 1 - \Omega_{m,\alpha}(\tau^{40m}) < 1 - \tau^{50m}$, and contradiction. \square

Claim 7.18. *We have that $|g(y)| \geq \tau$ for all $y \in \Gamma$.*

Proof. Assume otherwise, i.e. that $|g(y^*)| \leq \tau$ for some y^* . Note that $\|h\|_\infty \leq \sqrt{m} \|h\|_2$ as the probability of each atom in h is at least $1/m$, so we get that $|g(y^*)h(z)| \leq \sqrt{m}\tau$ for all z . Note that the support of the distribution $\nu'_{x,z}$ is $\Sigma' \times \Gamma$; this is because it contains $\{(b(x), y) \mid (x, y) \in \text{supp}(\mu_{x,y})\}$ and $\mu_{x,y}$ is full, so this is the same as $\Sigma' \times \Gamma$. Thus, for every $x \in \Sigma'$ we may find $z \in \Phi$ such that $(x, y^*, z) \in \text{supp}(\tilde{\nu})$, hence by Claim 7.17

$$|f(x)| \leq |g(y^*)h(z)| + \tau^{20m} \leq O_m(\tau).$$

On the other hand, by Claim 7.16 we have $\mathbb{E}_{x \in \Sigma'} [|f(x)|^2] \gtrsim_{m, \alpha} \tau$, so we may find x such that the left hand side is at least $\Omega_{m, \alpha}(\sqrt{\tau})$. We thus get a contradiction to the fact that $\tau \leq c$, for small enough c depending only on m and α . \square

Consider the interval $(0, 1)$, and in it define the intervals $I_j = [\tau^{3(j+1)}, \tau^{3j}]$ for $j = 0, 1, \dots$. We say an interval I_j is free if it doesn't contain any point from either $\text{Image}(f)$ or $\text{Image}(h)$. Note that as the intervals I_j are disjoint and each one of these sets has size at most m , we may find $j \in \{0, 1, \dots, 2m\}$ such that I_j is free, and we fix such j henceforth.

Claim 7.19. *We have that $|f(x)| \geq \tau^{3j+1.5}$ for all $x \in \Sigma_{\text{fin}}$ and $|h(z)| \geq \tau^{3j+1.5}$ for all $z \in \Phi$.*

Proof. Assume otherwise, and define

$$\sigma(x) = 1_{|f(x)| \geq \tau^{3j+1.5}}, \quad \phi(z) = 1_{|h(z)| \geq \tau^{3j+1.5}}.$$

Then by our assumption, at least one of σ, ϕ are not identically 1, say σ without loss of generality. Note that since $\|f\|_{2, \tilde{\nu}_x} = 1$, there is some x such that $|f(x)| \geq 1 > \tau$, so σ is also not constantly 0, hence σ is not constant. We next show that $\sigma(x) + \phi(z) = 0 \pmod{2}$ for all $(x, y, z) \in \text{supp}(\tilde{\nu})$, hence we conclude that $\text{supp}(\nu)$ has a non-trivial Abelian embedding in $(\mathbb{F}_2, +)$, and contradiction to the fact $\tilde{\nu}$ cannot be linearly embedded.

The case that $\sigma(x) = 1$. Suppose that $(x, y, z) \in \text{supp}(\tilde{\nu})$ are such that $\sigma(x) = 1$. It follows from Claim 7.17 that

$$|g(y)h(z)| \geq |f(x)| - \tau^{10m} \geq \tau^{3j+1.5} - \tau^{10m},$$

and as $|g(y)| \leq \sqrt{m} \|g\|_{2, \tilde{\nu}_y} = \sqrt{m}$, it follows that

$$|h(z)| \geq \frac{\tau^{3j+1.5} - \tau^{10m}}{\sqrt{m}} \geq \tau^{3j+2},$$

where we used the fact that $\tau < c$ is small enough, and $j \leq 2m$. Thus, as I_j is free, it follows that $|h(z)| \geq \tau^{3j}$, and so $\phi(z) = 1$.

The case that $\sigma(x) = 0$. Suppose that $(x, y, z) \in \text{supp}(\tilde{\nu})$ are such that $\sigma(x) = 0$. It follows from Claim 7.17 that

$$|g(y)h(z)| \leq |f(x)| + \tau^{10m} \leq \tau^{3j+1.5} + \tau^{10m},$$

and as $|g(y)| \geq \tau$ by Claim 7.18, it follows that

$$|h(z)| \leq \frac{\tau^{3j+1.5} + \tau^{10m}}{\tau} \leq \tau^{3j+0.4},$$

where we used the fact that $\tau < c$ is small enough, and $j \leq 2m$. Thus, as I_j is free, it follows that $|h(z)| < \tau^{3(j+1)}$, and so $\phi(z) = 0$. \square

Claim 7.20. *For all $(x, y, z) \in \text{supp}(\tilde{\nu})$ we have that $\text{sign}(f(x)) = \text{sign}(g(y))\text{sign}(h(z))$.*

Proof. Assume otherwise; then we may find $(x, y, z) \in \text{supp}(\tilde{\nu})$ where these signs are different, and so

$$|f(x) - g(y)h(z)| \geq |f(x)| \geq \tau^{3j+1.5} > \tau^{10m},$$

where we used Claim 7.20. This contradicts Claim 7.17. \square

Claim 7.20 gives a candidate linear embedding of $\text{supp}(\tilde{\nu})$ to $(\{-1, 1\}, \cdot)$. We know though that there doesn't exist such embedding consisting of non-constant functions, and therefore the sign of each one of f, g, h is constant, and without loss of generality we assume they are all positive. Define

$$f'(x) = \log(f(x)), \quad g'(y) = \log(g(y)), \quad h'(z) = \log(h(z)).$$

Then by the Mean Value Theorem

$$|f'(x) - g'(y) - h'(z)| = \log(f(x)) - \log(g(y)h(z)) = \log'(t) |f(x) - g(y)h(z)|$$

for some t between $f(x)$ and $g(y)h(z)$, and in particular $t \geq \tau^{3j+1.5}$. Thus, by Claim 7.17 we get

$$|f'(x) - g'(y) - h'(z)| \leq \frac{1}{\tau^{3j+1.5}} \tau^{10m} \leq \tau^{2m} \quad (30)$$

for all $(x, y, z) \in \text{supp}(\tilde{\nu})$. Thus, f', g' and h' are approximate embeddings, and we next show how to extract proper embeddings from them in an Abelian group. First though, we argue that at least one of them is far from constant; indeed, as there are $x, x' \in \Sigma'$ such that $|f(x) - f(x')| \geq \sqrt{\tau}$ – say that $f(x) \geq f(x') + \sqrt{\tau}$, it follows that

$$f'(x) = f'(x') + \log\left(\frac{f(x)}{f(x')}\right) \geq f'(x') + \log\left(1 + \frac{\sqrt{\tau}}{f(x')}\right) \geq f'(x') + \log(1 + \sqrt{\alpha\tau}) \geq f'(x') + \frac{\sqrt{\alpha\tau}}{10},$$

where we used the fact that $0 < f(x') \leq \alpha^{-1/2} \|f\|_2 = \alpha^{-1/2}$.

We now extract the proper embeddings σ, γ, ϕ . Looking at $S = \text{Image}(f') \cup \text{Image}(g') \cup \text{Image}(h')$, we have that $|S| \leq 3m$. We take $N = (\alpha\tau)^{-9m}$; from Dirichlet's approximation theorem it follows that there are $\sigma: \text{Image}(f') \rightarrow \mathbb{Z}$, $\gamma: \text{Image}(g') \rightarrow \mathbb{Z}$ and $\phi: \text{Image}(h') \rightarrow \mathbb{Z}$ such that for some integer $1 \leq q \leq N$ we have

$$\left|f'(x) - \frac{\sigma(x)}{q}\right| \leq \frac{1}{qN^{1/|S|}}, \quad \left|g'(y) - \frac{\gamma(y)}{q}\right| \leq \frac{1}{qN^{1/|S|}}, \quad \left|h'(z) - \frac{\phi(z)}{q}\right| \leq \frac{1}{qN^{1/|S|}}.$$

Claim 7.21. σ is not constant.

Proof. Consider x, x' such that $f'(x) \geq f'(x') + \frac{\sqrt{\alpha\tau}}{10}$ as before. We get that

$$\left|\frac{\sigma(x)}{q} - \frac{\sigma(x')}{q}\right| \geq \frac{\sqrt{\alpha\tau}}{10} - \frac{1}{qN^{1/|S|}} \geq \frac{\sqrt{\alpha\tau}}{10} - (\alpha\tau)^3 > 0$$

$\sigma(x) \neq \sigma(x')$. □

Claim 7.22. $\sigma(x) - \gamma(y) - \phi(z) = 0$ for all $(x, y, z) \in \text{supp}(\tilde{\nu})$.

Proof. From the choice of σ, γ, ϕ and (30) it follows that for all $(x, y, z) \in \text{supp}(\tilde{\nu})$ it holds that

$$\left|\frac{\sigma(x) - \gamma(y) - \phi(z)}{q}\right| \leq \frac{3}{qN^{1/|S|}} + \tau^{20m},$$

so

$$|\sigma(x) - \gamma(y) - \phi(z)| \leq \frac{3}{N^{1/|S|}} + \tau^{20m}q \leq \frac{3}{N^{1/|S|}} + \tau^{20m}N < 1$$

from the choice of N . It follows that $\sigma(x) - \gamma(y) - \phi(z)$ is an integer smaller than 1 in absolute value, hence it is 0. □

Combining Claims 7.21, 7.22 we get an Abelian embedding of $\text{supp}(\tilde{\nu})$. Indeed, we can pick a large enough Q so that images of σ, γ, ϕ are all contained in the interval $[-Q, Q]$, and then consider them as maps from $\Sigma_{\text{fin}}, \Gamma, \Phi$ to $([0, 3Q], + \pmod{3Q})$. Then by Claim 7.21 we conclude that σ is not constant (as we get that there are x, x' such that $\sigma(x) \neq \sigma(x')$ as they are both in $[-Q, Q]$ so they are different mod $3Q$), and from Claim 7.22 we have that $\sigma(x) - \gamma(y) - \phi(z) = 0 \pmod{3Q}$ in $\text{supp}(\tilde{\nu})$. This is a contradiction to the fact $\tilde{\nu}$ has no linear embedding, and we are done.

8 Applications

In this section, we give a few applications of our main analytical lemma.

8.1 Hardness of approximation of CSP's

In this section we use our main analytical lemma to get optimal dictatorship tests with completeness 1 for a large class of 3-ary predicates.

Definition 8.1. *A dictatorship test for a predicate $P : \Sigma^k \rightarrow \{0, 1\}$ can query a function $f : \Sigma^n \rightarrow \Sigma$. The test picks a random $k \times n$ matrix by letting every column to be a random satisfying assignment to P (i.e., in $P^{-1}(1)$, with some fixed distribution μ on $P^{-1}(1)$) and letting $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k \in \Sigma^n$ be the rows of the matrix. The test accepts if $(f(\mathbf{x}_1), f(\mathbf{x}_2), \dots, f(\mathbf{x}_k))$ is also a satisfying assignment to P .*

We now describe the dictatorship test that was studied in [5]. The test is given in Figure 8.1. The starting point is an instance ϕ of P -CSP and let the value (i.e., maximum fraction of the constraints that can be satisfied by an assignment) of this instance be s . The distribution μ in the test depends on the SDP solution for ϕ and we only consider instances whose SDP value is 1.¹³ The SDP solution consists of vectors as well as local distribution for each constraint. Since the SDP value is 1, all these local distributions are supported on the satisfying assignments to P . Let μ_i be the local distribution corresponding to the i^{th} constraint of the instance. The test is as follows. Here $\varepsilon > 0$ is a small constant independent of n .

Let $P : \Sigma^k \rightarrow \{0, 1\}$ be the predicate. Given $f : \Sigma^n \rightarrow \Sigma$,

1. Select a constraint from ϕ according to the weights of the constraints. Let i be the selected constraint.
2. Construct a $k \times n$ matrix by setting each column of the matrix independently according to the following distribution: sample the column using μ_i .
3. Check if $P(f(\mathbf{x}_1), f(\mathbf{x}_2), \dots, f(\mathbf{x}_k)) = 1$.

Figure 1: Dictatorship test for the predicate P .

If f is a dictator function, then the test accepts with probability 1. This follows because for every i , the distribution μ_i is supported on the satisfying assignments to P and therefore every column of the matrix is from $P^{-1}(1)$. A challenging task is to compute the acceptance probability when f is far from dictator functions.

¹³We refer the readers to [21, 5] for detailed information on the semidefinite program, its value and the local distributions.

This test is exactly the same as the one given in [5]. If we use our main analytical lemma, Lemma 1.4, to analyze the above dictatorship test, then we have the following theorem on the soundness of the above test.

Theorem 8.2 (Restatement of Theorem 1.5). *Let $P: \Sigma^3 \rightarrow \{0, 1\}$ be any predicate that satisfies the following conditions. (1) P does not satisfy any linear embedding, and (2) there exists an instance of P -CSP that has a $(1, s)$ -integrality gap for the basic SDP relaxation and every local distribution is not linearly embeddable. Then for every $\varepsilon > 0$, there is a dictatorship test for P that has perfect completeness and soundness $s + \varepsilon$.*

The proof of this theorem is identical to the proof of [5, Theorem1.1]. The only difference is that in the proof of [5, Theorem1.1], Lemma 1.4 with the added condition that the distribution μ is semi-rich was used. As we get of the semi-richness condition in our main analytical lemma, we get the above improved theorem that applies for a rather large class of 3-ary predicates. As the proof is identical to the proof of [5, Theorem1.1], we skip the proof of Theorem 8.2 in this version.

8.2 Counting Lemmas

Theorem 8.3. *Suppose μ is a distribution over $\Sigma \times \Gamma \times \Phi$ such that $\text{supp}(\mu)$ cannot be linearly embedded. Then for all $\delta > 0$, there exist $d \in \mathbb{N}$, $\tau > 0$, $\varepsilon > 0$ and $N \in \mathbb{N}$ such that for $n \geq N$, if $f: \Sigma^n \rightarrow [0, 1]$, $g: \Gamma^n \rightarrow [0, 1]$, $h: \Phi^n \rightarrow [0, 1]$ are functions with average at least δ and $\max_i(I_i[f^{\leq d}], I_i[g^{\leq d}], I_i[h^{\leq d}]) \leq \tau$, then*

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \geq \varepsilon.$$

Proof. Let $0 \ll \tau \ll d^{-1} \ll \xi \ll \nu \ll \kappa \ll \eta \ll \varepsilon \ll \delta$; first, we argue that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] - \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [T_{1-\xi}f(\mathbf{x})T_{1-\xi}g(\mathbf{y})T_{1-\xi}h(\mathbf{z})] \right| \leq \eta.$$

Here, it is understood that the operator $T_{1-\xi}$ applied on each one of the functions refers to the standard noise operator with respect to the marginal distribution of μ on that coordinate. This is done by a hybrid argument, wherein we switch at each time a single function to a noisy version of it and bound the difference. For example, we argue that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [(I - T_{1-\xi})f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right| \leq \frac{\eta}{3}.$$

Indeed, note that

$$\text{Stab}_{1-\nu}((I - T_{1-\xi})f) = \|T_{1-\nu}(I - T_{1-\xi})f\|_2^2 \leq \max_j (1 - \nu)^j (1 - (1 - \xi)^j),$$

as these are the eigenvalues of $T_{1-\nu}(I - T_{1-\xi})$. As $\xi \ll \nu$, these eigenvalues smaller than κ , and the bound follows from Lemma 1.4.

Consider the distribution μ' defined as follows:

1. Sample $(x, y, z) \sim \mu$;
2. sample x' by taking $x' = x$ with probability $\sqrt{1 - \nu}$ and otherwise resample it according to μ_x ;

3. sample y' by taking $y' = y$ with probability $\sqrt{1-\nu}$ and otherwise resample it according to μ_y ;
4. sample z' by taking $z' = z$ with probability $\sqrt{1-\nu}$ and otherwise resample it according to μ_z ;
5. output (x', y', z') .

Note that

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [\mathbb{T}_{1-\xi} f(\mathbf{x}) \mathbb{T}_{1-\xi} g(\mathbf{y}) \mathbb{T}_{1-\xi} h(\mathbf{z})] = \mathbb{E}_{(\mathbf{x}', \mathbf{y}', \mathbf{z}') \sim \mu'^{\otimes n}} [\mathbb{T}_{\sqrt{1-\xi}} f(\mathbf{x}') \mathbb{T}_{\sqrt{1-\xi}} g(\mathbf{y}') \mathbb{T}_{\sqrt{1-\xi}} h(\mathbf{z}')].$$

Also note that the distribution μ' is connected and each atom has probability $\Omega_{\nu, \alpha}(1)$, and also that the individual influences are at most $\tau + (1-\xi)^d$. Hence by [17, Theorem 1.14] it follows that this expectation is at least ε , provided τ is small enough. \square

Using regularity lemma for low-degree influences, one may remove the assumption on influences in some cases.

Lemma 8.4. *For all $\alpha > 0$, $m \in \mathbb{N}$, if μ is a distribution over Σ in which each atom has probability at least α , $|\Sigma| \leq m$, then the following holds. For all $\varepsilon > 0$, $d \in \mathbb{N}$ and $\tau > 0$ there exists $D \in \mathbb{N}$ such for every $f: \Sigma^n \rightarrow [0, 1]$, there exists a decision tree \mathcal{T} of depth at most D such that sampling a root to path leaf in it $(\mathbf{I}, \mathbf{x}')$ yields*

$$\Pr_{(\mathbf{I}, \mathbf{x}')} \left[I_i^{\leq d} [f_{\mathbf{I} \rightarrow \mathbf{x}'}; \mu] \leq \tau \forall i \in [n] \setminus \mathbf{I} \right] \geq 1 - \varepsilon.$$

Proof. We omit the full details of the proof, as it is virtually identical to the proof of Jones' regularity lemma [14] (see also [9] for details). \square

Theorem 8.5 (Restatement of Theorem 1.6). *Suppose μ is a distribution over Σ^3 such that (1) the marginal distributions μ_x, μ_y, μ_z are identical, (2) $\{(x, x, x) \mid x \in \Sigma\} \subseteq \text{supp}(\mu)$, and (3) $\text{supp}(\mu)$ cannot be linearly embedded. Then for all $\delta > 0$, there exists $\varepsilon > 0$ and $N \in \mathbb{N}$ such that for $n \geq N$ and $S \subseteq \Sigma^n$ with $|S| \geq \delta |\Sigma|^n$,*

$$\Pr_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [\mathbf{x} \in S, \mathbf{y} \in S, \mathbf{z} \in S] \geq \varepsilon.$$

Proof. Let $f = 1_S$ and $0 \ll \varepsilon \ll D^{-1} \ll \tau \ll d^{-1} \ll \xi \ll \nu \ll \kappa \ll \eta \ll \delta$. By Lemma 8.4 we may find a decision tree \mathcal{T} of depth at most $D(d, \tau, \delta)$ such that sampling a path on it according μ_x , i.e. a subset I of at most D variables and $\mathbf{x}' \sim \mu_x^I$, we get that $I_i^{\leq d} [f_{\mathbf{I} \rightarrow \mathbf{x}'}] \leq \tau$ for all except with probability $\delta/100$. We denote the process that samples a path on it by $(\mathbf{I}, \mathbf{x}')$.

Note that by an averaging argument, $\mu(f_{\mathbf{I} \rightarrow \mathbf{x}'}) \geq \delta/2$ with probability at least $\delta/2$, hence we get that with probability at least $\delta/4$ we have that all influences are small and the average is at least $\delta/2$; we refer to this event by E . Thus, we get that

$$\begin{aligned} & \Pr_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [\mathbf{x} \in S, \mathbf{y} \in S, \mathbf{z} \in S] \\ & \geq \mathbb{E}_{(\mathbf{I}, \mathbf{x}')} \left[1_E \mathbb{E}_{(\mathbf{y}, \mathbf{z}) \in \mu_{\mathbf{y}, \mathbf{z}}^{\mathbf{I}}} \left[1_{\mathbf{y}=\mathbf{z}=\mathbf{x}'} \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{[n] \setminus \mathbf{I}}} [f_{\mathbf{I} \rightarrow \mathbf{x}'}(\mathbf{x}) f_{\mathbf{I} \rightarrow \mathbf{x}'}(\mathbf{y}) f_{\mathbf{I} \rightarrow \mathbf{x}'}(\mathbf{z})] \mid \mathbf{x}' \right] \right] \\ & \geq \frac{\delta}{4} \alpha^D \mathbb{E}_{(\mathbf{x}', \mathbf{I})} \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{[n] \setminus \mathbf{I}}} [f_{\mathbf{I} \rightarrow \mathbf{x}'}(\mathbf{x}) f_{\mathbf{I} \rightarrow \mathbf{x}'}(\mathbf{y}) f_{\mathbf{I} \rightarrow \mathbf{x}'}(\mathbf{z}) \mid E] \\ & \geq \varepsilon, \end{aligned}$$

where the last inequality is by Theorem 8.3. \square

For example, Theorem 1.6 may be applied to find progressions of the form $(\mathbf{x}, \mathbf{x} + \mathbf{a}, \mathbf{x} + \mathbf{a}^2)$ in dense subsets of \mathbb{F}_p^n ; we omit the details.

References

- [1] Michael Bateman and Nets Katz. New bounds on cap sets. *Journal of the American Mathematical Society*, 25(2):585–613, 2012.
- [2] Vitaly Bergelson and Terence Tao. Multiple recurrence in quasirandom groups. *Geometric and Functional Analysis*, 24(1):1–48, 2014.
- [3] Amey Bhangale, Prahladh Harsha, and Sourya Roy. Mixing of 3-Term Progressions in Quasirandom Groups. In *13th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 215, pages 20:1–20:9, Dagstuhl, Germany, 2022.
- [4] Amey Bhangale and Subhash Khot. Optimal inapproximability of satisfiable k-lin over non-abelian groups. In *Proceedings of the 53rd annual Symposium on Theory of Computing (STOC)*, pages 1615–1628, New York, NY, USA, 2021. ACM.
- [5] Amey Bhangale, Subhash Khot, and Dor Minzer. On approximability of satisfiable k-csp: I. In *Proceedings of the 54th annual Symposium on Theory of Computing (STOC)*, pages 976–988, New York, NY, USA, 2022. ACM.
- [6] Mark Braverman, Subhash Khot, Noam Lifshitz, and Dor Minzer. An invariance principle for the multi-slice, with applications. In *Proceedings of the 62nd annual symposium on Foundations of Computer Science (FOCS)*, pages 228–236. IEEE, 2021.
- [7] Mark Braverman, Subhash Khot, and Dor Minzer. On Rich 2-to-1 Games. In *12th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 185, pages 27:1–27:20, Dagstuhl, Germany, 2021.
- [8] T.C Brown and J.P Buhler. A density version of a geometric ramsey theorem. *Journal of Combinatorial Theory, Series A*, 32(1):20–34, 1982.
- [9] Gilad Chase, Yuval Filmus, Dor Minzer, Elchanan Mossel, and Nitin Saurabh. Approximate polymorphisms. In *Proceedings of the 54th annual Symposium on Theory of Computing (STOC)*, pages 195–202, New York, NY, USA, 2022. ACM.
- [10] Ernie Croot, Vsevolod F. Lev, and Péter Pál Pach. Progression-free sets in F_4^n are exponentially small. *Annals of Mathematics*, 185(1):331–337, 2017.
- [11] Jordan S Ellenberg and Dion Gijswijt. On large subsets of with no three-term arithmetic progression. *Annals of Mathematics*, pages 339–343, 2017.
- [12] W.T. Gowers. Quasirandom groups. *Combinatorics, Probability and Computing*, 17(3):363–387, 2008.
- [13] Jan Hązła, Thomas Holenstein, and Elchanan Mossel. Product space models of correlation: Between noise stability and additive combinatorics. *Discrete Analysis*, page 6513, 2018.

- [14] Chris Jones. A noisy-influence regularity lemma for boolean functions. *CoRR*, abs/1610.06950, 2016.
- [15] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the 34th annual Symposium on Theory of Computing (STOC)*, pages 767–775. ACM, 2002.
- [16] Roy Meshulam. On subsets of finite abelian groups with no 3-term arithmetic progressions. *Journal of Combinatorial Theory, Series A*, 71(1):168–172, 1995.
- [17] Elchanan Mossel. Gaussian bounds for noise correlation of functions. *Geometric and Functional Analysis*, 19(6):1713–1756, 2010.
- [18] Elchanan Mossel, Ryan O’Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. In *Proceedings of the 46th annual symposium on Foundations of Computer Science (FOCS)*, pages 21–30. IEEE, 2005.
- [19] Ryan O’Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.
- [20] Sarah Peluse. Mixing for three-term progressions in finite simple groups. *Mathematical Proceedings of the Cambridge Philosophical Society*, 165(2):279–286, 2018.
- [21] Prasad Raghavendra. Optimal algorithms and inapproximability results for every csp? In *Proceedings of the 14th annual Symposium on Theory of Computing (STOC)*, pages 245–254, 2008.
- [22] Prasad Raghavendra. *Approximating NP-hard problems efficient algorithms and their limits*. University of Washington, 2009.
- [23] K. F. Roth. On certain sets of integers. *Journal of the London Mathematical Society*, s1-28(1):104–109, 1953.
- [24] Endre Szemerédi. On sets of integers containing k elements in arithmetic progression. *Acta Arithmetica*, 27(1):199–245, 1975.
- [25] Terence Tao. Mixing for progressions in nonabelian groups. In *Forum of Mathematics, Sigma*, volume 1, page e2. Cambridge University Press, 2013.

A Missing proofs from Section 3

In this section, we prove the implications from (8).

A.1 Proof that Lemma 3.2 implies Lemma 1.10

We use

$$0 \leq \delta \ll \xi \ll \varepsilon \ll c, \quad M^{-1} \ll \alpha, \quad m^{-1} \leq 1.$$

Note that

$$\begin{aligned} \left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})g(\mathbf{y})h(\mathbf{z})] \right| &\leq \left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})T_{1-\xi}g(\mathbf{y})h(\mathbf{z})] \right| \\ &+ \sum_{j=0}^{\infty} \left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x})(T_{1-2^{-j-1}\xi} - T_{1-2^{-j}\xi})g(\mathbf{y})h(\mathbf{z})] \right|. \end{aligned}$$

For the first term, we have

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x}) \mathbb{T}_{1-\xi} g(\mathbf{y}) h(\mathbf{z})] \right| \leq \|\mathbb{T}_{1-\xi} g\|_2 = \sqrt{\text{Stab}_{(1-\xi)^2}(g)} \leq \sqrt{\delta} \leq \frac{\varepsilon}{10},$$

so we focus on bounding the second term; fix some j , and denote $\xi' = 2^{-j}\xi$, $\mathbb{T}' = \mathbb{T}_{1-\xi'/2} - \mathbb{T}_{1-\xi'}$, and $\xi'' = M\xi'/\log(\xi')^3$. Clearly,

$$\begin{aligned} & \left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f(\mathbf{x}) \mathbb{T}' g(\mathbf{y}) h(\mathbf{z})] \right| \\ & \leq \underbrace{\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [(I - \mathbb{T}_{1-\xi''}) f(\mathbf{x}) \mathbb{T}' g(\mathbf{y}) h(\mathbf{z})] \right|}_{(I)} + \underbrace{\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [\mathbb{T}_{1-\xi''} f(\mathbf{x}) \mathbb{T}' g(\mathbf{y}) h(\mathbf{z})] \right|}_{(II)}. \end{aligned}$$

Upper bounding (II). From Lemma 3.2 we get that

$$(II) \lesssim_{M, m, \alpha} \frac{1}{\log(1/\xi')^6}.$$

Upper bounding (I). Clearly, we may bound it by

$$\begin{aligned} & \left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [(I - \mathbb{T}_{1-\xi''}) f(\mathbf{x}) \mathbb{T}_{1-\xi'} g(\mathbf{y}) h(\mathbf{z})] \right| \\ & + \left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [(I - \mathbb{T}_{1-\xi''}) f(\mathbf{x}) \mathbb{T}_{1-\xi'/2} g(\mathbf{y}) h(\mathbf{z})] \right|, \end{aligned}$$

and as the bound for each one of these is similar, we only explain the upper bound for the first. We can write

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [(I - \mathbb{T}_{1-\xi''}) f(\mathbf{x}) \mathbb{T}_{1-\xi'} g(\mathbf{y}) h(\mathbf{z})] = \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu'^{\otimes n}} [(I - \mathbb{T}_{1-\xi''}) f(\mathbf{x}) g(\mathbf{y}) h(\mathbf{z})],$$

where a sample in the distribution μ' is drawn by first sampling $(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}$, and then re-sampling each coordinate of \mathbf{y} independently with probability ξ' . Write $f = f^{\leq d} + f^{> d}$ where $d = \sqrt{M} \log(1/\xi')/\xi'$.

Claim A.1. $\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu'^{\otimes n}} [(I - \mathbb{T}_{1-\xi''})(f^{\leq d})(\mathbf{x}) g(\mathbf{y}) h(\mathbf{z})] \right| \leq \frac{M^{3/2}}{\log(1/\xi')^2}.$

Proof. By boundedness we have

$$\begin{aligned} & \left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu'^{\otimes n}} [(I - \mathbb{T}_{1-\xi''})(f^{\leq d})(\mathbf{x}) g(\mathbf{y}) h(\mathbf{z})] \right| \leq \|(I - \mathbb{T}_{1-\xi''})(f^{\leq d})\|_1 \\ & \leq \|(I - \mathbb{T}_{1-\xi''})(f^{\leq d})\|_2 \\ & \leq (1 - (1 - \xi'')^d) \\ & \leq d\xi'' \\ & \leq \frac{M^{3/2}}{\log(1/\xi')^2}. \end{aligned}$$

□

Claim A.2. $\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu'^{\otimes n}} \left[(I - T_{1-\xi''})(f^{>d})(\mathbf{x})g(\mathbf{y})h(\mathbf{z}) \right] \right| \leq \xi'^{100}$

Proof. The left hand side is at most

$$\mathbb{E}_{\mathbf{y} \sim \mu'_y} \left[\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}', \mathbf{z}) \sim \mu'^{\otimes n}} \left[(I - T_{1-\xi''})f^{>d}(\mathbf{x})h(\mathbf{z}) \mid \mathbf{y}' = \mathbf{y} \right] \right| \right]$$

Note that for each $i \in [n]$, the marginal distribution of $\mu' \mid \mathbf{y}'_i = y_i$ on x_i, z_i is $\xi' \mu_{x,z} + (1 - \xi') \mu_{x,z} \mid y_i$. We consider the following alternative way of sampling from μ' ; first, choose $I \subseteq [n]$ randomly by including each $i \in [n]$ in it with probability ξ' . On $i \in I$, sample $(x_i, z_i) \sim \mu_{x,z}$ and $y_i \sim \mu_y$ independently. On $i \notin I$, sample $(x_i, y_i, z_i) \sim \mu$. We thus have that the above expression can be written as

$$\begin{aligned} & \mathbb{E}_{\mathbf{y} \sim \mu'_y} \left[\left| \mathbb{E}_{I, (\mathbf{x}_{\bar{I}}, \mathbf{y}'_{\bar{I}}, \mathbf{z}_{\bar{I}})} \left[\mathbb{E}_{I, \mathbf{x}_I, \mathbf{z}_I} \left[(I - T_{1-\xi''})f^{>d}(\mathbf{x})h(\mathbf{z}) \mid \mathbf{y}_I = \mathbf{y}'_I \right] \right] \right| \right] \\ & \leq \mathbb{E}_{I, \mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}} \left[\left| \mathbb{E}_{\mathbf{x}_I, \mathbf{z}_I} \left[(I - T_{1-\xi''})f^{>d}(\mathbf{x})h(\mathbf{z}) \right] \right| \right] \\ & \leq \mathbb{E}_{I, \mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}} \left[\left| \mathbb{E}_{\mathbf{x}_I, \mathbf{z}_I} \left[f^{>d}(\mathbf{x})h(\mathbf{z}) \right] \right| \right] + \mathbb{E}_{I, \mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}} \left[\left| \mathbb{E}_{\mathbf{x}_I, \mathbf{z}_I} \left[T_{1-\xi''}(f^{>d})(\mathbf{x})h(\mathbf{z}) \right] \right| \right] \\ & = \mathbb{E}_{I, \mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}} \left[\left| \mathbb{E}_{\mathbf{x}_I, \mathbf{z}_I} \left[f^{>d}(\mathbf{x})h(\mathbf{z}) \right] \right| \right] + \mathbb{E}_{I, \mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}} \left[\left| \mathbb{E}_{\mathbf{x}_I, \mathbf{z}_I} \left[(T_{1-\xi''}f^{>d})(\mathbf{x})h(\mathbf{z}) \right] \right| \right]. \end{aligned}$$

We show that if f' is a function of 2-norm at most 1 such that $(f')^{\leq d} = 0$, then

$$\mathbb{E}_{I, \mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}} \left[\left| \mathbb{E}_{\mathbf{x}_I, \mathbf{z}_I} \left[f'(\mathbf{x})h(\mathbf{z}) \right] \right| \right] \leq \frac{1}{2} \xi'^{100}.$$

Since both $f^{>d}$ and $T_{1-\xi''}f^{>d}$ are such functions, this would yield the statement of the claim. Indeed, if f' are such function, then considering the operator $S_I: L_2(\Sigma^I, \mu_x^I) \rightarrow L_2(\Phi^I, \mu_z^I)$ we have

$$\mathbb{E}_{I, \mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}} \left[\left| \mathbb{E}_{\mathbf{x}_I, \mathbf{z}_I} \left[f'(\mathbf{x})h(\mathbf{z}) \right] \right| \right] = \mathbb{E}_{I, \mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}} \left[\left| \langle S_I(f'_{x_{\bar{I}} \rightarrow \mathbf{x}_{\bar{I}}}), h_{z_{\bar{I}} \rightarrow \mathbf{z}_{\bar{I}}} \rangle \right| \right] \leq \mathbb{E}_{I, \mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}} \left[\|S_I^* S_I(f'_{x_{\bar{I}} \rightarrow \mathbf{x}_{\bar{I}}})\|_2 \right].$$

By Jensen's inequality, we have that this is at most

$$\begin{aligned} \mathbb{E}_{I, \mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}} \left[\|S_I^* S_I(f'_{x_{\bar{I}} \rightarrow \mathbf{x}_{\bar{I}}})\|_2^2 \right]^{1/2} & \leq \mathbb{E}_{I, \mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}} \left[\left(W_{\leq d'}(f'_{x_{\bar{I}} \rightarrow \mathbf{x}_{\bar{I}}}) + (1 - c(m, \alpha))^{2d'} \right)^{1/2} \right] \\ & \leq (1 - c)^{d'} + \sqrt{\mathbb{E}_{I, \mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}} \left[W_{\leq d'}(f'_{x_{\bar{I}} \rightarrow \mathbf{x}_{\bar{I}}}) \right]}, \end{aligned}$$

where $d' = \xi' d/2$. In the second inequality, we used the fact that the operator $S_I^* S_I$ is connected and the probability of each transition is at least $\Omega_{\alpha, m}(1)$, so $\|S_I^* S_I f''\|_2 \leq (1 - c)^{d'} \|f''\|_2$ if $(f'')^{\leq d'} \equiv 0$. Note that

$$\mathbb{E}_{I, \mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}} \left[W_{\leq d'}(f'_{x_{\bar{I}} \rightarrow \mathbf{x}_{\bar{I}}}) \right] \leq \Pr_I \left[|[d] \cap I| \leq d' \right] \leq 2^{-\Omega(d')}$$

by Chernoff's bound. Overall, we get that

$$\mathbb{E}_{I, \mathbf{x}_{\bar{I}}, \mathbf{y}_{\bar{I}}, \mathbf{z}_{\bar{I}}} \left[\left| \mathbb{E}_{\mathbf{x}_I, \mathbf{z}_I} \left[f'(\mathbf{x})h(\mathbf{z}) \right] \right| \right] \leq 2^{-\Omega_{m, \alpha}(d')} + 2^{-\Omega(d')} \leq 2^{-\Omega_{m, \alpha}(\sqrt{M} \log(1/\xi'))} \leq \frac{1}{2} \xi'^{100}$$

□

We thus get that

$$(I) + (II) \lesssim_{M,m,\alpha} \frac{1}{\log(1/\xi')^2} + \xi'^{100} + \frac{1}{\log(1/\xi')^6} \lesssim \frac{1}{j^2 + \log(1/\xi)^2} + 2^{-j/2} \sqrt{\xi},$$

so

$$\sum_j (I) + (II) \lesssim_{M,m,\alpha} \sqrt{\xi} + \sum_{j=\log(1/\xi)}^{\infty} \frac{1}{j^2} \lesssim \sqrt{\xi} + \frac{1}{\log(1/\xi)},$$

hence $\sum_j (I) + (II) \leq \frac{\varepsilon}{2}$ provided ξ is small enough. This finishes the deduction of Lemma 1.10 from Lemma 3.2. \square

A.2 Softly truncating the effective degree of f from below: Lemma 3.3 implies Lemma 3.2

We choose

$$0 \leq \xi_0 \ll c, \quad M^{-1} \ll \alpha, \quad m^{-1} \leq 1,$$

and show that Lemma 3.3 implies Lemma 3.2. Letting $f' = T_{1-M\xi/\log(1/\xi)^3} f$, $g' = (T_{1-\xi/2} - T_{1-\xi})g$ and $\xi' = M\xi \log(1/\xi)^{100}$, we write

$$\begin{aligned} \left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f'(\mathbf{x})g'(\mathbf{y})h(\mathbf{z})] \right| &\leq \underbrace{\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [(I - E_{\xi'})f'(\mathbf{x})g'(\mathbf{y})h(\mathbf{z})] \right|}_{(I)} \\ &\quad + \underbrace{\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [E_{\xi'}f'(\mathbf{x})g'(\mathbf{y})h(\mathbf{z})] \right|}_{(II)}. \end{aligned}$$

Using Lemma 3.3, we have that $(I) \lesssim_{m,\alpha,M} \frac{1}{\log^6(1/\xi)}$, and we next bound (II) . We may write

$$(II) = \left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu'^{\otimes n}} [f'(\mathbf{x})g'(\mathbf{y})h(\mathbf{z})] \right|,$$

where the distribution μ' is defined by first sampling $(x, y, z) \sim \mu^{\otimes n}$, and then for each i such that $x_i \in \Sigma'$, with probability ξ' re-sampling it according to $\mu_x \mid \Sigma'$. We consider the following alternative way of sampling according to μ' . Let $p = \Pr_{x \sim \mu_x} [x \in \Sigma']$, and note that $p \geq \alpha$. We sample $R \subseteq [n]$ by independently including each $i \in [n]$ in R with probability p , and then sub-sample $I \subseteq R$ by including each $i \in R$ in I with probability ξ' . We then sample $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ as: for $i \in I$, we sample $\mathbf{x}_i \sim \mu_x \mid \Sigma'$, $y_i, z_i \sim \mu_{y,z} \mid x \in \Sigma'$ independently; for $i \in R \setminus I$ we sample $(x_i, y_i, z_i) \sim \mu \mid x_i \in \Sigma'$; for $i \notin R$, we sample $(x_i, y_i, z_i) \sim \mu \mid x_i \notin \Sigma'$. Thus, denoting $\nu = \mu_{y,z} \mid x \in \Sigma'$, we have

$$\begin{aligned} (II) &\leq \mathbb{E}_{I,R,\mathbf{x},\mathbf{y}_{\bar{I}},\mathbf{z}_{\bar{I}}} \left[|f'(\mathbf{x})| \left| \mathbb{E}_{(\mathbf{y}_I, \mathbf{z}_I) \sim \nu^I} [g'(\mathbf{y})h(\mathbf{z})] \right| \right] \leq \sqrt{\mathbb{E}_{\substack{I,R,\mathbf{x},\mathbf{y}_{\bar{I}},\mathbf{z}_{\bar{I}}} \\ (\mathbf{y}_I, \mathbf{z}_I) \sim \nu^I, (\mathbf{y}'_I, \mathbf{z}'_I) \sim \nu^I}} [g'(\mathbf{y})h(\mathbf{z})g'(\mathbf{y}')h(\mathbf{z}')] } \\ &= \sqrt{\mathbb{E}_{\substack{I,R,\mathbf{x},\mathbf{y}_{\bar{I}},\mathbf{z}_{\bar{I}}} \\ (\mathbf{y}_I, \mathbf{z}_I) \sim \nu^I, (\mathbf{y}'_I, \mathbf{z}'_I) \sim \nu^I}} [g'(\mathbf{y})L(\mathbf{y}', \mathbf{z}, \mathbf{z}')]} \end{aligned}$$

Let

$$\tilde{g} = (g')_{y_{\bar{I}} \rightarrow y_{\bar{I}}}, \quad \tilde{L} = L_{y'_{\bar{I}}, z'_{\bar{I}}, z'_{\bar{I}} \rightarrow y_{\bar{I}}, z_{\bar{I}}, z'_{\bar{I}}},$$

so that we may write the last expression as $\sqrt{\mathbb{E}_{I, R, y_{\bar{I}}, z_{\bar{I}}} [\langle S_I \tilde{g}, \tilde{L} \rangle]}$, where $S_I: L_2(\Gamma^n, \nu_y) \rightarrow L_2(\Gamma^n, \nu_z)$ is defined as

$$S_I \tilde{g}(z) = \mathbb{E}_{(\mathbf{y}, \mathbf{z}) \sim \nu} [\tilde{g}(y) \mid \mathbf{z} = z].$$

We may bound

$$\left| \langle S_I \tilde{g}, \tilde{L} \rangle \right| \leq \|S_I \tilde{g}\|_2 \leq \|S_I^* S_I \tilde{g}\|_2^{1/2} \|\tilde{g}\|_2^{1/2} \lesssim \|S_I^* S_I \tilde{g}\|_2^{1/2},$$

so combining everything so far and using Jensen, we get that

$$(II) \lesssim \mathbb{E}_{I, y_{\bar{I}}} \left[\|S_I^* S_I \tilde{g}\|_{2, \nu_y}^2 \right]^{1/8}.$$

Next, note that as $S_I^* S_I$ is connected and the probability of each atom is at least $\Omega_{\alpha, m}(1)$, we get that choosing $d = M \log(1/\xi)$ we have from Lemma 2.9 that

$$\|S_I^* S_I \tilde{g}\|_{2, \nu_y}^2 \leq W_{\leq d}[\tilde{g}; \nu_y] + (1-c)^d.$$

We next bound the expectation of the first term on the right hand side using Lemma 2.14, and argue that it follows from the lemma that

$$\mathbb{E}_{I, y_{\bar{I}}} [W_{\leq d}[\tilde{g}; \nu_y]] \lesssim \mathbb{E}_{I, y_{\bar{I}}} [\text{Stab}_{1-1/d}(\tilde{g}; \nu_y)] \lesssim \text{Stab}_{1-c\xi'/d}(g'; \mu_y),$$

for some $c = c(m, \alpha) > 0$. Indeed, in the setting of the lemma we have $1-s = \xi'$; the distribution ν_2 is ν_y , and the distribution ν_1 is $\frac{p'}{1-\xi'} \nu'_1 + \frac{p''}{1-\xi'} \nu''_1$ where $p' = 1-p$ and $p'' = 1-\xi' - p' = p - \xi'$, ν'_1 is the distribution of \mathbf{y} where $(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu \mid \mathbf{x} \notin \Sigma'$, and ν''_1 is the distribution of \mathbf{y} where $(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu \mid \mathbf{x} \in \Sigma'$. It follows that

$$\begin{aligned} \mathbb{E}_{I, y_{\bar{I}}} [W_{\leq d}[\tilde{g}; \nu_y]] &\leq \sum_S \left(1 - \frac{c\xi'}{d}\right)^{|S|} \|(g')^{=S}\|_2^2 \leq \sum_S \left(1 - \frac{c\xi'}{d}\right)^{|S|} \left(\left(1 - \frac{\xi}{2}\right)^{|S|} - (1-\xi)^{|S|} \right) \|g^{=S}\|_2^2 \\ &\leq \max_{k \in \mathbb{N}} \left(1 - \frac{c\xi'}{d}\right)^k \left(\left(1 - \frac{\xi}{2}\right)^k - (1-\xi)^k \right). \end{aligned}$$

If $k \leq \frac{1}{\xi \log(1/\xi)^{50}}$, the second factor is at most $\lesssim \xi k \lesssim \frac{1}{\log(1/\xi)^{50}}$. If $k > \frac{1}{\xi \log(1/\xi)^{50}}$, the first factor is at most

$$2^{-\Omega_{m, \alpha}(\frac{\xi' k}{d})} \leq 2^{-\Omega_{M, m, \alpha}(\log(1/\xi)^{50})} \leq \xi.$$

It follows that

$$(II) \lesssim \mathbb{E}_{I, y_{\bar{I}}} \left[\|S_I^* S_I \tilde{g}\|_{2, \nu_y}^2 \right]^{1/8} \lesssim (1-c)^d + \xi^{1/8} + \frac{1}{\log(1/\xi)^6} \lesssim \frac{1}{\log(1/\xi)^6}.$$

□

A.3 Getting the functions to be homogenous: Lemma 3.1 implies Lemma 3.3

In this section, we prove that Lemma 3.1 implies Lemma 3.3. We begin by stating a few basic properties of the operators T and E .

Claim A.3. *Let $\chi \in B_1 \cup B_2$. Then*

1. *If $\chi \neq \chi_{\text{const}}$, then $T_{1-\xi'}\chi = (1 - \xi')\chi$, and if $\chi = \chi_{\text{const}}$, then $T_{1-\xi'}\chi = \chi$.*
2. *If $\chi \in B_1$ then $E_{1-\xi'}\chi = \chi$.*
3. *If $\chi \in B_2$, then $E_{1-\xi'}\chi = (1 - \xi')\chi$.*

Proof. For the first item, if $\chi \neq \chi_{\text{const}}$, then the average of it is 0, hence

$$T_{1-\xi'}\chi(x) = (1 - \xi')\chi(x) + \xi' \cdot 0 = (1 - \xi')\chi(x).$$

The second part of the first item is clear. The second item is also clear, since $E_{1-\xi'}$ may only change x if $x \in \Sigma'$, and $\chi \in B_1$ gets the same value on all elements of Σ' . For the third item, note that if $\chi \in B_2$, then $\mathbb{E}_{x \sim \mu_x} [1_{x \in \Sigma'} \chi(x)] = 0$, so

$$E_{1-\xi'}\chi(x) = (1 - \xi')\chi(x) + \xi' \frac{\mathbb{E}_{x \sim \mu_x} [1_{x \in \Sigma'} \chi(x)]}{\mathbb{E}_{x \sim \mu_x} [1_{x \in \Sigma'}]} = (1 - \xi')\chi(x). \quad \square$$

Claim A.4. *Let $f: \Sigma^n \rightarrow \mathbb{R}$.*

1. *If f is homogenous of degree d , then $T_{1-\xi'}f = (1 - \xi')^d f$.*
2. *If f is effectively homogenous of degree d , then $E_{1-\xi'}f = (1 - \xi')^d f$.*

Proof. Both items are immediate from Claim A.3 by writing f as linear combination of monomials of degree d , or effectively degree d in the case of the second item. \square

Proof that Lemma 3.1 implies Lemma 3.3. Towards this end, fix f, g, h and α, m, M, ξ_0 and ξ as in Lemma 3.1; we shall assume that f, g and h have 2-norms equal to 1. Denote

$$f' = (I - E_{1-M\xi \log(1/\xi)^{100}})T_{1-M\xi/\log(1/\xi)^3}f, \quad g' = (T_{1-\xi/2} - T_{1-\xi})g.$$

We write $f' = f'_1 + f'_2 + f'_3$ where f'_1 is the part of f' of effective degree of at most $d_1 = \frac{1}{\xi \log(1/\xi)^{200}}$, f'_3 is the part of f' of degree at least $d_2 = \frac{\log(1/\xi)^6}{\xi}$ and effective degree more than d_1 , and f'_2 is the part of f' of effective degree more than d_1 and degree less than d_2 . We write

$$\begin{aligned} \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f'(\mathbf{x})g'(\mathbf{y})h(\mathbf{z})] &= \underbrace{\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f'_1(\mathbf{x})g'(\mathbf{y})h(\mathbf{z})]}_{(I)} + \underbrace{\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f'_2(\mathbf{x})g'(\mathbf{y})h(\mathbf{z})]}_{(II)} \\ &\quad + \underbrace{\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f'_3(\mathbf{x})g'(\mathbf{y})h(\mathbf{z})]}_{(III)}, \end{aligned}$$

and bound each one of these separately. First note that

$$|(I)|^2 \leq \|f'_1\|_2^2 \|g\|_2^2 \|h\|_2^2 \leq \|f'_1\|_2^2 = \sum_{\chi: \text{eff-deg}(\chi) \leq d_1} \widehat{f}_1(\chi)^2.$$

Using Claim A.4, we get that the last sum is equal to

$$\begin{aligned} & \sum_{\chi: \text{eff-deg}(\chi) \leq d_1} \left(1 - (1 - M\xi \log(1/\xi)^{100})^{\text{eff-deg}(\chi)}\right) \left(1 - \frac{M\xi}{\log(1/\xi)^3}\right)^{\text{deg}(\chi)} \widehat{f}(\chi)^2 \\ & \leq \left(1 - (1 - M\xi \log(1/\xi)^{100})^{d_1}\right) \sum_{\chi} \widehat{f}(\chi)^2, \end{aligned}$$

which is at most $M\xi \log(1/\xi)^{100} d_1 \leq \frac{1}{\log(1/\xi)^{50}}$. Similarly, we have that

$$|(III)|^2 \leq \|f'_3\|_2^2 \|g\|_2^2 \|h\|_2^2 \leq \|f'_3\|_2^2 = \sum_{\chi: \text{deg}(\chi) \geq d_2} \widehat{f}'_3(\chi)^2.$$

Using Claim A.4, we get that the last sum is equal to

$$\begin{aligned} & \sum_{\chi: \text{deg}(\chi) \geq d_2} \left(1 - (1 - M\xi \log(1/\xi)^{100})^{\text{eff-deg}(\chi)}\right) \left(1 - \frac{M\xi}{\log(1/\xi)^3}\right)^{\text{deg}(\chi)} \widehat{f}(\chi)^2 \\ & \leq \left(1 - \frac{M\xi}{\log(1/\xi)^3}\right)^{d_2} \sum_{\chi} \widehat{f}(\chi)^2, \end{aligned}$$

which is at most $2^{-\frac{M\xi}{\log(1/\xi)^3} d_2} \leq \frac{1}{\log(1/\xi)^{50}}$.

And we next bound (II). Let g'_1 be the part of g' of degree at most $d'_1 = \frac{1}{\xi \log(1/\xi)^{10}}$, g'_3 be the part of g' of degree at least $d'_2 = \frac{\log(1/\xi)^{10}}{\xi}$, and g'_2 be the part of g' of degree more than d'_1 and less than d'_2 . We write

$$(II) = \underbrace{\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f'_2(\mathbf{x}) g'_1(\mathbf{y}) h(\mathbf{z})]}_{(IV)} + \underbrace{\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f'_2(\mathbf{x}) g'_2(\mathbf{y}) h(\mathbf{z})]}_{(V)} + \underbrace{\mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f'_2(\mathbf{x}) g'_3(\mathbf{y}) h(\mathbf{z})]}_{(VI)}$$

Similar arguments to before show that $|(IV)|, |(VI)| \leq \frac{1}{\log(1/\xi)^8}$, and we next bound (V). Write $f'_2 = \sum_{i,j=d_1}^{d_2} f'_{i,j}$ where $f'_{i,j}$ is the part of f' homogenous of degree i and effectively homogenous j , and $g'_2 = \sum_{i'=d'_1}^{d'_2} g'_{2,i'}$ where $g'_{2,i'}$ is the part of g' homogenous of degree i' , and $h = \sum_{d=0}^n h^d$ where h^d is the homogenous degree d part of h . We have

$$|(V)| \leq \sum_{i,j,i',d} \left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f'_{i,j}(\mathbf{x}) g'_{2,i'}(\mathbf{y}) h^d(\mathbf{z})] \right|.$$

We note that if $d > i + i'$, then the expectation is 0. For other d , we have by Lemma 3.1 that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f'_{i,j}(\mathbf{x}) g'_{2,i'}(\mathbf{y}) h^d(\mathbf{z})] \right| \leq (1 - \delta)^{\frac{i}{\log^C(i)}}$$

for some $C = C(m, \alpha) > 0$ and $\delta = \delta(m, \alpha) > 0$. Thus, the above sum is at most

$$(d'_2)^4 (1 - \delta)^{\frac{1}{\xi \log^C(1/\xi)}} \leq \frac{1}{\log(1/\xi)^6}$$

where the last inequality holds as $0 < \xi \leq \xi_0$ and ξ_0 is sufficiently small. We conclude that $|(II)| \lesssim \frac{1}{\log(1/\xi)^6}$. Combining the bounds on $|(I)|$, $|(II)|$ and $|(III)|$ yields that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mu^{\otimes n}} [f'(\mathbf{x})g'(\mathbf{y})h(\mathbf{z})] \right| \lesssim \frac{1}{\log(1/\xi)^6}.$$

thereby establishing the assertion of Lemma 3.3. \square

B Missing proofs from Section 4

In this section, we prove the claims establishing the SVD decompositions for homogenous and non-homogenous functions.

B.1 Proof of Claim 4.8

Proof. We think of g as a matrix M in $\mathbb{R}^{\Gamma^I \times \Gamma^J}$, whose (\mathbf{a}, \mathbf{b}) entry is $g(\mathbf{y}_I = \mathbf{a}, \mathbf{y}_J = \mathbf{b})$. The decomposition stated by the claim is an appropriately chosen singular-value decomposition of M ; below are the details for completeness

Looking at $M^t M \in \mathbb{R}^{\Gamma^J \times \Gamma^J}$, we see that it is an $m \times m$ symmetric real-valued matrix, hence we may find an eigenbasis g'_1, \dots, g'_m of \mathbb{R}^{Γ^J} with eigenvalues $\lambda_1, \dots, \lambda_m$. We note that the all 1 vector is an eigenvector of $M^t M$; indeed:

$$\begin{aligned} (M^t M \vec{1})_{\mathbf{a}} &= \sum_{\mathbf{b}} (M^t M)_{\mathbf{a}, \mathbf{b}} = \sum_{\mathbf{b}} \sum_{\mathbf{y}_I \in \Gamma^I} M^t[\mathbf{a}, \mathbf{y}_I] M[\mathbf{y}_I, \mathbf{b}] = \sum_{\mathbf{b}} \sum_{\mathbf{y}_I \in \Gamma^I} g(\mathbf{y}_I, \mathbf{b}) g(\mathbf{y}_I, \mathbf{a}) \\ &= \sum_{\mathbf{y}_I \in \Gamma^I} g(\mathbf{y}_I, \mathbf{a}) \sum_{\mathbf{b}} g(\mathbf{y}_I, \mathbf{b}). \end{aligned}$$

Consider the function $A: \Gamma^I \rightarrow \mathbb{R}$ defined by $A(\mathbf{y}_I) = \sum_{\mathbf{b}} g(\mathbf{y}_I, \mathbf{b})$. Note that for each monomial in g that contains the variable from J sums up to 0 when we sum over \mathbf{b} , and any other monomial is multiplied by m . Thus, $A(\mathbf{y}_I) = m \cdot \tilde{g}$, where \tilde{g} is the part of g which does not include a variable from J , and in particular it is a homogenous function of degree d_1 .

Thus, the sum we look at is proportional to $\langle g_{J \rightarrow \mathbf{a}}, \tilde{g} \rangle = \langle \tilde{g}_{J \rightarrow \mathbf{a}}, \tilde{g} \rangle$. The last equality holds since any monomial in $(g - \tilde{g})_{J \rightarrow \mathbf{a}}$ has degree at most $d - 1$ and hence has inner product 0 with \tilde{g} . In particular, as $\langle \tilde{g}_{J \rightarrow \mathbf{a}}, \tilde{g} \rangle$ does not depend on \mathbf{a} , we get that the quantity $(M^t M \vec{1})_{\mathbf{a}}$ is the same for all \mathbf{a} , and hence $\vec{1}$ is an eigenvector of $M^t M$.

This means that in choosing the eigenbasis g'_1, \dots, g'_m , we can ensure that $g'_1 \equiv 1$. This gives us the third and fourth bullets, and next we choose g_1, \dots, g_m . Define $\tilde{g}_r = M g'_r$; first we note that \tilde{g}_r are orthogonal. Indeed,

$$\langle \tilde{g}_r, \tilde{g}_{r'} \rangle = \langle M g'_r, M g'_{r'} \rangle = \langle M^t M g'_r, g'_{r'} \rangle = \lambda_r \langle g'_r, g'_{r'} \rangle = \lambda_r \mathbf{1}_{r \neq r'}$$

This means that if we look only at the set R of i 's such that $\lambda_i \neq 0$, then we get that $\{\tilde{g}_i\}_{i \in R}$ is orthogonal, and we choose $g_r = \tilde{g}_r / \sqrt{\lambda_r}$ (which has 2-norm equal to 1)

We prove that

$$M = \sum_{r \in R} \sqrt{\lambda_r} g_r(\mathbf{y}_I) g'_r(\mathbf{y}_J).$$

To see that, define $M' = \sum_{r \in R} \lambda_r g_r(\mathbf{y}_I) g'_r(\mathbf{y}_J)$, and note that for all $r \in R$ we have $M g'_r = M' g'_r = \sqrt{\lambda_r} g_r \neq 0$, and for $r \notin R$ we have $M g'_r = 0$ as well as $M' g'_r = 0$, and hence $M g'_r = M' g'_r$ for all r . This implies that $M = M'$, and we have thus established the decomposition of g as well as the first item (we can freely define g_r for $r \notin R$ by noting that as $\lambda_r = 0$ there, it doesn't change anything).

For the fifth item, we observe that

$$\sum_r \sqrt{\lambda_r}^2 = \sum_r \lambda_r = \frac{\text{Tr}(M^t M)}{|\Gamma|^n} = \frac{1}{|\Gamma|^n} \sum_b (M^t M)_{b,b} = \frac{1}{|\Gamma|^n} \sum_{b,y} g(y,b)^2 = \|g\|_2^2 = 1.$$

Finally, we argue for second item. If $1 \in R$, then we get that $g_1(\mathbf{y}_I) = \lambda_1^{-1/2} \mathbb{E}_{\mathbf{y}_J} [g(\mathbf{y})]$, from which it is clear that g_1 is homogenous of degree d . For $r \neq 1$ in R , we get that

$$g_r(\mathbf{y}_I) = \lambda_r^{-1/2} \mathbb{E}_{\mathbf{y}_J} [g(\mathbf{y}) g'_r(\mathbf{y}_J)],$$

and we note that on the right hand side only monomials in g that contain the variable from J can contribute (the rest give 0), hence we get a homogenous function of degree $d - 1$. \square

B.2 Proof of Claim 4.10

Proof. We run the same argument as in the proof of Claim 4.8, replacing λ_r 's with γ_t 's, and get that

$$\sum_{t \in T} \sqrt{\gamma_t} F_t(\mathbf{y}_I, \mathbf{z}_I) F'_t(\mathbf{y}_J, \mathbf{z}_J).$$

This gives the above decomposition and shows that it satisfies all items except for the third and fourth. For $t = 1$, we note as in the proof of Claim 4.8 that

$$F_1(\mathbf{y}_I, \mathbf{z}_I) = \gamma_1^{-1/2} \mathbb{E}_{\mathbf{y}_J, \mathbf{z}_J} [F(\mathbf{y}, \mathbf{z})],$$

from which it is clear that F_1 has effective degree at least d' and that it is constant on connected components.

For $t \geq 2$ we have that

$$F_t(\mathbf{y}_I, \mathbf{z}_I) = \gamma_t^{-1/2} \mathbb{E}_{\mathbf{y}_J, \mathbf{z}_J} [F(\mathbf{y}, \mathbf{z}) F'_t(\mathbf{y}_J, \mathbf{z}_J)],$$

and expanding F in as linear combination of monomials, we see that monomials from F not involving the variable from J retain their effective degree, while those monomials that involve the variable from J may drop their effective degree by at most 1 (that happens if the monomial has a character from WB_2 on that coordinate). It also follows from this representation that as F is constant on connected components, F_t is also constant on connected components.

Finally, we also note that

$$F'_t(\mathbf{y}_J, \mathbf{z}_J) = \gamma_t^{-1/2} \mathbb{E}_{\mathbf{y}_I, \mathbf{z}_I} [F(\mathbf{y}, \mathbf{z}) F_t(\mathbf{y}_I, \mathbf{z}_I)],$$

from which it is clear that F'_t is constant on connected components. This completes the proof of the third item, and therefore of Claim 4.10. \square

B.3 Proof of Claim 4.11

Proof. We think of g as a matrix M in $\mathbb{R}^{\Gamma^I \times \Gamma^J}$, whose (\mathbf{a}, \mathbf{b}) entry is $g(\mathbf{y}_I = \mathbf{a}, \mathbf{y}_J = \mathbf{b})$. The decomposition stated by the claim is an appropriately chosen singular-value decomposition of M ; below are the details for completeness

Looking at $M^t M \in \mathbb{R}^{\Gamma^J \times \Gamma^J}$, we see that it is an $m \times m$ symmetric real-valued matrix, hence we may find an eigenbasis g'_1, \dots, g'_m of \mathbb{R}^{Γ^J} with eigenvalues $\lambda_1, \dots, \lambda_m$. We define $\tilde{g}_r = M g'_r$, set $R = \{r \mid \lambda_r \neq 0\}$ and normalize $g_r = \tilde{g}_r / \sqrt{\lambda_r}$ for $r \in R$, and as in Claim 4.8 we get that

$$M = \sum_{r \in R} \sqrt{\lambda_r} g_r(\mathbf{y}_I) g'_r(\mathbf{y}_J),$$

as well as the first, third and fourth items.

Finally, we argue for second item. For $r \in R$, then we get that $g_r(\mathbf{y}_I) = \lambda_r^{-1/2} \mathbb{E}_{\mathbf{y}_J} [g(\mathbf{y}) g'_r(\mathbf{y}_J)]$, and as for each \mathbf{y}_J , $g(\mathbf{y}) g'_r(\mathbf{y}_J)$ contains only monomials of degree $d - 1$ onwards, we get that the same is true for g_r . \square