

On Approximability of Satisfiable k-CSPs: IV

Amey Bhangale*

Subhash Khot†

Dor Minzer‡

Abstract

We prove a stability result for general 3-wise correlations over distributions satisfying mild connectivity properties. More concretely, we show that if Σ, Γ and Φ are alphabets of constant size, and μ is a distribution over $\Sigma \times \Gamma \times \Phi$ satisfying: (1) the probability of each atom is at least $\Omega(1)$, (2) μ is pairwise connected, and (3) μ has no Abelian embeddings into $(\mathbb{Z}, +)$, then the following holds. Any triplets of 1-bounded functions $f: \Sigma^n \rightarrow \mathbb{C}, g: \Gamma^n \rightarrow \mathbb{C}, h: \Phi^n \rightarrow \mathbb{C}$ satisfying

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

must arise from an Abelian group associated with the distribution μ . More specifically, we show that there is an Abelian group $(H, +)$ of constant size such that for any such f, g and h , the function f (and similarly g and h) is correlated with a function of the form $\tilde{f}(x) = \chi(\sigma(x_1), \dots, \sigma(x_n))L(x)$, where $\sigma: \Sigma \rightarrow H$ is some map, $\chi \in \hat{H}^{\otimes n}$ is a character, and $L: \Sigma^n \rightarrow \mathbb{C}$ is a low-degree function with bounded 2-norm.

En route we prove a few additional results that may be of independent interest, such as an improved direct product theorem, as well as a result we refer to as a “restriction inverse theorem” about the structure of functions that, under random restrictions, with noticeable probability have significant correlation with a product function.

In companion papers, we show applications of our results to the fields of Probabilistically Checkable Proofs, as well as various areas in discrete mathematics such as extremal combinatorics and additive combinatorics.

1 Introduction

1.1 Studying 3-wise Correlations with Respect to a Distribution

Let Σ, Γ and Φ be alphabets of constant size, suppose μ is a distribution over $\Sigma \times \Gamma \times \Phi$, and let $f: \Sigma^n \rightarrow \mathbb{C}, g: \Gamma^n \rightarrow \mathbb{C}, h: \Phi^n \rightarrow \mathbb{C}$ be 1-bounded functions. What sort of triplets of functions f, g and h have a significant 3-wise correlation with respect to μ ? In other words, what can be said about the functions f, g and h in the case that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f(x)g(y)h(z)] \right| \geq \varepsilon, \tag{1}$$

where $\varepsilon > 0$ is thought of as a small constant? In [30], it is shown that if μ is connected, then this can only be the case if each one of f, g and h is correlated with a low-degree function. Here, we say that a distribution

*Department of Computer Science and Engineering, University of California, Riverside.

†Department of Computer Science, Courant Institute of Mathematical Sciences, New York University. Supported by the NSF Award CCF-1422159, NSF CCF award 2130816, and the Simons Investigator Award.

‡Department of Mathematics, Massachusetts Institute of Technology. Supported by a Sloan Research Fellowship, NSF CCF award 2227876 and NSF CAREER award 2239160.

μ over $\Sigma_1 \times \Sigma_2 \times \Sigma_3$ is connected if for any partition of $\{1, 2, 3\}$ into two sets $I \cup J$, the bi-partite graph between $\text{supp}(\mu_I)$ and $\text{supp}(\mu_J)$ whose edges are all (a, b) if $(a, b) \in \text{supp}(\mu)$, is connected (μ_I is the marginal distribution of μ on the coordinates of I). In [4, 5], a strengthening of this result is proved, and it is shown that it suffices that the distribution μ does not admit any *Abelian embeddings*.

Definition 1.1. *An Abelian embedding of a distribution μ over $\Sigma \times \Gamma \times \Phi$ consists of an Abelian group $(H, +)$ and 3 maps $\sigma: \Sigma \rightarrow H$, $\gamma: \Gamma \rightarrow H$ and $\phi: \Phi \rightarrow H$ such that $\sigma(x) + \gamma(y) + \phi(z) = 0$ for all $(x, y, z) \in \text{supp}(\mu)$. We say that the embedding (σ, γ, ϕ) is non-trivial if at least one of the maps is not constant.*

Definition 1.2. *We say a distribution μ admits an Abelian embedding if it has a non-trivial Abelian embedding.*

In this language, the main result of [4, 5] asserts that if μ does not admit an Abelian embedding and the probability of each atom in μ is at least α thought of as a constant, then each one of f , g and h must be correlated with a low-degree function. As it can easily be seen, this result is strictly stronger than the corresponding result in [30] since any distribution μ which is connected does not admit an Abelian embedding. Moreover, as explained in [4, 5] this result is an if and only if, in the sense that in the presence of Abelian embedding one could design 1-bounded functions f , g and h for which (1) holds while at least one of the functions f , g and h only has $o(1)$ -correlation with any low-degree function.

The main goal of this paper is to extend this understanding beyond the realm of distributions which do not have Abelian embeddings and prove structural results on functions f , g and h satisfying (1) in this more general setting. At a high level, one would like to say that such functions f , g and h could only arise as a result of using Abelian embeddings, using low-degree functions, or both. To prove such result however, we must focus our attention on distributions μ satisfying a very mild form of connectivity, which we refer to as pairwise connectedness.

Definition 1.3. *Let $\Sigma_1, \Sigma_2, \Sigma_3$ be finite alphabets, and let $P \subseteq \Sigma_1 \times \Sigma_2 \times \Sigma_3$. For a pair of distinct coordinates $i, j \in \{1, 2, 3\}$, we say P is $\{i, j\}$ connected if the bipartite graph $G = (\Sigma_i \cup \Sigma_j, E_{i,j})$, where $E_{i,j}$ is the set of label pairs that appear in some element of P , is connected.*

We say P is pairwise connected if it is pairwise connected for any two distinct $i, j \in \{1, 2, 3\}$.

We say a distribution μ is pairwise connected if $\text{supp}(\mu)$ is pairwise connected.

At a high level, the notion of pairwise connectedness stems from the fact that if $\text{supp}(\mu)$ is not pairwise connected, then there are examples of functions satisfying (1) without any useful structure for our purposes. Indeed, if $\text{supp}(\mu)$ is not pairwise connected – without loss of generality it is not $\{1, 2\}$ -connected, then we may find a non-trivial partition $\Sigma = \Sigma' \cup \Sigma''$ and $\Gamma = \Gamma' \cup \Gamma''$ so that in the support of μ there can only be pairs from $\Sigma' \times \Gamma'$ and $\Sigma'' \times \Gamma''$ on the first two coordinates. In this case, we may pick any pair of functions $s, s': \{1, 2\}^n \rightarrow \mathbb{C}$ such that $s(a)s'(a) = 1$ for all $a \in \{1, 2\}^n$ (for example, one can take s whose absolute value is always 1, and s' to be its conjugate) and construct f, g, h as follows. For f , we set $f(x) = s(x')$ where for each i , $x'_i = 1$ if $x_i \in \Sigma'$ and otherwise $x'_i = 2$. For g , we similarly set $g(y) = s'(y')$ where for each i , $y'_i = 1$ if $y_i \in \Gamma'$ and otherwise $y'_i = 2$. For h , we take $h \equiv 1$. Thus, for any $(x, y, z) \in \text{supp}(\mu)$ we have that

$$f(x)g(y)h(z) = s(x')s'(y') = 1,$$

as we have that $x' = y'$ by construction.

Henceforth, we will focus our attention on distributions μ which are pairwise connected. With this in mind, as explained earlier there are two ways of constructing functions f , g and h satisfying (1):

1. If $\text{supp}(\mu)$ admits a linear embedding, say for simplicity a cyclic group $(H, +) = (\mathbb{Z}_p, +)$ and maps $\sigma: \Sigma \rightarrow H$, $\gamma: \Gamma \rightarrow H$ and $\phi: \Phi \rightarrow H$ not all constant such that $\sigma(x) + \gamma(y) + \phi(z) = 0$, then one can take

$$f(x_1, \dots, x_n) = e^{\frac{2\pi i}{|H|}(\sigma(x_1) + \dots + \sigma(x_n))}, \quad g(y_1, \dots, y_n) = e^{\frac{2\pi i}{|H|}(\gamma(y_1) + \dots + \gamma(y_n))},$$

$$h(z_1, \dots, z_n) = e^{\frac{2\pi i}{|H|}(\phi(z_1) + \dots + \phi(z_n))},$$

and note that $f(x)g(y)h(z) = 1$ pointwise, hence (1) holds. More generally, for a general Abelian group $(H, +)$ one can pick non-trivial characters $\chi_1, \dots, \chi_n \in \hat{H}$, define

$$f(x_1, \dots, x_n) = \prod_{j=1}^n \chi_j(\sigma(x_j)), \quad g(y_1, \dots, y_n) = \prod_{j=1}^n \chi_j(\gamma(y_j)), \quad h(z_1, \dots, z_n) = \prod_{j=1}^n \chi_j(\phi(z_j)),$$

and note again that $f(x)g(y)h(z) = 1$ pointwise hence (1) holds.

2. In general, it may also be the case that for a distribution μ , low-degree functions also satisfy (1). Indeed, in that case one may try to find univariate 1-bounded functions $u: \Sigma \rightarrow \mathbb{C}$, $v: \Gamma \rightarrow \mathbb{C}$ and $w: \Phi \rightarrow \mathbb{C}$ for which $|\mathbb{E}_{(x,y,z) \sim \mu} [u(x)v(y)w(z)]| \geq \Omega(1)$, and then tensorize them to get

$$f(x_1, \dots, x_n) = \prod_{j=1}^d u(x_j), \quad g(y_1, \dots, y_n) = \prod_{j=1}^d v(y_j), \quad h(z_1, \dots, z_n) = \prod_{j=1}^d w(z_j),$$

which get value of $2^{-O(d)}$ in (1).

1.2 Main Results

With the above discussion in mind, one is tempted to conjecture that if μ is pairwise connected, then the only possible examples of triplets of functions f, g and h satisfying (1) must come from the above template.

1.2.1 The Stability Result

The main result of this paper is a stability result that formalizes this intuition, saying that under some mild assumptions on the distribution μ , if f, g and h are 1-bounded functions achieving significant 3-wise correlation as in (1), then f (and similarly g and h) must be correlated with a product of an embedding type function as in the first recipe, with a low-degree function as in the second recipe. The mild assumptions on μ correspond to it being pairwise connected (which is necessary, otherwise the statement is simply false), and for technical reasons we also need an additional assumption, namely that μ cannot be embedded in the Abelian group $(\mathbb{Z}, +)$. We remark though that this additional assumption is, as far as we know, not necessary, but removing it seems to require more ideas. With this in mind, a precise formulation of our main result is:

Theorem 1.4. *For all $m \in \mathbb{N}$, $\alpha > 0$ and $\varepsilon > 0$, there are $d \in \mathbb{N}$ and $\delta > 0$ such that the following holds. Suppose that μ is a distribution over $\Sigma \times \Gamma \times \Phi$ such that*

1. *The probability of each atom in μ is at least α .*
2. *The size of each one of Σ, Γ, Φ is at most m .*

3. The distribution μ is pairwise connected.
4. μ does not admit an Abelian embedding into $(\mathbb{Z}, +)$.

Then, if $f: \Sigma^n \rightarrow \mathbb{C}$, $g: \Gamma^n \rightarrow \mathbb{C}$ and $h: \Phi^n \rightarrow \mathbb{C}$ are 1-bounded functions such that

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

then there are 1-bounded functions $u_1, \dots, u_n: \Sigma \rightarrow \mathbb{C}$ and a function $L: \Sigma^n \rightarrow \mathbb{C}$ of degree at most d and 2-norm at most 1 such that

$$\left| \mathbb{E}_{x \sim \mu_{\tilde{x}}^{\otimes n}} \left[f(x) \cdot L(x) \prod_{i=1}^n u_i(x_i) \right] \right| \geq \delta.$$

Furthermore, there is $r \in \mathbb{N}$ depending only on m and an Abelian embedding (σ, γ, ϕ) of μ into an Abelian group $(H, +)$ of size at most r such that for all i , $u_i(x_i) = \chi_i(\sigma(x_i))$ where $\chi_i \in \widehat{H}$ is a character of H .

Quantitatively, we have that

$$d = \text{poly}_{\alpha, m}(1/\varepsilon), \quad \delta = 2^{-\text{poly}_{\alpha, m}(1/\varepsilon)}.$$

The proof of Theorem 1.4 is quite long, and in Section 1.4 we give an overview of the steps we take in the proof. Some of the steps require ingredients that may be of independent interest, and which we explain next.

1.2.2 The Restriction Inverse Theorem

The proof of Theorem 1.4 uses a result which we refer to as the Restriction Inverse Theorem and present next.

Restrictions and Random Restrictions. Restrictions and random restrictions are vital to our argument to go through, and the Restriction Inverse Theorem can be thought of as a statement about them of independent interest. Given a function $f: (\Sigma^n, \mu^{\otimes n}) \rightarrow \mathbb{C}$, a set of coordinates $I \subseteq [n]$ and $\tilde{x} \in \Sigma^{\bar{I}}$, the restricted function $f_{\bar{I} \rightarrow \tilde{x}}$ is a function from Σ^I to \mathbb{C} defined as

$$f_{\bar{I} \rightarrow \tilde{x}}(x') = f(x_I = x', x_{\bar{I}} = \tilde{x}),$$

where $(x_I = x', x_{\bar{I}} = \tilde{x})$ is the point in Σ^n whose I -coordinates are set according to x' , and whose \bar{I} -coordinates are set according to \tilde{x} .

Random restrictions are restrictions in which either I , \tilde{x} or both are chosen randomly. A typical setting we use is one where we have a parameter $\rho > 0$, and we pick $I \subseteq_{\rho} [n]$, by which we mean that we include each $i \in [n]$ in I with probability ρ ; we then choose $\tilde{x} \sim \mu^{\bar{I}}$. For the purposes of this paper it is necessary to consider other (less standard) settings of random restrictions, but we will limit ourselves to this more typical setting for the purposes of this introduction; we refer the reader to Section 2.3 for a discussion on the other settings we use.

Product functions. A function $f: \Sigma^n \rightarrow \mathbb{C}$ is called a product function if there are 1-bounded functions $f_1, \dots, f_n: \Sigma \rightarrow \mathbb{C}$ such that

$$f(x_1 \dots, x_n) = \prod_{i=1}^n f_i(x_i).$$

It is clear that if f is a product function, then any restriction of it is still a product function. Thus, with probability 1, taking a random restriction of f yields a function which has perfect correlation with a product function. The Restriction Inverse Theorem is a statement about a deduction in the reverse direction: suppose f is a function that after random restriction it has a significant correlation with a product function. Is it necessarily the case that f itself is correlated with a product function?

As is usually the case with inverse-type questions, there are multiple regimes of parameters one may consider, and for us the most relevant regime is the so-called 1% regime. In this case, we have a parameter $\rho > 0$ (which is small but bounded away from 0) and a function $f: \Sigma^n \rightarrow \mathbb{C}$ such that

$$\Pr_{I \subseteq \rho[n], \tilde{x} \sim \mu^I} [\exists p: \Sigma^I \rightarrow \mathbb{C} \text{ a product function such that } |\langle f_{I \rightarrow \tilde{x}}, p \rangle| \geq \varepsilon] \geq \varepsilon, \quad (2)$$

and we wish to deduce a structural result about f . As discussed, such situations may arise when f is a product function – or more generally when it is correlated with a product function. However, if f is a low-degree function (or even if it is just correlated with a low-degree function), a random restriction of f will be correlated with a constant function with noticeable probability, and hence with a product function. The Restriction Inverse Theorem essentially says that these are the only two ways that (2) can come about:

Theorem 1.5 (The Restriction Inverse Theorem, Informal). *For all $\varepsilon, \rho, \alpha > 0$ and $m \in \mathbb{N}$ there are $d \in \mathbb{N}$ and $\delta > 0$ such that the following holds. Suppose Σ is a finite alphabet of size at most m , μ is a distribution over Σ in which the probability of each atom is at least α , and $f: (\Sigma^n, \mu^{\otimes n}) \rightarrow \mathbb{C}$ is a 1-bounded function satisfying (2). Then there is a product function $p: \Sigma^n \rightarrow \mathbb{C}$ and a function $L: \Sigma^n \rightarrow \mathbb{C}$ of degree at most d and $\|L\|_2 \leq 1$ such that*

$$|\langle f, pL \rangle| \geq \delta.$$

We refer the reader to Section 12 for a more formal and general version of the Restriction Inverse Theorem. We remark that among other things, we also give explicit dependency of d and δ on ε and ρ . These quantitative aspects are important if one wishes to get decent quantitative bounds in Theorem 1.4, and we think they are also interesting in their own right.

1.2.3 The Direct Product Theorem

The proof of Theorem 1.5 (and thus, in turn, of Theorem 1.4) hinges on a direct product testing result, which may also be of independent interest. The problem of direct product testing has its roots in the field of probabilistic checkable proofs and in particular in hardness amplification. In this setting, one wishes to encode a function $f: [n] \rightarrow [R]$ (where n is thought of as very large) by local pieces that, on the one hand allows for local access to values of f . On the other hand, the encoding should be testable, in the sense that there is a test that only looks at a handful of locations of the encoding and determines whether it is an encoding of an actual function $f: [n] \rightarrow [R]$, or whether it is far from the encoding of any such function.

Our application calls for a particular direct product tester that has been extensively studied in the literature [14, 12, 28, 27, 16, 11, 6]. In this tester, the function f is encoded via its table of restrictions to sub-cubes of certain dimension. Namely, given a parameter $k \in \mathbb{N}$ (which for us will be equal to ρn , where ρ should

be thought of as a very small constant), the direct product encoding of f is the mapping $F: \binom{[n]}{k} \rightarrow [R]^k$ defined by

$$F[A] = f|_A$$

for all $A \subseteq [n]$ of size k .

The test we associate with this encoding is determined by two parameters, $\alpha, \beta \in (0, 1)$ that also should be thought of as small constants. Given a supposed table of restrictions $G: \binom{[n]}{k} \rightarrow [R]^k$, the test, which we call $\text{DP}(\rho, \alpha, \beta)$, proceeds in the following way:

1. Sample $C \subseteq [n]$ of size αk and sample $A, B \in \binom{[n]}{k}$ independently containing C .
2. Sample $T \subseteq [n]$ of size βn .
3. Query $G[A], G[B]$ and check that $G[A]|_{C \cap T} = G[B]|_{C \cap T}$.

In other words, the tester selected two sets A, B that intersect on a sizable number of elements (at least αk), then a random subset of their shared elements and checks that the local assignments $G[A]$ and $G[B]$ agree on this random subset of shared elements.

Note that this test is complete, in the sense that if G is a legitimate direct product encoding, then it passes the test with probability 1. Thus, as is usually the case, the interesting aspect of this test is the soundness, which is equivalent to the following question. Suppose that the tester accepts a table $G: \binom{[n]}{k} \rightarrow [R]^k$ with probability at least s ; is it necessarily the case that G is somewhat close to a legitimate direct product testing codewords?

In the so-called 99% regime, where the probability $s = 1 - \varepsilon$ is thought of close to 1, this problem is completely understood, and in [16, 11] it is shown that in this case there is a function $f: [n] \rightarrow [R]$ such that for at least $1 - O(\varepsilon)$ fraction of $A \in \binom{[n]}{k}$ it holds that $G[A] = f|_A$.

For us, the most so-called 1% regime is more relevant, wherein the probability $s = \varepsilon$ is thought of as close to 0. In this case, one can no longer expect a strong conclusion as in the 99% regime. Instead, naturally one would expect that in this case, there would have to be a function $f: [n] \rightarrow [R]$ such that for at least $\delta = \delta(\rho, \alpha, \beta, \varepsilon) > 0$ fraction of $A \in \binom{[n]}{k}$ it holds that $G[A] = f|_A$, but this is also too much to expect. Indeed, to see that take any $g: [n] \rightarrow [R]$, and for each A take $G[A]$ uniformly from $[R]^k$ with probability $1 - \varepsilon$, and otherwise take it to be a string in $[R]^k$ of Hamming distance $r = \Theta(\log(1/\varepsilon))$ from $g|_A$. Using Chernoff's bound, one can prove that with high probability there is no function $f: [n] \rightarrow [R]$ satisfying the natural conclusion one expects, yet the tester passes with probability at least $\varepsilon^2(1 - \beta)^{2r} = \text{poly}(\varepsilon)$. The reason for that is that looking at two locations A, B queried by the tester, with probability ε^2 both of them get assigned strings close to $g|_A$ and $g|_B$ respectively, in which case with probability at least $(1 - \beta)^{2r}$ the subset T excludes all coordinates on which either $G[A]$ and $g|_A$, or $G[B]$ and $g|_B$, disagree on.

Due to a rather versatile set of examples, results in the 1% regime are often more challenging to prove. Indeed, earlier results by [12, 16] managed to show that in this case there is a function $f: [n] \rightarrow [R]$ such that for at least $\delta = \delta(\rho, \alpha, \beta, \varepsilon, \eta) > 0$ fraction of $A \in \binom{[n]}{k}$ it holds that $\Delta(G[A], f|_A) \leq \eta k$. Here and throughout, $\Delta(x, y)$ represents the Hamming distance between strings x and y . The main drawback of this result is that the distance between $G[A]$ and $f|_A$ is linear in k , which is not good enough for our purposes. Indeed, for our application we need a result that gets a Hamming distance which is a constant $r = r(\rho, \alpha, \beta, \varepsilon) \in \mathbb{N}$ as opposed to a constant fraction.

In [6], such result was proved for a more specialized version of this test in the case of $\beta = 1$ and $R = 2$. Therein, both of the parameters α and ρ are thought of as constant, and it is proved that there are $r = r(\varepsilon, \alpha, \rho) \in \mathbb{N}$ and $\delta = \delta(\varepsilon, \alpha, \rho) > 0$ such that if G passes the test $\text{DP}(\rho, \alpha, \beta = 1)$ with probability

at least ε , then there is a function $f: [n] \rightarrow \{0, 1\}$ such that for at least δ fraction of $A \in \binom{[n]}{k}$ it holds that $\Delta(G[A], f|_A) \leq r$. Besides being a natural question of interest, the motivation of this result therein was to establish an earlier, less general version of the Restriction Inverse Theorem, Theorem 1.5 herein.

In this paper, we are once again in a situation that our proof of a restriction inverse theorem requires a direct product testing result, and the relevant test for us is the test $\text{DP}(\rho, \alpha, \beta)$ above. Moreover, as herein we are concerned with getting good quantitative bounds, we no longer think of the parameters ρ, α, β as constants and thus try to get reasonable dependencies of r and δ on these parameters. For the purposes of this introductory section however, we do not mention these quantitative aspects and defer the interested reader to Section 13. Thus, without a concern for these quantitative aspects our result reads:

Theorem 1.6 (The Direct Product Testing Theorem, Informal). *For all $\varepsilon, \rho, \alpha, \beta > 0$ there are $r \in \mathbb{N}$ and $\delta > 0$ such that the following holds for all $R \in \mathbb{N}$. For $k = \rho n$, if $G: \binom{[n]}{k} \rightarrow [R]^k$ is a function that passes the test $\text{DP}(\rho, \alpha, \beta)$ with probability at least ε , then there is a function $f: [n] \rightarrow [R]$ such that*

$$\Pr_{A \in \binom{[n]}{k}} [\Delta(f|_A, G[A]) \leq r] \geq \delta.$$

1.3 Applications and Motivations

In this section, we discuss some applications and motivating fields and type of problems Theorem 1.4 (and possible extensions of it) are likely to be related to. For some of them, we already have initial leads (and pursue them in subsequent papers as the current paper is already long enough as is), while for others the connection is more speculative.

1.3.1 Hardness of Approximation

Recall Mossel’s result [30], asserting that in the case that μ is a connected distribution only the low-degree part of functions contributes to (1). For low-degree functions, one has the invariance principle of [31], and thus the combination of these two results can be seen as transforming expectations as in (1) to expectations over Gaussian space. This result has a few notable striking consequences in the field of hardness of approximation. Most notably, Raghavendra [34] uses precisely such ideas to show the relationship between dictatorship tests and Gaussian rounding scheme to semi-definite relaxations.

In this light, the result proved in this paper shows that only functions that are “characters times low-degree functions” can contribute to (1), and this suggests that an invariance principle that extends the invariance principle of [31] should exist. Indeed, in a future work [3] we are exploring this direction and will prove a more general such invariance principle, and discuss its relation to rounding schemes that combine semi-definite programming relaxations as well as linear programming relaxations. We believe such invariance principles will be crucial in the journey of understanding the approximability of satisfiable constraint satisfaction problems.

1.3.2 Higher Arity Predicates

The original motivation behind the question considered in this paper is the non-Abelian embedding hypothesis of [4], which is the following statement. Suppose $k \geq 3$ is an integer, $\Sigma_1, \dots, \Sigma_k$ finite alphabets and μ is a distribution over $\Sigma_1 \times \dots \times \Sigma_k$ in which the probability of each atom is at least $\alpha > 0$. We say μ admits an Abelian embedding if there is an Abelian group $(H, +)$ and maps $\sigma_i: \Sigma_i \rightarrow H$ for $i = 1, \dots, k$

such that $\sum_{i=1}^k \sigma_i(x_i) = 0$ for all $(x_1, \dots, x_k) \in \text{supp}(\mu)$. We say μ admits a non-trivial Abelian embedding if at least one of the maps σ_i is non-constant.

Hypothesis 1.7. *In the above setting, if μ admits no non-trivial Abelian embeddings, then for all $\varepsilon > 0$ there is $\delta > 0$ such that if $f_i: \Sigma_i^n \rightarrow \mathbb{C}$ are 1-bounded functions with $\text{Stab}_{1/2}(f_i; \mu_i^{\otimes n}) \leq \delta$ for at least one of the i 's, then*

$$\left| \mathbb{E}_{(x_1, \dots, x_k) \sim \mu^{\otimes n}} \left[\prod_{i=1}^k f_i(x_i) \right] \right| \leq \delta.$$

In [4] a special case of this hypothesis is proved for a class of $k = 3$ -ary distributions, and in [5] this hypothesis is proved in general for all $k = 3$ -ary distributions. In these terms, the current paper does not signify any further progress towards establishing Hypothesis 1.7 beyond the case of 3-ary predicates, however we believe that the stability version proved herein will be crucial towards making further progress in this direction.

1.3.3 Gowers' Norms

Theorem 1.4 can be seen as an analog of the U_2 -inverse theorem for Gowers uniformity norms [22] for general distributions. In the context of Gowers uniformity norms, the U_2 -inverse theorem is a simple Fourier analytic computation only involving Fourier coefficients. Interestingly, at a point in our argument we too have to carry out such a computation (this is, however, a small part of the proof). It is tempting to speculate, and we have initial leads for this fact, that there should be higher order analogs of Gowers inverse theorems in the much more general setting of Theorem 1.4.

If true, such statements could be very useful to make progress on multiple problems in extremal combinatorics, and in particular in Szemerédi-type theorems [38]. This is so because it appears they are strong enough to facilitate density increment arguments. Indeed, as we explain next, in a companion paper we have used Theorem 1.4 to give effective bounds for the problem of finding restricted 3-arithmetic progressions in dense sets in \mathbb{F}_p^n , for a prime p .

1.3.4 Extremal Combinatorics

A set $A \subseteq \mathbb{F}_p^n$ is called somewhat restricted 3-AP free if it does not contain an arithmetic progression $x, x + a, x + 2a$ where $x \in \mathbb{F}_p^n$ and $a \in \{0, 1, 2\}^n \setminus \{\vec{0}\}$. In a companion paper [2], we use Theorem 1.4 to give effective bounds on the density of restricted 3-AP sets:

Theorem 1.8. *There are absolute constants $C > 0$ and $1 \leq k \leq 10$ such that if $A \subseteq \mathbb{F}_p^n$ is a restricted 3-AP set, then*

$$\mu(A) = \frac{|A|}{p^n} \leq \frac{C}{\log^{(k)} n},$$

where $\log^{(k)} n$ is the k -fold iterated logarithm function.

Previously, the best known bound was $O(1/\log^* n)$, achieved by appealing to a quantitative version of the density Hales-Jewett theorem [33]. Theorem 1.8 makes progress on a question of Green [23] and on a question of Haszla, Holenstein and Mossel [30].

1.3.5 Multi-Player Parallel Repetition Theorems

Parallel repetition is a basic building block in the area of interactive protocols and in particular in applications in the field of hardness of approximation. In the setting of k -player games, we have a basic game Ψ involving a verifier and k players. The game consists of a set of questions X that are supposed to get labels from a finite alphabet Σ , a predicate $P: X^k \times \Sigma^k \rightarrow \{0, 1\}$ that gives k -challenges and answers to them dictates whether these answers are deemed satisfactory or not, and a distribution μ over k -tuples of challenges. In the basic game Ψ , the verifier samples a challenge $(x_1, \dots, x_k) \sim \mu$, sends the question x_i to the i th player, receives an answer $a_i \in \Sigma$ from player i , and then accepts if and only if $P(x_1, \dots, x_k, a_1, \dots, a_k) = 1$. The value of the game, $\text{val}(\Psi)$, is defined to be the maximum probability the verifier accepts under the best strategy for the players.

The t -fold repeated game, $\Psi^{\otimes t}$, is a game in which the verifier samples t sets of challenges, say $(x_{1,j}, \dots, x_{k,j}) \sim \mu$ for $j = 1, \dots, t$ independently, sends $(x_{i,1}, \dots, x_{i,t})$ to player i , receives from them answers $(a_{i,1}, \dots, a_{i,t})$ and accepts if and only if $P(x_{1,j}, \dots, x_{k,j}, a_{1,j}, \dots, a_{k,j}) = 1$ for all $j = 1, \dots, t$. In words, the game is repeated for t -times, but in parallel, and the verifier confirms that each one of the executions of the basic game was accepting. It is clear that $\text{val}(\Psi^{\otimes t}) \geq \text{val}(\Psi)^t$, and the main question of interest in parallel repetition theorems is regarding the rate of decay of $\text{val}(\Psi^{\otimes t})$ as a function of t ; in particular is this decay exponential?

For 2-player games, i.e. for the case that $k = 2$, this problem is by now well understood, and it is known that the value of $\Psi^{\otimes t}$ is indeed exponentially decaying in t (however not in the most obvious manner); see [36, 25, 35, 8, 15]. The techniques that go into these proofs are either information theoretical, or analytical. In a sense, the analytical proofs are based on the well-known fact that the eigenvalues of a matrix tensorize when one tensorizes the matrix, as it turns out that, in a sense, the value of a game can be vaguely viewed as eigenvalues of a matrix.

For $k \geq 3$, the situation is much more complicated, and the only known bound for general games is due to Verbitsky [39] and gives rather weak bounds (as, once again, it relies on the density Hales-Jewett theorem).

Recently, the work of [13] identified a class of games referred to as “connected games” for which the information theoretic techniques from the setting of 2-player games still work, which sparked renewed interest in multi-player parallel repetition theorems. We remark that the notion of “connectedness” therein is very much similar to the notion of connectedness of distribution in our setting (which is much stronger than pairwise connectedness). This motivated a recent line of works [26, 21, 19, 20] that studied parallel repetition of 3-player games over binary questions. This line of work started with studying a game known as the GHZ game (which is well known in the physics literature and is a bottleneck to the techniques of [13]), proving polynomial decay for it, and using this as a stepping stone to prove polynomial decay parallel repetition theorems for more general classes of games.

We believe that the notion of Abelian embeddability should have a fundamental connection to the problem of parallel repetition in multiplayer games. In a sense, this question too is about “tensorization” of some value, but in this time one has to deal with k -dimensional tensors as opposed to matrices. Some evidence to that has been given in [10], wherein the authors give a very simple proof for the fact that the value of the GHZ game is exponentially vanishing with t (as opposed to just polynomial) which is inspired by Abelian embeddability. In a sense, the proof proceeds by identifying that the GHZ game actually entails within it a $(\mathbb{Z}_4, +)$ -type additive structure. Then, using this fact along with powerful theorems from additive combinatorics, the authors give a structural result on the set of strategies for the players that perform well, which are then analyzed directly.

While being speculative, we believe that such connection should indeed exist, and in it the quantitative

aspects of Theorem 1.4 should be highly relevant. At the current state, the quantitative bounds we get are not very good, but we believe that with more effort these could be improved to results that would be able to show $2^{-t^{\Omega(1)}}$ rate of decay in parallel repetition.

1.4 High Level Overview of the Proof of Theorem 1.4

In this section we give a high level overview of the proof of Theorem 1.4. As such, we often omit details, make simplifying assumptions and appeal to intuition in order to concentrate on the main ideas. We also point out the sections relevant to each part of the argument.

At its core, our argument relies on the following intuition: if μ does not admit any Abelian embedding, then Theorem 1.4 is just equivalent to the main result of [4, 5]. Thus, one idea is to try to identify all Abelian embeddings of μ , define partial basis for $L_2(\Sigma^n; \mu_x^{\otimes n})$, $L_2(\Gamma^n; \mu_y^{\otimes n})$ and $L_2(\Phi^n; \mu_z^{\otimes n})$ based on these Abelian embeddings and then show that for f, g and h to satisfy (1), it must be the case that they correlated with a function from the span of this partial basis. The intuition is completing the partial bases into complete bases, once we “peel off” these embeddings based functions the rest of the functions in the bases are “oblivious” to the fact that μ admits Abelian embeddings. So, once we “peel off” these embedding functions, the situation is analogous to the case that μ does not have any Abelian embeddings, in which case the result of [4, 5] kick in.

Much of the effort in our proof goes into formalizing this rough idea, and once one is able to do that the rest of the proof is more streamline (but still requires a significant effort). Below, we give step by step description of the way we formalize this intuition.

1.4.1 Step 1: Master Embedding

The first issue is a distribution μ may have multiple linear embeddings, and they may interact in a non-trivial way. Indeed, given an Abelian embedding (σ, γ, ϕ) of μ into H , one can define a partial basis by composing characters of H with the embedding functions. But how do we know that different embeddings give us different basis elements? How do we combine these partial bases into something that makes sense and is convenient to work with?

Our first step is to identify that one may define a single embedding, which we refer to as the master embedding, that encapsulates within it all of the Abelian embeddings of μ . Indeed, we show that if μ does not have any $(\mathbb{Z}, +)$ embedding, then there is a size M such that any Abelian embedding of μ “comes from” an Abelian embedding of μ into an Abelian group of size at most M . Hence, to include all Abelian embeddings it suffices to only look into embeddings of μ into Abelian groups of size at most M , and as there are only finitely many such embeddings we can just tensorize them. That is, letting $\sigma_i: \Sigma \rightarrow H_i$ be all possible σ 's in linear embeddings of μ , where $(H_i, +)$ are Abelian groups, the master embedding of x is

$$\sigma_{\text{master}}: \Sigma \rightarrow \prod_{i=1}^R H_i \text{ defined by}$$

$$\sigma_{\text{master}}(x) = (\sigma_1(x), \dots, \sigma_R(x)),$$

and similarly one may define γ_{master} and ϕ_{master} . With the master embeddings in hand we now have a sensible way of defining a partial basis for functions in x, y and z by considering compositions of characters from H with the master embeddings.

At the present state, this partial basis is not very convenient. For example, it may well be the case that there are distinct $\chi, \chi' \in H$ such that $\chi \circ \sigma_{\text{master}} = \chi' \circ \sigma_{\text{master}}$. Indeed, this would be the case if the image of σ_{master} was a strict subgroup of H . More generally, linear dependencies within $\{\chi \circ \sigma_{\text{master}}\}_{\chi \in \hat{H}}$ already

start appearing as soon as the image of σ_{master} is not the entire group H , and this presents issues which we wish to avoid.

This part of the argument appears in Section 3.

1.4.2 Step 2: Saturating the Master Embeddings

Our goal is therefore to arrange for the master embeddings σ_{master} , γ_{master} and ϕ_{master} to be *saturated*, meaning that the image of each one of them is the entire group H . To do so, we must change the distribution μ into a distribution μ' such that (a) on μ' the master embeddings are saturated, (b) there is a good enough relationship between 3-wise correlations over μ and 3-wise correlations over μ' , and (c) we can deduce the conclusion of Theorem 1.4 on μ from the conclusion of Theorem 1.4 on μ' .

This transformation is achieved via the path trick, introduced in [4], which is ultimately just an application of the Cauchy-Schwarz inequality. The path trick is used in our arguments extensively, and often time the structure we need is quite subtle thereby requiring a very careful application of the path-trick. Nevertheless, below we explain at a high level the intuition behind the path trick and what it achieves.

Given a distribution μ , the path trick distribution (of length $2t + 1$) with respect to x can be described as the following distribution μ' :

1. Sample $(x_1, y_1, z_1) \sim \mu$.
2. Make a step from y : sample $(x'_1, y'_1, z'_1) \sim \mu$ conditioned on $y'_1 = y_1$.
3. Make a step from z : sample (x_2, y_2, z_2) conditioned on $z_2 = z'_1$.
4. Repeat make a step from y / make a step from z for t times.

Thus, the sequences $(y_1, y'_1, y_2, y'_2, \dots, y_t, y'_t, y_{t+1})$ of y 's and $(z_1, z'_1, z_2, z'_2, \dots, z_t, z'_t, z_{t+1})$ z 's are generated (where $z_{i+1} = z'_i$ and $y'_i = y_i$), as well as a sequence $\vec{x} = (x_1, x'_1, \dots, x_t, x'_t, x_{t+1})$ of x 's. The output of the distribution ν is (\vec{x}, y_{t+1}, z_1) , and it is thought of as a 3-ary distribution over $\Sigma' \times \Gamma \times \Phi$ where $\Sigma' \subseteq \Sigma^{2t+1}$ is the set of feasible tuples of x in the procedure.

We refer to this procedure as the path trick since one may consider the bi-partite graph $G = (\Gamma \cup \Phi, E)$ whose edges are $(y, z) \in \Gamma \times \Phi$ that are in the support of $\mu|_{\Gamma \times \Phi}$. Thus, thinking of the x 'es as labeling the edges of G , namely labeling an edge (y, z) by x if $(x, y, z) \in \text{supp}(\mu)$, one gets that the above procedure generates a random path of length $2t + 1$ in the graph and record the labels of the edges that it traversed on.

Moving from the distribution μ to μ' has several benefits that have been used in our earlier papers:

1. **Improving connectivity:** if μ is $\{2, 3\}$ -connected, then for large enough t the support of μ' on the last two coordinates is full. Indeed, taking the random path view of the path trick, it is clear that as the graph G is connected, for sufficiently large t the same graph corresponding to μ' would be a complete bipartite clique.
2. **Preserving properties of μ :** the distribution μ' preserves much of the properties of the distribution μ . In particular, if μ is pairwise connected then so is μ' , and if μ does not admit any Abelian embeddings, then so does μ' .
3. **The 3-wise correlations relations:** 3-wise correlations of functions over μ can be upper bounded by 3-wise correlations of functions related to the original functions over μ' . Indeed, assume for simplicity that the functions are real valued. If f, g and h achieve large 3-wise correlation in μ , then

for $(x, y, z) \sim \mu^{\otimes n}$ one has that the values $h(z)$ and $f(x)g(y)$ are correlated, so looking at the above path we get that $h(z_{i+1}) \approx f(x_{i+1})g(y_{i+1})$ and $g(y_{i+1}) \approx f(x'_{i+1})h(z'_{i+1})$ and combining these we get that

$$h(z_{t+1}) \approx f(x'_{t+1})h(z'_{t+1}) \approx f(x'_{t+1})f(x_t)g(y_t) \approx \dots \approx f(x'_{t+1})f(x_t) \cdots f(x_2)f(x_1)g(y_1),$$

and hence we expected g, h and $F(\vec{x}) = f(x'_{t+1})f(x_t) \cdots f(x_2)f(x_1)$ to achieve a significant correlation in μ' . Indeed, this can be proved via an appropriate application of the Cauchy-Schwarz inequality.

For the purposes of this paper we need additional properties of the path trick transformations, which we explain next.

1. **Abelian embeddings of μ lift to Abelian embeddings of μ' :** not only does the path trick preserve lack of Abelian embeddings, but in fact if μ does admit Abelian embeddings, then μ' does not introduce any new ones. To be more precise, suppose that $\sigma: \Sigma \rightarrow H$, $\gamma: \Gamma \rightarrow H$ and $\phi: \Phi \rightarrow H$ are Abelian embeddings of μ . Then, these embeddings give rise to an Abelian embedding $\sigma_t: \Sigma' \rightarrow H$ with γ and ϕ of ν , as follows:

$$\sigma_t(\vec{x}) = \sum_{i=1}^t \sigma(x_i) - \sigma(x'_i) + \sigma(x_{t+1}). \quad (3)$$

With the notation above, we have that $\sigma(x_i) + \gamma(y_i) + \phi(z_i) = 0$, $\sigma(x'_i) + \gamma(y'_i) + \phi(z'_i) = 0$, and doing a proper addition/ subtraction one gets that

$$\sigma_t(\vec{x}) + \gamma(y_{t+1}) + \phi(z_1) = 0,$$

hence (σ_t, γ, ϕ) form an Abelian embedding of μ' into H .

2. **The only Abelian embeddings of μ' are lifts of Abelian embeddings of μ :** all Abelian embeddings of μ' are precisely of this form. Namely, for any Abelian embedding (σ', γ, ϕ) of μ' into an Abelian group $(H, +)$ there is an Abelian embedding (σ, γ, ϕ) of μ into $(H, +)$ where σ satisfies a relation as in (3) where σ' plays the role of σ_t . (see Lemma 3.10). This result has a few important consequences, and in particular it says that the path trick preserves master embeddings. Namely, if we start with a master embedding of μ , apply the path trick and the above transformation corresponding to it on the embeddings, then we will get the master embedding of μ' .
3. **Saturating the embeddings:** it can be easily observed that if $(\sigma_{\text{master}}, \gamma_{\text{master}}, \phi_{\text{master}})$ is a master embedding of μ (or for this purpose, any embedding of μ), then after the path trick we get the embedding $(\sigma_{\text{master}, t}, \gamma_{\text{master}}, \phi_{\text{master}})$ that clearly satisfies that $\text{Image}(\sigma_{\text{master}, t}) \subseteq \text{Image}(\sigma_{\text{master}})$; this follows by looking at trivial paths that traverse the same edge back and fourth and use the same label of x all of the time. Moreover, it is clear that if $\text{Image}(\sigma_{\text{master}})$ was a sub-group of H then we would have that $\text{Image}(\sigma_{\text{master}, t}) = \text{Image}(\sigma_{\text{master}})$. It stands to reason that unless $\text{Image}(\sigma_{\text{master}})$ is indeed a subgroup, then for large enough t we would have that $\text{Image}(\sigma_{\text{master}, t}) \subsetneq \text{Image}(\sigma_{\text{master}})$, in which case we enlarged the image of σ_{master} via the path trick.

Indeed, something along these lines is true. Namely, we show that by combination of path tricks along different directions (not only x) one can indeed always enlarge the image of an embedding so long as it is not a subgroup.¹

¹In our formal proof this has to be done rather carefully as we wish to preserve the property that the alphabet of x is always a power of the original alphabet Σ .

In conclusion, using the path trick multiple times we can pass to a new distribution ν on which the embeddings are all saturated, 3-wise correlations over μ are upper bounded by 3-wise correlations over ν , and ν has improved connectivity – say that its support on the last two coordinates is full. It can be easily shown that in that case, the images of all of the components must be the same sub-group, and without loss of generality we assume it is the group H itself.

Note that in particular, the above properties mean that if $(\sigma_{\text{master}}, \gamma_{\text{master}}, \phi_{\text{master}})$ is a saturated master embedding of ν , then the distribution of $(\sigma_{\text{master}}(x), \gamma_{\text{master}}(y), \phi_{\text{master}}(z))$ where $(x, y, z) \sim \nu$ has a full support on

$$\{(a, b, c) \in H^3 \mid a + b + c = 0\},$$

which intuitively says that by moving from μ to ν we have “exposed” all of the Abelian structure in the distribution μ .

This part of the argument appears in Section 3.

1.4.3 Step 3: Setting up a Basis Consisting of Embedding and Non-embedding Functions

Fix distributions μ over $\Sigma \times \Gamma \times \Phi$ and ν over $\Sigma' \times \Gamma' \times \Phi'$ as we have done so far, and suppose that (a power of) the 3-wise correlation of f, g and h over μ is upper bounded by the 3-wise correlation of F, G and H over ν .

Now that we have saturated the master embeddings in ν we can set up a partial for functions in $x \in \Sigma'$ as basis as before $\tilde{\chi}(x) = \chi(\sigma_{\text{master}}(x))$ for all $\chi \in \hat{H}$ as before and get that now these functions are indeed linearly independent. We can then complete it to a basis of $L_2(\Sigma'; \nu_x)$ by adding to it the functions ψ_1, \dots, ψ_s that are orthogonal to all functions in $\text{Span}(B_1)$, so that writing $B_1 = \{\tilde{\chi} \mid \chi \in \hat{H}\}$ and $B_2 = \{\psi_1, \dots, \psi_s\}$ we have a basis $B_1 \cup B_2$ for $L_2(\Sigma'; \nu_x)$. Tensorizing, we get that $\{v_{\vec{b}}\}_{\vec{b} \in (B_1 \cup B_2)^{\otimes n}}$ where $v_{\vec{b}}: \Sigma'^n \rightarrow \mathbb{C}$ is defined by $v_{\vec{b}}(x) = \prod_{i=1}^n v_{b_i}(x_i)$, is a basis for $L_2(\Sigma'^n; \nu_x^{\otimes n})$. Thus, we can write

$$F(x_1, \dots, x_n) = \sum_{\alpha \in (B_1 \cup B_2)^n} \hat{F}(\alpha) v_{\alpha}(x).$$

We can define analogous bases for $L_2(\Gamma'; \nu_y)$ and $L_2(\Phi'; \nu_z)$. Now, each one of the functions F, G and H has an “embedding part”, which is the parts of the monomials that use functions from B_1 , and “non-embedding parts”, which are monomials using functions from B_2 . Intuitively, it should be the case that the more mass the functions have on the non-embedding parts, the smaller the 3-wise correlations would be; this is because that for uni-variate functions $u: \Sigma' \rightarrow \mathbb{C}$, $v: \Gamma' \rightarrow \mathbb{C}$ and $w: \Phi' \rightarrow \mathbb{C}$ of 2-norm 1, to achieve perfect 3-wise correlation it must be the case that $u(x) = v(y)w(z)$ in the support of ν , in which case u, v and w behave like an embedding function. We remark that there is a serious leap in this last step, which causes complications in later points of the argument. Later on, we refer to this gap as the Horn-SAT obstruction, and we will explain how it arises and how to overcome it later on.

In light of the above, it makes sense to define two notions of degrees for our partial basis. The first of which is the embedding degree of a monomial $v_{\vec{b}}$, $\text{embeddeg}(v_{\vec{b}})$, which is the number of components v_{b_i} that come from the partial embedding basis B_1 . The second of which is the non-embedding degree of a monomial $v_{\vec{b}}$, $\text{nedeg}(v_{\vec{b}})$, which is the number of components of v_{b_i} that come from B_2 .

This part of the argument appears in Section 4.

1.4.4 Step 4: Analyzing the Contribution of High Non-embedded Degree Components

The above discussion suggests that the parts of F , G and H of high non-embedding degree should contribute very little to their 3-wise correlation according to ν . Formally showing this, however, is quite tricky and this is where a considerable amount of effort in this paper is devoted to. Our argument here builds on an argument from [5] and it is the main subject in Sections 4, 5, 6, 7, 8, 9.

To give some intuition for the argument we make several simplifying assumptions (some of which can be ensured, while other are not necessary). Assume that the marginal distribution of ν over y, z is uniform, and that the distribution of $(\sigma_{\text{master}}(x), \gamma_{\text{master}}(y), \phi_{\text{master}}(z))$ where $(x, y, z) \sim \nu$ is uniform over on

$$\{(a, b, c) \in H^3 \mid a + b + c = 0\}.$$

Further assume that the functions F , G and H are embedding homogenous and non-embedding homogenous functions, by which we mean that the embedding degree of each monomial of F is the same, and also the non-embedding degree of each monomial in F is the same; the same goes for G and H . Our argument here will be inductive on the number of coordinates n , and we show that the 3-wise correlation of functions F , G and H as above can be upper bounded by either the 3-wise correlation of $n - 1$ variate functions of the same non-embedding degree, or by $(1 - \Omega(1))$ times the 3-wise correlation of $n - 1$ variate functions with non-embedding degree smaller by 1. Thus, iterating we would ultimately get a bound of $(1 - \Omega(1))^{\text{ndeg}(F)}$ on the 3-wise correlations, which is small if the non-embedding degree of F is high.

In fact, we have two separate inductive arguments depending on if n is much larger than the non-embedding degree of F , or if it is comparable to it; we refer to this last case as the ‘‘near linear non-embedding degree case’’, and we now elaborate on each one of these cases.

The case that n is much larger than the non-embedding degree of F . In this case there is a variable, say the n th variable, such that in almost all of the mass of F lies on monomials in which the component of x_n is an embedding function. Using the homogeneity of F we can use find a decomposition of F as

$$\sum_{t \in T} \psi_t F_t(x_1, \dots, x_{n-1}) F'_t(x_n)$$

where each F'_t is either from B_1 or from B_2 , and $\{F_t\}, \{F'_t\}$ form orthonormal sets and $\sum_{t \in T} |\psi_t|^2 = 1$.

Similarly, we can find analogous decompositions form G and H as

$$\sum_{r \in R} \kappa_r G_r(y_1, \dots, y_{n-1}) G'_r(y_n), \quad \sum_{s \in S} \rho_s H_s(z_1, \dots, z_{n-1}) H'_s(z_n).$$

Moreover, if F'_t is a function from B_1 then F_t has the same non-embedding degree as F , and if F'_t is from B_2 then F_t has one smaller non-embedding degree. The same goes for G and H , so to simplify presentation we consider the specialized case where

$$\begin{aligned} F(x) &= \psi_1 F_1(x_1, \dots, x_{n-1}) F'_1(x_n) + \psi_2 F_2(x_1, \dots, x_{n-1}) F'_2(x_n), \\ G(y) &= \kappa_1 G_1(y_1, \dots, y_{n-1}) G'_1(y_n) + \kappa_2 G_2(y_1, \dots, y_{n-1}) G'_2(y_n), \\ H(y) &= \rho_1 H_1(z_1, \dots, z_{n-1}) G'_1(z_n) + \rho_2 H_2(z_1, \dots, z_{n-1}) H'_2(z_n), \end{aligned}$$

where F'_1, G'_1 and H'_1 are embedding functions and F'_2, G'_2 and H'_2 are non-embedding functions. Thus, the coefficient ψ_2 is related to the mass x_n has on non-embedding functions and by choice is therefore small,

and similarly we can expect it to be the case that κ_2 and ρ_2 are also small (which is true, but requires some preparatory work). Thus, the 3-wise correlation of F , G and H according to ν can be written as

$$\begin{aligned} \mathbb{E}_{\nu^n} [FGH] &= \psi_1 \kappa_1 \rho_1 \mathbb{E}_{\nu^{\otimes n-1}} [F_1 G_1 H_1] \mathbb{E}_{\nu} [F'_1 G'_1 H'_1] + \psi_1 \kappa_1 \rho_2 \mathbb{E}_{\nu^{\otimes n-1}} [F_1 G_1 H_2] \mathbb{E}_{\nu} [F'_1 G'_1 H'_2] \\ &+ \psi_1 \kappa_2 \rho_1 \mathbb{E}_{\nu^{\otimes n-1}} [F_1 G_2 H_1] \mathbb{E}_{\nu} [F'_1 G'_2 H'_1] + \psi_1 \kappa_2 \rho_2 \mathbb{E}_{\nu^{\otimes n-1}} [F_1 G_2 H_2] \mathbb{E}_{\nu} [F'_1 G'_2 H'_2] \\ &+ \psi_2 \kappa_1 \rho_1 \mathbb{E}_{\nu^{\otimes n-1}} [F_2 G_1 H_1] \mathbb{E}_{\nu} [F'_2 G'_1 H'_1] + \psi_2 \kappa_1 \rho_2 \mathbb{E}_{\nu^{\otimes n-1}} [F_2 G_1 H_2] \mathbb{E}_{\nu} [F'_2 G'_1 H'_2] \\ &+ \psi_2 \kappa_2 \rho_1 \mathbb{E}_{\nu^{\otimes n-1}} [F_2 G_2 H_1] \mathbb{E}_{\nu} [F'_2 G'_2 H'_1] + \psi_2 \kappa_2 \rho_2 \mathbb{E}_{\nu^{\otimes n-1}} [F_2 G_2 H_2] \mathbb{E}_{\nu} [F'_2 G'_2 H'_2]. \end{aligned} \quad (4)$$

It turns out that terms the only term involving F'_1 that does not vanish is $\mathbb{E}_{\nu} [F'_1 G'_1 H'_1]$. Indeed, as F'_1 is a function from B_1 we may write it as a product of a function on y with a function on z , and thus expectations such as $\mathbb{E}_{\nu} [F'_1 G'_1 H'_2]$ can be written as expectation of product of a function of y and a function of z . Using independence, this product can be further written as the product of two expectations where at least one of these expectations is 0.

Thus, if the terms involving F_2 were not existent, then we would get the upper bound

$$\left| \mathbb{E}_{\nu^n} [FGH] \right| \leq |\psi_1 \kappa_1 \rho_1| \left| \mathbb{E}_{\nu^{\otimes n-1}} [F_1 G_1 H_1] \right| \left| \mathbb{E}_{\nu} [F'_1 G'_1 H'_1] \right| \leq \left| \mathbb{E}_{\nu^{\otimes n-1}} [F_1 G_1 H_1] \right|,$$

and we have decreased the number of variables n by 1 (while keeping the non-embedding degree. In a sense, as ψ_2 is small this term indeed should constitute the majority of the contribution to $\mathbb{E}_{\nu^{\otimes n}} [FGH]$, but we cannot just ignore the contribution from the other terms.

A naive attempt at bounding the other term (and using the Cauchy-Schwarz in a favorable way) can show that $|\mathbb{E}_{\nu^{\otimes n}} [FGH]|$ is at most the maximum of $|\mathbb{E}_{\nu^{\otimes n-1}} [F_i G_j H_k]|$ over i, j and k , however this bound is not good enough for us; indeed, if this maximum is achieved at anywhere other than $i = j = k = 1$ then the non-embedding degrees decrease, and in that case we must gain a factor of $(1 - \Omega(1))$ for our argument to go through.

The key to improve upon this naive attempt lies in what we refer to as the ‘‘additive base case’’. The additive base case is a statement about univariate functions that helps us to control the contribution from terms involving F'_2 in a favorable way. Stated simply, the additive base case we use is the statement that if F' is a univariate non-embedding function, and G', H' are any univariate functions, then

$$\left| \mathbb{E}_{\nu} [F'(G' + H')] \right| \leq (1 - \Omega(1)) \|F'\|_2 \|G' + H'\|.$$

The intuition for this inequality is that otherwise, the value of F' would be very close to the value of $\overline{G'} + \overline{H'}$, but this is only possible for embedding functions.

The point of the additive base case is that except for $\mathbb{E}_{\nu} [F'_2 G'_2 H'_2]$, the contribution of the terms not involving F'_1 in (4) may be re-casted as an expectation of the form dealt with in the additive base case. Indeed, if F'_1, G'_1 and H'_1 were the simplest of embedding functions – namely constant functions – then this is rather clear, as these terms can be written as

$$\mathbb{E}_{\nu} [F'_2(aG_2 + bH_2)]$$

for some coefficients a and b . In the case F'_1, G'_1 and H'_1 are not the constant functions more effort is needed, and in particular one needs to guarantee that they ‘‘come’’ from the same character of H . Namely, that there

is $\chi \in \hat{H}$ such that F'_1, G'_1 and H'_1 are multiples of $\chi \circ \sigma_{\text{master}}, \chi \circ \gamma_{\text{master}}$ and $\chi \circ \phi_{\text{master}}$ respectively. As we are able to guarantee this fact, the contribution of these terms can still be associated with the additive base case, as essentially $G'_1 = \overline{F'_1 H'_1}$ and $H'_1 = \overline{F'_1 G'_1}$. Hence, that contribution can be re-written as

$$\mathbb{E}_{\nu} [F'_1 F'_2 (aG_2 + bH_2)]$$

for some coefficients a and b , and this is still an expectation of the form handled by the additive base case.

Making an effective enough use of the additive base case, one can use 4 to either reduce n by 1 and keep the non-embedding degrees the same, or else reduce both n and the non-embedding degrees by 1 and then gain a factor of $1 - \Omega(1)$.

This part of the argument is presented in Section 8.

The near linear non-embedding degree case. Once the non-embedding degree of F, G and H is comparable to n , the above argument no-longer works, and the reason is that the last term in (4) is no longer negligible, and at the same time we do not know how to give an effective upper bound on it using only the additive base case. Thus we must have a new base case that handles this last term, and intuitively one would expect the following base case to hold. Suppose that $\tilde{F}: \Sigma' \rightarrow \mathbb{C}$ is a function from B_2 , and $\tilde{G}: \Gamma' \rightarrow \mathbb{C}$, $\tilde{H}: \Phi' \rightarrow \mathbb{C}$ are any functions of 2-norm 1. Then

$$\left| \mathbb{E}_{(x,y,z) \sim \nu} [\tilde{F}(x)\tilde{G}(y)\tilde{H}(z)] \right| \leq 1 - \Omega(1).$$

The reason we expect this to be true is that otherwise, by compactness we would be able to find 3 such functions satisfying $\tilde{F}(x) = \tilde{G}(y)\tilde{H}(z)$ on the support of ν . Thus, the logs of these functions form an Abelian embedding so $\log(\tilde{F})$ must be an embedding function, in which case \tilde{F} is also an embedding function in contradiction.

As stated, this argument is not quite correct, since it may be the case that the function F' gets the value 0 sometimes, in which case we cannot take logs and get away with it. This obstruction has already appeared in [5] wherein it was referred to as the ‘‘Horn-SAT obstruction’’, and here too we have to face it. In fact, as in our scenario we need to maintain many more properties of the distribution ν (compared to what was necessary in [5]), more care is needed to handle the Horn-SAT obstruction. Ultimately, the Horn-SAT obstruction is dealt with by stating a more complicated base case statement which we can guarantee to hold for the distribution ν while being useful enough to make our argument go through. For the simplicity of presentation however, we ignore this obstruction for now and explain the argument in the case the ideal base case holds.

Equipped with the ideal base case, we can give effective enough bounds on the last term in (4). In particular, if all of the contribution came from it, we would have been able to conclude that

$$\left| \mathbb{E}_{\nu^{\otimes n}} [FGH] \right| \leq (1 - \Omega(1)) \left| \mathbb{E}_{\nu^{\otimes n-1}} [F_2 G_2 H_2] \right|,$$

and iterating would finish the proof. One again however, there are other terms in (4) that need to be accounted for (the other terms involving F_2). To do so, ideally we would have liked (as in the additive base case) to re-arrange these terms so as to view their total contribution as an instantiation of the ideal base case, but this is not possible. Using a similar (but more complicated) argument we can still show that it is in fact the case that

$$\left| \mathbb{E}_{\nu^{\otimes n}} [FGH] \right| \leq (1 - \Omega(1)) \left| \mathbb{E}_{\nu^{\otimes n-1}} [\tilde{F}_2 \tilde{G}_2 \tilde{H}_2] \right|,$$

where \tilde{F}_2, \tilde{G}_2 and \tilde{H}_2 are functions of non-embedding degree at most 1 less of F_2, G_2 and H_2 .

This part of the argument is presented in Section 9.

Overcoming the Horn-SAT Obstruction. The bulk of Sections 4, 5, 6, 7 is devoted to gaining additional properties of ν , as well as other crucial reductions (for example, to allow us to assume homogeneity of F, G and H as above). One of the key properties achieved in this section is the so-called relaxed base case, which is a replacement for the ideal base case above that we are able to ensure.

A triplet of functions $u: \Sigma' \rightarrow \mathbb{C}$, $v: \Gamma' \rightarrow \mathbb{C}$ and $w: \Phi' \rightarrow \mathbb{C}$ is called a Horn-SAT embedding if $u(x) = v(y)w(z)$ in the support of ν . If u never vanishes, a Horn-SAT embedding can be transformed into an Abelian embedding, and thus (simply put) the Horn-SAT obstruction is really about the possible 0-patterns non-embedding functions may have. By careful manipulations (once again utilizing the path trick) we are able to find a set $\Sigma_{\text{modest}} \subseteq \Sigma'$ of size at least 2 such no Horn-SAT embedding can vanish on. Thus, we get that if u doesn't vanish on Σ_{modest} then it can never be a part of a Horn-SAT embedding. Therefore, it is natural to expect that if u has variance at least τ on Σ_{modest} , then

$$\left| \mathbb{E}_{\nu} [uvw] \right| \leq 1 - \theta(\tau)$$

where $\theta(\tau) > 0$ is some function of τ . This turns out to be true and useful, but there are many subtleties. For once, we need additional properties from Σ_{modest} to make this relaxed base case useful, and most important we need the symbols in Σ_{modest} to be mapped to the same group element in H by the master embedding. Secondly (and this has already appeared in [5]) we need a decent dependency between τ and $\theta(\tau)$.

1.4.5 Step 5: Reducing to Functions over H

We now return our functions F, G and H , armed with the knowledge that the contribution of high non-embedding degree parts is small. Thus, taking their parts of small non-embedding degree F', G' and H' , we are able to conclude that $|\mathbb{E}_{\nu^{\otimes n}} [FGH] - \mathbb{E}_{\nu^{\otimes n}} [F'G'H']| \leq \frac{\epsilon}{100}$, and so the 3-wise correlation of F', G' and H' according to ν is still significant.

We remark that as in our actual argument we will need the functions F', G' and H' to be bounded, so harsh truncations as we described do not fit the bill. Thus, we use a softer notion of truncations given by the *non-embedding noise operator*. For $\rho \in [0, 1]$ consider the Markov chain $T_{\text{non-embed}, \rho}$ on Σ' that on x , with probability ρ takes $x' = x$, and otherwise samples $x' \sim \nu|_x$ conditioned on $\sigma_{\text{master}}(x') = \sigma_{\text{master}}(x)$. When ρ is not specified, that is, when we write $T_{\text{non-embed}}$, we mean that ρ is taken to be 0. Given such Markov chain one may consider the corresponding averaging operator on $L_2(\Sigma', \nu_x)$ given as

$$T_{\text{non-embed}, \rho} f(x) = \mathbb{E}_{x' \sim T_{\text{non-embed}, \rho} x} [f(x')],$$

and tensorize it to get an averaging operator $T_{\text{non-embed}, \rho}^{\otimes n}: L_2(\Sigma'^n, \nu_x^{\otimes n}) \rightarrow L_2(\Sigma'^n, \nu_x^{\otimes n})$. Similarly, we can get averaging operators on $L_2(\Gamma'^n, \nu_y^{\otimes n})$ and $L_2(\Phi'^n, \nu_z^{\otimes n})$, which by abuse of notation we also denote by $T_{\text{non-embed}, \rho}$. These averaging operators can be shown to essentially kill monomials of high non-embedding degree, hence serve as a replacement for harsh truncation arguments as above.

With these operators in hand, we can replace the harsh truncations above by $F' = T_{\text{non-embed}, \rho}^{\otimes n} F$, $G' = T_{\text{non-embed}, \rho}^{\otimes n} G$ and $H' = T_{\text{non-embed}, \rho}^{\otimes n} H$ (for suitably chosen ρ) and effectively be in the same situation as before, wherein we have functions F', G' and H' that have almost all of their mass on monomials with small non-embedding degrees, and also that $|\mathbb{E}_{\nu^{\otimes n}} [FGH] - \mathbb{E}_{\nu^{\otimes n}} [F'G'H']| \leq \frac{\epsilon}{100}$

We wish to transform the functions F' , G' and H' into related bounded functions with non-embedding degree 0 for which the 3-wise correlation over ν is still significant. For that, we use a combination of random restrictions (so as the mass of F' , G' , and H' of small but not 0 non-embedding degree would almost all collapse to level 0), followed by averaging (to get rid of all monomials of positive non-embedding degree). Thus, we get functions $F''' = T_{\text{non-embed}}(F'')$, $G''' = T_{\text{non-embed}}(G'')$ and $H''' = T_{\text{non-embed}}(H'')$ where F'' , G'' and H'' are random restrictions of F' , G' and H' , so that with noticeable probability we have that

$$\mathbb{E}_{(x,y,z) \sim \nu^{\otimes n'''}} [F'''(x)G'''(y)H'''(z)] \geq \frac{\varepsilon}{2}$$

where n''' is the number of coordinates left alive after the random restriction. Now the functions F''' , G''' and H''' can be viewed as functions defined over $H^{n'''}$, so the above expectation should be amenable to standard tools from discrete Fourier analysis.

This part of the argument appears in Section 10.

1.4.6 Step 6: Applying the Linearity Testing Argument

Indeed, we re-cast the functions F''' , G''' and H''' above as $F^\sharp: H^{n'''} \rightarrow \mathbb{C}$, $G^\sharp: H^{n'''} \rightarrow \mathbb{C}$ and $H^\sharp: H^{n'''} \rightarrow \mathbb{C}$ defined in the natural way (for example, $F^\sharp(a) = F(x)$ for x such that $\sigma_{\text{master}}(x_i) = a_i$ for each coordinate, where we know that the specific choice of x doesn't matter). Thus, from the distribution ν we get a corresponding distribution ν^\sharp over $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$ whose support is full, and

$$\left| \mathbb{E}_{(a,b,c) \sim (\nu^\sharp)^{\otimes n'''} } [F^\sharp(a)G^\sharp(b)H^\sharp(c)] \right| \geq \frac{\varepsilon}{2}.$$

We now use random restrictions again, but for a different reason. Namely, we use random restrictions to shift from the distribution ν^\sharp to the uniform over $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$, and get from F^\sharp , G^\sharp and H^\sharp restrictions $F^{\sharp'}$, $G^{\sharp'}$ and $H^{\sharp'}$ so that with noticeable probability

$$\left| \mathbb{E}_{a,b \in H^{n''''}} [F^{\sharp'}(a)G^{\sharp'}(b)H^{\sharp'}(-a-b)] \right| \geq \frac{\varepsilon}{4},$$

where n'''' is the number of coordinates left alive. In this case, a straightforward, classical Fourier analytic computation can be applied to relate the left hand side to the Fourier coefficients of $F^{\sharp'}$, $G^{\sharp'}$ and $H^{\sharp'}$ so that we get

$$\left| \sum_{\chi \in \hat{H}^{\otimes n''''}} \widehat{F^{\sharp'}}(\chi) \widehat{G^{\sharp'}}(\chi) \widehat{H^{\sharp'}}(\chi) \right| \geq \frac{\varepsilon}{4},$$

from which one can quickly conclude that there is χ such that $\left| \widehat{F^{\sharp'}}(\chi) \right| \geq \frac{\varepsilon}{4}$. In words, after a sequence of random restrictions, averaging and further random restriction, the function F is correlated with a function of the form $\chi \circ \sigma_{\text{master}}$. This is the type of result we are after, except that we wish to have such result for F and not for F after this sequence of operations.

This part of the argument appears in Section 10.

1.4.7 Step 7: Going back to F via the Restriction Inverse Theorem

We have thus concluded that after random restriction, F^\sharp is correlated with a function of the form $\chi \circ \sigma_{\text{master}}$ where $\chi \in \hat{H}^{n''''}$, and we wish to unravel the steps we took to get from F to F^\sharp and conclude a structural result about F .

Noting that $\chi \circ \sigma_{\text{master}}$ is a product function, this is precisely a situation in which the restriction inverse theorem kicks in, and using a modified version of Theorem 1.5 we are able to conclude that F^\sharp is correlated with a function of the form $L \circ \sigma_{\text{master}} \cdot \chi \circ \sigma_{\text{master}}$ where $\chi \in \hat{H}^{n''''}$ and L is a low-degree function of 2-norm at most 1. Thus, the same conclusion holds for F''' (as it is essentially the same function as F^\sharp).

Recalling that $F''' = T_{\text{non-embed},0} F''$, we get that

$$\left| \langle T_{\text{non-embed},0} F'', L \circ \sigma_{\text{master}} \cdot \chi \circ \sigma_{\text{master}} \rangle \right| \geq \frac{\varepsilon}{4},$$

but on the other hand we also have that

$$\begin{aligned} \left| \langle T_{\text{non-embed},0} F'', L \circ \sigma_{\text{master}} \cdot \chi \circ \sigma_{\text{master}} \rangle \right| &= \left| \langle F'', T_{\text{non-embed},0}^* (L \circ \sigma_{\text{master}} \cdot \chi \circ \sigma_{\text{master}}) \rangle \right| \\ &= \left| \langle F'', T_{\text{non-embed},0} (L \circ \sigma_{\text{master}} \cdot \chi \circ \sigma_{\text{master}}) \rangle \right| \\ &= \left| \langle F'', L \circ \sigma_{\text{master}} \cdot \chi \circ \sigma_{\text{master}} \rangle \right|, \end{aligned}$$

where we used the fact that $T_{\text{non-embed},0}$ is self-adjoint. Hence, we conclude that F'' is correlated with $L \circ \sigma_{\text{master}} \cdot \chi \circ \sigma_{\text{master}}$.

We now wish to unravel the last step of random restriction (that goes from F to F''), and for that we once again want to appeal to the restriction inverse theorem. However, the correlations we are talking about now are not quite about correlations with product functions. Amusingly, to circumvent this issue we apply more random restrictions. Intuitively, after a suitably chosen random restriction, the function L becomes close to constant, hence one expects the fact that F'' is correlated with $L \circ \sigma_{\text{master}} \cdot \chi \circ \sigma_{\text{master}}$ to convert to the fact that a random restriction of F'' is correlated with a restriction of $\chi \circ \sigma_{\text{master}}$ (which is a product function), and we show that this is indeed the case. Thus, we conclude that a random restriction of F'' is correlated with a function of the form $\chi \circ \sigma_{\text{master}}$. Noting that a random restriction of F'' is (overall) a random restriction of F (with different parameters), we are thus able to conclude from the restriction inverse theorem that F itself is correlated with a function of the form $L \cdot \chi \circ \sigma_{\text{master}}$.

This part of the argument appears in Section 10.

1.4.8 Step 8: Going back to f via Properties of the Master Embedding

The last step in the proof of Theorem 1.4 is to use the structural result obtained for the function F to deduce a similar structural result for f . For that, we recall that (ignoring complex conjugates) the value of $F(x_1, \dots, x_s)$ is $f(x_1) \cdots f(x_s)$, and, ignoring the low-degree part L for now, we know that F is correlated with $\chi \circ \sigma_{\text{master}}$ for some $\chi \in \hat{H}^n$. Recalling the relation 3, one quickly gets from it that

$$\chi \circ \sigma_{\text{master}}(x_1, \dots, x_s) = \prod_{i=1}^s \chi \circ \sigma_{\text{master}}(x_i),$$

where, by abuse of notation, σ_{master} on the right hand side is the master embedding of μ (which is the original distribution, prior to any application of the path trick). Hence, the correlation between F and $\chi \circ \sigma_{\text{master}}$

translates to the fact that

$$\left| \mathbb{E}_{(x_1, \dots, x_s) \sim \nu_x^{\otimes s}} \left[\prod_{i=1}^s f(x_i) \chi \circ \sigma_{\text{master}}(x_i) \right] \right| \geq \varepsilon' = \varepsilon'(\varepsilon) > 0$$

As discussed earlier, in [30] it is shown that if ν_x is a connected distribution, a correlation such as in the above can be noticeable only if $f \cdot \chi \circ \sigma_{\text{master}}$ is correlated with a low-degree function. Thus, the proof would be concluded if we are able to ensure connectivity of ν_x , which we are indeed able to. This requires some care in some of our earlier steps, and most notably in the way we apply the path trick. In fact, we are able to guarantee that the support of ν_x is full, that is, Σ^s .

Bringing the low-degree part L back, essentially the same argument works except that we need to apply a suitable random restriction beforehand to get rid of the low-degree part. Thus, the previous argument gives that with noticeable probability, a random restriction of $f \cdot \chi \circ \sigma_{\text{master}}$ is correlated with a low-degree function. Hence, after more random restrictions, we conclude that with noticeable probability a random restriction of $f \cdot \chi \circ \sigma_{\text{master}}$ is correlated with a constant function. Re-phrasing, this means that with noticeable probability a random restriction of f is correlated with a function of the form $\chi \circ \sigma_{\text{master}}$, and a final invocation of the restriction inverse theorem finishes the proof.

This part of the argument appears in Section 11.

2 Preliminaries

Notations. We denote $[n] = \{1, \dots, n\}$. For a vector $x \in \Sigma^n$ and a subset $I \subseteq [n]$ of coordinates, we denote by x_I the vector in Σ^I which results by dropping from x all coordinates outside I . We denote by x_{-I} the vector in $\Sigma^{n-|I|}$ resulting from dropping from x all coordinates from I ; if $I = \{i\}$ we often simplify the notation and write it as x_{-i} . For $I \subseteq [n]$, $a \in \Sigma^I$ and $b \in \Sigma^{n-|I|}$ we denote by $(x_I = a, x_{-I} = b)$ the point in Σ^n whose I -coordinates are filled according to a , and whose \bar{I} -coordinates are filled according to b . For two strings $x, y \in \Sigma^n$ we denote by $\Delta(x, y)$ the Hamming distance between x and y , that is, the number of coordinates $i \in [n]$ such that $x_i \neq y_i$.

We denote $A \lesssim B$ to refer to the fact that $A \leq C \cdot B$ for some absolute constant $C > 0$, and $A \gtrsim B$ 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 , the corresponding notation is $A \lesssim_m B$. We will also use standard big- O notations: we denote $A = O(B)$ if $A \lesssim B$, $A = \Omega(B)$ if $A \gtrsim B$; if there is dependency of the hidden constant on some auxiliary parameter, say m , we denote $A = O_m(B)$ and $A = \Omega_m(B)$.

We denote by i the complex root of -1 , and by \bar{a} the complex conjugate of the number $a \in \mathbb{C}$. For a matrix M , we denote by M^* the conjugate transpose matrix of M .

2.1 Product Spaces

Let $(\Sigma^n, \mathcal{D}^{\otimes n})$ be a probability space. We often work with the space $L_2(\Sigma^n; \mathcal{D}^{\otimes n})$ of complex valued functions with finite values. We think of this space as an inner product space, where the inner product of $f, f': \Sigma^n \rightarrow \mathbb{C}$ is defined by

$$\langle f, f' \rangle_{L_2(\Sigma^n; \mathcal{D}^{\otimes n})} = \mathbb{E}_{x \sim \mathcal{D}^{\otimes n}} \left[f(x) \overline{f'(x)} \right].$$

Often times, when the measure $\mathcal{D}^{\otimes n}$ is clear from context, we will omit the $L_2(\Sigma^n; \mathcal{D}^{\otimes n})$ subscript and denote the inner product between f and f' by $\langle f, f' \rangle$.

2.2 The Degree Decomposition and the Efron-Stein Decomposition

Given an inner product space, one may associate with it orthogonal decompositions of $L_2(\Sigma^n; \mathcal{D}^{\otimes n})$. In this section we present two such decompositions, the degree decomposition and its refinement the Efron-Stein decomposition. We will only present the basic notions and facts we need about them, and refer the reader to [32] to a more comprehensive treatment.

2.2.1 Juntas, Degrees and the Degree Decomposition

To define the notion of degrees, it is most convenient to start with the notion of juntas, which are functions that depend only on few of their input coordinates.

Definition 2.1. For $D \subseteq [n]$, a function $f: \Sigma^n \rightarrow \mathbb{C}$ is called a D -junta if there exists $f': \Sigma^D \rightarrow \mathbb{C}$ such that $f(x) = f'(x_D)$ for all $x \in \Sigma^n$.

For an integer $0 \leq d \leq n$, a function $f: \Sigma^n \rightarrow \mathbb{C}$ is called a d -junta if there exists a set $D \subseteq [n]$ of size d such that f is a D -junta.

Equipped with the notion of juntas, we may define the degree decomposition in the following way:

Definition 2.2. For an inner product space $L_2(\Sigma^n; \mathcal{D}^{\otimes n})$ as above and an integer $0 \leq d \leq n$, we define the space $V_{\leq d}(\Sigma^n; \mathcal{D}^{\otimes n}) \subseteq L_2(\Sigma^n; \mathcal{D}^{\otimes n})$ as the space spanned by all d -junta. We also define

$$V_{=d}(\Sigma^n; \mathcal{D}^{\otimes n}) = V_{\leq d}(\Sigma^n; \mathcal{D}^{\otimes n}) \cap V_{\leq d-1}(\Sigma^n; \mathcal{D}^{\otimes n})^\perp.$$

It is clear by definition that the spaces $V_{=d}(\Sigma^n; \mathcal{D}^{\otimes n})$ are mutually orthogonal and

$$L_2(\Sigma^n; \mathcal{D}^{\otimes n}) = \bigoplus_{d=0}^n V_{=d}(\Sigma^n; \mathcal{D}^{\otimes n}),$$

so any function $f: \Sigma^n \rightarrow \mathbb{C}$ can be uniquely written as $f(x) = \sum_{d=0}^n f^{=d}(x)$ where $f^{=d} \in V_{=d}(\Sigma^n; \mathcal{D}^{\otimes n})$ is called the degree d component of f . With these notations, Plancherel's equality states that for any pair of functions $f, g: (\Sigma^n, \mathcal{D}^{\otimes n}) \rightarrow \mathbb{C}$ one has that

$$\langle f, g \rangle = \sum_{d=0}^n \langle f^{=d}, g^{=d} \rangle.$$

Parseval's equality is the specialized statement where $f = g$, in which case one get that $\|f\|_2^2 = \sum_{d=0}^n \|f^{=d}\|_2^2$.

We often refer to the function $f^{\leq d}(x) = \sum_{i=0}^d f^{=i}(x)$ as the part of f of level at most d , and refer to the quantity $\|f^{\leq d}\|_2^2$ as the level d weight of f :

Definition 2.3. The weight of f on level up to d is defined as $W_{\leq d}[f] = \|f^{\leq d}\|_2^2 = \sum_{i=0}^d \|f^{=i}\|_2^2$.

2.2.2 The Efron-Stein Decomposition

The Efron-Stein decomposition is a refinement of the degree decomposition, which we will make use of a handful of times.

Definition 2.4. For an integer $0 \leq d \leq n$ and $S \subseteq [n]$ of size d , we define

$$V_{=S}(\Sigma^n; \mathcal{D}^{\otimes n}) = V_{=d}(\Sigma^n; \mathcal{D}^{\otimes n}) \cap \text{Span}(\{f: \Sigma^n \rightarrow \mathbb{C} \mid f \text{ is an } S\text{-junta}\}).$$

It can be shown that the spaces $V_{=S}(\Sigma^n; \mathcal{D}^{\otimes n})$ are mutually orthogonal, thus

$$V_{=d}(\Sigma^n; \mathcal{D}^{\otimes n}) = \bigoplus_{|S|=d} V_{=S}(\Sigma^n; \mathcal{D}^{\otimes n}).$$

In particular, given any function $f: \Sigma^n \rightarrow \mathbb{C}$ and $0 \leq d \leq n$, we may further decompose the degree d component of f , namely $f^{=d}$, and uniquely write it as $f^{=d}(x) = \sum_{|S|=d} f^{=S}(x)$ where $f^{=S} \in V_{=S}(\Sigma^n; \mathcal{D}^{\otimes n})$. Thus, we get the Efron-Stein decomposition of f : $f(x) = \sum_{S \subseteq [n]} f^{=S}(x)$, where $f^{=S} \in V_{=S}(\Sigma^n; \mathcal{D}^{\otimes n})$. Once again, with these notations Plancherel's equality states that for any pair of functions $f, g: (\Sigma^n, \mathcal{D}^{\otimes n}) \rightarrow \mathbb{C}$, one has that

$$\langle f, g \rangle = \sum_{S \subseteq [n]} \langle f^{=S}, g^{=S} \rangle,$$

and Parseval's equality is the specialized statement where $f = g$, in which case we get $\|f\|_2^2 = \sum_{S \subseteq [n]} \|f^{=S}\|_2^2$.

2.3 Random Restrictions

In this section we define the notions of restrictions and random restrictions of functions, which are used in this paper extensively. We use two types of random restrictions. The first type is very common in the area of analysis of Boolean functions; one selects a set of coordinates randomly, fixes them according to the marginal distribution there and thinks of the rest of the coordinates as variables. The underlying measure of the input space stays the same. The second type is a much less common type of restrictions, and it is crucial for our arguments. In this second type, we still chooses a random set of variables and a fixing for them, but not according to these the marginal distribution on this coordinates. Rather, the fixing for this set of variables is chosen according to a different measure, and to balance this out the underlying measure of the rest of the coordinates changes. Below is a more formal description.

2.3.1 Restrictions that Preserve the Underlying Measure

For a finite alphabet Σ and a probability measure μ over it, a function $f: (\Sigma^n, \mu^{\otimes n}) \rightarrow \mathbb{C}$, a set of coordinates $I \subseteq [n]$ and a partial input $z \in \Sigma^I$, the restricted function $f_{I \rightarrow z}: \Sigma^{[n] \setminus I} \rightarrow \mathbb{C}$ is defined as

$$f_{I \rightarrow z}(y) = f(x_I = z, x_{\bar{I}} = y).$$

A random restriction of a function $f: (\Sigma^n, \mu^{\otimes n}) \rightarrow \mathbb{C}$ refers to a restriction in which either (or both) I and z are chosen randomly. Typically, random restrictions are associated with a parameter $\rho \in (0, 1)$: we first choose $I \subseteq [n]$ by including each element $i \in [n]$ independently with probability ρ , then choose $z \sim \mu^I$ and then consider the function $f_{I \rightarrow z}$ as a function from $(\Sigma^{\bar{I}}, \mu^{\bar{I}})$ to \mathbb{C} . we often denote by $I \subset_{\rho} [n]$ the distribution of I which is sampled in such a way.

2.3.2 Restrictions that Do Not Preserve the Underlying Measure

An important utility of restrictions for us will be that they allow us to change the underlying measure of our probability space. Suppose that the measure μ can be written as $\mu = \rho\mathcal{D}_1 + (1 - \rho)\mathcal{D}_2$, where \mathcal{D}_1 and \mathcal{D}_2 are distributions and $\rho \in (0, 1)$. In such situations we will often consider the following random restriction process: choose $I \subseteq_\rho [n]$, choose $z \sim \mathcal{D}_2^I$, and consider the function $f_{\bar{I} \rightarrow z}$ as a function from $(\Sigma^I, \mathcal{D}_1^I)$ to \mathbb{C} . Note that under these random choices, choosing $y \sim \mathcal{D}_1^I$, the distribution of the point $(x_I = y, x_{\bar{I}} = z)$ is still μ , hence this restriction process makes sense. In particular, the expected average, as well as the expected 2-norm squared of $f_{\bar{I} \rightarrow z}$ over the choice of z are the average and the 2-norm squared of f .

Such random restrictions are used extensively in the paper. An example case where this can be useful is the case that in the distribution μ the probability of each atom is at least α , and we wish to switch from it to the uniform distribution over Σ . In that case, we may write $\mu = \frac{\alpha}{2}U + (1 - \frac{\alpha}{2})\mu'$ where U is the uniform distribution over Σ and μ' is some distribution. Following the above procedure for random restrictions, we may thus change the underlying measure of our space from μ to U by approximately fixing $1 - \frac{\alpha}{2}$ randomly chosen fraction of the coordinates according to μ' .

2.4 Markov Chains

Given a probability space (Σ, μ) , we will often consider Markov chains over Σ that have μ as a stationary distribution. We often denote these Markov Chain by T , and abusing notations we will also think of T as an averaging operator from $L_2(\Sigma, \mu)$ to $L_2(\Sigma, \mu)$ defined as

$$Tf(x) = \mathbb{E}_{y \sim Tx} [f(y)].$$

We say a Markov chain T is connected if the graph, whose vertices are Σ and the edges are (a, b) if there is a transition from a to b in T , is connected. We need a few well known basic properties of Markov chains that we record below.

Fact 2.5. *If T is connected and the probability of each atom is at least α , then $\lambda_2(T) \leq 1 - \Omega(\alpha^2)$.*

Proof. The proof is by an application of Cheeger's inequality on the graph associated with T , and we refer the reader to [30] for a formal proof. \square

Given an averaging operator T acting on univariate functions, we often think of its n -fold tensor $T^{\otimes n}$. Again, we will think of $T^{\otimes n}$ both as a Markov chain over Σ^n (on which, on each coordinate the Markov chain T is applied independently), as well as an averaging operator acting on $L_2(\Sigma^n, \mu^{\otimes n})$. In the case that μ is a stationary distribution of T , it is easily shown that the spaces $V_{=S}$ are invariant under $T^{\otimes n}$, and for each $g \in V_{=S}$ it holds that $\|T^{\otimes n}g\|_2 \leq \lambda_2(T)^{|S|}\|g\|_2$.

The following lemma asserts that if $f, g: \Sigma^n \rightarrow \mathbb{C}$ are 1-bounded and $\langle f, T^{\otimes n}g \rangle$ is significant, then f and g must have significant mass of the low levels.

Lemma 2.6. *Suppose (Σ, μ) is a finite domain and T is a connected Markov chain with stationary distribution μ , in which the probability of each atom is at least α . Then for all $\varepsilon > 0$ there is $d = O_\alpha(\log(1/\varepsilon))$ such that if $f, g: (\Sigma^n, \mu^{\otimes n}) \rightarrow \mathbb{C}$ are 1-bounded and $|\langle f, T^{\otimes n}g \rangle| \geq \varepsilon$, then $W_{\leq d}[f] \geq \frac{\varepsilon^2}{4}$.*

Proof. Decomposing f, g according to the Efron-Stein decomposition of $(\Sigma^n, \mu^{\otimes n})$ as $f = \sum_S f^{=S}$ and $g = \sum_S g^{=S}$, we get that

$$\langle f, \mathbb{T}^{\otimes n} g \rangle = \sum_{S, Q} \langle f^{=S}, \mathbb{T}^{\otimes n} g^{=Q} \rangle = \sum_S \langle f^{=S}, \mathbb{T}^{\otimes n} g^{=S} \rangle = \sum_{S \neq \emptyset} \langle f^{=S}, \mathbb{T}^{\otimes n} g^{=S} \rangle.$$

The contribution from $|S| \leq d$ is at most

$$\begin{aligned} \sum_{|S| \leq d} |\langle f^{=S}, \mathbb{T}^{\otimes n} g^{=S} \rangle| &\leq \sum_{|S| \leq d} \|f^{=S}\|_2 \|\mathbb{T}^{\otimes n} g^{=S}\|_2 \leq \sum_{|S| \leq d} \|f^{=S}\|_2 \|g^{=S}\|_2 \leq \sqrt{W_{\leq d}[f] W_{\leq d}[g]} \\ &\leq \sqrt{W_{\leq d}[f]}. \end{aligned}$$

For $|S| > d$, we have by Fact 2.5 that

$$\|\mathbb{T}^{\otimes n} g^{=S}\|_2 \leq \lambda_2(\mathbb{T})^{|S|} \|g^{=S}\|_2 \leq (1 - \Omega_\alpha(1))^d \|g^{=S}\|_2 \leq \frac{\varepsilon}{2} \|g^{=S}\|_2$$

for d chosen suitably as in the statement. Thus, using Cauchy-Schwarz the contribution from $|S| > d$ is at most

$$\begin{aligned} \sum_{|S| > d} |\langle f^{=S}, \mathbb{T}^{\otimes n} g^{=S} \rangle| &\leq \sum_{|S| > d} \|f^{=S}\|_2 \|\mathbb{T}^{\otimes n} g^{=S}\|_2 \leq \frac{\varepsilon}{2} \sum_{|S| > d} \|f^{=S}\|_2 \|g^{=S}\|_2 \\ &\leq \frac{\varepsilon}{2} \sqrt{\sum_{|S| > d} \|f^{=S}\|_2^2} \sqrt{\sum_{|S| > d} \|g^{=S}\|_2^2}, \end{aligned}$$

which is at most $\frac{\varepsilon}{2}$ by Parseval. Combining, we get that $\sqrt{W_{\leq d}[f]} + \frac{\varepsilon}{2} \geq \varepsilon$, and the statement follows by re-arranging. \square

2.5 Some Markov Chain Lemmas

In this section, we collect a few basic results regarding Markov chains that we will make use of repeatedly.

2.5.1 The Eigenvalues of a Markov Chain

The first statement gives a description of the eigenvalues of a Markov chain (both upper and lower bounds) as a function of the minimum probability of each atom and the minimum probability that the Markov chain stays in the same state. It will help us to analyze noise-type operators that are similar to the standard noise operator on product spaces (but are not quite the same).

Lemma 2.7. *For all $\alpha > 0$ and $m \in \mathbb{N}$ there are $c > 0$ and $C > 0$ such that the following holds. Let Σ be an alphabet of size at most m , let μ be a distribution over Σ in which the probability of each atom is at least α , and let \mathbb{T} be a Markov chain over Σ in which μ is a stationary distribution. Let k be the number of connected components in \mathbb{T} , and let $\lambda_1(\mathbb{T}) \geq \dots \geq \lambda_m(\mathbb{T})$ be the eigenvalues of \mathbb{T} .*

1. We have $\lambda_1(\mathbb{T}) = \dots = \lambda_k(\mathbb{T}) = 1$.
2. If the probability of each transition in \mathbb{T} is at least ξ , then $\lambda_i(\mathbb{T}) \leq 1 - c\xi$ for all $i \geq k + 1$.

3. If for all $x \in \Sigma$ we have that $\Pr_{x' \sim T x} [x' = x] \geq 1 - \xi$, then $\lambda_i(T) \geq 1 - C\xi$.

Proof. For the first item, write $\Sigma = \Sigma_1 \cup \dots \cup \Sigma_k$ where $\Sigma_1, \dots, \Sigma_k$ are the connected components of T , and take g_i which is 1 on Σ_i and 0 on the rest. Note that the g_i 's are linearly independent and all have eigenvalues 1 in T .

For the second item, let g be an eigenvector of T perpendicular to g_1, \dots, g_k with eigenvalue λ , and normalize it so that $\|g\|_2^2 = 1$. Then there is i such that

$$\sum_{x \in \Sigma_i} \mu(x) |g(x)|^2 \geq \frac{1}{k},$$

so either for $h = \operatorname{Re}(g)$ or $h = \operatorname{Im}(g)$ we have that $\sum_{x \in \Sigma_i} \mu(x) h(x)^2 \geq \frac{1}{2k}$. Therefore there must be $x \in \Sigma'$ such that either $h(x) \geq \sqrt{1/2k}$ or $h(x) \leq -\sqrt{1/2k}$; without loss of generality we assume the former. As $\langle g, g_i \rangle = 0$ we get $\sum_{x \in \Sigma_i} \mu(x) g(x) = 0$ hence $\sum_{x \in \Sigma_i} \mu(x) h(x) = 0$, and in particular there is $x' \in \Sigma_i$ such that $h(x) \leq 0$. Summarizing, $|h(x) - h(x')| \geq 1/\sqrt{2k}$. As Σ_i is a connected component of T , there is a path $x = x_0 \rightarrow \dots \rightarrow x_\ell = x'$ between x and x' , and so

$$\frac{1}{\sqrt{2k}} \leq |h(x) - h(x')| \leq \sum_{i=0}^{\ell-1} |h(x_i) - h(x_{i+1})|,$$

so there is i such that $|h(x_i) - h(x_{i+1})| \geq \frac{1}{\sqrt{2k\ell}} \geq \frac{1}{m^{3/2}\sqrt{2}}$. Therefore $|g(x_i) - g(x_{i+1})| \geq \frac{1}{m^{3/2}\sqrt{2}}$, an it follows that

$$1 - \lambda = 1 - \langle g, Tg \rangle = \frac{1}{2} \mathbb{E}_{\substack{z \sim \mu \\ z' \sim Tz}} \left[|g(z) - g(z')|^2 \right] \geq \frac{\xi}{2} |g(x_i) - g(x_{i+1})|^2 \geq \frac{\xi}{4m^3},$$

finishing the proof of the second item.

For the third item, taking g as before we have that

$$1 - \lambda = 1 - \langle g, Tg \rangle = \frac{1}{2} \mathbb{E}_{\substack{z \sim \mu \\ z' \sim Tz}} \left[|g(z) - g(z')|^2 \right] \leq \frac{1}{2} \Pr_{z \sim \mu, z' \sim Tz} [z \neq z'] 4\|g\|_\infty^2,$$

which is at most $\xi \frac{2}{\alpha} \|g\|_2^2 = \frac{4}{\alpha} \xi$, and the third item is proved. \square

2.5.2 Markov Chains and Random Restrictions

We will often measure various notion of degrees via Markov Chains based notions (as opposed to precise degrees), so as to preserve boundedness of functions. As such, we often want to make assertions of the form: ‘‘if a function f has high degree, then a random restriction of f also has high degree’’ (where again, the notion of degree is not necessarily the standard notion). Such statements are quite straightforward when dealing with the standard notion of degree, but less so with Markov Chain based notions. For our purposes, the following lemma will play the role of such statement for our softer notion of high-degreeeness.

In the statement below one should think of the quantity $\langle f, T_{1-c\beta\xi, \mu^{\otimes n}} f \rangle$ as small, and of the fact that it is small as saying that if we write f in basis of eigenvectors of the operator $T_{1-c\beta\xi, \mu^{\otimes n}}$, then most of the L_2 -mass of f will lie on monomials involving many eigenvectors whose eigenvalue is not 1; the number of

such eigenvectors will need to $\Omega\left(\frac{1}{\beta\xi}\right)$. Morally, the lemma says that after a random restriction leaving β fraction of the coordinates alive, the degree with respect to such eigenvectors is at least $\Omega\left(\frac{1}{\xi}\right)$, but some care is needed as the operator with respect to which we measure degree, changes.

Lemma 2.8. *For all $\alpha, \beta > 0$ there is $c > 0$ such that the following holds. Let Σ be an alphabet of size at most m , let μ, ν_1, ν_2 be distributions over Σ in which the probability of each atom is at least α and $\mu = \beta\nu_1 + (1 - \beta)\nu_2$. Let $G = (\Sigma, E)$ be a graph on Σ , and consider the Markov chains $T_{1-\xi, \mu}$ and $T_{1-\xi, \nu_1}$ defined as: for $T_{1-\xi, \mu}$, on $x \in \Sigma$, a sample $x' \sim T_{1-\xi, \mu}x$ is generated by taking $x' = x$ with probability $1 - \xi$, and otherwise take $x' \sim \mu$ conditioned on x' being a neighbour of x in G . We similarly define $T_{1-\xi, \nu_1}$. Then for all $f: \Sigma^n \rightarrow \mathbb{C}$,*

$$\mathbb{E}_{\substack{I \subseteq \beta[n] \\ z \sim \nu_2^I}} \left[\langle f_{\bar{I} \rightarrow z}, T_{1-\xi, \nu_1^I} f_{\bar{I} \rightarrow z} \rangle_{\nu_1^I} \right] \leq \langle f, T_{1-c\beta\xi, \mu^{\otimes n}} f \rangle.$$

Proof. Expanding, the left hand side is equal to $\mathbb{E}_{(x, x') \sim \mathcal{D}} \left[f(x) \overline{f(x')} \right]$, where the distribution \mathcal{D} is defined as: for each i independently, with probability $1 - \beta$ we take $x_i = x'_i$ sampled according to ν_2 , and otherwise we sample $x_i \sim \nu$ and then $x'_i \sim T_{1-\xi, \nu} x_i$. Noting that marginally, x and x' are distributed according to μ , we may view \mathcal{D} as a reversible Markov chain $T^{\otimes n}$ whose stationary distribution is μ , so that the left hand side is equal to $\langle f, T^{\otimes n} f \rangle$. We show that it is at most the right hand side, and to do so we examine the eigenvalues and eigenvectors of these operators. We begin by remarking that the two operators are symmetric and positive semi-definite (as $\langle f, T_{1-\xi, \mu^{\otimes n}} f \rangle = \langle T_{\sqrt{1-\xi}, \mu^{\otimes n}} f, T_{\sqrt{1-\xi}, \mu^{\otimes n}} f \rangle \geq 0$), hence these eigenvalues are non-negative.

Let k be the number of connected components in G , and note that if g is a function which is constant on the connected components of G , then both operators act on it as the identity. We choose a basis for $L_2(\Sigma)$, say $g_1, \dots, g_k, g_{k+1}, \dots, g_m$, where g_1, \dots, g_k are constant on all of the connected components of G and g_{k+1}, \dots, g_m are perpendicular to g_1, \dots, g_k with respect to $\langle \cdot, \cdot \rangle_\mu$. Then we have that $Tg_i = T_{1-\xi, \mu}g_i$ for $i = 1, \dots, k$. Also, the space $\text{Span}(g_{k+1}, \dots, g_m)$ is invariant under both T and $T_{1-\xi, \mu}$, and by Lemma 2.7 all eigenvalues of T are at most $1 - s\beta\xi$ for some $s(m, \alpha) > 0$ (as the probability of each transition is at least $\xi\beta$), and all eigenvalues of $T_{1-c\xi\beta, \mu}$ are at least $1 - cC\xi\beta > 1 - s\beta\xi$. Thus, for every $g \in \text{Span}(g_{k+1}, \dots, g_m)$ we have that $\langle g, Tg \rangle \leq \langle g, T_{1-c\xi\beta, \mu}g \rangle$, and it follows that this inequality holds for every $g: \Sigma \rightarrow \mathbb{C}$.

For multi-variate functions, writing $V = \text{Span}(g_1, \dots, g_k)$ and $V' = \text{Span}(g_{k+1}, \dots, g_m)$, we may decompose

$$L_2(\Sigma; \mu^{\otimes n}) = \oplus_{S \subseteq [n]} \text{Span}(V^{\otimes S} \otimes V'^{\otimes [n] \setminus S})$$

and thus write any $g: \Sigma^n \rightarrow \mathbb{C}$ as $g = \sum_S c_S g_S$ where $g_S \in \text{Span}(V^{\otimes S} \otimes V'^{\otimes [n] \setminus S})$ has 2-norm equal to 1.

A Computation for T . For T we now get that:

$$\langle g, T^{\otimes n} g \rangle = \sum_S |c_S|^2 \langle g_S, T^{\otimes n} g_S \rangle \leq \sum_S |c_S|^2 \|g_S\|_2 \|T^{\otimes n} g_S\| \leq \sum_S (1 - s\beta\xi)^{n-|S|} |c_S|^2 \|g_S\|_2^2.$$

In the last inequality, we used the fact that $\|T^{\otimes n} g_S\| \leq (1 - s\beta\xi)^{n-|S|} c_S^2 \|g_S\|_2$, which may be observed as follows. Further decomposing V and V' into eigenspaces of T , we may write $g_S = \sum_{a \in A_S} c'_a v_a$ where

$v_a \in V^{\otimes S} \otimes V'^{\otimes [n] \setminus S}$ are orthogonal unit vectors. Moreover, v_a is an eigenvector of T with eigenvalue $0 \leq \lambda_a \leq (1 - s\beta\xi)^{2(n-|S|)}$, and so

$$\|T^{\otimes n} g_S\|_2^2 = \sum_{a \in A_S} |c'_a|^2 \lambda_a^2 \leq \sum_{a \in A_S} |c'_a|^2 (1 - s\beta\xi)^{2(n-|S|)} = (1 - s\beta\xi)^{2(n-|S|)} \|g_S\|_2^2.$$

A Computation for T' . Similarly, for $T_{1-\xi, \mu^{\otimes n}}$, we further decompose V and V' into eigenspaces of $T_{1-\xi, \mu^{\otimes n}}$ and write $g_S = \sum_{a \in A_S} c'_a v_a$ where the $v_a \in V^{\otimes S} \otimes V'^{\otimes [n] \setminus S}$'s are orthogonal, have 2-norm equal to 1 and are each tensor of eigenvectors of $T_{1-\xi, \mu}$. Thus, $T_{1-\xi, \mu^{\otimes n}} v_a = \lambda_a v_a$ for $\lambda_a \geq (1 - cC\xi\beta)^{n-|S|}$ and so

$$\begin{aligned} \langle g, T_{1-\xi, \mu^{\otimes n}} g \rangle &= \sum_S |c_S|^2 \langle g_S, T_{1-\xi, \mu^{\otimes n}} g_S \rangle = \sum_{S, a \in A_S} |c_S|^2 c_a^2 \lambda_a \|v_a\|_2^2 \\ &\geq \sum_S |c_S|^2 (1 - cC\xi\beta)^{n-|S|} \sum_{a \in A_S} |c'_a|^2, \end{aligned}$$

which is at least $\sum_S (1 - s\beta\xi)^{n-|S|} c_S^2 \|g_S\|_2^2 \geq \langle g, T^{\otimes n} g \rangle$, as required. \square

2.5.3 Comparing Two Markov Chains

The next lemma is tailored to handle the following case. Suppose that T and T' are two Markov chains on Σ that both have μ as stationary distribution, and whose corresponding averaging operators $T, T' : L_2(\Sigma, \mu) \rightarrow L_2(\Sigma, \mu)$ are both positive semi-definite. Intuitively, if the Markov chain T is richer than T' , then the averaging operator T does more averaging than T' . Hence, if we know that for some function $f : (\Sigma^n, \mu^{\otimes n}) \rightarrow \mathbb{C}$ of 2-norm equal to 1 it holds that $\|T'^{\otimes n} f\|_2$ is small, then $\|T^{\otimes n} f\|_2$ should also be small. More formally:

Lemma 2.9. *For all $m \in \mathbb{N}$ and $\alpha, \beta > 0$ there is $s > 0$ such that the following holds. Suppose that μ is a distribution over Σ in which the probability of each atom is at least α and $|\Sigma| \leq m$, and let T, T' be reversible Markov chains on Σ with μ as stationary distribution. Further suppose that there is $\delta > 0$ such that the following holds:*

1. *For all $x \in \Sigma$, $\Pr_{x' \sim T_x} [x' = x] \geq 1 - \delta$ and $\Pr_{x' \sim T'_x} [x' = x] \geq 1 - \delta$.*
2. *For any distinct $x, y \in \Sigma$ such that $T'(x, y) > 0$ it holds that $T(x, y) > 0$.*
3. *The probability of each transition in T, T' is at least $\beta\delta$.*

Then for all $f : (\Sigma^n, \mu^{\otimes n}) \rightarrow \mathbb{C}$ of 2-norm at most 1 it holds that

$$\|T^{\otimes n} f\|_2 \leq \|T'^{\otimes n} f\|_2^s.$$

Proof. Suppose for simplicity of notation that $|\Sigma| = m$. Let $C_1 \cup \dots \cup C_k$ be a partition of Σ into the connected components of T , and let $V \subseteq L_2(\Sigma)$ be the sub-space of functions that are constant on each C_i ; we take $V' = V^\perp$, so that $L_2(\Sigma^n, \mu^{\otimes n}) = \bigoplus_{S \subseteq [n]} \text{Span}(V_S)$ where $V_S = V^{\otimes S} \otimes V'^{\otimes [n] \setminus S}$.

We note that as T is self adjoint (by reversibility) and V is an invariant space of T , it follows that V' is also an invariant space of T . We thus write $f = \sum_{S \subseteq [n]} c_S f_S$ where $f_S \in \text{Span}(V_S)$ has 2-norm equal to

1, and get that Tf_S is also in $\text{Span}(V_S)$. Similarly, $T'f_S$ is in $\text{Span}(V_S)$, as T' preserves V and as each connected component of T' is contained in a connected component of T . Thus, we get that

$$\|T^{\otimes n} f\|_2^2 = \sum_S |c_S|^2 \|T^{\otimes n} f_S\|_2^2, \quad \|T'^{\otimes n} f\|_2^2 = \sum_S |c_S|^2 \|T'^{\otimes n} f_S\|_2^2.$$

We now argue that $\|T^{\otimes n} f_S\|_2^2 \leq (1 - c(\alpha, \beta, m)\delta)^{2(n-|S|)}$ and $\|T'^{\otimes n} f_S\|_2^2 \geq (1 - C(\alpha, \beta, m)\delta)^{2(n-|S|)}$ where $c, C > 0$. Indeed, for the first inequality we further decompose V and V' into eigenspaces of T (as in Lemma 2.8) and proceed with the same computation as there; we use Lemma 2.7 to upper bound the eigenvalues of T . For the second inequality, we decompose V and V' into eigenspaces of T' , and proceed with the same computation as there; we use Lemma 2.7 to lower bound the eigenvalues of T' . Thus, we conclude that there is a constant $A(\alpha, \beta, m)$ such that

$$\|T^{\otimes n} f\|_2^2 \leq \sum_S |c_S|^2 (1 - c(\alpha, \beta, m)\delta)^{2(n-|S|)}, \quad \|T'^{\otimes n} f\|_2^2 \geq \sum_S |c_S|^2 (1 - c(\alpha, \beta, m)\delta)^{2A(n-|S|)}.$$

By Hölder's inequality, we get that $\|T^{\otimes n} f\|_2^{2A}$ can be upper bounded as

$$\begin{aligned} \left(\sum_S |c_S|^2 (1 - c(\alpha, \beta, m)\delta)^{2(n-|S|)} \right)^A &= \left(\sum_S |c_S|^{2(A-1)/A} \cdot |c_S|^{2/A} (1 - c(\alpha, \beta, m)\delta)^{2(n-|S|)} \right)^A \\ &\leq \left(\sum_S |c_S|^2 \right)^{A-1} \sum_S |c_S|^2 (1 - c(\alpha, \beta, m)\delta)^{2A(n-|S|)} \\ &\leq \|f\|_2^{2(A-1)} \|T'^{\otimes n} f\|_2^2, \end{aligned}$$

where we used Parseval. As $\|f\|_2 \leq 1$, taking $2A$ -th root finishes the proof. \square

3 On the Master Embedding, Path Trick, and Their Interaction

In this section, we present master embeddings as well as the path trick from [4, 5], and prove some properties of them that will be crucial to us. The path trick is an idea which was used in previous works in this series, and for multiple reasons. In [4] it was used to enrich the distribution μ so as to gain pairwise independence, and in [5] it was used to gain a more limited form of pairwise independence as well as for overcoming a certain technical challenge referred to as the ‘‘Horn-SAT obstruction’’ therein.² One important property of the path trick used in both of these works, is that if a distribution μ does not admit Abelian embeddings, then applying the path trick on it results in a distribution that also does not admit Abelian embeddings

In the current work, however, we have to deal with distribution admitting Abelian embeddings, and thus we have to dig deeper. In particular, we have to study the interaction between the path trick and Abelian embeddings of μ , which is the primary topic of this section. To do that, we first define the master embedding, which is an embedding of μ that ‘‘encapsules’’ within it all Abelian embeddings on μ . We then show two properties of the path trick and master embeddings:

1. **The Master is Preserved Under Path Tricks.** We show that the path trick never ‘‘introduces’’ new Abelian embeddings. By that, we mean that from any Abelian embedding of the original distribution

²This obstruction will also appear in the present work and here too we will make an essential use of the path trick to resolve it; this will be the topic of discussion in Section 6.

μ one can construct an Abelian embedding of distribution after the path trick, and furthermore (and this is the important part) there are no other Abelian embeddings. This means that the a master embedding for μ remains a master embedding for the distribution after the path trick (after an appropriate transformation).

2. **The Master is Saturated after Path Tricks.** For reasons that were discussed in the introduction and will be further discussed below, it is desirable for us that the image of our master embeddings will be a whole group. A-priori, there is no reason this will be the case, and indeed it is most often not. We show that, by applying the path trick in a certain way, we are able to enrich the predicate so that the image of the master embeddings become complete Abelian groups.

3.1 Defining the Master Embedding

Let Σ , Γ , and Φ be finite alphabets and let μ be a distribution over $\Sigma \times \Gamma \times \Phi$. In this section, we wish to show that there is a single Abelian group and a single embedding of μ that within it “encapsulates” all Abelian embeddings of μ , which we often refer to as a master embedding of μ .

To start getting some intuition, note that every Abelian embedding of μ given by $\sigma: \Sigma \rightarrow (H, +)$, $\gamma: \Gamma \rightarrow (H, +)$ and $\phi: \Phi \rightarrow (H, +)$, where $(H, +)$ is a finite Abelian group can be thought about as partitions of each one of the alphabets Σ , Γ and Φ , and labeling each part of each partition by a group element. Thus, as the number of these partitions is a finite number depending only on the alphabet sizes, it makes sense that we would not need to look into too large of groups to find proper labelings of that partition by group elements so as to get an Abelian embedding (if such one exists). This requires some care, and towards this end we define the notion of equivalent embeddings.

Definition 3.1. *Let μ be a distribution over $\Sigma \times \Gamma \times \Phi$. We say an embedding $\sigma: \Sigma \rightarrow (G, +)$, $\gamma: \Gamma \rightarrow (G, +)$ and $\phi: \Phi \rightarrow (G, +)$ is a linear reduction of $\sigma': \Sigma \rightarrow (H, +)$, $\gamma': \Gamma \rightarrow (H, +)$ and $\phi': \Phi \rightarrow (H, +)$ if G is a subgroup of H , and there are injective linear maps $m_1, m_2, m_3: G \rightarrow H$, such that $\sigma'(x) = m_1(\sigma(x))$, $\gamma'(y) = m_2(\gamma(y))$ and $\phi'(z) = m_3(\phi(z))$.*

As a concrete example for the notion of linear refinements, we note that if σ , γ and ϕ form an embedding of μ into $(\mathbb{Z}_p, +)$, then $p\sigma$, $p\gamma$ and $p\phi$ form an embedding of μ into $(\mathbb{Z}_{p^2}, +)$. In essence though, these two embeddings are “the same”, thus to encapsulate all of the Abelian embeddings of μ it suffices to only take into account one of them. Indeed, one can observe that (σ, γ, ϕ) is a linear reduction of $(p\sigma, p\gamma, p\phi)$.

Below, we consider a distribution μ that does not admit non-trivial Abelian embeddings into $(\mathbb{Z}, +)$ and gradually construct a master embedding of it. Towards this end, we consider the basic building block of all Abelian groups, namely cyclic groups, and show that modulo equivalences, μ can admit only $O_m(1)$ many Abelian embeddings, where m is an upper bound on the alphabet sizes.

3.1.1 Embeddings into Cyclic Groups

We start by showing that if μ does not admit any $(\mathbb{Z}, +)$ embedding, then μ cannot have Abelian embedding into arbitrarily large Abelian groups of prime order.

Lemma 3.2. *Let $m \in \mathbb{N}$, and suppose that Σ , Γ and Φ are finite alphabets of size at most m . Then there exists $r \in \mathbb{N}$, such that for any distribution μ over $\Sigma \times \Gamma \times \Phi$, if μ does not admit an Abelian embedding into $(\mathbb{Z}, +)$, then μ does not admit any Abelian embedding to $(\mathbb{F}_p, +)$ for $p > r$.*

Proof. Write $\Sigma = \{x_1, \dots, x_{m_1}\}$, $\Gamma = \{y_1, \dots, y_{m_2}\}$ and $\Phi = \{z_1, \dots, z_{m_3}\}$, and associate with each symbol Σ a variable V_{x_i} and similarly for symbols in Γ and Φ . We will think of these variables as representing the values of an embedding of μ , and so the conditions that they form an embedding can be written as the system of linear equations $V_x + V_y + V_z = 0$ for all $(x, y, z) \in \text{supp}(\mu)$. Thus, the fact that there are no embedding of μ into $(\mathbb{Z}, +)$ is equivalent to the fact that over integers, the only solution to this system of equations are trivial constant (say, assigning all x -variables the value 17, all y -variables the value -6 , and all z -variables the value -11). We pick some $(x^*, y^*, z^*) \in \text{supp}(\mu)$ and assign V_{x^*} , V_{y^*} and V_{z^*} the value 0, each solution of the original system corresponds (over any Abelian group) corresponds to a shift of a solution of the new system. Thus, the new system only has the trivial all 0 solution over $(\mathbb{Z}, +)$.

We write this system in matrix form as $MV = 0$, and bring it to diagonal form without performing divisions. Namely, each time we pick a unpivoted equation from our system, then a variable from it, say V_i . We multiply all equations by appropriate constants so that the coefficients of V_i in each equation is the same (or 0 if V_i doesn't appear in that equation), then subtract the chosen equation from all equations in which V_i appears, and then declare the equation as pivoted. In the end of the process, we will end up with a system of equations of the form $c_i V_i = 0$ for $c_i \in \mathbb{Z}$, where all c_i are bounded by a universal constant $C(m)$ (as the coefficients grow by at most a constant factor depending only on m in each step, and the number of steps is at most m^3). We write this system as $M'V = 0$, and now the fact that there are no solutions over \mathbb{Z} means that the rank of M' is full, namely $3m$.

Thus, for primes $p > C(m)$, we note that an \mathbb{F}_p solution of the system is also a solution over integers. Indeed, for the equation $c_i V_i = 0 \pmod{p}$ to hold for $V_i \in \mathbb{F}_p$, it must be the case that either $V_i = 0$, or else c_i must be divisible by p . It follows that there are no non-trivial embeddings over $(\mathbb{F}_p, +)$. \square

Next, we discuss Abelian embeddings into cyclic groups of prime power order. Here, the situation is slightly trickier, as if μ has an embedding into $(\mathbb{F}_p, +)$, then by multiplying it by p^{k-1} one automatically gets an embedding into $(\mathbb{Z}_{p^k}, +)$. In the following lemma, we show that there is $k_0 = k_0(m) \in \mathbb{N}$ such that to "exhaust" all embeddings into groups of the form $(\mathbb{Z}_{p^k}, +)$, it suffices to consider $k \leq k_0$, in the sense that an embedding into $(\mathbb{Z}_{p^k}, +)$ for $k > k_0$ is equivalent to an embedding into $(\mathbb{Z}_{p^{k_0}}, +)$.

Lemma 3.3. *Let $m \in \mathbb{N}$, and suppose that Σ , Γ and Φ are finite alphabets of size at most m . Then there exists $p_0 \in \mathbb{N}$ and $k_0 \in \mathbb{N}$, such that for any distribution μ over $\Sigma \times \Gamma \times \Phi$ that doesn't admit any $(\mathbb{Z}, +)$ embedding, if μ has a non-trivial Abelian embedding (σ, γ, ϕ) into $(\mathbb{Z}_{p^k}, +)$, then $p \leq p_0$ and there is $k' \leq k$ and an Abelian embedding $(\sigma', \gamma', \phi')$ into $(\mathbb{Z}_{p^{k'}}, +)$, such that $(\sigma', \gamma', \phi')$ is a linear reduction of (σ, γ, ϕ) .*

Proof. We write a system of linear equations and bring it to a diagonal form as in the proof of Lemma 3.2, and let $C(m) \in \mathbb{N}$ be an upper bound on the size of all of the coefficients there. We prove the statement for $p_0 = C(m)$ and $k_0 = \lceil \log C(m) \rceil$. If $k \leq k_0$ we are done, so assume otherwise.

By applying appropriate affine shift, we can assume that $\sigma(x^*) = \gamma(y^*) = \phi(z^*) = 0$. Then the fact that $p \leq p_0$ for some $p_0 = O_m(1)$ follows from the argument in Lemma 3.2, and we next argue about k . Consider an equation $c_i V_i = 0 \pmod{p^k}$ therein, and assume that $k > k'$. Then p^k cannot divide c_i , so we may write $c_i = p^{a_i} c'_i$ where $a_i < k_0$ and c'_i is relatively prime to p . This means that $V_i = 0 \pmod{p^{k-a_i}}$, and so p divides V_i . This means that all values of σ, γ, ϕ are divisible by p^{k-k_0} , hence we may look at $\sigma' = \sigma/p^{k-k_0}$, $\gamma' = \gamma/p^{k-k_0}$ and $\phi' = \phi/p^{k-k_0}$ and get that $(\sigma', \gamma', \phi')$ form an embedding of μ into $(\mathbb{Z}_{p^{k_0}}, +)$, as required. \square

3.1.2 Finding Small Equivalent Embeddings on Finite Abelian Groups

With Lemmas 3.2, 3.3 in hand, we can now address general Abelian embeddings, and show that any Abelian embedding is equivalent to an Abelian embedding into a group of bounded size.

Lemma 3.4. *For all $m \in \mathbb{N}$ there is r such that for alphabets Σ, Γ, Φ of size at most m , if μ is a distribution over $\Sigma \times \Gamma \times \Phi$ that does not admit non-trivial Abelian embeddings over $(\mathbb{Z}, +)$, and (σ, γ, ϕ) is an Abelian embedding of μ , then there is an Abelian embedding $(\sigma', \gamma', \phi')$ into a group of size at most r which is a linear reduction of (σ, γ, ϕ) .*

Proof. Let σ, γ and ϕ be an Abelian embedding of μ into an Abelian group $(H, +)$. By the fundamental theorem of finite Abelian groups there are primes p_1, \dots, p_k and integers $r_1, \dots, r_k \geq 1$ such that

$$H = \prod_{i=1}^k \mathbb{Z}_{p_i^{r_i}},$$

and with this identification we can write $\sigma: \Sigma \rightarrow H$ as $\sigma = (\sigma_1, \dots, \sigma_k)$ where $\sigma_i: \Sigma \rightarrow \mathbb{Z}_{p_i^{r_i}}$, and similarly write $\gamma = (\gamma_1, \dots, \gamma_k)$ and $\phi = (\phi_1, \dots, \phi_k)$. We assume each $(\sigma_i, \gamma_i, \phi_i)$ is non-trivial, otherwise we may drop it altogether.

Note that each $(\sigma_i, \gamma_i, \phi_i)$ forms a cyclic Abelian embedding of μ , so by Lemma 3.2 it follows that $p_i \leq O_m(1)$ for all i . By Lemma 3.3 we get that $(\sigma_i, \gamma_i, \phi_i)$ is equivalent to an Abelian embedding $(\sigma'_i, \gamma'_i, \phi'_i)$ into $\mathbb{Z}_{p_i^{r'_i}}$ for $r'_i \leq O_m(1)$. Thus (σ, γ, ϕ) is equivalent to $(\sigma', \gamma', \phi')$ where $\sigma' = (\sigma'_1, \dots, \sigma'_k)$, $\gamma' = (\gamma'_1, \dots, \gamma'_k)$ and $\phi' = (\phi'_1, \dots, \phi'_k)$. Now, $(\sigma', \gamma', \phi')$ is an embedding of μ into $H' = \prod_{i \leq k} \mathbb{Z}_{p_i^{r'_i}}$. In other words, we

managed to reduce each $p_i^{r_i}$ to $p_i^{r'_i}$ that is bounded and get an equivalent embedding, and to simplify notation we drop the primes and assume the embeddings σ, γ and ϕ as well as H are of this form to begin with.

Next, we clean up redundancies. We say a coordinate i is Σ -redundant if the partition of Σ induced by $\sigma = (\sigma_1, \dots, \sigma_k)$ is the same as the partition induced by $\sigma_{-i} = (\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_k)$; similarly we define the notions of Γ -redundant and Φ -redundant coordinates. Note that if i is Σ -redundant, Γ -redundant and Φ -redundant, then (σ, γ, ϕ) is equivalent to $(\sigma_{-i}, \gamma_{-i}, \phi_{-i})$. Also note that there are at most $|\Sigma|$ non Σ -redundant coordinates, as well as at most $|\Gamma|$ non Γ -redundant coordinates and at most $|\Phi|$ non Φ -redundant coordinates. Thus, we can eliminate all but at most $k' = |\Sigma| + |\Gamma| + |\Phi|$ of the coordinates using this process, and thus get an embedding into a group of size at most $O_m(1)^{k'}$, which is a size depending only on the alphabet size, as required. \square

3.1.3 The Master Embedding

In this section we formally define the notion of master embeddings, which is crucial for our arguments. We begin by formally defining the notion of master embeddings.

Definition 3.5. *Let Σ, Γ and Φ be finite alphabets, let μ be a distribution over $\Sigma \times \Gamma \times \Phi$, let $(H, +) = \prod_{i=1}^s (H_i, +)$ be an Abelian group and let $\sigma: \Sigma \rightarrow H, \gamma: \Gamma \rightarrow H$ and $\phi: \Phi \rightarrow H$ be an Abelian embedding of μ . We say that (σ, γ, ϕ) is a master embedding if any Abelian embedding $(\sigma', \gamma', \phi')$ of μ there is an $i \in \{1, \dots, s\}$, such that $(\sigma_i, \gamma_i, \phi_i)$ is a linear refinement of $(\sigma', \gamma', \phi')$.*

Using Lemma 3.4, we can construct a master embedding for μ as follows. Choose r as in the lemma therein and consider all embeddings $(\sigma_i, \gamma_i, \phi_i)$ of μ into Abelian groups of size at most r for $i = 1, \dots, s$; and note that r and s are some finite numbers depending only on the alphabet sizes. We can thus define the master embedding as follows:

Definition 3.6. For all finite alphabets Σ , Γ and Φ take $r \in \mathbb{N}$ from Lemma 3.4 and k to be the number of Abelian groups of size at most r . For a distribution μ over $\Sigma \times \Gamma \times \Phi$, there are at most k Abelian embeddings of μ into Abelian groups of size at most r , and letting $(\sigma_1, \gamma_1, \phi_1), \dots, (\sigma_s, \gamma_s, \phi_s)$ into $(H_1, +), \dots, (H_s, +)$ be an enumeration of all of these embeddings, we define the embedding

$$\sigma_{\text{master}}(x) = (\sigma_1(x), \dots, \sigma_s(x)), \quad \gamma_{\text{master}}(y) = (\gamma_1(y), \dots, \gamma_s(y)), \quad \phi_{\text{master}}(z) = (\phi_1(z), \dots, \phi_s(z)),$$

$$\text{into } (H, +) = \left(\prod_{i=1}^s H_i, + \right).$$

In the following lemma, we prove that $(\sigma_{\text{master}}, \gamma_{\text{master}}, \phi_{\text{master}})$ is a master embedding of μ . We remark though that it is not necessarily unique.

Lemma 3.7. We have that $(\sigma_{\text{master}}, \gamma_{\text{master}}, \phi_{\text{master}})$ is a master embedding of μ .

Proof. By Lemma 3.4, for any Abelian embedding (σ, γ, ϕ) of μ there is an Abelian embedding $(\sigma', \gamma', \phi')$ into an Abelian group of size at most r such that $(\sigma', \gamma', \phi')$ is a linear refinement of (σ, γ, ϕ) , and the result follows from the definition of $(\sigma_{\text{master}}, \gamma_{\text{master}}, \phi_{\text{master}})$ (as it includes all such embeddings). \square

For technical reasons, it will be convenient for us to assume that for each $(\sigma_i, \gamma_i, \phi_i)$ in the definition of the master embedding, we have that 0 is in the image of each one of them; this can easily be arranged by a proper affine shift. We assume henceforth that all Abelian embeddings we are dealing with have 0 in their image.

3.2 The Path Trick

In this section, we formally define the path trick from [4, 5], and recall some basic properties of it from these works. We then begin discussing the interactions between the path trick and Abelian embeddings, and prove that, in a sense, the path trick preserves the structure of Abelian embeddings.

3.2.1 The Definition of the Path Trick Distribution

Suppose Σ , Γ and Φ are finite alphabets, and μ is a distribution over $\Sigma \times \Gamma \times \Phi$. Below, we think of $(x, y, z) \sim \mu$. The path trick of length $\ell = 2^t - 1$ with respect to x is a distribution μ_ℓ over $\Sigma' \times \Gamma \times \Phi$ where $\Sigma' \subseteq \Sigma^\ell$, and there are several equivalent ways of defining it. Below we present the two ways we use: one of them will be more intuitive to think about (namely, the path definition), whereas the other one will be more convenient to work with when we apply Cauchy-Schwarz.

The Path Definition. For the path definition, consider the bipartite graph $G = (\Gamma \cup \Phi, E)$ wherein the edges of the graph are between (y, z) for which there is an $x \in \Sigma$ such that $(x, y, z) \in \text{supp}(\mu)$. The edges of the graph are labeled by the x that produced them, as well as weighted according to the distribution μ ; we allow for parallel edges.

With this in mind, a sample from the path trick distribution μ_ℓ is generated as follows: sample a starting point $y \sim \mu_y$ and then proceed by taking a random walk of length ℓ , collecting all the labels x on edges encountered during this walk. Thus, in the end one has the labels $x, \vec{x} = (x_1, \dots, x_\ell)$ of edges encountered, the starting point of the walk $y \in \Gamma$, and the endpoint of the walk $z \in \Phi$. The output of the process is (\vec{x}, y, z) .

The Inductive Definition. For the inductive definition it is more convenient to define it for ℓ 's that are power of 2 rather than powers of 2 minus 1. We define ν_2 as the following distribution over $\Sigma' \times \Gamma' \times \Phi$ where $\Sigma' \subseteq \Sigma^2$, $\Gamma' \subseteq \Gamma^2$.

1. Sample $z \sim \mu_z$.
2. Sample (x_1, y_1, z_1) and (x_2, y_2, z_2) according to μ conditioned on $z_1 = z$ and $z_2 = z$.
3. Output $(x_1, x_2), (y_1, y_2)$ and z .

Note that the path $y_1 \rightarrow z \rightarrow y_2$ with the labels x_1 and x_2 corresponds to a path as sampled in the previous definition (of length 2). Once ν_ℓ has been defined, we define $\nu_{2\ell}$ as:

1. Sample $y \sim \mu_y$.
2. Sample $(x_1, \dots, x_\ell), (y_1, \dots, y_\ell)$ and $(z_1, \dots, z_{\ell-1})$ according to ν_ℓ conditioned on $y_1 = y$.
3. Independently sample $(x'_1, \dots, x'_\ell), (y'_1, \dots, y'_\ell)$ and $(z'_1, \dots, z'_{\ell-1})$ according to ν_ℓ conditioned on $y'_1 = y$.
4. Output $(x_\ell, x_{\ell-1}, \dots, x_1, x'_1, x'_2, \dots, x'_\ell), (y_\ell, \dots, y_1, y'_2, y'_3, \dots, y'_\ell)$ and $(z_{\ell-1}, \dots, z_1, z'_1, \dots, z'_{\ell-1})$.

In words, we sample $y \sim \mu_y$, take two independent walks of length ℓ starting from it, and then concatenate them. The distribution $\tilde{\mu}_\ell$ is naturally derived from the distribution ν_ℓ by just picking the end-points of the y and z , which in the notations above are y_ℓ and $z'_{\ell-1}$.

The following lemma from [5] shows that the two ways of describing the path trick distribution are equivalent. For the sake of completeness we give a sketch of the proof.

Lemma 3.8. *For all $\ell = 2^t - 1$, it holds that $\tilde{\mu}_\ell = \mu_\ell$.*

Proof. We prove by induction that the distribution $\tilde{\mu}_\ell$ is over (\vec{x}, y, z) wherein $y \sim \mu_y$, and conditioned on that \vec{x} and z are a result of a random walk as in the path definition. Indeed, for $\ell = 1$ and $\ell = 2$ this is clear by inspection.

Assume the statement for ℓ and prove for 2ℓ ; note that the process above may be thought of as first sampling a midpoint y , generate from it two paths of length ℓ independently, concatenate them and then chop off the final step in the path (so as that one ends up in the Φ -part as opposed to the Γ -part). Note that for every length ℓ , the distribution over paths from Γ to Γ of length ℓ is invariant under path reversals, and that the distribution over paths of length $r_1 + r_2$ is the concatenation of a distribution of a path of length r_1 with a path of length r_2 conditioned on the endpoint of the first path being the starting point of the second path. Thus, the distribution we end up with can be viewed as $\text{reverse}(p_1) \circ p_2$ where p_1, p_2 are path of length ℓ from Γ to Γ conditioned on them starting at the first point. Thus, the distribution of $\text{reverse}(p_1)$ and p_2 are of length ℓ conditioned on p_2 starting where $\text{reverse}(p_1)$ ended. \square

3.2.2 The Path Trick Distribution and 3-wise Correlations

One important feature of the path trick distribution that was already explored in [4, 5] is that it allows one to reduce upper bound powers of the 3-wise correlations as in Theorem 1.4 with respect to μ by 3-wise correlations with respect to μ_ℓ . Below is a formal statement.

Lemma 3.9. For all $t \geq 0$, and all 1-bounded functions $f: \Sigma^n \rightarrow \mathbb{C}$, $g: \Gamma^n \rightarrow \mathbb{C}$ and $h: \Phi^n \rightarrow \mathbb{C}$ we have that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h(z)] \right|^{2^t} \leq \left| \mathbb{E}_{(\vec{x},y,z) \sim \mu_{2^t}} [F(\vec{x})g(y)h'(z)] \right|,$$

where $F(\vec{x}) = f(x_{2^t-1}) \prod_{j=1}^{2^t-1} f(x_{2^j-1}) \overline{f(x_{2^j})}$ and $h'(z) = \mathbb{E}_{(x',y',z') \sim \mu} [f(x')g(y') \mid z' = z]$.

Proof. First, we note that

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

which proves the statement for $t = 1$, in which case the function F is just f . We now proceed by a sequence of Cauchy-Schwarz inequalities. To illustrate, note that

$$\begin{aligned} \left| \mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h'(z)] \right|^2 &= \left| \mathbb{E}_{z'} [h'(z') \mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y) \mid z = z']] \right|^2 \\ &\leq \mathbb{E}_{(x_1,x_2),(y_1,y_2),z \sim \nu_2} [f(x_1) \overline{f(x_2)} g(y_1) \overline{g(y_2)}]. \end{aligned}$$

Thus,

$$\begin{aligned} \left| \mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h'(z)] \right|^4 &\leq \mathbb{E}_{(x_1,x_2),(y_1,y_2),z \sim \nu_2} [f(x_1) \overline{f(x_2)} g(y_1) \overline{g(y_2)}]^2 \\ &= \mathbb{E}_{y'_2 \sim \mu_{y_2}} \left[\overline{g(y_2)} \mathbb{E}_{(x_1,x_2),(y_1,y_2),z \sim \nu_2} [f(x_1) \overline{f(x_2)} g(y_1) \mid y_2 = y'_2] \right]^2 \\ &\leq \mathbb{E}_{y'_2 \sim \mu_{y_2}} \left[\left| \mathbb{E}_{(x_1,x_2),(y_1,y_2),z \sim \nu_2} [f(x_1) \overline{f(x_2)} g(y_1) \mid y_2 = y'_2] \right|^2 \right] \\ &= \mathbb{E}_{\vec{x}, \vec{y} \sim \nu_4} [F(x_1, x_2, x_3, x_4) g(y_1) \overline{g(y_3)}], \end{aligned}$$

and inductively we get that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h'(z)] \right|^{2^t} \leq \mathbb{E}_{\vec{x}, \vec{y}, \vec{z} \sim \nu_{2^t}} [F'(\vec{x})g(y_1) \overline{g(y_{2^t-1+1})}]$$

where $F'(\vec{x}) = \prod_{j=1}^{2^t-1} f(x_{2^j-1}) \overline{f(x_{2^j})}$ is given as in the statement of the lemma. To finish the proof, note that

$$\begin{aligned} \mathbb{E}_{\vec{x}, \vec{y}, \vec{z} \sim \nu_{2^t}} [F'(\vec{x})g(y_1) \overline{g(y_{2^t})}] &= \mathbb{E}_{\vec{x}, \vec{y}, \vec{z} \sim \nu_{2^t}} [F'(\vec{x})g(y_1) \mathbb{E}_{x_{2^t}, y_{2^t-1}} [f(x_{2^t})g(y_{2^t-1+1}) \mid z_{2^t-1}]] \\ &= \mathbb{E}_{\vec{x}, y, z \sim \mu_{2^t}} [F'(\vec{x})g(y_1)h'(z_{2^t-1})], \end{aligned}$$

where in the last transition we used the definition of h' . □

Some remarks regarding Lemma 3.9 that are essential to make an effective use of it, are in order. Lemma 3.9 allows us to relate upper bound correlations with respect to μ with correlations with respect to μ_ℓ , but this comes at the expense of several complications:

1. **The role of y and z .** Note that in the above lemma, the role played by the y and z function in the premise are interchangeable, but this is not the case in the conclusion of the lemma. Namely, in the above formulation we kept the y function to be the same but replaced the z -function from h to h' . It is possible however, and this will be important for us, that we might as well have kept the h function in place and changed the function g into a function g' analogously to the above.
2. **Applying the path-trick with respect to other directions.** We could also apply the path trick with respect to y (or z) instead of on x as in the above formulation, in which case one gets an analogous statement to the above. It is always the case that the variable we apply the path trick on results in the function on that variable becoming a more complicated “product version” of the previous function. As for the other two functions, one of them stays put, whereas on the last one we have essentially no control over and it may change altogether. As explained above, there is flexibility for us in the choice of which function stays and on which we lose control over, and we will utilize this.
3. **The relationship between f and F .** The function on x , f , becomes a more complicated function F . In our argument for the proof of Theorem 1.4 we will make use of this lemma several times, and on each invocation the x -function will get more and more complicated, until eventually we will reach a distribution μ' on which we will do a direct analysis. At that point, we will be able to conclude a structural result on the x -function in that correlation, and from that deduce a structural result regarding f . In that respect, it is important to keep in mind the relationship between the x -functions; at each invocation of the path trick we will either keep the x function the same (this is the role played by the y -function in the formulation above), or else it will be a multiplication of the previous function on several inputs that are correlated in some way.

3.2.3 The Structure of Embeddings after the Path Trick

In [4, 5] it is shown that if μ does not admit an Abelian embedding, then the path trick distribution μ_ℓ also does not admit any Abelian embedding. In this section, we extend this connection by observing that, in a sense, even in the presence of Abelian embeddings, the path trick distribution has the “same” Abelian embeddings as the original distribution. More precisely:

Lemma 3.10. *Let Σ , Γ and Φ be finite alphabets, let μ be a distribution over $\Sigma \times \Gamma \times \Phi$, let $\ell \in \mathbb{N}$ be odd and let μ_ℓ be the path trick distribution applied on μ with respect to x , which is a distribution over $\Sigma' \times \Gamma \times \Phi$ where $\Sigma' \subseteq \Sigma^\ell$. If $\sigma: \Sigma' \rightarrow (H, +)$, $\gamma: \Gamma \rightarrow (H, +)$ and $\phi: \Phi \rightarrow (H, +)$ is an Abelian embedding of μ_ℓ in $(H, +)$, then there exists an $\sigma': \Sigma \rightarrow (H, +)$ such that*

1. (σ', γ, ϕ) form an embedding of μ in $(H, +)$.
2. For all $(x_1, \dots, x_\ell) \in \Sigma'$ we have that

$$\sigma(x_1, \dots, x_\ell) = \sum_{j=1}^{\ell} (-1)^{j-1} \sigma'(x_j).$$

Proof. Consider a sample $(\vec{x}, \vec{y}, \vec{z})$ as in ν_ℓ ; that is, we consider a sample according to μ_ℓ but record all the y 's and z 's generated in the process, so that (x_{2j-1}, y_j, z_j) and (x_{2j}, y_{j+1}, z_j) are in $\text{supp}(\mu)$ for all j .

Let $v_j \in \Sigma^\ell$ be the vector whose coordinates are all equal to x_{2j-1} , and $u_j \in \Sigma^\ell$ be the vector whose coordinates are all x_{2j} . Note that as $(v_j, y_j, z_j) \in \text{supp}(\mu_\ell)$ (as there is a path from y_j to z_j of length ℓ that just goes back and forth on the edge labeled by x_{2j-1}) and $(u_j, y_{j+1}, z_j) \in \text{supp}(\mu_\ell)$ (for similar reasons), it follows that

$$\sigma(v_j) + \gamma(y_j) + \phi(z_j) = 0, \quad \sigma(u_j) + \gamma(y_{j+1}) + \phi(z_j) = 0 \quad (5)$$

for all j . Denote $\sigma'(x) = \sigma(x, x, \dots, x)$, and note that (σ', γ, ϕ) form an Abelian embedding of μ into $(H, +)$. In the rest of the proof, we show that the formula in the second item of the statement holds. For that, we look at (5) for all $j = 1, \dots, (\ell + 1)/2$, multiply the left equations by 1 and the second one by -1 and add up to get that

$$\sum_{j=1}^{\ell} (-1)^{j-1} \sigma'(x_j) + \sigma(y) + \sigma(z) = 0,$$

where $y = y_1$ is the starting point of the walk and $z = z_{(\ell+1)/2}$ is the end point of the walk. We also have that $\sigma(\vec{x}) + \gamma(y) + \phi(z) = 0$ as (σ, γ, ϕ) form an embedding of μ_ℓ , so we get that

$$\sigma(\vec{x}) = \sum_{j=1}^{\ell} (-1)^{j-1} \sigma'(x_j),$$

finishing the proof. □

Lemma 3.10 has a few important consequences. Later on in Section 11, we will use it in order to convert structural results for F to structural results for f .³For now though, we shall use Lemma 3.10 in a different way in order to achieve the “master embedding is saturated” property as explained earlier. Towards this end, we introduce the following convenient notation.

Definition 3.11. For $\sigma: \Sigma \rightarrow (H, +)$, we define $\sigma^{\#\ell}: \Sigma^\ell \rightarrow (H, +)$ by

$$\sigma^{\#\ell}(x_1, \dots, x_\ell) = \sum_{j=1}^{\ell} (-1)^{j+1} \sigma(x_j).$$

Often times, σ will be part of some embedding of μ , in which case it will be more natural to view the domain of $\sigma^{\#\ell}$ as $\Sigma' \subseteq \Sigma^\ell$ where Σ' is the support of the first coordinate of μ_ℓ .

3.3 Saturating the Master Embedding via the Path Trick

Recalling the definition of the master embedding from Definition 3.6, there is no reason that its image would be the entire group $(H, +)$ (and in fact, typically it would not be). The goal of this section is to use the path trick to move from the distribution μ to a related distribution μ' such that the master embedding in μ' has as image a group. We moreover assert that it “suffices” to prove Theorem 1.4 for μ' , in the sense that then we would be able to deduce it for μ . For the sake of this section, we will focus on the first point – namely that the master embeddings have full images, and ignore the second point for now; this last deduction is covered in Section 11.

³Indeed, the above formula suggests that if F is correlated with an embedding function on μ_ℓ , then the embedding function can be broken to a product function over the x_i 's just the same way as F can be, and one thus expects to be able to argue that f itself is correlated with an embedding function.

3.3.1 The Evolution of the Master Embeddings under Path Tricks

The following lemma explains the way the master embeddings of a distribution evolve after an application of the path trick.

Lemma 3.12. *Let Σ , Γ and Φ be finite alphabets, let μ be a distribution over $\Sigma \times \Gamma \times \Phi$, let $\ell \in \mathbb{N}$ be odd and let μ_ℓ be the path trick distribution applied on μ with respect to x , which is a distribution over $\Sigma' \times \Gamma \times \Phi$ where $\Sigma' \subseteq \Sigma^\ell$.*

If $\sigma_{\text{master}}, \gamma_{\text{master}}, \phi_{\text{master}}$ is a master embedding with respect to μ , then $\sigma_{\text{master}}^{\#\ell}, \gamma_{\text{master}}, \phi_{\text{master}}$ is a master embedding for μ_ℓ .

Proof. Let $(\sigma', \gamma', \phi')$ be an Abelian embedding of μ_ℓ . By Lemma 3.10 we can find an embedding $(\sigma'', \gamma', \phi')$ of μ where $\sigma' = \sigma''^{\#\ell}$ on the domain of σ' . By Lemma 3.7, there is an i such that $(\sigma_{\text{master},i}, \gamma_{\text{master},i}, \phi_{\text{master},i})$ is a linear refinement of $(\sigma'', \gamma', \phi')$, and it follows that $(\sigma_{\text{master},i}^{\#\ell}, \gamma_{\text{master},i}, \phi_{\text{master},i})$ is a linear refinement of $(\sigma''^{\#\ell}, \gamma', \phi') = (\sigma', \gamma', \phi')$, as required. \square

In words, Lemma 3.12 says that once we have a master embedding for a distribution and we apply the natural operations on it, then it is also a master embedding for μ_ℓ . Hence, we are not losing anything with respect to the master embedding while performing path tricks.

3.3.2 The Path Trick Preserves Pairwise Connectedness

The following simple lemma shows that path tricks also preserve pairwise connectedness, and in fact improve pairwise connectedness with respect to 2 of the coordinates.

Lemma 3.13. *Let Σ , Γ and Φ be finite alphabets, let μ be a distribution over $\Sigma \times \Gamma \times \Phi$, let $\ell \in \mathbb{N}$ be odd and let μ_ℓ be the path trick distribution applied on μ with respect to x , which is a distribution over $\Sigma' \times \Gamma \times \Phi$ where $\Sigma' \subseteq \Sigma^\ell$.*

If μ is pairwise connected, then μ_ℓ is pairwise connected. Furthermore, for large enough ℓ depending only on the alphabet sizes, the graph between Γ and Φ becomes complete.

Proof. First, note that for every $x \in \Sigma$ and $y \in \Gamma$ such that there is $z \in \Phi$ for which (x, y, z) is in the support of μ , it holds that (\vec{x}, y, z) is in the support of μ_ℓ , where $\vec{x} = (x, \dots, x)$. Thus, by connectedness of μ it follows that there is a path between any \vec{x} and $y \in \Gamma$ in the graph of μ_ℓ , and as any $\vec{x}' \in \Sigma'$ is connected to some $y \in \Gamma$, it follows that the Σ', Γ graph in μ_ℓ is connected. Analogously, we have that the Σ', Φ graph is also connected.

As for the Γ, Φ graph, note that it is connected in μ , hence there is an odd number w such that between any $y \in \Gamma$ and $z \in \Phi$ there is a path of length at most w (of odd length). Note that in that case there will also be a path of length exactly w (as one can always do steps that go from z to some neighbour of it and back), hence we get that for $\ell = w$ the graph between Γ and Φ in μ_ℓ is complete. \square

3.3.3 The Path Trick Helps in Saturating the Master Embeddings

It is easy to observe that if ℓ is odd, then the image of $\sigma^{\#\ell}$ always contains the image of σ (by considering inputs of the form (x, \dots, x)). Additionally, the image of the other two components remains the same (and in particular does not shrink). In light of the formula in Definition 3.11, intuition suggests that as long as the image of σ is not a subgroup, the image of $\sigma^{\#\ell}$ would have to be larger (as we are considering signed sums). In this context, once the image becomes a sub-group it can never further increase. A natural hypothesis

therefore would be that by applying enough path tricks, the image of σ would eventually have to become a sub-group, in which point we will refer to σ as *saturated*.

Strictly speaking, this does not have to be the case if one applies the path trick in the naive way. Nevertheless, if one is willing to apply alternating path tricks on all coordinates, then eventually σ does become saturated. To show that, we first establish the following lemma, asserting that if the image of σ is not a sub-group, then one may enlarge it by applying the path trick 3 times:

Lemma 3.14. *Let Σ, Γ and Φ be finite alphabets, let μ be a pairwise connected distribution over $\Sigma \times \Gamma \times \Phi$. Then there are constants ℓ_1, ℓ_2 and $\ell_3 = 3$ depending only on the alphabet sizes such that if (σ, γ, ϕ) is a master embedding for μ into $(H, +)$, wherein each one of $\text{Image}(\sigma)$, $\text{Image}(\gamma)$, $\text{Image}(\phi)$ contains 0, then the following holds. Consider the distributions:*

1. μ' which is the result of the application of the path trick on μ with respect to z of length ℓ_1 , producing the master embedding $(\sigma', \gamma', \phi') = (\sigma, \gamma, \phi^{\#\ell_1})$.
2. μ'' which is the result of the application of the path trick on μ' with respect to y of length ℓ_2 , producing the master embedding $(\sigma'', \gamma'', \phi'') = (\sigma', \gamma^{\#\ell_2}, \phi')$.
3. μ''' which is a result of the application of the path trick on μ'' with respect to x of length ℓ_3 , producing the master embedding $(\sigma''', \gamma''', \phi''') = (\sigma''^{\#\ell_3}, \gamma'', \phi'')$.

then if $\text{Image}(\sigma)$ is not a subgroup of H , then $\text{Image}(\sigma) \subsetneq \text{Image}(\sigma''')$. Furthermore, the support of μ''' on its first coordinate (namely, its x -coordinate) is full, that is, Σ^{ℓ_3} .

Proof. Assume that $\text{Image}(\sigma)$ is not a subgroup of H ; then it is not closed under addition, so there are $x_1, x_2 \in \Sigma$ such that $\sigma(x_1) + \sigma(x_2) \notin \text{Image}(\sigma)$. We will show that after we choose ℓ_1 and ℓ_2 appropriately large, we could take $\ell_3 = 3$ and get that the support on Σ^3 is full, hence taking x_1, x_2 and some $x^* \in \Sigma$ for which $\sigma(x^*) = 0$, we would get that $\sigma'''(x_1, x^*, x_2) = \sigma(x_1) - \sigma(x^*) + \sigma(x_2)$ is in $\text{Image}(\sigma''')$ and not in $\text{Image}(\sigma)$.

Denote the alphabets in question as: μ' is a distribution over $\Sigma \times \Gamma \times \Phi'$, μ'' is a distribution over $\Sigma \times \Gamma' \times \Phi'$ and μ''' is a distribution over $\Sigma' \times \Gamma' \times \Phi'$. By Lemma 3.13, we may take ℓ_1 large enough so that the support of μ' on $\Sigma \times \Gamma$ is full, and fixing ℓ_1 we may choose ℓ_2 large enough so that the support of μ'' on $\Sigma \times \Phi'$ is full. We then pick $\ell_3 = 3$.

We show that any $(x, x', x'') \in \Sigma^3$ is in Σ' . Pick some $y \in \Gamma$ and look at $\vec{y} \in \Gamma^{\ell_2}$ which has all of its coordinates equal to y , and note that it is in Γ' . Since the support of μ' on $\Sigma \times \Gamma$ is full, we get that there are some $\vec{z}, \vec{z}' \in \Phi'$ such that (x, y, \vec{z}) and (x', y, \vec{z}') are in the support of μ' . It follows that (x, \vec{y}, \vec{z}) and (x', \vec{y}, \vec{z}') are both in the support of μ'' , so in the graph of μ'' we get a path from \vec{z} to \vec{y} to \vec{z}' labeled by x, x' . As the support of $\Sigma \times \Phi'$ in μ'' is full, we may continue this path by an edge labeled by x'' to get to some \vec{y}' . Overall, we get a path from \vec{z} to \vec{y} to \vec{z}' to \vec{y}' labeled by x, x', x'' , and this path means that (x, x', x'') is in the support of μ''' . \square

Using Lemma 3.14 iteratively, as long as our master embeddings do not have images which are sub-groups, we may enlarge them via consecutive applications of the path trick (while keeping all of the properties assumed for our original distribution, such as probability of atoms being $\Omega(1)$, pairwise connectedness and so on). It will be important for us to do this in a more careful manner, and maintain the fact that the alphabet of x is Σ^T for some $T \in \mathbb{N}$. With respect to that, iterating Lemma 3.14 directly gives saturation with respect to the embedding of x :

Lemma 3.15. *Let Σ, Γ and Φ be finite alphabets, let μ be a pairwise connected distribution over $\Sigma \times \Gamma \times \Phi$. Then there are constants T and ℓ depending only on the alphabet sizes such that if (σ, γ, ϕ) is a master embedding for μ into $(H, +)$, wherein each one of $\text{Image}(\sigma)$, $\text{Image}(\gamma)$, $\text{Image}(\phi)$ contains 0, then repeating the process in Lemma 3.14 T times one gets a distribution μ' over $\Sigma' \times \Gamma' \times \Phi'$ times where*

1. $\Sigma' = \Sigma^\ell$.

2. Letting $(\sigma', \gamma', \phi')$ be the induced master embedding of μ' , we have that $\text{Image}(\sigma')$ is a subgroup of H .

Proof. As long as $\text{Image}(\sigma)$ is not a sub-group, applying the process in Lemma 3.14 enlarges it, and the new alphabet of x is a power of the older alphabet of x . Hence repeating the process $T = |H|$ times gives eventually μ' on which σ' is saturated. \square

The next lemma is an analogous statement for y and z , and we show that the same procedure – applied sufficiently many times – also works. The argument is a bit more subtle, as we cannot afford ourselves to apply long path tricks on x ; such operation may not preserve the fact that the alphabet of x would remain a power of Σ . We remark that the role of y and z is symmetric, and thus while we formulate the statement in terms of z (as well as in Lemma 3.14), by flipping the roles of y and z one gets an analogous statement for y , which we shall also use.

Lemma 3.16. *Let Σ, Γ and Φ be finite alphabets, let μ be a pairwise connected distribution over $\Sigma \times \Gamma \times \Phi$. Then there are constants T, T' and ℓ depending only on the alphabet sizes such that if (σ, γ, ϕ) is a master embedding for μ into $(H, +)$, wherein each one of $\text{Image}(\sigma)$, $\text{Image}(\gamma)$, $\text{Image}(\phi)$ contains 0, then the following holds. Consider the distributions:*

1. Apply the transformation in Lemma 3.14 T times to get a distribution μ' . Let $(\sigma', \gamma', \phi')$ be the induced master embedding.
2. Apply the path trick on μ' with respect to y for T' times to get a distribution μ'' .
3. Apply the path trick on μ' with respect to z for 3 steps to get a distribution μ''' . Let $(\sigma''', \gamma''', \phi''')$ be the induced master embedding.

If $\text{Image}(\phi)$ is not a subgroup of H , then $\text{Image}(\phi) \subsetneq \text{Image}(\phi''')$. Furthermore, the support of μ''' on its first coordinate (namely, its x -coordinate) is full, that is, Σ^ℓ .

Proof. Consider the transformation in Lemma 3.14: given a distribution μ over $\Sigma \times \Gamma \times \Phi$, it outputs a distribution ν over $\Sigma^3 \times \Gamma' \times \Phi'$ where $\Gamma' \subseteq \Gamma^{\ell_2}$ and $\Phi' \subseteq \Phi^{\ell_3}$ for some ℓ_2, ℓ_3 depending only on the alphabet sizes. Let $\Gamma_{\text{original}} \subseteq \Gamma^{\ell_2}$ and $\Phi_{\text{original}} \subseteq \Phi^{\ell_3}$ be copies of Γ and Φ in ν , namely

$$\Gamma_{\text{original}} = \{(y, \dots, y) \mid y \in \Gamma\}, \quad \Phi_{\text{original}} = \{(z, \dots, z) \mid z \in \Phi\},$$

and note that $\Gamma_{\text{original}} \subseteq \Gamma'$ and $\Phi_{\text{original}} \subseteq \Phi'$. Note that if there is a path of length at most 3 between y and z in μ , then there is an edge between (y, \dots, y) and (z, \dots, z) in ν . Indeed, if there was a path y to z' to y' to z labeled by x_1, x_2, x_3 , then $((x_1, x_2, x_3), (y, \dots, y), (z, \dots, z))$ would be in $\text{supp}(\nu)$. Repeating this transformation twice get that if there was a path from y to z of length at most 4, then we would have an edge between (y, \dots, y) and (z, \dots, z) in the new distribution.

Thus, as μ is pairwise connected we can apply the transformation in Lemma 3.14 T times, where T only depends on the alphabet sizes, to get a distribution μ' whose alphabet on x is Σ^ℓ for some ℓ and its support

on $\Gamma_{\text{original}} \times \Phi_{\text{original}}$ is full. Let $(\sigma', \gamma', \phi')$ be the induced master embedding as in Definition 3.11. We then apply the path trick on μ' with respect to y for T' steps so that the support of the distribution over x and z is full (thanks to Lemma 3.13), and we think of μ'' as a distribution over $\Sigma'' \times \Gamma'' \times \Phi''$. We remark that the support of μ'' on $\Gamma_{\text{original}} \times \Phi_{\text{original}}$ is still full (this is preserved under the last application of the path trick), and also that $\Sigma'' = \Sigma^{3T}$.

We now apply the path trick on μ'' for length 3 on z to get the distribution μ''' over $\Sigma''' \times \Gamma''' \times \Phi'''$, where $\Sigma''' = \Sigma'' = \Sigma^{3T}$, and we have an induced master embedding of μ''' , denoted by $(\sigma''', \gamma''', \phi''')$ as in Definition 3.11. As in Lemma 3.14 we have that $\text{Image}(\phi) \subseteq \text{Image}(\phi''')$, and we next argue that if $\text{Image}(\phi)$ is not a sub-group, then this is a strict containment.

If $\text{Image}(\phi)$ is not a sub-group of H , then we may find $z_1, z_3 \in \Phi$ such that $\phi(z_1) + \phi(z_3) \notin \text{Image}(\phi)$, and we pick $z_2 \in \Phi$ such that $\phi(z_2) = 0$. Below we show that $((z_1, \dots, z_1), (z_2, \dots, z_2), (z_3, \dots, z_3))$ is in the support of μ''' on Φ''' , and we now argue that this would give the strict containment. Indeed, by Definition 3.11 we have that $\phi'''((z_1, \dots, z_1), (z_2, \dots, z_2), (z_3, \dots, z_3))$ is equal to

$$\phi'(z_1, \dots, z_1) - \phi'(z_2, \dots, z_2) + \phi'(z_3, \dots, z_3)$$

and as $\phi'(z_1, \dots, z_1) = \phi(z_1)$, $\phi'(z_2, \dots, z_2) = \phi(z_2) = 0$ and $\phi'(z_3, \dots, z_3) = \phi(z_3)$, we get that $\text{Image}(\phi''') \subsetneq \text{Image}(\phi)$ as desired.

To show that $((z_1, \dots, z_1), (z_2, \dots, z_2), (z_3, \dots, z_3))$ is in the support of μ''' on Φ''' , take any $y \in \Gamma$ so that since the support of μ'' on the set $\Gamma_{\text{original}} \times \Phi_{\text{original}}$ is full, we get that there are $x_1, x_2 \in \Sigma''$ such that $(x_1, (y, \dots, y), (z_1, \dots, z_1))$ and $(x_2, (y, \dots, y), (z_2, \dots, z_2))$ are in $\text{supp}(\mu'')$. As in μ'' we have that the support on $\Sigma'' \times \Phi''$ is full, it follows that there is \vec{y} such that $(x_2, \vec{y}, (z_3, \dots, z_3))$ is in $\text{supp}(\mu'')$. We note that now in μ'' we have a path from x_1 to (y, \dots, y) to x_2 to \vec{y} , whose labels are (z_1, \dots, z_1) , (z_2, \dots, z_2) and (z_3, \dots, z_3) , and therefore $((z_1, \dots, z_1), (z_2, \dots, z_2), (z_3, \dots, z_3))$ is in the support of μ''' on Φ''' as required. \square

We now combine Lemmas 3.15, 3.16 to get our overall transformation that saturates the master embeddings.

Lemma 3.17. *Let Σ, Γ and Φ be finite alphabets of size at most m , and let μ be a pairwise connected distribution in which the probability of each atom is at least $\alpha > 0$, and let $(\sigma_{\text{master}}, \gamma_{\text{master}}, \phi_{\text{master}})$ be a master embedding for μ into an Abelian group $(H, +)$. Then, there are $\alpha' = \alpha'(m, \alpha) > 0$, $\ell = \ell(m) \in \mathbb{N}$ and a distribution μ' over $\Sigma' \times \Gamma' \times \Phi'$ which results from μ by a sequence of applications of the path trick, such that:*

1. μ' is pairwise connected and the probability of each atom is at least α' .
2. The number of applications of the path trick is at most some $T \in \mathbb{N}$ depending only on m . Furthermore, the new alphabet sizes $|\Sigma'|$, $|\Gamma'|$ and $|\Phi'|$ are bounded by some function of m .
3. There is a master embedding $(\sigma'_{\text{master}}, \gamma'_{\text{master}}, \phi'_{\text{master}})$ of μ' into $(H, +)$ such that:
 - (a) The image of each one of σ'_{master} , γ'_{master} and ϕ'_{master} is H .
 - (b) This master embedding is given by a master embedding of μ by the transformations described in Definition 3.11 following the applications of the path trick.
 - (c) The support of $(\sigma'_{\text{master}}(x), \gamma'_{\text{master}}(y), \phi'_{\text{master}}(z))$ where $(x, y, z) \sim \mu'$ is full, namely it is

$$\{(a, b, c) \in H \mid a + b + c = 0\}.$$

(d) The support of μ' on the first coordinate is $\Sigma' = \Sigma^\ell$.

Proof. We proceed by an iterative process. Starting with the distribution μ and a master embedding for it into an Abelian group $(H, +)$, so long as the image of one of the embedding's component is not a sub-group, we apply either Lemma 3.15 or Lemma 3.16 to enlarge it (while clearly not decreasing the size of the image of the other two components), so eventually we get to a distribution ν and master embeddings σ' , γ' and ϕ' such that $\text{Image}(\sigma') = H_1$, $\text{Image}(\gamma') = H_2$ and $\text{Image}(\phi') = H_3$ where H_1, H_2 and H_3 are subgroups of H . We then take the distribution μ' which is a result of applying the path trick of ν with respect to z for $T = T(m) \in \mathbb{N}$ times for sufficiently large T , so that by Lemma 3.13 the distribution of μ' over the first two coordinates is full. Let the alphabets of μ' be Σ', Γ' and Φ' ; then by Lemmas 3.15, 3.16 we see that $\Sigma' = \Sigma^\ell$ for some $\ell = \ell(m) \in \mathbb{N}$.

We now argue that $H_1 = H_2 = H_3$. Letting the alphabets of μ' be Σ', Γ' and Φ' , note that by Lemma 3.13 we may apply the path trick on μ' to get μ'' in which the support on $\Sigma' \times \Gamma'$ is full, and we argue that this means that $H_1 + H_2 \subseteq H_3$. Indeed, we could pick any $h_1 \in H_1, h_2 \in H_2$ and find $x \in \Sigma'$ and $y \in \Gamma'$ such that $\sigma'_{\text{master}}(x) = -h_1$ and $\gamma'_{\text{master}}(y) = -h_2$, and thus find $\vec{z} \in \Phi''$ such that (x, y, \vec{z}) is in the support of μ'' , so

$$-h_1 - h_2 + \phi'_{\text{master}}(\vec{z}) = \sigma'_{\text{master}}(x) + \gamma'_{\text{master}}(y) + \phi'_{\text{master}}(\vec{z}) = 0$$

by definition of embeddings, so $\phi'_{\text{master}}(\vec{z}) = h_1 + h_2$, hence $h_1 + h_2$ are in the image of $\phi'_{\text{master}}(\vec{z})$, which is the same as the image of ϕ' (as it is already a subgroup), so $h_1 + h_2 \in H_3$.

Thus, $H_1 + H_2 \subseteq H_3$ and analogously $H_1 + H_3 \subseteq H_2$ and $H_2 + H_3 \subseteq H_1$, and it follows that $H_1 = H_2 = H_3$. Thus, we can view $(\sigma'_{\text{master}}, \gamma'_{\text{master}}, \phi'_{\text{master}})$ as an embedding into $(H_1, +)$, so now it is a master embedding satisfying the third bullet. The first two bullets are clear, as the number of path trick applications is some constant depending only on the alphabet sizes (and the size of H , which also only depends on the alphabet sizes of μ). \square

3.4 Conclusion of Section 3

Using Lemma 3.17 and Lemma 3.9 together, one gets that an expectation as in (1) over μ can be upper bounded by (some power bounded away from 0) of a similar looking expectation over μ' in which the master embeddings are saturated; by applying a few more path tricks and using Lemma 3.13, we can also ensure further connectedness properties of μ' (which we will need in the future). Thus, we have gained further important properties of our μ , at the expense of:

1. The y -function and z -function may become completely different as a result of these operations (but they remain bounded).
2. The x -function becomes more complicated. Indeed, our x -alphabet will be some $\Sigma' \subseteq \Sigma^\ell$, and our x -function will be given as

$$F(x_1, \dots, x_\ell) = \prod_{i=1}^{\ell} f_i(x_i),$$

where each f_i is either the function f or its complex conjugate \bar{f} .

Thus, to prove Theorem 1.4, it suffices to (1) prove a version of that theorem under the additional assumptions we gained on μ' ; (2) prove that a structural result for F as in Theorem 1.4 implies a similar structural result for f . The majority of our effort will be to establish the first step: this part of the argument is contained

in Sections 4, 5, 6, 7, 8, 9 and 10. The second step will be a relatively easy consequence: this part of the argument is contained in Section 11. Thus, we arrive at the following statement which is the same as the statement of Theorem 1.4, except that we have additional assumptions on the distribution μ :

Theorem 3.18. *For all $m \in \mathbb{N}$, $\alpha > 0$ and $\varepsilon > 0$, there exists $d \in \mathbb{N}$ and $\varepsilon' > 0$ such that the following holds. Suppose that μ is a distribution over $\Sigma \times \Gamma \times \Phi$ such that:*

1. *The probability of each atom is at least α .*
2. *The size of each one of Σ, Γ, Φ is at most m .*
3. *$\text{supp}(\mu)$ is pairwise connected.*
4. *There is a master embedding (σ, γ, ϕ) of μ into an Abelian group $(H, +)$ which is saturated, and the distribution of $(\sigma(x), \gamma(y), \phi(z))$ where $(x, y, z) \sim \mu$ has full support on $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$.*

Then, if $f: \Sigma^n \rightarrow \mathbb{C}$, $g: \Gamma^n \rightarrow \mathbb{C}$ and $h: \Phi^n \rightarrow \mathbb{C}$ are 1-bounded functions such that

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

then there are 1-bounded functions $u_1, \dots, u_n: \Sigma \rightarrow \mathbb{C}$ and a function $L: \Sigma^n \rightarrow \mathbb{C}$ of degree at most d and 2-norm at most 1 such that

$$\left| \mathbb{E}_{x \sim \mu_x^{\otimes n}} \left[f(x) \cdot L(x) \prod_{i=1}^n u_i(x_i) \right] \right| \geq \varepsilon'.$$

Furthermore, there are $\chi_i \in \hat{H}$ such that for all i , $u_i(x_i) = \chi_i(\sigma(x_i))$. Quantitatively, we have $d = \text{poly}_{m,\alpha}(\frac{1}{\varepsilon})$ and $\varepsilon' = 2^{-\text{poly}_{m,\alpha}(\frac{1}{\varepsilon})}$.

3.5 Embeddings Into the Infinite Cyclic Group

So far we have discussed embeddings of a distribution into finite Abelian groups, however it also makes sense to consider embeddings into infinite groups. Specifically, we will need to consider embeddings of a distribution into the infinite cyclic group $([0, 1), (\text{mod } 1))$. Using approximation arguments (and more specifically, Dirichlet's Approximation Theorem), we show in the following lemma that any embedding of a distribution into $([0, 1), (\text{mod } 1))$ is equivalent to an embedding into a finite Abelian group, hence there is nothing particularly special about them. For this, we first define the notion of equivalence.

Definition 3.19. *Let μ be a distribution over $\Sigma \times \Gamma \times \Phi$, and let (σ, γ, ϕ) and $(\sigma', \gamma', \phi')$ be Abelian embeddings of μ . We say (σ, γ, ϕ) is equivalent to $(\sigma', \gamma', \phi')$ if there are bijective maps $m_1: \text{Image}(\sigma) \rightarrow \text{Image}(\sigma')$, $m_2: \text{Image}(\gamma) \rightarrow \text{Image}(\gamma')$ and $m_3: \text{Image}(\phi) \rightarrow \text{Image}(\phi')$ such that $\sigma'(x) = m_1(\sigma(x))$, $\gamma'(y) = m_2(\gamma(y))$ and $\phi'(z) = m_3(\phi(z))$ for all $x \in \Sigma$, $y \in \Gamma$, $z \in \Phi$.*

Lemma 3.20. *Let Σ, Γ and Φ be finite alphabets and let μ be a distribution over $\Sigma \times \Gamma \times \Phi$. If $\sigma: \Sigma \rightarrow [0, 1)$, $\gamma: \Gamma \rightarrow [0, 1)$ and $\phi: \Phi \rightarrow [0, 1)$ is an embedding of μ into $([0, 1), + (\text{mod } 1))$, then (σ, γ, ϕ) is equivalent to an Abelian embedding of μ into a finite Abelian group.*

Proof. 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', γ', ϕ' 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 $\frac{3}{qN^{1/r}} < \frac{1}{q}$. 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 $\frac{1}{q}$ far from all integers. It follows that it is an integer, so $\sigma'(x) + \phi'(y) + \gamma'(z) = 0 \pmod{1}$.

2. Second, we argue that $(\sigma', \gamma', \phi')$ is equivalent to (σ, γ, ϕ) . For that, we have to argue that $\sigma(x) \neq \sigma(x')$ if and only if $\sigma'(x) \neq \sigma'(x')$. If $\sigma(x) = \sigma(x')$ then it is clear that $\sigma'(x) = \sigma'(x')$ by definition. If $\sigma(x) \neq \sigma(x')$, then by the definition of α we get that $\sigma(x) - \sigma(x')$ is at least α -far from all integers, and as $|\sigma(x) - \sigma'(x')| \leq \frac{\alpha}{3}$, it follows that $\sigma'(x) - \sigma'(x')$ is at least $\alpha/3$ far from all integers, and in particular from 0, so $\sigma'(x) \neq \sigma'(x')$.

In conclusion, we get that (σ, γ, ϕ) and $(\sigma', \gamma', \phi')$, and noting that, after multiplying by q , the latter is an embedding into $(\mathbb{Z}_q, +)$, the proof is concluded. \square

4 Non-embedding Degrees and Partial Bases

In this section we make progress towards the proof of Theorem 3.18, and state Theorem 4.13 which is a related by weaker form. The proof of Theorem 4.13 then spans Sections 6, 7, 8, 9, and the derivation of Theorem 3.18 from Theorem 4.13 is done in Section 10.

4.1 A Motivating Case

Let μ be a distribution over $\Sigma \times \Gamma \times \Phi$ as in Theorem 3.18, and let σ, γ, ϕ be a saturated master embedding of μ into $(H, +)$. To motivate the discussion, below we begin by considering a motivating example in which the master embeddings partition the alphabets into singletons.

Namely, suppose that for each $h \in H$ each one of $\sigma^{-1}(h)$, $\gamma^{-1}(h)$ and $\phi^{-1}(h)$ has size exactly 1. In that case the master embeddings form an identification between our alphabets and the group H , hence what we really have in our hands is 3 functions, $f^\sharp: H^n \rightarrow \mathbb{C}$, $g^\sharp: H^n \rightarrow \mathbb{C}$ and $h^\sharp: H^n \rightarrow \mathbb{C}$ defined as

$$\begin{aligned} f^\sharp(h_1, \dots, h_n) &= f(\sigma^{-1}(h_1), \dots, \sigma^{-1}(h_n)), & g^\sharp(h_1, \dots, h_n) &= g(\gamma^{-1}(h_1), \dots, \gamma^{-1}(h_n)), \\ h^\sharp(h_1, \dots, h_n) &= h(\phi^{-1}(h_1), \dots, \phi^{-1}(h_n)). \end{aligned}$$

Thus, considering the distribution ν over H^3 which is the distribution of $(\sigma(x), \gamma(y), \phi(z))$ where $(x, y, z) \sim \mu$, we get that

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

Thus, we have transformed our question into an equivalent question over Abelian groups. As μ is pairwise connected, ν is also pairwise connected, and by definition of the master embedding it follows that its support is contained in $S = \{(x^\#, y^\#, z^\#) \in H^3 \mid x^\# + y^\# + z^\# = 0\}$. Combining the pairwise connectedness and the fact that the master embedding is saturated, it follows that the support of ν is precisely S . Using other ideas (based on random restrictions) we can ensure that the distribution ν is actually uniform over S , in which case we have reduced the problem to a well-known Fourier analytic computation (which appears in many places, such as Roth's theorem [37, 29] as well as in theoretical computer science [7, 24]). In particular, one can show that there is a Fourier character $\chi \in \hat{H}^n$ such that $|\widehat{f^\#}(\chi)| \geq \varepsilon$, and translating this back into information about the function f one gets the conclusion of Theorem 3.18 with the low-degree part L being the constant 1 function. We remark that even in this simplistic argument, the presence of the low-degree function ultimately comes from the step in which we switched from the distribution ν to the uniform distribution over S . Nevertheless, we encourage the reader to ignore this point for now.

Our goal in this, and in the several subsequent sections will be to show that while in general, it need not be the case that σ , γ and ϕ completely partition their respective alphabets, the only functions f , g and h for which the expectation in Theorem 3.18 may be have a special property. Specifically, we show that such the function f “hardly distinguish” between two input symbols x and x' that are mapped to the same group element by the master embedding component σ (and similarly for g and h). Towards this end, in this section we first define a partial basis for the set of functions composed of functions that only depend on the values of the master embeddings, and then complete them to bases. We then define the notions of “embedding degree” and “non-embedding degree”. These are notions that capture how well does our function f distinguish between inputs that are mapped to the same group element by the master embedding. With these notions, we show that f, g, h for which the expectation in Theorem 3.18 is large, must have small non-embedding degree.

4.2 Setting Up a Partial Basis via Saturated Embeddings, and Non-embedding Degrees

Let μ be a distribution over $\Sigma \times \Gamma \times \Phi$ as in Theorem 3.18 and let σ, γ, ϕ be saturated master embeddings into $(H, +)$. In this section, we explain how to use these embeddings to define useful partial bases for the spaces of functions we are dealing with, as well as how to define the notion the related notion of non-embedding degree. For the sake of concreteness, we shall phrase everything in the language of functions of x and the alphabet Σ , however everything holds for the other two variables and alphabets as well.

Consider the space $L_2(\Sigma, \mu_x)$, and note that we may set up a partial basis for them using characters over H and the master embeddings.

Definition 4.1. *Given a finite set Σ , an Abelian group $(H, +)$ and $\sigma: \Sigma \rightarrow H$, for each $\chi \in \hat{H}$, we define $\chi_\sigma: \Sigma \rightarrow \mathbb{C}$ by $\chi_\sigma(x) = \chi(\sigma(x))$.*

We note that in our setting, the set $\{\chi_\sigma\}_{\chi \in \hat{H}}$ is a linearly independent set. Indeed, to observe note that as σ is saturated, it is enough to show that there is a distribution \mathcal{D} over Σ in which χ_σ is an orthonormal set, and we consider a distribution \mathcal{D} over Σ such that $\sigma(x)$ is distributed uniformly in H when $x \sim \mathcal{D}$ (this is clearly possible). In that case, for all $\chi, \chi' \in \hat{H}$ we have that $\langle \chi_\sigma, \chi'_\sigma \rangle_{\mathcal{D}} = \langle \chi, \chi' \rangle = 1_{\chi=\chi'}$, as required.

Definition 4.2. Given a finite set Σ , an Abelian group $(H, +)$ and $\sigma: \Sigma \rightarrow H$, we define $\text{Embed}_\sigma(\mu) \subseteq \{f: \Sigma \rightarrow \mathbb{C}\}$ by

$$\text{Embed}_\sigma(\mu) = \text{Span} \left(\left\{ \chi_\sigma \mid \chi \in \hat{H} \right\} \right).$$

Thus, in our setting we have $\text{Embed}_\sigma(\mu) \subseteq L_2(\Sigma; \mu_x)$, $\text{Embed}_\gamma(\mu) \subseteq L_2(\Gamma; \mu_y)$ and $\text{Embed}_\phi(\mu) \subseteq L_2(\Phi; \mu_z)$. These spaces capture the space of embedding functions; an x function f is called an embedding function if $f(x)$ only depends on $\sigma_{\text{master}}(x)$.

Claim 4.3. Suppose that μ is a distribution over $\Sigma \times \Gamma \times \Phi$ in which the master embeddings are saturated, and let $f: \Sigma \rightarrow \mathbb{C}$, $g: \Gamma \rightarrow \mathbb{C}$ and $h: \Phi \rightarrow \mathbb{C}$ be functions such that

$$f(x) + g(y) + h(z) = 0$$

for all $(x, y, z) \in \text{supp}(\mu)$. Then $f \in \text{Embed}_\sigma(\mu)$.

Proof. Taking real and imaginary parts separately, it suffices to prove the statement for real valued functions. Multiplying f, g, h by small enough constant, we may assume that $|f(x)|, |g(y)|, |h(z)| \leq 1$ for all x, y, z , hence f, g, h form an embedding of μ into $([-1, 1], + \pmod{2})$. By Lemma 3.20, it follows that there are m_1, m_2, m_3 injectives such that $f' = m_1(f)$, $g' = m_2(g)$ and $h' = m_3(h)$ is an embedding of μ into a finite Abelian group. By the definition of the master embedding it follows that f' is constant on each part of the partition on Σ induced by the master embedding, and so $f'(x) = f'(x')$ if $\sigma_{\text{master}}(x) = \sigma_{\text{master}}(x')$. Since m_1 is injective, it follows that the same is true for f , and so $f \in \text{Embed}_\sigma(\mu)$. \square

The above motivating example is just the case that the 3 containments $\text{Embed}_\sigma(\mu) \subseteq L_2(\Sigma; \mu_x)$, $\text{Embed}_\gamma(\mu) \subseteq L_2(\Gamma; \mu_y)$ and $\text{Embed}_\phi(\mu) \subseteq L_2(\Phi; \mu_z)$ are in fact equalities. This need not be necessarily the case for us, hence we may need to complete these sets to get all of $L_2(\Sigma; \mu_x)$. More precisely, let $B_1 \subseteq \text{Embed}_\sigma(\mu)$ be an orthonormal basis for $\text{Embed}_\sigma(\mu)$ (with respect to the inner product in $L_2(\Sigma, \mu_x)$). We complete it to an orthonormal basis for $L_2(\Sigma, \mu_x)$ by adding the set $B_2 \subseteq L_2(\Sigma; \mu_x)$.

Definition 4.4. A monomial in $L_2(\Sigma; \mu_x)$ is one of the basis functions from $B_1 \cup B_2$. A monomial in $L_2(\Sigma^n, \mu_x^{\otimes n})$ is $\prod_{i=1}^n u_i$ where $u_i \in B_1 \cup B_2$ for all i .

Definition 4.5. The non-embedding degree of u , $\text{nedeg}(u)$, is the number of i 's for which $u_i \in B_2$.

With these notions, we may write any $f: \Sigma^n \rightarrow \mathbb{C}$ as

$$f(x) = \sum_{\vec{u}=(u_1, \dots, u_n) \in (B_1 \cup B_2)^n} \hat{f}(\vec{u}) \prod_{i=1}^n u_i(x_i),$$

where $\hat{f}(\vec{u}) = \langle f, \vec{u} \rangle$. Each monomial of f has its non-embedding degree, and we will want to define a notion of non-embedding degree which captures the mass of f on low non-embedding degree monomials and is convenient to work with, and towards this end we define the non-embedding noise stability of a function.

4.3 The Non-embedding Noise Stability of a Function

In this section, we define the notion of non-embedding stability, which will be a crucial tool for us to measure the degree of a function with respect to “non-embedding functions”. We also state a few basic properties of it that will be used later on in our arguments.

For a parameter $\xi > 0$ and a distribution \mathcal{D} over Σ , consider the Markov chain $T_{\text{non-embed},1-\xi,\mathcal{D},\mu}$ on Σ that on $x \in \Sigma$, with probability $1 - \xi$ stays in x , and otherwise samples $x' \sim \mathcal{D}$ conditioned on $\sigma(x') = \sigma(x)$. It will most often be the case for us that $\mathcal{D} = \mu_x$, however this operator depends on μ as a whole (as it depends on the master embedding of it), we chose to include both in the notations. There will be some rare exceptions though, in which case we will make the notations explicit. Otherwise, to simplify notations we will often drop \mathcal{D} from the notation (with the understanding that it is just μ_x).

Observe that μ_x is a stationary distribution for $T_{\text{non-embed},1-\xi,\mu}$. Thus we can think of $T_{\text{non-embed},1-\xi,\mu}$ as an operator acting on $L_2(\Sigma; \mu_x)$ as

$$T_{\text{non-embed},1-\xi,\mu}f(x) = \mathbb{E}_{x' \sim T_{\text{non-embed},1-\xi}x} [f(x')].$$

With this in mind, we may define the non-embedding stability of f as follows:

Definition 4.6. *The non-embedding ξ -noise stability of $f: (\Sigma^n; \mathcal{D}^{\otimes n}) \rightarrow \mathbb{C}$ is defined as*

$$\text{NEStab}_{1-\xi,\mu}(f; \mathcal{D}^{\otimes n}) = \langle f, T_{\text{non-embed},1-\xi,\mathcal{D},\mu}^{\otimes n} f \rangle.$$

We note that the operator $T_{\text{non-embed},1-\xi,\mu}$ depends on the distribution μ itself and not only on its marginal on x , as it is defined using the master embedding of μ ; the same goes for the non-embedding noise stability of a function. Nevertheless, and to simplify notations we will often omit μ from notations when it is clear, and denote the operator by $T_{\text{non-embed},1-\xi}$ and the corresponding notion of noise stability by $\text{NEStab}_{1-\xi}(f; \mu_x^{\otimes n})$.

4.3.1 Diagonalizing the Non-embedding Stability Operator

The basis functions B_1, B_2 defined earlier are eigenfunctions of the operator $T_{\text{non-embed},1-\xi,\mu}$, and the following fact gives us their eigenvalues:

Fact 4.7. *Suppose that $u: \Sigma^n \rightarrow \mathbb{C}$ is a monomial of non-embedding degree equal to d . Then*

$$T_{\text{non-embed},1-\xi,\mu}^{\otimes u} = (1 - \xi)^d u.$$

Proof. It suffices to show that in the 1-dimensional case, for $u \in B_1 \cup B_2$, if $u \in B_1$ then $T_{\text{non-embed},1-\xi,\mu}u = u$, and if $u \in B_2$ then $T_{\text{non-embed},1-\xi,\mu}u = (1 - \xi)u$.

For $u \in B_1$ this is clear, since for every $x \in \Sigma$ and every $x' \in \text{supp}(T_{\text{non-embed},1-\xi,\mu}x)$ it holds that $\sigma(x) = \sigma(x')$, and so $u(x) = u(x')$.

Fix $u \in B_2$ and fix $x \in \Sigma$. We have that

$$T_{\text{non-embed},1-\xi,\mu}u(x) = (1 - \xi)u(x) + \xi \mathbb{E}_{x' \sim \mu_x} [u(x') \mid \sigma(x) = \sigma(x')].$$

Let $s = \sigma(x)$, and let $p_s = \Pr_{x'} [\sigma(x') = s]$. Then the expectation on the right hand side is equal to

$$\mathbb{E}_{x' \sim \mu_x} [u(x') \mid \sigma(x') = s] = p_s^{-1} \mathbb{E}_{x' \sim \mu_x} [u(x') 1_{\sigma(x')=s}] = p_s^{-1} \langle u, 1_{\sigma(\cdot)=s} \rangle_{\mu_x}.$$

Note that for all $s \in H$, the function $1_{\sigma(\cdot)=s}$ is in the span of B_1 , and as $u \in B_2$ it is orthogonal to it, and so the last expression is 0. We conclude that $T_{\text{non-embed},1-\xi,\mu}u(x) = (1 - \xi)u(x)$. \square

4.3.2 Changing Noise Rates in Non-embedding Stability

A basic property of the non-embedding noise stability is that it decreases as a result for increasing the noise rate:

Claim 4.8. *Suppose that $0 \leq \rho_1 \leq \rho_2 \leq 1$. Then for every $f: (\Sigma^n; \mu_x^{\otimes n}) \rightarrow \mathbb{C}$ we have that*

$$\text{NEStab}_{\rho_1}(f; \mu_x^{\otimes n}) \leq \text{NEStab}_{\rho_2}(f; \mu_x^{\otimes n}).$$

Proof. Writing $f(x) = \sum_{u \in (B_1 \cup B_2)^n} \widehat{f}(u) \prod_{i=1}^n u_i(x_i)$, we have by Fact 4.7 that

$$\begin{aligned} \text{NEStab}_{\rho_1}(f; \mu_x^{\otimes n}) &= \langle f, T_{\text{non-embed}, \rho_1} f \rangle = \left\langle \sum_{u \in (B_1 \cup B_2)^n} \widehat{f}(u) u, \sum_{u \in (B_1 \cup B_2)^n} \widehat{f}(u) \rho_1^{\text{non-embed-deg}(u)} u \right\rangle \\ &= \sum_{u \in (B_1 \cup B_2)^n} \rho_1^{\text{non-embed-deg}(u)} \left| \widehat{f}(u) \right|^2, \end{aligned}$$

and the result follows as all $\left| \widehat{f}(u) \right|^2$ are non-negative. \square

4.3.3 Random Restrictions and Non-embedding Stability

The following claim is an instantiation of Lemma 2.8 (and in fact our primary application for that lemma), asserting that if we have a function f that has small noise non-embedding stability, then in expectation after random restrictions it still has small noise non-embedding stability.

Claim 4.9. *For all $m \in \mathbb{N}$, $\alpha > 0$ there is $c > 0$ such that the following holds. Let ν be a distribution over $\Sigma \times \Gamma \times \Phi$ and let \mathcal{D} , \mathcal{D}' and \mathcal{D}'' be distributions over Σ such that:*

1. *The probability of each atom in \mathcal{D} , \mathcal{D}' , \mathcal{D}'' is at least α .*
2. *$\mathcal{D} = \beta \mathcal{D}' + (1 - \beta) \mathcal{D}''$.*

Then, for all $f: (\Sigma^n, \mathcal{D}^{\otimes n}) \rightarrow \mathbb{C}$ we have

$$\mathbb{E}_{\substack{J \subseteq [n] \\ x'' \sim \mathcal{D}''^J}} [\text{NEStab}_{1-\delta, \nu}(f_{\mathcal{J} \rightarrow x''}; \mathcal{D}'^J)] \leq \text{NEStab}_{1-c\beta\delta, \nu}(f; \mathcal{D}).$$

Proof. This is an immediate consequence of Lemma 2.8. In the notation therein, the vertex set of the graph G is Σ , and x, x' are adjacent if $\sigma(x) = \sigma(x')$ where σ is the master embedding of x for the distribution ν . The left hand side and the right hand side in the above claim are precisely the left hand side and the right hand in Lemma 2.8. \square

4.4 Non-Embedding Influences of a Function

We will need the notion of non-embedding influences of a function defined as follows.

Definition 4.10. *Let μ be a distribution over $\Sigma \times \Gamma \times \Phi$, and let $\sigma: \Sigma \rightarrow H$, $\gamma: \Gamma \rightarrow H$ and $\phi: \Phi \rightarrow H$ be a master embedding. For a function $f: \Gamma^n \rightarrow \mathbb{C}$ and a coordinate $j \in [n]$, we define the non-embedding influence of f to be*

$$I_{j, \text{non-embed}}[f] = \mathbb{E}_{x' \sim \mu_x^{\otimes n}, a, b \sim \mu_x} \left[\left| f(x_{-j} = x', x_j = a) - f(x_{-j} = x', x_j = b) \right|^2 \mid \sigma(a) = \sigma(b) \right].$$

Non-embedding degrees are defined analogously for functions over y and z .

Definition 4.11. *In the setting of Definition 4.10, the total non-embedding influence of f is*

$$I_{\text{non-embed}}[f] = \sum_{j=1}^n I_{j,\text{non-embed}}[f].$$

We have the following easy fact.

Fact 4.12. *In the setting of Definition 4.10, we have*

1. $I_{j,\text{non-embed}}[f] = 2 \sum_{\chi: \chi_j \notin B_{\text{embed}}} \left| \widehat{f}(\chi) \right|^2.$
2. $I_{\text{non-embed}}[f] = \sum_{\chi} \text{nedeg}(\chi) \left| \widehat{f}(\chi) \right|^2.$

Proof. For the first bullet, we note that

$$\begin{aligned} f(x_{-j} = x', x_j = a) - f(x_{-j} = x', x_j = b) &= \sum_{\chi} \widehat{f}(\chi) (\chi(x', a) - \chi(x', b)) \\ &= \sum_{\chi} \widehat{f}(\chi) (\chi_j(b) - \chi_j(a)) \prod_{i \neq j} \chi_i(x'_i). \end{aligned}$$

For $\chi_j \in B_{\text{embed}}$, $\chi_j(b) - \chi_j(a) = 0$ whenever $\sigma(a) = \sigma(b)$, hence such terms give no contribution to the non-embedding influence. We thus get

$$\begin{aligned} &\mathbb{E}_{x',a,b} \left[\left| f(x_{-j} = x', x_j = a) - f(x_{-j} = x', x_j = b) \right|^2 \mid \sigma(a) = \sigma(b) \right] \\ &= \sum_{\substack{\chi, \chi' \\ \chi_j, \chi'_j \notin B_{\text{embed}}}} \widehat{f}(\chi) \overline{\widehat{f}(\chi')} \left(\prod_{i \neq j} 1_{\chi_i = \chi'_i} \right) \mathbb{E}_{a,b} \left[(\chi_j(b) - \chi_j(a)) \overline{(\chi'_j(b) - \chi'_j(a))} \mid \sigma(a) = \sigma(b) \right] \end{aligned}$$

We claim that if $\chi_j \neq \chi'_j$, then the expectation is 0. Indeed, expanding we get terms such as $\chi_j(b) \overline{\chi'_j(a)}$, and we have that their expectation can be written as

$$\sum_{h \in H} \Pr_{a \sim \mu_x} [\sigma(a) = h] \mathbb{E}_b [\chi_j(b) \mid \sigma(b) = h] \mathbb{E}_a [\overline{\chi'_j(a)} \mid \sigma(a) = h],$$

which is 0 as $\mathbb{E}_b [\chi_j(b) \mid \sigma(b) = h]$ is proportional to the inner product between χ_j (a function orthogonal to embedding functions) and $1_{\sigma(b)=h}$ (an embedding function). Other terms are $\chi_j(b) \overline{\chi'_j(b)}$, which give the inner product between χ_j and χ'_j which is 0.

We conclude that to give non-zero contribution we must have that $\chi_j = \chi'_j$ and so the non-embedding influence of j is equal to

$$\sum_{\chi: \chi_j \notin B_{\text{embed}}} \left| \widehat{f}(\chi) \right|^2 \mathbb{E}_{a,b} \left[\left| \chi_j(b) - \chi_j(a) \right|^2 \mid \sigma(a) = \sigma(b) \right].$$

Expanding the square and repeating the above computation, we get that the expectation is 2, and the claim is proved.

The second bullet follows immediately by summing up the first bullet over all $j = 1, \dots, n$. \square

4.5 A Non-Embedding Stability Formulation of Theorem 3.18

With these notions, we can now state the result asserting that the only functions f, g, h for which expectations as in Theorem 3.18 may be large, are functions for which the non-embedding stability is significant. This result by itself is very much in the spirit of Theorem 3.18, but morally it is strictly weaker. Later, in Section 10, we will show how Theorem 3.18 is implied by Theorem 4.13 below. Another difference in the formulation is that below the dependency between the parameters is more explicit, and we do so as it is necessary for our proof to go through.

Theorem 4.13. *For all $m \in \mathbb{N}$, $\alpha > 0$ there are $M \in \mathbb{N}$, $\delta_0 > 0$ and $\eta > 0$ such that the following holds for all $0 < \delta \leq \delta_0$. Suppose that μ is a distribution over $\Sigma \times \Gamma \times \Phi$ satisfying:*

1. *The probability of each atom is at least α .*
2. *The size of each one of Σ, Γ, Φ is at most m .*
3. *$\text{supp}(\mu)$ is pairwise connected.*
4. *There are master embeddings σ, γ, ϕ for μ into an Abelian group $(H, +)$ that are saturated, and the distribution of $(\sigma(x), \gamma(y), \phi(z))$ where $(x, y, z) \sim \mu$ has full support on $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$.*

Then, if $f: \Sigma^n \rightarrow \mathbb{C}$, $g: \Gamma^n \rightarrow \mathbb{C}$ and $h: \Phi^n \rightarrow \mathbb{C}$ are 1-bounded functions such that $\text{NEStab}_{1-\delta}(g; \mu_y^{\otimes n}) \leq \delta$, then

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

5 Maximality, Merging Symbols and Some Motivating Examples

We do not know how to prove Theorem 4.13 directly, and our argument instead proceeds by further reducing this statement to a similar looking statement in which the distribution μ has additional useful properties. In this section we present two important ideas/ properties that are crucial in this reduction, which are called “merging symbols” and “maximality”. We then give a few examples of arguments that could be carried out using these notions, often making additional assumptions on μ (which we are not going to have in our formal argument in Sections 6, 7). We do this so as to demonstrate typical scenarios in which these ideas are useful while avoiding gory technicalities. As such, the language in this section will be informal at times, and we often appeal to intuition instead of making precise arguments.

Once we have explained these concepts, we will turn our attention into discussing the so-called “base case” of Theorem 4.13. By that, we mean a specialized statement in the setting of Theorem 4.13 for univariate functions, in which one manages to prove that the expectation is consideration is significantly smaller than 1 (for functions whose 2-norm is at most 1). We will discuss the “ideal base case”, which is a hypothetical scenario that we are not actually able to ensure; nevertheless, if such scenario were to hold, our argument would greatly simplify, and intuitively Theorem 4.13 would follow from the base case form some tensorization argument. Once again, our focus in this section will be to explain how “merges” and “maximality” facilitate such arguments.

Finally, after exploring the “ideal base case” scenario we will explain the issue that may arise, which we refer to as the “Horn-SAT” obstruction. We will explain the high level idea of how this issue is dealt with via what we call the “relaxed base case” and the intuition to why this relaxed base case should suffice for

the purpose of proving Theorem 4.13. Once again, our focus here will be in explaining how the concept of “merges” and “maximality” fit together with the relaxed base case.

5.1 Merging Symbols

The first operation we discuss is the merge operation. Suppose that we have a distribution μ as in Theorem 4.13 in which there are distinct symbols x and x' such that there are common $y \in \Gamma$ and $z \in \Phi$ for which (x, y, z) and (x', y, z) are both in the support of μ . Intuitively, this means that in coordinates wherein the g function gets y and the h functions gets z , some non-trivial averaging of the f function still occurs (as both x and x' are still possible). Naturally, averaging a function decreases its 2-norm, and we expect there to be a constant fraction of the coordinates in which even after fixing y and z there is still uncertainty whether x or x' occur in the corresponding coordinate in the function f . It follows that if the value of the function f “heavily distinguishes” between the symbols x and x' (in the sense that its value changes drastically if we change some coordinates in which x occurs to be x') then the expectation in Theorem 4.13 is small based solely on the fact that some non-trivial averaging occurs over f .

Following this line of reasoning leads one to speculate that one may assume that the function f does not distinguish between the two symbols x and x' (as otherwise the statement is trivial), in which case one may as well treat them as the same symbol. The goal of the merge operation is to precisely capture this idea, and we formally present it below.

Definition 5.1. *Let Σ, Γ, Φ be finite alphabets and let $P \subseteq \Sigma \times \Gamma \times \Phi$. We say that y, z imply x in P if for all $y \in \Gamma, z \in \Phi$ there is at most a single $x \in \Sigma$ such that $(x, y, z) \in P$. We say that y, z imply x in a distribution μ over $\Sigma \times \Gamma \times \Phi$ if y, z imply x in $\text{supp}(\mu)$.*

If the value of any two coordinates implies the third, we say a distribution μ is fully merged:

Definition 5.2. *Let Σ, Γ, Φ be finite alphabets and let μ be a distribution over $\Sigma \times \Gamma \times \Phi$. We say that μ is fully merged if (1) y, z imply x , (2) x, z imply y , and (3) x, y imply z .*

Fix finite alphabets Σ, Γ, Φ and a distribution μ over $\Sigma \times \Gamma \times \Phi$. To define the merge operation more precisely consider the graph $G_{x, \mu, \text{merge}}$ over Σ , wherein x and x' are adjacent if there are common $y \in \Gamma$ and $z \in \Phi$ such that (x, y, z) and (x', y, z) are both in $\text{supp}(\mu)$. Note that $G_{x, \mu, \text{merge}}$ has $|\Sigma|$ connected components if and only if y, z imply x (in which case the merge operation will do nothing), and by the above logic one should think of each connected component of $G_{x, \mu, \text{merge}}$ as a single symbol. More formally, given μ we may define the distribution μ' as follows:

Definition 5.3. *Let Σ, Γ, Φ be finite alphabets, let μ be a distribution over $\Sigma \times \Gamma \times \Phi$ in which y, z does not imply x , and let $G_{x, \mu, \text{merge}}$ be the graph above. We define the x -merged distribution μ' over $\Sigma' \times \Gamma \times \Phi$ where $\Sigma' \subsetneq \Sigma$ as: let the connected components of $G_{x, \mu, \text{merge}}$ be C_1, \dots, C_ℓ and choose a distinguished element x_i^* from each connected component. To sample according to μ' , we sample $(x, y, z) \sim \mu$, take i to be the connected components of x , and then output (x_i^*, y, z) .*

In words, the merge distribution should be thought of as combining each connected component of $G_{x, \mu, \text{merge}}$ into a single alphabet symbol.

Remark 5.4. *A few remarks are in order.*

1. *In Definition 5.3 we have defined the x -merge operation, and we will also use the y -merge operation and z -merge operation that are defined in an analogous way.*

2. The distributions μ and μ' are very closely related, and it is easy to observe that if μ satisfies all of the conditions of Theorem 4.13, then μ' also satisfies all of the conditions of that theorem. We omit the straightforward proof, but remark that a master embedding of μ is translated to a master embedding of μ' in the obvious way (restriction).
3. The merge operation has already made an appearance in [4, 5], however in that context the notions of degree/ noise stability are different. Here, our notions are more intricate and thus working with the merge operation requires a bit more care.

5.1.1 Simplifying Distributions via Merges

The following two lemmas allow us to reduce the proof of Theorem 4.13 for a distribution μ , to the proof of the same statement for a simpler distribution μ' which is a merge of μ . The first lemma handles x -merges, and analogously z -merges (as x and z are symmetric in the statement of Theorem 4.13).

Lemma 5.5. *Let Σ, Γ and Φ be finite alphabets, and let μ be a distribution over $\Sigma \times \Gamma \times \Phi$, and let μ' be the x -merged distribution coming from μ . If the conclusion of Theorem 4.13 holds for μ' , then it also holds for μ .*

Proof. Deferred to Section A.1. □

The second lemma handles y -merges; the proof is similar to the proof of Lemma 5.5 but some additional care is needed.

Lemma 5.6. *Let Σ, Γ and Φ be finite alphabets, and let μ be a distribution over $\Sigma \times \Gamma \times \Phi$, and let μ' be the y -merged distribution coming from μ . If the conclusion of Theorem 4.13 holds for μ' , then it also holds for μ .*

Proof. Deferred to Section A.2. □

Lemmas 5.5, 5.6 will be important for us later on, when we reduce the statement of Theorem 4.13 to a similar statement in which the distribution μ is fully merged. We note that for symbols $x, x' \in \Sigma$ that are mapped to different group elements via the master embedding, that is, $\sigma(x) \neq \sigma(x')$, we could never identify x and x' by a merge. The reason is that if this was possible, then there would be distinct tuples (x, y, z) and (x', y, z) in the support of μ , so that by the definition of the master embeddings we have

$$\sigma(x) + \gamma(y) + \phi(z) = 0, \quad \sigma(x') + \gamma(y) + \phi(z) = 0,$$

and it follows that it must be the case that $\sigma(x) = \sigma(x')$. Thus, merges will never decrease the alphabet sizes below $|H|$, and at some point it will be important for us to consider how large are the alphabet sizes in comparison to the size of H .

5.1.2 What are Merges Good For?

As explained, the idea of merges will be crucial for us and to illustrate it we consider a special case that nevertheless illustrates some important concepts. Besides the fact that if x and x' are merge-able then $\sigma(x) = \sigma(x')$, we are not aware of any other clear obstructions to merges. Thus, a special case that one may consider is what happens when the merge operation successfully reduces some of the alphabet sizes all the way down to $|H|$.

Suppose that after merging, the alphabet Σ and Φ both have size exactly $|H|$ so that each one of them could be identified with the Abelian group H . In that case, the partial basis we have set for x -functions and z -functions is not partial but rather a full basis, and one expects that multiplying $f(x)h(z)$ one would get an embedding function over y . As the non-embedding stability of g is small, it has small mass on monomials which are embedding functions, and thus it should be the case that the correlation of $g(y)$ and $f(x)h(z)$ is small. Indeed, arguments along these lines can be made – and are indeed crucial in Section 6. To illustrate such ideas, below we show an argument along these lines under several additional assumptions on the distribution μ .

Lemma 5.7. *Suppose that μ is a distribution as in Theorem 4.13, and further suppose that $|\Sigma| = |\Phi| = H$ and that the distribution of $(\sigma(x), \gamma(y), \phi(z))$ is uniform over $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$, where $(x, y, z) \sim \mu$. Then the conclusion of Theorem 4.13 holds.*

Proof. Relabeling Σ and Φ , we may assume that they are both equal to H and that the master embeddings σ and ϕ are the identity. Define $G(y) = \mathbb{E}_{(x', y', z') \sim \mu^{\otimes n}} [f(x')h(z') \mid y' = y]$, and expand

$$f(x) = \sum_{\chi \in \hat{H}^{\otimes n}} \hat{f}(\chi) \chi(x), \quad h(z) = \sum_{\chi' \in \hat{H}^{\otimes n}} \hat{h}(\chi') \chi'(z).$$

Then

$$G(y) = \sum_{\chi, \chi' \in \hat{H}^{\otimes n}} \hat{f}(\chi) \hat{h}(\chi') \mathbb{E}_{(x', y', z') \sim \mu^{\otimes n}} [\chi(x') \chi'(z') \mid y' = y].$$

Note that conditioned on y , the distribution over x' and z' is uniform such that $x' + \gamma(y) + z' = 0$. Thus, the expectation is 0 if $\chi \neq \chi'$ and otherwise is equal to $\chi(x') \chi'(z') = \chi(x' + z') = \chi(-\gamma(y))$. Thus,

$$G(y) = \sum_{\chi \in \hat{H}^{\otimes n}} \hat{f}(\chi) \hat{h}(\chi) \chi(-\gamma(y)).$$

It follows that G is an embedding function, that is, it is from $\text{Embed}_\gamma(\mu^{\otimes n})$, and hence it cannot have large correlation with g . More precisely, the left hand side of Theorem 4.13 is equal to

$$\langle G, \bar{g} \rangle = \langle T_{\text{non-embed}, \sqrt{1-\delta}} G, \bar{g} \rangle = \langle G, T_{\text{non-embed}, \sqrt{1-\delta}}^* \bar{g} \rangle = \langle G, T_{\text{non-embed}, \sqrt{1-\delta}} \bar{g} \rangle.$$

Using Cauchy-Schwarz, this is at most $\|G\|_2 \|T_{\text{non-embed}, \sqrt{1-\delta}} \bar{g}\|_2 \leq \text{NEStab}_{1-\delta}(\bar{g})^{1/2} \leq \sqrt{\delta}$. \square

Lemma 5.7 should be thought of as saying that if we managed to reduce the alphabet sizes of x and z all the way to their minimal point, which is the size of H , then we managed to prove Theorem 4.13. There are several ways to think about this; one way is as a sort of base case of an inductive statement, in which we are trying to prove Theorem 4.13 by induction on several parameters, one of which are the alphabet sizes of x and z (or rather their sum).

Our presentation below will be somewhat different (but morally equivalent), and we will use the idea of Lemma 5.7 to be able to assume that the size of the alphabet of x exceeds H , so that (by the pigeonhole principle) there are two distinct x symbols that are mapped to the same group element in H . Instead of directly trying to argue inductively based on this parameter, we will use these distinct elements, and the fact that μ is fully merged, to relate expectations as in Theorem 4.13 to expectations over different distributions μ' that have richer supports. In other words, in some scenarios we will be able to use Lemma 5.7 to enrich the support of the distribution μ . Intuitively, this marks significant progress since if we were able to enrich the support of μ indefinitely, we would eventually reach a distribution rich enough so that the statement of Theorem 4.13 becomes obvious. To capture this idea (and avoiding explicit iterative arguments of this nature), in the next section we define the notion of maximality of a distribution.

5.2 Maximality of Distributions

The next concept we describe is maximality of distributions. There are several equivalent ways of thinking about it, and ultimately it is just a form of induction. However, one good way to think about it is by assuming that there is a distribution μ which serves as a counter example for Theorem 4.13, and then trying to come up with a distribution ν which is a “maximal” counter example for the statement, in the sense that it is still a counter example to the statement but adding any additional atoms to the support of ν would lead to a distribution for which the statement is true. More formally:

Definition 5.8. *Let Σ , Γ and Φ , and let μ be a distribution over $\Sigma \times \Gamma \times \Phi$. We say that μ is maximal if for all distributions ν over $\Sigma \times \Gamma \times \Phi$ such that $\text{supp}(\mu) \subsetneq \text{supp}(\nu)$, the statement in Theorem 4.13 holds.*

The notions of maximality and merges often work in synergy together. Merges allow us to argue that certain atoms do not exist in μ . For example, if we know about distinct atoms (x, y, z) and (x', y', z') that are both in the support of μ where $x \neq x'$, and we know that in μ the value of y, z implies x (because we already applied some merges to guarantee that), then we automatically can conclude that (x', y, z) is not in the support of μ . Thus, whenever we are facing similar looking expectation to Theorem 4.13 which involves a distribution ν that contains the support of μ and additionally (x', y, z) , we can appeal to the maximality of μ and upper bound it as in Theorem 4.13.

5.2.1 What is Maximality Good For?

Below, we give a concrete example of such synergy, and once again the argument below is informal in nature.⁴ In Theorem 4.13 the assumption that the non-embedding stability of g is small really amounts to saying that g has almost all of its mass on monomials whose non-embedding degree is at least $\Theta(1/\delta)$. Intuitively, one expects that then the part of f of non-embedding degree significantly less – say $O(1/\delta^{0.99})$, would not contribute much to the expectation. Ignoring the h function for a moment, this would be clear as then there is simply a mis-match of degrees and one could appeal to the orthogonality of non-embedding functions to embedding functions. Now taking the h function into account once again, there doesn't seem to be a way for it to compensate for the large mis-match of the non-embedding degrees coming from the f function and those coming from the g function.

Indeed, in the following lemma we show that such expectations must indeed be small, and to do so we appeal to the notion of maximality; once again, we make some additional simplifying assumptions.

Lemma 5.9. *Suppose that μ is a distribution as in Theorem 4.13, and further suppose that μ is maximal and $|\Sigma| > |H|$. Then for all $\xi > 0$ there is $\eta = \eta(m, \alpha, \xi) > 0$ such that*

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} \left[\mathbb{T}_{\text{non-embed}, 1-\delta^{1-\xi}}^{\otimes n} f(x)g(y)h(z) \right] \right| \leq M\delta^\eta.$$

Proof. We may re-interpret the expectation in consideration as $\mathbb{E}_{(x,y,z) \sim \mu'^{\otimes n}} [f(x)g(y)h(z)]$ wherein the distribution μ' is defined by taking $(x, y, z) \sim \mu$, then $x' \sim \mathbb{T}_{\text{non-embed}, 1-\delta^{1-\xi}} x$, $y' = y$ and $z' = z$ and outputting (x', y', z') . The distribution μ' has small atoms with probability $\rho = \alpha\delta^{1-\xi}$, and to circumvent that we use random restrictions.

We may write $\mu' = \frac{\rho}{2}\mu'' + (1 - \frac{\rho}{2})\mu'''$ wherein μ'' , μ''' have the same support as μ' and in μ'' the probability of each atom is at least $\alpha' = \alpha'(\alpha) > 0$, and now use random restrictions to switch from

⁴Similar instantiations of this synergy appear numerous times in Sections 6, 7, and we refer the reader to there for a more formal presentation.

the probability measure μ' to the probability measure μ'' . Sample $J \subseteq_{\rho/2} [n]$ and choose the value of coordinates in \bar{J} according to $\mu''^{\bar{J}}$. That is, sample $(\tilde{x}, \tilde{y}, \tilde{z}) \sim \mu''^{\bar{J}}$ so that the expectation in hand can be seen as an averaging over these restrictions of a similar looking expectation over μ'' :

$$\mathbb{E}_{(x,y,z) \sim \mu'^{\otimes n}} [f(x)g(y)h(z)] = \mathbb{E}_{\substack{J \subseteq_{\rho/2} [n] \\ (\tilde{x}, \tilde{y}, \tilde{z}) \sim \mu''^{\bar{J}}}} \left[\mathbb{E}_{(x',y',z') \sim \mu''^J} [f'(x')g'(y')h'(z')] \right].$$

As most of the mass of g lies on monomials of non-embedding degree at least $\Theta(1/\delta)$, one expects most mass of g' to lie on non-embedding degree at least $\Theta(\rho/\delta) = \Theta(\delta^\xi)$; this is indeed true and can be argued by appealing to Claim 4.9, but we do not elaborate on it for now. Hence the inner expectation is an expectation of the same type as in Theorem 4.13, except that it is over the distribution μ'' .

We are now going to appeal to the maximality of μ to argue that this inner expectation must be small, and for that we must argue that the support of μ'' strictly contains the support of μ . By definition, the support of μ is contained in the support of μ' , and as the support of μ'' is the same as that of μ' , we get that $\text{supp}(\mu) \subseteq \text{supp}(\mu')$. Next, we argue that this is a strict containment. As $|\Sigma| > |H|$, by the pigeonhole principle there are distinct $a, a' \in \Sigma$ such that $\sigma(a) = \sigma(a')$, and we take $b' \in \Gamma$ and $c' \in \Phi$ such that $(a', b', c') \in \text{supp}(\mu)$. By definition of μ' it follows that (a, b', c') is in $\text{supp}(\mu')$, and hence it is in the support of μ'' . Note that $(a, b', c') \notin \text{supp}(\mu)$, as otherwise the symbols a and a' could be merged, but by assumption the distribution μ is already fully merged.

We may thus appeal to the maximality of μ and get the inner expectation is upper bounded by $M\delta^{\xi\eta}$ where $\eta > 0$ and M only depend on α and m , concluding the proof. \square

The main benefit of Lemma 5.9 is that it allows us to assume that not only the g function in Theorem 4.13 has large non-embedding degree, but rather also the f and h functions. Indeed, by Lemma 5.9 it follows that the expectation in Theorem 4.13 is very close to

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} \left[(I - \mathbb{T}_{\text{non-embed}, 1-\delta^{1-\xi}}^{\otimes n}) f(x)g(y)h(z) \right] \right|,$$

and now the new f function, which is $(I - \mathbb{T}_{\text{non-embed}, 1-\delta^{1-\xi}}^{\otimes n})f$, can be seen to have high non-embedding degree (while importantly also remaining bounded). As f and g now have high non-embedding degree, one could make a direct argument saying that the part of h of small non-embedding degree also has negligible contribution to the above expectation, and hence one is reduced to upper bounding an expectation of the form

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} \left[(I - \mathbb{T}_{\text{non-embed}, 1-\delta^{1-\xi}}^{\otimes n}) f(x)g(y)(I - \mathbb{T}_{\text{non-embed}, 1-\delta^{1-\xi}}^{\otimes n}) h(z) \right] \right|.$$

Proof by tensorization: motivating the base case. We will attempt to prove this statement by a tensorization argument, inducting on n . Ideally, we would have liked to argue that the worst-case high-degree non-embedding functions for the above expectation are simply product functions $\prod_{i=1}^{1/\delta} u_i(x_i)$, $\prod_{i=1}^{1/\delta} v_i(y_i)$ and

$\prod_{i=1}^{1/\delta} w_i(z_i)$, in which case the task of proving an upper bound as above reduces to a problem about univariate functions. There is one important distinction, however, which is that while the f, g and h functions are

bounded (in ℓ_∞ -norm), the functions u_i, v_i and w_i need not be bounded in ℓ_∞ , and instead we will only be able to guarantee ℓ_2 boundedness. Thus, while the “base case” of the tensorization statement seems much easier than the original statement, we have to prove it for a more general class of functions.

Below, we begin the discussion regarding this base case statement, and present the “ideal base case” scenario in which this logic is much simpler, however which we are not able to guarantee. Nevertheless, exploring the ideal base case scenario carries with it a lot of useful intuition, and in particular further relations between the utilization of a base case statement and the concepts of “merges” and “maximality”.

5.3 An Ideal Base Case Scenario in Theorem 4.13

With the notions of maximality and merging in hand, we can now present an ideal setting in which case the intuition behind the proof of Theorem 4.13 is relatively simple.

Consider a distribution μ as in Theorem 4.13, and suppose that it is fully merged as well as maximal. Further suppose that the marginal distribution $\mu_{y,z}$ is uniform over $\Gamma \times \Phi$. In this case, one notes that for $f: \Sigma \rightarrow \mathbb{C}$, $g: \Gamma \rightarrow \mathbb{C}$ and $h: \Phi \rightarrow \mathbb{C}$ of 2-norm equal to 1, one has that

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

where we used Cauchy-Schwarz and the fact that $\mu_{y,z}$ is uniform.

Inspecting equality cases for Cauchy-Schwarz, one notes that equality holds if and only if there is a constant $\theta \in \mathbb{C}$ with absolute value 1 such that $f(x) = \theta g(y)h(z)$, and for simplicity we assume that $\theta = 1$. Thus, equality holds if and only if the triplet f, g and h satisfy an embedding-like equation, except that it is not clear which group one should take. Taking the principle branch of the logarithm, one gets that $\log(f(x)) = \log(g(y)) + \log(h(z)) \pmod{2\pi\mathbf{i}}$, and now this is indeed an equation over an Abelian group (albeit infinite, but we already saw in Lemma 3.20 that one can convert such embeddings into finite Abelian group embeddings). Thus, it follows that the logs form an embedding of μ into an Abelian group, and by Claim 4.3 one may conclude that each one of them is an embedding function, hence $f \in \text{Embed}_\sigma(\mu)$, $g \in \text{Embed}_\gamma(\mu)$ and $h \in \text{Embed}_\phi(\mu)$.

In words, we have argued that if univariate functions achieve perfect value of $|\mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h(z)]|$, then they are embedding functions. This motivates the following statement, which we refer to as the “ideal base case”:

Statement 5.10. *Let μ be a distribution as above. Then for all $\tau > 0$ there is $\lambda > 0$ such that if $f: \Sigma \rightarrow \mathbb{C}$, $g: \Gamma \rightarrow \mathbb{C}$ and $h: \Phi \rightarrow \mathbb{C}$ are functions with 2-norm equal to 1 and $\|\text{Proj}_{\text{Embed}_\sigma(\mu)}(f)\|_2 \leq 1 - \tau$, then*

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

In words, Statement 5.10 asserts that if f is somewhat far from all embedding functions, then the value of $|\mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h(z)]|$ must be bounded away from 1. Indeed, in light of the above analysis this is something which is natural to expect; we examined the equality case, and by compactness type argument it follows that near equality cases can be characterized as “near embedding functions”.⁵

⁵We remark that in our actual argument we are going to need a decent quantitative dependency between the parameters τ and λ , typically a polynomial dependency. Thus, we will not be able to directly use compactness arguments and we will have to unravel them.

We take a moment to clarify that Statement 5.10 as stated is false in general. The issue in the above logic is that f , g and h may take the value 0 sometime, in which case we cannot apply the log function, and this turns out to be a rather serious obstacle referred to as the Horn-SAT obstruction. To resolve this issue we introduce the so-called *Relaxed Base Case*, which we give intuition to in Section 5.4 and which is the primary topic of Section 6.

Having said that, considering the class of distributions μ satisfying Statement 5.10 is helpful, and we now work under the assumption that it holds. In that case, given μ and f , g and h as in Theorem 4.13, we know that g has almost all of its ℓ_2 mass on monomials of non-embedding degree at least $1/\delta$. Using Lemma 5.9 we may also truncate the low-degree non-embedding degrees of both f and h , so that eventually we need to upper bound an expectation of the form (for simplicity of notation we ignored the $\xi > 0$ therein)

$$\left| \mathbb{E}_{(x,y,z) \sim \mu} \left[(I - T_{\text{non-embed}, 1-\delta}^{\otimes n}) f(x) g(y) (I - T_{\text{non-embed}, 1-\delta}^{\otimes n}) h(z) \right] \right|. \quad (6)$$

In conclusion, we are now reduced to working with the functions $f' = (I - T_{\text{non-embed}, 1-\delta}^{\otimes n}) f$, g and $h' = (I - T_{\text{non-embed}, 1-\delta}^{\otimes n}) h$, which all have almost all of their mass on monomials with non-embedding degree at least $1/\delta$.

As the base case gives some a gain of $1 - \lambda$ over the trivial bound when we have a univariate function with some non-embedding components, we expect to make this gain $1/\delta$ times, once for each non-embedding component in f' , g and h' ; in total, this would yield a bound of $(1 - \lambda)^{1/\delta}$, which is satisfactory for us (and even much better than what we're shooting for).

If Statement 5.10 was true, this argument would not be too far from the truth, and in fact can be made rigorous to work. Alas, as we said it could be the case that there are non-embedding functions f , g and h such that $f(x) = g(y)h(z)$, but then it is necessarily the case that the function f must vanish somewhere; we refer to such illegitimate-looking embeddings as Horn-SAT embeddings, and to the existence of which as the Horn-SAT obstruction.

To bypass the Horn-SAT obstruction we must study the possible vanishing patterns of the function f . We do not know how to argue about this for the distribution μ itself, and hence we have to once again move to a closely related distribution μ' (which is obtained from μ from a combination of more path tricks and merges), in which we are able to assert non-trivial information about the 0-sets of Horn-SAT embeddings. We defer the precise description of this reduction to Section 6, and in the next section we give some high level overview of the relaxed base case we are able to guarantee, how to work with it and the way that it fits in together with the notions of “merges” and “maximality”.

5.4 On the Relaxed Base Case Scenario in Theorem 4.13

5.4.1 A Naive Relaxed Base Case

Let μ be a distribution over $\Sigma \times \Gamma \times \Phi$ as above, and assume that $|\Sigma| > |H|$.⁶ After a suitable transformation of the distribution μ into a distribution $\tilde{\mu}$ over $\tilde{\Sigma} \times \tilde{\Gamma} \times \tilde{\Phi}$, we are (morally) able to make the following guarantee:

Statement 5.11. *There exists $\Sigma' \subseteq \tilde{\Sigma}$ of size larger than $|H|$, such that if $f: \tilde{\Sigma} \rightarrow \mathbb{C}$, $g: \tilde{\Gamma} \rightarrow \mathbb{C}$ and $h: \tilde{\Phi} \rightarrow \mathbb{C}$ are functions such that $f(x) = g(y)h(z)$ in the support of $\tilde{\mu}$, then $f|_{\Sigma'} \equiv 0$.*

⁶We remark that to justifying this assumption is precisely where results in the spirit of Lemma 5.7 come in handy.

In words, Statement 5.11 tells us that in $\tilde{\mu}$, all Horn-SAT embeddings must be 0 on Σ' . This motivates to attempt to formulate an analog of Statement 5.10 that instead of assuming non-trivial projection outside the subspace of embedding functions, assumes some variance of Σ' . And indeed, such statement can be proved to be true, but as is it is not very useful for us.

To be more specific, if we took that route and tried to write down an analog of (6), we would have to define a notion of degree that corresponds to not being 0 on Σ' and attempt to reduce ourselves to the case where this new notion of degree for the function f is large (so that we will be able to assert that we are avoiding the Horn-SAT obstruction on many coordinates, hence gaining some $1 - \lambda$ factor). At that point it is important though to keep the function f bounded, and hence to execute this logic we would need to define an averaging operator corresponding to the Markov chain that mixes inside the set Σ' and stays put on elements in $\Sigma \setminus \Sigma'$; this is so that functions that are not constant on Σ' would have their 2-norm decreased as a result of applying this averaging operator, so as to truncate of the part of f that has low degree with respect to the new notion.

This averaging operator however is incompatible with non-embedding degrees and the non-embedding averaging operator. The reason is that symbols in Σ' may be mapped to different group elements in H , in which case embedding functions also have variance on Σ' and hence get their 2-norm decreased by this averaging operator. We have no hope of gaining any $1 - \lambda$ factor from embedding functions, meaning that while identifying a property of Horn-SAT embeddings that we can ensure not to happen, we would re-introduce embedding functions into the mix and thus still not have a proper base case to induct on. There are other manifestations of this issue down the line if one pursues this direction, but ultimately they all boil down to the fact that the above averaging operator is not necessary a “sub-averaging operator” of our non-embedding operator $T_{\text{non-embed}, 1-\delta}$. By that, we mean that there are functions which the proposed averaging operator contracts, whereas $T_{\text{non-embed}, 1-\delta}$ keeps in place.

5.4.2 The Relaxed Base Case and Effective Non-embedding degrees

To resolve this issue, we take a subset of Σ' on which the master embedding is constant. As $|\Sigma'| > |H|$, by the pigeonhole principle there are distinct $x, x' \in \Sigma'$ such that $\sigma(x) = \sigma(x')$, so that we can take $\Sigma_{\text{modest}} = \{x, x'\} \subseteq \Sigma'$. We then have the following relaxed form of our ideal base case from above:

Statement 5.12. *For all $\tau > 0$ there is $\lambda > 0$ such that if $f: \tilde{\Sigma} \rightarrow \mathbb{C}$, $g: \tilde{\Gamma} \rightarrow \mathbb{C}$ and $h: \tilde{\Phi} \rightarrow \mathbb{C}$ are functions with 2-norm equal to 1 such that $\mathbb{E}_{x, x' \in \Sigma_{\text{modest}}} [|f(x) - f(x')|^2] \geq \tau$, then*

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

In words, for univariate functions, if our function f has a little bit of variance on Σ_{modest} , then we immediately get a gain of $1 - \lambda$ over the trivial bound. We remark that the property of having variance on Σ_{modest} immediately prohibits f from being an embedding function (as any embedding function is constant on Σ_{modest}), as well as from being part of a Horn-SAT embedding (as any Horn-SAT embedding must vanish on Σ' and hence on $\Sigma_{\text{modest}} \subseteq \Sigma'$).

This motivates defining a certain notion of degree, which we refer to as “effective non-embedding degree”; we often abbreviate this and just say effective degree instead. The effective non-embedding degree of a monomial over x is the number of coordinates on which the corresponding component has variance over Σ_{modest} , and intuitively this measures the number of times we will gain a factor of $1 - \lambda$ by appealing to

Statement 5.12. To make this definition more precise, we have to refine the basis we constructed consisting of embedding functions and non-embedding functions, and set-up an orthonormal basis of $L_2(\tilde{\Sigma}, \tilde{\mu}_x)$ composed on: (1) embedding functions, (2) non-embedding functions that are constant on Σ_{modest} , and (3) non-embedding functions that have variance on Σ_{modest} . Then, the effective (non-embedding) degree of a monomial is the number of components in it of functions not constant on Σ_{modest} .

5.4.3 Working with the Relaxed Base Case

Taking inspiration from the above discussion, one is tempted to argue that just like in (6) we managed to argue that the non-embedding degree of f can be assumed to be large, we should also manage to assume that the effective degree of f is large. This is indeed possible, and to do so we identify a proper Markov chain that captures effective degree; we refer to this Markov chain as “the modest Markov chain”, and it is defined as follows. On $a \in \tilde{\Sigma} \setminus \Sigma_{\text{modest}}$ the chain stays in place, and on $a \in \Sigma_{\text{modest}}$ the chain re-samples a symbol from Σ_{modest} according to the marginal distribution of $\tilde{\mu}_x$ on Σ_{modest} .

With the modest Markov chain in hand we can define a corresponding averaging operator, $E_{\text{non-embed}, 1-\delta}$ from $L_2(\tilde{\Sigma}, \tilde{\mu}_x)$ to $L_2(\tilde{\Sigma}, \tilde{\mu}_x)$ defined as $E_{\text{non-embed}, 1-\delta} f(x) = \mathbb{E}_{x' \sim E_{\text{non-embed}, 1-\delta} x} [f(x')]$, where in $x' \sim E_{\text{non-embed}, 1-\delta} x$ we take $x' = x$ with probability $1 - \delta$, and otherwise we sample x' according to the modest Markov chain on x . This Markov chain can be shown to precisely capture the notion of effective degrees, and hence our task now is to justify that we can assume that f has high effective degree, in the sense that we can reduce the task of proving that $|\mathbb{E}_{(x,y,z) \sim \tilde{\mu}^{\otimes n}} [f(x)g(y)h(z)]|$ is small to an analog of (6) of the form:

$$\left| \mathbb{E}_{(x,y,z) \sim \mu} \left[(I - E_{\text{non-embed}, 1-\delta}^{\otimes n}) f(x) g(y) (I - T_{\text{non-embed}, 1-\delta}^{\otimes n}) h(z) \right] \right|, \quad (7)$$

is small. Towards this end we must argue that the contribution from the part of f of small effective degree is small, and this is once again where maximality and merges come into play.

Lemma 5.13. *Suppose that $\tilde{\mu}$ is a distribution as above, and further suppose that $\tilde{\mu}$ is maximal and fully merged. Then for all $\xi > 0$ there is $\eta = \eta(m, \alpha, \xi) > 0$ such that*

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} \left[E_{\text{non-embed}, 1-\delta^{1-\xi}}^{\otimes n} f(x) g(y) h(z) \right] \right| \leq M \delta^\eta.$$

Proof. The proof is almost identical to the proof of Lemma 5.9, and we only sketch it. Re-interpreting this expectation as $\mathbb{E}_{(x,y,z) \sim \mu'^{\otimes n}} [f(x)g(y)h(z)]$ where the distribution μ' is the distribution in which we first sample $(x', y', z') \sim \tilde{\mu}$, then $x \sim E_{\text{non-embed}, 1-\delta^{1-\xi}}^{\otimes n} x'$, take $y = y'$, $z = z'$ and output (x, y, z) . As in Lemma 5.9, the support of μ' strictly contains the support of $\tilde{\mu}$ and hence it makes sense to try to appeal to the maximality of $\tilde{\mu}$. The only issue is that in $\tilde{\mu}$ there are atoms with small probability $\rho = \alpha \delta^{1-\xi}$, and to bypass that we write $\mu' = \frac{\rho}{2} \mu'' + (1 - \frac{\rho}{2}) \mu'''$ where μ'' and μ''' have the same supports as μ' and in μ'' the probability of each atom is at least $\alpha'(\alpha) > 0$, and then use random restrictions as in Lemma 5.9. \square

5.5 Some Additional Remarks on Combining These Ingredients

Throughout this section we have proved some useful lemmas regarding the interaction of expectations as in Theorem 4.13 and the notions of merges, maximality and how the relaxed base case fits in. Ideally, we would have liked to have a distribution μ that possess all of the properties that we needed (on top of the ones

assumed in Theorem 4.13): (1) the alphabet of x has size larger than $|H|$, (2) $\mu_{y,z}$ is uniform, (3) μ is fully merged, (4) μ admits a relaxed base case statement as in Statement 5.12, (5) μ is maximal.

We are not going to be able to ensure that all of these properties simultaneously occur for μ . Instead, we will argue that the distribution μ “contains within it” some other distribution ν (possibly on different alphabets) on which some of these properties hold. More specifically, we are not going to be able to guarantee that μ is maximal, and instead we will be able to argue that “within it” there is a maximal distribution. By that, we mean that there are Σ' and Φ' , a distribution ν over $\Sigma' \times \Gamma \times \Phi'$ which is maximal and maps $a: \Sigma' \rightarrow \Sigma$ and $c: \Phi' \rightarrow \Phi$ such that:

1. **Containment:** $\{(a(x), y, c(z)) \mid (x, y, z) \in \text{supp}(\nu)\} \subseteq \text{supp}(\mu)$;
2. **Alignment of Master Embeddings:** taking (σ, γ, ϕ) to be a master embedding of μ , we have that $(\sigma \circ a, \gamma, \phi \circ c)$ is a master embedding of ν .

Intuitively, the reason that this is useful is that, after suitable random restrictions, we can relate expectations with respect to μ to expectations with respect to ν . Indeed, letting μ' be the condition distribution of μ on $\{(a(x), y, c(z)) \mid (x, y, z) \in \text{supp}(\nu)\}$ and writing $\mu = \beta\mu' + (1 - \beta)\mu''$ for some distribution μ'' and $\beta = \beta(\alpha) > 0$, we can switch from the distribution μ to the distribution μ' (as in the proof of Lemma 5.9), and functions over this domain can be lifted to functions over the domain of ν :

$$f'(x') = f(a(x'_1), \dots, a(x'_n)), \quad g'(y') = g(y') \quad h'(z') = h(c(z'_1), \dots, c(z'_n)).$$

We thus managed to reduce the problem of bounding some expectation with respect to μ to the task of bounding some expectation with respect to ν . Often times, this line of reasoning (on top of arguments as above) will allow us to appeal to the maximality of ν (and get a result which is qualitatively the same as if we could assume that μ itself is maximal).

6 Arranging for a Base Case for Theorem 4.13

In this section, we begin the proof of Theorem 4.13. As explained earlier, the core of our argument will ultimately be by induction on n , and as such we are going to need a base case statement for functions over a single variable. Our inductive process though will be unable to preserve 1-boundedness and will only be able to give us L_2 -bound guarantees. Therefore, the base case we are looking for has to address functions with a bounded L_2 -norm.

The most naive attempts at arriving at such base case lead one to a difficulty referred to as the “Horn-SAT” obstruction, which refers to the possibility of a existence of a triplet of functions f, g and h that satisfy $f(x) = g(y)h(z)$ on the support on μ but that do not necessarily yield an Abelian embedding; this difficulty arises due to the fact that f may be 0 on some inputs.

In this section, our goal is to state a result that implies Theorem 4.13, and which is more amenable to a proof by induction along the lines of [5]. To do so, we will first have to go through some reductions and present intermediate statements which imply Theorem 4.13; roughly speaking, these statements will all be similar to the statement of Theorem 4.13 with additional assumptions on the distribution μ .

6.1 Further Preprocessing of the Distribution μ : Pushing Counter-examples to the Extreme

At a high level, the goal of our preprocessing step is to arrive at a distribution μ' (which may be different from μ) such that if Theorem 4.13 is false for μ , then it is also false for μ' , but moreover μ' is the “richest”

distribution on which the statement remains false. That is, if we consider any distribution ν satisfying the conditions of Theorem 4.13 whose support strictly contains the support of μ' , then the conclusion of that theorem is satisfied for ν . To be more precise, assuming the statement is false for a distribution μ , the distribution μ' will be a result of applying a sequence of the following two operations, so long as it is possible.

1. **Merging Symbols:** if there are distinct symbols $x, x' \in \Sigma$ for which there are $y \in \Gamma$ and $z \in \Phi$ such that (x, y, z) and (x', y, z) are both in $\text{supp}(\mu)$, then we can define a distribution μ' over $\Sigma' \times \Gamma \times \Phi$ where $\Sigma' \subsetneq \Sigma$, such that Theorem 4.13 holds for μ if and only if it holds for μ' . This is done via Lemma 5.5, and by symmetry of the roles of x and z the same goes for symbols in Φ .
2. **Enlarging the support:** looking at μ for which the statement is false, we ask ourselves whether there are additional atoms from $\Sigma \times \Gamma \times \Phi$ that can be inserted to $\text{supp}(\mu)$ so that the statement remains false. If so, we pass from μ to another distribution over $\Sigma \times \Gamma \times \Phi$ for which the statement is still false and whose support strictly contains the support of μ .

Repeating the above steps so long that it is possible (noting that it eventually terminates as we are either reducing the alphabet sizes or enlarging the size of $\text{supp}(\mu)$ at each step), we reach a distribution μ on which Theorem 4.13 is false and is extremal in these regards. More precisely:

Lemma 6.1. *For all $m \in \mathbb{N}$ there is $m' \in \mathbb{N}$ such that the following holds. Suppose that Σ, Γ and Φ are alphabets of size at most m and μ is a distribution over $\Sigma \times \Gamma \times \Phi$ for which Theorem 4.13 fails. Then there exist alphabets Σ', Γ' and Φ' of size at most m' and a distribution μ' over $\Sigma' \times \Gamma' \times \Phi'$ such that*

1. *The distribution μ' is fully merged as per Definition 5.2.*
2. *The distribution μ' is maximal as per Definition 5.8.*
3. *The conclusion of Theorem 4.13 fails for μ' .*

Proof. If μ is not merged, we perform a merge operation and use either Lemma 5.5 or Lemma 5.6, noting that all of the conditions of Theorem 4.13 continue to hold for the new distribution. If μ is not maximal, we keep greedily form a sequence of distributions by adding elements from $\Sigma \times \Gamma \times \Phi$ to the support of the distribution so long as Theorem 4.13 fails for it, and eventually reach a maximal distribution. We then iterate. Note that this process terminates, as each invocation of merge decreases the size of the alphabet Σ by at least 1.

A delicate point to notice, is that if we add an element to μ to form a distribution ν , then all of the conditions of Theorem 4.13 continue to hold. Indeed, if σ, γ, ϕ is an Abelian embedding of ν , then it is also an Abelian embedding of μ hence it is a coordinate of the master embedding of μ . Thus we can keep a subset of coordinates of the master embeddings of μ (which correspond to embeddings of ν) and have that these form a master embedding of ν ; for simplicity say $\sigma_{\text{master}, \mu} = (\sigma_1, \dots, \sigma_r)$ and $\sigma_{\text{master}, \nu} = (\sigma_1, \dots, \sigma_{r'})$ where $r' < r$ and $H = H_1 \times \dots \times H_r$.

Note that $\sigma_{\text{master}, \nu}$ is saturated on $H_{\leq r'} = H_1 \times \dots \times H_{r'}$. We define $\gamma_{\text{master}, \nu}$ and $\phi_{\text{master}, \nu}$ similarly and notice that they are also saturated on $H_{\leq r'}$. The distribution of $(\sigma_{\text{master}, \nu}(x), \gamma_{\text{master}, \nu}(y), \phi_{\text{master}, \nu}(z))$ where $(x, y, z) \sim \nu$ has full support on $\{(a, b, c) \in H_{\leq r'}^3 \mid a + b + c = 0\}$. \square

Remark 6.2. *Another important feature of a distribution ν that results from a distribution μ by adding elements to its support, is that the non-embedding noise operator of ν is “weaker” than that of μ . Intuitively, this follows as the master embedding of μ is a refinement of the master embedding of ν , hence the operator*

of ν does less averaging. We will use this fact in the future (formalized appropriately) to argue that if a function g has small noise stability with respect to the non-embedding noise operator of μ , then it also has small stability with respect to the non-embedding noise operator of ν .

In Section 4.1 we have seen that if each one of the alphabets Σ, Φ is equal to H , then the proof Theorem 4.13 is a rather simple, requiring a change of distribution and a basic Fourier analytic computation. As explained in Section 5 some of our arguments require either the size of Σ or of Φ to be strictly larger than the size of the group H , and thus we have to separately deal with the case that both have the same size as H . This is a slightly more general case than the case handled in Section 4.1, and in the following lemma we show that argument to the argument therein works:

Lemma 6.3. *Suppose that μ is a distribution as in Theorem 4.13, that $\text{NEStab}_{1-\delta}(g) \leq \delta$ and that $|\Sigma| = |\Phi| = |H|$. Then the conclusion of Theorem 4.13 holds.*

Proof. We first argue that it suffices to prove Theorem 4.13 in our case under the additional assumption that the distribution of $(\sigma_{\text{master}}(x), \gamma_{\text{master}}(y), \phi_{\text{master}}(z))$ where $(x, y, z) \sim \mu$ is uniform over $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$, and then show a direct argument for that case.

Reduction to the case the master embeddings are uniform. Consider the distribution ν over the support of μ which is defined by first sampling (a, b, c) uniformly from $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$ and then sampling $(x, y, z) \sim \mu$ conditioned on $\sigma_{\text{master}}(x) = a, \gamma_{\text{master}}(y) = b, \phi_{\text{master}}(z) = c$, and note that we may write $\mu = \alpha'\nu + (1 - \alpha')\nu'$ for some $\alpha' = \alpha/2$ where ν' is some distribution. We now think of generating $(x, y, z) \sim \mu^{\otimes n}$ as first choosing $I \subseteq_{\alpha'} [n]$, sampling $(x', y', z') \sim \nu^I, (x'', y'', z'') \sim \nu'^{\bar{I}}$ and taking $x = (x', x''), y = (y', y''), z = (z', z'')$. Then

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f(x)g(y)h(z)] \right| \leq \mathbb{E}_{I, (x'', y'', z'')} \left[\left| \mathbb{E}_{(x', y', z') \sim \nu^I} [f_{\bar{I} \rightarrow x''}(x')g_{\bar{I} \rightarrow y''}(y')h_{\bar{I} \rightarrow z''}(z')] \right| \right] \quad (8)$$

Note that in ν, ν' the probability of each atom is at least $\alpha/2$, and we choose $c(\alpha/2, \alpha')$ from Lemma 2.8. Let $c' = \alpha c/2$ and let E be the event that $\text{NEStab}_{1-c'-1\delta}(g_{\bar{I} \rightarrow y''}; \nu_y) \geq \sqrt{\delta}$. Then by Lemma 2.8 we have that

$$\mathbb{E}_{I, y''} \left[\text{NEStab}_{1-c'-1\delta}(g_{\bar{I} \rightarrow y''}; \nu_y^I) \right] \leq \text{NEStab}_{1-\delta}(g; \mu_y^{\otimes n}) \leq \delta,$$

so by Markov's inequality it follows that $\Pr[E_1] \leq \sqrt{\delta}$. Thus, looking at (8) we get that

$$\begin{aligned} \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f(x)g(y)h(z)] \right| &\leq \mathbb{E}_{I, (x'', y'', z'')} \left[\mathbb{1}_{E_1} \left| \mathbb{E}_{(x', y', z') \sim \nu^I} [f_{\bar{I} \rightarrow x''}(x')g_{\bar{I} \rightarrow y''}(y')h_{\bar{I} \rightarrow z''}(z')] \right| \right] \\ &+ \mathbb{E}_{I, (x'', y'', z'')} \left[\mathbb{1}_{\bar{E}_1} \left| \mathbb{E}_{(x', y', z') \sim \nu^I} [f_{\bar{I} \rightarrow x''}(x')g_{\bar{I} \rightarrow y''}(y')h_{\bar{I} \rightarrow z''}(z')] \right| \right] \\ &\leq \sqrt{\delta} + \mathbb{E}_{I, (x'', y'', z'')} \left[\mathbb{1}_{\bar{E}_1} \left| \mathbb{E}_{(x', y', z') \sim \nu^I} [f_{\bar{I} \rightarrow x''}(x')g_{\bar{I} \rightarrow y''}(y')h_{\bar{I} \rightarrow z''}(z')] \right| \right]. \end{aligned}$$

Whenever \bar{E}_1 holds, we may take $\delta' = \sqrt{\delta}$ and get that $\text{NEStab}_{1-\delta'}(g_{\bar{I} \rightarrow y''}; \nu_y) \leq \delta'$ (provided that δ is small enough). Moreover, in ν we now have the additional property that the master embeddings are uniform, hence the last expectation is at most $M\delta'^\eta \leq M\delta^{\eta/2}$ where η, M depend only on α and the alphabet sizes, hence we get that $|\mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f(x)g(y)h(z)]| \leq M'\delta^{\eta'}$ for $M' = 2M$ and $\eta' = \eta/2$.

The argument in the case the master embeddings are uniform. We now assume that the distribution of $(\sigma_{\text{master}}(x), \gamma_{\text{master}}(y), \phi_{\text{master}}(z))$ where $(x, y, z) \sim \mu$ is uniform over $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$. As $|\Sigma| = |\Phi| = |H|$, we may re-label the symbols in Σ to be group elements (and similarly Φ), so that the master embeddings of x (and similarly z) will be the identity. Define $\tilde{g}(y) = \mathbb{E}_{(x', y', z') \sim \mu} [f(x')h(z') \mid y' = y]$ and note that \tilde{g} is in $\text{Embed}_\gamma(\mu)$. Indeed, write $f = \sum_{\chi \in \hat{H}^{\otimes n}} \hat{f}(\chi)\chi(x)$ and $h = \sum_{\chi' \in \hat{H}^{\otimes n}} \hat{h}(\chi')\chi'(z)$. Multiplying out, one observes that for χ, χ' we have

$$\mathbb{E}_{(x', y', z') \sim \mu} [\chi(x')\chi'(z') \mid y' = y] = \chi(-\gamma(y))\mathbb{E}_{(x', y', z') \sim \mu} [(\chi'\bar{\chi})(z') \mid y' = y] = \chi(-\gamma(y))1_{\chi'=\bar{\chi}},$$

where we used the fact that conditioned on y' , the distribution over (x', z') is uniform over $x' + \gamma(y') + z' = 0$. Thus,

$$\left| \mathbb{E}_{(x, y, z) \sim \mu^{\otimes n}} [f(x)g(y)h(z)] \right| = \langle g, \bar{\tilde{g}} \rangle = \langle g, T_{\sqrt{1-\delta}, \text{non-embed}} \bar{\tilde{g}} \rangle = \langle T_{\sqrt{1-\delta}, \text{non-embed}} g, \bar{\tilde{g}} \rangle,$$

and by Cauchy-Schwarz this is at most $\|T_{\sqrt{1-\delta}, \text{non-embed}} g\|_2 = \sqrt{\text{NEStab}_{1-\delta}(g)} \leq \sqrt{\delta}$. \square

With Lemmas 6.3, 6.1 in hand, it quickly follows that it suffices to prove Theorem 4.13 in the case that non of these results apply, and hence it suffices to prove the following variant of Theorem 4.13:

Theorem 6.4. *For all $m \in \mathbb{N}$, $\alpha > 0$ there are $M \in \mathbb{N}$, $\delta_0 > 0$ and $\eta > 0$ such that the following holds for all $0 < \delta \leq \delta_0$. Suppose that μ is a distribution over $\Sigma \times \Gamma \times \Phi$ satisfying:*

1. *The probability of each atom is at least α .*
2. *The size of each one of Σ, Γ, Φ is at most m .*
3. *$\text{supp}(\mu)$ is pairwise connected.*
4. *There are master embeddings σ, γ, ϕ for μ into an Abelian group $(H, +)$ that are saturated, and the distribution of $(\sigma(x), \gamma(y), \phi(z))$ where $(x, y, z) \sim \mu$ has full support on $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$.*
5. *The distribution μ is maximal as per Definition 5.8.*
6. *$|\Sigma| > |H|$ or $|\Phi| > |H|$.*
7. *In μ , the value of y, z implies the value of x . Also, the value of x, z implies y .*

Then, if $f: \Sigma^n \rightarrow \mathbb{C}$, $g: \Gamma^n \rightarrow \mathbb{C}$ and $h: \Phi^n \rightarrow \mathbb{C}$ are 1-bounded functions such that $\text{NEStab}_{1-\delta}(g; \mu_y^{\otimes n}) \leq \delta$ then

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

6.1.1 Theorem 6.4 implies Theorem 4.13

By Lemma 6.1, Theorem 4.13 is equivalent to the same statement under the additional assumptions of maximality and that the distribution is fully merged. Letting μ be a distribution over $\Sigma \times \Gamma \times \Phi$ satisfying these properties, if $|\Sigma| = |\Phi| = |H|$ then the validity of the statement follows from Lemma 6.3. Otherwise, we also have that either $|\Sigma| > |H|$ or $|\Phi| > |H|$ and the validity follows from Theorem 6.4. \square

6.2 Restating Theorem 6.4 via Role Symmetry of x, z

Noting that the roles of x and z are symmetric in Theorem 6.4 we have that it is equivalent to the following statement, in which the same condition has been replaced by $|\Sigma| > |H|$.

Theorem 6.5. *For all $m \in \mathbb{N}$, $\alpha > 0$ there are $M \in \mathbb{N}$, $\delta_0 > 0$ and $\eta > 0$ such that the following holds for all $0 < \delta \leq \delta_0$. Suppose that μ is a distribution over $\Sigma \times \Gamma \times \Phi$ satisfying:*

1. *The probability of each atom is at least α .*
2. *The size of each one of Σ, Γ, Φ is at most m .*
3. *$\text{supp}(\mu)$ is pairwise connected.*
4. *There are master embeddings σ, γ, ϕ for μ into an Abelian group $(H, +)$ that are saturated, and the distribution of $(\sigma(x), \gamma(y), \phi(z))$ where $(x, y, z) \sim \mu$ has full support on $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$.*
5. *The distribution μ is maximal as per Definition 5.8.*
6. *$|\Sigma| > |H|$.*
7. *In μ , the value of any two coordinates implies the value of the third.*

Then, if $f: \Sigma^n \rightarrow \mathbb{C}$, $g: \Gamma^n \rightarrow \mathbb{C}$ and $h: \Phi^n \rightarrow \mathbb{C}$ are 1-bounded functions such that $\text{NEStab}_{1-\delta}(g; \mu_y^{\otimes n}) \leq \delta$ then

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

With Theorem 6.5 in hand and following the logic of Section 5, we should now strive to achieve a relaxed base case for distributions as in Theorem 6.5. This is the content of the next section.

6.3 Establishing a Relaxed Base Case

With the statement of Theorem 6.5 in hand, we are now ready to begin addressing the so-called Horn-SAT obstruction. The bulk of our proof will eventually be an inductive proof for a statement similar in spirit to Theorem 6.5, and to facilitate this induction we must have a base case. The goal of this section is to design such base case.

We begin by describing an ideal scenario, also discussed in Section 5, in which a simple base case statement holds for the distribution μ ; if this could be achieved for all distributions μ as in Theorem 6.5 our argument would simplify considerably. Alas, there are distributions μ for which this ideal base case fails, and we call such obstructions as Horn-SAT embeddings. After explaining what Horn-SAT embeddings are, we will turn to the question of how to overcome the Horn-SAT obstruction. To do so we will switch from the distribution μ to a related distribution μ' (using the path trick and merges), for which we have a relaxed form of this ideal base case (and is sufficient for our purposes).

6.3.1 The Ideal Base Case

Ideally, we would have liked to have a base case statement as follows. For all $\tau > 0$, there is $\lambda < 1$ such that if $f: \Sigma \rightarrow \mathbb{C}$, $g: \Gamma \rightarrow \mathbb{C}$ and $h: \Phi \rightarrow \mathbb{C}$ satisfy that $\text{NEStab}_{1-\xi}(g; \mu_y) \leq (1 - \tau)\|g\|_2^2$, then

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

There are several issues with this statement.

1. First, at the present scenario it is not even clear that the statement should be true for $\lambda = 1$, let alone $\lambda < 1$. The reason is that trivially, we can only bound the left hand side using Hölder's inequality by $\|f\|_3 \|g\|_3 \|h\|_3$ (or some other product of q_1, q_2, q_3 norms where $\frac{1}{q_1} + \frac{1}{q_2} + \frac{1}{q_3} = 1$), which may be much larger than $\|f\|_2 \|g\|_2 \|h\|_2$. To address this issue, we will ensure that in $(x, y, z) \sim \mu$ we have that y and z are independent, in which case we may now use Cauchy-Schwarz to always argue that

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

so we can at least take $\lambda = 1$ in the ideal base case.

2. Second, even with the above transformation, it turns out that one still may not be able to take $\lambda < 1$. Naively, one is tempted argue that it is possible to take some $\lambda < 1$ by considering the equality cases of Cauchy-Schwarz and argue as follows: if we cannot take $\lambda < 1$, then it means that there are functions f, g, h of 2-norm 1 for which the above Cauchy-Schwarz is tight, hence $f(x) = g(y)h(z)$ in the support of μ , and such f, g, h constitute an Abelian embedding, so $g \in \text{Embed}_\gamma(\mu)$ and in particular $\text{NEStab}_{1-\xi}(g; \mu_y) = \|g\|_2^2$.

Taking a closer inspection though reveals that there are issues if f sometimes gets the value 0; indeed, if f is always non-zero this argument goes through, but we do not know how to ensure that. This issue was referred to as the ‘‘Horn-SAT Obstruction’’ in [5], and here too we have to circumvent it.

6.3.2 The Relaxed Base Case

As the ideal base case may fail, we resort to a more relaxed form of it, referred to as the ‘‘relaxed base case’’; below is a formal definition.

Definition 6.6. *Let Σ, Γ and Φ be finite alphabets of size at most m , let μ be a distribution over $\Sigma \times \Gamma \times \Phi$ in which the probability of each atom is at least α , and let (σ, γ, ϕ) be master embeddings of μ into an Abelian group $(H, +)$. We say that μ satisfies the relaxed base case if there is $\Sigma_{\text{modest}} \subseteq \Sigma$ of size at least 2 such that the following holds:*

1. *For all $x, x' \in \Sigma_{\text{modest}}$ it holds that $\sigma(x) = \sigma(x')$.*
2. *There are $c(\alpha, m), C(\alpha, m) > 0$ such that the following holds. For all $\tau > 0$, for all $f: \Sigma \rightarrow \mathbb{C}$, $g: \Gamma \rightarrow \mathbb{C}$ and $h: \Phi \rightarrow \mathbb{C}$ such that $\mathbb{E}_{x, x' \in \Sigma_{\text{modest}}} [|f(x) - f(x')|^2] \geq \tau$ it holds that*

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

We now state a variant of Theorem 6.4 for distributions that have the relaxed base case, and then show it implies Theorem 6.4. In comparison to Theorem 6.5, below we have the additional assumption that the distribution of $(\sigma(x), \gamma(y), \phi(z))$ is uniform over elements $(a, b, c) \in H^3$ such that $a + b + c = 0$, that $\mu_{y,z}$ is uniform and that the relaxed base case holds.

Theorem 6.7. *For all $m \in \mathbb{N}$, $\alpha > 0$ there are $K, M \in \mathbb{N}$, $\delta_0 > 0$ and $\eta > 0$ such that the following holds for all $0 < \delta \leq \delta_0$. Suppose that ν is a distribution over $\Sigma \times \Gamma \times \Phi$ satisfying:*

1. *The probability of each atom is at least α .*
2. *The size of each one of Σ, Γ, Φ is at most m .*
3. *$\text{supp}(\nu)$ is pairwise connected.*
4. *There are master embeddings σ, γ, ϕ for μ into an Abelian group $(H, +)$ that are saturated, and the distribution of $(\sigma(x), \gamma(y), \phi(z))$ where $(x, y, z) \sim \nu$ is uniform over $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$.*
5. *The support of ν on $\Gamma \times \Phi$ is full, and the marginal distribution of ν on y, z is uniform.*
6. *There are $\Sigma_{\text{modest}} \subseteq \Sigma' \subseteq \Sigma$, $\Gamma' \subseteq \Gamma$ and $\Phi' \subseteq \Phi$ such that*
 - (a) **Relaxed base case:** *the distribution ν satisfies the relaxed base case with Σ_{modest} .*
 - (b) **Containing a maximal, fully merged distribution:** *There is an alphabet Σ'' of size at most K , a map $a: \Sigma'' \rightarrow \Sigma'$ and a distribution $\tilde{\nu}$ supported on $\Sigma'' \times \Gamma' \times \Phi'$ such that:*
 - i. *a is surjective.*
 - ii. *$\{(a(x), y, z) \mid (x, y, z) \in \text{supp}(\tilde{\nu})\} \subseteq \text{supp}(\nu)$.*
 - iii. *$\tilde{\nu}$ is maximal as per Definition 5.8.*
 - iv. *In $\tilde{\nu}$, the value of any two coordinates implies the third.*
 - v. *$(\sigma \circ a, \gamma, \phi)$ is a master embedding of $\tilde{\nu}$, and it is saturated.*
 - vi. *There are distinct $v, v' \in \Gamma'$ such that $\gamma(v) = \gamma(v')$.*
 - (c) **Full support on restriction of the first two coordinates:** *$\Sigma' \times \Gamma \subseteq \text{supp}(\nu_{x,y})$. In words, for all $x \in \Sigma'$ and $y \in \Gamma$ there is $z \in \Phi$ such that $(x, y, z) \in \text{supp}(\nu)$.*
7. *In ν , the value of y, z implies the value of x .*

Then, if $f: \Sigma^n \rightarrow \mathbb{C}$, $g: \Gamma^n \rightarrow \mathbb{C}$ and $h: \Phi^n \rightarrow \mathbb{C}$ are 1-bounded functions such that $\text{NEStab}_{1-\delta}(g; \nu_y^{\otimes n}) \leq \delta$ then

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

At first reading, we encourage the reader to think of condition 6b above as saying that μ itself is maximal and fully merged. We do not know how to ensure that however, at least without losing some other properties of μ which are necessary for us (or without introducing further complications). Instead, as explained in Section 5.5, we just say that within μ we can find a distribution ν which satisfies these additional properties; for our purposes, this is just as good as the distribution μ itself having these properties.

6.4 Achieving the Relaxed Base Case: Theorem 6.7 implies Theorem 6.5

In this section we show that Theorem 6.7 implies Theorem 6.5. For that, we start with a distribution μ as in Theorem 6.5 and construct from it a distribution ν as in Theorem 6.7, such that correlations with respect to μ can be upper bounded by a similar looking correlations over ν . We begin by noting that we may assume that in the notation of Theorem 6.5, there are distinct $v, v' \in \Gamma$ such that $\gamma(v) = \gamma(v')$. Indeed, otherwise the statement holds vacuously, as then $\|g\|_2^2 = \text{NEStab}_{1-\delta}(g; \mu_y^{\otimes n}) \leq \delta$ and the expectation would clearly be bounded by above by $\|g\|_2 \leq \sqrt{\delta}$.

6.4.1 The Construction of $\nu, \Sigma_{\text{modest}}, \Sigma'$ and Φ'

Take ℓ_1 large enough with respect to m , and then ℓ_2 large enough with respect to ℓ_1, m .

1. Take μ' to be the path trick distribution with respect to z on μ for length ℓ_1 , and choose ℓ_1 large enough so that the support of μ' on x, y is full (using Lemma 3.13). Then μ' is a distribution over $\Sigma \times \Gamma \times \tilde{\Phi}$ where $\tilde{\Phi} \subseteq \Phi^{\ell_1}$.
2. Take μ'' to be the path trick distribution with respect to x on μ' for length ℓ_2 , and choose ℓ_2 to be large enough so that μ'' has full support on y, z (again, using Lemma 3.13). Then μ'' is a distribution over $\tilde{\Sigma} \times \Gamma \times \tilde{\Phi}$ where $\tilde{\Sigma} \subseteq \Sigma^{\ell_2}$.
3. Take ν' to be the x -merge of μ'' , so that it is a distribution over $\tilde{\Sigma}_{\text{fin}} \times \Gamma \times \tilde{\Phi}$ where $\tilde{\Sigma}_{\text{fin}} \subseteq \tilde{\Sigma}$. For a symbol $x \in \tilde{\Sigma}$, we denote by $a(x) \in \tilde{\Sigma}_{\text{fin}}$ the distinguished element from the connected component of x as per Definition 5.3.
4. Duplicate symbols: let σ, γ and ϕ be a saturated master embedding for ν' into an Abelian group H . For each $a \in H$, let $d_3(a)$ be the number of $z \in \tilde{\Phi}$ such that $\phi(z) = a$ (and note that $d_3(a) \geq 1$ for all $a \in H$ as ϕ is saturated). Similarly, let $d_2(a)$ be the number of $y \in \Gamma$ such that $\gamma(y) = a$.

We are going to duplicate each symbol $y \in \Gamma$, and $z \in \tilde{\Phi}$ symbol multiple times. That is, define

$$\Gamma_{\text{dup}} = \left\{ (y, i) \mid y \in \Gamma, 1 \leq i \leq \frac{\prod_{t \in H} d_2(t)}{d_2(\gamma(y))} \right\}, \quad \tilde{\Phi}_{\text{dup}} = \left\{ (z, j) \mid z \in \tilde{\Phi}, 1 \leq j \leq \frac{\prod_{t \in H} d_3(t)}{d_3(\phi(z))} \right\}.$$

5. Take ν to be the following distribution: sample $(b, c) \in H^2$ uniformly, sample $y \in \Gamma$ and $z \in \tilde{\Phi}$ uniformly such that $\gamma(y) = b$ and $\phi(z) = c$, sample $1 \leq i \leq \frac{\prod_{t \in H} d_2(t)}{d_2(b)}$ and $1 \leq j \leq \frac{\prod_{t \in H} d_3(t)}{d_3(c)}$ uniformly and then sample $(x', y', z') \sim \nu'$ conditioned on $y' = y$ and $z' = z$. Output $(x', (y, i), (z, j))$.

We take $\Sigma' = \{a(x, \dots, x) \mid x \in \Sigma\}$, $\Gamma' = \{(y, 1) \mid y \in \Gamma\}$ and $\Phi' = \{((z, \dots, z), 1) \mid z \in \tilde{\Phi}\}$. As $|\Sigma| > |H|$, by the pigeonhole principle there are distinct $x^*, x^{*'} \in \Sigma$ such that $\sigma(x^*) = \sigma(x^{*'})$, and we take

$$\Sigma_{\text{modest}} = \{a(x^*, \dots, x^*), a(x^{*'}, \dots, x^{*'})\}.$$

There are two cases, depending on the size of Σ_{modest} :

1. If $|\Sigma_{\text{modest}}| = 1$, we show that the validity of Theorem 6.4 for μ follows from its maximality. The idea here is that if Σ_{modest} has size 1, it means that the symbols x and x' get “mixed up” when we look at the distribution ν ; however ν contains inside it a copy of μ , hence we are able to relate correlations over ν to correlations over distribution μ of random restrictions, and this mix-up means that we will actually look at a distribution richer than μ .

2. If $|\Sigma_{\text{modest}}| = 2$, we show that the relaxed base case for ν holds, and then Theorem 6.7 quickly implies Theorem 6.5.

The two cases are addressed in the following two lemmas:

Lemma 6.8. *If $|\Sigma_{\text{modest}}| = 1$, then μ satisfies the conclusion of Theorem 6.5.*

Proof. Deferred to Section 6.6. □

Lemma 6.9. *If $|\Sigma_{\text{modest}}| > 1$, then the distribution ν satisfies the relaxed base case and, furthermore, with $\Sigma_{\text{modest}}, \Sigma'$ and Φ' defined as above, it satisfies the conditions in Theorem 6.7.*

Proof. We discuss the master embedding $(\sigma', \gamma', \phi')$ of ν :

1. Let (σ, γ, ϕ) be a master embedding of μ . In Lemma 3.12 we have seen how master embedding evolve under the path trick, hence we get a master embedding for μ'' ; abusing notation we denote it also by (σ, γ, ϕ) . Note that if two symbols \vec{x} and \vec{x}' are to be merged in ν' , then they are mapped to the same group element by σ . Indeed, if there are $y \in \Gamma$ and $z \in \tilde{\Phi}$ such that (\vec{x}, y, z) and (\vec{x}', y, z) are both in the support of μ'' , then

$$\sigma(\vec{x}) + \gamma(y) + \phi(z) = 0 = \sigma(\vec{x}') + \gamma(y) + \phi(z),$$

hence $\sigma(\vec{x}) = \sigma(\vec{x}')$. It follows that σ is constant on each connected component of $G_{x, \mu'', \text{merge}}$, hence we may define $\sigma(a(\vec{x})) = \sigma(\vec{x})$ unambiguously and get an embedding for ν' .

Lastly, a master embedding for ν follows. The map σ stays the same, and γ and ϕ simply ignore i and j , that is,

$$\gamma'(y, i) = \gamma(y), \quad \phi'(z, j) = \phi(z).$$

We note that (σ, γ', ϕ') is a master embedding of ν . To see that, first note there is a 1-to-1 correspondence between embedding of ν and embedding of ν' : indeed, note that for all (y, i) and (y, i') there are x and (z, j) such that $(x, (y, i), (z, j))$ and $(x, (y, i'), (z, j))$ are in the support of ν' it follows that any embedding of the second symbol must ignore i , thereby essentially be an embedding of ν' . Thus, any embedding of ν must ignore i and j , hence is also an embedding of ν' , and thereby must be equivalent to some coordinate of (σ, γ, ϕ) (as it is a master embedding of ν').

2. For any $b, c \in H$, sampling $(x, (y, i), (z, j)) \sim \nu$ we get that the probability that $\gamma'(y, i) = b$ and $\phi'(z, j) = c$ is $1/|H|^2$ by definition, and then $\sigma(x)$ must be $-b - c$. Hence, the distribution of $(\sigma(x), \gamma(y, i), \phi(z, j))$ is uniform over $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$.
3. We note that it follows that the master embedding of ν maps the symbols in Σ_{modest} to the same group elements. Indeed, in μ we have that $\sigma(x^*) = \sigma(x^{*'})$, and following the evolution of the master embedding above (using Lemmas 3.10, 3.12 again) it follows that the master embedding of ν maps $a(x^*, \dots, x^*)$ and $a(x^{*'}, \dots, x^{*'})$ to the same group element.
4. Finally, we note that as above we found distinct $v, v' \in \Gamma$ such that $\gamma(v) = \gamma(v')$, it follows that $(v, 1)$ and $(v', 1)$ are distinct elements in Γ' such that $\gamma'(v, 1) = \gamma'(v', 1)$.

Next, we note that the marginal distribution of ν on $\Gamma_{\text{dup}} \times \tilde{\Phi}_{\text{dup}}$ is uniform. Indeed, by definition it is clear this distribution is a product distribution, so it suffices to argue that its marginal on each one of Γ_{dup} and $\tilde{\Phi}_{\text{dup}}$ is uniform. To see that, note that for all $t = (t', i') \in \Gamma_{\text{dup}}$

$$\Pr_{(x, (y, i), (z, j)) \sim \nu} [(y, i) = t] = \frac{1}{|H|} \cdot \frac{1}{d_2(t')} \cdot \frac{d_2(t')}{\prod_{a \in H} d_2(a)} = \frac{1}{|H|} \cdot \frac{1}{\prod_{a \in H} d_2(a)},$$

wherein the first factor counts the probability we chose $b = \gamma(t')$ in the process, the second factor counts the probability we chose a specific pre-image y of t' , and the third factor counts the probability that $i = i'$. This probability is independent of t , hence the marginal of ν on Γ_{dup} is uniform. The same argument shows that the marginal of ν on $\tilde{\Phi}_{\text{dup}}$ is uniform.

We now argue that $\Sigma' \times \Gamma_{\text{dup}} \subseteq \text{supp}(\nu_{x,y})$. Indeed, by construction we have that $\Sigma \times \Gamma \subseteq \text{supp}(\mu'_{x,y})$ and hence $\{(x, \dots, x) \mid x \in \Sigma\} \times \Gamma \subseteq \text{supp}(\mu''_{x,y})$. Thus, after the x -merge we get that $\Sigma' \times \Gamma \subseteq \text{supp}(\nu'_{x,y})$, and after duplicating symbols we get that $\Sigma' \times \Gamma_{\text{dup}} \subseteq \text{supp}(\nu_{x,y})$.

We defer the proof of the rest of properties of ν to Section 6.5. \square

To finish this section, we quickly explain how to derive Theorem 6.5 from Theorem 6.7.

6.4.2 Proof that Theorem 6.7 implies Theorem 6.5

With the above set-up above, if $|\Sigma_{\text{modest}}| = 1$ then we are done by Lemma 6.8, so assume otherwise. Thus, $|\Sigma_{\text{modest}}| > 1$, and by Lemma 6.9 it follows that ν satisfies all of the conditions of Theorem 6.7. By Lemma 3.9 we may upper bound

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f(x)g(y)h(z)] \right|^\ell \leq \left| \mathbb{E}_{(x,y,z) \sim \mu''^{\otimes n}} [F(x)g(y)H(z)] \right|^\ell,$$

where F and H are some bounded functions, $\ell = O(1)$ and the marginal of μ''_y is the same as μ_y . Hence, it suffices to establish the conclusion of Theorem 6.5 for μ'' . Using Lemma 5.5, it suffices to prove the statement for the x -merge of μ'' , namely for ν' .

For the distribution ν , by Theorem 6.7 we know that it satisfies the conclusion of Theorem 4.13. Consider the distribution ν_{pre} over $\tilde{\Sigma}_{\text{final}} \times \Gamma \times \tilde{\Phi}$ defined as: sample $(x, (y, i), (z, j)) \sim \nu$ and output (x, y, z) . Given functions $f: \tilde{\Sigma}_{\text{final}}^n \rightarrow \mathbb{C}$, $g: \Gamma^n \rightarrow \mathbb{C}$ and $h: \tilde{\Phi}^n \rightarrow \mathbb{C}$ we may define $f': \tilde{\Sigma}_{\text{final}}^n \rightarrow \mathbb{C}$, $g': \Gamma_{\text{dup}}^n \rightarrow \mathbb{C}$ and $h': \tilde{\Phi}_{\text{dup}}^n \rightarrow \mathbb{C}$ by $f' = f$ and

$$g'((y_1, i_1), \dots, (y_n, i_n)) = g(y_1, \dots, y_n), \quad h'((z_1, j_1), \dots, (z_n, j_n)) = h(z_1, \dots, z_n).$$

Then

$$\mathbb{E}_{(x,y,z) \sim \nu_{\text{pre}}^{\otimes n}} [f(x)g(y)h(z)] = \mathbb{E}_{(x',y',z') \sim \nu^{\otimes n}} [f'(x)g'(y)h'(z)]$$

and the master embeddings of ν and ν_{pre} are essentially the same, so $\text{NEStab}_{1-\delta}(g') = \text{NEStab}_{1-\delta}(g)$. Thus, as ν satisfies the conclusion of Theorem 6.5, it follows that ν_{pre} also satisfies the conclusion of Theorem 6.5.

We now note that ν' and ν'_{pre} are two distributions over $\tilde{\Sigma}_{\text{final}} \times \Gamma \times \tilde{\Phi}$ with the same support and in which the probability of each atom is at least $\alpha' = \alpha'(\alpha, m) > 0$, hence the conclusion of Theorem 6.5 is equivalent for them. Indeed, we may write $\nu' = \beta \nu'_{\text{pre}} + (1 - \beta) \mathcal{D}$ where $\beta = \beta(\alpha, m) > 0$ and \mathcal{D} is some

distribution. Thus, given functions $f: \tilde{\Sigma}_{\text{final}}^n \rightarrow \mathbb{C}$, $g: \Gamma^n \rightarrow \mathbb{C}$ and $h: \tilde{\Phi}^n \rightarrow \mathbb{C}$ that are 1-bounded and $\text{NEStab}_{1-\delta}(g; \nu'^{\otimes n}) \leq \delta$ we may choose $J \subseteq_{\beta} [n]$, sample $(\tilde{x}, \tilde{y}, \tilde{z}) \sim \mathcal{D}^{\tilde{J}}$, set

$$\tilde{f} = f_{\tilde{J} \rightarrow \tilde{x}}, \quad \tilde{g} = g_{\tilde{J} \rightarrow \tilde{y}}, \quad \tilde{h} = h_{\tilde{J} \rightarrow \tilde{z}},$$

and get that

$$\mathbb{E}_{(x,y,z) \sim \nu'^{\otimes n}} [f(x)g(y)h(z)] = \mathbb{E}_{J, (\tilde{x}, \tilde{y}, \tilde{z}) \sim \mathcal{D}^{\tilde{J}}} \left[\mathbb{E}_{(x', y', z') \sim \nu'_{\text{pre}}^J} \left[\tilde{f}(x') \tilde{g}(y') \tilde{h}(z') \right] \right].$$

We denote the inner expectation by $\phi_{\nu'_{\text{pre}}}(\tilde{f}, \tilde{g}, \tilde{h})$, and let E be the event that $\text{NEStab}_{1-\sqrt{\delta}}(\tilde{g}; \nu'_{\text{pre}}^I) \leq \sqrt{\delta}$. By Claims 4.9, 4.8 it follows that

$$\mathbb{E}_{\substack{J \subseteq_{\beta} [n] \\ \tilde{y} \sim \mathcal{D}_{\tilde{y}}^{\tilde{J}}}} \left[\text{NEStab}_{1-\sqrt{\delta}}(\tilde{g}; \nu'_{\text{pre}}) \right] \leq \text{NEStab}_{1-c(m, \alpha)\beta\sqrt{\delta}}(g; \nu') \leq \text{NEStab}_{1-\delta}(g; \nu') \leq \delta.$$

Thus, by Markov's inequality we have that $\Pr[E] \leq \sqrt{\delta}$, and so

$$\begin{aligned} \left| \mathbb{E}_{J, (\tilde{x}, \tilde{y}, \tilde{z}) \sim \mathcal{D}^{\tilde{J}}} \left[\phi_{\nu'_{\text{pre}}}(\tilde{f}, \tilde{g}, \tilde{h}) \right] \right| &\leq \left| \mathbb{E}_{J, (\tilde{x}, \tilde{y}, \tilde{z}) \sim \mathcal{D}^{\tilde{J}}} \left[1_E \phi_{\nu'_{\text{pre}}}(\tilde{f}, \tilde{g}, \tilde{h}) \right] \right| + \left| \mathbb{E}_{J, (\tilde{x}, \tilde{y}, \tilde{z}) \sim \mathcal{D}^{\tilde{J}}} \left[1_{\overline{E}} \phi_{\nu'_{\text{pre}}}(\tilde{f}, \tilde{g}, \tilde{h}) \right] \right| \\ &\leq M\delta^{s/2} + \Pr[E] \leq (M+1)\delta^{s/2}, \end{aligned}$$

where we used the fact that if E holds, then we may bound $\left| \phi_{\nu'_{\text{pre}}}(\tilde{f}, \tilde{g}, \tilde{h}) \right| \leq M\delta^s$ for some $M, s > 0$ depending only on m and α , and also we used the upper bound on $\Pr[E]$. \square

6.5 The case that $|\Sigma_{\text{modest}}| = 2$: Proof of Lemma 6.9

6.5.1 On the Supports of Horn-SAT Embeddings

A triplet of functions $f: \tilde{\Sigma}_{\text{fin}} \rightarrow \mathbb{C}$, $g: \Gamma_{\text{dup}} \rightarrow \mathbb{C}$ and $h: \tilde{\Phi}_{\text{dup}} \rightarrow \mathbb{C}$ is called a Horn-SAT embedding if $f(\vec{x}) = g(y)h(\vec{z})$; we say it is non-trivial if $f \notin \text{Embed}_{\sigma}(\nu)$. Consider the collection of subsets

$$\mathcal{F} = \{F \mid \exists (f, g, h) \text{ a non-trivial Horn-SAT embedding of } \nu, F = \text{supp}(f)\}.$$

Let $F_{\text{union}} = \bigcup_{F' \in \mathcal{F}} F'$.

Claim 6.10. *For $x \in \Sigma$ we have that $a(x, \dots, x) \notin F_{\text{union}}$.*

Proof. Assume towards contradiction otherwise, and let x^* be such x . Then there is $F' \in \mathcal{F}$ containing $a(x^*, \dots, x^*)$, hence there is a non-trivial Horn-SAT embedding f, g and h such that $f(a(x^*, \dots, x^*)) \neq 0$.

Define $f': \Sigma \rightarrow \mathbb{C}$ by $f'(x) = f(a(x, \dots, x))$, and note that for all $(x, y, z) \in \text{supp}(\mu')$ we get that

$$f'(x) = g(y)h(\vec{z}),$$

and also that $f'(x^*) \neq 0$. In μ' the support of the distribution on $\Sigma \times \Gamma$ is full, and hence we conclude that $g(y) \neq 0$ for all y . Thus, for all $(x, y, \vec{z}) \in \text{supp}(\mu')$ we get that $f'(x) = 0$ if and only if $h(\vec{z}) = 0$, and

since μ' is pairwise connected it follows that the functions f' and h must either be trivially 0, or else never 0. As $f'(x^*) \neq 0$, we get that f' and h never vanish.

Thus, g and h never vanish, and as f, g, h is a Horn-SAT embedding it follows that f also never vanishes. Consider the principal branch of the complex logarithm function \log , and define $f''(x) = \log(f(x))$, $g''(y) = \log(g(y))$ and $h''(z) = \log(h(z))$; by Claim 4.3 we get that $f'' \in \text{Embed}_\sigma(\nu)$, and as \log is injective it follows that $f \in \text{Embed}_\sigma(\nu)$, in contradiction to the fact that (f, g, h) is a non-trivial Horn-SAT embedding. \square

We have the following immediate corollary, asserting that Horn-SAT embeddings must be constant on Σ_{modest} (thereby giving some sense that the relaxed base case holds for ν).

Corollary 6.11. *If $f: \tilde{\Sigma}_{\text{fin}} \rightarrow \mathbb{C}$, $g: \Gamma_{\text{dup}} \rightarrow \mathbb{C}$ and $h: \tilde{\Phi}_{\text{dup}} \rightarrow \mathbb{C}$ form an Horn-SAT embedding, then $f|_{\Sigma_{\text{modest}}}$ is constant.*

Proof. Suppose towards contradiction that $f|_{\Sigma_{\text{modest}}}$ is not constant. Note that the master embedding of ν assigns that same group element to both members of Σ_{modest} , hence we get that $f \notin \text{Embed}_\sigma(\nu)$. Thus, f, g, h is a non-trivial Horn-SAT embedding, hence by Claim 6.10 f must vanish on Σ' , and as $\Sigma_{\text{modest}} \subseteq \Sigma'$ it follows that f must vanish on Σ_{modest} , and contradiction. \square

6.5.2 Proof of Lemma 6.9: Compactness

We begin by establishing a weak form of the relaxed base case via a compactness argument. As such, this argument does not produce quantitative bounds (which are crucial for our application), however it serves as a good warm-up for the actual argument proving the relaxed base case (which is very similar in spirit).

Namely, we show that for all $\tau > 0$ there is $\tau' > 0$ such that if $f: \tilde{\Sigma}_{\text{final}} \rightarrow \mathbb{C}$, $g: \Gamma_{\text{dup}} \rightarrow \mathbb{C}$ and $h: \tilde{\Phi}_{\text{dup}} \rightarrow \mathbb{C}$ are functions such that $\mathbb{E}_{x, x' \in \Sigma_{\text{modest}}} [|f(x) - f(x')|^2] \geq \tau$, then

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

Indeed, otherwise we could find a sequence f_m, g_m, h_m with 2-norm equal to 1 such that the left hand side approaches 1, and by limiting we could find f, g, h as above, in which $|\mathbb{E}_{(\vec{x}, y, z) \sim \nu} [f(x)g(y)h(z)]| = 1$. By Cauchy-Schwarz we have

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

hence Cauchy-Schwarz is an equality and so $f(x) = \theta g(y)h(z)$ in the support of ν , where θ has absolute value equal to 1. Multiplying g by θ we get that $f(x) = g(y)h(z)$, hence this triplet forms an Horn-SAT embedding. By limiting we have that $\mathbb{E}_{x, x' \in \Sigma_{\text{modest}}} [|f(x) - f(x')|^2] \geq \tau$, and this contradicts Corollary 6.11.

6.5.3 Proof of Lemma 6.9: Unraveling Compactness

In this section we show that ν satisfies the relaxed base case. Fix f, g, h and $\tau > 0$ as in the definition of the relaxed base case; by normalizing, we may assume that the 2-norms of each one of f, g and h is equal to 1. We will assume that $\tau \leq c$ where $c = c(m, \alpha) > 0$ is a constant to be determined, as otherwise we may lower τ .

Assume towards contradiction that $|\mathbb{E}_{(x,y,z)\sim\nu} [f(x)g(y)h(z)]| \geq 1 - \tau^{50m}$. Thus there is a complex number θ of absolute value 1 such that $\theta\mathbb{E}_{(x,y,z)\sim\nu} [f(x)g(y)h(z)] \geq 1 - \tau^{50m}$, and to simplify notation we multiply f by θ so that this inequality becomes $\mathbb{E}_{(x,y,z)\sim\nu} [f(x)g(y)h(z)] \geq 1 - \tau^{50m}$; we note we may multiply f by a constant without loss of generality, as this does not affect any of the assumptions on f .

Claim 6.12. $\mathbb{E}_{x \in \Sigma_{\text{modest}}} [|f(x)|^2] \geq \frac{\tau}{4}$.

Proof. By $|a - b|^2 \leq 2|a|^2 + 2|b|^2$ for all $a, b \in \mathbb{C}$ we get

$$\mathbb{E}_{x \in \Sigma_{\text{modest}}} [|f(x)|^2] \geq \frac{1}{4} \mathbb{E}_{x, x' \in \Sigma'} [|f(x) - f(x')|^2] \geq \frac{\tau}{4}. \quad \square$$

Claim 6.13. For all $(x, y, z) \in \text{supp}(\nu)$ we have $|\overline{f(x)} - g(y)h(z)| \leq \tau^{20m}$.

Proof. We have

$$\begin{aligned} \mathbb{E}_{(x,y,z)\sim\nu} \left[|\overline{f(x)} - g(y)h(z)|^2 \right] &= 2 - \mathbb{E}_{(x,y,z)\sim\nu} [f(x)g(y)h(z)] - \mathbb{E}_{(x,y,z)\sim\nu} [\overline{f(x)g(y)h(z)}] \\ &= 2 - 2 \mathbb{E}_{(x,y,z)\sim\nu} [f(x)g(y)h(z)] \\ &\leq 2\tau^{50m}, \end{aligned}$$

and as the probability of each atom in ν is least $\alpha'(\alpha, m) > 0$ we get that $|\overline{f(x)} - g(y)h(z)|^2 \leq \alpha'^{-1}\tau^{50m} \leq \tau^{40m}$ for c sufficiently small, giving the claim. \square

Claim 6.14. For all $y \in \Gamma_{\text{dup}}$ we have $|g(y)| \gtrsim_{\alpha, m} \sqrt{\tau}$.

Proof. Assume this is not the case, so that there is $y^* \in \Gamma_{\text{dup}}$ such that $|g(y^*)| \leq \eta(\alpha, m)\sqrt{\tau}$ where $\eta(\alpha, m) > 0$ is a small function of α and m . By Claim 6.12 there is $\vec{x} = a(x, \dots, x) \in \Sigma_{\text{modest}}$ such that $|f(\vec{x})| \geq \sqrt{\tau}/2$ (where $x \in \Sigma$). As the support of μ' on $\Sigma \times \Gamma$ is full, there is $\vec{z} \in \tilde{\Phi}$ such that (x, y^*, z) is in $\text{supp}(\mu')$, hence (\vec{x}, y^*, \vec{z}) is in the support of ν , so by Claim 6.13

$$\sqrt{\tau}/2 \leq |f(\vec{x})| \leq |g(y^*)| |h(\vec{z}, 1)| + \tau^{20m} \leq \eta\sqrt{\tau}(\alpha')^{-1/2} + \tau^{20m}$$

where we used the fact that $|h(\vec{z})| \leq (\alpha')^{-1/2} \|h\|_2^2 = (\alpha')^{-1/2}$, where $\alpha' = \alpha'(\alpha, m) > 0$ is a lower bound on the probability of any atom in ν . This is a contradiction for a small enough η . \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 6.15. We have that $|f(\vec{x})| \geq \tau^{3j+1.5}$ for all $\vec{x} \in \Sigma_{\text{modest}}$ and $|h(\vec{z})| \geq \tau^{3j+1.5}$ for all $\vec{z} \in \tilde{\Phi}_{\text{dup}}$.

Proof. Assume otherwise, and define

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

Then by our assumption, at least one of A, B is not identically 1, say A without loss of generality. Note that since $\|f\|_2 = 1$, there is some x such that $|f(x)| \geq 1 > \tau$, so a is also not constantly 0, hence A is not constant. We next show that $A(x) + B(z) = 0 \pmod{2}$ for all $(x, y, z) \in \text{supp}(\nu)$, hence conclude that ν is not pairwise connected in contradiction to the fact that μ is pairwise connected and Lemma 3.13.

The case that $A(x) = 1$. Suppose that $(x, y, z) \in \text{supp}(\nu)$ are such that $A(x) = 1$. It follows from Claim 6.13 that

$$|g(y)h(z)| \geq |f(x)| - \tau^{10m} \geq \tau^{3j+1.5} - \tau^{20m},$$

and as $|g(y)| \leq \alpha^{-1} \|g\|_{2, \tilde{\nu}_y} = \alpha^{-1}$, it follows that

$$|h(z)| \geq \frac{\tau^{3j+1.5} - \tau^{20m}}{1/\alpha} \geq \tau^{3j+2},$$

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 $B(z) = 1$.

The case that $A(x) = 0$. Suppose that $(x, y, z) \in \text{supp}(\tilde{\nu})$ are such that $A(x) = 0$. It follows from Claim 6.13 that

$$|g(y)h(z)| \leq |f(x)| + \tau^{20m} \leq \tau^{3j+1.5} + \tau^{20m},$$

and as $|g(y)| \geq \tau$ by Claim 6.14, it follows that

$$|h(z)| \leq \frac{\tau^{3j+1.5} + \tau^{20m}}{\tau} \leq \tau^{3j+0.4},$$

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 $B(z) = 0$. \square

Claim 6.15 implies that the functions f, g, h never get the value 0, and also that in the support of ν ,

$$\left| \frac{\overline{f(x)}}{g(y)h(z)} - 1 \right| = \frac{|\overline{f(x)} - g(y)h(z)|}{|g(y)h(z)|} \leq \frac{\tau^{20m}}{\Omega_{\alpha,m}(\sqrt{\tau})\tau^{3j+1.5}} \leq \tau^{10m}.$$

This motivates the definition of an approximate Abelian embedding by taking the principal branch of the logarithm:

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

Define $d(a, b) = \min_{k \in \mathbb{Z}} |a - b - 2\pi i k|$ for $a, b \in \mathbb{C}$. Then in the support of ν we have:

$$d(f'(x), g'(y) + h'(z)) = d\left(\log\left(\frac{\overline{f(x)}}{g(y)h(z)}\right), 0\right) \leq \left|\frac{\overline{f(x)}}{g(y)h(z)} - 1\right| \leq \tau^{10m}. \quad (9)$$

We now use the logic of Claim 4.3 (and the underlying application of Dirichlet's Approximation Theorem) to slightly change f', g' and h' to get a proper Abelian embedding.

We work separately with the real and imaginary part of f', g', h' . Looking at $S = \text{Image}(\text{Re}(f')) \cup \text{Image}(\text{Re}(g')) \cup \text{Image}(\text{Re}(h'))$, we have that $|S| \leq 3m$. We take $N = (\alpha\tau)^{-9m}$; from Dirichlet's Approximation Theorem it follows that there are $\sigma_{\text{Re}}: \Sigma_{\text{fin}} \rightarrow \mathbb{Z}$, $\gamma_{\text{Re}}: \Gamma \rightarrow \mathbb{Z}$ and $\phi_{\text{Re}}: \Phi \rightarrow \mathbb{Z}$ such that for some integer $1 \leq q \leq N$ we have

$$\begin{aligned} \left| \text{Re}(f'(x)) - \frac{\sigma_{\text{Re}}(x)}{q} \right| &\leq \frac{1}{qN^{1/|S|}}, & \left| \text{Re}(g'(y)) - \frac{\gamma_{\text{Re}}(y)}{q} \right| &\leq \frac{1}{qN^{1/|S|}}, \\ \left| \text{Re}(h'(z)) - \frac{\phi_{\text{Re}}(z)}{q} \right| &\leq \frac{1}{qN^{1/|S|}}. \end{aligned}$$

Similarly we find $\sigma_{\text{Im}}: \Sigma_{\text{fin}} \rightarrow \mathbb{Z}$, $\gamma_{\text{Im}}: \Gamma \rightarrow \mathbb{Z}$ and $\phi_{\text{Im}}: \Phi \rightarrow \mathbb{Z}$ such that

$$\begin{aligned} \left| \text{Im}(f'(x)) - \frac{\sigma_{\text{Im}}(x)}{q} \right| &\leq \frac{1}{qN^{1/|S|}}, & \left| \text{Im}(g'(y)) - \frac{\gamma_{\text{Im}}(y)}{q} \right| &\leq \frac{1}{qN^{1/|S|}}, \\ \left| \text{Im}(h'(z)) - \frac{\phi_{\text{Im}}(z)}{q} \right| &\leq \frac{1}{qN^{1/|S|}}. \end{aligned}$$

Define

$$\tilde{\sigma}(x) = \sigma_{\text{Re}}(x) + 2\pi\mathbf{i}\sigma_{\text{Im}}(x), \quad \tilde{\gamma}(y) = \gamma_{\text{Re}}(y) + 2\pi\mathbf{i}\gamma_{\text{Im}}(y), \quad \tilde{\phi}(z) = \phi_{\text{Re}}(yz) + 2\pi\mathbf{i}\phi_{\text{Im}}(z).$$

Claim 6.16. $\tilde{\sigma}(x) - \tilde{\gamma}(y) - \tilde{\phi}(z) = 0 \pmod{2\pi\mathbf{i}}$ for all $(x, y, z) \in \text{supp}(\nu)$.

Proof. From the choice of $\tilde{\sigma}, \tilde{\gamma}, \tilde{\phi}$ and (9) it follows that for all $(x, y, z) \in \text{supp}(\nu)$ it holds that

$$d\left(\frac{\tilde{\sigma}(x)}{q}, \frac{\tilde{\gamma}(y) + \tilde{\phi}(z)}{q}\right) \leq \frac{6}{qN^{1/|S|}} + d(f'(x), g'(y) + h'(z)) \leq \frac{6}{qN^{1/|S|}} + \tau^{10m}$$

where we used the triangle inequality and (9). Multiplying by q we get that $d(\sigma(x), \gamma(y) + \phi(z)) \leq \frac{6}{N^{1/|S|}} + q\tau^{10m} < 1$, where we used the fact that $q \leq N$, the choice of N and the fact that τ is sufficiently small. As $\tilde{\sigma}(x) - \tilde{\gamma}(y) - \tilde{\phi}(z)$ is a complex number of the form $a + \mathbf{b}\mathbf{i}$ for $a, b \in \mathbb{Z}$, it must be the case that $a = 0$ and b is a multiple of 2π , and the claim follows. \square

Looking at the real parts of $\tilde{\sigma}, \tilde{\gamma}$ and $\tilde{\phi}$, we get from Claim 4.3 that $\text{Re}(\tilde{\sigma}) \in \text{Embed}_\sigma(\nu)$ and looking at their imaginary part we get that they form an embedding into an infinite cyclic group, hence by Lemma 3.20 they are equivalent to an embedding into a finite Abelian group, and again by Claim 4.3 we get that $\text{Im}(\tilde{\sigma}) \in \text{Embed}_\sigma(\nu)$. It follows that $\tilde{\sigma}$ is in $\text{Embed}_\sigma(\nu)$, hence $e^{\tilde{\sigma}}$ is in $\text{Embed}_\sigma(\nu)$. Hence $e^{\tilde{\sigma}}$ is constant on Σ_{modest} and so

$$\begin{aligned} \mathbb{E}_{x, x' \in \Sigma_{\text{modest}}} \left[|f(x) - f(x')|^2 \right] &= \mathbb{E}_{x, x' \in \Sigma_{\text{modest}}} \left[\left| (f(x) - e^{\tilde{\sigma}(x)/q}) - (f(x') - e^{\tilde{\sigma}(x')/q}) \right|^2 \right] \\ &\leq 4 \mathbb{E}_{x \in \Sigma_{\text{modest}}} \left[\left| f(x) - e^{\tilde{\sigma}(x)/q} \right|^2 \right] \\ &= 4 \mathbb{E}_{x \in \Sigma_{\text{modest}}} \left[|f(x)|^2 \left| e^{\tilde{\sigma}(x)/q - f'(x)} - 1 \right|^2 \right] \end{aligned}$$

For each x , $\frac{\tilde{\sigma}(x)}{q} - f'(x)$ is a complex number $a + \mathbf{b}\mathbf{i}$ for $a, b \in \mathbb{R}$ which are at most $\frac{1}{qN^{1/|S|}}$ in absolute value, hence

$$\left| e^{\tilde{\sigma}(x)/q - f'(x)} - 1 \right| \leq |e^a \cos(b) - 1| + |e^a \sin(b)| \leq O\left(\frac{1}{qN^{1/|S|}}\right) \leq \tau^8.$$

Thus, we get that

$$\mathbb{E}_{x, x' \in \Sigma_{\text{modest}}} \left[|f(x) - f(x')|^2 \right] \leq 4\tau^{16} \mathbb{E}_{x \in \Sigma_{\text{modest}}} \left[|f(x)|^2 \right] \leq 4\tau^{16}(\alpha')^{-1} \|f\|_2^2 \leq \tau^{15},$$

and contradiction to the assumption that $\mathbb{E}_{x, x' \in \Sigma_{\text{modest}}} \left[|f(x) - f(x')|^2 \right] \geq \tau$.

6.5.4 Proof of Lemma 6.9: ν Satisfies the Other Conditions of Theorem 6.7

The first, second, third and seventh properties of ν are all direct by the construction. We have already argued about the fourth and fifth properties of ν in the beginning of the proof of Lemma 6.9. Property 6a was established in Section 6.5.3, and now we only need to verify property 6b. For that we describe a distribution $\tilde{\nu}$ satisfying the properties therein.

Take $\Sigma'' = \{(x, \dots, x) \mid x \in \Sigma\}$, recall that $\Gamma' = \{(y, 1) \mid y \in \Gamma\}$ and let $\tilde{\nu}$ be the distribution of

$$((x, \dots, x), (y, 1), ((z, \dots, z), 1))$$

where $(x, y, z) \sim \mu$. The map a that we take is the same as the map a we defined to make the x -merge. Property 6b(i) is clear by the definition of Σ' and Σ'' , and property 6b(ii) is clear by the properties of the path trick. Property 6b(iv) follows as in μ the value of any two coordinates implies the last one, and property 6b(v) is immediate by the construction. Property 6b(vi) was verified above in the body of the proof of Lemma 6.9. We now argue about maximality, property 6b(iii).

Suppose we have $f: \Sigma'^n \rightarrow \mathbb{C}$, $g: \Gamma'^n \rightarrow \mathbb{C}$ and $h: \Phi'^n \rightarrow \mathbb{C}$ that are 1-bounded, and define $f': \Sigma^n \rightarrow \mathbb{C}$, $g': \Gamma^n \rightarrow \mathbb{C}$ and $h': \Phi^n \rightarrow \mathbb{C}$ by

$$\begin{aligned} f'(x_1, \dots, x_n) &= f((x_1, \dots, x_1), \dots, (x_n, \dots, x_n)), & g'(y_1, \dots, y_n) &= g((y_1, 1), \dots, (y_n, 1)), \\ h'(z_1, \dots, z_n) &= h(((z_1, \dots, z_1), 1), \dots, ((z_n, \dots, z_n), 1)). \end{aligned}$$

Then

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

The point is now that if we take a distribution \mathcal{D} over $\Sigma'' \times \Gamma' \times \Phi'$ whose support strictly contains the support of $\tilde{\nu}$, then we could repeat a similar reasoning above and relate our expectation to an expectation with respect to a distribution $\tilde{\mathcal{D}}$ whose support strictly contains μ . Indeed, let \mathcal{D} be such distribution and define $\tilde{\mathcal{D}}$ by taking $(\vec{x}, (y, 1), (\vec{z}, 1)) \sim \mathcal{D}$ and outputting (x, y, z) . Clearly the support of $\tilde{\mathcal{D}}$ contains the support of $\tilde{\mu}$, and hence the support of μ . If it was the case that $\text{supp}(\tilde{\mathcal{D}}) = \text{supp}(\mu)$, then we would get that

$$\text{supp}(\mathcal{D}) \subseteq \{(\vec{x}, (y, 1), (\vec{z}, 1)) \mid \exists x \in \Sigma, z \in \Phi, \vec{x} = (x, \dots, x), \vec{z} = (z, \dots, z), (x, y, z) \in \text{supp}(\mu)\},$$

but this is contained in the support of $\tilde{\nu}$ in contradiction. Thus we get that taking f' , g' and h' as above we have that

$$\mathbb{E}_{(\vec{x}, y, \vec{z}) \sim \mathcal{D}^{\otimes n}} [f(\vec{x})g(y)h(\vec{z})] = \mathbb{E}_{(x, y, z) \sim \tilde{\mathcal{D}}^{\otimes n}} [f'(x)g'(y)h'(z)],$$

and $\tilde{\mathcal{D}}$ is a distribution whose support strictly contains $\text{supp}(\mu)$. If $\text{NEStab}_{1-\delta}(g; \mathcal{D}) \leq \delta$, then by Lemma 2.9 we get that $\text{NEStab}_{1-\delta}(g'; \tilde{\mathcal{D}}) \leq \delta^s$, where $s = s(m, \alpha) > 0$ (as the master embedding of y in \mathcal{D} is a refinement of the master embedding of y in $\tilde{\mathcal{D}}$, we get that the corresponding averaging operators satisfy the properties of the lemma). Thus, by Claim 4.8 we get that $\text{NEStab}_{1-\delta^s}(g; \mathcal{D}) \leq \delta^s$, and as μ is maximal we get that

$$\left| \mathbb{E}_{(x, y, z) \sim \tilde{\mathcal{D}}^{\otimes n}} [f'(x)g'(y)h'(z)] \right| \leq M\delta^{s\eta}$$

for some M, η depending only on m, α , concluding the proof. \square

6.6 The Case that $|\Sigma_{\text{modest}}| = 1$: Proof of Lemma 6.8

Let f, g, h be functions as in Theorem 6.5. Using Lemma 3.9 repeatedly, we get that there is $\ell = O_{m,\alpha}(1)$ such that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f(x)g(y)h(z)] \right|^\ell \leq \left| \mathbb{E}_{(x,y,z) \sim \mu'^{\otimes n}} [F(x)g(y)H(z)] \right|,$$

where $F: \tilde{\Sigma}^n \rightarrow \mathbb{C}$ and $H: \tilde{\Phi}^n \rightarrow \mathbb{C}$ are some 1-bounded functions (and g remains the same). Our goal is to show that this quantity is at most $M\delta^\eta$ where $M \in \mathbb{N}$ and $\eta > 0$ depend only on m, α . By Lemma 5.5, this would follow if we show the statement for the x -merge distribution of μ'' , which is nothing but ν' , hence it suffices to show that there are M and $\eta > 0$ such that for $F: \tilde{\Sigma}_{\text{final}} \rightarrow [-1, 1]$, $g: \Gamma^n \rightarrow [-1, 1]$ and $h: \tilde{\Phi}^n \rightarrow [-1, 1]$ such that $\text{NEStab}_{1-\delta}[g; \nu_y^{\otimes n}] \leq \delta$ it holds that

$$\left| \mathbb{E}_{(\tilde{x}, y, \tilde{z}) \sim \nu'^{\otimes n}} [F(\tilde{x})g(y)H(\tilde{z})] \right| \leq M\delta^\eta.$$

We focus on this task henceforth. Recall $\Sigma' = \{a(x, \dots, x) \mid x \in \Sigma\}$ and define $\Phi'_{\text{pre}} = \{(z, \dots, z) \mid z \in \Gamma\}$ as well as $\Sigma'_{\text{pre}} = \{(x, \dots, x) \mid x \in \Sigma\}$. Let $\tilde{\nu}$ be the distribution of $(\tilde{x}, y, \tilde{z}) \sim \nu$ conditioned on $\tilde{x} \in \Sigma'$ and $\tilde{z} \in \Phi'$, and write $\nu' = \beta\tilde{\nu} + (1-\beta)\nu''$ where $\beta > 0$ depends only on α and ν'' is some distribution. Choose $J \subseteq_\beta [n]$, $(\tilde{x}', y', \tilde{z}') \sim \nu''^J$ and define $\tilde{F}: \Sigma'^J \rightarrow \mathbb{C}$, $\tilde{g}: \Gamma^J \rightarrow \mathbb{C}$ and $\tilde{H}: \tilde{\Phi}^J \rightarrow \mathbb{C}$ by

$$\tilde{F} = F_{\tilde{J} \rightarrow \tilde{x}'}, \quad \tilde{g} = g_{\tilde{J} \rightarrow y'}, \quad \tilde{H} = H_{\tilde{J} \rightarrow \tilde{z}'}$$

Denote $\phi_{\nu'}(F, g, H) = \mathbb{E}_{(\tilde{x}, y, \tilde{z}) \sim \nu'^{\otimes n}} [F(\tilde{x})g(y)H(\tilde{z})]$ and analogously define $\phi_{\tilde{\nu}}(\tilde{F}, \tilde{g}, \tilde{H})$. Then

$$\phi_{\nu'}(F, g, H) = \mathbb{E}_{(\tilde{x}', y', \tilde{z}') \sim \nu''^J} \left[\phi_{\tilde{\nu}}(\tilde{F}, \tilde{g}, \tilde{H}) \right] = \underbrace{\mathbb{E}_{(\tilde{x}', y', \tilde{z}') \sim \nu''^J} \left[1_E \phi_{\tilde{\nu}}(\tilde{F}, \tilde{g}, \tilde{H}) \right]}_{(I)} + \underbrace{\mathbb{E}_{(\tilde{x}', y', \tilde{z}') \sim \nu''^J} \left[1_{\bar{E}} \phi_{\tilde{\nu}}(\tilde{F}, \tilde{g}, \tilde{H}) \right]}_{(II)}, \quad (10)$$

where E is the event that $\text{NEStab}_{1-c^{-1}\beta^{-1}\delta, \nu'}(\tilde{g}; \tilde{\nu}_y^J) \leq \sqrt{\delta}$; here, $c > 0$ is from Claim 4.9 (with $\mathcal{D} = \nu'_y$, $\mathcal{D}' = \tilde{\nu}_y$ and $\mathcal{D}'' = \nu''_y$ in the notations therein). By Claim 4.9 we have

$$\mathbb{E}_{J, y'} [\text{NEStab}_{1-c^{-1}\beta^{-1}\delta, \nu'}(\tilde{g}; \tilde{\nu}_y^J)] \leq \text{NEStab}_{1-\delta, \nu'}(g; \nu_y^{\otimes n}) \leq \delta,$$

so by Markov's inequality $\Pr[\bar{E}] \leq \sqrt{\delta}$. Hence, as \tilde{F}, \tilde{g} and \tilde{H} are 1-bounded, we have that $|(II)| \leq \Pr[\bar{E}] \leq \sqrt{\delta}$. In the rest of the argument we bound (I).

Bounding (I). Fix J and $\tilde{x}', y', \tilde{z}'$ so that the event E holds; we show that then $\left| \phi_{\tilde{\nu}}(\tilde{F}, \tilde{g}, \tilde{H}) \right| \leq M\delta^\eta$, where $M \in \mathbb{N}$ and $\eta > 0$ only depend on m and α . The main idea is to re-interpret $\phi_{\tilde{\nu}}(\tilde{F}, \tilde{g}, \tilde{H})$ as an expectation with respect to a distribution over $\Sigma \times \Gamma \times \Phi$ whose support strictly contains the support of μ , and then use the maximality of μ .

Without loss of generality, we assume for notational convenience that $J = \{1, \dots, n'\}$. Consider the functions $F^\sharp: \Sigma^J \rightarrow \mathbb{C}$ and $H^\sharp: \Phi^J \rightarrow \mathbb{C}$ defined as

$$F^\sharp(x) = \tilde{F}(a(x_1, \dots, x_1), \dots, a(x_{n'}, \dots, x_{n'})), \quad H^\sharp(z) = \tilde{H}((z_1, \dots, z_1), \dots, (z_{n'}, \dots, z_{n'})).$$

Consider the distribution \mathcal{D} over $\Sigma \times \Gamma \times \Phi$ that results from sampling (X, y, Z) according to $\tilde{\nu}$, writing $Z = (z, \dots, z)$ and sampling $(x, \dots, x) \sim \mu''|_{\Sigma}$ conditioned on $a(x, \dots, x) = X$, and then outputting (x, y, z) . Then:

$$\phi_{\tilde{\nu}}(\tilde{F}, \tilde{g}, \tilde{H}) = \mathbb{E}_{(x,y,z) \sim \mathcal{D}^J} \left[F^\sharp(x) \tilde{g}(y) H^\sharp(z) \right]. \quad (11)$$

It is easily seen that the probability of each atom in \mathcal{D} is at least $\alpha' > 0$ that depends only on m and α . We argue that $\text{supp}(\mu)$ is strictly contained in $\text{supp}(\mathcal{D})$. First, containment is clear by construction, and we next observe the strict containment. In μ we have that y, z implies x , and in the construction of ν we first took distinct $x^*, x^{*'} \in \Sigma$ (satisfying some other property that is not important for now), and took $\Sigma_{\text{modest}} = \{a(x^*, \dots, x^*), a(x^{*'}, \dots, x^{*'})\}$. Thus, there are $y, y' \in \Gamma$ and $z, z' \in \Phi$ such that $(x^*, y, z), (x^{*'}, y', z') \in \text{supp}(\mu)$, and it must be the case that (x^*, y', z') is not in $\text{supp}(\mu)$ (otherwise, this would mean that the value of the y and z coordinate does not imply the value of the x coordinate in μ). We argue that (x^*, y', z') is in $\text{supp}(\mathcal{D})$. Indeed, as by assumption $|\Sigma_{\text{modest}}| = 1$ we have that $a(x^*, \dots, x^*) = a(x^{*'}, \dots, x^{*'})$, and we have that for $X = a(x^{*'}, \dots, x^{*'})$, $Z = (z', \dots, z')$ it holds that (X, y', Z) is in $\text{supp}(\nu)$, hence by definition of \mathcal{D} and the fact that $a(x^*, \dots, x^*) = X$ we get that (x^*, y', z') is in $\text{supp}(\mathcal{D})$.

In conclusion, we get that the expectation on the right hand side of (11) is an expectation with respect to a distribution \mathcal{D} whose support strictly contains the support of μ and hence we can use the maximality of μ . One subtle point is that the ‘‘high degreeeness’’ of g is not phrased in quite the appropriate language; as the event E holds we know that $\text{NEStab}_{1-c-1\beta-1\delta, \nu'}(\tilde{g}; \tilde{\nu}_y^J) \leq \sqrt{\delta}$ and we need to conclude from this high-degreeeness with respect to non-embeddability in \mathcal{D} .

As $\tilde{\nu}_y = \mathcal{D}_y$ we have $\text{NEStab}_{1-c-1\beta-1\delta, \nu'}(\tilde{g}; \tilde{\nu}_y^J) = \text{NEStab}_{1-c-1\beta-1\delta, \nu'}(\tilde{g}; \mathcal{D}_y^J)$. Next, note that any Abelian embedding of ν' can be used to define an Abelian embedding of \mathcal{D} . This is done by mapping x to $a(x, \dots, x)$ and then applying the embedding of ν' on the first coordinate, applying the embedding of ν' on the second coordinate on y , and mapping z to (z, \dots, z) and then applying the embedding of ν' on the third coordinate. Thus, the partition of Γ defined by master embedding of ν' is a refinement of the partition of Γ defined by the master embedding of \mathcal{D} . Applying Lemma 2.9 we conclude that

$$\text{NEStab}_{1-c-1\beta-1\delta, \mathcal{D}}(\tilde{g}; \mathcal{D}_y^J) \leq \text{NEStab}_{1-c-1\beta-1\delta, \nu'}(\tilde{g}; \mathcal{D}_y^J)^s \leq \delta^{s/2}.$$

where $s = s(m, \alpha) > 0$. By Claim 4.8 we get that $\text{NEStab}_{1-\delta^{s/2}, \mathcal{D}}(\tilde{g}; \mathcal{D}_y^J) \leq \delta^{s/2}$ provided that δ_0 is small enough, hence by maximality of μ it follows that

$$\left| \mathbb{E}_{(x,y,z) \sim \mathcal{D}^J} \left[F^\sharp(x) \tilde{g}(y) H^\sharp(z) \right] \right| \leq M' \delta^{s\eta'/2}$$

where M' and $\eta' > 0$ depends only on α and m . Plugging this into (11) we have that $|\phi_{\nu'}(\tilde{F}, \tilde{g}, \tilde{H})| \leq \delta^{s\eta'/2}$, and so $|(I)| \leq \delta^{s\eta'/2}$.

Combining the bounds on (I), (II). Plugging the bounds on (I), (II) into (10) yields that

$$|\phi_{\nu}(F, g, H)| \leq M' \delta^{s\eta'/2} + \sqrt{\delta} \leq M'' \delta^{\eta''}$$

for $M'' = M' + 1$ and $\eta'' = s\eta'/2$, and we are done. \square

7 Reducing to the Homogenous Statement

In this section, our goal is to reduce Theorem 6.7 into a result that relaxes the assumption that our functions are 1-bounded to the assumption that they are bounded 2-norm (which is therefore more amendable to a proof by induction). See Theorem 7.23 for a precise statement.

7.1 Degree, Non-embedding degree and Effective Non-embedding Degree

Recall that in Section 4.2 we showed that given a distribution μ that has saturated master embeddings, one may construct a basis $B_1 \cup B_2$ for $L_2(\Sigma; \mu_x)$ in which B_1 consists of embedding functions and B_2 consists of functions that are orthogonal to all embeddings functions. The goal of this section is to refine this further so as to be more compatible with the relaxed base case.

Definition 7.1. Let μ be a distribution over $\Sigma \times \Gamma \times \Phi$ as in Theorem 6.7. We define an orthonormal basis $B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}}$ for $L_2(\Sigma; \mu_x)$, as follows:

1. Consider the space of embedding functions, $\text{Embed}_\sigma(\mu) = \text{span} \left(\left\{ \chi \circ \sigma \mid \chi \in \widehat{H} \right\} \right)$, and pick an orthonormal basis $B_{\text{embed}} = \left\{ \chi \circ \sigma \mid \chi \in \widehat{H} \right\}$ for it.
2. Consider the orthogonal space to $\text{Embed}_\sigma(\mu)$, namely $\text{Embed}_\sigma(\mu)^\perp$, and consider the subspace of it of functions that are constant on Σ_{modest} :

$$\text{Non-Embed-Non-Modest}_\sigma(\mu) = \left\{ f: \Sigma \rightarrow \mathbb{C} \mid f \in \text{Embed}_\sigma(\mu)^\perp, f|_{\Sigma_{\text{modest}}} \text{ is constant} \right\}.$$

Pick $B_{\text{non-embed}}$ an orthonormal basis of $\text{Non-Embed-Non-Modest}_\sigma(\mu)$.

3. Consider the space $\text{Non-Embed-Non-Modest}_\sigma(\mu)^\perp \cap \text{Embed}_\sigma(\mu)^\perp$, and take B_{modest} to be an orthonormal basis for it.

With the basis $B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}}$ in hand, we can now construct an orthonormal basis for $L_2(\Sigma^n, \mu_x^{\otimes n})$ by tensorizing. Namely, we take $(B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^{\otimes n}$, as an orthonormal basis of $L_2(\Sigma^n, \mu_x^{\otimes n})$, and thus we may write any $f: \Sigma^n \rightarrow \mathbb{C}$ as

$$f(x) = \sum_{\chi \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^{\otimes n}} \widehat{f}(\chi) \prod_{i=1}^n \chi_i(x_i), \quad \text{where } \widehat{f}(\chi) = \langle f, \chi \rangle.$$

Definition 7.2. A function χ of the form $\chi(x) = \prod_{i=1}^n \chi_i(x_i)$ where $\chi_i \in B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}}$ for all i is called a monomial.

There are also several important notions of degree that may be associated with monomials, which extend Definition 4.5.

Definition 7.3. Let $\chi = \prod_{i=1}^n \chi_i$ be a function in $(B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^{\otimes n}$.

1. For a character $a \in \widehat{H}$, the a -embedding degree of χ , denoted by $\text{embeddeg}_a(\chi)$ is the number of coordinates i on which $\chi_i = a$.

2. The embedding degree of χ , denoted by $\text{embeddeg}(\chi)$, is the number of coordinates i on which $\chi_i \in B_{\text{embed}}$. In other words, $\text{embeddeg}(\chi) = \sum_a \text{embeddeg}_a(\chi)$.
3. The non-embedding degree of χ , denoted by $\text{nedeg}(\chi)$, is the number of i on which $\chi_i \notin B_{\text{embed}}$, that is, $|\{i \in [n] \mid \chi_i \notin B_{\text{embed}}\}|$.
4. The effective non-embedding degree of χ , denoted by $\text{eff-nedeg}(\chi)$, is the number of i on which $\chi_i \in B_{\text{modest}}$, that is, $|\{i \in [n] \mid \chi_i \in B_{\text{modest}}\}|$.

We note that clearly, for every monomial χ it holds that $\text{eff-nedeg}(\chi) \leq \text{nedeg}(\chi) \leq n$, and that $\text{nedeg}(\chi) + \text{embeddeg}(\chi) = n$.

7.1.1 The Modest Markov Chain

Recalling the non-embedding noise operator $T_{\text{non-embed},1-\delta}^{\otimes n}: L_2(\Sigma^n, \mu_x^{\otimes n}) \rightarrow L_2(\Sigma^n, \mu_x^{\otimes n})$, we have by Fact 4.7 that $T_{\text{non-embed},1-\delta}^{\otimes n} \chi = (1 - \delta)^{\text{nedeg}(\chi)} \chi$. Next, we design the *effective non-embedding* noise operator, which will be helpful for us later on. Towards this end, we define the modest Markov chain associated with μ on x .

Definition 7.4. *The modest Markov chain, on Σ , denoted by Modest , is the Markov chain that on $x \in \Sigma$ takes $x' = x$ if $x \notin \Sigma_{\text{modest}}$, and otherwise, if $x \in \Sigma_{\text{modest}}$, samples $x' \sim \mu_x$ conditioned on $x' \in \Sigma_{\text{modest}}$.*

7.1.2 The Effective Noise Operator and Effective Non-embedding Degree

The modest Markov chain will be useful for us to define several notions. The first of which is the effective noise operator, and we first define the $E_{\text{non-embed},1-\delta}$.

Definition 7.5. *For $\delta > 0$, the Markov chain $E_{\text{non-embed},1-\delta}$ on Σ is the Markov chain that on $x \in \Sigma$, takes $x' = x$ with probability $1 - \delta$, and otherwise samples a neighbour x' of x according to the modest Markov chain.*

Clearly, μ_x is a stationary distribution of $E_{\text{non-embed},1-\delta}$, and as usual we associated this Markov chain an averaging operator acting. Abusing notation, we denote it by $E_{\text{non-embed},1-\delta}: L_2(\Sigma; \mu_x) \rightarrow L_2(\Sigma; \mu_x)$ and define it as

$$E_{\text{non-embed},1-\delta} f(x) = \mathbb{E}_{x' \sim E_{\text{non-embed},1-\delta} x} [f(x')].$$

The following lemma gives the most basic properties of $E_{\text{non-embed},1-\delta}$. In words, any function in $B_{\text{embed}} \cup B_{\text{non-embed}}$ is an eigenfunction of it with eigenvalue 1 and any function in B_{modest} is an eigenfunction with eigenvalue $1 - \delta$.

Lemma 7.6. *For all $\delta > 0$ we have that for any $\chi \in B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}}$,*

1. If $\chi \in B_{\text{modest}}$, then $E_{\text{non-embed},1-\delta} \chi = (1 - \delta) \chi$.
2. Else, $E_{\text{non-embed},1-\delta} \chi = \chi$.

Consequently, for a monomial $\chi: \Sigma^n \rightarrow \mathbb{C}$ we have $E_{\text{non-embed},1-\delta}^{\otimes n} \chi = (1 - \delta)^{\text{eff-nedeg}(\chi)} \chi$.

Proof. For the first item, note that for all $x \in \Sigma$,

$$\mathbb{E}_{\text{non-embed}, 1-\delta} \chi(x) = (1 - \delta)\chi(x) + \delta \mathbb{E}_{x' \sim \text{Modest } x} [\chi(x')],$$

and we show the last expectation is equal to 0. We may write this expectation as a sum $\sum_{x' \in \Sigma} p(x')\chi(x')$ where $p(x')$ is the probability that a set on the modest Markov chain reaches x' from x . Clearly, we may write this sum as $\langle \chi, \frac{p}{\mu_x} \rangle_{\mu_x}$. We note that p/μ_x is constant on Σ_{modest} . Indeed, if $a, b \in \Sigma_{\text{modest}}$, then $p(a) = \frac{\mu_x(a)}{\mu_x(\Sigma_{\text{modest}})}$ and $p(b) = \frac{\mu_x(b)}{\mu_x(\Sigma_{\text{modest}})}$. Thus, we have that p/μ_x is in the span of $B_{\text{embed}} \cup B_{\text{non-embed}}$ (as these contain all functions that are constant on Σ_{modest}), hence p/μ_x is perpendicular to B_{modest} and so $\langle \chi, \frac{p}{\mu_x} \rangle_{\mu_x} = 0$.

For the second item, note that any such χ is constant on Σ_{modest} and the Markov chain $\mathbb{E}_{\text{non-embed}, 1-\delta}$ stays at any $x \notin \Sigma_{\text{modest}}$, and otherwise stays inside Σ_{modest} . \square

7.1.3 Modest Influences

The modest Markov chain will also be useful for us to define the notion of modest influence of a coordinate as well as the modest total influence.

Definition 7.7. *In the setting of Theorem 6.7, for a function $f: (\Sigma^n, \mu_x^{\otimes n}) \rightarrow \mathbb{C}$ and a coordinate $i \in [n]$, we define the modest influence of f at coordinate i as*

$$I_{i, \text{modest}}[f] = \mathbb{E}_{x' \sim \mu_x^{n-1}, a \sim \mu_x, b \sim \text{Modest } a} \left[|f(x_{-i} = x', x_i = a) - f(x_{-i} = x', x_i = b)|^2 \right].$$

The modest total influence of f is $I_{\text{modest}}[f] = \sum_{i=1}^n I_{i, \text{modest}}[f]$.

As usual, using Parseval's equality we get an analytical formula for influences in terms of the orthogonal decomposition of f .

Fact 7.8. *In the setting of Theorem 6.7, for a function $f: (\Sigma^n, \mu_x^{\otimes n}) \rightarrow \mathbb{C}$ and a coordinate $i \in [n]$, we have*

$$I_{i, \text{modest}}[f] = 2 \sum_{\chi: \chi_i \in B_{\text{modest}}} |\widehat{f}(\chi)|^2.$$

Subsequently, the modest total influence of f is $I_{\text{modest}}[f] = 2 \sum_{\chi} \text{eff-nedeg}(\chi) |\widehat{f}(\chi)|^2$.

7.2 Embedding Homogenous functions and Effectively Homogenous functions

Equipped with the notions of embedding degree and effective embedding degree, we may define homogeneity with respect to them.

Definition 7.9. *We say f is completely embedding homogenous of degree D if there are integers $\{D_a\}_{a \in \hat{H}}$ that sum up to D such that for all $\chi \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^{\otimes n}$ such that $\widehat{f}(\chi) \neq 0$, we have that $\text{embeddeg}_a(\chi) = D_a$.*

We say that f is completely embedding homogenous if it is completely embedding homogenous of degree D for some D .

Definition 7.10. A function $f: \Sigma^n \rightarrow \mathbb{C}$ is called *non-embedding homogenous of non-embedding degree d* if for all $\chi \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^{\otimes n}$ such that $\widehat{f}(\chi) \neq 0$ it holds that $\text{nedeg}(\chi) = d$.

We say f is *non-embedding homogenous* if it is non-embedding homogenous of non-embedding degree d for some d .

Definition 7.11. We say a function $f: \Sigma^n \rightarrow \mathbb{C}$ has *effective non-embedding degree at least d* if for all $\chi \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^{\otimes n}$ such that $\widehat{f}(\chi) \neq 0$ it holds that $\text{eff-nedeg}(\chi) \geq d$.

With the notions of homogenous functions, we may now formulate a version of Theorem 6.7 for homogenous functions of bounded L_2 norm. We first define the collections of these homogenous functions:

Definition 7.12. Let $n \in \mathbb{N}$.

1. For d, d' we define the class $\mathcal{F}'_{n,d,d'}$ to be the collection of functions $f: \Sigma^n \rightarrow \mathbb{C}$ that are completely embedding homogenous, non-embedding homogenous of degree d and have effective degree at least d' .
2. We define the class \mathcal{G}_n to be the class of functions $g: \Gamma^n \rightarrow \mathbb{C}$ that are completely embedding homogenous and non-embedding homogenous.
3. We define the class \mathcal{H}_n to be the class of all functions $h: \Phi^n \rightarrow \mathbb{C}$ that are completely embedding homogenous and non-embedding homogenous.

We can now define the parameter β :

Definition 7.13. For integers $n \geq d \geq d'$, finite alphabets Σ, Γ, Φ and a distribution μ over $\Sigma \times \Gamma \times \Phi$ as in Theorem 6.7, we define

$$\beta'_{n,d,d'}[\mu] = \sup_{\substack{f \in \mathcal{F}'_{n,d,d'} \\ g \in \mathcal{G}_n \\ h \in \mathcal{H}_n}} \frac{|\mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f(x)g(y)h(z)]|}{\|f\|_2 \|g\|_2 \|h\|_2}.$$

When the distribution μ is clear from context, we often drop it from the notation and denote the parameter simply by $\beta'_{n,d,d'}$. We are now ready to formulate the homogenous version of Theorem 6.7.

Theorem 7.14. For all $\alpha > 0$ and $m \in \mathbb{N}$ there are $\xi > 0$ and $c > 0$ such that the following holds. Let Σ, Γ and Φ be alphabets of size at most m , and let μ be a distribution over $\Sigma \times \Gamma \times \Phi$ as in Theorem 6.7. Suppose that $d, d' \in \mathbb{N}$ satisfy that $d' \geq d^{1-\xi}$. Then

$$\beta'_{n,d,d'}[\mu] \leq (1+c)^{-d^{1-2\xi}}.$$

We have the following claim asserting that Theorem 7.14 implies Theorem 6.7. The proof uses ‘‘soft-truncation’’ and ‘‘truncation’’ type argument, and is deferred to Section B.

Claim 7.15. Theorem 7.14 implies Theorem 6.7.

Proof. Deferred to Section B. □

7.3 Reformulating Theorem 7.14: Functions that are Constant on Connected Components

As for the proof of Theorem 7.14, for technical reasons it will be more convenient for us to view the parameter β'_{n,d_1,d'_1} in a different but equivalent way, and prove an analogous statement for it. Fix a distribution μ as in Theorem 7.14; the distribution $\mu_{y,z}$ is uniform and hence is very nice to work with, but the distribution μ_x may be more complicated, hence it will be more convenient for us to switch to a statement that is only concerned with $\mu_{y,z}$. For that, we are going to use the fact that as in μ it holds that the value of y, z implies x , there is a natural identification between functions over x , and a certain class of functions over (y, z) .

Given a function $f: \Sigma^n \rightarrow \mathbb{C}$, we may define the function $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{C}$ by

$$F(y, z) = f(x)$$

where $x \in \Sigma^n$ is the unique point such that (x_i, y_i, z_i) is in the support of μ for all i (recall that this x is unique as in μ , it holds that y, z implies x). We view this transformation as a mapping $W: L_2(\Sigma^n; \mu_x^{\otimes n}) \rightarrow L_2(\Gamma^n \times \Phi^n; \mu_{y,z}^{\otimes n})$. A function in the image of W is not an arbitrary function, and we refer to such functions as functions that are constant on connected components:

Definition 7.16. *We say a function $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{C}$ is constant on connected components if there is $f: \Sigma^n \rightarrow \mathbb{C}$ such that $F = Wf$.*

The reason for this terminology is that we consider a graph (V, E) whose vertex set is $V = \Gamma \times \Phi$, and the vertices (y, z) and (y', z') are adjacent if there is an x such that (x, y, z) and (x, y', z') are both in $\Gamma \times \Phi$; then a function $F: \Gamma \times \Phi \rightarrow \mathbb{C}$ that is constant on connected components as per the above definition precisely corresponds to a function that is constant on each connected component of this graph:

Lemma 7.17. *The following conditions are equivalent for $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{C}$:*

1. F is constant on connected components as per Definition 7.16.
2. F is constant on each connected component of $(V, E)^{\otimes n}$.
3. $F \in \text{span}(\{W\chi \mid \chi \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^{\otimes n}\})$.

Proof. We show that the first item is equivalent to the second item and that the first item is equivalent to the third item.

The first item is equivalent to the second item. Given F as the second item, for $x \in \Sigma^n$ we take $y \in \Gamma^n$ and $z \in \Phi^n$ such that (x_i, y_i, z_i) is in the support of μ for all i , and define $f(x) = F(y, z)$. We note that this is well defined, as for such potential y, z the value of $F(y, z)$ is the same (as they are all in the same connected component of $(V, E)^{\otimes n}$). Thus, $F = Wf$ and so F is constant on connected components as per Definition 7.16.

Given F as in the first item, we have that $F = Wf$ for some $f: \Sigma^n \rightarrow \mathbb{C}$. If (y, z) and (y', z') are adjacent in the graph $(V, E)^{\otimes n}$, then there is $x \in \Sigma^n$ such that (x_i, y_i, z_i) and (x_i, y'_i, z'_i) are in the support for μ for all i , hence $F(y, z) = f(x)$ and $F(y', z') = f(x)$ by the definition of W . It follows that the value of F is the same vertices that are adjacent in $(V, E)^{\otimes n}$, hence F is constant on the connected components of $(V, E)^{\otimes n}$.

The first item is equivalent to the third item. Taking F as in the first item, writing $F = Wf$ and expanding f according to the basis $(B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^{\otimes n}$ and using the linearity of W , the third item follows. In the reverse, given

$$F = \sum_{\chi} a_{\chi} W\chi,$$

we can write $F = Wf$ for $f = \sum_{\chi} a_{\chi} \chi$. □

We note that the set $\{W\chi \mid \chi \in B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}}\}$ is an orthonormal set in $L_2(\Gamma \times \Phi, \mu_{y,z})$, as

$$\langle W\chi, W\chi' \rangle_{\mu_{y,z}} = \langle \chi, \chi' \rangle_{\mu_x} = 1_{\chi=\chi'},$$

hence $\{W\chi \mid \chi \in B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}}\}^{\otimes n}$ is an orthonormal basis to the space of functions $F \in L_2(\Gamma^n \times \Phi^n; \mu_{y,z}^{\otimes n})$ that are constant on connected components. Thus, to translate Theorem 7.14 to the language of functions that are constant on connected components, it remains to discuss the analog of degrees, non-embedding degrees and effective non-embedding degrees. These are all very natural analogs of the notions we have already seen.

Definition 7.18. A function $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{C}$ which is constant on connected components is said to be completely embedding homogenous of degree D if there are integers $\{D_a\}_{a \in \hat{H}}$ summing to D such that for all $\chi \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^{\otimes n}$ for which $\hat{F}(W\chi) \neq 0$ it holds that $\text{embeddeg}_a(\chi) = D_a$.

We say F is completely embedding homogenous if it is completely embedding homogenous of degree D for some D .

Definition 7.19. A function $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{C}$ which is constant on connected components is called non-embedding homogenous of non-embedding degree d if for all $\chi \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^{\otimes n}$ such that $\hat{F}(W\chi) \neq 0$ it holds that $\text{nedeg}(\chi) = d$.

Definition 7.20. We say a function $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{C}$ which is constant on connected components has effective non-embedding degree at least d if for all $\chi \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^{\otimes n}$ such that $\hat{F}(W\chi) \neq 0$ it holds that $\text{eff-nedeg}(\chi) \geq d$.

Lastly, we define the analog of the class $\mathcal{F}'_{n,d,d'}$.

Definition 7.21. For $n \geq d \geq d'$ we define the class $\mathcal{F}_{n,d,d'}$ to be the collection of functions $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{C}$ that are constant on connected components, are completely embedding homogenous, are non-embedding homogenous of degree d and effective non-embedding degree at least d' .

With all of this, we now define the analog of the parameter $\beta'_{n,d,d'}[\mu]$ for functions that are constant on connected components.

Definition 7.22. For integers $n \geq d \geq d'$, finite alphabets Σ, Γ, Φ and a distribution μ over $\Sigma \times \Gamma \times \Phi$ as in Theorem 6.7, we define

$$\beta_{n,d,d'}[\mu] = \sup_{\substack{F \in \mathcal{F}_{n,d,d'} \\ g \in \mathcal{G}_n \\ h \in \mathcal{H}_n}} \frac{|\mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [F(y,z)g(y)h(z)]|}{\|F\|_2 \|g\|_2 \|h\|_2}.$$

Often the distribution μ will be clear from context, in which case we may omit it from the notation and simply write $\beta_{n,d,d'}$. The following result is an equivalent formulation of Theorem 7.14 in the language of functions that are constant on connected components; its proof is given in Sections 8 and 9.

Theorem 7.23. *For all $\alpha > 0$ and $m \in \mathbb{N}$ there are $\xi > 0$ and $c > 0$ such that the following holds. Let Σ, Γ and Φ be alphabets of size at most m , and let μ be a distribution over $\Sigma \times \Gamma \times \Phi$ as in Theorem 6.7. Suppose that $n \geq d \geq d'$ satisfy that $d' \geq d^{1-\xi}$. Then*

$$\beta_{n,d,d'}[\mu] \leq (1+c)^{-d^{1-2\xi}}.$$

7.3.1 Theorem 7.23 implies Theorem 7.14

We first show that Theorem 7.23 implies Theorem 7.14.⁷ Let f, g, h be as in the definition of $\beta_{n,d,d'}[\mu]$. Taking $F = Wf$, we see that $\|F\|_2 = \|f\|_2$ and

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

Moreover, by Lemma 7.17 and the definitions it follows that F is completely embedding homogenous, non-embedding homogenous of degree d and effective non-embedding of degree at least d' . It follows that the supremum defining $\beta'_{n,d,d'}[\mu]$ is at most the supremum defining $\beta_{n,d,d'}[\mu]$, hence by Theorem 7.23 we get that $\beta'_{n,d,d'}[\mu] \leq \beta_{n,d,d'}[\mu] \leq (1+c)^{-d^{1-2\xi}}$.

7.4 Influences for Functions Constant on Connected Components

As the map W is a 1-to-1 correspondence between functions over x , and functions over y, z that are constant on connected components, we get analogs of the various notions of influences (non-embedding and modest) for functions that are constant on connected components. Below we define them formally.

Definition 7.24. *In the setting of Theorem 6.7, let $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{C}$ be a function which is constant on connected components, and let $f: \Sigma^n \rightarrow \mathbb{C}$ be the unique function such that $F = Wf$. For a coordinate $i \in [n]$:*

1. *The non-embedding influence of i on F , denoted by $I_{i,\text{non-embed}}[F]$, is defined to be $I_{i,\text{non-embed}}[f]$. The total non-embedding influence of F is $I_{\text{non-embed}}[F] = \sum_{i=1}^n I_{i,\text{non-embed}}[F]$.*
2. *The modest influence of i on F , denoted by $I_{i,\text{modest}}[F]$, is defined to be $I_{i,\text{modest}}[f]$. The total modest influence of F is $I_{\text{modest}}[F] = \sum_{i=1}^n I_{i,\text{modest}}[F]$.*

8 Proof of Theorem 7.23: Reducing to Near Linear Degrees

In this section we begin the proof of Theorem 7.23 and show the following lemma.

Lemma 8.1. *For all $m \in \mathbb{N}$ and $\alpha > 0$ there are $\varepsilon > 0$ and $L > 0$ such that the following holds. Let $n \geq d \geq d'$ be integers and let μ be as in Theorem 7.23. If $n \geq Ld$, then*

$$\beta_{n,d,d'} \leq \max(\beta_{n-1,d,d'}, (1-\varepsilon)\beta_{n-1,d-1,d'-1}).$$

⁷The reverse direction is also true and the proof proceeds via similar lines, but is omitted since it is unnecessary for us.

In words, Lemma 8.1 says that when we look at β , we can either decrease n and leave the non-embedding degree d , and the lower bound on the effective degree d' to be the same and get a bound by a similar parameter β ; else we drop d and d' by 1 but to compensate for that we gain a factor of $(1 - \varepsilon)$.

After proving Lemma 8.1 we are going to iterate it and conclude that either we are done (by gaining sufficiently many factors of $1 - \varepsilon$), or else we have reduced n all the way down to Ld , in which case the number of variables is linear in the non-embedding degree of the x -function and thus nearly linear in the effective non-embedding degree. The case that $n \approx Ld$ will be the subject of discussion in Section 9.

The rest of this section is devoted to the proof of Lemma 8.1, and our argument closely follows [5, Section 5], with a few additional complication due to the greater generality in our case. In a sense, one could think of the argument in [5, Section 5] as addressing the case that there is only the trivial embedding functions corresponding to the all 1 character, hence the constant all 1 function plays a special role therein. In our context embedding functions play the role of the constant functions in the setting of [5], and non-embedding functions play the role of functions with average 0 in the setting of [5]. The presence of multiple Abelian embeddings however makes the argument here a bit more tricky, and this is ultimately the reason we needed the notion of completely embedding homogenous functions.

8.1 Preliminaries: the Additive Base Case and Some Simple Orthogonalities

In our argument we are going to need to upper bound the absolute value of expectations of the form

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

for univariate functions. We may of course apply Cauchy-Schwarz to bound this, but importantly the following claim shows that if f is perpendicular to all embedding functions, then we can gain a factor of $(1 - c)$ over the trivial bound given by Cauchy-Schwarz.

Claim 8.2. *Let Σ , Γ and Φ be alphabets of size at most m , and let μ be a distribution as in Theorem 6.7. Then there exists $c = c(m, \alpha) > 0$ such that for all functions $f: \Sigma \rightarrow \mathbb{C}$, $g: \Gamma \rightarrow \mathbb{C}$ and $h: \Phi \rightarrow \mathbb{C}$ such that $f \in \text{Embed}_\sigma(\mu)^\perp$ we have that*

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

Proof. Assume this is not the case. Thus, we may find a sequence of functions (f_m, g_m, h_m) and measures μ_m such that $\|f_m\|_2 = 1$, $\|g_m + h_m\|_2 = 1$, $f_m \in \text{Embed}_\sigma(\mu_m)^\perp$, and

$$\left| \mathbb{E}_{(x,y,z) \sim \mu_m} [f_m(x)(g_m(y) + h_m(z))] \right| \geq 1 - \frac{1}{m}.$$

Passing to subsequences, we may assume that $\text{supp}(\mu_m)$ is the same for all m , and passing to further 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 and μ_m converges to a measure μ . We get that $f \in \text{Embed}_\sigma(\mu)^\perp$, $\|f\|_2 = \|g + h\|_2 = 1$ and

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

On the other hand, by Cauchy-Schwarz we get that

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

hence Cauchy-Schwarz is tight. Therefore there is $\theta \in \mathbb{C}$ of absolute value 1 such that $f(x) = \theta \overline{(g(y) + h(z))}$ for all $(x, y, z) \in \mathbb{C}$. Claim 4.3 implies now that $f \in \text{Embed}_\sigma(\mu)$, and this means that f must be identically 0 in contradiction to the fact that $\|f\|_2 = 1$. \square

Next, we have the following claim asserts that embedding functions in x are perpendicular to non-embedding functions in either y or z .

Claim 8.3. *Let Σ, Γ and Φ be alphabets of size at most m , and let μ be a distribution as in Theorem 6.7.*

1. *If $f \in \text{Embed}_\sigma(\mu)$, $g \in \text{Embed}_\gamma(\mu)^\perp$ and h is any function, then $\mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h(z)] = 0$.*
2. *If $f \in \text{Embed}_\sigma(\mu)$ and $h \in \text{Embed}_\phi(\mu)^\perp$ and g is any function, then $\mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h(z)] = 0$.*
3. *If $f \in \text{Embed}_\sigma(\mu)$, $g \in \text{Embed}_\gamma(\mu)$ and $h \in \text{Embed}_\phi(\mu)$ each have average 0, then*

$$\mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)] = \mathbb{E}_{(x,y,z) \sim \mu} [f(x)h(z)] = \mathbb{E}_{(x,y,z) \sim \mu} [g(y)h(z)] = 0.$$

Proof. For the first bullet, it suffices to prove this for any f which is a basis element of $\text{Embed}_\sigma(\mu)$. Thus, let $\chi \in \hat{H}$ and consider $f(x) = \chi(\sigma(x))$. Then $f(x) = \chi(-\gamma(y) - \phi(z))$, so

$$\begin{aligned} \mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h(z)] &= \mathbb{E}_{(x,y,z) \sim \mu} [\chi(-\gamma(y))g(y)\chi(-\phi(z))h(z)] \\ &= \mathbb{E}_{(x,y,z) \sim \mu} [\chi(-\gamma(y))g(y)] \mathbb{E}_{(x,y,z) \sim \mu} [\chi(-\phi(z))h(z)], \end{aligned}$$

where in the last transition we used the fact that $\mu_{y,z}$ is uniform. As g is perpendicular to all embedding functions in y we get that $\mathbb{E}_{(x,y,z) \sim \mu} [\chi(-\gamma(y))g(y)] = 0$.

The second bullet is identical.

For the third bullet, we argue that $\mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)] = 0$ and the other two are identical. Note that f and g are functions only of $\sigma(x)$ and $\gamma(y)$. As $\sigma(x)$ and $\gamma(y)$ are independent under $(x, y, z) \sim \mu$, we get that the values $f(x)$ and $g(y)$ are independent and so $\mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)] = \mathbb{E}_{(x,y,z) \sim \mu} [f(x)]\mathbb{E}_{(x,y,z) \sim \mu} [g(y)] = 0$. \square

8.2 The Singular-Value Decomposition

Next, our argument requires an appropriate singular value decomposition with respect to our notions of embedding degrees and non-embedding degrees, and we have the following lemma. Let I, J be a partition of $[n]$ where $|I| = n - 1$ and $|J| = 1$.

Claim 8.4. *Let μ be as in Theorem 7.23. If $g: \Gamma^n \rightarrow \mathbb{C}$ is a completely embedding homogenous, non-embedding homogenous and $\|g\|_2 = 1$, then we may write*

$$g(y) = \sum_{\chi \in \hat{H}} \kappa_\chi g_\chi(y_I) g'_\chi(y_J) + \sum_{r=1}^s \kappa_r g_r(y_I) g'_r(y_J),$$

where letting $P = \{1 \leq r \leq s \mid \kappa_r \neq 0\} \cup \{\chi \in \hat{H} \mid \kappa_\chi \neq 0\}$, we have $\kappa_r \geq 0$ for $r \in P$ and:

1. For $\chi \in \hat{H}$, $g'_\chi \in \text{span}(\{\chi \circ \gamma\})$ and g_χ is completely homogenous embedding function and non-embedding homogenous of degree d .
2. For $1 \leq r \leq s$, g_r is completely homogenous embedding function and non-embedding homogenous of degree $d - 1$. Also, g'_r is in $\text{Embed}_\gamma(\mu)^\perp$.
3. The set $\{g_r: \Gamma^I \rightarrow \mathbb{C}\}_{r \in P}$ is orthonormal.
4. The set $\{g'_r: \Gamma^J \rightarrow \mathbb{C}\}_{r \in P}$ is orthonormal.
5. $\sum_{r \in P} \kappa_r^2 = 1$.

Proof. The proof is deferred to Section C. □

We next state an SVD decomposition statement that addresses the function F . The idea is the same as in Claim 8.4, and the only difference is that we have an additional notion of effective non-embedding degree.

Claim 8.5. *Suppose $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{C}$ is a function which is constant on connected components, is completely embedding homogenous, is non-embedding homogenous of degree d , and has effective degree at least d' . If $\|F\|_2 = 1$, then we may write*

$$F(y, z) = \sum_{\chi \in \hat{H}} \psi_\chi F_\chi(y_I, z_I) F'_\chi(y_J, z_J) + \sum_{t=1}^s \psi_t F_t(y_I, z_I) F'_t(y_J, z_J),$$

and $P = \{1 \leq t \leq s \mid \psi_t \neq 0\} \cup \{\chi \in \hat{H} \mid \psi_\chi \neq 0\}$, where $\psi_t \geq 0$ for $t \in P$ as well as:

1. For $\chi \in \hat{H}$, $F'_\chi \in \text{span}(\{W \circ \chi \circ \sigma\})$ and $F_\chi: \Gamma^I \times \Phi^I \rightarrow \mathbb{C}$ is a completely embedding homogenous, non-embedding homogenous of degree d and effective non-embedding degree at least d' .
2. For $1 \leq t \leq s$, F_t is completely embedding homogenous, non-embedding homogenous of degree $d - 1$ and the effective non-embedding degree is at least $d' - 1$. Also, F'_t is orthogonal to all embedding functions.
3. The functions F_t and F'_t are constant on connected components for all t .
4. The set $\{F_t: \Gamma^I \times \Phi^I \rightarrow \mathbb{C}\}_{t \in P}$ is orthonormal.
5. The set $\{F'_t: \Gamma^J \times \Phi^J \rightarrow \mathbb{C}\}_{t \in P}$ is orthonormal.
6. $\sum_{t \in P} \psi_t^2 = 1$.

Proof. The proof is deferred to Section C. □

We will need the following claim, stating a connection between the coefficients ψ_t in the SVD decomposition and our notions of non-embedding influence.

Claim 8.6. Let $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{C}$ be a function which is constant on connected components. Suppose that F is completely embedding homogenous, non-embedding homogenous of degree d , the effective degree is at least d' , and $\|F\|_2 = 1$. Write

$$F(y, z) = \sum_{\chi \in \hat{H}} \psi_\chi F_\chi(y_I, z_I) F'_\chi(y_J, z_J) + \sum_{t \in T} \psi_t F_t(y_I, z_I) F'_t(y_J, z_J).$$

as in Claim 8.5. Then if j is the unique variable in the set J in the partition $[n] = I \cup J$, then

$$\sum_{t \in T} \psi_t^2 = \frac{1}{2} I_{j, \text{non-embed}}[F].$$

Proof. Write $F = Wf$ for some $f: \Sigma^n \rightarrow \mathbb{C}$ so that $I_{j, \text{non-embed}}[F] = I_{j, \text{non-embed}}[f]$, and note that from the decomposition of F we get an analogous decomposition of f where $F_t = Wf_t$, $F'_t = Wf'_t$ for all t .

Consider $I_{j, \text{non-embed}}[F]$; sampling $x \sim \mu_x^{n-1}$ and $a, b \sim \mu_x$ conditioned on $\sigma(a) = \sigma(b)$ we get that

$$I_{j, \text{non-embed}}[f] = \sum_{t \in T \cup \hat{H}} \kappa_t^2 \mathbb{E}_{a, b \sim \mu_x} \left[(f'_t(a) - f'_t(b)) \overline{(f'_t(a) - f'_t(b))} \mid \sigma(a) = \sigma(b) \right].$$

For $t \in \hat{H}$ we have that $f'_t(a) = f'_t(b)$ and so the expectation is 0. For $t \in T$, computing the expectation by expanding it we get it is equal to

$$2\|f'_t\|_2^2 - 2\mathbb{E}_{a, b \sim \mu} \left[f'_t(a) \overline{f'_t(b)} \mid \sigma(a) = \sigma(b) \right] = 2 - 2\mathbb{E}_{a, b \sim \mu} \left[f'_t(a) \overline{f'_t(b)} \mid \sigma(a) = \sigma(b) \right],$$

and we argue that last expectation is 0. Indeed, let \mathcal{D} be the distribution over the group H of $\sigma(a)$ where we sample $a \sim \mu_x$, we have that the expectation is equal to

$$\mathbb{E}_{w \sim \mathcal{D}} \left[\left| \mathbb{E}_a [f'_t(a) \mid \sigma(a) = w] \right|^2 \right] = \mathbb{E}_{w \sim \mathcal{D}} \left[\frac{|\langle f'_t, 1_{\sigma(\cdot)=w} \rangle|^2}{|\mathbb{E}_a [1_{\sigma(a)=w}]|^2} \right] = \mathbb{E}_{w \sim \mathcal{D}} \left[\frac{0}{|\mathbb{E}_a [1_{\sigma(a)=w}]|^2} \right] = 0,$$

where we used the fact that f'_t is orthogonal to the embedding function $1_{\sigma(\cdot)=w}$. \square

8.3 Proof of Lemma 8.1

In this section, we give the formal proof of Lemma 8.1. Let F, g and h be functions of 2-norm equal to 1 achieving the value $\beta_{n,d,d'}$. Using the singular-value decomposition we may write F, g and h as a sum of functions satisfying some orthogonality properties (see Claim 8.4 and Claim 8.5 for precise statements):

$$\begin{aligned} F(y, z) &= \sum_{\chi \in \hat{H}} \psi_\chi F_\chi(y_I, z_I) F'_\chi(y_J, z_J) + \sum_{t \in T} \psi_t F_t(y_I, z_I) F'_t(y_J, z_J), \\ g(y) &= \sum_{\pi \in \hat{H}} \kappa_\pi g_\pi(y_I) g'_\pi(y_J) + \sum_{r \in R} \kappa_r g_r(y_I) g'_r(y_J), \\ h(z) &= \sum_{\lambda \in \hat{H}} \rho_\lambda h_\lambda(z_I) h'_\lambda(z_J) + \sum_{s \in S} \rho_s h_s(y_I) h'_s(y_J). \end{aligned} \tag{12}$$

Here, each one of the sets $\{F_t\}_{t \in T} \cup \{F_\chi\}_{\chi \in \hat{H}}, \{F'_t\}_{t \in T} \cup \{F'_\chi\}_{\chi \in \hat{H}}, \{g_r\}_{r \in R} \cup \{g_\pi\}_{\pi \in \hat{H}}, \{g'_r\}_{r \in R} \cup \{g'_\pi\}_{\pi \in \hat{H}}, \{h_s\}_{s \in S} \cup \{h_\lambda\}_{\lambda \in \hat{H}}$ and $\{h'_s\}_{s \in S} \cup \{h'_\lambda\}_{\lambda \in \hat{H}}$ is orthonormal. By multiplying each one of the coefficients ψ_t, κ_r and ρ_s by an appropriate complex number of absolute value 1, we will assume henceforth that $F'_\chi = W \circ \chi \circ \sigma, g'_\pi = \pi \circ \gamma$ and $h'_\lambda = \lambda \circ \phi$. Thus, we have that

$$\mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [F(y,z)g(y)h(z)] = \sum_{\substack{r \in R \cup \hat{H} \\ s \in S \cup \hat{H} \\ t \in T \cup \hat{H}}} \psi_t \kappa_r \rho_s \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [F_t(y,z)g_r(y)h_s(z)] \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [F'_t(y,z)g'_r(y)h'_s(z)].$$

It will be convenient for us to denote

$$\widehat{F}_t(r,s) = \mathbb{E}_{(x,y,z) \sim \mu^{\otimes(n-1)}} [F_t(y,z)g_r(y)h_s(z)], \quad \widehat{F}'_t(r,s) = \mathbb{E}_{(x,y,z) \sim \mu} [F'_t(y,z)g'_r(y)h'_s(z)].$$

This is justified because $(\overline{g_r h_s})_{r \in R, s \in S}$ 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 $\overline{g_r h_s}$ in the representation of F_t . Thus, we get that

$$\mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [F(y,z)g(y)h(z)] = \sum_{\substack{r \in R \cup \hat{H} \\ s \in S \cup \hat{H} \\ t \in T \cup \hat{H}}} \psi_t \kappa_r \rho_s \widehat{F}_t(r,s) \widehat{F}'_t(r,s), \quad (13)$$

8.3.1 The parameters and choosing the partition

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

$$0 \ll \varepsilon \ll \eta \ll \tau \ll c \ll m^{-1}, \alpha \leq 1. \quad (14)$$

We need to choose the partition I, J so that the mass of F on the non-embedding components is small, namely

$$\sum_{t \in T} |\psi_t|^2 \leq \eta, \quad (15)$$

and next show that as long as n is much larger than d this is possible. Indeed, if $n \geq \frac{10}{\eta} d$, then choosing the partition randomly we have by Claim 8.6 and Fact 4.12 that

$$\mathbb{E}_{I,J} \left[\sum_{t \in T} |\psi_t|^2 \right] = \frac{1}{2} \mathbb{E}_{J=\{j\}} [I_{j, \text{non-embed}}[F]] \leq \frac{d}{n} \leq \frac{\eta}{10},$$

so by Markov's inequality we get that

$$\Pr_{I,J} [(15) \text{ fails}] \leq \frac{1}{10} < 1.$$

We may thus find a partition $[n] = I \cup J$ with $|J| = 1$ such that (15) holds. We fix this partition henceforth.

8.3.2 The main inductive argument

We are going to need to split the expression from (13) into various sums and towards this end we start with a few observations that asserts that various terms in (13) are 0.

Observation 8.7. *With the setup above:*

1. Consider $\chi, \pi, \lambda \in \hat{H}$, and note that $\widehat{F}'_\chi(\pi, \lambda) = 1$ if $\chi = \pi = \lambda$, and $\widehat{F}'_\chi(\pi, \lambda) = 0$ otherwise. Indeed,

$$\widehat{F}'_\chi(\pi, \lambda) = \mathbb{E}_{(x,y,z) \sim \mu} [F'_\chi(y, z)g'_\pi(y)h'_\phi(z)] = \mathbb{E}_{(x,y,z) \sim \mu} [\chi(\sigma(x))\pi(\gamma(y))\lambda(\phi(z))].$$

As $(\sigma(x), \gamma(y), \phi(z))$ is distributed uniformly among all $(a, b, c) \in H^3$ such that $a + b + c = 0$, we get that expectation is 0 if not all of χ, π, λ are equal, and 1 if they are equal.

2. For $\chi \in \hat{H}$ and r, s such that either $r \in R$ or $s \in S$ we have that $\widehat{F}'_\chi(r, s) = 0$. This is thanks to Claim 8.3.

3. For $t \in T$, we have that $\widehat{F}'_t(\pi, \lambda) = 0$ if $\pi = \lambda$. Indeed, write $F'_t = W \circ f'_t$, then

$$\begin{aligned} \widehat{F}'_\chi(\pi, \lambda) &= \mathbb{E}_{(x,y,z) \sim \mu} [f'_t(x)\pi(\gamma(y))\lambda(\phi(z))] = \mathbb{E}_{(x,y,z) \sim \mu} [f'_t(x)\pi(\gamma(y) + \phi(z))] \\ &= \mathbb{E}_{(x,y,z) \sim \mu} [f'_t(x)\pi(-\sigma(x))] \\ &= \langle f'_t, \pi \circ \sigma \rangle, \end{aligned}$$

which is 0 as f'_t is orthogonal to embedding functions.

With these observations in hand, we expand out the sum in (13) and drop off terms which are 0 to get that

$$\begin{aligned} \beta_{n,d,d'} &= \sum_{\chi \in \hat{H}} \psi_\chi \kappa_\chi \rho_\chi \widehat{F}'_\chi(\chi, \chi) + \sum_{t \in T, \pi, \lambda \in \hat{H}, \pi \neq \lambda} \psi_t \kappa_\pi \rho_\lambda \widehat{F}'_t(\pi, \lambda) \widehat{F}'_t(\pi, \lambda) \\ &+ \sum_{t \in T, r \in R, \lambda \in \hat{H}} \psi_t \kappa_r \rho_\lambda \widehat{F}'_t(r, \lambda) \widehat{F}'_t(r, \lambda) + \sum_{t \in T, \pi \in \hat{H}, s \in S} \psi_t \kappa_\pi \rho_s \widehat{F}'_t(\pi, s) \widehat{F}'_t(\pi, s) \\ &+ \sum_{t \in T, r \in R, s \in S} \psi_t \kappa_r \rho_s \widehat{F}'_t(r, s) \widehat{F}'_t(r, s). \end{aligned} \tag{16}$$

8.3.3 The Embedding Masses of F , g and h Must Align

We start off with handling the case where the mass of F , g and h is mis-aligned on the embedding part of the decomposition. By the choice of I, J we know that most of the mass of F lies on components corresponding to embedding functions, and as we noted above among these only the diagonal terms $\chi = \pi = \lambda$ survive. Hence, it makes sense that unless there is a single χ on which most of the embedding mass of F , g and h lies, we will already be able to use (16) to argue that $\beta_{n,d,d'}$ must be considerably smaller than $\beta_{n-1,d-1,d'-1}$. This is the content of the following lemma.

Lemma 8.8. *If $\sum_{\chi \in \hat{H}} |\psi_\chi \kappa_\chi \rho_\chi| \leq 1 - \tau$, then $\beta_{n,d,d'} \leq (1 - \varepsilon) \max(\beta_{n-1,d-1,d'-1}, \beta_{n-1,d,d'})$ (and hence we are done).*

Proof. Using (16) we get

$$\beta_{n,d,d'} \leq (1 - \tau)\beta_{n-1,d,d'} + O_m(\sqrt{\eta}\beta_{n-1,d-1,d'-1}),$$

where we used the fact that for $t \in T$ we have that $|\psi_t| \leq \sqrt{\sum_{t \in T} |\psi_t|^2} \leq \sqrt{\eta}$, as well as

$$\left| \widehat{F}_\chi(\chi, \chi) \right| \leq \beta_{n-1,d,d'}, \quad \left| \widehat{F}_t(r, s) \right| \leq \beta_{n-1,d-1,d'-1}$$

for all $t \in T, r \in R \cup \widehat{H}, s \in S \cup \widehat{H}$. We get that

$$\frac{\beta_{n,d,d'}}{\max(\beta_{n-1,d-1,d'-1}, \beta_{n-1,d,d'})} \leq 1 - \tau + O_m(\sqrt{\eta}) \leq 1 - \tau/2 \leq 1 - \varepsilon,$$

and we are done. \square

By Lemma 8.8, we assume henceforth that $\sum_{\chi \in \widehat{H}} |\psi_\chi \kappa_\chi \rho_\chi| > 1 - \tau$, and we note that

$$\begin{aligned} \sum_{\chi \in \widehat{H}} |\psi_\chi \kappa_\chi \rho_\chi| &\leq \max_{\chi} |\psi_\chi \kappa_\chi \rho_\chi|^{1/3} \sum_{\chi \in \widehat{H}} |\psi_\chi \kappa_\chi \rho_\chi|^{2/3} \\ &\leq \max_{\chi} |\psi_\chi \kappa_\chi \rho_\chi|^{1/3} \left(\sum_{\chi \in \widehat{H}} |\psi_\chi|^2 \right)^{1/3} \left(\sum_{\chi \in \widehat{H}} |\kappa_\chi|^2 \right)^{1/3} \left(\sum_{\chi \in \widehat{H}} |\rho_\chi|^2 \right)^{1/3} \\ &\leq \max_{\chi} |\psi_\chi \kappa_\chi \rho_\chi|^{1/3}, \end{aligned}$$

where we used Hölder's inequality. It follows that there is χ such that $|\psi_\chi \kappa_\chi \rho_\chi|^{1/3} \geq 1 - \tau$, and we call it χ^* . In particular it follows that

$$|\psi_{\chi^*}|, |\kappa_{\chi^*}|, |\rho_{\chi^*}| \geq 1 - 3\tau. \quad (17)$$

8.3.4 Upper Bounding Terms in (16)

We now proceed to upper bounding terms in (16). The first sum is handled by the following lemma:

Lemma 8.9. *It holds that:*

$$\left| \sum_{\chi \in \widehat{H}} \psi_\chi \kappa_\chi \rho_\chi \widehat{F}_\chi(\chi, \chi) \right| \leq \beta_{n-1,d,d'} \sqrt{\sum_{\chi \in \widehat{H}} |\psi_\chi|^2} \sqrt{\sum_{\chi \in \widehat{H}} |\kappa_\chi|^2 |\rho_\chi|^2}.$$

Proof. By definition $\left| \widehat{F}_\chi(\chi, \chi) \right| \leq \beta_{n-1,d,d'}$ for all $\chi \in \widehat{H}$, and the result follows from the triangle inequality and Cauchy-Schwarz. \square

The next lemma handles summands from the second, third and fourth sums in which χ^* does not appear.

Lemma 8.10. *It holds that:*

$$\left| \sum_{\substack{t \in T, r \in R \cup \hat{H}, s \in S \cup \hat{H} \\ r, s \neq \chi^*}} \psi_t \kappa_r \rho_s \widehat{F}_t(r, s) \widehat{F}'_t(r, s) \right| \lesssim_m \beta_{n-1, d-1, d'-1} \sqrt{\sum_{t \in T} |\psi_t|^2} \sqrt{1 - |\kappa_{\chi^*}|^2} \sqrt{1 - |\rho_{\chi^*}|^2}.$$

Proof. For all $t \in T$ it holds that $|\psi_t| \leq \sqrt{\sum_{t \in T} |\psi_t|^2}$, for $r \neq \chi^*$ we have $|\kappa_r| \leq \sqrt{\sum_{r' \neq \chi^*} |\kappa_{r'}|^2} = \sqrt{1 - |\kappa_{\chi^*}|^2}$, and similarly for $s \neq \chi^*$ we have $|\rho_s| \leq \sqrt{1 - |\rho_{\chi^*}|^2}$. Also, for all such t, r, s we have that $|\widehat{F}_t(r, s)| \leq \beta_{n-1, d-1, d'-1}$ and $|\widehat{F}'_t(r, s)| \leq 1$, and as the number of summands is at most m^3 the bound follows. \square

The following lemma handles the rest of the summands from the second, third and fourth sums, namely those in which χ^* does appear (we recall that if both r and s are equal to χ^* , then such summands are 0, hence we are dealing with summands in which exactly one of r and s is equal to χ^*).

Lemma 8.11. *It holds that:*

$$\left| \sum_{\substack{t \in T \\ r \in R \cup \hat{H}, s \in S \cup \hat{H} \\ r = \chi^* \text{ or } s = \chi^*}} \psi_t \kappa_r \rho_s \widehat{F}_t(r, s) \widehat{F}'_t(r, s) \right| \leq (1 - c) \beta_{n-1, d-1, d'-1} \sqrt{\sum_{t \in T} |\psi_t|^2} \sqrt{|\kappa_{\chi^*}|^2 \sum_{s \neq \chi^*} |\rho_s|^2 + |\rho_{\chi^*}|^2 \sum_{r \neq \chi^*} |\kappa_r|^2}.$$

Proof. We write the sum in discussion as

$$\sum_{t \in T, r \neq \chi^*} \psi_t \kappa_r \rho_{\chi^*} \widehat{F}_t(r, \chi^*) \widehat{F}'_t(r, \chi^*) + \sum_{t \in T, s \neq \chi^*} \psi_t \kappa_{\chi^*} \rho_s \widehat{F}_t(\chi^*, s) \widehat{F}'_t(\chi^*, s). \quad (18)$$

Fix $t \in T$, write $F'_t = W f'_t$ and consider the corresponding terms. We note that

$$\widehat{F}'_t(r, \chi^*) = \mathbb{E}_{(x, y, z) \sim \mu} [f'_t(x) g'_r(y) \chi^*(\phi(z))] = \mathbb{E}_{(x, y, z) \sim \mu} [f'_t(x) \chi^*(-\sigma(x)) g'_r(y) \chi^*(-\gamma(y))]$$

and similarly

$$\widehat{F}'_t(\chi^*, s) = \mathbb{E}_{(x, y, z) \sim \mu} [f'_t(x) \chi^*(\phi(y)) h'_s(z)] = \mathbb{E}_{(x, y, z) \sim \mu} [f'_t(x) \chi^*(-\sigma(x)) h'_s(z) \chi^*(-\phi(z))].$$

Thus, we define

$$\tilde{g}(y) = \sum_{r \neq \chi^*} \widehat{F}_t(r, \chi^*) \kappa_r \rho_{\chi^*} g'_r(y) \chi^*(-\gamma(y)), \quad \tilde{h}(z) = \sum_{s \neq \chi^*} \widehat{F}_t(\chi^*, s) \kappa_{\chi^*} \rho_s h'_s(z) \chi^*(-\phi(z)),$$

and get that the contribution of t to (18) is equal to

$$\mathbb{E}_{(x,y,z) \sim \mu} \left[f'_t(x) \chi^*(-\sigma(x)) (\tilde{g}(y) + \tilde{h}(z)) \right].$$

We note that \tilde{g}, \tilde{h} are orthogonal, and

$$\|\tilde{g}\|_2^2 = \sum_{r, r' \neq \chi^*} \widehat{F}_t(r, \chi^*) \overline{\widehat{F}_t(r', \chi^*)} |\rho_{\chi^*}|^2 \kappa_r \overline{\kappa_{r'}} \langle g'_r, g'_{r'} \rangle = |\rho_{\chi^*}|^2 \sum_{r \neq \chi^*} \left| \widehat{F}_t(r, \chi^*) \right|^2 |\kappa_r|^2,$$

and similarly for \tilde{h} . Also, $f'_t(x) \chi^*(-\sigma(x))$ is perpendicular to all embedding functions. Applying Claim 8.2 we conclude that

$$\begin{aligned} & \left| \mathbb{E}_{(x,y,z) \sim \mu} \left[f'_t(x) \chi^*(-\sigma(x)) (\tilde{g}(y) + \tilde{h}(z)) \right] \right| \\ & \leq (1-c) \sqrt{|\kappa_{\chi^*}|^2 \sum_{s \neq \chi^*} |\rho_s|^2 \left| \widehat{F}_t(\chi^*, s) \right|^2 + |\rho_{\chi^*}|^2 \sum_{r \neq \chi^*} |\kappa_r|^2 \left| \widehat{F}_t(r, \chi^*) \right|^2}. \end{aligned}$$

Plugging this into (18) gives that

$$\begin{aligned} |(18)| & \leq (1-c) \sum_{t \in T} |\psi_t| \sqrt{|\kappa_{\chi^*}|^2 \sum_{s \neq \chi^*} |\rho_s|^2 \left| \widehat{F}_t(\chi^*, s) \right|^2 + |\rho_{\chi^*}|^2 \sum_{r \neq \chi^*} |\kappa_r|^2 \left| \widehat{F}_t(r, \chi^*) \right|^2} \\ & \leq (1-c) \sqrt{\sum_{t \in T} |\psi_t|^2} \sqrt{|\kappa_{\chi^*}|^2 \sum_{s \neq \chi^*} |\rho_s|^2 \sum_{t \in T} \left| \widehat{F}_t(\chi^*, s) \right|^2 + |\rho_{\chi^*}|^2 \sum_{r \neq \chi^*} |\kappa_r|^2 \sum_{t \in T} \left| \widehat{F}_t(r, \chi^*) \right|^2}, \quad (19) \end{aligned}$$

where we used Cauchy-Schwarz. For fixed $s \in S \cup \hat{H}$, note that letting $\tilde{F} = \frac{\sum_{t \in T} \widehat{F}_t(\chi^*, s) F_t}{\sqrt{\sum_{t \in T} \left| \widehat{F}_t(\chi^*, s) \right|^2}}$ we have that \tilde{F} has 2-norm equal to 1, is completely embedding homogenous and non-embedding homogenous of degree $d-1$. Thus,

$$\sum_{t \in T} \left| \widehat{F}_t(\chi^*, s) \right|^2 = \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n-1}} \left[\tilde{F}(y, z) g_{\chi^*}(y) h_s(z) \right] \right|^2 \leq \beta_{n-1, d-1, d'-1}^2.$$

Similarly, $\sum_{t \in T} \left| \widehat{F}_t(r, \chi^*) \right|^2 \leq \beta_{n-1, d-1, d'-1}^2$ for all $r \in R \cup \hat{H}$, and plugging this into (19) yields that

$$|(18)| \leq (1-c) \beta_{n-1, d-1, d'-1} \sqrt{\sum_{t \in T} |\psi_t|^2} \sqrt{|\kappa_{\chi^*}|^2 \sum_{s \neq \chi^*} |\rho_s|^2 + |\rho_{\chi^*}|^2 \sum_{r \neq \chi^*} |\kappa_r|^2}.$$

□

Plugging Lemmas 8.9 8.10, 8.11 into (16) we get that

$$\begin{aligned}\beta_{n,d,d'} &\leq \beta_{n-1,d,d'} \sqrt{\sum_{\chi \in \hat{H}} |\psi_\chi|^2} \sqrt{\sum_{\chi \in \hat{H}} |\kappa_\chi|^2 |\rho_\chi|^2} \\ &\quad + (1-c)\beta_{n-1,d-1,d'-1} \sqrt{\sum_{t \in T} |\psi_t|^2} \sqrt{|\kappa_{\chi^*}|^2 \sum_{s \neq \chi^*} |\rho_s|^2 + |\rho_{\chi^*}|^2 \sum_{r \neq \chi^*} |\kappa_r|^2} \\ &\quad + C_m \beta_{n-1,d,d'} \sqrt{\sum_{t \in T} |\psi_t|^2} \sqrt{1 - |\kappa_{\chi^*}|^2} \sqrt{1 - |\rho_{\chi^*}|^2}.\end{aligned}$$

Denoting $\beta'' = \max((1-c/2)\beta_{n-1,d-1,d'-1}, \beta_{n-1,d,d'})$ we get by Cauchy-Schwarz that

$$\begin{aligned}\beta_{n,d,d'} &\leq \beta'' \sqrt{\sum_{\chi \in \hat{H}} |\psi_\chi|^2 + \sum_{t \in T} |\psi_t|^2} \sqrt{\sum_{\chi \in \hat{H}} |\kappa_\chi|^2 |\rho_\chi|^2 + (1-c/2) |\kappa_{\chi^*}|^2 \sum_{s \neq \chi^*} |\rho_s|^2 + (1-c/2) |\rho_{\chi^*}|^2 \sum_{r \neq \chi^*} |\kappa_r|^2} \\ &\quad + C_m \beta'' \sqrt{\sum_{t \in T} |\psi_t|^2} \sqrt{1 - |\kappa_{\chi^*}|^2} \sqrt{1 - |\rho_{\chi^*}|^2}.\end{aligned}$$

We have that

$$\sum_{\chi \in \hat{H}} |\kappa_\chi|^2 |\rho_\chi|^2 + |\kappa_{\chi^*}|^2 \sum_{s \neq \chi^*} |\rho_s|^2 + |\rho_{\chi^*}|^2 \sum_{r \neq \chi^*} |\kappa_r|^2 \leq \sum_{r,s} |\kappa_r|^2 |\rho_s|^2 = 1,$$

so we conclude that

$$\begin{aligned}\beta_{n,d,d'} &\leq \beta'' \sqrt{1 - \frac{c}{2} \left(|\kappa_{\chi^*}|^2 \sum_{s \neq \chi^*} |\rho_s|^2 + |\rho_{\chi^*}|^2 \sum_{r \neq \chi^*} |\kappa_r|^2 \right)} \\ &\quad + C_m \beta'' \sqrt{\sum_{t \in T} |\psi_t|^2} \sqrt{1 - |\kappa_{\chi^*}|^2} \sqrt{1 - |\rho_{\chi^*}|^2}.\end{aligned}$$

Using $\sum_{s \neq \chi^*} |\rho_s|^2 = 1 - |\rho_{\chi^*}|^2$ and $\sum_{r \neq \chi^*} |\kappa_r|^2 = 1 - |\kappa_{\chi^*}|^2$ we get

$$\begin{aligned}\beta_{n,d,d'} &\leq \beta'' \sqrt{1 - \frac{c}{4} \left(|\kappa_{\chi^*}|^2 + |\rho_{\chi^*}|^2 - 2 |\kappa_{\chi^*}|^2 |\rho_{\chi^*}|^2 \right)} \\ &\quad + C_m \beta'' \sqrt{\sum_{t \in T} |\psi_t|^2} \sqrt{1 - |\kappa_{\chi^*}|^2} \sqrt{1 - |\rho_{\chi^*}|^2}.\end{aligned}$$

As $\sqrt{1-a} \leq 1 - a/2$ and $\sqrt{ab} \leq \frac{1}{2}(a+b)$ for non-negative a, b and $\sum_{t \in T} |\psi_t|^2 \leq \eta$, we get that

$$\frac{\beta_{n,d,d'}}{\beta''} \leq \left(1 - \frac{c}{8} \left(|\kappa_{\chi^*}|^2 + |\rho_{\chi^*}|^2 - 2 |\kappa_{\chi^*}|^2 |\rho_{\chi^*}|^2 \right)\right) + C_m \sqrt{\eta} \left(2 - |\kappa_{\chi^*}|^2 - |\rho_{\chi^*}|^2\right).$$

We write the right hand side as $1 - \frac{c}{8} P(|\kappa_{\chi^*}|^2, |\rho_{\chi^*}|^2)$, where $P: [0, 1]^2 \rightarrow \mathbb{R}$ is defined as

$$P(a, b) = a + b - 2ab - \frac{8C_m}{c} \sqrt{\eta} (2 - a - b).$$

Claim 8.12. Given (14), we have that $P(a, b) \geq 0$ for all $a, b \geq 1 - 6\tau$.

Proof. Looking at partial derivatives, we see that $\frac{\partial P}{\partial a} = 1 - 2b + \frac{8C_m}{c}\sqrt{\eta}$, and similarly $\frac{\partial P}{\partial b} = 1 - 2a + \frac{8C_m}{c}\sqrt{\eta}$. It follows that P is decreasing in both variables in the range that $a, b \geq 1/2 + \frac{4C_m}{c}\sqrt{\eta}$, hence for all $a, b \geq 1 - 6\tau$ we have that $P(a, b) \geq P(1, 1) = 0$. \square

Combining Claim 8.12 and (17), it follows that $\frac{\beta_{n,d,d'}}{\beta''} \leq 1 - \frac{c}{8}P(|\kappa_{\chi^*}|^2, |\rho_{\chi^*}|^2) \leq 1$, concluding the proof. \square

8.4 Conclusion of the Reduction to Near Linear Degree: Iterating Lemma 8.1

To conclude this section, we iterate Lemma 8.1 and get that we can either get an exponential upper bound on our parameter β , or else we can reduce to the case the number of variables n is proportional to the non-embedding degree d , and furthermore the non-embedding degree and effective non-embedding degree stay roughly the same. Formally:

Lemma 8.13. For all $m \in \mathbb{N}$ and $\alpha > 0$ there are $L > 0$ and $\varepsilon > 0$ such that the following holds for all $\xi > 0$. Let μ be as in Theorem 7.23, and let $n \geq d \geq d'$ be integers such that $d' \geq d^{1-\xi}$. Then either:

1. $\beta_{n,d,d'}[\mu] \leq (1 - \varepsilon)^{d'/2}$.
2. Else, there are integers $e \geq d/2$, $e' \geq d'/2$ and $n' \leq 2Le$. such that $\beta_{n,d,d'}[\mu] \leq \beta_{n',e,e'}$.

Proof. We apply Lemma 8.1 so long as possible. When we can no longer do that, let s be the number of times it was the case that the upper bound was $(1 - \varepsilon)\beta_{n-1,d-1,d'-1}$. If $s \geq d'/2$, then we clearly got a factor of $1 - \varepsilon$ s times, and using the trivial bound of 1 on all β 's it follows that the first item holds.

Else, $s \leq d'/2$ and hence the non-embedding degree parameter e in the end has dropped by at most s , so $e \geq d - s \geq d/2$. Also, the effective non-embedding degree lower bound parameter e' also drops by at most s , hence we get that $e' \geq d' - s \geq d'/2$. Since we can no longer apply Lemma 8.1 it follows that the number of variables we reached to is $n' \leq L \cdot d \leq 2Le$, and also that

$$\beta_{n,d,d'}[\mu] \leq (1 - \varepsilon)^s \beta_{n',e,e'}[\mu] \leq \beta_{n',e,e'}[\mu]. \quad \square$$

9 Proof of Theorem 7.23: Proof for Near Linear Degrees

In this section we start off where Section 8 ended, and handle the case that the number of variables n is as in Lemma 8.13. In that case, we present a different inductive argument, similar to the one in [5, Section 6], that establishes an exponential upper bound on β . This is the part in our argument in which we use the relaxed base case.

9.1 The Parameter δ

In contrast to Section 8, the argument we present herein will not be able to preserve homogeneity and instead will reduce degrees in a controlled manner. To facilitate that, we define a modification of the class $\mathcal{F}_{n,d,d'}$ as follows:

Definition 9.1. We define the class $\mathcal{F}_{n,d'}$ as the class of functions $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{C}$ that are constant on connected components and have effective non-embedding degree at least d' .

We now define a variant of the parameter β that is central to this section.

Definition 9.2. For integers $n \geq d'$, finite alphabets Σ, Γ, Φ and a distribution μ over $\Sigma \times \Gamma \times \Phi$ as in Theorem 6.7, we define

$$\delta_{n,d'}[\mu] = \sup_{\substack{F \in \mathcal{F}_{n,d'} \\ g: \Gamma^n \rightarrow \mathbb{C} \\ h: \Phi^n \rightarrow \mathbb{C}}} \frac{|\mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [F(y,z)g(y)h(z)]|}{\|F\|_2 \|g\|_2 \|h\|_2}.$$

When the distribution μ is clear from context, we will omit it from the notation and simply write $\delta_{n,d'}$. With the parameter $\delta_{n,d'}$ in hand, we can now state the main inductive statement of this section.

Lemma 9.3. For all $m \in \mathbb{N}$ and $\alpha > 0$ there is $C > 0$ such that the following holds. For integers $n \geq d'$ such that $n \leq d' A$ and $A \geq 2$, finite alphabets Σ, Γ, Φ and a distribution μ over $\Sigma \times \Gamma \times \Phi$ as in Theorem 6.7, we have that

$$\delta_{n,d'}[\mu] \leq \left(1 - \frac{1}{A^C}\right) \delta_{n-1,d'-1}[\mu].$$

The rest of this section is devoted to the proof of Lemma 9.3. We first establish a singular-value decomposition stated similar to Claim 8.5 but for the case that F is not homogenous, and relate this decomposition with the modest influences similarly to Claim 8.6.

9.2 The SVD Decomposition and Relation to Modest Influences

We are going to need the following two singular-value decomposition statements, and as usual we are going to have a partition I, J of $[n]$ into $|I| = n - 1$ and $|J| = 1$. The first statement addresses the functions g and h :

Claim 9.4. Suppose $g: \Gamma^n \rightarrow \mathbb{C}$ has 2-norm equal to 1. Then we may write

$$g(y) = \sum_{r \in R} \kappa_r g_r(y_I) g'_r(y_J),$$

where each $\kappa_r \geq 0$ and

1. For $r \in R$, $F_r: \Gamma^I \rightarrow \mathbb{C}$ is an orthonormal set of functions.
2. For $r \in R$, $F'_r: \Gamma^J \rightarrow \mathbb{C}$ is an orthonormal set of functions.
3. $\sum_{r \in R} \kappa_r^2 = 1$.

Proof. The proof is similar to the proof of Claim 8.4 and is simpler (no discussion of invariant spaces is needed), and we omit the details. \square

The second singular-value decomposition statement addresses the function F :

Claim 9.5. Suppose $F: \Gamma^n \times \Phi^n \rightarrow \mathbb{C}$ has 2-norm equal to 1, is constant on connected components and has effective non-embedding degree at least d' . Then we may write

$$F(y, z) = \sum_{t \in T} \psi_t F_t(y_I, z_I) F'_t(y_J, z_J),$$

where each $\psi_t \geq 0$ and

1. For $t \in T$, $F_t: \Gamma^I \times \Phi^I \rightarrow \mathbb{C}$ is an orthonormal set of functions.
2. For $t \in T$, $F'_t: \Gamma^J \times \Phi^J \rightarrow \mathbb{C}$ is an orthonormal set of functions.
3. For $t \in T$, F_t has effective non-embedding degree at least $d' - 1$.
4. For $t \in T$, the functions F_t and F'_t are constant on connected components.
5. $\sum_{t \in T} \psi_t^2 = 1$.

Proof. The proof is similar to the proof of Claim 8.5 and is simpler (no discussion of invariant spaces is needed), and we omit the details. \square

The next statement relates the SVD decomposition of F and the variance the F'_t have on Σ_{modest} with the notion of modest influences. For that, we define the variance a function on Σ_{modest} :

Definition 9.6. For a function $f: \Sigma \rightarrow \mathbb{C}$ we denote $\text{var}_{\Sigma_{\text{modest}}}(f) = \mathbb{E}_{a,b \in \Sigma_{\text{modest}}} [|f(a) - f(b)|^2]$.

With this definition, the following claim asserts that if a function has large modest influence, then the corresponding parts F'_t have significant variance on Σ_{modest} . More precisely:

Claim 9.7. If $J = \{j\}$, then letting $f'_t: \Sigma \rightarrow \mathbb{C}$ be such that $F'_t = W f'_t$ we have that

$$\sum_t \psi_t^2 \text{var}_{\Sigma_{\text{modest}}}(f'_t) \gtrsim_{m,\alpha} I_{j,\text{modest}}[F].$$

Proof. Write $F = Wf$ for $f: \Sigma^n \rightarrow \mathbb{C}$, so that from the SVD decomposition of F we get a similar SVD decomposition for f with f_t, f'_t satisfying $F_t = W f_t, F'_t = W f'_t$. Consider $I_{j,\text{modest}}[F] = I_{j,\text{modest}}[f]$, and consider the distribution over (a, b) where $a \sim \mu_x$, and $b \sim \text{Modest } a$. Then by definition

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

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

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

9.3 Proof of Lemma 9.3

We are going to have the following hierarchy of parameters in this section (we recall that A is in Lemma 9.3):

$$\begin{aligned} 0 \ll R_4^{-1} \ll R_3^{-1} \ll R_2^{-1} \ll R_1^{-1} \ll c \ll \alpha, m^{-1} \leq 1, \\ 0 < \eta = \frac{1}{AR_4} < \varepsilon = \frac{1}{AR_3} < \zeta = \frac{1}{AR_2} < \tau = \frac{1}{AR_1} \leq 1. \end{aligned} \quad (20)$$

Fix n and d' as in Lemma 9.3, and let $F \in \mathcal{F}_{n,d'}$ and g, h be functions of 2-norm 1 so that

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

As $I_{\text{modest}}[F] \geq 2d'$ by Fact 7.8, it follows that there is $j \in [n]$ such that

$$I_{j,\text{modest}}[F] \geq \frac{I_{\text{modest}}[F]}{n} \geq \frac{2d'}{n} \geq \frac{1}{A}, \quad (21)$$

and we fix such j henceforth. We choose the partition $I = [n] \setminus \{j\}$ and $J = \{j\}$, and then use the SVD decompositions for g and h from Claim 9.4 and for F as in Claim 9.5 to write

$$g(y) = \sum_{r \in R} \kappa_r g_r(y_I) g'_r(y_J), \quad h(z) = \sum_{s \in S} \rho_s h_s(y_I) h'_s(z_J), \quad F(y, z) = \sum_{t \in T} \psi_t F_t(y_I, z_J) F'_t(y_I, z_J).$$

Thus, using the notations of the previous section we have that

$$\delta_{n,d'} = \sum_{t \in T, r \in R, s \in S} \psi_t \kappa_r \rho_s \widehat{F}_t(r, s) \widehat{F}'_t(r, s). \quad (22)$$

9.3.1 Separating Out Coefficients Bounded Away From 0 and Coefficients Close to 0

The role of the parameters η and τ above will be that none of the coefficients κ_r, ρ_s and ψ_t will be in the interval $[\eta, \tau)$. Formally, we show

Claim 9.8. *We may find parameters as in (20) such that none of the coefficients κ_r, ρ_s and ψ_t are in the interval $[\eta, \tau)$.*

Proof. We start with some parameters as in (20). As long as there is a coefficient in the interval $[\eta, \tau)$, we take a new set of R 's by taking $R'_1 = R_4, R'_4 \gg R'_3 \gg R'_2 \gg R'_1$ and take the collection $\{R'_i\}_{i=1}^4$. Note that as $R'_1 > R_4$ the intervals we consider at each step are disjoint, and therefore after at most $3m$ iterations we will get that none of the coefficients lie in the interval $[\eta, \tau)$, and we are done. \square

We assume henceforth that $\eta \ll \tau$ that none of κ_r, ρ_s, ψ_t are in $[\eta, \tau)$. Define

$$T' = \{t \in T \mid \psi_t \geq \tau\}, \quad R' = \{r \in R \mid \kappa_r \geq \tau\}, \quad S' = \{s \in S \mid \rho_s \geq \tau\}.$$

Using (22) we write

$$\delta_{n,d'} = \underbrace{\sum_{t \in T', r \in R', s \in S'} \psi_t \kappa_r \rho_s \widehat{F}_t(r, s) \widehat{F}'_t(r, s)}_{(I)} + \underbrace{\sum_{\substack{t \in T, r \in R, s \in S \\ t \notin T' \text{ or } r \notin R' \text{ or } s \notin S'}} \psi_t \kappa_r \rho_s \widehat{F}_t(r, s) \widehat{F}'_t(r, s)}_{(II)}. \quad (23)$$

The majority of our effort will go into bounding the contribution of (I) (which we defer to the rest of this section), and we first bound the contribution of (II) by the following claim.

Claim 9.9. *It holds that $|(II)| \lesssim_m \eta \delta_{n-1, d'-1}$.*

Proof. By the triangle inequality, it suffices to upper bound each summand. For each r, s, t in the sum we have that $|\psi_t \kappa_r \rho_s| \leq \eta$, as at least one of them is at most η in absolute value, and the other two are at most 1. Also, $|\widehat{F}_t(r, s)| \leq \delta_{n-1, d'-1}$ and $|\widehat{F}'_t(r, s)| \leq 1$, and the claim follows since the number of terms in the sum is at most m^3 . \square

9.3.2 A Naive Cauchy-Schwarz Bound on (I)

We now present a naive Cauchy-Schwarz based argument showing that $|(I)| \leq \delta_{n-1,d'-1}$ which can be used to recover the trivial bound $\delta_{n,d'} \leq \delta_{n-1,d'-1}$. The upshot of this argument is that by inspecting near equality cases, the argument will allow us to say that unless near equalities in the pursuing applications of Cauchy-Schwarz argument hold, we will successfully have shown that $\delta_{n,d'} \leq (1 - \varepsilon)\delta_{n-1,d'-1}$. Hence, we will be able to assume in the rest of the argument that all of the Cauchy-Schwarz applications were nearly tight. To be more precise, by Cauchy-Schwarz

$$\begin{aligned} |(I)| &\leq \sqrt{\sum_{t \in T', r \in R', s \in S'} \kappa_r^2 \rho_s^2 |\widehat{F}_t(r, s)|^2} \sqrt{\sum_{t \in T', r \in R', s \in S'} \psi_t^2 |\widehat{F}'_t(r, s)|^2} \\ &= \sqrt{\sum_{r \in R', s \in S'} \kappa_r^2 \rho_s^2 \sum_{t \in T'} |\widehat{F}_t(r, s)|^2} \sqrt{\sum_{t \in T'} \psi_t^2 \sum_{r \in R', s \in S'} |\widehat{F}'_t(r, s)|^2}. \end{aligned}$$

We note that $\sum_{t \in T'} \psi_t^2 \sum_{r \in R', s \in S'} |\widehat{F}'_t(r, s)|^2 \leq \sum_{t \in T'} \psi_t^2$ as $\sum_{r \in R', s \in S'} |\widehat{F}'_t(r, s)|^2 \leq 1$, and that

$$\sum_{t \in T'} |\widehat{F}_t(r, s)|^2 = \mathbb{E}_{(x,y,z) \sim \mu^{\otimes(n-1)}} \left[\widehat{F}(x) g_r(y) h_s(z) \right]^2$$

where $\widehat{F}(x) = \frac{\sum_{t \in T'} \widehat{F}_t(r, s) F_t}{\sqrt{\sum_{t \in T'} |\widehat{F}_t(r, s)|^2}}$, which by the definition is at most $\delta_{n-1,d'-1}^2$. Combining, we get that

$$|(I)| \leq \sqrt{\sum_{r \in R', s \in S'} \kappa_r^2 \rho_s^2 \delta_{n-1,d'-1}^2} \sqrt{\sum_{t \in T'} \psi_t^2} \leq \delta_{n-1,d'-1}.$$

Inspecting the proof, we see that if we had that

$$\begin{aligned} &\left| \sum_{t \in T', r \in R', s \in S'} \psi_t \kappa_r \rho_s \widehat{F}_t(r, s) \widehat{F}'_t(r, s) \right| \\ &\leq (1 - \zeta) \sqrt{\sum_{t \in T', r \in R', s \in S'} \kappa_r^2 \rho_s^2 |\widehat{F}_t(r, s)|^2} \sqrt{\sum_{t \in T', r \in R', s \in S'} \psi_t^2 |\widehat{F}'_t(r, s)|^2} \end{aligned}$$

then we would get that $|(I)| \leq (1 - \zeta)\gamma_{n-1,d'-1}$ and combining with Claim 9.9 we get that $\delta_{n,d'} \leq (1 - \varepsilon)\delta_{n-1,d'-1}$ and the proof would be concluded. We henceforth assume that

$$\begin{aligned} &\left| \sum_{t \in T', r \in R', s \in S'} \psi_t \kappa_r \rho_s \widehat{F}_t(r, s) \widehat{F}'_t(r, s) \right| \\ &> (1 - \zeta) \sqrt{\sum_{t \in T', r \in R', s \in S'} \kappa_r^2 \rho_s^2 |\widehat{F}_t(r, s)|^2} \sqrt{\sum_{t \in T', r \in R', s \in S'} \psi_t^2 |\widehat{F}'_t(r, s)|^2}. \end{aligned} \quad (24)$$

Also, we see that if we had that $\sum_{t \in T'} \psi_t^2 \sum_{r \in R', s \in S'} \left| \widehat{F}'_t(r, s) \right|^2 \leq (1 - \zeta) \sum_{t \in T'} \psi_t^2$, then we would get that $|(I)| \leq \sqrt{1 - \zeta} \delta_{n-1, d'-1}$ and combining with Claim 9.9 we get that $\delta_{n, d'} \leq (1 - \varepsilon) \delta_{n-1, d'-1}$ and the proof would be concluded. We henceforth assume that

$$\sum_{t \in T'} \psi_t^2 \sum_{r \in R', s \in S'} \left| \widehat{F}'_t(r, s) \right|^2 > (1 - \zeta) \sum_{t \in T'} \psi_t^2. \quad (25)$$

Lastly, if $\sum_{t \in T', r \in R', s \in S'} \kappa_r^2 \rho_s^2 \left| \widehat{F}_t(r, s) \right|^2 \leq (1 - \zeta) \delta_{n-1, d'-1}^2 \sum_{r \in R', s \in S'} \kappa_r^2 \rho_s^2$ then once again we will be able to conclude that $\delta_{n, d'} \leq (1 - \varepsilon) \delta_{n-1, d'-1}$, and thus henceforth we have that

$$\sum_{t \in T', r \in R', s \in S'} \kappa_r^2 \rho_s^2 \left| \widehat{F}_t(r, s) \right|^2 > (1 - \zeta) \delta_{n-1, d'-1}^2 \sum_{r \in R', s \in S'} \kappa_r^2 \rho_s^2. \quad (26)$$

9.3.3 The Random Perturbation Argument

Define $M_1 = \sum_{r \in R'} \kappa_r^2$ and $M_2 = \sum_{s \in S'} \rho_s^2$. Below, we will consider complex-valued variables π_r and θ_s satisfying that

$$\sum_{r \in R'} |\pi_r|^2 \kappa_r^2 = M_1, \quad \sum_{s \in S'} |\theta_s|^2 \rho_s^2 = M_2. \quad (27)$$

Eventually, we choose π_r, θ_s according to a distribution over the points satisfying these equalities. More precisely, we choose the distribution of π_r 's so that the vector $(\kappa_r \pi_r)_{r \in R'}$ is distributed uniformly over vectors in $\mathbb{C}^{R'}$ with 2-norm equal to $\sqrt{M_1}$; similarly, we choose the distribution of θ_s so that the vector $(\theta_s \pi_s)_{s \in R'}$ is distributed uniformly over vectors in $\mathbb{C}^{S'}$ with 2-norm equal to $\sqrt{M_2}$.

The point of these constraints is so that we can define the functions

$$\tilde{g}(y) = \sum_{r \in R'} \kappa_r \pi_r g_r(y), \quad \tilde{h}(z) = \sum_{s \in S'} \rho_s \theta_s h_s(z),$$

and have that their 2-norms squared are equal to M_1 and M_2 respectively.

The functions \tilde{g} , and \tilde{h} allow us to consider the lower order, $n - 1$ dimensional problems corresponding to \tilde{F} and \tilde{g}, \tilde{h} where $\tilde{F} = \frac{\sum_t \langle F_t, \tilde{g} \tilde{h} \rangle F_t}{\sqrt{\sum_t |\langle F_t, \tilde{g} \tilde{h} \rangle|^2}}$. Indeed, we define the function

$$p(\{\pi_r\}_{r \in R'}, \{\theta\}_{s \in S'}) = \left| \mathbb{E}_{(x, y, z) \sim \mu^{\otimes (n-1)}} \left[\tilde{F}(x) \tilde{g}(y) \tilde{h}(z) \right] \right|^2 = \sum_t \left| \langle F_t, \overline{\tilde{g} \tilde{h}} \rangle \right|^2.$$

By the first expression for p and definition, we know that $|p(\{\pi_r\}_{r \in R'}, \{\theta\}_{s \in S'})| \leq \delta_{n-1, d'-1}^2$. In our argument we will analyze the expectation of $p(\cdot)$ and show it is close to $\delta_{n-1, d'-1}^2$, from which it follows that the value of $p(\cdot)$ is roughly constant. On the other end, we will analyze the variance of $p(\cdot)$ and lower bound it, and combining these two facts will finish the proof. We now do each one of these steps in detail.

First, we establish the point-wise upper bound on p :

Claim 9.10. *For all inputs satisfying (27) we have that $p(\{\pi_r\}_{r \in R'}, \{\theta\}_{s \in S'}) \leq \delta_{n-1, d'-1}^2$.*

Proof. This is immediate by definition of $\delta_{n-1,d'-1}$. □

To compute the expectation and variance of p it will be useful for us to expand it out. We have:

$$\begin{aligned}
p(\{\pi_r\}_{r \in R'}, \{\theta_s\}_{s \in S'}) &= \sum_t \left| \langle F_t, \widetilde{g\hat{h}} \rangle \right|^2 \\
&= \sum_t \left| \sum_{r \in R', s \in S'} \kappa_r \rho_s \pi_r \theta_s \widehat{F}_t(r, s) \right|^2 \\
&= \sum_t \sum_{r, r' \in R', s, s' \in S'} \kappa_r \overline{\kappa_{r'}} \rho_s \overline{\rho_{s'}} \pi_r \overline{\pi_{r'}} \theta_s \overline{\theta_{s'}} \widehat{F}_t(r, s) \overline{\widehat{F}_t(r', s')} \\
&= \sum_{r, r' \in R', s, s' \in S'} \kappa_r \overline{\kappa_{r'}} \rho_s \overline{\rho_{s'}} \pi_r \overline{\pi_{r'}} \theta_s \overline{\theta_{s'}} \langle V_{r,s}, V_{r',s'} \rangle, \tag{28}
\end{aligned}$$

where the vector $V_{r,s} \in \mathbb{C}^{|T'|}$ is defined as $V_{r,s}(t) = \widehat{F}_t(r, s)$. We are also going to need the following claim that gives us bounds on the norms of the vectors $V_{r,s}$.

Claim 9.11. For all $r \in R'$ and $s \in S'$, $(1 - \sqrt{\zeta})\delta_{n-1,d'-1}^2 \leq \|V_{r,s}\|_2^2 \leq \delta_{n-1,d'-1}^2$.

Proof. Note that we may write (26) as $\sum_{r \in R', s \in S'} \kappa_r^2 \rho_s^2 \|V_{r,s}\|_2^2 \geq (1 - \zeta)\delta_{n-1,d'-1}^2 M_1 M_2$, and so

$$\sum_{r \in R', s \in S'} \kappa_r^2 \rho_s^2 (\delta_{n-1,d'-1}^2 - \|V_{r,s}\|_2^2) \leq \zeta \delta_{n-1,d'-1}^2.$$

Thus, all of the terms on the left hand side are non-negative and it follows that for all $r \in R'$ and $s \in S'$

$$\delta_{n-1,d'-1}^2 - \|V_{r,s}\|_2^2 \leq \frac{\zeta \delta_{n-1,d'-1}^2}{\kappa_r^2 \rho_s^2} \leq \frac{\zeta \delta_{n-1,d'-1}^2}{\tau^4} \leq \sqrt{\zeta} \delta_{n-1,d'-1}^2,$$

and re-arranging gives the lower bound. As for the upper bound, note that for all $r \in R'$ and $s \in S'$ we have that $\|V_{r,s}\|_2^2 \leq \delta_{n-1,d'-1}^2$, as this norm can be written as $\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes(n-1)}} [F'(x)g_r(y)h_s(z)] \right|^2$ where

$$F'(x) = \frac{\sum_{t \in T'} \widehat{F}_t(r,s) F_t}{\sqrt{\sum_{t \in T'} |\widehat{F}_t(r,s)|^2}}. \tag{28}$$

□

9.3.4 Lower Bounding the Expectation of p and Upper Bounding the Variance of p

The following claim bounds the expectation of p .

Claim 9.12. Consider the distribution over π_r and θ_s as defined after (27). Then

$$\mathbb{E} [p(\{\pi_r\}_{r \in R'}, \{\theta_s\}_{s \in S'})] \geq \delta_{n-1,d'-1}^2 (1 - 3\sqrt{\zeta}).$$

Proof. We use (28) and linearity of expectation. If $(r, s) \neq (r', s')$, say $r \neq r'$, then

$$\mathbb{E} [\pi_r \overline{\pi_{r'}} \theta_s \overline{\theta_{s'}}] = \mathbb{E} [\pi_r \overline{\pi_{r'}}] \mathbb{E} [\theta_s \overline{\theta_{s'}}] = 0 \mathbb{E} [\theta_s \overline{\theta_{s'}}] = 0,$$

as the distribution of π_r is invariant under multiplying any one of these numbers by a random sign. Thus, only $(r, s) = (r', s')$ contribute to the expectation, and so

$$\begin{aligned}\mathbb{E}[p(\{\pi_r\}_{r \in R'}, \{\theta_s\}_{s \in S'})] &= \mathbb{E}\left[\sum_{r \in R', s \in S'} \kappa_r^2 \rho_s^2 |\pi_r|^2 |\theta_s|^2 \|V_{r,s}\|_2^2\right] \\ &\geq \delta_{n-1, d'-1}^2 (1 - \sqrt{\zeta}) \mathbb{E}\left[\sum_{r \in R', s \in S'} \kappa_r^2 \rho_s^2 |\pi_r|^2 |\theta_s|^2\right],\end{aligned}$$

where we used Claim 9.11. Using (27) we get that the right hand side is at least $M_1 M_2 \delta_{n-1, d'-1}^2 (1 - \sqrt{\zeta})$, and lastly we note that

$$M_1 = 1 - \sum_{r \in T \setminus T'} \kappa_r^2 \geq 1 - m\eta^2 \geq 1 - \zeta,$$

and similarly $M_2 \geq 1 - \zeta$, and the claim follows. \square

Combining Claims 9.12 and 9.10 give a strong upper bound on the variance of p .

Claim 9.13. $\text{var}(p(\{\pi_r\}_{r \in R'}, \{\theta_s\}_{s \in S'})) \lesssim \sqrt{\zeta} \delta_{n-1, d'-1}^4$.

Proof. By definition of variance it is equal to

$$\mathbb{E}[p(\{\pi_r\}_{r \in R'}, \{\theta_s\}_{s \in S'})^2] - \mathbb{E}[p(\{\pi_r\}_{r \in R'}, \{\theta_s\}_{s \in S'})]^2 \leq \delta_{n-1, d'-1}^4 - \delta_{n-1, d'-1}^4 (1 - 3\sqrt{\zeta})^2,$$

where we used Claim 9.10 to upper bound the first expectation and Claim 9.12 to lower bound the second expectation. Simplifying finishes the proof. \square

9.3.5 Lower Bounding the Variance of p

The rest of the argument is devoted to lower bounding the variance of p . Write $p = p_1 + p_2 + p_3 + p_4$ where

$$\begin{aligned}p_1 &= \sum_{r \in R', s \neq s' \in S'} \kappa_r^2 \rho_s \rho_{s'} |\pi_r|^2 \theta_s \overline{\theta_{s'}} \langle V_{r,s}, V_{r,s'} \rangle, & p_2 &= \sum_{r \neq r' \in R', s \in S'} \kappa_r \kappa_{r'} \rho_s^2 \pi_r \overline{\pi_{r'}} |\theta_s|^2 \langle V_{r',s}, V_{r,s} \rangle, \\ p_3 &= \sum_{r \neq r' \in R', s \neq s' \in S'} \kappa_r \kappa_{r'} \rho_s \overline{\rho_{s'}} \pi_r \overline{\pi_{r'}} \theta_s \theta_{s'} \langle V_{r',s}, V_{r,s'} \rangle, & p_4 &= \sum_{r \in R', s \in S'} \kappa_r^2 \rho_s^2 |\pi_r|^2 |\theta_s|^2 \|V_{r,s}\|_2^2.\end{aligned}$$

Note that p_4 is almost constant, and close to $\mathbb{E}[p]$. Indeed, by Claim 9.11

$$\begin{aligned}|\mathbb{E}[p] - p_4| &\leq \sum_{r \in R', s \in S'} \kappa_r^2 \rho_s^2 |\pi_r|^2 |\theta_s|^2 (\delta_{n-1, d'-1}^2 - \|V_{r,s}\|_2^2) \leq \sqrt{\zeta} \delta_{n-1, d'-1}^2 \sum_{r \in R', s \in S'} \kappa_r^2 \rho_s^2 |\pi_r|^2 |\theta_s|^2 \\ &\leq \sqrt{\zeta} \delta_{n-1, d'-1}^2.\end{aligned}$$

Thus, we get that

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

We have that $\mathbb{E}[|p_4 - \mathbb{E}[p]|^2] \leq \zeta \delta_{n-1, d'-1}^4$, as well as

$$\left| \mathbb{E}[(p_1 + p_2 + p_3) \overline{(p_4 - \mathbb{E}[p])}] \right| \lesssim \mathbb{E}[\zeta^{1/2} |p_1 + p_2 + p_3|^2 + \zeta^{-1/2} |p_4 - \mathbb{E}[p]|^2]$$

and

$$\left| \mathbb{E}[\overline{(p_1 + p_2 + p_3)}(p_4 - \mathbb{E}[p])] \right| \lesssim \mathbb{E}[\zeta^{1/2} |p_1 + p_2 + p_3|^2 + \zeta^{-1/2} |p_4 - \mathbb{E}[p]|^2].$$

Therefore, we get that

$$\text{var}(p) \geq \left(1 - O(\sqrt{\zeta})\right) \mathbb{E}[|p_1 + p_2 + p_3|^2] - O(\sqrt{\zeta} \delta_{n-1, d'-1}^4). \quad (29)$$

Inspecting (29), we now calculate $\mathbb{E}[|p_1 + p_2 + p_3|^2]$. We get diagonal terms $\mathbb{E}[|p_1|^2]$, $\mathbb{E}[|p_2|^2]$, $\mathbb{E}[|p_3|^2]$ as well as off diagonal terms such as $\mathbb{E}[p_1 \overline{p_2}]$, and we claim the off diagonal terms are 0. Indeed, looking at $p_1 \overline{p_2}$ for example, if we expand it out we will get that each term will be of the a multiple of the pattern $|\theta_s|^2 \overline{\pi_r} \pi_{r'} \theta_{s'} \theta_{s''} |\pi_{r''}|^2$ where $r \neq r'$ (and $s' \neq s''$), and in particular either r or r' (or both) are different from r'' — say r — in which case the expectation of this pattern is 0 as the distribution is invariant under multiplying π_r by a random sign. Therefore (29) gives that

$$\text{var}(p) \geq \left(1 - O(\zeta^{1/2})\right) \mathbb{E}[|p_1|^2 + |p_2|^2 + |p_3|^2] - O(\sqrt{\zeta} \delta_{n-1, d'-1}^4). \quad (30)$$

To use (30) effectively we must prove that at least one of the monomials p_1, p_2 or p_3 has a significant 2-norm, and towards this end we must show that some inner product $\langle V_{r,s}, V_{r',s'} \rangle$ for $(r, s) \neq (r', s')$ is significant. This is the content of the following lemma, whose proof is deferred to Section 9.3.6.

Lemma 9.14. *There are $r, r' \in R'$ and $s, s' \in S$ such that $(r, s) \neq (r', s')$ and*

$$|\langle V_{r,s}, V_{r',s'} \rangle| \geq c \delta_{n-1, d'_1-1}^2.$$

Proof. The proof is deferred to Section 9.3.6. We remark that this is the place in the argument in which the relaxed base case is used. \square

We can now lower bound the variance of p .

Lemma 9.15. $\text{var}(p) \gtrsim_m \tau^9 \delta_{n-1, d'_1-1}^4$.

Proof. By Lemma 9.14 there are $(r^*, s^*) \neq (r^{*'}, s^{*'})$ such that $|\langle V_{r^*, s^*}, V_{r^{*'}, s^{*'}} \rangle| \gtrsim_m c \delta_{n-1, d'-1}^2$, and we split into cases.

The case $r^* \neq r^{*'}$, $s^* \neq s^{*'}$. In this case we have by (30) that

$$\text{var}(p) \geq \left(1 - O(\zeta^{1/2})\right) \mathbb{E}[|p_3|^2] - O(\sqrt{\zeta} \delta_{n-1, d'-1}^4),$$

and we lower bound $\mathbb{E}[|p_3|^2]$ by expanding. Indeed, it is equal to

$$\sum_{\substack{r \neq r' \in R', s \neq s' \in S' \\ r'' \neq r''' \in R', s'' \neq s''' \in S'}} \kappa_r \overline{\kappa_{r'}} \kappa_{r''} \overline{\kappa_{r'''}} \rho_s \overline{\rho_{s'}} \rho_{s''} \overline{\rho_{s'''}} \langle V_{r', s'}, V_{r, s} \rangle \langle \overline{V_{r''', s'''}}, \overline{V_{r'', s''}} \rangle \mathbb{E}[\overline{\pi_r} \overline{\pi_{r'}} \overline{\pi_{r''}} \overline{\pi_{r'''}} \theta_s \overline{\theta_{s'}} \theta_{s''} \overline{\theta_{s'''}}].$$

The expectation is 0 unless $r' = r'''$, $r = r''$ and $s' = s'''$, $s = s''$; this is because the distribution of each π_r is invariant under multiplying each one of them by a random complex number of absolute value 1. We thus get

$$\begin{aligned}
\mathbb{E}[|p_3|^2] &= \sum_{r \neq r' \in R', s \neq s' \in S'} \kappa_r^2 \kappa_{r'}^2 \rho_s^2 \rho_{s'}^2 |\langle V_{r',s}, V_{r,s'} \rangle|^2 \mathbb{E} \left[|\pi_r|^2 |\pi_{r'}|^2 |\theta_s|^2 |\theta_{s'}|^2 \right] \\
&\geq \kappa_{r^*}^2 \kappa_{r^{*'}}^2 \rho_{s^*}^2 \rho_{s^{*'}}^2 |\langle V_{r^{*'},s^*}, V_{r^*,s^{*'}} \rangle|^2 \mathbb{E} \left[|\pi_{r^*}|^2 |\pi_{r^{*'}}|^2 |\theta_{s^*}|^2 |\theta_{s^{*'}}|^2 \right] \\
&\geq c\tau^8 \delta_{n-1,d'-1}^4 \mathbb{E} \left[|\pi_{r^*}|^2 |\pi_{r^{*'}}|^2 |\theta_{s^*}|^2 |\theta_{s^{*'}}|^2 \right] \\
&= c\tau^8 \delta_{n-1,d'-1}^4 \mathbb{E} \left[|\pi_{r^*}|^2 |\pi_{r^{*'}}|^2 \right] \mathbb{E} \left[|\theta_{s^*}|^2 |\theta_{s^{*'}}|^2 \right].
\end{aligned}$$

We now argue that

$$\mathbb{E} \left[|\pi_{r^*}|^2 |\pi_{r^{*'}}|^2 \right], \mathbb{E} \left[|\theta_{s^*}|^2 |\theta_{s^{*'}}|^2 \right] \gtrsim_m 1$$

We show the argument for the first expectation and the second one is similar. Looking at $(\kappa_r \pi_r)_{r \in R'}$, we see that this is a random vector in $\mathbb{C}^{|R'|}$ of 2-norm equal to 1, and so $\mathbb{E} \left[|\kappa_r \pi_r|^2 |\kappa_{r'} \pi_{r'}|^2 \right] \gtrsim_m 1$ for all r, r' .

Hence, $\mathbb{E} \left[|\pi_{r^*}|^2 |\pi_{r^{*'}}|^2 \right] \gtrsim 1$.

Overall, we get that $\text{var}(p) \geq (\Omega_{c,m}(\tau^8) - \sqrt{c}) \delta_{n-1,d'-1}^4$, hence $\text{var}(p) \geq \tau^9 \delta_{n-1,d'-1}^4$.

The case $r^* = r^{*'}$, $s^* \neq s^{*'}$. Let E be the event that

$$|\kappa_{r^*} \pi_{r^*}| \geq \sqrt{M_1 - \frac{c^{50}}{m^{50}}}, \quad |\rho_{s^*} \theta_{s^*}| \geq \frac{1}{\sqrt{2}} \sqrt{M_2 - \frac{c^{50}}{m^{50}}}, \quad |\rho_{s^{*'}} \theta_{s^{*'}}| \geq \frac{1}{\sqrt{2}} - \frac{c^{50}}{m^{50}}.$$

Then $\Pr[E] \gtrsim_{c,m} 1$, and if E happens then for every $r \neq r^*$, $s \neq s^*$, $s^{*'}$ it holds that

$$|\kappa_r \pi_r| \leq \sqrt{M_1 - |\kappa_{r^*} \pi_{r^*}|^2} \leq \frac{c^{50}}{m^{50}}, \quad |\rho_s \theta_s| \leq \sqrt{M_2 - |\rho_{s^*} \theta_{s^*}|^2 - |\rho_{s^{*'}} \theta_{s^{*'}}|^2} \leq \frac{c^{50}}{m^{50}},$$

and it follows that

$$|p_1| \geq |\kappa_{r^*} \pi_{r^*}|^2 |\rho_{s^*} \theta_{s^*}| |\rho_{s^{*'}} \theta_{s^{*'}}| |\langle V_{r^*,s^*}, V_{r^*,s^{*'}} \rangle| - m^4 \frac{c^{50}}{m^{50}} \max_{r,s} \|V_{r,s}\|_2^2.$$

By the choice of r^* , s^* , $s^{*'}$ and Claim 9.11, the last expression is at least $\gtrsim_{c,m} \delta_{n-1,d-1}^2$. It follows that $\mathbb{E} \left[|p_1|^2 \right] \gtrsim_{c,m} \Pr[E] \delta_{n-1,d-1}^4 \gtrsim_{c,m} \delta_{n-1,d-1}^4$, and the proof is concluded by plugging this into (30).

The case $r^* \neq r^{*'}$, $s^* = s^{*'}$. This case is very similar to the case above, except that we work with p_2 instead of p_1 . We omit the details. \square

Combining Claim 9.13 and Lemma 9.15 gives contradiction to (20). The contradiction means that not all of (24), (25), (26) can hold, and therefore as explained above the proof of Lemma 9.3 is concluded.

9.3.6 Proof of Lemma 9.14: the Dimensionality Argument

Assume towards contradiction that Lemma 9.14 is false. We start with the following general lemma, which quantifying the fact that ℓ vectors in an $\ell - 1$ -dimensional space cannot be all orthogonal.

Fact 9.16. *Suppose that u_1, \dots, u_ℓ are unit vectors in \mathbb{C}^k where $k \leq \ell - 1$. Then there are i, j such that $|\langle u_i, u_j \rangle| \geq \frac{1}{k\ell}$.*

Proof. Construct the Gram matrix $M \in \mathbb{C}^{\ell \times \ell}$ defined as $M[i, j] = \langle u_i, u_j \rangle$. Then M is Hermitian and therefore it has non-negative real eigenvalues a_1, \dots, a_ℓ . As the rank of M is at most k , it follows that only k of these eigenvalues can be non-zero, and without loss of generality these are a_1, \dots, a_k . Thus, we have that

$$\sum_{i,j} |\langle u_i, u_j \rangle|^2 = \text{tr}(MM^*) = \sum_{i=1}^k a_i^2 \geq \frac{1}{k} \left(\sum_{i=1}^k a_i \right)^2 = \frac{1}{k} \text{tr}(M)^2 = \frac{\ell^2}{k}.$$

The summands where $i = j$ contribute ℓ to the left hand side, and so

$$\sum_{i \neq j} |\langle u_i, u_j \rangle|^2 \geq \frac{\ell(\ell - k)}{k} \geq \frac{\ell}{k}.$$

It follows that there are $i \neq j$ such that $|\langle u_i, u_j \rangle| \geq \frac{1}{k\ell}$. □

Next, we show that F'_t has almost all of its mass on $r \in R'$ and $s \in S'$.

Claim 9.17. *For all $t \in T'$ it holds that $\sum_{r \in R', s \in S'} |\widehat{F}'_t(r, s)|^2 \geq 1 - \sqrt{\zeta}$.*

Proof. First, note that

$$\sum_{t \in T'} \psi_t^2 = 1 - \sum_{t \in T \setminus T'} \psi_t^2 \geq 1 - m\eta^2 \geq 1 - \eta.$$

Combining with (25) we get that

$$\sum_{t \in T'} \psi_t^2 \left(1 - \sum_{r \in R', s \in S'} |\widehat{F}'_t(r, s)|^2 \right) \leq \zeta + \eta \leq 2\zeta.$$

As each summand on the left hand side is non-negative, it follows that for each $t \in T'$ it holds that

$$\sum_{r \in R', s \in S'} |\widehat{F}'_t(r, s)|^2 \geq 1 - \frac{2\zeta}{\psi_t^2} \geq 1 - \sqrt{\zeta}. \quad \square$$

We can now show that for each $r \in R'$, $s \in S'$, the function $g'_r h'_s$ is close to be a linear combination of F'_t for $t \in T'$. Formally:

Claim 9.18. *For all $r \in R'$ and $s \in S'$ it holds that $\sum_{t \in T'} |\widehat{F}'_t(r, s)|^2 \geq 1 - \zeta^{1/4}$.*

Proof. Assume this is not the case, so that for some r^* and s^* we have that $\sum_{t \in T'} \left| \widehat{F}'_t(r^*, s^*) \right|^2 \leq 1 - \zeta^{1/4}$, we remark that for any other r, s this sum is at most 1. Summing Claim 9.17 over $t \in T'$ gives:

$$(1 - \sqrt{\zeta}) |T'| \leq \sum_{r \in R', s \in S'} \sum_{t \in T'} \left| \widehat{F}'_t(r, s) \right|^2 \leq (|R'| |S'| - 1) + 1 - \zeta^{1/4}.$$

Thus,

$$|T'| \leq (|R'| |S'| - 1)(1 + O(\sqrt{\zeta})) + (1 - \zeta^{1/4})(1 + O(\sqrt{\zeta})) \leq |R'| |S'| + O_m(\sqrt{\zeta}) - \frac{1}{2} \zeta^{1/4} < |R'| |S'|,$$

and as $|T'|$ is an integer it follows that $|T'| \leq |R'| |S'| - 1$. From Fact 9.16 it follows that there are $(r, s) \neq (r', s')$ such that

$$\left| \left\langle \frac{V_{r,s}}{\|V_{r,s}\|_2}, \frac{V_{r',s'}}{\|V_{r',s'}\|_2} \right\rangle \right| \geq \frac{1}{m^2},$$

and by Claim 9.11 we have $|\langle V_{r,s}, V_{r',s'} \rangle| \geq c \delta_{n-1, d'-1}^2$, and contradiction. \square

Motivated by Claim 9.18, we define

$$F'_{r,s} = \frac{\sum_{t \in T'} \overline{\widehat{F}'_t(r, s)} F'_t}{\sum_{t \in T'} \left| \widehat{F}'_t(r, s) \right|^2}.$$

Thus, $F'_{r,s}$ has 2-norm 1 and by Claim 9.18

$$\mathbb{E}_{(x,y,z) \sim \mu} [F'_{r,s}(y, z) g'_r(y) h'_s(z)] = \sqrt{\sum_{t \in T'} \left| \widehat{F}'_t(r, s) \right|^2} \geq 1 - \zeta^{1/4}.$$

It follows that

$$\|F'_{r,s} - g'_r h'_s\|_2^2 \leq 2\zeta^{1/4}. \quad (31)$$

Also, as each F'_t is constant on connected components it follows that each $F'_{r,s}$ is constant on connected components and so it can be written as $F'_{r,s} = W f'_{r,s}$ for $f'_{r,s}: \Sigma \rightarrow \mathbb{C}$. Thus, as

$$\left| \mathbb{E}_{(x,y,z) \sim \mu} [f'_{r,s}(x) g'_r(y) h'_s(z)] \right| \geq \left| \mathbb{E}_{(x,y,z) \sim \mu} [F'_{r,s}(y, z) g'_r(y) h'_s(z)] \right| \geq 1 - \zeta^{1/4}$$

it follows by the relaxed base case (Definition 6.6) that $\text{var}_{\Sigma_{\text{modest}}}(f'_{r,s}) \leq \zeta^c$ for all $r \in R'$ and $s \in S'$. We are now going to use Claim 9.17 and (31) to argue that each F'_t is close to a linear combination of $F'_{r,s}$, and hence has small variance on Σ_{modest} .

Claim 9.19. *For all $t \in T'$, write $F'_t = W f'_t$ for $f'_t: \Sigma \rightarrow \mathbb{C}$. Then for all $t \in T'$ we have*

$$\text{var}_{\Sigma_{\text{modest}}}(f'_t) \lesssim_m \zeta^{c/2}.$$

Proof. By Claim 9.18, taking $\tilde{F}'_t = \sum_{r \in R', s \in S'} \overline{\widehat{F}'_t(r, s)} g'_r h'_s$ we have

$$\mathbb{E}_{(x, y, z) \sim \mu} \left[\tilde{F}'_t(y, z) F'_t(y, z) \right] = \sqrt{\sum_{r \in R', s \in S'} \left| \widehat{F}'_t(r, s) \right|^2} \geq 1 - \sqrt{\zeta},$$

hence $\|\tilde{F}'_t - F'_t\|_2 \lesssim \zeta^{1/4}$. Using (31) we get that

$$\|\tilde{F}'_t - \sum_{r \in R', s \in S'} \overline{\widehat{F}'_t(r, s)} F'_{r, s}\|_2 \lesssim_m \zeta^{1/4},$$

so by the triangle inequality

$$\|F'_t - \sum_{r \in R', s \in S'} \overline{\widehat{F}'_t(r, s)} F'_{r, s}\|_2 \lesssim_m \zeta^{1/4}.$$

From the above it follows that

$$\|f'_t - \sum_{r \in R', s \in S'} \overline{\widehat{F}'_t(r, s)} f'_{r, s}\|_2 \lesssim_m \zeta^{1/4},$$

and as $\text{var}_{\Sigma_{\text{modest}}}(f'_{r, s}) \leq \zeta^c$ for all $r \in R', s \in S'$ it follows that $\text{var}_{\Sigma_{\text{modest}}}(f'_t) \leq \zeta^{c/2}$, as desired. \square

By Claim 9.19 we get that

$$\sum_{t \in T'} \psi_t^2 \text{var}_{\Sigma_{\text{modest}}}(f'_t) \leq \zeta^{c/2},$$

which contradicts the choice of the partition I, J as in (21) and (20). This contradiction completes the proof of Lemma 9.14. \square

9.4 Finishing the Proof of Theorem 7.23

We are now ready to prove Theorem 7.23. We will take the parameters

$$0 \ll \xi \ll C^{-1} \ll L^{-1} \ll \alpha, m^{-1} \leq 1.$$

Let $n \geq d \geq d'$ and μ be as in Theorem 7.23. Applying Lemma 8.13, if the first item holds we are done, so assume otherwise. By the second item we get that there are n', e and e' satisfying $e \geq d/2$, $n' \leq Le$ and $e' \geq d'/2$ such that $\beta_{n, d, d'} \leq \beta_{n', e, e'}$. Clearly, $\beta_{n', e, e'} \leq \delta_{n', e, e'}$, and we now use Lemma 9.3. Now that $e' \geq \frac{e}{e\xi} \geq \frac{n}{Le\xi}$, and applying Lemma 9.3 $e'/2$ times we get that

$$\delta_{n', e, e'} \leq \left(1 - \frac{1}{LCe\xi C}\right) \delta_{n'-1, e-1, e'-1} \leq \dots \leq \left(1 - \frac{1}{2LCe\xi C}\right)^{e'/2} \delta_{n'-e'/2, e-e'/2, e'-e'/2},$$

which is at most $\left(1 - \frac{1}{2LCe\xi C}\right)^{e'/2}$ by the trivial bound $\delta_{n'-e'/2, e-e'/2, e'-e'/2} \leq 1$. It follows that

$$\delta_{n', e, e'} \leq 2^{-\Omega\left(\frac{e'/2}{2LCe\xi C}\right)} \leq 2^{-\Omega_{m, \alpha}(e^{1-\sqrt{\xi}})},$$

concluding the proof.

9.5 Finishing the Proof of Theorem 4.13: the Chain of Implications

We now note that Theorem 4.13 has been proved, and below is the chain of implications:

1. We completed the proof of Theorem 7.23 in Section 9.4.
2. By Section 7.3.1 it follows that Theorem 7.14 is true.
3. By Claim 7.15 (namely, the content of Section B) it follows that Theorem 6.7 is true.
4. By Section 6.4 it follows that Theorem 6.5 is true.
5. By Section 6.1.1 it follows that Theorem 4.13 is true.

10 A Structural Results for f After Path Tricks: Proof of Theorem 3.18

In this section, we prove Theorem 3.18. We start by giving a high level overview of the argument. Given a distribution μ and functions $f: \Sigma^n \rightarrow \mathbb{C}$, $g: \Gamma^n \rightarrow \mathbb{C}$ and $h: \Phi^n \rightarrow \mathbb{C}$ as in the statement of Theorem 3.18, we perform the following steps:

1. **Softly truncating the high non-embedding parts:** First, consider the high non-embedding degree parts of f , g and h , namely $(I - T_{\text{non-embed}, 1-\delta})f$, $(I - T_{\text{non-embed}, 1-\delta})g$ and $(I - T_{\text{non-embed}, 1-\delta})h$. We show that they give very little contribution to the 3-wise correlation, and hence conclude that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [T_{\text{non-embed}, 1-\delta} f(x) T_{\text{non-embed}, 1-\delta} g(y) T_{\text{non-embed}, 1-\delta} h(z)] \right| \geq \frac{\varepsilon}{2}. \quad (32)$$

Set $f' = T_{\text{non-embed}, 1-\delta} f$, $g' = T_{\text{non-embed}, 1-\delta} g$ and $h' = T_{\text{non-embed}, 1-\delta} h$, so that almost all of the mass of these functions lies on monomials with non-embedding degree which is constant.

2. **Reducing the non-embedding degree to be 0:** By random restriction, we show that with noticeable probability, after a random restriction almost all of the mass of f' , g' and h' collapses to non-embedding degrees 0 and (32) continues to hold (with some loss in parameters). We then use an averaging operator that annihilates all of the mass of f' , g' and h' of non-zero non-embedding degree, apply it on f' , g' and h' and show that (32) continues to hold (with an additional loss in parameters). We perform these steps (as opposed to harsher truncation-style operations) so as to arrive at functions that have embedding degree which is 0 which are also bounded.
3. **Shifting to the uniform distribution over an equation on H :** By another random restriction argument we switch from the distribution μ to a distribution ν , in which sampling $(x, y, z) \sim \nu$, the distribution of $(\sigma(x), \gamma(y), \phi(z))$ is uniform over $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$. We argue that moving to the distribution ν , (32) continues to hold for the restrictions (again, with some loss in parameters). Once we have shifted the distribution to be ν , we can use a standard Fourier-analytic computation showing that the restriction we chose for f' has a significant Fourier coefficient. In our case, this last assertion translates to the fact that the restriction of f' is correlated with a function of the form $\chi \circ \sigma$ for some character χ over H .

4. **Unraveling restrictions and averaging:** The last part in the argument is to invoke a result that we refer to as the restriction inverse theorem. Morally speaking, up until now we have shown that after random restriction, with noticeable probability f' is correlated with a function of the form $\chi \circ \sigma$. Now we would like to use this information to conclude a correlation result for the function f' itself. The restriction inverse theorem is a result asserting precisely that: for restrictions of a function f' to be correlated with a function of the form $\chi \circ \sigma$ with noticeable probability, it must be the case that f' itself is already correlated with a function of the form $L \cdot \chi \circ \sigma$, where L is a low-degree function.

The restriction inverse theorem will be the subject of discussion in Section 12, and for now we give a special case of it (which is the version necessary for the current argument).

Theorem 10.1. *For all $\alpha > 0$, $r, m \in \mathbb{N}$, $\rho \in (0, 1)$ and $\varepsilon, \delta > 0$ there are $d \in \mathbb{N}$ and $\varepsilon' > 0$ such that the following holds. Suppose that Σ is an alphabet of size at most m and μ is a distribution over Σ in which the probability of each atom is at least α . Suppose that $(H, +)$ is an Abelian group of size at most r and $\sigma: \Sigma \rightarrow H$ is a map that has 0 in its image.*

Suppose that $\mu = \rho\nu + (1 - \rho)\nu'$ for distributions ν, ν' in which the probability of each atom is at least α . If $f: (\Sigma^n, \mu^{\otimes n}) \rightarrow \mathbb{C}$ is a 1-bounded function such that

$$\Pr_{I \subseteq_{\rho} [n], \tilde{x} \sim \nu'^I} \left[\exists \chi \in \hat{H}^I, \left| \langle f_{I \rightarrow \tilde{x}}, \chi \circ \sigma \rangle_{L_2(\Sigma^I; \nu^I)} \right| \geq \varepsilon \right] \geq \delta,$$

then there exists $\chi' \in \hat{H}^n$ and $L: \Sigma^n \rightarrow \mathbb{C}$ a function of degree at most d and $\|L\|_2 \leq 1$ such that

$$\left| \langle f, L\chi' \circ \sigma \rangle \right| \geq \varepsilon'.$$

Quantitatively, $d = \text{poly}_{m, \alpha} \left(\frac{\log(1/\varepsilon) \log(1/\delta)}{\rho} \right)$ and $\varepsilon' = 2^{-\text{poly}_{m, \alpha} \left(\frac{\log(1/\varepsilon) \log(1/\delta)}{\rho} \right)}$.

Proof. Deferred to Section 12. □

10.1 Softly Truncating the High Non-embedding Parts: Applying Theorem 4.13

The following lemma applies Theorem 4.13 along with a standard replacement argument to truncate the high non-embedding degree parts of the functions f, g and h as explained above.

Lemma 10.2. *For all $m \in \mathbb{N}$ and $\alpha > 0$ there is $\xi > 0$ such that the following holds for sufficiently small $\varepsilon > 0$. Let μ be a distribution as in Theorem 3.18, and let $f: \Sigma^n \rightarrow \mathbb{C}$, $g: \Gamma^n \rightarrow \mathbb{C}$, $h: \Phi^n \rightarrow \mathbb{C}$ be 1-bounded functions such that*

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

Then for $\delta = (\varepsilon/20)^{3/\xi}$ we have

$$\left| \mathbb{E}_{(x, y, z) \sim \mu^{\otimes n}} [\mathbb{T}_{\text{non-embed}, 1-\delta} f(x) \mathbb{T}_{\text{non-embed}, 1-\delta} g(y) \mathbb{T}_{\text{non-embed}, 1-\delta} h(z)] \right| \geq \frac{\varepsilon}{2}.$$

Proof. Denoting $f' = \mathbb{T}_{\text{non-embed}, 1-\delta} f$, $g' = \mathbb{T}_{\text{non-embed}, 1-\delta} g$ and $h' = \mathbb{T}_{\text{non-embed}, 1-\delta} h$, we wish to show that f, g and h to f', g' and h' one step at a time. First, we claim that provided that ξ is small enough it holds that

$$\left| \mathbb{E}_{(x, y, z) \sim \mu^{\otimes n}} [(f - f')(x)g(y)h(z)] \right| \leq \frac{\varepsilon}{10}.$$

Indeed, defining $f'' = \frac{1}{2}(I - T_{\text{non-embed},1-\delta})f$, we get that f'' is 1-bounded and

$$\text{NEStab}_{1-\delta^{1/3}}(f'') \leq \|T_{\text{non-embed},1-\delta^{1/3}}(I - T_{\text{non-embed},1-\delta})f\|_2,$$

which is at most the largest eigenvalue of $T_{\text{non-embed},1-\delta^{1/3}}(I - T_{\text{non-embed},1-\delta})$. This eigenvalue can be easily seen to be at most $\delta^{1/3}$, so $\text{NEStab}_{1-\delta^{1/3}}(f'') \leq \delta^{1/3}$. Applying Theorem 4.13 we get that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f''(x)g(y)h(z)] \right| \leq \delta^{\xi/3} \leq \frac{\varepsilon}{20},$$

Thus $|\mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [(f - f')(x)g(y)h(z)]| \leq \frac{\varepsilon}{10}$ and by the triangle inequality

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f'(x)g(y)h(z)] \right| \geq \frac{9\varepsilon}{10}.$$

Continuing in this way, we apply the same argument to replace g into g' and h into h' to get that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f'(x)g'(y)h'(z)] \right| \geq \frac{7\varepsilon}{10},$$

and the proof is concluded. □

10.2 Reducing the Non-embedding Degree to be 0

Let f, g, h be as in the setting of Lemma 10.2, and take

$$f' = T_{\text{non-embed},1-\delta}f, \quad g' = T_{\text{non-embed},1-\delta}g, \quad h' = T_{\text{non-embed},1-\delta}h.$$

Choose $I \subseteq_{\kappa} [n]$ where $\kappa = \frac{\varepsilon^3 \delta}{100^3}$, sample $(\tilde{x}, \tilde{y}, \tilde{z}) \sim \mu^{\bar{I}}$ and define

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

Also, let

$$f^{\sharp} = T_{\text{non-embed},0}f'', \quad g^{\sharp} = T_{\text{non-embed},0}g'', \quad h^{\sharp} = T_{\text{non-embed},0}h''$$

in other words, f^{\sharp} is the part of f'' of non-embedding degree 0, g^{\sharp} is the part of g'' of non-embedding degree 0 and h^{\sharp} is the part of h'' of non-embedding degree 0

Claim 10.3. *In the above setting,*

$$\mathbb{E}_{I, \tilde{x}} [\|f'' - f^{\sharp}\|_2^2] \leq \frac{\kappa}{\delta}, \quad \mathbb{E}_{I, \tilde{y}} [\|g'' - g^{\sharp}\|_2^2] \leq \frac{\kappa}{\delta}, \quad \mathbb{E}_{I, \tilde{z}} [\|h'' - h^{\sharp}\|_2^2] \leq \frac{\kappa}{\delta}.$$

Proof. We prove the first inequality, and the other two are analogous.

For a fixed I and \tilde{x} , the norm $\|f'' - f^{\sharp}\|_2^2$ is the mass of f'' on monomials of non-embedding degree at least 1. Expanding

$$f'(x) = \sum_{v \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^n} \widehat{f}'(v)v(x),$$

we have that

$$(f'' - f^\sharp)(x) = \sum_{u \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^I} \left(\sum_{v \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^{\bar{I}}} \widehat{f}'(v, u) v(\tilde{x}) \right) u(x).$$

Thus, for fixed I we have

$$\begin{aligned} \mathbb{E}_{\tilde{x}} \left[\|f'' - f^\sharp\|_2^2 \right] &= \sum_{u \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^I} \mathbb{E}_{\tilde{x}} \left[\left(\sum_v \widehat{f}'(u, v) v(\tilde{x}) \right)^2 \right] \mathbf{1}_{\text{nedeg}(u) \geq 1} \\ &= \sum_{u \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^I} \sum_v \widehat{f}'(v, u)^2 \mathbf{1}_{\text{nedeg}(u) \geq 1} \\ &= \sum_{v \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^n} \widehat{f}'(v)^2 \mathbf{1}_{\text{nedeg}(v|_I) \geq 1} \\ &= \sum_{v \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^n} \widehat{f}(v)^2 (1 - \delta)^{2 \text{nedeg}(v)} \mathbf{1}_{\text{nedeg}(v|_I) \geq 1} \end{aligned}$$

Taking expectation over I gives

$$\begin{aligned} \mathbb{E}_{I, \tilde{x}} \left[\|f'' - f^\sharp\|_2^2 \right] &\leq \sum_{v \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^n} \widehat{f}(v)^2 (1 - \delta)^{2 \text{nedeg}(v)} \mathbb{E}_I \left[\mathbf{1}_{\text{nedeg}(v|_I) \geq 1} \right] \\ &\leq \sum_{v \in (B_{\text{embed}} \cup B_{\text{non-embed}} \cup B_{\text{modest}})^n} \widehat{f}(v)^2 (1 - \delta)^{2 \text{nedeg}(v)} \kappa \cdot \text{nedeg}(v), \end{aligned}$$

where we used the union bound. Looking at the function $P: [0, \infty) \rightarrow [0, \infty)$ defined by $P(s) = \kappa s (1 - \delta)^{2s}$, its maximum is achieved at $s = \frac{1}{2 \log(1/(1-\delta))} \leq \frac{1}{\delta}$ and hence is at most $\frac{\kappa}{\delta}$. \square

Next we analyze the triple correlation of f'' , g'' and h'' and show that with noticeable probability it remains significant.

Claim 10.4. *In the above setting*

$$\Pr_{I, \tilde{x}, \tilde{y}, \tilde{z}} \left[\left| \mathbb{E}_{(x, y, z) \sim \mu^I} [f''(x) g''(y) h''(z)] \right| \geq \frac{\varepsilon}{4} \right] \geq \frac{\varepsilon}{4}.$$

Proof. Consider the real valued random variable $V(I, \tilde{x}, \tilde{y}, \tilde{z}) = \left| \mathbb{E}_{(x, y, z) \sim \mu^I} [f''(x) g''(y) h''(z)] \right|$. Then by the triangle inequality

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

and as $0 \leq V \leq 1$ it follows by an averaging argument that $\Pr [V \geq \varepsilon/4] \geq \varepsilon/4$, as required. \square

We now combine the last two claims to conclude that the triple correlation of f^\sharp , g^\sharp and h^\sharp remains significant with noticeable probability.

Lemma 10.5. *In the above setting*

$$\Pr_{I, \tilde{x}, \tilde{y}, \tilde{z}} \left[\left| \mathbb{E}_{(x,y,z) \sim \mu^I} \left[f^\#(x)g^\#(y)h^\#(z) \right] \right| \geq \frac{\varepsilon}{8} \right] \geq \frac{\varepsilon}{8}.$$

Proof. Let E_1 be the event that $\left| \mathbb{E}_{(x,y,z) \sim \mu^I} [f''(x)g''(y)h''(z)] \right| \geq \frac{\varepsilon}{4}$ and let E_2 be the event that each one of $\|f'' - f^\#\|_2^2$, $\|g'' - h^\#\|_2^2$ and $\|g'' - h^\#\|_2^2$ is at most $\frac{100\kappa}{\varepsilon\delta}$. By the Union bound, Claim 10.3 and Markov's inequality we have that

$$\Pr[E_2] = 1 - \Pr[\overline{E_2}] \geq 1 - 3\frac{\varepsilon}{100},$$

so $\Pr[E_1 \cap E_2] \geq \frac{\varepsilon}{4} - \frac{3\varepsilon}{100} \geq \frac{\varepsilon}{8}$. We argue that when $E_1 \cap E_2$ holds, we have that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^I} \left[f^\#(x)g^\#(y)h^\#(z) \right] \right| \geq \frac{\varepsilon}{8}.$$

To do so, we start with $\left| \mathbb{E}_{(x,y,z) \sim \mu^I} [f''(x)g''(y)h''(z)] \right| \geq \frac{\varepsilon}{4}$ and replace f'' , g'' and h'' by $f^\#$, $g^\#$ and $h^\#$ one step at a time. Note that by the triangle inequality and 1-boundedness

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^I} \left[(f^\# - f'')(x)g''(y)h''(z) \right] \right| \leq \|f^\# - f''\|_1 \leq \|f^\# - f''\|_2 \leq \sqrt{\frac{100\kappa}{\varepsilon\delta}} \leq \frac{\varepsilon}{100},$$

hence $\left| \mathbb{E}_{(x,y,z) \sim \mu^I} [f^\#(x)g''(y)h''(z)] \right| \geq \frac{\varepsilon}{4} - \frac{\varepsilon}{100}$. Continuing in this way we get that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^I} \left[f^\#(x)g^\#(y)h^\#(z) \right] \right| \geq \frac{\varepsilon}{4} - 3\frac{\varepsilon}{100} \geq \frac{\varepsilon}{8},$$

as required. \square

The functions $f^\#$, $g^\#$ and $h^\#$ are now only functions of $\sigma(x)$, $\gamma(y)$ and $\phi(z)$, and abusing notations we will also think of them as functions defined over H^I . As the distribution of $(\sigma(x), \gamma(y), \phi(z))$ is fully supported on $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$ we expect the 3-wise correlation

$$\mathbb{E}_{(x,y,z) \sim \mu^I} \left[f^\#(x)g^\#(y)h^\#(z) \right] = \mathbb{E}_{(x,y,z) \sim \mu^I} \left[f^\#(\sigma(x))g^\#(\gamma(y))h^\#(\phi(z)) \right]$$

to be related to the correlations of $f^\#$, $g^\#$ and $h^\#$. Indeed, examples for this include Roth's theorem [37, 29] as well as the analysis of the Blum-Luby-Rubinfeld linearity test [7, 1, 24]. However, as the distribution of $(\sigma(x), \gamma(y), \phi(z))$ is not uniform over $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$ our case is more closely related to the analysis of biased versions of this test and the straightforward Fourier-analytic argument does not work. This point was part of our motivation in [6], and the argument below generalizes that argument to our setting.

10.3 Shifting to the Uniform Distribution over an Equation on H

Fix I and \tilde{x}, \tilde{y} and \tilde{z} for which the event in Lemma 10.5 holds.

We may write $\mu = \frac{\alpha}{2}U + (1 - \frac{\alpha}{2})\mu'$ where U, μ' are distributions over $\Sigma \times \Gamma \times \Phi$, and sampling $(x, y, z) \sim U$, the distribution of $(\sigma(x), \gamma(y), \phi(z))$ is uniform over $\{(a, b, c) \in H^3 \mid a + b + c = 0\}$. We then take $J \subseteq_{\alpha/2} I$, $(\tilde{x}', \tilde{y}', \tilde{z}') \sim \mu'^{I \setminus J}$ and set

$$f^{\sharp'} = f^{\sharp}_{I \setminus J \rightarrow \tilde{x}'}, \quad g^{\sharp'} = g^{\sharp}_{I \setminus J \rightarrow \tilde{x}'}, \quad h^{\sharp'} = h^{\sharp}_{I \setminus J \rightarrow \tilde{x}'}$$

The following claim asserts that the triple correlation remains large with noticeable probability after passing to $f^{\sharp'}, g^{\sharp'}$ and $h^{\sharp'}$.

Claim 10.6. *In the above setting*

$$\Pr_{J, \tilde{x}', \tilde{y}', \tilde{z}'} \left[\left| \mathbb{E}_{(x, y, z) \sim U^J} \left[f^{\sharp'}(x) g^{\sharp'}(y) h^{\sharp'}(z) \right] \right| \geq \frac{\varepsilon}{16} \right] \geq \frac{\varepsilon}{16}.$$

Proof. Consider the real valued random variable $V(J, \tilde{x}', \tilde{y}', \tilde{z}') = \left| \mathbb{E}_{(x, y, z) \sim U^J} \left[f^{\sharp'}(x) g^{\sharp'}(y) h^{\sharp'}(z) \right] \right|$. Then by the triangle inequality

$$\mathbb{E}_{J, \tilde{x}', \tilde{y}', \tilde{z}'} [|V|] \geq \left| \mathbb{E}_{J, \tilde{x}', \tilde{y}', \tilde{z}'} \left[\mathbb{E}_{(x, y, z) \sim U^J} \left[f^{\sharp'}(x) g^{\sharp'}(y) h^{\sharp'}(z) \right] \right] \right| = \left| \mathbb{E}_{(x, y, z) \sim \mu^I} \left[f^{\sharp}(x) g^{\sharp}(y) h^{\sharp}(z) \right] \right| \geq \frac{\varepsilon}{8},$$

and as $0 \leq V \leq 1$ it follows by an averaging argument that $\Pr[V \geq \varepsilon/16] \geq \varepsilon/16$, as required. \square

Next, we have the standard Fourier analytic computation that handles 3-wise over U , showing that they can be significant only when $f^{\sharp'}$ has a significant Fourier coefficient.

Claim 10.7. *If $J, \tilde{x}', \tilde{y}'$ and \tilde{z}' satisfy the event in Claim 10.6, then there exists $\chi \in \hat{H}^J$ such that*

$$\left| \widehat{f^{\sharp'}}(\chi) \right| \geq \frac{\varepsilon}{16}.$$

Proof. Expand $f^{\sharp'}, g^{\sharp'}$ and $h^{\sharp'}$ into Fourier basis over H :

$$f^{\sharp'}(x) = \sum_{\chi \in \hat{H}^J} \widehat{f^{\sharp'}}(\chi) \chi(\sigma(x)), \quad g^{\sharp'}(y) = \sum_{\chi' \in \hat{H}^J} \widehat{g^{\sharp'}}(\chi') \chi'(\gamma(y)), \quad h^{\sharp'}(z) = \sum_{\chi'' \in \hat{H}^J} \widehat{h^{\sharp'}}(\chi'') \chi''(\phi(z)),$$

and plug this into the expectation to get

$$\mathbb{E}_{(x, y, z) \sim U^J} \left[f^{\sharp'}(x) g^{\sharp'}(y) h^{\sharp'}(z) \right] = \sum_{\chi, \chi', \chi'' \in \hat{H}^J} \widehat{f^{\sharp'}}(\chi) \widehat{g^{\sharp'}}(\chi') \widehat{h^{\sharp'}}(\chi'') \mathbb{E}_{(x, y, z) \sim U^J} \left[\chi(\sigma(x)) \chi'(\gamma(y)) \chi''(\phi(z)) \right].$$

Using the fact that $\phi(z) = -\sigma(x) - \gamma(y)$ to write the expectation on the right hand side as

$$\mathbb{E}_{(x, y, z) \sim U^J} \left[(\chi \overline{\chi''})(\sigma(x)) (\chi' \overline{\chi''})(\gamma(y)) \right] = 1_{\chi = \chi' = \chi''},$$

we conclude that

$$\left| \mathbb{E}_{(x,y,z) \sim U^J} [f^{\sharp'}(x)g^{\sharp'}(y)h^{\sharp'}(z)] \right| = \left| \sum_{\chi \in \hat{H}^J} \widehat{f^{\sharp'}}(\chi)\widehat{g^{\sharp'}}(\chi)\widehat{h^{\sharp'}}(\chi) \right| \leq \max_{\chi} |\widehat{f^{\sharp'}}(\chi)| \sum_{\chi \in \hat{H}^J} |\widehat{g^{\sharp'}}(\chi)| |\widehat{h^{\sharp'}}(\chi)|,$$

and using the fact that $\sum_{\chi \in \hat{H}^J} |\widehat{g^{\sharp'}}(\chi)| |\widehat{h^{\sharp'}}(\chi)| \leq 1$ which follows by Cauchy-Schwarz and Parseval, we get that

$$\frac{\varepsilon}{16} \leq \left| \mathbb{E}_{(x,y,z) \sim U^J} [f^{\sharp'}(x)g^{\sharp'}(y)h^{\sharp'}(z)] \right| \leq \max_{\chi} |\widehat{f^{\sharp'}}(\chi)|,$$

concluding the proof. \square

10.4 Unraveling Restrictions and Averaging: Applying the Restriction Inverse Theorem

Summarizing, we have shown that after a sequence of restrictions, averaging and further restriction, our function f has correlation with a function of the form $\chi \circ \sigma$ with noticeable probability. We next unravel this operations to deduce a result about f itself.

Claim 10.8. *Suppose that f^{\sharp} , g^{\sharp} and h^{\sharp} satisfy the event in Lemma 10.5. Then there is $\chi \in \hat{H}^I$ and an embedding function $L: \Sigma^I \rightarrow \mathbb{C}$ with $\|L\|_2 \leq 1$ and degree at most d , such that*

$$|\langle f^{\sharp}, L \cdot \chi \circ \sigma \rangle| \geq \varepsilon'$$

where $d = \text{poly}_{m,\alpha}(\log(1/\varepsilon))$ and $\varepsilon' = 2^{-\text{poly}_{m,\alpha}(\log(1/\varepsilon))}$.

Proof. By Claims 10.6 10.7, we conclude that

$$\Pr_{J, \tilde{x}'} [\exists \chi \in \hat{H}, |\langle f^{\sharp}_{I \setminus J \rightarrow \tilde{x}'}, \chi \circ \sigma \rangle| \geq \frac{\varepsilon}{16}] \geq \frac{\varepsilon}{16},$$

where we think of functions as being defined over H . Applying Theorem 10.1 the conclusion follows. \square

Note that by definition of f^{\sharp} ,

$$|\langle f^{\sharp}, L \cdot \chi \circ \sigma \rangle| = |\langle T_{\text{non-embed},0} f'', L \cdot \chi \circ \sigma \rangle| = |\langle f'', T_{\text{non-embed},0}(L \cdot \chi \circ \sigma) \rangle| = |\langle f'', L \cdot \chi \circ \sigma \rangle|,$$

where we used the fact that $T_{\text{non-embed},0}$ is Hermitian and that $L \cdot \chi \circ \sigma$ is an embedding function and so it is an eigenfunction of $T_{\text{non-embed},0}$ of eigenvalue 1 (see Fact 4.7). Thus, we conclude from Claim 10.8 that if f'' , g'' and h'' satisfy the event in Claim 10.4, then there are L and χ as in Claim 10.8 such that $|\langle f'', L \cdot \chi \circ \sigma \rangle| \geq \varepsilon'$.

Define $f''' = f'' \overline{\chi \circ \sigma}$, and observe that the above means that $|\langle f''', L \rangle| \geq \varepsilon'$ provided that f'' , g'' and h'' satisfy the event in Claim 10.4. Thus, by Cauchy-Schwarz $W_{\leq d}[f'1''] \geq |\langle f''', L \rangle|^2 \geq \varepsilon'^2$, and we now use random restrictions to show that after a suitable restriction, f''' has significant average.

Lemma 10.9. *Suppose that $F: (\Sigma^I, \mu^I) \rightarrow \mathbb{C}$ is a 1-bounded function. Then choosing (I', x') by including each $i \in I$ in I' with probability $1/2d$ and sampling $x' \sim \mu^{I \setminus I'}$, we have that*

$$\Pr_{I', x'} \left[|\mathbb{E}[F_{I \setminus I' \rightarrow x'}]| \geq \sqrt{\frac{W_{\leq d}[F]}{2e}} \right] \geq \frac{W_{\leq d}[F]}{2e}.$$

Proof. Define the random variable $V_{I',x'} = \mathbb{E}[F_{I \setminus I' \rightarrow x'}]$, and note that

$$\mathbb{E}_{I',x'}[|V_{I',x'}|^2] = \mathbb{E}_{I'} \left[\mathbb{E}_{x' \sim \mu^{I \setminus I'}} \left[\left| \sum_{S \neq \emptyset} F^{=S}(x') 1_{S \subseteq I \setminus I'} \right|^2 \right] \right] = \mathbb{E}_{I'} \left[\sum_S \|F^{=S}\|_2^2 1_{S \subseteq I \setminus I'} \right] \geq \frac{1}{e} W_{\leq d}[F].$$

As $|V_{I',x'}| \leq 1$ always, it follows that with probability at least $\frac{\xi}{2e}$ we have $|V_{I',x'}| \geq \sqrt{\frac{\xi}{2e}}$. \square

Applying Lemma 10.9 on $F = f'''$, we conclude that provided that f'' , g'' and h'' satisfy the event in Claim 10.4, we have that

$$\Pr_{I',x'} \left[\left| \mathbb{E}[f'''_{I \setminus I' \rightarrow x'}] \right| \geq \frac{\varepsilon'}{\sqrt{2e}} \right] \geq \frac{\varepsilon'^2}{2e}.$$

Noting that $\mathbb{E}[f'''_{I \setminus I' \rightarrow x'}] = \langle f''_{I \setminus I' \rightarrow x'}, \chi \circ \sigma |_{I \setminus I' \rightarrow x'} \rangle$, we conclude that after random restriction f'' is correlated with a function $\chi' \circ \sigma$ for $\chi' \in \hat{H}^{I'}$. Thus, we are now in a position again to apply the restriction inverse theorem and conclude the proof of Theorem 3.18.

Claim 10.10. *Theorem 3.18 is true.*

Proof. Fixing f, g, h and μ as in Theorem 3.18, we conclude by Claim 10.4 that

$$\Pr_{I,\tilde{x},\tilde{y},\tilde{z}} \left[\left| \mathbb{E}_{(x,y,z) \sim \mu^I} [f''(x)g''(y)h''(z)] \right| \geq \frac{\varepsilon}{4} \right] \geq \frac{\varepsilon}{4},$$

which by the above discussion implies that

$$\Pr_{\substack{I,\tilde{x},\tilde{y},\tilde{z} \\ I',x',y',z'}} \left[\exists \chi' \in \hat{H}^{I'}, \left| \langle f''_{I \setminus I' \rightarrow x'}, \chi' \circ \sigma \rangle \right| \geq \frac{\varepsilon'}{\sqrt{2e}} \right] \geq \frac{\varepsilon \varepsilon'^2}{4 \cdot 2e}.$$

Recalling the definition of f'' , this means that

$$\Pr_{\substack{I,\tilde{x},\tilde{y},\tilde{z} \\ I',x',y',z'}} \left[\exists \chi' \in \hat{H}^{I'}, \left| \langle f_{I \setminus I' \rightarrow x'}, \chi' \circ \sigma \rangle \right| \geq \frac{\varepsilon'}{\sqrt{2e}} \right] \geq \frac{\varepsilon \varepsilon'^2}{8e},$$

and we now apply Theorem 10.1. First, note that the restriction we are doing above fits the pattern therein; indeed, another way of viewing this restriction is as writing $\mu = \frac{\kappa}{2d}\mu + (1 - \frac{\kappa}{2d})\mu$. Theorem 10.1 now implies that there is $\chi \in \hat{H}^n$ and $L: \Sigma^n \rightarrow \mathbb{C}$ with $\|L\|_2 \leq 1$ and degree at most $d' = \varepsilon^{-O_{m,\alpha}(1)}$ such that $|\langle f, L \cdot \chi \circ \sigma \rangle| \geq \varepsilon''$ where

$$\varepsilon'' = 2^{-\text{poly}_{m,\alpha}\left(\frac{\log(1/\varepsilon')}{\kappa/2d}\right)} = 2^{-\text{poly}_{m,\alpha}(1/\varepsilon)},$$

concluding the proof. \square

11 Deducing the Structural Result for f : Proof of Theorem 1.4

In this section we use Theorem 3.18 to prove Theorem 1.4. At a high level, the argument proceeds as follows:

1. **Applying Theorem 3.18.** Starting with a distribution μ and functions f, g, h as in Theorem 1.4, we use path tricks (and more specifically, Lemma 3.17) to upper bound the 3-wise correlation of f, g and h over μ by the 3-wise correlation of functions F, G and H over a distribution μ' , where values of F correspond to product of values of f (as in Lemma 3.9) and G and H are arbitrary (but 1-bounded) functions. We are then in a position to apply Theorem 3.18 to conclude that the function F is correlated with a function of the desired form, namely $L \cdot \chi \circ \sigma'_{\text{master}}$ where L is a low-degree function, $\chi \in \hat{H}$ is a character and σ'_{master} is part of the master embedding of μ' .
2. **Unraveling products.** Ignoring the low-degree part for a moment, we have that the function F is a product of values of the function f , and as (by Lemma 3.10) the master embedding σ'_{master} can be written as by an alternating sum of values of σ_{master} (which is part of the master embedding of μ), we have that $\chi \circ \sigma'_{\text{master}}$ can also be written as product of values of $\chi \circ \sigma_{\text{master}}$. Combining these facts, we conclude that $F \cdot \chi \circ \sigma'_{\text{master}}$ can be written as product of values of $f \cdot \chi \circ \sigma_{\text{master}}$, and we know that the expectation of this value over some distribution is significant. In applying the path tricks appropriately, we have made sure that the distribution of points on which we take product over is good enough (and more precisely, connected) so that this is only possible whenever $f \cdot \chi \circ \sigma_{\text{master}}$ have significant mass on the low-levels (see Lemma 2.6). Thus, we are able to conclude that $f \cdot \chi \circ \sigma_{\text{master}}$ is correlated with a low-degree function, which is the result we are aiming for in Theorem 1.4.
3. **Applying the restriction inverse theorem.** To formalize this more precisely we must address the “ignoring the low-degree function L ” part of the argument. For that, we apply random restrictions. Intuitively, a low-degree function becomes constant after random restrictions, and thus we would indeed be able to deduce that after random restriction, F is correlated with a function of the form $\chi \circ \sigma_{\text{master}}$. The rest of the argument proceeds in the same way, and we indeed manage to conclude that random restrictions of $f \cdot \chi \circ \sigma_{\text{master}}$ are correlated with low-degree polynomials with noticeable probability. After further random restrictions we conclude that f is correlated with a function of the form $\chi \circ \sigma_{\text{master}}$, and to lift this information back to information about the function f itself we use the restriction inverse theorem, namely Theorem 10.1.

We now proceed to the formal argument.

11.1 Applying Theorem 3.18

Let f, g, h and μ be as in Theorem 1.4, and let $(\sigma_{\text{master}}, \gamma_{\text{master}}, \phi_{\text{master}})$ be a master embedding of μ into an Abelian group $(H, +)$ of size at most $r = r(m) \in \mathbb{N}$ as constructed in Lemma 3.7. By Lemma 3.17 we may apply the path trick on μ at most $T = T(m)$ times to get a distribution μ' over $\Sigma' \times \Gamma' \times \Phi'$, where $\Sigma' = \Sigma^{T_1}$, $\Gamma' \subseteq \Gamma^{T_2}$, $\Phi' \subseteq \Phi^{T_3}$ and T_1, T_2, T_3 are odd numbers that depend only on m . In addition, for future reference we will look at the marginal distribution of μ' on its first coordinate as $(x(1), \dots, x(T_1)) \sim \mu'$, and remark that by the construction in Lemma 3.17 it follows that the marginal distribution of each $x(j)$ is μ_x .

Following the evolution of the master embedding as in Definition 3.11, we see that a master embedding

of μ' ($\sigma'_{\text{master}}, \gamma'_{\text{master}}, \phi'_{\text{master}}$) is given by

$$\begin{aligned}\sigma'_{\text{master}}(x(1), \dots, x(T_1)) &= \sum_{j=1}^{T_1} (-1)^{j+1} \sigma_{\text{master}}(x(j)), \\ \gamma'_{\text{master}}(y(1), \dots, y(T_2)) &= \sum_{j=1}^{T_2} (-1)^{j+1} \gamma_{\text{master}}(y(j)), \\ \phi'_{\text{master}}(z(1), \dots, z(T_3)) &= \sum_{j=1}^{T_3} (-1)^{j+1} \phi_{\text{master}}(z(j)).\end{aligned}\tag{33}$$

Using Lemma 3.9 we that for $F: \Sigma'^n \rightarrow \mathbb{C}$ defined by

$$F(\vec{x}_1, \dots, \vec{x}_n) = \prod_{j=1}^{T_1} C^j \circ f(x(j)_1, \dots, x(j)_n),\tag{34}$$

where C^j represents the operation of applying complex conjugate if j is even and else applying the identity, there are 1-bounded functions $G: \Gamma'^n \rightarrow \mathbb{C}$ and $H: \Phi'^n \rightarrow \mathbb{C}$ such that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^n} [f(x)g(y)h(z)] \right|^M \leq \left| \mathbb{E}_{(X,Y,Z) \sim \mu'^n} [F(X)G(Y)H(Z)] \right|,$$

where $M = M(m) \in \mathbb{N}$. By the premise of Theorem 1.4 we get $|\mathbb{E}_{(X,Y,Z) \sim \mu'^n} [F(X)G(Y)H(Z)]| \geq \varepsilon^M$, hence we are in a position to apply Theorem 3.18. Using it, we get that there are $\chi \in \hat{H}^n$ and $L: \Sigma'^n \rightarrow \mathbb{C}$ such that:

1. $|\langle F, L \cdot \chi \circ \sigma'_{\text{master}} \rangle| \geq \varepsilon'$ where $\varepsilon' = 2^{-\text{poly}_{m,\alpha}(\frac{1}{\varepsilon})}$.
2. $\|L\|_2 \leq 1$.
3. L has degree at most $d = \text{poly}_{m,\alpha}(\frac{1}{\varepsilon})$.

11.2 Unraveling Products

We now get rid of the low-degree part L via random restrictions, and then use the product structure of F and of $\chi \circ \sigma'_{\text{master}}$ to convert the information regarding the correlation between them to information about the correlation between restrictions of f and restrictions of $\chi \circ \sigma_{\text{master}}$. Let \mathcal{D} be the marginal distribution of μ' on Σ' . Choose (I, \tilde{x}) a random restriction by including in I each $i \in [n]$ with probability $1/2d$ and taking $\tilde{x} \sim \mathcal{D}^I$. By Cauchy-Schwarz we have that

$$W_{\leq d}[F \cdot \overline{\chi \circ \sigma'_{\text{master}}}] \geq \left| \langle F \cdot \overline{\chi \circ \sigma'_{\text{master}}}, L \rangle \right|^2 = |\langle F, L \cdot \chi \circ \sigma'_{\text{master}} \rangle|^2 \geq \varepsilon'^2$$

hence by Lemma 10.9 it follows that

$$\Pr_{I, \tilde{x} \sim \mathcal{D}^I} \left[\left| \mathbb{E}[F|_{\tilde{I} \rightarrow \tilde{x}} \cdot \overline{\chi \circ \sigma'_{\text{master}}|_{\tilde{I} \rightarrow \tilde{x}}}] \right| \geq \frac{\varepsilon'}{\sqrt{2e}} \right] \geq \frac{\varepsilon'^2}{2e}.\tag{35}$$

Whenever this event holds we get that

$$|\langle F|_{\bar{I} \rightarrow \tilde{x}}, \chi \circ \sigma'_{\text{master}}|_{\bar{I} \rightarrow \tilde{x}} \rangle| \geq \frac{\varepsilon'}{2e}. \quad (36)$$

We denote $F' = F|_{\bar{I} \rightarrow \tilde{x}}$, and thus get that if I, \tilde{x} are such that (35) holds, then there is $\chi' \in \hat{H}^I$ such that $|\langle F', \chi' \circ \sigma'_{\text{master}} \rangle| \geq \frac{\varepsilon'}{\sqrt{2e}}$. Each coordinate of the restriction of \tilde{x} , namely each \tilde{x}_i , is an element in $\Sigma' = \Sigma^{T_1}$ and we will view it as a vector length T_1 , $\tilde{x}_i = (\tilde{x}(1)_i, \dots, \tilde{x}(T_1)_i)$. Thus, for each $j = 1, \dots, T_1$ we denote

$$\tilde{x}(j) = (\tilde{x}(j)_i)_{i \in \bar{I}},$$

With this notation, we get from (34) that

$$F'(x) = \prod_{j=1}^{T_1} C^j \circ f_{\bar{I} \rightarrow \tilde{x}(j)}(x(j)). \quad (37)$$

The following claim asserts that after further random restriction, with noticeable probability the function $f_{\bar{I} \rightarrow \tilde{x}(1)}$ is correlated with a character.

Claim 11.1. *Let I and \tilde{x} be such (36) holds. Then there is $D = O_{\alpha, m}(\log(1/\varepsilon'))$ such that choosing $J \subseteq I$ by including each element with probability $1/2D$ and sampling $\tilde{x}' \sim \mathcal{D}^{I \setminus J}$, we have that*

$$\Pr_{J, \tilde{x}'} \left[\exists \chi \in \hat{H}^J, \left| \left\langle f_{\bar{I} \rightarrow \tilde{x}(1)}|_{I \setminus J \rightarrow \tilde{x}'(1)}, \chi \circ \sigma_{\text{master}} \right\rangle \right| \geq \frac{\varepsilon'^2}{100} \right] \geq \frac{\varepsilon'^4}{1000}.$$

Proof. Note that by (37), we may write F' as

$$F'(x) = \prod_{j=1}^{T_1} C^j \circ f_j(x(j)), \quad (38)$$

where $f_j = f_{\bar{I} \rightarrow \tilde{x}(j)}$. By (33) we have

$$\sigma'_{\text{master}}(x) = \sum_{j=1}^{T_1} (-1)^{j+1} \sigma'_{\text{master}}(x(j)),$$

and so

$$\chi \circ \sigma'_{\text{master}}(x) = \prod_{j=1}^{T_1} C^{j+1} \circ \chi \circ \sigma'_{\text{master}}(x(j)). \quad (39)$$

Plugging (38) and (39) into (36) yields that

$$\left| \mathbb{E}_{(x, x') \sim \mathcal{D}^I} \left[\prod_{j=1}^{T_1} C^j \circ f_j(x(j)) \cdot C^{j+1} \circ \chi \circ \sigma'_{\text{master}}(x(j)) \right] \right| \geq \frac{\varepsilon'}{\sqrt{2e}}.$$

Let $\tilde{\Sigma} = \Sigma^{T_1-1}$ and define $Q: \tilde{\Sigma}^I \rightarrow \mathbb{C}$ by

$$Q(w) = \prod_{j=2}^{T_1} C^j \circ f_j(w(j)) \cdot C^{j+1} \cdot \chi \circ \sigma'_{\text{master}}(w(j)),$$

then we get that $\left| \mathbb{E}_{(x,w) \sim \nu^I} [f_1(x) \cdot C \circ \chi \circ \sigma'_{\text{master}}(x) \cdot Q(w)] \right| \geq \frac{\varepsilon'}{\sqrt{2e}}$, where ν is the marginal distribution of \mathcal{D} on the first coordinate viewed as an element in $\Sigma \times \tilde{\Sigma}$. Consider the distribution ν' over Σ^2 where we sample $(x, w) \sim \nu$, then $(x', w') \sim \nu$ conditioned on $w' = w$, and then output (x, x') . By Cauchy-Schwarz it follows that

$$\begin{aligned} \frac{\varepsilon'^2}{2e} &\leq \left| \mathbb{E}_{(x,w) \sim \nu^I} [f_1(x) \cdot C \circ \chi \circ \sigma'_{\text{master}}(x) \cdot Q(w)] \right|^2 \\ &\leq \left| \mathbb{E}_{(x,x') \sim \nu^I} \left[f_1(x) \cdot C \circ \chi \circ \sigma'_{\text{master}}(x) \cdot \overline{f_1(x') \cdot C \circ \chi \circ \sigma'_{\text{master}}(x')} \right] \right|^2 \\ &= \left| \langle f_1 \cdot C \circ \chi \circ \sigma_{\text{master}}, R^I(f_1 \cdot C \circ \chi \circ \sigma_{\text{master}}) \rangle \right|, \end{aligned}$$

where $R: L_2(\Sigma; \nu_x) \rightarrow L_2(\Sigma; \nu_x)$ is the averaging operator corresponding to ν' defined by $Rp(a) = \mathbb{E}_{(x,x') \sim \nu} [p(x') \mid x = a]$.

As the support of μ' on the first coordinate is full we get that the support of ν' on Σ^2 is full, hence R is connected. Also, the probability of each atom is at least $\alpha' = \alpha'(\alpha, m) > 0$, so by Lemma 2.6 we conclude that for $D = O_{\alpha, m}(\log(1/\varepsilon'))$ we have that $W_{\leq D}[f_1 \chi \circ \sigma_{\text{master}}] \geq \frac{\varepsilon'^4}{4e^2}$. The proof is now concluded by Lemma 10.9. \square

11.3 Applying the Restriction Inverse Theorem

Combining (35) and Claim 11.1 gives that

$$\Pr_{\substack{I, \tilde{x} \sim \mathcal{D}^I \\ J, \tilde{x}' \sim \mathcal{D}^{I \setminus J}}} \left[\exists \chi \in \hat{H}^J, \left| \langle f \Big|_{\substack{I \rightarrow \tilde{x}(1) \\ I \setminus J \rightarrow \tilde{x}'(1)}}, \chi \circ \sigma''_{\text{master}} \rangle \right| \geq \frac{\varepsilon'^2}{100} \right] \geq \frac{\varepsilon'^2}{2e} \cdot \frac{\varepsilon'^4}{1000} = \frac{\varepsilon'^6}{2000e}.$$

We now appeal to the restriction inverse theorem, Theorem 10.1, to finish the proof. Towards this end, we note that the sequence of restrictions can be viewed as a standard restriction: note that the distribution of $\tilde{x}(1)$ where $\tilde{x} \sim \mathcal{D}$ is μ_x , so the above restriction amounts to choosing $K \subseteq [n]$ by including each element with probability $\frac{1}{4dD}$, and then restricting the coordinates outside K to be $x' \sim \mu_x^K$. Thus, we may appeal Theorem 10.1, and the result follows.

12 The Restriction Inverse Theorem

The main goal of this section is to prove the restriction inverse theorem, Theorem 10.1. We begin by presenting a few notions that will be necessary for the statement of the theorem, give a formal statement of a slight generalization of Theorem 10.1 and then present some tools necessary for the proof. Finally, in Section 12.4 we give the formal proof.

12.1 Product Functions and Classes of Product Functions

The restriction inverse theorem is concerned with functions that, after random restriction, are correlated with product functions, defined as follows.

Definition 12.1. We say $f': \Sigma^n \rightarrow \mathbb{C}$ is a product function if there are functions $p_1, \dots, p_n: \Sigma \rightarrow \mathbb{C}$ that are 1-bounded such that

$$f(x_1, \dots, x_n) = \prod_{i=1}^n p_i(x_i).$$

We denote by $\mathcal{P}(n, \Sigma)$ the collection of all product functions over Σ^n , and denote $\mathcal{P}(\Sigma) = \bigcup_{n \in \mathbb{N}} \mathcal{P}(n, \Sigma)$.

We will need the notion of a class of product functions, which is a sub-collection of functions closed under restrictions. For technical reasons, this closure will be up to multiplying by a complex number of absolute value 1.

Definition 12.2. We say $\mathcal{F}(\Sigma) \subseteq \mathcal{P}(\Sigma)$ is a class of product functions if for all $f' \in \mathcal{F}(\Sigma)$, say $f': \Sigma^n \rightarrow \mathbb{C}$, for all $I \subseteq [n]$ and for all $x \in \Sigma^I$ it holds that there is $\theta \in \mathbb{C}$ of absolute value 1 such that $\theta f'_{I \rightarrow x} \in \mathcal{F}(\Sigma)$.

We also need the notion of separateness of product functions. Intuitively, this says that any two univariate functions in the class either have correlation 1, or else the correlation is bounded away from 1.

Definition 12.3. Suppose Σ is a finite alphabet, μ is a distribution over Σ and consider the inner product spaces $L_2(\Sigma^n, \mu^{\otimes n})$. For $\tau > 0$, we say a collection of product functions $\mathcal{F}(\Sigma) \subseteq \mathcal{P}(\Sigma)$ is τ -separated if for any uni-variate functions $p, p' \in \mathcal{F}(\Sigma)$ it is either the case that $p = p'$, or else $|\langle p, p' \rangle| \leq 1 - \tau$.

An important class of functions for us will be the class of functions arising from Abelian embeddings. Suppose that we have an Abelian group $(H, +)$ and a map $\sigma: \Sigma \rightarrow H$. Then in this setting, we may define the collection

$$\mathcal{F}_\sigma(\Sigma, n) = \left\{ f(x_1, \dots, x_n) = \prod_{i=1}^n \chi_i(\sigma(x_i)) \mid \chi_i \in \hat{H} \forall i = 1, \dots, n \right\}.$$

The following fact asserts that the collection $\mathcal{F}_\sigma(\Sigma) = \bigcup_{n \in \mathbb{N}} \mathcal{F}_\sigma(\Sigma, n)$ is a class of product functions which is separated.

Fact 12.4. For all $m \in \mathbb{N}$ and $\alpha > 0$ there is $\tau > 0$ such that the following holds. If Σ is an alphabet of size at most m and μ is a distribution over Σ in which the probability of each atom is at least α , then collection $\mathcal{F}_\sigma(\Sigma)$ is a class of product functions which is τ -separated.

Proof. It is clear that each function in $\mathcal{F}_\sigma(\Sigma)$ is a product function. Also, when we restrict a set of variables, the corresponding terms give a constant factor θ with absolute value 1, hence $\mathcal{F}_\sigma(\Sigma)$ is closed under restrictions.

For the τ -separatedness, fix univariate functions $p, p' \in \mathcal{F}_\sigma(\Sigma)$ and suppose that $p \neq p'$. We note that $|\langle p, p' \rangle| < 1$: otherwise, Cauchy-Schwarz would be tight, hence p and p' would be proportional. However, as $0 \in H$ is in the image of σ_{master} we may find x such that $\sigma_{\text{master}}(x) = 0$, and so $p(x) = 1 = p'(x)$, and in conjunction with the fact they are proportional we would get that $p \equiv p'$.

Multiplying p by a constant of absolute value 1, we may assume that $\langle p, p' \rangle \geq 0$, hence $\langle p, p' \rangle = 1 - \frac{1}{2} \|p - p'\|_2^2$. By definition of \hat{H} , we may find $T = T(m) > 0$ and λ of absolute value 1 such that the values of p and p' take the forms $\lambda e^{\frac{L}{T} 2\pi i}$ and $e^{\frac{L'}{T} 2\pi i}$ respectively, where L, L' are integers.

As p and p' are not proportional there are $a, b \in \Sigma$ such that $p(a) \neq p'(a)$ and b such that $p(b) \neq e^{-\frac{1}{2T} 2\pi i} p'(b)$. Writing $\lambda = e^{2\pi i \theta}$ we have

$$p(a) = e^{\frac{L_1 + \theta T}{T} 2\pi i}, \quad p'(a) = e^{\frac{L'_1}{T} 2\pi i}, \quad p(b) = e^{\frac{L_2 + \theta T}{T} 2\pi i}, \quad p'(b) = e^{\frac{L'_2 + 1/2}{T} 2\pi i}$$

for integers L_1, L_2, L_3, L_4 . Then

$$\|p - p'\|_2^2 \geq \alpha |p(a) - p'(a)|^2 + \alpha |p(b) - p'(b)|^2,$$

and we argue that the right hand side is at least $\Omega(\alpha/T^2)$. Indeed,

$$|p(a) - p'(a)| = \left| e^{\frac{L_1 - L_1 - \theta T}{T} 2\pi i} - 1 \right|, \quad |p(b) - p'(b)| = \left| e^{\frac{L_2 - L_2 - \theta T + 1/2}{T} 2\pi i} - 1 \right|,$$

so if both are at most $\frac{1}{100T}$ then each one of $-\theta T$ and $-\theta T + \frac{1}{2}$ is $\frac{1}{10}$ -close to be an integer, but by the triangle inequality this is impossible. \square

12.2 Statement of the Restriction Inverse Theorem

With the above set up, we are now ready to state the restriction inverse theorem.

Theorem 12.5. *For all $\alpha, \tau > 0$, $m \in \mathbb{N}$, $\rho \in (0, 1)$ and $\varepsilon > 0$ there are $d \in \mathbb{N}$ and $\varepsilon' > 0$ such that the following holds. Suppose that Σ is an alphabet of size at most m and μ, ν, ν' are a distribution over Σ in which the probability of each atom is at least α and $\mu = \rho\nu + (1 - \rho)\nu'$. Suppose further that $\mathcal{F}(\Sigma)$ is a class of product functions that is τ separated.*

If $f: (\Sigma^n, \mu^{\otimes n}) \rightarrow \mathbb{C}$ is a 1-bounded function such that

$$\Pr_{I \subseteq_\rho [n], \tilde{x} \sim \mu^{\tilde{I}}} [\exists f' \in \mathcal{F}(\Sigma), |\langle f_{\tilde{I} \rightarrow \tilde{x}}, f' \rangle| \geq \varepsilon] \geq \varepsilon,$$

then there exist $f' \in \mathcal{F}(\Sigma)$, as well as $L: \Sigma^n \rightarrow \mathbb{C}$ a function of degree at most d and $\|L\|_2 \leq 1$, such that

$$|\langle f, Lf' \rangle| \geq \varepsilon'.$$

Quantitatively, $d = \text{poly}_{m, \alpha, \tau} \left(\frac{\log(1/\varepsilon)}{\rho} \right)$ and $\varepsilon' = 2^{-\text{poly}_{m, \alpha, \tau} \left(\frac{\log(1/\varepsilon)}{\rho} \right)}$.

We note that Theorem 12.5 immediately implies Theorem 10.1:

Claim 12.6. *Theorem 12.5 implies Theorem 10.1.*

Proof. Using Fact 12.4, the collection $\mathcal{F}_\sigma(\Sigma)$ satisfies the properties required by Theorem 12.5, and so applying Theorem 12.5 on it gives translates to the statement of Theorem 10.1. \square

12.3 Tools for the Proof of Theorem 12.5

In this section we give the key ingredients for the proof of Theorem 12.5. Throughout this section, we denote by $[R]^{\leq n}$ the set of vectors over $[R]$ of length at most n .

12.3.1 The Direct Product Theorem

First, we need a suitable direct product result. A function $F: (\{0, 1\}^n, \mu_\rho^{\otimes n}) \rightarrow [R]^{\leq n}$ is called a direct product function if there is $f: [n] \rightarrow [R]$ for which $F[A] = f|_A$ for all A . Here and throughout, μ_ρ represents the ρ -biased distribution, in which $A \sim \mu_\rho^{\otimes n}$ is sampled by including each element independently with probability ρ . The goal in direct product testing is to design a randomized test that queries a few locations at the function F , performs a test on them and then accept or reject accordingly. The tester should have the following properties:

1. **Completeness:** if F is a direct product function, the tester must accept with probability 1.
2. **Soundness:** if the tester accepts with noticeable probability, then F is somewhat correlated with a direct product function. By that, ideally one would like to say that if the tester accepts with probability ε , then there is a direct product function G such that $F[A] = G[A]$ with probability at least $\delta = \delta(\varepsilon) > 0$ over $A \sim \mu_\rho^{\otimes n}$. We will not be able to guarantee that type of soundness, and instead settle for something slightly weaker: with probability at least δ over the choice of A , we have that $\Delta(F[A], G[A]) \leq r$, where $\Delta(x, y)$ measures the Hamming distance between two strings $x, y \in [R]^k$ and $r = r(\varepsilon) \in \mathbb{N}$. In words, $F[A]$ and $G[A]$ agree on all but constantly many coordinates.⁸

For our application we need to consider a particular test, which is also the most natural direct product tester one may think of. Given parameters $\rho, \alpha, \beta \in (0, 1)$ and oracle access to $F: (\{0, 1\}^n, \mu_\rho^{\otimes n}) \rightarrow [R]^{\leq n}$ thought of as mapping A to $F[A] \in [R]^A$, perform the following test, referred to as the $\text{DP}(\rho, \alpha, \beta)$ test:

1. Sample $C \sim \mu_{\rho\alpha}^{\otimes n}$ and sample $A, B \sim \mu_\rho^{\otimes n}$ independently conditioned on $A, B \supseteq C$.
2. Sample $T \subseteq_\beta [n]$.
3. Check that $F[A]|_{C \cap T} = F[B]|_{C \cap T}$.

Roughly speaking, we sample A and B that have $\rho\alpha$ of their elements in common. Then, we choose a subset $T \subseteq_\beta [n]$ so that $T \cap C$ contains roughly β fraction of their shared elements, and check that the assignments $F[A]$ and $F[B]$ are consistent on $C \cap T$. It is clear that the tester $\text{DP}(\rho, \alpha, \beta)$ has the completeness property, and the following result addresses the soundness of $\text{DP}(\rho, \alpha, \beta)$:

Theorem 12.7. *For all $C > 0$ there is $c > 0$ such that the following holds for sufficiently large n and $\varepsilon \geq 2^{-n^c}$. Suppose that $\frac{1}{\log(1/\varepsilon)^C} \leq \alpha, \beta \leq \frac{9}{10}$ and $\rho \in (0, 1)$; then there are $r \in \mathbb{N}$ and $\varepsilon' > 0$ such that if $F: (\{0, 1\}^n, \mu_\rho^{\otimes n}) \rightarrow [R]^{\leq n}$ is an assignment as above that passes the $\text{DP}(\rho, \alpha, \beta)$ test with probability at least ε , then there exists $f: [n] \rightarrow [R]$ such that*

$$\Pr_{A \sim \mu_\rho^{\otimes n}} [\Delta(F[A], f|_A) \leq r] \geq \varepsilon'.$$

Quantitatively, we have $r = \rho^{-O_C(1)} \log(1/\varepsilon)^{O_C(1)}$ and $\varepsilon' = \varepsilon^{O_C(\log(1/\rho)^2)}$.

Proof. Deferred to Section 13. □

12.3.2 Stability and Level d Inequalities

We record here a few basic notions from analysis of Boolean functions over product domains that we need; we refer the reader to [32] for details. First is the notion of noise stability, for which we first define the standard noise operator:

Definition 12.8. *For a finite probability space (Σ, μ) , $\rho \in [0, 1]$ and $x \in \Sigma$, we define the distribution over ρ -correlated inputs with x , denoted by $\mathbb{T}_\rho x$, to be: take $y = x$ with probability ρ , otherwise sample $y \sim \mu$.*

⁸We remark that this is the best soundness one may hope to get in general. Indeed, for essentially all direct product testers and in particular the one we consider, it is the case that slight perturbations of legitimate direct product functions pass the test with noticeable probability. Namely, taking a direct product function F and taking F' such that $\Delta(F[A], F'[A]) \leq r$ for at least ε fraction of the A 's, one typically has that F' passes the direct product tester with probability at least $\varepsilon^2 2^{-O(r)}$.

As usual, we may associate with T_ρ an averaging operator $T_\rho: L_2(\Sigma, \mu) \rightarrow L_2(\Sigma, \mu)$, as well as tensor it to get an operator acting on n -variate functions.

Definition 12.9. For a finite probability space (Σ, μ) , $\rho \in [0, 1]$ and $f: \Sigma^n \rightarrow \mathbb{C}$, we define the ρ -noise stability of f as $\text{Stab}_\rho = \langle f, T_\rho f \rangle$.

The following fact asserts that if a function f has significant noise stability, then it has significant weight on the low-levels.

Fact 12.10. Suppose that (Σ, μ) is a finite probability space, $f: \Sigma^n \rightarrow \mathbb{C}$ is a function with 2-norm at most 1 and $\text{Stab}_{1-\varepsilon}(f) \geq \delta$. Then $W_{\leq \frac{2 \log(1/\delta)}{\varepsilon}}[f] \geq \frac{\delta}{2}$.

Proof. Writing $f = \sum_{S \subseteq [n]} f^{=S}$ according to the Efron Stein decomposition on $(\Sigma^n, \mu^{\otimes n})$ and noting that $T_\rho^{\otimes n} f^{=S} = \rho^{|S|} f^{=S}$, we get that

$$\text{Stab}_{1-\varepsilon}(f) = \sum_{S \subseteq [n]} (1-\varepsilon)^{|S|} \|f^{=S}\|_2^2.$$

The contribution from $|S| > 2 \log(1/\delta)/\varepsilon$ is at most

$$(1-\varepsilon)^{2 \log(1/\delta)/\varepsilon} \sum_{|S| > 2 \log(1/\delta)/\varepsilon} \|f^{=S}\|_2^2 \leq e^{-2 \log(1/\delta)} \|f\|_2^2 \leq \frac{\delta}{2},$$

and the contribution from $|S| \leq 2 \log(1/\delta)/\varepsilon$ is at most $W_{\leq 2 \log(1/\delta)/\varepsilon}[f]$. It follows that $W_{\leq 2 \log(1/\delta)/\varepsilon}[f] + \frac{\delta}{2} \geq \delta$, and the proof is concluded by re-arranging. \square

The following fact is known as the level d inequality; it asserts that a Boolean function with small average may only have very small weight on low levels.

Fact 12.11. Let Σ be a finite alphabet and let μ be a distribution over Σ in which the probability of each atom is at least α . If $F: \Sigma^n \rightarrow \{0, 1\}$ is a function with $\mathbb{E}_\mu[F] = s$, then $W_{\leq d}[F] \leq 2^{O_\alpha(d)} s^{3/2}$ for all d .

Proof. This is a standard consequence of the hypercontractive inequality, asserting that there is $C(\alpha) > 0$ such that $\|g\|_4 \leq C(\alpha)^d \|g\|_2$ for all functions g of degree at most d ; see [32, Theorem 10.21]. Thus, by Hölder's inequality

$$W_{\leq d}[F] = \langle F^{\leq d}, F \rangle \leq \|F^{\leq d}\|_4 \|F\|_{4/3} \leq C(\alpha)^d \|F^{\leq d}\|_2 \|F\|_{4/3} = C(\alpha)^d \sqrt{W_{\leq d}[F]} \|F\|_{4/3}$$

and re-arranging gives $W_{\leq d}[F] \leq C(\alpha)^{2d} \|F\|_{4/3}^2 = C(\alpha)^{2d} s^{3/2}$. \square

12.3.3 Some Averaging Arguments

Our argument makes use of several standard probabilistic facts which we collect here. The first of which asserts that if we have functions f and g that are somewhat correlated and g is 1-bounded, then with noticeable probability they remain somewhat correlated after a random restriction.

Fact 12.12. Let Σ be a finite alphabet, μ be a distribution over Σ and let $f, g: \Sigma^n \rightarrow \mathbb{C}$ be functions such that g is 1-bounded and $\|f\|_2 \leq 1$. For all $I \subseteq [n]$, if $|\langle f, g \rangle| \geq \eta$, then sampling $\tilde{x} \sim \mu^{\bar{I}}$, we have that

$$\Pr_{\tilde{x} \sim \mu^{\bar{I}}} \left[|\langle f_{\bar{I} \rightarrow \tilde{x}}, g_{\bar{I} \rightarrow \tilde{x}} \rangle| \geq \frac{\eta}{2} \right] \geq \frac{\eta^2}{4}.$$

Proof. Denote the event in question by E , and note that

$$\mathbb{E}_{\tilde{x}} [|\langle f_{\bar{T} \rightarrow \tilde{x}}, g_{\bar{T} \rightarrow \tilde{x}} \rangle|] \geq \left| \mathbb{E}_{\tilde{x}} [\langle f_{\bar{T} \rightarrow \tilde{x}}, g_{\bar{T} \rightarrow \tilde{x}} \rangle] \right| = |\langle f, g \rangle| \geq \eta.$$

On the other hand,

$$\mathbb{E}_{\tilde{x}} [|\langle f_{\bar{T} \rightarrow \tilde{x}}, g_{\bar{T} \rightarrow \tilde{x}} \rangle|] = \underbrace{\mathbb{E}_{\tilde{x}} [1_E |\langle f_{\bar{T} \rightarrow \tilde{x}}, g_{\bar{T} \rightarrow \tilde{x}} \rangle|]}_{(I)} + \underbrace{\mathbb{E}_{\tilde{x}} [1_{\bar{E}} |\langle f_{\bar{T} \rightarrow \tilde{x}}, g_{\bar{T} \rightarrow \tilde{x}} \rangle|]}_{(II)},$$

and we upper bound each term as follows. For (I) we have by Cauchy-Schwarz that

$$|(I)| \leq \sqrt{\Pr[E]} \sqrt{\mathbb{E}_{\tilde{x}} [|\langle f_{\bar{T} \rightarrow \tilde{x}}, g_{\bar{T} \rightarrow \tilde{x}} \rangle|^2]} \leq \sqrt{\Pr[E]} \sqrt{\mathbb{E}_{\tilde{x}} [\|f_{\bar{T} \rightarrow \tilde{x}}\|_2^2]} = \sqrt{\Pr[E]} \|f\|_2 \leq \sqrt{\Pr[E]}.$$

Here, we used the fact that g is 1-bounded. For (II) we have by definition of \bar{E} we have $|(II)| \leq \frac{\eta}{2}$. Combining, we get that $\eta \leq \sqrt{\Pr[E]} + \frac{\eta}{2}$, hence $\Pr[E] \geq \frac{\eta^2}{4}$. \square

The second fact asserts that if X, Y are independent random variables and $E(X, Y)$ is an event with noticeable probability, then sampling $x_1, \dots, x_k \sim X$ and $y_1, \dots, y_\ell \sim Y$ independently we have that all of $E(x_i, y_j)$ occur with noticeable probability.

Fact 12.13. *Let X, Y be independent random variables and let E be an event depending only on X, Y , and suppose that $\Pr[E] \geq \delta$. Then*

$$\Pr_{\substack{x_1, \dots, x_k \sim X \\ y_1, \dots, y_\ell \sim Y}} \left[\bigcap_{i,j} E(x_i, y_j) \right] \geq \delta^{k\ell}.$$

Proof. We have that $\mathbb{E}_{x \sim X, y \sim Y} [1_{E(x,y)}] = \Pr[E] \geq \delta$, so raising to the power k and using Hölder's inequality gives

$$\begin{aligned} \delta^k &\leq \mathbb{E}_{x \sim X, y \sim Y} [1_{E(x,y)}]^k = \mathbb{E}_{y \sim Y} \left[\mathbb{E}_{x \sim X} [1_{E(x,y)}] \right]^k \leq \mathbb{E}_{y \sim Y} \left[\mathbb{E}_{x \sim X} [1_{E(x,y)}^k] \right] \\ &= \mathbb{E}_{y \sim Y} \left[\mathbb{E}_{x_1, \dots, x_k \sim X} [1_{\bigcap_{i=1}^k E(x_i, y)}] \right]. \end{aligned}$$

Raising to the power ℓ and using Hölder's inequality again gives

$$\begin{aligned} \delta^{k\ell} &\leq \mathbb{E}_{y \sim Y} \left[\mathbb{E}_{x_1, \dots, x_k \sim X} [1_{\bigcap_{i=1}^k E(x_i, y)}] \right]^\ell = \mathbb{E}_{x_1, \dots, x_k \sim X} \left[\mathbb{E}_{y \sim Y} [1_{\bigcap_{i=1}^k E(x_i, y)}] \right]^\ell \\ &\leq \mathbb{E}_{x_1, \dots, x_k \sim X} \left[\mathbb{E}_{y \sim Y} [1_{\bigcap_{i=1}^k E(x_i, y)}]^\ell \right] \\ &= \mathbb{E}_{x_1, \dots, x_k \sim X} \left[\mathbb{E}_{y_1, \dots, y_\ell \sim Y} [1_{\bigcap_{i,j} E(x_i, y_j)}] \right], \end{aligned}$$

as desired. \square

12.4 Proof of Theorem 12.5

We now proceed to the proof of Theorem 12.5. The argument we present is similar to an argument from [6] with some differences. In the setting therein, the class of product functions was the collection of all multiplicative characters over \mathbb{F}_2 and the underlying measure was uniform, hence any two distinct product functions were orthogonal. Such orthogonality properties were used multiple times, and most importantly it implies that a given function $f: \Sigma^n \rightarrow \mathbb{C}$ with 2-norm at most 1 could be η -correlated with at most $\frac{1}{\eta^2}$ product functions. In the current setting we do not have these orthogonality properties. To circumvent that, we consider nets, which are small collections of product functions that in some sense capture all of the product functions correlated with f . Stated simply, while there could be many product functions that are correlated with f (their number could depend on the dimension n for example), we argue that one could choose a short list of product functions that are correlated with f , so that any other product function that is correlated with f must be close to a product function from the list. With this change, we can use the main ideas from argument in [6] (which still requires some non-trivial but more minor adaptations).

Throughout this section we will use the following notation:

Definition 12.14. For a product function $p: \Sigma^n \rightarrow \mathbb{C}$ given as $p(x_1, \dots, x_n) = \prod_{i=1}^n p_i(x_i)$ and $T \subseteq [n]$, we define $p|_T: \Sigma^T \rightarrow \mathbb{C}$ by $p|_T(y) = \prod_{i \in T} p_i(y_i)$.

Next, we define the action of this operation on a collection of functions in the natural way.

Definition 12.15. For a collection of product functions $W \subseteq \{p: \Sigma^n \rightarrow \mathbb{C} \mid p \text{ is a product function}\}$ and $T \subseteq [n]$ we define $W|_T = \{p|_T \mid p \in W\}$.

We also note that by the τ -separatedness of \mathcal{F} , it follows that the number of univariate functions in $\mathcal{F}(\Sigma)$ is at most some finite number $R = R(\tau)$, and we fix this R henceforth. We will identify between $[R]$ and univariate functions in $\mathcal{F}(\Sigma)$, and thus also between n -variate product functions in $\mathcal{F}(\Sigma)$ and $[R]^n$.

12.4.1 The Net of Product Functions and the Symbolic Distance

Fix Σ , a distribution μ over Σ and a class $\mathcal{F} = \mathcal{F}(\Sigma)$ as in Theorem 12.5. We will consider various domains, and to make the notations more precise we will denote by $\mathcal{F}(\Sigma, I)$ the subset of \mathcal{F} consisting of functions whose domain is Σ^I .

For a function $f: (\Sigma^n, \mu^n) \rightarrow \mathbb{C}$ and parameter $\varepsilon > 0$, denote

$$\text{List}_\varepsilon[f] = \{p \in \mathcal{F} \mid |\langle f, p \rangle| \geq \varepsilon\}.$$

As discussed earlier, there need not be a bound on the size of $\text{List}_\varepsilon[f]$ in terms of ε ; it may well be the case that its size grows with the dimension n . Nevertheless, the following lemma states that within $\text{List}_\varepsilon[f]$, we may find a short list whose size is bounded in terms of ε that essentially captures all of $\text{List}_\varepsilon[f]$.

Lemma 12.16. Fix Σ , a distribution μ over Σ and a class $\mathcal{F} = \mathcal{F}(\Sigma)$ as in Theorem 12.5, let $f: (\Sigma^n, \mu^n) \rightarrow \mathbb{C}$ be such that $\|f\|_2 \leq 1$, and let $\varepsilon, \delta > 0$ be parameters satisfying that $\delta < \varepsilon^2$. Then one may find $\text{ShortList}_{\varepsilon, \delta}[f] \subseteq \text{List}_\varepsilon[f]$ such that:

1. $|\text{ShortList}_{\varepsilon, \delta}[f]| \leq \frac{1}{\varepsilon^2 - \delta}$.
2. For all $p \in \text{List}_\varepsilon[f]$ there is $p' \in \text{ShortList}_{\varepsilon, \delta}[f]$ such that $|\langle p, p' \rangle| \geq \delta$.

Proof. The proof is by a greedy algorithm. Starting with $\text{ShortList}_{\varepsilon, \delta}[f] = \emptyset$, so long as there is $p \in \text{List}_{\varepsilon}[f]$ that is at most δ -correlated with all functions in $\text{ShortList}_{\varepsilon, \delta}[f]$, we add it to $\text{ShortList}_{\varepsilon, \delta}[f]$.

We show that the above process terminates after less than $k = \frac{1}{\varepsilon^2 - \delta}$ steps. Indeed, otherwise we would be able to find $p_1, \dots, p_{k+1} \in \text{List}_{\varepsilon}[f]$ whose pairwise correlations are at most δ . Write $\langle f, p_j \rangle = \theta_j \rho_j$ where $\rho_j \geq 0$ and θ_j has absolute value 1; then we have that $\rho_j \geq \varepsilon$. We get that

$$k\varepsilon \leq \sum_{j=1}^k \theta_j^{-1} \langle f, p_j \rangle = \langle f, \sum_{j=1}^k \theta_j^{-1} p_j \rangle \leq \|f\|_2 \left\| \sum_{j=1}^k \theta_j^{-1} p_j \right\|_2.$$

Using $\|f\|_2 \leq 1$ and the fact that

$$\left\| \sum_{j=1}^{k+1} \theta_j^{-1} p_j \right\|_2^2 = \sum_{j=1}^k \|p_j\|_2^2 + \sum_{j \neq j'} \|p_j\|_2 \|p_{j'}\|_2 \leq k + k(k-1)\delta$$

gives $k\varepsilon \leq \sqrt{k + k^2\delta}$, and simplifying finishes the proof. \square

Note that the short-list found in Lemma 12.16 is not unique, and indeed there may be several choices for it. We will want to think of some canonical short-list that is associated with a given function f . This may be achieved in several ways; for instance, we may fix a total ordering among all product functions, and then consider some induced ordering (say, lexicographic ordering) induced on collections of functions, and define the canonical short list of f to be the short list which is first according to this ordering. Thus, henceforth when we write $\text{ShortList}_{\varepsilon, \delta}[f]$, we refer to the canonical short list of f .

Lemma 12.16 gives us a rather satisfactory answer in the sense that we get a short list of product functions that, in a sense, encapsulates within it the entire list of product functions correlated with f . We will want to imagine this short list as the center of Hamming balls, and of this property as saying that any product function correlated with f is inside a ball of small radius around some product function from the short list. To facilitate that, we define the symbolic distance between product functions.

Definition 12.17. Fix Σ , a distribution μ over Σ and a class $\mathcal{F} = \mathcal{F}(\Sigma)$ as in Theorem 12.5. For two functions $p, p' \in \mathcal{F}$ over n -variables $p, p' : \Sigma^n \rightarrow \mathbb{C}$ written as $p(x) = \prod_{j=1}^n p_j(x_j)$ and $p'(x) = \prod_{j=1}^n p'_j(x_j)$, we define the symbolic distance between p and p' as:

$$\Delta_{\text{symbolic}}(p, p') = |\{j \in [n] \mid p_j \neq p'_j\}|.$$

The following lemma asserts that products functions that are correlated are close in symbolic distance.

Lemma 12.18. Fix Σ , a distribution μ over Σ and a class $\mathcal{F} = \mathcal{F}(\Sigma)$ as in Theorem 12.5 which is τ -separated. Then for any $n \in \mathbb{N}$ and any n -variable functions $p, p' \in \mathcal{F}$ we have that

$$|\langle p, p' \rangle| \leq (1 - \tau)^{\Delta_{\text{symbolic}}(p, p')}.$$

Proof. Writing $p(x) = \prod_{j=1}^n p_j(x_j)$ and $p'(x) = \prod_{j=1}^n p'_j(x_j)$, we have that

$$|\langle p, p' \rangle| = \left| \prod_{j=1}^n \langle p_j, p'_j \rangle \right| = \prod_{j=1}^n |\langle p_j, p'_j \rangle|.$$

For j such that $p_j = p'_j$ the inner product is 1, and for any other j the absolute value of the inner product is at most $1 - \tau$. The result follows. \square

12.4.2 Local Structure

Fix Σ , μ distribution over Σ , $\mathcal{F} = \mathcal{F}(\Sigma)$, f and the parameters in Theorem 12.5. We will use the parameters

$$0 \ll c_4 \ll c_3 \ll c_2 \ll c_1 \ll s, C^{-1} \ll \tau, m^{-1}, \alpha \leq 1,$$

take

$$0 \leq \eta = \varepsilon^{\frac{1}{c_3 \rho}} \leq \zeta = \varepsilon^{\frac{1}{c_2 \rho}} \leq \xi = \varepsilon^{1/c_1} \leq 1,$$

and

$$\varepsilon' = \zeta^{-\log(1/\rho)^2/c_4}. \quad (40)$$

Throughout, we denote by $I \subseteq [n]$ a set which is picked by including each $i \in [n]$ in it with probability ρ . For $I \subseteq [n]$ and $z \in \Sigma^I$ define

$$W_{I,z} = \text{List}_\varepsilon[f_{\bar{I} \rightarrow z}], \quad \tilde{W}_{I,z} = \text{List}_{\varepsilon/2}[f_{\bar{I} \rightarrow z}],$$

and

$$SW_{I,z} = \text{ShortList}_{\varepsilon, \varepsilon^2/10}[f_{\bar{I} \rightarrow z}], \quad \tilde{S}W_{I,z} = \text{ShortList}_{\varepsilon/2, \varepsilon^2/100}[f_{\bar{I} \rightarrow z}].$$

We will now be interested in looking at $I' \subseteq_{\rho/2} [n]$. An equivalent way of sampling such I' is by first taking $I \subseteq_\rho [n]$, and then taking $I' \subseteq_{1/2} I$. As for restrictions, we can also first take $z \sim \nu^{I'}$, then $z' \sim \nu^{I \setminus I'}$ and concatenate them to get $f_{\bar{I} \rightarrow z \circ z'}$.

For fixed I and z , if we have that $p(x) = \prod_{i \in I} p_i(x_i)$ is in $SW_{I,z}$, then by Fact 12.12 we get that

$$|\langle f_{\bar{I} \rightarrow z \circ z'}, p|_{I \setminus I' \rightarrow z'} \rangle| \geq \frac{\varepsilon}{2}$$

with probability at least $\frac{\varepsilon^2}{4}$, in which case we get that $p|_{I'} \in \tilde{W}_{I', z \circ z'}$, so that in a formula the above says that

$$\Pr_{\substack{I, I' \\ z, z'}} \left[p|_{I'} \in \tilde{W}_{I', z \circ z'} \mid p \in SW_{I,z} \right] \geq \frac{\varepsilon^2}{4}. \quad (41)$$

This gives us some interesting information: while the identity of the original product function p may depend on both I and z , the function $p|_{I'}$ depends only on I, I' and z and not on z' . This suggests that on average, the identity of p itself should also only depend on I (the set of live variables) and not really on the value that we restrict outside them. To formalize this we use (41) in conjunction with Fact 12.13 to get that

$$\Pr_{\substack{I, z \\ I', I'' \subseteq_{1/2} I \\ z'_1, z''_1, z'_2, z''_2}} \left[p|_{I'} \in \tilde{W}_{I', z \circ z'_1} \cap \tilde{W}_{I', z \circ z'_2}, \quad p|_{I''} \in \tilde{W}_{I'', z \circ z''_1} \cap \tilde{W}_{I'', z \circ z''_2} \mid p \in SW_{I,z} \right] \geq \frac{\varepsilon^{16}}{4^4}. \quad (42)$$

We note that in particular, this inequality means that with noticeable probability, the short-lists of $I', z \circ z'_1$ and $I', z \circ z'_2$ contain two elements that have symbolic distance at most $C \cdot \log(1/\varepsilon)$. Indeed, as $p|_{I''}$ is in the list of these two restrictions, it follows by definition of the short-lists that for each one of this restrictions, the corresponding short list contains a product function that has correlation at least $\varepsilon^2/100$ with it, at which point we can use Lemma 12.18 to say that $p|_{I''}$ is close in symbolic distance to an element in each short list.

Take $T \subseteq [n]$ randomly by including each element in it with probability $\beta = \frac{1}{C \log(1/\varepsilon)}$, and define the following events:

1. Let $E_1(p)$ be the event that $p|_{I' \cap T} \in \tilde{S}\tilde{W}_{I', z \circ z'_1}|_T \cap \tilde{S}\tilde{W}_{I', z \circ z'_2}|_T$.
2. Let $E_2(p)$ be the event that $p|_{I'' \cap T} \in \tilde{S}\tilde{W}_{I'', z \circ z''_1}|_T \cap \tilde{S}\tilde{W}_{I'', z \circ z''_2}|_T$.
3. Let $E(p)$ be the event that $p \in SW_{I, z}$.

With these notations, the inequality (42) gives that

$$\Pr_{\substack{I, z \\ I', I'' \subseteq_{1/2} I \\ z'_1, z''_1, z'_2, z''_2 \\ T \subseteq_{\beta} [n]}} [E_1(p) \cap E_2(p) \mid E(p)] \geq \frac{\varepsilon^{16}}{4^4 e}. \quad (43)$$

Define $\mathcal{D} = \frac{1-\rho}{1-\rho/2}\nu' + \frac{\rho/2}{1-\rho/2}\nu$, and note that choosing I, I' and z, z' as above, the distribution of $z \circ z'$ is exactly \mathcal{D} . Let

$$W_{I', T} = \left\{ p \in \mathcal{F}(\Sigma, I' \cap T) \mid \Pr_{z \sim \mathcal{D}^{I'}} [p|_{I' \cap T} \in \tilde{S}\tilde{W}_{I', z}|_T] \geq \zeta \right\}.$$

Claim 12.19. *For all I' and T we have that*

$$\Pr_{\substack{\bar{T} = \bar{T} \cup (I \setminus I') \\ z, z', z''}} [\exists p \in \tilde{W}_{I \cup I' \setminus I, z \circ z'}|_T \cap \tilde{W}_{I \cup I' \setminus I, z \circ z''}|_T, p \notin W_{I', T}] \leq \sqrt{\xi}.$$

Proof. For each $p \in \mathcal{F}(\Sigma, I' \cap T)$, define $X_p = \left\{ z \in \Sigma^{\bar{T}} \mid p \in \tilde{S}\tilde{W}_{I', z}|_T \right\}$. Note that the condition that $p \notin W_{I', T}$ is equivalent to $\mathcal{D}(X_p) < \zeta$, and also that

$$\sum_p \mathcal{D}(X_p) \leq \sum_{p \in \mathcal{F}(\Sigma, I' \cap T)} \sum_z \mathcal{D}(z) 1_{p \in \tilde{S}\tilde{W}_{I', z}|_T} \leq \sum_z \mathcal{D}(z) \left| \tilde{S}\tilde{W}_{I', z}|_T \right| \leq O\left(\frac{1}{\varepsilon^2}\right), \quad (44)$$

where we used the fact that the short-lists as defined above have size at most $O(1/\varepsilon^2)$. Now consider the distribution over $z \circ z'$ and $z \circ z''$ as in the claim, and note that it is a product distribution $\tilde{\mu}^{I'}$ which has full support and the probability of each atom is at least $\Omega(\rho)$. Thus, denoting by \tilde{T} the averaging operator corresponding to $\tilde{\mu}$, we get by Lemma 2.7 that $\lambda_2(\tilde{T}) \leq 1 - \Omega_{\alpha, m}(\rho)$. Thus, noting that \mathcal{D} is a stationary distribution for \tilde{T} and using the Efron-Stein decomposition we get that the probability of the left hand side of the claim can be written as

$$\sum_{p \notin W_{I', T}} \langle 1_{X_p}, \tilde{T}^{\otimes I'} 1_{X_p} \rangle = \sum_{p \notin W_{I', T}} \sum_{d=0}^n \langle 1_{X_p}^{\overline{=d}}, \tilde{T}^{\otimes I'} 1_{X_p}^{\overline{=d}} \rangle \leq \sum_{p \notin W_{I', T}} \sum_{d=0}^n \lambda_2(\tilde{T})^d \|1_{X_p}^{\overline{=d}}\|_2^2.$$

Setting $D = \frac{1}{\rho} \log(1/\xi)$, we get that for $d \leq D$ the contribution is at most

$$\sum_{p \notin W_{I', T}} \sum_{d=0}^D \|1_{X_p}^{\overline{=d}}\|_2^2 = \sum_{p \notin W_{I', T}} \|1_{X_p}^{\leq D}\|_2^2 \leq \sum_{p \notin W_{I', T}} 2^{O_{\alpha}(D)} \mu(X_p)^{3/2} \leq 2^{O_{\alpha}(D)} \sqrt{\zeta} \sum_{p \notin W_{I', T}} \mu(X_p) \leq \frac{\sqrt{\xi}}{2},$$

where we used Fact 12.11, Parseval and (44). For $d > D$, the contribution is at most

$$(1 - \Omega_{\alpha, m}(\rho))^D \sum_{p \notin W_{I', T}} \sum_{d=D}^n \|1_{X_p}^{\overline{=d}}\|_2^2 \leq (1 - \Omega_{\alpha, m}(\rho))^D \sum_{p \notin W_{I', T}} \|1_{X_p}\|_2^2 \leq (1 - \Omega_{\alpha, m}(\rho))^D O\left(\frac{1}{\varepsilon^2}\right) \leq \frac{\sqrt{\xi}}{2},$$

where we used Parseval and (44) again. The proof is thus concluded by summing up the two inequalities. \square

Note that the premise of the theorem says that with probability at least ε over the choice of I and z we have that $W_{I,z}$ is non-empty, and hence by the definition of short lists we get that $\Pr[\exists p, E(p)] \geq \varepsilon$. In conjunction with (43) we conclude that

$$\Pr_{\substack{I \subseteq_{\rho}[n] \\ I', I'' \subseteq_{1/2} I \\ T \subseteq_{\beta}[n]}} [\exists p, E_1(p) \cap E_2(p) \cap E(p)] \geq \frac{\varepsilon^{17}}{4^4 e}.$$

By Claim 12.19, sampling $T \subseteq_{\beta}[n]$, the probability that $E_1(p)$ holds but $p|_{I' \cap T} \notin W_{I', T}$ is at most $\sqrt{\xi}$ and similarly for $E_2(p)$, so we get that

$$\Pr_{\substack{I \subseteq_{\rho}[n], z \sim \nu^I \\ I', I'' \subseteq_{1/2} I \\ T \subseteq_{\beta}[n]}} [\exists p, E(p), p|_{I' \cap T} \in W_{I', T}, p|_{I'' \cap T} \in W_{I'', T}] \geq \frac{\varepsilon^{17}}{4^4 e} - 2\sqrt{\xi} \geq \varepsilon^{18}. \quad (45)$$

12.4.3 Designing the Direct Product Function

We are going to use (45) to define a strategy for the direct product test. Towards this end, we first define the collection of I' 's on which we are going to assign a value which will give us a decent acceptance probability.

Definition 12.20. We say I' is good if there is $p \in \mathcal{F}(\Sigma, I')$ such that $\Pr_{T \subseteq_{\beta}[n]} [p|_{I' \cap T} \in W_{I', T}] \geq \xi$.

The following claim asserts that there is a significant fraction of good I' 's.

Claim 12.21. $\Pr_{I' \subseteq_{\rho/2}[n]} [I' \text{ is good}] \geq \frac{\varepsilon^{18}}{2}$.

Proof. Note that by (45), with probability at least $\varepsilon^{18}/2$ over the choice of I, I' and z we get that

$$\Pr_{T \subseteq_{\beta}[n]} [\exists p, E(p), p|_{I' \cap T} \in W_{I', T}] \geq \frac{\varepsilon^{18}}{2},$$

and as there are at most $O(1/\varepsilon^2)$ product functions p for which $E(p)$ holds we get that there is p such that

$$\Pr_{T \subseteq_{\beta}[n]} [p|_{I' \cap T} \in W_{I', T}] \geq \Omega(\varepsilon^{20}) \geq \xi,$$

hence I' is good. □

Define the following randomized strategy F for the direct product test. For each good I' :

- Choose $z' \sim \mathcal{D}^{I'}$.
- Sample $p \in \tilde{S}W_{I', z'}$ uniformly among the p 's satisfying that $\Pr_{T \subseteq_{\beta}[n]} [p|_{I' \cap T} \in W_{I', T}] \geq \xi$ and set $F[I'] = p$.

For I' which is not good, we choose $F[I']$ as a random string in $[R]^{I'}$. We now show two claims. The first of which asserts that in expectation, the above randomized strategy passes the direct product test with significant probability.

Claim 12.22. $\mathbb{E} [\Pr [F \text{ passes DP}(\rho, 1/2, \beta)]] \geq \Omega(\varepsilon^{20} \zeta^2) \geq \zeta^3$.

Proof. Note that the left hand side of the claim can be written as

$$\mathbb{E} \left[\mathbb{E}_{\substack{I \subseteq_\rho [n], T \subseteq_\beta [n] \\ I', I'' \subseteq_{1/2} I \\ z' \sim \mathcal{D}^{I'}, z'' \sim \mathcal{D}^{I''}}} \left[\mathbb{E}_{\substack{p' \in \overline{SW}_{I', z'} \\ p'' \in \overline{SW}_{I'', z'}}} \left[\mathbb{1}_{I', I'' \text{ are good}} \mathbb{1}_{p'|_T = p''|_T} \right] \right] \right].$$

Looking at (45), we get that with probability at least $\varepsilon^{18}/2$ over the choice of I, z and I', I'' it holds that

$$\Pr_{T \subseteq_\beta [n]} \left[\exists p, E(p), p|_{I' \cap T} \in W_{I', T}, p|_{I'' \cap T} \in W_{I'', T} \right] \geq \frac{\varepsilon^{18}}{2},$$

so as in the proof of Claim 12.21 there is p such that

$$\Pr_{T \subseteq_\beta [n]} \left[p|_{I' \cap T} \in W_{I', T}, p|_{I'' \cap T} \in W_{I'', T} \right] \geq \Omega(\varepsilon^{20}),$$

and in particular I' and I'' are good. Also, it follows that p is a candidate for both $F[I']$ and $F[I'']$ in the above randomized strategy, and by definition of $W_{I', T}$ we get that p' chosen in the randomized strategy satisfies that $p'|_T = p|_{I' \cap T}$ with probability at least ζ ; the same goes for p'' . Hence under the randomness of the choice of F the test passes with probability at least ζ^2 . The conclusion follows. \square

The second claim asserts that the contribution of the assignment F on non-good parts to the agreement of F with any direct product function is small.

Claim 12.23. For all $r \leq n/2$,

$$\Pr_F \left[\exists g: [n] \rightarrow [R], \Pr_{I'} \left[I' \text{ not good and } \Delta(g|_{I'}, F[I']) \leq r \right] \geq \frac{1}{2} \varepsilon' \right] \leq 2^{-n}.$$

Proof. Fix $g: [n] \rightarrow [R]$. Note that for each I' which is not good, the probability that $\Delta(g|_{I'}, F[I']) \leq r$ is at most $n^r R^{r-n}$. Let $Z_{I'}$ be the indicator random variable of this event, so that $Z_{I'}$ are independent and each of them has expectation at most $n^r m^{r-n}$. Thus, we get by Chernoff's inequality that

$$\Pr \left[\sum_{I' \text{ not good}} Z_{I'} \geq \frac{1}{2} \varepsilon' N \right] \leq 2^{-\Omega(\varepsilon' N)}$$

where N is the number of I' which are not good. Taking a union bound over g , we get that the probability in question is at most $R^n 2^{-\Omega(\varepsilon' N)} \leq 2^{-\Omega(\varepsilon' N)} \leq 2^{-n}$ as required. \square

12.4.4 Applying the Direct Product Theorem

We are now in a position to invoke the direct product theorem to finish the proof.

By Claim 12.22 and an averaging argument, with probability at least $\zeta^3/2$ over the choice of F we have that F passes the direct product test with probability at least $\zeta^3/2$. Hence, by Claim 12.23 and the union bound we have that with probability at least $\zeta^3/2 - 2^{-n} \geq \zeta^3/4$ it holds that F passes the direct product test with probability at least $\zeta^3/2$ and the contribution of the not-good I' 's to the agreement of F with any direct product function is at most $\frac{1}{2} \varepsilon'$. We fix such choice of F .

Using Theorem 12.7 we conclude that there is $g: [n] \rightarrow [R]$ such that

$$\Pr_{I'} [\Delta(g|_{I'}, F[I']) \leq r] \geq \varepsilon'$$

for $r = \text{poly}\left(\frac{\log(1/\zeta)}{\rho}\right)$ and ε' as in (40). Thus, we get that $\Pr_{I'} [I' \text{ is good, } \Delta(g|_{I'}, F[I']) \leq r] \geq \varepsilon'/2$.

Let $p_g \in \mathcal{F}(\Sigma, [n])$ be the product function corresponding to g . We note that for any good I' such that $\Delta(g|_{I'}, F[I']) \leq r$, by the definition of the randomized strategy we have that choosing $T \subseteq_{\beta} [n]$ and $T' \subseteq_{1/r} T$ it follows that $F[I'] \in W_{I', T}$ with probability at least ξ , and that conditioned on T we have that $g|_{I' \cap T'} = F[I']|_{T'}$ with probability at least $\Omega(1)$. Hence both events hold with probability at least $\Omega(\xi)$, and we fix such choices of T and T' henceforth.

By definition of $W_{I', T}$ we get that as $F[I'] \in W_{I', T}$, choosing $z \sim \mathcal{D}^{\bar{I}'}$ we have $F[I']|_T \equiv p_z|_T$ for some $p_z \in S\tilde{W}_{I', z}$ with probability at least ζ . For such z it holds that $|\langle p_z, f_{\bar{I}' \rightarrow z} \rangle| \geq \varepsilon/2$, and so by Fact 12.12 we get that choosing $z' \sim \nu^{I' \setminus T}$ it holds that

$$\left| \langle p_z|_{I' \cap T}, f_{I' \setminus T \rightarrow z'} \rangle \right| \geq \frac{\varepsilon}{4}$$

with probability at least $\varepsilon^2/16$. In this case we get

$$\left| \langle F[I']|_T, f_{I' \setminus T \rightarrow z'} \rangle \right| \geq \frac{\varepsilon}{4}.$$

Choosing $z'' \sim \mu^{T \setminus T'}$, we get by Fact 12.12 that

$$\left| \langle F[I']|_{T \setminus T' \rightarrow z''}, f_{I' \setminus T \rightarrow z'} \rangle \right| \geq \frac{\varepsilon}{8}$$

with probability at least $\varepsilon^2/64$, in which case it follows that

$$\left| \langle p_g|_{T'}, f_{I' \setminus T \rightarrow z'} \rangle \right| \geq \frac{\varepsilon}{8}.$$

In conclusion, we get that

$$\mathbb{E}_{\substack{I', T, T' \\ z, z', z''}} \left[\left| \langle p_g|_{T'}, f_{I' \setminus T \rightarrow z'} \rangle \right|^2 \right] \gtrsim \varepsilon' \xi \zeta \varepsilon^2.$$

Thus, looking at $G = p_g f$ we get that

$$\mathbb{E}_{\substack{I', T, T' \\ z, z', z''}} \left[\text{Stab}_{1/2} \left(G_{\substack{\bar{I}' \rightarrow z \\ I' \setminus T \rightarrow z' \\ T \setminus T' \rightarrow z''}} \right) \right] \geq \mathbb{E}_{\substack{I', T, T' \\ z, z', z''}} \left[\left| \mathbb{E}[G_{\substack{\bar{I}' \rightarrow z \\ I' \setminus T \rightarrow z' \\ T \setminus T' \rightarrow z''}}] \right|^2 \right] \gtrsim \varepsilon' \xi \zeta \varepsilon^2.$$

We now consider the distribution over the restriction, and note that it is equivalent to a restriction that chooses $J \subseteq_{\frac{\rho}{2}, \beta, \frac{1}{r}} [n]$ and fixes \bar{J} according to $(\mathcal{D}')^{\bar{J}}$, where \mathcal{D}' is some mixture of the distributions \mathcal{D} and ν . Thus, re-writing the above inequality we get that

$$\mathbb{E}_{\substack{J \subseteq_{\frac{\rho}{2}, \beta, \frac{1}{r}} [n] \\ y \sim (\mathcal{D}')^{\bar{J}}}} \left[\text{Stab}_{1/2} \left(G_{\bar{J} \rightarrow y} \right) \right] \gtrsim \varepsilon' \xi \zeta \varepsilon^2.$$

Using Lemma 2.8 we get that the left hand side is at most $\text{Stab}_{1-A, \rho\beta/r}(G)$ where A depends only on α and $m = |\Sigma|$, so $\text{Stab}_{1-A, \rho\beta/r}(G) \gtrsim \varepsilon' \xi \zeta \varepsilon^2$. It follows from Fact 12.10 that $W_{\leq A', \frac{r \log(1/\varepsilon' \xi \zeta \varepsilon)}{\rho\beta}} [G] \gtrsim \varepsilon' \xi \zeta \varepsilon^2$,

where A' depends only on α and m . Thus, $L = (p_g f)^{\leq A', \frac{r \log(1/\varepsilon' \xi \zeta \varepsilon)}{\rho\beta}}$ is a function of degree at most $A' \cdot \frac{r \log(1/\varepsilon' \xi \zeta \varepsilon)}{\rho\beta}$ and 2-norm 1 such that

$$\langle f, \bar{p}_g L \rangle = \langle G, L \rangle = W_{\leq A', \frac{r \log(1/\varepsilon' \xi \zeta \varepsilon)}{\rho\beta}} [G] \gtrsim \varepsilon' \xi \zeta \varepsilon^2,$$

and we are done. □

13 Proof of the Direct Product Theorem

In this section, we prove Theorem 12.7. This result improves can be seen as a quantitative improvement over the direct product theorem from [6]. To prove this result, we will first present and analyze a uniform version of this direct product tester; by that, we mean that we are going to be given an assignment to sets of size precisely $k = \rho n$, and that the various intersection sizes in the set are all replaced with exact intersection sizes. In Sections 13.3 and 13.4, we will show how to derive the direct product theorem when the underlying distribution is a product distribution as required for Theorem 12.7, using the well known trick of “going to infinity and back”.

Notation. Fix an alphabet $[m] = \{0, 1, 2, \dots, m-1\}$. Given a string $x \in [m]^n$ and a subset $S \subseteq [n]$, we use the notation $x|_S$ to denote the part of the string x restricted to the set S . Given $x, y \in [m]^n$, we use the notation $x \stackrel{\geq t}{\neq} y$ to denote that the set $\{i \in [n] \mid x_i \neq y_i\}$ is of size at least t . Similarly, we use $x \stackrel{\leq t}{\neq} y$ to denote that the strings x and y differ in at most t locations.

The set up for the direct product testing. Fix $q, q' \in (0, 1)$ such that $q' < q$ and an integer $t \in \mathbb{N}$. Let $q'' = q - q'$. Suppose we are given a table $F : \binom{[n]}{qn} \rightarrow [m]^{q''n}$ where $F[S] \in [m]^{q''n}$ can be thought of as assigning a symbol from $[m]$ to every element in S (by associating some fixed ordering on the elements of $[n]$). Consider the agreement test (*Agreement-Test*) parameterized by (q, q', t) given in Figure 13. Let $\mathcal{D}_{q, q'}$ be the distribution associated with the pair $(A_0 \cup B_0, A_0 \cup B_1)$ in the test.

It is clear that if the table F comes from a global string $a \in [m]^n$, in the sense that there is a vector a such that $F[S] = a|_S$ for all S , then the test accepts with probability 1 (even when $t = 0$). The following result is an inverse type result to this statement in the small soundness regime, and as discussed in the introduction in this case there are several challenging examples. In the following theorem, we prove that the type of assignments discussed in the introduction are essentially the only assignments that pass the test with non-negligible probability.

Given $F : \binom{[n]}{qn} \rightarrow [m]^{qn}$,

- Pick a random set $A_0 \cup B_0$ of size qn where $|A_0| = q'n$
- Select a random set $B_1 \subseteq [n] \setminus A_0$ of size $q''n$
- Check if $F[A_0 \cup B_0]|_{A_0} \stackrel{\leq t}{\neq} F[A_0 \cup B_1]|_{A_0}$

Figure 1: *Agreement-Test* with parameters (q, q', t) .

Theorem 13.1. *There exists $c > 0$ such that the following holds for sufficiently large $n \in \mathbb{N}$ and $\varepsilon \geq 2^{-n^c}$. Fix an alphabet $[m]$ and $c_0 > 0$ be any constant. For all $\varepsilon > 0$, $0 < q' < q < 1$ such that $q' \leq \frac{9q}{10}$ and $T = \left(\frac{\log(1/\varepsilon)}{q'}\right)^{c_0}$, suppose that $F : \binom{[n]}{qn} \rightarrow [m]^{qn}$ satisfies*

$$\Pr_{(S_1, S_2) \sim \mathcal{D}_{q, q'}} \left[F[S_1]|_{S_1 \cap S_2} \stackrel{\leq T}{\neq} F[S_2]|_{S_1 \cap S_2} \right] \geq \varepsilon.$$

Then there exists a function $g : [n] \rightarrow [m]$ such that for at least an $\varepsilon^{O(\log(1/q')^2)}$ fraction of $S \in \binom{[n]}{qn}$, we have $|\{i \in S \mid F[S]_i \neq g(i)\}| \leq \left(\frac{T \log(1/\varepsilon)}{q'}\right)^{O(1)}$.

We start by giving a high level overview of the proof of Theorem 13.1. The overall argument is similar to the one from [6], and in order to improve upon the quantitative bounds therein we require a more careful analysis, as well as more explicit quantitative bounds small set expansion result on the multi-cube graph (which plays an important role in our proof)

1. **Getting a local structure:** Suppose the table F passes the *Agreement-Test* with probability at least ε . For a pair (A, B) , we define a function $g_{A, B} : [n] \rightarrow [m]$ by taking the plurality vote among $F[A \cup B']|_i$ where (A, B') is such that $F[A \cup B]|_A$ agrees with $F[A \cup B']|_A$ on all but t coordinates. We then show that for a typical (A_0, B_0) , the function g_{A_0, B_0} agrees with the table $F[A_0 \cup B']$, on all but $\text{poly}(t)$ coordinates, for at least $\varepsilon^{O(1)}$ fraction of the B' s. Since in this step we can only show that g_{A_0, B_0} only agrees sets of the form $A_0 \cup B'$, which only constitutes $\exp(-n)$ fraction of the sets from $\binom{[n]}{qn}$, we somehow need to make sure that these different g_{A_0, B_0} s indeed agree with each other.
2. **Establishing consistency between various local views:** In this step, we show that for a typical pairs (A_0, B_0) and (A'_0, B'_0) , the functions g_{A_0, B_0} and $g_{A'_0, B'_0}$ agree with each other on all but $t^{O(1)}$ many coordinates. We show this in two steps. In the first step, we study a slightly different test that we call the *Modified-Test*. The purpose of analyzing this test is to conclude that for a typical pair (A_0, B_0) , if we select a pair (\tilde{A}_0, B'_0) , where \tilde{A}_0 is a correlated copy of A_0 and B'_0 is independent of (A_0, B_0) , then with noticeable probability the functions g_{A_0, B_0} and $g_{\tilde{A}_0, B'_0}$ agree on all but $t^{O(1)}$ many coordinates. In the second step, we show how to break the correlation between \tilde{A}_0 and A_0 . Here, we appeal to the small set expansion property of a graph defined over a multi-slice in [3]ⁿ. Towards this end, we think of (A, B) as an element $x \in [3]^n$ in the following way: $x_i = 1$ if $i \in A$, $x_i = 2$ if $i \in B$ and $x_i = 0$ otherwise. The edges of the graph are given by the distribution on the pairs (A_0, B_0) and (\tilde{A}_0, B'_0) . We show that this graph is a small-set expander using techniques from [9]. As every small set in this

graph expands, we use this to show that for typical pairs (A_0, B_0) and (A'_0, B'_0) , the functions g_{A_0, B_0} and $g_{A'_0, B'_0}$ agree with each other on all but $\text{poly}(\eta)$ many coordinates. Finally, this means that there exists a pair (A^*, B^*) such that g_{A^*, B^*} satisfies the conclusion of Theorem 13.1.

3. **From uniform setting to a product distribution:** We take N which is significantly larger than n , and attempt to simulate the q -biased distribution over $\{0, 1\}^n$ by the uniform distribution over qN sized subsets of $[N]$. More precisely, given a function $G : (P[n], \mu^{\otimes n}) \rightarrow [m]^{\leq n}$, we define a map $\tilde{G} : \binom{[N]}{qN} \rightarrow [m]^{qN}$ as follows: for $S \in \binom{[N]}{qN}$, define $\tilde{G}(S)|_{S \cap [n]} = G(S \cap [n])$ and $\tilde{G}(S)|_{S \setminus [n]} = 0^{|S \setminus [n]|}$. We note that taking a random qN sized set S from $[N]$, the distribution of $A = S \cap [n]$ is very close to being of a random subset of $[n]$ in which each coordinate is included with probability q . Thus, we are able to relate the performance of the (uniform sized set) direct product tester of \tilde{G} , and the performance of the (q -biased) direct product tester of G . Applying the uniform sized direct product testing result on \tilde{G} , we are quick able then to conclude Theorem 12.7.

13.1 Preliminaries

13.1.1 A Sampling Lemma

Consider the bipartite inclusion graph $G(n, \ell) = G([n] \cup Y, E)$ between $[n]$ and $Y = \binom{[n]}{\ell}$ for some $1 \leq \ell < n$, in which the edge set consists of pairs (i, A) such that $i \in A$. The following sampling lemma from [28] will be useful for our analysis.

Lemma 13.2. *Let $G(n, \ell)$ be the inclusion graph for $1 \leq \ell < n$. Let $Y' \subseteq Y$ be any subset of measure $\rho < 1/2$. For any constant $0 < \nu < 1$, we have that for all but at most $O_\nu\left(\frac{\log 1/\rho}{\ell}\right)$ fraction of vertices $x \in [n]$,*

$$\left| \Pr_{y \in N(x)} [y \in Y'] - \rho \right| \leq \nu \rho.$$

Here, $N(x)$ is the neighbors of the vertex x in G .

13.1.2 A Small Set Expansion Result

For a graph $G(V, E)$, let $\mathcal{T}(G)$ be the Markov operator associated with G . Also, let

$$\phi_G(\mu) := \min_{\substack{S \subseteq V(G), \\ |S| \leq \mu |V(G)|}} \Pr_{(u, v) \in E(G)} [v \notin S \mid u \in S].$$

Note that if every subset of size at most μ in G expands, then $\phi_G(\mu)$ is large. For any linear operator T , its $p \rightarrow q$ norm is defined as $\|T\|_{p \rightarrow q} := \max_{v \neq 0} \frac{\|Tv\|_q}{\|v\|_p}$. We will need the small set expansion property of a graph \mathcal{G}_n defined below.

The graph \mathcal{G}_n . Consider the graph \mathcal{G}_n induced on the set of vertices $\{(A, B) \mid A, B \subseteq [n], A \cap B = \emptyset, |A| = q'n, |B| = (q - q')n\}$ as follows, defined using a parameter c (that is to be thought of as an absolute but small constant). A random neighbor (A', B') of (A, B) in this graph is a pair where A' is distributed uniformly conditioned on $|A'| = q'n$, $|A \cap A'| = cq'n$ and $B' \subseteq [n] \setminus A'$ is a uniformly random set of size $(q - q')n$. An alternative view of this game (which will be crucial for us in order to derive small-set expansion results) proceeds by viewing it as a graph over the multi-slice – namely the set of strings in

$\{0, 1, 2\}^n$ with prescribed number of coordinates equal to 0, 1 and 2. Indeed, we map the vertex (A, B) to the string $x \in \{0, 1, 2\}^n$ where $x_i = 1$ if $i \in A$, $x_i = 2$ if $i \in B$ and $x_i = 0$ if $i \in [n] \setminus (A \cup B)$. The edges of the graph \mathcal{G}_n then naturally translate to edges over strings: a random neighbour of $x \in \{0, 1, 2\}^n$ is a string resulting from choosing a subset of size $cq'n$ of the 1's in x , and re-sampling the rest of the coordinates so that the number of coordinates equal to 0, 1 and 2 is as required. By abuse of notation we shall denote this graph also as \mathcal{G}_n .

The following result from [9] in order to give the following bounds on the expansion of \mathcal{G}_n .

Lemma 13.3. *For every $c > 0$, $0 < q' < q < 1$ such that $q' \leq \frac{9q}{10}$, and $\mu > 0$, the graph \mathcal{G}_n defined above has*

$$\phi_{\mathcal{G}_n}(\mu) \geq 1 - \mu^{\Omega\left(\frac{1}{\log(1/cq')^2}\right)}.$$

We give a proof of this lemma in Section 13.5.

13.1.3 From Uniform Size Distributions to Biased Distributions

We need the following two claims from [6] to move from the uniform setting to the product setting.

Claim 13.4. *Fix $q \in (0, 1)$, $t \leq n$ and $N = \omega(n^4)$. We have,*

$$\frac{\binom{N-n}{qN-t}}{\binom{N}{qN}} = q^t(1-q)^{n-t}(1 \pm o(1)).$$

where the $o(1)$ factor can be taken as $2^{-\Omega(n)}$.

Proof. The proof proceeds by a direct calculation and is deferred to Section D. □

With Claim 13.4 in hand, we can now show a coupling between the uniform distribution over subsets of $[N]$ of fixed size, and a distribution close to q -biased subsets of $[n]$.

Claim 13.5. *Fix $q \in (0, 1)$ and $N = \omega(n^4)$. Consider the following two distributions on $P([n])$:*

- \mathcal{D}_1 : *Select a subset $A \subseteq [n]$ by including $i \in [n]$ to A with probability q for each i independently.*
- \mathcal{D}_2 : *Select a random subset $S \subseteq [N]$ of size qN and output $S|_{[n]}$.*

Then, the statistical distance between \mathcal{D}_1 and \mathcal{D}_2 is at most e^{-n} .

Proof. We will compare the point-wise probabilities $p_1, p_2 : P([n]) \rightarrow \mathbb{R}$ assigned by the two distributions \mathcal{D}_1 and \mathcal{D}_2 , respectively. Fix any set $A \subseteq [n]$ of size t . We have $p_1(A) = q^t(1-q)^{n-t}$. Now, in order to sample A from \mathcal{D}_2 , it must be the case that $S|_{[n]} = A$ and therefore, we have

$$p_2(A) = \frac{\binom{N-n}{qN-t}}{\binom{N}{qN}},$$

which is $q^t(1-q)^{n-t}(1 \pm e^{-n})$ using Claim 13.4. □

13.2 Direct Product in the Uniform Setting: Proof of Theorem 13.1

Throughout this section, we use ε to denote the passing probability of the *Agreement-Test*. Since there is almost a black-box reduction from non-binary alphabet setting to the binary alphabet setting, we first focus on the binary alphabet for simplicity and prove Theorem 13.1 when $m = 2$. In Section 13.2.6, we show how to generalize the result for non-binary alphabet by keeping all the parameters asymptotically the same.

13.2.1 Parameters

Throughout this section, we are going to use the parameters

$$\begin{aligned} 1 &\leq c_0 \ll C_1 \ll C_2 \ll C_3 \ll C_4 \\ 0 &< \varepsilon \leq \nu = \varepsilon^3 \leq \gamma = \varepsilon^{10}, \end{aligned} \quad (46)$$

$$T = \left(\frac{\log(1/\varepsilon)}{q'} \right)^{c_0} \leq h = \frac{C_1 T \log(1/\varepsilon)}{q'} \leq R = W = \frac{C_2 h}{q'} \leq \tilde{h} = \frac{C_3 W^2 \log(1/\varepsilon)}{q'} \leq \tilde{R} = \tilde{W} = \frac{C_4 \tilde{h}}{q'} \quad (47)$$

and $q'' = q - q'$.

13.2.2 Local structure

In this section, we prove the local structure stated in Lemma 13.9 below. We need a few definitions to state the lemma.

Consider selecting a random set of size qn as follow. First sample a subset $A \subseteq [n]$ of size $q'n$ and then select a set $B \subseteq [n] \setminus A$ of size $(q - q')n$ uniformly at random. Output (A, B) . We need the following few definitions that are similar to the definitions from [28], adapted towards analyzing the *Agreement-Test*.

Definition 13.6. (*consistency*) Fix a set (A, B) . A subset $B' \subseteq [n] \setminus A$ is said to be t -consistent with (A, B) if $F[A, B]|_A \stackrel{\leq t}{\neq} F[A, B']|_A$. Let $\text{Cons}_t(A, B)$ be the set of all the sets that are t -consistent with (A, B) .

We say that (A, B) is good if $\text{Cons}_t(A, B)$ has a significant size. More precisely,

Definition 13.7. (*goodness*) A set (A, B) is called $(\varepsilon/2, t)$ -good if

$$\Pr_{B' \subseteq [n] \setminus A} [B' \in \text{Cons}_t(A, B)] \geq \varepsilon/2.$$

We next define the notion of excellence, for which we need to describe an auxiliary distribution over sets (E, D_1, D_2) . Given $A \subseteq [n]$, select two random subsets $B_1, B_2 \subseteq [n] \setminus A$ independently, each of size $q''n$ and let E be a random subset of $B_1 \cap B_2$ of size $q''n/R$. We take $D_1 = B_1 \setminus E$ and $D_2 = B_2 \setminus E$. We remark that the probability that $|B_1 \cap B_2|$ is smaller than $q''n/R$ is at most $\exp(-q''^2 n)$, hence so long as $\gamma = \omega(\exp(-q''^2 n))$ it will be absorbed into γ in the definition below.

Definition 13.8. (*excellence*) A set (A, B) is called $(\varepsilon/2, t, R, h, \gamma)$ -excellent if it is $(\varepsilon/2, t)$ -good and

$$\Pr_{E, D_1, D_2} [(E, D_i) \in \text{Cons}_t(A, B) \text{ for } i=1,2 \wedge F[A, E, D_1]|_E \stackrel{>h}{\neq} F[A, E, D_2]|_E] \leq \gamma.$$

Fix any $(\varepsilon/2, T, R, h, \gamma)$ -excellent pair (A_0, B_0) . We define a function $g_{A_0, B_0} : [n] \rightarrow \{0, 1\}$ based on the majority vote of the table F restricted to the sets in $\text{Cons}_t(A_0, B_0)$. More formally, for $x \in [n] \setminus A_0$, we set

$$g_{A_0, B_0}(x) := \text{Majority}_{B \in \text{Cons}_t(A_0, B_0) | B \ni x} F[A_0, B]|_x.$$

If there is no such B that contains x then we set $g_{A_0, B_0}(x) = 0$. We also set $g_{A_0, B_0}(A_0) = F[A_0, B_0]|_{A_0}$.

Based on these definitions, we prove the following local structure, which is the main lemma from this subsection. This is called a local structure as the functions g_{A_0, B_0} enjoy strong consistency (similar to what we need for the global function in Theorem 13.1) but it is weaker: the consistent is only guaranteed to be *local*, namely within $\text{Cons}_t(A_0, B_0)$.

Lemma 13.9. *For all $\varepsilon > 0$, if (A_0, B_0) is $(\varepsilon/2, T, R, h, \gamma)$ -excellent then*

$$\Pr_{B \in \text{Cons}_T(A_0, B_0)} [F[A_0, B]|_B \stackrel{>W^2}{\neq} g_{A_0, B_0}(B)] \leq \nu.$$

Furthermore, a random pair (A_0, B_0) is $(\varepsilon/2, T, R, h, \gamma)$ -excellent with probability at least $\frac{\varepsilon}{2} - \frac{2^{-\Omega(h)}}{\gamma} \geq \frac{\varepsilon}{4}$.

For notational convenience, in the remaining part of this subsection, we call a pair (A, B) good if it is $(\varepsilon/2, T)$ -good. Similarly, we call a pair (A, B) excellent if it is $(\varepsilon/2, T, R, h, \gamma)$ -excellent. We start by showing that a random pair is good and excellent with noticeable probability.

Claim 13.10. *If $\Pr_{A_0, B_0, B_1} [F[A_0, B_0]|_{A_0} \stackrel{\leq T}{\neq} F[A_0, B_1]|_{A_0}] \geq \varepsilon$, then a random (A_0, B_0) is $(\varepsilon/2, T)$ -good with probability at least $\varepsilon/2$.*

Proof. The proof is by a simple averaging argument. □

The next claim shows that almost all the good pairs are excellent.

Claim 13.11. *It holds that*

$$\Pr_{A_0, B_0} [(A_0, B_0) \text{ is } (\varepsilon/2, T)\text{-good but not } (\varepsilon/2, T, R, h, \gamma)\text{-excellent}] \leq \frac{2^{-\Omega(h)}}{\gamma}.$$

Proof. Consider the following two events.

1. Event Z_1 : (A_0, B_0) is good but

$$\Pr_{E, D_1, D_2} [(E, D_i) \in \text{Cons}_T(A_0, B_0) \text{ for } i = 1, 2 \wedge F[A_0, E, D_1]|_E \stackrel{>h}{\neq} F[A_0, E, D_2]|_E] > \gamma.$$

2. Event Z_2 : (A_0, B_0) is good, $(E, D_i) \in \text{Cons}_T(A_0, B_0)$ for $i = 1, 2$ and

$$F[A_0, E, D_1]|_{A_0 \cup E} \stackrel{>h/2}{\neq} F[A_0, E, D_2]|_{A_0 \cup E}.$$

We wish to upper bound $\Pr[Z_1]$, and towards this end we write:

$$\Pr[Z_1] = \Pr[Z_1 \wedge Z_2] / \Pr[Z_2 | Z_1] \leq \Pr[Z_2] / \Pr[Z_2 | Z_1],$$

and we give an upper bound on the numerator as well as a lower bound on the denominator. Note that $\Pr[Z_2 | Z_1] \geq \gamma$, and we now upper bound $\Pr[Z_2]$. The sets from the event Z_2 can be equivalently sampled as follows. First, sample a random subset A' of $[n]$ of size $q'n + \frac{q''n}{R}$ (which is to be thought of as $A_0 \cup E$), and then pick random sets $D_1, D_2 \subseteq [n] \setminus A'$ of size $q''n - \frac{q''n}{R}$, conditioned on the event $F[A', D_1]_{A'} \neq F[A', D_2]_{A'}$. Let $A'' \subseteq A'$ be the set of coordinates $x \in A'$ where $F[A', D_1](x) \neq F[A', D_2](x)$, so that $|A''| \geq h/2$. Taking A_0 to be a random subset of A' of size $q'n$ and setting $E = A' \setminus A_0$. If $(E, D_i) \in \text{Cons}_T(A_0, B_0)$ for $i = 1, 2$, we get that $F[A_0, E, D_1]_{A_0} \neq F[A_0, E, D_2]_{A_0}$ and hence $|A'' \cap A_0| \leq 2T$. On the other hand, the expected size of $|A'' \cap A_0|$ is at least $h/4$, and so by Chernoff's inequality the probability that $|A'' \cap A_0| \leq 2T$ is at most $2^{-\Omega(h)}$. \square

Next, the following claim asserts that for an excellent pairs (A_0, B_0) , the function g_{A_0, B_0} enjoys strong agreement with F inside $\text{Cons}_T(A_0, B_0)$.

Claim 13.12. *If (A_0, B_0) is excellent then $\Pr_{B \in \text{Cons}_T(A_0, B_0)} [F[A_0, B]_B \stackrel{>W^2}{\neq} g_{A_0, B_0}(B)] \leq \nu$.*

Before proving Claim 13.12, we first note that it implies Lemma 13.9.

Proof of Lemma 13.9. The lemma is immediate by Claims 13.10, 13.11 and 13.12. \square

The rest of this section is devoted to the proof Claim 13.12, and for that we introduce additional notations and prove some auxiliary claims. Fix an excellent pair (A_0, B_0) , and let

$$\text{Cons}_T^x = \{B \in \text{Cons}_T(A_0, B_0) \mid x \in B\};$$

we note that this set is used to define $g_{A_0, B_0}(x)$ in the majority voting step. Also, for $x \in [n]$ we denote by \mathcal{B}_x the collection of all $B \subseteq [n] \setminus A_0$ such that $|B| = (q - q')n$ and $x \in B$

Claim 13.13. *For at least $1 - O\left(\frac{\ln 1/\varepsilon}{q''n}\right)$ fraction of $x \in [n] \setminus A_0$, we have $2\varepsilon|\mathcal{B}_x| \geq |\text{Cons}_T^x| \geq \frac{\varepsilon}{6}|\mathcal{B}_x|$.*

Proof. This claim follows from the sampler property of the inclusion graph by invoking Lemma 13.2 for the graphs $G((1 - q')n, q''n)$. \square

We are now ready to prove Claim 13.12.

Proof of Claim 13.12. Fix A_0, B_0 as in the claim and assume towards contradiction the statement is false; then as $\text{Cons}_T(A_0, B_0)$ consists of at least $\varepsilon/2$ fraction of $B \subseteq [n] \setminus A_0$ of size $(q - q')n$, we get that

$$\Pr_{B \subseteq [n] \setminus A_0} [B \in \text{Cons}_T(A_0, B_0) \wedge F[A_0, B]_B \stackrel{>\zeta^2}{\neq} g_{A_0, B_0}(B)] > \frac{\nu\varepsilon}{2}. \quad (48)$$

Denote the above event by Z , and sample B conditioned on Z . Let $X_B \subseteq B$ be the set of coordinates $x \in B$ for which $F[A_0, B]_x \neq g_{A_0, B_0}(x)$. We have the following claim:

Claim 13.14. *For all but $O\left(\frac{\ln 1/\varepsilon}{q''n}\right)$ many $x \in [n] \setminus A_0$, we have*

$$\Pr_{B' \in \text{Cons}_T(A_0, B_0)} [x \in B' \wedge g_{A_0, B_0}(x) = F[A_0, B']_x] \geq \frac{q''}{10}.$$

Proof. Immediate by Claim 13.13. \square

Let $X'_B \subseteq X_B$ be the set of coordinates for the inequality in Claim 13.14 holds. Note that given the event Z , $|X'_B| \geq |X_B| - O\left(\frac{\ln 1/\varepsilon}{q''}\right) \geq W^2/2$, so we may pick a subset of X'_B of size $W^2/2$; without loss of generality we assume X'_B is already of that size. Picking $B' \in \text{Cons}_T(A_0, B_0)$ randomly, by Claim 13.14 the expected size of $Y_{B'} := \{x \in B' \mid x \in X'_B \ \& \ g_{A_0, B_0}(x) = F[A_0, B']|_x\}$ is at least $\frac{q''}{10}W^2$. Therefore, by an averaging argument, as $|Y_{B'}| \leq \frac{W^2}{2}$ always, with probability at least $\frac{q''}{20}$ we have $|Y_{B'}| \geq \frac{q''}{20}W^2$.

Also, picking $B' \in \text{Cons}_T(A_0, B_0)$ at random, we have by Chernoff's bound that $|B \cap B'| \geq \frac{q''^2 n}{10^{10}}$ with probability at least $1 - \exp(-\Omega(q''^2 n))$. Therefore, by union bound, with probability at least $\frac{q''}{30}$, we have

1. $|\{x \in B' \mid x \in X'_B \ \& \ g_{A_0, B_0}(x) = F[A_0, B']|_x\}| \geq \frac{q''}{100}W^2$.
2. $|B \cap B'| \geq \frac{q''^2 n}{10^{10}}$.

Taking a random $E \subseteq B \cap B'$ of size $q''n/R$, with probability at least $1/2$, we get that $|E \cap Y_{B'}| \geq \frac{q''}{200R}W^2$. Removing the conditioning on $B' \in \text{Cons}_T(A_0, B_0)$ and the above three conditions, we get

$$\Pr_{\substack{B, B' \subseteq [n] \setminus A_0 \\ E \subseteq B \cap B'}} [B, B' \in \text{Cons}_T(A_0, B_0) \ \wedge \ F[A_0, B]|_E \neq \frac{q''}{200R}W^2 \ \& \ F[A_0, B']|_E \geq \frac{\nu\varepsilon}{2} \cdot \frac{q''}{30} \cdot \frac{\varepsilon}{2} \cdot \frac{1}{2}] > \gamma.$$

Since $\frac{q''}{200R}W^2 \geq h$, this is a contradiction to the fact that (A_0, B_0) is $(\varepsilon/2, T, R, h, \gamma)$ -excellent. \square

13.2.3 Global Structure: the Modified Test

Now that we have a function g_{A_0, B_0} for every excellent pair (A_0, B_0) , the last step is to show that these functions are similar to each other and hence there is a global function g that (almost) agrees with at least $\delta(\varepsilon)$ fraction of the entries from the table $F[\cdot]$. We follow the same proof strategy as appeared in [6], however, we will use an explicit bound on the small-set expansion property of a certain graph.

For $c \in (0, 1)$, consider selecting a random set of size qn as follows. First, sample a subset $A \subseteq [n]$ of size $q'n$ and then select a set $B \subseteq [n] \setminus A$ of size $(q - q')n$ uniformly at random. Select a random subset of A of size $cq'n$ and call it D . Let $E = A \setminus D$. Output (D, E, B) , where $A = D \cup E$. Consider the modified agreement test (*Modified-Test*) given in Figure 13.2.3.

Note that *Modified-Test* is similar to the agreement test *Agreement-Test* that we wish to analyze, except that we change the parameters from (q', q, T) to $(cq'n, q, 2T)$ for $c \in (0, 1)$. Another (minor) difference is that we require the sets E and E' to be disjoint in the above distribution, whereas in the *Agreement-Test*, the sets B_0 and B_1 are uncorrelated. As the distribution of (E', B') depends on (D, E) , for notational convenience, we denote this marginal distribution by $\mathcal{D}(D, E)$.

We now relate the two tests *Agreement-Test* and *Modified-Test* in order to show the consistency between the functions $g_{A, B}$. Towards this, we define the notion of consistency, goodness, and excellence tailored to *Modified-Test*.

Definition 13.15. (*consistency*) Fix a set (D, E, B) where $A = D \cup E$. A subset (E', B') in the support of $\mathcal{D}(D, E)$ is said to be consistent with (D, E, B) if $F[D, E, B]|_D \stackrel{\leq 2T}{\equiv} F[D, E', B']|_D$. Also, we let $\text{newCons}_{2T}(D, E, B)$ be the set of all the sets (E', B') that are consistent with (D, E, B) .

Given $F : \binom{[n]}{qn} \rightarrow \{0, 1\}^{qn}$,

- Pick a random set $(D \cup E \cup B)$ of size qn where $|D| = cq'n$, $|E| = (1 - c)q'n$ and $|B| = q'n = (q - q')n$.
- Select a random subset $E' \subseteq [n] \setminus (D \cup E)$ of size $(1 - c)q'n$.
- Select a random set $B' \subseteq [n] \setminus (D \cup E')$ u.a.r. where $|B'| = (q - q')n$.
- Check if $F[D \cup E \cup B]|_D \stackrel{\leq 2T}{\equiv} F[D \cup E' \cup B']|_D$.

Figure 2: Modified agreement test *Modified-Test*

Definition 13.16. (*goodness*) A set (D, E, B) is called $(\varepsilon^2/2, 2T)$ -good if

$$\Pr_{(E', B') \sim \mathcal{D}(D, E)} [(E', B') \in \text{newCons}_{2T}(D, E, B)] \geq \varepsilon^2/2.$$

Definition 13.17. (*excellence*) A set (D, E, B) is called $(\varepsilon^2/2, 2T, R, h, \gamma)$ -excellent if it is $(\varepsilon^2/2, 2T)$ -good and

$$\Pr_{\substack{(E_1, B_1), (E_2, B_2) \sim \mathcal{D}(D, E) \\ E' \subseteq (E_1 \cup B_1) \cap (E_2 \cup B_2) \\ |E'| = q'n/R}} \left[\begin{array}{l} F[D \cup E_1 \cup B_1]|_{E'} \stackrel{\geq h}{\neq} F[D \cup E_2 \cup B_2]|_{E'} \text{ and} \\ (E_i, B_i) \in \text{newCons}_{2T}(D, E, B) \text{ for } i = 1, 2 \end{array} \right] \leq \gamma.$$

The following claim shows that if the *Agreement-Test* passes with probability at least ε , then the test *Modified-Test* passes with probability at least $0.99\varepsilon^2$.

Claim 13.18. If $\Pr_{A_0, B_0, B_1} [F[A_0, B_0]|_{A_0} \stackrel{\leq T}{\equiv} F[A_0, B_1]|_{A_0}] \geq \varepsilon$, then there is $c \in \left[\frac{q'}{2q}, \frac{2q'}{q} \right]$ such that the test *Modified-Test* passes with probability at least $0.99\varepsilon^2$. Consequently, a random triple (D, E, B) , with $|D| = cq'n$, is $(\varepsilon^2/2, 2T)$ -good with probability at least $0.49\varepsilon^2$.

Proof. Consider the following distribution.

- Select $S \subseteq [n]$ of size qn u.a.r.
- Select $A, A' \subseteq S$ each of size $q'n$ u.a.r.
- Select $B \subseteq [n] \setminus A$ of size $(q - q')n$ u.a.r.
- Select $B' \subseteq [n] \setminus A'$ of size $(q - q')n$ u.a.r.

We observe the following properties of the above distribution.

1. The pairs $(A, S \setminus A)$ and (A, B) are distributed according to the test distribution *Agreement-Test*. The same holds for the pairs $(A', S \setminus A')$ and (A', B')
2. For a fixed S , the pairs (A, B) and (A', B') are independent.

Note that

$$\begin{aligned}
& \mathbb{E}_{S,(A,B),(A',B')} \left[\mathbf{1}_{F[S]|_A \not\leq^T F[A,B]|_A} \cdot \mathbf{1}_{F[S]|_{A'} \not\leq^T F[A',B']|_{A'}} \right] \\
&= \mathbb{E}_S \left[\mathbb{E}_{(A,B),(A',B')} \left[\mathbf{1}_{F[S]|_A \not\leq^T F[A,B]|_A} \cdot \mathbf{1}_{F[S]|_{A'} \not\leq^T F[A',B']|_{A'}} \right] \right] \\
&= \mathbb{E}_S \left[\mathbb{E}_{(A,B)} \left[\mathbf{1}_{F[S]|_A \not\leq^T F[A,B]|_A} \right]^2 \right] \tag{Property 2. above} \\
&\geq \left(\mathbb{E}_S \left[\mathbb{E}_{(A,B)} \left[\mathbf{1}_{F[S]|_A \not\leq^T F[A,B]|_A} \right] \right] \right)^2 \tag{Jensen's inequality} \\
&= \varepsilon^2.
\end{aligned}$$

Note that the events $F[S]|_A \not\leq^T F[A,B]|_A$ and $F[S]|_{A'} \not\leq^T F[A',B']|_{A'}$ together imply that $F[A,B]|_{A \cap A'} \not\leq^{2T} F[A',B']|_{A \cap A'}$. Therefore, we have

$$\Pr_{(A,B),(A',B')} [F[A,B]|_{A \cap A'} \not\leq^{2T} F[A',B']|_{A \cap A'}] \geq \varepsilon^2.$$

Based on how the sets A and A' are distributed, we have with $1 - \exp(-n)$ probability, the size of $A \cap A'$ lies in $\left[\frac{q'^2 n}{2q}, \frac{2q'^2 n}{q} \right]$. Thus, there exist $c \in \left[\frac{q'}{2q}, \frac{2q'}{q} \right]$ such that

$$\Pr_{(A,B),(A',B') \mid |A \cap A'| = cq'n} [F[A,B]|_{A \cap A'} \not\leq^{2T} F[A',B']|_{A \cap A'}] \geq \varepsilon^2 - \exp(-n) \geq 0.99\varepsilon^2.$$

Now, if we let $D = A \cap A'$, $E = A \setminus D$ and $E' = A' \setminus D$, then the pairs (D, E, B) and (D, E', B') are distributed according to the distribution in *Modified-Test*. Hence, the acceptance probability of *Modified-Test* is at least $0.99\varepsilon^2$. The claim now follows from the averaging argument similar to the one in the proof of Claim 13.10. \square

The following claim shows that a random good set is excellent with high probability.

Claim 13.19. *A random $(\varepsilon^2/2, 2T)$ -good set (D, E, B) is $(\varepsilon^2/2, 2T, R, h, \gamma)$ -excellent with probability at least $1 - \frac{2^{-\Omega(h)}}{\gamma}$.*

Proof. The is along the same lines as the proof of Claim 13.11, and we prove it here again for completeness. Consider the following two events.

1. Event Z_1 : (D_0, E_0, B_0) is good but

$$\Pr_{\substack{(E_1, B_1), (E_2, B_2) \sim \mathcal{D}(D, E) \\ E' \subseteq (E_1 \cup B_1) \cap (E_2 \cup B_2) \\ |E'| = q''n/R}} \left[\begin{array}{l} (E_i, B_i) \in \text{newCons}_{2T}(D_0, E_0, B_0) \text{ for } i = 1, 2 \\ \wedge F[D_0 \cup E_1 \cup B_1]|_{E'} \not\stackrel{\geq h}{\neq} F[D_0 \cup E_2 \cup B_2]|_{E'} \end{array} \right] > \gamma.$$

2. Event Z_2 : (D_0, E_0, B_0) is good, $(E_i, B_i) \in \text{newCons}_{2T}(D_0, E_0, B_0)$ for $i = 1, 2$ such that $|(E_1 \cup B_1) \cap (E_2 \cup B_2)| \geq q''n/R$, $E' \subseteq (E_1 \cup B_1) \cap (E_2 \cup B_2)$ of size $q''n/R$ and

$$F[D_0, E_1, B_1]|_{D_0 \cup E'} \stackrel{> h/2}{\neq} F[D_0, E_2, B_2]|_{D_0 \cup E'}.$$

We wish to upper bound $\Pr[Z_1]$, and towards this end we write:

$$\Pr[Z_1] = \Pr[Z_1 \wedge Z_2] / \Pr[Z_2 \mid Z_1] \leq \Pr[Z_2] / \Pr[Z_2 \mid Z_1],$$

and we give an upper bound on the numerator as well as a lower bound on the denominator. Note that $\Pr[Z_2 \mid Z_1] \geq \gamma$, and we now upper bound $\Pr[Z_2]$. The sets from the event Z_2 can be equivalently sampled as follows. First, sample a random subset D' of $[n]$ of size $cq'n + \frac{q''n}{R}$ (which is to be thought of as $D_0 \cup E'$), and then pick random sets $H_1, H_2 \subseteq [n] \setminus D'$ of size $qn - cq'n - \frac{q''n}{R}$, conditioned on the event $F[D' \cup H_1] \stackrel{>h/2}{\neq} F[D' \cup H_2]$ and the event that the distribution of $(D' \cup H_1), (D' \cup H_2)$ is consistent with the distribution of $(D_0 \cup E_1 \cup B_1), (D_0 \cup E_2 \cup B_2)$ from event Z_2 . Let $D'' \subseteq D'$ be the set of coordinates $x \in D'$ where $F[D', H_1]_x \neq F[D', H_2]_x$, so that $|D''| \geq h/2$. Take D_0 to be a random subset of D' of size $cq'n$ and set $E' = D' \setminus D_0$. If $(E_i, B_i) \in \text{newCons}_{2T}(D_0, E_0, B_0)$ for $i = 1, 2$, we get that $F[D_0, E_1, B_1] \stackrel{\leq 4T}{\neq} F[D_0, E_2, B_2]$ and hence $|D'' \cap D_0| \leq 4T$. On the other hand, the expected size of $|D'' \cap D_0|$ is at least $h/4$, and so by Chernoff's inequality the probability that $|D'' \cap D_0| \leq 4T$ is at most $2^{-\Omega(h)}$. \square

Similar to the previous analysis, for an $(\varepsilon^2/2, 2T, R, h, \gamma)$ -excellent triple (D_0, E_0, B_0) , we define a function $g_{D_0, E_0, B_0} : [n] \rightarrow \{0, 1\}$ based on the majority vote of the table F restricted to the sets in $\text{newCons}_{2T}(D_0, E_0, B_0)$. More formally, for $x \in [n] \setminus D_0$, we set

$$g_{D_0, E_0, B_0}(x) := \underset{\substack{(E, B) \in \text{newCons}_{2T}(D_0, E_0, B_0) \\ E \cup B \ni x}}{\text{Majority}} F[D_0, E, B]_x.$$

If there is no such $E \cup B$ that contains x then we set $g_{D_0, E_0, B_0}(x) := 0$. We also set $g_{D_0, E_0, B_0}(x) = F[D_0, E_0, B_0]_x$ for $x \in D_0$. The following claim is analogous to Claim 13.12, saying that the local function g_{D_0, E_0, B_0} strongly agrees with the table F on $\text{newCons}_{2T}(D_0, E_0, B_0)$.

Claim 13.20. *If (D_0, E_0, B_0) is $(\varepsilon^2/2, 2T, R, h, \gamma)$ -excellent then*

$$\Pr_{(E, B) \in \text{newCons}_{2T}(D_0, E_0, B_0)} [F[D_0, E, B] \stackrel{>W^2}{\neq} g_{D_0, E_0, B_0}(E, B)] \leq \nu.$$

Proof. The proof of this claim is analogous to the proof of Claim 13.12. Fix an excellent pair (D_0, E_0, B_0) . We write newCons_{2T} for $\text{newCons}_{2T}(D_0, E_0, B_0)$ for convenience. We begin with the following claim.

Claim 13.21. *For at least $1 - O\left(\frac{\ln 1/\varepsilon}{q'n}\right)$ fraction of $x \in [n] \setminus D_0$,*

$$\Pr_{\substack{(E, B) \sim \mathcal{D}(D_0, E_0), \\ x \in E \cup B}} [(E, B) \in \text{newCons}_{2T}] \geq \frac{1}{100} \Pr_{(E, B) \sim \mathcal{D}(D_0, E_0)} [(E, B) \in \text{newCons}_{2T}].$$

Proof. The proof follows essentially from the sampler property of an inclusion graph. In this case, the distribution of (E, B) is not uniform among all the sets of size $(q - cq')n$ of $[n] \setminus D_0$. Therefore we cannot apply Lemma 13.2 directly. However, similar proof works here and we include it below.

For succinctness, let \mathcal{D} be the distribution $\mathcal{D}(D_0, E_0)$ and let \mathcal{D}_x be the distribution \mathcal{D} conditioned on $x \in E \cup B$. The way we sample $(E, B) \sim \mathcal{D}$, the marginal distribution on $x \in E \cup B$ is not the same. However, for every $x \in [n] \setminus (D_0 \cup E_0)$, the marginal distribution is identical. Same is true for every $x \in E_0$.

We will show that at most $O\left(\frac{\ln 1/\tilde{\varepsilon}}{q''}\right)$ many $x \in [n] \setminus (D_0 \cup E_0)$, the claimed inequality does not hold. Similar proof shows that for at most $O\left(\frac{\ln 1/\tilde{\varepsilon}}{q''}\right)$ many $x \in E_0$, the inequality does not hold.

We know that $\Pr_{(E,B) \sim \mathcal{D}}[(E,B) \in \text{newCons}_{2T}] = \tilde{\varepsilon} \geq \varepsilon^2/2$. Let $n' = n - |D_0| - |E_0| = (1 - q')n$ and $Z = \{x \in [n] \setminus (D_0 \cup E_0) \mid \Pr_{(E,B) \sim \mathcal{D}_x}[S \in \text{newCons}_{2T}] < \tilde{\varepsilon}/100\}$. We need to show that $|Z| = O\left(\frac{\ln 1/\tilde{\varepsilon}}{q''}\right)$. For $x \in [n] \setminus (D_0 \cup E_0)$, let p be the probability that $x \in (E \cup B)$ when (E, B) is sampled according to \mathcal{D} . Note that $p = \Omega(q'')$. Suppose for contradiction, $|Z| = C\frac{\ln 1/\tilde{\varepsilon}}{p}$ for a large $C > 0$. By the definition of the set Z we have,

$$\Pr_{\substack{x \in [n] \setminus (D_0 \cup E_0) \\ (E,B) \sim \mathcal{D}_x}} [x \in Z \wedge (E, B) \in \text{newCons}_{2T}] < \frac{|Z|}{n'} \cdot \frac{\tilde{\varepsilon}}{100} = \frac{C \ln(1/\tilde{\varepsilon}) \cdot \tilde{\varepsilon}}{100p(1 - q')n}. \quad (49)$$

If we sample a random $(E, B) \sim \mathcal{D}$, then by Chernoff's inequality, $E \cup B$ contains at least $\frac{C}{2} \ln(1/\tilde{\varepsilon})$ elements from Z with probability at least $1 - \exp(-\Omega(C \ln 1/\tilde{\varepsilon})) \geq 1 - \tilde{\varepsilon}/10$. Therefore, we have

$$\Pr_{(E,B) \sim \mathcal{D}} \left[(E, B) \in \text{newCons}_{2T} \wedge |(E \cup B) \cap Z| \geq \frac{C}{2} \ln(1/\tilde{\varepsilon}) \right] \geq \frac{9\tilde{\varepsilon}}{10}.$$

Conditioned on the above event, if we sample a random $x \in (E \cup B) \setminus E_0$, then $x \in Z$ with probability at least $\frac{1}{2} \frac{C/2 \ln(1/\tilde{\varepsilon})}{pn}$. Therefore, combining with the above event, we have

$$\Pr_{\substack{(E,B) \sim \mathcal{D} \\ x \in (E \cup B) \setminus E_0}} [x \in Z \wedge (E, B) \in \text{newCons}_{2T}] \geq \frac{9C \ln(1/\tilde{\varepsilon})\tilde{\varepsilon}}{40pn}.$$

Assuming $q' \leq 9/10$, this contradicts (49) □

We now proceed to proving the claim. Assume towards contradiction the statement of the claim is false; then as $(E, B) \in \text{newCons}_{2T}(D_0, E_0, B_0)$ with probability at least $\varepsilon^2/2$ when sampled according to $\mathcal{D}(D_0, E_0)$, we get that

$$\Pr_{(E,B) \sim \mathcal{D}(D_0, E_0)} [(E, B) \in \text{newCons}_{2T}(D_0, E_0, B_0) \wedge F[D_0, E, B]|_{E,B} \stackrel{>W^2}{\neq} g_{D_0, E_0, B_0}(E, B)] > \frac{\nu \varepsilon^2}{2}. \quad (50)$$

Denote the above event by Z , and sample (E, B) conditioned on Z . Let $X_{E,B} \subseteq E \cup B$ be the set of coordinates $x \in E \cup B$ for which $F[D_0, E, B]|_x \neq g_{D_0, E_0, B_0}(x)$. We have the following claim:

Claim 13.22. For all but $O\left(\frac{\ln 1/\varepsilon}{q''}\right)$ many $x \in [n] \setminus D_0$, we have

$$\Pr_{(E', B') \in \text{newCons}_{2T}(D_0, E_0, B_0)} [x \in E' \cup B' \wedge g_{D_0, E_0, B_0}(x) = F[D_0, E', B']|_x] \geq \frac{q''}{2000}.$$

Proof. Immediate by Claim 13.21 and the fact that $g_{D_0, E_0, B_0}(x)$ was set according to the majority vote. □

Let $X'_{E,B} \subseteq X_{E,B}$ be the set of coordinates for the inequality in Claim 13.22 holds. Note that given the event Z , $|X'_{E,B}| \geq |X_{E,B}| - O\left(\frac{\ln 1/\varepsilon}{q''}\right) \geq W^2/2$, so we may pick a subset of $X'_{E,B}$ of size $W^2/2$; without loss of generality we assume $X'_{E,B}$ is already of that size. Picking $(E', B') \in \text{newCons}_{2T}(D_0, E_0, B_0)$

randomly, by Claim 13.22 the expected size of $Y_{E',B'} := \{x \in E' \cup B' \mid x \in X'_{E,B} \ \& \ g_{D_0,E_0,B_0}(x) = F[D_0, E', B']|_x\}$ is at least $\frac{q''}{2000}W^2$. Therefore, by an averaging argument, as $|Y_{E',B'}| \leq \frac{W^2}{2}$ always, with probability at least $\frac{q''}{4000}$ we have $|Y_{E',B'}| \geq \frac{q''}{4000}W^2$.

Also, picking $(E', B') \sim \mathcal{D}(D_0, E_0)$ at random conditioned on $(E', B') \in \text{newCons}_{2T}(D_0, E_0, B_0)$, we have by Chernoff's bound that $|(E \cup B) \cap (E' \cup B')| \geq \frac{q''n}{10^{10}}$ with probability at least $1 - \exp(-\Omega(q''^2n))$, as $\varepsilon^2/2 \ll \exp(-\Omega(q''^2n))$. Therefore, by union bound, with probability at least $\frac{q''}{5000}$, we have

1. $|\{x \in E' \cup B' \mid x \in X'_{E,B} \ \& \ g_{D_0,E_0,B_0}(x) = F[D_0, E', B']|_x\}| \geq \frac{q''}{4000}W^2$.
2. $|B \cap B'| \geq \frac{q''n}{10^{10}}$.

Taking a random $E \subseteq (E \cup B) \cap (E' \cup B')$ of size $q''n/R$, with probability at least $1/2$, we get that $|E \cap Y_{E',B'}| \geq \frac{q''}{8000R}W^2$. Removing the conditioning on $(E', B') \in \text{newCons}_{2T}(D_0, E_0, B_0)$ and the above three conditions, we get

$$\Pr_{\substack{(E,B),(E',B') \sim \mathcal{D}(D_0,E_0) \\ E \subseteq (E \cup B) \cap (E' \cup B')}} \left[\begin{array}{l} (E,B),(E',B') \in \text{newCons}_{2T}(D_0,E_0,B_0) \\ \wedge F[D_0,E,B]|_E > \frac{q''}{8000R}W^2 \neq F[D_0,E',B']|_E \end{array} \right] \geq \frac{\nu\varepsilon^2}{2} \cdot \frac{q''}{5000} \cdot \frac{\varepsilon^2}{2} \cdot \frac{1}{2} > \gamma.$$

Since $\frac{q''}{8000R}W^2 \geq h$, this is a contradiction to the fact that (D_0, E_0, B_0) is $(\varepsilon^2/2, 2T, R, h, \gamma)$ -excellent. \square

13.2.4 Consistency Between the Functions $g_{A,B}$ via newCons

In this section, we use the local structure result proved in the previous section (but with different parameters) and show a relationship between these local functions. To do so, we introduce a variant of the ‘‘Cons’’ set which allows us to compare the local functions $g_{A,B}$ and $g_{A',B'}$ for A and A' that are correlated (but not the same).

Let c_0 be the constant from Theorem 13.1 and take $c \in \left[\frac{q'}{2q}, \frac{2q'}{q}\right]$ from Claim 13.18.

Definition 13.23. For an $(\varepsilon^2/2, 2T, R, h, \gamma)$ -excellent triple (D_0, E_0, B_0) , let $\text{newCons}_{W^2}^*(D_0, E_0, B_0) \subseteq \text{newCons}_{2T}(D_0, E_0, B_0)$ consist of all (E, B) 's such that $F[D_0, E, B]|_{E,B} \stackrel{\leq W^2}{\neq} g_{D_0,E_0,B_0}(E, B)$. Similarly, for an $(\varepsilon^2/8, 2W^2, \tilde{R}, \tilde{h}, \gamma)$ -excellent pair (A_0, B_0) , let $\text{Cons}_{2W^2}^*(A_0, B_0) \subseteq \text{Cons}_{2W^2}(A_0, B_0)$ consist of B 's such that $F[A_0, B]|_B \stackrel{\leq W^2}{\neq} g_{A_0,B_0}(B)$.

Remark 13.24. Note that in the definition above we look at $\text{Cons}_{2W^2}(A_0, B_0)$, which are the same as the ‘‘Cons’’ sets considered in previous sections except that the parameter T is replaced by the parameter $2W^2$. To compensate for that, we have adjusted the excellence parameters accordingly and have taken \tilde{R} and \tilde{h} instead of R and h from the previous section.

Note that by Claims 13.18, 13.20, for every excellent triple (D_0, E_0, B_0) , we have

$$\Pr_{(E',B') \sim \mathcal{D}(D_0,E_0)} [(E', B') \in \text{newCons}_{W^2}^*(D_0, E_0, B_0)] \geq \varepsilon^2/4. \quad (51)$$

Definition 13.25. Fix (D_0, E_0, B_0) that is $(\varepsilon^2/2, 2T)$ -good. A pair $(E, B) \in \text{newCons}_{W^2}^*(D_0, E_0, B_0)$ is called a dense pair, if $\Pr_{B' \subseteq [n] \setminus (D_0 \cup E)} [(E, B') \in \text{newCons}_{W^2}^*(D_0, E_0, B_0)] \geq \varepsilon^2/8$.

The next claim shows that many sets (E, B) in $\text{newCons}_{W^2}^*(D_0, E_0, B_0)$ when we pair it with D_0 as $(D_0 \cup E, B)$, then the pair $(D_0 \cup E, B)$ is excellent for *Agreement-Test*.

Claim 13.26. *Let (D_0, E_0, B_0) be an $(\varepsilon^2/2, 2T, R, h, \gamma)$ -excellent triple; then,*

$$\Pr_{(E,B) \in \text{newCons}_{W^2}^*(D_0, E_0, B_0)} [(E, B) \text{ is dense} \wedge (D_0 \cup E, B) \text{ is } (\varepsilon^2/8, 4W^2, \tilde{R}, \tilde{h}, \gamma)\text{-excellent}] \geq \varepsilon^2/16.$$

Proof. Since (D_0, E_0, B_0) is excellent, by using (51) and an averaging argument we have that at least $\varepsilon^2/8$ fraction of $(E, B) \in \text{newCons}_{W^2}^*(D_0, E_0, B_0)$ are dense pairs. We will show that for such pair (E, B) , $(D_0 \cup E, B)$ is $(\varepsilon^2/8, 2W^2)$ -good.

Fix any such (E, B) . By the definition of a dense pair, we have

$$\Pr_{B' \subseteq [n] \setminus D_0 \cup E} \left[\begin{array}{l} g_{D_0, E_0, B_0}(D_0 \cup E) \stackrel{\leq 2W^2}{\neq} F[D_0, E, B]|_{D_0 \cup E} \text{ and} \\ g_{D_0, E_0, B_0}(D_0 \cup E) \stackrel{\leq 2W^2}{\neq} F[D_0, E, B']|_{D_0 \cup E} \end{array} \right] \geq \varepsilon^2/8.$$

Note that the first event in the above probability follows from the fact that $(E, B) \in \text{newCons}_{W^2}^*(D_0, E_0, B_0)$ (W^2 accounts for the possible disagreement on E and another $W^2 \geq 2T$ accounts for the disagreement on D_0) and the second event follows from the fact that for a random B' , $(E, B') \in \text{newCons}_{W^2}^*(D_0, E_0, B_0)$ with probability at least $\varepsilon^2/8$, as (E, B) is a dense pair. It follows that

$$\Pr_{B' \subseteq [n] \setminus D_0 \cup E} \left[F[D_0, E, B]|_{D_0 \cup E} \stackrel{\leq 4W^2}{\neq} F[D_0, E, B']|_{D_0 \cup E} \right] \geq \varepsilon^2/8,$$

and hence $(D_0 \cup E, B)$ is $(\varepsilon^2/8, 4W^2)$ -good. By Claim 13.11 at least $1 - \nu$ fraction of the $(\varepsilon^2/8, 4W^2)$ -good pairs are $(\varepsilon^2/8, 4W^2, \tilde{R}, \tilde{h}, \gamma)$ -excellent, it follows that at least $\varepsilon^2/8 - O(\nu/\varepsilon^2) \geq \varepsilon^2/16$ of the pairs $(E, B) \in \text{newCons}_{W^2}^*(D_0, E_0, B_0)$ it holds that $(D_0 \cup E, B)$ is excellent, concluding the proof. \square

Claim 13.26 will help us to relate the functions g_{D_0, E_0, B_0} and $g_{D_0 \cup E, B}$ where (D_0, E_0, B_0) is excellent and (E, B) is a typical pair from $\text{newCons}_{W^2}^*(D_0, E_0, B_0)$. Since the functions g_{D_0, E_0, B_0} and $g_{D_0 \cup E, B}$ are defined based on their respective ‘‘Cons’’ set, we need to make sure that the sets in $\text{Cons}_{2W^2}^*$ and $\text{newCons}_{W^2}^*$ are correlated. The following claim (and in particular the third item) shows that this is indeed true.

Claim 13.27. *Fix any $(\varepsilon^2/2, 2T, R, W, \gamma)$ -excellent triple (D_0, E_0, B_0) . When (E, B) is sampled according to the distribution $\mathcal{D}(D_0, E_0)$, then with probability at least $\Omega(\varepsilon^4)$, we have:*

1. $(E, B) \in \text{newCons}_{W^2}^*(D_0, E_0, B_0)$.
2. $(D_0 \cup E, B)$ is $(\varepsilon^2/8, 4W^2, \tilde{R}, \tilde{h}, \gamma)$ -excellent.
3. *There is at least an $\Omega(\varepsilon^2)$ fraction of $B' \subseteq [n] \setminus (D_0 \cup E)$ such that $(E, B') \in \text{newCons}_{W^2}^*(D_0, E_0, B_0)$ and $B' \in \text{Cons}_{4W^2}^*(D_0 \cup E, B)$.*

Proof. Using Claim 13.26 and (51), for a random $(E, B) \sim \mathcal{D}(D_0, E_0)$, with probability at least $\Omega(\varepsilon^4)$, we have

- (i) $(E, B) \in \text{newCons}_{W^2}^*(D_0, E_0, B_0)$,
- (ii) $(D_0 \cup E, B)$ is $(\varepsilon^2/8, 4W^2, \tilde{R}, \tilde{h}, \gamma)$ -excellent, and

(iii) $\Pr_{B' \subseteq [n] \setminus (D_0 \cup E)} [(E, B') \in \text{newCons}_{W^2}^*(D_0, E_0, B_0)] = \Omega(\varepsilon^2)$.

The condition (i) above implies

$$g_{D_0, E_0, B_0}(D_0 \cup E \cup B) \stackrel{\leq 2W^2}{\neq} F[D_0 \cup E \cup B], \quad (52)$$

where the additional $W^2 \geq 2T$ accounts for the possible disagreement on D_0 . Using the condition (iii) above, we also have

$$g_{D_0, E_0, B_0}(D_0 \cup E \cup B') \stackrel{\leq 2W^2}{\neq} F[D_0 \cup E \cup B'] \quad (53)$$

for at least $\Omega(\varepsilon^2)$ fraction of $B' \subseteq [n] \setminus (D_0 \cup E)$. Using (52) and (53), for at least $\Omega(\varepsilon^2)$ fraction of B' we get that $F[D_0 \cup E \cup B]|_{D_0 \cup E} \stackrel{\leq 4W^2}{\neq} F[D_0 \cup E \cup B']|_{D_0 \cup E}$, in which case $B' \in \text{Cons}_{4W^2}(D_0 \cup E, B)$ and the claim follows. \square

Using the above claim, we show that for every (D_0, E_0, B_0) that is excellent, the functions g_{D_0, E_0, B_0} and $g_{D_0 \cup E, B}$ are very close to each other in hamming distance for many pairs (E, B) .

Claim 13.28. Fix any $(\varepsilon^2/2, 2T, R, h, \gamma)$ -excellent triple (D_0, E_0, B_0) . Then

$$\Pr_{(E, B) \sim \mathcal{D}(D_0, E_0)} [g_{D_0, E_0, B_0} \stackrel{\leq \tilde{W}^3}{\neq} g_{D_0 \cup E, B}] \gtrsim \varepsilon^4.$$

Proof. Select a pair (E, B) according to the distribution $\mathcal{D}(D_0, E_0)$. Using Claim 13.27 with probability at least $\Omega(\varepsilon^4)$, we have $g_{D_0, E_0, B_0}(D_0 \cup E \cup B) \stackrel{\leq 2W^2}{\neq} F[D_0 \cup E \cup B]$, $(D_0 \cup E, B)$ is $(\varepsilon^2/8, 4W^2, \tilde{R}, \tilde{h}, \gamma)$ -excellent and $g_{D_0 \cup E, B}(D_0 \cup E \cup B) \stackrel{\leq 3\tilde{W}^2}{\neq} F[D_0 \cup E \cup B]$. Note that in the last condition, we used the fact that if $(D_0 \cup E, B)$ is excellent and if $B' \in \text{Cons}_{2W^2}^*(D_0 \cup E, B)$, then $(D_0 \cup E, B')$ is also excellent and furthermore the functions $g_{D_0 \cup E, B}$ and $g_{D_0 \cup E, B'}$ are the same as they are defined using the set $\text{Cons}_{4W^2}^*(D_0 \cup E, B) = \text{Cons}_{4W^2}^*(D_0 \cup E, B')$.

Combining these events and Claim 13.27, we get that over the randomness of (E, B) , with probability at least $\Omega(\varepsilon^4)$, the following happens:

1. $(D_0 \cup E, B)$ is $(\varepsilon^2/8, 4W^2, \tilde{R}, \tilde{h}, \gamma)$ -excellent,
2. $g_{D_0, E_0, B_0}(D_0 \cup E \cup B) \stackrel{\leq 4\tilde{W}^2}{\neq} g_{D_0 \cup E, B}(D_0 \cup E \cup B)$, and
3. $\Pr_{B' \subseteq [n] \setminus (D_0 \cup E)} [(E, B') \in \text{newCons}_{W^2}^*(D_0, E_0, B_0) \wedge B' \in \text{Cons}_{4W^2}^*(D_0 \cup E, B)] \gtrsim \varepsilon^2$.

From the third point we can conclude that $\Pr_{B' \subseteq [n] \setminus (D_0 \cup E)} [g_{D_0, E_0, B_0}(B') \stackrel{\leq 4W^2}{\neq} g_{D_0 \cup E, B}(B')] \gtrsim \varepsilon^2$. Using

Claim 13.29, we get that $g_{D_0, E_0, B_0}|_{[n] \setminus (D_0 \cup E)} \stackrel{\leq 16W^2/q''}{\neq} g_{D_0 \cup E, B}|_{[n] \setminus (D_0 \cup E)}$. Combined with the second item above saying that g_{D_0, E_0, B_0} and $g_{D_0 \cup E, B}$ differ on at most $4\tilde{W}^2$ elements from $D_0 \cup E$, we get that g_{D_0, E_0, B_0} and $g_{D_0 \cup E, B}$ differ on at most \tilde{W}^3 elements. \square

Claim 13.29. For any $q \in (0, 1)$, $a \geq 1$ and $\delta > e^{-qa/8}$, given two functions $f, g : [n] \rightarrow \{0, 1\}$ such that $\Pr_{S \subseteq [n] ||S|=qn} [f(S) \stackrel{\leq a}{\neq} g(S)] \geq \delta$, then $|\{i \in [n] \mid f(i) \neq g(i)\}| \leq 4a/q$.

Proof. Let $\mathcal{E} = \{i \in [n] \mid f(i) \neq g(i)\}$, and suppose towards contradiction $|\mathcal{E}| \geq 4a/q$. Sample a subset $S \subseteq_{q/2} [n]$; by the Chernoff's bound $|S| \leq qn$ with probability at least $1 - \exp(-\Omega(q^2n))$. Furthermore, the expected size of $S \cap \mathcal{E}$ is $2a$. Now, again by the Chernoff's bound, the probability that $|S \cap \mathcal{E}|$ is at most a is at most $e^{-qa/4}$. This implies $\Pr_{S \subseteq [n] \mid |S|=qn} [f(S) \stackrel{\leq a}{\neq} g(S)] \leq e^{-qa/8}$ which is a contradiction. \square

Claim 13.28 allows us to relate functions defined using ‘‘Cons’’ and functions defined using ‘‘newCons’’, and using it along with an application of the Cauchy-Schwarz inequality we conclude relations between functions defined using ‘‘Cons’’ (but with different A_0 's).

Claim 13.30. *There exists a constant $d_0 \in (0, 1)$ such that for $c' = d_0c$, we have the following. For at least $\Omega(\varepsilon^{12})$ fraction of excellent pairs (A_0, B_0) ,*

$$\Pr_{(\tilde{A}_0, B)} [g_{A_0, B_0} \stackrel{\leq 2\tilde{W}^3}{\neq} g_{\tilde{A}_0, B}] \gtrsim \varepsilon^{12}.$$

Here, \tilde{A}_0 is distributed uniformly conditioned on $|\tilde{A}_0| = q'n$, $|A_0 \cap \tilde{A}_0| = c'q'n$ and $B \subseteq [n] \setminus \tilde{A}_0$ of size $(q - q')n$.

Proof. Since $\Omega(\varepsilon^2)$ fraction of the triples (D_0, E_0, B_0) are $(\varepsilon^2/2, 2T, R, h, \gamma)$ -excellent, we have the following from Claim 13.28.

$$\Omega(\varepsilon^6) \leq \mathbb{E}_{(D_0, E_0, B_0)} \mathbb{E}_{(E, B) \sim \mathcal{D}(D_0, E_0)} \left[\mathbf{1}_{g_{D_0, E_0, B_0} \stackrel{\leq \tilde{W}^3}{\neq} g_{D_0 \cup E, B}} \right].$$

Note that $(D_0 \cup E, B)$ being $(\varepsilon^2/8, 4W^2, \tilde{R}, \tilde{h}, \gamma)$ -excellent is implicit in the event. Squaring and applying Cauchy-Schwarz we get that

$$\begin{aligned} \Omega(\varepsilon^{12}) &\leq \mathbb{E}_{(D_0, E_0, B_0)} \mathbb{E}_{\substack{(E, B) \sim \mathcal{D}(D_0, E_0) \\ (E', B') \sim \mathcal{D}(D_0, E_0)}} \left[\mathbf{1}_{g_{D_0, E_0, B_0} \stackrel{\leq \tilde{W}^3}{\neq} g_{D_0 \cup E, B}} \mathbf{1}_{g_{D_0, E_0, B_0} \stackrel{\leq \tilde{W}^3}{\neq} g_{D_0 \cup E', B'}} \right] \\ &\leq \mathbb{E}_{(D_0, E_0, B_0)} \mathbb{E}_{\substack{(E, B) \sim \mathcal{D}(D_0, E_0) \\ (E', B') \sim \mathcal{D}(D_0, E_0)}} \left[\mathbf{1}_{g_{D_0 \cup E, B} \stackrel{\leq 2\tilde{W}^3}{\neq} g_{D_0 \cup E', B'}} \right] \end{aligned}$$

Now letting $A_0 = D_0 \cup E'$, $B_0 = B'$, $\tilde{A}_0 = D_0 \cup E$ and $B = B$, we see that with probability at least $1 - \exp(-\Omega((cq')^2n))$, we have $|E \cap E'| = O(q^2n)$. Therefore, there exists a constant d_0 such that when we condition the above distribution on $|A_0 \cap \tilde{A}_0| = d_0cq'n$, the expectation is still at least $\Omega(\varepsilon^{12})$. The claim now follows from an averaging argument. \square

13.2.5 Applying Small-set Expansion

Intuitively, Claim 13.30 asserts that after re-randomizing B_0 and perturbing A_0 , the functions g_{A_0, B_0} remain basically the same with noticeable probability. In the next lemma we appeal to small set expansion type results to show that there is a single function g such that g_{A_0, B_0} is very close to g with noticeable probability, hence concluding the proof of Theorem 13.1 for $m = 2$.

Lemma 13.31. *If $\Pr_{(A_0 \cup B_0, A_0 \cup B_1) \sim \mathcal{D}_{q,q'}} [F[A_0, B_0]|_{A_0} \stackrel{\leq T}{\neq} F[A_0, B_1]|_{A_0}] \geq \varepsilon$, then there exists a pair (A_0^*, B_0^*) which is $(\varepsilon^2/8, 2W^2, \tilde{R}, \tilde{h}, \gamma)$ -excellent such that*

$$\Pr_{T \in \binom{[n]}{qn}} \left[g_{A_0^*, B_0^*}(T) \stackrel{\leq \tilde{W}^{O(1)}}{\neq} F[T] \right] \geq \delta,$$

where $\delta = \varepsilon^{O(\log(1/q')^2)}$.

Proof. We first note that if two functions $f, g : [n] \rightarrow \{0, 1\}$ differ at $O(a)$ locations and $a \leq \sqrt{n}$, then for a random subset $S \subseteq [n]$ of size n/a the probability that $f(S) = g(S)$ is at least $\Omega(1)$. Using this fact and Claim 13.30, and by letting $W' = 2\tilde{W}^3$, we have the following

$$\mathbb{E}_{S \subseteq [n] \mid |S| = \frac{n}{W'}} \left[\mathbb{E}_{(A_0, B_0), (\tilde{A}_0, B)} \left[\mathbf{1}_{g_{A_0, B_0}(S) = g_{\tilde{A}_0, B}(S)} \right] \right] \gtrsim \varepsilon^{\Theta(1)}.$$

where (\tilde{A}_0, B) is distributed as in Claim 13.30. Note that the event $\mathbf{1}_{g_{A_0, B_0}(S) = g_{\tilde{A}_0, B}(S)}$ implicitly implies that the pairs $(A_0, B_0), (\tilde{A}_0, B)$ are excellent. By an averaging argument, at least $\Omega(\varepsilon^{\Theta(1)})$ fraction of the sets S are such that

$$\mathbb{E}_{(A_0, B_0), (\tilde{A}_0, B)} \left[\mathbf{1}_{g_{A_0, B_0}(S) = g_{\tilde{A}_0, B}(S)} \right] \gtrsim \varepsilon^{\Theta(1)}.$$

Given such a set S , we define a partition of $\binom{[n]}{qn}$ based on the value $g_{A, B}(S)$. In other words, we have parts identified by strings in $\{0, 1\}^{n/W'^2}$; (A, B) and (A', B') belong to the part \mathcal{S}_τ if $g_{A, B}(S) = g_{A', B'}(S) = \tau$. We claim that there exists a τ such that $|\mathcal{S}_\tau| = \eta \binom{n}{qn}$, for some $\eta = \eta(\varepsilon, q') = \varepsilon^{O(\log(1/q')^2)}$.

Consider the graph \mathcal{G}_n induced on the set of vertices $\{(A, B) \mid |A| = q'n, |B| = (q - q')n\}$ as follows: a random neighbor (A', B') of (A, B) in this graph is sampled conditioned on the fact that A' is distributed uniformly conditioned on $|A'| = q'n$, $|A \cap A'| = c'q'n$ and $B' \subseteq [n] \setminus A'$ is a uniformly random set of size $(q - q')n$. Using the small set expansion property of \mathcal{G}_n from Lemma 13.3, if all the parts were of size at most $\delta \binom{n}{qn}$, then

$$\Pr_{(A_0, B_0), (\tilde{A}_0, B)} \left[(A_0, B_0), (\tilde{A}_0, B) \in \mathcal{S}_\tau \text{ for some } \tau \right] \leq \delta^{O(1/\log(1/q')^2)}.$$

We conclude that there must be a part $\tau^*(S)$ whose fractional size is at least η . Therefore, for any such S we get that

$$\mathbb{E}_{(A_0, B_0), (A'_0, B'_0)} \left[\mathbf{1}_{g_{A_0, B_0}(S) = g_{A'_0, B'_0}(S)} \right] \geq \left(\mathbb{E}_{(A_0, B_0), (A'_0, B'_0)} \left[\mathbf{1}_{g_{A_0, B_0}(S) = \tau^*(S)} \right] \right)^2 \geq \eta^2.$$

Taking expectation over S yields that

$$\mathbb{E}_{S \subseteq [n] \mid |S| = \frac{n}{c'^2}} \left[\mathbb{E}_{(A_0, B_0), (A'_0, B'_0)} \left[\mathbf{1}_{g_{A_0, B_0}(S) = g_{A'_0, B'_0}(S)} \right] \right] \gtrsim \varepsilon^{\Theta(1)} \eta^2.$$

This implies,

$$\mathbb{E}_{(A_0, B_0), (A'_0, B'_0)} \left[\mathbb{E}_{S \subseteq [n] \mid |S| = \frac{n}{c'^2}} \left[\mathbf{1}_{g_{A_0, B_0}(S) = g_{A'_0, B'_0}(S)} \right] \right] \gtrsim \varepsilon^{\Theta(1)} \eta^2,$$

and so by an averaging argument for at least $\Omega(\varepsilon^{\Theta(1)}\eta^2)$ of pairs $(A_0, B_0), (A'_0, B'_0)$ we have that

$$\mathbb{E}_{S \subseteq [n] \mid |S| = \frac{n}{\zeta^{1/2}}} \left[\mathbf{1}_{g_{A_0, B_0}(S) = g_{A'_0, B'_0}(S)} \right] \gtrsim \varepsilon^{\Theta(1)}\eta^2.$$

Using Claim 13.29 we get that

$$\Pr_{(A_0, B_0), (A'_0, B'_0)} \left[g_{A_0, B_0} \stackrel{\leq W'^{O(1)}}{\neq} g_{A'_0, B'_0} \right] \gtrsim \varepsilon^{\Theta(1)}\eta^2.$$

This means that there exists an excellent pair (A_0^*, B_0^*) such that

$$\Pr_{(A'_0, B'_0)} \left[g_{A_0^*, B_0^*} \stackrel{\leq W'^{O(1)}}{\neq} g_{A'_0, B'_0} \right] \gtrsim \varepsilon^{\Theta(1)}\eta^2.$$

Now consider selecting a random set T of size qn and select a random subset A of T of size $q'n$ and let $B = T \setminus A$. Select a random set $B' \subseteq [n] \setminus A$. Using the above inequality, with probability $\Omega(\varepsilon^{\Theta(1)}\eta^2)$, $g_{A_0^*, B_0^*} \stackrel{\leq W'^{O(1)}}{\neq} g_{A, B'}$ and with probability at least $\Omega(\varepsilon^2)$, $B \in \text{Cons}_{2W^2}^*(A, B')$ and hence $g_{A, B'}(T) \stackrel{\leq W'}{\neq} F[T]$. Combining all these events, we have

$$\Pr_{T \in \binom{[n]}{qn}} \left[g_{A_0^*, B_0^*}(T) \stackrel{\leq W'^{O(1)}}{\neq} F[T] \right] \gtrsim \varepsilon^{\Theta(1)}\eta^2 \cdot \varepsilon^2 = \varepsilon^{O(\log(1/q')^2)},$$

and the claim follows. \square

13.2.6 Extending the Result to Large Alphabets

In this section, we prove the direct product result with similar conclusion but for larger alphabet.

Theorem 13.32 (Restatement of Theorem 13.1). *There exists $c > 0$ such that the following holds for sufficiently large $n \in \mathbb{N}$ and $\varepsilon \geq 2^{-nc}$. Fix an alphabet $[m]$ and $c_0 > 0$ be any constant. For all $\varepsilon > 0$, $0 < q' < q < 1$ and $T = \left(\frac{\log(1/\varepsilon)}{q'}\right)^{c_0}$, suppose that $F : \binom{[n]}{qn} \rightarrow [m]^{qn}$ satisfies*

$$\Pr_{(S_1, S_2) \sim \mathcal{D}_{q, q'}} \left[F[S_1] \stackrel{\leq T}{\neq} F[S_2] \mid S_1 \cap S_2 \right] \geq \varepsilon.$$

Then there exists a function $g : [n] \rightarrow [m]$ such that for at least an $\varepsilon^{O(\log(1/q')^2)}$ fraction of $S \in \binom{[n]}{qn}$, we have $|\{i \in S \mid F[S]_i \neq g(i)\}| \leq \left(\frac{T \log(1/\varepsilon)}{q'}\right)^{O(1)}$.

The proof of this theorem follows the same lines as the proof of the binary case. In the binary case, the only place we used the fact that the alphabet is binary is when deriving the local structure, i.e. Lemma 13.9. Given such a local structure for larger alphabet the rest of the proof goes through as is. Thus, to prove the case $m \geq 3$ it suffices to prove a local structure lemma in this setting, and we show it via a reduction to the binary case (Lemma 13.9). We remark that this type of argument already appeared in [28], and we reproduce it here for the sake of completeness.

The definitions of goodness and excellence still makes sense, and we use the same parameters as in the binary case. It is also easy to observe that Claim 13.10 and Claim 13.11 follow in this setting as is. Fix an excellent pair (A_0, B_0) , we can define the function $g_{A_0, B_0} : [n] \rightarrow [m]$ as follows:

$$g_{A_0, B_0}(x) := \text{Plurality}_{B \in \text{Cons}_T(A_0, B_0) | B \ni x} F[A_0, B]|_x.$$

We have the following claim, which is a substitution for Lemma 13.9:

Claim 13.33. *Fix an alphabet $[m]$ and $c_0 > 0$ be any constant. For every $\varepsilon > 0$, and $q, q' \in (0, 1)$, suppose a table $F : \binom{[n]}{qn} \rightarrow [m]^{qn}$ passes the Agreement-Test with parameters (q, q', η) with probability at least ε . If (A_0, B_0) is excellent then*

$$\Pr_{B \in \text{Cons}_T(A_0, B_0)} [F[A_0, B]|_B \stackrel{>10W^2}{\neq} g_{A_0, B_0}(B)] \leq 2\nu.$$

Proof. Let $Z \subseteq \text{Cons}_T(A_0, B_0)$ be the such sets B such that $g_{(A_0, B_0)}(B) \stackrel{>10W^2}{\neq} F[A_0, B]|_B$. Select n random functions $f_1, f_2, \dots, f_n : [m] \rightarrow \{0, 1\}$. Define a table $G_f : \binom{[n]}{qn} \rightarrow \{0, 1\}^{qn}$ as follows: For a given tuple $S = (i_1, i_2, \dots, i_{qn}) \in \binom{[n]}{qn}$, if $F(S) = (a_1, a_2, \dots, a_{qn})$ then set

$$G_f(S) = (f_{i_1}(a_1), f_{i_2}(a_2), \dots, f_{i_{qn}}(a_{qn})).$$

The reason for doing this is that whenever sets (A, B) and (A, B') agree on A w.r.t. the table F , then they also agree on A w.r.t. the map G_f . Therefore, for any fixed set of functions f_i 's, if (A_0, B_0) is excellent with respect to the table F then, the same pair (A_0, B_0) is also excellent with respect to the map G_f . Therefore, by Lemma 13.9, there is a function g_f defined by the majority votes from Cons_T such that for at most a ν fraction of the sets in $B \in \text{Cons}_T$, we have $G_f[A_0, B] \stackrel{\geq W^2}{\neq} g_f(B)$. Let Z_f be the collection of such sets in $B \in \text{Cons}_T$. We will show that if $B \in Z$, then almost surely $B \in Z_f$. Since we know that the density of Z_f in Cons_T is at most ν , this will show that the density of Z in Cons_T is also roughly ν and this will conclude the proof.

Towards showing this, fix a B , fix a $x \in B$ and consider a $u \in [m]$ such that $u \neq g_{A_0, B_0}(x)$. We first show that with probability at least $1/4$ it holds that $g_f(x) \neq f_x(u)$. Indeed, let us change the value of f_x at u and at $g_{A_0, B_0}(x)$ for a moment to both be \perp , and define

$$b = \text{Majority}_{\substack{B \in \text{Cons}_T(A_0, B_0) | B \ni x \\ G_f[A_0, B]|_x \neq \perp}} G_f[A_0, B]|_x.$$

Now consider defining f_x once again on u and $g_{A_0, B_0}(x)$ by uniform $\{0, 1\}$ independently chosen assignments, and the affect of that on the value of b . With probability $1/4$, f_x maps $g_{A_0, B_0}(x)$ to b and u to $1 - b$, in which case as $g_{A_0, B_0}(x)$ is more popular than u , the majority value above will be unchanged and equal to $g_f(x)$. In this case we get that $g_f(x) \neq f_x(u)$.

Fix $B \in Z$ so that $g_{(A_0, B_0)}(B) \stackrel{>10W^2}{\neq} F[A_0, B]|_B$. On each coordinate $x \in B$ where the disagreement occurs, letting $u = F[A_0, B]|_B(x)$ we have $g_f(x) \neq f_x(u)$ with probability at least $1/4$. However, $G_f[A_0, B]|_x = f_x(u)$ by definition and hence $g_f(x) \neq G_f[A_0, B]|_x$ with probability at least $1/4$. For each coordinate $x \in B$ of disagreement these events are independent (as f_x are independently chosen), so

by Chernoff's bound with probability $1 - 2^{-\Omega(W^2)}$ the function g_f disagrees with $G_f[A_0, B]|_B$ on at least $2W^2$ coordinates, in which case $B \in Z_f$.

By linearity of expectation we conclude that the expected density of Z_f inside Cons_T is at least half the density of Z ; thus, as the density of Z_f in Cons_T is at most ν it follows that the density of Z in Cons_T is at most 2ν and the proof is concluded. \square

13.3 Direct Product Testing: From Sets to Product Distributions

In this section, we move from the uniform setting to a product setting. Consider the q -biased measure over $P([n])$, i.e. $\mu_q^{\otimes n}(A) = q^{|A|}(1-q)^{n-|A|}$, and let $G: (P[n], \mu_q^{\otimes n}) \rightarrow [m]^{\leq n}$ be an assignment that to each $A \in P([n])$ assigns a string $G[A] \in [m]^{|A|}$ in a *locally consistent manner*. Namely, for $\alpha \in (0, 1)$, consider the distribution $\mathcal{D}_{q,\alpha}$ over $A, A' \subseteq [n]$ that results from by taking, for each $i \in [n]$ independently, i to in $A \cap A'$ with probability αq , to be in $A \setminus A'$ with probability $(1-\alpha)q$ and to be in $A' \setminus A$ with probability $(1-\alpha)q$. Suppose that

$$\Pr_{(A,A') \sim \mathcal{D}_{q,\alpha}} \left[G[A]|_{(A \cap A')} \stackrel{\leq T}{\neq} G[A']|_{(A \cap A')} \right] \geq \varepsilon.$$

Using Theorem 13.1, we show that G must be correlated with a global $S \in [m]^n$.

Corollary 13.34. *For all $C, c_0 > 0$ there is $c > 0$ such that the following holds for sufficiently large $n \in \mathbb{N}$ and $\varepsilon \geq 2^{-n^c}$.*

Fix the alphabet $[m]$, $q \in (0, 1)$, $T = \left(\frac{\log(1/\varepsilon)}{q}\right)^{c_0}$ and suppose that $0.9 > \alpha \geq \frac{1}{\log^C(1/\varepsilon)}$. If $G: (P[n], \mu_q^{\otimes n}) \rightarrow [m]^{\leq n}$ satisfies

$$\Pr_{(A,A') \sim \mathcal{D}_{q,\alpha}} \left[G[A]|_{(A \cap A')} \stackrel{\leq T}{\neq} G[A']|_{(A \cap A')} \right] \geq \varepsilon.$$

Then there exists a string $S \in [m]^n$ such that $\Pr_{A \sim_q [n]} [\Delta(G[A], S|_A) \leq r] \geq \delta$, where $\delta = \varepsilon^{O(\log(1/q)^2)}$ and $r = \left(\frac{\log(1/\varepsilon)}{q}\right)^{O(1)}$.

Proof. Fix $N = \omega(n^4)$. Given a function $G: (P[n], \mu_q^{\otimes n}) \rightarrow [m]^{\leq n}$ where $G(A)$ can be thought of as a string in $[m]^{|A|}$ by specifying a fixed arbitrary ordering on $[n]$, we define a map $\tilde{G}: \binom{[N]}{qN} \rightarrow [m]^{qN}$ as follows. For a set $S \in \binom{[N]}{qN}$, define $\tilde{G}(S)|_{S \cap [n]} = G(S \cap [n])$ and $\tilde{G}(S)|_{S \setminus [n]} = 0^{|S \setminus [n]|}$. We know that

$$\Pr_{(A,A') \sim \mathcal{D}_{q,\alpha}} [G[A]|_{(A \cap A')} \stackrel{\leq T}{\neq} G[A']|_{(A \cap A')}] \geq \varepsilon.$$

Instead of checking consistency on $A \cap A'$, we select a set $A'' \subseteq A \cap A'$ by independently including $i \in A \cap A'$ to A'' with probability $q'/\alpha q$. This can only increase the acceptance probability, and hence

$$\Pr_{\substack{(A,A') \sim \mathcal{D}_{q,\alpha} \\ A'' \sim_{q'/\alpha q} A \cap A'}} [G[A]|_{A''} \stackrel{\leq T}{\neq} G[A']|_{A''}] \geq \varepsilon. \quad (54)$$

We denote the overall distribution on (A, A', A'') from the above probability by \mathcal{D} . Now consider selecting the pairs (A_0, B_0) and (A_0, B_1) according to the *Agreement-Test* with parameters (q, q', η) for the table \tilde{G} . Let $\tilde{A} = A_0 \cup B_0 \cap [n]$, $\tilde{A}' = A_0 \cup B_1 \cap [n]$ and $\tilde{A}'' = A_0 \cap [n]$. We use $\tilde{\mathcal{D}}$ to denote the distribution on $(\tilde{A}, \tilde{A}', \tilde{A}'')$. We show that the statistical distance between the distributions \mathcal{D} and $\tilde{\mathcal{D}}$ is at most $o(1)$.

Claim 13.35. *The statistical distance between the distributions \mathcal{D} and $\tilde{\mathcal{D}}$ is at most $e^{-\Omega(n)}$ when $\alpha q = q' + \frac{(q-q')^2}{(1-q')}$.*

Proof. Deferred to Section D. □

Using Claim 13.35 and (54), we conclude $\Pr[G[\tilde{A}]|_{\tilde{A}'} \stackrel{\leq T}{\neq} G[\tilde{A}']|_{\tilde{A}'}] \geq \varepsilon - 2^{-\Omega(n)}$. Since $\tilde{G}(S)|_i = 0$ for every $i > n$ and $S \ni i$, we have

$$\Pr_{(A_0, B_0), (A_0, B_1)} [\tilde{G}[A_0, B_0]|_{A_0} \stackrel{\leq T}{\neq} \tilde{G}[A_0, B_1]|_{A_0}] \geq \varepsilon - 2^{-\Omega(n)},$$

Therefore, using Theorem 13.1, we conclude that there exists a global function $\tilde{g} : [N] \rightarrow [m]$ such that

$$\Pr_{S \in \binom{[N]}{qN}} [\tilde{G}[S] \stackrel{\leq r}{\neq} \tilde{g}(S)] \geq \delta,$$

where r and δ are as in Theorem 13.1. Furthermore, based on how we came up with the global function \tilde{g} , we have $\tilde{g}(i) = 0$ for all $i \in (n, N]$ as $\tilde{G}(S)|_i = 0$ for all $S \in \binom{[N]}{qN}$ and $i \in S$. If we let $g : [n] \rightarrow [m]$ be the function \tilde{g} restricted to the domain $[n]$, we have

$$\Pr_{A \sim \mu_q^{\otimes n}} [G[A] \stackrel{\leq r}{\neq} g(A)] \geq \delta - 2^{-\Omega(n)}.$$

Here, we used Claim 13.5 that shows the statistical distance between the distribution $\mu_q^{\otimes n}$ and the distribution on $S|_{[n]}$ where S is a uniformly random set from $\binom{[N]}{qN}$ is at most $2^{-\Omega(n)}$. This concludes the proof. □

13.4 Getting the Final Direct Product Theorem: Proof of Theorem 12.7

In this section we analyze the main direct product test from Theorem 12.7, which is a slight variant of the test given in Corollary 13.34. Let $G : (P[n], \mu_q^{\otimes n}) \rightarrow [m]^{\leq n}$ be an assignment that to each $A \in P([n])$ assigns a string $G[A] \in [m]^{|A|}$. We show that if G passes this test from Theorem 12.7 with probability at least $\varepsilon > 0$, it also passes the test from Corollary 13.34 with probability at least $\varepsilon/2$ and hence we get the same global structure.

Given $G : (P[n], \mu_q^{\otimes n}) \rightarrow [m]^{\leq n}$,

- Sample (A, A') from the distribution $\mathcal{D}_{q, \alpha}$.
- Select a random subset $\tilde{A}_0 \subseteq A \cap A'$ by adding each $i \in A \cap A'$ to \tilde{A}_0 with probability β independently.
- Check if $G[A]|_{\tilde{A}_0} = G[A']|_{\tilde{A}_0}$

Figure 3: *Subset-Agreement-Test* with parameters (q, α, β) .

We first analyze the *Subset-Agreement-Test* given in Figure 13.4. In comparison to Corollary 13.34, in this test we check for complete agreement (i.e., $\eta = 0$), but only on a subset of coordinates from $A \cap A'$.

Theorem 13.36. For all $C, c_0 > 0$ there is $c > 0$ such that the following holds for sufficiently large $n \in \mathbb{N}$ and $\varepsilon \geq 2^{-n^c}$.

Suppose that $\alpha, \beta \geq \frac{1}{\log^C(1/\varepsilon)}$, $\alpha \leq \frac{9}{10}$ and further suppose that $G: (P[n], \mu_q^{\otimes n}) \rightarrow [m]^{\leq n}$ passes the Subset-Agreement-Test with parameters (q, α, β) with probability at least ε , then there exists a string $S \in [m]^n$ such that

$$\Pr_{A \sim_q [n]} [\Delta(G[A], S|_A) \leq r] \geq \delta,$$

where $\delta = \varepsilon^{O(\log(1/q)^2)}$ and $r = \left(\frac{\log(1/\varepsilon)}{q}\right)^{O(1)}$.

Proof. Let $c_0 \gg C$. We show that if G passes the *Subset-Agreement-Test* with parameters (q, α, β) with probability ε , then it also passes the test from Corollary 13.34, with respect to the distribution $\mathcal{D}_{q, \alpha}$ and $\eta = \left(\frac{1}{q\beta}\right)^{c_0}$, with probability at least $\varepsilon/2$. Let $E_{>\eta}$ be the event that $G[A]|_{A \cap A'} \stackrel{>\eta}{\neq} G[A']|_{A \cap A'}$, and $E_{\leq \eta}$ be the event that $G[A]|_{A \cap A'} \stackrel{\leq \eta}{\neq} G[A']|_{A \cap A'}$. Let $p_{>\eta}$ and $p_{\leq \eta}$ be the probability of events $E_{>\eta}$ and $E_{\leq \eta}$ respectively. We have

$$\begin{aligned} \varepsilon &\leq \Pr_{\substack{(A, A') \sim \mathcal{D}_{q, \alpha}, \\ \tilde{A}_0 \sim_{\beta} A \cap A'}} [G[A]|_{\tilde{A}_0} = G[A']|_{\tilde{A}_0}] = p_{>\eta} \cdot \Pr_{\substack{(A, A') \sim \mathcal{D}_{q, \alpha}, \\ \tilde{A}_0 \sim_{\beta} A \cap A'}} [G[A]|_{\tilde{A}_0} = G[A']|_{\tilde{A}_0} \mid E_{>\eta}] \\ &\quad + p_{\leq \eta} \cdot \Pr_{\substack{(A, A') \sim \mathcal{D}_{q, \alpha}, \\ \tilde{A}_0 \sim_{\beta} A \cap A'}} [G[A]|_{\tilde{A}_0} = G[A']|_{\tilde{A}_0} \mid E_{\leq \eta}] \\ &\leq 1 \cdot (1 - \beta)^\eta + p_{\leq \eta} \cdot 1. \\ &\leq \varepsilon/2 + p_{\leq \eta}. \end{aligned}$$

This shows that $p_{\leq \eta} \geq \varepsilon/2$ and hence G passes the test from Corollary 13.34 with probability at least $\varepsilon/2$. We can now apply Corollary 13.34 on G to get the conclusion.

We are now ready to prove Theorem 12.7.

Proof of Theorem 12.7: The distribution in the test $\text{DP}(\rho, \alpha, \beta)$ can be simulated by the *Subset-Agreement-Test* with parameters $(q, \tilde{\alpha}, \tilde{\beta})$ where $q = \rho$, $\tilde{\alpha}q = \rho\alpha + (\rho - \alpha\rho)^2$ and $\tilde{\beta}\tilde{\alpha}q = \beta\alpha\rho$. Therefore, the proof follows from Theorem 13.36.

13.5 Small-set Expansion Property of the Graphs over a Multi-slice

In this section, we show the small set expansion property of the graph as stated in Lemma 13.3.

Recall the graph \mathcal{G}_n that we defined on the set of vertices $\{(A, B) \mid A, B \subseteq [n], A \cap B = \emptyset, |A| = q'n, |B| = (q - q')n\}$. A random neighbor (A', B') of (A, B) in this graph is sampled conditioned on the fact that A' is distributed uniformly conditioned on $|A'| = q'n$, $|A \cap A'| = cq'n$ and $B' \subseteq [n] \setminus A'$ is a uniformly random set of size $(q - q')n$. In this section, we deduce the small-set expansion property of the graph \mathcal{G}_n . We can view the vertices of the above graph \mathcal{G}_n as the multi-slice of $\{0, 1, 2\}^n$ – map the vertex (A, B) to $x \in \{0, 1, 2\}^n$ where $x_i = 1$ if $i \in A$, $x_i = 2$ if $i \in B$ and $x_i = 0$ if $i \in [n] \setminus (A \cup B)$. Let us denote the multi-slice by \mathcal{U}_n .

One can view the multi-slice \mathcal{U}_n as a quotient space $S_n / (S_{(1-q)n} \times S_{q'n} \times S_{(q-q')n})$, which is useful in lifting the standard representation-theoretic decomposition of functions over S_n to decompositions of functions over \mathcal{U}_n . In order to state the relevant lemmas from [9], we need the following few definitions.

Definition 13.37. A function $f : S_n \rightarrow \mathbb{R}$ is called a d -junta if there exists a set of coordinate $A \subseteq [n]$ of size at most d such that $f(\pi) = g(\pi(A))$ for some function $g : [n]^A \rightarrow \mathbb{R}$.

For two function $f, g : S_n \rightarrow \mathbb{R}$, define the inner product $\langle f, g \rangle$ as $\mathbb{E}_\pi [f(\pi)g(\pi)]$.

Definition 13.38. For $d = 0, 1, \dots, n$ we denote by $V_d(S_n) \subseteq \{f : S_n \rightarrow \mathbb{R}\}$ the span of d -juntas. Also, define $V_{=d}(S_n) = V_d(S_n) \cap V_{d-1}(S_n)^\perp$.

Therefore, we can write the space of real-valued functions as $V_{=0}(S_n) \oplus V_{=1}(S_n) \oplus \dots \oplus V_{=n-1}(S_n)$, and thus write any $f : S_n \rightarrow \mathbb{R}$ uniquely as $f = \sum_{i=0}^{n-1} f^{=i}$ where $f^{=i} \in V_{=i}(S_n)$. Let $V_{\geq d}(\mathcal{U}_n)$ ($V_{\leq d}(\mathcal{U}_n)$) be the span of functions over \mathcal{U}_n whose degree is at least (at most) d .

We say a distribution μ over $([3] \times [3])^n$ commutes with the action of S_n if the following distributions are the same for all $\pi \in S_n$ and $x \in [3]^n$: a) x' , where $(x, x') \sim \mu$ conditioned on $x = \pi(x)$, and b) $\pi(x')$, where $(x, x') \sim \mu$ conditioned on $x = x$. The following claim shows that the operator \mathcal{T} that commutes with S_n preserves the degree of the functions.

Claim 13.39. (Claim 3.6 from [9]) Suppose \mathcal{T} is an operator that commutes with the action of S_n on functions over the multi-slice \mathcal{U}_n . Then for each $0 \leq d < n$, we have that $\mathcal{T}(V_{=d}(\mathcal{U}_n)) \subseteq V_{=d}(\mathcal{U}_n)$.

We observe the following few properties of the multi-slice \mathcal{U}_n and the operator $\mathcal{T}(\mathcal{G}_n)$.

1. In a multi-slice, every symbol appears the same number of times in every element of the multi-slice. If we let k_i be the number of times the symbol i appears, then the multi-slice is called α -balanced if $k_i \geq \alpha n$ for every i . The multi-slice \mathcal{U}_n is α -balanced for $\alpha = \min\{q', (q - q'), (1 - q)\}$.
2. The edge distribution of the graph \mathcal{G}_n is a distribution on the multi-slices $\mathcal{U}_n \times \mathcal{U}_n$. A distribution μ on $\mathcal{U}_n \times \mathcal{U}_n$ is called α -admissible if a) the distribution is symmetric under S_n , and b) for all $(a, b) \in \{0, 1, 2\} \times \{0, 1, 2\}$, the quantity $\Pr_{(x,y) \sim \mu, i \in [n]} [x_i = a \ \& \ y_i = b]$ is either at least α or 0. It can be easily observed that the edge distribution of \mathcal{G}_n is α -admissible for $\alpha = \Omega((cq')^2)$.
3. A distribution μ on $\mathcal{U}_n \times \mathcal{U}_n$ also called *connected* iff the bipartite graph $(V_1 \cup V_2, E)$ where a) V_i is the corresponding support of the marginal distribution of μ , and b) $(x, y) \in E$ iff (x, y) is in the support of μ , is connected. Here again, it is easy to observe the connectedness property of the edge distribution of the graph \mathcal{G}_n .
4. Finally, the operator $\mathcal{T}(\mathcal{G}_n)$ commutes with the action of S_n .

One of the important characteristic of α -admissible and connected distributions, as shown in [9], is that it can be replaced by a certain product distribution $\nu^{\otimes n}$ as far as low-degree functions are concerned. Thus, this gives a way to prove analytical results for a multi-slice by invoking the corresponding results over a product distribution.

The following lemma from [18, 17] gives an upper bound on $\|\mathcal{P}_{\leq d}\|_{2 \rightarrow 4}$, where $\mathcal{P}_{\leq d}$ is the projector operator into the space $V_{\leq d}(\mathcal{U}_n)$. It crucially uses the fact that \mathcal{U}_n is α -balanced.

Lemma 13.40 (Lemma 29 from [18]). For all $c > 0$ and $0 < q' < q < 1$ and $d \in \mathbb{N}$, if $f : \mathcal{U}_n \rightarrow \mathbb{R}$ is a function of degree at most d , then $\|f\|_4 \leq \left(\frac{1}{cq'}\right)^{O(d)} \|f\|_2$.

Finally, we need the following lemma that bounds the eigenvalues of the operator $\mathcal{T}(\mathcal{G}_n)$. This lemma uses the fact the the edge distribution is α -admissible and connected.

Lemma 13.41 (Lemma 3.11 from [9]). *There are constants $C > 0$ and $\delta > 0$ such that for all $c > 0$ and $0 < q' < q < 1$ such that $q' = \Omega(q)$, for all $d \in \mathbb{N}$, if $f \in V_{>d}(\mathcal{U}_n)$, we have $\|\mathcal{T}(\mathcal{G}_n)f\|_2 \leq C(1 + \delta)^{-\frac{d}{\log(1/cq')}} \|f\|_2$.*

We note that the Lemma 3.11 from [9] gives a bound of the form $\|\mathcal{T}(\mathcal{G}_n)f\|_2 \leq C(1 + \delta)^{-d} \|f\|_2$, where the parameters C and δ depend on the parameters q', q , or more generally the parameter α from the α -balancedness and α -admissible property. The dependence of α on δ can be as bad as $\delta \sim \alpha^{-1}$ and such a bound will give us $\|\mathcal{T}(\mathcal{G}_n)f\|_2 \leq C(1 + \delta)^{-d \cdot (cq')^2} \|f\|_2$ which is not sufficient for our purpose. We observe that a slight adjustment of their proof gives a better quantitative bound as stated in the lemma in our current setting. We first describe the difference between our setting and the setting in [9] and then briefly sketch the proof of Lemma 13.41.

The α -admissible property is used to get bound the second eigenvalue of a graph H defined as follows: The vertex set is $\{0, 1, 2\}$ and for $a, b \in \{0, 1, 2\}$, the weight of the edge (a, b) is

$$w(a, b) = \Pr_{(x, x') \sim \mathcal{G}_n, i \in [n]} [x'_i = b \mid x_i = a].$$

Therefore, if the edge distribution of \mathcal{G}_n is α -admissible, then the non-zero weight of an edge in the graph H is at least α , and hence $\lambda_2(H) \leq 1 - \alpha^2/2$. In our case, the stationary distribution of H is given by ν where $\nu(0) = q', \nu(1) = (q - q')$ and $\nu(2) = (1 - q)$. As $q' = \Omega(q)$, in our case, we have $\lambda_2(H) = 1 - \Omega(1)$. Therefore, we do not lose much in the final bound. We now sketch the proof and assume readers to be familiar with the concepts from [9].

Proof sketch of Lemma 13.41: For convenience we write $T := \mathcal{T}(\mathcal{G}_n)$ and let $\alpha = \Omega((cq')^2)$. Suppose $f \in V_{>d}(\mathcal{U}_n)$. We can write the function f in as the sum of eigenvectors f_1, f_2, \dots of the operator T . Therefore, it is enough to given an upper bound the eigenvalue of T corresponding to the functions f_i , $\|Tf_i\|_2 \leq C(1 + \delta)^{-d} \|f_i\|_2$. Let θ be the eigenvalue corresponding to the function f_i . In [9], the authors show $\theta \leq C(1 + \delta)^{-d}$ using the trace method as follows: suppose the multiplicity of θ is m , then using the fact that the trace of an operator is the sum of its eigenvalues, we have

$$m\theta \leq \text{Tr}(T).$$

This gives a bound $\theta \leq \frac{\text{Tr}(T)}{m}$. In order to get a reasonable bound, instead of bounding the trace of T , in [9], the authors work with T^h for some $h \geq 1$ and get

$$\theta^h \leq \frac{\text{Tr}(T^h)}{m}.$$

Since the degree of f was at least d to begin with, using this fact one can show that $m \geq c_0^d$, for some absolute constant $c_0 > 1$. Therefore, it is enough to bound $\text{Tr}(T^h)$. Note that if we show $\text{Tr}(T^h) \leq (1 + \xi)^d$, ξ is an absolute constant and $1 + \xi < c_0$, then we get $\theta = (1 + \delta)^{-d/h}$. Therefore, if we can show that $\text{Tr}(T^h) \leq (1 + \xi)^d$ when $h = \Theta(\log(1/cq'))$, then this will be enough to prove Lemma 13.41. We show this is indeed the case when $\lambda_2(H) \leq 1 - \Omega(1)$.

The crucial point is that as $\lambda_2(H) \leq 1 - \Omega(1)$, we have the following claim which gives a stronger quantitative bound on h with respect to ε .

Claim 13.42. *For all $\varepsilon > 0$ and $h = O(\log(1/\varepsilon))$ such that for all $v_1 \in [3]$*

$$\left| \Pr_{\substack{v_2, v_3, \dots, v_h \\ \text{random walk on } H \text{ from } v_1}} [v_h = a] - \nu(a) \right| \leq \varepsilon.$$

Proof. Using the standard spectral argument, the left-hand side quantity can be upper bounded by $\lambda_2(H)^h$ and since $\lambda_2(H) \leq 1 - \Omega(1)$, the claim follows. \square

In order to give a concrete bound on the trace of T^h using combinatorial analysis, the proof goes by writing the operator T^h as a sum of operators of the form $R_{\vec{r}}$ where $\vec{r} = (r_{a,b})_{a,b \in [3]}$ that sum to n . The operator $R_{\vec{r}}$ corresponds to the distribution on (\mathbf{x}, \mathbf{y}) where $\mathbf{x} \in \mathcal{U}_n$ is uniform and \mathbf{y} is sampled according to T^h conditioned on the statistics if (\mathbf{x}, \mathbf{y}) being \vec{r} . Thus, we can write

$$T^h = \sum_{\vec{r}} p_{\vec{r}} R_{\vec{r}}.$$

Call \vec{r} ε -reasonable $|r_{a,b} - n\nu(a)\nu(b)| \leq 3\varepsilon n$ for all $a, b \in [3]$; otherwise call \vec{r} unreasonable. Claim 3.15 from [9] shows that $\sum_{\text{unreasonable } \vec{r}} p_{\vec{r}} \leq 100\alpha^{-1}e^{-\alpha\varepsilon^2 n}$. Therefore, we can focus on reasonable \vec{r} . The actual argument is based on whether $d \geq \gamma n$ for some constant γ of $d \leq \gamma n$. Since we do not care about the constants in the exponent, we can focus on the latter case when $d < \gamma n$.

As shown in [9], when $d \leq \gamma n$, one can assume that $n = 3d$ in order to compute the trace of the operator $R_{\vec{r}}$ (and hence the trace of T^h). Unraveling the quantitative bounds, the Claim 3.24 from the paper shows that

$$\text{Tr}(R_{\vec{r}}) \leq \left(1 + O\left(\frac{\varepsilon}{\alpha^2}\right)\right)^d.$$

Therefore if we set $\varepsilon = O_{c_0}(\alpha^2)$, then using Claim 13.42 we get $\text{Tr}(T^h) \leq (1 + \xi)^d$ for $h = \Theta(\log(1/cq^l))$ as required, ignoring the unreasonable \vec{r} as their total weight is $\ll (1 + \xi)^d$. \square

We are now ready to prove the small-set expansion property of the graph \mathcal{G}_n .

Lemma 13.43 (Restatement of Lemma 13.3). *For every $c > 0$, $0 < q^l < q < 1$ and $\mu > 0$, the graph \mathcal{G}_n defined above has*

$$\phi_{\mathcal{G}_n}(\mu) \geq 1 - \mu^{\Omega\left(\frac{1}{\log(1/cq^l)^2}\right)}.$$

Proof. Fix any set $S \subseteq V(\mathcal{G}_n)$ of density at most μ . Let $f : \mathcal{U}_n \rightarrow \{0, 1\}$ be the indicator function of S . As stated at the beginning of the section, we can write f as $f = \sum_{i=0}^{n-1} f^{=i}$ where $f^{=i} \in V_{=i}(\mathcal{U}_n)$. Let $f = f_1 + f_2$, where the component $f_1 = \sum_{i=0}^d f^{=i}$ and $f_2 = \sum_{i=d+1}^{n-1} f^{=i}$ for some d to be fixed later. By letting $\mathcal{T} := \mathcal{T}(\mathcal{G}_n)$, we have

$$\phi_{\mathcal{G}_n}(S) = 1 - \frac{\langle f, \mathcal{T}f \rangle}{\mu}.$$

By letting $\tau := C(1 + \delta)^{-\frac{d}{\log(1/cq^l)}}$ from Lemma 13.41, we can bound

$$\begin{aligned} \langle f, \mathcal{T}f \rangle &= \langle f_1, \mathcal{T}f_1 \rangle + \langle f_2, \mathcal{T}f_2 \rangle && \text{(Using Claim 13.39 and orthogonality of spaces } V_{=i}(\mathcal{U}_n)) \\ &\leq \|f_1\|_2^2 + \tau \|f_2\|_2^2 && \text{(Using Claim 13.41)} \\ &\leq \|f_1\|_2^2 + \tau \mu. \end{aligned}$$

Therefore,

$$\phi_{\mathcal{G}_n}(S) \geq 1 - \tau - \frac{\|f_1\|_2^2}{\mu}. \quad (55)$$

If we let $\mathcal{P}_{\leq d}$ be the projector operator into the subspace $V_{\leq d}(\mathcal{U}_n)$, then we have

$$\|\mathcal{P}_{\leq d}\|_{4/3 \rightarrow 2} = \max_{g \neq 0} \frac{\|\mathcal{P}_{\leq d}g\|_2}{\|g\|_{4/3}} \geq \frac{\|\mathcal{P}_{\leq d}f\|_2}{\|f\|_{4/3}} = \frac{\|f_1\|_2}{\|f\|_{4/3}}.$$

Since $\|f\|_{4/3} = \mu^{3/4}$, we have $\|f_1\|_2^2 \leq \|\mathcal{P}_{\leq d}\|_{4/3 \rightarrow 2}^2 \cdot \mu^{3/2}$. Using the fact that $\|\mathcal{P}_{\leq d}\|_{4/3 \rightarrow 2}^2 = \|\mathcal{P}_{\leq d}\|_{2 \rightarrow 4}^2$, we get $\|f_1\|_2^2 \leq \|\mathcal{P}_{\leq d}\|_{2 \rightarrow 4}^2 \cdot \mu^{3/2}$. Therefore,

$$\phi_{\mathcal{G}_n}(S) \geq 1 - \tau - \|\mathcal{P}_{\leq d}\|_{2 \rightarrow 4}^2 \cdot \mu^{1/2}. \quad (56)$$

Finally, using Lemma 13.40, we can bound $\|\mathcal{P}_{\leq d}\|_{2 \rightarrow 4}^2$ as follows. For any function g , let g_1 be the component of g from $V_{\leq d}$ and $g_2 := g - g_1$ be orthogonal to g_1 .

$$\|\mathcal{P}_{\leq d}\|_{2 \rightarrow 4}^2 = \max_{g:=g_1+g_2 \neq 0} \frac{\|\mathcal{P}_{\leq d}g\|_4^2}{\|g\|_2^2} = \max_{g:=g_1+g_2} \frac{\|g_1\|_4^2}{\|g_1\|_2^2 + \|g_2\|_2^2} \leq \max_{g_1 \in V_{\leq d}} \frac{\|g_1\|_4^2}{\|g_1\|_2^2} \leq (1/cq')^{O(d)},$$

where the last inequality uses Lemma 13.40. Plugging this into (56), we get

$$\phi_{\mathcal{G}_n}(\mu) \geq 1 - C(1 + \delta)^{-\frac{d}{\log(1/cq')}} - (1/cq')^{O(d)} \cdot \mu^{1/2}.$$

If we choose $d = O\left(\frac{\log(1/\mu)}{\log(1/cq')}\right)$, then it is easy to observe that $\phi_{\mathcal{G}_n}(\mu) \geq 1 - \mu^{\Omega\left(\frac{1}{\log(1/cq')^2}\right)}$. \square

References

- [1] Mihir Bellare, Don Coppersmith, Johan Håstad, Marcos A. Kiwi, and Madhu Sudan. Linearity testing in characteristic two. *IEEE Trans. Inf. Theory*, 42(6):1781–1795, 1996.
- [2] Amey Bhangale, Subhash Khot, Noam Lifshitz, and Dor Minzer. Effective bounds for restricted 3-arithmetic progressions in \mathbb{F}_p^n .
- [3] Amey Bhangale, Subhash Khot, and Dor Minzer. A new invariance principle and relations to rounding schemes.
- [4] Amey Bhangale, Subhash Khot, and Dor Minzer. On approximability of satisfiable k -csps: I. In *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 976–988, 2022.
- [5] Amey Bhangale, Subhash Khot, and Dor Minzer. On approximability of satisfiable k -csps: II. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 632–642, 2023.
- [6] Amey Bhangale, Subhash Khot, and Dor Minzer. On approximability of satisfiable k -csps: III. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 643–655, 2023.
- [7] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *J. Comput. Syst. Sci.*, 47(3):549–595, 1993.

- [8] Mark Braverman and Ankit Garg. Small value parallel repetition for general games. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 335–340. ACM, 2015.
- [9] Mark Braverman, Subhash Khot, Noam Lifshitz, and Dor Minzer. An invariance principle for the multi-slice, with applications. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 228–236. IEEE, 2022.
- [10] Mark Braverman, Subhash Khot, and Dor Minzer. Parallel repetition for the GHZ game: Exponential decay. *CoRR*, abs/2211.13741, 2022.
- [11] Irit Dinur, Yuval Filmus, and Prahladh Harsha. Analyzing boolean functions on the biased hypercube via higher-dimensional agreement tests. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2124–2133. SIAM, 2019.
- [12] Irit Dinur and Elazar Goldenberg. Locally testing direct product in the low error range. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 613–622. IEEE, 2008.
- [13] Irit Dinur, Prahladh Harsha, Rakesh Venkat, and Henry Yuen. Multiplayer parallel repetition for expanding games. In *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, volume 67 of *LIPICs*, pages 37:1–37:16, 2017.
- [14] Irit Dinur and Omer Reingold. Assignment testers: Towards a combinatorial proof of the pcg theorem. *SIAM Journal on Computing*, 36(4):975–1024, 2006.
- [15] Irit Dinur and David Steurer. Analytical approach to parallel repetition. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 624–633. ACM, 2014.
- [16] Irit Dinur and David Steurer. Direct product testing. In *2014 IEEE 29th Conference on Computational Complexity (CCC)*, pages 188–196. IEEE, 2014.
- [17] Yuval Filmus, Ryan O’Donnell, and Xinyu Wu. A log-sobolev inequality for the multislice, with applications. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPICs*, pages 34:1–34:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [18] Yuval Filmus, Ryan ODonnell, and Xinyu Wu. Log-sobolev inequality for the multislice, with applications. *Electronic Journal of Probability*, 27:1–30, 2022.
- [19] Uma Girish, Justin Holmgren, Kunal Mittal, Ran Raz, and Wei Zhan. Parallel repetition for the GHZ game: A simpler proof. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2021, August 16-18, 2021, University of Washington, Seattle, Washington, USA (Virtual Conference)*, volume 207 of *LIPICs*, pages 62:1–62:19, 2021.
- [20] Uma Girish, Justin Holmgren, Kunal Mittal, Ran Raz, and Wei Zhan. Parallel repetition for all 3-player games over binary alphabet. In *STOC ’22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 998–1009. ACM, 2022.
- [21] Uma Girish, Kunal Mittal, Ran Raz, and Wei Zhan. Polynomial bounds on parallel repetition for all 3-player games with binary inputs. In *Approximation, Randomization, and Combinatorial Optimization*.

- Algorithms and Techniques, APPROX/RANDOM 2022, September 19-21, 2022, University of Illinois, Urbana-Champaign, USA (Virtual Conference)*, volume 245 of *LIPICs*, pages 6:1–6:17, 2022.
- [22] William T Gowers. A new proof of Szemerédi’s theorem. *Geometric & Functional Analysis GAFA*, 11(3):465–588, 2001.
- [23] Ben Green. 100 open problems. *manuscript*.
- [24] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.
- [25] Thomas Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory Comput.*, 5(1):141–172, 2009.
- [26] Justin Holmgren and Ran Raz. A parallel repetition theorem for the GHZ game. *CoRR*, abs/2008.05059, 2020.
- [27] Russell Impagliazzo, Ragesh Jaiswal, Valentine Kabanets, and Avi Wigderson. Uniform direct product theorems: Simplified, optimized, and derandomized. *SIAM J. Comput.*, 39(4):1637–1665, 2010.
- [28] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. New direct-product testers and 2-query pcps. *SIAM Journal on Computing*, 41(6):1722–1768, 2012.
- [29] 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.
- [30] Elchanan Mossel. Gaussian bounds for noise correlation of functions. *Geometric and Functional Analysis*, 19(6):1713–1756, 2010.
- [31] Elchanan Mossel, Ryan O’Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: Invariance and optimality. *Annals of Mathematics*, 171(1):295–341, 2010.
- [32] Ryan O’Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.
- [33] DHJ Polymath. A new proof of the density Hales-Jewett theorem. *Annals of Mathematics*, pages 1283–1327, 2012.
- [34] Prasad Raghavendra. Optimal algorithms and inapproximability results for every csp? In *Proceedings of the fortieth annual ACM symposium on Theory of computing (STOC)*, pages 245–254, 2008.
- [35] Anup Rao. Parallel repetition in projection games and a concentration bound. *SIAM J. Comput.*, 40(6):1871–1891, 2011.
- [36] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998.
- [37] Klaus F Roth. On certain sets of integers. *J. London Math. Soc.*, 28(104-109):3, 1953.
- [38] Endre Szemerédi. On sets of integers containing no k elements in arithmetic progression. *Acta Arith.*, 27(299-345):21, 1975.
- [39] Oleg Verbitsky. Towards the parallel repetition conjecture. *Theoretical Computer Science*, 157(2):277–282, 1996.

A Missing Proofs

A.1 Merging Symbols: Proof of Lemma 5.5

Consider $T: L_2(\Sigma; \mu_x) \rightarrow L_2(\Gamma \times \Phi; \mu_{y,z})$ defined as

$$Tf(y, z) = \mathbb{E}_{(x', y', z') \sim \mu} [f(x') \mid y' = y, z' = z],$$

and let $S = T^*T$. Note that S can also be viewed as a Markov chain over Σ , and for $a \in \Sigma$ we denote by Sa the distribution over Σ of the neighbours of a according to this Markov chain. Let $G = (\Sigma, E)$ be the graph in which the edges are the support of the Markov chain S , and note that y, z imply z if and only if the number of connected components of G is exactly $|\Sigma|$. Thus, we assume that the number of connected components is $\ell < |\Sigma|$, and we write them as $C_1 \cup \dots \cup C_\ell$.

Note that S is a symmetric operator, hence we may diagonalize it on each connected components separately. Namely, we may find functions $\chi_{i,j}: \Sigma \rightarrow \mathbb{R}$ for $i = 1, \dots, \ell$ and $j = 0, \dots, |C_i| - 1$ that are an orthogonal basis of $L_2(\Sigma; \mu_x)$ and furthermore:

1. $\chi_{i,j}$ is only supported on C_i .
2. $\chi_{i,0}$ is constant on C_i .
3. $\chi_{i,j}$ are eigenfunctions of S with eigenvalue $\lambda_{i,j}$.

Thus, for all i we have $\lambda_{i,0} = 1$, and for $j > 0$ we have as in Lemma 2.7 that $\lambda_{i,j} \leq 1 - \Omega_{\alpha,m}(1)$, and $\lambda_{i,j} \geq 0$ as S is positive semi-definite.

We consider the orthonormal basis $\chi_{\vec{i}, \vec{j}}$ now over $L_2(\Sigma^n, \mu_x^{\otimes n})$ defined by $\chi_{\vec{i}, \vec{j}}(x) = \prod_{k=1}^n \chi_{i_k, j_k}(x_k)$.

The non-merged degree of a monomial $\chi_{\vec{i}, \vec{j}}$ is defined to be the number of k 's such that $j_k > 0$. Thus, we can write

$$f(x) = \sum_{\vec{i}, \vec{j}} \widehat{f}(\vec{i}, \vec{j}) \chi_{\vec{i}, \vec{j}}(x), \quad \text{where} \quad \widehat{f}(\vec{i}, \vec{j}) = \langle f, \chi_{\vec{i}, \vec{j}} \rangle_{L_2(\Sigma^n; \mu_x^n)}.$$

Take $d = W \log(1/\delta)$ for a parameter $W = W(\alpha, m) > 0$ to be chosen later, and write $f = f_1 + f_2$ where f_1 is the part of f with non-merged degree less than d , and f_2 is the part of f with non-merged at least d . Then

$$\mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h(z)] = \underbrace{\mathbb{E}_{(x,y,z) \sim \mu} [f_1(x)g(y)h(z)]}_{(I)} + \underbrace{\mathbb{E}_{(x,y,z) \sim \mu} [f_2(x)g(y)h(z)]}_{(II)}, \quad (57)$$

and we upper bound each term on the right hand side separately.

Bounding (II). For f_2 , we have by Cauchy-Schwarz that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f_2(x)g(y)h(z)] \right|^2 = |\langle \overline{gh}, Tf_2 \rangle_{\mu_{y,z}}|^2 \leq \|Tf_2\|_{2; \mu_{y,z}}^2 = \langle Tf_2, Tf_2 \rangle_{\mu_{y,z}} = \langle Sf_2, f_2 \rangle_{\mu_x}.$$

To upper bound the last expression, we note that

$$\langle Sf_2, f_2 \rangle_{\mu_x} = \sum_{\substack{\vec{i}, \vec{j} \\ \text{non-merge degree at least } d}} \left| \widehat{f}(\vec{i}, \vec{j}) \right|^2 \prod_{k=1}^n \lambda_{i_k, j_k} \leq (1 - \Omega_{\alpha,m}(1))^d \leq \delta,$$

for appropriately chosen W . Hence, we get that (II) $\leq \sqrt{\delta}$.

Bounding (I). Here, we are going to use random restrictions so that almost all of the mass of f_1 will collapse to monomials of non-merged degree 0, at which point we could truncate off the part of non-merged degree exceeding 0 and get a function that does not distinguish between the distribution μ and the distribution μ' . More precisely, let $s = \delta^{1/3}/d$ and choose $J \subseteq [n]$ by including each element in it with probability s , sample $(\tilde{x}, \tilde{y}, \tilde{z}) \sim \mu^J$ and define

$$\tilde{f}_1 = (f_1)_{\tilde{J} \rightarrow \tilde{x}}, \quad \tilde{g} = (g)_{\tilde{J} \rightarrow \tilde{y}}, \quad \tilde{h} = (h)_{\tilde{J} \rightarrow \tilde{z}}.$$

Let $\eta > 0$ be from Theorem 4.13 for μ' , and define the events:

1. E_1 : $\|\tilde{f}_1\|_2 \geq \delta^{-\eta/10}$.
2. E_2 : $\text{NEStab}_{1-K\delta/s}(\tilde{g}; \mu_y) \leq \delta^{2/3}$, where $K = K(m, \alpha) > 0$.
3. E_3 : The mass of \tilde{f}_1 on monomials of non-merged degree more than 0 exceeds $\delta^{1/6}$.

Let $E = \overline{E_1} \cap E_2 \cap \overline{E_3}$; we show that $\Pr[E] \geq 1 - 3\delta^{\eta'}$ for $\eta' = \min(\eta/5, 1/6)$. We do so using the union bound. For E_1 , we have that the expected value of $\|\tilde{f}_1\|_2^2$ is $\|f_1\|_2^2 \leq 1$, hence by Markov's inequality $\Pr[E_1] \leq \delta^{\eta/5}$. For E_2 , we have by Lemma 2.8 that $\mathbb{E}[\text{NEStab}_{1-K\delta/s}(\tilde{g}; \mu_y)] \leq \text{NEStab}_{1-\delta}(g; \mu_y) \leq \delta$ for suitably chosen K , hence by Markov's inequality $\Pr[E_2] \geq 1 - \delta^{1/3}$. For E_3 , we have

$$\begin{aligned} \mathbb{E}_{J, \tilde{x}} \left[\sum_{\vec{i} \in [\ell]^J, \vec{j}} \widehat{f}_1(\vec{i}, \vec{j})^2 1_{\chi_{\vec{i}, \vec{j}} \text{ has non-merge degree at least 1}} \right] &= \mathbb{E}_J \left[\sum_{\vec{i} \in [\ell]^n, \vec{j}} \widehat{f}_1(\vec{i}, \vec{j})^2 1_{\{k \in [n] \mid j_k > 1\} \cap J \neq \emptyset} \right] \\ &= \sum_{\vec{i} \in [\ell]^n, \vec{j}} \widehat{f}_1(\vec{i}, \vec{j})^2 \mathbb{E}_J [1_{\{k \in [n] \mid j_k > 1\} \cap J \neq \emptyset}] \\ &\leq \sum_{\vec{i} \in [\ell]^n, \vec{j}} \widehat{f}_1(\vec{i}, \vec{j})^2 ds \\ &\leq \delta^{1/3}, \end{aligned}$$

and so by Markov's inequality $\Pr[E_3] \leq \delta^{1/6}$. We write

$$(I) = \underbrace{\mathbb{E}_{J, \tilde{x}, \tilde{y}, \tilde{z}} [1_{\overline{E}} \phi_\mu(\tilde{f}_1, \tilde{g}, \tilde{h})]}_{(III)} + \underbrace{\mathbb{E}_{J, \tilde{x}, \tilde{y}, \tilde{z}} [1_E \phi_\mu(\tilde{f}_1, \tilde{g}, \tilde{h})]}_{(IV)}$$

where $\phi_\mu(\tilde{f}_1, \tilde{g}, \tilde{h}) = \mathbb{E}_{(x, y, z) \sim \mu^J} [\tilde{f}_1(x) \tilde{g}(y) \tilde{h}(z)]$.

Bounding (III). For (III) we have that

$$|(III)| \leq \sqrt{\Pr[E]} \sqrt{\mathbb{E}_{J, \tilde{x}, \tilde{y}, \tilde{z}} [\phi_\mu(\tilde{f}_1, \tilde{g}, \tilde{h})^2]} \leq \sqrt{\Pr[E]} \sqrt{\mathbb{E}_{J, \tilde{x}, \tilde{y}, \tilde{z}} [\|\tilde{f}_1\|_2^2]} = \sqrt{\Pr[E]} \|f_1\|_2,$$

which is at most $\sqrt{3}\delta^{\eta'/2}$.

Bounding (IV). For (IV), the point is that after random restriction the function \tilde{f}_1 barely notices the difference between the distributions μ and μ' , and so we can try to upper bound $\phi(\tilde{f}_1, \tilde{g}, \tilde{h})$ by appealing to Theorem 4.13 over the distribution μ' . The only two issues is that \tilde{f}_1 still has a slight mass on monomials with non-merge degree greater than 0, and secondly that \tilde{f}_1 is not bounded, and we next address these issues.

Let \tilde{f}_0 be the part of \tilde{f}_1 of non-merge degree 0. Then

$$(IV) = \underbrace{\mathbb{E}_{J, \tilde{x}, \tilde{y}, \tilde{z}} \left[\mathbb{1}_E \phi_\mu(\tilde{f}_1 - \tilde{f}_0, \tilde{g}, \tilde{h}) \right]}_{(V)} + \underbrace{\mathbb{E}_{J, \tilde{x}, \tilde{y}, \tilde{z}} \left[\mathbb{1}_E \phi_\mu(\tilde{f}_0, \tilde{g}, \tilde{h}) \right]}_{(VI)},$$

Bounding (V). For (V) we clearly have that

$$|(V)| \leq \mathbb{E}_{J, \tilde{x}, \tilde{y}, \tilde{z}} \left[\mathbb{1}_E \|\tilde{f}_1 - \tilde{f}_0\|_2 \right] \leq \delta^{1/12}$$

as the event E_3 fails.

Bounding (VI). For (VI), we first note that $\phi_\mu(\tilde{f}_0, \tilde{g}, \tilde{h}) = \phi_{\mu'}(\tilde{f}_0, \tilde{g}, \tilde{h})$, namely we can switch from the distribution μ to its x -merge μ' . Secondly, note that for fixed J, \tilde{x}, \tilde{y} and \tilde{z} , defining

$$f''(x) = \mathbb{E}_{(x', y', z') \sim \mu'^J} \left[\tilde{g}(y') \tilde{h}(z') \mid x' = x \right]$$

it holds that

$$\left| \phi_{\mu'}(\tilde{f}_0, \tilde{g}, \tilde{h}) \right|^2 = \left| \langle \tilde{f}_0, f'' \rangle \right|^2 \leq \|\tilde{f}_0\|_2^2 \|f''\|_2^2 \leq 2\delta^{-\eta/10} \phi_{\mu'}(\overline{f''}, \tilde{g}, \tilde{h}),$$

Thus, if the event E holds we may use Theorem 4.13 on μ' to get that $\phi_{\mu'}(\overline{f''}, \tilde{g}, \tilde{h}) \leq M\delta^\eta$, hence $\left| \phi_{\mu'}(\tilde{f}_0, \tilde{g}, \tilde{h}) \right|^2 \leq 2M\delta^{9\eta/10}$ and so $|(VI)| \leq \sqrt{2M}\delta^{9\eta/20}$.

Therefore, we get that $|(IV)| \leq M'\delta^{\eta'}$ for $M' > 0$ depending only on M and $\eta' > 0$ depending only on η , hence $|(I)|$ is upper bounded by the same type of bound, and using (57) the proof is concluded.

A.2 Merging Symbols: Proof of Lemma 5.6

The proof here is very similar to the proof of Lemma 5.5, but as the roles of x and y in Theorem 4.13 are not symmetric we give it in detail.

Consider $T: L_2(\Gamma; \mu_y) \rightarrow L_2(\Sigma \times \Phi; \mu_{x,z})$ defined as

$$Tg(x, z) = \mathbb{E}_{(x', y', z') \sim \mu} \left[g(y') \mid x' = x, z' = z \right],$$

and let $S = T^*T$. Note that S can also be viewed as a Markov chain over Γ , and for $a \in \Gamma$ we denote by Sa the distribution over Γ of the neighbours of a according to this Markov chain. Let $G = (\Gamma, E)$ be the graph in which the edges are the support of the Markov chain S , and note that x, z imply y if and only if the number of connected components of G is exactly $|\Gamma|$. Thus, we assume that the number of connected components is $\ell < |\Gamma|$, and we write them as $C_1 \cup \dots \cup C_\ell$.

Note that S is a symmetric operator, and we may diagonalize it on each connected components separately. Namely, we may find functions $\chi_{i,j}: \Sigma \rightarrow \mathbb{R}$ for $i = 1, \dots, \ell$ and $j = 0, \dots, |C_i| - 1$ that are an orthogonal basis of $L_2(\Gamma; \mu_y)$ and furthermore:

1. $\chi_{i,j}$ is only supported on C_i .
2. $\chi_{i,0}$ is constant on C_i .
3. $\chi_{i,j}$ are eigenfunctions of S with eigenvalue $\lambda_{i,j}$.

Thus, for all i we have $\lambda_{i,0} = 1$, and for $j > 0$ we have as in Lemma 2.7 that $\lambda_{i,j} \leq 1 - \Omega_{\alpha,m}(1)$, and $\lambda_{i,j} \geq 0$ as S is positive semi-definite.

We consider the orthonormal basis $\chi_{\vec{i},\vec{j}}$ now over $L_2(\Gamma^n, \mu_y^{\otimes n})$ defined by $\chi_{\vec{i},\vec{j}}(y) = \prod_{k=1}^n \chi_{i_k, j_k}(y_k)$. The non-merged degree of a monomial $\chi_{\vec{i},\vec{j}}$ is defined to be the number of k 's such that $j_k > 0$ and denoted by $\text{non-merge-deg}(\chi_{\vec{i},\vec{j}})$. Thus, we can write

$$g(y) = \sum_{\vec{i},\vec{j}} \widehat{g}(\vec{i},\vec{j}) \chi_{\vec{i},\vec{j}}(y), \quad \text{where} \quad \widehat{g}(\vec{i},\vec{j}) = \langle g, \chi_{\vec{i},\vec{j}} \rangle_{\mu_y}.$$

Take $d = W\delta^{-1/2}$ for a parameter $W = W(\alpha, m) > 0$ to be chosen later. We now would like to carry out the argument which splits g into high and low non-merged degrees, however we need to be more careful now so as to preserve boundedness.⁹ Indeed, the argument below is morally the same, except that we apply a softer type of such split.

Consider the Markov chain R on Γ that on $y \in \Gamma$, takes $y' = y$ with probability $1 - 1/d$, and otherwise we take the connected component C_i in which y lies, and then sample $y' \sim \mu_y$ conditioned on $y' \in C_i$. Then $R\chi_{i,0} = \chi_{i,0}$ for all i , and for $j > 0$ we have that $R\chi_{i,j} = (1 - \frac{1}{d})\chi_{i,j}$. Write $g = g_1 + g_2$ where $g_1 = R^{\otimes n}g$ and $g_2 = (I - R^{\otimes n})g$. Then

$$\mathbb{E}_{(x,y,z) \sim \mu} [f(x)g(y)h(z)] = \underbrace{\mathbb{E}_{(x,y,z) \sim \mu} [f(x)g_1(y)h(z)]}_{(I)} + \underbrace{\mathbb{E}_{(x,y,z) \sim \mu} [f(x)g_2(y)h(z)]}_{(II)}, \quad (58)$$

and we upper bound each term on the right hand side separately.

Bounding (II). For g_2 , we have by Cauchy-Schwarz that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f(x)g_2(y)h(z)] \right|^2 = |\langle \overline{f}h, \mathbb{T}g_2 \rangle_{\mu_{x,z}}|^2 \leq \|\mathbb{T}g_2\|_{2;\mu_{x,z}}^2 = \langle \mathbb{T}g_2, \mathbb{T}g_2 \rangle_{\mu_{x,z}} = \langle \mathbb{S}g_2, g_2 \rangle_{\mu_y}.$$

To upper bound the last expression, we note that

$$\begin{aligned} \langle \mathbb{S}g_2, g_2 \rangle_{\mu_y} &= \sum_{\vec{i},\vec{j}} |\widehat{g}(\vec{i},\vec{j})|^2 \left(1 - \left(1 - \frac{1}{d} \right)^{\text{non-merged-deg}(\chi_{\vec{i},\vec{j}})} \right) \prod_{k=1}^n \lambda_{i_k, j_k} \\ &\leq \max_r \left(1 - \left(1 - \frac{1}{d} \right)^r \right) (1 - \Omega_{\alpha,m}(1))^r. \end{aligned}$$

For $r > W \log(1/\delta)$, the second term is at most δ (for appropriately chosen W), and for $r \leq W \log(1/\delta)$ the first term is at most $r/d \leq \delta^{1/4}$ for sufficiently small δ_0 , hence $|(II)| \leq \delta^{1/4}$.

⁹In Section A.1, we didn't really care about boundedness as we were able to gain it back later on by changing the function f . However, in this case we cannot afford to change the function g because we need to preserve its non-embedding stability to be small.

Bounding (I). Here, we are going to use random restrictions so that almost all of the mass of g_1 will collapse to monomials of non-merged degree 0, at which point we could truncate off the part of non-merged degree exceeding 0 and get a function that does not distinguish between the distribution μ and the distribution μ' . More precisely, let $s = \delta^{1/3}/d$ and choose $J \subseteq [n]$ by including each element in it with probability s , sample $(\tilde{x}, \tilde{y}, \tilde{z}) \sim \mu^J$ and define

$$\tilde{f} = (f)_{\overline{J} \rightarrow \tilde{x}}, \quad \tilde{g}_1 = (g_1)_{\overline{J} \rightarrow \tilde{y}}, \quad \tilde{h} = (h)_{\overline{J} \rightarrow \tilde{z}}.$$

Let $M \in \mathbb{N}$ and $\eta > 0$ be from Theorem 4.13 for μ' . We will prove the statement for $\eta' = \eta/20$ and $M' = M + 3$. and define the events:

1. E_1 : $\text{NEStab}_{1-K\delta/s}(\tilde{g}_1; \mu_y) \leq \delta^{2/3}$, where $K = K(m, \alpha) > 0$.
2. E_2 : The mass of \tilde{g}_1 on monomials of non-merged degree more than 0 exceeds $\delta^{1/6}$.

We show that the probability of $E = E_1 \cap \overline{E_2}$ is at least $1 - 2\delta^{1/6} \log(1/\delta)$; to do that, we use the union bound and bound the probability of $\overline{E_1}$ and of E_2 .

The event E_1 . For the event E_1 , let V_t be the space of functions spanned by monomial of non-embedding degree exactly t . We claim that V_t is an invariant space of $\mathbb{R}^{\otimes n}$. Indeed, to see that it suffices to show that if $\chi: \Gamma \rightarrow \mathbb{C}$ is orthogonal to all embedding functions (that is, univariate functions in $\text{Embed}_\gamma(\mu)$), then $\mathbb{R}\chi$ is orthogonal to all embedding functions. Indeed, if χ' is an embedding function then

$$\langle \mathbb{R}\chi, \chi' \rangle = \langle \chi, \mathbb{R}^* \chi' \rangle = \langle \chi, \mathbb{R}\chi' \rangle = \langle \chi, \chi' \rangle = 0,$$

where we used the fact that \mathbb{R} is self adjoint as it is an averaging operator corresponding to a reversible Markov chain. Thus, writing $g = \sum_t g_{=t}$ where $g_{=t}$ is in V_t , we get that

$$\begin{aligned} \text{NEStab}_{1-\delta}(g_1) &= \text{NEStab}_{1-\delta}\left(\sum_t \mathbb{R}^{\otimes n} g_{=t}\right) = \sum_{t,t'} \langle \mathbb{R}^{\otimes n} g_{=t'}, \mathbb{T}_{\text{non-embed}, 1-\delta} \mathbb{R}^{\otimes n} g_{=t} \rangle \\ &= \sum_{t,t'} \langle \mathbb{R}^{\otimes n} g_{=t'}, (1-\delta)^t \mathbb{R}^{\otimes n} g_{=t} \rangle, \end{aligned}$$

where we used the fact that $\mathbb{R}^{\otimes n} g_{=t} \in V_t$ and the fact that V_t is an eigenspace of $\mathbb{T}_{\text{non-embed}, 1-\delta}$ of eigenvalue $(1-\delta)^t$. The inner product is 0 for all $t \neq t'$, hence we get that this is equal to

$$\sum_t (1-\delta)^t \|\mathbb{R}^{\otimes n} g_{=t}\|_2^2 \leq \sum_t (1-\delta)^t \|g_{=t}\|_2^2 = \text{NEStab}_{1-\delta}(g) \leq \delta,$$

so $\text{NEStab}_{1-\delta}(g_1) \leq \delta$. Thus, by Lemma 2.8 it holds that $\mathbb{E} [\text{NEStab}_{1-K\delta/s}(\tilde{g}_1; \mu_y)] \leq \text{NEStab}_{1-\delta}(g_1; \mu_y) \leq \delta$ for suitably chosen K , hence by Markov's inequality $\Pr [E_2] \geq 1 - \delta^{1/3}$.

The event E_2 . The expected weight of \tilde{g}_1 on monomials of non-merged degree more than 0 is at most

$$\begin{aligned} & \mathbb{E}_{J, \vec{y}} \left[\sum_{\vec{i} \in [\ell]^J, \vec{j}} \widehat{g}_1(\vec{i}, \vec{j})^2 1_{\chi_{\vec{i}, \vec{j}} \text{ has non-merge degree at least 1}} \right] \\ &= \mathbb{E}_J \left[\sum_{\vec{i} \in [\ell]^n, \vec{j}} \widehat{g}_1(\vec{i}, \vec{j})^2 1_{\{k \in [n] \mid j_k > 1\} \cap J \neq \emptyset} \right] \\ &= \sum_{\vec{i} \in [\ell]^n, \vec{j}} \widehat{g}(\vec{i}, \vec{j})^2 \left(1 - \frac{1}{d}\right)^{|\{k \in [n] \mid j_k > 1\}|} \mathbb{E}_J [1_{\{k \in [n] \mid j_k > 1\} \cap J \neq \emptyset}]. \end{aligned}$$

Denoting $S = |\{k \in [n] \mid j_k > 1\}|$, we have that the expectation in consideration is equal to $1 - (1 - s)^S$ hence the last sum is at most

$$\max_S \left(1 - \frac{1}{d}\right)^S \left(1 - (1 - s)^S\right).$$

For $S \geq d \log(1/\delta)$, the first term is at most δ hence it is at most δ , and for $S < d \log(1/\delta)$ the second term is at most $sS \leq \log(1/\delta) \delta^{1/3}$. Overall we get that the expected weight of \tilde{g}_1 on monomials of non-merged degree more than 0 is at most $\log(1/\delta) \delta^{1/3}$. Hence by Markov's inequality $\Pr[E_2] \leq \delta^{1/6} \log(1/\delta)$.

We write

$$(I) = \underbrace{\mathbb{E}_{J, \vec{x}, \vec{y}, \vec{z}} [1_{\overline{E}} \phi_\mu(\tilde{f}, \tilde{g}_1, \tilde{h})]}_{(III)} + \underbrace{\mathbb{E}_{J, \vec{x}, \vec{y}, \vec{z}} [1_E \phi_\mu(\tilde{f}, \tilde{g}_1, \tilde{h})]}_{(IV)}$$

where $\phi_\mu(\tilde{f}, \tilde{g}_1, \tilde{h}) = \mathbb{E}_{(x, y, z) \sim \mu^J} [\tilde{f}(x) \tilde{g}_1(y) \tilde{h}(z)]$.

Bounding (III). For (III) we have that $|(III)| \leq \Pr[\overline{E}]$ which is at most $\delta^{\eta'}$.

Bounding (IV). For (IV), let $\tilde{g}_1': \Gamma^J \rightarrow \mathbb{C}$ be defined by

$$\tilde{g}_1'(y) = \mathbb{E}_{y' \sim y} [\tilde{g}_1(y')],$$

where y' is distributed as: for each coordinate i , sample y'_i according to μ_{y_i} conditioned on it being in the connected component of y_i . In words, we average over the connected components. Then the value of \tilde{g}_1' is constant on all connected components, and we also have that $\|\tilde{g}_1 - \tilde{g}_1'\|_2^2$ is the mass of \tilde{g}_1 on monomials of non-merge degree at least 1, hence it is at most $\delta^{1/6}$. It follows that

$$\text{NEStab}_{1-K\delta/s}(\tilde{g}_1'; \mu_y) \leq \text{NEStab}_{1-K\delta/s}(\tilde{g}_1; \mu_y) + O(\|\tilde{g}_1 - \tilde{g}_1'\|_2) \leq O(\delta^{1/12}) \leq \delta^{1/20},$$

hence by Claim 4.8 we have $\text{NEStab}_{1-\delta^{1/20}}(\tilde{g}_1'; \mu_y) \leq \delta^{1/20}$. We get that

$$\begin{aligned} |(IV)| &\leq \mathbb{E}_{J, \vec{x}, \vec{y}, \vec{z}} \left[1_E \left| \phi_\mu(\tilde{f}, \tilde{g}_1, \tilde{h}) \right| \right] \leq \mathbb{E}_{J, \vec{x}, \vec{y}, \vec{z}} \left[1_E \left(\left| \phi_\mu(\tilde{f}, \tilde{g}_1', \tilde{h}) \right| + O(\|\tilde{g}_1 - \tilde{g}_1'\|_2) \right) \right] \\ &= \mathbb{E}_{J, \vec{x}, \vec{y}, \vec{z}} \left[1_E \left(\left| \phi_{\mu'}(\tilde{f}, \tilde{g}_1', \tilde{h}) \right| + O(\|\tilde{g}_1 - \tilde{g}_1'\|_2) \right) \right], \end{aligned}$$

where in the last transition we used the fact that \tilde{g}_1' is constant on connected components hence $\phi_\mu(\tilde{f}, \tilde{g}_1', \tilde{h}) = \phi_{\mu'}(\tilde{f}, \tilde{g}_1', \tilde{h})$. As $\left| \phi_{\mu'}(\tilde{f}, \tilde{g}_1', \tilde{h}) \right| \leq M \delta^{\eta/20}$ and $\|\tilde{g}_1 - \tilde{g}_1'\|_2 \leq \delta^{1/12}$, we get that $|(IV)| \leq (M + 1) \delta^{\eta/20}$.

Finishing the proof. Together, we get that $|(I)| \leq (M + 2)\delta^{\eta'}$; combining with the bound on (II) and plugging into (58) finishes the proof.

B Theorem 7.14 Implies Theorem 6.7

In this section we show that Theorem 7.14 Implies Theorem 6.7. The argument here is analogous to the argument in [5, Section A], with one important difference. Therein, we used the fact that μ has some sub-distribution that has no Abelian embedding, and this was used in the “soft truncation” steps (to argue that terms with mis-match of degrees have negligible contribution). Here, we no longer have this property, and instead we appeal to the maximality property of the same sub-distribution of μ (which we have ensured in our case). This is ultimately the reason we have to introduce the additional complication that comes to accommodate the notion of maximality.

The argument proceeds by first splitting g into soft non-embedding homogenous parts, which are functions that have most of their ℓ_2 mass of characters of roughly the same non-embedding degree and are still bounded. Then, we softly truncate the non-embedding degree of f to be of at most roughly of the same order (while keeping boundedness), and truncate the effective non-embedding to be roughly at least that same order (while keeping boundedness). Finally, once all of this is done we will be able to apply harsh truncations on the resulting replacements of f and g function and consider the homogenous parts, then do the same for h , and then appeal to Theorem 7.14 to upper bound each one of these homogenous terms.

Let f, g, h be as in Theorem 6.7, and let $\rho_j = 1 - 2^{-j}\delta$. We take $\xi = \xi(m, \alpha) > 0$ to be small enough.

B.1 Softly Splitting g into Softly (non-embedding) Homogenous Parts

By the triangle inequality

$$\begin{aligned} \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f(x)g(y)h(z)] \right| &\leq \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f(x)\mathbb{T}_{\text{non-embed},\rho_0}g(y)h(z)] \right| \\ &+ \sum_{j=0}^{\infty} \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f(x)(\mathbb{T}_{\text{non-embed},\rho_{j+1}} - \mathbb{T}_{\text{non-embed},\rho_j})g(y)h(z)] \right|. \end{aligned} \quad (59)$$

For the first term, as f, h are 1-bounded we have that it is at most

$$\mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [|\mathbb{T}_{\text{non-embed},\rho_0}g(y)|] \leq \sqrt{\mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [|\mathbb{T}_{\text{non-embed},\rho_0}g(y)|^2]} = \sqrt{\text{NEStab}_{\rho_0^2}(g)} \leq \sqrt{\text{NEStab}_{\rho_0}(g)},$$

which is at most $\sqrt{\delta}$ by assumption; in the last inequality we used Claim 4.8. For the second term on the right hand side of (59) we show:

Claim B.1. *There are $M, \delta_0 > 0$ and $\eta > 0$ depending only on m and α such that for all $0 < \delta \leq \delta_0$*

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f(x)(\mathbb{T}_{\text{non-embed},\rho_{j+1}} - \mathbb{T}_{\text{non-embed},\rho_j})g(y)h(z)] \right| \leq M(1 - \rho_j)^\eta.$$

Given Claim B.1, the proof is quickly concluded. Summing up over j yields that the second term above is at most

$$M \sum_{j=0}^{\infty} 2^{-j\eta} \delta^\eta \leq M \frac{1}{1-2^{-\eta}} \delta^\eta,$$

hence $|\mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f(x)g(y)h(z)]| \leq M' \delta^{\eta'}$ as desired.

The rest of this section is devoted to the proof of Claim B.1, and we fix some $j \geq 0$.

B.2 Softly Truncating the (non-embedding) Degree of f from Above

For the simplicity of notation we write $\rho = \rho_j$ and $T' = T_{\text{non-embed}, \rho_{j+1}} - T_{\text{non-embed}, \rho_j}$. Thus,

$$\begin{aligned} & \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f(x)T'g(y)h(z)] \right| \leq \\ & \underbrace{\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [(I - T_{\text{non-embed}, 1-\rho^{1+\xi}})f(x)T'g(y)h(z)] \right|}_{(I)} + \underbrace{\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [T_{\text{non-embed}, 1-\rho^{1+\xi}}f(x)T'g(y)h(z)] \right|}_{(II)}. \end{aligned} \quad (60)$$

Claim B.2. *There are $M, \delta_0 > 0$ and $\eta > 0$ depending only on m, α and ξ such that for $0 < \delta \leq \delta_0$ it holds that $(I) \leq M\rho^\eta$.*

Proof. We may upper bound $|(I)|$ by

$$\begin{aligned} & \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [(I - T_{\text{non-embed}, 1-\rho^{1+\xi}})f(x)T_{\text{non-embed}, 1-\rho}g(y)h(z)] \right| \\ & + \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [(I - T_{\text{non-embed}, 1-\rho^{1+\xi}})f(x)T_{\text{non-embed}, 1-\rho/2}g(y)h(z)] \right|. \end{aligned}$$

Each one of these terms is upper bounded in the same way, and we demonstrate it on the first. We re-interpret it as

$$\left| \mathbb{E}_{(x,y,z) \sim \mu'^{\otimes n}} [(I - T_{\text{non-embed}, 1-\rho^{1+\xi}})f(x)g(y)h(z)] \right|,$$

where the distribution μ' is defined by first sampling $(x', y', z') \sim \mu$, then $y \sim T_{\text{non-embed}, 1-\rho}y'$ and outputting (x', y, z') . A somewhat annoying feature of μ' is that there are atoms whose probability is of the order of ρ , and to circumvent it we write $\mu' = \frac{\rho}{2}\mu'' + (1 - \frac{\rho}{2})\mu'''$ where μ'' and μ''' are distributions as follows: (1) Recalling property 6b from Theorem 6.7, we take Σ', Γ' and Φ' from there; the support of μ'' is $\{(x, y, z) \in \text{supp}(\mu) \mid x \in \Sigma', y \in \Gamma', z \in \Phi'\}$, (2) each atom in μ'' has probability at least $\alpha'(\alpha) > 0$. Write $f' = (I - T_{\text{non-embed}, 1-\rho^{1+\xi}})f$.

Pick $J \subseteq_{\rho/2} [n]$, let $(\tilde{x}, \tilde{y}, \tilde{z}) \sim \mu'''^{\bar{J}}$ and define

$$f'' = (f')_{\bar{J} \rightarrow \tilde{x}}, \quad g' = (g)_{\bar{J} \rightarrow \tilde{y}}, \quad h' = (h)_{\bar{J} \rightarrow \tilde{z}}.$$

Then sampling $(x', y', z') \sim \mu''^J$ we have

$$\mathbb{E}_{(x,y,z) \sim \mu'^{\otimes n}} [f'(x)g(y)h(z)] = \mathbb{E}_{J,(\tilde{x},\tilde{y},\tilde{z}) \sim \mu''^{\bar{J}}} \left[\mathbb{E}_{(x',y',z') \sim \mu''^J} [f''(x')g'(y)h'(z)] \right],$$

so

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

Define the event E that $\text{NEStab}_{1-\rho^{\xi/4}}(f'') \leq \rho^{\xi/8}$, and write

$$\begin{aligned} & \mathbb{E}_{J,(\tilde{x},\tilde{y},\tilde{z}) \sim \mu''^{\bar{J}}} \left[\left| \mathbb{E}_{(x',y',z') \sim \mu''^J} [f''(x')g'(y)h'(z)] \right| \right] = \\ & \underbrace{\mathbb{E}_{J,(\tilde{x},\tilde{y},\tilde{z}) \sim \mu''^{\bar{J}}} \left[\mathbb{1}_E \left| \mathbb{E}_{(x',y',z') \sim \mu''^J} [f''(x')g'(y)h'(z)] \right| \right]}_{(III)} + \underbrace{\mathbb{E}_{J,(\tilde{x},\tilde{y},\tilde{z}) \sim \mu''^{\bar{J}}} \left[\mathbb{1}_{\bar{E}} \left| \mathbb{E}_{(x',y',z') \sim \mu''^J} [f''(x')g'(y)h'(z)] \right| \right]}_{(IV)}. \end{aligned}$$

Upper bounding (IV). We will show that the probability of \bar{E} is close to 0. First, we estimate the non-embedding stability of f :

$$\text{NEStab}_{1-\rho^{1+\xi/2}}(f') \leq \|\mathbb{T}_{\text{non-embed},1-\rho^{1+\xi/2}}(I - \mathbb{T}_{\text{non-embed},1-\rho^{1+\xi}})f\|_2,$$

and as the eigenvalues of $\mathbb{T}_{\text{non-embed},1-\rho^{1+\xi/2}}(I - \mathbb{T}_{\text{non-embed},1-\rho^{1+\xi}})$ are $(1 - \rho^{1+\xi/2})^d(1 - (1 - \rho^{1+\xi})^d)$, it follows that the last 2-norm is at most the maximum of that over $d \in \mathbb{N}$. If $d \geq \rho^{-(1+0.75\xi)}$, then $(1 - \rho^{1+\xi/2})^d \leq e^{-\Omega(\rho^{-\xi/4})} \leq \rho^{\xi/4}$ for small enough δ_0 . If $d \leq \rho^{-(1+0.75\xi)}$, then by Bernoulli's inequality $1 - (1 - \rho^{1+\xi})^d \leq \rho^{1+\xi}d \leq \rho^{\xi/4}$. Thus we get that $\text{NEStab}_{1-\rho^{1+\xi/2}}(f') \leq \rho^{\xi/4}$.

Appealing to Claim 4.9 we get that

$$\mathbb{E}_{J,\tilde{x}} \left[\text{NEStab}_{1-\rho^{\xi/4}}(f'') \right] \leq \text{NEStab}_{1-c\rho^{1+\xi/4}}(f')$$

where $c > 0$ is a constant depending only on m and α , and as $c\rho^{1+\xi/4} \geq \rho^{1+\xi/2}$ for sufficiently small δ_0 we get from Claim 4.8 that $\text{NEStab}_{1-c\rho^{1+\xi/4}}(f') \leq \text{NEStab}_{1-\rho^{1+\xi/2}}(f') \leq \rho^{\xi/4}$. Hence, by Markov's inequality $\Pr[\bar{E}] \leq \rho^{\xi/8}$, and by 1-boundedness of the functions

$$|(IV)| \leq \Pr[\bar{E}] \leq \rho^{\xi/8}.$$

Upper bounding (III). Fix J and $(\tilde{x}, \tilde{y}, \tilde{z})$ for which E holds. Here, we are going to appeal to maximality conditions. Recall property 6b in Theorem 6.7, the condition there roughly said that we can embed in a subset of

$$\{(x, y, z) \in \text{supp}(\mu) \mid x \in \Sigma', y \in \Gamma', z \in \Phi'\}$$

a maximal distribution ν , and the point is that in our μ'' we will be able to embed a distribution ν' on $\Sigma' \times \Gamma \times \Phi'$ whose support strictly contains ν (and in which the probability of each atom is at least $\alpha''(\alpha, m) > 0$).

Indeed, take the mapping a and the distribution $\tilde{\nu}$ as in property 6b in Theorem 6.7, which we denote by ν for notational convenience. Let ν' be the distribution of (x'', y'', z'') sampled as follows: first sample

$(x', y', z') \sim \mu''$, sample $x'' \sim a^{-1}(x')$ uniformly and take $y'' = y'$ and $z'' = z'$. Also, define $f'''(x'') = f''(a(x''))$; then

$$(III) = \mathbb{E}_{(x'', y'', z'') \sim \nu'^J} [f'''(x'')g'(y'')h'(z'')]. \quad (61)$$

Thus, the right hand side is an expectation as in Theorem 4.13 over a distribution ν' . It is clear that ν' satisfies all of the conditions of Theorem 4.13. Also, if $(u, v, w) \in \text{supp}(\nu)$, then $(a(u), v, w) \in \text{supp}(\mu'')$, and so by definition of ν' we have that $(u, v, w) \in \text{supp}(\nu')$, hence $\text{supp}(\nu) \subseteq \text{supp}(\nu')$. Next, we argue that this is a strict containment. Indeed, by property 6b(vi) there are distinct $v, v' \in \Gamma$ such that $\gamma(v) = \gamma(v')$ (here γ is the master embedding of y). We take $u, u' \in \Sigma''$ and $w, w' \in \Phi'$ such that $(u, v, w), (u', v', w') \in \text{supp}(\nu)$, and note that as in ν the value of any two coordinates implies the last, it follows that $(u, v', w) \notin \text{supp}(\nu)$. However, by definition of μ'' we have that as $(a(u), v, w)$ is in the support of μ , and as $\gamma(v) = \gamma(v')$ it follows that $(a(u), v', w)$ is in the support of μ , hence $(a(u), v', w)$ is in the support of μ'' and so (u, v', w) is in the support of ν' .

It follows that the expectation on the right hand side of (61) is an expectation with respect to a distribution ν' whose support strictly contains the support of a maximal distribution. Also, by Claim 4.8

$$\text{NEStab}_{1-\rho^{\xi/8}, \nu'}(f'''; \nu'_x) \leq \text{NEStab}_{1-\rho^{\xi/4}, \nu'}(f'''; \nu'_x) = \text{NEStab}_{1-\rho^{\xi/4}, \mu''}(f'''; \mu''_x) \leq \rho^{\xi/8},$$

hence we get that we may apply Theorem 4.13 and get that $|(III)| \leq M\rho^{\eta\xi/8}$ for some M and $\eta > 0$ depending only on m and α . \square

B.3 Softly Truncating the Effective (non-embedding) Degree of f from Below

Looking at (II), we write $f' = T_{\text{non-embed}, 1-\rho^{1+\xi}} f$ and $g' = T'g$. By the triangle inequality we get

$$(II) = \underbrace{\left| \mathbb{E}_{(x, y, z) \sim \mu^{\otimes n}} \left[\mathbb{E}_{\text{non-embed}, 1-\rho^{1-\xi}} f'(x)g'(y)h(z) \right] \right|}_{(V)} + \underbrace{\left| \mathbb{E}_{(x, y, z) \sim \mu^{\otimes n}} \left[(I - \mathbb{E}_{\text{non-embed}, 1-\rho^{1-\xi}}) f'(x)g'(y)h(z) \right] \right|}_{(VI)}. \quad (62)$$

Claim B.3. *There are $M, \delta_0 > 0$ and $\eta > 0$ depending only on m, α and ξ such that for $0 < \delta \leq \delta_0$ it holds that $(V) \leq M\rho^\eta$.*

Proof. We re-interpret the expectation in (V) as

$$\mathbb{E}_{(x', y', z') \sim \mu'^{\otimes n}} [f'(x')g'(y')h(z')],$$

where the distribution μ' is defined as first sampling $(x, y, z) \sim \mu$, then $x' \sim \mathbb{E}_{\text{non-embed}, 1-\rho^{1-\xi}} x$ and outputting (x', y', z') where $y' = y$ and $z' = z$. Again, a somewhat annoying feature of μ' is that there are atoms with probability $\rho^{1-\xi}$, and to circumvent it we again use random restrictions.

More precisely, we write $\mu' = \frac{1}{2}\rho^{1-\xi}\mu'' + (1 - \frac{1}{2}\rho^{1-\xi})\mu'''$ where μ'' and μ''' are distributions as follows: (1) Recalling property 6b from Theorem 6.7, we take Σ' and Φ' from there; the support of μ'' is $\{(x, y, z) \in \text{supp}(\mu) \mid x \in \Sigma', z \in \Phi'\}$, (2) each atom in μ'' has probability at least $\alpha'(\alpha) > 0$.

Pick $J \subseteq_{\rho^{1-\xi/2}} [n]$, let $(\tilde{x}, \tilde{y}, \tilde{z}) \sim \mu^{\prime\prime J}$ and define

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

Then sampling $(x', y', z') \sim \mu^{\prime\prime J}$ we have

$$\mathbb{E}_{(x,y,z) \sim \mu^{\prime \otimes n}} [f'(x)g'(y)h(z)] = \mathbb{E}_{J,(\tilde{x},\tilde{y},\tilde{z}) \sim \mu^{\prime\prime J}} \left[\mathbb{E}_{(x',y',z') \sim \mu^{\prime\prime J}} [f''(x')g''(y)h'(z)] \right],$$

so

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\prime \otimes n}} [f'(x)g'(y)h(z)] \right| \leq \mathbb{E}_{J,(\tilde{x},\tilde{y},\tilde{z}) \sim \mu^{\prime\prime J}} \left[\left| \mathbb{E}_{(x',y',z') \sim \mu^{\prime\prime J}} [f''(x')g''(y)h'(z)] \right| \right]. \quad (63)$$

Our goal is to use the maximality property from 6b in Theorem 6.7 to bound this. Towards this end, first note that

$$\text{NEStab}_{1-\rho^{1-\xi/2}}(g') = \langle g', \text{T}_{\text{non-embed}, 1-\rho^{1-\xi/2}} g' \rangle \leq \|\text{T}_{\text{non-embed}, 1-\rho^{1-\xi/2}} g'\|_2,$$

where the last inequality is Cauchy-Schwarz (and using the fact that g' has 2-norm at most 1). Recalling the definition of g' ,

$$\|\text{T}_{\text{non-embed}, 1-\rho^{1-\xi/2}} g'\|_2 = \|\text{T}_{\text{non-embed}, 1-\rho^{1-\xi/2}} (\text{T}_{\text{non-embed}, 1-\rho/2} - \text{T}_{\text{non-embed}, 1-\rho}) g\|_2,$$

which is at most the largest eigenvalue of $\text{T}_{\text{non-embed}, 1-\rho^{1-\xi/2}} (\text{T}_{\text{non-embed}, 1-\rho/2} - \text{T}_{\text{non-embed}, 1-\rho})$. The eigenvalues of this operator are $(1 - \rho^{1-\xi/2})^r ((1 - \rho/2)^r - (1 - \rho)^r)$ for $r \in \mathbb{N}$, and we argue they are all at most $\rho^{\xi/4}$. Indeed, if $r \geq \rho^{1-\xi/4}$ then the first term is at most that (and the second one is at most 1), and otherwise the first term is at most 1 and the second one is at most $(1 - (1 - r\rho)) \leq r\rho = \rho^{\xi/4}$.

Hence we get that $\text{NEStab}_{1-\rho^{1-\xi/2}}(g') \leq \rho^{\xi/4}$. Let E be the event that $\text{NEStab}_{1-\rho^{\xi/4}}(g'') \leq \rho^{\xi/8}$; by Claim 4.9 we get that

$$\mathbb{E}_{J,\tilde{y}} \left[\text{NEStab}_{1-\rho^{\xi/4}}(g'') \right] \leq \text{NEStab}_{1-c\rho^{1-3\xi/4}}(g') \leq \text{NEStab}_{1-\rho^{1-\xi/2}}(g') \leq \rho^{\xi/4}$$

(where we used Claim 4.8 and also that ρ is sufficiently small as δ_0 is sufficiently small), so by Markov's inequality $\Pr[E] \geq 1 - \rho^{\xi/8}$. Hence, splitting the right hand side of (63) according to whether E holds or not, we get that it is at most

$$\underbrace{\mathbb{E}_{J,(\tilde{x},\tilde{y},\tilde{z}) \sim \mu^{\prime\prime J}} \left[\mathbb{1}_E \left| \mathbb{E}_{(x',y',z') \sim \mu^{\prime\prime J}} [f''(x')g''(y)h'(z)] \right| \right]}_{(VII)} + \underbrace{\mathbb{E}_{J,(\tilde{x},\tilde{y},\tilde{z}) \sim \mu^{\prime\prime J}} \left[\mathbb{1}_{\overline{E}} \left| \mathbb{E}_{(x',y',z') \sim \mu^{\prime\prime J}} [f''(x')g''(y)h'(z)] \right| \right]}_{(VIII)}.$$

We have $|(VIII)| \leq 4\Pr[\overline{E}] \leq 4\rho^{\xi/8}$ (as f'' and g'' are 2-bounded and h' is 1-bounded).

For (VII) we use maximality. Take the mapping a and the distribution ν as in property 6b in Theorem 6.7, and let ν' be the distribution of (x'', y'', z'') sampled as follows: first sample $(x', y', z') \sim \mu''$, sample $x'' \sim a^{-1}(x')$ uniformly and take $y'' = y'$ and $z'' = z'$. Also, define $f'''(x'') = f''(a(x''))$; then

$$(VII) = \mathbb{E}_{(x'',y'',z'') \sim \nu^{\prime\prime J}} [f'''(x'')g'(y'')h'(z'')]. \quad (64)$$

Thus, the right hand side is an expectation as in Theorem 4.13 over a distribution ν' . It is clear that ν' satisfies all of the conditions of Theorem 4.13. Also, if $(u, v, w) \in \text{supp}(\nu)$, then $(a(u), v, w) \in \text{supp}(\mu'')$, and so by definition of ν' we have that $(u, v, w) \in \text{supp}(\nu')$, hence $\text{supp}(\nu) \subseteq \text{supp}(\nu')$. Next, we argue that this is a strict containment.

Take distinct $u, u' \in \Sigma_{\text{modest}}$ and $u_{\text{pre}} \in a^{-1}(u)$, $u'_{\text{pre}} \in a^{-1}(u')$, and take $v, v' \in \Gamma$ and $w, w' \in \Phi'$ such that (u_{pre}, v, w) and $(u'_{\text{pre}}, v', w')$ are in $\text{supp}(\nu)$. Then $(u'_{\text{pre}}, v, w) \notin \text{supp}(\nu)$ as in ν it is the case that the value of two coordinates implies the third. On the other hand, it holds that $(u', v, w) \in \text{supp}(\mu'')$ and so $(u'_{\text{pre}}, v, w) \in \text{supp}(\nu')$.

Thus, for every J and \tilde{x}, \tilde{y} and \tilde{z} such that E hold we have that (VII) is an expectation with respect to a distribution satisfying the conditions of Theorem 4.13 whose support strictly contains the support of a maximal distribution, hence $|(VII)| \leq M' \rho^{\eta\xi/8}$ for $\eta = \eta(\alpha, m) > 0$. \square

B.4 Harsh Truncations

Claim B.4. *There are M and $\xi_0, \delta_0 > 0$ depending only on m and α such that if $0 < \xi \leq \xi_0$ and $0 < \delta \leq \delta_0$, then $|(VI)| \leq M 2^{-\frac{1}{\rho^{1-10\xi}}}$.*

Proof. Set $f'' = (I - \mathbb{E}_{\text{non-embed}, 1-\rho^{1-\xi}})f'$. Let g'_1 be the part of g' of non-embedding degree at most $d_1 = \frac{1}{\rho^{1-\xi}}$, g'_2 be the part of g' of non-embedding degree at least $d_2 = \frac{1}{\rho^{1+\xi}}$ and g'_3 be the part of g' of non-embedding degree between d_1 and d_2 . Then $\|g'_1\|_2 \leq (1 - \rho/2)^{d_1} - (1 - \rho)^{d_1} \leq d_1 \rho \leq \rho^\xi$ and $\|g'_2\|_2 \leq (1 - \rho/2)^{d_2} - (1 - \rho)^{d_2} \leq \rho^\xi$, hence

$$|(VI)| \leq 2\rho^\xi + \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f''(x)g'_3(y)h(z)] \right|.$$

Similarly, let f''_1 be the part of f'' of effective non-embedding degree at most $d_3 = \frac{1}{\rho^{1-2\xi}}$, f''_2 be the part of f'' of non-embedding degree more than $d_4 = \frac{1}{\rho^{1+2\xi}}$ and f''_3 be the part of f'' of effective non-embedding degree at least d_3 and non-embedding degree at most d_4 . We have $\|f''_1\|_2, \|f''_2\|_2 \leq \rho^\xi$ hence

$$\begin{aligned} \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f''(x)g'_3(y)h(z)] \right| &\leq \|f''_1\|_2 \|g'_3\|_2 + \|f''_2\|_2 \|g'_3\|_2 + \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f''_3(x)g'_3(y)h(z)] \right| \\ &\leq 2\rho^\xi + \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f''_3(x)g'_3(y)h(z)] \right|, \end{aligned}$$

hence

$$|(VI)| \leq 4\rho^\xi + \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f''_3(x)g'_3(y)h(z)] \right|.$$

We now partition f''_3, g'_3 and h according to embedding degrees, non-embedding degrees and effective degrees. Denoting by $\vec{D} = (D_a)_{a \in \hat{H}}$ a sequence representing embedding degrees, we write $f''_3 = \sum_{\vec{D}, i, i'} f''_{3, \vec{D}, i, i'}$

where $f''_{3, \vec{D}, i, i'}$ is the part of f'' of a -embedding degree equal to D_a for all $a \in \hat{H}$, effective non-embedding degree exactly i and non-embedding degree exactly i' . We also write $g'_3 = \sum_{\vec{D}', j} g'_{3, J, j}$ where $g'_{3, \vec{D}', j}$ is

the part of g'_3 of a -embedding degree exactly D_a for all a and non-embedding degree exactly j . Finally, we write $h = \sum_{\bar{D}''_k} h_{\bar{D}''_k}$ where $h_{\bar{D}''_k}$ is the part of h of a -embedding degree equal to D''_a for all $a \in \hat{H}$ and non-embedding degree exactly k . Then

$$\mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f''_3(x)g'_3(y)h(z)] = \sum_{\bar{D}, \bar{D}', \bar{D}''} \sum_{i, i'=d_3}^{d_4} \sum_{j=d_1}^{d_2} \sum_{k=0}^n \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f''_{3, \bar{D}, i, i'}(x)g'_{3, \bar{D}', j}(y)h_{\bar{D}''_k}(z)].$$

In the rest of the argument, we upper bound the absolute value of the sum on the right hand side.

The contribution is zero unless embedding degrees nearly match. We first claim that unless $|D_a - D'_a|$ and $|D_a - D''_a|$ are each at most $10d_4$ for all a , this expectation is 0. To do so, we are going to expand each function according to the monomials, multiply this out and consider a term in this resulting expression.

Suppose that there is an $a \in \hat{H}$ such that $|D_a - D'_a| \geq 10d_4$, say $D_a \geq D'_a + 10d_4$, then there must be at least $D'_a - (D'_a + d_2) - d_4 \geq 1$ coordinates j such that the term contains $\chi_j(x_j)\chi'_j(y_j)\chi''_j(z_j)$ where $\chi_j = a$, $\chi'_j = a'$ for some $a' \in \hat{H}$ different from a , and χ''_j is an embedding function (may be constant). By independence it suffices to argue that

$$\mathbb{E}_{(x_j, y_j, z_j) \sim \mu} [\chi_j(x_j)\chi'_j(y_j)\chi''_j(z_j)] = 0,$$

and this is clear as it is equal to

$$\mathbb{E}_{(x_j, y_j, z_j)} [a(\sigma(x_j))a'(\gamma(y_j))\chi''_j(z_j)] = \mathbb{E}_{(x_j, y_j, z_j)} [\overline{a(\gamma(y))}a'(\gamma(y_j))\overline{a(\phi(z))}\chi''_j(z_j)] = \langle a', a \rangle \langle \chi''_j, a \rangle,$$

where in the last equality we used independence of y and z . The case $D'_a \geq D_a + 10d_4$ is nearly identical as well as the case that $|D_a - D''_a| \geq 10d_4$.

Next, note that if $k > i' + j$, then $\mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f''_{3, I, i, i'}(x)g'_{3, J, j}(y)h_{K, k}(z)] = 0$; this is true because looking at $h'(z) = \mathbb{E}_{(x', y', z') \sim \mu^{\otimes n}} [f''_{3, i, i'}(x)g'_{3, j}(y) \mid z' = z]$, we get that h' is a function of non-embedding degree at most $i' + j$, and the expectation is $\langle h, \bar{h}' \rangle$.

Denoting by \mathcal{D} the set of tuples of degree sequences $(\bar{D}, \bar{D}', \bar{D}'')$ such that $|D_a - D'_a|, |D_a - D''_a| \leq 10d^4$, we get that

$$\begin{aligned} & \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f''_3(x)g'_3(y)h(z)] \right| \\ & \leq \sum_{(\bar{D}, \bar{D}', \bar{D}'') \in \mathcal{D}} \sum_{i, i'=d_3}^{d_4} \sum_{j=d_1}^{d_2} \sum_{k=0}^{d_2+d_4} \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f''_{3, \bar{D}, i, i'}(x)g'_{3, \bar{D}', j}(y)h_{\bar{D}''_k}(z)] \right| \\ & \leq \sum_{(\bar{D}, \bar{D}', \bar{D}'') \in \mathcal{D}} \sum_{i, i'=d_3}^{d_4} \sum_{j=d_1}^{d_2} \sum_{k=0}^{d_2+d_4} \beta_{n, i, i'}[\mu] \|f''_{3, \bar{D}, i, i'}\|_2 \|g'_{3, \bar{D}', j}\|_2 \|h_{\bar{D}''_k}\|_2, \end{aligned}$$

where by abuse of notation we denote by D, D' and D'' the sum of entries in \bar{D}, \bar{D}' and \bar{D}'' respectively. By Theorem 7.14 we get that provided that ξ is small enough it holds that $\beta_{n, i, i'}[\mu] \leq (1+c)^{-i^{1-2\xi}}$, and

plugging that above yields that

$$\begin{aligned}
& \left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f_3''(x)g_3'(y)h(z)] \right| \\
& \leq \sum_{(\vec{D}, \vec{D}', \vec{D}'') \in \mathcal{D}} \sum_{i, i'=d_3}^{d_4} \sum_{j=d_1}^{d_2} \sum_{k=0}^{d_2+d_4} (1+c)^{-i^{1-2\xi}} \|f_{3, \vec{D}, i, i'}''\|_2 \|g'_{3, \vec{D}', j}\|_2 \|h_{\vec{D}'', k}\|_2 \\
& \leq \sum_{(\vec{D}, \vec{D}', \vec{D}'') \in \mathcal{D}} \sum_{i, i'=d_3}^{d_4} \sum_{j=d_1}^{d_2} \sum_{k=0}^{d_2+d_4} (1+c)^{-i^{1-2\xi}} \|f_{3, \vec{D}}''\|_2 \|g'_{3, \vec{D}'}\|_2 \|h_{\vec{D}''}\|_2,
\end{aligned}$$

where $f_{3, \vec{D}}''$ is the part of f_3'' of embedding degrees \vec{D} , and similarly $g'_{3, \vec{D}'}$ and $h_{\vec{D}''}$; here used Parseval. We get that the above is at most

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f_3''(x)g_3'(y)h(z)] \right| \leq \sum_{(\vec{D}, \vec{D}', \vec{D}'') \in \mathcal{D}} \|f_{3, \vec{D}}''\|_2 \|g'_{3, \vec{D}'}\|_2 \|h_{\vec{D}''}\|_2 \sum_{i, i'=d_3}^{d_4} \sum_{j=d_1}^{d_2} \sum_{k=0}^{d_2+d_4} (1+c)^{-i^{1-2\xi}};$$

the inner sum is at most $M(1+c)^{-d_3^{1-\xi}}$ for M depending only on c , and

$$\begin{aligned}
\sum_{(\vec{D}, \vec{D}', \vec{D}'') \in \mathcal{D}} \|f_{3, \vec{D}}''\|_2 \|g'_{3, \vec{D}'}\|_2 \|h_{\vec{D}''}\|_2 & \leq \sum_{\vec{D}} \|f_{3, \vec{D}}''\|_2 \sum_{\vec{D}', \vec{D}'' : (\vec{D}, \vec{D}', \vec{D}'') \in \mathcal{D}} \|g'_{3, \vec{D}'}\|_2 \\
& \leq \sqrt{\sum_{\vec{D}} \|f_{3, \vec{D}}''\|_2^2} \sqrt{\sum_{\vec{D}} \left(\sum_{\vec{D}', \vec{D}'' : (\vec{D}, \vec{D}', \vec{D}'') \in \mathcal{D}} \|g'_{3, \vec{D}'}\|_2 \right)^2} \\
& \leq (20d_4)^{|H|} \sqrt{\sum_{\vec{D}} \|f_{3, \vec{D}}''\|_2^2} \sqrt{\sum_{(\vec{D}, \vec{D}', \vec{D}'') \in \mathcal{D}} \|g'_{3, \vec{D}'}\|_2^2} \\
& \leq (20d_4)^{2|H|} \sqrt{\sum_{\vec{D}} \|f_{3, \vec{D}}''\|_2^2} \sqrt{\sum_{\vec{D}} \|g'_{3, \vec{D}}\|_2^2} \\
& = (20d_4)^{2|H|} \|f_3''\|_2 \|g_3'\|_2,
\end{aligned}$$

which is at most $(20d_4)^{2|H|}$. We used Cauchy-Schwarz multiple times, the fact that once we fix \vec{D} there are at most $(20d_4)^{2|H|}$ degree vectors \vec{D}', \vec{D}'' such that $(\vec{D}, \vec{D}', \vec{D}'') \in \mathcal{D}$; in the end we also used Parseval. Plugging above, we conclude that

$$\left| \mathbb{E}_{(x,y,z) \sim \mu^{\otimes n}} [f_3''(x)g_3'(y)h(z)] \right| \leq (20d_4)^{2m} \cdot M''(1+c)^{-d_3^{1-\xi}} \leq M'''(1+c')^{-d_3^{1-\xi}}$$

where M''' , ξ and c depend only on m and α , and we are done. \square

B.5 Proof of Claim B.1

We are now ready to prove Claim B.1. Fixing m and α , we take M_1, ξ_1 and δ_1 from Claim B.4. We take $\xi = \min(\xi_1, 1/100)$, then M_2, δ_2 and η_1 from Claim B.2 and M_3, δ_3, η_2 from Claim B.3. Finally we pick $M = M_1 + M_2 + M_3$, $\eta = \min(\eta_1, \eta_2, 1/100)$ and $\delta_0 = \min(\delta_1, \delta_2, \delta_3, 1/100)$.

Using (60), we upper bound the absolute value of left hand side of Claim B.1 by $(I) + (II)$, and by Claim B.3 we have that $(I) \leq M_3\rho^\eta$. Using (62) we bound $(II) \leq (V) + (VI)$ and by Claim B.3 we have $(V) \leq M_2\rho^\eta$. Finally, by Claim B.4 we have that $(VI) \leq M_1\rho^\eta$ and the proof is complete. \square

C SVD Decompositions Proofs

In this section, we prove the claims establishing the SVD decompositions. We remind the reader that throughout, I, J is a partition of $[n]$ where $|I| = n - 1$ and $|J| = 1$.

C.1 Homogeneity and Singular Value Decompositions

The proof of our singular value decomposition will require a basic connection between such decompositions and the various notions of homogeneity. This will be used multiple times, and therefore we abstract below.

Definition C.1. Let $B_1, \dots, B_\ell \subseteq \{g: \Gamma \rightarrow \mathbb{C}\}$ be orthonormal sets and suppose that $B_1 \cup \dots \cup B_\ell$ is an orthonormal basis. Then we may write any $g: \Gamma^n \rightarrow \mathbb{C}$ as

$$g = \sum_{v_{i_1}, \dots, v_{i_n} \in B_1 \cup \dots \cup B_\ell} \alpha_{i_1, \dots, i_n} v_{i_1} \otimes v_{i_2} \otimes \dots \otimes v_{i_n}$$

Then the degree of $v_{i_1} \otimes v_{i_2} \otimes \dots \otimes v_{i_n}$ with respect to B_j is defined to be the number of i_k 's such that $v_{i_k} \in B_j$. Furthermore, we say that g is degree d homogenous with respect to B_j if for each monomial $v_{i_1} \otimes v_{i_2} \otimes \dots \otimes v_{i_n}$ in it with non-zero coefficient it holds that the degree of that monomial with respect to B_j is d .

Claim C.2. Let $B_1, \dots, B_\ell \subseteq \{g: \Gamma^J \rightarrow \mathbb{C}\}$ be orthonormal sets and suppose that $B_1 \cup \dots \cup B_\ell$ is an orthonormal basis. Given a function $g: \Gamma^I \times \Gamma^J \rightarrow \mathbb{C}$ define the matrix $M \in \mathbb{C}^{\Gamma^J \times \Gamma^I}$ as $M(a, b) = g(a, b)$, and consider the matrix $MM^* \in \mathbb{C}^{\Gamma^J \times \Gamma^J}$.

If g is degree d homogenous with respect to B_j , then $\text{Span}(B_j)$ is an invariant space of MM^* .

Proof. For notational convenience, we assume that $J = \{n\}$ and $m = |\Gamma|$. Let $v \in B_j$, and write

$$\begin{aligned} (MM^*v)_a &= \sum_b (MM^*)_{a,b} v(b) = \sum_b \sum_{y_I \in \Gamma^I} M[a, y_I] M^*[y_I, b] v(b) = \sum_b \sum_{y_I \in \Gamma^I} \overline{g(y_I, b)} g(y_I, a) v(b) \\ &= \sum_{y_I \in \Gamma^I} g(y_I, a) \sum_b \overline{g(y_I, b)} v(b). \end{aligned}$$

Expanding g according to the basis $(B_1 \cup \dots \cup B_\ell)^{\otimes n}$ we may write

$$g(y_I, b) = \sum_{v' \in B_j} \tilde{g}_{v'}(y_I) v'(b) + \sum_{v'' \notin B_j} \tilde{g}_{v''}(y_I) v''(b),$$

where $\tilde{g}_{v'}$ is the part of g in which the coordinate in J contributes the vector v' . Plugging this above, we get that

$$\sum_b \overline{g(y_I, b)} v(b) = m \sum_{v' \in B_j} \overline{\tilde{g}_{v'}(y_I)} \langle v', v \rangle + \sum_{v'' \notin B_j} \overline{\tilde{g}_{v''}(y_I)} \langle v'', v \rangle = m \sum_{v' \in B_j} \overline{\tilde{g}_{v'}(y_I)} \langle v', v \rangle,$$

as v and v'' are orthogonal. Plugging this above further, we conclude that

$$(MM^*v)_a = m \sum_{y_I \in \Gamma^I} g(y_I, a) \sum_{v' \in B_j} \overline{\tilde{g}_{v'}(y_I)} \langle v', v \rangle = m^n \sum_{v' \in B_j} \langle g_{J \rightarrow a}, \tilde{g}_{v'} \rangle \langle v', v \rangle.$$

Expanding g again, we have $g_{J \rightarrow a} = \sum_{\tilde{v}' \in B_j} \tilde{g}_{\tilde{v}'} \tilde{v}'(a) + \sum_{\tilde{v}'' \notin B_j} \tilde{g}_{\tilde{v}''} \tilde{v}''(a)$, and so

$$\langle g_{J \rightarrow a}, \tilde{g}_{v'} \rangle = \sum_{\tilde{v}' \in B_j} \langle \tilde{g}_{\tilde{v}'}, \tilde{g}_{v'} \rangle \tilde{v}'(a) + \sum_{\tilde{v}'' \notin B_j} \langle \tilde{g}_{\tilde{v}''}, \tilde{g}_{v'} \rangle \tilde{v}''(a).$$

Note that for $v' \in B_j$, the function $\tilde{g}_{v'}$ is degree $d - 1$ homogenous with respect to B_j , whereas for $\tilde{v}'' \notin B_j$ the function $\tilde{g}_{\tilde{v}''}$ is degree d homogenous with respect to B_j , and so $\langle \tilde{g}_{\tilde{v}''}, \tilde{g}_{v'} \rangle = 0$. Thus, $\langle g_{J \rightarrow a}, \tilde{g}_{v'} \rangle = \sum_{\tilde{v}' \in B_j} \langle \tilde{g}_{\tilde{v}'}, \tilde{g}_{v'} \rangle \tilde{v}'(a)$, and plugging that above (and re-arranging) we get that

$$MM^*v = \sum_{\tilde{v}' \in B_j} \left(m^n \sum_{v' \in B_j} \langle \tilde{g}_{\tilde{v}'}, \tilde{g}_{v'} \rangle \langle v', v \rangle \right) \tilde{v}' \in \text{Span}(B_j) \quad \square$$

C.2 Proof of Claim 8.4

Proof. Recall that marginal distribution of μ on y is uniform, and we denote $m = |\Gamma|$. We think of g as a matrix M in $\mathbb{C}^{\Gamma^J \times \Gamma^I}$, whose (a, b) entry is $g(y_J = a, y_I = b)$. The decomposition stated by the claim is an appropriately chosen singular-value decomposition of M ; below are the details for completeness

Looking at $MM^* \in \mathbb{C}^{\Gamma^J \times \Gamma^J}$, we see that it is an $m \times m$ Hermitian matrix, hence we may find an eigenbasis g'_1, \dots, g'_m of \mathbb{R}^{Γ^J} with real non-negative eigenvalues $\lambda_1, \dots, \lambda_m \geq 0$. We will use Claim C.2 to find a more structured eigenbasis.

By Claim C.2 we conclude that each one of $\text{span}(\{\chi \circ \gamma\})$ for $\chi \in \hat{H}$, as well as $\text{span}(B_{\text{non-embed}})$, are invariant spaces of MM^* . Therefore we may choose an orthonormal eigenbasis g'_1, \dots, g'_m in which each g'_i is from one of these spaces, and we choose this. We also let κ_r be the eigenvalue corresponding to g'_r . To choose g_1, \dots, g_m , define $\tilde{g}_r = M^*g'_r$; first we note that \tilde{g}_r are orthogonal:

$$\langle \tilde{g}_r, \tilde{g}_{r'} \rangle = \langle M^*g'_r, M^*g'_{r'} \rangle = \langle MM^*g'_r, g'_{r'} \rangle = \kappa_r \langle g'_r, g'_{r'} \rangle = \kappa_r \mathbf{1}_{r \neq r'}$$

This means that if we look only at the set R of r 's such that $\kappa_r \neq 0$, then we get that $\{\tilde{g}_r\}_{r \in R}$ is orthogonal, and we choose $g_r = \tilde{g}_r / \sqrt{\kappa_r}$ (which has 2-norm equal to 1).

We prove that

$$M = \sum_{r \in R} \sqrt{\kappa_r} g_r(y_I) \overline{g'_r(y_J)}.$$

Define $M' = \sum_{r \in R} \sqrt{\kappa_r} g_r(y_I) \overline{g'_r(y_J)}$, and note that for all $r \in R$ we have $Mg'_r = M'g'_r = \sqrt{\kappa_r} g_r \neq 0$. As for $r \notin R$, g'_r is in the kernel of M and also in the kernel of M' , and so $Mg'_r = 0 = M'g'_r$. This implies that $M = M'$.

Next, we observe that

$$\sum_r \sqrt{\kappa_r}^2 = \sum_r \kappa_r = \frac{\text{Tr}(MM^*)}{m^n} = \frac{1}{m^n} \sum_b (M^t M)_{b,b} = \|g\|_2^2 = 1.$$

Now we argue about the homogeneity properties of the g_r . All of these arguments are basically the same. For $\chi \in \hat{H}$, if r is such that g'_r is in $\text{span}(\chi \circ \gamma)$, then by definition again

$$g_r(y_I) = \frac{m}{\sqrt{\kappa_r y_J}} \mathbb{E} \left[\overline{g(y)} g'_r(y_J) \right],$$

and expanding g we note that on the right hand side only monomials in g in which the variable from J gives the embedding function $\chi \circ \gamma$ give non-zero contribution, and we see that g_r is completely embedding homogenous and non-embedding homogenous.

If r is such that g'_r is in $\text{span}(B_{\text{non-embed}})$, then by definition again

$$g_r(y_I) = \frac{m}{\sqrt{\kappa_r y_J}} \mathbb{E} \left[\overline{g(y)} g'_r(y_J) \right],$$

and expanding g we note that on the right hand side only monomials in g on which the variable from J give a non-embedding function can contribute (the rest give 0), hence g_r is completely embedding homogenous and non-embedding homogenous. \square

C.3 Proof of Claim 8.5

We run the same argument as in the proof of Claim 8.4, and we sketch it below. Let $P = \Gamma \times \Phi$. Recall that marginal distribution of μ on y, z is uniform, and we denote $m = |P|$. We think of F as a matrix M in $\mathbb{C}^{P^J \times P^I}$, whose (a, b) entry is $F((y_J, z_J) = a, (y_I, z_I) = b)$.

Looking at $MM^* \in \mathbb{C}^{P^J \times P^J}$, we see that it is an $m \times m$ Hermitian matrix, and by Claim C.2 we conclude that each one of $\text{span}(\{W \circ \chi \circ \sigma\})$ for $\chi \in \hat{H}$, as well as $\text{span}(B_{\text{non-embed}})$, are invariant spaces of MM^* . Hence we may choose an orthonormal eigenbasis F'_1, \dots, F'_m in which each F'_i is from one of these spaces, and we choose this. All of the items follow exactly in the same way, except that we also argue about the non-embedding degrees.

Indeed, defining F_t analogously to there, we get that

$$F_t(y_I, z_I) = \frac{m}{\sqrt{\psi_t y_J, z_J}} \mathbb{E} \left[\overline{F_t(y_I, z_I)} F'_t(y_J, z_J) \right].$$

Expanding F we note that if $F_t \in \text{span}(\chi \circ \sigma)$ for some $\chi \in \hat{H}$, then only monomials in F in which the variable from J gives the embedding function $\chi \circ \sigma$ give non-zero contribution, so F_t is completely embedding homogenous and non-embedding homogenous of degree d . Otherwise, if F_t is a non-embedding function, then only monomials in F in which the variable from J gives a non-embedding function give non-zero contribution, so F_t is completely embedding homogenous and non-embedding homogenous of degree $d - 1$. Also, the effective non-embedding degree drops by at most 1, hence the effective non-embedding degree of F_t is at least $d' - 1$. \square

D Missing Proofs: the Direct Product Theorem

D.1 Proof of Claim 13.4

Expanding the left-hand side,

$$\begin{aligned}
\frac{\binom{N-n}{qN-t}}{\binom{N}{qN}} &= \frac{(N-n)!(qN)!(N-qN)!}{N!(qN-t)!(N-qN-n+t)!} \\
&= \frac{qN \cdot (qN-1) \cdots (qN-t+1) \cdot (N-qN) \cdot (N-qN-1) \cdots (N-qN-n+t+1)}{N \cdot (N-1) \cdots (N-n+1)} \\
&= q^t (1-q)^{n-t} \frac{N(N-1/q) \cdots (N-(t-1)/q) \cdot N(N-1/(1-q)) \cdots (N-(n-t-1)/(1-q))}{N(N-1) \cdots (N-(n-1))} \\
&= q^t (1-q)^{n-t} \frac{(1-1/qN) \cdots (1-(t-1)/qN) \cdot (1-1/(1-q)N) \cdots (1-(n-t-1)/(1-q)N)}{(1-1/N) \cdots (1-(n-1)/N)} \\
&= q^t (1-q)^{n-t} \cdot \frac{e^{-\Theta(t^2/qN)} \cdot e^{-\Theta((n-t)^2/(1-q)N)}}{e^{-\Theta(n^2/N)}} \\
&= q^t (1-q)^{n-t} (1 \pm o(1)),
\end{aligned}$$

and the claim follows.

D.2 Proof of Claim 13.35

For simplicity, consider the following distribution which is a refinement of the distribution \mathcal{D} . Fix $0 \leq p_1, p_2, p_3, p_4 \leq 1$ such that $\sum_i p_i \leq 1$. For each $i \in [n]$ independently, $i \in A \setminus A'$ with probability p_1 and $i \in A' \setminus A$ with probability p_2 , $i \in (A \cap A') \setminus A''$ with probability p_3 , and $i \in A''$ with probability p_4 . Note that by choosing $p_1 = p_2 = (1-\alpha)q$, $p_4 = q'$ and $p_3 = \alpha q - q'$, we recover the given distribution \mathcal{D} and hence we will fix these values of p_i throughout the claim.

Fix a triple (A, A', A'') . Let p and \tilde{p} be the probability masses given to (A, A', A'') by the distributions \mathcal{D} and $\tilde{\mathcal{D}}$, respectively. Let $a = |A \setminus A'|$, $b = |A' \setminus A|$, $c = |(A \cap A') \setminus A''|$ and $d = |A''|$. We have

$$p = p_1^a \cdot p_2^b \cdot p_3^c \cdot p_4^d \cdot (1 - (p_1 + p_2 + p_3 + p_4))^{n-(a+b+c+d)}$$

We can compute \tilde{p} as follows:

$$\begin{aligned}
\tilde{p} &= \Pr_{A_0, B_0, B_1} [A_0|_{[n]} = A'', B_0|_{[n]} = A \setminus A'', B_1|_{[n]} = A' \setminus A''] \\
&= \Pr_{A_0} [A_0|_{[n]} = A''] \cdot \Pr_{A_0, B_0} [B_0|_{[n]} = A \setminus A'' \mid A_0|_{[n]} = A''] \Pr_{A_0, B_1} [B_1|_{[n]} = A' \setminus A'' \mid A_0|_{[n]} = A''] \\
&= \frac{\binom{N-n}{q'N-d}}{\binom{N}{q'N}} \cdot \frac{\binom{N-q'N-(n-d)}{(q-q')N-(a+c)}}{\binom{N-q'N}{(q-q')N}} \cdot \frac{\binom{N-q'N-(n-d)}{(q-q')N-(b+c)}}{\binom{N-q'N}{(q-q')N}} \\
&= (1 \pm 2^{-\Omega(n)}) \cdot q'^d (1-q')^{n-d} \cdot \left(\frac{q-q'}{1-q'}\right)^{a+c} \left(\frac{1-q}{1-q'}\right)^{n-(a+c+d)} \cdot \left(\frac{q-q'}{1-q'}\right)^{b+c} \left(\frac{1-q}{1-q'}\right)^{n-(b+c+d)},
\end{aligned}$$

where the last equality follows from Claim 13.4. It can be shown that $p = \tilde{p}$ with the following setting of p_i s

$$p_1 = p_2 = \frac{(q-q')(1-q)}{(1-q')}, p_3 = \frac{(q-q')^2}{(1-q')}, \text{ and } p_4 = q'.$$

Therefore, when $\alpha q = q' + \frac{(q-q')^2}{(1-q')}$, the above identities hold along with $p_1 = p_2 = (1 - \alpha)q$, $p_4 = q'$ and $p_3 = \alpha q - q'$. This finishes the proof of this claim. \square