

Dimension Reduction for Polynomials over Gaussian Space and Applications

Badih Ghazi *

Pritish Kamath †

Prasad Raghavendra ‡

January 24, 2018

Abstract

We introduce a new technique for reducing the dimension of the ambient space of low-degree polynomials in the Gaussian space while preserving their relative correlation structure. As an application, we obtain an explicit upper bound on the dimension of an ε -optimal noise-stable Gaussian partition. In fact, we address the more general problem of upper bounding the number of samples needed to ε -approximate any joint distribution that can be *non-interactively simulated* from a correlated Gaussian source. Our results significantly improve (from Ackermann-like to “merely” exponential) the upper bounds recently proved on the above problems by De, Mossel & Neeman (CCC 2017, SODA 2018 resp.), and imply decidability of the larger alphabet case of the *gap non-interactive simulation* problem posed by Ghazi, Kamath & Sudan (FOCS 2016).

Our technique of dimension reduction for low-degree polynomials is simple and can be seen as a generalization of the Johnson-Lindenstrauss lemma and could be of independent interest.

*MIT, Supported in parts by NSF CCF-1650733 and CCF-1420692. Email: badih@mit.edu

†MIT, Supported in parts by NSF CCF-1420956, CCF-1420692, CCF-1218547 and CCF-1650733. Email: pritish@mit.edu

‡UC Berkeley. Research supported by Okawa Research Grant and NSF CCF-1408643. Email: raghavendra@berkeley.edu

Contents

1	Introduction	1
1.1	Gaussian Isoperimetry & Noise Stability	1
1.2	Non-Interactive Simulation from Correlated Gaussian Sources	3
1.3	Extension: Non-Interactive Simulation from General Discrete Sources	4
1.4	Dimension Reduction for Polynomials over Gaussian Space	5
1.5	Comparisons with recent works of De, Mossel & Neeman	7
1.6	Other Related Work and Future Directions	8
1.7	Organization of the Paper	9
2	Preliminaries	9
2.1	Gaussian Probability Spaces	9
2.2	Hermite Analysis	10
2.3	Vector-valued functions	10
3	Dimension Reduction for Low-Degree Multilinear Polynomials	11
3.1	Proof Sketch of Lemma 3.3	12
4	Transformation to Low-Degree Multilinear Polynomials	13
5	Non-Interactive Simulation from Correlated Gaussian Sources	14
6	Acknowledgments	15
A	Proofs of Mean and Variance Bounds in Dimension Reduction	19
A.1	Product of Inner Products of Normalized Correlated Gaussian Vectors	20
A.2	Mean & Variance Bounds for Multilinear Monomials	25
A.3	Mean & Variance Bounds for Multilinear Polynomials	28
B	Proof of Low-Degree Multilinear Transformation Lemma	30
B.1	Transformation to Low-Degree	31
B.2	Transformation to Multi-linear	33

1 Introduction

1.1 Gaussian Isoperimetry & Noise Stability

Isoperimetric problems over the Gaussian space have become central in various areas of theoretical computer science such as hardness of approximation and learning theory. In its simplest and classic form, the central question in isoperimetry is to determine what is the smallest possible surface area for a body of a given volume. Alternately, isoperimetric problems can also be formulated in terms of the notion of *Noise stability*.

Fix a real number $\rho \in [0, 1]$ and let $f : \mathbb{R}^n \rightarrow \{0, 1\}$ denote the indicator function of a subset (say \mathcal{A}_f) of the n -dimensional Gaussian space (\mathbb{R}^n with the Gaussian measure γ_n given by the density function $d\gamma_n/d\mathbf{X} = \exp(-\|\mathbf{X}\|_2^2/2)/(2\pi)^{n/2}$). The noise stability $\text{Stab}_\rho(f)$ is the probability that two ρ -correlated Gaussians \mathbf{X}, \mathbf{Y} both fall inside or outside \mathcal{A}_f . More generally, the Gaussian “noise operator” U_ρ (also known as the Ornstein-Uhlenbeck operator), defined for each $\rho \in [0, 1]$, acts on any $f : \mathbb{R}^n \rightarrow [0, 1]$ as

$$(U_\rho f)(\mathbf{X}) := \mathbb{E}_{\mathbf{Z} \sim \gamma_n} \left[f \left(\rho \mathbf{X} + \sqrt{1 - \rho^2} \cdot \mathbf{Z} \right) \right] .$$

The noise stability is then defined as $\text{Stab}_\rho(f) := \mathbb{E}_{\mathbf{X} \sim \gamma_n} [f(\mathbf{X}) \cdot U_\rho f(\mathbf{X}) + (1 - f(\mathbf{X})) \cdot (1 - U_\rho f(\mathbf{X}))]$.

Reformulated in terms of noise stability, the isoperimetric problem is to determine the largest possible value of $\text{Stab}_\rho(f)$ for a function $f : \mathbb{R}^n \rightarrow [0, 1]$ with a given expectation $\mathbb{E}[f] = \mu$. The seminal isoperimetric theorem of Borell [Bor85] shows that indicator functions of halfspaces are the most noise-stable among all functions $f : \mathbb{R}^n \rightarrow [0, 1]$ with a given expectation over the Gaussian measure. Borell’s theorem (along with the invariance principle [MOO05, Mos10]) has had fundamental applications in theoretical computer science, e.g., in the hardness of approximation for Max-Cut under the Unique Games conjecture [KKMO07] and in voting theory [Mos10].

In this work, we are interested in analogues of Borell’s theorem for partitions of the Gaussian space *into more than two subsets*, or equivalently noise stability of functions f taking values over $[k] := \{1, \dots, k\}$. Towards stating these analogues, let’s state Borell’s theorem formally in a more general notation. Let Δ_k be the probability simplex in \mathbb{R}^k (i.e., convex hull of the basis vectors $\{e_1, \dots, e_k\}$). The Ornstein-Uhlenbeck operator naturally extends to vector valued functions $f : \mathbb{R}^n \rightarrow \mathbb{R}^k$ as $U_\rho f := (U_\rho f_1, \dots, U_\rho f_k)$ where $f = (f_1, \dots, f_k)$. The noise stability of functions $f : \mathbb{R}^n \rightarrow \Delta_k$ is now defined as $\text{Stab}_\rho(f) := \mathbb{E}_{\mathbf{X} \sim \gamma_n} [\langle f(\mathbf{X}), U_\rho f(\mathbf{X}) \rangle]$ where $\langle \cdot, \cdot \rangle$ denotes the inner product over \mathbb{R}^k . We can similarly define the noise stability of a function $f : \mathbb{R}^n \rightarrow [k]$ by embedding $[k]$ in Δ_k , i.e., identifying coordinate $i \in [k]$ with the standard basis vector $e_i \in \Delta_k$. Borell’s theorem can be formally stated in this notation as follows:

Borell’s Theorem [Bor85]. *For any $f : \mathbb{R}^n \rightarrow \Delta_2$, consider the halfspace function $h = (h_1, h_2) : \mathbb{R}^n \rightarrow \Delta_2$ given by $h_1(\mathbf{X}) = \mathbf{1}\{\langle a, \mathbf{X} \rangle \geq b\}$ and $h_2(\mathbf{X}) = 1 - h_1(\mathbf{X})$, for suitable $a \in \mathbb{R}^n, b \in \mathbb{R}$ such that $\mathbb{E}[f] = \mathbb{E}[h]$. Then, $\text{Stab}_\rho(f) \leq \text{Stab}_\rho(h)$.*

While Borell’s theorem deals with the case of $k = 2$, it is natural to consider the question of maximal noise stability for $k > 2$, stated as follows.

Question 1. [Maximum Noise Stability (MNS)] Given a positive integer $k \geq 2$ and $\alpha \in \Delta_k$, what is the maximum noise stability of a function $f : \mathbb{R}^n \rightarrow \Delta_k$ satisfying $\mathbb{E}[f] = \alpha$?

Question 1 remains open even for $k = 3$. In the particular case where $\alpha = (\frac{1}{k}, \dots, \frac{1}{k})$, the *Standard Simplex Conjecture* posits that the maximum noise stability is achieved by a “standard simplex partition” (this is equivalent to the *Plurality is Stablest* conjecture) [KKMO07, IM12]. Even in the special case when $k = 3$ and $\alpha = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, the answer is still tantalizingly open. In fact, a surprising result of [HMN16] shows that when the α_i ’s are not all equal, the standard simplex partition (and any variant thereof) *does not* achieve the maximum noise stability. This indicates that the case $k \geq 3$ is fundamentally different than the case of $k = 2$. On the positive side, if we consider $0 < \rho < \rho_0(k, n)$ (for some $\rho_0(k, n)$ that goes

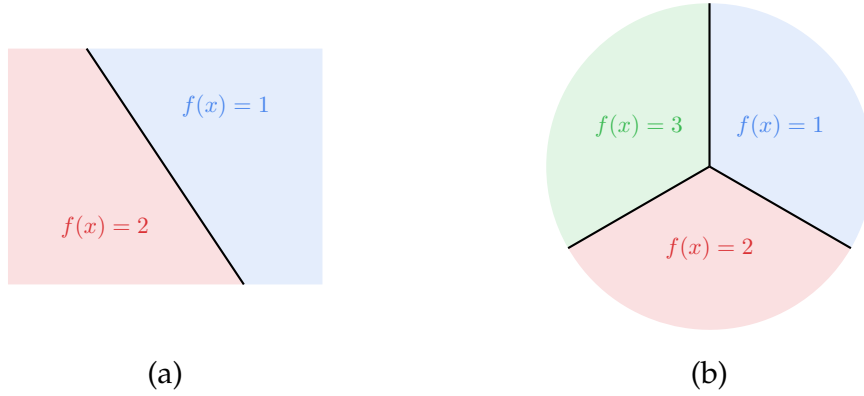


Figure 1: (a) Borell's Theorem: Halfspaces are most noise stable (b) Standard Simplex Partition for $k = 3$ conjectured to be most noise stable (also known as the "Peace Sign Conjecture").

to 0 for large n), then the Standard Simplex Conjecture has been shown to hold [Hei12]. However, this result is not applicable in the setting where ρ is fixed and $n \rightarrow \infty$.

The fact that we do not understand optimal partitions for $k \geq 3$, led De, Mossel & Neeman [DMN17] to ask whether the optimal partition is realized in any finite dimension. More formally:

Question 2. Given $k \geq 2$, $\rho \in (0, 1)$, and $\alpha \in \Delta_k$, let $S_n(\alpha)$ be the optimal noise stability of a function $f : \mathbb{R}^n \rightarrow \Delta_k$ subject to $\mathbb{E}[f] = \alpha$. Is there an n_0 such that $S_n(\alpha) = S_{n_0}(\alpha)$ for all $n \geq n_0$?

Even Question 2 is open as of now! In this light, De, Mossel & Neeman [DMN17] ask whether one can obtain an *explicitly computable* $n_0 = n_0(k, \rho, \varepsilon)$ such that $S_{n_0}(\alpha) \geq S_n(\alpha) - \varepsilon$ for all $n \in \mathbb{N}$ (in other words, there exists a function $f : \mathbb{R}^{n_0} \rightarrow \Delta_k$ that comes ε -close to achieving the optimal noise stability). Note that the challenge is really about n_0 being *explicit*, since some $n_0(k, \rho, \varepsilon)$ always exists, as $S_n(\alpha)$ is a converging sequence as $n \rightarrow \infty$.

A natural approach to proving such an explicit bound is the idea of *dimension reduction*. Basically, it suffices to obtain an $n_0 = n_0(k, \rho, \varepsilon)$ such that for any n and any given function $f : \mathbb{R}^n \rightarrow \Delta_k$, there exists a function $\tilde{f} : \mathbb{R}^{n_0} \rightarrow \Delta_k$ with $\mathbb{E}[\tilde{f}] = \mathbb{E}[f]$ and $\text{Stab}_\rho(\tilde{f}) \geq \text{Stab}_\rho(f) - \varepsilon$. Instantiating f with an optimal (or near-optimal) partition in \mathbb{R}^n , for arbitrarily large n , then gives an ε -optimal partition \tilde{f} in \mathbb{R}^{n_0} .

Indeed, De, Mossel and Neeman follow such an approach and obtain an *explicitly computable* bound on n_0 . To do so, they use and build on the theory of *eigenregular polynomials* that were previously studied in [DS14], which in turn uses other tools such as Malliavin calculus.

In this work, we introduce fundamentally different, but more elementary techniques (elaborated on shortly), thereby significantly improving the bound in [DMN17]. In particular, we show the following.

Theorem 1.1 (Dimension Bound on Approximately Optimal Noise-Stable Function). *Given parameters $k \geq 2$, $\rho \in [0, 1]$ and $\varepsilon > 0$, there exists an explicitly computable $n_0 = n_0(k, \rho, \varepsilon)$ such that the following holds: For any $n \in \mathbb{N}$ and $f : \mathbb{R}^n \rightarrow \Delta_k$, there exists $\tilde{f} : \mathbb{R}^{n_0} \rightarrow \Delta_k$ such that,*

1. $\|\mathbb{E}[f] - \mathbb{E}[\tilde{f}]\|_1 \leq \varepsilon$.
2. $\text{Stab}_\rho(\tilde{f}) \geq \text{Stab}_\rho(f) - \varepsilon$.

In particular, the explicit choice of n_0 can be upper bounded by $\exp\left(\text{poly}\left(k, \frac{1}{1-\rho}, \frac{1}{\varepsilon}\right)\right)$.

Remarks.

- (i) In contrast to our theorem, the bound on n_0 in [DMN17] has an Ackermann-type growth.

(ii) It is a slight technicality that we get $\|\mathbb{E}[f] - \mathbb{E}[\tilde{f}]\|_1 \leq \varepsilon$ instead of $\mathbb{E}[f] = \mathbb{E}[\tilde{f}]$ as was required. However, it is possible to slightly modify \tilde{f} to make $\mathbb{E}[f] = \mathbb{E}[\tilde{f}]$, if we allow n_0 to depend on $\alpha = \mathbb{E}[f]$ (which is the case in Question 2).

Theorem 1.1 has an immediate application to showing that approximately most-stable voting schemes (among all low-influential voting schemes) can be computed efficiently. We refer the reader to [DMN17] for the details of this application.

In order to prove **Theorem 1.1**, we in fact turn to the more general setting of *non-interactive simulation*.

1.2 Non-Interactive Simulation from Correlated Gaussian Sources

Consider a more general setting where instead of a single function f , we have two players, Alice and Bob, with corresponding functions $A : \mathbb{R}^n \rightarrow \Delta_k$ and $B : \mathbb{R}^n \rightarrow \Delta_k$. They apply A and B on the sequence of random variables $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)$ and $\mathbf{Y} = (\mathbf{Y}_1, \dots, \mathbf{Y}_n)$ respectively, where $(\mathbf{X}, \mathbf{Y}) \sim \mathcal{G}_\rho^{\otimes n}$, which is the distribution of ρ -correlated Gaussians in n dimensions, i.e. each coordinate $(\mathbf{X}_i, \mathbf{Y}_i)$ is independently sampled from $\mathcal{G}_\rho := \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}\right)$. The goal is to choose A and B such that $\mathbb{E}[A] = \mathbb{E}[B] = \alpha$, which is a pre-specified vector in Δ_k , while maximizing $\mathbb{E}_{(\mathbf{X}, \mathbf{Y})}[\langle A(\mathbf{X}), B(\mathbf{Y}) \rangle]$. Note that, this quantity is same as $\mathbb{E}_{\mathbf{X} \sim \gamma_n}[\langle A(\mathbf{X}), U_\rho B(\mathbf{X}) \rangle]$, and hence in the restricted setting of $A = B = f$ this quantity is exactly the noise stability of f .

We can interpret Alice as observing \mathbf{X} and outputting the symbol $i \in [k]$ with probability $A_i(\mathbf{X})$, similarly Bob observes \mathbf{Y} and outputs $j \in [k]$ with probability $B_j(\mathbf{Y})$. In this sense, Alice and Bob wish to maximize their “agreement probability”, i.e., their probability of outputting the same symbol. The dimension reduction mentioned in **Theorem 1.1** generalized to this setup would require obtaining an $n_0(k, \rho, \varepsilon)$ and a dimension reduction of A and B that approximately preserves the marginals and does not decrease the agreement probability by more than ε .

However, in this language, it is more natural to ask for a much stronger dimension reduction that preserves the entire joint distribution of symbols that Alice and Bob output, up to ε in total variation distance. We denote the joint distribution of Alice and Bob’s outputs as $(A(\mathbf{X}), B(\mathbf{Y}))_{(\mathbf{X}, \mathbf{Y}) \sim \mathcal{G}_\rho^{\otimes n}}$, which is interpreted as the distribution over $(i, j) \in [k] \times [k]$ given as $\Pr[\text{Alice outputs } i \text{ and Bob outputs } j] = \mathbb{E}_{\mathbf{X}, \mathbf{Y}}[A_i(\mathbf{X})B_j(\mathbf{Y})]$. In the case of $k = 2$, such a dimension reduction follows from Borell’s theorem with in fact $n_0 = 1$! Our main result is indeed such a dimension reduction for all $k \geq 2$.

Theorem 1.2 (NIS from correlated Gaussian source). *Given parameters $k \geq 2$, $\rho \in (0, 1)$ and $\varepsilon > 0$, there exists an explicitly computable $n_0 = n_0(k, \rho, \varepsilon)$ such that the following holds:*

For any n and $A : \mathbb{R}^n \rightarrow \Delta_k$ and $B : \mathbb{R}^n \rightarrow \Delta_k$, there exist $\tilde{A} : \mathbb{R}^{n_0} \rightarrow \Delta_k$ and $\tilde{B} : \mathbb{R}^{n_0} \rightarrow \Delta_k$ such that,

$$d_{\text{TV}}\left(\left(A(\mathbf{X}), B(\mathbf{Y})\right)_{(\mathbf{X}, \mathbf{Y}) \sim \mathcal{G}_\rho^{\otimes n}}, \left(\tilde{A}(\mathbf{a}), \tilde{B}(\mathbf{b})\right)_{(\mathbf{a}, \mathbf{b}) \sim \mathcal{G}_\rho^{\otimes n_0}}\right) \leq \varepsilon.$$

In particular, the explicit choice of n_0 is upper bounded as $\exp\left(\text{poly}\left(k, \frac{1}{1-\rho}, \frac{1}{\varepsilon}\right)\right)$.

The transformation satisfies a stronger property that there exists an “oblivious” randomized transformation (with a shared random seed) to go from A to \tilde{A} and from B to \tilde{B} , which works with probability at least $1 - \varepsilon$. Since the same transformation is applied on A and B with the same random seed, if $A = B$, then $\tilde{A} = \tilde{B}$ as well.

Theorem 1.1 follows immediately from **Theorem 1.2**, by simply setting $A = B = f$ to obtain $\tilde{f} = \tilde{A} = \tilde{B}$. In fact, following up on [DMN17], De, Mossel & Neeman were able to extend their techniques to prove **Theorem 1.2** [DMN18] (again with Ackerman-type bounds on n_0). To do so, they build on the tools developed in [DMN17] along with a new smoothing argument inspired by boosting procedures in learning theory and potential function arguments in complexity theory and additive combinatorics. As we shall present shortly, our approach gets directly to **Theorem 1.2** in a much more elementary way.

1.3 Extension: Non-Interactive Simulation from General Discrete Sources

The *Non-Interactive Simulation of Joint Distributions* is quite well studied in Information Theory and more recently in Theoretical Computer Science. Two players, Alice and Bob, observe the sequences of random variables (x_1, \dots, x_n) and (y_1, \dots, y_n) respectively, where each pair (x_i, y_i) is independently drawn from a known *source* distribution μ . The fundamental question here is to understand which other *target* joint distributions ν can Alice and Bob simulate, without communicating with each other? How many samples from μ are needed to do so, or in other words, what is the *simulation rate*?

The history of this problem goes back to the classical works of Gács and Körner [GK73] and Wyner [Wyn75]. Specifically, consider the distribution Eq over $\{0, 1\} \times \{0, 1\}$ where both marginals are $\text{Ber}(1/2)$ and the bits identical with probability 1. Gács and Körner studied the special case of this problem corresponding to the target distribution $\nu = \text{Eq}$. They characterized the simulation rate in this case, showing that it is equal to what is now known as the *Gács-Körner common information* of μ . On the other hand, Wyner studied the special case corresponding to the source distribution $\mu = \text{Eq}$. He characterized the simulation rate in this case, showing that it is equal to what is now known as *Wyner common information* of ν . Another particularly important work was by Witsenhausen [Wit75] who studied the case where the target distribution $\nu = \mathcal{G}_\rho$. In this case, he showed that the largest correlation ρ that can be simulated is exactly the well-known “maximal correlation coefficient”¹ $\rho(\mu)$ which was first introduced by Hirschfeld [Hir35] and Gebelein [Geb41] and then studied by Rényi [Rén59]. Witsenhausen also considered the case where the target distribution $\nu = \text{DSBS}_\rho$ is a “doubly symmetric binary source”, which is a pair of ρ -correlated bits (i.e., a pair of ± 1 random variables with correlation ρ), and gave an approach to simulate correlated bits by first simulating \mathcal{G}_ρ starting with samples from μ , and then applying half-space functions to get outputs in $\{\pm 1\}$. Starting with μ , such a approach simulates $\text{DSBS}_{\rho'}$ where $\rho' = 1 - \frac{2 \arccos \rho(\mu)}{\pi}$. Indeed, this is very similar to the rounding technique employed in Goemans-Williamson’s approximation algorithm [GW95] for MAXCUT 20 years later!

We will consider the modern formulation of the NIS question as defined in [KA16]. This formulation ignores the simulation rate, and only focuses on whether simulation is even possible or not, given infinitely many samples from μ – that is, whether the simulation rate is non-zero or not.

Definition 1.3 (Non-interactive Simulation of Joint Distributions [KA16]). *Let $(\mathcal{Z} \times \mathcal{Z}, \mu)$ and $([k] \times [k], \nu)$ be two joint probability spaces. We say that the distribution ν can be non-interactively simulated from distribution μ if there exists a sequence of functions $\{A^{(n)} : \mathcal{Z}^n \rightarrow \Delta_k\}_{n \in \mathbb{N}}$ and $\{B^{(n)} : \mathcal{Z}^n \rightarrow \Delta_k\}_{n \in \mathbb{N}}$ such that the joint distribution $\nu_n = (A^{(n)}(\mathbf{x}), B^{(n)}(\mathbf{y}))_{(\mathbf{x}, \mathbf{y}) \sim \mu^{\otimes n}}$ over $[k] \times [k]$ is such that $\lim_{n \rightarrow \infty} d_{\text{TV}}(\nu_n, \nu) = 0$.*

A central question that was left open following the work of Witsenhausen is: given distributions μ and ν , can ν be non-interactively simulated from μ ? Can this be even decided algorithmically? Even when μ and ν are extremely simple, e.g., μ is uniform on the triples $\{(0, 0), (0, 1), (1, 0)\}$ and ν is the doubly symmetric binary source $\text{DSBS}_{0.49}$, it is open if μ can simulate ν ! This problem was formalized as a natural gap-problem in a work by a subset of the authors along with Sudan [GKS16b]. Here we state a slightly more general version.

Problem 1.4 (GAP-NIS($(\mathcal{Z} \times \mathcal{Z}, \mu), \mathcal{P}, k, \varepsilon$), cf. [GKS16b]). *Given a joint probability space $(\mathcal{Z} \times \mathcal{Z}, \mu)$, a family of joint probability spaces \mathcal{P} supported over $[k] \times [k]$, and an error parameter $\varepsilon > 0$, distinguish between the following cases:*

- (i) *there exists n and $A : \mathcal{Z}^n \rightarrow \Delta_k$ and $B : \mathcal{Z}^n \rightarrow \Delta_k$, s.t. the distribution $\nu' = (A(\mathbf{x}), B(\mathbf{y}))_{(\mathbf{x}, \mathbf{y}) \sim \mu^{\otimes n}}$ satisfies $d_{\text{TV}}(\nu', \nu) \leq \varepsilon$ for some $\nu \in \mathcal{P}$.*
- (ii) *for all n and all $A : \mathcal{Z}^n \rightarrow \Delta_k$ and $B : \mathcal{Z}^n \rightarrow \Delta_k$, the distribution $\nu' = (A(\mathbf{x}), B(\mathbf{y}))_{(\mathbf{x}, \mathbf{y}) \sim \mu^{\otimes n}}$ satisfies $d_{\text{TV}}(\nu', \nu) > 2\varepsilon$ for all $\nu \in \mathcal{P}$.²*

¹We skip this definition as it is not central to our paper. The definition can be found in e.g. [GKS16b].

²the choice of constant 2 is arbitrary. Indeed, we could replace it by any constant greater than 1.




Figure 2: Non-Interactive Simulation, e.g., as studied in [KA16]

In prior work [GKS16b], it was shown that GAP-NIS for discrete distributions μ and ν is decidable, in the special case where $k = 2$. This was done by introducing a framework, which reduced the problem to understanding GAP-NIS for the special case where $\mu = \mathcal{G}_\rho$. Indeed, the reason why the case of $k = 2$ was easier was precisely because Borell’s theorem [Bor85] gives an exact characterization of the distributions over $[2] \times [2]$ that can be simulated from \mathcal{G}_ρ . The lack of understanding of the distributions over $[k] \times [k]$ that can be simulated from \mathcal{G}_ρ was suggested in [GKS16b] as a barrier for extending their result to $k > 2$. With Theorem 1.2 in hand, it is possible to extend the framework in [GKS16b] of using a Regularity Lemma and Invariance principle, to yield the following theorem (as also done in [DMN18], but with Ackerman-type bounds).

Theorem 1.5 (NIS from Discrete Sources). *Let $(\mathcal{Z} \times \mathcal{Z}, \mu)$ be a joint probability space. Given parameters $k \geq 2$ and $\varepsilon > 0$, there exists an explicitly computable $n_0 = n_0(k, \mu, \varepsilon)$ such that the following holds: For any n and $A : \mathcal{Z}^n \rightarrow \Delta_k$ and $B : \mathcal{Z}^n \rightarrow \Delta_k$, there exist $\tilde{A} : \mathcal{Z}^{n_0} \rightarrow \Delta_k$ and $\tilde{B} : \mathcal{Z}^{n_0} \rightarrow \Delta_k$ such that,*

$$d_{\text{TV}} \left((A(\mathbf{x}), B(\mathbf{y}))_{(\mathbf{x}, \mathbf{y}) \sim \mu^{\otimes n}}, (\tilde{A}(\mathbf{a}), \tilde{B}(\mathbf{b}))_{\mathbf{a}, \mathbf{b} \sim \mu^{\otimes n_0}} \right) \leq \varepsilon.$$

In particular, the explicit choice of n_0 is upper bounded as $\exp \left(\text{poly} \left(k, \frac{1}{\varepsilon}, \frac{1}{1-\rho}, \log \left(\frac{1}{\alpha} \right) \right) \right)$, where $\alpha = \alpha(\mu)$ is the smallest atom in μ and $\rho = \rho(\mu)$ is the maximal correlation coefficient of μ .

The above theorem immediately suggests a brute force algorithm to decide $\text{GAP-NIS}((\mathcal{Z} \times \mathcal{Z}, \mu), \mathcal{P}, k, \varepsilon)$. We do not provide details of the proof of the above theorem in this extended abstract. The interested reader is referred to the full version of this paper [GKR17] (available online) for details.

1.4 Dimension Reduction for Polynomials over Gaussian Space

We now describe the main technique of “dimension reduction for low-degree polynomials” that we introduce in this work, which could be of independent interest. We highlight that this technique is the main contribution of this paper.

Let’s start with Theorem 1.2, and explain the main ideas behind its proof. We are given two vector-valued functions $A : \mathbb{R}^n \rightarrow \Delta_k$ and $B : \mathbb{R}^n \rightarrow \Delta_k$. We wish to reduce the dimension n of the Gaussian space on which A and B act while preserving the joint distribution $(A(\mathbf{X}), B(\mathbf{Y}))_{(\mathbf{X}, \mathbf{Y}) \sim \mathcal{G}_\rho^{\otimes n}}$ over $[k] \times [k]$. Recall that $\mathbb{E}_{(\mathbf{X}, \mathbf{Y}) \sim \mathcal{G}_\rho^{\otimes n}} [A_i(\mathbf{X}) \cdot B_j(\mathbf{Y})]$ is the probability of the event [Alice outputs i and Bob outputs j]. For succinctness, we write this expectation as $\langle A_i, B_j \rangle_{\mathcal{G}_\rho^{\otimes n}}$. In order to approximately preserve the joint distribution $(A(\mathbf{X}), B(\mathbf{Y}))_{(\mathbf{X}, \mathbf{Y}) \sim \mathcal{G}_\rho^{\otimes n}}$, it suffices to approximately preserve $\langle A_i, B_j \rangle_{\mathcal{G}_\rho^{\otimes n}}$ for each $(i, j) \in$

$[k] \times [k]$ upto an additive ε/k^2 . Thus, to prove [Theorem 1.2](#), we wish to find an explicit constant $n_0 = n_0(\rho, k, \varepsilon)$, along with functions $\tilde{A} : \mathbb{R}^{n_0} \rightarrow \Delta_k$ and $\tilde{B} : \mathbb{R}^{n_0} \rightarrow \Delta_k$ such that

$$\left| \langle \tilde{A}_i, \tilde{B}_j \rangle_{\mathcal{G}_\rho^{\otimes n_0}} - \langle A_i, B_j \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| \leq \frac{\varepsilon}{k^2}.$$

Achieving this directly is highly unclear, since a priori, we have no structural information about A and B ! To get around this, we show that it is possible to first apply a structural transformation on A and B to convert them to low-degree and multilinear polynomials (see [Section 2.2](#) for formal definitions). Such transformations are described in [Section 4](#). This however creates a new problem that the transformed A and B no longer map to Δ_k . Nevertheless, we will show that after the said transformation, we still have that the outputs of A and B are close to Δ_k in expected ℓ_2^2 distance, that is, $\text{dist}(A, \Delta_k) := (\mathbb{E}_{\mathbf{X}} \|\mathcal{R}(A(\mathbf{X})) - A(\mathbf{X})\|_2^2)^{1/2}$ is small (where $\mathcal{R} : \mathbb{R}^k \rightarrow \Delta_k$ denotes the *rounding operator* that maps any $v \in \mathbb{R}^k$ to its closest point in Δ_k). This will ensure that rounding the outputs of A and B to Δ_k will approximately preserve the correlations $\langle A_i, B_j \rangle_{\mathcal{G}_\rho^{\otimes n}}$.

We are now able to revise our objective as follows: Given two (vector-valued) degree- d polynomials $A : \mathbb{R}^n \rightarrow \mathbb{R}^k$ and $B : \mathbb{R}^n \rightarrow \mathbb{R}^k$, does there exist an explicit function $n_0 = n_0(k, d, \delta)$, along with polynomials $\tilde{A} : \mathbb{R}^{n_0} \rightarrow \mathbb{R}^k$ and $\tilde{B} : \mathbb{R}^{n_0} \rightarrow \mathbb{R}^k$ that δ -approximately preserve (i) the correlation $\langle A_i, B_j \rangle_{\mathcal{G}_\rho^{\otimes n}}$ for all $(i, j) \in [k] \times [k]$ and (ii) closeness of the outputs of A and B to Δ_k , that is, $\text{dist}(A, \Delta_k)$ and $\text{dist}(B, \Delta_k)$.

We introduce a very simple and natural dimension-reduction procedure for low-degree multilinear polynomials over Gaussian space. Specifically, for M that is a randomly sampled $n \times n_0$ matrix with i.i.d. standard Gaussian entries, we set

$$\tilde{A}(a) := A \left(\frac{Ma}{\|a\|_2} \right) \quad \text{and} \quad \tilde{B}(b) := B \left(\frac{Mb}{\|b\|_2} \right) \quad \text{for } a, b \in \mathbb{R}^{n_0} \setminus \{0\}. \quad (1)$$

We leave \tilde{A} and \tilde{B} undefined on $0 \in \mathbb{R}^{n_0}$. This is inconsequential as $\{0\}$ is a measure zero set under γ_n . Our main dimension-reduction theorem for polynomials is stated as follows.

Theorem 1.6 (Dimension Reduction Over Gaussian Space). *Given parameters $k \geq 2$, $d \in \mathbb{Z}_{\geq 0}$, $\rho \in (0, 1)$ and $\delta > 0$, there exists an explicitly computable $n_0 = n_0(d, k, \delta)$ such that the following holds:*

Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^k$ and $B : \mathbb{R}^n \rightarrow \mathbb{R}^k$ be degree- d multilinear polynomials. Additionally, suppose that $\text{dist}(A, \Delta_k), \text{dist}(B, \Delta_k) \leq \delta$. Consider the functions $\tilde{A} : \mathbb{R}^{n_0} \rightarrow \mathbb{R}^k$ and $\tilde{B} : \mathbb{R}^{n_0} \rightarrow \mathbb{R}^k$ as defined in [Equation \(1\)](#). With probability at least $1 - O(\delta)$ over the choice of $M \sim \gamma_1^{\otimes (n \times n_0)}$, the following holds:

1. For every $i, j \in [k]$: $\left| \langle A_i, B_j \rangle_{\mathcal{G}_\rho^{\otimes n}} - \langle \tilde{A}_i, \tilde{B}_j \rangle_{\mathcal{G}_\rho^{\otimes n_0}} \right| \leq \delta$.
2. $\text{dist}(\tilde{A}, \Delta_k) \leq \sqrt{\delta}$ and $\text{dist}(\tilde{B}, \Delta_k) \leq \sqrt{\delta}$.

In particular, the explicit choice of n_0 is upper bounded as $\exp(\text{poly}(d, \log k, \log(\frac{1}{\delta})))$.

It is clear from the construction of \tilde{A} and \tilde{B} that this theorem is giving us an “oblivious” randomized transformation, as also remarked in [Theorem 1.2](#). The proof of [Theorem 1.6](#) is obtained by combining [Theorem 3.1](#) and [Proposition 3.2](#) in [Section 3](#).

Proof outline and analogy with the Johnson-Lindenstrauss lemma. We will now highlight a few parallels between our proof of [Theorem 1.6](#) and the proof of the Johnson-Lindenstrauss (JL) lemma (cf. [\[JL84, DG03\]](#)), which has been extremely influential in computer science with numerous applications including compressed sensing, manifold learning, unsupervised learning and graph embedding.

Suppose that we have two unit vectors $u, v \in \mathbb{R}^n$. We wish to obtain a randomized transformation $\Psi_s : \mathbb{R}^n \rightarrow \mathbb{R}^{n_0}$ (for some random seed s) that approximately preserves the inner product, that is,

$\langle \Psi_s(u), \Psi_s(v) \rangle \approx_\delta \langle u, v \rangle$ holds with probability $1 - \delta$, over the randomness of seed s ; note that here $\langle \cdot, \cdot \rangle$ denotes the inner product over \mathbb{R}^n and \mathbb{R}^{n_0} respectively. Indeed, there is such a transformation, namely, $\Psi_M(u) = \frac{M \cdot u}{\sqrt{n_0}}$ where $M \sim \gamma_1^{\otimes n_0 \times n}$. Let $F(M) = \langle \Psi_M(u), \Psi_M(v) \rangle$. Such a transformation satisfies,

$$\mathbb{E}_M[F(M)] = \langle u, v \rangle \quad \text{and} \quad \text{Var}_M(F(M)) = \frac{\langle u, v \rangle^2 + \|u\|_2^2 \|v\|_2^2}{n_0} \leq \frac{2}{n_0},$$

where we use that u and v are unit vectors. Thus, if we choose $n_0 = 2/\delta^3$, then we can make the variance smaller than δ^3 . Thereby, using Chebyshev's inequality, we get that with probability at least $1 - \delta$, the inner product $\langle u, v \rangle$ is preserved, that is, $|\langle \Psi_M(u), \Psi_M(v) \rangle - \langle u, v \rangle| \leq \delta$. Thus, we have a *oblivious* randomized dimension reduction that reduces the dimension of any pair of unit vectors to $O(1/\delta^3)$, independent of n . Note that, instead of using Chebyshev's inequality, we could use a much sharper concentration bound to show that $n_0 = O(1/\varepsilon^2 \log(1/\delta))$ suffices to preserve the inner product up to an additive ε , with probability $1 - \delta$. However, we described the Chebyshev's inequality version as this is similar to our proof of [Theorem 1.6](#).

The problem we are facing, although morally similar, is technically entirely different. For simplicity, let's first consider the task of reducing the dimension of the domain of a *single* pair of polynomials $A : \mathbb{R}^n \rightarrow \mathbb{R}$ and $B : \mathbb{R}^n \rightarrow \mathbb{R}$. And for the moment, consider the transformation such that $\Psi_M A : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$ is given by $A(Ma/\sqrt{n_0})$. Similarly, $\Psi_M B(b) = B(Mb/\sqrt{n_0})$. Our proof of [Theorem 1.6](#) proceeds along similar lines as the above proof of the JL Lemma, that is, by considering $F(M) = \langle \Psi_M A, \Psi_M B \rangle_{\mathcal{G}_\rho^{\otimes n_0}}$, and proving bounds on $\mathbb{E}_M[F(M)]$ and $\text{Var}(F(M))$. This turns out to be quite delicate! Unlike the JL case, we don't even have $\mathbb{E}_M[F(M)] = \langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}}$. What we do show however is that,

$$\left| \mathbb{E}_M[F(M)] - \langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| \leq o_{n_0}(1) \quad \text{and} \quad \text{Var}_M(F(M)) \leq o_{n_0}(1),$$

that is, both are converging to 0 at an explicit rate determined by n_0 (with some dependence on the degree d of A and B). Interestingly however, in the case of $d = 1$, it turns out that $F(M)$ is in fact an unbiased estimator of $\langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}}$. Indeed, this is not a coincidence! We leave it to the interested reader to figure out that in the case of $d = 1$, our transformation is in fact identical to the above described JL transformation on the n -dimensional space of Hermite coefficients of A and B .

Our actual transformation is slightly different, namely $\Psi_M A(a) = A(Ma/\|a\|_2)$. This is to ensure [Item 2](#), about preserving the closeness of the output of A to Δ_k . The proof gets a little more technical due to this change, but is intuitively similar to the above transformation since $\|a\|_2$ is tightly concentrated around $\sqrt{n_0}$. It is important to note that [Item 2](#) is quite critical to the entire approach. If it were not for this restriction, [Item 1](#) is very easy to satisfy on its own by other more direct dimension reduction operations on the Hermite coefficients.

The mean and variance bounds on $F(M)$ are presented as [Lemma 3.3](#). This is the most technical part of this work, but we stress that the main ideas are conceptually simple and elementary (for the most part). We provide a brief sketch of the proof in [Section 3.1](#) that illustrates all the main ideas in under a page, and defer all details to [Appendix A](#). To prove these mean and variance bounds, we first analyze the case when A and B are multi-linear monomials ([Appendix A.2](#)) and then combine these monomial calculations to obtain bounds for general multilinear polynomials ([Appendix A.3](#)).

1.5 Comparisons with recent works of De, Mossel & Neeman

Our main theorems [Theorems 1.1, 1.2 and 1.5](#) significantly improve the bounds in the versions proved by De, Mossel & Neeman [[DMN17](#), [DMN18](#)]. Our work was inspired by [[DMN17](#), [DMN18](#)] through several high-level ideas, such as the use of the transformation to low-degree and multilinear polynomials (although these transformations are technically different in our case). However, it seems that the key

insight into “*why dimension reduction is possible*” provided by the works of De Mossel & Neeman and the current work are fundamentally different.

The key insight for dimension reduction in the work of De, Mossel & Neeman is (quoting [DMN17]): “*the fact that a collection of homogeneous polynomials can be replaced by polynomials in bounded dimensions is a tensor analogue of the fact that for any k vectors in \mathbb{R}^n , there exist k vectors in \mathbb{R}^k with the same matrix of inner products*”. By contrast, the main intuition in our work is an “oblivious” dimension reduction technique, very similar to the Johnson-Lindenstrauss Lemma, as described in Section 1.4.

Also, we point out a minor difference in our versions of Theorem 1.1. In [DMN17] the function \tilde{f} maps to $[k]$, while in our theorem \tilde{f} maps to Δ_k . Interestingly however, this is not a major difference and it follows from a thresholding lemma in [DMN17, Lemma 15 & 16] that any such \tilde{f} can be modified to have range $[k]$, while preserving $\mathbb{E}[\tilde{f}]$ without decreasing the noise stability.

1.6 Other Related Work and Future Directions

Information Theory. Several previous works in information theory and theoretical computer science study “non-interactive simulation” type of questions. For instance, the non-interactive simulation of joint distributions question studied in this work is a generalization of the “non-interactive correlation distillation” problem³ which was studied by [MO04, MOR⁺06]. Moreover, recent works in the information theory community [KA16, BG15] derive analytical tools (based on hypercontractivity and the so-called *strong data processing constant*) to prove impossibility results for NIS. While these results provide stronger bounds for some sources, they are not tight in general. Finally, the “non-interactive agreement distillation” problem studied by [BM11] can also be viewed as a particular case of the NIS setup.

Randomness in Computation. As discussed in [GKS16b], one motivation for studying NIS problems stems from the study of the role of randomness in distributed computing. Specifically, recent works in cryptography [AC93, AC98, BS94, CN00, Mau93, RW05], quantum computing [Nie99, CDS08, DB14] and communication complexity [BGI14, CGMS14, GKS16a, GJ18] study how the ability to solve various computational tasks gets affected by weakening the source of shared randomness. In this context, it is very natural to ask how well can a source of randomness be transformed into another (more structured) one, which is precisely the setup of non-interactive simulation.

The classic Newman’s theorem [New91] tells us that any communication protocol with n -bit inputs and 0-1 outputs can be simulated with only $O(\log n)$ bits of randomness. On the other hand, if we consider the setting where Alice and Bob run a communication protocol with correlated randomness, such as those defined in [BGI14, CGMS14], then we currently don’t know *any* analog of Newman’s theorem, i.e. starting with a protocol using unbounded amount of (correlated) randomness, we don’t know how to reduce the amount of randomness used. Theorem 1.5 implies randomness reduction for zero-communication or even simultaneous message protocols, and hence can be seen as a first step towards understanding the randomness requirements of arbitrary (one or two way) communication protocols with access to correlated randomness.

Tensor Power problems. Another motivation comes from the fact that NIS belongs to the class of *tensor power* problems, which have been very challenging to analyze. In such questions, the goal is to understand how some combinatorial quantity behaves in terms of the dimensionality of the problem as the dimension tends to infinity. A famous instance of such problems is the *Shannon capacity of a graph* [Sha56, Lov79] where the aim is to understand how the independence number of the power of a graph behaves in terms of the exponent. The question of showing the computability of the Shannon capacity remains open to this day [AL06]. Other examples of such open problems (which are more closely related to NIS) arise in the problems of *local state transformation of quantum entanglement* [Bei12, DB13], the problem of computing the *entangled value of a 2-prover 1-round game* (see for, e.g., [KKM⁺11] and also

³which considered the problem of maximizing agreement on a single bit, in various multi-party settings.

the open problems [ope]). Another example is the problem of computing the *amortized value of parallel repetitions of a 2-prover 1-round game* [Raz98, Hol09, Rao11, Raz11, BHH⁺08]. While we don't have computability results for the amortized value, there has been a recent work that tries to characterize it in terms of an information theoretic quantity [BK18]. Yet another example of a tensor-power problem is the task of computing the *amortized communication complexity of a communication problem*. Braverman-Rao [BR11] showed that this equals the information complexity of the communication problem, however the computability of information complexity was shown only recently [BS15].

We hope that the recent progress on the Non-Interactive Simulation problem would stimulate progress on these other notable tensor-power problems. A concrete question is whether the techniques used for NIS (regularity lemma, invariance principle, etc.) can be translated to any of the above mentioned setups.

Deterministic Approximate Counting. We also point out that the notions of eigenregularity used in [DMN17, DMN18] were originally introduced and used in [DS14] to give the only known fixed-polynomial time *deterministic* approximate counting algorithm for polynomial threshold functions (PTFs). Our randomized techniques don't seem directly applicable to the PTF counting problem, as the emphasis there is on being *deterministic*. However, it will be interesting if our techniques could yield some further insights into approximate counting problems.

1.7 Organization of the Paper

In Section 2, we summarize some useful definitions and provide a simple lemma that will be useful later. In Section 3, we state our main technique of dimension reduction for polynomials (Theorem 1.6) and provide a brief sketch of the proof, with most details deferred to Appendix A. In Section 4, we describe the transformations to make functions low-degree and multilinear, with proofs deferred to Appendix B. Finally, in Section 5, we prove Theorem 1.2 (which implies Theorem 1.1 as a corollary).

2 Preliminaries

2.1 Gaussian Probability Spaces

Throughout this paper, we deal with the n -dimensional Gaussian space, i.e. \mathbb{R}^n equipped with Gaussian measure γ_n given by the density function

$$\frac{d\gamma_n}{d\mathbf{X}} := \frac{1}{(2\pi)^{n/2}} \cdot \exp\left(-\frac{1}{2} \cdot \|\mathbf{X}\|_{\mathbb{R}^n}^2\right).$$

where $\|\cdot\|_{\mathbb{R}^n}$ denotes the ℓ_2 norm of a vector. We use letters such as X, Y to denote points in \mathbb{R}^n , bold symbols such as \mathbf{X}, \mathbf{Y} to denote random variables, subscripts such as X_i or \mathbf{X}_i denote the i -th coordinate.

The ℓ_2 -norm of a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is defined as $\|f\|_2 := \left[\mathbb{E}_{\mathbf{X} \sim \gamma_n} f(\mathbf{X})^2 \right]^{1/2}$. We use $L^2(\mathbb{R}^n, \gamma_n)$ to denote the space of all ℓ_2 -integrable functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$, i.e. $\|f\|_2 < \infty$. All functions we consider will be ℓ_2 -integrable. The inner product of $f, g \in L^2(\mathbb{R}^n, \gamma_n)$ is defined as $\langle f, g \rangle_{\gamma_n} := \mathbb{E}_{\mathbf{X} \sim \gamma_n} [f(\mathbf{X})g(\mathbf{X})]$.

The joint distribution of ρ -correlated Gaussians is denoted as \mathcal{G}_ρ , which is a 2-dimensional Gaussian distribution (\mathbf{X}, \mathbf{Y}) , where \mathbf{X} and \mathbf{Y} are marginally distributed according to γ_1 , with $\mathbb{E}[\mathbf{X}\mathbf{Y}] = \rho$. For $A, B \in L^2(\mathbb{R}^n, \gamma_n)$, the *noisy correlation between A and B over $\mathcal{G}_\rho^{\otimes n}$* is defined as,

$$\langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}} := \mathbb{E}_{(\mathbf{X}, \mathbf{Y}) \sim \mathcal{G}_\rho^{\otimes n}} [A(\mathbf{X}) \cdot B(\mathbf{Y})]$$

Finally, the total variation distance between two distributions μ and ν over domain Ω is defined as,

$$d_{\text{TV}}(\mu, \nu) := \sup_{S \subseteq \Omega} |\mu(S) - \nu(S)|.$$

2.2 Hermite Analysis

The set of Hermite polynomials $\{H_r : \mathbb{R} \rightarrow \mathbb{R} : r \in \mathbb{Z}_{\geq 0}\}$ form an orthonormal basis for functions in $L^2(\mathbb{R}, \gamma_1)$ with respect to the inner product $\langle \cdot, \cdot \rangle_{\gamma_1}$. The r -th Hermite polynomial $H_r : \mathbb{R} \rightarrow \mathbb{R}$ (for $r \in \mathbb{Z}_{\geq 0}$) is defined as,

$$H_0(x) = 1; \quad H_1(x) = x; \quad H_r(x) = \frac{(-1)^r}{\sqrt{r!}} e^{x^2/2} \cdot \frac{d^r}{dx^r} e^{-x^2/2}.$$

Hermite polynomials can also be obtained via the generating function, $e^{xt - \frac{t^2}{2}} = \sum_{r=0}^{\infty} \frac{H_r(x)}{\sqrt{r!}} \cdot t^r$.

For any $\sigma = (\sigma_1, \dots, \sigma_n) \in \mathbb{Z}_{\geq 0}^n$, define $H_\sigma : \mathbb{R}^n \rightarrow \mathbb{R}$ as $H_\sigma(\mathbf{X}) = \prod_{i=1}^n H_{\sigma_i}(\mathbf{X}_i)$. It easily follows that the set $\{H_\sigma : \sigma \in \mathbb{Z}_{\geq 0}^n\}$ forms an orthonormal basis for $L^2(\mathbb{R}^n, \gamma_n)$. Thus, every $A \in L^2(\mathbb{R}^n, \gamma_n)$ has a *Hermite expansion* given by $A(\mathbf{X}) = \sum_{\sigma \in \mathbb{Z}_{\geq 0}^n} \widehat{A}(\sigma) \cdot H_\sigma(\mathbf{X})$, where the $\widehat{A}(\sigma)$'s are the *Hermite coefficients* of A obtained as $\widehat{A}(\sigma) = \langle A, H_\sigma \rangle_{\gamma_n}$. The degree of σ is defined as $|\sigma| := \sum_{i \in [n]} \sigma_i$, and the degree of A is the largest $|\sigma|$ for which $\widehat{A}(\sigma) \neq 0$. We say that $A \in L^2(\mathbb{R}^n, \gamma_n)$ is *multilinear* if $\widehat{A}(\sigma)$ is non-zero only if $\sigma_i \in \{0, 1\}$ for all $i \in [n]$.

We list several useful facts about Hermite coefficients:

- (1) Parseval's identity: $\|A\|_2^2 = \sum_{\sigma \in \mathbb{Z}_{\geq 0}^n} \widehat{A}(\sigma)^2$ and $\text{Var}(A) = \sum_{\mathbf{0} \neq \sigma \in \mathbb{Z}_{\geq 0}^n} \widehat{A}(\sigma)^2$.
- (2) Plancherel's identity: $\langle A, A' \rangle_{\gamma_n} = \sum_{\sigma \in \mathbb{Z}_{\geq 0}^n} \widehat{A}(\sigma) \widehat{A}'(\sigma)$.
- (3) Ornstein-Uhlenbeck operator: $U_\rho A(X) = \sum_{\sigma \in \mathbb{Z}_{\geq 0}^n} \rho^{|\sigma|} \cdot \widehat{A}(\sigma) \cdot H_\sigma(X)$.
- (4) Noisy Correlation: $\langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}} = \langle A, U_\rho B \rangle_{\gamma_n} = \sum_{\sigma \in \mathbb{Z}_{\geq 0}^n} \rho^{|\sigma|} \widehat{A}(\sigma) \widehat{B}(\sigma)$

For convenience, $U_\rho(X)$ denotes the distribution $(\rho X + \sqrt{1 - \rho^2} \mathbf{Z})$ where $\mathbf{Z} \sim \gamma_n$, for any $X \in \mathbb{R}^n$.

2.3 Vector-valued functions

For any function $A : \mathbb{R}^n \rightarrow \mathbb{R}^k$, we will interpret A as a vector of functions (A_1, \dots, A_k) , where $A_i : \mathbb{R}^n \rightarrow \mathbb{R}$ is the i -th coordinate of the output of A . The definitions of Hermite analysis extend naturally to vector-valued functions as follows. For $A : \mathbb{R}^n \rightarrow \mathbb{R}^k$, the Hermite coefficient $\widehat{A}(\sigma)$ is $(\widehat{A}_1(\sigma), \dots, \widehat{A}_k(\sigma)) \in \mathbb{R}^k$. We can extend the definition of ℓ_2 -norm as $\|A\|^2 := \mathbb{E}_{\mathbf{X} \sim \gamma_n} \|A(\mathbf{X})\|^2$ or equivalently $\|A\|^2 = \|A_1\|^2 + \dots + \|A_k\|^2 = \sum_{\sigma \in \mathbb{Z}_{\geq 0}^n} \|\widehat{A}(\sigma)\|^2$. Also, $\text{deg}(A) := \max_{i \in [k]} \text{deg}(A_i)$. Again, all the vector-valued functions with domain \mathbb{R}^n that we consider will be such that the function in each coordinate is in $L^2(\mathbb{R}^n, \gamma_n)$.

For $k \in \mathbb{N}$ and $i \in [k]$, let e_i be the unit vector along coordinate i in \mathbb{R}^k . The simplex Δ_k is defined as the convex hull formed by $\{e_i : i \in [k]\}$. Equivalently, $\Delta_k = \{v \in \mathbb{R}^k : \|v\|_1 = 1\}$ is the set of probability distributions over $[k]$. While we will consider vector-valued functions mapping to \mathbb{R}^k , we will be primarily interested in functions which map to Δ_k . The rounding operator $\mathcal{R}^{(k)} : \mathbb{R}^k \rightarrow \Delta_k$ maps any $v \in \mathbb{R}^k$ to its closest point in Δ_k . In particular, it is the identity map on Δ_k . We will drop the superscript on \mathcal{R} , as k is fixed throughout this paper. Similar to our notation for vector-valued functions, \mathcal{R}_i denotes the i -th coordinate of \mathcal{R} . Thus, while the i -th coordinate of A is denoted by A_i , the i -th coordinate of $\mathcal{R}(A)$ is denoted by $\mathcal{R}_i(A)$.

As mentioned already, an important relaxation in our work is to consider functions that do not map to Δ_k , but instead map to \mathbb{R}^k . For such functions to be meaningful, we will require that the outputs are *usually close* to Δ_k , in which case, we will be rounding them to the simplex Δ_k . Towards this end, the following simple proposition will be very useful, which says that if we modify the strategies of Alice and Bob slightly (in ℓ_2 -distance), then the correlation between the strategies does not change significantly. The proof follows by a simple triangle inequality and the Cauchy-Schwarz inequality.

Proposition 2.1 (Close strategies, have similar correlations). *Let $A, \tilde{A}, B, \tilde{B} \in L^2(\mathbb{R}^n, \gamma_n)$ such that $\|A\|_2, \|\tilde{A}\|_2, \|B\|_2, \|\tilde{B}\|_2 \leq 1$. If $\|A - \tilde{A}\|_2 \leq \varepsilon$ and $\|B - \tilde{B}\|_2 \leq \varepsilon$, then it holds that,*

$$\left| \langle \tilde{A}, \tilde{B} \rangle_{\mathcal{G}_\rho^{\otimes n}} - \langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| \leq 2\varepsilon$$

3 Dimension Reduction for Low-Degree Multilinear Polynomials

In this section, we present our main technique of dimension reduction for low-degree multilinear polynomials over Gaussian space. [Theorem 1.6](#) is obtained immediately as a combination of [Theorem 3.1](#) and [Proposition 3.2](#) stated below.

Theorem 3.1. *Given $d \in \mathbb{Z}_{>0}$, $\rho \in [0, 1]$ and $\delta > 0$, there exists an explicitly computable $n_0 = n_0(d, \delta)$, such that the following holds:*

Let $A : \mathbb{R}^n \rightarrow \mathbb{R}$ and $B : \mathbb{R}^n \rightarrow \mathbb{R}$ be degree- d multilinear polynomials, such that $\|A\|_2, \|B\|_2 \leq 1$.

For $M \in \mathbb{R}^{n \times n_0}$ with entries i.i.d. sampled from γ_1 , define the functions⁴ $A_M : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$ and $B_M : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$ as

$$A_M(a) = A\left(\frac{Ma}{\|a\|_2}\right) \quad \text{and} \quad B_M(b) = B\left(\frac{Mb}{\|b\|_2}\right) \quad \text{for } a, b \in \mathbb{R}^{n_0} \setminus \{0\}.$$

Then, with probability at least $1 - \delta$ (over the choice of M), it holds that,

$$\left| \langle A_M, B_M \rangle_{\mathcal{G}_\rho^{\otimes n_0}} - \langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| < \delta.$$

In particular, the explicit choice of n_0 is upper bounded as $\frac{d^{O(d)}}{\delta^4}$.

In other words, for a typical choice of $M \sim \gamma_1^{\otimes(n \times n_0)}$, the correlation between A and B is approximately preserved if we replace $(X, Y) \sim \mathcal{G}_\rho^{\otimes n}$ by $(Ma/\|a\|_2, Mb/\|b\|_2)$, where $(a, b) \sim \mathcal{G}_\rho^{\otimes n_0}$. Intuitively, M can be thought of as a means to “stretch” n_0 coordinates of \mathcal{G}_ρ into effectively n coordinates of \mathcal{G}_ρ , while “fooling” correlations between degree- d multilinear polynomials.

Before we prove the above theorem, we prove a simple proposition that completely handles [Item 2](#) of [Theorem 1.6](#) by showing that if this dimension reduction were applied to vector-valued functions whose outputs lie close to the simplex Δ_k , then with high probability, even the dimension-reduced functions will have outputs close to the simplex. More formally,

Proposition 3.2. *Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^k$ and $B : \mathbb{R}^n \rightarrow \mathbb{R}^k$, such that $\|\mathcal{R}(A) - A\|_2, \|\mathcal{R}(B) - B\|_2 \leq \delta$. Then, with probability at least $1 - 2\delta$ (over choice of M), it holds that,*

$$\|\mathcal{R}(A_M) - A_M\|_2 \leq \sqrt{\delta} \quad \text{and} \quad \|\mathcal{R}(B_M) - B_M\|_2 \leq \sqrt{\delta}.$$

Proof. Observe that even for a fixed non-zero $a \in \mathbb{R}^{n_0}$, the distribution of $\frac{Ma}{\|a\|_2}$ is identical to that of a standard n -variate Gaussian distribution γ_n . Thus, we immediately have that,

$$\begin{aligned} \mathbb{E}_M \mathbb{E}_a \left\| \mathcal{R}\left(A\left(\frac{Ma}{\|a\|_2}\right)\right) - A\left(\frac{Ma}{\|a\|_2}\right) \right\|_2^2 &= \mathbb{E}_X \left\| \mathcal{R}(A(X)) - A(X) \right\|_2^2 \\ \text{Alternately, } \mathbb{E}_M \left\| \mathcal{R}(A_M) - A_M \right\|_2^2 &= \left\| \mathcal{R}(A) - A \right\|_2^2 \leq \delta^2 \end{aligned}$$

Thus, by Markov’s inequality, $\|\mathcal{R}(A_M) - A_M\|_2 \leq \sqrt{\delta}$ holds with probability at least $1 - \delta$. We can similarly argue for B_M , and a union bound completes the proof. \square

⁴we leave A_M and B_M undefined on $0 \in \mathbb{R}^{n_0}$. This is inconsequential as $\{0\}$ is a measure zero set under γ_n .

To prove [Theorem 3.1](#), we primarily use the second moment method (i.e., Chebyshev's inequality). In particular, let $F(\mathbf{M})$ be defined as,

$$F(\mathbf{M}) \stackrel{\text{def}}{=} \langle A_{\mathbf{M}}, B_{\mathbf{M}} \rangle_{\mathcal{G}_\rho^{\otimes n_0}}$$

The most technical part of this work is to show sufficiently good bounds on the mean and variance of $F(\mathbf{M})$ for a random choice of $\mathbf{M} \sim \gamma_1^{\otimes(n \times n_0)}$, given by the following lemma.

Lemma 3.3. (Mean & Variance Bound). *Given d and δ , there exists an explicitly computable $n_0 := n_0(d, \delta)$ such that for $\mathbf{M} \sim \gamma_1^{\otimes(n \times n_0)}$,*

$$\begin{aligned} \left| \mathbb{E}_{\mathbf{M}} F(\mathbf{M}) - \langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| &\leq \delta && \textbf{(Mean bound)} \\ \text{Var}_{\mathbf{M}}(F(\mathbf{M})) &\leq \delta && \textbf{(Variance bound)} \end{aligned}$$

In particular, one may take $n_0 = \frac{d^{O(d)}}{\delta^2}$.

We provide a little sketch of the proof of [Lemma 3.3](#) below, with the full details in [Appendix A](#). Assuming [Lemma 3.3](#), we can easily prove [Theorem 3.1](#).

Proof of Theorem 3.1. We invoke [Lemma 3.3](#) with parameters d and $\delta^2/2$, to get a choice of $n_0 = \frac{d^{O(d)}}{\delta^4}$. Using Chebyshev's inequality and the Variance bound in [Lemma 3.3](#), we have that for any $\eta > 0$,

$$\Pr_{\mathbf{M}} \left[|F(\mathbf{M}) - \mathbb{E}_{\mathbf{M}} F(\mathbf{M})| > \eta \right] \leq \frac{\delta^2}{2\eta}.$$

Using the triangle inequality, and the Mean bound in [Lemma 3.3](#), we get

$$\begin{aligned} \Pr_{\mathbf{M}} \left[\left| F(\mathbf{M}) - \langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| > \delta \right] &\leq \Pr_{\mathbf{M}} \left[|F(\mathbf{M}) - \mathbb{E}_{\mathbf{M}} F(\mathbf{M})| + \left| \mathbb{E}_{\mathbf{M}} F(\mathbf{M}) - \langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| > \delta \right] \\ &\leq \Pr_{\mathbf{M}} \left[|F(\mathbf{M}) - \mathbb{E}_{\mathbf{M}} F(\mathbf{M})| > \delta - \delta^2 \right] \leq \delta. \quad \square \end{aligned}$$

3.1 Proof Sketch of [Lemma 3.3](#)

While the proof of [Lemma 3.3](#) is somewhat technical as a whole, the main driver of the entire lemma is a simple combinatorial fact that if we sample d times with replacement from a bag with n_0 items, then the probability of not sampling distinct items is at most $O(d^2/n_0) = o_{n_0}(1)$. We briefly illustrate this idea at play by proving a *simpler* version of the mean bound. For this section, let's consider a different dimension reduction of setting $A_{\mathbf{M}}$ and $B_{\mathbf{M}}$ as, $A_{\mathbf{M}}(a) = A(\mathbf{M}a/\sqrt{n_0})$ and $B_{\mathbf{M}}(b) = B(\mathbf{M}b/\sqrt{n_0})$, where $\mathbf{M} \sim \gamma_1^{\otimes(n \times n_0)}$. Let $\mathbf{m}_i \in \mathbb{R}^{n_0}$ denote the vector corresponding to the i -th row of \mathbf{M} . Consider the mean of $F(\mathbf{M}) = \langle A_{\mathbf{M}}, B_{\mathbf{M}} \rangle_{\mathcal{G}_\rho^{\otimes n_0}}$:

$$\mathbb{E}_{\mathbf{M}} F(\mathbf{M}) = \mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} A \left(\frac{\mathbf{M}\mathbf{a}}{\sqrt{n_0}} \right) B \left(\frac{\mathbf{M}\mathbf{b}}{\sqrt{n_0}} \right) = \mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} \sum_{\sigma, \kappa} \frac{\hat{A}(\sigma) \hat{B}(\kappa)}{n_0^{(|\sigma|+|\kappa|)/2}} \cdot \prod_{i: \sigma_i=1} \langle \mathbf{m}_i, \mathbf{a} \rangle \cdot \prod_{j: \kappa_j=1} \langle \mathbf{m}_j, \mathbf{b} \rangle$$

where, recall that A and B are multilinear, and so the relevant σ and κ are in $\{0, 1\}^n$, with $|\sigma|, |\kappa| \leq d$. Next, observe that $\mathbb{E} \mathbf{m}_i \mathbf{m}_i^T = I_{n_0 \times n_0}$, and hence we get that,

$$\mathbb{E}_{\mathbf{M}} \prod_{i: \sigma_i=1} \langle \mathbf{m}_i, \mathbf{a} \rangle \cdot \prod_{j: \kappa_j=1} \langle \mathbf{m}_j, \mathbf{b} \rangle = \begin{cases} \langle \mathbf{a}, \mathbf{b} \rangle^{|\sigma|} & \text{if } \sigma = \kappa \\ 0 & \text{if } \sigma \neq \kappa \end{cases}.$$

Finally, we observe that if we expand $\langle \mathbf{a}, \mathbf{b} \rangle^d$ as $\sum_{i_1, \dots, i_d \in [n_0]} \mathbf{a}_{i_1} \mathbf{b}_{i_1} \dots \mathbf{a}_{i_d} \mathbf{b}_{i_d}$, then from the combinatorial fact above, except for a $O(d^2) \cdot n_0^{d-1}$ out of total n_0^d terms, the indices i_1, \dots, i_d are all distinct. It is immediate to see that if all the i_j 's are distinct then $\mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbf{a}_{i_1} \mathbf{b}_{i_1} \dots \mathbf{a}_{i_d} \mathbf{b}_{i_d} = \rho^d$. Additionally, we show that if the i_j 's are not all distinct then $|\mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbf{a}_{i_1} \mathbf{b}_{i_1} \dots \mathbf{a}_{i_d} \mathbf{b}_{i_d}| \leq d^{O(d)}$ (this follows from the fact that for the d -th moments of γ_1 are at most $d^{O(d)}$). Putting this together we get for any σ (with $|\sigma| \leq d$) that,

$$\mathbb{E}_{\mathbf{a}, \mathbf{b}} \frac{\langle \mathbf{a}, \mathbf{b} \rangle^{|\sigma|}}{n_0^{|\sigma|}} = \rho^{|\sigma|} \pm \frac{d^{O(d)}}{n_0}$$

Putting everything together we get,

$$\mathbb{E}_{\mathbf{M}} F(\mathbf{M}) = \sum_{\sigma} \widehat{A}(\sigma) \widehat{B}(\sigma) \cdot \left(\rho^{|\sigma|} \pm \frac{d^{O(d)}}{n_0} \right) = \langle A, B \rangle_{\mathcal{G}_{\rho}^{\otimes n}} \pm \sum_{\sigma} \widehat{A}(\sigma) \widehat{B}(\sigma) \cdot \frac{d^{O(d)}}{n_0}$$

And hence,

$$\left| \mathbb{E}_{\mathbf{M}} F(\mathbf{M}) - \langle A, B \rangle_{\mathcal{G}_{\rho}^{\otimes n}} \right| \leq \frac{d^{O(d)}}{n_0} \cdot \sum_{\sigma} \widehat{A}(\sigma) \widehat{B}(\sigma) \leq \frac{d^{O(d)}}{n_0} \cdot \|A\|_2 \cdot \|B\|_2 \leq \delta,$$

where we use the Cauchy-Schwarz inequality and that $n_0 \geq d^{O(d)}/\delta$. This completes a proof sketch of the mean bound in [Lemma 3.3](#). Replacing $\sqrt{n_0}$ by $\|a\|_2$ introduces a minor technicality, but still works because $\|a\|_2$ is tightly concentrated around $\sqrt{n_0}$. The variance bound is slightly more complicated with the use of a hypercontractive inequality instead of Cauchy-Schwarz. The full details of the proof are in [Appendix A](#).

4 Transformation to Low-Degree Multilinear Polynomials

While [Theorem 1.6](#) applies only for low-degree multilinear polynomials, we can extend it for all functions by using the following lemma that transforms k -dimensional ℓ_2 -integrable functions $A : \mathbb{R}^n \rightarrow \mathbb{R}^k$ and $B : \mathbb{R}^n \rightarrow \mathbb{R}^k$ into low-degree multilinear polynomials while approximately preserving all correlations and also not deviating much from the simplex Δ_k (although slightly increasing the number of variables).

Lemma 4.1 (Low-Degree Multilinear Transformation). *Given parameters $\rho \in [0, 1]$, $\delta > 0$, $k \in \mathbb{N}$, there exists an explicit $d = d(k, \rho, \delta)$ and $t := t(k, d, \delta)$ such that the following holds:*

Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^k$ and $B : \mathbb{R}^n \rightarrow \mathbb{R}^k$, such that, for any $i \in [k]$, it holds that $\text{Var}(A_i), \text{Var}(B_i) \leq 1$.

Then, there exist functions $\widetilde{A} : \mathbb{R}^{nt} \rightarrow \mathbb{R}^k$ and $\widetilde{B} : \mathbb{R}^{nt} \rightarrow \mathbb{R}^k$ such that the following statements hold.

1. \widetilde{A} and \widetilde{B} are multilinear with degree at most d .
2. For any $i \in [k]$, it holds that $\text{Var}(\widetilde{A}_i) \leq \text{Var}(A_i) \leq 1$ and $\text{Var}(\widetilde{B}_i) \leq \text{Var}(B_i) \leq 1$.
3. $\left\| \mathcal{R}(\widetilde{A}) - \widetilde{A} \right\|_2 \leq \left\| \mathcal{R}(A) - A \right\|_2 + \delta$ and $\left\| \mathcal{R}(\widetilde{B}) - \widetilde{B} \right\|_2 \leq \left\| \mathcal{R}(B) - B \right\|_2 + \delta$
4. For every $i, j \in [k]$,

$$\left| \left\langle \widetilde{A}_i, \widetilde{B}_j \right\rangle_{\mathcal{G}_{\rho}^{\otimes nt}} - \left\langle A_i, B_j \right\rangle_{\mathcal{G}_{\rho}^{\otimes n}} \right| \leq \frac{\delta}{\sqrt{k}}$$

In particular, one may take $d = O\left(\frac{\sqrt{k} \log^2(k/\delta)}{\delta(1-\rho)}\right)$ and $t = O\left(\frac{kd^2}{\delta^2}\right)$.

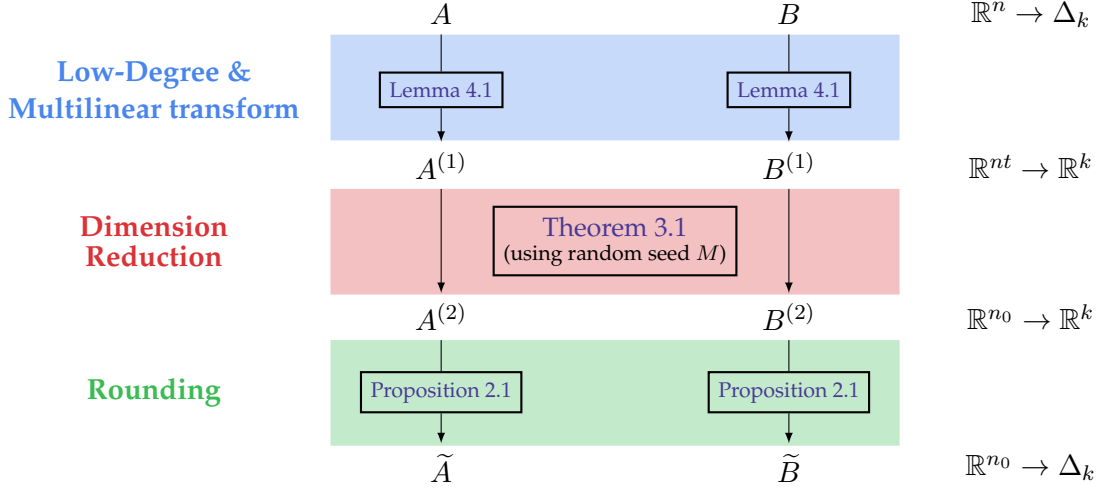


Figure 3: Transformations for Non-interactive simulation from Correlated Gaussian Sources

This lemma is itself proved in two stages. The first stage transforms general functions to low-degree polynomials by applying a small noise operator (making the functions have “decaying Hermite tails”) followed by truncation of the higher degree terms. The second stage transforms low-degree polynomials into multilinear ones, by replacing each variable by a normalized sum of new variables (making the functions have low mass on non-multilinear terms) followed by truncation of the non-multilinear terms.

These techniques are quite standard in literature. For the use of noise operator in the first stage see e.g. [MOO05, Mos10]. For the substitution of variables in the second stage see e.g. [DMN17]. However, since we are stating particular quantitative versions of the lemmas, we provide the proofs in Appendix B for completeness.

5 Non-Interactive Simulation from Correlated Gaussian Sources

In this section, we complete the proof of our main theorem regarding non-interactive simulation from correlated Gaussian sources, i.e. Theorem 1.2. Recall that it immediately implies Theorem 1.1 by setting $A = B = f$ and obtaining $\tilde{f} = \tilde{A} = \tilde{B}$.

Proof of Theorem 1.2. Starting with functions $A : \mathbb{R}^n \rightarrow \Delta_k$ and $B : \mathbb{R}^n \rightarrow \Delta_k$, we first apply Lemma 4.1 to transform A and B to low-degree and multilinear polynomials, and subsequently apply Theorem 3.1. Unfortunately after these transformations, the range is no longer restricted to Δ_k . Nevertheless, we do have that these transformations ensure that the functions still output something “close” to the simplex Δ_k . This allows us to apply the rounding operator and get the range as Δ_k again (using Proposition 2.1). An overview of the transformations done is presented in Figure 3.

We thus transform A and B through each of the following steps. At each step, we approximately preserve the correlation $\langle A_i, B_j \rangle$ for every $i, j \in [k]$. Additionally, in each step $\|\mathcal{R}(A) - A\|_2$ and $\|\mathcal{R}(B) - B\|_2$ doesn’t increase significantly (note that, to begin with, the range of A and B is Δ_k and hence we start with $\|\mathcal{R}(A) - A\|_2 = \|\mathcal{R}(B) - B\|_2 = 0$).

1. **Transformation to Low-Degree & Multilinear:** We apply Lemma 4.1 on A and B with parameter δ (chosen later), setting $d = d(\rho, k, \delta)$ and $t = t(d, k, \delta)$ as required, to get degree- d and multilinear $A^{(1)} : \mathbb{R}^n \rightarrow \mathbb{R}^k$ and $B^{(1)} : \mathbb{R}^n \rightarrow \mathbb{R}^k$. Moreover, we have that for every $i, j \in [k]$,

$$\left| \langle A_i^{(1)}, B_j^{(1)} \rangle_{\mathcal{G}_\rho^{\otimes nt}} - \langle A_i, B_j \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| \leq \delta \quad (2)$$

Additionally, we have that $\|\mathcal{R}(A^{(1)}) - A^{(1)}\|_2 \leq \|\mathcal{R}(A) - A\|_2 + \delta \leq \delta$ and similarly for $B^{(1)}$.

2. **Dimension reduction:** We apply [Theorem 3.1](#) with parameter δ/k^2 , setting $n_0 = n_0(d, \rho, \delta/k^2)$ as required, on individual coordinates of $A^{(1)}$ and $B^{(1)}$ to obtain functions $A^{(2)} : \mathbb{R}^{n_0} \rightarrow \mathbb{R}^k$ and $B^{(2)} : \mathbb{R}^{n_0} \rightarrow \mathbb{R}^k$. Taking a union bound, we have that with probability at least $1 - \delta$, it holds for every $i, j \in [k]$ that,

$$\left| \langle A_i^{(2)}, B_j^{(2)} \rangle_{\mathcal{G}_\rho^{\otimes n_0}} - \langle A_i^{(1)}, B_j^{(1)} \rangle_{\mathcal{G}_\rho^{\otimes nt}} \right| \leq \delta \quad (3)$$

From [Proposition 3.2](#), we have that with probability $1 - 4\delta$,

$$\|\mathcal{R}(A^{(2)}) - A^{(2)}\|_2 \leq \sqrt{\|\mathcal{R}(A^{(1)}) - A^{(1)}\|_2} \leq \sqrt{\delta}$$

$$\|\mathcal{R}(B^{(2)}) - B^{(2)}\|_2 \leq \sqrt{\|\mathcal{R}(B^{(1)}) - B^{(1)}\|_2} \leq \sqrt{\delta}$$

Note that this is the only randomized step in the entire transformation, and it succeeds in obtaining the above three statements with probability at least $1 - 5\delta$.

3. **Rounding to Δ_k :** Finally, we set $\tilde{A} = \mathcal{R}(A^{(2)})$ and $\tilde{B} = \mathcal{R}(B^{(2)})$. Thus, assuming the previous step succeeds, we have that $\|\tilde{A}_i - A_i^{(2)}\|_2 \leq \sqrt{\delta}$ and $\|\tilde{B}_j - B_j^{(2)}\|_2 \leq \sqrt{\delta}$. Hence we can invoke [Proposition 2.1](#), to conclude that,

$$\left| \langle \tilde{A}_i, \tilde{B}_j \rangle_{\mathcal{G}_\rho^{\otimes n_0}} - \langle A_i^{(2)}, B_j^{(2)} \rangle_{\mathcal{G}_\rho^{\otimes n_0}} \right| \leq 2\sqrt{\delta}. \quad (4)$$

Thus we started with functions $A : \mathbb{R}^n \rightarrow \Delta_k$ and $B : \mathbb{R}^n \rightarrow \Delta_k$ and ended with functions $\tilde{A} : \mathbb{R}^{n_0} \rightarrow \Delta_k$ and $\tilde{B} : \mathbb{R}^{n_0} \rightarrow \Delta_k$ such that for every $i, j \in [k]$ (by combining [Equations \(2\)](#) to [\(4\)](#)) it holds that,

$$\left| \langle \tilde{A}_i, \tilde{B}_j \rangle_{\mathcal{G}_\rho^{\otimes n_0}} - \langle A_i, B_j \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| \leq O(\sqrt{\delta})$$

Thus, more strongly, if we instantiate $\delta = O(\varepsilon^2/k^4)$, then we get that our entire transformation succeeds with probability $1 - \varepsilon$ in obtaining \tilde{A} and \tilde{B} such that,

$$d_{\text{TV}} \left((A(\mathbf{X}), B(\mathbf{Y}))_{(\mathbf{X}, \mathbf{Y}) \sim \mathcal{G}_\rho^{\otimes n}}, (\tilde{A}(\mathbf{a}), \tilde{B}(\mathbf{b}))_{(\mathbf{a}, \mathbf{b}) \sim \mathcal{G}_\rho^{\otimes n_0}} \right) \leq \varepsilon.$$

It is easy to see that n_0 works out to be

$$n_0 = \frac{d^{O(d)}}{\delta^4} = \exp \left(\tilde{O} \left(\frac{k^{4.5}}{\varepsilon^2(1-\rho)} \right) \right) = \exp \left(\text{poly} \left(k, \frac{1}{\varepsilon}, \frac{1}{1-\rho} \right) \right).$$

□

6 Acknowledgments

The authors are extremely grateful to Madhu Sudan for helpful guidance throughout all the stages of this project. The authors would also like to thank Anindya De, Elchanan Mossel and Joe Neeman for clarifying explanations of their papers and helpful discussions. PK would like to thank Irit Dinur, Mika Göös, Raghu Meka and Li-Yang Tan for helpful discussions.

References

- [AC93] Rudolf Ahlswede and Imre Csiszár. Common randomness in information theory and cryptography. part i: secret sharing. *IEEE Transactions on Information Theory*, 39(4), 1993. 8
- [AC98] Rudolf Ahlswede and Imre Csiszár. Common randomness in information theory and cryptography. ii. cr capacity. *Information Theory, IEEE Transactions on*, 44(1):225–240, 1998. 8
- [AL06] Noga Alon and Eyal Lubetzky. The shannon capacity of a graph and the independence numbers of its powers. *Information Theory, IEEE Transactions on*, 52(5):2172–2176, 2006. 8
- [Bei12] Salman Beigi. A new quantum data processing inequality. *CoRR*, abs/1210.1689, 2012. 8
- [BG15] Salman Beigi and Amin Gohari. On the duality of additivity and tensorization. *arXiv preprint arXiv:1502.00827*, 2015. 8
- [BGI14] Mohammad Bavarian, Dmitry Gavinsky, and Tsuyoshi Ito. On the role of shared randomness in simultaneous communication. In *Automata, Languages, and Programming*, pages 150–162. Springer, 2014. 8
- [BHH⁺08] Boaz Barak, Moritz Hardt, Ishay Haviv, Anup Rao, Oded Regev, and David Steurer. Rounding parallel repetitions of unique games. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 374–383, 2008. 9
- [BK18] Mark Braverman and Young Kun-Ko. Information value of two-prover games. In *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, pages 12:1–12:15, 2018. 9
- [BM11] Andrej Bogdanov and Elchanan Mossel. On extracting common random bits from correlated sources. *Information Theory, IEEE Transactions on*, 57(10):6351–6355, 2011. 8
- [Bor85] Christer Borell. Geometric bounds on the ornstein-uhlenbeck velocity process. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 70(1):1–13, 1985. 1, 5
- [BR11] Mark Braverman and Anup Rao. Information equals amortized communication. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 748–757. IEEE, 2011. 9
- [BS94] Gilles Brassard and Louis Salvail. Secret-key reconciliation by public discussion. In *advances in Cryptology—EUROCRYPT’93*, pages 410–423. Springer, 1994. 8
- [BS15] Mark Braverman and Jon Schneider. Information complexity is computable. *arXiv preprint arXiv:1502.02971*, 2015. 9
- [CDS08] Eric Chitambar, Runyao Duan, and Yaoyun Shi. Tripartite entanglement transformations and tensor rank. *Physical review letters*, 101(14):140502, 2008. 8
- [CGMS14] Clement Canonne, Venkat Guruswami, Raghu Meka, and Madhu Sudan. Communication with imperfectly shared randomness. *ITCS*, 2014. 8
- [CN00] Imre Csiszár and Prakash Narayan. Common randomness and secret key generation with a helper. *Information Theory, IEEE Transactions on*, 46(2):344–366, 2000. 8
- [DB13] Payam Delgosha and Salman Beigi. Impossibility of local state transformation via hypercontractivity. *CoRR*, abs/1307.2747, 2013. 8

- [DB14] Payam Delgosha and Salman Beigi. Impossibility of local state transformation via hypercontractivity. *Communications in Mathematical Physics*, 332(1):449–476, 2014. 8
- [DG03] Sanjoy Dasgupta and Anupam Gupta. An elementary proof of a theorem of Johnson and Lindenstrauss. *Random Structures & Algorithms*, 22(1):60–65, 2003. 6
- [DMN17] Anindya De, Elchanan Mossel, and Joe Neeman. Noise stability is computable and approximately low-dimensional. In *32nd Computational Complexity Conference, CCC 2017, July 6-9, 2017, Riga, Latvia*, pages 10:1–10:11, 2017. 2, 3, 7, 8, 9, 14
- [DMN18] Anindya De, Elchanan Mossel, and Joe Neeman. Non interactive simulation of correlated distributions is decidable. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 2728–2746, 2018. 3, 5, 7, 9
- [DS14] Anindya De and Rocco A. Servedio. Efficient deterministic approximate counting for low-degree polynomial threshold functions. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 832–841. ACM, 2014. 2, 9
- [Geb41] Hans Gebelein. Das statistische problem der korrelation als variations-und eigenwertproblem und sein zusammenhang mit der ausgleichsrechnung. *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik*, 21(6):364–379, 1941. 4
- [GJ18] Badih Ghazi and T. S. Jayram. Resource-efficient common randomness and secret-key schemes. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1834–1853, 2018. 8
- [GK73] Peter Gács and János Körner. Common information is far less than mutual information. *Problems of Control and Information Theory*, 2(2):149–162, 1973. 4
- [GKR17] Badih Ghazi, Pritish Kamath, and Prasad Raghavendra. Dimension reduction for polynomials over gaussian space and applications. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:125, 2017. 5
- [GKS16a] Badih Ghazi, Pritish Kamath, and Madhu Sudan. Communication complexity of permutation-invariant functions. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1902–1921, 2016. 8
- [GKS16b] Badih Ghazi, Pritish Kamath, and Madhu Sudan. Decidability of non-interactive simulation of joint distributions. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 545–554, 2016. 4, 5, 8
- [GW95] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, 1995. 4
- [Hei12] Steven Heilman. Euclidean partitions optimizing noise stability. *CoRR*, abs/1211.7138, 2012. 2
- [Hir35] Hermann O. Hirschfeld. A connection between correlation and contingency. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 31, pages 520–524. Cambridge Univ Press, 1935. 4

- [HMN16] Steven Heilman, Elchanan Mossel, and Joe Neeman. Standard simplices and pluralities are not the most noise stable. *Israel Journal of Mathematics*, 213(1):33–53, 2016. 1
- [Hol09] Thomas Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory of Computing*, 5(1):141–172, 2009. 9
- [IM12] Marcus Isaksson and Elchanan Mossel. Maximally stable gaussian partitions with discrete applications. *Israel Journal of Mathematics*, 189(1):347–396, 2012. 1
- [JL84] William Johnson and Joram Lindenstrauss. Extensions of lipschitz maps into a hilbert space. 26:189–206, 01 1984. 6
- [KA16] Sudeep Kamath and Venkat Anantharam. On non-interactive simulation of joint distributions. *IEEE Trans. Information Theory*, 62(6):3419–3435, 2016. 4, 5, 8
- [KKM⁺11] Julia Kempe, Hirotada Kobayashi, Keiji Matsumoto, Ben Toner, and Thomas Vidick. Entangled games are hard to approximate. *SIAM J. Comput.*, 40(3):848–877, 2011. 8
- [KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. Optimal inapproximability results for max-cut and other 2-variable csps? *SIAM Journal on Computing*, 37(1):319–357, 2007. 1
- [Lov79] László Lovász. On the shannon capacity of a graph. *Information Theory, IEEE Transactions on*, 25(1):1–7, 1979. 8
- [Mau93] Ueli M Maurer. Secret key agreement by public discussion from common information. *Information Theory, IEEE Transactions on*, 39(3):733–742, 1993. 8
- [MO04] Elchanan Mossel and Ryan O’Donnell. Coin flipping from a cosmic source: On error correction of truly random bits. *arXiv preprint math/0406504*, 2004. 8
- [MOO05] Elchanan Mossel, Ryan O’Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, pages 21–30. IEEE, 2005. 1, 14
- [MOR⁺06] Elchanan Mossel, Ryan O’Donnell, Oded Regev, Jeffrey E Steif, and Benny Sudakov. Non-interactive correlation distillation, inhomogeneous markov chains, and the reverse bonami-beckner inequality. *Israel Journal of Mathematics*, 154(1):299–336, 2006. 8
- [Mos10] Elchanan Mossel. Gaussian bounds for noise correlation of functions. *Geometric and Functional Analysis*, 19(6):1713–1756, 2010. 1, 14
- [New91] Ilan Newman. Private vs. common random bits in communication complexity. *Information processing letters*, 39(2):67–71, 1991. 8
- [Nie99] Michael A Nielsen. Conditions for a class of entanglement transformations. *Physical Review Letters*, 83(2):436, 1999. 8
- [ope] OpenQIPProblemsWiki - All the Bell Inequalities. <http://qiq.itp.uni-hannover.de/qiproblems/1>. Accessed: 2016-07-12. 9
- [Rao11] Anup Rao. Parallel repetition in projection games and a concentration bound. *SIAM J. Comput.*, 40(6):1871–1891, 2011. 9
- [Raz98] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998. 9

- [Raz11] Ran Raz. A counterexample to strong parallel repetition. *SIAM J. Comput.*, 40(3):771–777, 2011. 9
- [Rén59] Alfréd Rényi. On measures of dependence. *Acta mathematica hungarica*, 10(3-4):441–451, 1959. 4
- [RW05] Renato Renner and Stefan Wolf. Simple and tight bounds for information reconciliation and privacy amplification. In *Advances in cryptology-ASIACRYPT 2005*, pages 199–216. Springer, 2005. 8
- [Sha56] Claude E Shannon. The zero error capacity of a noisy channel. *Information Theory, IRE Transactions on*, 2(3):8–19, 1956. 8
- [Wit75] Hans S Witsenhausen. On sequences of pairs of dependent random variables. *SIAM Journal on Applied Mathematics*, 28(1):100–113, 1975. 4
- [Wol07] Pawel Wolff. Hypercontractivity of simple random variables. *Studia Mathematica*, 180(3):219–236, 2007. 29
- [Wyn75] Aaron D. Wyner. The common information of two dependent random variables. *IEEE Transactions on Information Theory*, 21(2):163–179, 1975. 4

A Proofs of Mean and Variance Bounds in Dimension Reduction

In this section, we provide the proof of [Lemma 3.3](#). This is the main new technical component introduced in this paper. Even though the calculations might seem cumbersome, they involve mostly elementary steps. To understand the high level picture, we recommend the reader to go through a short proof sketch presented in [Section 3.1](#).

Recall that starting with degree d multilinear polynomials $A : \mathbb{R}^n \rightarrow \mathbb{R}$ and $B : \mathbb{R}^n \rightarrow \mathbb{R}$, we defined functions $A_M : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$ and $B_M : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$, for $M \sim \gamma_1^{\otimes(n \times n_0)}$, as

$$A_M(a) = A\left(\frac{Ma}{\|a\|_2}\right) \quad \text{and} \quad B_M(b) = B\left(\frac{Mb}{\|b\|_2}\right) \quad \text{for } a, b \in \mathbb{R}^{n_0} \setminus \{0\} .$$

and we defined their correlation as $F(M) \stackrel{\text{def}}{=} \langle A_M, B_M \rangle_{\mathcal{G}_\rho^{\otimes n_0}}$. [Lemma 3.3](#) proves bounds on the mean and variance of $F(M)$, which we restate below for convenience.

Lemma 3.3. (Mean & Variance Bound). *Given d and δ , there exists an explicitly computable $n_0 := n_0(d, \delta)$ such that for $M \sim \gamma_1^{\otimes(n \times n_0)}$,*

$$\begin{aligned} \left| \mathbb{E}_M F(M) - \langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| &\leq \delta && \text{(Mean bound)} \\ \text{Var}_M(F(M)) &\leq \delta && \text{(Variance bound)} \end{aligned}$$

In particular, one may take $n_0 = \frac{d^{O(d)}}{\delta^2}$.

We break down the full proof into the following three modular steps.

1. In [Appendix A.1](#), we prove a *meta-lemma* ([Lemma A.1](#)) that will help us prove both the mean and variance bounds; indeed this meta-lemma is at the heart of why [Theorem 3.1](#) holds. Morally, this lemma says that if we have an expectation of a product of a small number of inner products of normalized correlated Gaussian vectors, then, we can exchange the product and the expectations while incurring only a small additive error. [Lemma A.1](#) is the main take away from this section, and the reader may skip to [Appendices A.2](#) and [A.3](#) to see the rest of the proof.

2. In [Appendix A.2](#), we prove bounds on the mean and co-variances of degree- d multilinear monomials, under the above transformation of replacing $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^n$ (inputs to A and B) by $\frac{M\mathbf{a}}{\|\mathbf{a}\|_2}$ and $\frac{M\mathbf{b}}{\|\mathbf{b}\|_2}$ respectively.
3. In [Appendix A.3](#), we finally use the above bounds on mean and co-variances of degree- d multilinear monomials in order to prove [Lemma 3.3](#).

Remark. To make our notations convenient, we will often write equations such as $\alpha = \beta \pm \varepsilon$ which is to be interpreted as $|\alpha - \beta| \leq \varepsilon$.

A.1 Product of Inner Products of Normalized Correlated Gaussian Vectors

The following is the main lemma in this subsection (this is the *meta-lemma* alluded to earlier).

Lemma A.1. *Given $d, D \in \mathbb{Z}_{\geq 0}$ and $\delta > 0$ (with D sufficiently larger than d), let $(\mathbf{u}_1, \dots, \mathbf{u}_d, \mathbf{v}_1, \dots, \mathbf{v}_d)$ be a $2dD$ -dimensional multivariate Gaussian distribution such that,*

- each $\mathbf{u}_i, \mathbf{v}_i \in \mathbb{R}^D$ are marginally distributed as standard D -dimensional Gaussians, i.e. γ_D .
- for each $j \in [D]$, the joint distribution of the j -th coordinates, i.e., $(\mathbf{u}_{1,j}, \dots, \mathbf{u}_{d,j}, \mathbf{v}_{1,j}, \dots, \mathbf{v}_{d,j})$, is independent across different values of j .

Then,

$$\left| \mathbb{E}_{\{\mathbf{u}_i, \mathbf{v}_i\}_i} \left[\prod_{i=1}^d \frac{\langle \mathbf{u}_i, \mathbf{v}_i \rangle}{\|\mathbf{u}_i\|_2 \|\mathbf{v}_i\|_2} \right] - \prod_{i=1}^d \mathbb{E}_{\{\mathbf{u}_i, \mathbf{v}_i\}_i} \left[\frac{\langle \mathbf{u}_i, \mathbf{v}_i \rangle}{D} \right] \right| \leq \frac{d^{O(d)}}{D}.$$

We point out that there are two steps taking place in [Lemma A.1](#):

- (i) the replacement of $\|\mathbf{u}_i\|_2$ (and $\|\mathbf{v}_i\|_2$) by \sqrt{D} (around which it is tightly concentrated), and
- (ii) the interchanging of the expectation and the product.

We will handle each of these changes one by one.

Product of Negative Moments of ℓ_2 -norm of Correlated Gaussian vectors

In order to handle the replacement of $\|\mathbf{u}_i\|_2$ (and $\|\mathbf{v}_i\|_2$) by \sqrt{D} , we will prove some bounds on the mean and variance of products of negative powers of the ℓ_2 -norm of a standard Gaussian vector.

Lemma A.2. *Let $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_\ell$ be (possibly correlated) multivariate Gaussians where each $\mathbf{w}_i \in \mathbb{R}^D$ is marginally distributed as γ_D , and let d_1, d_2, \dots, d_ℓ be non-negative integers with $d := \sum_{i=1}^{\ell} d_i$. Then,*

$$\left| \mathbb{E} \left[\prod_{i=1}^{\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right] - \frac{1}{D^{d/2}} \right| \leq O\left(\frac{d^5}{D^{\frac{d}{2}+1}}\right),$$

$$\text{Var} \left[\prod_{i=1}^{\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right] \leq O\left(\frac{d^5}{D^{d+1}}\right).$$

Remark A.3. *It is conceivable that the bounds in [Lemma A.2](#) could be improved in terms of the dependence on d . However, this was not central to our application, so we go ahead with the stated bounds. The main point to note in the above lemma is the extra factor of D in the denominator.*

We start out by first proving the base case where we have a single vector \mathbf{w} , that is, $\ell = 1$.

Proposition A.4. *There exists an absolute constant C such that for sufficiently large $d, D \in \mathbb{Z}_{>0}$ satisfying $D > Cd^2$, we have that for $\mathbf{w} \sim \gamma_D$,*

$$\left| \mathbb{E}_{\mathbf{w}} \left[\frac{1}{\|\mathbf{w}\|_2^d} \right] - \frac{1}{D^{d/2}} \right| \leq C \cdot \left(\frac{d^2}{D^{d/2+1}} \right), \quad (5)$$

$$\text{Var}_{\mathbf{w}} \left[\frac{1}{\|\mathbf{w}\|_2^d} \right] \leq 8C \cdot \left(\frac{d^2}{D^{d+1}} \right). \quad (6)$$

Proof. It is well-known that the distribution of $\|\mathbf{w}\|_2$ follows a χ -distribution with parameter D , and whose probability density function is given by

$$f_D(x) = \frac{x^{D-1} \cdot e^{-\frac{x^2}{2}}}{2^{\frac{D}{2}-1} \cdot \Gamma(\frac{D}{2})}, \quad (x \in \mathbb{R}_{\geq 0})$$

where $\Gamma(\cdot)$ denotes the Gamma function. Thus, we have that

$$\begin{aligned} \mathbb{E}_{\mathbf{w}} \left[\frac{1}{\|\mathbf{w}\|_2^d} \right] &= \int_0^\infty \frac{1}{x^d} \cdot f_D(x) dx = \int_0^\infty \frac{x^{D-d-1} \cdot e^{-\frac{x^2}{2}}}{2^{\frac{D}{2}-1} \cdot \Gamma(\frac{D}{2})} dx \\ &= \frac{2^{\frac{D-d-1}{2}} \cdot \Gamma(\frac{D-d}{2})}{2^{\frac{D}{2}-1} \cdot \Gamma(\frac{D}{2})} = \frac{1}{D^{d/2}} \cdot \left(1 \pm O\left(\frac{d^2}{D}\right) \right), \end{aligned}$$

where the last equality follows from the Stirling's approximation of the Gamma function, which holds for every real number $z > 0$:

$$\Gamma(z+1) = \sqrt{2\pi z} \cdot \left(\frac{z}{e}\right)^z \cdot \left(1 \pm O\left(\frac{1}{z}\right)\right).$$

This completes the proof of Equation (5), for the explicit constant C that can be derived from the Stirling's approximation. Now, Equation (6) immediately follows as:

$$\begin{aligned} \text{Var}_{\mathbf{w}} \left[\frac{1}{\|\mathbf{w}\|_2^d} \right] &= \mathbb{E}_{\mathbf{w}} \left[\frac{1}{\|\mathbf{w}\|_2^{2d}} \right] - \mathbb{E}_{\mathbf{w}} \left[\frac{1}{\|\mathbf{w}\|_2^d} \right]^2 \\ &= \left(\frac{1}{D^d} \pm C \cdot \left(\frac{(2d)^2}{D^{d+1}} \right) \right) - \left(\frac{1}{D^{d/2}} \pm C \cdot \left(\frac{d^2}{D^{d/2+1}} \right) \right)^2 \\ &\leq 8C \cdot \left(\frac{d^2}{D^{d+1}} \right), \end{aligned}$$

where, we use that D is sufficiently large that $C^2 \left(\frac{d^4}{D^{d+2}} \right) < 2C \cdot \left(\frac{d^2}{D^{d+1}} \right)$, i.e. $D > Cd^2$. □

We now show how to generalize the above to prove Lemma A.2.

Proof of Lemma A.2. More specifically, we will show that,

$$\left| \mathbb{E} \left[\prod_{i=1}^{\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right] - \frac{1}{D^{d/2}} \right| \leq C \cdot \ell^3 \cdot \left(\frac{d^2}{D^{d/2+1}} \right) \quad (7)$$

$$\text{Var} \left[\prod_{i=1}^{\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right] \leq 8C \cdot \ell^3 \cdot \left(\frac{d^2}{D^{d+1}} \right) \quad (8)$$

where C is the absolute constant (as obtained in Proposition A.4). This implies the lemma since $\ell \leq d$.

We proceed by induction on ℓ (more specifically on $\log \ell$). For $\ell = 1$, the bound immediately follows from [Proposition A.4](#). For the inductive step, we assume that the bound in [Equations \(7\) and \(8\)](#) holds for ℓ , and we prove that the bound also holds for 2ℓ . While it may seem that our bounds are being proven only when ℓ is a power of 2, it is not hard to see that our proof could be done for non powers of 2 as well, giving a bound that is monotonically increasing in ℓ and hence it suffices having proved it for ℓ that are powers of 2. Let $d_1, d_2, \dots, d_{2\ell}$ be non-negative integers with $d := \sum_{i=1}^{2\ell} d_i$. For notational convenience, let $s_1 = \sum_{i=1}^{\ell} d_i$ and $s_2 = \sum_{i=\ell+1}^{2\ell} d_i$, and so $d = s_1 + s_2$.

We will first prove [Equation \(7\)](#) inductively by using the following idea: for any two random variables \mathbf{X} and \mathbf{Y} , we have $\mathbb{E}[\mathbf{X}\mathbf{Y}] = \mathbb{E}[\mathbf{X}]\mathbb{E}[\mathbf{Y}] + \text{Cov}[\mathbf{X}, \mathbf{Y}]$ and $|\text{Cov}[\mathbf{X}, \mathbf{Y}]| \leq \sqrt{\text{Var}[\mathbf{X}] \cdot \text{Var}[\mathbf{Y}]}$ (by Cauchy-Schwarz inequality), and hence $\mathbb{E}[\mathbf{X}\mathbf{Y}] = \mathbb{E}[\mathbf{X}]\mathbb{E}[\mathbf{Y}] \pm \sqrt{\text{Var}[\mathbf{X}] \cdot \text{Var}[\mathbf{Y}]}$. Thus, we get,

$$\mathbb{E} \left[\prod_{i=1}^{2\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right] = \mathbb{E} \left[\prod_{i=1}^{\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right] \cdot \mathbb{E} \left[\prod_{i=\ell+1}^{2\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right] \pm \sqrt{\text{Var} \left[\prod_{i=1}^{\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right] \cdot \text{Var} \left[\prod_{i=\ell+1}^{2\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right]}. \quad (9)$$

Using the inductive assumption w.r.t. ℓ , we get that,

$$\mathbb{E} \left[\prod_{i=1}^{\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right] = \frac{1}{D^{s_1/2}} \left(1 \pm C \cdot \ell^3 \cdot \left(\frac{s_1^2}{D} \right) \right) \quad (10)$$

$$\mathbb{E} \left[\prod_{i=\ell+1}^{2\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right] = \frac{1}{D^{s_2/2}} \left(1 \pm C \cdot \ell^3 \cdot \left(\frac{s_2^2}{D} \right) \right) \quad (11)$$

and

$$\text{Var} \left[\prod_{i=1}^{\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right] \leq \frac{1}{D^{s_1}} \cdot 8C \cdot \ell^3 \cdot \left(\frac{s_1^2}{D} \right) \quad (12)$$

$$\text{Var} \left[\prod_{i=\ell+1}^{2\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right] \leq \frac{1}{D^{s_2}} \cdot 8C \cdot \ell^3 \cdot \left(\frac{s_2^2}{D} \right) \quad (13)$$

Plugging [Equations \(10\) to \(13\)](#) in [Equation \(9\)](#), it is not hard to see that,

$$\mathbb{E} \left[\prod_{i=1}^{2\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right] = \frac{1}{D^{d/2}} \left(1 \pm C \cdot (2\ell)^3 \cdot \left(\frac{d^2}{D} \right) \right).$$

This completes the proof of [Equation \(7\)](#). Now, [Equation \(8\)](#) follows easily as,

$$\begin{aligned} \text{Var} \left[\prod_{i=1}^{2\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right] &= \mathbb{E} \left[\prod_{i=1}^{2\ell} \frac{1}{\|\mathbf{w}_i\|_2^{2d_i}} \right] - \mathbb{E} \left[\prod_{i=1}^{2\ell} \frac{1}{\|\mathbf{w}_i\|_2^{d_i}} \right]^2 \\ &= \left(\frac{1}{D^d} \pm C \cdot (2\ell)^3 \cdot \left(\frac{(2d)^2}{D^{d+1}} \right) \right) - \left(\frac{1}{D^{d/2}} \pm C \cdot (2\ell)^3 \cdot \left(\frac{d^2}{D^{d/2+1}} \right) \right)^2 \\ &\leq 8C \cdot (2\ell)^3 \cdot \left(\frac{d^2}{D^{d+1}} \right). \end{aligned}$$

□

Interchanging Product and Expectation

In order to handle the interchanging of the product and expectation operations, we will show the following lemma.

Lemma A.5. Let $(\mathbf{u}_1, \dots, \mathbf{u}_d, \mathbf{v}_1, \dots, \mathbf{v}_d)$ be distributed as in [Lemma A.1](#). Then,

$$\left| \mathbb{E}_{\{\mathbf{u}_i, \mathbf{v}_i\}_i} \left[\prod_{i=1}^d \langle \mathbf{u}_i, \mathbf{v}_i \rangle \right] - \prod_{i=1}^d \mathbb{E}_{\{\mathbf{u}_i, \mathbf{v}_i\}_i} [\langle \mathbf{u}_i, \mathbf{v}_i \rangle] \right| \leq d^{O(d)} \cdot D^{d-1}.$$

Remark A.6. The $d^{O(d)}$ term has an explicit expression, although we only highlight its qualitative nature for clarity. Again, it is conceivable that the bounds in [Lemma A.5](#) could be improved in terms of the dependence on d , although we suspect that it is tight upto constant factors in the exponent. Anyhow, this was not central to our application, so we go ahead with the stated bounds. The main point to note in the above lemma is that the exponent of D is $(d-1)$ instead of d .

To prove the lemma, we first obtain the following proposition on moments of a multivariate Gaussian.

Proposition A.7. Let $\mathbf{w} \in \mathbb{R}^\ell$ be any multivariate Gaussian vector with each coordinate marginally distributed according to γ_1 . Let d_1, d_2, \dots, d_ℓ be non-negative integers such that $d := \sum_{i=1}^\ell d_i$. Then,

$$\left| \mathbb{E} \left[\prod_{i=1}^\ell \mathbf{w}_i^{d_i} \right] \right| \leq (2d)^{3d}. \quad (14)$$

Proof. More specifically we will show that when ℓ is a power of 2,

$$\left| \mathbb{E} \left[\prod_{i=1}^\ell \mathbf{w}_i^{d_i} \right] \right| \leq 2^{\ell-1} (\ell d)^d. \quad (15)$$

It is easy to see that this immediately implies the bound of $2^d \cdot d^{2d}$ in the main lemma, since $\ell \leq d$. However if ℓ is not a power of 2 we can round it up to the nearest power of 2, which amounts to substituting $\ell \leq 2d$ in the above, obtaining a bound of $2^{3d} \cdot d^{2d} \leq (2d)^{3d}$.

We proceed by induction on ℓ (more specifically on $\log \ell$). For $\ell = 1$, we use the well-known fact that for $w \sim \gamma_1$,

$$|\mathbb{E}[w^d]| = \begin{cases} 0 & \text{if } d \text{ is odd} \\ (d-1)!! & \text{if } d \text{ is even} \end{cases} \leq d^d,$$

where $(d-1)!!$ denotes the double factorial of $(d-1)$, i.e., the product of all integers from 1 to $d-1$ that have the same parity as $d-1$. For the inductive step, we assume that the bound in (15) holds for ℓ and we show that it also holds for 2ℓ . For notational convenience, let $s_1 = \sum_{i=1}^\ell d_i$ and $s_2 = \sum_{i=\ell+1}^{2\ell} d_i$, and so $d = s_1 + s_2$.

The main idea that we use to prove the inductive step is Cauchy-Schwarz inequality. Additionally, we use that $\text{Var}[\mathbf{X}] \leq \mathbb{E}[\mathbf{X}^2]$. Thus, we get,

$$\left| \mathbb{E} \left[\prod_{i=1}^{2\ell} \mathbf{w}_i^{d_i} \right] \right| \leq \sqrt{\mathbb{E} \left[\prod_{i=1}^\ell \mathbf{w}_i^{2d_i} \right] \cdot \mathbb{E} \left[\prod_{i=\ell+1}^{2\ell} \mathbf{w}_i^{2d_i} \right]} \leq \sqrt{2^{\ell-1} (2\ell s_1)^{2s_1} \cdot 2^{\ell-1} (2\ell s_2)^{2s_2}} \leq 2^{2\ell-1} (2\ell d)^d,$$

where, we use the inductive assumption regarding product of ℓ terms and that $s_1 + s_2 = d$. \square

Using the above proposition, we are now able to prove [Lemma A.5](#).

Proof of Lemma A.5. Let $S \subseteq [D]^d$ be the set of all tuples $c \in [D]^d$ such that $c_j \neq c_k$ for all $j \neq k \in [d]$. Let \bar{S} denote the complement of S in $[D]^d$. Note that $|\bar{S}| \leq d^2 \cdot D^{d-1}$. We have that

$$\mathbb{E} \left[\prod_{i=1}^d \langle \mathbf{u}_i, \mathbf{v}_i \rangle \right] = \mathbb{E} \left[\prod_{i=1}^d \sum_{k=1}^D \mathbf{u}_{i,k} \mathbf{v}_{i,k} \right] = \sum_{c \in [D]^d} \mathbb{E} \left[\prod_{i=1}^d \mathbf{u}_{i,c_i} \mathbf{v}_{i,c_i} \right]$$

$$\begin{aligned}
&= \sum_{c \in S} \mathbb{E} \left[\prod_{i=1}^d \mathbf{u}_{i,c_i} \mathbf{v}_{i,c_i} \right] + \sum_{c \in \bar{S}} \mathbb{E} \left[\prod_{i=1}^d \mathbf{u}_{i,c_i} \mathbf{v}_{i,c_i} \right] \\
&= \sum_{c \in S} \prod_{i=1}^d \mathbb{E}[\mathbf{u}_{i,c_i} \mathbf{v}_{i,c_i}] + \sum_{c \in \bar{S}} \mathbb{E} \left[\prod_{i=1}^d \mathbf{u}_{i,c_i} \mathbf{v}_{i,c_i} \right], \tag{16}
\end{aligned}$$

where the last equality follows from the assumption that the distribution of $(\mathbf{u}_{1,j}, \dots, \mathbf{u}_{d,j}, \mathbf{v}_{1,j}, \dots, \mathbf{v}_{d,j})$ is independent across $j \in [D]$. On the other hand, we have that

$$\begin{aligned}
\prod_{i=1}^d \mathbb{E}[\langle \mathbf{u}_i, \mathbf{v}_i \rangle] &= \prod_{i=1}^d \mathbb{E} \left[\sum_{k=1}^D \mathbf{u}_{i,k} \mathbf{v}_{i,k} \right] = \sum_{c \in [D]^d} \prod_{i=1}^d \mathbb{E}[\mathbf{u}_{i,c_i} \mathbf{v}_{i,c_i}] \\
&= \sum_{c \in S} \prod_{i=1}^d \mathbb{E}[\mathbf{u}_{i,c_i} \mathbf{v}_{i,c_i}] + \sum_{c \in \bar{S}} \prod_{i=1}^d \mathbb{E}[\mathbf{u}_{i,c_i} \mathbf{v}_{i,c_i}] \tag{17}
\end{aligned}$$

Combining Equations (16) and (17), we get

$$\begin{aligned}
\left| \mathbb{E} \left[\prod_{i=1}^d \langle \mathbf{u}_i, \mathbf{v}_i \rangle \right] - \prod_{i=1}^d \mathbb{E}[\langle \mathbf{u}_i, \mathbf{v}_i \rangle] \right| &= \left| \sum_{c \in \bar{S}} \left(\mathbb{E} \left[\prod_{i=1}^d \mathbf{u}_{i,c_i} \mathbf{v}_{i,c_i} \right] - \prod_{i=1}^d \mathbb{E}[\mathbf{u}_{i,c_i} \mathbf{v}_{i,c_i}] \right) \right| \\
&\leq |\bar{S}| \cdot \max_{c \in \bar{S}} \left| \mathbb{E} \left[\prod_{i=1}^d \mathbf{u}_{i,c_i} \mathbf{v}_{i,c_i} \right] - \prod_{i=1}^d \mathbb{E}[\mathbf{u}_{i,c_i} \mathbf{v}_{i,c_i}] \right| \\
&\leq d^2 \cdot D^{d-1} \cdot \left((2d)^{3d} + 1 \right) \leq d^{O(d)} \cdot D^{d-1},
\end{aligned}$$

where the second last inequality follows from the fact that $|\bar{S}| \leq d^2 \cdot D^{d-1}$ and from Proposition A.7. \square

Putting things together to prove Lemma A.1

Proof of Lemma A.1. We show the following bounds, which immediately imply Lemma A.1.

$$\left| \mathbb{E}_{\{\mathbf{u}_i, \mathbf{v}_i\}} \left[\prod_{i=1}^d \frac{\langle \mathbf{u}_i, \mathbf{v}_i \rangle}{\|\mathbf{u}_i\|_2 \|\mathbf{v}_i\|_2} \right] - \mathbb{E}_{\{\mathbf{u}_i, \mathbf{v}_i\}} \left[\prod_{i=1}^d \frac{\langle \mathbf{u}_i, \mathbf{v}_i \rangle}{D} \right] \right| \leq \frac{d^{O(d)}}{D}. \tag{18}$$

$$\left| \mathbb{E}_{\{\mathbf{u}_i, \mathbf{v}_i\}} \left[\prod_{i=1}^d \frac{\langle \mathbf{u}_i, \mathbf{v}_i \rangle}{D} \right] - \prod_{i=1}^d \mathbb{E}_{\{\mathbf{u}_i, \mathbf{v}_i\}} \left[\frac{\langle \mathbf{u}_i, \mathbf{v}_i \rangle}{D} \right] \right| \leq \frac{d^{O(d)}}{D}. \tag{19}$$

Note that Equation (19) is simply a restatement of Lemma A.5. To prove Equation (18), we define the random variables

$$\mathbf{W} := \prod_{i=1}^d \langle \mathbf{u}_i, \mathbf{v}_i \rangle \quad \text{and} \quad \mathbf{Z} := \prod_{i=1}^d \frac{1}{\|\mathbf{u}_i\|_2 \|\mathbf{v}_i\|_2} - \frac{1}{D^d}.$$

Note that Equation (18) is equivalent to showing bounds on $|\mathbb{E}[\mathbf{W} \cdot \mathbf{Z}]|$. In order to do so, we use the following four bounds:

1. $|\mathbb{E}[\mathbf{W}]| \leq D^d + d^{O(d)} \cdot D^{d-1}$. Since, by Lemma A.5, we have that

$$|\mathbb{E}[\mathbf{W}]| \leq \left| \prod_{i=1}^d \mathbb{E}[\langle \mathbf{u}_i, \mathbf{v}_i \rangle] \right| + d^{O(d)} \cdot D^{d-1} \leq D^d + d^{O(d)} \cdot D^{d-1}$$

2. $\text{Var}[\mathbf{W}] \leq d^{O(d)} \cdot D^{2d-1}$. Since,

$$\begin{aligned} \text{Var}[\mathbf{W}] &= \mathbb{E}[\mathbf{W}^2] - [\mathbb{E} \mathbf{W}]^2 \\ &= \mathbb{E} \left[\prod_{i=1}^d \langle \mathbf{u}_i, \mathbf{v}_i \rangle^2 \right] - \left[\mathbb{E} \prod_{i=1}^d \langle \mathbf{u}_i, \mathbf{v}_i \rangle \right]^2 \\ &\leq d^{O(d)} \cdot D^{2d-1} \quad \dots \text{(from Lemma A.5)} \end{aligned}$$

3. $|\mathbb{E}[\mathbf{Z}]| = O\left(\frac{d^5}{D^{2d+1}}\right)$ (follows exactly from Lemma A.2).

4. $\text{Var}[\mathbf{Z}] = O\left(\frac{d^5}{D^{2d+1}}\right)$ (follows exactly from Lemma A.2).

Thus, we can bound $|\mathbb{E}[\mathbf{W} \cdot \mathbf{Z}]|$ as,

$$|\mathbb{E}[\mathbf{W} \cdot \mathbf{Z}]| \leq |\mathbb{E}[\mathbf{W}]| \cdot |\mathbb{E}[\mathbf{Z}]| + \sqrt{\text{Var}[\mathbf{W}] \cdot \text{Var}[\mathbf{Z}]} \leq \frac{d^{O(d)}}{D}.$$

This completes the proof of Equation (18) and hence of Lemma A.1. \square

A.2 Mean & Variance Bounds for Multilinear Monomials

For the rest of this section, we simplify our notations as follows:

- For $(\mathbf{a}, \mathbf{b}) \sim \mathcal{G}_\rho^{\otimes n_0}$, we will use $\tilde{\mathbf{a}}$ and $\tilde{\mathbf{b}}$ to denote the normalized vectors $\frac{\mathbf{a}}{\|\mathbf{a}\|_2}$ and $\frac{\mathbf{b}}{\|\mathbf{b}\|_2}$ respectively.
- We will use $\mathbf{U} \in \mathbb{R}^n$ to denote $M\tilde{\mathbf{a}}$ and similarly $\mathbf{V} \in \mathbb{R}^n$ to denote $M\tilde{\mathbf{b}}$. We will also have independent variables $(\mathbf{a}', \mathbf{b}') \sim \mathcal{G}_\rho^{\otimes n_0}$, for which we use $\mathbf{U}' = M\tilde{\mathbf{a}}'$ and $\mathbf{V}' = M\tilde{\mathbf{b}}'$.
- U_i denotes the i -th coordinate of \mathbf{U} . Similarly, $m_i \in \mathbb{R}^{n_0}$ is the i -th row of M . Note that $U_i = \langle m_i, \tilde{\mathbf{a}} \rangle$. For $S \subseteq [n]$, let \mathbf{U}_S denote $\prod_{i \in S} U_i = \prod_{i \in S} \langle m_i, \tilde{\mathbf{a}} \rangle$. Similarly for \mathbf{V}_S .
- We will take expectations over random variables $M, \mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}'$. It will be understood that we are sampling $M \sim \gamma_1^{\otimes (n \times n_0)}$. Also, (\mathbf{a}, \mathbf{b}) and $(\mathbf{a}', \mathbf{b}')$ are independently sampled from $\mathcal{G}_\rho^{\otimes n_0}$.

Lemma A.8 (Mean bounds for monomials). *Given parameter d and δ , there exists an explicitly computable $n_0 := n_0(d, \delta)$ such that the following holds: For any subsets $S, T \subseteq [n]$ satisfying $|S|, |T| \leq d$, it holds that,*

$$\mathbb{E}_M \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbf{U}_S \mathbf{V}_T = \begin{cases} 0 & \text{if } S \neq T \\ \rho^{|S|} \pm \delta & \text{if } S = T \end{cases}.$$

In particular, one may take $n_0 = \frac{d^{O(d)}}{\delta}$.

Proof. We have that

$$\begin{aligned} \mathbb{E}_M \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbf{U}_S \mathbf{V}_T &= \mathbb{E}_M \mathbb{E}_{\mathbf{a}, \mathbf{b}} \left[\prod_{i \in S} U_i \cdot \prod_{i \in T} V_i \right] \\ &= \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbb{E}_M \left[\prod_{i \in S \cap T} U_i V_i \cdot \prod_{i \in S \setminus T} U_i \cdot \prod_{i \in T \setminus S} V_i \right] \\ &= \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbb{E}_M \left[\prod_{i \in S \cap T} \langle m_i, \tilde{\mathbf{a}} \rangle \langle m_i, \tilde{\mathbf{b}} \rangle \cdot \prod_{i \in S \setminus T} \langle m_i, \tilde{\mathbf{a}} \rangle \cdot \prod_{i \in T \setminus S} \langle m_i, \tilde{\mathbf{b}} \rangle \right] \end{aligned}$$

$$\mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbf{U}_S \mathbf{V}_T = \mathbb{E}_{\mathbf{a}, \mathbf{b}} \left[\prod_{i \in S \cap T} \mathbb{E}_{\mathbf{m}_i} \langle \mathbf{m}_i, \tilde{\mathbf{a}} \rangle \langle \mathbf{m}_i, \tilde{\mathbf{b}} \rangle \cdot \prod_{i \in S \setminus T} \mathbb{E}_{\mathbf{m}_i} \langle \mathbf{m}_i, \tilde{\mathbf{a}} \rangle \cdot \prod_{i \in T \setminus S} \mathbb{E}_{\mathbf{m}_i} \langle \mathbf{m}_i, \tilde{\mathbf{b}} \rangle \right], \quad (20)$$

where the last equality follows from the independence of the \mathbf{m}_i 's.

If $S \neq T$, one of $\prod_{i \in S \setminus T} \mathbb{E}_{\mathbf{m}_i} [\langle \mathbf{m}_i, \tilde{\mathbf{a}} \rangle]$ or $\prod_{i \in T \setminus S} \mathbb{E}_{\mathbf{m}_i} [\langle \mathbf{m}_i, \tilde{\mathbf{b}} \rangle]$ is 0. This is because even for any fixed vector a and for each $i \in [n]$, the random variable $\langle \mathbf{m}_i, \tilde{\mathbf{a}} \rangle$ has zero-mean (and similarly for $\langle \mathbf{m}_i, \tilde{\mathbf{b}} \rangle$). The first part of the lemma now follows from Equation (20).

If $S = T$, Equation (20) becomes

$$\begin{aligned} \mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbf{U}_S \mathbf{V}_T &= \mathbb{E}_{\mathbf{a}, \mathbf{b}} \left[\prod_{i \in S} \mathbb{E}_{\mathbf{m}_i} \frac{\langle \mathbf{m}_i, \mathbf{a} \rangle \langle \mathbf{m}_i, \mathbf{b} \rangle}{\|\mathbf{a}\|_2 \|\mathbf{b}\|_2} \right] \\ &= \mathbb{E}_{\mathbf{a}, \mathbf{b}} \left[\prod_{i \in S} \frac{\langle \mathbf{a}, \mathbf{b} \rangle}{\|\mathbf{a}\|_2 \|\mathbf{b}\|_2} \right] \left[\text{since } \mathbb{E}_{\mathbf{m}_i} \mathbf{m}_i \cdot \mathbf{m}_i^T = I_{n_0 \times n_0} \right] \\ &= \prod_{i \in S} \left[\frac{\mathbb{E}_{\mathbf{a}, \mathbf{b}} \langle \mathbf{a}, \mathbf{b} \rangle}{n_0} \right] \pm \delta \left[\text{from Lemma A.1, for } n_0 = \frac{d^{O(d)}}{\delta} \right] \\ &= \rho^{|S|} \pm \delta. \end{aligned} \quad \square$$

Lemma A.9 (Covariance bounds for monomials). *Given parameters d and δ , there exists an explicitly computable $n_0 := n_0(d, \delta)$ such that the following holds: For any subsets $S, T, S', T' \subseteq [n]$ satisfying $|S|, |T|, |S'|, |T'| \leq d$, it holds that,*

$$\begin{aligned} \text{if } S \Delta T \Delta S' \Delta T' \neq \emptyset : & \left| \mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} [\mathbf{U}_S \mathbf{V}_T \mathbf{U}_{S'} \mathbf{V}_{T'}] - \left(\mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} [\mathbf{U}_S \mathbf{V}_T] \right) \cdot \left(\mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} [\mathbf{U}_{S'} \mathbf{V}_{T'}] \right) \right| = 0, \\ \text{if } S \Delta T \Delta S' \Delta T' = \emptyset : & \left| \mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} [\mathbf{U}_S \mathbf{V}_T \mathbf{U}_{S'} \mathbf{V}_{T'}] - \left(\mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} [\mathbf{U}_S \mathbf{V}_T] \right) \cdot \left(\mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} [\mathbf{U}_{S'} \mathbf{V}_{T'}] \right) \right| \leq \delta. \end{aligned}$$

Here, $S \Delta T \Delta S' \Delta T'$ is the symmetric difference of the sets S, T, S', T' , equivalently, the set of all $i \in [n]$ which appear an odd number of times in the multiset $S \sqcup T \sqcup S' \sqcup T'$.

In particular, one may take $n_0 = \frac{d^{O(d)}}{\delta^2}$.

In order to prove Lemma A.9, we need the following lemma.

Lemma A.10. *For $\mathbf{m} \sim \gamma_{n_0}$,*

$$\mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}'} \left[\left(\mathbb{E}_{\mathbf{m}} [\langle \mathbf{m}, \tilde{\mathbf{a}} \rangle \langle \mathbf{m}, \tilde{\mathbf{b}} \rangle \langle \mathbf{m}, \tilde{\mathbf{a}}' \rangle \langle \mathbf{m}, \tilde{\mathbf{b}}' \rangle] - \mathbb{E}_{\mathbf{m}} [\langle \mathbf{m}, \tilde{\mathbf{a}} \rangle \langle \mathbf{m}, \tilde{\mathbf{b}} \rangle] \cdot \mathbb{E}_{\mathbf{m}} [\langle \mathbf{m}, \tilde{\mathbf{a}}' \rangle \langle \mathbf{m}, \tilde{\mathbf{b}}' \rangle] \right)^2 \right] \leq O\left(\frac{1}{n_0}\right)$$

and

$$\mathbb{E}_{\mathbf{a}, \mathbf{a}'} \left[\left(\mathbb{E}_{\mathbf{m}} [\langle \mathbf{m}, \tilde{\mathbf{a}} \rangle \langle \mathbf{m}, \tilde{\mathbf{a}}' \rangle] - \mathbb{E}_{\mathbf{m}} [\langle \mathbf{m}, \tilde{\mathbf{a}} \rangle] \cdot \mathbb{E}_{\mathbf{m}} [\langle \mathbf{m}, \tilde{\mathbf{a}}' \rangle] \right)^2 \right] \leq O\left(\frac{1}{n_0}\right).$$

Proof. To prove the first part of the lemma, consider the quantity

$$\begin{aligned} T(\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}') &:= \mathbb{E}_{\mathbf{m}} [\langle \mathbf{m}, \tilde{\mathbf{a}} \rangle \langle \mathbf{m}, \tilde{\mathbf{b}} \rangle \langle \mathbf{m}, \tilde{\mathbf{a}}' \rangle \langle \mathbf{m}, \tilde{\mathbf{b}}' \rangle] - \mathbb{E}_{\mathbf{m}} [\langle \mathbf{m}, \tilde{\mathbf{a}} \rangle \langle \mathbf{m}, \tilde{\mathbf{b}} \rangle] \cdot \mathbb{E}_{\mathbf{m}} [\langle \mathbf{m}, \tilde{\mathbf{a}}' \rangle \langle \mathbf{m}, \tilde{\mathbf{b}}' \rangle] \\ &= \langle \tilde{\mathbf{a}}, \tilde{\mathbf{b}} \rangle \langle \tilde{\mathbf{a}}', \tilde{\mathbf{b}}' \rangle + \langle \tilde{\mathbf{a}}, \tilde{\mathbf{a}}' \rangle \langle \tilde{\mathbf{b}}, \tilde{\mathbf{b}}' \rangle + \langle \tilde{\mathbf{a}}, \tilde{\mathbf{b}}' \rangle \langle \tilde{\mathbf{a}}', \tilde{\mathbf{b}} \rangle - \langle \tilde{\mathbf{a}}, \tilde{\mathbf{b}} \rangle \langle \tilde{\mathbf{a}}', \tilde{\mathbf{b}}' \rangle \\ &= \langle \tilde{\mathbf{a}}, \tilde{\mathbf{a}}' \rangle \langle \tilde{\mathbf{b}}, \tilde{\mathbf{b}}' \rangle + \langle \tilde{\mathbf{a}}, \tilde{\mathbf{b}}' \rangle \langle \tilde{\mathbf{a}}', \tilde{\mathbf{b}} \rangle. \end{aligned}$$

where we use that for any $j \in [n_0]$, it holds that $\mathbb{E}_m[m_j^4] = 3$ and $\mathbb{E}_m[m_j^2] = 1$. Thus,

$$\begin{aligned} \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}'} [T(\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}')^2] &= \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}'} \left[\left[\langle \tilde{\mathbf{a}}, \tilde{\mathbf{a}}' \rangle \langle \tilde{\mathbf{b}}, \tilde{\mathbf{b}}' \rangle + \langle \tilde{\mathbf{a}}, \tilde{\mathbf{b}}' \rangle \langle \tilde{\mathbf{a}}', \tilde{\mathbf{b}} \rangle \right]^2 \right] \\ &\leq 2 \cdot \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}'} \left[\langle \tilde{\mathbf{a}}, \tilde{\mathbf{a}}' \rangle^2 \langle \tilde{\mathbf{b}}, \tilde{\mathbf{b}}' \rangle^2 \right] + 2 \cdot \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}'} \left[\langle \tilde{\mathbf{a}}, \tilde{\mathbf{b}}' \rangle^2 \langle \tilde{\mathbf{a}}', \tilde{\mathbf{b}} \rangle^2 \right] \\ &\leq O\left(\frac{1}{n_0}\right), \end{aligned}$$

where the last step follows by two applications of [Lemma A.1](#) (with $d = 4$). This completes the proof of the first part of the lemma. The second part of the lemma similarly follows from [Lemma A.1](#) (with $d = 2$) along with the fact that $\mathbb{E}_m[\langle \mathbf{m}, \tilde{\mathbf{a}} \rangle] = 0$. \square

Proof of Lemma A.9. Let $\mathbf{1}(E)$ denote the 0/1 indicator function of an event E . We have that

$$\begin{aligned} \mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} [U_S V_T U_{S'} V_{T'}] &= \mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} \left[\prod_{i \in S \cup T \cup S' \cup T'} U_i^{\mathbf{1}(i \in S)} V_i^{\mathbf{1}(i \in T)} U_i^{\mathbf{1}(i \in S')} V_i^{\mathbf{1}(i \in T')} \right] \\ &= \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} \left[\prod_{i \in S \cup T \cup S' \cup T'} \mathbb{E}_{\mathbf{m}_i} \left[U_i^{\mathbf{1}(i \in S)} V_i^{\mathbf{1}(i \in T)} U_i^{\mathbf{1}(i \in S')} V_i^{\mathbf{1}(i \in T')} \right] \right]. \end{aligned} \quad (21)$$

On the other hand, we have that

$$\mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} [U_S V_T] = \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbb{E}_{\mathbf{M}} \left[\prod_{i \in S \cup T} U_i^{\mathbf{1}(i \in S)} V_i^{\mathbf{1}(i \in T)} \right] = \mathbb{E}_{\mathbf{a}, \mathbf{b}} \left[\prod_{i \in S \cup T} \mathbb{E}_{\mathbf{m}_i} \left[U_i^{\mathbf{1}(i \in S)} V_i^{\mathbf{1}(i \in T)} \right] \right], \quad (22)$$

$$\text{and similarly, } \mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} [U_{S'} V_{T'}] = \mathbb{E}_{\mathbf{a}', \mathbf{b}'} \left[\prod_{i \in S' \cup T'} \mathbb{E}_{\mathbf{m}_i} \left[U_i^{\mathbf{1}(i \in S')} V_i^{\mathbf{1}(i \in T')} \right] \right]. \quad (23)$$

If there exists $i \in S \cup T \cup S' \cup T'$ that appears in an odd number of S, T, S' and T' , then it can be seen that the expectation in [Equation \(21\)](#) is equal to 0, and that at least one of the expectations in [Equations \(22\)](#) and [\(23\)](#) is equal to 0. This already handles the case that $S \Delta T \Delta S' \Delta T' \neq \emptyset$.

Henceforth, we assume that each $i \in S \cup T \cup S' \cup T'$ appears in an even number of S, T, S' and T' . Assume for ease of notation that $S \cup T \cup S' \cup T' \subseteq [4d]$. Define

$$g_i(\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}') := \mathbb{E}_{\mathbf{m}_i} \left[U_i^{\mathbf{1}(i \in S)} V_i^{\mathbf{1}(i \in T)} U_i^{\mathbf{1}(i \in S')} V_i^{\mathbf{1}(i \in T')} \right] \quad (24)$$

$$h_i(\mathbf{a}, \mathbf{b}) := \mathbb{E}_{\mathbf{m}_i} \left[U_i^{\mathbf{1}(i \in S)} V_i^{\mathbf{1}(i \in T)} \right]. \quad (25)$$

$$h'_i(\mathbf{a}', \mathbf{b}') := \mathbb{E}_{\mathbf{m}_i} \left[U_i^{\mathbf{1}(i \in S')} V_i^{\mathbf{1}(i \in T')} \right]. \quad (26)$$

Combining [Equations \(21\)](#) to [\(23\)](#) along with the definitions in [Equations \(24\)](#) to [\(26\)](#), we get

$$\begin{aligned} &\left| \mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} [U_S V_T U_{S'} V_{T'}] - \mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} [U_S V_T] \cdot \mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} [U_{S'} V_{T'}] \right| \\ &= \left| \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} \left[\prod_{i=1}^{4d} g_i(\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}') - \prod_{i=1}^{4d} h_i(\mathbf{a}, \mathbf{b}) \cdot h'_i(\mathbf{a}', \mathbf{b}') \right] \right| \\ &= \left| \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} \left[\sum_{j=1}^{4d} \left[\prod_{i=1}^{j-1} h_i(\mathbf{a}, \mathbf{b}) \cdot h'_i(\mathbf{a}', \mathbf{b}') \prod_{i=j}^{4d} g_i(\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}') - \prod_{i=1}^j h_i(\mathbf{a}, \mathbf{b}) \cdot h'_i(\mathbf{a}', \mathbf{b}') \prod_{i=j+1}^{4d} g_i(\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}') \right] \right] \right| \end{aligned}$$

$$\begin{aligned} &\leq \sum_{j=1}^{4d} \left| \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} \left[\prod_{i=1}^{j-1} h_i(\mathbf{a}, \mathbf{b}) \cdot h'_i(\mathbf{a}', \mathbf{b}') \prod_{i=j+1}^{4d} g_i(\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}') \cdot [g_j(\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}') - h_j(\mathbf{a}, \mathbf{b}) \cdot h'_j(\mathbf{a}', \mathbf{b}')] \right] \right| \\ &\leq 4 \cdot d \cdot \sqrt{\tau \cdot \kappa}, \end{aligned}$$

where the last inequality follows from the Cauchy-Schwarz inequality with

$$\begin{aligned} \tau &:= \max_{j \in [4d]} \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} \left[\prod_{i=1}^{j-1} h_i(\mathbf{a}, \mathbf{b})^2 \cdot h_i(\mathbf{a}', \mathbf{b}')^2 \prod_{i=j+1}^{4d} g_i(\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}')^2 \right] \\ \kappa &:= \max_{j \in [4d]} \mathbb{E}_{\mathbf{a}, \mathbf{b}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} [g_j(\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}') - h_j(\mathbf{a}, \mathbf{b}) \cdot h'_j(\mathbf{a}', \mathbf{b}')]^2 \end{aligned}$$

Lemma A.10 implies that $\kappa \leq O(1/n_0)$. We now show that $\tau \leq 2^{O(d)}$. Note that for any $i \in [n_0]$, it holds that,

$$\begin{aligned} h_i(\mathbf{a}, \mathbf{b}) &= \begin{cases} \langle \tilde{\mathbf{a}}, \tilde{\mathbf{b}} \rangle & \text{if } i \in S \text{ and } i \in T \\ 1 & \text{if } i \notin S \text{ and } i \notin T \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad h'_i(\mathbf{a}', \mathbf{b}') = \begin{cases} \langle \tilde{\mathbf{a}'}, \tilde{\mathbf{b}'} \rangle & \text{if } i \in S' \text{ and } i \in T' \\ 1 & \text{if } i \notin S' \text{ and } i \notin T' \\ 0 & \text{otherwise} \end{cases} \\ g_i(\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}') &= \begin{cases} \langle \tilde{\mathbf{a}}, \tilde{\mathbf{b}} \rangle \langle \tilde{\mathbf{a}'}, \tilde{\mathbf{b}'} \rangle + \langle \tilde{\mathbf{a}}, \tilde{\mathbf{a}'} \rangle \langle \tilde{\mathbf{b}}, \tilde{\mathbf{b}'} \rangle + \langle \tilde{\mathbf{a}}, \tilde{\mathbf{b}'} \rangle \langle \tilde{\mathbf{a}'}, \tilde{\mathbf{b}} \rangle & \text{if } i \in S \cap T \cap S' \cap T' \\ \langle \tilde{\mathbf{a}}, \tilde{\mathbf{b}} \rangle & \text{if } i \in S \cap T \text{ and } i \notin S' \cup T' \\ \langle \tilde{\mathbf{a}}, \tilde{\mathbf{a}'} \rangle & \text{if } i \in S \cap S' \text{ and } i \notin T \cup T' \\ \langle \tilde{\mathbf{a}}, \tilde{\mathbf{b}'} \rangle & \text{if } i \in S \cap T' \text{ and } i \notin S' \cup T \\ \langle \tilde{\mathbf{a}'}, \tilde{\mathbf{b}} \rangle & \text{if } i \in S' \cap T \text{ and } i \notin S \cup T' \\ \langle \tilde{\mathbf{a}'}, \tilde{\mathbf{b}'} \rangle & \text{if } i \in S' \cap T' \text{ and } i \notin S \cup T \\ \langle \tilde{\mathbf{b}}, \tilde{\mathbf{b}'} \rangle & \text{if } i \in T \cap T' \text{ and } i \notin S \cup S' \\ 1 & \text{otherwise} \end{cases} \end{aligned}$$

Thus, if we expand out a single term $\prod_{i=1}^{j-1} h_i(\mathbf{a}, \mathbf{b})^2 \cdot h_i(\mathbf{a}', \mathbf{b}')^2 \prod_{i=j+1}^{4d} g_i(\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}')^2$, we get at most 3^{8d} terms (since each g_i can increase the number of terms by a factor of at most 3). Each of these terms is the expectation of the product of inner product of some correlated Gaussian vectors. We have from **Lemma A.1** that each such term is at most $1 + \delta$ and thus $\tau \leq 2^{O(d)}$. Thus, for an explicit choice of n_0 that is upper bounded by $d^{O(d)}/\delta^2$, we get that $4d\sqrt{\tau\kappa} \leq \delta$, which concludes the proof of the lemma. \square

A.3 Mean & Variance Bounds for Multilinear Polynomials

We are now ready to prove **Lemma 3.3**. Recall again that,

$$F(M) = \mathbb{E}_{\mathbf{a}, \mathbf{b}} [A(\mathbf{U}) \cdot B(\mathbf{V})] \quad \text{where, } \mathbf{U} = \frac{M\mathbf{a}}{\|\mathbf{a}\|_2} \text{ and } \mathbf{V} = \frac{M\mathbf{b}}{\|\mathbf{b}\|_2}.$$

We wish to bound the mean and variance of $F(M)$. These proofs work by considering the Hermite expansions of A and B given by,

$$A(\mathbf{X}) = \sum_{S \subseteq [n]} \hat{A}_S \mathbf{X}_S \quad \text{and} \quad B(\mathbf{X}) = \sum_{T \subseteq [n]} \hat{B}_T \mathbf{Y}_T.$$

The basic definitions and facts related to Hermite polynomials were given in **Section 2**.

Proof of Lemma 3.3. We start out by proving the bound on $\left| \mathbb{E}_{\mathbf{M}} F(\mathbf{M}) - \langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}} \right|$. To this end, we will use Lemma A.8 with parameters d and δ . Thus, for a choice of $n_0 = d^{O(d)}/\delta^2$, we have that,

$$\begin{aligned}
& \left| \mathbb{E}_{\mathbf{M}} F(\mathbf{M}) - \langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| \\
&= \left| \mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} [A(\mathbf{U}) \cdot B(\mathbf{V})] - \mathbb{E}_{\mathbf{X}, \mathbf{Y} \sim \mathcal{G}_\rho^{\otimes n}} [A(\mathbf{X}) \cdot B(\mathbf{Y})] \right| \\
&= \left| \sum_{S, T \subseteq [n]} \widehat{A}_S \widehat{B}_T \cdot \left(\mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} [U_S \cdot V_T] - \mathbb{E}_{\mathbf{X}, \mathbf{Y} \sim \mathcal{G}_\rho^{\otimes n}} [X_S \cdot Y_T] \right) \right| \\
&= \left| \sum_{S \subseteq [n]} \widehat{A}_S \widehat{B}_S \cdot \left(\mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} [U_S \cdot V_S] - \rho^{|S|} \right) \right| \quad \dots \text{(terms corresponding to } S \neq T \text{ are 0.)} \\
&\leq \sum_{S \subseteq [n]} \left| \widehat{A}_S \widehat{B}_S \right| \cdot \delta \quad \dots \dots \text{(using Lemma A.8)} \\
&\leq \|A\|_2 \cdot \|B\|_2 \cdot \delta \quad \dots \dots \text{(Cauchy-Schwarz inequality)} \\
&\leq \delta \quad \dots \dots (\|A\|_2, \|B\|_2 \leq 1)
\end{aligned}$$

We now move to proving the bound on $\text{Var}_{\mathbf{M}}(F(\mathbf{M}))$. To this end, we will use Lemma A.9 with parameters d and $\delta/9^d$. Thus, for a choice of $n_0 = d^{O(d)}/\delta^2$, we have that,

$$\begin{aligned}
& \mathbb{E}_{\mathbf{M}} \left(\mathbb{E}_{\mathbf{a}, \mathbf{b}} A(\mathbf{U}) \cdot B(\mathbf{V}) \right)^2 - \left(\mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} A(\mathbf{U}) \cdot B(\mathbf{V}) \right)^2 \\
&= \left| \mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}'} [A(\mathbf{U})B(\mathbf{V})A(\mathbf{U}')B(\mathbf{V}')] - \left(\mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} [A(\mathbf{U})B(\mathbf{V})] \right) \cdot \left(\mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} [A(\mathbf{U}')B(\mathbf{V}')] \right) \right| \\
&\leq \sum_{\substack{S, T \subseteq [n] \\ S', T' \subseteq [n]}} \left| \widehat{A}_S \widehat{B}_T \widehat{A}_{S'} \widehat{B}_{T'} \right| \cdot \left| \mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}, \mathbf{a}', \mathbf{b}'} [U_S V_T U_{S'} V_{T'}] - \left(\mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}, \mathbf{b}} [U_S V_T] \right) \cdot \left(\mathbb{E}_{\mathbf{M}} \mathbb{E}_{\mathbf{a}', \mathbf{b}'} [U_{S'} V_{T'}] \right) \right| \\
&\leq \frac{\delta}{9^d} \cdot \sum_{\substack{S, T, S', T' \subseteq [n] \\ S \Delta T \Delta S' \Delta T' = \emptyset}} \left| \widehat{A}_S \widehat{B}_T \widehat{A}_{S'} \widehat{B}_{T'} \right|.
\end{aligned}$$

To finish the proof, we will show that,

$$\sum_{\substack{S, T, S', T' \subseteq [n] \\ S \Delta T \Delta S' \Delta T' = \emptyset}} \left| \widehat{A}_S \widehat{B}_T \widehat{A}_{S'} \widehat{B}_{T'} \right| \leq 9^d \cdot \|A\|_2^2 \cdot \|B\|_2^2.$$

Define functions $f : \{1, -1\}^n \rightarrow \mathbb{R}$, $g : \{1, -1\}^n \rightarrow \mathbb{R}$ over the boolean hypercube as,

$$f(x) = \sum_{\substack{S \subseteq [n] \\ |S| \leq d}} \widehat{A}_S \mathcal{X}_S(x) \quad \text{and} \quad g(x) = \sum_{\substack{S \subseteq [n] \\ |S| \leq d}} \widehat{B}_S \mathcal{X}_S(x).$$

Hypercontractivity bounds [Wol07] for degree- d polynomials over the boolean hypercube imply that,

$$\mathbb{E}_x [f(x)^4] \leq 9^d \left(\mathbb{E}_x [f(x)^2] \right)^2 \quad \text{and} \quad \mathbb{E}_x [g(x)^4] \leq 9^d \left(\mathbb{E}_x [g(x)^2] \right)^2.$$

We now finish the proof as follows,

$$\sum_{\substack{S, T, S', T' \subseteq [n] \\ S \Delta T \Delta S' \Delta T' = \emptyset}} \left| \widehat{A}_S \widehat{B}_T \widehat{A}_{S'} \widehat{B}_{T'} \right| = \mathbb{E}_x [f(x)^2 g(x)^2]$$

$$\begin{aligned}
&\leq \left(\mathbb{E}_x [f(x)^4] \right)^{1/2} \cdot \left(\mathbb{E}_x [g(x)^4] \right)^{1/2} \quad \dots \text{(Cauchy-Schwarz)} \\
&\leq 9^d \cdot \left(\mathbb{E}_x [f(x)^2] \right) \cdot \left(\mathbb{E}_x [g(x)^2] \right) \quad \dots \text{(Hypercontractivity)} \\
&= 9^d \cdot \|A\|_2^2 \cdot \|B\|_2^2.
\end{aligned}$$

Thus, overall we get that, $\text{Var}_M(F(M)) \leq \delta$.

This completes the proof of [Lemma 3.3](#) for an explicit choice of $n_0 \leq d^{O(d)}/\delta^2$. \square

B Proof of Low-Degree Multilinear Transformation Lemma

The goal of this section is to prove [Lemma 4.1](#), which follows immediately by putting together the following two lemmas. The first lemma transforms general functions to low-degree polynomials and second lemma subsequently transforms it to multilinear polynomials.

Lemma B.1 (Low Degree Transformation). *Given parameters $\rho \in [0, 1]$, $\delta > 0$, $k \in \mathbb{N}$, there exists an explicit $d = d(\rho, k, \delta)$ such that the following holds:*

Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^k$ and $B : \mathbb{R}^n \rightarrow \mathbb{R}^k$, such that, for any $j \in [k] : \text{Var}(A_j), \text{Var}(B_j) \leq 1$.

Then, there exist functions $\tilde{A} : \mathbb{R}^n \rightarrow \mathbb{R}^k$ and $\tilde{B} : \mathbb{R}^n \rightarrow \mathbb{R}^k$ such that the following hold.

1. \tilde{A} and \tilde{B} have degree at most d .
2. For any $i \in [k]$, it holds that $\text{Var}(\tilde{A}_i) \leq \text{Var}(A_i) \leq 1$ and $\text{Var}(\tilde{B}_i) \leq \text{Var}(B_i) \leq 1$.
3. $\|\mathcal{R}(\tilde{A}) - \tilde{A}\|_2 \leq \|\mathcal{R}(A) - A\|_2 + \delta$ and $\|\mathcal{R}(\tilde{B}) - \tilde{B}\|_2 \leq \|\mathcal{R}(B) - B\|_2 + \delta$
4. For every $i, j \in [k]$,

$$\left| \langle \tilde{A}_i, \tilde{B}_j \rangle_{\mathcal{G}_\rho^{\otimes n}} - \langle A_i, B_j \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| \leq \frac{\delta}{\sqrt{k}}$$

In particular, one may take $d = O\left(\frac{\sqrt{k} \log^2(k/\delta)}{\delta(1-\rho)}\right)$.

Lemma B.2 (Multi-linear Transformation). *Given parameters $\rho \in [0, 1]$, $\delta > 0$, $d, k \in \mathbb{Z}_{\geq 0}$, there exists an explicit $t = t(k, d, \delta)$ such that the following holds:*

Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^k$ and $B : \mathbb{R}^n \rightarrow \mathbb{R}^k$ be degree- d polynomials, such that, for any $j \in [k] : \text{Var}(A_j), \text{Var}(B_j) \leq 1$.

Then, there exist functions $\tilde{A} : \mathbb{R}^{nt} \rightarrow \mathbb{R}^k$ and $\tilde{B} : \mathbb{R}^{nt} \rightarrow \mathbb{R}^k$ such that the following hold:

1. \tilde{A} and \tilde{B} are multilinear with degree at most d .
2. For any $i \in [k]$, it holds that $\text{Var}(\tilde{A}_i) \leq \text{Var}(A_i) \leq 1$ and $\text{Var}(\tilde{B}_i) \leq \text{Var}(B_i) \leq 1$.
3. $\|\mathcal{R}(\tilde{A}) - \tilde{A}\|_2 \leq \|\mathcal{R}(A) - A\|_2 + \delta$ and $\|\mathcal{R}(\tilde{B}) - \tilde{B}\|_2 \leq \|\mathcal{R}(B) - B\|_2 + \delta$
4. For every $i, j \in [k]$,

$$\left| \langle \tilde{A}_i, \tilde{B}_j \rangle_{\mathcal{G}_\rho^{\otimes nt}} - \langle A_i, B_j \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| \leq \frac{\delta}{\sqrt{k}}$$

In particular, one may take $t = O\left(\frac{kd^2}{\delta^2}\right)$.

Simple Proposition for Rounding

Before getting to the proofs of the above lemmas, we present a simple proposition that will be useful. It says that if we have two strategies which are close in ℓ_2 -distance, and one of them is *close* to the simplex Δ_k , then so is the other. The proof follows by a straightforward triangle inequality.

Proposition B.3. For $A : \mathbb{R}^n \rightarrow \mathbb{R}^k$ and $\tilde{A} : \mathbb{R}^n \rightarrow \mathbb{R}^k$ such that $\|A\|_2, \|\tilde{A}\|_2 \leq 1$, it holds that,

$$\|\mathcal{R}(\tilde{A}) - \tilde{A}\|_2 \leq \|\mathcal{R}(A) - A\|_2 + \|A - \tilde{A}\|_2.$$

B.1 Transformation to Low-Degree

The key idea behind [Lemma B.1](#) is quite standard, that applying a “small” amount of noise (via the Ornstein-Uhlenbeck operator) to a pair of functions doesn’t hurt their correlation “significantly”. In particular, we have the following lemma.

Lemma B.4. Let $P, Q \in L^2(\mathbb{R}^n, \gamma_n)$ and $\varepsilon > 0$. There exists $\nu = \nu(\rho, \varepsilon)$ such that,

$$\left| \langle P, Q \rangle_{\mathcal{G}_\rho^{\otimes n}} - \langle U_\nu P, U_\nu Q \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| \leq \varepsilon \cdot \sqrt{\text{Var}[P] \text{Var}[Q]}$$

In particular, one may take $\nu := (1 - \varepsilon)^{\log \rho / (\log \varepsilon + \log \rho)}$, or even $\nu := 1 - C \frac{(1-\rho)\varepsilon}{\log(1/\varepsilon)}$ for some constant $C > 0$.

Proof. Consider the Hermite expansions of P and Q . That is,

$$P(\mathbf{X}) = \sum_{\sigma \in \mathbb{Z}_{\geq 0}^n} \hat{P}(\sigma) H_\sigma(\mathbf{X}) \quad \text{and} \quad Q(\mathbf{Y}) = \sum_{\sigma \in \mathbb{Z}_{\geq 0}^n} \hat{Q}(\sigma) H_\sigma(\mathbf{Y}).$$

Using properties of Hermite polynomials, namely, $U_\nu H_\sigma = \nu^{|\sigma|} H_\sigma$, we get that,

$$U_\nu P(\mathbf{X}) = \sum_{\sigma \in \mathbb{Z}_{\geq 0}^n} \nu^{|\sigma|} \hat{P}(\sigma) H_\sigma(\mathbf{X}) \quad \text{and} \quad U_\nu Q(\mathbf{Y}) = \sum_{\sigma \in \mathbb{Z}_{\geq 0}^n} \nu^{|\sigma|} \hat{Q}(\sigma) H_\sigma(\mathbf{Y}).$$

Our choice of ν was to ensure that $\rho^d (1 - \nu^{2d}) \leq \varepsilon$ for all $d \in \mathbb{N}$. Thus, we get that,

$$\begin{aligned} & \left| \langle P, Q \rangle_{\mathcal{G}_\rho^{\otimes n}} - \langle U_\nu P, U_\nu Q \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| \\ &= \left| \sum_{\sigma \neq \mathbf{0}} \rho^{|\sigma|} \cdot \hat{P}(\sigma) \hat{Q}(\sigma) \cdot (1 - \nu^{2|\sigma|}) \right| \\ &\leq \sum_{\sigma \neq \mathbf{0}} \left| \hat{P}(\sigma) \hat{Q}(\sigma) \right| \cdot \rho^{|\sigma|} (1 - \nu^{2|\sigma|}) \\ &\leq \varepsilon \cdot \sum_{\sigma \neq \mathbf{0}} \left| \hat{P}(\sigma) \hat{Q}(\sigma) \right| \quad \dots \text{(since, } \rho^d (1 - \nu^{2d}) \leq \varepsilon \text{ for all } d \in \mathbb{N}) \\ &\leq \varepsilon \cdot \sqrt{\text{Var}[P] \text{Var}[Q]} \quad \dots \text{(Cauchy-Schwarz inequality)} \end{aligned}$$

□

The above lemma transforms general functions into functions which are concentrated on low-degree. Thus, to complete the proof of [Lemma B.1](#), we consider the definition of *low-degree truncation*.

Definition B.5 (Low-degree truncation). Let $A \in L^2(\mathbb{R}^n, \gamma_n)$ is given by the Hermite expansion $A(\mathbf{X}) = \sum_{\sigma \in \mathbb{Z}_{\geq 0}^n} \hat{A}_\sigma H_\sigma(\mathbf{X})$. The degree- d truncation of A is defined as the function $A^{\leq d} \in L^2(\mathbb{R}^n, \gamma_n)$ given by

$$A^{\leq d}(\mathbf{X}) := \sum_{\substack{\sigma \in \mathbb{Z}_{\geq 0}^n \\ |\sigma| \leq d}} \hat{A}_\sigma H_\sigma(\mathbf{X}).$$

That is, $A^{\leq d}$ is obtained by retaining only the terms with degree at most d in the Hermite expansion of A , where recall that for $\sigma \in \mathbb{Z}_{\geq 0}^n$, its degree is defined as $|\sigma| = \sum_{i=1}^n \sigma_i$.

For convenience, define $A^{> d} := A - A^{\leq d}$. Also, for vector valued functions A , we define $A^{\leq d}$ as the function obtained by applying the above low-degree truncation on each coordinate.

Proof of Lemma B.1. We obtain \tilde{A} and \tilde{B} by first applying some suitable amount of noise to the functions such that the functions have decaying Hermite tails and then truncating the Hermite coefficients corresponding to terms larger than degree d .

In particular, given parameter δ , we first choose ε and ν in Lemma B.4, such that $\varepsilon = \frac{\delta}{2\sqrt{k}}$ and then $\nu = 1 - C \frac{(1-\rho)\varepsilon}{\log(1/\varepsilon)}$ as required. We choose d to be large enough such that $\nu^{2d} \leq \frac{\delta}{4\sqrt{k}}$, that is, $d = O\left(\frac{\log(k/\delta)}{\log(1/\nu)}\right) = O\left(\frac{\sqrt{k} \log^2(k/\delta)}{\delta(1-\rho)}\right)$. Finally, we let $\tilde{A} := (U_\nu A)^{\leq d}$ and $\tilde{B} := (U_\nu B)^{\leq d}$.

We now verify the four properties required of the lemma.

1. By definition, \tilde{A} and \tilde{B} have degree at most d .
2. $\text{Var}(\tilde{A}_i) = \sum_{\substack{\sigma \neq \mathbf{0} \\ |\sigma| \leq d}} \nu^{2|\sigma|} \cdot \hat{A}_i(\sigma)^2 \leq \text{Var}(A_i)$. Similarly, $\text{Var}(\tilde{B}_i) \leq \text{Var}(B_i)$.
3. For convenience, define $\bar{A} := U_\nu A$, and hence $\tilde{A} = \bar{A}^{\leq d}$. Observe that, since Δ_k is a convex body, $\|\mathcal{R}(v) - v\|_2^2$ is a convex function in $v \in \mathbb{R}^k$. Thus, we have that,

$$\begin{aligned}
\|\mathcal{R}(\bar{A}) - \bar{A}\|_2^2 &= \mathbb{E}_{\mathbf{X} \sim \gamma_n} \|\mathcal{R}(\bar{A}(\mathbf{X})) - \bar{A}(\mathbf{X})\|_2^2 \\
&= \mathbb{E}_{\mathbf{X} \sim \gamma_n} \left\| \mathcal{R} \left(\mathbb{E}_{\mathbf{X}' \sim U_\nu(\mathbf{X})} A(\mathbf{X}') \right) - \mathbb{E}_{\mathbf{X}' \sim U_\nu(\mathbf{X})} A(\mathbf{X}') \right\|_2^2 \\
&\leq \mathbb{E}_{\mathbf{X} \sim \gamma_n} \mathbb{E}_{\mathbf{X}' \sim U_\nu(\mathbf{X})} \|\mathcal{R}(A(\mathbf{X}')) - A(\mathbf{X}')\|_2^2 \quad \dots \text{(using convexity of } \|\mathcal{R}(v) - v\|_2^2 \text{)} \\
&= \mathbb{E}_{\mathbf{X}' \sim \gamma_n} \|\mathcal{R}(A(\mathbf{X}')) - A(\mathbf{X}')\|_2^2 \\
&= \|\mathcal{R}(A) - A\|_2^2.
\end{aligned}$$

Next, observe that, $\|\bar{A}^{> d}\|_2^2 = \sum_{|\sigma| > d} \nu^{2|\sigma|} \cdot \|\hat{A}(\sigma)\|_2^2 \leq \nu^{2d} \cdot \sqrt{k} \leq \frac{\delta}{4}$. Thus, we get that,

$$\begin{aligned}
\|\mathcal{R}(\tilde{A}) - \tilde{A}\|_2 &\leq \|\mathcal{R}(\bar{A}) - \bar{A}\|_2 + \|\bar{A} - \tilde{A}\|_2 \quad \dots \text{(Proposition B.3)} \\
&= \|\mathcal{R}(\bar{A}) - \bar{A}\|_2 + \|\bar{A}^{> d}\|_2 \\
&\leq \|\mathcal{R}(A) - A\|_2 + \delta/4.
\end{aligned}$$

Similar argument holds for \tilde{B} .

4. For every $i, j \in [k]$, we simply have from Lemma B.4 that $\left| \langle \bar{A}_i, \bar{B}_j \rangle_{\mathcal{G}_\rho^{\otimes n}} - \langle A_i, B_j \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| \leq \varepsilon = \frac{\delta}{2\sqrt{k}}$. Additionally, since $\|\tilde{A}_i - \bar{A}_i\|_2 \leq \frac{\delta}{4\sqrt{k}}$ and $\|\tilde{B}_j - \bar{B}_j\|_2 \leq \frac{\delta}{4\sqrt{k}}$, we get using Proposition 2.1 that $\left| \langle \tilde{A}_i, \tilde{B}_j \rangle_{\mathcal{G}_\rho^{\otimes n}} - \langle \bar{A}_i, \bar{B}_j \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| \leq \frac{\delta}{2\sqrt{k}}$. We get the desired statement by combining the two above statements. \square

B.2 Transformation to Multi-linear

The key idea behind [Lemma B.2](#) is similar to that of [Lemma B.1](#) in that, we first apply a transformation on our polynomials that makes it concentrated on multilinear terms, while slightly increasing the number of variables. Subsequently, we apply a *multi-linear truncation* defined as follows.

Definition B.6 (Multilinear truncation). *Suppose $A \in L^2(\mathbb{R}^n, \gamma_n)$ is given by the Hermite expansion $A(x) = \sum_{\sigma \in \mathbb{Z}_{\geq 0}^n} \widehat{A}_\sigma H_\sigma(x)$. The multilinear truncation of A is defined as the function $A^{\text{ml}} \in L^2(\mathbb{R}^n, \gamma_n)$ given by*

$$A^{\text{ml}}(x) := \sum_{\sigma \in \{0,1\}^n} \widehat{A}_\sigma H_\sigma(x).$$

That is, A^{ml} is obtained by retaining only the multilinear terms in the Hermite expansion of A .

For convenience, also define $A^{\text{nmml}} := A - A^{\text{ml}}$. Also, for vector valued functions A , we define A^{ml} as the function obtained by applying the above multilinear truncation on each coordinate.

Lemma B.7. *Given parameters $\rho \in [0, 1]$, $\delta > 0$ and $d \in \mathbb{Z}_{\geq 0}$, there exists $t = t(d, \delta)$ such that the following holds:*

Let $A, B \in L^2(\mathbb{R}^n, \gamma_n)$ be degree- d polynomials, such that $\|A\|_2, \|B\|_2 \leq 1$. Define polynomials $\bar{A}, \bar{B} \in L^2(\mathbb{R}^{nt}, \gamma_{nt})$ over variables $\bar{\mathbf{X}} := \{\mathbf{X}_j^{(i)} : (i, j) \in [n] \times [t]\}$ and $\bar{\mathbf{Y}} := \{\mathbf{Y}_j^{(i)} : (i, j) \in [n] \times [t]\}$ respectively, as,

$$\bar{A}(\bar{\mathbf{X}}) := A(\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(n)}) \quad \text{and} \quad \bar{B}(\bar{\mathbf{Y}}) := B(\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(n)})$$

where $\mathbf{X}^{(i)} = (\mathbf{X}_1^{(i)} + \dots + \mathbf{X}_t^{(i)}) / \sqrt{t}$ and $\mathbf{Y}^{(i)} = (\mathbf{Y}_1^{(i)} + \dots + \mathbf{Y}_t^{(i)}) / \sqrt{t}$. Since $(\mathbf{X}^{(i)}, \mathbf{Y}^{(i)})$ is distributed according to \mathcal{G}_ρ , this transformation doesn't change the "structure" of A and B . In particular, it is follows that,

$$\langle \bar{A}, \bar{B} \rangle_{\mathcal{G}_\rho^{\otimes nt}} = \langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}} \quad \text{and} \quad \|\bar{A}\|_2 = \|A\|_2 \quad \text{and} \quad \|\bar{B}\|_2 = \|B\|_2$$

Next, let $\bar{A}^{\text{ml}}, \bar{B}^{\text{ml}} \in L^2(\mathbb{R}^{nt}, \gamma_{nt})$ be the multilinear truncations of \bar{A} and \bar{B} respectively. Then the following hold,

1. \bar{A}^{ml} and \bar{B}^{ml} are multilinear with degree at most d .
2. $\text{Var}(\bar{A}^{\text{ml}}) \leq \text{Var}(A) \leq 1$ and $\text{Var}(\bar{B}^{\text{ml}}) \leq \text{Var}(B) \leq 1$.
3. $\|\bar{A}^{\text{ml}} - \bar{A}\|_2, \|\bar{B}^{\text{ml}} - \bar{B}\|_2 \leq \delta/2$.
4. $\left| \langle \bar{A}^{\text{ml}}, \bar{B}^{\text{ml}} \rangle_{\mathcal{G}_\rho^{\otimes nt}} - \langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}} \right| \leq \delta$.

In particular, one may take $t = O\left(\frac{d^2}{\delta^2}\right)$.

In order to prove [Lemma B.7](#), we will need the following multinomial theorem for Hermite polynomials. It can be proved quite easily using the generating function for Hermite polynomials.

Fact B.8 (Multinomial theorem for Hermite polynomials). *Let $\beta_1, \dots, \beta_t \in \mathbb{R}$ satisfying $\sum_{i=1}^t \beta_i^2 = 1$. Then, for any $d \in \mathbb{N}$, it holds that*

$$H_d(\beta_1 X_1 + \dots + \beta_t X_t) = \sum_{\substack{d_1, \dots, d_t \in \mathbb{Z}_{\geq 0} \\ d_1 + \dots + d_t = d}} \sqrt{\frac{d!}{d_1! \dots d_t!}} \cdot \prod_{i=1}^t \beta_i^{d_i} H_{d_i}(X_i).$$

Proof of Lemma B.7. Before we prove the theorem, we will first understand the effect of the transformation from X to \bar{X} for a univariate Hermite polynomial. Instantiating β_i 's in Fact B.8 with $1/\sqrt{t}$, we get that,

$$H_d\left(\frac{X_1 + \dots + X_t}{\sqrt{t}}\right) = \sum_{\substack{d_1, \dots, d_t \in \mathbb{Z}_{\geq 0} \\ d_1 + \dots + d_t = d}} \sqrt{\frac{d!}{d_1! \dots d_t!}} \cdot \frac{\prod_{i=1}^t H_{d_i}(X_i)}{t^{d/2}}.$$

We will split the terms into multilinear and non-multilinear terms, writing the above as $H_d^{\text{ml}} + H_d^{\text{nmml}}$. Note that there are at most $O\left(\frac{d^2 t^{d-1}}{d!}\right)$ non-multilinear terms (for $t \gg d^2$). Also, note that each coefficient $\frac{1}{t^{d/2}} \cdot \sqrt{\frac{d!}{d_1! \dots d_t!}}$ is at most $\sqrt{\frac{d!}{t^d}}$. Thus, we can bound $\|H_d^{\text{nmml}}\|_2$ as follows,

$$\|H_d^{\text{nmml}}\|_2^2 = \sum_{\substack{d_1, \dots, d_t \in \mathbb{Z}_{\geq 0} \\ d_1 + \dots + d_t = d \\ \exists i \ d_i \geq 2}} \left(\frac{1}{t^{d/2}} \cdot \sqrt{\frac{d!}{d_1! \dots d_t!}} \right)^2 \leq O\left(\frac{d^2 t^{d-1}}{d!}\right) \cdot \frac{d!}{t^d} \leq O\left(\frac{d^2}{t}\right) \quad (27)$$

More generally, if we consider a term $\bar{H}_\sigma(\bar{X}) = H_{\sigma_1}(X^{(1)}) \cdot H_{\sigma_2}(X^{(2)}) \dots H_{\sigma_n}(X^{(n)})$, where each $X^{(i)} = (X_1^{(i)} + \dots + X_t^{(i)})/\sqrt{t}$. Let's write $\bar{H}_\sigma(\bar{X}) = \bar{H}_\sigma^{\text{ml}}(\bar{X}) + \bar{H}_\sigma^{\text{nmml}}(\bar{X})$, that is, separating out the multilinear and non-multilinear terms. Similarly, for any i , let $H_{\sigma_i}(X^{(i)}) = H_{\sigma_i}^{\text{ml}}(X^{(i)}) + H_{\sigma_i}^{\text{nmml}}(X^{(i)})$. We wish to bound $\|\bar{H}_\sigma^{\text{nmml}}\|_2$, which can be done as follows,

$$\begin{aligned} \|\bar{H}_\sigma^{\text{nmml}}\|_2^2 &= \left\| \prod_{i=1}^n (H_{\sigma_i}^{\text{ml}} + H_{\sigma_i}^{\text{nmml}}) - \prod_{i=1}^n H_{\sigma_i}^{\text{ml}} \right\|_2^2 \\ &\leq \prod_{i=1}^n \left(1 + O\left(\frac{\sigma_i^2}{t}\right) \right) - 1 && \text{(from Equation (27))} \\ &\leq O\left(\frac{|\sigma|^2}{t}\right) && \text{(since, } t \gg |\sigma|^2\text{)} \end{aligned}$$

Thus, $\|\bar{H}_\sigma^{\text{nmml}}\|_2^2 < \delta^2/4$. (for $t = \Theta(d^2/\delta^2)$) (28)

We are now ready to prove the parts of Lemma B.7.

1. It holds by definition that \bar{A}^{ml} and \bar{B}^{ml} are multilinear. Also, note that the transformation from A to \bar{A} and finally to \bar{A}^{ml} does not increase the degree. So both \bar{A}^{ml} and \bar{B}^{ml} have degree at most d .
2. It is easy to see that $\text{Var}(\bar{A}) = \text{Var}(A)$. Since \bar{A}^{ml} is obtained by truncating certain Hermite coefficients of \bar{A} , it immediately follows that $\text{Var}(\bar{A}^{\text{ml}}) \leq \text{Var}(\bar{A}) = \text{Var}(A) \leq 1$. Similarly, $\text{Var}(\bar{B}^{\text{ml}}) \leq \text{Var}(B) \leq 1$.
3. Recall that $\bar{A}^{\text{nmml}} = \bar{A} - \bar{A}^{\text{ml}}$. We wish to bound $\|\bar{A}^{\text{nmml}}\|_2^2 \leq \delta^2/4$. Consider the Hermite expansion of A , namely $A(\mathbf{X}) = \sum_{\sigma \in \mathbb{Z}_{\geq 0}^n} \hat{A}(\sigma) \cdot H_\sigma(\mathbf{X})$. Note that, $\bar{A}^{\text{nmml}}(\bar{\mathbf{X}}) = \sum_{\sigma \in \mathbb{Z}_{\geq 0}^n} \hat{A}(\sigma) \cdot \bar{H}_\sigma^{\text{nmml}}(\bar{\mathbf{X}})$, where recall that $\bar{H}_\sigma^{\text{nmml}}$ is the non-multilinear part of $\bar{H}_\sigma(\bar{\mathbf{X}}) = H_{\sigma_1}(X^{(1)}) \cdot H_{\sigma_2}(X^{(2)}) \dots H_{\sigma_n}(X^{(n)})$, where each $X^{(i)} = (X_1^{(i)} + \dots + X_t^{(i)})/\sqrt{t}$.

From Equation 28, we have that for any $\sigma \in \mathbb{Z}_{\geq 0}^n$, it holds that $\|\bar{H}_\sigma^{\text{nmml}}\|_2^2 < \delta^2/4$. And hence we get that,

$$\|\bar{A}^{\text{nmml}}\|_2^2 = \sum_{\sigma} \hat{A}(\sigma)^2 \cdot \|\bar{H}_\sigma^{\text{nmml}}\|_2^2 \leq \sum_{\sigma} \hat{A}(\sigma)^2 \cdot (\delta^2/4) = (\delta^2/4) \|A\|_2^2 \leq (\delta^2/4).$$

Note that, here we use that $\overline{H}_\sigma(\overline{\mathbf{X}})$ are mutually orthogonal for different σ . Similarly, we can also get that $\left\| \overline{B}^{\text{nmml}} \right\|_2^2 \leq \delta^2/4$.

4. Note that we already have,

$$\langle \overline{A}, \overline{B} \rangle_{\mathcal{G}_\rho^{\otimes nt}} = \langle A, B \rangle_{\mathcal{G}_\rho^{\otimes n}} .$$

And combining Part 3 and [Proposition 2.1](#), we immediately get that

$$\left| \langle \overline{A}^{\text{ml}}, \overline{B}^{\text{ml}} \rangle_{\mathcal{G}_\rho^{\otimes nt}} - \langle \overline{A}, \overline{B} \rangle_{\mathcal{G}_\rho^{\otimes nt}} \right| \leq \delta$$

where we use that $\left\| \overline{B}^{\text{ml}} \right\|_2 \leq \left\| \overline{B} \right\|_2 \leq 1$ and $\left\| \overline{A}^{\text{ml}} \right\|_2 \leq \left\| \overline{A} \right\|_2 \leq 1$.

□

Proof of [Lemma B.2](#). We apply the transformation in [Lemma B.7](#), with parameter δ being δ/\sqrt{k} , to each of the k -coordinates of $A : \mathbb{R}^n \rightarrow \mathbb{R}^k$ and $B : \mathbb{R}^n \rightarrow \mathbb{R}^k$ to get polynomials $\tilde{A} : \mathbb{R}^{nt} \rightarrow \mathbb{R}^k$ and $\tilde{B} : \mathbb{R}^n \rightarrow \mathbb{R}^k$. Namely, for any $j \in [k]$, we set $\tilde{A}_j(\overline{\mathbf{X}}) = \overline{A}_j^{\text{ml}}(\overline{\mathbf{X}})$ and $\tilde{B}_j(\overline{\mathbf{Y}}) = \overline{B}_j^{\text{ml}}(\overline{\mathbf{Y}})$ as described in [Lemma B.7](#).

It is easy to see that parts 1, 2, 4 follow immediately from the conditions satisfied in [Lemma B.7](#). For part 3, we have that $\left\| \overline{A}_j^{\text{ml}} - \overline{A}_j \right\|_2 \leq \delta/\sqrt{k}$ for every $j \in [k]$, which implies that $\left\| \overline{A}^{\text{ml}} - \overline{A} \right\|_2 \leq \delta$. Using [Proposition B.3](#), we immediately get that,

$$\left\| \mathcal{R}(\overline{A}^{\text{ml}}) - \overline{A}^{\text{ml}} \right\|_2 \leq \left\| \mathcal{R}(\overline{A}) - \overline{A} \right\|_2 + \delta .$$

Finally, it is a simple observation that $\left\| \mathcal{R}(\overline{A}) - \overline{A} \right\|_2 = \left\| \mathcal{R}(A) - A \right\|_2$, and hence,

$$\left\| \mathcal{R}(\tilde{A}) - \tilde{A} \right\|_2 \leq \left\| \mathcal{R}(A) - A \right\|_2 + \delta .$$

Similarly, $\left\| \mathcal{R}(\tilde{B}) - \tilde{B} \right\|_2 \leq \left\| \mathcal{R}(B) - B \right\|_2 + \delta$. This concludes the proof.

□