THEORY OF COMPUTING, Volume 15 (6), 2019, pp. 1–47
www.theoryofcomputing.org

SPECIAL ISSUE: CCC 2017

Noise Stability is Computable and Approximately Low-Dimensional

Anindya De* Elchanan Mossel† Joe Neeman‡

Received September 13, 2017; Revised July 7, 2019; Published October 11, 2019

Abstract: The notion of Gaussian noise stability plays an important role in hardness of approximation in theoretical computer science as well as in the theory of voting. The Gaussian noise stability of a partition of $\mathbb{R}^n$ is simply the probability that two correlated Gaussian vectors both fall into the same part. In many applications, the goal is to find an optimizer of noise stability among all possible partitions of $\mathbb{R}^n$ to k parts with given Gaussian measures $\mu_1, \ldots, \mu_k$. We call a partition ε-optimal, if its noise stability is optimal up to an additive ε. In this paper, we give a computable function $n(\varepsilon)$ such that an ε-optimal partition exists in $\mathbb{R}^{n(\varepsilon)}$. This result has implications for the computability of certain problems in non-interactive simulation, which are addressed in a subsequent paper.

A conference version of this paper appeared in the Proceedings of the 32nd Computational Complexity Conference, 2017.

*Supported by NSF grant CCF 1926872 (transferred from CCF 1814706). Most of the work done as a faculty at Northwestern University.

†Supported by DMS-1737944 and ONR N00014-16-1-2227, N00014-17-1-2598.

‡Funded by a Deutsche Forschungs-gemeinschaft (DFG, German Research Foundation) under GermanyâĂŹs Excellence Strategy GZ 2047/1, Projekt-ID 390685813 and a Sloan Fellowship.

ACM Classification: G.3, G.1.6

AMS Classification: 60

Key words and phrases: noise stability, Gaussian surface area, computability

 DOI: 10.4086/toc.2019.v015a006

1 Introduction

Isoperimetric theory and noise stability. Isoperimetric problems have been studied in mathematics since antiquity, and they have become central to mathematics and theoretical computer science. Our interest in this work is in a modern version of the isoperimetric problem: that of maximizing Gaussian noise stability. This problem was originally studied by Borell in the context of stochastic processes [3]. Benjamini, Kalai, and Schramm [2] introduced Boolean noise stability (both the notion and the terminology) to computer science: roughly speaking, a Boolean function—that is, a function $\{0,1\}^n \to \{0,1\}$—is *noise sensitive* if slightly rerandomizing the inputs has the effect of completely rerandomizing the output. This notion naturally extends to functions $\{0,1\}^n \to \{0,\dots,k-1\}$, or, equivalently, to partitions of $\{0,1\}^n$ into k parts; this extension and its applications in theoretical computer science are discussed in [17].

Boolean noise stability has a natural Gaussian analogue, in which the Boolean hypercube $\{0,1\}^n$ is replaced by $\mathbb{R}^n$, and the Boolean noise is replaced by Gaussian noise. The *invariance principle* of [25] shows a deep connection between these two notions of noise stability. In particular, it allows one to translate results about Gaussian noise stability into results about Boolean noise stability, which are typically harder. This connection is responsible for many results in hardness of approximation [21]. For more on the applications of noise stability in the theory of computing, see the book by O'Donnell [29].

We equip $\mathbb{R}^n$ with the standard n-dimensional Gaussian measure, denoted by γ_n. In particular, for any $x \in \mathbb{R}^n$, the density $\gamma_n(x)$ is given by

$$\gamma_n(x) = \frac{1}{(2\pi)^{n/2}} \mathrm{e}^{-\|x\|_2^2/2}.$$

The Ornstein-Uhlenbeck operator P_t is defined for $t \in [0,\infty)$ and bounded, measurable $f : \mathbb{R}^n \to \mathbb{R}$ by

$$(P_t f)(x) = \int_{y\in\mathbb{R}^n} f(\mathrm{e}^{-t} \cdot x + \sqrt{1-\mathrm{e}^{-2t}} \cdot y) d\gamma_n(y).$$

We define the (Gaussian) noise stability of f by $\mathsf{Stab}_t(f) := \mathbf{E}[f \cdot P_t f + (1-f)\cdot(1-P_t f)]$, where the expectation is with respect to γ_n. If f denotes the indicator function of a set (call it A), then $\mathsf{Stab}_t(f)$ is the probability that two e^{-t}-correlated Gaussian random variables fall either both in A or both in A^c. In the limit $t \to 0^+$, noise stability is related to the Gaussian surface area of the set A [23].

Halfspaces are the most stable sets. Since Borell's work in the 1980s [3], halfspaces have been known to maximize noise stability subject to a measure constraint. We will state this fact in a slightly convoluted way, in order to more easily generalize it. Let Δ_k denote the standard simplex in $\mathbb{R}^k$ (i. e., the convex hull of the standard unit vectors $\{\mathbf{e}_1,\dots,\mathbf{e}_k\}$). Using $[k]$ to denote $\{1,\dots,k\}$, we observe that any function with range $[k]$ naturally embeds in Δ_k by identifying $i \in [k]$ with $\mathbf{e}_i$. Moreover, we extend the Ornstein-Uhlenbeck operator to act on vector-valued functions in the obvious way: if $f = (f_1,\dots,f_k) : \mathbb{R}^n \to \mathbb{R}^k$ then $P_t f = (P_t f_1,\dots,P_t f_k)$. Finally, we say that $f = (f_1, f_2) : \mathbb{R}^n \to \Delta_2$ is a *halfspace* if there exist $a, b \in \mathbb{R}^n$ such that $f_1(x) = 1_{\{\langle x-a,b\rangle \le 0\}}$ for all $x \in \mathbb{R}^n$. (Here, $\langle v,w\rangle$ denotes the Euclidean inner product $\sum_{i=1}^n v_i \cdot w_i$.)

Theorem 1.1 (Borell). *For any $g : \mathbb{R}^n \to \Delta_2$ and any $t \geq 0$, every halfspace $f : \mathbb{R}^n \to \Delta_2$ with $\mathbf{E}[f] = \mathbf{E}[g]$ satisfies*

$$\mathbf{E}[\langle f, P_t f \rangle] \geq \mathbf{E}[\langle g, P_t g \rangle].$$

Theorem 1.1 has many applications in computational complexity. Most famously, it can be combined with an *invariance principle* [25] in order to prove a number of tight Unique-Games-hardness of approximation results [20].

More parts? It is straightforward to extend the notion of noise stability to partitions with more than two parts. Namely, a partition of $\mathbb{R}^n$ can be described by $f : \mathbb{R}^n \to [k]$. Identifying $[k]$ with $\{\mathbf{e}_1, \dots, \mathbf{e}_k\}$ as above, we define the noise stability of such a partition by $\mathbf{E}[\langle f, P_t f \rangle]$. One may then ask for an analogue of Theorem 1.1 where Δ_2 is replaced by Δ_k for some $k \geq 3$.

Even the three-part version of Theorem 1.1 turns out to be significantly harder than the two-part one. We will try to explain why, by analogy with isoperimetric problems. First of all, the Euclidean, isoperimetric analogue for $k = 3$ is known as the "double bubble" problem. The well-known "double bubble conjecture" states that the minimum total surface area of two bodies separating and enclosing two given (Lebesgue) volumes is achieved by two spheres meeting at $120°$. After being open for more than a century, this problem was settled rather recently [15, 16, 32, 31]. For the Gaussian space, [5] showed that for some small positive constant $c > 0$, the Gaussian surface area of three-partitions is minimized by the "standard simplex partition" as long as the measure of each part is within $1/3 \pm c$.

Since halfspaces both maximize noise stability and minimize Gaussian surface area, and since the standard simplex partition is known to minimize multi-part Gaussian surface area in certain cases, it seems natural to guess that the standard simplex partition also maximizes multi-part noise stability. This was explicitly conjectured in [20], in the special case that all of the parts in the partition have equal measure. However, a somewhat surprising recent result [14] showed that the standard simplex partition fails to maximize multi-part noise stability *unless* all of the parts have equal measure. On the other hand, there is also some support for the conjecture in the equal-measure case: Heilman [13] showed that the conjecture is true if the noise rate t is larger than some explicit function of the ambient dimension n.

Approximate noise stability of multipartitions? In light of the uncertainty about optimal partitioning for $k \geq 3$, one can ask a more modest question. Given $k \geq 3$, $t > 0$, and prescribed measures $v \in \Delta_k$ for the k parts, let $\alpha_n = \alpha_n(k, v)$ be the maximal noise stability that can be obtained in $\mathbb{R}^n$ under these constraints. Since the Gaussian measure is a product measure, α_n is clearly non-decreasing in n. Since it is bounded above by one, it has a limit as $n \to \infty$.

Our main result is that there is a computable $n_0 = n_0(k, t, \varepsilon)$ such that $\alpha_{n_0} \geq \alpha_m - \varepsilon$ for all $m \in \mathbb{N}$. Although the bound on n_0 that we give is not particularly good (in fact, it is not primitive recursive), the key point is that it is computable. As a consequence, up to error ε, the noise-stable partition is also computable. We conclude the introduction by an open question:

Question 1.2. Fix $k \geq 3$ and measure constraints $v \in \Delta^k$. Does there exist n_0 such that $\alpha_{n_0} = \alpha_n$ for all $n \geq n_0$?

Our current techniques are not suitable for addressing the question above.

1.1 Acknowledgments

The authors started this research at the workshop on "Information theory and concentration phenomena" at the IMA, Minneapolis in 2015. The authors would like to thank the organizers of this workshop and the IMA for their hospitality.

2 Main theorem and overview of proof technique

In order to state the main theorem, we first need to recall the notion of a *polynomial threshold function.* A function $f : \mathbb{R}^n \to \{0,1\}$ is said to be a degree-d PTF if there exists a polynomial $p : \mathbb{R}^n \to \mathbb{R}$ of degree-d such that $f(x) = 1$ if and only if $p(x) > 0$. We will need a k-ary generalization of this definition. We note that there are several possible ways to generalize the notion of PTFs to k-ary PTFs and our particular choice is dictated by the convenience of using the relevant results from [9].

Definition 2.1. A function $f : \mathbb{R}^n \to [k]$ is said to be a multivariate PTF if there exist polynomials $p_1, \dots, p_k : \mathbb{R}^n \to \mathbb{R}$ such that

$$f(x) = \begin{cases} j & \text{if } p_j(x) > 0 \text{ and for } i \neq j,\ p_i(x) \leq 0, \\ 1 & \text{otherwise.} \end{cases}$$

In this case, we denote $f = \mathsf{PTF}(p_1, \dots, p_k)$. Further, f is said to be a degree-d multivariate PTF if $p_1, \dots, p_k$ are degree-d polynomials.

We now state the main theorem of this paper. We set the convention that, unless explicitly mentioned otherwise, the underlying distribution is γ_n, the standard n-dimensional Gaussian measure. Likewise, given any measurable function $f : \mathbb{R}^n \to \mathbb{R}^k$, $\mathbf{E}[f]$ denotes its (vector-valued) expectation (with respect to γ_n) and $\mathsf{Var}(f)$ denotes its covariance matrix. For a vector $v \in \mathbb{R}^k$ and a number $p \geq 1$, $\|v\|_p$ denotes the ℓ^p norm, $(\sum_i |v_i|^p)^{1/p}$.

Theorem 2.2. *There exist computable functions $n_0 = n_0(t,k,\varepsilon)$ and $d = d(t,k,\varepsilon)$ such that the following holds for all d, k, $t > 0$ and $\varepsilon > 0$. Let $f : \mathbb{R}^n \to [k]$ such that $\mathbf{E}[f] = \mu \in \mathbb{R}^k$. For $t > 0$ and $\varepsilon > 0$, there is a degree-d PTF $g : \mathbb{R}^{n_0} \to [k]$ such that*

1. $\|\mathbf{E}[f] - \mathbf{E}[g]\|_1 \leq \varepsilon$,
2. $\mathbf{E}[\langle g, P_t g\rangle] \geq \mathbf{E}[\langle f, P_t f\rangle] - \varepsilon$.

We note that while the functions $n_0(\cdot,\cdot,\cdot)$ and $d(\cdot,\cdot,\cdot)$ are computable, n_0 is not primitive recursive.

Note that the above theorem automatically implies that a function g satisfying the above properties can be computed up to some additional error ε. This is because the set of degree-d PTFs on $\mathbb{R}^{n_0}$ has a finite ε-cover (whose elements can be computed).

Remark 2.3. As explained above, we may interpret the expression $\mathbf{E}[\langle f, P_t f\rangle]$ in Theorem 1.1 in terms of Gaussian random variables taking values in $\mathbb{R}^n$. The correlation of these random variables is e^{-t} (which is always positive), but in many applications it is useful for allow for negative correlations also. For example, Theorem 1.1 remains true with negative correlations, but with the inequality reversed.

We do not know how to prove the analogue of Theorem 2.2 in the case of negative correlations, although we do discuss some related results in the follow-up paper [8]. The reason for this failure is that a certain trick which works in the $k = 2$ setting fails for $k \geq 3$: suppose instead of maximizing $\mathbf{E}[\langle f, P_t f\rangle]$, we attempt to minimize $\mathbf{E}[\langle f_1, P_t f_2\rangle]$ where f_1 and f_2 are not required to be the same function. In the setting of Theorem 1.1, the minimizers miraculously satisfy $f_1(x) = f_2(-x)$ for all x, and it follows that f_1 minimizes the noise stability for the *negative* correlation $-\mathrm{e}^{-t}$. As far as we know, this miracle does not occur for $k \geq 3$.

Proof sketch. The proof of Theorem 2.2 consists of the following main steps:

From general partitions to PTF. The first step in the proof is to show that given any $f : \mathbb{R}^n \to [k]$, there is a multivariate PTF $g' : \mathbb{R}^n \to [k]$ which meets the two criteria in Theorem 2.2 and has degree $d = d(t, k, \varepsilon)$, for some computable function $d(t, k, \varepsilon)$. In other words, g' satisfies $\|\mathbf{E}[f] - \mathbf{E}[g']\|_1 \leq \varepsilon$ and $\mathsf{Stab}_t(g') \geq \mathsf{Stab}_t(f') - \varepsilon$. This is done in Section 6. Note that main difference between the desired conclusion of Theorem 2.2 and what is accomplished in this step is that the ambient dimension remains n as opposed to a bounded dimension n_0.

Why is this true? The basic intuition is that if f is noise stable then it should have most of its Hermite expansion weight at low degree. Therefore we should be able to replace f with the PTF where the polynomial is the truncated expansion of f. There are a number of challenges in formalizing this intuition: 1. We cannot rule out that a positive fraction of the weight of f is at high degrees (perhaps as large as n). 2. It is not clear that the PTF obtained this way is noise stable nor that 3. It has the right expected value.

Our analysis proceeds as follows. We would like to construct g' from f by "rounding" $P_\delta f$ for some small $\delta > 0$. The advantage of $P_\delta f$ over f is that $P_\delta f$ is guaranteed to have Hermite coefficients that decay at a certain rate. The rounding of $P_\delta f$ can be performed given some $a \in \mathbb{R}^n$ by considering the function $g_a : \mathbb{R}^n \to [k]$ which takes the value i whenever i is the largest coordinate of $P_\delta f - a$. It is not hard to prove that it is possible to choose a such that $\mathbf{E}[g_a] = \mathbf{E}[f]$; moreover, one can show that this function g_a has better noise stability than f does. The main obstacle is that the function g_a is not a PTF. Unfortunately the Hermite decay of $P_\delta f$ does not translate to Hermite decay of g_a. Instead we use a randomized construction to show that for most a's, g_a has Hermite decay and can therefore be well approximated by a PTF. The analysis of this construction uses the co-area formula [11] and gradient bounds [23] in the Gaussian space and draws on ideas from [26, 22].

From PTF in dimension n to a small PTF of bounded-degree polynomials. Given the function $g' : \mathbb{R}^n \to [k]$ of degree $d = d(t, k, \varepsilon)$, our next goal is to show that it is possible to obtain a PTF g on some $n_0 = n_0(d, k, \varepsilon)$ variables such that (i) $\|\mathbf{E}[g] - \mathbf{E}[g']\|_1 \leq \varepsilon$ and (ii) $|\langle g, P_t g\rangle - \langle g', P_t g'\rangle| \leq \varepsilon$. This part builds on and extends the theory and results of [9]. The key notion introduced in [9] is that of an *eigenregular* polynomial. Namely, a polynomial is said to be δ-eigenregular if for the canonical tensor $\mathcal{A}_p$ associated with the polynomial, the ratio of the maximum singular value to its Frobenius norm is at most δ (the tensor notions are explicitly defined later).

The key advantage of this definition is that as shown in [9], when $\delta \to 0$, the distribution of p (under γ_n) converges to a normal. In other words, eigenregular polynomials obey a central limit theorem. In fact, given k polynomials $p_1, \dots, p_k$ which are δ-eigenregular, they also obey a multidimensional central limit theorem.

The *regularity lemma* from [9] implies that the polynomials $p_1, \dots, p_k$ can be jointly expressed as bounded-size polynomials in eigenregular homogenous polynomials. More precisely, for any $\delta > 0$ there is a collection of "inner" polynomials $\{\mathrm{In}(p_{s,q,\ell})\}$ for $1 \le s \le k$, $1 \le q \le d$ and $1 \le \ell \le \mathrm{num}(s,q)$ (where $\mathrm{num}(\cdot,\cdot)$ is an explicitly defined function), and a collection of "outer" polynomials $\{\mathrm{Out}(p_s)\}$ for $1 \le s \le k$ such that

$$p_s = \mathrm{Out}(p_s)(\{\mathrm{In}(p_{s,q,\ell})\}_{q\in[d],\ell\in[\mathrm{num}(s,q)]}).$$

Moreover, every inner polynomial is δ-eigenregular and homogeneous, and each outer polynomial $\mathrm{Out}(p_s)$ has arity $\mathrm{num}(s) = \sum_{q=1}^{d} \mathrm{num}(s,q)$ bounded by a computable function of d, k, and δ.

In [9], the regularity lemma was used to conclude that the joint distribution of $p_1, \dots, p_k$ can be approximated in a bounded dimension as we can replace each of the inner polynomials by a one dimensional Gaussian. For our application things are more delicate, as we are not only interested in the joint distribution of $p_1, \dots, p_k$ but also in the noise stability of $p_1, \dots, p_k$. For this reason it is important for us to maintain the degrees of the inner polynomials (each of which is homogenous) and not replace them with Gaussians.

A small PTF representation. In the final step of the proof, we maintain $\mathrm{Out}(p_s)$ and show how the polynomials $\{\mathrm{In}(p_{s,q,\ell})\}$ can be replaced by a collection of polynomials $\{\mathrm{In}(r_{s,q,\ell})\}$ in bounded dimensions (i. e., bounded in terms of d, k and ε). The fact that a collection of homogenous 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. Once such polynomials are found, it is not hard to construct eigenregular polynomials from them by averaging the polynomials over independent copies of random variables.

3 Applications

Given the wide applicability of Borell's isoperimetric result to combinatorics and theoretical computer science, we believe that Theorem 2.2 will also be widely applicable. We will now point out some applications of this theorem. First, by combining Theorem 2.2 with the invariance principle [25], we derive a weak k-ary analogue of "Majority is Stablest." The analogue of the Ornstein-Uhlenbeck operator is the so-called *Bonami operator* (also referred to as the Bonami-Beckner operator) defined as follows: for $\rho \in [-1,1]$ and $x \in \{-1,1\}^n$, let $\mathcal{D}_\rho(x)$ be the product distribution over $\{-1,1\}^n$ such that for $y \sim \mathcal{D}_\rho(x)$, for all $1 \le i \le n$, $\mathbf{E}[x_i y_i] = \rho$. Then, for any $f : \{-1,1\}^n \to \mathbb{R}^k$,

$$T_\rho f(x) = \mathbf{E}_{y\sim D_\rho(x)}[f(y)].$$

Likewise, for any $i \in [n]$ and $z \in \{-1,1\}^{n-1}$, let $f_{z,-i} : \{-1,1\} \to \mathbb{R}^k$ denote the function obtained by restricting all but the i-th coordinate to z. Then, $\mathsf{Var}(f_{z,-i}) = \mathbf{E}_{x\in\{-1,1\}}[\|f_{z,-i}(x) - \mathbf{E}_z[f_{z,-i}]\|_2^2]$. Define the influence of the i-th coordinate on f by

$$\mathsf{Inf}_i(f) = \mathbf{E}_{z\in\{-1,1\}^{n-1}}[\mathsf{Var}(f_{z,-i})].$$

Theorem 3.1. *There exist computable functions $n_0 = n_0(k,\varepsilon,\rho)$ and $C = C(k,\varepsilon,\rho)$ such that the following holds. Given any $k \in \mathbb{N}$ and $\rho \in [0,1], \varepsilon > 0$ and $\mu = (\mu_1,\ldots,\mu_k) \in \Delta_k$ and $n \geq n_0$, there is a function $g = g_{\mu,\rho,\varepsilon} : \{-1,1\}^n \to [k]$ such that*

1. $\max_i \mathsf{Inf}_i(g) \leq C/\sqrt{n}$;
2. $|\Pr_{x\in\{-1,1\}^{n_0}}[g(x) = i] - \mu_i| \leq C/\sqrt{n}$; *and*
3. *for any $f : \{-1,1\}^n \to [k]$ satisfying $|\Pr_{x\in\{-1,1\}^n}[f(x) = i] - \mu_i| \leq \varepsilon$ and $\max_i \mathsf{Inf}_i(f) \leq \varepsilon$,*

$$\mathbf{E}[\langle f, T_\rho f\rangle] \leq \mathbf{E}[\langle g, T_\rho g\rangle] + 2k\varepsilon.$$

Further, the function g can be computed given the parameters μ, ρ and ε.

In particular, the case where $(\mu_1,\ldots,\mu_k) = (1/k,\ldots,1/k)$ implies that in a tied election between k alternatives, we can find an ε-optimally robust noise-stable voting rule, where the stability is with respect of each voter randomizing their vote independently with probability $1-\rho$. Note that in the actual "Majority is Stablest" theorem, the function g is known explicitly (it is the "majority" function), whereas in Theorem 3.1 we only know that g is computable. An anonymous referee suggested that we call Theorem 3.1 a "Something is Stablest" theorem.

The proof of the theorem, which is omitted, is by now a standard reduction from a Gaussian noise stability result to a discrete one [25, 17, 6]. In one direction, the invariance principle immediately bounds the discrete stability of any partition by the Gaussian stability. In the other direction, starting from an ε-optimal Gaussian partition, each Gaussian can be replaced by a normalized sum of independent variables to obtain a discrete partition with nearly the same stability. The same result holds for other domains, for example for $f, g : [k]^n \to [k]$.

3.1 Relationship to rounding of SDPs

We do not know of any connection between our results and those of Raghavendra and Steurer [30] who showed that for any CSP, there is a rounding algorithm that is optimal up to ε, whose running time is polynomial in the instance size and doubly exponential in $1/\varepsilon$. It is natural to suspect that the two results are related, since there are well-known connections between SDP rounding and Gaussian noise stability. However, the usual analysis relating rounding to Gaussian noise stability seems to require that halfspaces maximize noise stability for all possible values of the noise. In our results, this property does not necessarily hold.

In the other direction, it is tempting to try to cast the noise stability problem as an optimization problem on a "Gaussian graph:" divide $\mathbb{R}^n$ into tiny pieces $A_1,\ldots,A_m$, and consider the weighted graph with vertices $1,\ldots,m$ and an edge between vertices i and j that has weight $\mathbf{E}[1_{A_i} P_t 1_{A_j}]$ (where 1_A denotes the indicator function of the set A). If the partition is fine enough, maximizing noise stability is approximately equivalent to finding an appropriate optimal partition of this graph, and so one might hope to apply the results of Steurer and Raghavendra to obtain computable bounds on the dimension where an almost optimal solution can be achieved. It is hard to implement this approach for two reasons: first, we do not know the SDP solution for the Gaussian graph; second, we are interested in the optimal solution and it is not clear what is the relation between the best integral solution and the SDP solution for the Gaussian graph.

While we do not see how to formally relate the two works, connecting the two (if possible) would be likely to yield important insights.

3.2 Non-interactive correlation distillation

Next, we discuss a basic problem in information theory and communication complexity which was recently considered in the work of Ghazi, Kamath and Sudan [12]. Let there be two non-communicating players Alice and Bob who have access to independent samples from a joint distribution $\mathbf{P}$ on the set $\mathcal{A} \times \mathcal{B}$. In other words, Alice (resp. Bob) has access to $(x_1, x_2, \ldots,)$ (resp. $(y_1, y_2, \ldots)$) such that $x_i \in \mathcal{A}$, $y_i \in \mathcal{B}$ and for each $i \in \mathbb{N}$, (x_i, y_i) is distributed according to $\mathbf{P}$, and the random variables $\{(x_i, y_i)\}_{i=1}^n$ are mutually independent (as i varies). Let $\mu = (\mu_1, \ldots, \mu_k) \in \Delta_k$ and $\nu = (\nu_1, \ldots, \nu_k) \in \Delta_k$. What is the maximum $\kappa \in [0,1]$ such that Alice and Bob can non-interactively jointly sample a distribution $\mathbf{Q}$ on $[k] \times [k]$ such that the distribution of the marginals of Alice and Bob are μ and ν, respectively, and they sample the same output with probability κ?

To formulate this problem more precisely, we introduce some notation. Let $x^n = (x_1, \ldots, x_n)$ and $y^n = (y_1, \ldots, y_n)$, where each pair (x_i, y_i) is drawn from $\mathbf{P}$, independently with respect to i. Now, note that a non-interactive protocol for Alice and Bob is equivalent to a pair (f, g) where $f : \mathcal{A}^n \to [k]$ and $g : \mathcal{B}^n \to [k]$ (for some $n \in \mathbb{N}$). In this terminology, the question now becomes the following: given μ, ν, do there exist n and $f : \mathcal{A}^n \to [k]$ and $g : \mathcal{B}^n \to [k]$ such that $f(x^n) \sim \mu$, $g(y^n) \sim \nu$, and $\Pr(f(x^n) = g(y^n)) = \kappa$? (Here, $f(x^n) \sim \mu$ means that μ is the law of $f(x^n)$.)

Before we state the main result of [12] and our extension, we consider a motivating example. Let $\mathcal{A} = \mathcal{B} = \mathbb{R}$ and let $\mathbf{P}$ be the law of two ρ-correlated standard Gaussians. Let $k = 2$ and $\mu = \nu$. Then, Borell's isoperimetric theorem (Theorem 1.1) states that the maximum achievable κ is given by

$$\kappa = \Pr_{x,y}[f(x) = f(y)],$$

where $f : \mathbb{R} \to \Delta_2$ is a halfspace with measure μ. Thus, in the above case, $n = 1$ suffices and $f = g$ is the halfspace whose measure is μ. We now state the main result of [12]. For a probability distribution $\mathbf{P}$, we let $|\mathbf{P}|$ denote the size of some standard encoding of $\mathbf{P}$.

Theorem 3.2 (Ghazi-Kamath-Sudan). *Let $(\mathcal{A} \times \mathcal{B}, \mathbf{P})$ be a probability space, fix $\delta > 0$, and let x^n and y^n be as above. There is an algorithm running in time $O_{|\mathbf{P}|,\delta}(1)$ such that given μ and ν in Δ_2 and a parameter $\kappa \in [0,1]$, it distinguishes between the following two cases:*

1. *There exist $n \in \mathbb{N}$, $f : \mathcal{A}^n \to \{0,1\}$, and $g : \mathcal{B}^n \to \{0,1\}$ such that $f(x^n) \sim \mu$, $g(y^n) \sim \nu$, and $\Pr(f(x^n) = g(y^n)) \geq \kappa - \delta$. Moreover, there is a computable function n_0 such that $n \leq n_0(\mathbf{P}, \delta)$. Similarly, f and g can also be computed given $\mathbf{P}$, μ, ν and δ as inputs.*

2. *For any $n \in \mathbb{N}$ and $f : \mathcal{A}^n \to \{0,1\}$ and $g : \mathcal{B}^n \to \{0,1\}$ that satisfy $f(x^n) \sim \mu$ and $g(y^n) \sim \nu$, $\Pr(f(x^n) = g(y^n)) \leq \kappa - 8\delta$.*

In other words, the above theorem states that there is an algorithm which, given $\mathbf{P}$, target marginals μ, ν, and correlation κ, can distinguish between two cases: (a) In the first case, it is possible for Alice and Bob to non-interactively simulate a distribution $\mathbf{R}$ which has the correct marginals and achieves the correlation κ up to δ. In this case, there is a computable bound on the number of copies of $\mathbf{P}$ required and the algorithm also outputs the functions f, g used for the non-interactive simulation. (b) In the second case, no non-interactive protocol between Alice and Bob can simulate a distribution $\mathbf{R}$ such that it has the correct marginals and the correlation is more than $\kappa - 8\delta$.

We remark that the requirement for the marginals to match *exactly* is unimportant. Indeed, any protocol where the marginals match approximately can be "fixed" to have exact marginals, with a small loss in the correlation.

The main restriction of Theorem 3.2 is that the output of the non-interactive protocol is a pair of bits. Theorem 2.2 immediately implies the following modification of Theorem 3.2: The distribution $\mathbf{P}$ is a Gaussian measure and $\{0,1\}$ is replaced by $[k]$. In the best of both worlds, we would be able to replace $\{0,1\}$ by $[k]$ without adding the assumption that $\mathbf{P}$ is a Gaussian measure. Using the methods we develop here, this also turns out to be possible; this is done in a follow-up work [8].

4 Preliminaries

We will start by defining some technical preliminaries which will be useful for the rest of the paper.

Definition 4.1. For $k \in \mathbb{N}$ and $1 \le i \le k$, let $\mathbf{e}_i \in \mathbb{R}^k$ be the unit vector along coordinate i and let Δ_k be the convex hull formed by $\{\mathbf{e}_i\}_{1 \le i \le k}$.

A function $f : \mathbb{R}^n \to \mathbb{R}^k$ is said to be in $L^2(\gamma_n)$ if $\int_x \|f(x)\|_2^2 \cdot \gamma_n(x) dx$ is finite. In this paper, unless explicitly mentioned otherwise, (a) the domain is always equipped with the standard Gaussian measure γ_n and (b) all functions are in $L^2(\gamma_n)$. A key property of such functions is that they admit the so-called Hermite expansion. To define this notion, let us first define a family of polynomials $H_q : \mathbb{R} \to \mathbb{R}$ (for $q \ge 0$) as

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

This family of polynomials is referred to as the *Hermite polynomials*. Let $\mathbb{Z}^*$ denote the subset of non-negative integers. For $S \in \mathbb{Z}^{*n}$, define $H_S : \mathbb{R}^n \to \mathbb{R}$ as

$$H_S(z) = \prod_{i=1}^{n} H_{S_i}(z_i).$$

It is well known that the set $\{H_S\}_{S \in \mathbb{Z}^{*n}}$ forms an orthonormal basis for $L^2(\gamma_n)$. In other words, every $f \in L^2(\gamma_n)$ may be written as

$$f = \sum_{S \in \mathbb{Z}^{*n}} \widehat{f}(S) \cdot H_S,$$

where the sum converges respect to $L^2(\gamma_n)$. We call $\widehat{f}(S) = (\widehat{f}_1(s) \ldots, \widehat{f}_k(S)) \in \mathbb{R}^k$ the *Hermite coefficients*, and we refer to the formula above as the *Hermite expansion* of f. Because the H_S form an orthonormal basis, we obtain *Parseval's identity*:

$$\int_{\mathbb{R}^n} \|f(x)\|_2^2 \, \gamma_n(x)dx = \sum_{S\in\mathbb{Z}^{*n}} \|\widehat{f}(S)\|_2^2. \tag{4.1}$$

Finally, for $f:\mathbb{R}^n\to\mathbb{R}^k$ and $d\in\mathbb{N}$, we define the degree-d truncation of f by "throwing away" all Hermite coefficients at levels higher than d:

$$f_{\le d}(x) = \sum_{S:|S|\le d} \widehat{f}(S)\cdot H_S(x).$$

We also define $\mathsf{W}^{\le d}[f] = \|f_{\le d}\|_2^2$ and $\mathsf{W}^{>d}[f] = \sum_{|S|>d} \|\widehat{f}(S)\|_2^2$.

Ornstein-Uhlenbeck operator

Definition 4.2. The Ornstein-Uhlenbeck operator P_t is defined for $t\in[0,\infty)$ and any bounded, measurable function $f:\mathbb{R}^n\to\mathbb{R}^k$ by

$$(P_t f)(x) = \int_{y\in\mathbb{R}^n} f(\mathrm{e}^{-t}\cdot x + \sqrt{1-\mathrm{e}^{-2t}}\cdot y)d\gamma_n(y).$$

Note that if f takes values in $\Delta_k\subset\mathbb{R}^k$, then so does $P_t f$ for every $t>0$. A basic fact about the Ornstein-Uhlenbeck operator is that the functions $\{H_S\}$ are eigenfunctions of this operator. We leave the proof of the next proposition to the reader.

Proposition 4.3. *For $S\in\mathbb{Z}^{*n}$, $P_t H_S = \mathrm{e}^{-t\cdot|S|}\cdot H_S$.*

We define the noise stability of $f:\mathbb{R}^n\to\Delta_k$ with noise rate $t\in[0,\infty)$ by

$$\mathsf{Stab}_t(f) = \mathbf{E}[\langle f, P_t f\rangle].$$

Multivariate polynomial threshold functions

We recall the definition of polynomial threshold functions from the introduction (primarily to define an associated quantity which shall be useful later).

Definition. A function $f:\mathbb{R}^n\to[k]$ is said to be a multivariate PTF if there exist polynomials $p_1,\dots,p_k:\mathbb{R}^n\to\mathbb{R}$ such that

$$f(x) = \begin{cases} j & \text{if } p_j(x)>0 \text{ and for } i\ne j,\ p_i(x)\le 0,\\ 1 & \text{otherwise.}\end{cases}$$

We then denote f as $f=\mathsf{PTF}(p_1,\dots,p_k)$. Further, f is said to be a degree-d multivariate PTF if $p_1,\dots,p_k$ are degree-d polynomials. For the rest of this paper, we will use "PTF" to mean "multivariate PTF," and the underlying k will be clear from the context. Also, let us define $\mathsf{Collision}(f)\subseteq\mathbb{R}^n$ defined as

$$\mathsf{Collision}(f) = \big\{x\in\mathbb{R}^n : |\{j : p_j(x)>0\}|\ne 1\big\}.$$

For technical reasons, we will also require the PTFs to satisfy a certain regularity property:

Definition 4.4. A degree-d multivariate PTF $f = \mathsf{PTF}(p^{(1)}, \dots, p^{(k)})$ is said to be (d,δ)-balanced if each $p^{(i)}$ has variance 1 and $|\mathbf{E}[p^{(i)}]| \le \log^{d/2}(k \cdot d/\delta)$.

Observe that the first condition (namely, $\mathsf{Var}(p^{(i)}) = 1$) can be achieved without loss of generality by simply rescaling all the polynomials to have variance 1. This rescaling does not change the value of the PTF at any point x. While the condition on expectation is non-trivial, the next proposition says that any multivariate PTF can be assumed to be (d,δ)-balanced while only changing the value of the PTF at δ-fraction of places.

Proposition 4.5. *Let $f : \mathbb{R}^n \to [k]$ satisfy $f = \mathsf{PTF}(p^{(1)}, \dots, p^{(k)})$ for some collection of degree-d polynomials $p^{(1)}, \dots, p^{(k)}$. For any $\delta > 0$, there is a (d,δ)-balanced PTF $g = \mathsf{PTF}(q^{(1)}, \dots, q^{(k)})$ defined on $\mathbb{R}^n$ such that $\Pr_x[g(x) \ne f(x)] \le \delta$ and $\Pr_x[x \in \mathsf{Collision}(g)] \le \Pr_x[x \in \mathsf{Collision}(f)] + \delta$.*

To prove this proposition, we will require the following standard concentration bound for low-degree polynomials which follows as a consequence of the well-known hypercontractive inequality—e. g., Lemma 2.2 in [10] proves the analogue of this statement for the Boolean hypercube. To get the statement for Gaussians, one can just substitute each one-dimensional Gaussian x_i by $((y_{i,1} + \cdots + y_{i,m})/\sqrt{m})$ where each $y_{i,j}$ is an unbiased $\{\pm 1\}$ random variable and then take the limit as $m \to \infty$.

Theorem 4.6. *Let $p : \mathbb{R}^n \to \mathbb{R}$ be a degree-d polynomial. Then, for any $t > 0$,*

$$\Pr_x \left[|p(x) - \mathbf{E}[p(x)]| \ge t \cdot \sqrt{\mathsf{Var}[p]} \right] \le d \cdot \mathrm{e}^{-t^{2/d}}.$$

Proof of Proposition 4.5. As we have observed, we can assume without loss of generality that the polynomials $p^{(i)}$ have variance 1. We define the polynomials $q^{(i)}$ as follows. If $|\mathbf{E}[p^{(i)}]| \le \log^{d/2}(d \cdot k/\delta)$, then we set $q^{(i)} = p^{(i)}$. Else, let $b_i = \mathsf{sign}(\mathbf{E}[p^{(i)}])$ and define

$$q^{(i)} = p^{(i)} - \mathbf{E}[p^{(i)}] + b_i \cdot \log^{d/2}(d \cdot k/\delta).$$

According to this definition, $\mathsf{sign}(q^{(i)}(x)) \ne \mathsf{sign}(p^{(i)}(x))$ implies that $|q^{(i)}(x) - \mathbf{E}[q^{(i)}]| \ge \log^{d/2}(dk/\delta)$. By Theorem 4.6,

$$\Pr_{x \sim \gamma_n} \left[\mathsf{sign}(q^{(i)}(x)) \ne \mathsf{sign}(p^{(i)}(x)) \right] \le \delta/k.$$

Note that if $f(x) \ne g(x)$ then there is at least one i for which $\mathsf{sign}(p^{(i)}(x)) \ne \mathsf{sign}(q^{(i)}(x))$. By a union bound, it follows that $\Pr_{x \sim \gamma_n}[f(x) \ne g(x)] \le \delta$. Similarly, if $x \in \mathsf{Collision}(g)$ but $x \notin \mathsf{Collision}(f)$ then there is at least one i for which $\mathsf{sign}(p^{(i)}(x)) \ne \mathsf{sign}(q^{(i)}(x))$. Using another union bound,

$$\Pr_{x \sim \gamma_n}[x \in \mathsf{Collision}(f)] \le \Pr_x[x \in \mathsf{Collision}(g)] + \delta. \qquad \square$$

Convention

At several places in the paper, we will establish bounds on one quantity in terms of others without the explicit mention of the dependence. For example, if we state that a quantity $d = d(k,\varepsilon)$, we mean that that d can be bounded as a computable function of k and ε (even if the function $d(\cdot,\cdot)$ is simple, in the interest of clarity, we do not explicitly write it).

5 Proof strategy for Theorem 2.2

We now formally state the main results required to prove Theorem 2.2. The first theorem says that given any k-ary function with a given set of measures for each of the k-partitions, there is a low-degree multivariate PTF which up to an error ε has (a) the same measures for the induced partitions and (b) is no less noise stable at a fixed noise rate t. In particular, we have the following theorem.

Theorem 5.1. *Let $f : \mathbb{R}^n \to [k]$ satisfy $\mathbf{E}[f] = \mu$ where $\mu = (\mu_1, \dots, \mu_k)$. Then, for every $\varepsilon > 0$ and $t > 0$, there exists a multivariate PTF $g' : \mathbb{R}^n \to [k]$ of degree $d = d(t, k, \varepsilon)$ such that*

- $\|\mathbf{E}[g'] - \mu\|_1 \leq \varepsilon$,
- $\mathbf{E}[\langle g', P_t g' \rangle] \geq \mathbf{E}[\langle f, P_t f \rangle] - \varepsilon$,
- $\Pr[x \in \mathsf{Collision}(g')] \leq \varepsilon$, *and*
- g' *is (d, ε)-balanced.*

The key difference between Theorem 5.1 and Theorem 2.2 is that the number of variables in PTF g' is the same as that in f. The above statement corresponds to the first step ("from general partitions to PTF") of the proof of Theorem 2.2 sketched in Section 2. The next theorem statement formalizes the second and third steps sketched in Section 2. In particular, it states that given a degree-d multivariate PTF over n variables, there is another multivariate PTF which induces approximately the same partition sizes and has approximately the same noise stability (at any fixed noise rate t) but the new PTF is only over some $O_{d,t}(1)$ variables.

Theorem 5.2. *Let $f : \mathbb{R}^n \to [k]$ be a degree-d, (d, ε)-balanced PTF with $\mathbf{E}_x[f(x)] = \mu$ where $\mu = (\mu_1, \dots, \mu_k)$. Further, let us assume that $\Pr[x \in \mathsf{Collision}(f)] \leq \varepsilon/(40k^2)$. Then, for every $\varepsilon > 0$ and $t > 0$, there exists a degree-d PTF $f_{\mathsf{junta}} : \mathbb{R}^{n_0} \to [k]$ such that,*

- $\|\mathbf{E}_x[f_{\mathsf{junta}}(x)] - \mu\|_1 \leq \varepsilon$,
- $\mathbf{E}[\langle f_{\mathsf{junta}}, P_t f_{\mathsf{junta}} \rangle] \geq \mathbf{E}[\langle f, P_t f \rangle] - \varepsilon$.

Further, $n_0 = n_0(d, k, \varepsilon, t)$ is a computable function.

We now see how Theorem 5.1 and Theorem 5.2 yield Theorem 2.2.

Proof of Theorem 2.2

Let us assume that given measure $\mu = (\mu_1, \dots, \mu_k)$ and noise rate $t > 0$, the most noise-stable partition is $f : \mathbb{R}^n \to [k]$. Then, applying Theorem 5.1 (with error parameter $\varepsilon/(80k^2)$), we obtain a PTF g_1 of degree $d = O_{t,k,\varepsilon}(1)$ such that it is $(d, \varepsilon/2)$ balanced and $\Pr[x \in \mathsf{Collision}(g_1)] \leq \varepsilon/(80k^2)$. Further, note that $\|\mathbf{E}[g_1] - \mathbf{E}[f]\|_1 \leq \varepsilon/2$ and $\mathbf{E}[\langle g_1, P_t g_1 \rangle] \geq \mathbf{E}[\langle f, P_t f \rangle] - \varepsilon/2$.

We next apply Theorem 5.2 to function g_1 to obtain f_{junta} which is a degree-d PTF on $n_0 = O_{t,k,\varepsilon}(1)$ variables such that $\|\mathbf{E}[g_1] - \mathbf{E}[f_{\mathsf{junta}}]\|_1 \leq \varepsilon/2$ and $\mathbf{E}[\langle f_{\mathsf{junta}}, P_t f_{\mathsf{junta}} \rangle] \geq \mathbf{E}[\langle g_1, P_t g_1 \rangle] - \varepsilon/2$.

Combining these two facts, we obtain $\|\mathbf{E}[g_1] - \mathbf{E}[f_{\mathsf{junta}}]\|_1 \leq \varepsilon$ and $\mathbf{E}[\langle f_{\mathsf{junta}}, P_t f_{\mathsf{junta}} \rangle] \geq \mathbf{E}[\langle f, P_t f \rangle] - \varepsilon$. Setting $g = f_{\mathsf{junta}}$ concludes the proof.

6 Reduction from arbitrary functions to PTFs: Proof of Theorem 5.1

In this section, we will prove Theorem 5.1 in two parts. The first step is to show that we can replace f by a function with a lower bound on the Hermite decay:

Lemma 6.1. *For every $f : \mathbb{R}^n \to [k]$, every $\varepsilon > 0$, and every $t > 0$, there exists a function $h : \mathbb{R}^n \to [k]$ satisfying the following:*

- $\mathbf{E}[h] = \mathbf{E}[f]$,
- $\mathbf{E}[\langle h, P_t h\rangle] \geq \mathbf{E}[\langle f, P_t f\rangle] - \varepsilon$,
- $\mathsf{W}^{>d}[h] \leq \varepsilon$ *for a computable* $d = d(t,k,\varepsilon)$.

The second step in the proof of Theorem 5.1 goes from the Hermite decay to an actual PTF.

Lemma 6.2. *Let $h : \mathbb{R}^n \to [k]$ satisfy $\mathsf{W}^{>d}[h] \leq \varepsilon$. Then, there is a PTF $g : \mathbb{R}^n \to [k]$ of degree d such that $\mathbf{E}_x[g(x) \neq h(x)] \leq k^2 \cdot \varepsilon$ and $\Pr_x[x \in \mathsf{Collision}(h)] \leq k^2 \cdot \varepsilon$.*

It is easy to see that Theorem 5.1 follows in a straightforward way by combining Lemma 6.1 and Lemma 6.2. Note that the condition of g being (d,ε)-balanced can be obtained by applying Proposition 4.5. We begin with the proof of Lemma 6.2, which is easier.

Proof of Lemma 6.2. By identifying $i \in [k]$ with $\mathbf{e}_i$ as before, we view h as a function taking values in $\mathbb{R}^k$. Define $h^{(0)} : \mathbb{R}^n \to \mathbb{R}^k$ by $h^{(0)}(x) = h(x) - (1/k)\cdot \mathbf{1}$, where $\mathbf{1} = (1,\dots,1) \in \mathbb{R}^k$. Note that for every $x \in \mathbb{R}^n$, $h^{(0)}(x)$ is a k-dimensional vector with the entry $(k-1)/k$ in one of the coordinates and $-1/k$ in every other coordinate. Call such a point *a k-lattice point*.

Let $h^{(0)}_{\leq d} : \mathbb{R}^n \to \mathbb{R}^k$ be the function obtained by truncating the Hermite expansion of $h^{(0)}$ at degree d. Consider the degree-d polynomials $p_1,\dots,p_k : \mathbb{R}^n \to \mathbb{R}$ where p_i is the i-th coordinate of $h^{(0)}_{\leq d}$, and define $g = \mathsf{PTF}(p_1,\dots,p_k)$. We claim that g satisfies the conditions of the lemma.

Consider any point $x \in \mathbb{R}^n$.

- If $x \in \mathsf{Collision}(g)$, we claim that $\|h^{(0)}(x) - h^{(0)}_{\leq d}(x)\|_2^2 \geq k^{-2}$. To see this, note that if $x \in \mathsf{Collision}(g)$, there are two possibilities: either $p_i(x) \leq 0$ for all $1 \leq i \leq k$ or there are two distinct i,j such that $p_i(x) > 0$ and $p_j(x) > 0$. In either case, it easily follows that the ℓ_2 distance from $h^{(0)}_{\leq d}(x)$ to any k-lattice point is at least $1/k$.

- If $x \notin \mathsf{Collision}(g)$ and $g(x) \neq h(x)$, then again we claim that $\|h^{(0)}(x) - h^{(0)}_{\leq d}(x)\|_2^2 \geq k^{-2}$. To see this, suppose that $g(x) = \mathbf{e}_i \neq \mathbf{e}_j = h(x)$. Then the jth coordinate of $h^{(0)}(x)$ is $(k-1)/k$, but on the other hand $p_j(x) \leq 0$ because $g(x) = \mathbf{e}_i$ and $x \notin \mathsf{Collision}(g)$. Hence, $\|h^{(0)}(x) - h^{(0)}_{\leq d}(x)\|_2^2 \geq (k-1)^2/k^2$.

Combining these two, we get

$$\mathbf{E}_x\big[\|h^{(0)}(x) - h^{(0)}_{\leq d}(x)\|_2^2\big] \geq \frac{1}{k^2}\cdot\big(\Pr_x[x \in \mathsf{Collision}(g)] + \Pr_x[x \notin \mathsf{Collision}(g) \ \wedge\ g(x) \neq h(x)]\big).$$

As $h^{(0)}$ and h differ only in the Hermite coefficient for $S = \emptyset$, the left hand side is the same as $\mathsf{W}^{\geq d}[h]$. Combining this with the above inequality achieves all the stated guarantees. ∎

The proof of Lemma 6.1 is longer, and so we begin with an outline. The first observation is that bounding $\mathbf{E}[|\nabla h|]$ implies a bound on $\mathsf{W}^{>d}[h]$. This follows a standard spectral argument, and is stated as Corollary 6.4. The second observation is that the function h obtained by thresholding $P_t f$ at a suitable value (chosen, for example, so that $\mathbf{E}[h] = \mathbf{E}[f]$) satisfies $\mathbf{E}[\langle h, P_t h\rangle] \geq \mathbf{E}[\langle f, P_t f\rangle]$. This is stated as Lemma 6.7.

Based on the previous paragraph, it seems like we would like to bound $\mathbf{E}[|\nabla h|]$ where h is obtained by thresholding $P_t f$; let $a \in \mathbb{R}^k$ be the desired threshold value, so that thresholding $P_t f$ at a produces a partition with the right measures. It turns out, unfortunately, that for h defined in this way, $\mathbf{E}[|\nabla h|]$ could be arbitrarily large. A key insight of [22] is that (using the co-area formula and gradient bounds on $P_t f$) the partition produced by thresholding at a random value near a has bounded expected surface area. In particular, although thresholding exactly at a might be a bad idea, there exist many good nearby values at which to threshold. Based on this observation, we construct h in two steps. In the first step, we define h by thresholding $P_t f$, but only on the set of $x \in \mathbb{R}^n$ for which $P_t f(x)$ is not too close to a. By choosing "not too close" in a suitable random way, the observation of [22] implies that this step only contributes a bounded amount to $\mathbf{E}[|\nabla h|]$. Since the first step is almost the same as just thresholding $P_t f$ at a, it is consistent with our desire that $\mathbf{E}[\langle h, P_t h\rangle] \geq \mathbf{E}[\langle f, P_t f\rangle] - \varepsilon$.

In the second step, we partition the remaining part of $\mathbb{R}^n$ by chopping it with halfspaces of the correct size. Since halfspaces have a bounded surface area, this also contributes a bounded amount to $\mathbf{E}[|\nabla h|]$. Crucially, this step does not destroy the value of $\langle h, P_t h\rangle$; fundamentally, this is because $P_t f$ is almost constant on the set we are partitioning. This finishes the outline of the proof of Lemma 6.1. We now start with some preliminaries required in the proof of this lemma.

6.1 Surface area and spectrum

In our outline of the proof of Lemma 6.1, we claimed that a control of $\mathbf{E}[|\nabla h|]$ implies a control of $\mathsf{W}^{>d}[h]$. Here we prove that claim, using a theorem of Ledoux [23] that gives a lower bound on the gradient in terms of the noise sensitivity.

Theorem 6.3 (Ledoux)**.** *For a differentiable function $f : \mathbb{R}^n \to [0,1]$ and any $t > 0$,*

$$\mathop{\mathbf{E}}_{x\sim\gamma_n}[f \cdot (f - P_t f)] \leq \frac{\arccos(\mathrm{e}^{-t})}{\sqrt{2\pi}} \cdot \mathop{\mathbf{E}}_{x\sim\gamma_n}[|\nabla f|].$$

Using this theorem, it is possible to establish an upper bound on $\mathsf{W}^{\geq d}[f]$ in terms of $|\nabla f|$ as done below.

Corollary 6.4. *For any differentiable $f : \mathbb{R}^n \to [0,1]$ be of bounded variation (in Gaussian measure) and any $d \in \mathbb{N}$,*

$$\mathsf{W}^{\geq d}[f] \leq O\left(\frac{1}{\sqrt{d}}\right) \cdot \mathbf{E}[|\nabla f|].$$

Proof.

$$\mathbf{E}[f(f - P_t f)] = \sum_{j=0}^{\infty} (1 - \mathrm{e}^{-t\cdot j})\mathsf{W}^j[f] \geq (1 - \mathrm{e}^{-d\cdot t}) \cdot \mathsf{W}^{\geq d}[f].$$

Choose $d = 1/t$ and apply Theorem 6.3, we get

$$\mathsf{W}^{\geq d}[f] \leq O(1) \cdot \mathbf{E}[f(f - P_{\frac{1}{d}} f)] \leq \frac{\arccos(\mathrm{e}^{-\frac{1}{d}})}{\sqrt{2\pi}} \cdot \mathop{\mathbf{E}}_{x \sim \gamma_n}[|\nabla f|] \leq O\left(\frac{1}{\sqrt{d}}\right) \cdot \mathbf{E}[|\nabla f|]. \qquad \square$$

6.2 Various lemmas for thresholding

Recall that our proof of Lemma 6.1 is based on thresholding $P_t f$. Before proving it, we introduce various properties of the thresholding procedure. The first lemma states that the "best" way to round a Δ_k-valued function to a $\{\mathbf{e}_1, \dots, \mathbf{e}_k\}$-valued function while preserving its expectation is simply to threshold it. Here, "best" means that we are trying to maximize the correlation between the original function and the rounded one.

Lemma 6.5. *Let $f : \mathbb{R}^n \to \Delta_k$, $g : \mathbb{R}^n \to \{\mathbf{e}_1, \dots, \mathbf{e}_k\}$, and $z \in \mathbb{R}^k$. Suppose that whenever $g(x) = \mathbf{e}_i$, we have*

$$i \in \operatorname{argmax}\{f_j(x) - z_j : j = 1, \dots, k\}.$$

Then g maximizes $\mathbf{E}[\langle f, h \rangle]$ among all $h : \mathbb{R}^n \to \{\mathbf{e}_1, \dots, \mathbf{e}_k\}$ satisfying $\mathbf{E}[h] = \mathbf{E}[g]$.

Proof. Let $h : \mathbb{R}^n \to \{\mathbf{e}_1, \dots, \mathbf{e}_k\}$ satisfy $\mathbf{E}[h] = \mathbf{E}[g]$. By the defining assumption of g, $\langle g, f - z \rangle \geq \langle h, f - z \rangle$ pointwise. Hence,

$$\mathbf{E}[\langle h, f \rangle] = \mathbf{E}[\langle h, f - z \rangle] + \langle \mathbf{E}[h], z \rangle \geq \mathbf{E}[\langle g, f - z \rangle] + \langle \mathbf{E}[g], z \rangle = \mathbf{E}[\langle g, f \rangle]. \qquad \square$$

Let us assign a notation to the sort of thresholding involved in Lemma 6.5.

Definition 6.6. For $f : \mathbb{R}^n \to \Delta_k$ and $z \in \mathbb{R}^k$, let $\mathcal{F}_z(f)$ denote the set of functions $g : \mathbb{R}^n \to \{\mathbf{e}_1, \dots, \mathbf{e}_k\}$ such that

$$\langle g(x), f(x) - z \rangle = \max_j \{f_j(x) - z_j\}$$

for every $x \in \mathbb{R}^n$.

We say that functions in $\mathcal{F}_z(f)$ are obtained by "thresholding" f around the value z.

The next step is to show that a function obtained by thresholding $P_t f$ is always at least as noise-stable (with parameter t) as the original function f.

Lemma 6.7. *Let $f : \mathbb{R}^n \to \{\mathbf{e}_1, \dots, \mathbf{e}_k\}$ and take $t > 0$. If $g \in \mathcal{F}_z(f)$ for some $z \in \mathbb{R}^k$ satisfies $\mathbf{E}[g] = \mathbf{E}[f]$ then $\mathsf{Stab}_t(g) \geq \mathsf{Stab}_t(f)$.*

Proof. Since $\mathbf{E}[g] = \mathbf{E}[f]$, Lemma 6.5 implies that $\mathbf{E}[\langle g, P_t f \rangle] \geq \mathbf{E}[\langle f, P_t f \rangle] = \mathsf{Stab}_t(f)$. On the other hand, the Cauchy-Schwarz inequality and the semi-group property of P_t imply that

$$\mathbf{E}[\langle g, P_t f \rangle] = \mathbf{E}[\langle P_{t/2} g, P_{t/2} f \rangle] \leq \sqrt{\mathsf{Stab}_t(g) \mathsf{Stab}_t(f)}.$$

The claim follows. $\square$

Next, we come to the importance of the thresholding parameter z: it allows us to tune the expectation of the thresholded functions. In particular, in order to make Lemma 6.7 useful we need to show that we can always find a thresholded function with the same expectation as the original function.

Lemma 6.8. *For any $f : \mathbb{R}^n \to \Delta_k$, there exists some $z \in \mathbb{R}^k$ and $g \in \mathcal{F}_z(f)$ such that $\mathbf{E}[g] = \mathbf{E}[f]$.*

It will actually be convenient to prove a more general Lemma about splitting up the mass of a general probability measure. Given $z \in \mathbb{R}^k$ and $i \in [k]$, let $A_i(z) \subset \Delta_k$ be the set $\{x \in \Delta_k : x_i - z_i = \max_j x_j - z_j\}$. The correspondence between this notion and the definition of $\mathcal{F}_z(f)$ is that $g \in \mathcal{F}_z(f)$ if and only if $g(x) = \mathbf{e}_i$ implies $f(x) \in A_i(z)$.

Lemma 6.9. *For every probability measure ζ on Δ_k and every $q \in \Delta_k$, there exists $z \in \mathbb{R}^k$ such that $\zeta(\bigcup_{i \in I} A_i(z)) \geq \sum_{i \in I} q_i$ for every $I \subset [k]$.*

Before proving Lemma 6.9, we will show how it implies Lemma 6.8. For this, we introduce one more ingredient: a weighted version of Hall's marriage theorem. It will be used to show that we can divide a certain amount of "missing" volume into pieces of the right size, while preserving certain constraints.

Lemma 6.10. *Let $G = (U,V,w,E)$ be a finite, vertex-weighted, bipartite graph. That is, U and V are finite sets, w is a weight function $w : U \sqcup V \to (0,\infty)$, and $E \subset U \times V$ is the set of edges. Suppose that for every $U' \subset U$, $\sum_{u \in U} w(u) \leq \sum_{v \sim U'} w(v)$, where $v \sim U'$ if there exists $u \in U'$ such that $(u,v) \in E$. Then there exists $p : V \to \Delta_{|U|}$ such that $p_u(v) > 0$ implies $u \sim v$, and such that*

$$\sum_{v \in V} p_u(v) w(v) = w(u)$$

for every $u \in U$.

In the case that w takes values in $\{0,1\}$, Lemma 6.10 follows from the usual formulation of Hall's marriage theorem (which guarantees in addition that $p : V \to \{\mathbf{e}_1, \dots, \mathbf{e}_{|U|}\}$). When w is integer-valued (or, by scaling, rational-valued), it follows by applying Hall's marriage theorem to the graph in which u is replicated $w(u)$ times. The general case follows by a simple approximation argument.

Proof of Lemma 6.8. Let ζ be the distribution of f and let $q = \mathbf{E}[f]$. Applying Lemma 6.9, we may take $z \in \mathbb{R}^k$ such that $\zeta(\bigcup_{i \in I} A_i(z)) \geq \sum_{i \in I} q_i$ for every $I \subset [k]$. For $I \subset [k]$, let B_I be the set of $x \in \mathbb{R}^n$ such that $\{i : f(x_i) - z_i = \max_j f(x_j) - z_j\} = I$. Then $\{B_I : I \subset [k]\}$ is a partition of $\mathbb{R}^n$. Moreover, if we define $w(I) = \gamma_n(B_I)$ then

$$\sum_{J : J \cap I \neq \emptyset} w(J) = \gamma_n\Big(\bigcup_{J : J \cap I \neq \emptyset} B_J\Big) = \zeta\Big(\bigcup_{i \in I} A_i(z)\Big) \geq \sum_{i \in I} q_i. \tag{6.1}$$

Now we construct a bipartite graph with $U = [k]$ and $V = \{I : I \subset [k]\}$. The pair (i,I) is an edge if $i \in I$. We define w on V as above, and we define w on U by $w(i) = p_i$. According to (6.1), the condition of Lemma 6.10 holds; hence, there exists $p : V \to \Delta_k$ such that $p_i(I) > 0$ only when $i \in I$, and such that $\sum_{I \ni i} p_i(I) \gamma_n(B_I) = q_i$. Now we will define g: for each $I \subset [k]$, partition B_I arbitrarily into sets $\{B_{I,i} : i \in I\}$, where $\gamma_n(B_{I,i}) = p_i(I)\gamma_n(B_I)$. This may be done because γ_n has no atoms. Finally, define g to be e_i on every $B_{I,i}$. Then the condition $\sum_{I \ni i} p_i(I) \gamma_n(B_I) = q_i$ ensures that $\gamma_n(\{g = e_i\}) = q_i$. Moreover, note that $g(x) = e_i$ implies that $x \in B_I$ for some I, which implies that $f_i(x) - z_i = \max_j (f_j(x) - z_j)$. This completes the construction of g. $\square$

Proof of Lemma 6.9. We will assume initially that ζ is absolutely continuous with respect to the uniform measure on Δ_k. As a consequence, the function $\psi : \Delta_k \to \mathbb{R}^k$ defined by $\psi_i(z) = \zeta(A_i(-kz))$ is continuous. Moreover, the image of ψ is in Δ_k, and we need prove that it is all of Δ_k.

For $I \subsetneq \{1,\dots,k\}$, let $F_I = \{x \in \Delta_k : x_i = 0 \text{ for all } i \in I\}$. Note that every face of Δ_k is of the form F_I for some $I \subsetneq \{1,\dots,k\}$. Next, we claim that ψ maps F_I into F_I for every I. Indeed, if $z \in F_I$ then there is at least one $j \notin I$ such that $z_j \geq 1/k$. For this j and any $i \in I$, if x is in the interior of Δ_k then

$$x_i + kz_i = x_i < x_j + kz_j.$$

It follows that $A_i(-kz)$ does not intersect the interior of Δ_k; hence, $\psi_i(z) = 0$. Since this holds for all $i \in I$, $\psi(z) \in F_I$. By Sperner's lemma, any map from Δ_k into itself that leaves all faces invariant must be onto, and so ψ is onto, as claimed.

To complete the proof, we must eliminate the assumption that ζ is absolutely continuous with respect to the uniform measure. For an arbitrary probability measure ζ, let ζ_n be a sequence of probability measures that are absolutely continuous with respect to the uniform measure, and which converge to ζ in distribution. (For example, we may construct ζ_n by first pushing ζ forward under the map $x \mapsto (1/2n)(1/k,\dots,1/k) + (1-1/n)x$, and by convolving the push-forward with a smooth bump function supported on $(1/2n)\Delta_2$.) Define $\psi_n : \Delta_k \to \Delta_k$ by $\psi_n(z) = (\zeta_n(A_1(z)),\dots,\zeta_n(A_k(z)))$. By the previous argument, for every n there is some $z_n \in \Delta_k$ such that $\psi(z_n) = q$. Since Δ_k is compact, we may pass to a subsequence and thereby assume that z_n converges to some limit z_∞. Note that

$$\bigcup_{i \in I} A_i(z_\infty) = \bigcap_{n=1}^{\infty} \Big(\bigcup_{i \in I} \bigcup_{m \geq n} A_i(z_m) \cup A_i(z_\infty) \Big).$$

Moreover, $\bigcup_{m \geq n} A_i(z_n) \cap A_i(z_\infty)$ is closed for every n, since z_∞ is the only limit point of the sequence z_n. It follows that

$$\mu\Big(\bigcup_{i \in I} A_i(z_\infty)\Big) = \lim_{n\to\infty} \mu\Big(\bigcup_{i \in I} \bigcup_{m \geq n} A_i(z_m) \cup A_i(z_\infty)\Big) \geq \lim_{n\to\infty} \lim_{\ell\to\infty} \mu_\ell\Big(\bigcup_{i \in I} \bigcup_{m \geq n} A_i(z_m) \cup A_i(z_\infty)\Big) \geq \sum_i q_i,$$

since $\sum_{i \in I} \mu_\ell(A_i(z_\ell)) = \sum_{i \in I} q_i$ for every I and ℓ (using the fact that μ has a density, and so it assigns no mass to $A_i(z_\ell) \cap A_j(z_\ell)$). $\square$

6.3 Proof of Lemma 6.1

We introduce two final ingredients before beginning the proof of Lemma 6.1: the first is a gradient bound on $P_t f$ that is due (in a much more precise form) to Bakry and Ledoux [1].

Theorem 6.11. *If $f : \mathbb{R}^n \to [0,1]$ then for any $t > 0$, $P_t f$ is differentiable and $\|\nabla P_t f\|_\infty \leq C/\sqrt{t}$.*

The second ingredient is the co-area formula in Gaussian space. To state this, let γ_n^+ denote the Gaussian surface area in n-dimensions. The co-area formula (see [11] for a reference) is stated next. For any locally Lipschitz $f : \mathbb{R}^n \to \mathbb{R}$ and for any continuous, compactly supported $\mu : \mathbb{R} \to \mathbb{R}$

$$\int_{\mathbb{R}} \mu(t) \cdot \gamma_n^+(\{x : f(x) \geq t\}) dt = \int_{\mathbb{R}^n} \mu(f(x)) |\nabla(f(x))| d\gamma(x).$$

The co-area formula relates the Gaussian surface of a Boolean function (specified by $\{x : f(x) \geq t\}$) to the gradient of f. We remark that the co-area formula is sometimes written with $\gamma_n^+(\{x : f(x) \geq t\})$ replaced by the Gaussian-weighted $(n-1)$-dimensional Hausdorff measure of $\{x : f(x) = t\}$. The two formulations are equivalent, because these two quantities are equal for almost every t (see, for example, [24, Chapter 13]).

Proof of Lemma 6.1. We begin by applying Lemma 6.8 to $P_t f$: let z and $g \in \mathcal{F}_z(P_t f)$ be such that $\mathbf{E}[g] = \mathbf{E}[P_t f] = \mathbf{E}[f]$. According to Lemma 6.7, $\mathsf{Stab}_t(g) \geq \mathsf{Stab}_t(f)$. However, $\mathbf{E}[|\nabla g|]$ cannot be controlled in general; therefore, we will move to an approximation of g.

Let μ be a probability measure on $[0,\varepsilon]$ with continuous density bounded by $2/\varepsilon$. By the co-area formula, for any $i \neq j$ we have

$$\begin{aligned}
&\int_{\mathbb{R}} \mu(t)\gamma^+(\{x : P_t f_i(x) - z_i > P_t f_j(x) - z_j + t\})\, dt \\
&= \int_{\mathbb{R}^n} \mu(P_t f_i(x) - z_i - P_t f_j(x) + z_j)|\nabla P_t f_i(x) - \nabla P_t f_j(x)|\, d\gamma(x) \\
&\leq \frac{2}{\varepsilon}\int_{\mathbb{R}^n} |\nabla P_t f_i(x) - \nabla P_t f_j(x)|\, d\gamma(x) \\
&\leq \frac{C}{\varepsilon\sqrt{t}},
\end{aligned}$$

where the last inequality follows from Theorem 6.11. In particular, there exist some y_{ij}^+ and y_{ij}^- in $[0,\varepsilon]$ such that if $A_{ij}^{\pm} = \{x \in \mathbb{R}^n : P_t f_i(x) - z_i > P_t f_j(x) - z_j \pm y_{ij^{\pm}}\}$ then

$$\gamma^+(A_{ij}^{\pm}) \leq C(\varepsilon)/\sqrt{t}.$$

We repeat this construction of y_{ij} and A_{ij} for every ordered pair (i, j) with $i \neq j$. Define $A_i = \bigcap_{j\neq i} A_{ij}^+$, and note that $A_i \subset \{x : g(x) = \mathbf{e}_i\}$. On the other hand, $A_{ij}^- \supset \{x : g(x) = \mathbf{e}_i\}$ for every i, j. Next, define

$$\begin{aligned}
C_i &= \bigcap_{j\neq i} A_{ij}^-, \\
C_I &= \bigcup_{i\in I} C_i, \\
B_I &= C_I \setminus \bigcup_{J\supsetneq I} C_J \setminus \bigcup_{i\in I} A_i.
\end{aligned}$$

The meaning of these sets is the following: A_i is the set where $f_i - z_i$ is significantly larger than any other $f_j - z_j$. On C_i, $f_i - z_i$ is almost $\max_j f_j - z_j$; on C_I, $f_i - z_i$ is almost maximal for *every* $i \in I$; and on B_I, the set of i for which $f_i - z_i$ is almost maximal is exactly I. Importantly, the collection of all A_i and B_I form a partition of $\mathbb{R}^n$. Our basic strategy will be to set h to be $\mathbf{e}_i$ on A_i, and then to define h on the remaining part of the space in order to satisfy two properties: $\mathbf{E}[h] = \mathbf{E}[f]$ and $h(x) = \mathbf{e}_i$ only if $x \in B_I$ for some $I \ni i$.

Since the Gaussian surface area obeys the inequalities $\gamma_n^+(A \cap B) \leq \gamma_n^+(A) + \gamma_n^+(B)$ and $\gamma_n^+(A \cup B) \leq \gamma_n^+(A) + \gamma_n^+(B)$, and since B_I and A_i are defined using a finite (depending on k) number of intersections and unions, it follows that

$$\begin{aligned}\gamma^+(A_i) &\le C(\varepsilon,k)/\sqrt{t},\\ \gamma^+(B_I) &\le C(\varepsilon,k)/\sqrt{t}.\end{aligned} \tag{6.2}$$

Now, for any $I \subset [k]$,

$$\bigcup_{J:J\cap I\neq\emptyset} B_J \cup \bigcup_{i\in I} A_i$$

contains the set of x for which $g(x) \in \{\mathbf{e}_i : i \in I\}$. It follows that

$$\sum_{J:J\cap I\neq\emptyset} \gamma_n(B_J) \ge \sum_{i\in I} \mathbf{E}[g_i] - \gamma_n(A_i). \tag{6.3}$$

Consider the bipartite graph where $U = [k]$ and $V = \{I : I \subset [k]\}$, and $(i,I) \in E$ if $i \in I$. We assign the weights $w(i) = \mathbf{E}[g_i] - \gamma_n(A_i)$ to $i \in U$ and $w(I) = \gamma_n(B_I)$ for $I \in V$ (note that $w(i) \ge 0$ because the construction of A_i ensures that $A_i \subset \{x : g(x) = \mathbf{e}_i)\}$). Equation (6.3) ensures that this weighted graph satisfies the hypothesis of Lemma 6.10, and so there exists $p : V \to \Delta_k$ with $p_i(I) > 0$ only if $i \in I$, and with

$$\sum_{I\ni i} p_i(I)\gamma_n(B_I) = \mathbf{E}[g_i] - \gamma_n(A_i).$$

Finally, we will use p to define h. First, for every I and $i \in I$, let $B_{I,i}$ be a set of the form $\{x \in \mathbb{R}^n : a \le x_1 \le b\}$ such that $\gamma_n(B_{I,i} \cap B_I) = p_i(I)\gamma_n(B_I)$; moreover, we choose $B_{I,i}$ such that $\gamma_n(B_{I,i} \cap B_{I,j}) = 0$ when $i \neq j$. Then we set $h(x)$ to equal $\mathbf{e}_i$ on the set $A_i \cup \bigcup_{I\ni i}(B_{I,i} \cap B_I)$. By the defining property of p, h satisfies $\mathbf{E}[h_i] = \mathbf{E}[g_i] = \mathbf{E}[f_i]$. Moreover, $h(x) = \mathbf{e}_i$ only on a subset of $A_i \cup \bigcup_{I\ni i} B_I$, and on this set $P_t f_i(x) - z_i \ge \max_j P_t f_j(x) - z_j - \varepsilon = \langle g(x), P_t f(x) - z\rangle - \varepsilon$. It follows that

$$\mathbf{E}[\langle h, P_t f\rangle] \ge \mathbf{E}[\langle g, P_t f\rangle] - \varepsilon \ge \mathsf{Stab}_t(f) - \varepsilon.$$

By the Cauchy-Schwarz inequality,

$$\sqrt{\mathsf{Stab}_t(h)} \ge \sqrt{\mathsf{Stab}_t(f)} - \frac{\varepsilon}{\sqrt{\mathsf{Stab}_t(f)}},$$

which implies that $\mathsf{Stab}_t(h) \ge \mathsf{Stab}_t(f) - 2\varepsilon$.

Finally, we address the surface area of h. Recall that

$$\{x : h(x) = \mathbf{e}_i\} = A_i \cup \bigcup_{I,i} B_{I,i} \cap B_I.$$

The number of terms on the right is some constant depending on k. By (6.2), $\gamma_n^+(A_i)$ and $\gamma_n^+(B_I)$ are bounded by $C(\varepsilon,t,k)$. Since $B_{I,i}$ is an intersection of two halfspaces, its Gaussian surface area is at most a constant. It follows that $\gamma_n^+(\{x : h(x) = \mathbf{e}_i\}) \le C(\varepsilon,k)/\sqrt{t}$ for every i. Hence, $\mathbf{E}[|\nabla h|] \le C(\varepsilon,k)/\sqrt{t}$.

After applying Corollary 6.4, this completes the proof except that $d = d(k,\varepsilon,t)$ instead of the claimed $d(k,\varepsilon)$, where $d(k,\varepsilon,t)$ blows up as $t \to 0$. To eliminate this dependence on t, it suffices to note that the claim is trivial if $t \ll (\varepsilon/k)^2$. Indeed, one can easily construct h with $\mathbf{E}[h] = \mathbf{E}[f]$, $\mathsf{Stab}_t(h) \ge \mathbf{E}[|h|^2] - O(k\sqrt{t})$ and $\mathbf{E}[|\nabla h|] \le O(k)$. For example, we could take the pieces $\{x : h(x) = \mathbf{e}_i\}$ to be parallel slabs of the form $\{x : a_i \le x_1 \le b_i\}$, where $\{a_i\}$ and $\{b_i\}$ are chosen so that the slabs have the correct volumes. If $t \ll (\varepsilon/k)^2$ then such an example will satisfy the claim of the lemma. $\square$

7 Reduction from PTFs to PTFs on a constant number of variables

In this section, we are going to prove Theorem 5.2. We restate it here again for the convenience of the reader.

Theorem 5.2. *Let $f : \mathbb{R}^n \to [k]$ be a degree-d, (d,ε)-balanced PTF with $\mathbf{E}_x[f(x)] = \mu$ where $\mu = (\mu_1,\dots,\mu_k)$. Further, let us assume that $\Pr[x \in \mathsf{Collision}(f)] \le \varepsilon/(40k^2)$. Then, for every $\varepsilon > 0$ and $t > 0$, there exists a degree-d PTF $f_{\mathsf{junta}} : \mathbb{R}^{n_0} \to [k]$ such that*

- $\|\mathbf{E}_x[f_{\mathsf{junta}}(x)] - \mu\|_1 \le \varepsilon$, *and*
- $\mathbf{E}[\langle f_{\mathsf{junta}}, P_t f_{\mathsf{junta}}\rangle] \ge \mathbf{E}[\langle f, P_t f\rangle] - \varepsilon$.

Our main workhorse for this section is going to be two structural theorems for low-degree polynomials proven in [9]. In order to state these theorems, we will need to introduce several notions related to tensors and polynomial decomposition.

7.1 Polynomials and tensors

We give a brief description of the connection between symmetric tensors and polynomials under the Gaussian distribution. The interested reader may consult the book [18] for a detailed background. Let $\mathcal{H}$ denote the Hilbert space $\mathbb{R}^n$ and let $\mathcal{H}^{\otimes q}$ be used to denote the q-ary tensor product of $\mathcal{H}$, i. e., the space of all multilinear functions $\mathcal{H}^q \to \mathbb{R}$. For $f_1,\dots,f_q \in \mathcal{H}$, their tensor product $f_1 \otimes \cdots \otimes f_q$ is the element of $\mathcal{H}^{\otimes q}$ defined by

$$(f_1 \otimes \cdots \otimes f_q)(h_1,\dots,h_q) = \prod_{i=1}^{q} \langle f_i, h_i\rangle.$$

It is straightforward to check that $\mathcal{H}^{\otimes q}$ is spanned by elements of the form $f_1 \otimes \cdots \otimes f_q$; we therefore define an inner product on $\mathcal{H}^{\otimes q}$ by setting

$$\langle f_1 \otimes \cdots \otimes f_q, g_1 \otimes \cdots \otimes g_q\rangle = \prod_{i=1}^{q} \langle f_i, g_i\rangle$$

and extending to $\mathcal{H}^{\otimes q}$ by bilinearity. The norm of a tensor is defined by $\|f\|_F^2 = \langle f, f\rangle$.

An element $f \in \mathcal{H}^{\otimes q}$ is said to be symmetric if it is invariant under any permutation $\sigma : [q] \to [q]$, in the sense that $f(h_1,\dots,h_q) = f(h_{\sigma(1)},\dots h_{\sigma(q)})$ for every $h_1,\dots,h_q \in \mathcal{H}$. Let $\mathcal{H}^{\odot q}$ denote the symmetric subspace of $\mathcal{H}^{\otimes q}$. We now describe a map between the space $\mathcal{H}^{\odot q}$ and polynomials. Recall that H_q denotes a Hermite polynomial.

Definition 7.1. The iterated Itô integral I_q maps $\mathcal{H}^{\odot q}$ as follows: Let $h \in \mathcal{H}$ be a unit vector and note that $h^{\otimes q} \in \mathcal{H}^{\odot q}$. Then, $I_q(h^{\otimes q}) = H_q(\langle h, x\rangle)$ where $x \in \mathbb{R}^n$. The map I_q is extended linearly to $\mathcal{H}^{\odot q}$.

For the convenience of the reader, here we describe a basis of the space $\mathcal{H}^{\odot q}$. Consider an unordered multiset $S = \{s_1,\dots,s_q\} \subseteq [n]$ of size q, and define $\Phi_S \in \mathcal{H}^{\odot q}$ by

$$\Phi_S = \sum_{\sigma \in \mathsf{Sym}_q} e_{s_{\sigma(1)}} \otimes \cdots \otimes e_{s_{\sigma(q)}}.$$

It is easy to check (and we omit the proof) that the tensors Φ_S form a basis for $\mathcal{H}^{\odot q}$:

Proposition 7.2. *Let $\mathcal{S}$ be the set of unordered multisets of $[n]$ of size q. Then, $\{\Phi_S\}_{S\in\mathcal{S}}$ forms an orthogonal basis of $\mathcal{H}^{\odot q}$.*

A fundamental property of the map I_q is that it is an isometry between the space of symmetric tensors and polynomials endowed with the standard normal measure (see [18] for a proof).

Proposition 7.3. $\mathbf{E}_{x\sim\gamma_n}[I_q(f)\cdot I_p(g)] = \delta_{p=q}\cdot\langle f,g\rangle$.

We will refer to the range of I_q as the Wiener chaos $\mathcal{W}^q$. Based on Proposition 7.3, I_q is a bijective map from $H^{\odot q}$ to $\mathcal{W}^q$. On the other hand, it is easy to show that any polynomial $p:\mathbb{R}^n\to\mathbb{R}$ of degree at most d can be expressed uniquely as

$$p = \sum_{q=0}^{d} I_q(f_q) \ \text{ where } \ f_q \in \mathcal{H}^{\odot q}. \tag{7.1}$$

In order to apply the results from [9], we will often restrict our attention to multilinear polynomials, i. e., those for which every variable appears with degree at most 1. We call a tensor $f\in\mathcal{H}^{\odot q}$ *square-free* if $I_q(f)$ is a multilinear polynomial. Square-free tensors in $\mathcal{H}^{\odot q}$ may also be characterized as the set of tensors f for which $f(h_1,\dots,h_q)=0$ whenever $h_1,\dots,h_q$ contains a repeated element. Alternatively, the space of square-free tensors in $\mathcal{H}^{\odot q}$ is the span of $\{\Phi_S\}$ where S ranges over all *sets* (not multisets). It is easy to check that if $p:\mathbb{R}^n\to\mathbb{R}$ is a multilinear polynomial, then the tensors $\{f_q\}_{0\le q\le d}$ appearing in the decomposition of p in (7.1) are all square-free.

7.2 Itô's multiplication formula

We now state the formula for product of two polynomials $I_p(f)$ and $I_q(g)$ in terms of f and g. For a reference, see Nourdin's survey [28]. First, we define tensor contraction: for $f\in\mathcal{H}^{\otimes p}$, $g\in\mathcal{H}^{\otimes q}$, and $r\le p\wedge q$, define $f\otimes_r g\in\mathcal{H}^{\otimes p+q-2r}$ by

$$(f\otimes_r g)(t_1,\dots,t_{p+q-2r}) = \sum_{1\le z_1,\dots,z_r\le n} f(t_1,\dots,t_{p-r},z_1,\dots,z_r)\cdot g(t_{p-r+1},\dots,t_{p+q-r},z_1,\dots,z_r).$$

The symmetrized contraction product, denoted by $f\,\widetilde{\otimes}_r\,g\in\mathcal{H}^{p+q-2r}$ is defined as the symmetrization of the tensor $f\otimes_r g$:

$$(f\,\widetilde{\otimes}_r\,g)(t_1,\dots,t_{p+q-2r}) = \frac{1}{(p+q-2r)!}\sum_{\sigma}(f\ \otimes_r\ g)(t_{\sigma(1)},\dots,t_{\sigma(p+q-2r)}),$$

where the sum ranges over all permutations on $[p+q-2r]$.

Proposition 7.4 (Itô's multiplication formula)**.**

$$I_p(f)\cdot I_q(g) = \sum_{r=0}^{p\wedge q} r!\cdot\binom{p}{r}\cdot\binom{q}{r}\cdot\frac{\sqrt{(p+q-2r)!}}{\sqrt{p!\cdot q!}}\cdot I_{p+q-2r}(f\,\widetilde{\otimes}_r\,g).$$

We now list a basic fact about contraction products of tensors which shall be helpful later. In interpreting this fact, it is useful to think of tensors as tables which can be "flattened" into matrices. That is, if we fix an orthonormal basis $e_1,\dots,e_n$ of $\mathcal{H}$ then the n^q elements $\{e_{i_1}\otimes\cdots\otimes e_{i_q} : 1 \le 1_1,\dots,i_q \le n\}$ form a basis for $\mathcal{H}^{\otimes q}$. For any $1 \le r < q$, we can view $f \in \mathcal{H}^{\otimes q}$ as a $n^r \times n^{q-r}$ matrix, where the rows are indexed by $[n]^r$, the columns are indexed by $[n]^{q-r}$, and the entry in position $(i_1,\dots,i_r)$, $(j_1,\dots,j_{q-r})$ is the coefficient in f of the basis element $e_{i_1}\otimes\cdots\otimes e_{i_r}\otimes e_{j_1}\otimes\cdots\otimes e_{j_{q-r}}$. The advantage of this point of view is that contraction can be viewed as matrix multiplication: if M_f is the $n^{p-r}\times n^r$ matrix representing $f \in \mathcal{H}^{\otimes p}$ and M_g is the $n^r \times n^{q-r}$ matrix representing $g \in \mathcal{H}^{\otimes q}$ then $M_f M_g$ is the $n^{p-r}\times n^{q-r}$ matrix representing $f \otimes_r g$. (Note that each of these matrix representations depends on the choice of a basis for $\mathcal{H}$, but that the contraction itself does not.)

The following fact follows immediately from the previous paragraph and the Cauchy-Schwarz inequality:

Fact 7.5. *For $f \in \mathcal{H}^{\otimes p}$ and $g \in \mathcal{H}^{\otimes q}$ and $0 \le r \le p \wedge q$, $\|f \otimes_r g\|_F \le \|f\|_F \cdot \|g\|_F$. Consequently, $\|f \widetilde{\otimes}_r g\|_F \le \|f\|_F \cdot \|g\|_F$.*

The following polynomial anticoncentration theorem is due to Carbery and Wright.

Theorem 7.6. [4] *Let $p : \mathbb{R}^n \to \mathbb{R}$ be a degree-d polynomial. Then, for $\varepsilon > 0$,*

$$\sup_{\theta\in\mathbb{R}} \Pr_x[|p(x)-\theta| \le \varepsilon \cdot \sqrt{\mathsf{Var}[p]}] = O(d \cdot \varepsilon^{1/d}).$$

An immediate consequence is that $\Pr_x[|p(x)| > d^{-O(d)}] \ge 1/2$ for a degree-d polynomial p with $\mathsf{Var}[p] = 1$. With this observation and Theorem 4.6, it is easy to deduce the following bound. (See Lemma 5 in [9].)

Theorem 7.7. *Let $a, b : \mathbb{R}^n \to \mathbb{R}$ be degree-d polynomials. Further, $\mathbf{E}_x[a(x) - b(x)] = 0$ and $\mathsf{Var}[a-b] \le (\tau/d)^{3d} \cdot \mathsf{Var}[a]$. Then, $\Pr_x[\mathsf{sign}(a(x)) \neq \mathsf{sign}(b(x))] = O(\tau)$.*

Before we proceed further, we will make a minor simplifying assumption, namely that the polynomials involved in defining f in Theorem 5.2 can be assumed to be multilinear. As we have said earlier, this is because the results in [9] are stated for multilinear polynomials (though it should be easily possible to carry it over to non-multilinear polynomials). In order to show this, we will use the following simple lemma.

Lemma 7.8. *Let $f : \mathbb{R}^n \to [k]$ be a degree-d, (d,ε)-balanced PTF. Then, for any $\varepsilon > 0$ and $t > 0$, there exists a degree-d, $(d, 2\varepsilon)$-balanced multilinear PTF $f_{\mathsf{multi}} : \mathbb{R}^\ell \to \mathbb{R}$ such that*

- *$\|\mathbf{E}_x[f(x)] - \mathbf{E}_x[f_{\mathsf{multi}}(x)]\|_1 \le \varepsilon$,*
- *$|\mathbf{E}_x[\langle f, P_t f\rangle] - \mathbf{E}_x[\langle f_{\mathsf{multi}}, P_t f_{\mathsf{multi}}\rangle]| \le \varepsilon$.*
- *$\Pr_x[x \in \mathsf{Collision}(f_{\mathsf{multi}})] \le \Pr_x[x \in \mathsf{Collision}(f)] + \varepsilon$.*

Here $\ell = n \cdot (k^2/d\varepsilon)^{3d} \cdot d^2$.

Proof. Let $f = \mathsf{PTF}(p^{(1)}, \ldots, p^{(k)})$, where $p^{(1)}, \ldots, p^{(k)}$ are (d, ε)-balanced. Let $p^{(i)} = \sum_{q=0}^{d} I_q(f_q^{(i)})$ for $f_q^{(i)} \in \mathcal{H}^{\otimes q}$. For some $T \in \mathbb{N}$ which will be fixed later, consider a collection of $\ell = n \cdot T$ variables denoted by $\{x_{i,j}\}_{1 \le i \le n, 1 \le j \le T}$. For $s = 1, \ldots, k$, we define the polynomial $r^{(s)}$ (in the variables $\{x_{i,j}\}$) by

$$r^{(s)}(\{x_{i,j}\}) = p^{(s)}\left(\frac{x_{1,1} + \cdots + x_{1,T}}{\sqrt{T}}, \ldots, \frac{x_{n,1} + \cdots + x_{n,T}}{\sqrt{T}}\right).$$

In terms of the tensor representations, if $f_q^{(i)} = \sum_{1 \le j_1, \ldots, j_i \le n} \alpha_{j_1, \ldots, j_i}^{(q)} e_{j_1} \otimes \cdots \otimes e_{j_i}$ and we define $g_q^{(i)}$ by $r^{(i)} = \sum_{q=0}^{d} I_q(g_q^{(i)})$ then

$$g_q^{(i)} = \sum_{1 \le j_1, \ldots, j_q \le n} \alpha_{j_1, \ldots, j_q}^{(i)} \frac{e_{j_1,1} + \cdots + e_{j_1,T}}{\sqrt{T}} \otimes \cdots \otimes \frac{e_{j_q,1} + \cdots + e_{j_q,T}}{\sqrt{T}}.$$

Next, define $f' : \mathbb{R}^\ell \to [k]$ as $f' = \mathsf{PTF}(r^{(1)}, \ldots, r^{(k)})$. Since the joint distribution of $r^{(1)}, \ldots, r^{(k)}$ under γ_ℓ equals the joint distribution of $p^{(1)}, \ldots, p^{(k)}$ under γ_n, we have that for $1 \le i \le k$,

$$\begin{aligned} \Pr_{x \sim \gamma_n}[f(x) = i] &= \Pr_{x \sim \gamma_\ell}[f'(x) = i], \\ \mathbf{E}_{x \sim \gamma_n}[\langle f, P_t f \rangle] &= \mathbf{E}_{x \sim \gamma_\ell}[\langle f', P_t f' \rangle], \\ \text{and } \Pr_{x \sim \gamma_n}[x \in \mathsf{Collision}(f)] &= \Pr_{x \sim \gamma_\ell}[x \in \mathsf{Collision}(f')]. \end{aligned} \tag{7.2}$$

To construct a multilinear polynomial, let $h_q^{(i)}$ be the orthogonal projection of $g_q^{(i)}$ on the space of square-free tensors, and define the polynomials $w^{(1)}, \ldots, w^{(k)}$ by $w^{(i)} = \sum_{q=0}^{d} I_q(h_q^{(i)})$. Finally, we set $f_{\mathsf{multi}} : \mathbb{R}^\ell \to [k]$ as $f_{\mathsf{multi}} = \mathsf{PTF}(w^{(1)}, \ldots, w^{(k)})$ and we claim that it has the desired properties. By its construction, f_{multi} is a degree-d multilinear PTF. Next, for $1 \le j_1, \ldots, j_q \le n$, let us define $\mathcal{S}_{j_1, \ldots, j_q}$ as

$$\mathcal{S}_{j_1, \ldots, j_q} = \{(s_1, \ldots, s_q) \in [T]^q : |\cup_{b=1}^{q} (j_b, s_b)| < q\}.$$

With this definition, observe that

$$r^{(i)} - w^{(i)} = \sum_{q=1}^{d} \sum_{1 \le j_1, \ldots, j_q \le n} \sum_{(s_1, \ldots, s_q) \in \mathcal{S}_{j_1, \ldots, j_q}} \alpha_{j_1, \ldots, j_q}^{(i)} \cdot \frac{1}{T^{q/2}} \cdot I_q\big(\otimes_{u=1}^{q} e_{j_u, s_u}\big).$$

As the tensors $\{\otimes_{u=1}^{T} e_{j_u, s_u}\}_{j_u \in [n], s_u \in [T]}$ are orthonormal, using Proposition 7.3, we get that

$$\mathsf{Var}[r^{(i)} - w^{(i)}] = \sum_{q=1}^{d} \sum_{1 \le j_1, \ldots, j_q \le n} (\alpha_{j_1, \ldots, j_q}^{(i)})^2 \cdot \frac{|\mathcal{S}_{j_1, \ldots, j_q}|}{T^q}.$$

On the other hand,

$$\mathsf{Var}[r^{(i)}] = \sum_{q=1}^{d} \sum_{1 \le j_1, \ldots, j_q \le n} (\alpha_{j_1, \ldots, j_q}^{(i)})^2.$$

It is easy to see that $|\mathcal{S}_{j_1,\dots,j_q}| \le |\mathcal{S}_{1,\dots,1}| \le q^2 T^{q-1} \le d^2 T^{q-1}$. Thus,

$$\mathsf{Var}[r^{(i)} - w^{(i)}] \le \mathsf{Var}[r^{(i)}] \cdot \frac{d^2}{T}.$$

By applying Theorem 7.7, we get that if $T \ge (\tau/d)^{3d} \cdot d^2$, then

$$\Pr_{x \sim \gamma_\ell}[\mathsf{sign}(r^{(i)}) \neq \mathsf{sign}(w^{(i)})] \le \tau.$$

By a union bound, it follows that $\Pr_{x\sim\gamma_\ell}[f_{\mathsf{multi}}(x) \neq f'(x)] \le k \cdot \tau$ and

$$\Pr_{x\sim\gamma_\ell}[x \in \mathsf{Collision}(f_{\mathsf{multi}})] \le \Pr_{x\sim\gamma_\ell}[x \in \mathsf{Collision}(f')] + k \cdot \tau.$$

Combining this with (7.2) and setting $\tau = \varepsilon/k$ gives us the claim. Also, by an application of Cauchy-Schwarz inequality, we have that

$$\sqrt{\mathsf{Var}(w^{(i)})} \ge \sqrt{\mathsf{Var}(r^{(i)})} \cdot \left(1 - \frac{d}{\sqrt{T}}\right) \ge \sqrt{\mathsf{Var}(r^{(i)})} \cdot \big(1 - \varepsilon\big).$$

Hence, by rescaling the polynomials $w^{(i)}$, we can ensure that f_{multi} is a $(d, 2\varepsilon)$-balanced PTF. ☐

7.3 Eigenvalues of tensors and polynomials

We will now define the notion of eigenvalues of tensors and the corresponding polynomials.

Definition 7.9. Let $f \in \mathcal{H}^{\odot d}$ and $d \ge 2$. For a partition of $[d]$ into $S, \overline{S}$, define

$$\lambda_{S,\overline{S}}(f) = \max_{g \in \mathcal{H}^{\otimes |S|}, h \in \mathcal{H}^{\otimes |\overline{S}|}} \frac{\langle f, g \otimes h\rangle}{\|g\|_F \cdot \|h\|_F}.$$

Define $\lambda_{\max}(f)$ by

$$\lambda_{\max}(f) = \max_{S: 0 < |S| < d} \lambda_{S,\overline{S}}(f).$$

Definition 7.10. Let $p = \sum_{q=0}^{d} I_q(f_q)$ be a multilinear polynomial of degree $d > 0$. and let $(f_0, \dots, f_d) \in \mathcal{H}^{\odot 0} \times \cdots \times \mathcal{H}^{\odot d}$ be the tensor associated with p. Define

$$\lambda_{\max}(p) = \max_{1 < q \le d} \lambda_{\max}(f_q).$$

(Note that we exclude $\lambda_{\max}(f_1)$.) Further, we say that p is δ-eigenregular if $\frac{\lambda_{\max}(p)}{\sqrt{\mathrm{Var}(p)}} \le \delta$.

It is easy to check that eigenvalues of tensors are invariant under unitary transformations, from which it follows that eigenregularity is also unaffected by certain changes of variables:

Fact 7.11. *If $p(x)$ is δ-eigenregular then for any $0 < \rho < 1$, $p(\rho \cdot x + \sqrt{1-\rho^2} y)$ is also δ-eigenregular.*

The next fact states that the contraction product with an eigenregular tensor is significantly contractive.

Fact 7.12. *Let $f \in \mathcal{H}^{\odot d}$ satisfy $\lambda_{\max}(f) \leq \kappa$. For any $g \in \mathcal{H}^{\otimes d'}$ and any $0 < r \leq \min\{d-1, d'\}$, $\|f \otimes_r g\|_F \leq \kappa$.*

Proof. For a fixed orthonormal basis of $\mathcal{H}$, we represent f and g with the matrices M_f and M_g as discussed above, where M_f is an $n^{d-r} \times n^r$ matrix and M_g is $n^r \times n^{d'-r}$. The condition $\lambda_{\max}(f) \leq \kappa$ implies that $\|M_f\|_{op} \leq \kappa$, where $\|\cdot\|_{op}$ denotes the operator norm. For $\mathcal{S} \in [n]^{d'-r}$, let $M_{g,\mathcal{S}}$ denote the corresponding column of M_g. Then the corresponding column of $M_f \cdot M_g$ is given by $M_f M_{g,\mathcal{S}}$. Thus,

$$\|M_f \cdot M_g\|_F^2 = \sum_{\mathcal{S} \in [n]^{d'-r}} \|M_f \cdot M_{g,\mathcal{S}}\|_F^2 \leq \sum_{\mathcal{S} \in [n]^{d'-r}} \kappa^2 \cdot \|M_{g,\mathcal{S}}\|_F^2 = \kappa^2 \cdot \|M_g\|_F^2.$$

This finishes the proof. □

We next have the following easy proposition which says the sum of copies of the same polynomial over disjoint variables is eigenregular.

Fact 7.13. *Let $p : \mathbb{R}^n \to \mathbb{R}$ and define $\tilde{p} : \mathbb{R}^{nk} \to \mathbb{R}$ by*

$$\tilde{p}(X_1, \ldots, X_k) = \frac{1}{\sqrt{k}} \cdot \big(p(X_1) + \cdots + p(X_k)\big),$$

where each $X_i = (x_{i,1}, \ldots, x_{i,n})$ is a disjoint "block" of n variables. Then $\tilde{p}$ is $\frac{1}{\sqrt{k}}$-eigenregular.

Proof. Let d be the degree of $\tilde{p}$ and define the tensors f_q by $q(X_1, \ldots, X_k) = \sum_{q=0}^{d} I_q(f_q)$. Let $\mathcal{H} = \mathbb{R}^{nk}$ and let $\mathcal{H} = \mathcal{H}_1 \oplus \cdots \oplus \mathcal{H}_k$ where $\mathcal{H}_i = \mathbb{R}^n$ corresponds to the coordinates in X_i. Note that f_q has a block diagonal structure on $\mathcal{H}_1 \oplus \cdots \oplus \mathcal{H}_k$, in the sense that $f_q(h_1, \ldots, h_q) = 0$ whenever two of the h_i belong to different components $\mathcal{H}_i$.

For any $g = \mathcal{H}^{\otimes q}$, let $g^{(i)} \in \mathcal{H}_i^{\otimes q}$ denote the ith "diagonal block" of g, defined by

$$g^{(i)}(h_1, \ldots, h_q) = g(\iota(h_1), \ldots, \iota(h_q)),$$

where $\iota : \mathcal{H}_i \to \mathcal{H}$ is the inclusion map. The definition of $\tilde{p}$ ensures that $\|f_q^{(1)}\|_F^2 = \cdots = \|f_q^{(q)}\|_F^2$ for every q, and the fact that f_q is block diagonal ensures that $\|f_q\|_F^2 = \sum_{i=1}^{k} \|f_q^{(i)}\|_F^2 = k\|f_q^{(1)}\|_F^2$. Assuming without loss of generality that $\mathrm{Var}(\tilde{p}) = \mathrm{Var}(p) = 1$, we have $\|f_q^{(i)}\|_F \leq k^{-1/2}\|f_q\|_F^2 \leq k^{-1/2}$ for every i.

For any non-trivial partition $S,\overline{S}$ of $[q]$, we have

$$\begin{aligned}
\lambda_{S,\overline{S}}(f_q) &= \max_{g\in\mathcal{H}^{\odot S},h\in\mathcal{H}^{\odot\overline{S}}} \frac{\langle f_q, g\otimes h\rangle}{\|g\|_F\cdot\|h\|_F} \\
&= \max_{g\in\mathcal{H}^{\odot S},h\in\mathcal{H}^{\odot\overline{S}}} \frac{\sum_{i=1}^q \langle f_q^{(i)}, g^{(i)}\otimes h^{(i)}\rangle}{\|g\|_F\cdot\|h\|_F} \\
&\le \max_{g\in\mathcal{H}^{\odot S},h\in\mathcal{H}^{\odot\overline{S}}} \frac{\sum_{i=1}^q \langle f_q^{(i)}, g^{(i)}\otimes h^{(i)}\rangle}{\sqrt{\sum_{i=1}^q \|g^{(i)}\|_F^2}\cdot\sqrt{\sum_{i=1}^q \|h^{(i)}\|_F^2}} \\
&\le \max_{g\in\mathcal{H}^{\odot S},h\in\mathcal{H}^{\odot\overline{S}}} \frac{\sum_{i=1}^q \frac{1}{\sqrt{k}}\|g^{(i)}\otimes h^{(i)}\|_F}{\sqrt{\sum_{i=1}^q \|g^{(i)}\|_F^2}\cdot\sqrt{\sum_{i=1}^q \|h^{(i)}\|_F^2}} \\
&= \max_{g\in\mathcal{H}^{\odot S},h\in\mathcal{H}^{\odot\overline{S}}} \frac{\sum_{i=1}^q \frac{1}{\sqrt{k}}\|g^{(i)}\|_F\cdot\|h^{(i)}\|_F}{\sqrt{\sum_{i=1}^q \|g^{(i)}\|_F^2}\cdot\sqrt{\sum_{i=1}^q \|h^{(i)}\|_F^2}} \le \frac{1}{\sqrt{k}},
\end{aligned}$$

where the last equality uses the fact that $\|A\otimes B\|_F = \|A\|_F\cdot\|B\|_F$, and the last inequality is a consequence of Cauchy-Schwarz. This implies that $\lambda_{\max}(f_q)\le 1/\sqrt{k}$ for all q, and hence $\lambda_{\max}(\tilde{p})\le 1/\sqrt{k}$. □

One of the two main results that we will use from [9] is a Central Limit Theorem for Gaussian polynomials:

Theorem 7.14. *Let $d>1$ and let $p_1,\dots,p_t:\mathbb{R}^n\to\mathbb{R}$ be t degree-d polynomials such that for $1\le i\le t$, $\mathbf{E}[p_i]=0$ and $\mathsf{Var}(p_i)\le 1$ and each p_i is ε-eigenregular. Let $F=(p_1(x),\dots,p_t(x))$ where $x\sim\gamma_n$ and let C denote the covariance matrix of F, i. e., $C[i,j]=\mathbf{E}_{x\sim\gamma_n}[p_i(x)\cdot p_j(x)]$. Let $Z=(Z_1,\dots,Z_t)$ be a t-dimensional Gaussian with mean zero and covariance C. Then, for any $\alpha:\mathbb{R}^t\to\mathbb{R}$ such that $\alpha\in\mathcal{C}^2$,*

$$\left|\mathbf{E}_{x\sim\gamma_n}\left[\alpha(p_1(x),\dots,p_t(x))\right]-\mathbf{E}_Z\left[\alpha(Z_1,\dots,Z_t)\right]\right|\le 2^{O(d\log d)}\cdot t^2\cdot\sqrt{\varepsilon}\cdot\|\alpha''\|_\infty.$$

7.4 Polynomial regularity lemma

The next result we will need from [9] is a regularity lemma for low-degree polynomials. Suppose that we begin with a collection of $k(d+1)$ polynomials $\{p_{s,q}:1\le s\le k, 0\le q\le d\}$, where $p_{s,q}\in\mathcal{W}^q$ and $\mathrm{Var}[p_{s,q}]=1$ for all s and q. We will consider "decompositions" of these polynomials consisting of the following objects:

- for every $1\le s\le k$ and $0\le q\le d$, an "outer" polynomial $\mathrm{Out}(p_{s,q})$ of arity $\mathrm{num}(s,q)\in\mathbb{N}$; and
- for every $q\le s\le k$, $0\le q\le d$, and $1\le\ell\le\mathrm{num}(s,q)$, an "inner" polynomial $\mathrm{In}(p_{s,q})_\ell$.

We also introduce the following statistics of these polynomials:

- the "total arity," $\mathrm{Num}=\sum_{s=1}^k\sum_{q=0}^d\mathrm{num}(s,q)$;
- $\mathrm{Coeff}(p_{s,q})$, defined to be the sum of the absolute values of the coefficients of $\mathrm{Out}(p_{s,q})$; and

- $\mathrm{Coeff} = \sum_{s=1}^{k} \sum_{q=0}^{d} \mathrm{Coeff}(p_{s,q})$.

Definition 7.15. Consider polynomials $\{p_{s,q}\}$ and a function $\beta : [1,\infty) \to (0,1)$ satisfying $\beta(x) \le 1/x$. Given collections of polynomials $\{\mathrm{Out}(p_{s,q})$ and $\{\mathrm{In}(p_{s,q})_\ell\}$ as above, we say that these polynomials are a (β, M, N) decomposition of $\{p_{s,q}\}$ if the following conditions hold:

1. for every $1 \le s \le k$, every $0 \le q \le d$, and every $x \in \mathbb{R}^n$,
$$p_{s,q}(x) = \mathrm{Out}(p_{s,q})\left(\mathrm{In}(p_{s,q})_1(x), \ldots, \mathrm{In}(p_{s,q})_{\mathrm{num}(s,q)}(x)\right); \tag{7.3}$$
2. for every $1 \le s \le k$, every $1 \le q \le d$, and every $1 \le \ell \le \mathrm{num}(s,q)$, the polynomial $\mathrm{In}(p_{s,q})_\ell$ lies in $\mathcal{W}^j$ for some $1 \le j \le q$ and satisfies $\mathrm{Var}[\mathrm{In}(p_{s,q})_\ell] = 1$;
3. for every $1 \le s \le k$ and every $1 \le q \le d$, $\mathrm{Out}(p_{s,q})$ is a multilinear polynomial of degree at most d;
4. $\mathrm{Coeff} \le M$ and $\mathrm{Num} \le N$; and
5. for every $1 \le s \le k$, every $1 \le q \le d$, and every $1 \le \ell \le \mathrm{num}(s,q)$, $\mathrm{In}(p_{s,q})_\ell$ is $\beta(\mathrm{Num} + \mathrm{Coeff})$-eigenregular.

Intuitively, the above definition states the following: given a decreasing function $\beta(\cdot)$ and a collection of k multilinear polynomials of degree d, a (β, M, N) decomposition of expresses the polynomials as a "outer polynomial" composed with "inner polynomials." The quality of this decomposition is captured by the following parameters:

- the sum of the arities (denoted by Num, and bounded by N) and the sum of absolute values of coefficients (denoted by Coeff, and bounded by M) of the outer polynomials; and
- the eigenregularity of the inner polynomial (bounded in terms of β).

The main decomposition theorem of [9] says that approximate (β, M, N) decompositions exist and can be computed.

Theorem 7.16. *Fix $d \ge 2$ and fix any non-increasing computable function $\beta : [1,\infty) \to (0,1)$ that satisfies $\beta(x) \le 1/x$. There is an algorithm* **MultiRegularize-Many-Wieners**$_{d,\beta}$ *with the following properties. The algorithm takes as input:*

- *a collection $\{p_{s,q} : 1 \le s \le k, 0 \le q \le d\}$ of multilinear polynomials satisfying $p_{s,q} \in \mathcal{W}^q$ and $\mathrm{Var}[p_{s,q}] = 1$ for every s, q; and*
- *a parameter $\tau > 0$.*

The output of the algorithm consists of

- *a collection $\{\tilde{p}_{s,q} : 1 \le s \le k, 0 \le q \le d\}$ of polynomials satisfying (for every s and q) $\tilde{p}_{s,q} \in \mathcal{W}^q$, $\tilde{p}_{s,0} = p_{s,0}$, and $\mathrm{Var}[p_{s,q} - \tilde{p}_{s,q}] \le \tau$;*
- *two numbers $M = M_\beta(k, d, \tau)$ and $N = N_\beta(k, d, \tau)$; and*

- *a* (β, M, N) *decomposition of* $\{\tilde{p}_{s,q}\}$.

The functions $M_\beta(k,d,\tau)$ *and* $N_\beta(k,d,\tau)$ *are computable but not necessarily primitive recursive.*

We remark that our use of Theorem 7.16 from [9] is precisely the reason why the quantity n_0 in Theorem 2.2 is not primitive recursive.

7.5 Second moments and approximate equivalence

As we mentioned in the outline of the proof of Theorem 5.2, one of the main steps in the proof is (after applying a (β, M, N) decomposition) to swap out one collection of inner polynomials for a different one. The main result of this section essentially says that in order to do so, it suffices to preserve the second moments of all the inner polynomials:

Lemma 7.17. *Let* $\Psi : \mathbb{R}^{\mathsf{Num}_\Psi} \to \mathbb{R}$ *be a degree-d polynomial, and let* Coeff_Ψ *be the sum of the absolute values of its coefficients. Let* $\{A_i\}_{1\le i\le \mathsf{Num}_\Psi}$ *and* $\{B_i\}_{1\le i\le \mathsf{Num}_\Psi}$ *be centered families of polynomials of variance* 1 *such that for all* $1 \le i \le m$, $\exists 0 < j \le d$ *such that* $A_i, B_i \in \mathcal{W}^j$. *Assume, moreover, that for all* $1 \le i \le j \le \mathsf{Num}_\Psi$, $\mathbf{E}[A_iA_j] = \mathbf{E}[B_iB_j]$. *If each* A_i, B_i *is* ζ*-eigenregular where* $\zeta \cdot \mathsf{Coeff}_\Psi \cdot 2^{\mathsf{Num}_\Psi \cdot (\mathsf{Num}_\Psi \cdot d+1)} \le \kappa$, *then* $|\mathbf{E}[\Psi(A_1,\dots,A_{\mathsf{Num}})] - \mathbf{E}[\Psi(B_1,\dots,B_{\mathsf{Num}})]| \le \kappa$.

Our main tool in the proof of Lemma 7.17 is the following lemma, which says that the expectation of a product of eigenregular polynomials is essentially determined by their second moments.

Lemma 7.18. *Let* $p_1,\dots,p_t \in \mathbb{R}^n \to \mathbb{R}$ *be* κ*-eigenregular polynomials with variance 1, such that* $p_i \in \mathcal{W}^{q_i}$ *for every i. Define* $C \in \mathbb{R}^{t\times t}$ *by* $C(i,j) = \langle p_i, p_j \rangle$. *There exists a function* $F : \mathbb{R}^{t\times t} \times \mathbb{Z}^t \to \mathbb{R}$ *such that*

$$\left| \mathbf{E}\Big[\prod_{i=1}^t p_i\Big] - F(C, q_1,\dots,q_t) \right| \le 2^{t(q_1+\cdots+q_t+1)} \cdot \kappa.$$

In other words, up to the error term of $2^{t(q_1+\cdots+q_t)} \cdot \kappa$, *the expectation of the product* $\prod_{i=1}^t p_i$ *is just dependent on the degrees of the polynomials and the covariance matrix of the polynomials.*

Proof of Lemma 7.17. Applying Lemma 7.18 and triangle inequality, we have

$$|\mathbf{E}[\Psi(A_1,\dots,A_{\mathsf{Num}})] - \mathbf{E}[\Psi(B_1,\dots,B_{\mathsf{Num}})]| \le \zeta \cdot \mathsf{Coeff}_\Psi \cdot 2^{\mathsf{Num}_\Psi \cdot (\mathsf{Num}_\Psi \cdot d+1)}.$$

Applying the assumption on ζ completes the proof. □

Before proving Lemma 7.18, we will explore some further consequences. The following lemma is along the same lines as Lemma 7.17, but it concerns the sign of Ψ instead of just its expectation.

Lemma 7.19. *Fix* $\varepsilon > 0$ *and* $\eta > 0$. *Let* Ψ, A_i, *and* B_i *satisfy the assumptions of the previous lemma, but assume that* A_i *and* B_i *are* ζ*-eigenregular for* ζ *satisfying*

$$\zeta \cdot \mathsf{Coeff}_\Psi^2 \cdot 2^{\mathsf{Num}_\Psi \cdot (\mathsf{Num}_\Psi \cdot 2d+1)} \cdot 2^{O(d^3)} \le \frac{\eta}{4}, \;\; 2^{O(d \log d)} \cdot \mathsf{Num}_\Psi^2 \cdot (4c^2) \cdot \zeta \le \varepsilon,$$

where c is defined by

$$B = \Omega\left(\ln \frac{\mathsf{Num}_\Psi \cdot d}{\varepsilon}\right)^{d/2};\ \delta = \left(\frac{\varepsilon \cdot \eta^{1/2d}}{d \cdot B \cdot \sqrt{\mathsf{Num}_\Psi} \cdot \mathsf{Coeff}_\Psi^{1/d}}\right)^d;\ c = \frac{\mathsf{Num}_\Psi}{\delta \cdot \sqrt{\varepsilon}}.$$

If $\mathsf{Var}(\Psi(A_1,\ldots,A_{\mathsf{Num}_\Psi})) \geq \eta$ *and* $|\mathbf{E}[\Psi(A_1,\ldots,A_{\mathsf{Num}_\Psi})]| \leq \eta$ *then*

$$|\mathbf{E}[\mathsf{sign}(\Psi(A_1,\ldots,A_{\mathsf{Num}}))] - \mathbf{E}[\mathsf{sign}(\Psi(B_1,\ldots,B_{\mathsf{Num}}))]| \leq O(\varepsilon).$$

One step in the proof of Lemma 7.19 is a kind of mollification for the sign function:

Lemma 7.20. *Fix $\varepsilon > 0$, $\eta > 0$. Let Ψ be a degree-d polynomial of arity Num_Ψ, and let B, δ, and c be defined as in Lemma 7.19.*

There is a function $\tilde{g}_c : \mathbb{R}^{\mathsf{Num}_\Psi} \to [0,1]$ satisfying $\|\tilde{g}_c''\|_\infty \leq 4c^2$ such that for every collection $\{A_i\}_{1\leq i\leq \mathsf{Num}_\Psi}$ of centered, unit-variance polynomials in n variables that satisfy $\mathsf{Var}(\Psi(A_1,\ldots,A_m)) \geq \eta$, we have

$$\mathop{\mathbf{E}}_{x\sim\gamma_n}[|\mathsf{sign}(\Psi(A_1,\ldots,A_m)) - \tilde{g}_c(A_1,\ldots,A_m)|] \leq O(\varepsilon).$$

Proof. Let R denote the set $\{x : \mathsf{sign}(\Psi(x)) = 1\}$. By a standard argument based on convolution with a mollifier, there is a function $\tilde{g}_c : \mathbb{R}^{\mathsf{Num}_\Psi} \to [0,1]$ such that $\|\tilde{g}_c''\|_\infty \leq 4c^2$ and

$$\left|\mathsf{sign}(\Psi(x)) - \tilde{g}_c(x)\right| \leq \min\left\{1, \frac{\mathsf{Num}_\Psi^2}{c^2 \cdot \mathsf{dist}^2(x,\partial R)}\right\} \tag{7.4}$$

for all x, where $\mathsf{dist}(x,\partial R)$ denotes the Euclidean distance of x from the boundary of R.

It is not hard to check (and follows immediately from Claim 49 in [9]) that if $x \in \mathbb{R}^{\mathsf{Num}_\Psi}$ satisfies $\|x\|_\infty \leq B$ and $|\phi(x)| \geq d\cdot(2B)^d \cdot \delta \cdot \mathsf{Coeff}_\Psi \cdot \mathsf{Num}_\Psi^{d/2}$ then $\mathsf{dist}(x,\partial R) \geq \delta$, in which case (7.4) implies that $|\mathsf{sign}(\Psi(x)) - \tilde{g}_c(x)| \leq \mathsf{Num}_\Psi^2/(c^2\delta^2)$. Rearranging the logic, if $|\mathsf{sign}(\Psi(x)) - \tilde{g}_c(x)| > \mathsf{Num}_\Psi^2/(c^2\delta^2)$ then either $\max_j |A_j(x)| > B$ or $\Psi(A_1(x),\ldots A_{\mathsf{Num}_\Psi}(x)) \leq d(2B)^d\delta\cdot\mathsf{Coeff}_\Psi\cdot\mathsf{Num}_\Psi^{d/2}$. Hence,

$$\begin{aligned}
&\mathbf{E}_{x\sim\gamma_n}[|\mathsf{sign}(\phi(A_1,\ldots,A_{\mathsf{Num}_\Psi})) - \tilde{g}_c(A_1,\ldots,A_{\mathsf{Num}_\Psi})|] \\
&\quad \leq \Pr_{x\sim\gamma_n}\left[\max_j |A_j| > B\right] + \Pr_{x\sim\gamma_n}\left[|\phi(A_1,\ldots,A_{\mathsf{Num}_\Psi})| \leq d\cdot(2B)^d\cdot\delta\cdot\mathsf{Coeff}_\Psi\cdot\mathsf{Num}_\Psi^{d/2}\right] + \frac{\mathsf{Num}_\Psi^2}{c^2\delta^2}.
\end{aligned}$$

We now bound each term on the right hand side. For the first term, each A_j is a centered, unit-variance, degree-d polynomial and so Theorem 4.6 gives

$$\Pr_{x\sim\gamma_n}\left[\max_j |A_j(x)| > B\right] \leq \mathsf{Num}_\Psi \cdot d \cdot \exp\left(-B^{2/d}\right).$$

To bound the second term, apply Theorem 7.6 and the assumption that $\mathsf{Var}(\Psi(A_1,\ldots,A_{\mathsf{Num}_\Psi})) \geq \eta$:

$$\Pr_{x\sim\gamma_n}\left[|\phi(A_1,\ldots,A_m)| \leq d\cdot(2B)^{d/2}\cdot\delta\cdot\mathsf{Coeff}_\Psi\cdot\mathsf{Num}_\Psi^{d/2}\right] \leq \frac{d\cdot(2B)\cdot\delta^{1/d}\cdot\mathsf{Coeff}_\Psi^{1/d}\cdot\sqrt{\mathsf{Num}_\Psi}}{\eta^{1/2d}}.$$

Plugging in our choice of parameters completes the proof. □

Proof of Lemma 7.19. Consider the polynomial $\Upsilon(z_1,\ldots,z_{\mathsf{Num}_\Psi}) = (\Psi(z_1,\ldots,z_{\mathsf{Num}_\Psi}))^2$; let Num_Υ be its arity and let Coeff_Υ be the sum of the absolute values of its coefficients. Observe that Υ is a degree-$2d$ polynomial, and that $\mathsf{Coeff}_\Upsilon \le \mathsf{Coeff}_\Psi^2$ and $\mathsf{Num}_\Upsilon = \mathsf{Num}_\Psi$. Now, applying Lemma 7.17 to the function Υ,

$$\left|\mathbf{E}[\Psi^2(A_1,\ldots,A_{\mathsf{Num}})] - \mathbf{E}[\Psi^2(B_1,\ldots,B_{\mathsf{Num}})]\right| \le \eta/4. \tag{7.5}$$

By Lemma 7.17 applied to the function Ψ,

$$\begin{aligned}\left|\mathbf{E}[\Psi(A_1,\ldots,A_{\mathsf{Num}})] - \mathbf{E}[\Psi(B_1,\ldots,B_{\mathsf{Num}})]\right| &\le \zeta\cdot\mathsf{Coeff}_\Psi\cdot 2^{\mathsf{Num}_\Psi\cdot(\mathsf{Num}_\Psi\cdot d+1)}\\ &\le \frac{\eta}{4\cdot\mathsf{Coeff}_\Psi\cdot 2^{\mathsf{Num}_\Psi^2\cdot d}\cdot 2^{O(d^3)}}.\end{aligned} \tag{7.6}$$

Hölder's inequality and the Gaussian hypercontractive inequality (see Theorem 5.10 in [18]) imply that for any monomial in A_i, $\mathbf{E}[\prod_{j=1}^d A_{i_j}] \le \prod_{j=1}^d \mathbf{E}[|A_{i_j}|^d]^{1/d} \le d^{2d} \le 2^{O(d^2)}$; by the triangle inequality, it follows that $\mathbf{E}[\Psi(A_1,\ldots,A_{\mathsf{Num}})] \le \mathsf{Coeff}_\Psi 2^{O(d^3)}$. The same bound applies with B in place of A, and together with (7.6) this yields

$$\left|\left(\mathbf{E}[\Psi(A_1,\ldots,A_{\mathsf{Num}})]\right)^2 - \left(\mathbf{E}[\Psi(B_1,\ldots,B_{\mathsf{Num}})]\right)^2\right| \le \eta/4. \tag{7.7}$$

Combining (7.5) and (7.7), we get that $|\mathsf{Var}(\Psi(A_1,\ldots,A_{\mathsf{Num}})) - \mathsf{Var}(\Psi(A_1,\ldots,A_{\mathsf{Num}}))| \le \eta/2$ and thus,

$$\mathsf{Var}(\Psi(B_1,\ldots,B_{\mathsf{Num}})) \ge \eta/2.$$

Applying Lemma 7.20, we obtain a function $\tilde{g}_c : \mathbb{R}^{\mathsf{Num}} \to [0,1]$ such that $\|\tilde{g}_c''\|_\infty \le 4c^2$,

$$\begin{aligned}&\left|\mathbf{E}[\mathsf{sign}(\Psi(A_1,\ldots,A_{\mathsf{Num}}))] - \mathbf{E}[\tilde{g}_c(\Psi(A_1,\ldots,A_{\mathsf{Num}}))]\right| \le O(\varepsilon), \text{ and}\\ &\left|\mathbf{E}[\mathsf{sign}(\Psi(B_1,\ldots,B_{\mathsf{Num}}))] - \mathbf{E}[\tilde{g}_c(\Psi(B_1,\ldots,B_{\mathsf{Num}}))]\right| \le O(\varepsilon).\end{aligned}$$

Finally, applying Theorem 7.14, we have

$$\left|\mathbf{E}[\tilde{g}_c(\Psi(B_1,\ldots,B_{\mathsf{Num}}))] - \mathbf{E}[\tilde{g}_c(\Psi(A_1,\ldots,A_{\mathsf{Num}}))]\right| \le 2^{O(d\log d)}\cdot\mathsf{Num}_\Psi^2\cdot(4c^2)\cdot\zeta \le \varepsilon.$$

The last inequality uses the second condition on ζ. Combining these three inequalities, we obtain the proof. □

In the remainder of this section, we prove Lemma 7.18. Let p_i and q_i be as in the statement of Lemma 7.18, and let $h_i \in \mathcal{H}^{\otimes q_i}$ be defined by $p_i = I_{q_i}(h_i)$. To prove the lemma, we will express the product $p_1\cdot\dots\cdot p_t$ in terms of the tensors $\{h_1,\ldots,h_t\}$ using successive applications of Itô's multiplication formula. Towards this, we first have the following definition.

Definition 7.21. Given $p_1,\ldots,p_t$ as above, a tuple $\mathbf{r} = (r_1,\ldots,r_{t-1})$ is said to be a contraction sequence if it satisfies the following definition:

$$\text{For all } 1 \le i \le t-1,\ 0 \le r_i \le \min\left\{q_{i+1}, q_i + \sum_{j=1}^{i-1}(q_j - 2r_j)\right\}. \tag{7.8}$$

Let $\mathcal{S}_{\mathsf{valid}}$ be the set of all contraction sequences. Also for a contraction sequence $\mathbf{r} = (r_1, \ldots, r_t)$, define the quantity

$$\gamma_{\mathbf{r}} = \prod_{i=1}^{t-1} r_i! \cdot \binom{q_{i+1}}{r_i} \cdot \binom{q_i + \sum_{j=1}^{i-1}(q_j - 2r_j)}{r_i} \cdot \frac{\sqrt{(q_{i+1} + \sum_{j=1}^{i}(q_j - 2r_j))!}}{\sqrt{q_{i+1}!} \cdot \sqrt{(q_i + \sum_{j=1}^{i-1}(q_j - 2r_j))!}}.$$

Also, let us define $g_1^{\mathbf{r}} = h_1$ and for $1 < i \le t$, define $g_i^{\mathbf{r}}$ iteratively as $g_i^{\mathbf{r}} = g_{i-1}^{\mathbf{r}} \widetilde{\otimes}_{r_{i-1}} h_i$. Define $m_i^{\mathbf{r}} = q_{i+1} + \sum_{j=1}^{i}(q_j - 2r_j)$ and observe that $g_{i+1}^{\mathbf{r}} \in \mathcal{H}^{\odot m_i^{\mathbf{r}}}$. Applying Itô's multiplication formula (Proposition 7.4) iteratively, we obtain that

$$\prod_{i=1}^{t} p_i = \sum_{\mathbf{r}=(r_1,\ldots,r_{t-1}) \in \mathcal{S}_{\mathsf{valid}}} \gamma_{\mathbf{r}} I_{m_{t-1}^{\mathbf{r}}}(g_t^{\mathbf{r}}). \tag{7.9}$$

The next claim establishes a simple upper bound on $\gamma_{\mathbf{r}}$.

Claim 7.22. $\sum_{\mathbf{r}=(r_1,\ldots,r_{t-1}) \in \mathcal{S}_{\mathsf{valid}}} |\gamma_{\mathbf{r}}| \le 2^{t(q_1+\cdots+q_t+1)}$.

Proof. By applying the fact that $\binom{n}{x} \le 2^n$, it is easy to show that

$$r_i! \cdot \binom{q_{i+1}}{r_i} \cdot \binom{q_i + \sum_{j=1}^{i-1}(q_j - 2r_j)}{r_i} \cdot \frac{\sqrt{(q_{i+1} + \sum_{j=1}^{i}(q_j - 2r_j))!}}{\sqrt{q_{i+1}!} \cdot \sqrt{(q_i + \sum_{j=1}^{i-1}(q_j - 2r_j))!}} \le 2^{\sum_{j=1}^{i}(q_j - 2r_j) + q_{i+1} + r_i}.$$

This implies that

$$|\gamma_{\mathbf{r}}| \le 2^{\sum_{i=1}^{t-1}(\sum_{j=1}^{i}(q_j - 2r_j) + q_{i+1} + r_i)} \le 2^{t(q_1+\cdots+q_t) - \sum_{i=1}^{t-1} r_i}.$$

As a result, we obtain

$$\sum_{\mathbf{r}=(r_1,\ldots,r_{t-1}) \in \mathcal{S}_{\mathsf{valid}}} |\gamma_{\mathbf{r}}| \le \sum_{\mathbf{r}=(r_1,\ldots,r_{t-1}) \in \mathcal{S}_{\mathsf{valid}}} 2^{t(q_1+\cdots+q_t) - \sum_{i=1}^{t-1} r_i} \le 2^{t(q_1+\cdots+q_t+1)}.$$ □

Our strategy for proving Lemma 7.18 is as follows. We will first partition the sum in (7.9) into two sets and show that the terms belonging to the first set are all bounded by κ. We will then further partition the terms in the second set into two sets: Again, the terms in the first set will be bounded by κ whereas terms in the second set will just be a function of $(C, q_1, \ldots, q_t)$ which will conclude our proof. Towards this, let us partition $\mathcal{S}_{\mathsf{valid}}$ into two partitions, $\mathcal{S}_z$ and $\mathcal{S}_{\mathsf{valid}} \setminus \mathcal{S}_z$ where the first set is defined as $\mathcal{S}_z = \{\mathbf{r} \in \mathcal{S}_{\mathsf{valid}} : \text{For all } i < t, (r_i = 0) \vee (r_i = q_{i+1})\}$. Next, prove the following proposition.

Proposition 7.23. *Let $h_1, \ldots, h_t$ be as above such that for all $1 \le i \le t$, $\|h_1\|_F = \ldots = \|h_t\|_F = 1$. Further, each h_i is κ-eigenregular. If $\mathbf{r} \notin \mathcal{S}_z$, then $\|g_t^{\mathbf{r}}\|_F \le \kappa$.*

Proof. If $\mathbf{r} \notin \mathcal{S}_z$, then there is a coordinate $1 \le j \le t-1$ such that $r_j \ne 0$ and $r_j \ne q_{j+1}$. Appealing to Fact 7.5, it is easy to see that $\|g_j^{\mathbf{r}}\|_F \le 1$. Next, we appeal to Fact 7.12, we get $\|g_j^{\mathbf{r}} \otimes_{r_j} h_{j+1}\|_F \le \kappa$. By the triangle inequality, symmetrization decreases norm and so we obtain that $\|g_j^{\mathbf{r}} \widetilde{\otimes}_{r_j} h_{j+1}\|_F \le \kappa$. Finally, again using that $\|h_{j+2}\|_F, \ldots, \|h_t\|_F \le 1$ and applying Fact 7.5 iteratively, we obtain $\|g_t^{\mathbf{r}}\|_F \le \kappa$. This finishes the proof. □

To define further partitions of $\mathbb{S}_z$, we will need a somewhat more elaborate definition. To do this, consider any $\mathbf{r} \in \mathbb{S}_z$. Recall that $m_i^{\mathbf{r}} = q_{i+1} + \sum_{j=1}^{i}(q_j - 2r_j)$ and $g_{i+1} \in \mathcal{H}^{\odot m_i^{\mathbf{r}}}$. For any $\ell \in \mathbb{N}$, let $\mathbb{S}_\ell$ denote the symmetric group on ℓ elements. Consider a tuple of permutations $\sigma = (\sigma_1, \ldots, \sigma_{t-1})$ such that $\sigma_i \in \mathbb{S}_{m_i^{\mathbf{r}}}$. For σ and $\mathbf{r}$, we define $g_i^{\mathbf{r},\sigma}$ iteratively as $g_1^{\mathbf{r},\sigma} = h_1$ and $g_i^{\mathbf{r},\sigma} = \sigma_{i-1}(g_{i-1}^{\mathbf{r},\sigma} \otimes_{r_{i-1}} h_i)$. Let $\Lambda_{\mathbf{r}} = \{\sigma : \sigma = (\sigma_1, \ldots, \sigma_{t-1})\}$. With this notation, we have that $g_t^{\mathbf{r}} = (1/|\Lambda_{\mathbf{r}}|) \cdot \sum_{\sigma \in \Lambda_{\mathbf{r}}} g_t^{\mathbf{r},\sigma}$. We will now partition pairs $(\mathbf{r}, \sigma)$ (where $\mathbf{r} \in \mathbb{S}_z$). For this subsequent partitioning, we will need to associate a graph with the pair $(\mathbf{r}, \sigma)$ where $\mathbf{r} \in \mathbb{S}_z$ and then partition based on this graph. To understand what this graph is meant to capture, note that to obtain $g_{i+1}^{\mathbf{r},\sigma}$ from $g_i^{\mathbf{r},\sigma}$, one of the following two events happen: (a) If $r_i = 0$, then we first compute $g_i^{\mathbf{r},\sigma} \otimes h_{i+1}$ and then permute the indices of the resulting tensor by σ_i. Note that each time an index is added, we can associate a unique $1 \le j \le t$ with this index. (b) If $r_i = q_{i+1}$, we compute the contraction product of $g_i^{\mathbf{r},\sigma}$ with h_{i+1} and subsequently permute the indices of the resulting tensor by σ_i. Note that crucially, the "contraction" is along the last r_i indices of $g_i^{\mathbf{r},\sigma}$. If we do a contraction at $i = i_0$, we can associate i_0 with all the indices that got collapsed at $i = i_0$. In other words, when an index gets created, we associate an element of $[t]$ and another one, when it gets annihilated. The graph structure is meant to capture this relation as formalized below.

Description of graph

1. Initialize bipartite graph with $L = R = \phi$. The vertex set L will be ordered at all points in the algorithm.
2. Add q_1 vertices (in order), each colored as '1' to L.
3. For $i = 2$ to $i = t$,
   (a) If $r_{i-1} = 0$, then add q_i vertices to L with color 'i' (in order). Permute the 'unmatched' vertices in L with permutation σ_{i-1}.
   (b) If $r_{i-1} = q_i$, then add q_i vertices to R with color 'i' (in order) and match them to the 'last' q_i vertices in L. Permute the unmatched vertices in L with permutation σ_{i-1}.

The correspondence between the graph constructed above and the contraction process used to generate $g_{i+1}^{\mathbf{r},\sigma}$ is obvious. For example, we have the following easy observation.

Observation 7.24. The dimension of the tensor $g_t^{\mathbf{r},\sigma}$ is precisely $n^{|U|}$ where U is the set of unmatched vertices in L. Consequently, $\mathbf{E}[g_t^{\mathbf{r},\sigma}] \neq 0$ only if $U \neq \phi$.

Since we are interested in the expectation of $\mathbf{E}[\prod_{i=1}^{t} p_i]$, thus from now onwards we only focus on $(\mathbf{r}, \sigma)$ such that there are no unmatched vertices in the graph generated above. Note that the property of having no unmatched vertices is just a property of $\mathbf{r}$ and not σ. The next definition forms the crux of our basis of partitioning the pairs $(\mathbf{r}, \sigma)$.

Definition 7.25. For a pair $(\mathbf{r}, \sigma)$ such that the corresponding graph has no unmatched vertices, we say that $(\mathbf{r}, \sigma)$ is "aligned" if and only each color 'i', there is a unique color 'j' such that vertices of color i are only matched to vertices of color 'j' and vice-versa. Else, it is said to be "unaligned."

We will now prove the following claim.

Claim 7.26. *Let* $(\mathbf{r},\sigma)$ *be an unaligned pair. Then,* $\|g_t^{\mathbf{r},\sigma}\| \leq \kappa$.

Proof. Note that we are assuming that there are no unmatched vertices in the graph. If the pair is $(\mathbf{r},\sigma)$ is unaligned, then there is a unique smallest integer $1 < N \leq t$ such that at the N-th step of the algorithm to form the graph, the pair becomes unaligned. Namely, N is the smallest integer such that one of the following events happen:

Event A: At the beginning of Step (3), when $i = N$, we have $r_{i-1} = q_i$ and the last q_i vertices in L are not of the same color.

Event B: At the beginning of Step (3), when $i = N$, we have $r_{i-1} = q_i$, the last q_i vertices in L are of the same color (say j) but there is also another unmatched vertex in L colored j.

We ask the reader to verify that $(\mathbf{r},\sigma)$ is unaligned if and only if one of Event A or B happen. We will now prove that in both cases A and B, our conclusion holds.

Proof for Event A: Let $\mathcal{J}_1 = \{1\} \cup \{N > i > 1 : r_{i-1} = 0\}$ and $\mathcal{J}_2 = \{N > i > 1 : r_{i-1} \neq 0\}$. By definition of N, it is easy to verify that there is a one-one mapping ν from $\mathcal{J}_2$ to $\mathcal{J}_1$ such that for any $a \in \mathcal{J}_1$, vertices of color a are matched with vertices of color $\nu(a)$ and vice-versa. Let us define $\mathcal{J}_3 = \mathcal{J}_1 \setminus \nu(\mathcal{J}_2)$. As a consequence, we have

$$g_{N-1}^{\mathbf{r},\sigma} = \prod_{a\in\mathcal{J}_2} \langle h_a, h_{\nu(a)}\rangle \cdot \mu(\otimes_{j\in\mathcal{J}_3} h_j),$$

where μ is a permutation of the indices of $\otimes_{j\in\mathcal{J}_3} h_j$. Further, since at $i = N$, the last q_N vertices of L are of at least two different colors, we have that the last q_N indices of $\mu(\otimes_{j\in\mathcal{J}_3} h_j)$ belong to more than one h_j (for $j \in \mathcal{J}_3$). Let A be the set of all indices of $\otimes_{j\in\mathcal{J}_3} h_j$ which do not map to the last q_N positions under μ. Consider some assignment ρ for these indices. Then, we have that

$$\|g_N^{\mathbf{r},\sigma}\|_F^2 = \sum_{\rho} \|\prod_{a\in\mathcal{J}_2} \langle h_a, h_{\nu(a)}\rangle \cdot \mu(\otimes_{j\in\mathcal{J}_3} h_{j,\rho}) \otimes_{q_N} h_N\|_F^2.$$

Here $h_{j,\rho}$ is tensor obtained by restricting the indices in A to ρ. Since $\lambda_{\max}(h_N) \leq \kappa$, we get

$$\|\mu(\otimes_{j\in\mathcal{J}_3} h_{j,\rho}) \otimes_{q_N} h_N\|_F^2 \leq \kappa^2 \cdot \|\otimes_{j\in\mathcal{J}_3} h_{j,\rho}\|_F^2.$$

This implies that

$$\|g_N^{\mathbf{r},\sigma}\|_F^2 \leq \kappa^2 \sum_{\rho} \prod_{a\in\mathcal{J}_2} \langle h_a, h_{\nu(a)}\rangle^2 \cdot \|\otimes_{j\in\mathcal{J}_3} h_{j,\rho}\|_F^2 \leq \kappa^2 \prod_{a\in\mathcal{J}_2} \langle h_a, h_{\nu(a)}\rangle^2 \leq \kappa^2.$$

The penultimate inequality uses the fact that $\sum_{\rho} \|\otimes_{j\in\mathcal{J}_3} h_{j,\rho}\|_F^2 = 1$. Using the fact that $\|g_j\|_F = 1$ for all $j > N$ and applying Fact 7.5, we get that $|g_t^{\mathbf{r},\sigma}| \leq \kappa$. This finishes the proof.

Proof for Event B: Let us define $\mathcal{J}_1$, $\mathcal{J}_2$ and $\mathcal{J}_3$ as in Event A. As before, we have

$$g_{N-1}^{\mathbf{r},\sigma} = \prod_{a\in\mathcal{J}_2} \langle h_a, h_{\nu(a)}\rangle \cdot \mu(\otimes_{j\in\mathcal{J}_3} h_j),$$

where μ is a permutation of the indices of $\otimes_{j\in\mathcal{J}_3} h_j$. The crucial difference is that there exists $j_0 \in \mathcal{J}_3$ such that if $\mathcal{B}$ denotes the set of positions where the indices of h_{j_0} get mapped under μ, then this includes the last q_N indices as a proper subset. Using $\lambda_{\max}(h_{j_0}) \leq \kappa$, we get that $\|\mu(\otimes_{j\in\mathcal{J}_3} h_j) \otimes_{q_N} h_N\|_F^2 \leq \kappa^2 \cdot \|\otimes_{j\in\mathcal{J}_3\setminus\{j_0\}} h_j\|_F^2$. However, $\|\otimes_{j\in\mathcal{J}_3\setminus\{j_0\}} h_j\|_F^2 = 1$, and hence we get $\|g_N^{\mathbf{r},\sigma}\|_F^2 \leq \kappa^2$. Using the fact that $\|g_j\|_F = 1$ for all $j > N$ and applying Fact 7.5, we get that $|g_t^{\mathbf{r},\sigma}| \leq \kappa$. This finishes the proof. □

Claim 7.27. *Let $(\mathbf{r},\sigma)$ be an aligned pair. Then, $g_N^{\mathbf{r},\sigma}$ is just dependent on the covariance matrix C (and the pair $(\mathbf{r},\sigma)$).*

Proof. Since $(\mathbf{r},\sigma)$ is an aligned pair, it means that the set $\{1,\ldots,t\}$ can be partitioned into two sets given by L and R and a bijection $\pi : L \to R$ such that $g_t^{\mathbf{r},\sigma} = \prod_{a\in L}\langle h_a, h_{\pi(a)}\rangle$. Note that the mapping π is just a function of $(\mathbf{r},\sigma)$ and once π is fixed, $g_t^{\mathbf{r},\sigma}$ is just a function of the matrix C. □

Proof of Lemma 7.18. Let $B = \{(\mathbf{r},\sigma) : \mathbf{r} \in \mathcal{S}_z \text{ and } (\mathbf{r},\sigma) \text{ is aligned}\}$. Applying (7.9), we get that

$$\left|\mathbf{E}[\prod_{i=1}^{t} p_i] - \frac{1}{|\Lambda_{\mathbf{r}}|}\sum_{(\mathbf{r},\sigma)\in B} \gamma_{\mathbf{r}} I_{m_{t-1}^{\mathbf{r}}}(g_t^{\mathbf{r},\sigma})\right| \leq \sum_{\mathbf{r}\notin\mathcal{S}_z} \|\gamma_{\mathbf{r}} \cdot g_t^{\mathbf{r}}\| + \frac{1}{|\Lambda_{\mathbf{r}}|}\sum_{(\mathbf{r},\sigma)\notin B:\mathbf{r}\in\mathcal{S}_z} \|\gamma_{\mathbf{r}} \cdot g_t^{\mathbf{r},\sigma}\|.$$

Note that $(1/|\Lambda_{\mathbf{r}}|)\sum_{(\mathbf{r},\sigma)\in B} \gamma_{\mathbf{r}} g_t^{\mathbf{r},\sigma}$ is dependent just on C and $\mathbf{r}$. Further, applying Claim 7.26, Proposition 7.23 and Claim 7.22, we see that the right hand side can be bound by $2^{t(q_1+\cdots+q_t+1)} \cdot \kappa$. This finishes the proof. □

7.6 Proof of Theorem 5.2

Let us assume that $f = \mathsf{PTF}(p_1,\ldots,p_k)$ where f is the PTF appearing in the statement of Theorem 5.2 After applying Lemma 7.8, we can pretend that the PTF $f = \mathsf{PTF}(p_1,\ldots,p_k)$ in Theorem 5.2 is multilinear. After scaling the polynomials $p_1,\ldots,p_k$, we can assume that $\mathsf{Var}(p_s) = 1$ for all $1 \leq s \leq k$. With this, we note that there are constants $c_{s,q}$ for $1 \leq s \leq k$ and $1 \leq q \leq d$ such that

$$p_s = p_{s,0} + \sum_{q=1}^{d} c_{s,q} p_{s,q},$$

where $p_{s,q} \in \mathcal{W}^q$, $\mathsf{Var}(p_{s,q}) = 1$ for all $s \in [k]$, $1 \leq q \leq d$ and $\sum_{q=1}^{d} c_{s,q}^2 = 1$.

Let $\beta : [0,\infty) \to [0,1)$ be a quickly decreasing function, with properties that will be specified later. Apply Theorem 7.16 (with $\tau = (\varepsilon/(kd))^{3d}$), and let $\tilde{p}_{s,q}$ (for $0 \leq s \leq k$ and $1 \leq q \leq d$) be the resulting polynomials. Defining $\tilde{p}_s = \tilde{p}_{s,0} + \sum_{q=1}^{d} c_{s,q}\tilde{p}_{s,q}$, Theorem 7.16 implies that $\mathbf{E}[p_s - \tilde{p}_s] = 0$ and

$$\mathsf{Var}(p_s - \tilde{p}_s) = \sum_{q=1}^{d} c_{s,q}^2 \mathsf{Var}(p_{s,q} - \tilde{p}_{s,q}) \leq \tau.$$

Define the PTF $\tilde{f} = \mathsf{PTF}(\tilde{p}_1,\ldots,\tilde{p}_k)$. By Theorem 7.7 and a union bound,

$$\begin{aligned}
\|\mathbf{E}_{x\sim\gamma_n}[\tilde{f}(x)] - \mathbf{E}_{x\sim\gamma_\ell}[f(x)]\| &\leq \varepsilon, \\
\left|\mathbf{E}_{x\sim\gamma_n}[\langle\tilde{f}, P_t\tilde{f}\rangle] - \mathbf{E}_{x\sim\gamma_\ell}[\langle f, P_t f\rangle]\right| &\leq \varepsilon, \\
\Pr_{x\sim\gamma_\ell}[x \in \mathsf{Collision}(\tilde{f})] &\leq \Pr_{x\sim\gamma_n}[x \in \mathsf{Collision}(f)] + \varepsilon.
\end{aligned}$$

From now on, we will forget about f and work with $\tilde{f}$. Recalling (from Theorem 7.16) that

$$\tilde{p}_{s,q} = \mathrm{Out}(\tilde{p}_{s,q})\left(\mathrm{In}(\tilde{p}_{s,q})_1(x), \ldots, \mathrm{In}(\tilde{p}_{s,q})_{\mathrm{num}(s,q)}(x)\right),$$

our next task is to "swap out" the inner polynomials for a new collection. Let $\mathcal{P}$ denote the collection of polynomials $\{\mathrm{In}(\tilde{p}_{s,q})_\ell : 1 \le s \le k, 1 \le q \le d, 1 \le \ell \le \mathrm{num}(s,q)\}$, and we will introduce another family of polynomials $\mathcal{R} = \{\mathrm{In}(r_{s,q})_\ell : 1 \le s \le k, 1 \le q \le d, 1 \le \ell \le \mathrm{num}(s,q)\}$ that satisfies certain conditions:

Condition 7.28.

- Each polynomial $\mathrm{In}(r_{s,q})_\ell$ has arity $n_0 = O_{s,k,\beta,\varepsilon}(1)$.
- For every s, q, and ℓ, $\mathrm{In}(r_{s,q})_\ell \in \mathcal{W}^j$ if and only if $\mathrm{In}(p_{s,q})_\ell \in \mathcal{W}^j$.
- $\mathcal{P}$ and $\mathcal{R}$ have the same covariances: for every $1 \le s_1, s_2 \le k$, $1 \le q_1, q_2 \le d$, and $1 \le \ell_i \le \mathrm{num}(s_i, q_i)$ (for $i \in \{1,2\}$),

$$\mathbf{E}[\mathrm{In}(\tilde{p}_{s_1,q_1})_{\ell_1}(x) \cdot \mathrm{In}(\tilde{p}_{s_2,q_2})_{\ell_2}(x)] = \mathbf{E}[\mathrm{In}(r_{s_1,q_1})_{\ell_1}(x) \cdot \mathrm{In}(r_{s_2,q_2})_{\ell_2}(x)].$$

- Every polynomial in the family $\mathcal{R}$ is $\beta(\mathrm{Num} + \mathrm{Coeff})$-eigenregular.

It turns out to be possible to satisfy these conditions:

Lemma 7.29. *There exists a family $\mathcal{R}$ of polynomials on $n_0 = \mathsf{poly}(d, \mathrm{Num}, \beta(\mathrm{Num} + \mathrm{Coeff}))$ variables that meets the requirements of Condition 7.28.*

Proof. For $1 \le i \le d$, let $\mathcal{P}_i = \{\mathrm{In}(p_{s,q})_\ell\}_{\ell=1,\ldots,\mathrm{num}(s,q)} \cap \mathcal{W}^i$. Let m_i denote the size of $\mathcal{P}_i$. We will construct $\mathcal{R}$ by constructing the corresponding set $\mathcal{R}_i$ for each $1 \le i \le d$. Note that individually constructing $\mathcal{R}_i$ for each $1 \le i \le d$ suffices for our construction, because if $r_1 \in \mathcal{W}^i$, $r_2 \in \mathcal{W}^j$ and $i \ne j$ then $\mathbf{E}[r_1(x) \cdot r_2(x)] = 0$.

For convenience of notation, let us enumerate the elements of $\mathcal{P}_i$ as $\{p_1, \ldots, p_{m_i}\}$, and we will construct $\mathcal{R}_i = \{r_1, \ldots, r_{m_i}\}$. Define $h_j \in \mathcal{H}^{\odot i}$ by $p_j = I_i(h_j)$, and note that $h_1, \ldots, h_{m_i}$ span a space of dimension at most m_i. Setting $\mathcal{H}_1 = \mathbb{R}^{m_i}$, Proposition 7.2 implies that $\dim(\mathcal{H}_1^{\odot i}) \ge m_i$, and so there exist $v_1, \ldots, v_{m_i} \in \mathcal{H}_1^{\odot i}$ satisfying $\langle v_j, v_{j'} \rangle = \langle h_j, h_{j'} \rangle$ for all j, j'.

Let $\delta = \beta(\mathsf{Num} + \mathsf{Coeff})$, $\kappa = \lceil 1/\delta^2 \rceil$, and $n_0 = m_i \kappa$; note that n_0 is indeed polynomial in d, Num, and $\beta(\mathsf{Num} + \mathsf{Coeff})$. For $1 \le j \le m_i$, define $r_j : \mathbb{R}^{n_0} \to \mathbb{R}$ as follows: divide $x \in \mathbb{R}^{n_0}$ into κ blocks of size m_i each (call the blocks $X_1, \ldots, X_\kappa$) and define

$$r_j = \frac{1}{\sqrt{\kappa}} \cdot \left(q_j(X_1) + \cdots + q_j(X_\kappa)\right).$$

Then

$$\mathbf{E}[r_{j_1} \cdot r_{j_2}] = \sum_{j=1}^{\kappa} \frac{1}{\kappa} \cdot \mathbf{E}[q_{j_1}(X_j) \cdot q_{j_2}(X_j)] = \sum_{j=1}^{\kappa} \frac{1}{\kappa} \cdot \langle v_{j_1}, v_{j_2} \rangle = \sum_{j=1}^{\kappa} \frac{1}{\kappa} \cdot \langle h_{j_1}, h_{j_2} \rangle = \mathbf{E}[p_{j_1} \cdot p_{j_2}].$$

The first equality relies on the observation that $q_{j_1}, q_{j_2} \in \mathcal{W}^i$ for $i > 1$ and thus $\mathbf{E}[q_{j_1}(X_j) \cdot q_{j_2}(X_\ell)] = \mathbf{E}[q_{j_1}(X_j)]\mathbf{E}[q_{j_2}(X_\ell)] = 0$ if $j \ne \ell$; the second and fourth equality use Proposition 7.3. Finally, Fact 7.13 implies that the polynomials $\{r_j\}_{1 \le j \le m_i}$ are $\delta = \beta(\mathrm{Num} + \mathrm{Coeff})$-eigenregular. $\square$

Having constructed the family $\mathcal{R}$, for $1 \le s \le k$ and $1 \le q \le d$, we define

$$r_{s,q} = \mathrm{Out}(\tilde{p}_{s,q})\left(\mathrm{In}(r_{s,q})_1(x), \ldots, \mathrm{In}(r_{s,q})_{\mathrm{num}(s,q)}(x)\right).$$

That is, compared to $\tilde{p}_{s,q}$, the inner polynomials have changed but the outer polynomial is the same. For $1 \le s \le k$, we define

$$r_s = \tilde{p}_{s,0} + \sum_{q=1}^{d} c_{s,q} r_{s,q},$$

Finally, we define $f_{\mathsf{junta}} = \mathsf{PTF}(r_1, \ldots, r_k)$. Note that $f_{\mathsf{junta}} : \mathbb{R}^{n_0} \to \{0,1\}$.

We will now show that this construction has all the properties stated in Theorem 5.2. First of all, note that as long as the function $\beta(\cdot)$ is computable, n_0 is a computable function of d and ε. To examine the noise stability of f_{junta}, we introduce a new family of polynomials which will essentially be "noise-attenuated" versions of the existing ones. These polynomials will be over either $2n$ or $2n_0$ variables. For $1 \le s \le k$, $1 \le q \le d$ and $1 \le \ell \le \mathrm{num}(s,q)$, we define

$$\mathrm{In}(u_{s,q})_\ell(x,y) = \mathrm{In}(\tilde{p}_{s,q})_\ell(\mathrm{e}^{-t}x + \sqrt{1-\mathrm{e}^{-2t}}y),$$
$$\mathrm{In}(v_{s,q})_\ell(x,y) = \mathrm{In}(r_{s,q})_\ell(\mathrm{e}^{-t}x + \sqrt{1-\mathrm{e}^{-2t}}y),$$

and

$$u_{s,q} = \mathrm{Out}(\tilde{p}_{s,q})\left(\mathrm{In}(u_{s,q})_1(x,y), \ldots, \mathrm{In}(u_{s,q})_{\mathrm{num}(s,q)}(x,y)\right),$$
$$v_{s,q} = \mathrm{Out}(\tilde{p}_{s,q})\left(\mathrm{In}(v_{s,q})_1(x,y), \ldots, \mathrm{In}(v_{s,q})_{\mathrm{num}(s,q)}(x,y)\right),$$

and finally

$$u_s = u_{s,0} + \sum_{q=1}^{d} c_{s,q} u_{s,k},$$

$$v_s = v_{s,0} + \sum_{q=1}^{d} c_{s,q} v_{s,k},$$

where $u_{s,0} = v_{s,0} = p_{s,0} = r_{s,0}$. To help keep the notation straight, note that every polynomial involving "u" is on $2n$ variables, while every polynomial involving "v" is on $2n_0$ variables. The PTFs corresponding to u and v will be denoted $f_u = \mathsf{PTF}(u_1, \ldots, u_k)$ and $f_v = \mathsf{PTF}(v_1, \ldots, v_k)$.

The purpose of introducing these new polynomials is to have the following expressions for noise stability:

$$\mathop{\mathbf{E}}_{x\sim\gamma_n}[\langle \tilde{f}(x), P_t \tilde{f}(x)\rangle] = \mathop{\mathbf{E}}_{x\sim\gamma_n, y\sim\gamma_n}[\langle \tilde{f}(x), f_u(x,y)\rangle],$$

$$\mathop{\mathbf{E}}_{x\sim\gamma_n}[\langle f_{\mathsf{junta}}(x), P_t f_{\mathsf{junta}}(x)\rangle] = \mathop{\mathbf{E}}_{x\sim\gamma_n, y\sim\gamma_n}[\langle f_{\mathsf{junta}}(x), f_v(x,y)\rangle].$$

The next claim establishes relations between the pairwise correlations of polynomials in the families $\{\mathrm{In}(u_{s,q})_\ell\}$ and $\{\mathrm{In}(v_{s,q})_\ell\}$.

Claim 7.30. *For all* $1 \le s_1, s_2 \le k$, $1 \le q_1, q_2 \le d$ *and* $1 \le \ell_i \le \mathsf{num}(s_i, q_i)$ *(for* $i \in \{1,2\}$*), we have*

$$\mathbf{E}[\mathrm{In}(u_{s_1,q_1})_{\ell_1}(x,y) \cdot \mathrm{In}(u_{s_2,q_2})_{\ell_2}(x,y)] = \mathbf{E}[\mathrm{In}(v_{s_1,q_1})_{\ell_1}(x,y) \cdot \mathrm{In}(v_{s_2,q_2})_{\ell_2}(x,y)]$$

and

$$\mathbf{E}[\mathrm{In}(p_{s_1,q_1})_{\ell_1}(x) \cdot \mathrm{In}(u_{s_2,q_2})_{\ell_2}(x,y)] = \mathbf{E}[\mathrm{In}(r_{s_1,q_1})_{\ell_1}(x) \cdot \mathrm{In}(v_{s_2,q_2})_{\ell_2}(x,y)].$$

Proof. Note that $\mathrm{In}(u_{s_1,q_1})_{\ell_1}$ and $\mathrm{In}(u_{s_2,q_2})_{\ell_2}$ are obtained by applying a unitary transformation (on the space of variables) to the polynomials $\mathrm{In}(p_{s_1,q_1})_{\ell_1}$ and $\mathrm{In}(p_{s_2,q_2})_{\ell_2}$. Since the standard Gaussian measure is rotationally invariant,

$$\mathbf{E}[\mathrm{In}(u_{s_1,q_1})_{\ell_1}(x,y) \cdot \mathrm{In}(u_{s_2,q_2})_{\ell_2}(x,y)] = \mathbf{E}[\mathrm{In}(\tilde{p}_{s_1,q_1})_{\ell_1}(x) \cdot \mathrm{In}(\tilde{p}_{s_2,q_2})_{\ell_2}(x)].$$

Likewise,

$$\mathbf{E}[\mathrm{In}(v_{s_1,q_1})_{\ell_1}(x) \cdot \mathrm{In}(v_{s_2,q_2})_{\ell_2}(x)] = \mathbf{E}[\mathrm{In}(r_{s_1,q_1})_{\ell_1}(x,y) \cdot \mathrm{In}(r_{s_2,q_2})_{\ell_2}(x,y)].$$

Since the right hand sides above are equal (by the construction of the family $\{\mathrm{In}(r_{s,q})_\ell\}$), the first claim follows.

To prove the second claim, note that that there exist $j_1, j_2 \in \mathbb{N}$, such that

$$\mathrm{In}(p_{s_1,q_1})_{\ell_1}(x),\ \mathrm{In}(r_{s_1,q_1})_{\ell_1}(x) \in \mathcal{W}^{j_1} \text{ and } \mathrm{In}(u_{s_2,q_2})_{\ell_2}(x,y),\ \mathrm{In}(v_{s_2,q_2})_{\ell_2}(x,y) \in \mathcal{W}^{j_2}.$$

If $j_1 \neq j_2$, then the second claim holds trivially because both sides are zero; we will assume, therefore, that $j_1 = j_2 = j$. Next, we observe that

$$\mathbf{E}[\mathrm{In}(p_{s_1,q_1})_{\ell_1}(x) \cdot \mathrm{In}(u_{s_2,q_2})_{\ell_2}(x,y)] = \langle \mathrm{In}(p_{s_1,q_1})_{\ell_1}(x),\ P_t\ \mathrm{In}(p_{s_2,q_2})_{\ell_2}(x)\rangle,$$

and recall that $P_t w = \mathrm{e}^{-jt} w$ for all $w \in \mathcal{W}^j$. Hence,

$$\begin{aligned}
\mathbf{E}[\mathrm{In}(p_{s_1,q_1})_{\ell_1}(x) \cdot \mathrm{In}(u_{s_2,q_2})_{\ell_2}(x,y)] &= \mathrm{e}^{-jt}\mathbf{E}[\mathrm{In}(p_{s_1,q_1})_{\ell_1}(x) \cdot \mathrm{In}(p_{s_2,q_2})_{\ell_2}(x)],\\
\mathbf{E}[\mathrm{In}(r_{s_1,q_1})_{\ell_1}(x) \cdot \mathrm{In}(v_{s_2,q_2})_{\ell_2}(x,y)] &= \mathrm{e}^{-jt}\mathbf{E}[\mathrm{In}(r_{s_1,q_1})_{\ell_1}(x) \cdot \mathrm{In}(r_{s_2,q_2})_{\ell_2}(x)],
\end{aligned}$$

and as before, the right hand sides are equal. $\square$

We will now list the conditions that we will require on the function $\beta(\cdot)$. These conditions look somewhat intimidating, but it is easy to see that they will be satisfied as long as β decreases sufficiently quickly.

Condition 7.31. The function $\beta : [1,\infty) \to [0,1]$ should satisfy the following conditions:

1. For $\xi = \varepsilon/(40k^2)$ and for $B^{(1)}, \delta^{(1)}, c_{(1)}$ defined as

$$B^{(1)} = \Omega\left(\ln \frac{\mathsf{Num} \cdot d}{\xi}\right)^{d/2};\ \delta^{(1)} = \left(\frac{\xi}{d \cdot B^{(1)} \cdot \sqrt{\mathsf{Num}} \cdot \mathsf{Coeff}^{1/d}}\right)^{d};\ c_{(1)} = \frac{\mathsf{Num}}{\delta^{(1)} \cdot \sqrt{\xi}},$$

   we require

$$\beta(\mathsf{Coeff} + \mathsf{Num}) \cdot \mathsf{Coeff}^2 \cdot 2^{\mathsf{Num}\cdot(\mathsf{Num}\cdot 2d+1)} \cdot 2^{O(d^3)} \le \frac{1}{4} \text{ and}$$
$$2^{O(d\log d)} \cdot \mathsf{Num}^2 \cdot 4c_{(1)}^2 \cdot \beta(\mathsf{Coeff} + \mathsf{Num}) \le \xi.$$

2. Define $\xi = \varepsilon/(40k^2)$, $L = 4 \cdot 9^{d+1} \cdot (d+1)^2 \cdot \log^{d/2}(k \cdot d/\varepsilon)$ and $\vartheta = (L/2) \cdot L^{-2^{2d}}$. Further, define $B^{(2)}$, $\delta^{(2)}$ and $c_{(2)}$ as

$$B^{(2)} = \Omega\left(\ln \frac{2 \cdot \mathsf{Num} \cdot 2d}{\xi}\right)^d; \; \delta^{(2)} = \left(\frac{\xi \cdot \vartheta^{1/2d}}{2d \cdot B^{(2)} \cdot \sqrt{2 \cdot \mathsf{Num}} \cdot \mathsf{Coeff}^{2/d}}\right)^{2d}; \; c_{(2)} = \frac{2 \cdot \mathsf{Num}}{\delta^{(2)} \cdot \sqrt{\xi}}.$$

We require

$$\beta(\mathsf{Coeff}^2 + 2 \cdot \mathsf{Num}) \cdot \mathsf{Coeff}^4 \cdot 2^{2\mathsf{Num} \cdot (\mathsf{Num} \cdot 4d+1)} \cdot 2^{O(d^3)} \leq \frac{\vartheta}{4} \text{ and}$$
$$2^{O(d \log d)} \cdot \mathsf{Num}^2 \cdot 4c_{(2)}^2 \cdot \beta(\mathsf{Coeff}^2 + 2 \cdot \mathsf{Num}) \leq \xi.$$

Applying the above conditions on β will allow us to conclude that our pairs of polynomials have approximately the same joint distribution of signs. For the rest of this section, set $\xi = \varepsilon/(40k^2)$.

Corollary 7.32. *If β satisfies Condition 7.31 then for all $1 \leq s, s' \leq k$,*

(a) $\left|\Pr_{x \sim \gamma_n}[\tilde{p}_s \geq 0] - \Pr_{x \sim \gamma_{n_0}}[r_s \geq 0]\right| \leq \xi$,

(b) $\left|\Pr_{x,y \sim \gamma_n}[u_s \geq 0] - \Pr_{x,y \sim \gamma_{n_0}}[v_s \geq 0]\right| \leq \xi$,

(c) $\left|\Pr_{x \sim \gamma_n}[\tilde{p}_s \cdot \tilde{p}_{s'} \geq 0] - \Pr_{x \sim \gamma_{n_0}}[r_s \cdot r_{s'} \geq 0]\right| \leq \xi$,

(d) $\left|\Pr_{x,y \sim \gamma_n}[u_s \cdot u_{s'} \geq 0] - \Pr_{x,y \sim \gamma_{n_0}}[v_s \cdot v_{s'} \geq 0]\right| \leq \xi$,

(e) $\left|\Pr_{x,y \sim \gamma_n}[\tilde{p}_s \cdot u_{s'} \geq 0] - \Pr_{x,y \sim \gamma_{n_0}}[r_s \cdot v_{s'} \geq 0]\right| \leq \xi$.

Proof. We first note that it suffices to prove (a), (c) and (e). This is because the polynomials $\{u_s\}$ (resp. $\{v_s\}$) are obtained by a unitary transformation of variables on the polynomials $\{\tilde{p}_s\}$ (resp. $\{r_s\}$). For $1 \leq s \leq k$, note that there is a $\Psi_s : \mathbb{R}^{\mathsf{Num}} \to \mathbb{R}$ such that

$$\tilde{p}_s = \Psi_s(\{\mathrm{In}(\tilde{p}_{s,q})_1(x), \ldots, \mathrm{In}(\tilde{p}_{s,q})_{\mathrm{num}(s,q)}(x)\}_{1 \leq q \leq d}),$$

$$r_s = \Psi_s(\{\mathrm{In}(r_{s,q})_1(x), \ldots, \mathrm{In}(r_{s,q})_{\mathrm{num}(s,q)}(x)\}_{1 \leq q \leq d}).$$

Proof of Item (a): Recall that $\mathsf{Var}(\tilde{p}_s) = 1$. Define $B^{(1)}, \delta^{(1)}, c_{(1)}$ as in Condition 7.31. We now apply Lemma 7.19 with $\Psi = \Psi_s$, $\eta = 1$, $\{A_i\} = \mathcal{P}$, and $\{B_i\} = \mathcal{R}$. Condition 7.31 implies that the B polynomials (i. e., $\mathrm{In}(r_{s,q})_\ell$) satisfy the eigenregularity assumption in Lemma 7.19. The conclusion of Lemma 7.19 proves (a).

Proof of Item (c): For $1 \leq s, s' \leq k$, note that there is a function $\Psi_{s,s'} : \mathbb{R}^{2 \cdot \mathsf{Num}} \to \mathbb{R}$ such that

$$\tilde{p}_s \cdot \tilde{p}_{s'} = \Psi_{s,s'}(\{\mathrm{In}(\tilde{p}_{s,q})_\ell(x)\}_{q,\ell}, \{\mathrm{In}(\tilde{p}_{s',q})(x)\}_{q,\ell}),$$
$$r_s \cdot r_{s'} = \Psi_{s,s'}(\{\mathrm{In}(r_{s,q})_\ell(x)\}_{q,\ell}, \{\mathrm{In}(r_{s',q})(x)\}_{q,\ell}).$$

Note that $\Psi_{s,s'}$ is a degree-$2d$ polynomial with $\mathsf{Coeff}_{\Psi_{s,s'}} \le \mathsf{Coeff}^2$ and $\mathsf{Num}_{\Psi_{s,s'}} \le 2\mathsf{Num}$. In order to apply Lemma 7.19, we will need a lower bound on the variance of $\tilde{p} \cdot \tilde{p}_s$:

$$\mathsf{Var}(\tilde{p}_s \cdot \tilde{p}_{s'}) \ge \frac{L}{2} \cdot L^{-2^{2d}},$$

where $L = 4 \cdot 9^{d+1} \cdot (d+1)^2 \cdot \log^{d/2}(k \cdot d/\varepsilon)$. (This lower bound will be proven as Lemma 7.33.) Note that the right hand side above is exactly what we called ϑ in Condition 7.31. Now we will apply Lemma 7.19 with $\eta = \vartheta$: Claim 7.30 implies that the second moments match, while the second part of Condition 7.31 ensures that the polynomials $\mathrm{In}(r_{t,q})_\ell$ are sufficiently eigenregular. The conclusion of Lemma 7.19 proves item (c).

Proof of Item (e): The proof of item (e) is the same as the proof of item (c) with $u_{s'}$ taking over the role of $\tilde{p}_{s'}$. We omit the details. □

Lemma 7.33. *Let $p, q : \mathbb{R}^n \to \mathbb{R}$ be degree-d polynomials such that $\mathsf{Var}(p) = \mathsf{Var}(q) = 1$. Let $T = \max\{\mathbf{E}[p^2], \mathbf{E}[q^2]\}$ and $L = 4 \cdot T \cdot 9^{d+1} \cdot (d+1)^2$. Then*

$$\mathsf{Var}(p \cdot q) \ge \frac{L}{2} \cdot (L)^{2^{-2d}}.$$

We will begin with an easier lower bound, in terms of the "highest-level weights:"

Lemma 7.34. *If $p = \sum_{i=0}^{d_1} I_i(f_i)$ and $q = \sum_{i=0}^{d_2} I_i(g_i)$ for $f_i \in \mathcal{H}^{\odot i}$ and $g_i \in \mathcal{H}^{\odot i}$ then*

$$\mathsf{Var}(p \cdot q) \ge \|f_{d_1}\|_F^2 \cdot \|g_{d_2}\|_F^2.$$

Proof.

$$\begin{aligned} p \cdot q &= \sum_{r_1=0}^{d_1} \sum_{r_2=0}^{d_2} I_{r_1}(f_{r_1}) \cdot I_{r_2}(g_{r_2}) \\ &= \sum_{r_1=0}^{d_1} \sum_{r_2=0}^{d_2} \sum_{r=0}^{r_1 \wedge r_2} r! \cdot \binom{r_1}{r} \cdot \binom{r_2}{r} \cdot \frac{\sqrt{(r_1+r_2-2r)!}}{\sqrt{r_1!}\sqrt{r_2!}} \cdot I_{r_1+r_2-2r}(f_{r_1} \widetilde{\otimes}_r g_{r_2}). \end{aligned}$$

The second equality uses Itô's multiplication formula (Proposition 7.4). The term with $r = 0$, $r_1 = d_1$ and $r_2 = d_2$ is orthogonal to all the other terms because it is the only one belonging to the $(r_1 + r_2)$th Wiener chaos; hence,

$$\mathsf{Var}(p \cdot q) \ge \mathsf{Var}\left(\frac{\sqrt{(d_1+d_2)!}}{\sqrt{d_1!}\sqrt{d_2!}} \cdot I_{d_1+d_2}(f_{d_1} \widetilde{\otimes}_0 g_{d_2}) \right) = \frac{(d_1+d_2)!}{d_1! \cdot d_2!} \|f_{d_1} \widetilde{\otimes}_r g_{d_2}\|_F^2.$$

Finally, Neuberger [27] showed that

$$\|f_{d_1} \widetilde{\otimes}_r g_{d_2}\|_F^2 \ge \frac{1}{\binom{d_1+d_2}{d_1}} \|f_{d_1}\|_F^2 \|g_{d_2}\|_F^2,$$

which completes the proof. □

Proof of Lemma 7.33. Let us express p and q in terms of iterated Itô integrals.

$$p = \sum_{r=0}^{d} I_r(f_r) \text{ and } q = \sum_{r=0}^{d} I_r(g_r).$$

Let $\alpha = \mathbf{E}[p^2]$ and $\beta = \mathbf{E}[q^2]$. Let $\Gamma : \mathbb{N} \to [0,1]$ be a function (which we shall fix later). Consider the following iterative process:

- Start with $i = d$ and $j = d$. If $\|f_i\|_F \cdot \|g_j\|_F \geq \Gamma(i+j)$, then stop the process.
- If $\|f_i\|_F \leq \sqrt{\Gamma(i+j)}$, then $i \leftarrow i-1$, else $j \leftarrow j-1$.
- If any of i or j reaches zero, terminate the process.

Recall that $T = \max\{\alpha, \beta\}$ and $L = 4 \cdot T \cdot 9^{d+1} \cdot (d+1)^2$. We let $\Gamma(x) = L \cdot L^{-2^x}$. First, we claim that for $\Gamma(\cdot)$ as chosen here, the iterative process above terminates with (i,j) such that neither of them is zero. Towards a contradiction, assume that the above process terminates with $i = 0$ and $j > 0$. Further, for $d \leq \ell \leq 1$, when the value of i drops from ℓ to $\ell - 1$, let the value of $i + j$ be κ_ℓ. Note that $\kappa_1 < \cdots < \kappa_d$ and $\kappa_1 \geq 2$. Thus, if the process terminates with $i = 0$, then

$$\mathsf{Var}(p) = \sum_{\ell=1}^{d} \|f_\ell\|_F^2 \leq \sum_{\ell=1}^{d} \Gamma(\kappa_\ell) \leq 2 \cdot \Gamma(\kappa_1) < 1.$$

This results in a contradiction which means that both i and j must be non-zero at the end of the process. Here the penultimate inequality uses that for our choice of $\Gamma(\cdot)$, $\Gamma(x+1) \leq \Gamma(x)/2$ and the last inequality uses that $\Gamma(\kappa_1) < 1/2$.

Let us assume that (i_0, j_0) is the pair returned by the above iterative process and define $\tilde{p} = \sum_{r=0}^{i_0} I_r(f_r)$ and $\tilde{q} = \sum_{r=0}^{j_0} I_r(g_r)$. Applying Lemma 7.34, we have $\mathsf{Var}(\tilde{p} \cdot \tilde{q}) \geq \Gamma^2(i_0 + j_0)$. Observe that $p \cdot q = \tilde{p} \cdot \tilde{q} + (p - \tilde{p}) \cdot q + \tilde{p} \cdot (q - \tilde{q})$. Thus, it follows that $\mathsf{Var}(p \cdot q - \tilde{p} \cdot \tilde{q}) = \mathsf{Var}((p - \tilde{p}) \cdot q + \tilde{p} \cdot (q - \tilde{q}))$. Applying Jensen's inequality, we get

$$\sqrt{\mathsf{Var}(p \cdot q)} \geq \sqrt{\mathsf{Var}(\tilde{p} \cdot \tilde{q})} - \sqrt{\mathsf{Var}(\tilde{p} \cdot (q - \tilde{q}))} - \sqrt{\mathsf{Var}(q \cdot (p - \tilde{p}))}.$$

Hence, it suffices to bound $\mathsf{Var}(\tilde{p} \cdot (q - \tilde{q}))$ and $\mathsf{Var}(q \cdot (p - \tilde{p}))$. To do this, notice that by definition of the process,

$$\mathsf{Var}(q - \tilde{q}),\ \mathsf{Var}(p - \tilde{p}) < \sum_{\ell = i_0 + j_0 + 1}^{d} \Gamma(\ell) < 2 \cdot \Gamma(i_0 + j_0 + 1).$$

Now, Cauchy-Schwarz and the Gaussian hypercontractive inequality imply that $\mathbf{E}[(r \cdot s)^2] \leq 9^d \mathbf{E}[r^2] \mathbf{E}[s^2]$ for any degree-d polynomials r and s. Hence,

$$\mathsf{Var}(\tilde{p} \cdot (q - \tilde{q})),\ \mathsf{Var}(q \cdot (p - \tilde{p})) \leq 2 \cdot \Gamma(i_0 + j_0 + 1) \cdot d \cdot (d+1) \cdot 8^d \cdot T.$$

Thus,

$$\sqrt{\mathsf{Var}(p \cdot q)} \geq \Gamma(i_0 + j_0) - 2\sqrt{2} \cdot (d+1) \cdot (2\sqrt{2})^d \cdot \sqrt{T} \cdot \sqrt{\Gamma(i_0 + j_0 + 1)}.$$

However, note that for any $x \geq 0$, $\sqrt{\Gamma(x+1)} = \Gamma(x)/\sqrt{L}$. Plugging the value of L, we obtain that

$$\sqrt{\mathsf{Var}(p \cdot q)} \geq \Gamma(i_0 + j_0)/2 \geq \frac{L}{2} \cdot L^{2^{-2d}}. \qquad \square$$

Before continuing with the proof, let us summarize what we have shown so far:

Theorem 7.35. *Fix $\varepsilon > 0$ and $t > 0$, and let $p_1, \dots, p_k : \mathbb{R}^n \to \mathbb{R}$ be degree-d polynomials satisfying* $\mathrm{Var}(p_s) = 1$ *and* $|\mathbf{E}[p_s]| \le \log^{d/2}(k \cdot d/\varepsilon)$ *for all* $1 \le s \le k$. *There exists a computable* $n_0 = n_0(k, d, t, \varepsilon)$ *and polynomials* $r_1, \dots, r_k : \mathbb{R}^{n_0} \to \mathbb{R}$ *such that for every* $1 \le s, s' \le k$:

(a) $\left|\Pr_{x\sim\gamma_n}[p_s \ge 0] - \Pr_{x\sim\gamma_n}[r_s \ge 0]\right| \le \varepsilon$,

(b) $\left|\Pr_{x,y\sim\gamma_n}[u_s \ge 0] - \Pr_{x,y\sim\gamma_{n_0}}[v_s \ge 0]\right| \le \varepsilon$,

(c) $\left|\Pr_{x\sim\gamma_n}[p_s \cdot p_{s'} \ge 0] - \Pr_{x\sim\gamma_{n_0}}[r_s \cdot r_{s'} \ge 0]\right| \le \varepsilon$,

(d) $\left|\Pr_{x,y\sim\gamma_n}[u_s \cdot u_{s'} \ge 0] - \Pr_{x,y\sim\gamma_{n_0}}[v_s \cdot v_{s'} \ge 0]\right| \le \varepsilon$,

(e) $\left|\Pr_{x,y\sim\gamma_n}[p_s \cdot u_{s'} \ge 0] - \Pr_{x,y\sim\gamma_{n_0}}[r_s \cdot v_{s'} \ge 0]\right| \le \varepsilon$,

where $u_s(x,y) = p_s(\mathrm{e}^{-t}x + \sqrt{1-\mathrm{e}^{-2t}}y)$ *and* $v_s(x,y) = r_s(\mathrm{e}^{-t}x + \sqrt{1-\mathrm{e}^{-2t}}y)$.

Indeed, Theorem 7.35 follows directly from Corollary 7.32 and the fact that $\mathsf{sign}(p_s)$ and $\mathsf{sign}(\tilde{p}_s)$ disagree with probability at most $O(\varepsilon)$.

Returning to the notation of Corollary 7.32, our next goal is to prove that $\Pr_{x\sim\gamma_{n_0}}[x \in \mathsf{Collision}(f_{\mathsf{junta}})]$ is small.

Claim 7.36. $\Pr_{x\sim\gamma_{n_0}}[x \in \mathsf{Collision}(f_{\mathsf{junta}})] \le (k^2+1) \cdot \Pr[x \in \mathsf{Collision}(f)] + k \cdot \xi + \dfrac{3k^2\xi}{2}$.

Proof. For a multivariate PTF $f = \mathsf{PTF}(p_1, \dots, p_k)$, define $\mathsf{Unique}(f,i) = \overline{\mathsf{Collision}(f)} \cap \{x : p_i(x) \ge 0\}$. In other words, it is the set of all points such that i is the unique index such that $p_i(x) \ge 0$. Since we haven't mentioned f and f_{junta} for a while, recall that $f = \mathsf{PTF}(p_1, \dots, p_k)$ and $f_{\mathsf{junta}} = \mathsf{PTF}(r_1, \dots, r_k)$.

Next, observe that for real numbers $a, b \ne 0$,

$$\mathbf{I}\,[a \ge 0 \wedge b \ge 0] = \frac{1}{2}\big(\mathbf{I}\,[a \cdot b \ge 0] + \mathbf{I}\,[a \ge 0] + \mathbf{I}\,[b \ge 0] - 1\big) \tag{7.10}$$

(where $\mathbf{I}(P)$ is 1 if P is true and 0 otherwise). Since $p_s(x) \ne 0$ almost surely for every s, the identity above implies that

$$\begin{aligned}
&|\Pr[p_s(x) \ge 0 \wedge p_{s'}(x) \ge 0] - \Pr[r_s(x) \ge 0 \wedge r_{s'}(x) \ge 0]| \\
&\le \frac{1}{2}\Big(\big|\Pr[p_s(x) \ge 0] - \Pr[r_s(x) \ge 0]\big| + \big|\Pr_{x\sim\gamma_n}[p_{s'}(x) \ge 0] - \Pr_{x\sim\gamma_n}[r_{s'}(x) \ge 0]\big|\Big) \\
&\quad + \frac{1}{2}\Big(\big|\Pr_{x\sim\gamma_n}[p_{s'}(x) \cdot p_s(x) \ge 0] - \Pr_{x\sim\gamma_n}[r_{s'}(x) \cdot r_s(x) \ge 0]\big|\Big) \\
&\le \frac{3\xi}{2},
\end{aligned} \tag{7.11}$$

where the last inequality follows from the first and third items of Corollary 7.32. Now observe that

$$\Pr_{x\sim\gamma_n}[p_s(x)\geq 0]-\sum_{s'\neq s}\Pr_{x\sim\gamma_n}[p_s(x)\geq 0\wedge p_{s'}(x)\geq 0]\leq \Pr_{x\sim\gamma_n}[x\in \mathsf{Unique}(f,s)]\leq \Pr_{x\sim\gamma_n}[p_s(x)\geq 0],$$

$$\Pr_{x\sim\gamma_n}[r_s(x)\geq 0]-\sum_{s'\neq s}\Pr_{x\sim\gamma_n}[r_s(x)\geq 0\wedge r_{s'}(x)\geq 0]\leq \Pr_{x\sim\gamma_n}[x\in \mathsf{Unique}(f_{\mathsf{junta}},s)]\leq \Pr_{x\sim\gamma_n}[r_s(x)\geq 0].$$

Subtracting these two and applying (7.11), we obtain

$$\begin{aligned}\Big|\Pr_{x\sim\gamma_n}[x\in \mathsf{Unique}(f,s)]-\Pr_{x\sim\gamma_n}[x\in \mathsf{Unique}(f_{\mathsf{junta}},s)]\Big| &\leq \Big|\Pr_{x\sim\gamma_n}[p_s(x)\geq 0]-\Pr_{x\sim\gamma_n}[r_s(x)\geq 0]\Big| \\ &+ \Big|\sum_{s'\neq s}\Pr_{x\sim\gamma_n}[p_s(x)\geq 0\wedge p_{s'}(x)\geq 0]\Big|+\frac{3k\xi}{2}.\end{aligned}$$

Adding over all s and applying the first item of Corollary 7.32, we obtain

$$\sum_{s=1}^{k}\Big|\Pr_{x\sim\gamma_n}[x\in \mathsf{Unique}(f,s)]-\Pr_{x\sim\gamma_n}[x\in \mathsf{Unique}(f_{\mathsf{junta}},s)]\Big|\leq k\cdot\xi+\frac{3k^2\xi}{2}+\sum_{s'\neq s}\Pr_{x\sim\gamma_n}[p_s(x)\geq 0\wedge p_{s'}(x)\geq 0]$$

Noting that for all (s,s'), $\Pr[p_s(x)\geq 0\wedge p_{s'}(x)\geq 0]\leq \Pr[x\in \mathsf{Collision}(f)]$, we obtain

$$\sum_{s=1}^{k}\Big|\Pr_{x\sim\gamma_n}[x\in \mathsf{Unique}(f,s)]-\Pr_{x\sim\gamma_n}[x\in \mathsf{Unique}(f_{\mathsf{junta}},s)]\Big|\leq k^2\cdot\Pr[x\in \mathsf{Collision}(f)]+k\cdot\xi+\frac{3k^2\xi}{2}.$$

This however immediately implies that

$$\Pr_x[x\in \mathsf{Collision}(f_{\mathsf{junta}})]\leq (k^2+1)\cdot\Pr[x\in \mathsf{Collision}(f)]+k\cdot\xi+\frac{3k^2\xi}{2}.$$ □

Lemma 7.37.

$$\big|\mathbf{E}[\langle f,P_tf\rangle]-\mathbf{E}[\langle f_{\mathsf{junta}},P_tf_{\mathsf{junta}}\rangle]\big|\leq 2\Pr[x\in \mathsf{Collision}(f)]+2\Pr[x\in \mathsf{Collision}(f_{\mathsf{junta}})]+\frac{3k\xi}{2}.$$

Proof. We begin by noting the following two inequalities :

$$\Big|\mathbf{E}[\langle f,P_tf\rangle]-\sum_{s=1}^{k}\Pr[p_s(x)\geq 0\wedge u_s(x,y)\geq 0]\Big|\leq 2\Pr[x\in \mathsf{Collision}(f)],$$

$$\Big|\mathbf{E}[\langle f_{\mathsf{junta}},P_tf_{\mathsf{junta}}\rangle]-\sum_{s=1}^{k}\Pr[r_s(x)\geq 0\wedge v_s(x,y)\geq 0]\Big|\leq 2\Pr[x\in \mathsf{Collision}(f_{\mathsf{junta}})].$$

Thus,

$$\begin{aligned}\big|\mathbf{E}[\langle f,P_tf\rangle]-\mathbf{E}[\langle f_{\mathsf{junta}},P_tf_{\mathsf{junta}}\rangle]\big| &\leq 2\Pr[x\in \mathsf{Collision}(f)]+2\Pr[x\in \mathsf{Collision}(f_{\mathsf{junta}})] \\ &+\sum_{s=1}^{k}\big|\Pr[p_s(x)\geq 0\wedge u_s(x,y)\geq 0]-\Pr[r_s(x)\geq 0\wedge v_s(x,y)\geq 0]\big|.\end{aligned} \tag{7.12}$$

We now seek to bound $\Pr[p_s(x) \geq 0 \wedge u_s(x) \geq 0] - \Pr[r_s(x) \geq 0 \wedge v_s(x) \geq 0]$. As before, using that the polynomials p_s, u_s, r_s and v_s are zero only on a measure zero set and using the identity from (7.10), we obtain that

$$\begin{aligned} & \left|\Pr[p_s(x) \geq 0 \wedge u_s(x,y) \geq 0] - \Pr[r_s(x) \geq 0 \wedge v_s(x,y) \geq 0]\right| \\ \leq \; & \frac{1}{2} \cdot \left|\Pr[p_s(x) \geq 0] - \Pr[r_s(x) \geq 0]\right| + \frac{1}{2} \cdot \left|\Pr[u_s(x,y) \geq 0] - \Pr[v_s(x,y) \geq 0]\right| \\ + \; & \frac{1}{2} \cdot \left|\Pr[p_s(x) \cdot u_s(x,y) \geq 0] - \Pr[r_s(x) \cdot v_s(x,y) \geq 0]\right|. \end{aligned}$$

Applying Corollary 7.32 to bound all the terms, we obtain that

$$\left|\Pr[p_s(x) \geq 0 \wedge u_s(x,y) \geq 0] - \Pr[r_s(x) \geq 0 \wedge v_s(x,y) \geq 0]\right| \leq \frac{3\xi}{2}.$$

Plugging this bound back to (7.12), we obtain that

$$\left|\mathbf{E}[\langle f, P_t f\rangle] - \mathbf{E}[\langle f_{\mathsf{junta}}, P_t f_{\mathsf{junta}}\rangle]\right| \leq 2\Pr[x \in \mathsf{Collision}(f)] + 2\Pr[x \in \mathsf{Collision}(f_{\mathsf{junta}})] + \frac{3k\xi}{2}. \qquad \square$$

We are now in a position to finish the proof of Theorem 5.2 and show that the construction f_{junta} indeed satisfies the required properties. Note that by construction, f_{junta} is a degree-d PTF over n_0 variables where as we have said before, once $\beta(\cdot)$ is a computable function, $n_0 = n_0(k, \varepsilon, d, t)$ is a computable function. Further, it is easy to see that $\beta(\cdot)$ can be chosen to be computable; we just need to decrease sufficiently quickly to satisfy Condition 7.31.

We first bound $\|\mathbf{E}[f_{\mathsf{junta}}] - \mu\|_1 = \|\mathbf{E}[f_{\mathsf{junta}}] - \mathbf{E}[f]\|_1$. To do this, note that

$$\begin{aligned} \|\mathbf{E}[f_{\mathsf{junta}}] - \mathbf{E}[f]\|_1 \leq \sum_{s=1}^{k} |\Pr[\tilde{p}_s(x) \geq 0] - \Pr[r_s(x) \geq 0]| + \Pr[x \in \mathsf{Collision}(f)] \\ + \Pr[x \in \mathsf{Collision}(f_{\mathsf{junta}})] \\ \leq k \cdot \xi + \Pr[x \in \mathsf{Collision}(f)] + (k^2+1) \cdot \Pr[x \in \mathsf{Collision}(f)] + k \cdot \xi + \frac{3k^2 \cdot \xi}{2}. \end{aligned}$$

Here the last inequality follows by applying the first item of Corollary 7.32 and Claim 7.36. Plugging in the value of $\xi = \varepsilon/(20k^2)$ and the upper bound on $\Pr[x \in \mathsf{Collision}(f)]$, we obtain that $\|\mathbf{E}[f_{\mathsf{junta}}] - \mathbf{E}[f]\|_1 \leq \varepsilon$.

For the second part, note that applying Lemma 7.37, we obtain

$$\langle f_{\mathsf{junta}}, P_t f_{\mathsf{junta}}\rangle \geq \langle f, P_t f\rangle - 2\Pr[x \in \mathsf{Collision}(f)] - 2\Pr[x \in \mathsf{Collision}(f_{\mathsf{junta}})] - \frac{3k\xi}{2}.$$

Again plugging in the value $\xi = \varepsilon/(20k^2)$ and the upper bound on $\Pr[x \in \mathsf{Collision}(f)]$, we obtain that

$$\langle f_{\mathsf{junta}}, P_t f_{\mathsf{junta}}\rangle \geq \langle f, P_t f\rangle - \varepsilon.$$

This completes the proof of Theorem 5.2.

References

[1] DOMINIQUE BAKRY AND MICHEL LEDOUX: Lévy–Gromov's isoperimetric inequality for an infinite dimensional diffusion generator. *Inventiones Math.*, 123(1):259–281, 1996. [doi:10.1007/BF01232376] 17

[2] ITAI BENJAMINI, GIL KALAI, AND ODED SCHRAMM: Noise sensitivity of Boolean functions and applications to percolation. *Publ. Math. Inst. Hautes Études Sci.*, 90:5–43, 1999. [doi:10.1007/BF02698830, arXiv:math/9811157] 2

[3] CHRISTER BORELL: Geometric bounds on the Ornstein-Uhlenbeck velocity process. *Zeitschrift für Wahrscheinlichkeitstheorie und Verw. Gebiete*, 70(1):1–13, 1985. [doi:10.1007/BF00532234] 2

[4] ANTHONY CARBERY AND JAMES WRIGHT: Distributional and L^q norm inequalities for polynomials over convex bodies in $\mathbb{R}^n$. *Mathematical Research Letters*, 8(3):233–248, 2001. [doi:10.4310/MRL.2001.v8.n3.a1] 22

[5] JOE CORNELI, IVAN CORWIN, STEPHANIE HURDER, VOJISLAM SESUM, YA XU, ELIZABETH ADAMS, DIANA DAVIS, MICHELLE LEE, REGINA VISOCCHI, NEIL HOFFMAN, AND ROBERT HARDT: Double bubbles in Gauss space and spheres. *Houston J. Math.*, 34(1):181–204, 2008. 3

[6] ANINDYA DE, ELCHANAN MOSSEL, AND JOE NEEMAN: Majority is stablest: Discrete and SoS. *Theory of Computing*, 12(4):1–50, 2016. Preliminary version in STOC'13. [doi:10.4086/toc.2016.v012a004, arXiv:1211.1001] 7

[7] ANINDYA DE, ELCHANAN MOSSEL, AND JOE NEEMAN: Noise stability is computable and approximately low-dimensional. In *Proc. 32nd Computational Complexity Conf. (CCC'17)*, pp. 10:1–10:11. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017. (conference version of this paper). [doi:10.4230/LIPIcs.CCC.2017.10, arXiv:1701.01483]

[8] ANINDYA DE, ELCHANAN MOSSEL, AND JOE NEEMAN: Non interactive simulation of correlated distributions is decidable. In *Proc. 29th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'18)*, pp. 2728–2746. ACM Press, 2018. [doi:10.1137/1.9781611975031.174, arXiv:1701.01485] 5, 9

[9] ANINDYA DE AND ROCCO A. SERVEDIO: Efficient deterministic approximate counting for low-degree polynomial threshold functions. In *Proc. 46th STOC*, pp. 832–841. ACM Press, 2014. [doi:10.1145/2591796.2591800, arXiv:1311.7178] 4, 5, 6, 20, 21, 22, 26, 27, 28, 29

[10] IRIT DINUR, EHUD FRIEDGUT, GUY KINDLER, AND RYAN O'DONNELL: On the Fourier tails of bounded functions over the discrete cube. *Israel J. Math.*, 160(1):389–412, 2007. Preliminary version in STOC'06. [doi:10.1007/s11856-007-0068-9] 11

[11] HERBERT FEDERER: *Geometric Measure Theory*. Springer, 1969. [doi:10.1007/978-3-642-62010-2] 5, 17

[12] BADIH GHAZI, PRITISH KAMATH, AND MADHU SUDAN: Decidability of non-interactive simulation of joint distributions. In *Proc. 57th FOCS*, pp. 545–554. IEEE Comp. Soc. Press, 2016. [doi:10.1109/FOCS.2016.65] 8

[13] STEVEN HEILMAN: Euclidean partitions optimizing noise stability. *Electron. J. Probab.*, 19(71):1–37, 2014. [doi:10.1214/EJP.v19-3083, arXiv:1211.7138] 3

[14] STEVEN HEILMAN, ELCHANAN MOSSEL, AND JOE NEEMAN: Standard simplices and pluralities are not the most noise stable. *Israel J. Math.*, 213(1):33–53, 2016. Preliminary version in ITCS'15. [doi:10.1007/s11856-016-1320-y, arXiv:1403.0885] 3

[15] MICHAEL HUTCHINGS, FRANK MORGAN, MANUEL RITORÉ, AND ANTONIO ROS: Proof of the double bubble conjecture. *Electron. Res. Announc. Amer. Math. Soc.*, 6:45–49, 2000. [doi:10.1090/S1079-6762-00-00079-2, arXiv:math/0406017] 3

[16] MICHAEL HUTCHINGS, FRANK MORGAN, MANUEL RITORÉ, AND ANTONIO ROS: Proof of the double bubble conjecture. *Ann. of Math.*, 155(2):459–489, 2002. [doi:10.2307/3062123, arXiv:math/0406017] 3

[17] MARCUS ISAKSSON AND ELCHANAN MOSSEL: Maximally stable Gaussian partitions with discrete applications. *Israel J. Math.*, 189(1):347–396, 2012. [doi:10.1007/s11856-011-0181-7, arXiv:0903.3362] 2, 7

[18] SVANTE JANSON: *Gaussian Hilbert Spaces*. Cambridge University Press, 1997. [doi:10.1017/CBO9780511526169] 20, 21, 30

[19] GIL KALAI: A Fourier-theoretic perspective on the Condorcet paradox and Arrow's theorem. *Adv. in Appl. Math.*, 29(3):412–426, 2002. [doi:10.1016/S0196-8858(02)00023-4]

[20] SUBHASH KHOT, GUY KINDLER, ELCHANAN MOSSEL, AND RYAN O'DONNELL: Optimal inapproximability results for Max-Cut and other 2-variable CSPs? In *Proc. 45th FOCS*, pp. 146–154. IEEE Comp. Soc. Press, 2004. [doi:10.1109/FOCS.2004.49] 3, 45

[21] SUBHASH KHOT, GUY KINDLER, ELCHANAN MOSSEL, AND RYAN O'DONNELL: Optimal inapproximability results for Max-Cut and other 2-variable CSPs? *SIAM J. Comput.*, 37(1):319–357, 2007. Preliminary version [20]. [doi:10.1137/S0097539705447372] 2

[22] PRAVESH KOTHARI, AMIR NAYYERI, RYAN O'DONNELL, AND CHENGGANG WU: Testing surface area. In *Proc. 25th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'14)*, pp. 1204–1214. ACM Press, 2014. [doi:10.1137/1.9781611973402.89] 5, 14

[23] MICHEL LEDOUX: Semigroup proofs of the isoperimetric inequality in Euclidean and Gauss space. *Bull. des Sciences Mathématiques*, 118(6):485–510, 1994. 2, 5, 14

[24] FRANCESO MAGGI: *Sets of Finite Perimeter and Geometric Variational Problems: An Introduction to Geometric Measure Theory*. Cambridge University Press, 2012. [doi:10.1017/CBO9781139108133] 18

[25] ELCHANAN MOSSEL, RYAN O'DONNELL, AND KRZYSZTOF OLESZKIEWICZ: Noise stability of functions with low influences: Invariance and optimality. *Ann. of Math.*, 171(1):295–341, 2010. Preliminary version in FOCS'05. [doi:10.4007/annals.2010.171.295, arXiv:math/0503503] 2, 3, 6, 7

[26] JOE NEEMAN: Testing surface area with arbitrary accuracy. In *Proc. 46th STOC*, pp. 393–397. ACM Press, 2014. [doi:10.1145/2591796.2591807, arXiv:1309.1387] 5

[27] JOHN W. NEUBERGER: Norm of symmetric product compared with norm of tensor product. *Linear and Multilinear Algebra*, 2(2):115–121, 1974. [doi:10.1080/03081087408817047] 39

[28] IVAN NOURDIN: Lectures on Gaussian approximations with Malliavin calculus. In *Séminaire de Probabilités XLV*, pp. 3–89. Springer, 2013. [doi:10.1007/978-3-319-00321-4_1, arXiv:1203.4147] 21

[29] RYAN O'DONNELL: *Analysis of Boolean Functions*. Cambridge Univ. Press, 2014. 2

[30] PRASAD RAGHAVENDRA AND DAVID STEURER: How to round any CSP. In *Proc. 50th FOCS*, pp. 586–594. IEEE Comp. Soc. Press, 2009. [doi:10.1109/FOCS.2009.74] 7

[31] BEN W. REICHARDT: Proof of the double bubble conjecture in $\mathbb{R}^n$. *J. Geom. Anal.*, 18(1):172–191, 2008. [doi:10.1007/s12220-007-9002-y, arXiv:0705.1601] 3

[32] BEN W. REICHARDT, CORY HEILMANN, YUAN Y. LAI, AND ANITA SPIELMAN: Proof of the double bubble conjecture in $\mathbb{R}^4$ and certain higher dimensional cases. *Pacific J. Math.*, 208(2):347–366, 2003. [PJM]. 3

AUTHORS

Anindya De
Assistant professor
University of Pennsylvania
anindyad@seas.upenn.edu
http://www.seas.upenn.edu/~anindya

Elchanan Mossel
Professor
Massachusetts Institute of Technology
elmos@mit.edu
http://www.math.mit.edu/~elmos/

Joe Neeman
Assistant professor
University of Texas, Austin
jneeman@math.utexas.edu
https://web.ma.utexas.edu/users/jneeman/

ABOUT THE AUTHORS

ANINDYA DE is an assistant professor of computer science at the University of Pennsylvania. He received his Ph. D. from U. C. Berkeley in 2013, advised by Luca Trevisan. He was a postdoc at the Institute for Advanced Study and DIMACS from 2013 to 2015 and an assistant professor at Northwestern University from 2015 to 2018. His research interests are in complexity theory, learning theory, discrete analysis and applied probability.

ELCHANAN MOSSEL is a professor of mathematics at MIT. He received his Ph. D. in mathematics from Hebrew University in 2000, advised by Yuval Peres. After his Ph. D., he was a postdoc in the Theory Group of Microsoft Research, Redmond, and a Miller fellow in Statistics and Computer Science at U. C. Berkeley. His research interests include combinatorial statistics, discrete Fourier analysis and influences, randomized algorithms, computational complexity, MCMC, Markov random fields, social choice, game theory and evolution.

JOE NEEMAN is a professor of mathematics at the University of Bonn and an assistant professor of mathematics at the University of Texas, Austin. He received his Ph. D. in 2013 from U. C. Berkeley, advised by Elchanan Mossel, after which he was a postdoc at U. T. Austin. He works in various areas of probability and on applications to computer science.