

An Invariance Principle for the Multi-slice, with Applications

Mark Braverman^{*} Subhash Khot[†] Noam Lifshitz[‡] Dor Minzer[§]

Abstract

Given an alphabet size $m \in \mathbb{N}$ thought of as a constant, and $\vec{k} = (k_1, \dots, k_m)$ whose entries sum of up n , the $\vec{k}$ -multi-slice is the set of vectors $x \in [m]^n$ in which each symbol $i \in [m]$ appears precisely k_i times. We show an invariance principle for low-degree functions over the multi-slice, to functions over the product space $([m]^n, \mu^n)$ in which $\mu(i) = k_i/n$. This answers a question raised by [22].

As applications of the invariance principle, we show:

1. An analogue of the “dictatorship test implies computational hardness” paradigm for problems with perfect completeness, for a certain class of dictatorship tests. Our computational hardness is proved assuming a recent strengthening of the Unique-Games Conjecture, called the Rich 2-to-1 Games Conjecture.

Using this analogue, we show that assuming the Rich 2-to-1 Games Conjecture, (a) there is an r -ary CSP $\mathcal{P}_r$ for which it is NP-hard to distinguish satisfiable instances of the CSP and instances that are at most $\frac{2r+1}{2^r} + o(1)$ satisfiable, and (b) hardness of distinguishing 3-colorable graphs, and graphs that do not contain an independent set of size $o(1)$.

2. A reduction of the problem of studying expectations of products of functions on the multi-slice to studying expectations of products of functions on correlated, product spaces. In particular, we are able to deduce analogues of the Gaussian bounds from [40] for the multi-slice.
3. In a companion paper, we show further applications of our invariance principle in extremal combinatorics, and more specifically to proving removal lemmas of a wide family of hypergraphs H called ζ -forests, which is a natural extension of the well-studied case of matchings.

1 Introduction

The invariance principle of Mossel, O’Donnell and Oleszkiewicz [42] is a powerful analytical tool with wide range of applications throughout theoretical computer science, discrete mathematics and combinatorics.¹ Initially motivated by proving optimal hardness of approximation result for Max-Cut [35] (assuming the Unique-Games Conjecture [34]), this technique has become ubiquitous in analysis of PCP constructions, and also has many other significant applications in different areas of discrete mathematics.

^{*}Department of Computer Science, Princeton University. Research supported in part by the NSF Alan T. Waterman Award, Grant No. 1933331, a Packard Fellowship in Science and Engineering, and the Simons Collaboration on Algorithms and Geometry.

[†]Department of Computer Science, Courant Institute of Mathematical Sciences, New York University. Supported by the NSF Award CCF-1422159, the Simons Collaboration on Algorithms and Geometry, and the Simons Investigator Award.

[‡]Einstein Institute of Mathematics, Hebrew University, Jerusalem, Israel. Research supported in part by ERC advanced grant 834735.

[§]Department of Mathematics, Massachusetts Institute of Technology, Cambridge, USA. Supported by a Sloan Research Fellowship.

¹Earlier versions of the invariance principle have been proved earlier by Rotar [46] as well as by Chatterjee [10], however they are less relevant with respect to the applications to PCP’s, which are a key point of this paper.

The invariance principle has been since extended and generalized in several ways: for general product spaces and operators [40], under weaker assumptions [41], and to a limited extent also beyond product spaces [22, 23]. These extensions are crucial in the analysis of many PCP constructions (especially based on the Unique-Games Conjecture), a prime example being Raghavendra’s result regarding the optimality of semi-definite algorithms for CSPs [44]. They are also useful in other areas such as coding theory [27], learning theory [20], property testing [7], as well as in combinatorics [11, 38, 33].

The main goal of this paper is to establish an invariance principle for the multi-slice, a generalization of the Boolean slice for which an invariance principle was recently established [22, 23]. We also give applications of this invariance principle to hardness of approximation, and extremal combinatorics.

1.1 The multi-slice

Much of this paper is devoted to the study of multi-slices, defined below.

Definition 1.1. Let $n \in \mathbb{N}$ and $\Sigma = [m]$ be a finite alphabet. For a tuple of non-negative integers $\vec{k} = (k_1, \dots, k_m)$ that add up to n , we define the multi-slice $\mathcal{U}_{\vec{k}} \subseteq [m]^n$ as the set of vectors $x \in [m]^n$ that for all $i \in [m]$, have exactly k_i of their coordinates equal to i .

We will mostly be concerned with α -balanced multi-slices, which are multi-slices $\mathcal{U}_{\vec{k}} \subseteq [m]^n$ in which $k_i \geq \alpha n$ for all $i \in [m]$. Here and throughout, the alphabet m and the parameter α are thought of as constants.² Given a multi-slice $\mathcal{U}_{\vec{k}}$, it is natural to consider its product analogue, i.e. $[m]^n$ along with the product measure $\nu_{\vec{k}}^{\otimes n}$ defined by $\nu_{\vec{k}}(i) = \frac{k_i}{n}$, and ask how different do the domains $(\mathcal{U}_{\vec{k}}, \text{Uniform})$ and $([m]^n, \nu_{\vec{k}}^{\otimes n})$ behave. More precisely, we will be interested in the following type of invariance between the two domains: given a function $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$, can one associate with f a function $\tilde{f}: [m]^n \rightarrow \mathbb{R}$, such that f and $\tilde{f}$ are very close to each other? For the function to make sense, we must of course define what “close to each other” even means – these are functions over two different domains, so distance measures such as the L^2 distance do not make sense. To define things more precisely, we note that a typical point $y \sim \nu_{\vec{k}}^{\otimes n}$ is quite close to a point in the multi-slice – its distance from it is typically $O_m(\sqrt{n})$. Thus, it makes sense to consider a coupling between the two domains, say (x, y) , such that x is (marginally) a uniform point from the multi-slice, y is (marginally) distributed according to $\nu_{\vec{k}}^{\otimes n}$, and typically the Hamming distance between x and y is $O_m(\sqrt{n})$. In this language, we will say that f and $\tilde{f}$ are close if $f(x) \approx \tilde{f}(y)$ where we sample x, y according to the above coupled distribution.³

Unfortunately, one cannot hope for a statement of this great generality. The issue is that “high-degree” functions may be able to distinguish the multi-slice from its product analogue. To be more concrete, take any $\mathcal{A} \subseteq \binom{[n]}{k_1}$ of size $\frac{1}{2} \binom{n}{k_1}$, and consider the function $f(x) = \sum_{A \in \mathcal{A}} 1_{x_A=1^{k_1}}$. Then there is no $\tilde{f}$ that is close to f in the above sense – the issue is that there is no good choice for $\tilde{f}(y)$ when the number of 1’s in $\tilde{f}$ is different from k_1 (which happens with probability $1 - o(1)$).

Our first result asserts that “high degree” functions are in fact the only obstruction to an invariance between the multi-slice and its product analogue. Namely, we show that if f is a low-degree function, then one may find $\tilde{f}$ as above. To make a formal statement, we require a few definitions. First, we formally define the degree of a function on the multi-slice.

²There are natural extensions of our results to the case that m is not necessarily a constant but is allowed to grow with n , but we omit them from this paper.

³For our applications, we will need additional natural properties from the associated function $\tilde{f}$. These properties are mainly concerned with the notion of “influence” and “noisy-influence” of coordinates, and we will want them to be roughly preserved.

Definition 1.2. For a multi-slice $\mathcal{U}_{\vec{k}}$ over alphabet $[n]$, we say a function $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ is a d -junta if there exists a set of coordinates $A \subseteq [n]$ of size at most d , such that $f(x) = g(x_A)$ for some function $g: [m]^A \rightarrow \mathbb{R}$.

We can now define the space of degree d functions. We will use expectation inner products throughout the paper, i.e. for $f, g: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ we define $\langle f, g \rangle = \mathbb{E}_{\mathbf{x} \in_R \mathcal{U}_{\vec{k}}} [f(\mathbf{x})g(\mathbf{x})]$.

Definition 1.3. For $d = 0, 1, \dots, n$, we denote by $V_d(\mathcal{U}_{\vec{k}}) \subseteq \{f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}\}$ the span of d -juntas. We often refer to this space as the space of degree d functions. We also define $V_{=d}(\mathcal{U}_{\vec{k}}) = V_d(\mathcal{U}_{\vec{k}}) \cap V_{d-1}(\mathcal{U}_{\vec{k}})^\perp$.

Thus, slightly informally, we will often refer to functions from $V_d(\mathcal{U}_{\vec{k}})$ for small d as low-degree functions, and refer to functions perpendicular to that space, i.e. in $V_d(\mathcal{U}_{\vec{k}})^\perp$, as high-degree functions.

Next, we wish to define the notion of couplings that will be central in our paper, and we first need to define symmetric distributions.

Definition 1.4. For $r \in \mathbb{N}$, we say a distribution μ over $([m_1] \times \dots \times [m_r])^n$ is symmetric under the action of S_n (or simply symmetric) if for all $\pi \in S_n$, the distribution of $(\pi(\mathbf{x}(1)), \dots, \pi(\mathbf{x}(r)))$, where $(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu$, is the same as μ .

We will often identify the domains $([m_1] \times \dots \times [m_r])^n$ and $[m_1]^n \times \dots \times [m_r]^n$, choosing the precise view depending on the context. Next, we define a specific type of couplings that will be useful for us in this paper.

Definition 1.5. Let $\alpha, \zeta > 0$, $m \in \mathbb{N}$. For symmetric distributions ν_1, ν_2 over $[m]^n$, a (α, ζ) -coupling between $([m]^n, \nu_1)$ and $([m]^n, \nu_2)$ is a jointly distributed $\mathcal{C} = (\mathbf{x}, \mathbf{y})$ satisfying the following properties.

1. The marginal distribution of $\mathbf{x}$ is ν_1 , and the marginal distribution of $\mathbf{y}$ is ν_2 .
2. The distribution of $\mathcal{C}$ is symmetric under the action of S_n .
3. For all $i \in [n]$, $\Pr_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}} [\mathbf{x}_i \neq \mathbf{y}_i] \leq \zeta$.
4. Tail bounds: for all $\varepsilon > 0$, it holds that $\Pr_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}} [|\{i \mid x_i \neq y_i\}| \geq \varepsilon n] \leq \frac{1}{\alpha} e^{-\alpha \varepsilon^2 n}$.

A basic example is given by Boolean slices. For example, setting $m = 2$ and taking ν_1 to be the uniform distribution over $\{0, 1\}^n$, ν_2 to be the uniform distribution over the vectors from $\{0, 1\}^n$ with Hamming weight $n/2$. Indeed, a valid coupling in this case is simple to construct: $(\mathbf{x}, \mathbf{y})$ are sampled by taking $\mathbf{x} \sim \nu_1$; if the Hamming weight of $\mathbf{x}$ is at least $n/2$, we take $\mathbf{y} \leq \mathbf{x}$ of Hamming weight $n/2$ randomly, and otherwise we take $\mathbf{y} \geq \mathbf{x}$ of Hamming weight $n/2$ randomly. One may show that this is an (α, ζ) -coupling for $\alpha = \Omega(1)$ and $\zeta = O\left(\frac{1}{\sqrt{n}}\right)$. We omit the proof of this fact as in Proposition 4.6, we give a more general construction that also applies to this case and gives the same parameters.

Earlier, we have seen that couplings allow us to define distance measures between functions on different domains. More importantly however, they also allow us to naturally lift functions from one domain to another, as described below. Suppose $\mathcal{C}$ is a coupling between $(\mathcal{U}_{\vec{k}}, \text{Uniform})$ and $([m]^n, \nu_{\vec{k}}^{\otimes n})$. We introduce the operators $T_{\mathcal{C}}: L^2(\mathcal{U}_{\vec{k}}, \text{Uniform}) \rightarrow L^2([m]^n, \nu_{\vec{k}}^{\otimes n})$ and its adjoint operator $T_{\mathcal{C}}^*: L^2([m]^n, \nu_{\vec{k}}^{\otimes n}) \rightarrow L^2(\mathcal{U}_{\vec{k}}, \text{Uniform})$, defined as

$$T_{\mathcal{C}}f(y) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}} [f(\mathbf{x}) \mid \mathbf{y} = y], \quad T_{\mathcal{C}}^*g(x) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}} [g(\mathbf{y}) \mid \mathbf{x} = x].$$

With this in mind, given a function $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$, our associated function on the product space will be $T_{\mathcal{C}}f$, where $\mathcal{C}$ is a coupling with sufficiently good parameters. Our invariance principle for low-degree functions reads:

Lemma 1.6. *Let $d, m, n \in \mathbb{N}$, $\alpha, \zeta > 0$, and let $\mathcal{U}_{\vec{k}}$ be multi-slice over alphabet $[m]$. Suppose $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ is a function of degree at most d , and $\mathcal{C}, \mathcal{C}'$ are (α, ζ) -coupling between $(\mathcal{U}_{\vec{k}})$ and $([m]^n, \nu_{\vec{k}}^{\otimes n})$. Then*

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}'} [(f(\mathbf{x}) - T_{\mathcal{C}}f(\mathbf{y}))^2] \leq 8\sqrt{d\zeta} \|f\|_2^2.$$

Next, to go beyond low-degree functions, we must consider settings in which the contributions of high-degree terms are vanishing (in which case the question is, in essence, about low-degree functions). In [42, 40], it was shown that connected product distributions provide an important such example, and we show that an appropriate version of this result continues to hold in the multi-slice. To formally define this setting, we require a few definitions. First, we define the notion of admissible distributions.

Definition 1.7. *For a distribution μ over $([m_1] \times \dots \times [m_r])^n$ and $\vec{a} \in [m_1] \times \dots \times [m_r]$, we denote*

$$\mu_{\vec{a}} = \Pr_{\substack{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu \\ i \in [n]}} [\mathbf{x}(1)_i = a_1, \dots, \mathbf{x}(r)_i = a_r].$$

Definition 1.8. *Let $\alpha \in (0, 1)$, $r \in \mathbb{N}$ and let $\mathcal{U}_{\vec{k}(1)}, \dots, \mathcal{U}_{\vec{k}(r)}$ be multi-slices over alphabets $[m_1], \dots, [m_r]$ respectively. We say a distribution μ over $\mathcal{U}_{\vec{k}(1)} \times \dots \times \mathcal{U}_{\vec{k}(r)}$ is α -admissible if it satisfies the following properties.*

1. *The distribution μ is symmetric under S_n .*
2. *For all $\vec{a} \in [m_1] \times \dots \times [m_r]$ it either holds that $\mu_{\vec{a}} = 0$ or $\mu_{\vec{a}} \geq \alpha$.*

Next, we define the notion of connectedness for a distribution μ .

Definition 1.9. *Let $r, m_1, \dots, m_r \in \mathbb{N}$, and let μ be a distribution over $[m_1]^n \times \dots \times [m_r]^n$. We say μ is connected, if for all i , the graph $H_i = (V_1 \cup V_2, E)$ defined as follows is connected: (1) $V_1 \subseteq [m_i]^n$, $V_2 = \prod_{j \neq i} [m_j]^n$ are the support of the corresponding marginal distributions of μ , and (2) (v_1, v_2) is an edge if $x \in [m_1]^n \times \dots \times [m_r]^n$ whose i th component is in v_1 , and the rest of its components are according to v_2 , is in the support of μ .*

Finally, we define the product analogue $\tilde{\mu}$ of a given distribution μ over the multi-slice.

Definition 1.10. *Let $\mathcal{U}_{\vec{k}(1)}, \dots, \mathcal{U}_{\vec{k}(r)}$ be multislices over alphabets $[m_1], \dots, [m_r]$, and let μ be a distribution over $\mathcal{U}_{\vec{k}(1)} \times \dots \times \mathcal{U}_{\vec{k}(r)}$. For each $i \in [n]$, define the distribution $\tilde{\mu}_i$ over $[m_1] \times \dots \times [m_r]$ as*

$$\tilde{\mu}_i(\vec{a}) = \Pr_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu} [\mathbf{x}(1)_i = a_1, \dots, \mathbf{x}(r)_i = a_r].$$

Then, the product analogue of μ is defined as $\tilde{\mu} = \prod_{i=1}^n \tilde{\mu}_i$.

We remark that we will mostly be interested with admissible distributions μ , hence the marginal distributions μ_i (as well as of $\tilde{\mu}_i$) will be the same for all $i \in [n]$. We can now state our main invariance principle.

Theorem 1.11. *For all $\alpha \in (0, 1)$, $M, r \in \mathbb{N}$, $m_1, \dots, m_r \in \mathbb{N}$, $\varepsilon > 0$ there are $\zeta > 0$, $N \in \mathbb{N}$ such that the following holds for $n \geq N$. Suppose that*

- *For $i = 1, \dots, r$, the multi-slice $\mathcal{U}_{\vec{k}(i)} \subseteq [m_i]^n$ is α -balanced, and $\mathcal{C}_i$ is a (α, ζ) -coupling between $(\mathcal{U}_{\vec{k}(i)}, \text{Uniform})$ and $([m_i]^n, \nu_{\vec{k}(i)}^{\otimes n})$;*
- *μ is a connected, α -admissible distribution over $\prod_{i=1}^r \mathcal{U}_{\vec{k}(i)}$;*
- *$\tilde{\mu}$ the product version of μ as in Definition 1.10, and there is an (α, ζ) -coupling between μ and $\tilde{\mu}$.*

Then for all $f_i: \mathcal{U}_{\vec{k}(i)} \rightarrow \mathbb{R}$ such that $\|f_i\|_{2r} \leq M$ for all i , it holds that

$$\left| \mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu} \left[\prod_{i=1}^r f_i(\mathbf{x}(i)) \right] - \mathbb{E}_{(\mathbf{y}(1), \dots, \mathbf{y}(r)) \sim \tilde{\mu}} \left[\prod_{i=1}^r \text{T}_{\mathcal{C}_i} f_i(\mathbf{y}(i)) \right] \right| \leq \varepsilon.$$

A simple example application for Theorem 1.11 is given in the Boolean slice and the Kneser graph on it. Namely, we consider the case that $r = 2$, $m_1 = m_2 = 2$, $\vec{k}(1) = \vec{k}(2) = \vec{k} = (\ell, n - \ell)$ for $\ell = 0.49n$ and μ is the uniform distribution on (x, y) from $\{(x, y) \in \mathcal{U}_{\vec{k}} \times \mathcal{U}_{\vec{k}} \mid (x_i, y_i) \neq (2, 2) \forall i \in [n]\}$. The conditions of the theorem are easily seen to hold in this case, and hence Theorem 1.11 applies. In this case, if f_1, f_2 are indicator functions of subsets of $\mathcal{U}_{\vec{k}}$, the first expectation in Theorem 3.10 counts the number of edges in the Kneser graph that go between these two subsets, and the theorem allows us to transfer this problem into a related problem over a product domain. In the latter domain good bounds are often already known (for example in [40] bounds for such expectations were proved by further moving to Gaussian space).

Remark 1.12. *A few remarks are in order.*

1. *First, our result is more general, in the sense that it works for all distributions for which the contribution of high-degree functions to expectations such as $\mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu} \left[\prod_{i=1}^r f_i(\mathbf{x}(i)) \right]$ is negligible. See Section 5 for details. Thus, the majority of the effort in the proof of Theorem 1.11 lies in showing that high-degree terms do not contribute much to the expectation of the product provided μ is connected and admissible.*
2. *We remark that for most distributions μ of interest, the existence of couplings as in Theorem 1.11 is easy to establish, and thus encourage the reader to think of these conditions as trivially holding. See Proposition 4.6 for a general construction of couplings (which captures all of the cases considered herein).*

We note that our invariance principle is different from the one proved in [22, 23], even when specialized to the Boolean slice, i.e. the alphabet $m = 2$. The reason is that therein, the authors find a canonical basis of the space of functions $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$, which then allows them to interpret them as functions over $\{0, 1\}^n$ without any ambiguity. The authors then show this extension then behaves very similarly to the original function f . Our associated function with f on the other hand, namely $\text{T}_{\mathcal{C}} f$, can be thought of as mild averaging of the function f . Typically, the parameter ζ of our coupling would be $\Theta(1/\sqrt{n})$, thus the points $(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}$ are typically of Hamming distance $\Theta(\sqrt{n})$ apart, and so $\text{T}_{\mathcal{C}} f(\mathbf{y})$ could be thought of as averaging over points x in the multi-slice that are at distance $\Theta(\sqrt{n})$. In this language, our invariance principle for low-degree functions (Lemma 1.6) asserts that this averaging barely affects the function f if its degree is

significantly smaller than $\sqrt{n}$, which is precisely the range of degrees that [22, 23] are able to handle. We think however that our version of the invariance principle is perhaps more natural. Two points to support that are as follows:

1. Our proofs are much more “generic”, i.e. less domain specific than of [22, 23], and can probably be extended to other S_n -modules and not only the multi-slice.
2. The functions f and $T_C f$ enjoy many of the same properties that we will be concerned with. For example, we have the following important properties:
 - (a) if all of the noisy influences of f are small (a notion that we only define in Section 6, but should be thought of as an appropriate analogue of the noisy influences from the Boolean cube), then all of the noisy influences of $T_C f$ are small.
 - (b) If f has a coordinate i with significant noisy influence, then the same coordinate would have significant noisy influence in $T_C f$.
 - (c) If f is pseudo-random in the sense that fixing constantly many coordinates does not change the average of f by too much, then the same holds for $T_C f$.

These two properties are very important for us in our applications, and can be used to convert any of the Gaussian bounds of [40] to the multi-slice (we do not do it explicitly in this paper).

1.2 Applications to Hardness of Approximation

In this section, we discuss an application of our invariance principle to the field of hardness of approximation. More specifically, we show that it enables us to lift a certain class dictatorship tests with perfect completeness to hardness results, assuming a conjecture of [9] that we discuss below.

1.2.1 The basic PCP Theorem

The primary goal in hardness of approximation is to show that approximate solutions to optimization problems within a certain approximation ratio is NP-hard. For some problems, the basic PCP theorem combined with the long-code based paradigm developed over the last 20 years, yield tight inapproximability results (such as the results in [30, 29]). This paradigm however fails to establish tight inapproximability results for many problems, such as Max-Cut and Vertex-Cover. Towards this end, the Unique-Games Conjecture was proposed [34], and is now known to imply tight inapproximability result for a wide range of optimization problems. A prime example is Raghavendra’s result [44] regarding constraint satisfaction problems, who showed a generic SDP algorithm for the class of CSPs, and proved that it is essentially optimal, assuming the Unique-Games Conjecture. We remark that despite recent progress establishing a weaker variant of the Unique-Games Conjecture, called the 2-to-1 Games Conjecture [36, 13, 12, 37, 4] (which can be used to prove a wide range of weaker inapproximability results, see [5]), proving Unique-Games Conjecture or some its consequences independently remains an open challenge.

To present the Unique-Games Conjecture, we first define the 2-Prover-1-Round problem.

Definition 1.13. A 2PIR Games instance $\Psi = (L \cup R, E, \Sigma_L, \Sigma_R, \Phi)$ consists of a regular, bipartite graph $(L \cup R, E)$, the alphabet Σ_L for the vertex set L , the alphabet Σ_R for the vertex set R , and a set of constraints $\Phi = \{\phi_e\}_{e \in E}$, one for each edge. Each vertex is supposed to receive a label from the respective alphabet. The constraint ϕ_e for an edge $e = (u, v) \in E$, $u \in L, v \in R$ is defined by a relation $\phi_e \subseteq \Sigma_L \times \Sigma_R$, thought of as the set of label-pairs to the vertices u and v that satisfy the constraint.

For $0 < s < c \leq 1$, we denote by $\text{Gap-2P1R}[c, s]$ the promise problem, where given an instance Ψ of 2P1R Games, one has to distinguish between the case that there is an assignment satisfying at least c fraction of the constraints in Ψ . Combining the basic PCP Theorem [19, 3, 2] with the Parallel Repetition Theorem [45], one can show that for all $s > 0$, the problem $\text{Gap-2P1R}[1, s]$ is NP-hard on alphabets of size $n = \text{poly}(1/\delta)$. This result serves as the starting point of many hardness of approximation results. However, for many other results, one needs a stronger structure from the constraints.

1.2.2 The Unique-Games Conjecture and Imperfect Completeness

Definition 1.14. A d -to-1 Games instance is a 2P1R Games instance $\Psi = (L \cup R, E, \Sigma_L, \Sigma_R, \Phi)$, in which $|\Sigma_L| = d |\Sigma_R|$ and each constraint $\phi_{u,v}: \Sigma_L \rightarrow \Sigma_R$ is a d -to-1 map.

Khot conjectured that for $d \geq 2$, the problem $\text{Gap-}d\text{-to-1Games}[1, s]$ remains NP-hard provided the alphabets are sufficiently large constant (depending on s). For $d = 1$ however, i.e. for 1-to-1 Games problem often called the Unique-Games problem, one cannot hope for such a strong hardness result: it is easy to see that given a satisfiable instance of Unique-Games, one can efficiently find a satisfying assignment. For that reason, to deal with Unique-Games, one has to give up the perfect completeness, and indeed the Unique-Games Conjecture reads:

Conjecture 1.15. For all $\varepsilon, s > 0$, there is $n \in \mathbb{N}$, such that given a Unique-Games instance Ψ with alphabet size n , it is NP-hard to distinguish between the following two cases:

YES case: there is an assignment to Ψ satisfying at least $1 - \varepsilon$ fraction of the constraints.

NO case: there is no assignment to Ψ satisfying more than s fraction of the constraints.

It is thus clear that when working with the Unique-Games Conjecture, one cannot hope to prove hardness for satisfiable instances of problems, and indeed all of the results based on the Unique-Games Conjecture (such as Raghavendra's result [44]) have imperfect completeness. We stress here that even though the difference between perfect and imperfect completeness is often thought of as minor, it is critical for some problems and affects their complexity dramatically: the problem of solving systems of linear equations of over finite fields is such problem [30], and the Unique-Games problem is likely to be such problem, to name a few.

This raises the following issue: how can one prove general hardness results for satisfiable instances, assuming some feasible complexity assumption? One option is to work with 2-to-1 Games for which the perfect completeness version is conjectured to be hard (and the imperfect completeness version was already proved to be hard). Alas, the structure of the constraints in 2-to-1 Games is often not strong enough to use the powerful reductions and tools used in the context of Unique-Games.

1.2.3 The Rich 2-to-1 Games Conjecture

Recently, a stronger variant of the 2-to-1 Games Conjecture, called the Rich 2-to-1 Games conjecture was proposed [9] in the hope to shed more light on Unique-Games and the complexity of satisfiable instances, and we present this conjecture below.

Let $\Psi = (L \cup R, E, \Sigma_L, \Sigma_R, \Phi)$ be a 2-to-1-Game, with $|\Sigma_L| = n$ and $|\Sigma_R| = n/2$. Fix a vertex $u \in L$. Let $e = (u, v) \in E$ be an edge incident on u and let π_e be the 2-to-1 projection defining that constraint. The map defines a partition of Σ_L as

$$\Sigma_L = \bigcup_{\rho \in \Sigma_R} \pi_e^{-1}(\rho)$$

into disjoint sets of size 2. Let us denote by $\mathcal{P}(u)$ the distribution over partitions of Σ_L into sets of size 2, given by first sampling a uniformly random edge $e = (u, v)$ incident on u and then outputting the partition of Σ_L as above.

Definition 1.16. *An instance of Rich 2-to-1 Games is an instance of 2-to-1 Games with the additional property that for every vertex $u \in L$, the distribution $\mathcal{P}(u)$ is uniform over all partitions of Σ_L into sets of size 2.*

For $0 < s < c \leq 1$, we denote by $\text{Gap-Rich-2-to-1}_n[c, s]$ the promise problem, where given an instance Ψ of Rich 2-to-1 Games, one has to distinguish between the case that there is an assignment satisfying at least c fraction of the constraints in Ψ , or all assignments satisfy at most s fraction of the constraints in Ψ . The following conjecture was made in [9]:

Conjecture 1.17. *For all $\delta > 0$, there is a sufficiently large even $n \in \mathbb{N}$ such that $\text{Gap-Rich-2-to-1}_n[1, \delta]$ is NP-hard.*

A few remarks are in order. First, one may consider the imperfect completeness variant of Conjecture 1.17, i.e. the question of whether $\text{Gap-Rich-2-to-1}_n[1 - \varepsilon, \delta]$ is NP-hard for all $\varepsilon, \delta > 0$ provided $n(\varepsilon, \delta)$ is large enough. It turns out that this problem is equivalent to the Unique-Games Conjecture [9], and thus Conjecture 1.17 is strictly stronger than the Unique-Games Conjecture, and thus can be thought of as a variant of it with perfect completeness. It is thus natural to ask whether one can base on it hardness of satisfiable instances of problems, the prime example being constraint satisfaction problems.

1.2.4 Our application: hardness from dictatorship tests with perfect completeness

Using our invariance principle, we are able to convert certain dictatorship tests to NP-hardness results, assuming Conjecture 1.17. This result generalizes a theme from the context of Unique-Games Conjecture based hardness results, for which one has an automatic machinery to convert dictatorship tests with imperfect completeness to hardness of approximation results (a theme that has started with in [35] and plays an important role in Raghavendra's work [44]). Thus, the task of proving hardness of approximation result for a problem with imperfect completeness reduces to the design of a certain dictatorship test.

Our result extends this theme to the realm of perfect completeness, however we are not able to handle dictatorship tests in full generality and require a certain technical condition that is nevertheless natural in existing dictatorship tests. We hope that as more general dictatorship tests emerge in the literature, the ideas in the work herein will be useful in converting these more general tests into hardness results.

A precise statement of our result (given in Theorem 6.7 and slightly extended in Remark 6.28) requires quite a bit of set-up, and is deferred to Section 6. Instead, we give two instantiations of it, using existing dictatorship tests from the literature.

For a collection of r -ary predicates $\mathcal{P}_r \subseteq \{P: \{0, 1\}^r \rightarrow \{0, 1\}\}$, an instance of $\text{CSP-}\mathcal{P}_r$ consists of a set of variables $X = \{x_1, \dots, x_m\}$, and a collection of constraints E of the form $P(x_{i_1}, \dots, x_{i_r}) = 1$ where $P \in \mathcal{P}_r$. The goal is to find a Boolean assignment to the variables that satisfies as many constraints as possible. We also define the gap version, $\text{Gap-CSP-}\mathcal{P}_r$ in the natural way. Using the dictatorship test of [6], we have:⁴

⁴Strictly speaking, in [6] only a single predicate $P_r: \{0, 1\}^r \rightarrow \{0, 1\}$ is used. However, they use the folding technique, and our collection $\mathcal{P}_r$ replaces this folding. For all $\vec{a} \in \{0, 1\}^r$, we define $P_{r, \vec{a}}: \{0, 1\}^r \rightarrow \{0, 1\}$ as $P_{r, \vec{a}}(x) = P_r(x_1 + a_1, \dots, x_r + a_r)$ (addition is done mod 2), and our collection $\mathcal{P}_r$ is $\{P_{r, \vec{a}} \mid \vec{a} \in \{0, 1\}^r\}$.

Corollary 1.18. *Assuming Conjecture 1.17, for all $r \in \mathbb{N}$ of the form $2^m - 1$ there is a collection of $\mathcal{P}_r \subseteq \{P: \{0, 1\}^r \rightarrow \{0, 1\}\}$ such that for all $\varepsilon > 0$, $\text{Gap-CSP-}\mathcal{P}_r[1, \frac{2r+1}{2^r} + \varepsilon]$ is NP-hard.*

Unconditionally, it is known that there is a predicate P for which $\text{Gap-CSP-}\mathcal{P}_r[1, \frac{2\tilde{O}(r^{1/3})}{2^r} + \varepsilon]$ is NP-hard [31]. Corollary 1.18 significantly improves upon this result, and is nearly tight (however, it is conditioned on Conjecture 1.17).

Next, we show that assuming Conjecture 1.17, one can prove strong NP-hardness results for 3-colorable graphs. Recently, it was shown that for all $d \geq 2$, the d -to-1 Games Conjecture implies the hardness of coloring 3-colorable graphs with constantly many colors [28]. We show that the Rich 2-to-1 Games Conjecture implies the same result with a stronger soundness guarantee asserting that there are no independent sets of significant size. More precisely, using the dictatorship test of [14], we get:

Corollary 1.19. *Assuming Conjecture 1.17, for $\delta > 0$, given a graph G it is NP-hard to distinguish between the cases:*

1. **YES case:** G is 3-colorable;
2. **NO case:** G has no independent set of fractional size δ .⁵

1.3 Applications to Combinatorics

In a subsequent paper [8], we show several applications of our invariance principle in extremal combinatorics. Below, we explain the sort of problems that will be addressed in that paper.

We write $[n] = \{1, \dots, n\}$ and $\binom{[n]}{k}$ for the family of all k -subsets of $[n]$. The following problems are fundamental in hypergraph theory.

1. *The hypergraph Turán problem:* Given $H \subseteq \binom{[n]}{k}$. How large can a family $\mathcal{F} \subseteq \binom{[n]}{k}$ be if it does not contain a copy of H ?
2. *The Ramsey problem:* What is the largest r , such that every r -colouring of $\binom{[n]}{k}$ contains a copy of H ?
3. *Removal lemmas:* Is it true that every family that contains few copies of H is close to a family without any copies of H ?

In the subsequent paper [8], we focus on the above problems in the case where k is linear in n . This regime was extensively studied when H is an s -matching, i.e. the hypergraph consisting of s pairwise disjoint edges. The methods used for tackling the problems above are varied and include shifting [25], algebraic topology [39, 1], Regularity [26], sharp threshold results [17], and agreement testing [11].

In the companion paper we throw our invariance principal into the mix. We are then able to extend these results from the case where H is an s -matching, to the case where H belongs to a wide class of hypergraphs we call ζ -forests.

Definition 1.20. *We say that a hypergraph H is a ζ -forest if its edges can be ordered as $A_1, \dots, A_r$, where $|A_i \setminus \bigcup_{j=1}^{i-1} A_j| > \zeta n$.*

⁵In particular, the chromatic number of G is at least $\lceil 1/\delta \rceil$.

We show that every family that contains few copies of a ζ -forest H is close to an H -free family, provided that $|\bigcup_{A \in H} A| \leq (1 - \zeta)n$. We also show that every r -colouring of $\binom{[n]}{k}$ contains a monochromatic copy of H , provided that $n \geq n_0(r, \zeta)$. Finally, we make progress on the Turán problem by reducing it to a special case where the H -free family $\mathcal{F} \subseteq \binom{[n]}{k}$ is a junta. This means that there exists a set J of size $O(1)$, such that the question whether A is in $\mathcal{F}$, depends only on $A \cap J$.

We remark that such a reduction was shown to be useful Ellis et al. [17], who considered the $\frac{1}{4}$ -forest $M_{2,t}$ consisting of two edges that intersect in t -elements. The Turán problem for $M_{2,t}$ corresponds to an old problem of Erdős and Sós [18]. One of the main ingredients in [17] was a reduction of the Erdős-Sós forbidden intersection problem to the special case where the $M_{2,t}$ -free family is a junta.

1.4 Our techniques

1.4.1 The proof of our invariance principles

Invariance principle for low-degree functions. The proof of our invariance principle for low-degree functions mainly involves representation-theoretic arguments, and builds on the approach of [21]. To prove Lemma 3.25, we reduce the problem into the problem of understanding eigenvalues of a certain operator associated with our couplings. For simplicity of presentation, let us consider the case that the couplings $\mathcal{C}$ and $\mathcal{C}'$ in Lemma 3.25 are identical, in which case we study the eigenvalues of the operator $S = T_{\mathcal{C}}^* T_{\mathcal{C}}: L^2(\mathcal{U}_{\vec{k}}) \rightarrow L^2(\mathcal{U}_{\vec{k}})$. Since this is a symmetric operator, we can find an eigenbasis of $L^2(\mathcal{U}_{\vec{k}})$ consisting of eigenvectors, and our claim reduces to showing that eigenvalues of eigenfunctions which are low-degree functions are close to 1.

Using the symmetry of $\mathcal{C}$, we show that the operator S preserve *juntas*. Namely, if $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ depends only on coordinates from a set $J \subseteq [n]$, then the value of Sf at a point x also only depends on x_J . This allows to further reduce the question of studying eigenvalues corresponding to low-degree functions, to studying eigenvalues corresponding to juntas. The last part is easy to do the, using fact that the probability the coupling changes a specific coordinate is small (this is the parameter ζ).

Invariance principle for admissible distributions. For simplicity of presentation, let us assume we are dealing with the case that $r = 2$ and that the two multi-slices are identical, in which case we denote the functions from Theorem 1.11 by f and g . Roughly speaking, the proof proceeds by analyzing the contribution of the high-degree parts of f and g and the low-degree parts separately, say $f = f^{\leq d} + f^{> d}$ and $g = g^{\leq d} + g^{> d}$. Lemma 1.6 allows us to handle the low-degree parts, and the main effort in our proof goes into upper bounding the contribution of the high-degree parts, e.g. $\mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \mu} [f^{\leq d}(\mathbf{x})g^{> d}(\mathbf{y})]$. We now turn this problem again into the problem of understanding eigenvalues of a certain operator. More specifically, we consider the operator $T_{\mu}: L^2(\mathcal{U}_{\vec{k}}) \rightarrow L^2(\mathcal{U}_{\vec{k}})$ defined as $T_{\mu}f(x') = \mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \mu} [f(\mathbf{x}) \mid \mathbf{x}' = x']$ so that the expectation that we wish to study is $\langle f^{\leq d}, T_{\mu}g^{> d} \rangle$, and using the Cauchy-Schwarz inequality it suffices to bound $\|T_{\mu}^* T_{\mu}g^{> d}\|$. In other words, we want to prove that the operator $T_{\mu}^* T_{\mu}$ contracts high degree functions very strongly, and to do that we study its eigenvalues.

To study the eigenvalues of $T_{\mu}^* T_{\mu}$ we use the trace method: first, note that this operator defines a random walk on $L^2(\mathcal{U}_{\vec{k}})$. By considering long (but constant length) walks, we show that this random walk mixes rather well – in the sense that the trace of the resulting operator is $(1 + \varepsilon)^n$. Intuitively, it is easy to observe that this would happen in the product analogue of μ , namely $\tilde{\mu}$, and we know by the coupling $\mathcal{C}$ that we may consider the coupling along the random walk, so that we expect the two traces to be same. Thus, using the trace method, we are able to bound eigenvalues whose multiplicity is exponentially large, say at least $(1 + 2\varepsilon)^n$, and this way we are able to handle functions of degree $\Omega(n)$.

To handle functions of degree $o(n)$, we also use the eigenvalues approach, and in this case show that for each eigenvalue θ of degree d function, we may find a d -junta which is an eigenfunction with eigenvalue θ , say on a set of variables $[d] \subseteq [n]$. At this point, we essentially project the problem on a set coordinates slightly larger than J – say $[3d]$, ignoring the rest of the coordinates. That is, we consider the long random walk $T_\mu^* T_\mu$ induces on $[n]^{3d}$, and show again the trace of this operator is small, namely $(1+\varepsilon)^{3d}$; intuitively this happens for similar reasons as before, except that n is effectively $3d$. Thus it remains to show that the multiplicity of our eigenvalue θ is at least $(1+2\varepsilon)^{3d}$, and for that we again use symmetry: we note that for each $\pi \in S_{3d}$, the function ${}^\pi f: \mathcal{U}_k \rightarrow \mathbb{R}$ defined as ${}^\pi f(x) = f(\pi(x))$ is an eigenfunction of $T_\mu^* T_\mu$ with eigenvalue θ , and then show that these functions span a subspace of dimension at least $(1+2\varepsilon)^{3d}$.

1.4.2 Converting dictatorship tests into hardness results

Our proof of Theorem 6.7, which converts a dictatorship test into a hardness result (and implies Corollaries 1.18 and 1.19) uses roughly the same reduction as Raghavendra [44]. We begin by outlining a simple adaptation of this reduction, which we do not know how to analyze, and then explain our adaptations and in particular how multi-slices enter the picture.

For simplicity of presentation, let us assume that a dictatorship test for a predicate $P: \Sigma^r \rightarrow \{0, 1\}$ is a distribution $\mathcal{D}$ over Σ^r such that:

1. If $f: \Sigma^n \rightarrow \Sigma$ is a dictatorship (i.e. $f(x) = x_i$ for some i), then

$$\mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mathcal{D}^{\otimes n}} [P(f(\mathbf{x}(1)), \dots, f(\mathbf{x}(r)))] = 1.$$

2. If $f: \Sigma^n \rightarrow \Sigma$ is far from dictatorship, i.e. has all of low-degree influences being small, then

$$\mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mathcal{D}^{\otimes n}} [P(f(\mathbf{x}(1)), \dots, f(\mathbf{x}(r)))] \leq s.$$

Given a Rich 2-to-1 Games instance, $\Psi = (L \cup R, E, [n], [n/2], \Phi)$, we wish to encode a label to $u \in L$ using the long-code of u , and encode a label to $v \in R$ using the long code of v . That is, we replace each $u \in L$ with the cloud $\{u\} \times \{0, 1\}^n$ and each $v \in R$ by $\{v\} \times \{0, 1\}^{n/2}$. The intention is that a label $\sigma \in [n]$ to u would be encoded by the dictatorship assignment $f_u: \Sigma^n \rightarrow \Sigma$ defined as $f_u(x) = x_\sigma$. Our goal is to test, using the predicate P and the dictatorship test $\mathcal{D}$, that the assignments f_u 's (and f_v 's) are correlated with legitimate codewords.

The idea then is to sample $v \in R$, and virtually perform the dictatorship test on its cloud. Namely, sample queries $(\mathbf{x}(1), \dots, \mathbf{x}(r))$ according to the given dictatorship test, then sample $u_1, \dots, u_r$ neighbours of v independently, lift the points $\mathbf{x}(i)$ to $\mathbf{x}(i)'$ on the long-code of u_i using the projection map ϕ_{v, u_i} and then test that $f_{u_1}(\mathbf{x}(1)'), \dots, f_{u_r}(\mathbf{x}(r)')$ satisfies the predicate P .

While the completeness of this reduction is straight-forward, the soundness analysis of this reduction is more tricky. For one, the distribution of the points $\mathbf{x}(i)'$ that we consider in the reduction is not very nice: ideally, we would have liked the distribution to be close to the distribution of $\mathbf{y}(i)$ if we sampled $(\mathbf{y}(1), \dots, \mathbf{y}(r))$ according to the dictatorship test in the cloud of u_i ; alas this is not true. The distribution of $\mathbf{x}(i)'$ is not a product distribution, its support is not even full (for example, each symbol appears even number of times) and in general it is not convenient to work with. This is an artifact of the fact the constraints maps ϕ_{v, u_i} are only 2-to-1, as opposed to 1-to-1 as in Raghavendra's theorem (in which case, these maps preserve the distributions on the long-codes up to relabeling the coordinates).

To bypass this issue, we note that if instead of the full long-code of u , we used a multi-slice $\mathcal{U}_{\vec{k}} \subseteq \Sigma^n$ with an appropriate $\vec{k}$ (actually, we will use several multislices, each matching the statistics of one of $\mathbf{y}(1), \dots, \mathbf{y}(r)$ above) and the same for v , our distributions would align perfectly. Indeed, we use this idea, and use our invariance principle to argue that the soundness of the multi-slice analogue of the dictatorship test $\mathcal{D}^{\otimes n}$ from above is roughly the same as in the product version. To finish up the proof, we must then define analogues of low-degree influences on the multi-slice (it will actually be more convenient for us to work with noisy influence), and prove several basic properties of them – see Section 6.1 for more details.

2 Preliminaries

Notations. Throughout the paper, we denote random variables by boldface letters (such as $\mathbf{x}, \mathbf{y}, \boldsymbol{\pi}$). We denote $[n] = \{1, \dots, n\}$, and let S_n be the symmetric group, i.e. the set of all permutations $\pi: [n] \rightarrow [n]$. For a vector x of length n and $A \subseteq [n]$, we denote by x_A the vector whose length is $|A|$ corresponding to the entries of x from A . For a vector x of length n and $\pi \in S_n$, we denote by ${}^\pi x$ the vector of length n whose i th entry is $x_{\pi(i)}$.

In this section, we present standard background regarding various decompositions of functions the multi-slice. More precisely, we will use the degree decomposition, as well as a refinement of it given by representation theory. Towards that end, it will be useful for us to think of the multi-slice $\mathcal{U}_{k_1, \dots, k_m}$ as a quotient space $S_n / (S_{k_1} \times S_{k_2} \times \dots \times S_{k_m})$, which will allow us to lift the standard representation-theoretic decomposition of functions over S_n to decompositions of functions over $\mathcal{U}_{k_1, \dots, k_m}$.

2.1 The degree decomposition and representation-theoretic decomposition over S_n

We begin by presenting the coarse degree decomposition of functions $f: S_n \rightarrow \mathbb{R}$.

Definition 2.1. We say a function $f: S_n \rightarrow \mathbb{R}$ is a d -junta if there exists a set of coordinates $A \subseteq [n]$ of size at most d , such that $f(\pi) = g(\pi(A))$ for some function $g: [n]^A \rightarrow \mathbb{R}$.

We can now define the space of degree d functions. For that, we introduce the expectation inner product: for $f, g: S_n \rightarrow \mathbb{R}$, we define $\langle f, g \rangle = \mathbb{E}_{\boldsymbol{\pi}} [f(\boldsymbol{\pi})g(\boldsymbol{\pi})]$.

Definition 2.2. For $d = 0, 1, \dots, n$, we denote by $V_d(S_n) \subseteq \{f: S_n \rightarrow \mathbb{R}\}$ the span of d -juntas. We often refer to this space as the space of degree d functions. We also define $V_{=d}(S_n) = V_d(S_n) \cap V_{d-1}(S_n)^\perp$.

Thus, one can write the space of real-valued functions as $V_{=0}(S_n) \oplus V_{=1}(S_n) \oplus \dots \oplus V_{=n-1}(S_n)$, and thus write any $f: S_n \rightarrow \mathbb{R}$ uniquely as $f = \sum_{i=0}^{n-1} f^{=i}$ where $f^{=i} \in V_{=i}(S_n)$. We next present a refinement of this decomposition, given by representation theory of S_n .

Partitions and the decomposition by partitions

A partition of n , often denoted as $\lambda \vdash n$, is a monotonically decreasing sequence of positive integers, $\lambda = (\lambda_1, \dots, \lambda_r)$, that sum up to n . It is well known that partitions index the equivalence classes of representations of S_n , and thereby give rise to a decomposition of real-valued functions over S_n .

Given a partition λ , a λ -tabloid is a partition of $[n]$ into sets $A_1, \dots, A_k$ such that $|A_i| = \lambda_i$. Thus, for λ -tabloids $A = (A_1, \dots, A_k)$ and $B = (B_1, \dots, B_k)$, we define

$$T_{A,B} = \{ \pi \in S_n \mid \pi(A_i) = B_i \ \forall i = 1, \dots, k \},$$

and refer to any such $T_{A,B}$ as a λ -coset.

Definition 2.3. For a partition λ of n , we define the space $V_\lambda(S_n)$ as the linear span of indicator functions of all λ -cosets.

It is easily seen that if $\lambda_1 = n - k$, then $V_\lambda(S_n) \subseteq V_k(S_n)$, thus the $V_\lambda(S_n)$ are a refinement of the spaces $V_d(S_n)$. Next, we define a natural ordering on partitions which will allow us to further refine the space $V_\lambda(S_n)$, so that they will be refinements of the pure-degree spaces $V_{=d}(S_n)$.

Definition 2.4. Let $\lambda = (\lambda_1, \dots, \lambda_k)$, $\mu = (\mu_1, \dots, \mu_s)$ be partitions of $[n]$. We say that λ dominates μ , and denote $\lambda \succeq \mu$, if for all $j = 1, \dots, k$ it holds that $\sum_{i=1}^j \lambda_i \geq \sum_{i=1}^j \mu_i$.

Definition 2.5. For a partition λ of n , we define $V_{=\lambda}(S_n) = V_\lambda(S_n) \cap \bigcap_{\mu \succ \lambda} V_\mu(S_n)^\perp$.

It is well known that the spaces $V_{=\lambda}(S_n)$ are orthogonal, and that this decomposition is a refinement of the degree decomposition, namely that $V_{=\lambda}(S_n) \subseteq V_{=k}(S_n)$ where $\lambda_1 = n - k$.

An important parameter of a partition $\lambda \vdash n$ is its dimension, denoted by $\dim(\lambda)$, which is the dimension of the smallest irreducible sub-representation of a representation from the equivalence class of representations of λ . The dimension of λ will allow us to bound eigenvalues of certain operators using the trace method, and for that we first define the action of S_n on $f: S_n \rightarrow \mathbb{R}$. For a function $f: S_n \rightarrow \mathbb{R}$ and $\pi \in S_n$, define the function ${}^\pi f(\sigma) = f(\pi \circ \sigma)$; this way, we will think of the space of real-valued functions over S_n as a left module. With this action in hand, a situation that would often arise for us is the following: we are given an operator $T: L^2(S_n) \rightarrow L^2(S_n)$ which is positive semi-definite, and has a small trace; given an eigenvector $f \in V_{=\lambda}(S_n)$ of T , we would like to argue that the eigenvalue α_f of f is small. By the trace method, $\alpha_f \leq \text{Tr}(T)/m_{\alpha_f}$, where m_{α_f} is the multiplicity of α_f , and thus we would get a decent bound on α_f provided the multiplicity m_f is large. Here comes the importance of the parameter $\dim(\lambda)$: our operators of interest T will always commute with the action of S_n , and thus each ${}^\pi f$ is also an eigenvector of T with eigenvalue α_f . Thus, $m_{\alpha_f} \geq \dim(\text{span}(\{{}^\pi f \mid \pi \in S_n\}))$, and the latter quantity is well known to be at least $\dim(\lambda)$, as the following claim asserts.

Claim 2.6. Suppose $f \in V_{=\lambda}(S_n)$ is not identically 0. Then $\dim(\text{span}(\{{}^\pi f \mid \pi \in S_n\})) \geq \dim(\lambda)$.

Proof. Let $\rho: S_n \rightarrow V_{=\lambda}(S_n)$ be a representation, and let $W = \text{span}(\{{}^\pi f \mid \pi \in S_n\})$. Then since W is closed under the action of S_n , we have that W is a sub-representation. Since f is not identically 0, W is a non-trivial subspace, and therefore it follows that $\dim(W) \geq \dim(\lambda)$. $\square$

Thus, to get meaningful bounds on eigenvalues this way, it suffices to have effective bounds on the dimension of partitions, and for that we use the well-known Hook formula (or rather a corollary of it).

Lemma 2.7 (Claim 1, Theorem 19 in [16]). For all $c > 0$, there exists $\delta > 0$ such that the following holds. Let $\lambda \vdash n$ be given as $\lambda = (\lambda_1, \dots, \lambda_k)$, and denote $d = \min(n - \lambda_1, k)$.

1. If $\lambda = (n)$, then $\dim(\lambda) = 1$.
2. If $d > 0$, then $\dim(\lambda) \geq \left(\frac{n}{de}\right)^d$.
3. If $d > c \cdot n$, then $\dim(\lambda) \geq (1 + \delta)^n$.

2.2 The degree decomposition and the representation-theoretic decomposition over $\mathcal{U}_{\vec{k}}$

In this section, we refine the degree decomposition of functions on $\mathcal{U}_{\vec{k}}$ presented in the introduction using the representation-theoretic decomposition.

2.2.1 The maximal degree on the multi-slice

Recalling the spaces $V_0(\mathcal{U}_{\vec{k}}), \dots, V_n(\mathcal{U}_{\vec{k}})$ from the introduction, we note that it is clear that $V_n(\mathcal{U}_{\vec{k}})$ contains all real-valued functions on $\mathcal{U}_{\vec{k}}$. The following claim shows that something stronger happens, and one already has that V_{n-k_1} contains all real valued functions on $\mathcal{U}_{\vec{k}}$.

Claim 2.8. *For a multi-slice $\mathcal{U}_{\vec{k}}$, one has that $V_n = V_{n-k_1}$.*

Proof. Consider the set of functions

$$\mathcal{A} = \left\{ 1_{x_A=\alpha} \mid |A| = n - k_1, \alpha \in \{2, \dots, m\}^A, |\{i : \alpha_i = j\}| = k_j \forall j \in \{2, \dots, m\} \right\}. \quad (1)$$

We note that each one of the functions in this set is a $n - k_1$ -junta, so their span is contained in V_{n-k_1} . Thus, it is enough to prove that this set spans V_n . Indeed, let $B \subseteq [n]$, and $\beta \in [m]^n$ and partition $B = B' \cup B''$ where B' is the set of coordinates i of β for which $\beta_i = 1$, and B'' is the rest. We prove by induction on $|B'|$ that $1_{x_B=\beta}$ is in $\text{span}(\mathcal{A})$.

Base case. For $|B'| = 0$, we have $|B''| \leq n - k_1$ and the statement is clear. For $|B'| = 1$, we have

$$1_{x_B=\beta} = 1_{x_{B''}=\beta''} - \sum_{\gamma \in \{2, \dots, m\}} 1_{x_{B'}=\gamma, x_{B''}=\beta''},$$

we note that if $|B' \cup B''| > n - k_1$, then all of the functions in the sum on the right hand side are 0 and so $1_{x_B=\beta} = 1_{x_{B''}=\beta''} \in \mathcal{A}$. Otherwise, we clearly get $1_{x_B=\beta} \in \text{span}(\mathcal{A})$.

Inductive step. Write $B' = \{i\} \cup B'''$ and note that

$$1_{x_B=\beta} = 1_{x_{B''}=\beta'', x_{B'''}=\beta'''} - \sum_{\gamma \in \{2, \dots, m\}} 1_{x_i=\gamma, x_{B''}=\beta'', x_{B'''}=\beta''' }.$$

The statement now follows as in the base case using the inductive hypothesis. $\square$

Thus, letting $k = \max_i k_i$, we have that $V_\ell = \{0\}$ for all $\ell > n - k$, and thus the space of real-valued functions over $\mathcal{U}_{\vec{k}}$ can be written as $V_{=0}(\mathcal{U}_{\vec{k}}) \oplus V_{=1}(\mathcal{U}_{\vec{k}}) \oplus \dots \oplus V_{=n-k}(\mathcal{U}_{\vec{k}})$.

2.2.2 The representation-theoretic decomposition over $\mathcal{U}_{\vec{k}}$

In this section we lift the decomposition according to partitions from S_n to the multi-slice. The basic observation is that any function $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ may be identified with a function $\tilde{f}: S_n \rightarrow \mathbb{R}$ in the following way. Partition $[n]$ into m sets, $K_1, \dots, K_m$ where K_j has size k_j , and define $\tilde{f}(\pi) = f(x)$ where x is defined as $x_i = j$ for all $i \in \pi(K_j)$. It is easy to note that the mapping $f \mapsto \tilde{f}$ is linear, and sends d -juntas to d -juntas; it thus makes sense to define

$$V_\lambda(\mathcal{U}_{\vec{k}}) = \left\{ f \mid \tilde{f} \in V_\lambda(S_n) \right\}, \quad V_{=\lambda}(\mathcal{U}_{\vec{k}}) = \left\{ f \mid \tilde{f} \in V_{=\lambda}(S_n) \right\}.$$

This allows us to establish an analogue of Claim 2.6 for the multi-slice, and for that we first define the action of S_n on functions over the multi-slice as $(\pi, f) \rightarrow^\pi f$, where $^\pi f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ is given by $^\pi f(x) = f(\pi x)$. Here and throughout, for $x \in [m]^n$ and $\pi \in S_n$ we denote by $^\pi x$ the vector whose i th coordinate is $x_{\pi(i)}$.

Claim 2.9. *Suppose $f \in V_{=\lambda}(\mathcal{U}_{\vec{k}})$ is not identically 0. Then $\dim(\text{span}(\{^\pi f \mid \pi \in S_n\})) \geq \dim(\lambda)$.*

Proof. Let $W = \text{span}(\{^\pi f \mid \pi \in S_n\})$, and let $\tilde{W} = \left\{ \tilde{f} \mid f \in W \right\}$. Noting that $\tilde{W} = \left\{ ^\pi \tilde{f} \mid \pi \in S_n \right\}$, we have by Claim 2.6 that $\dim(\tilde{W}) \geq \dim(\lambda)$. As $\dim(\tilde{W}) \leq \dim(W)$ the claim follows. $\square$

2.2.3 An alternative description

Lastly, we need an alternative, equivalent description of the spaces $V_\lambda(\mathcal{U}_{\vec{k}})$ and $V_{=\lambda}(\mathcal{U}_{\vec{k}})$ in terms of cosets. For $\lambda \vdash n$, a λ -tabloid is a partition of $[n]$ into $A = (A_1, \dots, A_r)$, and a partition of the multiset $\{1, \dots, 1, \dots, m, \dots, m\}$ (where j appears k_j times) into $B = (B_1, \dots, B_r)$, where $|A_i| = |B_i| = \lambda_i$. Given a λ -tabloid (A, B) , we define the corresponding coset as

$$\tilde{T}_{A,B} = \left\{ x \in \mathcal{U}_{\vec{k}} \mid \{x_\ell \mid \ell \in A_i\} = B_i \quad \forall i = 1, \dots, r \right\}.$$

Following the map $f \mapsto \tilde{f}$ directly, one sees that it sends λ -tabloids over $\mathcal{U}_{\vec{k}}$ into λ -tabloids over S_n , and so $V_\lambda(\mathcal{U}_{\vec{k}})$ is the span of $1_{\tilde{T}_{A,B}}$ for all λ -tabloids (A, B) . It also follows that

$$V_{=\lambda}(\mathcal{U}_{\vec{k}}) = V_\lambda(\mathcal{U}_{\vec{k}}) \cap \bigcap_{\mu \triangleright \lambda} V_\mu(\mathcal{U}_{\vec{k}})^\perp.$$

We will sometimes want to show that the space V_λ is an invariant space for some operators, and for that it will be useful to further decompose V_λ . Towards this end, for a partition $A = (A_1, \dots, A_r)$ of $[n]$ into r sets where $|A_i| = \lambda_i$, we define

$$V_A(\mathcal{U}_{\vec{k}}) = \left\{ f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R} \mid ^\pi f = f \quad \forall \pi \in S_n \text{ such that } \pi(A_1) = A_1, \dots, \pi(A_r) = A_r \right\}.$$

Claim 2.10. *For all partitions $\lambda = (\lambda_1, \dots, \lambda_r)$ of n we have*

$$V_\lambda(\mathcal{U}_{\vec{k}}) = \sum_{\substack{A=(A_1, \dots, A_r) \\ |A_i|=\lambda_i}} V_A(\mathcal{U}_{\vec{k}}).$$

Proof. The direction $\subseteq$ is trivial: it is enough to observe that $1_{\tilde{T}_{A,B}}$ is in $V_A(\mathcal{U}_{\vec{k}})$ for all λ -cosets $T_{A,B}$, which is clear.

For the other direction, fix $A = (A_1, \dots, A_r)$, let $f \in V_A(\mathcal{U}_{\vec{k}})$, and fix a sequence of multi-sets $B_1, \dots, B_r \subseteq [m]$ for which $|B_i| = \lambda_i$. We claim that f is constant on the coset $\tilde{T}_{A,B}$. Indeed, let $x, y \in \tilde{T}_{A,B}$ then for all j , the multi-sets $\{x_i \mid i \in A_j\}$ and $\{y_i \mid i \in A_j\}$ are equal to B_j and hence the same, so we may find a permutation π_j that maps A_j to itself and has $[n] \setminus A_j$ as fixed point, such that $y_{\pi_j(i)} = x_i$ for all $i \in A_j$. Thus, letting $\pi = \pi_1 \circ \pi_2 \circ \dots \circ \pi_m$ we have that $x = \pi(y)$ and $\pi(A_j) = A_j$ for all j , and so $f(x) = f(\pi(y)) = ^\pi f(y) = f(y)$. Thus, we get that f is constant on any coset $\tilde{T}_{A,B}$ for any multi-sets $B_1, \dots, B_r \subseteq [m]$ such that $|B_i| = \lambda_i$, and we denote this value by $f_{A,B}$. Since these co-sets are disjoint, we get that

$$f = \sum_B f_{A,B} 1_{\tilde{T}_{A,B}} \in V_\lambda(\mathcal{U}_{\vec{k}}). \quad \square$$

2.3 Hypercontractivity

We will need the following variant of the hypercontractive inequality over general product space. For that, given a probability measure $\mathcal{D}$ on a finite space Ω , we define expectation L^q norms as $\|f\|_{q,\mathcal{D}} = (\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [|f(\mathbf{x})|^q])^{1/q}$. We often drop the $\mathcal{D}$ subscript from the notation whenever it is clear from context.

Theorem 2.11. *[[43, Theorem 10.21]] Let $(\Omega = \Omega_1 \times \dots \times \Omega_n, \nu = \nu_1 \times \dots \times \nu_n)$ be a product space. Suppose $\alpha \in (0, 1/2)$ is such that $\nu_i(\omega) \geq \alpha$ for all $i \in [n]$ and $\omega \in \Omega_i$. Then for all $d \in \mathbb{N}$, $p \geq 2$ and $f: \Omega \rightarrow \mathbb{R}$ of degree at most d we have $\|f\|_q \leq \left(\frac{10q}{\alpha}\right)^d \|f\|_2$.*

We will also need the following variant of the hypercontractive inequality on the multi-slice from [21]. Again, we use here expectation norms according to the uniform measure on $\mathcal{U}_{\vec{k}}$.

Theorem 2.12. *For all $c > 0$, $d \in \mathbb{N}$, and $q \in \mathbb{N}$ there is $N = N(c, d, q) > 0$ and $C = C(c, d, q) > 0$ such that the following holds. Let $n \geq N$, $\mathcal{U}_{\vec{k}} \subseteq [m]^n$ be a c -balanced multi-slice and let $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ be a function of degree at most d . Then $\|f\|_q \leq C \|f\|_2$.*

2.4 Tail bounds

We need the following standard tail bound.

Theorem 2.13. *Suppose $\mathbf{Z}_1, \dots, \mathbf{Z}_n$ are independent random bits with the same expectation p . Then*

$$\Pr \left[\left| \sum_{i=1}^n \mathbf{Z}_i - pn \right| \geq \varepsilon n \right] \leq 2e^{-2\varepsilon^2 n}.$$

We also need a tail bound for negatively associated random variables, introduced in [32].

Definition 2.14. *We say a collection of random variables $\mathbf{Z}_1, \dots, \mathbf{Z}_n$ over $\mathbb{R}$ is negatively associated if for each pair of disjoint sets $I, J \subset [n]$ and any two increasing functions $f_I: \mathbb{R}^I \rightarrow \mathbb{R}$, $f_J: \mathbb{R}^J \rightarrow \mathbb{R}$ it holds that the random variables $f_I((\mathbf{Z}_i)_{i \in I})$ and $f_J((\mathbf{Z}_j)_{j \in J})$ have non-positive covariance.*

Remark 2.15. *An important example of negatively associated random that we will use proceeds as follows. Suppose an urn contains n balls of different colors $i = 1, \dots, n$, and we pick $k < n$ balls from it randomly without replacement. Then, letting $\mathbf{Z}_i$ be the random variable indicating that the ball of color i has been chosen, we have that the collection $\mathbf{Z}_i$ for $i = 1, \dots, n$ is negatively associated [32].*

We will make use of this fact throughout the paper as follows. Suppose $\mathcal{U}_{\vec{k}}$ is a multi-slice over $[m]$. We will be concerned with distributions $(x, y) \in \mathcal{U}_{\vec{k}} \times \mathcal{U}_{\vec{k}}$ which is uniform over all pairs such that the number of coordinates i such that $(x_i, y_i) = (a, b)$ is some fixed number $n_{a,b}$ for all $a, b \in [m]$. It is thus clear that fixing a, b , defining for each $i \in [n]$ the random variable $\mathbf{Z}_i$ indicating that $(x_i, y_i) = (a, b)$, the collection $\mathbf{Z}_i$ is of the above form, and hence is negatively associated.

The following result is [15, Proposition 5].

Theorem 2.16. *Suppose $\mathbf{Z}_1, \dots, \mathbf{Z}_n$ are negatively correlated random bits with the same expectation p . Then*

$$\Pr \left[\sum_{i=1}^n \mathbf{Z}_i \geq (p + \varepsilon)n \right] \leq e^{-2\varepsilon^2 n}.$$

3 A basic invariance principle: the bilinear case

Definition 3.1. We say a distribution μ over $([m] \times [m'])^n$ commutes with the action of S_n if the following distributions are the same for all $\pi \in S_n$ and $x \in [m]^n$: (a) $\mathbf{x}'$, where $(\mathbf{x}, \mathbf{x}') \sim \mu$ conditioned on $\mathbf{x} = \pi(x)$, (b) $\pi(\mathbf{x}')$, where $(\mathbf{x}, \mathbf{x}') \sim \mu$ conditioned on $\mathbf{x} = x$.

Claim 3.2. Suppose a distribution μ over $[m]^n \times [m']^n$ is symmetric under the action of S_n . Then μ commutes with the action of S_n .

Proof. Fix $x, y \in [m]^n$ and $\pi \in S_n$. Then

$$\Pr_{(\mathbf{x}, \mathbf{y})} [\mathbf{y} = y \mid \mathbf{x} = \pi(x)] = \frac{\mu(\pi(x), y)}{\Pr_{(\mathbf{x}, \mathbf{y}) \sim \mu} [\mathbf{x} = \pi(x)]}, \quad \Pr_{(\mathbf{x}, \mathbf{y})} [\pi(\mathbf{y}) = y \mid \mathbf{x} = x] = \frac{\mu(x, \pi^{-1}(y))}{\Pr_{(\mathbf{x}, \mathbf{y}) \sim \mu} [\mathbf{x} = x]}.$$

By the symmetry of μ , we have that

$$\mu(\pi(x), y) = \Pr_{(\mathbf{x}, \mathbf{y}) \sim \mu} [\mathbf{x} = \pi(x), \mathbf{y} = y] = \Pr_{(\mathbf{x}, \mathbf{y}) \sim \mu} [\pi^{-1}(\mathbf{x}) = x, \pi^{-1}(\mathbf{y}) = \pi^{-1}(y)] = \mu(x, \pi^{-1}(y)),$$

which shows that the numerators of the two fractions are equal; summing this over y also shows that the denominators of the two fractions are equal. $\square$

3.1 Couplings between measures on $[m]^n$

Consider the multi-slice $\mathcal{U}_k$ over alphabet $[m]$, and consider the distribution over $[m]$ defined by $\nu_k(i) = k_i/n$. We will show that in several senses, the multi-slice $\mathcal{U}_k$ behaves very similarly to the product domain $([m]^n, \nu_k^{\otimes n})$, and for that we will use couplings, which we defined earlier in the introduction and recall for convenience below.

Definition 3.3. Let $\alpha, \zeta > 0$, $m \in \mathbb{N}$. For symmetric distributions ν_1, ν_2 over $[m]^n$, a (α, ζ) -coupling between $([m]^n, \nu_1)$ and $([m]^n, \nu_2)$ is a jointly distributed $\mathcal{C} = (\mathbf{x}, \mathbf{y})$ satisfying the following properties.

1. The marginal distribution of $\mathbf{x}$ is ν_1 , and the marginal distribution of $\mathbf{y}$ is ν_2 .
2. The distribution of $\mathcal{C}$ is symmetric under the action of S_n .
3. For all $i \in [n]$, $\Pr_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}} [\mathbf{x}_i \neq \mathbf{y}_i] \leq \zeta$.
4. Tail bounds: for all $\varepsilon > 0$, it holds that $\Pr_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}} [|\{i \mid x_i \neq y_i\}| \geq \varepsilon n] \leq \frac{1}{\alpha} e^{-\alpha \varepsilon^2 n}$.

We also recall that given a coupling $\mathcal{C}$ between $(\mathcal{U}_k, \text{Uniform})$ and $([m]^n, \nu_k^{\otimes n})$, we may define lifting operators on functions between the two spaces, namely $T_{\mathcal{C}}: L^2(\mathcal{U}_k) \rightarrow L^2([m]^n, \nu_k^{\otimes n})$ and its adjoint operator $T_{\mathcal{C}}^*: L^2([m]^n, \nu_k^{\otimes n}) \rightarrow L^2(\mathcal{U}_k)$, as:

$$T_{\mathcal{C}}f(y) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}} [f(\mathbf{x}) \mid \mathbf{y} = y], \quad T_{\mathcal{C}}^*g(x) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}} [f(\mathbf{y}) \mid \mathbf{x} = x].$$

3.2 Invariance for low-degree functions: Proof of Lemma 1.6

A key property of couplings is that at least for low-degree functions $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$, we have that $f(\mathbf{x}) \approx \text{T}_{\mathcal{C}} f(\mathbf{y})$ where $(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}$. This is essentially the content of Lemma 1.6, restated below. We remark that due to technical reasons, we will use two different couplings $\mathcal{C}$ and $\mathcal{C}'$. Here $\mathcal{C}$, will be used to lift functions from the multi-slice to the product space, and $\mathcal{C}'$ will be used to sample the inputs $(\mathbf{x}, \mathbf{y})$.

Lemma 1.6 (Restated) . *Let $d, m, n \in \mathbb{N}$, $\alpha, \zeta > 0$, and let $\mathcal{U}_{\vec{k}}$ be multi-slice over alphabet $[m]$. Suppose $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ is a function of degree at most d , and $\mathcal{C}, \mathcal{C}'$ are (α, ζ) -coupling between $(\mathcal{U}_{\vec{k}})$ and $([m]^n, \nu_{\vec{k}}^{\otimes n})$. Then*

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}'} [(f(\mathbf{x}) - \text{T}_{\mathcal{C}} f(\mathbf{y}))^2] \leq 8\sqrt{d\zeta} \|f\|_2^2.$$

The rest of this section is devoted to the proof of Lemma 1.6.

3.2.1 Proof overview

It will be more convenient for us to consider operators that map a space to itself, and more specifically we will consider $\mathcal{S}: L^2(\mathcal{U}_{\vec{k}}) \rightarrow L^2(\mathcal{U}_{\vec{k}})$ defined as $\mathcal{S} = \text{T}_{\mathcal{C}'}^* \text{T}_{\mathcal{C}}$. With these definitions, expanding the difference in Lemma 1.6 we see that

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}'} [(f(\mathbf{x}) - \text{T}_{\mathcal{C}} f(\mathbf{y}))^2] = \|f\|_2^2 + \|\text{T}_{\mathcal{C}} f\|_2^2 - 2\langle f, \mathcal{S}f \rangle \leq 2(\|f\|_2^2 - \langle f, \mathcal{S}f \rangle) = 2\langle f, (\text{I} - \mathcal{S})f \rangle. \quad (2)$$

We would like to reason about the eigenvalues of $\text{I} - \mathcal{S}$ and say they are close to 0 for low-degree functions, however this may be problematic since $\mathcal{S}$ may not be diagonalizable. We thus use a standard trick by Cauchy-Schwarz to bound

$$\begin{aligned} (2) &\leq 2\|f\|_2 \|\text{I} - \mathcal{S}\|_2 = 2\|f\|_2 \sqrt{\langle (\text{I} - \mathcal{S})f, (\text{I} - \mathcal{S})f \rangle} = 2\|f\|_2 \sqrt{\langle f, (\text{I} - \mathcal{S})^*(\text{I} - \mathcal{S})f \rangle} \\ &\leq 2\|f\|_2^{3/2} \|(\text{I} - \mathcal{S})^*(\text{I} - \mathcal{S})f\|_2^{1/2}. \end{aligned} \quad (3)$$

Thus, we reduced our question to understanding the operator $\mathcal{S}^* = (\text{I} - \mathcal{S})^*(\text{I} - \mathcal{S})$, for which we may study eigenvalues since it is self-adjoint. Indeed, we show that it strongly contracts low degree functions, and for that we follow the following plan:

1. We argue that $\mathcal{S}$ and its adjoint preserves juntas. Therefore $\mathcal{S}^*$ preserves juntas.
2. We deduce that for each eigenvalue θ of $\mathcal{S}^*$ corresponding to a low-degree eigenfunction f , there is a junta g which is also an eigenfunction of $\mathcal{S}^*$ with eigenvalue θ .
3. Finally, we note that the coupling operator is unlikely to touch a small set of coordinates (given that n is sufficiently large in comparison to d), so we must have that any eigenvalue θ corresponding to a junta must be close to 0.
4. Combining the above two items with Parseval's equality establishes Lemma 1.6.

We begin with the formal proof of Lemma 1.6 following the above outline.

3.2.2 $\mathcal{S}$ preserves juntas

For a set J , we denote by $V_J(\mathcal{U}_{\vec{k}})$ the set of J -juntas. Note that we may write $V_{\leq d}(\mathcal{U}_{\vec{k}})$ as the direct sum of all $V_J(\mathcal{U}_{\vec{k}})$. Recall the action of S_n on functions over the multi-slice is defined as: for $\pi \in S_n$, $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ define ${}^\pi f = f(\pi(x))$. To prove that $\mathcal{S}$ preserves juntas, it will be more convenient for us to use the following equivalent criteria.

Fact 3.4. *A function $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ is a J -junta if and only if ${}^\pi f = f$ for all π that fix J .*

Proof. The $\Rightarrow$ direction is clear, and we prove the $\Leftarrow$ direction. Suppose ${}^\pi f = f$ for all π that fix J , and define $g: [m]^J \rightarrow \mathbb{R}$ as $g(\alpha) = \mathbb{E}_{\substack{x \in \mathcal{U}_{\vec{k}} \\ x_J = \alpha}} [f(x)]$; it suffices to argue that $f(x) = g(x_J)$ for all x . Indeed, fix x , and denote α_J . If $x' \in \mathcal{U}_{\vec{k}}$ is a point such that $x'_J = \alpha$, then we have that for all $\ell = 1, \dots, m$, the sets

$$X_\ell = \{i \in [n] \setminus J \mid x_i = \ell\}, \quad X'_\ell = \{i \in [n] \setminus J \mid x'_i = \ell\}$$

have the same size, and so there is $\pi \in S_{[n] \setminus J}$ that maps each X_ℓ to X'_ℓ . Thus, there is π fixing J such that $x' = \pi(x)$, and so $f(x') = {}^\pi f(x) = f(x)$. $\square$

Claim 3.5. *For all $\pi \in S_n$ and $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$, we have that $\mathcal{S}({}^\pi f) = {}^\pi (\mathcal{S}f)$, as well as $\mathcal{S}^*({}^\pi f) = {}^\pi (\mathcal{S}^* f)$.*

Proof. We prove the first assertion – the proof of the second assertion is analogous. By Claim 3.2, we get that $\mathcal{C}, \mathcal{C}'$ commute with the action of S_n . Fix $x \in \mathcal{U}_{\vec{k}}$ and sample $\mathbf{y}$ as $(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}$ conditioned on $\mathbf{x} = x$, and then $(\mathbf{x}', \mathbf{y}') \sim \mathcal{C}'$ conditioned on $\mathbf{y}' = \mathbf{y}$. Then by definition of $\mathcal{S}$ and the definition on the adjoint operator we have

$$\mathcal{S}({}^\pi f)(x) = \mathbb{E}_{\mathbf{x}' \sim \mathcal{S}x} [{}^\pi f(\mathbf{x}')] = \mathbb{E}_{\mathbf{x}' \sim \mathcal{S}x} [f(\pi(\mathbf{x}'))].$$

Consider the points $x'' = \pi(x)$, $\mathbf{y}'' = \pi(\mathbf{y}')$. Since $\mathcal{C}$ commutes with S_n , the distribution of $\mathbf{y}''$ is the same as of $\mathbf{b}$ where $(\mathbf{a}, \mathbf{b}) \sim \mathcal{C}$ conditioned on $\mathbf{a} = x''$; also, the distribution of $\pi(\mathbf{x}')$ the same as of $\mathbf{a}'$ where $(\mathbf{a}', \mathbf{b}') \sim \mathcal{C}'$ conditioned on $\mathbf{b}' = \mathbf{b}$. Together, these two facts imply that the distribution of $\pi(\mathbf{x}')$ is $\mathcal{S}(\pi(x))$, and the above average is equal to $\mathcal{S}f(\pi(x)) = {}^\pi (\mathcal{S}f)(x)$, as required. $\square$

Claim 3.6. *Suppose $\mathcal{R}$ is an operator that commutes with the action of S_n on functions over the multi-slice. Then for each $J \subseteq [n]$, we have that $\mathcal{R}(V_J(\mathcal{U}_{\vec{k}})) \subseteq V_J(\mathcal{U}_{\vec{k}})$. Subsequently, $\mathcal{R}$ preserves $V_d(\mathcal{U}_{\vec{k}})$.*

Proof. Let $\pi \in S_n$ be a permutation that fixes J , and let f be a J -junta. Then, by Claim 3.5 we have $\pi(\mathcal{S}f) = \mathcal{S}({}^\pi f) = \mathcal{S}(f)$, where the last transition is since f is a J -junta and π fixes J . Since this holds for all $\pi \in S_n$ that fixes J , we get that $\mathcal{S}(f)$ is a J -junta.

The “subsequently” part follows as $V_d(\mathcal{U}_{\vec{k}}) = \sum_{|J| \leq d} V_J(\mathcal{U}_{\vec{k}})$. $\square$

From Claims 3.5, 3.6 it follows that $V_J(\mathcal{U}_{\vec{k}})$ is invariant under $\mathcal{S}^*$, and since $\mathcal{S}^*$ is self-adjoint we may decompose this space as sum of eigenspaces. Let us denote by $V_J^\theta(\mathcal{U}_{\vec{k}}) \subseteq V_J(\mathcal{U}_{\vec{k}})$ the subspace of eigenvectors with eigenvalue θ , and write $V_{\leq d}(\theta) = \bigoplus_{|J| \leq d} V_J^\theta(\mathcal{U}_{\vec{k}})$.

Claim 3.7. $V_{\leq d}(\mathcal{U}_{\vec{k}}) = \bigoplus_{\theta} V_{\leq d}(\theta)$.

Proof. To see that the space on the left hand side is the sum of the spaces on the right hand side, note that

$$\sum_{\theta} V_{\leq d}(\theta) = \sum_{\theta} \sum_{|J| \leq d} V_J^\theta(\mathcal{U}_{\vec{k}}) = \sum_{|J| \leq d} \sum_{\theta} V_J^\theta(\mathcal{U}_{\vec{k}}) = \sum_{|J| \leq d} V_J(\mathcal{U}_{\vec{k}}) = V_{\leq d}(\mathcal{U}_{\vec{k}}).$$

The statement now follows as the spaces $V_{\leq d}(\theta)$ are orthogonal for different θ 's. $\square$

3.2.3 Bounding the eigenvalues of $\mathcal{S}^*$

Next, we will estimate the eigenvalues of $\mathcal{S}^*$ corresponding to low-degree functions. We begin by the following easy claim.

Claim 3.8. *Let $\mathcal{U}_k^-$ be a c -balanced multi-slice over alphabet $[m]$, let $J \subseteq [n]$ be of size at most d , and let $x \in \mathcal{U}$. Then*

$$\Pr_{\mathbf{x}' \sim \mathcal{S}x} [\mathbf{x}'_J \neq x_J] \leq 2d\zeta$$

Proof. The claim is immediate by the third property in Definition 1.5 and the union bound. $\square$

Claim 3.9. *Let θ be such that $V_{\leq d}(\theta) \neq \{0\}$. Then $\theta \leq 16d\zeta$.*

Proof. By assumption, $V_{\leq d}(\theta) \neq \{0\}$ so by definition there is J of size at most d such that $V_J^\theta(\mathcal{U}_k^-) \neq \{0\}$, and we take $g \in V_J^\theta(\mathcal{U}_k^-)$. Let $x \in \mathcal{U}_k^-$ be a point achieving the maximum of $|g(x)|$ (note that it is necessarily positive); we may assume without loss of generality that $g(x) > 0$, otherwise we work with the function $-g$. Let $\mathbf{x}' \sim \mathcal{S}x$, then

$$\begin{aligned} \mathcal{S}g(x) &= \mathbb{E}_{\mathbf{x}'} [g(\mathbf{x}')] \geq \Pr_{\mathbf{x}'} [\mathbf{x}'_J = x_J] \mathbb{E}_{\mathbf{x}'} [g(\mathbf{x}') \mid \mathbf{x}'_J = x_J] - (1 - \Pr_{\mathbf{x}'} [\mathbf{x}'_J = x_J])g(x) \\ &= (2\Pr_{\mathbf{x}'} [\mathbf{x}'_J = x_J] - 1)g(x), \end{aligned}$$

which by Claim 3.8 is at least $(1 - 4d\zeta)g(x)$. Since clearly $\mathcal{S}g(x) \leq g(x)$, it follows that $|\mathcal{S}g(x) - g(x)| \leq 4d\zeta$. A similar argument shows that $|\mathcal{S}^*g(x) - g(x)| \leq 4d\zeta g(x)$, and also that $|\mathcal{S}^*\mathcal{S}g(x) - g(x)| \leq 8d\zeta g(x)$. We thus get that

$$\theta g(x) = |\mathcal{S}^*g(x)| = |g(x) - \mathcal{S}g(x) - \mathcal{S}^*g(x) + \mathcal{S}^*\mathcal{S}g(x)| \leq 16d\zeta g(x),$$

and dividing by $g(x)$ finishes the proof. $\square$

3.2.4 Proof of Lemma 1.6

Let $f \in V_{\leq d}(\mathcal{U}_k^-)$. From (2) and (3) we get that

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}'} [(f(\mathbf{x}) - \text{T}_C f(\mathbf{y}))^2] \leq 2 \|f\|_2^{3/2} \|\mathcal{S}^* f\|_2^{1/2},$$

and we upper bound the last norm. By Claim 3.7 we may write $f = \sum_{\theta} f^{\theta}$ where $f^{\theta} \in V_{\leq d}(\theta)$, and so

$$\|\mathcal{S}^* f\|_2^2 = \sum_{\theta} \|\mathcal{S}^* f^{\theta}\|_2^2 = \sum_{\theta} \theta^2 \|f^{\theta}\|_2^2 \leq (16d\zeta)^2 \sum_{\theta} \|f^{\theta}\|_2^2 = (16d\zeta)^2 \|f\|_2^2$$

where we used Claim 3.9. Plugging this into the inequality above finishes the proof.

3.3 Invariance for admissible distributions: the bilinear case

Next, we state and prove an invariance principle for admissible distributions as in Definition 1.8. Let μ be an admissible distribution on $\mathcal{U}_{\vec{k}} \times \mathcal{U}_{\vec{k}'}$ where the multi-slices are over alphabets m, m' respectively, and let $\tilde{\mu}$ be the product version of μ as per Definition 1.10. In this section, we will prove Theorem 1.11 for the case that $r = 2$, which captures the main ideas of the argument, and in the subsequent section we generalize it for all $r \in \mathbb{N}$. Throughout, we will denote samples of μ by $(\mathbf{x}, \mathbf{x}') \sim \mu$, and samples of $\tilde{\mu}$ by $(\mathbf{y}, \mathbf{y}') \sim \tilde{\mu}$. We will be interested in couplings between the distribution μ and $\tilde{\mu}$, as in Definition 1.5 – namely, a jointly distributed $\mathcal{C} = ((x, x'), (y, y'))$ satisfying the properties in Definition 1.5.

Theorem 3.10. *For all $\alpha \in (0, 1)$, $m, m' \in \mathbb{N}$, $\varepsilon > 0$ there is $\zeta > 0$ such that the following holds. Suppose $\mathcal{U}_{\vec{k}}$ and $\mathcal{U}_{\vec{k}'}$ are α -balanced multi-slices over alphabets $[m]$ and $[m']$ respectively, and μ is a connected, α -admissible distribution over $\mathcal{U}_{\vec{k}} \times \mathcal{U}_{\vec{k}'}$. Suppose $\mathcal{C}$ is a (α, ζ) -coupling between $\mathcal{U}_{\vec{k}}$ and $\nu_{\vec{k}}^{\otimes n}$, $\mathcal{C}'$ is a (α, ζ) -coupling between $\mathcal{U}_{\vec{k}'}$ and $\nu_{\vec{k}'}^{\otimes n}$, and that there is a (α, ζ) -coupling $\mathcal{C}''$ between μ and $\tilde{\mu}$.*

Then for all $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ and $g: \mathcal{U}_{\vec{k}'} \rightarrow \mathbb{R}$ it holds that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \mu} [f(\mathbf{x})g(\mathbf{x}')] - \mathbb{E}_{(\mathbf{y}, \mathbf{y}') \sim \tilde{\mu}} [\mathsf{T}_{\mathcal{C}} f(\mathbf{y}) \mathsf{T}_{\mathcal{C}'} g(\mathbf{y}')] \right| \leq \varepsilon \|f\|_2 \|g\|_2.$$

The rest of this section is devoted to the proof of Theorem 3.10, and we begin by outlining the idea of the proof.

3.3.1 Proof overview

We begin by noting that Lemma 1.6 implies Theorem 3.10 for low degree functions. Indeed, by adding and subtracting $\mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \mathcal{C}''} [f(\mathbf{x}) \mathsf{T}_{\mathcal{C}'} g(\mathbf{y}')]$, the left hand side in Theorem 3.10 may be upper bounded by

$$\begin{aligned} & \left| \mathbb{E}_{\substack{(\mathbf{x}, \mathbf{x}') \\ (\mathbf{y}, \mathbf{y}') \sim \mathcal{C}''}} [f(\mathbf{x})(g(\mathbf{x}') - \mathsf{T}_{\mathcal{C}'} g(\mathbf{y}'))] \right| + \left| \mathbb{E}_{\substack{(\mathbf{x}, \mathbf{x}') \\ (\mathbf{y}, \mathbf{y}') \sim \mathcal{C}''}} [(f(\mathbf{x}) - \mathsf{T}_{\mathcal{C}} f(\mathbf{y})) \mathsf{T}_{\mathcal{C}'} g(\mathbf{y}')] \right| \\ & \leq \|f\|_2 \sqrt{\mathbb{E}_{\substack{(\mathbf{x}, \mathbf{x}') \\ (\mathbf{y}, \mathbf{y}') \sim \mathcal{C}''}} [(g(\mathbf{x}') - \mathsf{T}_{\mathcal{C}'} g(\mathbf{y}'))^2]} + \|\mathsf{T}_{\mathcal{C}'} g\|_2 \sqrt{\mathbb{E}_{\substack{(\mathbf{x}, \mathbf{x}') \\ (\mathbf{y}, \mathbf{y}') \sim \mathcal{C}''}} [(f(\mathbf{x}) - \mathsf{T}_{\mathcal{C}} f(\mathbf{y}))^2]}, \end{aligned}$$

where we used Cauchy-Schwarz. By Jensen's inequality we have that $\|\mathsf{T}_{\mathcal{C}'} g\|_2 \leq \|g\|_2$, and by Lemma 1.6 the expectations are at most $8\sqrt{d\zeta} \|g\|_2^2$ and $8\sqrt{d\zeta} \|f\|_2^2$ respectively, yielding the bound $8(d\zeta)^{1/4} \|f\|_2 \|g\|_2$.

This observation suggests that we should consider the contribution from the high-degree parts and the low-degree parts in Theorem 3.10 separately, and show that the contribution from the high-degree parts is small. More precisely, we will decompose f and g into their high degree and low-degree parts as $f = f^{\leq d} + f^{> d}$ and $g = g^{\leq d} + g^{> d}$, and analyze the terms corresponding to $f^{\leq d} g^{\leq d}$, $f^{> d} g^{\leq d}$ and $f g^{> d}$ separately. We show that only the contribution from the first term is meaningful, i.e. of $f^{\leq d} g^{\leq d}$, which by Lemma 1.6 is very close to $\mathbb{E} [\mathsf{T}_{\mathcal{C}}(f^{\leq d})(\mathbf{y}) \mathsf{T}_{\mathcal{C}'}(g^{\leq d})(\mathbf{y}')]$.

To complete the proof, we have to show that the last quantity is very close to $\mathbb{E} [\mathsf{T}_{\mathcal{C}} f(\mathbf{y}) \mathsf{T}_{\mathcal{C}'} g(\mathbf{y})]$. The first step in this argument is to observe yet another important property of $\mathsf{T}_{\mathcal{C}}$, $\mathsf{T}_{\mathcal{C}'}$, namely that they nearly commute with degree truncations. This comes in the form of Lemma 3.25, asserting that $\mathsf{T}_{\mathcal{C}}(f^{\leq d})$ and

$(T_C f)^{\leq d}$ are close in 2-norm, provided that $\zeta \ll d^{-7}$ (and similarly for $T_C(g^{\leq d})$ is close to $(T_{C'} g)^{\leq d}$). Thus, $\mathbb{E} [T_C(f^{\leq d})(\mathbf{y}) T_{C'}(g^{\leq d})(\mathbf{y}')]]$ is close to $\mathbb{E} [(T_C f)^{\leq d}(\mathbf{y}) (T_{C'} g)^{\leq d}(\mathbf{y}')]]$, and it remains to show that the contribution from the high-degree parts is small. That is, we must show that the expectation of each one of the terms $T_C(f)(T_{C'} g)^{> d}$ and $(T_C f)^{> d}(T_C g)^{\leq d}$ is small in magnitude. Such bounds have been established in previous works, e.g. in [40], and we reproduce the argument in the appendix for completeness.

Finally, we elaborate on how we show that the expectations corresponding products involving a high-degree functions over the multi-slice are small. To do so, it will be convenient to us to view the distribution μ as an operator $T_\mu: L^2(\mathcal{U}_{\vec{k}}) \rightarrow L^2(\mathcal{U}_{\vec{k}'})$ defined as $T_\mu f(y) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mu} [f(\mathbf{x}) | \mathbf{y} = y]$. Thus, we are interested in studying expressions of the form $\langle T_\mu f^{> d}, g^{\leq d} \rangle$, which by Cauchy-Schwarz is at most $\|T_\mu f^{> d}\|_2 \|g^{\leq d}\|_2$. Again, it will be more convenient for us to work with operators from a space to itself, which we may do again by defining $S_\mu = T_\mu^* T_\mu$ and using Cauchy-Schwarz to bound

$$\|T_\mu f^{> d}\|_2^2 = \langle f^{> d}, S_\mu f^{> d} \rangle \leq \|f^{> d}\|_2 \|S_\mu f^{> d}\|_2.$$

Thus, it suffices to show that the eigenvalues of S_μ corresponding to high-degree functions are small. We prove the latter statement using the trace method: we show that on the one hand, the multiplicity of each one of them is exponentially large in n , say at least $(1 + \delta)^n$, and that the sum of some constant power of them is small. For the latter part, we consider a large enough (but constant) power h , and show that the operator S_μ^h mixes well enough so that its trace is at most $(1 + \xi(h))^n$, where $\lim_{h \rightarrow \infty} \xi(h) = 0$.

The above argument works for $d \geq \Omega(n)$, and does not immediately apply in the case that $d = o(n)$. The reason is that the upper bound we have on the trace, $(1 + o(1))^n$, is too weak in comparison to the multiplicity we expect an eigenvalue corresponding to degree d functions to have (which we expect to be $2^{\Theta(d)}$). Fortunately, there is an elegant fix. We show that we may assume the eigenvalue we are interested in bounding, say θ , corresponds to a d -junta (for the same reasons as in the proof of Lemma 1.6), and thus we may “restrict our view” to a set of $3d$ coordinates that contains the d -coordinates of the junta. Morally, this view allows us to assume that $n = 3d$, in which case the previous argument appears to be applicable. Due to technical reasons, however, we use a different and more direct argument once we have “focused our view” on the set of $3d$ relevant coordinates.

3.4 Strong contraction for high-degrees functions

Throughout this section, we analyze operators such as S_μ defined above, and prove that they contract high degree functions very strongly.

Lemma 3.11. *For all $\alpha > 0$, $m \in \mathbb{N}$ there exist $\delta > 0$ and $C > 0$ such that the following holds. Let $\mathcal{U}_{\vec{k}} \subseteq [m]^n$ be α -balanced, and μ be a α -admissible, connected distribution over $\mathcal{U}_{\vec{k}} \times \mathcal{U}_{\vec{k}'}$, and suppose there is a (α, ζ) -coupling $\mathcal{C}$ between μ and $\tilde{\mu}$.*

Then for all $d \in \mathbb{N}$, if $f \in V_{> d}(\mathcal{U}_{\vec{k}'})$, we have $\|T_\mu f\|_2 \leq C(1 + \delta)^{-d} \|f\|_2$.

The proof of Lemma 3.11 proceeds differently for d 's that are at least linear in n , and d 's that are sub-linear. We begin by presenting the argument for d 's that are at least γn , where $\gamma > 0$ is to be thought of as a small constant (depending on m, m' and α).

Remark 3.12. *We remark that we may, and will, assume that T_μ is self-adjoint (i.e. symmetric as a matrix). We can do that without loss of generality as one has, by a standard Cauchy-Schwarz argument as before, that $\|T_\mu f\|_2^2 \leq \|f\|_2^{3/2} \|T_\mu^* T_\mu f\|_2^{1/2}$. The operator $T_\mu^* T_\mu$ is self-adjoint and equal to $T_{\mu'}$ for a distribution μ' which is sampled by taking $(\mathbf{x}, \mathbf{y}) \sim \mu$, then $(\mathbf{x}', \mathbf{y}') \sim \mu$ conditioned on $\mathbf{y}' = \mathbf{y}$ and outputting $(\mathbf{x}, \mathbf{x}')$.*

Clearly, if μ is α -admissible, then μ' is α^2 -admissible, if μ is connected then μ' is connected, and a (α, ζ) -coupling between μ and $\tilde{\mu}$ naturally induces a $(\alpha/2, 2\zeta)$ -coupling between μ' and $\tilde{\mu}'$.

3.4.1 The case that $d \geq \gamma n$

Claim 3.13. *For all $m \in \mathbb{N}$ and $\alpha, \gamma > 0$, there are $\delta > 0$, $C > 0$ such that the following holds. Let $\mathcal{U}_{\vec{k}} \subseteq [m]^n$ be α -balanced, and μ be a α -admissible, connected distribution over $\mathcal{U}_{\vec{k}} \times \mathcal{U}_{\vec{k}}$, such that T_μ is self-adjoint.*

Then for all $d \geq \gamma n$, if $f \in V_{=d}(\mathcal{U}_{\vec{k}})$, then $\|T_\mu f\|_2 \leq C(1 + \delta)^{-d} \|f\|_2$.

In the rest of this section, we will fix $m \in \mathbb{N}$, $\alpha, \gamma > 0$ and μ as in the above claim. Consider the graph H , whose vertex set is $[m]$ and its edges are weighted as follows. For $a, b \in [m]$, the weight of the edge (a, b) is

$$w(a, b) = \Pr_{\substack{(\mathbf{x}, \mathbf{x}') \sim \mu \\ i \in [n]}} [\mathbf{x}'_i = b \mid \mathbf{x}_i = a].$$

We note that as μ is connected, it follows that H is connected, and from the third property in Definition 1.8 it follows that for each a, b , we either have $w(a, b) = 0$ or $w(a, b) \geq \alpha$. Therefore, by Claim B.1 it follows that H has constant spectral gap, i.e. $\lambda_2(H) \leq 1 - \Omega_\alpha(1)$. Note that the stationary distribution of H is precisely $\nu_{\vec{k}}$. It follows from a standard spectral argument that random walks of length $O_{\alpha, \varepsilon}(1)$ from any vertex reach distribution that is ε -close to $\nu_{\vec{k}}$:

Claim 3.14. *For all $\varepsilon > 0$, there is $h = h(m, \alpha, \varepsilon) \in \mathbb{N}$, such that for all $v_1 \in V(H)$,*

$$\left| \Pr_{\substack{v_2, \dots, v_h \\ \text{random walk on } H \text{ from } v_1}} [v_h = a] - \nu_{\vec{k}}(a) \right| \leq \varepsilon.$$

Let $\varepsilon = \varepsilon(m, \alpha) > 0$ to be determined later, and take $h = h(m, \alpha, \varepsilon)$ from Claim 3.14. We now consider T_μ^h , and want to argue it is almost mixing in the following sense.

Claim 3.15. *For all $a, b \in [m]$, it holds that*

$$\Pr_{\substack{\mathbf{x} \in \mathcal{U}_{\vec{k}} \\ \mathbf{x}' \sim T_\mu^h \mathbf{x}}} [|\#\{i \mid \mathbf{x}_i = a, \mathbf{x}'_i = b\} - n\nu_{\vec{k}}(a)\nu_{\vec{k}}(b)| \geq 3\varepsilon n] \leq 4\alpha^{-1}e^{-\alpha\varepsilon^2 n}.$$

Proof. Let $\mathbf{x}(0) \in_R \mathcal{U}_{\vec{k}}$, and sample $x(1)$ as $\mathbf{x}'$ where $(\mathbf{x}, \mathbf{x}') \sim \mu$ conditioned on $\mathbf{x} = \mathbf{x}(0)$. Inductively for $j \geq 0$, once $\mathbf{x}(j)$ have been defined, sample $\mathbf{x}(j+1)$ as $\mathbf{x}'$ where $(\mathbf{x}, \mathbf{x}') \sim \mu$ conditioned on $\mathbf{x} = \mathbf{x}(j)$. Then the distribution of $(\mathbf{x}(0), \mathbf{x}(h))$ is the same as $(\mathbf{x}, \mathbf{x}')$ in the statement of the claim, and we wish to study the number of coordinates on which it is equal to (a, b) . This will be much more convenient to do over product domains, and indeed we next use our coupling $\mathcal{C}$ to move to that scenario.

More precisely, let $\mathbf{y}(0), \mathbf{y}(1)$ be distributed as $(\mathbf{y}, \mathbf{y}')$ where $((\mathbf{x}, \mathbf{x}'), (\mathbf{y}, \mathbf{y}')) \sim \mathcal{C}$ conditioned on $(\mathbf{x}, \mathbf{x}') = (\mathbf{x}(0), \mathbf{x}(1))$. Then inductively, for all $j > 0$, once $\mathbf{y}(\ell-1)$ has been drawn, sample $\mathbf{y}(j)$ as $\mathbf{y}'$ where $((\mathbf{x}, \mathbf{x}'), (\mathbf{y}, \mathbf{y}')) \sim \mathcal{C}$ conditioned on $(\mathbf{x}, \mathbf{x}', \mathbf{y}) = (\mathbf{x}(j-1), \mathbf{x}(j), \mathbf{y}(j-1))$.

Consider the distribution over $(\mathbf{y}(0), \dots, \mathbf{y}(h))$. An equivalent way to sample it, is by first taking $\mathbf{y}(0) \sim \nu_{\vec{k}}^{\otimes n}$, then taking $\mathbf{y}(1)$ as $\mathbf{y}'$ where $(\mathbf{y}, \mathbf{y}') \sim \tilde{\mu}$ conditioned on $\mathbf{y} = \mathbf{y}(0)$, and continuing inductively.

This shows that the random variables $(\mathbf{y}(0)_i, \mathbf{y}(h)_i)$ are independent. Thus, letting $\mathbf{Z}_i$ be the indicator of the event that $(\mathbf{y}(0)_i, \mathbf{y}(h)_i) = (a, b)$, we get that $\mathbf{Z}_i$ are independent. Also, from Claim 3.14 we have that

$$\left| \mathbb{E}[\mathbf{Z}_i] - \nu_{\vec{k}}(a)\nu_{\vec{k}}(b) \right| \leq \varepsilon,$$

so by Chernoff bound (Theorem 2.13) we have

$$\Pr \left[\left| \sum_{i=1}^n \mathbf{Z}_i - n\nu_{\vec{k}}(a)\nu_{\vec{k}}(b) \right| \geq \varepsilon n \right] \leq 2e^{-2\varepsilon^2 n}.$$

Finally, we relate the statistics of $(\mathbf{y}(0), \mathbf{y}(h))$ to that of $(\mathbf{x}(0), \mathbf{x}(h))$. For each j , the distribution of $(\mathbf{x}(j-1), \mathbf{x}(j), \mathbf{y}(j-1), \mathbf{y}(j))$ is $\mathcal{C}$. Thus, using the tail bound property, we see that the probability that $\mathbf{x}(0)$ and $\mathbf{y}(0)$ differ in at least εn coordinates, is at most $\alpha^{-1}e^{-\alpha\varepsilon^2 n}$. Also, $\mathbf{x}(h), \mathbf{y}(h)$ differ in at least εn coordinates with probability at most $\alpha^{-1}e^{-\alpha\varepsilon^2 n}$. We thus get

$$\begin{aligned} & \Pr \left[\left| \# \{i \mid \mathbf{x}(0)_i = a, \mathbf{x}(h)_i = b'\} - n\nu_{\vec{k}}(a)\nu_{\vec{k}}(b) \right| \geq 3\varepsilon n \right] \\ & \leq 2\alpha^{-1}e^{-\alpha\varepsilon^2 n} + \Pr \left[\left| \sum_{i=1}^n \mathbf{Z}_i - n\nu_{\vec{k}}(a)\nu_{\vec{k}}(b) \right| \geq \varepsilon n \right] \\ & \leq 4\alpha^{-1}e^{-\alpha\varepsilon^2 n}. \end{aligned}$$

□

Claim 3.15 motivates us to write the operator T_μ^h as a convex combination of operators according to the statistics of $(\mathbf{x}, \mathbf{y})$ where $\mathbf{x} \in_R \mathcal{U}_{\vec{k}}$ and $\mathbf{y} \sim T_\mu^h \mathbf{x}$. That is, for each list of non-negative integers $\vec{r} = (r_{a,b})_{a,b \in [m]}$ that sum of to n , let $\mathcal{R}_{\vec{r}}$ be the operator corresponding to the distribution of $(\mathbf{x}, \mathbf{y})$ where $\mathbf{x} \in \mathcal{U}_{\vec{k}}$ is uniform and $\mathbf{y} \sim T_\mu^h \mathbf{x}$, conditioned on the statistics of $(\mathbf{x}, \mathbf{y})$ being $\vec{r}$. Then we may write $T_\mu^h = \sum_{\vec{r}} p_{\vec{r}} \mathcal{R}_{\vec{r}}$ where $p_{\vec{r}}$ is the probability that the statistics $\vec{r}$ is achieved by $(\mathbf{x}, T_\mu^h \mathbf{x})$.

Definition 3.16. We say $\vec{r}$ is ε -reasonable if $|r_{a,b} - n\nu_{\vec{k}}(a)\nu_{\vec{k}}(b)| \leq 3\varepsilon n$ for all $a', b' \in [m']$.

By Claim 3.15 and the union bound, we have $\sum_{\vec{r} \text{ unreasonable}} p_{\vec{r}} \leq 4m^2\alpha^{-1}e^{-\alpha\varepsilon^2 n}$, so it will be enough for us to show that the operators $\mathcal{R}_{\vec{r}}$ strongly contract high degree functions for reasonable $\vec{r}$.

By the symmetry of μ , it follows that for each $\vec{r}$, the distribution of $(\mathbf{x}, \mathbf{y})$, where $\mathbf{x} \in \mathcal{U}_{\vec{k}}, \mathbf{y} \sim \mathcal{R}_{\vec{r}} \mathbf{x}$ is also symmetric, and hence uniform over vectors $(x, y) \in \mathcal{U}_{\vec{k}} \times \mathcal{U}_{\vec{k}}$ achieving the statistics $\vec{r}$. In particular, the operator $\mathcal{R}_{\vec{r}}$ is self-adjoint. Also, from Claim 3.2 we get that the distribution $(\mathbf{x}, \mathcal{R}_{\vec{r}} \mathbf{x})$ commutes with the action of S_n , and so the operator $\mathcal{R}_{\vec{r}}$ commutes with the action of S_n on functions. Thus, the following claim finds invariant subspaces of $\mathcal{R}_{\vec{r}}$.

Claim 3.17. Suppose $\mathcal{R}: L^2(\mathcal{U}_{\vec{k}}) \rightarrow L^2(\mathcal{U}_{\vec{k}})$ is an operator such that both $\mathcal{R}$ and $\mathcal{R}^*$ commute with the action of S_n . Then

1. for all $\lambda \vdash n$, the space $V_{=\lambda}(\mathcal{U}_{\vec{k}})$ is an invariant space of $\mathcal{R}$;
2. consequently, for all d the space $V_{=d}(\mathcal{U}_{\vec{k}})$ is an invariant space of $\mathcal{R}$.

Proof. Let $\lambda \vdash n$, and $A = (A_1, \dots, A_r)$ be a partition of $[n]$ such that $|A_i| = \lambda_i$. Recall the spaces $V_A(\mathcal{U}_{\vec{k}})$ from Section 2.2.3; then for all $f \in V_A(\mathcal{U}_{\vec{k}})$ and $\pi \in S_n$ such that $\pi(A_i) = A_i$ it holds that

$$\pi(\mathcal{R}f) = \mathcal{R}(\pi f) = \mathcal{R}f,$$

and so $\mathcal{R}f \in V_A(\mathcal{U}_{\vec{k}})$. Thus, each $V_A(\mathcal{U}_{\vec{k}})$ is an invariant space of $\mathcal{R}$, and by Claim 2.10 it follows that $V_\lambda(\mathcal{U}_{\vec{k}})$ is an invariant space of $\mathcal{R}$; the same argument applies to $\mathcal{R}^*$.

Now, letting $f \in V_{=\lambda}(\mathcal{U}_{\vec{k}})$, $\lambda' \triangleright \lambda$ and $g \in V_{=\lambda'}(\mathcal{U}_{\vec{k}})$, we have that $\mathcal{R}^*g \in V_{=\lambda}(\mathcal{U}_{\vec{k}})$ and so

$$\langle \mathcal{R}f, g \rangle = \langle f, \mathcal{R}^*g \rangle = 0,$$

so by the alternative definition from Section 2.2.3 we get that $\mathcal{R}f \in V_{=\lambda}(\mathcal{U}_{\vec{k}})$.

The second item follows by summing the first item over $\lambda \vdash n$ such that $\lambda_1 = n - d$. $\square$

Claim 3.18. *For all $\alpha, \gamma > 0$, $m \in \mathbb{N}$ there exist $\delta > 0$ and $\varepsilon > 0$ such that the following holds. Suppose $\vec{r}$ is ε -reasonable. Then for all $d \geq \gamma n$ and $f \in V_{=d}(\mathcal{U}_{\vec{k}})$ we have that $\|\mathcal{R}_{\vec{r}}f\|_2 \leq (1 + \delta_2)^{-d} \|f\|_2$.*

Proof. We prove that for all $\lambda \vdash n$ such that $\lambda_1 = n - d$, and $f \in V_{=\lambda}(\mathcal{U}_{\vec{k}})$ it holds that $\|\mathcal{R}_{\vec{r}}f\|_2 \leq (1 + \delta)^{-d} \|f\|_2$, from which the statement follows as $V_{=d}(\mathcal{U}_{\vec{k}}) = \bigoplus_{\lambda \vdash n, \lambda_1 = n-d} V_{=\lambda}(\mathcal{U}_{\vec{k}})$.

Fix $\lambda \vdash n$ such that $\lambda_1 = n - d$ and $V_{=\lambda}(\mathcal{U}_{\vec{k}}) \neq \{0\}$. By Claim 3.17, the space $V_{=\lambda}(\mathcal{U}_{\vec{k}})$ is an invariant under $\mathcal{R}_{\vec{r}}$, so we may decompose $V_{=\lambda}(\mathcal{U}_{\vec{k}})$ as $\bigoplus_{\theta} V_{=\lambda}^{\theta}(\mathcal{U}_{\vec{k}})$ where each $V_{=\lambda}^{\theta}(\mathcal{U}_{\vec{k}})$ is an eigenspace of $\mathcal{R}_{\vec{r}}$ with eigenvalue θ . Therefore, to establish our claim it suffices to show that for any θ such that $V_{=\lambda}^{\theta}(\mathcal{U}_{\vec{k}}) \neq \{0\}$, it holds that $|\theta| \leq (1 - \delta_2)^n$ for $\delta_2 = \delta_2(\alpha, \gamma) > 0$.

Fix such θ , take a non-zero $f \in V_{=\lambda}^{\theta}(\mathcal{U}_{\vec{k}})$, and let m_f be the multiplicity of θ . Then since the trace of an operator is the sum of its eigenvalues, we have that

$$\text{Tr}(\mathcal{R}_{\vec{r}}^2) \geq m_f \theta^2, \quad (4)$$

and as $m_f \geq \dim(\text{span}(\{\pi f \mid \pi \in S_n\}))$ (since each πf is an eigenvector of $\mathcal{R}_{\vec{r}}$ with eigenvalue α) we get by Claim 2.9 that $m_f \geq \dim(\lambda)$. By Claim 2.7, the latter is at least $(1 + \delta_1)^n$ for some $\delta_1 = \delta_1(\alpha, \gamma) > 0$, where we used $d \geq \gamma n$ (which we assumed to hold) and $\lambda_1 \leq (1 - \alpha)n$ (which holds by Claim 2.8). Thus, it is enough to prove that $\text{Tr}(\mathcal{R}_{\vec{r}}^2) \leq (1 + \delta_1/2)^n$, and this is done by a direct computation. Let $\mathbf{x} \in \mathcal{U}_{\vec{k}}$, $\mathbf{y} \sim \mathcal{R}_{\vec{r}}\mathbf{x}$ and $\mathbf{z} \sim \mathcal{R}_{\vec{r}}\mathbf{y}$, then by symmetry the probability that $\mathbf{z} = \mathbf{x}$ is $\frac{1}{\prod_{a \in [m]} \binom{k_a}{r_{a,1}, \dots, r_{a,m}}}$, so by Fact A.1

$$\text{Tr}(\mathcal{R}_{\vec{r}}^2) = \frac{\binom{n}{k_1, \dots, k_m}}{\prod_{a \in [m]} \binom{k_a}{r_{a,1}, \dots, r_{a,m}}} \leq n^{m^2} \frac{2^{H\left(\frac{k_1}{n}, \dots, \frac{k_m}{n}\right)n}}{\prod_{a \in [m]} 2^{H\left(\frac{r_{a,1}}{k_a}, \dots, \frac{r_{a,m}}{k_a}\right)k_a}} = n^{m^2} \frac{2^{H(\nu_{\vec{k}}(1), \dots, \nu_{\vec{k}}(m))n}}{\prod_{a \in [m]} 2^{H\left(\frac{r_{a,1}}{k_a}, \dots, \frac{r_{a,m}}{k_a}\right)k_a}}. \quad (5)$$

Since $\vec{r}$ is reasonable, for all $a, b \in [m]$ we have

$$\left| \frac{r_{a,b}}{k_a} - \nu_{\vec{k}}(b) \right| = \frac{|r_{a,b} - n\nu_{\vec{k}}(a)\nu_{\vec{k}}(b)|}{k_a} \leq \frac{3\varepsilon n}{k_a} \leq \frac{3}{\alpha}\varepsilon.$$

It follows from Fact A.2 that

$$\left| H(\nu_{\vec{k}'}(1), \dots, \nu_{\vec{k}'}(m')) - H\left(\frac{r_{a,1}}{k_a}, \dots, \frac{r_{a,m}}{k_a}\right) \right| = \tilde{O}_{m,\alpha}(\varepsilon),$$

and plugging this into (5) yields that $\text{Tr}(\mathcal{R}_{\vec{r}}^2) \leq n^{m^2} 2^{\tilde{O}_{m,\alpha}(\varepsilon)n} \leq (1 + \delta_1/2)^n$ provided ε is small enough with respect to m and α . Plugging everything into (4) yields that $(1 + \delta_1/2)^n \geq (1 + \delta_1)^n \theta^2$, and so $|\theta| \leq (1 - \delta_2)^n$ for $\delta_2(\alpha, \gamma) > 0$. $\square$

We are now ready to prove Claim 3.13.

Proof of Claim 3.13. Fix $d \geq \gamma n$, pick $\delta(\alpha, \gamma, m), \varepsilon(\alpha, \gamma, m) > 0$ from Claim 3.18 and then $h = h(\alpha, \varepsilon) > 0$ from Claim 3.14. By the symmetry of μ , it follows from Claims 3.2, 3.17 that $V_{=d}(\mathcal{U}_{\vec{k}})$ is invariant under T_μ . We may therefore decompose it as a sum of eigenspaces of T_μ , and so by Parseval it suffices to prove the claim for each one of these eigenspaces.

Let $f \in V_{=d}(\mathcal{U}_{\vec{k}})$ be a non-zero eigenvector of T_μ with eigenvalue θ , then

$$|\theta|^h \|f\|_2 = \left\| T_\mu^h f \right\|_2 \leq \sum_{\vec{r}} p_{\vec{r}} \|\mathcal{R}_{\vec{r}} f\|_2 = \sum_{\substack{\vec{r} \\ \text{reasonable}}} p_{\vec{r}} \|\mathcal{R}_{\vec{r}} f\|_2 + \sum_{\substack{\vec{r} \\ \text{unreasonable}}} p_{\vec{r}} \|\mathcal{R}_{\vec{r}} f\|_2.$$

For the first sum, we bound $\|\mathcal{R}_{\vec{r}} f\|_2 \leq (1 + \delta_2)^{-n} \|f\|_2$ by Claim 3.18, and the sum of the $p_{\vec{r}}$'s by 1. For the second sum, we bound $\|\mathcal{R}_{\vec{r}} f\|_2 \leq \|f\|_2$ by Jensen's inequality and the sum of $p_{\vec{r}}$'s by $4m^2 \alpha^{-1} e^{-\alpha \varepsilon^2 n}$. Overall, we get that

$$|\theta|^h \|f\|_2 \leq (4m^2 \alpha^{-1} e^{-\alpha \varepsilon^2 n} + (1 + \delta_2)^{-n}) \|f\|_2,$$

and dividing by $\|f\|_2$, simplifying and taking h -th root finishes the proof. $\square$

3.4.2 The case that $d < \gamma n$

In this section, we prove the following claim.

Claim 3.19. *For all $\alpha > 0$, $m \in \mathbb{N}$ there are $\gamma > 0$, $\delta > 0$ and $C > 0$ such that the following holds. Let $\mathcal{U}_{\vec{k}} \subseteq [m]^n$ be α -balanced, and μ be a α -admissible, connected distribution over $\mathcal{U}_{\vec{k}} \times \mathcal{U}_{\vec{k}}$, such that T_μ is self-adjoint.*

For all $d \leq \gamma n$, if $f \in V_{=d}(\mathcal{U}_{\vec{k}})$, then $\|T_\mu f\|_2 \leq C(1 + \delta)^{-d} \|f\|_2$.

For the rest of this section, we fix $\alpha > 0$, $m \in \mathbb{N}$ and μ as in the above claim. The proof follows the same sequence of claims as in the previous section, except for Claim 3.18 whose formulation and proof is adjusted to this case.

Claim 3.20. *For all $\alpha > 0$, $m \in \mathbb{N}$ there are $\varepsilon > 0$, $\gamma > 0$, $\delta > 0$ such that the following holds. Suppose $\vec{r}$ be ε -reasonable. Then for all $d \leq \gamma n$ and $f \in V_{=d}(\mathcal{U}_{\vec{k}})$ we have that $\|\mathcal{R}_{\vec{r}} f\|_2 \leq (1 + \delta)^{-d} \|f\|_2$.*

Proof overview. We begin with an outline of the argument. Let $\varepsilon > 0$ be sufficiently small to be determined later, and take $\gamma = \min(\alpha/2, \varepsilon)$. Then the α -balancedness of our multi-slices implies that for any $S \subseteq [n]$ of size γn , the support of x_S where $x \in \mathcal{U}_{\vec{k}}$ is full, i.e. $[m]^S$.

As before, we will reduce Claim 3.20 to understanding the eigenvalues of $\mathcal{R}_{\vec{r}}$ on the space $V_{=d}(\mathcal{U}_{\vec{k}})$. Here however, we will show that for any eigenvalue $\theta \neq 0$ of $\mathcal{R}_{\vec{r}}$ in this space, there is a d -junta $f \in V_{=d}(\mathcal{U}_{\vec{k}})$ for which $\mathcal{R}_{\vec{r}} f = \theta f$. Assume without loss of generality that the junta is $S = \{1, \dots, d\}$. Roughly speaking, the idea then is to take a set of $N = 3d$ coordinates containing the junta coordinates, then “project everything” onto these coordinates, reducing us to case very similar to the $d \geq \gamma n$ case.

More precisely, we will consider the induced measure ν on $[m]^N$, and the natural projection of the operator $\mathcal{R}_{\vec{r}}$ on this space, i.e. $\mathcal{R}_{\vec{r}, N}: L^2([m]^N, \nu) \rightarrow L^2([m]^N, \nu)$. We will do so while keeping f as an eigenvector of $\mathcal{R}_{\vec{r}, N}$ with eigenvalue θ . We will then study the operator $\mathcal{R}_{\vec{r}, N}$, and in particular show (similarly to before) that: (a) $\text{Tr}(\mathcal{R}_{\vec{r}, N}^2) \leq (1 + \xi(h))^d$ where $\lim_{h \rightarrow \infty} \xi(h) = 0$, and (b) the multiplicity of the eigenvalue θ is large, namely at least $(1 + \Omega(1))^d$. Combining these two facts yields, as before, that $|\theta| \leq (1 - \delta')^d$ for some $\delta' = \delta(\alpha, \gamma, m) > 0$.

We remark that the proof of (a) is again by a direct computation (as can perhaps be expected). As for (b), here we use the fact that the multiplicity of f is at least $\dim(\text{span}(\{\pi f \mid \pi \in S_N\}))$ (as each πf will automatically be an eigenvector of $\mathcal{R}_{\vec{r},N}$), so our task reduces to showing that this is a space of large dimension. To evaluate this dimension, we construct an invertible linear map $\psi: L^2([m]^N, \nu) \rightarrow L^2([m]^N, \text{Uniform})$ that commutes with the action of S_N , and show that the space $\psi(\text{span}(\{\pi f \mid \pi \in S_N\}))$ has a large dimension by orthogonality considerations (which are much easier to carry out in a product space).

We now proceed to the formal proof, and we fix an ε -reasonable statistics $\vec{r}$ for the rest of the proof. By Claims 3.2, 3.17 the space $V_{=d}(\mathcal{U}_{\vec{k}})$ is invariant under $\mathcal{R}_{\vec{r}}$. Thus, we may decompose it as $V_{=d}(\mathcal{U}_{\vec{k}}) = \bigoplus_{\theta} V_{=d}^{\theta}(\mathcal{U}_{\vec{k}})$ where $V_{=d}^{\theta}(\mathcal{U}_{\vec{k}})$ is an eigenspace of $\mathcal{R}_{\vec{r}}$ with eigenvalue θ . Thus, by Parseval, to prove Claim 3.2 it is enough to show that if $V_{=d}^{\theta}(\mathcal{U}_{\vec{k}}) \neq \{0\}$, then $|\theta| \leq C(1 - \delta)^d$.

Recall the junta spaces $V_J(\mathcal{U}_{\vec{k}})$. We further refine them by considering $V_{=J}(\mathcal{U}_{\vec{k}}) = V_J(\mathcal{U}_{\vec{k}}) \cap V_{d-1}(\mathcal{U}_{\vec{k}})^{\perp}$ for all J of size d .

Claim 3.21. *Let θ be such that $V_{=d}^{\theta}(\mathcal{U}_{\vec{k}}) \neq \{0\}$. Then there is a non-zero d -junta $f \in V_{=d}^{\theta}(\mathcal{U}_{\vec{k}})$.*

Proof. By Claim 3.6 the operator $\mathcal{R}_{\vec{r}}$ preserves the spaces $V_J(\mathcal{U}_{\vec{k}})$ as well as $V_{d-1}(\mathcal{U}_{\vec{k}})$. Since it is self-adjoint, it follows that $\mathcal{R}_{\vec{r}}$ preserves $V_{=J}(\mathcal{U}_{\vec{k}})$. Thus, we may write each such space as $V_{=J}(\mathcal{U}_{\vec{k}}) = \bigoplus_{\theta} V_{=J}^{\theta}(\mathcal{U}_{\vec{k}})$ where $V_{=J}^{\theta}(\mathcal{U}_{\vec{k}})$ is an eigenspace of $\mathcal{R}_{\vec{r}}$ with eigenvalue θ . Thus, we get that

$$\bigoplus_{\theta} V_{=d}^{\theta}(\mathcal{U}_{\vec{k}}) = V_{=d}(\mathcal{U}_{\vec{k}}) = \bigoplus_J V_J(\mathcal{U}_{\vec{k}}) = \bigoplus_J \bigoplus_{\theta} V_{=J}^{\theta}(\mathcal{U}_{\vec{k}}) = \bigoplus_{\theta} \bigoplus_J V_{=J}^{\theta}(\mathcal{U}_{\vec{k}}).$$

By uniqueness of the decomposition into eigenvalues, we get that for all θ , $V_{=d}^{\theta}(\mathcal{U}_{\vec{k}}) = \bigoplus_J V_{=J}^{\theta}(\mathcal{U}_{\vec{k}})$, and the claim follows. $\square$

Fix θ such that $V_{=d}^{\theta}(\mathcal{U}_{\vec{k}}) \neq \{0\}$, and pick $f^* \neq 0$ a J -junta from Claim 3.21 for $|J| = d$. Without loss of generality, assume that $J = [d]$, and set $N = 3d$. Let $\mathcal{D}$ be the uniform measure over $\mathcal{U}_{\vec{k}}$, and let $\mathcal{D}_{\downarrow}$ be the marginal distribution of $x_{[N]}$ where $x \sim \mathcal{D}$. Given a function $g: [m]^N \rightarrow \mathbb{R}$, we may lift it to $\tilde{g}: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ as $\tilde{g}(z) = g(z_{[N]})$. Define the operator $\mathcal{R}_{\vec{r},N}: L^2([m]^n, \nu) \rightarrow L^2([m]^n, \nu)$ as follows:

$$(\mathcal{R}_{\vec{r},N}g)(x) = \sum_{y \in [m]^{n-N}} \frac{\mathcal{D}(x, y)}{\mathcal{D}_{\downarrow}(x)} (\mathcal{R}_{\vec{r}}\tilde{g})(x, y).$$

Abusing notation, since f^* is a $[d]$ -junta and $[d] \subseteq [N]$, we may think of f^* as a function over $[m]^N$. Let $F = \text{span}(\{\pi f^* \mid \pi \in S_N\})$.

Claim 3.22. *For all $f \in F$, we have $\mathcal{R}_{\vec{r},N}f = \theta f$.*

Proof. Clearly, it is enough to prove that each πf^* is an eigenvector. The proof is virtually identical for all π (replacing the role of $[d]$ with $\pi([d])$), and we thus assume that π is the identity permutation. By definition,

$$\mathcal{R}_{\vec{r},N}f^*(x) = \sum_{y \in [m]^{n-N}} \frac{\mathcal{D}(x, y)}{\mathcal{D}_{\downarrow}(x)} (\mathcal{R}_{\vec{r}}\tilde{f}^*)(x, y) = \sum_{y \in [m]^{n-N}} \frac{\mathcal{D}(x, y)}{\mathcal{D}_{\downarrow}(x)} \theta \tilde{f}^*(x, y) = \theta \sum_{y \in [m]^{n-N}} \frac{\mathcal{D}(x, y)}{\mathcal{D}_{\downarrow}(x)} f^*(x),$$

which is equal to $\theta f^*(x)$. $\square$

Thus, using the trace method we see that $\text{Tr}(\mathcal{R}_{\vec{r},N}^2) \geq \theta^2 \dim(F)$, i.e. that

$$|\theta| \leq \sqrt{\frac{\text{Tr}(\mathcal{R}_{\vec{r},N}^2)}{\dim(F)}}, \quad (6)$$

and in the rest of the proof we prove an upper bound on the above trace and lower bound on the dimension of F .

Claim 3.23. $\dim(F) \geq \binom{2d}{d}$.

Proof. Using the expectation inner product according to $\mathcal{D}_\downarrow$, we may define the spaces $V_{=j}([m]^N, \mathcal{D}_\downarrow)$ for $j = 0, 1, \dots, N$. It is easily seen then that $F \subseteq V_{=d}([m]^N, \mathcal{D}_\downarrow)$.

Consider the linear map $\Psi: V_{=d}([m]^N, \mathcal{D}_\downarrow) \rightarrow V_{=d}([m]^N, \text{Uniform})$ defined as follows. Given a function $g \in V_{=d}([m]^N, \mathcal{D}_\downarrow)$, we think of it as a function over $[m]^N$ with the uniform measure, and then take its projection g_d onto $V_{=d}([m]^N, \text{Uniform})$; i.e. $\Psi(g) = g_d$.

We claim that Ψ commutes with the action of S_n . Indeed, let $\pi \in S_n$, and write ${}^\pi g$ as $h_0 + \dots + h_d$ where h_j is the projection of ${}^\pi g$ onto $V_{=j}([m]^N, \text{Uniform})$, so that $\Psi({}^\pi g) = h_d$. Also, we have that $g(x) = {}^{\pi^{-1}} h_0 + \dots + {}^{\pi^{-1}} h_d$, and since the spaces $V_{=j}([m]^N, \text{Uniform})$ are invariant under the action of S_n , we get that ${}^{\pi^{-1}} h_d$ is the projection of g onto $V_{=d}([m]^N, \text{Uniform})$. We get that ${}^\pi \Psi(g) = {}^\pi ({}^{\pi^{-1}} h_d) = h_d$, and $\Psi({}^\pi g) = h_d$, and so Ψ commutes with π .

We claim that Ψ is invertible. Indeed, we show that if $\Psi(g) = 0$, then $g = 0$. As $g \in V_{=d}([m]^N, \mathcal{D}_\downarrow)$, we get that g is a linear combination of d -juntas so $g \in V_d([m]^N, \text{Uniform})$, but as $\Psi(g) = 0$ we have that its projection onto $V_{=d}([m]^N, \text{Uniform})$ is 0, so $g \in V_{d-1}([m]^N, \text{Uniform})$. Thus, g is a linear combination of $(d-1)$ -juntas, and so $g \in V_{d-1}([m]^N, \mathcal{D}_\downarrow)$; it follows that $g = 0$. We conclude that $\dim(F) = \dim(\Psi(F))$, and in the rest of the proof we lower bound the latter quantity.

We note that as f^* is a $[d]$ -junta, $g^* = \Psi(f^*)$ is also a $[d]$ -junta. Thus, since Ψ commutes with the action of S_n , for all $\pi \in S_n$, the function $\Psi({}^\pi f^*) = {}^\pi g^*$ is a $\pi([d])$ -junta. We now consider $B = \{{}^\pi g^* \mid \pi \in S_n\}$, and we show that B contains a large orthogonal set.

We claim that if π, π' are permutations such that $\pi([d]) \neq \pi'([d])$, then ${}^\pi g^*$ and ${}^{\pi'} g^*$ are orthogonal. This is equivalent to showing that if $\pi^{-1}([d]) \neq [d]$, then ${}^\pi g^*$ and g^* are orthogonal. To see that, let $I = \pi^{-1}([d])$; then we have that

$$\langle {}^\pi g^*, g^* \rangle = \mathbb{E}_{x \in [m']^I, y \in [m']^{[N] \setminus I}} [g^*(\pi(x, y)) g^*(x, y)] = \mathbb{E}_{x \in [m']^I} \left[g^*(\pi(x, 0)) \mathbb{E}_y [g^*(x, y)] \right].$$

Letting I' be the intersection of $[d]$ and I , we see that I' has size at most $d-1$ and $\mathbb{E}_y [g^*(x, y)] = \mathbb{E}_{z \in [m]^{[N] \setminus I'}} [g^*(x_{I'}, z)] = 0$, so $\langle {}^\pi g^*, g^* \rangle = 0$.

Consider the collection $\tilde{L}$ of injections $\tilde{\pi}: [d] \rightarrow [N] \setminus [d]$, and extend each $\tilde{\pi} \in \tilde{L}$ to a permutation $\pi \in S_n$ arbitrarily, to form a collection of permutations L . We get that any two permutations from L disagree on $[d]$, and by the above paragraph it follows that

$$\dim(\Psi(F)) \geq |L| = |\tilde{L}| = \frac{(2d)!}{d!^2} = \binom{2d}{d}. \quad \square$$

Claim 3.24. $\text{Tr}(\mathcal{R}_{\vec{r},N}^2) \leq (1 + O_\alpha(\varepsilon))^d$.

Proof. Let $x \in [m]^N$, and let $1_x \in \{0, 1\}^{m^N}$ be the indicator vector of x . Denote the diagonal entry of $\mathcal{R}_{\vec{r}, N}^2$ corresponding to x by $a_{x,x}$; then $a_{x,x} = 1_x^t \cdot \mathcal{R}_{\vec{r}, N}^2 1_x$. Expanding the definition of $\mathcal{R}_{\vec{r}, N}$, the vector $\mathcal{R}_{\vec{r}, N}^2 1_x$ is a probability vector whose x 's entry is equal to

$$\Pr_{\substack{\mathbf{w}' \sim \mathcal{D} \\ \mathbf{w}' \sim \mathcal{R}_{\vec{r}}^2 \mathbf{w}}} [\mathbf{w}'_{[N]} = x' \mid \mathbf{w}_{[N]} = x], \text{ so } a_{x,x} = \Pr_{\substack{\mathbf{w}' \sim \mathcal{D} \\ \mathbf{w}' \sim \mathcal{R}_{\vec{r}}^2 \mathbf{w}}} [\mathbf{w}'_{[N]} = x \mid \mathbf{w}_{[N]} = x].$$

Fix any w such that $w_{[N]} = x$. In the rest of the proof, we show that $\Pr_{\mathbf{w}' \sim \mathcal{R}_{\vec{r}}^2 w} [\mathbf{w}'_{[N]} = x] \leq \mathcal{D}_{\downarrow}(x)(1 + O_{\alpha}(\varepsilon))^N$, from which the claim quickly follows: it implies that $a_{x,x} \leq \mathcal{D}_{\downarrow}(x)(1 + O_{\alpha}(\varepsilon))^N$, which by summing over x and using $N = 3d$ gives us $\text{Tr}(\mathcal{R}_{\vec{r}, N}^2) \leq (1 + O_{\alpha}(\varepsilon))^d$.

Indeed, using conditional probabilities we have that

$$\Pr_{\mathbf{w}' \sim \mathcal{R}_{\vec{r}}^2 w} [\mathbf{w}'_{[N]} = x] = \prod_{i=1}^N \Pr_{\mathbf{w}' \sim \mathcal{R}_{\vec{r}}^2 w} [\mathbf{w}'_i = x_i \mid \mathbf{w}'_{<i} = x_{<i}].$$

By definition of the operator $\mathcal{R}_{\vec{r}}$, we have for $i = 1$ that $\Pr_{\mathbf{w}' \sim \mathcal{R}_{\vec{r}}^2 w} [\mathbf{w}'_1 = x_1] = \frac{r_{x_1, x_1}}{n}$, and for all $i > 1$

$$\Pr_{\mathbf{w}' \sim \mathcal{R}_{\vec{r}}^2 w} [\mathbf{w}'_i = x_i \mid \mathbf{w}'_{<i} = x_{<i}] \leq \frac{r_{x_i, x_i}}{k_{x_i} - i} \leq \frac{r_{x_i, x_i}}{k_{x_i}} \frac{k_{x_i}}{k_{x_i} - N} = \frac{r_{x_i, x_i}}{k_{x_i}} \left(1 + \frac{N}{k_{x_i} - N}\right),$$

which is at most $\frac{r_{x_i, x_i}}{k_{x_i}}(1 + O_{\alpha}(\varepsilon))$ as $N \leq \gamma n \leq \varepsilon n, \alpha n/2$ and $k_{x_i} \geq \alpha n$. Thus we get that

$$\Pr_{\mathbf{w}' \sim \mathcal{R}_{\vec{r}}^2 w} [\mathbf{w}'_{[N]} = x] \leq (1 + O_{\alpha}(\varepsilon))^N \prod_{i=1}^N \frac{r_{x_i, x_i}}{k_{x_i}}.$$

Next, note that

$$\mathcal{D}_{\downarrow}(x) \geq \prod_{i=1}^N \frac{k_{x_i} - i}{n} \geq (1 - \varepsilon)^N \prod_{i=1}^N \frac{k_{x_i}}{n},$$

so we get that

$$\frac{\Pr_{\mathbf{w}' \sim \mathcal{R}_{\vec{r}}^2 w} [\mathbf{w}'_{[N]} = x]}{\mathcal{D}_{\downarrow}(x)} \leq (1 + O_{\alpha}(\varepsilon))^N \prod_{i=1}^N \frac{r_{x_i, x_i}}{k_{x_i}^2/n}. \quad (7)$$

Since r is ε -reasonable,

$$\frac{r_{x_i, x_i}}{k_{x_i}^2/n} = 1 + \frac{r_{x_i, x_i} - n\nu_k^-(x_i)^2}{k_{x_i}^2/n} \leq 1 + \frac{3\varepsilon n}{k_{x_i}^2/n} \leq 1 + \frac{3}{\alpha^2}\varepsilon,$$

where we used $k_{x_i} \geq \alpha n$. Plugging this into (7) shows that $\frac{\Pr_{\mathbf{w}' \sim \mathcal{R}_{\vec{r}}^2 w} [\mathbf{w}'_{[N]} = x]}{\mathcal{D}_{\downarrow}(x)} \leq (1 + O_{\alpha}(\varepsilon))^N$, finishing the proof. $\square$

Plugging Claims 3.23, 3.24 into (6) bounds all eigenvalues of $\mathcal{R}_{\vec{r}}$ for reasonable $\vec{r}$, and the rest of the proof of Claim 3.20 is identical to the proof of Claim 3.13 and we do not repeat it.

3.4.3 Proof of Lemma 3.11

By Remark 3.12, we assume T_μ is self-adjoint. Fix $\alpha > 0$ and $m \in \mathbb{N}$. Take $\gamma_1, \delta_1 > 0$ and $C_1 > 0$ from Claim 3.19. Now take $\delta_2, C_2 > 0$ from Claim 3.13 for c, m and γ_1 ; we show that the lemma holds for $C = \max(C_1, C_2)$ and $\delta = \min(\delta_1, \delta_2)$.

Let $f \in V_{>d}(\mathcal{U}_{\vec{k}})$, and write $f = \sum_{j>d+1} f^{=j}$ where $f^{=j} \in V_{=j}(\mathcal{U}_{\vec{k}})$. By Claims 3.2, 3.17, the spaces $V_{=j}(\mathcal{U}_{\vec{k}})$ are invariant under T_μ , so we get that $(T_\mu f)^{=j} = T_\mu(f^{=j})$. Thus by Parseval $\|T_\mu f\|_2^2 = \sum_{j>d} \|T_\mu(f^{=j})\|_2^2$. By Claims 3.19, 3.13 we have that $\|T_\mu f^{=j}\|_2 \leq C(1 - \delta)^j \|f^{=j}\|_2$, so summing and using Parseval again gives us the statement of the lemma. $\square$

3.5 Proof of Theorem 3.10

Let $d \in \mathbb{N}$ to be chosen later, and write $f = f^{\leq d} + f^{>d}$, $g = g^{\leq d} + g^{>d}$. Then we have that

$$\mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \mu} [f(\mathbf{x})g(\mathbf{x}')] = \mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \mu} [f^{\leq d}(\mathbf{x})g^{\leq d}(\mathbf{x}')] + \mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \mu} [f^{\leq d}(\mathbf{x})g^{>d}(\mathbf{x}')] + \mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \mu} [f^{>d}(\mathbf{x})g(\mathbf{x}')], \quad (8)$$

and we first show that the second expectation is small in absolute value. Indeed, writing it as an inner product we have that it is equal to

$$\left| \langle T_\mu f^{\leq d}, g^{>d} \rangle \right| \leq \|T_\mu f^{\leq d}\|_2 \|g^{>d}\|_2 = \langle T_\mu^* T_\mu f^{\leq d}, f^{>d} \rangle^{1/2} \|g^{>d}\|_2 \leq \|T_\mu^* T_\mu f^{\leq d}\|_2^{1/2} \|f^{>d}\|_2^{1/2} \|g^{>d}\|_2.$$

Consider the distribution ν defined as: pick $(\mathbf{x}, \mathbf{x}') \sim \mu$, pick $(\mathbf{x}'', \mathbf{x}''')$ conditioned on $\mathbf{x}''' = \mathbf{x}'$, and output $(\mathbf{x}, \mathbf{x}'')$. Note that as μ is connected, ν is connected, as μ is α -admissible, ν is α^2 -admissible, and that the coupling between μ and $\tilde{\mu}$ naturally induces a coupling between ν and $\tilde{\nu}$ with similar parameters. Finally, note that $T_\nu = T_\mu^* T_\mu$.

Write $f^{>d} = \sum_{\ell>d} f^{=\ell}$ for $f^{=\ell} \in V_{=\ell}(\mathcal{U}_{\vec{k}})$. By Claim 3.17, the operator T_ν preserves the spaces $V_{=\ell}(\mathcal{U}_{\vec{k}})$, so by Parseval we have $\|T_\nu f^{>d}\|_2^2 = \sum_{\ell>d} \|T_\nu f^{=\ell}\|_2^2$. Using Lemma 3.11, we get that

$$\|T_\nu f^{>d}\|_2^2 \leq \sum_{\ell>d} C^2(1 + \delta)^{-2\ell} \|f^{=\ell}\|_2^2 \leq C^2(1 + \delta)^{-2d} \|f\|_2^2,$$

so the third expectation in (8) is at most $\sqrt{C}(1 + \delta)^{-d/2} \|f\|_2 \|g\|_2$. The same argument applies to show that the absolute value of the second expectation is also upper bounded by that. We thus get that

$$\left| \mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \mu} [f(\mathbf{x})g(\mathbf{x}')] - \mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \mu} [f^{\leq d}(\mathbf{x})g^{\leq d}(\mathbf{x}')] \right| \leq 2\sqrt{C}(1 + \delta)^{-d/2} \|f\|_2 \|g\|_2. \quad (9)$$

We now argue that

$$\left| \mathbb{E}_{(\mathbf{y}, \mathbf{y}') \sim \tilde{\mu}} [T_c f(\mathbf{y}) T_{c'} g(\mathbf{y}')] - \mathbb{E}_{(\mathbf{y}, \mathbf{y}') \sim \tilde{\mu}} [(T_c f)^{\leq d}(\mathbf{y}) (T_{c'} g)^{\leq d}(\mathbf{y}')] \right| \leq 2\sqrt{C'}(1 + \delta')^{-d/2} \|f\|_2 \|g\|_2. \quad (10)$$

Indeed, we show that by proving upper bounding the expectation of each one of $(T_C f)^{>d}(\mathbf{y})(T_{C'} g)(\mathbf{y}')$ and $(T_C f)^{\leq d}(\mathbf{y})(T_{C'} g)^{>d}(\mathbf{y}')$. The arguments are identical, and we demonstrate on the first one. Writing it as an inner product and using Cauchy-Schwarz we get that

$$\begin{aligned} \mathbb{E}_{(\mathbf{y}, \mathbf{y}') \sim \tilde{\mu}} \left[(T_C f)^{>d}(\mathbf{y})(T_{C'} g)(\mathbf{y}') \right] &= \langle T_{\tilde{\mu}}^* (T_C f)^{>d}, T_{C'} g \rangle \leq \left\| T_{\tilde{\mu}} (T_C f)^{>d} \right\|_2 \|T_{C'} g\|_2 \\ &\leq \left\| T_{\tilde{\mu}}^* T_{\tilde{\mu}} (T_C f)^{>d} \right\|_2^{1/2} \left\| (T_C f)^{>d} \right\|_2^{1/2} \|T_{C'} g\|_2. \end{aligned}$$

Consider the distribution μ' on $[m] \times [m]$ corresponding to picking $(\mathbf{a}, \mathbf{b}) \sim \tilde{\mu}_1$, and then $(\mathbf{a}', \mathbf{b}') \sim \tilde{\mu}_1$ conditioned on $\mathbf{b}' = \mathbf{b}$, outputting $\mathbf{a}, \mathbf{a}'$. By properties of μ we get that μ' is connected and the probability of each atom is at least α^2 . Also, $T_{\mu'} = T_{\tilde{\mu}}^* T_{\tilde{\mu}}$, so from Lemma B.3

$$\left\| T_{\tilde{\mu}}^* T_{\tilde{\mu}} (T_C f)^{>d} \right\|_2 \leq (1 - \Omega_\alpha(1))^d \left\| (T_C f)^{>d} \right\|_2 \leq (1 - \Omega_\alpha(1))^d \|f\|_2.$$

This completes the proof of (10).

By Lemma 3.25,

$$\left\| (T_C f)^{\leq d} - T_C(f^{\leq d}) \right\|_2 = O_{\alpha,d}(\zeta^{1/4}) \|f\|_2, \quad \left\| (T_{C'} g)^{\leq d} - T_{C'}(g^{\leq d}) \right\|_2 \leq O_{\alpha,d}(\zeta^{1/4}) \|g\|_2$$

(we used $\|f^{\leq d}\|_2 \leq \|f\|_2$ and similarly for g), so by another hybrid argument and Cauchy-Schwarz

$$\left| \mathbb{E}_{(\mathbf{y}, \mathbf{y}') \sim \tilde{\mu}} \left[(T_C f)^{\leq d}(\mathbf{y})(T_{C'} g)^{\leq d}(\mathbf{y}') \right] - \mathbb{E}_{(\mathbf{y}, \mathbf{y}') \sim \tilde{\mu}} \left[T_C(f^{\leq d})(\mathbf{y}) T_{C'}(g^{\leq d})(\mathbf{y}') \right] \right| \leq O_{\alpha,d}(\zeta^{1/4}) \|f\|_2 \|g\|_2. \quad (11)$$

Using Lemma 1.6 the same way we did as in Section 3.3.1, we get that

$$\begin{aligned} \left| \mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \mu} \left[f^{\leq d}(\mathbf{x}) g^{\leq d}(\mathbf{x}') \right] - \mathbb{E}_{(\mathbf{y}, \mathbf{y}') \sim \tilde{\mu}} \left[T_C(f^{\leq d})(\mathbf{y}) T_{C'}(g^{\leq d})(\mathbf{y}') \right] \right| &\leq O_d(\zeta^{1/4}) \|f^{\leq d}\|_2 \|g^{\leq d}\|_2 \\ &\leq O_d(\zeta^{1/4}) \|f\|_2 \|g\|_2. \end{aligned} \quad (12)$$

Combining (9), (10), (11) and (12) with the triangle inequality, and choosing d large enough so that $2\sqrt{C'}(1 + \delta')^{-d/2}, 2\sqrt{C}(1 + \delta)^{-d/2} \leq \varepsilon/4$, and then ζ small enough so that $O_d(\sqrt{\zeta}), O_{\alpha,d}(\zeta^{1/4}) \leq \varepsilon/4$, finishes the proof. $\square$

3.6 T_C almost commutes with degree truncations

In this section, we prove the following lemma, asserting that if $\mathcal{C}$ is a coupling with good enough parameters in comparison to d , then $(T_C f)^{\leq d} \approx T_C(f^{\leq d})$. More precisely:

Lemma 3.25. *Suppose $\mathcal{C}$ is a (α, ζ) -coupling between $(\mathcal{U}_{\vec{k}}, \text{Uniform})$ and $([m]^n, \nu_{\vec{k}}^{\otimes n})$ as in Definition 1.5. If $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ is a function, and $d \in \mathbb{N}$, then $\left\| (T_C f)^{\leq d} - T_C(f^{\leq d}) \right\|_2 \leq O(d^{7/4} \alpha^{-1/4} \zeta^{1/4} \|f\|_2)$.*

The rest of this section is devoted to the proof of Lemma 3.25.

Let T_ρ be the standard noise operator on $L^2([m]^n, \nu_{\vec{k}})$ from Section B.3. Our main goal is to show that for all $\rho > 0$ and $f \in V_{=\ell}(\mathcal{U}_{\vec{k}})$, it holds that

$$\left\| T_\rho T_C f - \rho^\ell T_C f \right\|_2^2 \leq 3\sqrt{\frac{4\ell}{\alpha}} \zeta \|f\|_2^2. \quad (13)$$

Intuitively, this would say that $T_C f$ is nearly of purely degree ℓ , and indeed once established, we will quickly deduce Lemma 3.25 from (13).

Claim 3.26. *Suppose that for every $\rho > 0$, $\ell \in \mathbb{N}$ and $g \in V_{=\ell}(\mathcal{U}_{\vec{k}})$ we have that*

$$\left\| (T_C^* T_\rho T_C - \rho^\ell I) f \right\|_2^2 \leq \frac{4\ell}{\alpha} \zeta \|f\|_2^2. \quad (14)$$

Then (13) holds for all $\rho > 0$ and $\ell \in \mathbb{N}$.

Proof.

$$\begin{aligned} \left\| T_\rho T_C f - \rho^\ell T_C f \right\|_2^2 &= \|T_\rho T_C f\|_2^2 + \rho^{2\ell} \|T_C f\|_2^2 - 2\rho^\ell \langle T_\rho T_C f, T_C f \rangle \\ &= \langle T_C^* T_\rho T_C f, f \rangle + \rho^{2\ell} \langle T_C^* T_C f, f \rangle - 2\rho^\ell \langle T_C^* T_\rho T_C f, f \rangle \\ &= \langle (T_C^* T_\rho T_C - \rho^{2\ell} I) f, f \rangle + \rho^{2\ell} \langle (T_C^* T_C - I) f, f \rangle - 2\rho^\ell \langle (T_C^* T_\rho T_C - \rho^\ell I) f, f \rangle \\ &\leq 3\sqrt{\frac{4\ell}{\alpha}} \zeta \|f\|_2^2, \end{aligned}$$

where in the last transition we used Cauchy-Schwarz and our assumption. $\square$

Let $S_\rho = T_C^* T_\rho T_C - \rho^\ell I$, and think of $S_\rho : L^2(\mathcal{U}_{\vec{k}}) \rightarrow L^2(\mathcal{U}_{\vec{k}})$. By symmetry of $\mathcal{C}$, it follows that T_C, T_C^* commute with the action of S_n on functions, and so S_ρ commutes with the action of S_n on functions. Thus, by Claim 3.6 it preserves juntas. Clearly, S_ρ is self-adjoint.

By Claim 3.17, S_ρ preserves the spaces $V_{=\ell}(\mathcal{U}_{\vec{k}})$, hence we may decompose it as a sum of eigenspaces $\bigoplus_\theta V_{=\ell}^\theta(\mathcal{U}_{\vec{k}})$. Thus, to establish (14), it suffices to show that if $V_{=\ell}^\theta(\mathcal{U}_{\vec{k}}) \neq \{0\}$, then $|\theta| \leq \frac{2\ell}{\alpha} \zeta$. We fix θ^* such that $V_{=\ell}^{\theta^*}(\mathcal{U}_{\vec{k}}) \neq \{0\}$ henceforth. Recall the spaces $V_{=J}(\mathcal{U}_{\vec{k}})$ from Section 3.4.2; it follows that these spaces are invariant under S_ρ , so we may decompose each one of them as a sum of eigenspaces $\bigoplus_\theta V_{=J}^\theta(\mathcal{U}_{\vec{k}})$. Denote $V_\ell(\theta) = \bigoplus_{|J|=\ell} V_{=J}^\theta(\mathcal{U}_{\vec{k}})$, then

$$\bigoplus_\theta V_\ell(\theta) = \bigoplus_\theta \bigoplus_{|J|=\ell} V_{=J}^\theta(\mathcal{U}_{\vec{k}}) = \bigoplus_{|J|=\ell} \bigoplus_\theta V_{=J}^\theta(\mathcal{U}_{\vec{k}}) = \bigoplus_{|J|=\ell} \bigoplus_\theta V_{=J}(\mathcal{U}_{\vec{k}}) = V_{=\ell}(\mathcal{U}_{\vec{k}}),$$

so $V_\ell(\theta)$ are the eigenspaces of S_ρ . Therefore, as $V_{=\ell}^{\theta^*}(\mathcal{U}_{\vec{k}}) \neq \{0\}$, there is $|J| = \ell$ such that $V_{=J}^{\theta^*}(\mathcal{U}_{\vec{k}}) \neq \{0\}$. We thus find a non-zero $f^* \in V_{=J}^{\theta^*}(\mathcal{U}_{\vec{k}})$.

Next, define an auxiliary operator $R_{\rho,J}$. For $x \in \mathcal{U}_{\vec{k}}$, we choose a set $\mathbf{S} \subseteq J$ randomly by including each $i \in [n]$ in $\mathbf{S}$ with probability ρ . We then pick $\mathbf{y} \in_R \mathcal{U}_{\vec{k}}$ conditioned on $\mathbf{y}_{\mathbf{S}} = x_{\mathbf{S}}$. We finish the proof with the following two claims.

Claim 3.27. $R_{\rho,J} f^* = \rho^\ell f^*$.

Proof. Write $R_{\rho,J} = \sum_{T \subseteq J} p_T R_{\rho,J,T}$, where p_T is the probability that $\mathbf{S} = T$, and $R_{\rho,J,T}$ is the action of the operator conditioned on $\mathbf{S} = T$. Note that $(R_{\rho,J,T} f^*)(y)$ is just the average of $f^*(\mathbf{x})$ over $\mathbf{x} \in \mathcal{U}_{\vec{k}}$ conditioned on $\mathbf{x}_T = y_T$, so it is proportional to the inner product $\langle f^*, 1_{x_T=y_T} \rangle$. If $|T| \leq \ell - 1$, then the last inner product is 0 as $1_{x_T=y_T} \in V_{\ell-1}(\mathcal{U}_{\vec{k}})$. Thus, $R_{\rho,J} f^* = p_J \cdot R_{\rho,J} f^* = \rho^\ell f^*$, where the last transition is since f^* is a J junta, and $p_J = \rho^{|J|} = \rho^\ell$. $\square$

Claim 3.28. $|\theta^*| \leq \frac{4\ell}{\alpha} \zeta$.

Proof. By the previous claim, $S_\rho f^* = (T_C^* T_\rho T_C - R_{\rho,J}) f^*$. Take x that maximizes $|f^*(x)|$. We now sample $\mathbf{y}$ according to $T_C^* T_\rho T_C x$, and $\mathbf{y}'$ according to $R_{\rho,J} x$ in a coupled way so that $\Pr[\mathbf{y}'_J \neq \mathbf{y}_J] \leq \frac{2|J|}{\alpha} \zeta$.

To do that, first sample $\mathbf{T} \subseteq [n]$ by including each $i \in [n]$ in $\mathbf{T}$ with probability ρ independently. Sample $\mathbf{y}(0) \sim T_C x$, and sample $\mathbf{y}(1)$ by taking $\mathbf{y}(1)_{\mathbf{T}} = \mathbf{y}(0)_{\mathbf{T}}$ and resampling $\mathbf{y}(1)_i \sim \nu_k^*$ for each $i \notin \mathbf{T}$ independently, and finally sample $\mathbf{y} \sim T_C^* \mathbf{y}(1)$. To sample $\mathbf{y}'$, run the sampling procedure of $R_{\rho,J} x$ with $S = \mathbf{T} \cap J$.

We note that if none of the coupling operators T_C, T_C^* change the coordinates in J , then we indeed have $\mathbf{y}'_J = \mathbf{y}_J$. Since $\mathcal{C}$ is a (α, ζ) -coupling, it follows from the union bound that this probability is at least $1 - \frac{2\ell\zeta}{\alpha}$.⁶ It follows that

$$|\theta^*| |f^*(x)| = |S_\rho f^*(x)| = \left| \mathbb{E}_{\mathbf{y}, \mathbf{y}'} [f^*(\mathbf{y}) - f^*(\mathbf{y}')] \right| \leq \Pr[\mathbf{y}'_J \neq \mathbf{y}_J] 2 |f^*(x)| \leq \frac{4\ell\zeta}{\alpha} |f^*(x)|,$$

and dividing by $|f^*(x)|$ finishes the proof. $\square$

It thus follows that (14) holds, and by Claim 3.26 we have that (13) holds. Finally, to establish Lemma 3.25, we require the following claim, asserting that T_C cannot decrease degrees.

Claim 3.29. $f \in V_d(\mathcal{U}_k)^\perp$, then $T_C f \in V_d([m]^n, \nu_k^{\otimes n})^\perp$.

Proof. To show that, it is enough to show that if $g \in V_d([m]^n, \nu_k^{\otimes n})$ then $T_C^* g \in V_d(\mathcal{U}_k)$, as that would say that

$$\langle T_C f, g \rangle = \langle f, T_C^* g \rangle = 0,$$

implying $T_C f \in V_d([m]^n, \nu_k^{\otimes n})^\perp$.

By linearity of the operator T_C^* , it suffices to show that if $|A| \leq d$, and g is a A -junta, then $T_C^* g$ is also a A junta. Fix A ; then as $\mathcal{C}$ is symmetric it follows that T_C^* commutes with the action of S_n on functions, so for all $\pi \in S_n$ that fixes the coordinates of A we have that $\pi(T_C^* g) = T_C^* \pi g = T_C^* g$, so by Fact 3.4 we get that $T_C^* g$ is indeed a A -junta. $\square$

We are now ready to prove Lemma 3.25.

Proof of Lemma 3.25. Let $\rho > 0$ to be determined later. By Claim 3.29 we get that $(T_C f)^{\leq d} = (T_C(f^{\leq d}))^{\leq d}$.

Let $g_j = T_C(f^{\leq j})$, and write $g_j = \sum_{i=0}^n g_j^{\leq i}$. Then by Fact B.10 it follows that $T_\rho g_j^{\leq i} = \rho^i g_j^{\leq i}$, and so by linearity $T_\rho g_j = \sum_{i=0}^n \rho^i g_j^{\leq i}$. On the other hand, by (13) $\|T_\rho g_j - \rho^j g_j\| \leq 6\sqrt{\frac{j}{\alpha}\zeta} \|f^{\leq j}\|_2^2$, and plugging in the expansion of g_j we get that

$$\sum_{i=0}^n (\rho^j - \rho^i)^2 \|g_j^{\leq i}\|_2^2 \leq 6\sqrt{\frac{j}{\alpha}\zeta} \|f^{\leq j}\|_2^2.$$

It follows that $\sum_{i \neq j} \|g_j^{\leq i}\|_2^2 \leq 6 \frac{\sqrt{j\zeta}}{\rho^{2j}(1-\rho)^2\sqrt{\alpha}} \|f^{\leq j}\|_2^2$, and therefore by Parseval

$$\left\| \sum_{i \neq j} g_j^{\leq i} \right\|_2 \leq \frac{\sqrt{6}(j\zeta)^{1/4}}{\rho^j(1-\rho)\alpha^{1/4}} \|f^{\leq j}\|_2,$$

⁶Strictly speaking, the second property in Definition 1.5 asserts that this is true when we sample $(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}$. However the symmetry of $\mathcal{C}$ and the balancedness of our multi-slices imply that for all i and x , when we sample $(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}$ conditioned on $\mathbf{x} = x$, we have that $\Pr[\mathbf{y}_i \neq x_i] \leq \zeta/\alpha$.

Summarizing, we get that

$$\begin{aligned}
\left\| (\mathsf{T}_C f)^{\leq d} - \mathsf{T}_C(f^{\leq d}) \right\|_2 &= \left\| (\mathsf{T}_C f^{\leq d})^{\leq d} - \mathsf{T}_C(f^{\leq d}) \right\|_2 \leq \sum_{j \leq d} \left\| (\mathsf{T}_C f^{=j})^{\leq d} - \mathsf{T}_C(f^{=j}) \right\|_2 \\
&= \sum_{j \leq d} \left\| \sum_{i > d} g_j^{=i} \right\|_2 \\
&\leq \sum_{j \leq d} \left\| \sum_{i \neq j} g_j^{=i} \right\|_2 \\
&\leq \sum_{j \leq d} \frac{\sqrt{6}(j\zeta)^{1/4}}{\rho^j(1-\rho)\alpha^{1/4}} \|f^{=j}\|_2 \\
&\leq \frac{\sqrt{6}(d\zeta)^{1/4}}{(1-\rho)\alpha^{1/4}} \sqrt{\sum_{j \leq d} \rho^{-j}} \sqrt{\sum_{j \leq d} \|f^{=j}\|_2^2} \\
&\leq \frac{\sqrt{6}(d\zeta)^{1/4}}{(1-\rho)\alpha^{1/4}} \sqrt{\frac{\rho^{-(d+1)} - 1}{\rho^{-1} - 1}} \|f\|_2 \\
&= \frac{\sqrt{6}(d\zeta)^{1/4}}{(1-\rho)^{3/2}\alpha^{1/4}} \rho^{-d/2} \|f\|_2.
\end{aligned}$$

Choosing $\rho = 1 - \frac{1}{2d} > 0$ finishes the proof. $\square$

4 Invariance principle in the multi-linear case

In this section, we generalize Theorem 3.10 to the case that the distribution μ is over r -tuples, where $r > 2$.

4.1 Invariance principle for products

In this section we prove Theorem 1.11, restated below.

Theorem 1.11 (Restated) . *For all $\alpha \in (0, 1)$, $M, r \in \mathbb{N}$, $m_1, \dots, m_r \in \mathbb{N}$, $\varepsilon > 0$ there are $\zeta > 0$, $N \in \mathbb{N}$ such that the following holds for $n \geq N$. Suppose $\mathcal{U}_{\vec{k}(1)}, \dots, \mathcal{U}_{\vec{k}(r)}$ are α -balanced multi-slices over alphabets $[m_1], \dots, [m_r]$ respectively, μ is a connected, α -admissible distribution over $\prod_{i=1}^r \mathcal{U}_{\vec{k}(i)}$, and let $\tilde{\mu}$ the product version of μ as in Definition 1.10. Suppose $\mathcal{C}_1, \dots, \mathcal{C}_r$ are couplings such that $\mathcal{C}_i$ is a (α, ζ) -coupling between $\mathcal{U}_{\vec{k}(i)}$ and $([m_i]^n, \nu_{\vec{k}(i)}^{\otimes n})$, and that there is a (α, ζ) -coupling, say $\mathcal{C}$, between μ and $\tilde{\mu}$.*

Then for all $f_i: \mathcal{U}_{\vec{k}(i)} \rightarrow \mathbb{R}$ such that $\|f_i\|_{2r} \leq M$ for all i , it holds that

$$\left| \mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu} \left[\prod_{i=1}^r f_i(\mathbf{x}(i)) \right] - \mathbb{E}_{(\mathbf{y}(1), \dots, \mathbf{y}(r)) \sim \tilde{\mu}} \left[\prod_{i=1}^r \mathsf{T}_{\mathcal{C}_i} f_i(\mathbf{y}(i)) \right] \right| \leq \varepsilon.$$

The rest of this section is devoted to the proof of Theorem 1.11. We choose a sequence of degrees $d_1 \leq d_2 \leq \dots \leq d_r$, where d_1 is sufficiently large function of $\alpha, r, m_1, \dots, m_r$ and ε , and for each i , once

d_i has been picked, d_{i+1} is picked to be sufficiently large. Finally, N is picked to be sufficiently large, and ζ is picked to be sufficiently small.

Claim 4.1.

$$\left| \mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu} \left[\prod_{i=1}^r f_i(\mathbf{x}(i)) \right] - \mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu} \left[\prod_{i=1}^r f_i^{\leq d_i}(\mathbf{x}(i)) \right] \right| \leq \frac{\varepsilon}{3}.$$

Proof. Let $g_i = f_i^{\leq d_i}$. We intend to prove for all $j = 0, 1, \dots, r-1$ that

$$\left| \mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu} \left[\prod_{i=1}^j g_i(\mathbf{x}(i)) \prod_{i=j+1}^r f_i(\mathbf{x}(i)) \right] - \mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu} \left[\prod_{i=1}^{j+1} g_i(\mathbf{x}(i)) \prod_{i=j+2}^r f_i(\mathbf{x}(i)) \right] \right| \leq \frac{\varepsilon}{3r},$$

and the claim follows from summing this up and using the triangle inequality. Fix j , then the above expression is equal to

$$\left| \mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu} \left[h(\mathbf{x}(1), \dots, \mathbf{x}(j), \mathbf{x}(j+2), \dots, \mathbf{x}(r)) f_{j+1}^{> d_{j+1}}(\mathbf{x}(j+1)) \right] \right|, \quad (15)$$

where $h(x(1), \dots, x(j), x(j+2), \dots, x(r)) = \prod_{i=1}^j g_i(x(i)) \prod_{i=j+2}^r f_i(x(i))$. Next, we interpret this expectation as an inner product. Consider the operator $T_{j+1}: L^2(\mathcal{U}_{\vec{k}(j+1)}) \rightarrow L^2(\prod_{i \neq j+1} \mathcal{U}_{\vec{k}(i)})$ defined as

$$T_{j+1}f(z) = \mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu} [f(\mathbf{x}(j+1)) \mid (\mathbf{x}(1), \dots, \mathbf{x}(j), \mathbf{x}(j+2), \dots, \mathbf{x}(r)) = z].$$

Consider the distribution ν_{j+1} over $\mathcal{U}_{\vec{k}(j+1)} \times \mathcal{U}_{\vec{k}(j+1)}$ as the distribution of $(\mathbf{x}(j+1), \mathbf{x}(j+1)')$ where we first pick $(\mathbf{x}(1), \dots, \mathbf{x}(h)) \sim \mu$, and then $(\mathbf{x}(1)', \dots, \mathbf{x}(h')) \sim \mu$ conditioned on $\mathbf{x}(i)' = \mathbf{x}(i)$ for all $i \neq j+1$. Then note that $T_\nu = T_{j+1}^* T_{j+1}$, and so

$$\begin{aligned} (15) &= \langle h, T_{j+1} f_{j+1}^{> d_{j+1}} \rangle \leq \|h\|_2 \left\| T_{j+1} f_{j+1}^{> d_{j+1}} \right\|_2 = \|h\|_2 \sqrt{\langle T_{j+1} f_{j+1}^{> d_{j+1}}, T_{j+1} f_{j+1}^{> d_{j+1}} \rangle} \\ &\leq \|h\|_2 \left\| f_{j+1}^{> d_{j+1}} \right\|_2^{1/2} \left\| T_\nu f_{j+1}^{> d_{j+1}} \right\|_2^{1/2}. \end{aligned} \quad (16)$$

First, we use Lemma 3.11 to bound $\left\| T_\nu f_{j+1}^{> d_{j+1}} \right\|_2$. Indeed, the connectedness of μ implies that ν is connected, and the α -admissibility of μ implies that ν is α^2 admissible. Also, using the coupling $\mathcal{C}$ one may easily construct a $(\alpha/2, 2\zeta)$ -coupling $\mathcal{C}'$ between ν and $\tilde{\nu}$. We thus have by Lemma 3.11 that

$$\left\| T_\nu f_{j+1}^{> d_{j+1}} \right\|_2 \leq C(m, \alpha)(1 + \delta)^{-d_{j+1}} \left\| f_{j+1}^{> d_{j+1}} \right\|_2 \leq C(m, \alpha)(1 + \delta)^{-d_{j+1}} \|f_{j+1}\|_2,$$

which is at most $C(m, M, \alpha)(1 + \delta)^{-d_{j+1}}$ as $\|f_{j+1}\|_2 \leq \|f_{j+1}\|_{2r} \leq M$. Here $\delta > 0$ depends only on α and m . Plugging this into (16) we get that (15) $\leq C(m, \alpha)(1 + \delta)^{-d_{j+1}} \|h\|_2$, and we next upper bound the norm of h . By Hölder's inequality we get that

$$\|h\|_2^2 \leq \mathbb{E} \left[\prod_{i=1}^j g_i(\mathbf{x}(i))^2 \prod_{i'=j+2}^r f_{i'}(\mathbf{x}(i'))^2 \right] \leq \prod_{i=1}^j \|g_i^2\|_{r-1} \prod_{i'=j+2}^r \|f_{i'}^2\|_{r-1} = \prod_{i=1}^j \|g_i\|_{2(r-1)}^2 \prod_{i'=j+2}^r \|f_{i'}\|_{2(r-1)}^2,$$

and we argue that $\|g_i\|_{2(r-1)}, \|f_{i'}\|_{2(r-1)} = O_{r,M,d_1,\dots,d_j}(1)$. Indeed, using monotonicity of the norms and then hypercontractivity (Theorem 2.12) we get

$$\|g_i\|_{2(r-1)} \leq \|g_i\|_{2r} \leq O_{r,d_1,\dots,d_j}(1) \|g_i\|_2 \leq O_{r,d_1,\dots,d_j}(1) M \leq O_{r,M,d_1,\dots,d_j}(1),$$

where we used Parseval to bound $\|g_i\|_2 \leq \|f_i\|_2 \leq \|f_i\|_{2r} \leq M$. For $f_{i'}$, we have $\|f_{i'}\|_{2(r-1)} \leq \|f_{i'}\|_{2r} \leq M$.

Plugging everything into (16) we get that (15) $\leq O_{\alpha,m,r,d_1,\dots,d_j}(1)(1+\delta)^{-d_{j+1}}$, so picking d_{j+1} large enough gives that (15) $\leq \frac{\varepsilon}{3r}$. $\square$

Claim 4.2.

$$\left| \mathbb{E}_{(\mathbf{x}(1),\dots,\mathbf{x}(r)) \sim \tilde{\mu}} \left[\prod_{i=1}^r T_{\mathcal{C}_i} f_i(\mathbf{x}(i)) \right] - \mathbb{E}_{(\mathbf{x}(1),\dots,\mathbf{x}(r)) \sim \tilde{\mu}} \left[\prod_{i=1}^r (T_{\mathcal{C}_i} f_i)^{\leq d_i}(\mathbf{x}(i)) \right] \right| \leq \frac{\varepsilon}{3}.$$

Proof. The proof is identical to the proof of Claim 4.1, except that we use Lemma B.3 instead of Lemma 3.11, and Theorem 2.11 instead of Theorem 2.12. $\square$

Claim 4.3.

$$\left| \mathbb{E}_{(\mathbf{x}(1),\dots,\mathbf{x}(r)) \sim \mu} \left[\prod_{i=1}^r f_i^{\leq d_i}(\mathbf{x}(i)) \right] - \mathbb{E}_{(\mathbf{x}(1),\dots,\mathbf{x}(r)) \sim \tilde{\mu}} \left[\prod_{i=1}^r (T_{\mathcal{C}_i} f_i)^{\leq d_i}(\mathbf{x}(i)) \right] \right| \leq \frac{\varepsilon}{3}.$$

Proof. Let $\mathcal{C}$ be a (α, ζ) -coupling between μ and $\tilde{\mu}$ guaranteed to exist by the assumption of the theorem. Applying the hybrid argument again, it is enough to show that for all $j = 0, \dots, r-1$, the absolute value of

$$\mathbb{E}_{\substack{(\mathbf{x}(1),\dots,\mathbf{x}(r)) \sim \mathcal{C} \\ (\mathbf{y}(1),\dots,\mathbf{y}(r)) \sim \mathcal{C}}} \left[\prod_{i=1}^j f_i^{\leq d_i}(\mathbf{x}(i)) \left(f_{j+1}^{\leq d_{j+1}}(\mathbf{x}(j+1)) - (T_{\mathcal{C}_{j+1}} f_{j+1})^{\leq d_{j+1}}(\mathbf{y}(j+1)) \right) \prod_{i=j+2}^r (T_{\mathcal{C}_i} f_i)^{\leq d_i}(\mathbf{y}(i)) \right] \quad (17)$$

is at most $\frac{\varepsilon}{3r}$. By Cauchy-Schwarz, this absolute value is at most

$$\|h\|_2 \sqrt{\mathbb{E}_{\substack{(\mathbf{x}(1),\dots,\mathbf{x}(r)) \sim \mathcal{C} \\ (\mathbf{y}(1),\dots,\mathbf{y}(r)) \sim \mathcal{C}}} \left[\left(f_{j+1}^{\leq d_{j+1}}(\mathbf{x}(j+1)) - (T_{\mathcal{C}_{j+1}} f_{j+1})^{\leq d_{j+1}}(\mathbf{y}(j+1)) \right)^2 \right]},$$

where $h = \prod_{i=1}^j f_i^{\leq d_i}(\mathbf{x}(i)) \prod_{i=j+2}^r (T_{\mathcal{C}_i} f_i)^{\leq d_i}(\mathbf{y}(i))$. The norm of h is bounded as in the previous two claims using hypercontractivity by $O_{r,M,m,\alpha,d_1,\dots,d_r}(1)$. For the other expectation, note first that by Lemma 3.25

$$\left\| (T_{\mathcal{C}_{j+1}} f_{j+1})^{\leq d_{j+1}} - T_{\mathcal{C}_{j+1}} (f_{j+1}^{\leq d_{j+1}}) \right\|_2 \leq O_{d_{j+1},\alpha}(\zeta^{1/4}),$$

so by $(a+b)^2 \leq 2(a^2 + b^2)$ the expectation is upper bounded by

$$O_{d_{j+1},\alpha}(\sqrt{\zeta}) + 2 \mathbb{E}_{\substack{(\mathbf{x}(1),\dots,\mathbf{x}(r)) \sim \mathcal{C} \\ (\mathbf{y}(1),\dots,\mathbf{y}(r)) \sim \mathcal{C}}} \left[\left(f_{j+1}^{\leq d_{j+1}}(\mathbf{x}(j+1)) - T_{\mathcal{C}_{j+1}} (f_{j+1}^{\leq d_{j+1}})(\mathbf{y}(j+1)) \right)^2 \right].$$

Finally, Lemma 1.6 implies that the last expectation is upper bounded by $8\sqrt{d_{j+1}\zeta}$. Thus, overall (17) is bounded by $O_{r,M,m,\alpha,d_1,\dots,d_r}(\zeta^{1/4})$, so taking ζ small enough gives it is at most $\frac{\varepsilon}{3r}$. Summing (17) over all j and using the triangle inequality finishes the proof. $\square$

Theorem 1.11 now follows from Claims 4.1, 4.2, 4.3 and the triangle inequality.

4.2 Invariance principle for label-assignments

In this section, we state a version of our invariance principle, Theorem 1.11, that will be useful in our PCP reduction. Let $\Sigma_i = [m_i]$ be alphabets for $i = 1, \dots, r$, and let $P: \prod_{i=1}^r \Sigma_i \rightarrow \{0, 1\}$ be a predicate. Let $f_i: \mathcal{U}_{\vec{k}(i)} \rightarrow [m_i]$ be functions where $\mathcal{U}_{\vec{k}(i)} \subseteq [m_i]^n$ are all α -balanced multislices. We will want to prove invariance for expressions of the form $\mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu} [P(f_1(\mathbf{x}(1)), \dots, f_r(\mathbf{x}(r)))]$. At the moment, it does not make any sense – the function $T_{C_i} f_i$ is not even well defined (as Σ does not have a sensible additive structure). We therefore view our, as is standard, the alphabets as standard simplices.

Let $\Delta_m = \{(t_0, \dots, t_{m-1}) | t_i \geq 0, \sum_{i=0}^{m-1} t_i = 1\}$ be the standard m -dimensional simplex. We identify a symbol $\sigma \in \Sigma_i$ with the standard basis element $e_\sigma \in \Delta_{m_i}$. This allows us to view a given function $f_i: \mathcal{U}_{\vec{k}(i)} \rightarrow [m_i]$ as a function $f_i: \mathcal{U}_{\vec{k}(i)} \rightarrow \Delta_{m_i}$, which we will do (slightly abusing notation by denoting the two functions the same way). Thus, now given a coupling C_i between $\mathcal{U}_{\vec{k}(i)}$ and $([m]^n, \nu_{\vec{k}(i)})$, we define

$$T_{C_i} f_i = (T_{C_i} f_{i,1}, \dots, T_{C_i} f_{i,m_i}),$$

where $f_{i,j}$ is the j th coordinate of f_i . We remark that clearly, if $f_i: \mathcal{U}_{\vec{k}(i)} \rightarrow \Delta_{m_i}$, then as T_{C_i} is an averaging operator and Δ_{m_i} is convex, we get that $T_{C_i} f_i: \mathcal{U}_{\vec{k}(i)} \rightarrow \Delta_{m_i}$.

Next, we define the natural extension of our predicate P , denoted by $\tilde{P}: \prod_{i=1}^r \Delta_{m_i} \rightarrow [0, 1]$. If for $i = 1, \dots, r$, each $y(i) \in \Delta_{m_i}$ is standard basis elements, say $y(i) = e_{\sigma_i}$, then we define $\tilde{P}(y(1), \dots, y(r)) = P(\sigma_1, \dots, \sigma_r)$. For the rest of the points, we extend $\tilde{P}$ multilinearly; let $y(i) \in \Delta_{m_i}$ for $i = 1, \dots, r$, and express each $y(i)$ using the standard basis as $y(i) = \sum_{\sigma \in \Sigma_i} y(i)_\sigma e_\sigma$. Then we define

$$\tilde{P}(y(1), \dots, y(r)) = \sum_{\sigma_1 \in \Sigma_1, \dots, \sigma_r \in \Sigma_r} P(\sigma_1, \dots, \sigma_r) \prod_{i=1}^r y(i)_{\sigma_i}.$$

Theorem 4.4. *For all $\alpha \in (0, 1)$, $r \in \mathbb{N}$, $m_1, \dots, m_r \in \mathbb{N}$, $\varepsilon > 0$ there are $\zeta > 0$, $N \in \mathbb{N}$ such that the following holds for $n \geq N$. Suppose $\mathcal{U}_{\vec{k}(1)}, \dots, \mathcal{U}_{\vec{k}(r)}$ are α -balanced multi-slices over alphabets $[m_1], \dots, [m_r]$ respectively, and μ is a connected, α -admissible distribution over $\prod_{i=1}^r \mathcal{U}_{\vec{k}(i)}$. Suppose for each $i \in [r]$, C_i is a (α, ζ) -coupling between $\mathcal{U}_{\vec{k}(i)}$ and $([m_i]^n, \nu_{\vec{k}(i)}^{\otimes n})$, and that there is a (α, ζ) -coupling between μ and $\tilde{\mu}$. Then for all $P: \prod_{i=1}^r [m_i] \rightarrow [0, 1]$ and $f_i: \mathcal{U}_{\vec{k}(i)} \rightarrow [m_i]$ it holds that*

$$\left| \mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu} [P(f_1(\mathbf{x}(1)), \dots, f_r(\mathbf{x}(r)))] - \mathbb{E}_{(\mathbf{y}(1), \dots, \mathbf{y}(r)) \sim \tilde{\mu}} [\tilde{P}(T_{C_1} f_1(\mathbf{y}(1)), \dots, T_{C_r} f_r(\mathbf{y}(r)))] \right| \leq \varepsilon.$$

Proof. Set $\varepsilon' = \varepsilon / (m_1 \cdots m_r)$, and pick δ, N from Theorem 1.11 for $\alpha, r, m_1, \dots, m_r$ and ε' . Then the left hand side is at most

$$\sum_{\sigma_1 \in [m_1], \dots, \sigma_r \in [m_r]} |P(\sigma_1, \dots, \sigma_r)| \left| \mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu} \left[\prod_{i=1}^r f_{i,\sigma_i}(\mathbf{x}(i)) \right] - \mathbb{E}_{(\mathbf{y}(1), \dots, \mathbf{y}(r)) \sim \tilde{\mu}} \left[\prod_{i=1}^r T_{C_i} f_{i,\sigma_i}(\mathbf{y}(i)) \right] \right|.$$

By Theorem 1.11, the second absolute value is at most ε' , hence we get the bound

$$\sum_{\sigma_1 \in [m_1], \dots, \sigma_r \in [m_r]} |P(\sigma_1, \dots, \sigma_r)| \varepsilon' \leq \varepsilon. \quad \square$$

4.3 Using our invariance principle: a construction of a useful coupling

In this section, we construct a useful coupling between admissible distributions over products of balanced multi-slices, and their product version.

Definition 4.5. For each $i = 1, \dots, r$, let $\mathcal{U}_{\vec{k}(1)}^-$ be a multi-slice over the alphabet $[m_i]$, and let μ be a distribution over $\mathcal{U}_{\vec{k}(1)}^- \times \dots \times \mathcal{U}_{\vec{k}(r)}^-$. We say μ is negatively correlated if for all $\vec{a} \in [m_1] \times \dots \times [m_r]$, defining the events $A_i = \{(x(1), \dots, x(r)) \mid (x(1)_i, \dots, x(r)_i) = \vec{a}\}$, the random variables 1_{A_i} for $i = 1, \dots, n$ are negatively associated (where the sampling of $(x(1), \dots, x(r))$ is done under μ).

Proposition 4.6. For all $\alpha > 0$, $r, m_1, \dots, m_r \in \mathbb{N}$ there are $\alpha', K > 0$ such that the following holds. Let μ be an α -admissible distribution, negatively correlated distribution over $\mathcal{U}_{\vec{k}(1)}^- \times \dots \times \mathcal{U}_{\vec{k}(r)}^-$, where for each i , the multi-slice $\mathcal{U}_{\vec{k}(i)}^-$ has alphabet $[m_i]$ and is α -balanced. Let $\tilde{\mu}$ be the product version of μ .

Then there is a $(\alpha', K \frac{1}{\sqrt{n}})$ -coupling between μ and $\tilde{\mu}$.

Proof. We define the coupling $\mathcal{C} = ((\mathbf{x}(1), \dots, \mathbf{x}(r)), (\mathbf{y}(1), \dots, \mathbf{y}(r)))$ as follows:

1. Sample a statistics of $(\mathbf{y}(1)', \dots, \mathbf{y}(r)') \sim \tilde{\mu}$. That is, sample $(\mathbf{y}(1)', \dots, \mathbf{y}(r)') \sim \tilde{\mu}$, and for each $\vec{a} = (a_1, \dots, a_r) \in [m_1] \times \dots \times [m_r]$, denote by $r_{\vec{a}}$ the number of coordinates $i \in [n]$ such that $\mathbf{y}(1)'_i = a_1, \dots, \mathbf{y}(r)'_i = a_r$.
2. Sample a statistics of $(\mathbf{x}(1)', \dots, \mathbf{x}(r)') \sim \mu$. That is, sample $(\mathbf{x}(1)', \dots, \mathbf{x}(r)') \sim \mu$, and for each $\vec{a} = (a_1, \dots, a_r) \in [m_1] \times \dots \times [m_r]$ denote by $k_{\vec{a}}$ the number of coordinates $i \in [n]$ such that $\mathbf{x}(1)'_i = a_1, \dots, \mathbf{x}(r)'_i = a_r$.
3. Let us consider the lexicographical ordering on $[m_1] \times \dots \times [m_r]$. Sample $\pi \in S_n$, and divide the integers $1, \dots, n$ into consecutive segments $(I_{\vec{a}})_{\vec{a} \in [m_1] \times \dots \times [m_r]}$ of lengths $k_{\vec{a}}$, according to the lexicographical ordering on $[m_1] \times \dots \times [m_r]$. Also, divide the integers $1, \dots, n$ into segments $(J_{\vec{a}})_{\vec{a} \in [m_1] \times \dots \times [m_r]}$ of lengths $r_{\vec{a}}$ again according to the lexicographical ordering on $[m_1] \times \dots \times [m_r]$.
4. Define $\mathbf{x}(1), \dots, \mathbf{x}(r)$ by setting $\mathbf{x}(1)_{\pi(\ell)} = a_1, \dots, \mathbf{x}(r)_{\pi(\ell)} = a_r$ for all $\vec{a} \in [m_1] \times \dots \times [m_r]$ and $\ell \in I_{\vec{a}}$. Also, define $\mathbf{y}(1), \dots, \mathbf{y}(r)$ by setting $\mathbf{y}(1)_{\pi(\ell)} = a_1, \dots, \mathbf{y}(r)_{\pi(\ell)} = a_r$ for all $\vec{a} \in [m_1] \times \dots \times [m_r]$ and $\ell \in J_{\vec{a}}$.

We now argue that $\mathcal{C}$ is a $(\Omega_{\alpha, r, m_1, \dots, m_r}(1), O_{\alpha, r, m_1, \dots, m_r}(1/\sqrt{n}))$ -coupling according to Definition 1.5. The first, second items are immediate, and we focus on the third and fourth items.

For the fourth item, let $R \subseteq [m_1] \times \dots \times [m_r]$ be the set of symbols that occur with positive probability in a coordinate of μ , and let $\vec{a} \in R$. Set $m = m_1 \cdots m_r$, $\varepsilon' = \frac{\varepsilon}{4m^3}$. Fix $\vec{a} \in R$, and note that for $i = 1, \dots, n$, the events $(\mathbf{x}(1)'_i, \dots, \mathbf{x}(r)'_i) = \vec{a}$ are negatively associated, so by Theorem 2.16 we have that $k_{\vec{a}} \leq n\mu_{\vec{a}} + \varepsilon'n$ with probability at least $1 - e^{-\Omega_m(\varepsilon^2 n)}$, so by the union bound $k_{\vec{a}} \leq n\mu_{\vec{a}} + \varepsilon'n$ for all $\vec{a} \in R$ with probability at least $1 - me^{-\Omega_m(\varepsilon^2 n)}$. We note that in this case, we have for all $\vec{a} \in R$ that

$$k_{\vec{a}} \geq n - \sum_{\substack{\vec{b} \in R \\ \vec{b} \neq \vec{a}}} k_{\vec{b}} \geq n - \sum_{\substack{\vec{b} \in R \\ \vec{b} \neq \vec{a}}} (n\mu_{\vec{b}} + \varepsilon'n),$$

which is at least $\mu_{\vec{a}}n - m\varepsilon'n$. We get that $|k_{\vec{a}} - \mu_{\vec{a}}n| \leq m\varepsilon'n$ for all $\vec{a} \in R$ with probability $1 - me^{-\Omega_m(\varepsilon^2n)}$. The same argument applies for $r_{\vec{a}}$ to show that $|r_{\vec{a}} - \mu_{\vec{a}}n| \leq m\varepsilon'n$ for all $\vec{a} \in R$ with probability $1 - me^{-\Omega_m(\varepsilon^2n)}$. Therefore, with probability $1 - me^{-\Omega_{\alpha,m}(\varepsilon^2n)}$ we have $|k_{\vec{a}} - r_{\vec{a}}| = 2m\varepsilon'n$, and we show that in this case $(\mathbf{x}(1), \dots, \mathbf{x}(r))$ and $(\mathbf{y}(1), \dots, \mathbf{y}(r))$ differ in at most εn coordinates. Indeed, note that

$$|I_{\vec{a}} \triangle J_{\vec{a}}| \leq \sum_{\vec{b} \preceq \vec{a}} 2 |k_{\vec{b}} - r_{\vec{b}}| \leq 4m^2\varepsilon'n,$$

and the set $B = \bigcup_{\vec{a} \in R} I_{\vec{a}} \triangle J_{\vec{a}}$ contains all i 's in which $(\mathbf{x}(1)_i, \dots, \mathbf{x}(r)_i) \neq (\mathbf{y}(1)_i, \dots, \mathbf{y}(r)_i)$. Thus, $|B| \leq 4m^3\varepsilon'n \leq \varepsilon n$, and the fourth item is established.

For the third item, again we see that the number of coordinates (as a random variable) is at most $2m \sum_{\vec{b} \preceq \vec{a}} |k_{\vec{b}} - r_{\vec{b}}|$, so its expectation is at mo

$$2m \sum_{\vec{b} \preceq \vec{a}} \mathbb{E}[|k_{\vec{b}} - r_{\vec{b}}|] \leq 2m \sum_{\vec{b} \preceq \vec{a}} \sqrt{\mathbb{E}[|k_{\vec{b}} - r_{\vec{b}}|^2]} \leq 2m \sum_{\vec{b} \preceq \vec{a}} O_m(\sqrt{n}) = O_m(\sqrt{n}).$$

The third item thus follows from symmetry. $\square$

5 Beyond connected distributions?

In this section, we remark that one can prove the invariance principle from the multi-slice to its product version continues to hold for a class of distributions μ in which the “contribution of high-degree functions” is always small. For such distributions, one only has to care about the contribution of the low-degree parts of the functions, and for we use Lemma 1.6. In this language, the main content of the previous section is showing that in admissible, connected distributions, the contribution of high-degree functions is very small, i.e. Lemma 3.11.

Definition 5.1. Let $r, m \in \mathbb{N}$ and let $([m]^n, \nu_{1,n}), \dots, ([m]^n, \nu_{r,n})$ be measure spaces for all $n \in \mathbb{N}$. Let $\mu = \{\mu_n\}_{n \in \mathbb{N}}$ be a sequence of distribution over $(\prod_{i=1}^r [m]^n)^n$ whose marginals match $\nu_{i,n}$. We say μ annihilates high-degree functions, if for all $\varepsilon > 0$, $M \in \mathbb{N}$, there are $q, d \in \mathbb{N}$ and $N \in \mathbb{N}$ such that the following holds for all $n \geq N$.

Whenever $f_1, \dots, f_r: [m]^n \rightarrow \mathbb{R}$ are functions such that:

1. for all $i = 1, \dots, r$ it holds that $\mathbb{E}_{\mathbf{x} \sim \nu_{i,n}} [f_i(\mathbf{x})^q] \leq M$, and
2. there is $i = 1, \dots, r$ such that $f_i \in V_{>d}([m]^n, \nu_{i,n})$,

it holds that $\left| \mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu_n} \left[\prod_{i=1}^r f_i(\mathbf{x}(i)) \right] \right| \leq \varepsilon$.

Thus one gets the following result.

Theorem 5.2. Let $\alpha > 0$, $r, m_1, \dots, m_r \in \mathbb{N}$. For each $n \in \mathbb{N}$, let $\mathcal{U}_{\vec{k}(1)}^n, \dots, \mathcal{U}_{\vec{k}(r)}^n$ be α -balanced multi-slices over alphabets $[m_1], \dots, [m_r]$ of vectors of length n , respectively, let $\mu = \{\mu_n\}$ be a sequence of measures on $\prod_{i=1}^r \mathcal{U}_{\vec{k}(i)}^n$ whose marginal on each i is uniform on the respective multi-slice, and let $\tilde{\mu} = \{\tilde{\mu}_n\}_{n \in \mathbb{N}}$ be the product version of μ .

Suppose, that μ and $\tilde{\mu}$ both annihilate high-degree functions. Then for all $\varepsilon > 0$ there are $\zeta > 0$, $N \in \mathbb{N}$ such that the following hold for $n \geq N$. If $\mathcal{C}_1, \dots, \mathcal{C}_r$ are couplings, where $\mathcal{C}_i$ is a (α, ζ) -coupling between $\mathcal{U}_{k(i)}^n$ and $\nu_{k(i)}^{\otimes n}$, and there is a (α, ζ) -coupling between μ_n and $\tilde{\mu}_n$, then for all $f_i: \mathcal{U}_{k(i)} \rightarrow [-1, 1]$ we have

$$\left| \mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu_n} \left[\prod_{i=1}^r f_i(\mathbf{x}(i)) \right] - \mathbb{E}_{(\mathbf{y}(1), \dots, \mathbf{y}(r)) \sim \tilde{\mu}_n} \left[\prod_{i=1}^r \mathbb{E}_{\mathcal{C}_i} f_i(\mathbf{y}(i)) \right] \right| \leq \varepsilon.$$

We omit the proof since it follows easily from the argument in the previous section by replacing the use of Lemmas 3.11 and Lemma B.3 with the property that μ and $\tilde{\mu}$ annihilate high degree functions. We believe it would be interesting to identify new classes of distributions that annihilate high degrees. We also remark that one immediately gets an analogue of Theorem 4.4 for distributions μ such that both μ and $\tilde{\mu}$ annihilates high-degree functions.

6 Applications to hardness of approximation

In this section, we use our invariance principle in order to convert dictatorship tests into NP-hardness results assuming the Rich 2-to-1 Games Conjecture. We begin with some definitions to formally set up terminology.

Definition 6.1. Let Σ be a finite alphabet, $\mathcal{P} \subseteq \{P: \Sigma^r \rightarrow \{0, 1\}\}$ be a collection of predicates. An instance of CSP- $\mathcal{P}$ is (Z, E) where $Z = \{z_1, \dots, z_n\}$ is a set of variables, and E is the set of constraints of the form $P(z_{i_1}, \dots, z_{i_r}) = 1$ where $P \in \mathcal{P}$.

Given an instance $\Psi = (Z, E)$ of CSP- $\mathcal{P}$, we define the value of Ψ as follows. An assignment is a mapping $A: Z \rightarrow \Sigma$, and the value of the assignment $\text{val}_\Psi(A)$ is the fraction of constraints that are satisfied by it, i.e. $\mathbb{E}_{e \in E} [1_{A \text{ satisfies } e}]$. The value of the instance, $\text{val}(\Psi)$, is defined to be $\max_{A: Z \rightarrow \Sigma} \text{val}_\Psi(A)$.

Definition 6.2. Let $0 < s < c \leq 1$, $r, m \in \mathbb{N}$ and let $\mathcal{P} \subseteq \{P: [m]^r \rightarrow \{0, 1\}\}$ be a collection of r -ary predicates over alphabet $[m]$. We denote by Gap-CSP- $\mathcal{P}[c, s]$ the promise problem where one has to distinguish between the following two cases, given an instance Ψ of CSP- $\mathcal{P}$:

1. **YES case:** $\text{val}(\Psi) \geq c$;
2. **NO case:** $\text{val}(\Psi) \leq s$.

Next, we wish to define dictatorship tests, and for that we first define regular functions.

Definition 6.3. Let $\Sigma = [m]$, let $\mathcal{D} = \nu_1 \times \dots \times \nu_n$ be a distribution over Σ^n , and let $f: \Sigma^n \rightarrow \mathbb{R}^m$. The influence of variable $i \in [n]$ is defined as

$$I_i[f; \mathcal{D}] = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\|f(\mathbf{x}) - \mathbb{E}_i f(\mathbf{x})\|_2^2 \right].$$

Here, $\mathbb{E}_i$ is the averaging operator over coordinate i defined as $\mathbb{E}_i f(\mathbf{x}) = \mathbb{E}_{\mathbf{y}_i \sim \nu_i} [f(\mathbf{x}_{-i}, \mathbf{y}_i)]$.

Note that denoting $f: \Sigma^n \rightarrow \mathbb{R}^m$ as $f = (f_1, \dots, f_m)$ for real-valued f_ℓ 's, we have $I_i[f] = \sum_{j \in [m]} I_i[f_j]$.

Definition 6.4. Let $\Sigma = [m]$, let $\mathcal{D} = \nu_1 \times \dots \times \nu_n$ be a distribution over Σ^n , $d \in \mathbb{N}$, and let $f: \Sigma^n \rightarrow \mathbb{R}^m$. The degree d influence of variable $i \in [n]$ is $I_i^{\leq d}[f; \mathcal{D}] = I_i[f^{\leq d}; \mathcal{D}]$. Here, $f^{\leq d} = (f_1^{\leq d}, \dots, f_m^{\leq d})$.

Definition 6.5. Let $\Sigma = [m]$ and $\mathcal{D}$ be a distribution over Σ^r with marginals $\mathcal{D}_1, \dots, \mathcal{D}_r$, and let $f: \Sigma^n \rightarrow \Sigma$. As in Section 4.2, view f as a function from Σ^n to Δ_m . We say f is (d, τ) -regular with respect to $\mathcal{D}$, if for all $i \in [r]$ and $j \in [n]$ we have that $I_j^{\leq d}[f; \mathcal{D}_i^{\otimes n}] \leq \tau$.

For a collection of distributions $\{\mathcal{D}(t)\}_{t \in \mathcal{T}}$ and a measure $p: \mathcal{T} \rightarrow [0, 1]$, we say $f: \Sigma^n \rightarrow \Sigma$ is (γ, d, τ) -regular with respect to $\mathcal{T}$ if choosing $\mathbf{t} \sim \mathcal{T}$, with probability at least $1 - \gamma$, we have that f is (d, τ) -regular with respect to $\mathcal{D}(\mathbf{t})$.

Definition 6.6. Let $r \in \mathbb{N}$, $0 < s < c \leq 1$, and $\mathcal{P} \subseteq \{P: \Sigma^r \rightarrow \{0, 1\}\}$ be a collection of r -ary predicates. A (c, s) -dictatorship test for $\mathcal{P}$ consists of a probability measure $w: \mathcal{P} \rightarrow [0, 1]$, a collection of distributions $\{\mathcal{D}(t, P)\}_{t \in \mathcal{T}, P \in \mathcal{P}}$ over Σ^r , and a measure $p: \mathcal{T} \rightarrow [0, 1]$ such that

1. *Completeness:* for all n , if $f: \Sigma^n \rightarrow \Sigma$ is a dictatorship, i.e. $f(x) = x_i$ for some i , then

$$\mathbb{E}_{\substack{\mathbf{t} \sim p_{\mathcal{T}} \\ \mathbf{P} \sim w}} \left[\mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mathcal{D}(\mathbf{t}, \mathbf{P})^{\otimes n}} [\mathbf{P}(f(\mathbf{x}(1)), \dots, f(\mathbf{x}(r)))] \right] \geq c.$$

2. *Soundness:* for all $\gamma, \varepsilon > 0$, there are $\tau > 0$ and $d, N \in \mathbb{N}$ such that the following holds. If $n \geq N$, and $f: \Sigma^n \rightarrow \Sigma$ is a (γ, d, τ) -regular function with respect to $\{\mathcal{D}(t)\}_{t \in \mathcal{T}}$, then

$$\mathbb{E}_{\substack{\mathbf{t} \sim p_{\mathcal{T}} \\ \mathbf{P} \sim w}} \left[\mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mathcal{D}(\mathbf{t}, \mathbf{P})^{\otimes n}} [\mathbf{P}(f(\mathbf{x}(1)), \dots, f(\mathbf{x}(r)))] \right] \leq s + \varepsilon.$$

Raghavendra [44] showed that a given (c, s) -dictatorship test for a predicate can be converted into a $(c - \varepsilon, s + \varepsilon)$ NP-hardness result for the predicate, assuming the Unique-Games Conjecture. When $c < 1$, this result is easily seen to be equivalent to a $(c, s + \varepsilon)$ NP-hardness result, however the situation is very different for the case of perfect completeness, i.e. $c = 1$, which we focus on henceforth. In this case, the loss of perfect completeness in Raghavendra's result stems from two different places, each one of which seems fundamental. First, in order to carry out his analysis, Raghavendra has to work with connected distributions μ ; this is easy to achieve if one allows a small change in the completeness parameter c , by mixing each one of the distributions μ_t with a small multiple of the uniform distribution (i.e. working with the distribution $\mu'_t(x) = (1 - \varepsilon)\mu_t(x) + \frac{\varepsilon}{|\Sigma|^r}$). Second, Raghavendra starts off the reduction from the Unique-Games problem, which inherently has to have imperfect completeness.

Roughly speaking, our result shows that if one is willing to assume the stronger Rich 2-to-1 Games Conjecture, then only the first of these issues persist. More precisely, we show:

Theorem 6.7. Let $r \in \mathbb{N}$, $m \in \mathbb{N}$, and suppose $\mathcal{P} \subseteq \{P: [m]^r \rightarrow \{0, 1\}\}$ is a collection of predicates. Suppose there is a $(1, s)$ dictatorship test, $(w, \{\mathcal{D}(t)\}_{t \in \mathcal{T}}, p)$, for $\mathcal{P}$. If each one $\mathcal{D}(t)$ is connected, then assuming Conjecture 1.17, for all $\varepsilon > 0$, the problem $\text{Gap-}\mathcal{P}[1, s + \varepsilon]$ is NP-hard.

The rest of this section is devoted to proving Theorem 6.7. Our reduction is very similar to Raghavendra's reduction; however, we do not know how to analyze the most natural version of it using product spaces. As is natural in many PCP reductions, the constraint maps of our starting initial Rich-2-to-1 Games instance define a projection operator from the "large long-code" to the "small long-code". The issue is that unlike in Unique-Games, these projection operators do not behave well with respect to the measure on these spaces, and this introduces several technical difficulties.

Our idea therefore is to use multi-slices instead of product spaces, for which we show that the above issues no longer exist. As is often the case, the completeness of this reduction is trivial, and one only has to

worry about the soundness of the test. This is where our invariance principle enters the picture: intuitively, it asserts that the performance of our dictatorship test, when applied on multi-slices, is very close to its performance of the original dictatorship test, which we know how to analyze. There are still technical issues that arise because of multi-slices and the fact we are working with 2-to-1 constraints (as opposed to 1-to-1 constraints), mainly in the decoding phase. For that, we define an appropriate notion of “noisy influences” that has similar properties to the standard noisy influences (e.g. the total noisy influence is small), and also interacts well with 2-to-1 projections and couplings.

6.1 An appropriate notion of noisy influences

Definition 6.8. Suppose $0 < \beta \leq \alpha < 1$, and $n \in \mathbb{N}$. Suppose $\mathcal{U}_{\vec{k}} \subseteq [m]^n$ is a α -balanced multi-slice, and βn is an integer. For each $x \in \mathcal{U}_{\vec{k}}$, define the distribution $\mathbf{y} \sim S_{1-\beta}x$ over $\mathcal{U}_{\vec{k}}$ as:

1. for each $j \in [m]$, choose $\mathbf{A}_j \subseteq \{i \in [n] \mid x_i = j\}$ of size βn randomly;
2. take $\mathbf{y} \in \mathcal{U}_{\vec{k}}$, where $\mathbf{y}_i = x_i$ if $i \notin \bigcup_{j \in [m]} \mathbf{A}_j$, $\mathbf{y}_i = j + 1$ if $i \in \mathbf{A}_j$, $j \neq m$ and $\mathbf{y}_i = 1$ if $i \in \mathbf{A}_m$.

We will assume henceforth that $\beta n/2$ is an integer, and the point is so that the operator $S_{1-\beta}$ is well defined on $\mathcal{U}_{\vec{k}}$ and $\mathcal{U}_{\vec{k}/2}$. We remark that by adjusting the definition of $S_{1-\beta}$ slightly, the statements below remain true without this assumption, however this only introduces more cumbersome notations.

Definition 6.9. Suppose $0 < \beta \leq \alpha < 1$, and $n \in \mathbb{N}$. Suppose $\mathcal{U}_{\vec{k}} \subseteq [m]^n$ is a α -balanced multi-slice, and βn is an integer. The noise operator $S_{1-\beta}: L^2(\mathcal{U}_{\vec{k}}) \rightarrow L^2(\mathcal{U}_{\vec{k}})$ is defined as

$$(S_{1-\beta}f)(x) = \mathbb{E}_{\mathbf{y} \sim S_{1-\beta}x} [f(\mathbf{y})].$$

We will also define influences of functions on the multi-slice.

Definition 6.10. Let $\mathcal{U}_{\vec{k}} \subseteq [m]^n$, $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$, and $i \in [n]$. The influence of coordinate i is $I_i[f] = \mathbb{E}_{j \in [n]} \left[|f(\pi_{i,j}x) - f(x)|^2 \right]$, where $\pi_{i,j}$ is the transposition permutation between i and j . The total influence of f is $I[f] = \sum_{i=1}^n I_i[f]$.

Definition 6.11. Suppose $0 < \beta \leq \alpha < 1$, and $n \in \mathbb{N}$. Suppose $\mathcal{U}_{\vec{k}} \subseteq [m]^n$ is a α -balanced multi-slice, and βn is an integer. For $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$, the β -noisy influence of a variable $i \in [n]$ is defined as $I_i^{(\beta)}[f] = I_i[S_{1-\beta}f]$. The total noisy influence of f is $I^{(\beta)}[f] = \sum_{i=1}^n I_i^{(\beta)}[f]$.

We next state several properties of noisy influences in multi-slices, mostly analogous to properties of noisy influences in product spaces, which will be crucial for us in the proof of Theorem 6.7. Namely:

1. Lemma 6.13: the total β -noisy influence of a bounded function is $O_\beta(1)$ (so the number of coordinates with β -noisy influence at least τ is $O_\beta(1/\tau)$).
2. Lemma 6.16: if $f: \mathcal{U}_{\vec{k}} \rightarrow [-1, 1]$ is a function in which all noisy influences are small, and $\mathcal{C}$ is a good enough coupling between $(\mathcal{U}_{\vec{k}}, \text{Uniform})$ and $([m]^n, \nu_{\vec{k}}^{\otimes n})$, then all of the low-degree influences of $T_{\mathcal{C}}f$ are small.

3. Lemma 6.21: in Definition 6.20 we consider the notion of projections of a function over the multi-slice, which plays a crucial role in our analysis. Given a vector $\vec{k} = (k_1, \dots, k_m)$ of even integers summing up to n and a 2-to-1 map $\pi: [n] \rightarrow [n/2]$, we define the projection of $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ along π as $f|_{\pi}: \mathcal{U}_{\vec{k}/2} \rightarrow \mathbb{R}$, given by $f|_{\pi}(x) = f(y)$ where y is the point in which $y_i = x_{\pi(i)}$. In Lemma 6.21 we prove that if f has a coordinate i with significant noisy influence, then for a random projection π , the coordinate $\pi(i)$ has significant β -noisy influence in $f|_{\pi}$ with significant probability.

6.1.1 The total noisy influence is constant

In this section we show that the sum of noisy influences of a function over the multi-slice is constant. We begin with the following variant of the well-known connection between noisy influences and the Fourier transform.

Claim 6.12. *For all $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ we have that $I[f] \leq \sum_{d>0} d \|f^{=d}\|_2^2$.*

Proof. We recall the mapping $f \rightarrow \tilde{f}$ from Section 2.2.2 that given a function over the multi-slice produces a function $\tilde{f}: S_n \rightarrow \mathbb{R}$. We note that as we have already observed, this mapping is symmetric with respect to the action of S_n on function, so it preserves degrees, i.e. sends $f^{=d}$ to $\tilde{f}^{=d}$. This mapping also clearly preserves norms, so $\|\tilde{f}^{=d}\|_2 = \|f^{=d}\|_2$.

Our claim will now easily follow from the results of [24]. To be more precise, defining the influence of $g: S_n \rightarrow \mathbb{R}$ on a transposition $\pi_{i,j} \in [n]$ as $I_{i,j}[g] = \mathbb{E}_{\pi} [(g(\pi^{i,j}\pi) - g(\pi))^2]$, we see that $I_i[f] = \frac{1}{n} \sum_{j=1}^n I_{i,j}[\tilde{f}]$. Thus, taking $I[g] = \sum_{i=1}^n I_i[g]$, and using [24, inequality (32)] (adapting the definition of Laplacians and d_{λ} from there), we get that

$$I[f] = \frac{1}{n} I[\tilde{f}] = \frac{\binom{n}{2}}{n} \langle \tilde{f}, L\tilde{f} \rangle \leq \frac{n-1}{2} \sum_{\lambda \vdash n} d_{\lambda} \|\tilde{f}^{\lambda}\|_2^2.$$

We now note that by [24, Corollary], one has that $d_{\lambda} \leq \frac{2}{n-1}(n - \lambda_1)$, so we get that

$$I[f] \leq \sum_{\lambda \vdash n} (n - \lambda_1) \|\tilde{f}^{\lambda}\|_2^2 = \sum_{d>0} d \sum_{\substack{\lambda \vdash n \\ \lambda_1 = n-d}} \|\tilde{f}^{\lambda}\|_2^2 = \sum_{d>0} d \|\tilde{f}^{=d}\|_2^2 = \sum_{d>0} d \|f^{=d}\|_2^2. \quad \square$$

Lemma 6.13. *Suppose $m \in \mathbb{N}$, $0 < \beta \leq \alpha < 1$ and $n \in \mathbb{N}$ are such that βn is an integer. If $\mathcal{U}_{\vec{k}} \subseteq [m]^n$ is a α -balanced multi-slice, then for all $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ we have $I^{(\beta)}[f] \leq O_{m,\beta}(\|f\|_2^2)$.*

Proof. From Claim 6.12

$$I^{(\beta)}[f] = I[S_{1-\beta}f] \leq \sum_{d>0} d \|(S_{1-\beta}f)^{=d}\|_2^2. \quad (18)$$

By symmetry of $S_{1-\beta}$, we by Claim 3.17 that it preserves degrees, so $(S_{1-\beta}f)^{=d} = S_{1-\beta}(f^{=d})$. Note that the operator $S_{1-\beta}$ is the same as the operator $T_{\mu_{\beta}}$ where μ_{β} is the distribution $(\mathbf{x}, S_{1-\beta}\mathbf{x})$, which is β -admissible and connected. Also, μ_{β} is negatively correlated as per Definition 4.5 (see Remark 2.15), so by Proposition 4.6 there is a $(\Omega_{m,\beta}(1), O_{m,\beta}(1/\sqrt{n}))$ -coupling between μ_{β} and $\tilde{\mu}_{\beta}$. Thus from Lemma 3.11 we get $\|S_{1-\beta}(f^{=d})\| \leq C(m, \beta)(1 + \delta(m, \beta))^{-d} \|f^{=d}\|_2 \leq C(m, \beta)(1 + \delta(m, \beta))^{-d} \|f\|_2$ for some $C(m, \beta), \delta(m, \beta) > 0$. Plugging this into (18) we get

$$I^{(\beta)}[f] \leq C \|f\|_2^2 \sum_{d>0} d(1 + \delta(m, \beta))^{-d} = O_{m,\beta}(\|f\|_2^2). \quad \square$$

For technical reasons, we will also need the following claim. In spirit, it is close to Lemma 6.13 since $T_C S_{1-\beta}$ may be thought of as a noise operator. The bounds we get are however much worse (and are very likely not optimal), but are still good enough for our purposes.

Claim 6.14. *Suppose $0 < \beta \leq \alpha < 1$, and $n \in \mathbb{N}$ such that βn is an integer. Suppose $\mathcal{U}_{\vec{k}} \subseteq [m]^n$ is a α -balanced multi-slice, and $\mathcal{C}$ is a (α, ζ) -coupling between $\mathcal{U}_{\vec{k}}$ and $\nu_{\vec{k}}$. Then for all $f: \mathcal{U}_{\vec{k}} \rightarrow [-1, 1]$ we have that $I[T_C S_{1-\beta} f] \leq O_{m, \beta, \alpha}(\log^2 n) + O_{m, \beta, \alpha}(\zeta n \log^2 n)$.*

Proof. Take $D = M(m, \beta) \log n$ for sufficiently large M depending on m, β . We will use several basic notions from Section B.2. By Fact B.5

$$I[T_C S_{1-\beta} f] = \langle T_C S_{1-\beta} f, L T_C S_{1-\beta} f \rangle = \langle S_{1-\beta}^* T_C^* L T_C S_{1-\beta} f, f \rangle.$$

Each one of the operators $S_{1-\beta}$, $S_{1-\beta}^*$, T_C , T_C^* commutes with the action of S_n on functions, and it is also easy to check that L commutes with the action of S_n . Thus, $R = S_{1-\beta}^* T_C^* L T_C S_{1-\beta}$ commutes with the action of S_n , and therefore by Claim 3.17 it preserves the spaces $V_{=d}(\mathcal{U}_{\vec{k}})$, so we may continue the last chain of equalities as

$$I[T_C S_{1-\beta} f] = \sum_{d=0}^n \langle (Rf)^{=d}, f^{=d} \rangle = \sum_{d=0}^n \langle R(f^{=d}), f^{=d} \rangle.$$

We bound terms corresponding to $d \geq D$ and $d < D$ separately.

Contribution from $d > D$. By Cauchy-Schwarz and Jensen

$$\sum_{d=D}^n \langle R(f^{=d}), f^{=d} \rangle \leq \sum_{d=D}^n \|R(f^{=d})\|_2 \|f^{=d}\|_2 \leq \sum_{d=D}^n \|L T_C S_{1-\beta}(f^{=d})\|_2 \|f^{=d}\|_2,$$

as $\|Lg\|_2 \leq n \|g\|_2$ (which follows for example from Fact B.7) we get that the last sum is upper bounded by

$$n \sum_{d=D}^n \|T_C S_{1-\beta}(f^{=d})\|_2 \|f^{=d}\|_2 \leq n \sum_{d=D}^n \|S_{1-\beta}(f^{=d})\|_2 \|f^{=d}\|_2. \text{ By Lemma 3.11, we have}$$

$$\|S_{1-\beta}(f^{=d})\|_2 \leq O_{m, \beta}(1)(1 + \Omega_{m, \beta}(1))^d \|f^{=d}\|_2,$$

so we get that the above sum is at most

$$n \cdot O_{m, \beta}(1)(1 + \Omega_{m, \beta}(1))^D \sum_{d=D}^n \|f^{=d}\|_2^2 = n O_{m, \beta}(1)(1 + \Omega_{m, \beta}(1))^D \|f\|_2^2 = O_{m, \beta}(1)$$

by the choice of D .

Contribution from $d \leq D$. By definition of R

$$\sum_{d \leq D} \langle R(f^{=d}), f^{=d} \rangle = \sum_{d \leq D} \langle L T_C S_{1-\beta}(f^{=d}), T_C S_{1-\beta}(f^{=d}) \rangle = \sum_{d \leq D} I[T_C S_{1-\beta}(f^{=d})]. \quad (19)$$

We now bound $I[T_C S_{1-\beta}(f^{=d})]$ for each d separately. Let $g_d = S_{1-\beta}(f^{=d})$, and note that $g_d \in V_d(\mathcal{U}_{\vec{k}})$ by Claim 3.6 as $S_{1-\beta}$ commute with the action of S_n . Thus, by Claim 6.15 we get that $I[T_C g_d] \leq \left(\frac{d^2}{\alpha^2} \zeta n + d^2\right) \|g_d\|_2^2$, and plugging this into (19) gives

$$(19) \leq \left(\frac{D^2}{\alpha^2} \zeta n + D^2\right) \sum_{d \leq D} \|g_d\|_2^2 = \left(\frac{D^2}{\alpha^2} \zeta n + D^2\right) \|S_{1-\beta}(f^{\leq D})\|_2^2 \leq \left(\frac{4D^2}{\alpha^2} \zeta n + D^2\right). \quad \square$$

Claim 6.15. Let $\mathcal{U}_{\vec{k}} \subseteq [m]^n$ be a α -balanced multi-slice, and let $\mathcal{C}$ be a (α, ζ) -coupling between $\mathcal{U}_{\vec{k}}$ and $([m]^n, \nu_{\vec{k}}^{\otimes n})$. Then for all $g \in V_d(\mathcal{U}_{\vec{k}})$, we have $I[\mathcal{T}_{\mathcal{C}}g] \leq \left(\frac{4d^2}{\alpha^2}\zeta n + d^2\right) \|g\|_2^2$.

Proof. Deferred to Section C. □

6.1.2 Relating noisy influences of f and low-degree influences of $\mathcal{T}_{\mathcal{C}}f$

The goal of this section is to show that if $\mathcal{C}$ is a coupling between our multi-slice and the product domain, then if a function f has all of its noisy influences small, then all low-degree influences of $\mathcal{T}_{\mathcal{C}}f$ are small. To be more precise, we prove:

Lemma 6.16. For all $\tau > 0$, $\alpha > 0$, $m, d \in \mathbb{N}$ there are $\beta_0, \tau' > 0$, such that for all $0 < \beta \leq \beta_0$ there is $N \in \mathbb{N}$ such that the following holds for all $n \geq N$, $\zeta > 0$ such that $\zeta \leq \frac{1}{\log^3 n}$.

Let $\mathcal{U}_{\vec{k}} \subseteq [m]^n$ be a α -balanced mutislice, and let $\mathcal{C}$ be a (α, ζ) -coupling between $\mathcal{U}_{\vec{k}}$ and $([m]^n, \nu_{\vec{k}}^{\otimes n})$. Then for all $f: \mathcal{U}_{\vec{k}} \rightarrow [-1, 1]$, if $\max_i I_i^{(\beta)}[f] \leq \tau'$, then $\max_i I_i^{\leq d}[\mathcal{T}_{\mathcal{C}}f] \leq \tau$.

The proof of the above lemma relies on two auxiliary claims. The first asserts that if all noisy influences of f are small, then $I_i[\mathcal{T}_{\mathcal{C}}S_{1-\beta}f]$ are all small.

Claim 6.17. Let $\mathcal{U}_{\vec{k}} \subseteq [m]^n$ be α -balanced, $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$, and let $\mathcal{C}$ be a (α, ζ) -coupling between $\mathcal{U}_{\vec{k}}$ and $([m]^n, \nu_{\vec{k}}^{\otimes n})$. Then for all $i \in [n]$ we have that $I_i[\mathcal{T}_{\mathcal{C}}f] \leq \frac{2}{\alpha} I_i[f] + \frac{2}{\alpha^2 n} I[\mathcal{T}_{\mathcal{C}}f]$.

Proof. First, consider the following definition of influences over product spaces, which will be easier for us to establish the claim on. For $g: ([m]^n, \nu_{\vec{k}}^{\otimes n}) \rightarrow \mathbb{R}$ define $Q_i(g) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}, \mathbf{j} \in [n]} [(g(\pi_{i,\mathbf{j}}\mathbf{x}) - g(\mathbf{x}))^2]$.

Sample $(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}$, and $\mathbf{j} \in [n]$. Note that by the symmetry of $\mathcal{C}$, the distribution of $\pi_{i,\mathbf{j}}\mathbf{x}$ is the same as of $\mathbf{x}'$ where $(\mathbf{x}', \mathbf{y}') \sim \mathcal{C}$ conditioned on $\mathbf{y}' = \pi_{i,\mathbf{j}}\mathbf{y}$. Thus

$$Q_i[\mathcal{T}_{\mathcal{C}}f] = \mathbb{E}_{\mathbf{y}, \mathbf{j}} [|\mathcal{T}_{\mathcal{C}}f(\pi_{i,\mathbf{j}}\mathbf{y}) - \mathcal{T}_{\mathcal{C}}f(\mathbf{y})|^2] = \mathbb{E}_{\mathbf{y}, \mathbf{j}} \left[\left| \mathbb{E}_{\mathbf{x}} [f(\pi_{i,\mathbf{j}}\mathbf{x}) - f(\mathbf{x}) \mid \mathbf{y}, \mathbf{j}] \right|^2 \right] \leq \mathbb{E}_{\mathbf{x}} [|f(\pi_{i,\mathbf{j}}\mathbf{x}) - f(\mathbf{x})|^2],$$

which is precisely $I_i[f]$.

We now argue that $Q_i(g) \geq \frac{\alpha}{2} I_i[g] - \frac{1}{\alpha n} I[g]$ to finish the proof (taking $g = \mathcal{T}_{\mathcal{C}}f$). Sample $\mathbf{x} \sim \nu_{\vec{k}}^{\otimes n}$ and $\mathbf{j} \in [n]$. Consider the point $\mathbf{z}$ which is the same as $\mathbf{x}$ on all coordinates except i , where it is equal to the $\mathbf{j}$ coordinate of $\mathbf{x}$. Then writing $(g(\pi_{i,\mathbf{j}}\mathbf{x}) - g(\mathbf{x})) = (g(\pi_{i,\mathbf{j}}\mathbf{x}) - g(\mathbf{z})) + (g(\mathbf{z}) - g(\mathbf{x}))$ and using $(a + b)^2 \geq \frac{1}{2}a^2 - b^2$, we get that

$$Q_i(g) \geq \frac{1}{2} \mathbb{E}_{\mathbf{x}, \mathbf{j}} [(g(\mathbf{z}) - g(\mathbf{x}))^2] - \mathbb{E}_{\mathbf{x}, \mathbf{j}} [(g(\pi_{i,\mathbf{j}}\mathbf{x}) - g(\mathbf{z}))^2].$$

For the first expectation, note that the points $\mathbf{z}, \mathbf{x}$ are equal on all coordinates except the i th coordinate and distributed according to $\nu_{\vec{k}}$ on these coordinates. As $\nu_{\vec{k}}(a) \geq \alpha$ for all $a \in [m]$ it follows that $\mathbb{E}_{\mathbf{x}, \mathbf{j}} [(g(\mathbf{z}) - g(\mathbf{x}))^2] \geq \alpha I_i[g]$. For the second expectation, conditioning on $\mathbf{j} = j$ we have by the same argument that $\mathbb{E}_{\mathbf{x}} [(g(\pi_{i,j}\mathbf{x}) - g(\mathbf{z}))^2] \leq \frac{1}{\alpha} I_j[g]$, and we are done. □

The second claim directly relates $\max_i I_i[\mathcal{T}_{\mathcal{C}}S_{1-\beta}f]$ to the low-degree influences of $\mathcal{T}_{\mathcal{C}}f$, showing that if the former are all small, then the latter are also small.

Claim 6.18. For all $\tau > 0$, $\alpha > 0$, $d \in \mathbb{N}$ there are $\beta_0, \tau' > 0$, such that for all $0 < \beta \leq \beta_0$ there are $N \in \mathbb{N}$, $\zeta > 0$ such that the following holds for $n \geq N$. Let $\mathcal{U}_k \subseteq [m]^n$ be a α -balanced mutlislice, and suppose $\mathcal{C}$ is a (α, ζ) -coupling between $(\mathcal{U}_k, \text{Uniform})$, and $([m]^n, \nu_k^{\otimes n})$. Then for all $f: \mathcal{U}_k \rightarrow [-1, 1]$, if $\max_i I_i[\mathcal{T}_C S_{1-\beta} f] \leq \tau'$, then $\max_i I_i^{\leq d}[\mathcal{T}_C f] \leq \tau$.

Proof. Fix $\tau, \alpha > 0$ and $d \in \mathbb{N}$. We choose $\beta_0 = \frac{\alpha}{8d}$, and $\tau' > 0$ small enough. Fix $0 < \beta \leq \beta_0$; our proof uses a parameter $D \in \mathbb{N}$ which is large enough with respect to $1/\beta, 1/\tau, 1/\alpha, d$, and then N large enough with respect to D .

Consider the distribution μ defined by $(\mathbf{x}, S_{1-\beta}\mathbf{x})$ where $\mathbf{x} \in_R \mathcal{U}_k$, and consider its product version $\tilde{\mu}$. By Proposition 4.6 we may find a $(\Omega_{\alpha,m}(1), \zeta')$ -coupling $\mathcal{C}' = ((\mathbf{x}(1), \mathbf{x}(2)), (\mathbf{y}(1), \mathbf{y}(2)))$ between μ and $\tilde{\mu}$, where $\zeta' = O_{\alpha,\beta,m}(1/\sqrt{n})$ (see Remark 2.15). We intend to show that the functions $\mathcal{T}_C S_{1-\beta} f$ and $\mathcal{T}_{\tilde{\mu}} \mathcal{T}_C f$ are very close to each other in the 2-norm, which will tell us that their individual influences are very close. From this, the statement of the claim quickly follows.

Consider $g_1 = \mathcal{T}_C((S_{1-\beta} f)^{\leq D})$, $g_2 = (\mathcal{T}_{\tilde{\mu}} \mathcal{T}_C f)^{\leq D}$. Note that the operators $\mathcal{T}_{\tilde{\mu}}, S_{(1-\beta)}$ and their adjoint operators preserve juntas, and therefore preserve $V_{=\ell}([m]^n, \nu_k)$ and $V_{=\ell}(\mathcal{U}_k)$ respectively for all ℓ . Thus, $g_2 = \mathcal{T}_{\tilde{\mu}}((\mathcal{T}_C f)^{\leq D})$, and $(S_{1-\beta} f)^{\leq D} = S_{1-\beta}(f^{\leq D})$. We therefore get that

$$\|\mathcal{T}_C S_{1-\beta} f - g_1\|_2 = \|\mathcal{T}_C S_{1-\beta}(f^{>D})\|_2 \leq \|S_{1-\beta} f^{>D}\|_2 \leq (1 + \Omega_{\beta,m}(1))^{-D} \|f^{>D}\|_2,$$

where the third transition is by Lemma 3.11; this is at most $(1 + \Omega_{\beta,m}(1))^{-D}$ as $\|f^{>D}\|_2 \leq \|f\|_2 \leq 1$. We also get that

$$\|\mathcal{T}_{\tilde{\mu}} \mathcal{T}_C f - g_2\|_2 = \|\mathcal{T}_{\tilde{\mu}}(\mathcal{T}_C f)^{>D}\|_2 \leq (1 + \Omega_{\beta,m}(1))^{-D} \|(\mathcal{T}_C f)^{>D}\|_2 \leq (1 + \Omega_{\beta,m}(1))^{-D} \|\mathcal{T}_C f\|_2,$$

which is at most $(1 + \Omega_{\beta,m}(1))^{-D}$ as f is $[-1, 1]$ -valued; in the second transition we used Lemma B.3. Thus,

$$\|\mathcal{T}_{\tilde{\mu}} \mathcal{T}_C f - \mathcal{T}_C S_{1-\beta} f\|_2 \leq \|g_1 - g_2\|_2 + 2(1 + \Omega_{\beta,m}(1))^{-D},$$

and we bound $\|g_1 - g_2\|_2 \leq O_{\alpha,m,D}(\sqrt{\zeta} + \zeta')$.

Claim 6.19. $\|g_1 - g_2\|_2 \leq O_{\alpha,D}(\sqrt{\zeta} + \zeta')$.

Proof. Let $\mathcal{C}'_{i,j}$ denote the marginal distribution of $\mathbf{x}(i)$ and $\mathbf{y}(j)$, and define $g'_1 = \mathcal{T}_{\mathcal{C}'_{1,2}}(f^{\leq D})$, $g'_2 = \mathcal{T}_{\tilde{\mu}} \mathcal{T}_C(f^{\leq D})$. Then

$$\|g_1 - g_2\|_2 \leq \|g_1 - g'_1\|_2 + \|g'_1 - g'_2\|_2 + \|g'_2 - g_2\|_2, \quad (20)$$

and we upper bound each one of the norms on the right hand side separately. For the second norm in the right hand side of (20), note that

$$g'_1(y) = \mathbb{E}_{\substack{(\mathbf{x}(1), \mathbf{x}(2)) \\ (\mathbf{y}(1), \mathbf{y}(2)) \sim \mathcal{C}'}} [f^{\leq D}(\mathbf{x}(2)) \mid \mathbf{y}(1) = y], \quad g'_2(y) = \mathbb{E}_{\substack{(\mathbf{x}(1), \mathbf{x}(2)) \\ (\mathbf{y}(1), \mathbf{y}(2)) \sim \mathcal{C}'}} [\mathcal{T}_C(f^{\leq D})(\mathbf{y}(2)) \mid \mathbf{y}(1) = y],$$

so by Cauchy-Schwarz

$$\|g'_1 - g'_2\|_2^2 \leq \mathbb{E}_{\substack{(\mathbf{x}(1), \mathbf{x}(2)) \\ (\mathbf{y}(1), \mathbf{y}(2)) \sim \mathcal{C}'}} [|f^{\leq D}(\mathbf{x}(2)) - \mathcal{T}_C(f^{\leq D})(\mathbf{y}(2))|^2] \leq O(\sqrt{D\zeta}),$$

where the last inequality is by Lemma 1.6. To bound $\|g_1 - g'_1\|_2$, consider the following distribution: sample $(\mathbf{x}(1), \mathbf{x}(2), \mathbf{y}(1), \mathbf{y}(2)) \sim \mathcal{C}'$, and then $\mathbf{y}$ as $(\mathbf{x}, \mathbf{y}) \sim \mathcal{C}$ conditioned on $\mathbf{x} = \mathbf{x}(1)$. Then

$$g_1(y) = \mathbb{E}_{\left(\begin{smallmatrix} \mathbf{x}(1), \mathbf{x}(2) \\ \mathbf{y}(1), \mathbf{y}(2) \end{smallmatrix}\right) \sim \mathcal{C}', \mathbf{y}} [f^{\leq D}(\mathbf{x}(2)) \mid \mathbf{y} = y].$$

The above sampling procedure naturally defines $(\Omega_{\alpha, m}(1), \zeta + \zeta')$ -couplings $\mathcal{C}''$ of $(\mathbf{x}(2), \mathbf{y}(1))$ and a coupling $\mathcal{C}'''$ of $(\mathbf{x}(2), \mathbf{y})$, and with this notation $g'_1(y) = T_{\mathcal{C}''} f^{\leq D}(y)$ and $g_1(y) = T_{\mathcal{C}'''} f^{\leq D}(y)$, so by Lemma 1.6

$$\mathbb{E}_{\mathcal{C}'''} [(g_1(\mathbf{y}) - f^{\leq D}(\mathbf{x}(2)))^2] \leq O(\sqrt{D(\zeta + \zeta')}), \quad \mathbb{E}_{\mathcal{C}''} [(g'_1(\mathbf{y}) - f^{\leq D}(\mathbf{x}(2)))^2] \leq O(\sqrt{D(\zeta + \zeta')}).$$

Thus, by the triangle inequality

$$\mathbb{E}_{\mathbf{y} \sim \nu_k^{\otimes n}} [(g_1(\mathbf{y}) - g'_1(\mathbf{y}))^2] \leq O(\sqrt{D(\zeta + \zeta')}),$$

i.e. $\|g'_1 - g_1\| \leq O((D(\zeta + \zeta'))^{1/4})$.

Finally, we bound the third norm in (20):

$$\begin{aligned} \|g_2 - g'_2\| &= \|(T_{\tilde{\mu}} T_{\mathcal{C}} f)^{\leq D} - T_{\tilde{\mu}} T_{\mathcal{C}} (f^{\leq D})\| = \|T_{\tilde{\mu}} ((T_{\mathcal{C}} f)^{\leq D} - T_{\mathcal{C}} (f^{\leq D}))\| \\ &\leq \|(T_{\mathcal{C}} f)^{\leq D} - T_{\mathcal{C}} (f^{\leq D})\| \\ &\leq O_{\alpha, D}(\zeta^{1/4}) \|f\|_2 \\ &\leq O_{\alpha, D}(\zeta^{1/4}), \end{aligned}$$

where we used Lemma 3.25. Thus by the triangle inequality $\|g_1 - g_2\|_2 \leq O_{\alpha, D}((\zeta + \zeta')^{1/4})$. $\square$

Combining all, we get that $\|T_{\mathcal{C}} S_{1-\beta} f - T_{\tilde{\mu}} T_{\mathcal{C}} f\|_2 \leq O_{\alpha, D}((\zeta + \zeta')^{1/4}) + 2(1 + \Omega_{\beta, m}(1))^{-D}$. Choosing D large enough so that $2(1 + \Omega_{\beta, m}(1))^{-D} \leq \tau'/2$, and then N large enough so that ζ' is small enough, and then $\zeta > 0$ small enough, we get that $\|T_{\mathcal{C}} S_{1-\beta} f - T_{\tilde{\mu}} T_{\mathcal{C}} f\|_2 \leq \tau'$. Thus,

$$\max_i I_i[T_{\tilde{\mu}} T_{\mathcal{C}} f] \leq 2 \max_i I_i[T_{\mathcal{C}} S_{1-\beta} f] + 2 \|T_{\mathcal{C}} S_{1-\beta} f - T_{\tilde{\mu}} T_{\mathcal{C}} f\|_2 = O(\tau').$$

Note that by definition of $\tilde{\mu}$, on each coordinate i , the probability that $\mathbf{y}(1)_i = \mathbf{y}(2)_i$ is at least $1 - \frac{\beta}{\alpha} \geq 1 - \frac{1}{8d}$, so from Claim B.8 it follows that $\max_i I_i^{\leq d}[T_{\mathcal{C}} f] \leq O(\tau')$. Choosing $\tau' = c\tau$ for sufficiently small absolute $c > 0$ finishes the proof. $\square$

We are now ready to prove Lemma 6.16.

Proof of Lemma 6.16. Choose β_0, τ'', N' from Claim 6.18 for τ, α, d . For $0 < \beta \leq \beta_0$ we have by Claim 6.14,

$$I[T_{\mathcal{C}} S_{1-\beta} f] = O_{m, \beta}(\log^2 n) + O_{m, \beta, \alpha}(\zeta n \log^2 n),$$

so by Claim 6.17 we get that

$$\begin{aligned} \max_i I_i[T_{\mathcal{C}} S_{1-\beta} f] &\leq \frac{2}{\alpha} I_i[f] + \frac{2}{\alpha^2 n} I[T_{\mathcal{C}} f] \leq \frac{2}{\alpha} \tau' + O_{m, \beta} \left(\frac{\log^2 n}{n} \right) + O_{m, \beta, \alpha}(\zeta \log^2 n) \\ &\leq \frac{2}{\alpha} \tau' + O_{m, \beta, \alpha} \left(\frac{1}{\log n} \right), \end{aligned}$$

where we used the upper bound on ζ . Choosing $\tau' = \tau''\alpha/4$ and $N'' = N''(\alpha, \beta, m, \tau, d)$ large enough, it follows that for $N = \max(N', N'')$, if $n \geq N$ then $\max_i I_i[\text{TC}S_{1-\beta}f] \leq \tau''$, and from Claim 6.18 that we conclude that $\max_i I_i^{\leq d}[\text{TC}f] \leq \tau$ (we also use the fact that $\zeta \leq 1/\log n \leq 1/\log N''$, so ζ is small enough to apply Claim 6.18 provided N'' is large enough). $\square$

6.1.3 Relation with projections

Suppose n is even. We denote by $S_{n:n/2}$ the set of all 2-to-1 maps $\pi: [n] \rightarrow [n/2]$. For any $x \in [m]^{n/2}$ and $\pi \in S_{n:n/2}$, we define the point $y = \pi^{-1}(x) \in [m]^n$ as $y_i = x_{\pi(i)}$ for all $i \in [n]$.

Definition 6.20. Suppose $\mathcal{U}_{\vec{k}}$ is a multi-slice and k_i is even for all i (so in particular n is even). Given $\pi \in S_{n:n/2}$, define the projection of $f: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ along π as $f|_{\pi}: \mathcal{U}_{\vec{k}/2} \rightarrow [m]^{n/2} \rightarrow \mathbb{R}$, $f|_{\pi}(x) = f(\pi^{-1}(x))$.

The main goal of this section is to prove the following lemma, asserting that if a function over the multi-slice has a coordinate i with large noisy influence, then for a random 2-to-1 map π , the coordinate $\pi(i)$ will have large noisy influence in $f|_{\pi}$ with constant probability. More precisely:

Lemma 6.21. For all $\tau', \alpha > 0$, $m \in \mathbb{N}$ there exist $\beta_1 > 0$, $\tau'' > 0$ and $N \in \mathbb{N}$ such that the following holds for $n \geq N$ and $0 < \beta \leq \beta_1$. Suppose $\vec{k}$ is a vector of m even integers summing up to n . If $f: \mathcal{U}_{\vec{k}} \rightarrow [-1, 1]$ and $i \in [n]$ are such that $I_i^{(\beta)}[f] \geq \tau'$, then

$$\Pr_{\pi \in S_{n:n/2}} \left[I_{\pi(i)}^{(\beta)}[f|_{\pi}] \geq \tau'' \right] \geq \tau''.$$

The rest of this section is devoted to the proof of Lemma 6.21.

It will be convenient for us to consider different notions of influence for functions on the multi-slice. For $i \in [n]$ and $g: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$, define $A_i^g(x) = \mathbb{E}_{j \in [n]} [g(\pi_{i,j} \mathbf{x})]$ and then $M_i(g) = \mathbb{E}_{\mathbf{x} \in \mathcal{U}_{\vec{k}}} [(A_i^g(\mathbf{x}) - g(\mathbf{x}))^2]$.

Claim 6.22. For all $g: \mathcal{U}_{\vec{k}} \rightarrow \mathbb{R}$ and $i \in [n]$ we have $M_i(g) \geq \frac{1}{6} I_i[g] - \frac{1}{2n} I[g]$.

Proof. Using $(a + b + c)^2 \leq 3(a^2 + b^2 + c^2)$ we get

$$\begin{aligned} I_i[g] &= \mathbb{E}_{\mathbf{x}, \mathbf{s}} [(g(\pi_{i,\mathbf{s}} \mathbf{x}) - g(\mathbf{x}))^2] \leq 3 \mathbb{E}_{\mathbf{x}, \mathbf{s}} [(A_i^g(\mathbf{x}) - g(\mathbf{x}))^2 + (A_i^g(\pi_{i,\mathbf{s}} \mathbf{x}) - g(\pi_{i,\mathbf{s}} \mathbf{x}))^2 + (A_i^g(\pi_{i,\mathbf{s}} \mathbf{x}) - A_i^g(\mathbf{x}))^2] \\ &= 6M_i(g) + 3 \mathbb{E}_{\mathbf{x}, \mathbf{s}} [(A_i^g(\pi_{i,\mathbf{s}} \mathbf{x}) - A_i^g(\mathbf{x}))^2], \end{aligned} \quad (21)$$

and we upper bound the last expectation by $I[g]/n$ as follows.

$$\mathbb{E}_{\mathbf{x}, \mathbf{s}} [(A_i^g(\pi_{i,\mathbf{s}} \mathbf{x}) - A_i^g(\mathbf{x}))^2] = \mathbb{E}_{\mathbf{x}, \mathbf{s}} \left[\left(\mathbb{E}_{\mathbf{j}} [g(\pi_{i,\mathbf{j}} \pi_{i,\mathbf{s}} \mathbf{x}) - g(\pi_{i,\mathbf{j}} \mathbf{x})] \right)^2 \right] \leq \mathbb{E}_{\mathbf{x}, \mathbf{s}, \mathbf{j}} [(g(\pi_{i,\mathbf{j}} \pi_{i,\mathbf{s}} \mathbf{x}) - g(\pi_{i,\mathbf{j}} \mathbf{x}))^2].$$

Make the change of variables $\mathbf{y} = \pi_{i,\mathbf{j}} \pi_{i,\mathbf{s}} \mathbf{x}$, and note that then we have that $\pi_{i,\mathbf{j}} \pi_{i,\mathbf{s}} \mathbf{x} = \pi_{i,\mathbf{j}} \mathbf{y}$. Thus

$$\mathbb{E}_{\mathbf{x}, \mathbf{s}} [(A_i^g(\pi_{i,\mathbf{s}} \mathbf{x}) - A_i^g(\mathbf{x}))^2] \leq \mathbb{E}_{\mathbf{y}, \mathbf{s}, \mathbf{j}} [(g(\pi_{i,\mathbf{j}} \mathbf{y}) - g(\mathbf{y}))^2] = \frac{1}{n} I[g].$$

Plugging this into (21) and rearranging finishes the proof. $\square$

We next define another notion of influence that will get us closer to the left hand of Lemma 6.21. Define

$$B_i^g(x) = \mathbb{E}_{\substack{\mathbf{i}' \in [n] \setminus \{i\} \\ \mathbf{j} \in R[n] \\ \mathbf{j}' \in [n] \setminus \{i, \mathbf{i}', \mathbf{j}\}}} [g(\pi_{\mathbf{i}', \mathbf{j}'}^{\pi_{i, \mathbf{j}}} x) \mid x_{\mathbf{i}'} = x_{\mathbf{i}}, x_{\mathbf{j}'} = x_{\mathbf{j}}], \quad Q_i(g) = \mathbb{E}_{\mathbf{x} \in \mathcal{U}_k} [(B_i^g(\mathbf{x}) - g(\mathbf{x}))^2].$$

(we stress that the distribution over $\mathbf{j}$ is uniform, and then the distribution over $\mathbf{j}'$ is uniform among all $\mathbf{j}'$ such that $x_{\mathbf{j}'} = x_{\mathbf{j}}$ and different from $i, \mathbf{i}', \mathbf{j}$)

Claim 6.23. For all $g: \mathcal{U}_k \rightarrow \mathbb{R}$ and $i \in [n]$ we have $Q_i(g) \geq \frac{1}{2}M_i(g) - O_\alpha\left(\frac{I[g]}{n}\right)$.

Proof. Using $(a+b)^2 \geq \frac{1}{2}a^2 - b^2$ we get $Q_i(g) \geq \frac{1}{2}M_i(g) - \mathbb{E}_{\mathbf{x}} [(A_i^g(\mathbf{x}) - B_i^g(\mathbf{x}))^2]$ and we upper bound the last expectation by $O_\alpha\left(\frac{I[g]}{n}\right)$. Indeed, note that

$$\begin{aligned} \mathbb{E}_{\mathbf{x}} [(A_i^g(\mathbf{x}) - B_i^g(\mathbf{x}))^2] &= \mathbb{E}_{\mathbf{x}} \left[\left(\mathbb{E}_{\substack{\mathbf{i}' \in [n] \setminus \{i\} \\ \mathbf{j} \in R[n] \\ \mathbf{j}' \in [n] \setminus \{i, \mathbf{i}', \mathbf{j}\}}} [g(\pi_{\mathbf{i}', \mathbf{j}'}^{\pi_{i, \mathbf{j}}} \mathbf{x}) - g(\pi_{i, \mathbf{j}} \mathbf{x}) \mid \mathbf{x}_{\mathbf{i}'} = \mathbf{x}_{\mathbf{i}}, \mathbf{x}_{\mathbf{j}'} = \mathbf{x}_{\mathbf{j}}] \right)^2 \right] \\ &\leq \mathbb{E}_{\mathbf{x}} \left[\mathbb{E}_{\substack{\mathbf{i}' \in [n] \setminus \{i\} \\ \mathbf{j} \in R[n] \\ \mathbf{j}' \in [n] \setminus \{i, \mathbf{i}', \mathbf{j}\}}} [(g(\pi_{\mathbf{i}', \mathbf{j}'}^{\pi_{i, \mathbf{j}}} \mathbf{x}) - g(\pi_{i, \mathbf{j}} \mathbf{x}))^2 \mid \mathbf{x}_{\mathbf{i}'} = \mathbf{x}_{\mathbf{i}}, \mathbf{x}_{\mathbf{j}'} = \mathbf{x}_{\mathbf{j}}] \right] \\ &\leq \frac{n}{\alpha n - 3} \left(1 + O\left(\frac{1}{n}\right) \right) \mathbb{E}_{\mathbf{x}} \left[\mathbb{E}_{\mathbf{i}', \mathbf{j}, \mathbf{j}'} [(g(\pi_{\mathbf{i}', \mathbf{j}'}^{\pi_{i, \mathbf{j}}} \mathbf{x}) - g(\pi_{i, \mathbf{j}} \mathbf{x}))^2] \right] \\ &= O\left(\frac{1}{\alpha}\right) \mathbb{E}_{\mathbf{x}} \left[\mathbb{E}_{\mathbf{i}', \mathbf{j}'} [(g(\pi_{\mathbf{i}', \mathbf{j}'} \mathbf{x}) - g(\mathbf{x}))^2] \right] \\ &= O_\alpha\left(\frac{I[g]}{n}\right). \square \end{aligned}$$

Finally, we relate the left hand side of Lemma 6.21 directly to the Q_i notion of influence.

Claim 6.24. For all $\alpha > 0$, there is $\beta_0 > 0$ such that if $0 < \beta \leq \beta_1$, then $f: \mathcal{U}_k \rightarrow [-1, 1]$ and $i \in [n]$ we have

$$\mathbb{E}_{\pi \in S_{n:n/2}} [I_{\pi(i)}^{(\beta)}[f\pi]] \geq \frac{1}{4}Q_i(S_{1-\beta}f) - O\left(\frac{\beta}{\alpha}\right).$$

Proof.

$$\mathbb{E}_{\pi \in S_{n:n/2}} [I_{\pi(i)}^{(\beta)}[f\pi]] = \mathbb{E}_{\mathbf{x} \in \mathcal{U}_{k/2}} \left[\mathbb{E}_{\mathbf{j} \in [n/2]} \left[\mathbb{E}_{\pi \in S_{n:n/2}} [(S_{1-\beta}f|_{\pi})^{(\pi_{i, \mathbf{j}} \mathbf{x})} - (S_{1-\beta}f|_{\pi})(\mathbf{x}))^2] \right] \right].$$

Next, we introduce $\mathbf{z} = \pi^{-1}(\mathbf{x})$, condition on it and then use Cauchy-Schwarz to see that the above expression is at least

$$\mathbb{E}_{\mathbf{z} \in \mathcal{U}_k^-} \left[\left(\mathbb{E}_{\substack{\mathbf{x} \in \mathcal{U}_{k/2}^- \\ \pi \in S_{n:n/2} \\ \mathbf{j} \in [n/2]}} [(S_{1-\beta}f|_{\pi})(\pi^{\pi(i),\mathbf{j}}\mathbf{x}) \mid \pi^{-1}(\mathbf{x}) = \mathbf{z}] - \mathbb{E}_{\substack{\mathbf{x} \in \mathcal{U}_{k/2}^- \\ \pi \in S_{n:n/2} \\ \mathbf{j} \in [n/2]}} [(S_{1-\beta}f|_{\pi})(\mathbf{x}) \mid \pi^{-1}(\mathbf{x}) = \mathbf{z}] \right)^2 \right]. \quad (22)$$

We say $\pi \in S_{n:n/2}$ is consistent with $z \in \mathcal{U}_k^-$ if $z \in \pi^{-1}(\mathcal{U}_{k/2}^-)$. Fix $\mathbf{z} = z$; in the second expectation, we first sample π consistent with z , then take $\mathbf{x} = \pi(\mathbf{z})$, then take $\mathbf{x}' \sim S_{1-\beta}\mathbf{x}$, and average $f_{\pi}(\mathbf{x}') = f(\pi^{-1}(\mathbf{x}'))$. Note that the distribution of $\pi^{-1}(\mathbf{x}')$ is the same as $S_{1-\beta}z$, so the second expectation is nothing but $S_{1-\beta}f(\mathbf{z})$.

We move on to the first expectation. Fix $\mathbf{z} = z$; we first sample π consistent with z , then set $\mathbf{x} = \pi(\mathbf{z})$, take $\mathbf{j} \in [n/2]$, look at the point $\mathbf{x}' = \pi^{\pi(i),\mathbf{j}}\mathbf{x}$, then sample $\mathbf{x}'' \sim S_{1-\beta}\mathbf{x}'$ to yield $\mathbf{z}' = \pi^{-1}(\mathbf{x}'')$. The first expectation is the average of $f(\mathbf{z}')$ sampled this way, and we analyze the distribution of $\mathbf{z}'$.

Note that the points $\mathbf{x}, \mathbf{x}'$ differ in either 0 or 2 coordinates. Thus, there is a natural coupling between $S_{1-\beta}\mathbf{x}$ and $S_{1-\beta}\mathbf{x}'$, call it $(\mathbf{u}, \mathbf{u}')$, in which $\mathbf{u}, \mathbf{u}'$ always differ in either 0 or 2 coordinates. To see that, consider for each $a \in [m]$ the set of coordinate of $\mathbf{x}$ equal to a , and the set of coordinates $\mathbf{x}'$ equal to a . Then either:

1. these sets are the same for all a – in which case we can make the sampling procedure $S_{1-\beta}$ to choose the same subsets of coordinates to change in each color class.
2. Else, these sets are the same for all but 2 of the a 's, say a_1 and a_2 . This, the sets of coordinates corresponding to a_1 (as well as the sets of coordinates corresponding to a_2) have symmetric difference of size 1. In this case, for each $a \neq a_1, a_2$ we may make the sampling procedure $S_{1-\beta}$ choose the subsets of coordinates, and for $a \in \{a_1, a_2\}$ we may choose subsets of coordinates that have symmetric difference at most 1.

Denote the set of coordinates in which $\mathbf{x}'$ and $\mathbf{x}''$ differ by $\mathbf{U}$. Thus, we get a coupling of $S_{1-\beta}z$ and $S_{1-\beta}z'$, call it $\mathbf{w}, \mathbf{w}'$, in which the points differ in either 0 or 4 coordinates; denote the set of these coordinates as $\mathbf{W} = \pi^{-1}(\mathbf{U})$. Denote by $\pi_{\mathbf{W}}$ a composition of two transpositions on $\mathbf{W}$, so that $\mathbf{w}' = \pi_{\mathbf{W}}\mathbf{w}$ (or the identity if $\mathbf{W} = \emptyset$). Then Plugging this into (22), we get that

$$(22) = \mathbb{E}_{\mathbf{z}} \left[\left(\mathbb{E}_{\mathbf{w} \sim S_{1-\beta}\mathbf{z}} \left[\mathbb{E}_{\mathbf{W}} [f(\pi_{\mathbf{W}}\mathbf{w}) - f(\mathbf{w})] \right] \right)^2 \right] = \mathbb{E}_{\mathbf{z}} \left[\left(\mathbb{E}_{\mathbf{W}} \left[\mathbb{E}_{\mathbf{w} \sim S_{1-\beta}\mathbf{z}} [f(\pi_{\mathbf{W}}\mathbf{w}) - f(\mathbf{w})] \right] \right)^2 \right],$$

where in the last transition we used the fact that the distribution of $\mathbf{w}$ and $\mathbf{W}$ are independent. When $\mathbf{W}$ is empty set get 0, so

$$(22) = \mathbb{E}_{\mathbf{z}} \left[\left(\mathbb{E}_{\mathbf{w} \sim S_{1-\beta}\mathbf{z}} \left[\mathbb{E}_{\mathbf{W}} [f(\pi_{\mathbf{W}}\mathbf{w}) - f(\mathbf{w})] \right] \right)^2 \right] = \mathbb{E}_{\mathbf{z}} \left[\left(\mathbb{E}_{\mathbf{W}} \left[1_{|\mathbf{W}|=4} \mathbb{E}_{\mathbf{w} \sim S_{1-\beta}\mathbf{z}} [\pi_{\mathbf{W}}f(\mathbf{w}) - g(\mathbf{w})] \right] \right)^2 \right].$$

Next, we note that the probability that $\mathbf{W}$ is non-empty is $(1 - O(\beta/\alpha)) \geq \frac{1}{2}$, and we need to inspect the distribution of $\mathbf{W}$ more closely when it is non-empty. In that case, $\mathbf{U}$ contains the coordinate $\pi(i)$ along

with a random coordinate $\mathbf{j} \in [n/2]$. Thus, the distribution of $\pi^{-1}(\mathbf{U})$ is $\{i, i', \mathbf{j}_1, \mathbf{j}_2\}$, where: i' is uniform on $i' \neq i$ such that $z_{i'} = z_i$, $\mathbf{j}_1 \in [n]$ is uniform, and conditioned on $\mathbf{j}_1 = j_1$ the distribution of $\mathbf{j}_2$ is uniform among $j_2 \neq j$ such that $z_{j_2} = z_{j_1}$. Thus, we get that

$$(22) \geq \frac{1}{4} \mathbb{E}_{\mathbf{z}} \left[\left(\mathbb{E}_{i', \mathbf{j}_1 \in [n], \mathbf{j}_2} [S_{1-\beta}(\pi_{i', \mathbf{j}_2}^{\pi_{i, \mathbf{j}_1}} f)(\mathbf{z}) - S_{1-\beta} f(\mathbf{z}) \mid \mathbf{z}_{i'} = \mathbf{z}_i, \mathbf{z}_{\mathbf{j}_2} = \mathbf{z}_{\mathbf{j}_1}] \right)^2 \right].$$

Next, observe that for all z and $\pi \in S_n$ which is a composition of two transpositions we have that $|S_{1-\beta}(\pi f)(z) - S_{1-\beta} f(\pi z)| \leq O\left(\frac{\beta}{\alpha}\right)$. This is seen by considering a natural coupling between $S_{1-\beta} \pi z$ and $\pi S_{1-\beta} z$ similar to the coupling of $\mathbf{w}, \mathbf{w}'$ above, and noting that the probability $S_{1-\beta}$ touches a coordinate ℓ which is not a fixed point of π is at most $O\left(\frac{\beta}{\alpha}\right)$, and if this doesn't happen then the sampled points are the same. Thus, the statistical distance between $S_{1-\beta} \pi z$ and $\pi S_{1-\beta} z$ is $O\left(\frac{\beta}{\alpha}\right)$, and as f is bounded our observation follows. Thus,

$$(22) \geq \frac{1}{4} \mathbb{E}_{\mathbf{z}} \left[\left(\mathbb{E}_{i', \mathbf{j}_1 \in [n], \mathbf{j}_2} [S_{1-\beta} f(\pi_{i', \mathbf{j}_2}^{\pi_{i, \mathbf{j}_1}} \mathbf{z}) - S_{1-\beta} f(\mathbf{z}) \mid \mathbf{z}_{i'} = \mathbf{z}_i, \mathbf{z}_{\mathbf{j}_2} = \mathbf{z}_{\mathbf{j}_1}] \right)^2 \right] - O\left(\frac{\beta}{\alpha}\right),$$

which is equal to $\frac{1}{4} Q_i(S_{1-\beta} g) - O\left(\frac{\beta}{\alpha}\right)$. □

We are now ready to prove Lemma 6.21

Proof. Fix $\alpha, \tau' > 0$ and $m \in \mathbb{N}$. We choose $\tau = \frac{\tau'}{1000}, \beta'_1$ from Claim 6.24, then $\beta_1 > 0$ small enough, and finally $N = N(\alpha, \tau', m, \beta)$ large enough.

Assume in the setting of the lemma we have $I_i^{(\beta)}[f] \geq \tau'$. By Lemma 6.13 we have $I^{(\beta)}[f] = O_\beta(1)$. Set $g = S_{1-\beta} f$ and apply Claim 6.22 to get that $M_i(g) \geq \frac{1}{6} \tau' - O_\beta\left(\frac{1}{N}\right) \geq \frac{1}{7} \tau'$ by the choice of N . Thus, by Claim 6.23 we get $Q_i(g) \geq \frac{1}{14} \tau' - O_\alpha\left(\frac{1}{N}\right) \geq \frac{1}{20} \tau'$, by choice of N . Thus, by Claim 6.24

$$\mathbb{E}_{\pi \in S_{n:n/2}} \left[I_{\pi(i)}^{(\beta)}[f_\pi] \right] \geq \frac{1}{4} Q_i(g) - O\left(\frac{\beta}{\alpha}\right) \geq \frac{\tau'}{80} - O\left(\frac{\beta}{\alpha}\right) \geq \frac{\tau'}{100},$$

provided β_1 is small enough. Now as f is bounded, all of its projections are bounded and so $I_{\pi(i)}^{(\beta)}[f_\pi] \leq 4$ for all π ; thus

$$\Pr_{\pi \in S_{n:n/2}} \left[I_{\pi(i)}^{(\beta)}[f_\pi] \geq \frac{\tau'}{1000} \right] \geq \frac{\tau'}{1000}. \quad \square$$

6.2 The reduction

Let $\mathcal{P}$ be a weighted collection of r -ary CSPs over the alphabet $\Sigma = [m]$. Given a $(1, s)$ dictatorship test for $\mathcal{P}$, denoted by $(w, \{\mathcal{D}_t\}_{t \in \mathcal{T}}, p)$, in which each $\mathcal{D}_t$ is connected, we describe a polynomial time reduction that takes as input an instance Φ of $\text{Gap-Rich-2-to-1}_n[1, \delta]$, and outputs an instance Ψ of $\text{CSP-}\mathcal{P}$ such that:

1. **Completeness:** if $\text{val}(\Phi) = 1$, then $\text{val}(\Psi) = 1$.
2. **Soundness:** for all $\varepsilon > 0$, there is sufficiently small $\delta > 0$ and sufficiently large even n , such that if $\text{val}(\Phi) \leq \delta$, then $\text{val}(\Psi) \leq s + \varepsilon$.

6.2.1 Parameters

Throughout the proof, we will have several parameters $\varepsilon, r, m, \alpha, \beta_0, \beta_1, \beta, \tau, \tau', \tau'', \zeta, \delta$ and n , as well as a finite set $\mathcal{T}'$ whose size we'll be interested in, and we state the hierarchy between them. We denote by $0 < a \ll b$ the relation that once we choose a , we take b to be sufficiently large in comparison to a . With this notation, our hierarchy is

$$0 < r, m, \varepsilon^{-1}, \alpha^{-1}, |\mathcal{T}'| \ll \tau^{-1}, \tau'^{-1}, \tau''^{-1} \ll \beta_0^{-1}, \beta_1^{-1} \ll \beta^{-1} \ll \delta^{-1} \ll N \ll n,$$

and $\zeta = n^{-1/4}$. Also, we will assume that $\beta n/2$ is an integer (otherwise we change β slightly to arrange that, while keeping the hierarchy of the parameters).

6.2.2 Setting up the multi-slices and the distributions $\mu_{t,P}$

We may find a finite $\mathcal{T}' \subseteq \mathcal{T}$ such that $\sum_{t \in \mathcal{T} \setminus \mathcal{T}'} p(t) \leq \varepsilon/2$, and we fix such $\mathcal{T}'$. First, as n is an even integer, we may assume that for all $t \in \mathcal{T}'$, $P \in \mathcal{P}$ and $\vec{a} \in \text{supp}(\mathcal{D}(t, P))$, we have that $n \cdot \mathcal{D}(t, P)(\vec{a})$ is an even integer. Otherwise, we may find distribution $(\mathcal{D}(t, P)')_{t \in \mathcal{T}'}$ such that $\mathcal{D}(t, P)', \mathcal{D}(t, P)$ have the same support and are $O_m(1/n^2)$ -close in KL-divergence, and so $\mathcal{D}(t, P)'^{\otimes n}, \mathcal{D}(t, P)^{\otimes n}$ are $O_m(1/n)$ -close in KL-divergence and by Pinsker's inequality the statistical distance between them is $O_m(1/\sqrt{n}) = o(1)$. In particular, this collection of distributions forms a $(1, s + o(1))$ dictatorship test, and we may proceed the argument with it. Denote

$$\alpha = \min_{t \in \mathcal{T}', P \in \mathcal{P}} \min_{\vec{a} \in \text{supp}(\mathcal{D}(t, P))} (\mathcal{D}(t, P))(\vec{a}) > 0.$$

For each $t \in \mathcal{T}'$, $P \in \mathcal{P}$ consider the distribution $\mathcal{D}(t, P)$ over $[m]^r$, and denote its marginal distributions by $\mathcal{D}(t, P)_1, \dots, \mathcal{D}(t, P)_r$. Denote $k(i, t, P)_a = n \cdot \mathcal{D}(t, P)_i(a)$, and note that $k(i, t, P)_a$ is an even integer. We thus have that $\vec{k}(i, t, P) = (k(i, t, P)_1, \dots, k(i, t, P)_m)$ is a sequence of even integers summing up to n .

Thus, we may consider the $r|\mathcal{T}'||\mathcal{P}|$ -mutli slices $\left(\mathcal{U}_{\vec{k}(i, t, P)} \right)_{\substack{i \in [r] \\ t \in \mathcal{T}', P \in \mathcal{P}}} \subseteq [m]^n$; note that each one of them is α -balanced. Next, we define for each $t \in \mathcal{T}'$, $P \in \mathcal{P}$ a distribution $\mu_{t,P}$ over $\prod_{i=1}^r \mathcal{U}_{\vec{k}(i, t, P)}$ that would “simulate” the dictatorship test $\mathcal{D}(t, P)^{\otimes n}$. The distribution $\mu_{t,P}$ is uniform on $(x(1), \dots, x(r))$ whose statistics match the expectation in $\mathcal{D}(t, P)^{\otimes n}$, i.e. on

$$\left\{ (x(1), \dots, x(r)) \in \prod_{i=1}^r \mathcal{U}_{\vec{k}(i, t, P)} \mid \forall \vec{a} \in [m]^r, \# \{i \mid x(1)_i = a_1, \dots, x(r)_i = a_r\} = n \cdot (\mathcal{D}(t, P))(\vec{a}) \right\}.$$

It is easily seen that $\mu_{t,P}$ is α -admissible as per Definition 1.8, and negatively correlated as per Definition 4.5. Also, we note that the product version of $\mu_{t,P}$ as per Definition 1.10, i.e. $\mu_{t,P}^{\otimes n}$, is nothing but $\mathcal{D}(t, P)^{\otimes n}$.

Remark 6.25. Below, we will assume that $\mu_{t,P_i}^{\otimes n}$ and $\mu_{t,P'_i}^{\otimes n}$ are either identical, or are $1 - o(1)$ far and hence the distribution can be recovered by a single sample except with probability $o(1)$. This may be guaranteed provided that n is large enough, as follows. We may find a constant $\beta > 0$ such that for any i, t, P, i', t', P' , either $\mathcal{D}(t, P)_i$ and $\mathcal{D}(t, P)_{i'}$ are identical, or else they are at least β far in total variation distance. Thus, the distributions $\mu_{t,P_i}^{\otimes n}$ and $\mu_{t,P'_i}^{\otimes n}$ are either identical, or are $1 - o(1)$ far. This means that provided that n is large enough, sampling $\mathbf{t} \sim \mathcal{T}$, $\mathbf{P} \sim w$ and $x \sim (\mu_{\mathbf{t}, \mathbf{P}})_i^{\otimes n}$ for some i , and outputting x , one may recover t', P' and i' such that $\mu_{\mathbf{t}, \mathbf{P}'_i}^{\otimes n} = \mu_{t', P'_i}^{\otimes n}$ except with probability $o(1)$.

6.2.3 The PCP construction

Write the Rich-2-to-1 Games instance Φ as $(L \cup R, E, [n], [n/2], \{\phi_e\}_{e \in E})$. We replace each vertex $v \in R$ with copies of the multi-slices $\left(\mathcal{U}_{\vec{k}(i,t,P)/2}\right)_{\substack{i \in [r] \\ t \in \mathcal{T}' \\ P \in \mathcal{P}}} \subseteq [m]^{n/2}$; we represent these as (v, i, t, P, x) where $x \in \mathcal{U}_{\vec{k}(i,t,P)/2}$, and denote by V the set of all of these tuples. Our instance Ψ of CSP- $\mathcal{P}$ will have a variable $z_{(v,i,t,x)}$ for each $(v, i, t, P, x) \in V$. It will also be useful for us to imagine that each $u \in L$ also has $r |\mathcal{T}'| |\mathcal{P}|$ multislices $\left(\mathcal{U}_{\vec{k}(i,t,P)}\right)_{\substack{i \in [r] \\ t \in \mathcal{T}' \\ P \in \mathcal{P}}} \subseteq [m]^n$ associated with it, but the reduction does not actually use these.

Folding. We define the relation $\sim$ on V by $(v(1), i(1), t(1), x(1)) \sim (v(2), i(2), t(2), x(2))$ if there is a common neighbour $u \in R$ of $v(1), v(2)$, as well as $i \in [r]$, $t \in \mathcal{T}'$, $P \in \mathcal{P}$ and $x \in \mathcal{U}_{\vec{k}(i,t,P)}$, such that $x = \phi_{u,v(1)}^{-1}(x(1))$ and $x = \phi_{u,v(2)}^{-1}(x(2))$. In words and slightly informally, we say two vertices are equivalent if they are the projection of the same point (u, i, t, P, x) in the imaginary copies of the multislices corresponding to L , along the corresponding constraint 2-to-1 maps $\phi_{u,v(1)}$ and $\phi_{u,v(2)}$. We extend $\sim$ to an equivalence relation on V , which, by abusing notation, we also denote by $\sim$. Thus, we will look at the equivalence classes of $\sim$, and denote the equivalence class of (v, i, t, P, x) by $[(v, i, t, P, x)]$. For each equivalence class C of $\sim$, we choose a representative and identify all the variables corresponding to (v, i, t, P, x) with the variable of the representative. Formally, let C be an equivalence class and suppose (v, i, t, P, x) is its representative. Then whenever we refer to a variable $z_{(v', i', t', P', x')}$ corresponding to (v', i', t', P', x') from the same equivalence class C , we replace its occurrence with $z_{(v, i, t, P, x)}$.

Constraints. The constraints of Ψ are described by the following sampling procedure. Sample $\mathbf{t} \sim p$, and if $\mathbf{t} \notin \mathcal{T}'$ put a trivial constraint (say, we add r variables $\ell_1, \dots, \ell_r$, pick $P \in \mathcal{P}$, and put the constraint that $P(\ell_1, \dots, \ell_r)$). Otherwise, $\mathbf{t} \in \mathcal{T}'$, and then we sample $u \in L$ with probability proportional to its degree, $\mathbf{P} \sim w$, and $(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu_{\mathbf{t}, \mathbf{P}}$. Then for each $i = 1, \dots, r$, we sample $\mathbf{v}(i)$ a neighbour of u conditioned on $\mathbf{x}(i) \in \phi_{u, \mathbf{v}(i)}^{-1}(\mathcal{U}_{\vec{k}(i, \mathbf{t}, \mathbf{P})/2})$. Denote by $\mathbf{x}(i)'$ the unique point in $\mathcal{U}_{\vec{k}(i, \mathbf{t}, \mathbf{P})/2}$ such that $\mathbf{x}(i) = \phi_{u, \mathbf{v}(i)}^{-1}(\mathbf{x}(i)'),$ and add the constraint

$$\mathbf{P}(z_{\mathbf{v}(1), 1, \mathbf{t}, \mathbf{x}(1)'}, \dots, z_{\mathbf{v}(r), r, \mathbf{t}, \mathbf{x}(r)'})$$

Remark 6.26. We note that in the sampling procedure above, one can take an arbitrary $\mathbf{v}(i)$ satisfying this property as opposed to a random one, without introducing any ambiguity. The reason is that any two possible such $\mathbf{v}_i$'s would yield two tuples in V , between which the relation $\sim$ holds, and in that case by the folding they correspond to the same variable.

The analysis of this reduction spans the following 2 subsections.

6.3 Completeness of the reduction

As is often the case, the completeness of the reduction is easy. Indeed, assume Φ is satisfiable, and let $A_L: L \rightarrow [n]$, $A_R: R \rightarrow [n/2]$ be labeling that satisfy all of the constraints of Φ . For each v , we define the assignment A' to the variables of Ψ as

$$A'(z_{v,i,t,x}) = x_{A_R(v)}.$$

First, we note that the variables in each equivalence class of $\sim$ get the same value. For that, it suffices to observe that for any $(v(1), i(1), t(1), x(1)) \sim (v(2), i(2), t(2), x(2))$ in the basic form of $\sim$, it holds that their variables get the same value by A' . This is true, since if (u, i, t, x) is the common neighbour demonstrating that the basic form of $\sim$ holds, then as x is the pullback of $x(1)$ along $\phi_{u,v(1)}$ we have that both coordinates in $\phi_{u,v(1)}^{-1}(A_R(v(1)))$ have the same value in x , and since the constraint is satisfied we have $A_L(u) \in \phi_{u,v(1)}^{-1}(A_V(v(1)))$. Thus, $x(1)_{A_R(v(1))} = x_{A_L(u)}$, and a similar argument shows that $x(2)_{A_R(v(2))} = x_{A_L(u)}$. We thus get that

$$A'(z_{v(1),i(1),t(1),x(1)}) = x(1)_{A_R(v)} = x_{A_L(u)} = x(2)_{A_R(v(2))} = A'(z_{v(2),i(2),t(2),x(2)}).$$

We may therefore think of the labeling A' as also giving values to the imaginary multi-slices of vertices in L , where (u, i, t, x) gets the value $A'(z_{v,i',t',x'})$ where $x = \phi_{u,v}^{-1}(x')$. Next, we show that the constraints hold. Let $P(z_{v_1,1,t,x(1)}, \dots, z_{v_r,r,t,x(r)})$ be a constraint in the instance, and let $u \in L$ be the vertex that samples $v_1, \dots, v_r$. The constraint being satisfied amounts to $P(x(1)_{A_R(v_1)}, \dots, x(r)_{A_R(v_r)}) = 1$, and as before we have that if $x(i)'$ is the pull-back of $x(i)$ along $\phi_{u,v(i)}$, then $x(i)_{A_R(v_i)} = x(i')_{A_L(u)}$, so we must show that $P(x(1)'_{A_L(u)}, \dots, x(r)'_{A_L(u)})$. Considering the dictatorship function $f: [m]^n \rightarrow [m]$ defined as $f(y) = y_{A_L(u)}$, the constraint being satisfied is equivalent to $P(f(x(1)'), \dots, f(x(r)')) = 1$, and as f is a dictatorship and $(x(1)', \dots, x(r)')$ is in the support μ_t , and hence of $\mathcal{D}_t^{\otimes n}$, we have that this is satisfied by the completeness property of the dictatorship test (the first item of Definition 6.6).

6.4 Soundness analysis

In this section we finish the analysis of the PCP construction by showing that if $\text{val}(\Phi) \leq \delta$, then $\text{val}(\Psi) \leq s + \varepsilon$. We prove that counter-positively.

Namely, suppose we have an assignment A' to the variables satisfying at least $s + \varepsilon$ fraction of the constraints; we show that it is possible to find an assignment to Φ satisfying more than δ fraction of the constraints.

For each v, i, t, P we will think of the assignment on the multi-slice of v as $f_{v,i,t,P}: \mathcal{U}_{\tilde{k}(i,t,P)/2} \rightarrow [m]$. Let $u \in L$, since for all i, t, P, x such that $x \in \mathcal{U}_{\tilde{k}(i,t,P)}$ we have that all projections of x along $\pi_{u,v}$ for neighbours v of u get the same value by A' , we may think that this value is also given to (u, i, t, P, x) , and thus think of a function $f_{u,i,t,P}: \mathcal{U}_{\tilde{k}(i,t,P)} \rightarrow [m]$ corresponding to these values. Thus, for all $u \in L$ and a neighbour v of u , we have that $f_{v,i,t,P} = f_{u,i,t,P}|_{\phi_{u,v}}$.

Below, when we write $\mathbf{u} \in L$ we mean that we sample $\mathbf{u}$ with probability proportional to its degree.

Claim 6.27. $\mathbb{E}_{\mathbf{u} \in U} \left[\mathbb{E}_{\mathbf{P} \sim w} \left[\mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu_{\mathbf{t}, \mathbf{P}}} [\mathbf{P}(f_{\mathbf{u},1,\mathbf{t},\mathbf{P}}(\mathbf{x}(1)), \dots, f_{\mathbf{u},r,\mathbf{t},\mathbf{P}}(\mathbf{x}(r)))] \right] \right] \geq s + \varepsilon.$

Proof. By construction of the functions $f_{\mathbf{u}}$, this expression is equal to the fraction of constraints satisfied by A' . $\square$

Note that $\mathbf{t} \notin \mathcal{T}'$ with probability at most $\varepsilon/2$, so we get that

$$\mathbb{E}_{\substack{\mathbf{u} \in U \\ \mathbf{t} \sim p}} \left[\mathbb{1}_{\mathbf{t} \in \mathcal{T}'} \mathbb{E}_{\mathbf{P} \sim w} \left[\mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu_{\mathbf{t}, \mathbf{P}}} [\mathbf{P}(f_{\mathbf{u},1,\mathbf{t},\mathbf{P}}(\mathbf{x}(1)), \dots, f_{\mathbf{u},r,\mathbf{t},\mathbf{P}}(\mathbf{x}(r)))] \right] \right] \geq s + \frac{1}{2}\varepsilon.$$

Denote

$$\text{Good}[L] = \left\{ u \in L \mid \mathbb{E}_{\mathbf{t} \sim p} \left[1_{\mathbf{t} \in \mathcal{T}'} \mathbb{E}_{\mathbf{P} \sim w} \left[\mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mu_{\mathbf{t}, \mathbf{P}}} [\mathbf{P}(f_{u,1,\mathbf{t},\mathbf{P}}(\mathbf{x}(1)), \dots, f_{u,r,\mathbf{t},\mathbf{P}}(\mathbf{x}(r)))] \right] \right] \geq s + \frac{1}{4}\varepsilon \right\},$$

then by an averaging argument $\Pr_{\mathbf{u} \in L} [\mathbf{u} \in \text{Good}[L]] \geq \frac{1}{4}\varepsilon$. Next, we devise a randomized labeling strategy for vertices in $\text{Good}[L]$ and their neighbours.

Fix $u \in \text{Good}[L]$, $t \in \mathcal{T}'$, $P \in \mathcal{P}$. Note that by Proposition 4.6 there is a $(\Omega_{\alpha,m}(1), \zeta)$ -coupling between $\mathcal{U}_{\vec{k}(i,t,P)}$ and $\nu_{\vec{k}(i,t,P)}$, denote it by $\mathcal{C}_{i,t,P}$, and we also have a $(\Omega_{\alpha,r,m}(1), \zeta)$ -coupling between $\mu_{t,P}$ and $\mu_{t,P}^{\otimes n} = \mathcal{D}(t, P)^{\otimes n}$ (see Remark 2.15). We now think of the functions $f_{u,i,t}$ as receiving values in Δ_m (as in Section 4.2). Thus, we get from Theorem 4.4 that

$$\mathbb{E}_{\mathbf{t} \sim p} \left[1_{\mathbf{t} \in \mathcal{T}'} \mathbb{E}_{\mathbf{P} \sim w} \left[\mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mathcal{D}(\mathbf{t}, \mathbf{P})^{\otimes n}} [\tilde{\mathbf{P}}(\text{T}_{\mathcal{C}_{1,\mathbf{t},\mathbf{P}}} f_{u,1,\mathbf{t},\mathbf{P}}(\mathbf{x}(1)), \dots, \text{T}_{\mathcal{C}_{1,\mathbf{t},\mathbf{P}}} f_{u,r,\mathbf{t},\mathbf{P}}(\mathbf{x}(r)))] \right] \right] \geq s + \frac{1}{8}\varepsilon.$$

We now define a single function g_u in order to appeal to the assumption regarding the dictatorship test. Upon receiving an input x , we guess t' , P' and i such that we believe x was sampled from $\mathcal{D}(\mathbf{t}, \mathbf{P})_i^{\otimes n}$ as in Remark 6.25. We then set $g_u(x) = \text{T}_{\mathcal{C}_{i,t',P'}} f_{u,i,t',P'}(x)$. Then the above inequality and the assertion of Remark 6.25 implies that

$$\mathbb{E}_{\mathbf{t} \sim p} \left[1_{\mathbf{t} \in \mathcal{T}'} \mathbb{E}_{\mathbf{P} \sim w} \left[\mathbb{E}_{(\mathbf{x}(1), \dots, \mathbf{x}(r)) \sim \mathcal{D}(\mathbf{t}, \mathbf{P})^{\otimes n}} [\tilde{\mathbf{P}}(g_u(\mathbf{x}(1)), \dots, g_u(\mathbf{x}(r)))] \right] \right] \geq s + \frac{1}{8}\varepsilon - o(1) \geq s + \frac{1}{16}\varepsilon.$$

Pick d, τ from the second item in Definition 6.6 for our dictatorship test for $\gamma = \varepsilon/16$ and $\varepsilon/16$. Then from the soundness property of our dictatorship test, we get that for each $u \in \text{Good}[L]$, there are $t \in \mathcal{T}'$, $i \in [r]$ and $j \in [n]$ such that $I_j^{\leq d}[g_u; \nu_{\vec{k}(i,t)}^{\otimes n}] \geq \tau$. Thus, by Remark 6.25 again we have that

$$I_j^{\leq d}[\text{T}_{\mathcal{C}_{i,t}} f_{u,i,t}; \nu_{\vec{k}(i,t)}^{\otimes n}] \geq \tau/2.$$

Now recall that $\text{T}_{\mathcal{C}_{i,t}} f_{u,i,t}$ is a function to Δ_m , so it follows that one of its m components has degree d influence at least $\frac{\tau}{2m}$; denote the ℓ th components by $(\text{T}_{\mathcal{C}_{i,t}} f_{u,i,t})_\ell$, and note that by definition of the action of $\text{T}_{\mathcal{C}_{i,t}}$ on vector valued functions, this ℓ th component is equal to $\text{T}_{\mathcal{C}_{i,t}}((f_{u,i,t})_\ell)$. Thus, taking $\beta_0, \tau' > 0$ from Lemma 6.16 for $\frac{\tau}{2m}$, we get that there is $j \in [n]$ such that $I_j^{(\beta)}[(f_{u,i,t})_\ell] \geq \tau'$ (here we use the choice of ζ). Define

$$\text{List}[u] = \left\{ j \in [n] \mid \exists t \in \mathcal{T}', \ell \in [m], i \in [r] \text{ such that } I_j^{(\beta)}[(f_{u,i,t})_\ell] \geq \tau' \right\},$$

then by the above reasoning $\text{List}[u] \neq \emptyset$ for all $u \in \text{Good}[L]$, and by Lemma 6.13 and the union bound we have that $|\text{List}[u]| \leq O_{\beta, |\mathcal{T}'|, m, r, \tau'}(1)$.

Take $\tau'', \beta_1 > 0$ from Lemma 6.21, and for each $v \in R$ denote

$$\text{List}[v] = \left\{ j \in [n/2] \mid \exists t \in \mathcal{T}', \ell \in [m], i \in [r] \text{ such that } I_j^{(\beta)}[f_{v,i,t}] \geq \tau'' \right\}.$$

As before, we have $|\text{List}[v]| \leq O_{\beta, |\mathcal{T}'|, m, r, \tau''}(1)$. Fix $u \in \text{Good}[L]$ and $j \in \text{List}[u]$. Choosing a neighbour $\mathbf{v}$ of u randomly, by the richness of the game, the constraint $\phi_{u,\mathbf{v}}$ is distributed uniformly on $S_{n:n/2}$, and so by

Lemma 6.21 we get that with probability at least τ'' , $I_{\phi_{u,v}(j)}[(f_{u,i,t})_\ell | \phi_{u,v}] \geq \tau''$. But by construction of the $f_{u,i,t}$ we have that $(f_{u,i,t})_\ell | \phi_{u,v} = (f_{v,i,t})_\ell$, so we get that $\phi_{u,v}(j) \in \text{List}[v]$ with probability at least τ'' . We have thus proved that for all $u \in \text{Good}[L]$, choosing a neighbour v of u randomly, with probability at least τ'' the lists $\text{List}[u]$, $\text{List}[v]$ contain a pair of labels that satisfy $\phi_{u,v}$.

We thus use the following randomized strategy: define the assignment to Φ by labeling each $u \in L$ with an element $\sigma \in \text{List}[u]$ chosen uniformly at random if this list is not empty (and otherwise an arbitrary $\sigma \in [n]$), and similarly pick the label of each $v \in R$ using $\text{List}[v]$. It follows that the expected fraction of constraints satisfied by this strategy is at least

$$\Pr_{\mathbf{u} \in U} [\mathbf{u} \in \text{Good}[L]] \cdot \Omega_{\beta, |\mathcal{T}'|, m, r, \tau', \tau''}(1) = \Omega_{\varepsilon, \beta, |\mathcal{T}'|, m, r, \tau', \tau''}(1) > \delta,$$

so there is an assignment to Φ satisfying more than δ of the constraints, and we are done.

Remark 6.28. Using Theorem 5.2 (or rather the analogue of Theorem 4.4 that follows from it) instead of Theorem 4.4, one may relax the connectedness assumption in Theorem 6.7 as in Theorem 6.29 below. Roughly speaking, it is enough that for all $t \in \mathcal{T}$, $P \in \mathcal{P}$, the sequence of distributions $\{\mathcal{D}(t, P)^{\otimes n}\}_{n \in \mathbb{N}}$ and its multi-slice analogue annihilate high degrees functions.

Theorem 6.29. Let $r \in \mathbb{N}$, $m \in \mathbb{N}$, and suppose $\mathcal{P} \subseteq \{P: [m]^r \rightarrow \{0, 1\}\}$ is a collection of predicates. Suppose there is a $(1, s)$ dictatorship test, $(w, \{\mathcal{D}(t, P)\}_{t \in \mathcal{T}}, P \in \mathcal{P}, p)$, for $\mathcal{P}$ satisfying the following properties:

1. for all $t \in \mathcal{T}$, $P \in \mathcal{P}$, the collection of distributions $\{\mathcal{D}(t, P)^{\otimes n}\}_{n \in \mathbb{N}}$ annihilates high degree functions as per Definition 5.1;
2. the multi-slice analogues of $\mathcal{D}(t, P)^{\otimes n}$ also annihilate high-degree functions. That is, for each $t \in \mathcal{T}$, $P \in \mathcal{P}$ and $n \in \mathbb{N}$, considering $k_n(\vec{t}, P) = (k_n(t, P)_{\vec{a}})_{\vec{a} \in [m]^r}$ a list of integers adding up to n such that $|k_n(t, P)_{\vec{a}} - n\mathcal{D}(t, P)(\vec{a})| < 1$ for all $\vec{a} \in [m]^r$, the sequence of uniform distributions on the $k_n(\vec{t}, P)$ multi-slice of $([m]^r)^n$, that is $\left(\mathcal{U}_{k_n(\vec{t}, P)}, \text{Uniform}\right)_{n \in \mathbb{N}}$, annihilates high degree functions.

Then, for all $\varepsilon > 0$, the problem $\text{Gap-}\mathcal{P}[1, s + \varepsilon]$ is NP-hard assuming Conjecture 1.17.

We omit the full proof, as it is a straightforward adaptation of the argument above.

6.5 Proof of Corollaries 1.18, 1.19

In this section, we instantiate Theorem 6.7 to establish Corollaries 1.18, 1.19.

Proof of Corollary 1.18. Fix $\varepsilon > 0$. In [6, Theorem 1.1], the authors have constructed an r -ary predicate $P_r: \{-1, 1\}^r \rightarrow \{0, 1\}$ with $2r + 1$ accepting assignments based on the Hadamard code, and constructed a discrete distribution $\mathcal{E}$ of δ and a distribution $\mathcal{D}_{r, \delta}$ for each δ in the support of $\mathcal{E}$, such that if

$$\mathbb{E}_{\delta \sim \mathcal{E}} \left[\mathbb{E}_{(x(1), \dots, x(r)) \sim \mathcal{D}_{r, \delta}} [P(f(x(1)), \dots, f(x(r)))] \right] \geq \frac{2r + 1}{2^r} + \varepsilon$$

for $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ that is folded⁷, then $\max_{i \in [n]} I_i^{\leq d}[f] \geq \tau$ for $d \in \mathbb{N}$, $\tau > 0$ depending only on ε . Considering the multi-linear expansion of P_r and thereby extending P_r to $[-1, 1]^r$, their proof actually

⁷I.e., $f(x) = -f(-x)$ for all x .

shows the same conclusion for $f: \{0, 1\}^n \rightarrow [-1, 1]$ that is folded, and we will make use of this. We also record here that by [6, Observation 4.1], the marginal of $x(j)_i$ is uniform for all i, j .

We now transform this into a dictatorship test as per Definition 6.6. The measure w is uniform over $\{-1, 1\}^r$, and for each $a \in \{-1, 1\}^r$ we consider the predicate $P_{r,a}(x) = P(a_1x_1, \dots, a_rx_r)$. We take $\mathcal{T} = \text{supp}(\mathcal{E})$, the measure p is the same as the distribution $\mathcal{E}$, and the distribution $\mathcal{D}(\delta, P_{r,a})$ is defined as $a\mathcal{D}_{r,\delta}$, i.e. $(a_1x(1), \dots, a_rx(r))$ where $(x(1), \dots, x(r)) \sim \mathcal{D}_{r,\delta}$.

As all dictators satisfy P_r with probability 1, it is easy to see that the completeness of the above test is 1.

As for the soundness, given a function $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ that passes the dictatorship test with probability at least $\frac{2r+3}{2^r} + \varepsilon$, we have

$$\begin{aligned} \frac{2r+3}{2^r} + \varepsilon &\leq \mathbb{E}_{a,\delta} \left[\mathbb{E}_{(x(1), \dots, x(r)) \sim a\mathcal{D}_{r,\delta}} [P_{a,r}(f(x(1)), \dots, f(x(r)))] \right] \\ &= \mathbb{E}_{a,\delta} \left[\mathbb{E}_{(x(1), \dots, x(r)) \sim \mathcal{D}_{r,\delta}} [P_r(a_1f(a_1x(1)), \dots, a_rf(a_rx(r)))] \right] \\ &= \mathbb{E}_{\delta} \left[\mathbb{E}_{(x(1), \dots, x(r)) \sim \mathcal{D}_{r,\delta}} [P_r(g(x(1)), \dots, g(x(r)))] \right], \end{aligned}$$

where $g(x) = \frac{1}{2}(f(x) - f(-x))$. Note that $g: \{-1, 1\}^n \rightarrow [-1, 1]$ is folded. Thus, by the main result of [6] it follows that there is $i \in [n]$ such that $I_i^{\leq d}[g] \geq \tau$ for d, τ depending only on ε . Looking at the Fourier expression for $I_i^{\leq d}[g]$, it follows that $I_i^{\leq d}[f] \geq \tau$, and thus we have shown that $(w, p, \mathcal{D}(\delta, P_{r,a}))$ forms a dictatorship test with completeness 1 and soundness $\frac{2r+3}{2^r}$.

The result now follows by Theorem 6.7 □

Proof sketch of Corollary 1.19. Consider $\Sigma = \{0, 1, 2\}$ and $P: \Sigma^2 \rightarrow \{0, 1\}$ defined as $P(x, y) = 1_{x \neq y}$. The distributions w, t are trivial, and the distribution $\mathcal{D}(t, P)$ is uniform over $(x, y) \in \{(x, y) \in \Sigma^2 \mid x \neq y\}$. We execute the reduction in Theorem 6.7 to construct a graph H whose vertices correspond to variables of the instance produced, and edges correspond to pairs of variables between which there is a constraint. The completeness is clear, as any dictatorship satisfies $P_a(f(x), f(y)) = 1$ if $(x, y) \in P_a^{-1}(1)^{\otimes n}$.

As for the soundness, if the resulting graph has an independent set of size at least δ fraction of the vertices, then we may consider its indicator function f , and note that for every edge (x, y) in H we have that either $f(x) = 0$ or $f(y) = 1$. As in the analysis therein, we consider the set of vertices $\mathcal{U}$ of the Rich-2-to-1 Games instances for which the average of f inside their cloud is at least $\delta/2$, so that the fractional size of $\mathcal{U}$ is at least $\delta/2$. In other words, letting $Q(x, y) = 1_{x=y=1}$, we have that $\mathbb{E}_{(x,y)} [Q(f_u(x), f_u(y))] = 0$ for all $u \in \mathcal{U}$. Applying our invariance principle, it follows that there is a coupling $\mathcal{C}$ between the multi-slice of u and Σ^n such that $\mathbb{E}_{(x,y) \sim P^{-1}(1)^{\otimes n}} [Q(\mathcal{T}_{\mathcal{C}}f_u(x), \mathcal{T}_{\mathcal{C}}f_u(y))] = o(1)$. From this and Theorem [14, Theorem 3.1] it follows that there is a coordinates i such that $I_i^{\leq d}[\mathcal{T}_{\mathcal{C}}f_u] \geq \tau$ where $d \in \mathbb{N}$, $\tau > 0$ depend only on $\delta > 0$. From here the analysis in the proof of Theorem 6.7 continues as is. □

7 Acknowledgements

We thank an anonymous referee for several helpful comments, and in particular for bringing [32, 15] to our attention.

References

- [1] Noga Alon, Peter Frankl, and László Lovász. The chromatic number of kneser hypergraphs. *Transactions of the American Mathematical Society*, 298(1):359–370, 1986.
- [2] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.
- [3] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998.
- [4] Boaz Barak, Pravesh K. Kothari, and David Steurer. Small-set expansion in shortcode graph and the 2-to-2 conjecture. In *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, pages 9:1–9:12, 2019.
- [5] Amey Bhangale and Subhash Khot. UG-hardness to NP-hardness by losing half. In *34th Computational Complexity Conference, CCC 2019, July 18-20, 2019, New Brunswick, NJ, USA*, pages 3:1–3:20, 2019.
- [6] Amey Bhangale, Subhash Khot, and Devanathan Thiruvengatachari. An improved dictatorship test with perfect completeness. In *37th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2017, December 11-15, 2017, Kanpur, India*, pages 15:1–15:23, 2017.
- [7] Eric Blais and Ryan O’Donnell. Lower bounds for testing function isomorphism. In *2010 IEEE 25th Annual Conference on Computational Complexity*, pages 235–246. IEEE, 2010.
- [8] Mark Braverman, Subhash Khot, Noam Lifshitz, and Dor Minzer. Removal, Ramsey and Turán number for hypergraphs of large uniformity. Forthcoming.
- [9] Mark Braverman, Subhash Khot, and Dor Minzer. On rich 2-to-1 games. In *12th Innovations in Theoretical Computer Science Conference, ITCS 2021, January 6-8, 2021, Virtual Conference*, pages 27:1–27:20, 2021.
- [10] Sourav Chatterjee. A simple invariance theorem. *arXiv preprint math/0508213*, 2005.
- [11] Irit Dinur, Ehud Friedgut, and Oded Regev. Independent sets in graph powers are almost contained in juntas. *Geometric and Functional Analysis*, 18(1):77–97, 2008.
- [12] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. On non-optimally expanding sets in grassmann graphs. In *STOC 2018*, pages 940–951, 2018.
- [13] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. Towards a proof of the 2-to-1 games conjecture? In *STOC 2018*, pages 376–389, 2018.
- [14] Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional hardness for approximate coloring. *SIAM J. Comput.*, 39(3):843–873, 2009.
- [15] Devdatt P Dubhashi and Desh Ranjan. Balls and bins: A study in negative dependence. *BRICS Report Series*, 3(25), 1996.

- [16] David Ellis, Ehud Friedgut, and Haran Pilpel. Intersecting families of permutations. *Journal of the American Mathematical Society*, 24(3):649–682, 2011.
- [17] David Ellis, Nathan Keller, and Noam Lifshitz. Stability for the complete intersection theorem, and the forbidden intersection problem of erdős and sós. *arXiv e-prints*, pages arXiv–1604, 2016.
- [18] Paul Erdős. Problems and results in graph theory and combinatorial analysis.
- [19] Uriel Feige, Shafi Goldwasser, Laszlo Lovász, Shmuel Safra, and Mario Szegedy. Interactive proofs and the hardness of approximating cliques. *J. ACM*, 43(2):268–292, March 1996.
- [20] Vitaly Feldman, Venkatesan Guruswami, Prasad Raghavendra, and Yi Wu. Agnostic learning of monomials by halfspaces is hard. *SIAM Journal on Computing*, 41(6):1558–1590, 2012.
- [21] Yuval Filmus, Guy Kindler, Noam Lifshitz, and Dor Minzer. Hypercontractivity on the symmetric group. *CoRR*, abs/2009.05503, 2020.
- [22] Yuval Filmus, Guy Kindler, Elchanan Mossel, and Karl Wimmer. Invariance principle on the slice. *ACM Trans. Comput. Theory*, 10(3):11:1–11:37, 2018.
- [23] Yuval Filmus and Elchanan Mossel. Harmonicity and invariance on slices of the Boolean cube. *Probability Theory and Related Fields*, 175(3–4):721–782, 2019.
- [24] Yuval Filmus, Ryan O’Donnell, and Xinyu Wu. A log-sobolev inequality for the multislice, with applications. In *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, pages 34:1–34:12, 2019.
- [25] Peter Frankl and Andrey Kupavskii. The Erdős matching conjecture and concentration inequalities. *arXiv preprint arXiv:1806.08855*, 2018.
- [26] Ehud Friedgut and Oded Regev. Kneser graphs are like swiss cheese. *arXiv preprint arXiv:1702.04073*, 2017.
- [27] Badih Ghazi, Pritish Kamath, and Madhu Sudan. Decidability of non-interactive simulation of joint distributions. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 545–554. IEEE, 2016.
- [28] Venkatesan Guruswami and Sai Sandeep. d-to-1 hardness of coloring 3-colorable graphs with $O(1)$ colors. In *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, pages 62:1–62:12, 2020.
- [29] Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. In *37th Annual Symposium on Foundations of Computer Science, FOCS ’96, Burlington, Vermont, USA, 14-16 October, 1996*, pages 627–636, 1996.
- [30] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.
- [31] Sangxia Huang. Approximation resistance on satisfiable instances for predicates with few accepting inputs. In *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 457–466, 2013.

- [32] Kumar Joag-Dev and Frank Proschan. Negative association of random variables with applications. *The Annals of Statistics*, pages 286–295, 1983.
- [33] Peter Keevash, Noam Lifshitz, Eoin Long, and Dor Minzer. Forbidden intersections for codes. *arXiv preprint arXiv:2103.05050*, 2021.
- [34] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the 17th Annual IEEE Conference on Computational Complexity, Montréal, Québec, Canada, May 21-24, 2002*, page 25, 2002.
- [35] 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.
- [36] Subhash Khot, Dor Minzer, and Muli Safra. On independent sets, 2-to-2 games, and Grassmann graphs. In *STOC 2017*, pages 576–589, 2017.
- [37] Subhash Khot, Dor Minzer, and Muli Safra. Pseudorandom sets in grassmann graph have near-perfect expansion. In *FOCS 2018*, pages 592–601, 2018.
- [38] Noam Lifshitz. Hypergraph removal lemmas via robust sharp threshold theorems. *arXiv preprint arXiv:1804.00328*, 2018.
- [39] László Lovász. Kneser’s conjecture, chromatic number, and homotopy. *Journal of Combinatorial Theory, Series A*, 25(3):319–324, 1978.
- [40] Elchanan Mossel. Gaussian bounds for noise correlation of functions. *Geometric and Functional Analysis*, 19(6):1713–1756, 2010.
- [41] Elchanan Mossel. Gaussian bounds for noise correlation of resilient functions. *arXiv preprint arXiv:1704.04745*, 2017.
- [42] Elchanan Mossel, Ryan O’Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: Invariance and optimality. *Annals of Mathematics*, pages 295–341, 2010.
- [43] Ryan O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.
- [44] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 245–254, 2008.
- [45] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998.
- [46] Vladimir I Rotar. Limit theorems for polylinear forms. *Journal of Multivariate analysis*, 9(4):511–530, 1979.

A Facts about multinomial coefficients

Let $H(p_1, \dots, p_m) = \sum_i p_i \log(1/p_i)$ be the entropy function. We need the following standard facts. The first estimates multinomial coefficients with the entropy function and is easy to derive from Stirling approximation, and we omit the proof

Fact A.1. For positive integers $v_1, \dots, v_r$ that sum up to n we have

$$\binom{n}{v_1, \dots, v_r} = \left(1/(2\pi)^{(r-1)/2} + o_n(1)\right) \sqrt{\frac{n}{v_1 \cdots v_r}} 2^{H(\frac{v_1}{n}, \dots, \frac{v_r}{n}) \cdot n}$$

Second, we need the following crude estimate on the difference between the entropy of two distributions that are close point-wise.

Fact A.2. Suppose $p_1, \dots, p_m$ and $q_1, \dots, q_m$ are distributions satisfying that $|p_i - q_i| \leq \varepsilon$ for all i . Then

$$|H(p_1, \dots, p_m) - H(q_1, \dots, q_m)| \leq 4m\varepsilon \log(1/\varepsilon).$$

Proof. We prove that $H(p_1, \dots, p_m) - H(q_1, \dots, q_m) \leq 4m\varepsilon \log(1/\varepsilon)$, and the claim follows by symmetry between p and q . Let $X = \{i \mid p_i \geq 2\varepsilon\}$. Then $\sum_{i \notin X} p_i \log(1/p_i) \leq 2m\varepsilon \log(1/\varepsilon)$, and for any $i \in X$ we have that $q_i \geq \varepsilon$. Thus,

$$\begin{aligned} H(p_1, \dots, p_m) - H(q_1, \dots, q_m) &\leq 2m\varepsilon \log\left(\frac{1}{\varepsilon}\right) + \sum_{i \in X} p_i \log\left(\frac{1}{p_i}\right) - q_i \log\left(\frac{1}{q_i}\right) \\ &\leq 2m\varepsilon \log\left(\frac{1}{\varepsilon}\right) + \sum_{i \in X} p_i \log\left(\frac{1}{p_i}\right) - (p_i - \varepsilon) \log\left(\frac{1}{q_i}\right) \\ &\leq 3m\varepsilon \log\left(\frac{1}{\varepsilon}\right) + \sum_{i \in X} p_i \log\left(\frac{q_i}{p_i}\right) \\ &\leq 3m\varepsilon \log\left(\frac{1}{\varepsilon}\right) + \sum_{i \in X} p_i \left|\frac{q_i}{p_i} - 1\right| \leq 4m\varepsilon \log\left(\frac{1}{\varepsilon}\right), \end{aligned}$$

where in the penultimate inequality we used $\log z \leq |z - 1|$. □

B Facts about product spaces

B.1 Contraction of high degrees

Claim B.1. [Lemma 2.9 in [40]] Let $\alpha > 0$. Suppose μ is a connected measure on $[m] \times [m]$ whose marginal distributions both are equal to ν , and the probability of each atom is at least α . Consider $T_\mu : L^2([m], \nu) \rightarrow L^2([m], \nu)$ defined as

$$T_\mu f(y) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mu} [f(\mathbf{x}) \mid \mathbf{y} = y].$$

Then $\lambda_2(T_\mu) \leq 1 - \alpha^2/2$.

Claim B.2. Let $\mu_1, \dots, \mu_n$ be measures on $[m] \times [m]$, such that

1. For each i , the marginal distribution of μ_i on both coordinates is the same, denote it by ν_i ;
2. $\lambda_2(T_{\mu_i}) \leq 1 - \beta$.

Let $n, d \in \mathbb{N}$, and consider $T : L^2([m]^n, \nu) \rightarrow L^2([m]^n, \nu)$ defined as $T = T_{\mu_1} \otimes \dots \otimes T_{\mu_n}$, where $\nu = \nu_1 \times \dots \times \nu_n$. Then for all $f \in V_{>d}([m]^n, \nu^{\otimes n})$ it holds that

$$\|Tf\|_{2, \nu}^2 \leq (1 - \beta)^d \|f\|_{2, \nu}^2.$$

Proof. Write f by its Efron-Stein decomposition with respect to the measure ν , $f = \sum_{S \subseteq [n], |S| > d} f^{=S}$. Then we have that $\|Tf\|_2^2 = \sum_{S \subseteq [n], |S| > d} \|Tf^{=S}\|_2^2$, so it is enough to show that $\|Tf^{=S}\|_2^2 \leq (1 - \beta)^{2d} \|f^{=S}\|_2^2$.

Assume without loss of generality that $S = \{1, \dots, s\}$. For each j , define $T_{\leq j} = \bigotimes_{i \leq j} T_{\mu_i} \otimes I^{\otimes(n-j)}$, and $g_j = T_{\leq j} f^{=S}$. We argue that $\|g_{j+1}\|_{2,\nu}^2 \leq (1 - \beta)^2 \|g_j\|_{2,\nu}^2$, which by induction finishes the proof. For each j and $z \in [m]^{[n] \setminus \{j+1\}}$, let $g_j[z]: [m] \rightarrow \mathbb{R}$ be the restriction of the function g_j wherein all coordinates, except $j + 1$, have been fixed according to z .

$$\|g_{j+1}\|_{2,\nu}^2 = \mathbb{E}_{\mathbf{z} \sim \prod_{i \neq j+1} \nu_i} \left[\|T_{\mu_{j+1}}(g_j[\mathbf{z}])\|_{2,\nu_{j+1}}^2 \right] \leq \mathbb{E}_{\mathbf{z} \sim \prod_{i \neq j+1} \nu_i} \left[(1 - \beta)^2 \|g_j[\mathbf{z}]\|_{2,\nu_{j+1}}^2 \right] = (1 - \beta)^2 \|g_j\|_{2,\nu}^2.$$

In the second transition, we used the fact that the average of $g_j[z]$ according to ν_{j+1} is 0 for all z . $\square$

Combining the two claims, we get the following lemma.

Lemma B.3. *For all $\alpha > 0$, there is $\varepsilon > 0$ such that the following holds. Let μ be a measure on $[m] \times [m]$, such that*

1. *The marginal distribution of μ on both coordinates is the same, denote it by ν .*
2. *The probability of each atom in μ is at least α .*
3. *μ is connected.*

Let $n, d \in \mathbb{N}$, and consider $T: L^2([m]^n, \nu^{\otimes n}) \rightarrow L^2([m]^n, \nu^{\otimes n})$ defined as $T = T_\mu^{\otimes n}$. Then for all $f \in V_{>d}([m]^n, \nu^{\otimes n})$ it holds that $\|Tf\|_{2,\nu^n}^2 \leq (1 - \beta)^d \|f\|_{2,\nu^n}^2$.

B.2 Influences

In this section, we define influences in product spaces, and state several well-known statements. Our exposition will be very brief and we defer the reader to [43, Chapter 8] for a more thorough presentation. Let $([m]^n, \nu^{\otimes n})$ be a product space. For $f: [m]^n \rightarrow \mathbb{R}$, we define the averaging operator E_i according to the i th coordinate as $E_i f(x) = \mathbb{E}_{\mathbf{y}_i \sim \nu_{\bar{i}}} [f(x_{-i}, \mathbf{y}_i)]$. We then define the Laplacian operator of the i th coordinate as

$$L_i f(x) = f(x) - E_i f(x), \text{ and the Laplacian } L \text{ as } Lf(x) = \sum_{i=1}^n L_i f(x).$$

Definition B.4. *The influence of variable i on $f: ([m]^n, \nu^{\otimes n}) \rightarrow \mathbb{R}$ is $I_i[f] = \|L_i f\|_2^2$. The total influence of f is $I[f] = \sum_{i=1}^n I_i[f]$.*

Fact B.5. $I[f] = \langle f, Lf \rangle$.

We will also use the Efron-Stein decomposition of product spaces. For S , we denote by $V_S([m]^n, \nu^{\otimes n})$ the subspace of functions $f: [m]^n \rightarrow \mathbb{R}$ depending only on the coordinates of S , and by $V_{=S}([m]^n, \nu^{\otimes n}) \subseteq V_S([m]^n, \nu^{\otimes n})$ the subspace of functions that are also orthogonal to all $V_T([m]^n, \nu^{\otimes n})$ for $T \subsetneq S$. Then it is well known that any $f: [m]^n \rightarrow \mathbb{R}$ may be uniquely written as $f = \sum_{S \subseteq [n]} f^{=S}$, where $f^{=S} \in V_{=S}([m]^n, \nu^{\otimes n})$. We also define the degree decomposition as $f^{=d} = \sum_{|S|=d} f^{=S}$, and $f^{\leq d} = f^{=0} + \dots + f^{=d}$.

Definition B.6. The degree d influence of variable i on $f: ([m]^n, \nu^{\otimes n}) \rightarrow \mathbb{R}$ is $I_i^{\leq d}[f] = \|(\mathbf{L}_i f)^{\leq d}\|_2^2$.

Fact B.7. Let $f: [m]^n \rightarrow \mathbb{R}$, $i \in [n]$ and $d \in \mathbb{N}$. Then $I_i[f] = \sum_{S \ni i} \|f^=S\|_2^2$, and $I_i^{\leq d}[f] = \sum_{\substack{S \ni i \\ |S| \leq d}} \|f^=S\|_2^2$.

Claim B.8. Suppose $d \in \mathbb{N}$, $\beta > 0$ are such that $d\beta \leq \frac{1}{8}$, and let μ be a connected distribution on $[m] \times [m]$ whose marginals are both equal to $\nu(a)$, such that $\mu(a, a) \geq (1 - \beta)\nu(a)$. Then for all $f: ([m]^n, \nu^{\otimes n}) \rightarrow \mathbb{R}$ and $i \in [n]$ it holds that $I_i^{\leq d}[f] \leq 2I_i[T_\mu f]$.

Proof. Consider the spaces $V_{=S}([m]^n, \nu^{\otimes n})$ defining the Efron-Stein decomposition of f , and note that both T_μ and $S_\mu = T_\mu^* T_\mu$ preserves each one of them. We may therefore write $V_{=S}([m]^n, \nu^{\otimes n}) = \bigoplus_{\theta} V_{=S}^\theta([m]^n, \nu^{\otimes n})$, where $V_{=S}^\theta([m]^n, \nu^{\otimes n})$ is an eigenspace of S_μ with eigenvalue θ , and note that all eigenvalues θ are non-negative. Thus, we write $f = \sum_{S \subseteq [n], \theta} f^=S, \theta$, where $f^=S, \theta \in V_{=S}^\theta([m]^n, \nu^{\otimes n})$ so by

Fact B.7 and Parseval

$$\begin{aligned} I_i[T_\mu f] &= \sum_{S \ni i} \|T_\mu f^=S\|_2^2 = \sum_{S \ni i} \langle f^=S, S_\mu f^=S \rangle = \sum_{S \ni i, \theta} \sum_{\theta} \langle f^=S, \theta f^=S, \theta \rangle = \sum_{S \ni i, \theta} \theta \|f^=S, \theta\|_2^2 \\ &\geq \sum_{\substack{S \ni i, \theta \\ |S| \leq d}} \theta \|f^=S, \theta\|_2^2, \end{aligned} \quad (23)$$

and we show that $\theta \geq \frac{1}{2}$ for all $|S| \leq d$ and θ such that $V_{=S}^\theta([m]^n, \nu^{\otimes n}) \neq \{0\}$. Indeed, let $g \in V_{=S}^\theta([m]^n, \nu^{\otimes n})$ be non-zero, and let $x \in [m]^n$ be the point maximizing $|g(x)|$. Denote by μ' the distribution corresponding to S_μ (i.e. taking $(\mathbf{a}, \mathbf{b}) \sim \mu$, then $(\mathbf{a}', \mathbf{b}') \sim \mu'$ conditioned on $\mathbf{b}' = \mathbf{b}$ and outputting $(\mathbf{a}, \mathbf{a}')$). Then as g is a S -junta, we have that

$$\begin{aligned} \theta |g(x)| &= |S_\mu g(x)| \\ &\geq \left| \Pr_{(\mathbf{x}, \mathbf{y}) \sim \mu'^{\otimes n}} [\mathbf{y}_S = x_S \mid \mathbf{x} = x] g(x) + \Pr_{(\mathbf{x}, \mathbf{y}) \sim \mu'^{\otimes n}} [\mathbf{y}_S \neq x_S \mid \mathbf{x} = x] \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mu'^{\otimes n}} [g(\mathbf{y}) \mid \mathbf{x} = x, \mathbf{y}_S \neq x_S] \right| \\ &\geq (1 - 2 \Pr_{(\mathbf{x}, \mathbf{y}) \sim \mu'^{\otimes n}} [\mathbf{y}_S \neq x_S \mid \mathbf{x} = x]) |g(x)|. \end{aligned}$$

As by the union bound $\Pr_{(\mathbf{x}, \mathbf{y}) \sim \mu'^{\otimes n}} [\mathbf{y}_S \neq x_S \mid \mathbf{x} = x] \leq 2\beta |S| \leq 2\beta d \leq \frac{1}{4}$, it follows that $\theta \geq \frac{1}{2}$. Plugging this into (23) and using

$$\sum_{\substack{S \ni i, \theta \\ |S| \leq d}} \|f^=S, \theta\|_2^2 = \sum_{\substack{S \ni i \\ |S| \leq d}} \|f^=S\|_2^2 = I_i^{\leq d}[f],$$

where the first transition is by Parseval and the second transition is by Fact B.7, finishes the proof. $\square$

B.3 The noise operator

Let $L^2([m]^n, \nu_1 \times \dots \times \nu_n)$ be a product space.

Definition B.9. For $\rho > 0$, the standard noise operator T_ρ is defined as follows. For each $x \in [m]^n$, we have a distribution $T_\rho x$ defined as: to sample $\mathbf{y} \sim T_\rho x$, for each $i \in [n]$ independently, take $\mathbf{y}_i = x_i$ with probability ρ , and otherwise sample $\mathbf{y}_i \sim \nu_i$.

We then define the operator $T_\rho: L^2([m]^n, \nu_1 \times \dots \times \nu_n) \rightarrow L^2([m]^n, \nu_1 \times \dots \times \nu_n)$ as

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

Fact B.10. Let $f \in V_{=S}([m]^n, \nu_1 \times \dots \times \nu_n)$. Then $T_\rho f = \rho^{|S|} f$.

C Deferred proofs

C.1 Proof of Claim 6.15

The main component in the proof of Claim 6.15 is the following claim.

Claim C.1. Let $J \subseteq [n]$ be of size at most d , let $\mathcal{U}_k \subseteq [m]^n$ be α -balanced, and let $\mathcal{C}$ be a (α, ζ) -coupling between $\mathcal{U}_k \subseteq [m]^n$ and $([m]^n, \nu_k^{\otimes n})$. Then for all $g \in V_{=J}(\mathcal{U}_k)$, we have $I[T_\rho g] \leq \left(\frac{4d^2}{\alpha^2} \zeta n + d^2\right) \|g\|_2^2$.

Proof. Fix i , and denote $h = T_\rho g$. For $i \in J$, we use the trivial bound $I_i[h] \leq \|h\|_2^2 \leq \|g\|_2^2$, and the rest of the argument is devoted to bounding the influence of $i \notin J$.

By Cauchy-Schwarz we have $I_i[h] \leq \mathbb{E}_{\mathbf{y}, \mathbf{y}' \sim \nu_k} [(h(\mathbf{y}_{-i}, \mathbf{y}'_i) - h(\mathbf{y}))^2]$. We sample $\mathbf{y}, \mathbf{y}'$, and then take $\mathbf{y}'' = (\mathbf{y}_{-i}, \mathbf{y}'_i)$. We now sample $\mathbf{x}$ as $(\mathbf{x}, \mathbf{y}(0)) \sim \mathcal{C}$ conditioned on $\mathbf{y}(0) = \mathbf{y}$ and $\mathbf{x}''$ as $(\mathbf{x}'', \mathbf{y}(0)) \sim \mathcal{C}$ conditioned on $\mathbf{y}(0) = \mathbf{y}''$. Then we get that

$$I_i[h] \leq \mathbb{E}_{\mathbf{y}, \mathbf{y}'} \left[\left(\mathbb{E}_{\mathbf{x}, \mathbf{x}''} [g(\mathbf{x}) - g(\mathbf{x}'') \mid \mathbf{y}, \mathbf{y}'] \right)^2 \right] \leq \mathbb{E}_{\substack{\mathbf{y}, \mathbf{y}' \\ \mathbf{x}, \mathbf{x}''}} [(g(\mathbf{x}) - g(\mathbf{x}''))^2].$$

Denote by $\mathbf{j}_1, \dots, \mathbf{j}_{2\ell}$ the set of coordinates on which $\mathbf{x}$ and $\mathbf{x}'$ differ. We stress here that ℓ is also a random variable, and by the third property of couplings its expectation is at most $2n\zeta$.

We may find a matching between them, say $\mathbf{j}_{2s+1}$ and $\mathbf{j}_{2s+2}$ for $s = 0, \dots, \ell - 1$, such that performing the transpositions $\pi_{\mathbf{j}_{2s+1}, \mathbf{j}_{2s+2}}$ in the order $s = 0, \dots, \ell - 1$ on $\mathbf{x}$ gives $\mathbf{x}'$. We then permute these pairs randomly. Denote $\mathbf{x}(0) = \mathbf{x}$, and inductively $\mathbf{x}(s+1) = \pi_{\mathbf{j}_{2s+1}, \mathbf{j}_{2s+2}} \mathbf{x}(s)$. We note that if $\mathbf{j}_{2s+1}, \mathbf{j}_{2s+2} \notin J$, then $g(\mathbf{x}(s+1)) = g(\mathbf{x}(s))$. We thus get

$$\begin{aligned} \mathbb{E}_{\substack{\mathbf{y}, \mathbf{y}' \\ \mathbf{x}, \mathbf{x}''}} [(g(\mathbf{x}) - g(\mathbf{x}''))^2] &= \mathbb{E}_{\substack{\mathbf{y}, \mathbf{y}' \\ \mathbf{x}, \mathbf{x}''}} \left[\left(\sum_{s=0}^{\ell-1} (g(\mathbf{x}(s)) - g(\mathbf{x}(s+1))) 1_{J \cap \{\mathbf{j}_{2s+1}, \mathbf{j}_{2s+2}\} \neq \emptyset} \right)^2 \right] \\ &\leq \mathbb{E}_{\substack{\mathbf{y}, \mathbf{y}' \\ \mathbf{x}, \mathbf{x}''}} \left[\left(\sum_{s=0}^{\ell-1} (g(\mathbf{x}(s)) - g(\mathbf{x}(s+1)))^2 \right) \left(\sum_{s=0}^{\ell-1} 1_{J \cap \{\mathbf{j}_{2s+1}, \mathbf{j}_{2s+2}\} \neq \emptyset} \right) \right] \\ &\leq d \mathbb{E}_{\substack{\mathbf{y}, \mathbf{y}' \\ \mathbf{x}, \mathbf{x}''}} \left[\sum_{s=0}^{\ell-1} (g(\mathbf{x}(s)) - g(\mathbf{x}(s+1)))^2 \right]. \end{aligned} \tag{24}$$

Conditioning, we get that

$$\begin{aligned}
(24) \leq & d \mathbb{E}_\ell \left[\sum_{s=0}^{\ell-1} \mathbb{E}_{\substack{\mathbf{y}, \mathbf{y}' \\ \mathbf{x}, \mathbf{x}''}} \left[(g(\mathbf{x}(s)) - g(\mathbf{x}(s+1)))^2 \mid \ell, \mathbf{j}_{2s+1}, \mathbf{j}_{2s+2} \neq i \right] \right] \\
& + d \mathbb{E}_\ell \left[\mathbb{E}_{\substack{\mathbf{y}, \mathbf{y}' \\ \mathbf{x}, \mathbf{x}''}} \left[\sum_{s=0}^{\ell-1} 1_{i \in \{\mathbf{j}_{2s+1}, \mathbf{j}_{2s+2}\}} (g(\mathbf{x}(s)) - g(\mathbf{x}(s+1)))^2 \mid \ell \right] \right], \tag{25}
\end{aligned}$$

and we analyze each expectation separately. For the first expectation, conditioned on ℓ , by symmetry, for each s , the marginal distribution of $\mathbf{x}(s)$ is uniform over $\mathcal{U}_{\vec{k}}$. Conditioned on $\mathbf{x}(s)$, $\mathbf{j}_{2s+1}, \mathbf{j}_{2s+2} \neq i$, by symmetry, for all j, j' we have that the probability that $\mathbf{j}_{2s+1} = j, \mathbf{j}_{2s+2} = j'$ is at most $\frac{1}{(\alpha n - 1)(\alpha n - 2)}$, thus the first conditional expectation in (25) is at most

$$\frac{2}{\alpha^2} \mathbb{E}_{\substack{\mathbf{j} \in [n], \mathbf{j}' \in [n] \\ \mathbf{x}}} \left[(g(\pi_{\mathbf{j}, \mathbf{j}'} \mathbf{x}) - g(\mathbf{x}))^2 \right] = \frac{2}{n\alpha^2} I[g] \leq \frac{2d}{n\alpha^2} \|g\|_2^2,$$

where we used Claim 6.12. Thus, the contribution of the first term in (25) is at most $\frac{2d^2}{\alpha^2 n} \|g\|_2^2 \mathbb{E}[\ell] \leq \frac{4d^2}{\alpha^2} \zeta \|g\|_2^2$. As for the second term, for at most one s we have that $1_{i \in \{\mathbf{j}_{2s+1}, \mathbf{j}_{2s+2}\}} = 1$; denote this as a random variable $\mathbf{s}$, and note that by symmetry it is uniform over ℓ , so we may as well assume that $\mathbf{s} = 1$. Thus, this expectation is at most $I_i[g]$. We therefore get that $I_i[h] \leq (25) \leq \frac{4d}{\alpha^2} \zeta \|g\|_2^2 + d I_i[g]$, and summing over i we get that

$$I[h] \leq \frac{4d^2}{\alpha^2} \zeta n \|g\|_2^2 + d I[g] \leq \left(\frac{4d^2}{\alpha^2} \zeta n + d^2 \right) \|g\|_2^2,$$

where we used Claim 6.12. □

We are now ready to prove Claim 6.15.

Proof of Claim 6.15. First, note that by Fact B.5

$$I[\text{TC}g] = \langle \text{LT}g, \text{TC}g \rangle = \langle \text{TC}^* \text{LT}g, g \rangle.$$

Denote $R = \text{TC}^* \text{LT}g$. Note that $\mathcal{C}$ is symmetric, and so by Claim 3.2 commutes with S_n , and L also commutes with S_n , we get that $R: L^2(\mathcal{U}_{\vec{k}}) \rightarrow L^2(\mathcal{U}_{\vec{k}})$ commutes with the action of S_n , and therefore by Claim 3.17 preserves degrees. As R is self-adjoint, we may then decompose $V_d(\mathcal{U}_{\vec{k}})$ as a sum of eigenspaces $\bigoplus_\theta V_d^\theta(\mathcal{U}_{\vec{k}})$. Thus, if θ^* is the maximum eigenvalue of R on $V_d(\mathcal{U}_{\vec{k}})$, then we get that $I[\text{TC}g] \leq \theta^* \|g\|_2^2$.

To bound θ^* , note that by Claim 3.29 we get that R preserves juntas and degrees, so $V_{=J}(\mathcal{U}_{\vec{k}})$ are invariant spaces under R . We may therefore write it as a sum of eigenspaces $\bigoplus_\theta V_{=J}^\theta(\mathcal{U}_{\vec{k}})$, and get as we have seen previously that the spaces $\bigoplus_{|J| \leq d} V_{=J}^\theta(\mathcal{U}_{\vec{k}})$ are the eigenspaces of R on $V_d(\mathcal{U}_{\vec{k}})$. Thus, there is $|J| \leq d$ such that $V_{=J}^{\theta^*}(\mathcal{U}_{\vec{k}}) \neq \{0\}$, and we take a non-zero $g^* \in V_{=J}^{\theta^*}(\mathcal{U}_{\vec{k}})$. We now get from Claim C.1 that $\theta^* \leq \frac{I[\text{TC}g^*]}{\|g^*\|_2^2} \leq \frac{4d^2}{\alpha^2} \zeta n + d^2$, and we are done. □