

Spectral asymptotics for contracted tensor ensembles*

Benson Au[†] Jorge Garza-Vargas[‡]

Abstract

Let $\mathbf{T}_{d,N} : \Omega \rightarrow \mathbb{R}^{N^d}$ be a random real symmetric Wigner-type tensor. For unit vectors $(u_N^{(i,j)})_{i \in I, j \in [d-2]} \subset \mathbb{S}^{N-1}$, we study the contracted tensor ensemble

$$\left(\frac{1}{\sqrt{N}} \mathbf{T}_{d,N} \left[u_N^{(i,1)} \otimes \cdots \otimes u_N^{(i,d-2)} \right] \right)_{i \in I}.$$

For large N , we show that the joint spectral distribution of this ensemble is well-approximated by a semicircular family $(s_i)_{i \in I}$ whose covariance $(\mathbf{K}_{i,i'}^{(N)})_{i,i' \in I}$ is given by the rescaled overlaps of the corresponding symmetrized contractions

$$\mathbf{K}_{i,i'}^{(N)} = \frac{1}{d(d-1)} \langle u_N^{(i,1)} \odot \cdots \odot u_N^{(i,d-2)}, u_N^{(i',1)} \odot \cdots \odot u_N^{(i',d-2)} \rangle,$$

which is the true covariance of the ensemble up to a $O_d(N^{-1})$ correction. We further characterize the extreme cases of the variance $\mathbf{K}_{i,i}^{(N)} \in [\frac{1}{d!}, \frac{1}{d(d-1)}]$. Our analysis relies on a tensorial extension of the usual graphical calculus for moment method calculations in random matrix theory, allowing us to access the independence in our random tensor ensemble.

Keywords: random tensor; random matrix; free probability.

MSC2020 subject classifications: Primary 60B20; 46L54, Secondary 46L53; 15B52.

Submitted to EJP on May 19, 2023, final version accepted on August 1, 2023.

Supersedes arXiv:2110.01652v3.

1 Introduction

The storied history of random matrix theory (RMT) has its beginnings in applications: first, in multivariate statistics [34]; then, famously, in nuclear physics [33]. The spectacular progress of RMT in the interim, documented for example by [21, 2, 6, 13], shows a flourishing mathematical discipline with ever-increasing ties to the applied sciences. Indeed, by now, RMT techniques comprise a robust and increasingly influential set of tools for a wide range of disparate fields [1, 7, 23].

In modern applications, one often encounters higher order structures. In this generality, a matrix is replaced by a tensor [18].

*JGV was supported in part by NSF grant CCF-2009011.

[†]University of California, Berkeley. E-mail: bensonau@berkeley.edu

[‡]California Institute of Technology. E-mail: jgarzav@caltech.edu

Definition 1.1 (Tensor). A (real) d -th order tensor $\mathbf{A}_{d,N_1,\dots,N_d}$ of dimension (N_1, \dots, N_d) is a tuple

$$(\mathbf{A}_{d,N_1,\dots,N_d}(k_1, \dots, k_d))_{(k_1,\dots,k_d) \in [N_1] \times \dots \times [N_d]} \in \mathbb{R}^{N_1 \dots N_d} \cong \mathbb{R}^{N_1} \otimes \dots \otimes \mathbb{R}^{N_d}.$$

A tensor of the form $\mathbf{A}_{d,N_1,\dots,N_d} = v_1 \otimes \dots \otimes v_d$ for $v_i \in \mathbb{R}^{N_i}$ is said to be pure. We write $\mathcal{T}_{d,N_1,\dots,N_d}$ for the inner product space of all d -th order tensors of dimension (N_1, \dots, N_d) . If $N_1 = \dots = N_d = N$, then we abbreviate the notation to $\mathbf{A}_{d,N}$ and $\mathcal{T}_{d,N}$ respectively.

For $\sigma \in \mathfrak{S}_d$ and $\mathbf{A}_{d,N} \in \mathcal{T}_{d,N}$, we define the permuted tensor $\mathbf{A}_{d,N}^\sigma \in \mathcal{T}_{d,N}$ by

$$\mathbf{A}_{d,N}^\sigma(k_1, \dots, k_d) := \mathbf{A}_{d,N}(k_{\sigma(1)}, \dots, k_{\sigma(d)}).$$

A tensor $\mathbf{A}_{d,N}$ is symmetric if it is invariant under permutation:

$$\mathbf{A}_{d,N}^\sigma = \mathbf{A}_{d,N}, \quad \forall \sigma \in \mathfrak{S}_d.$$

Similarly, a tensor $\mathbf{A}_{d,N}$ is antisymmetric if

$$\mathbf{A}_{d,N}^\sigma = \text{sgn}(\sigma) \mathbf{A}_{d,N}, \quad \forall \sigma \in \mathfrak{S}_d.$$

We write $\mathcal{S}_{d,N}$ for the subspace of symmetric tensors and $\text{Sym} : \mathcal{T}_{d,N} \rightarrow \mathcal{T}_{d,N}$ for the symmetrization operator

$$\text{Sym}(\mathbf{A}_{d,N}) = \frac{1}{d!} \sum_{\sigma \in \mathfrak{S}_d} \mathbf{A}_{d,N}^\sigma,$$

which is a projection $\text{Sym} = \text{Sym}^2 = \text{Sym}^*$ onto $\mathcal{S}_{d,N}$. For pure tensors, we use the notation

$$v_1 \odot \dots \odot v_d := \text{Sym}(v_1 \otimes \dots \otimes v_d) = \frac{1}{d!} \sum_{\sigma \in \mathfrak{S}_d} v_{\sigma(1)} \otimes \dots \otimes v_{\sigma(d)}.$$

The spectral theory of symmetric tensors for $d \geq 3$ can be quite different from the familiar case of symmetric matrices $d = 2$. Indeed, simply finding the right definitions in this setting already constitutes a major conceptual challenge. We assume hereafter that $d \geq 3$. We focus on one particular definition of eigenvalues for tensors, formulated independently by Lim under the name ℓ^2 -eigenvalues [17] and Qi under the name Z -eigenvalues [28]. We refer the reader to [29] for a comprehensive treatment of the general theory.

Definition 1.2 (Tensor operations). Let $\mathbf{A}_{d,N}$ be a symmetric tensor. For $p \leq d$ and $\mathbf{B}_{p,N} \in \mathcal{T}_{p,N}$, we define the contracted tensor $\mathbf{A}_{d,N}[\mathbf{B}_{p,N}] \in \mathcal{S}_{d-p,N}$ by

$$\mathbf{A}_{d,N}[\mathbf{B}_{p,N}](k_1, \dots, k_{d-p}) = \sum_{(l_1, \dots, l_p) \in [N]^p} \mathbf{A}_{d,N}(k_1, \dots, k_{d-p}, l_1, \dots, l_p) \mathbf{B}_{p,N}(l_1, \dots, l_p).$$

The symmetry of $\mathbf{A}_{d,N}$ implies that this definition does not depend on which coordinates we choose to contract. Consequently,

$$\mathbf{A}_{d,N}[\text{Sym}(\mathbf{B}_{p,N})] = \mathbf{A}_{d,N}[\mathbf{B}_{p,N}].$$

We say that $\mathbf{A}_{d,N}$ has an eigenvalue at $\lambda \in \mathbb{R}$ if there exists a unit vector $u_N \in \mathbb{S}^{N-1}$ such that

$$\mathbf{A}_{d,N}[u_N^{\otimes d-1}] = \lambda u_N,$$

in which case we call u_N an eigenvector.

Contracted tensor ensembles

This definition of an eigenpair $(\lambda, u_N) \in \mathbb{R} \times \mathbb{S}^{N-1}$ can be related to the usual definition for matrices through the following straightforward observation:

$$\mathbf{A}_{d,N}[u_N^{\otimes d-1}] = \lambda u_N \iff \mathbf{A}_{d,N}[u_N^{\otimes d-2}]u_N = \lambda u_N.$$

In other words, λ is an eigenvalue of the real symmetric tensor $\mathbf{A}_{d,N}$ with eigenvector u_N iff λ is an eigenvalue of the real symmetric matrix $\mathbf{A}_{d,N}[u_N^{\otimes d-2}]$ with eigenvector u_N . One can then hope to analyze the spectral properties of a tensor via appropriate contractions.

For a random symmetric tensor $\mathbf{T}_{d,N} : \Omega \rightarrow \mathcal{S}_{d,N}$, de Morais Goulart, Couillet, and Comon defined the *contraction ensemble of $\mathbf{T}_{d,N}$* as the family of random matrices

$$\mathcal{M}(\mathbf{T}_{d,N}) = (\mathbf{T}_{d,N}[u_N^{\otimes d-2}])_{u_N \in \mathbb{S}^{N-1}}.$$

For $d = 3$ and $\mathbf{T}_{3,N}$ distributed according to the Gaussian orthogonal tensor ensemble (GOTE), they proved that the empirical spectral distribution of $\frac{1}{\sqrt{N}}\mathbf{T}_{3,N}[u_N]$ converges weakly almost surely to the semicircle distribution of variance $\frac{1}{6}$ in the large N limit for any sequence of unit vectors $(u_N)_{N \in \mathbb{N}}$ [11, Theorem 6]. In particular, the variance of the limiting semicircle does not depend on the choice of $(u_N)_{N \in \mathbb{N}}$. Their proof relies on an analysis of the resolvent using Stein's lemma, exploiting the Gaussian nature of the entries and the tractability of small order $d = 3$. Naturally, we are led to consider the questions of

- (Q1) higher order $d \geq 4$;
- (Q2) universality for general tensor distributions;
- (Q3) general contractions $u_N^{(1)} \otimes \dots \otimes u_N^{(d-2)} \neq u_N^{\otimes d-2}$;
- (Q4) the joint spectral distribution of the full family of matrices in the contracted tensor ensemble with general contractions.

We emphasize the complications that already arise just in (Q1) and (Q2). While the entries of the GOTE are independent up to the symmetry constraint, the same is not true in general for the contracted GOTE $\mathbf{T}_{d,N}[u_N^{(1)} \otimes \dots \otimes u_N^{(d-2)}]$. For $d = 3$, the contraction operation can only introduce dependence between entries of $\mathbf{T}_{3,N}[u_N]$ belonging to the same row or the same column; for $d \geq 4$, any two entries of the contracted GOTE can now be dependent. Moreover, the resulting correlation structure does not fit into the universality scheme of [9], which would have allowed one to deduce the general result from the Gaussian case. See Section 1.2 for the precise details. To these ends, we consider an analogue of Wigner matrices for tensors.

Definition 1.3 (Wigner tensor). *Let $(X_N(k_1, \dots, k_d))_{N \in \mathbb{N}, 1 \leq k_1 \leq \dots \leq k_d \leq N}$ be a family of independent random variables such that*

- (i) *the off-diagonal entries $\#\(\{k_1, \dots, k_d\}) \neq 1$ are centered;*
- (ii) *the entries with at least three distinct indices $\#\(\{k_1, \dots, k_d\}) \geq 3$ have variance $\binom{d}{b_1, \dots, b_N}^{-1}$, where $b_l = \#\(\{s \in [d] : k_s = l\})$ is the multiplicity of $l \in [N]$ as an index;*
- (iii) *we have a strong uniform control on the moments:*

$$\sup_{N \in \mathbb{N}} \sup_{1 \leq k_1 \leq \dots \leq k_d \leq N} \mathbb{E}[|X_N(k_1, \dots, k_d)|^m] \leq C_m < \infty, \quad \forall m \in \mathbb{N}. \quad (1.1)$$

We call the random symmetric tensor defined by

$$\mathbf{T}_{d,N}(k_1, \dots, k_d) = X_N(k_{i_1}, \dots, k_{i_d}), \quad \forall k_1, \dots, k_d \in [N]$$

an unnormalized d -th order Wigner tensor, where $k_{i_1} \leq \dots \leq k_{i_d}$ is any nondecreasing ordering of the indices k_1, \dots, k_d . Hereafter, when we refer to a tensor $\mathbf{T}_{d,N}$, we implicitly refer to a sequence of tensors $(\mathbf{T}_{d,N})_{N \in \mathbb{N}}$.

Remark 1.4. The variance profile in (ii) is chosen to match the $\text{GOTE}(d, N)$, which is defined by the density

$$f_{d,N}(T) = \frac{1}{Z_d(N)} e^{-\|T\|_2^2/2}, \quad T \in \mathcal{S}_{d,N}.$$

Here, $\|\cdot\|_2$ is the usual Euclidean norm on \mathbb{R}^{N^d} and $Z_d(N)$ is a normalizing constant. The condition in (ii) omits the entries with at most two distinct indices because they do not contribute to the leading order term in the trace expansion. Our results extend to any variance profile that is constant on the integer partition of d determined by $(b_l)_{l \in [N]}$; however, in general, the resulting covariance $(\mathbf{K}_{i,i'}^{(N)})_{i,i' \in I}$ no longer admits a simple formulation as below (cf. Section 1.1.1).

For any sequence of families $(u_N^{(i,j)})_{i \in I, j \in [d-2]} \subset \mathbb{S}^{N-1}$, we define $\mathbf{K}^{(N)} = (\mathbf{K}_{i,i'}^{(N)})_{i,i' \in I}$ as the rescaled Gram matrix of the symmetrized pure tensors $(u_N^{(i,1)} \odot \dots \odot u_N^{(i,d-2)})_{i \in I}$:

$$\mathbf{K}_{i,i'}^{(N)} = \frac{1}{d(d-1)} \langle u_N^{(i,1)} \odot \dots \odot u_N^{(i,d-2)}, u_N^{(i',1)} \odot \dots \odot u_N^{(i',d-2)} \rangle. \quad (1.2)$$

We can now state our main result in answer to questions (Q1)–(Q4).

Theorem 1.5. Let $\mathbf{T}_{d,N}$ be a Wigner tensor. Then the corresponding contracted tensor ensemble $\left(\frac{1}{\sqrt{N}} \mathbf{T}_{d,N} \left[u_N^{(i,1)} \otimes \dots \otimes u_N^{(i,d-2)} \right] \right)_{i \in I}$ converges in noncommutative distribution almost surely iff the limits

$$\mathbf{K}_{i,i'} = \lim_{N \rightarrow \infty} \mathbf{K}_{i,i'}^{(N)}$$

exist for every $i, i' \in I$, in which case the limit object is a semicircular family $(s_i)_{i \in I}$ of covariance $(\mathbf{K}_{i,i'})_{i,i' \in I}$. In particular, the empirical spectral distribution of the single matrix model $\frac{1}{\sqrt{N}} \mathbf{T}_{d,N} \left[u_N^{(i,1)} \otimes \dots \otimes u_N^{(i,d-2)} \right]$ converges weakly almost surely to a semicircle distribution:

$$\frac{1}{N} \sum_{k=1}^N \delta_{\lambda_k} \left(\frac{1}{\sqrt{N}} \mathbf{T}_{d,N} \left[u_N^{(i,1)} \otimes \dots \otimes u_N^{(i,d-2)} \right] \right) \xrightarrow{w} \frac{1}{2\pi \mathbf{K}_{i,i}} (4\mathbf{K}_{i,i} - x^2)_+^{1/2} dx.$$

The proof of Theorem 1.5 relies on a tight control of the mixed moments of the contracted tensor ensemble. To formulate this precisely, we introduce some notation. For any $i_1, \dots, i_m \in I$, we define the product

$$P_N(i_1, \dots, i_m) := \left(\frac{1}{\sqrt{N}} \mathbf{T}_{d,N} \left[u_N^{(i_1,1)} \otimes \dots \otimes u_N^{(i_1,d-2)} \right] \right) \dots \left(\frac{1}{\sqrt{N}} \mathbf{T}_{d,N} \left[u_N^{(i_m,1)} \otimes \dots \otimes u_N^{(i_m,d-2)} \right] \right).$$

The following theorem quantifies the approximation of our tensor ensemble by a semicircular family.

Theorem 1.6. *Let $I_0 \subset I$ be a finite subset. For any $m_0 \in \mathbb{N}$ and $M, \varepsilon > 0$, there is a constant $C = C(d, \#(I_0), m_0, M, \varepsilon)$ independent of N and the choice of contracting vectors $(u_N^{(i,j)})_{i \in I, j \in [d-2]} \subset \mathbb{S}^{N-1}$ such that*

$$\mathbb{P} \left[\max_{m \leq m_0, i_1, \dots, i_m \in I_0} \left| \frac{1}{N} \text{Tr}(P_N(i_1, \dots, i_m)) - \varphi(s_{i_1}^{(N)} \cdots s_{i_m}^{(N)}) \right| > \varepsilon \right] \leq CN^{-M}, \quad (1.3)$$

where $\varphi(s_{i_1}^{(N)} \cdots s_{i_m}^{(N)})$ is the mixed moment of a semicircular family $(s_i^{(N)})_{i \in I}$ of covariance $\mathbf{K}^{(N)}$.

In particular, Theorem 1.6 does *not* assume that the limits $\lim_{N \rightarrow \infty} \mathbf{K}_{i,i'}^{(N)}$ exist (cf. Theorem 1.5). In any case, it is still true that the joint spectral distribution of the contracted tensor ensemble is approximately multivariate semicircular in high dimensions. The only possibility precluding convergence is the joint spectral distribution accumulating around semicircular families with distinct covariance structures (rather than a different type of distribution entirely). Nevertheless, one can use the Cauchy-Schwarz inequality and the fact that $\|\text{Sym}\| \leq 1$ to show that

$$|\mathbf{K}_{i,i'}^{(N)}| \leq \frac{1}{d(d-1)}. \quad (1.4)$$

Thus, one always has convergence to a specific semicircular family along some subsequence. See Figure 1 and Figure 2 for simulations.

Example 1.7 (Common contracting vectors). In the case of common contracting vectors $u_N^{(i,j)} \equiv u_N^{(i)}$ and $u_N^{(i',j)} \equiv u_N^{(i')}$, the inner product in (1.2) amounts to

$$\mathbf{K}_{i,i'}^{(N)} = \frac{1}{d(d-1)} \langle u_N^{(i)}, u_N^{(i')} \rangle^{d-2}.$$

In particular, the marginals of the limiting semicircular family $(s_i)_{i \in I}$ have variance $\mathbf{K}_{i,i} = \frac{1}{d(d-1)}$ regardless of the choice of contracting vectors $(u_N^{(i)})_{N \in \mathbb{N}}$ (cf. [11, Theorem 6]), and asymptotic orthogonality $\langle u_N^{(i)}, u_N^{(i')} \rangle = o(1)$ is equivalent to asymptotic freeness $\mathbf{K}_{i,i'} = 0$ (cf. [31, Theorem 2.6.2]).

Example 1.8 (Extreme cases for the variance). The previous example shows that the upper bound in (1.4) is tight. In fact, this is essentially the only way to achieve the upper bound. Indeed, since Sym is a projection onto $\mathcal{S}_{d-2,N}$,

$$\mathbf{K}_{i,i}^{(N)} = \frac{1}{d(d-1)} \iff u_N^{(i,1)} \odot \cdots \odot u_N^{(i,d-2)} = u_N^{(i,1)} \otimes \cdots \otimes u_N^{(i,d-2)}.$$

In other words, the upper bound is achieved iff the pure tensor is symmetric. For unit vectors $u_N^{(i,j)} \in \mathbb{S}^{N-1}$, this is equivalent to $u_N^{(i,j)} = \pm u_N^{(i,j')}$, a common contracting vector up to sign.

Given that the upper bound for $\mathbf{K}_{i,i}^{(N)}$ is achieved iff the pure tensor is symmetric, one naturally expects a lower bound that is achieved iff the pure tensor is maximally asymmetric in some sense. To obtain such a characterization, we use the permanental representation of the symmetrized inner product [22, Theorem 2.2]:

$$\langle u_N^{(i,1)} \odot \cdots \odot u_N^{(i,d-2)}, u_N^{(i',1)} \odot \cdots \odot u_N^{(i',d-2)} \rangle = \frac{1}{(d-2)!} \text{per} \left[\left(\langle u_N^{(i,j)}, u_N^{(i',k)} \rangle \right)_{j,k=1}^{d-2} \right].$$

When $i = i'$, this is the permanent of the Gram matrix

$$\mathbf{G}_N^{(i)} = \left(\langle u_N^{(i,j)}, u_N^{(i,k)} \rangle \right)_{j,k=1}^{d-2}.$$




Figure 1: Realizations of the empirical spectral distribution of the contracted tensor ensemble in the extreme cases of the variance covered in Example 1.8. We start with a single realization of a Wigner tensor $\mathbf{T}_{4,1000}$ whose upper triangular entries $k_1 \leq k_2 \leq k_3 \leq k_4$ are Rademacher distributed if k_1 is odd and normally distributed if k_1 is even. We use this to construct two contracted tensors $\mathbf{A}_{1000} = \frac{1}{\sqrt{1000}} \mathbf{T}_{4,1000}[u \otimes u]$ and $\mathbf{B}_{1000} = \frac{1}{\sqrt{1000}} \mathbf{T}_{4,1000}[v \otimes w]$, where u is the vector in (1.6), $v(k) \equiv \frac{1}{\sqrt{1000}}$, and $w(k) = (-1)^{k+1}v(k)$. The eigenvalue histogram of \mathbf{A}_{1000} (resp., \mathbf{B}_{1000}) is then plotted against the semicircle density of variance $\frac{1}{d(d-1)} \langle u \odot u, u \odot u \rangle = \frac{1}{12}$ on the left (resp., $\frac{1}{d(d-1)} \langle v \odot w, v \odot w \rangle = \frac{1}{24}$ on the right).

An inequality of Marcus [20, Theorem 2], valid for positive semidefinite matrices, states that

$$\text{per}(\mathbf{G}_N^{(i)}) \geq \prod_{j=1}^{d-2} \mathbf{G}_N^{(i)}(j, j)$$

with equality iff $\mathbf{G}_N^{(i)}$ has a zero row or $\mathbf{G}_N^{(i)}$ is diagonal. Since $\mathbf{G}_N^{(i)}(j, j) \equiv 1$, we see that $\mathbf{K}_N^{(i,i)} \geq \frac{1}{d!}$ with equality iff the contracting vectors $(u_N^{(i,j)})_{j \in [d-2]} \subset \mathbb{S}^{N-1}$ are orthonormal.

Remark 1.9 (Kolmogorov-Smirnov distance). In the single-matrix model $\#(I) = 1$, the uniform bound $\mathbf{K}_{i,i}^{(N)} \in [\frac{1}{d!}, \frac{1}{d(d-1)}]$ on the variances allows us to convert the moment approximation in Theorem 1.6 into the Kolmogorov-Smirnov distance d_{KS} . This follows from a Taylor series expansion of the Fourier transforms in the Berry-Esseen smoothing inequality [14, Chapter XVI.3, Lemma 2]. Consequently, for any $M, \varepsilon > 0$, there is a constant $C = C(d, M, \varepsilon)$ independent of N and the choice of contracting vectors $(u_N^{(j)})_{j \in [d-2]} \subset \mathbb{S}^{N-1}$ such that

$$\mathbb{P} \left[d_{\text{KS}}(\mu_N, \mu_{\text{sc}}^{\mathbf{K}^{(N)}}) > \varepsilon \right] \leq CN^{-M},$$

where μ_N is the empirical spectral distribution of $\frac{1}{\sqrt{N}} \mathbf{T}_{d,N}[u_N^{(1)} \otimes \dots \otimes u_N^{(d-2)}]$ and $\mu_{\text{sc}}^{\mathbf{K}^{(N)}}$ is the semicircle distribution of variance

$$\mathbf{K}^{(N)} = \frac{1}{d(d-1)} \|u_N^{(1)} \odot \dots \odot u_N^{(d-2)}\|_2^2.$$

The proof of Theorem 1.6 relies on the classical moment method in RMT; however, to accommodate the tensorial nature of our matrices, we augment the usual graphical framework for such calculations to include this information. Our diagrammatic approach is well-suited to the dependence structure that appears in our random matrix precisely because it allows us to access the independence in our random tensor. If one forgets the mechanism of this lineage and simply works with the matrix, then this property of the tensor is obscured. Our approach does come at the cost of additional combinatorial considerations, necessitating a bipartite representation, but we emphasize




Figure 2: Realizations of the joint behavior of the contracted tensor ensemble. Using the same matrices \mathbf{A}_{1000} and \mathbf{B}_{1000} from Figure 1, we construct the sum $\mathbf{C}_{1000} = \mathbf{A}_{1000} + \mathbf{B}_{1000}$ and the anticommutator $\mathbf{D}_{1000} = \mathbf{A}_{1000}\mathbf{B}_{1000} + \mathbf{B}_{1000}\mathbf{A}_{1000}$. The orthogonality $\langle u \odot u, v \odot w \rangle = 0$ of the symmetrized contracting vectors implies that the approximating semicircular family in Theorem 1.6 is a *free* semicircular family. The analytic machinery of free probability can then be used to deduce the approximations for \mathbf{C}_{1000} and \mathbf{D}_{1000} . Accordingly, the eigenvalue histogram of \mathbf{C}_{1000} (resp., \mathbf{D}_{1000}) is plotted against the semicircle density of variance $\frac{1}{12} + \frac{1}{24} = \frac{1}{8}$ on the left (resp., the symmetric Poisson distribution of variance $\frac{1}{144}$ [24] on the right).

some additional advantages as well. At a high level, the graphs provide a parallel between contracted Wigner tensors and Wigner matrices, explaining the universality of the semicircle distribution for all orders d . After the initial investment in the single-matrix case, the multi-matrix case follows essentially from the Cauchy-Schwarz inequality. The surprisingly elegant form of the covariance in (1.2) can then be easily read off from the graphs. In contrast, the dependence of the covariance on the shape of the contracting vectors greatly complicates the usual analytic approach via the resolvent, even in the case of $d = 3$ [11, Appendix A]. Finally, we expect that our framework can be applied to other structured tensor calculations. We present this framework and review the relevant aspects from free probability in Section 2. The proofs of the main results are given in Section 3 with the proofs of the technical results deferred to Section 4.

Notation

For convenience, we adopt the following convention for big O notation. Asymptotics will always be in the large N limit. If the implicit constant depends on some parameters, then we indicate this with subscripts. For example, the asymptotic $O_d(N^{-1})$ in the abstract indicates a dependence of the implicit constant on the order d of the tensor.

1.1 Extensions

For the sake of brevity, we restrict our proofs to Wigner tensors as defined in Definition 1.3 and contractions by pure tensors. Here, we collect a number of straightforward extensions.

1.1.1 General variance profiles

Instead of the GOTE variance profile in (ii), suppose that the entries with at least three distinct indices have variance σ_λ^2 , where $\lambda \vdash d$ is the integer partition of d determined by $(b_l)_{l \in [N]}$. To keep track of how the dependence of our tensor entries propagates through repeated contractions, we use a combinatorial object known as a uniform block permutation [26].

Definition 1.10 (Uniform block permutation). Consider two disjoint copies $[n]_L$ and $[n]_R$ of the set $[n] = \{1, \dots, n\}$. We think of $[n]_L$ and $[n]_R$ as a left side and a right side respectively. A uniform block permutation π of $[n]$ is a partition of the set $[n]_L \sqcup [n]_R$ such that each block $B = B_L \sqcup B_R \in \pi$ contains the same number of elements from each side $\#(B_L) = \#(B_R)$, where $B_L = B \cap [n]_L$ and $B_R = B \cap [n]_R$. We write $\mathcal{UBP}(n)$ for the set of uniform block permutations of $[n]$.

For a uniform block permutation $\pi = \{B_1, \dots, B_{\#(\pi)}\} \in \mathcal{UBP}(d-2)$, we define the partition

$$\lambda_\pi := \left(\frac{\#(B_1)}{2}, \dots, \frac{\#(B_{\#(\pi)})}{2}, 1, 1 \right) \vdash d,$$

where $\#(B_1) \geq \dots \geq \#(B_{\#(\pi)})$. Theorems 1.5 and 1.6 still hold for a general variance profile $(\sigma_\lambda^2)_{\lambda \vdash d}$, but with a less tractable formula for the covariance

$$\mathcal{K}_{i,i'}^{(N)} = \sum_{\pi \in \mathcal{UBP}(d-2)} \sigma_{\lambda_\pi}^2 \left(\sum_{\phi: \pi \rightarrow [N]} \prod_{B \in \pi} \prod_{l \in B_L} u_N^{(i,l)}(\phi(B)) \prod_{r \in B_R} u_N^{(i',r)}(\phi(B)) \right). \quad (1.5)$$

Example 1.11 (A constant variance profile). Even in the case of the constant variance profile $\sigma_\lambda^2 \equiv 1$, the covariance in (1.5) can be difficult to compute. If $d = 3$, then the covariance amounts to a single inner product

$$\mathcal{K}_{i,i'}^{(N)} = \langle u_N^{(i)}, u_N^{(i')} \rangle.$$

In particular, the marginals of the limiting family $(s_i)_{i \in I}$ are always standard semicircular $\mathcal{K}_{i,i} = 1$ and asymptotic orthogonality is still equivalent to asymptotic freeness. This is no longer the case for $d \geq 4$ even if we restrict to common contracting vectors. For example, if $d = 4$, the covariance can be written as

$$\mathcal{K}_{i,i'}^{(N)} = 2 \langle u_N^{(i)}, u_N^{(i')} \rangle^2 - \langle (u_N^{(i)})^{\circ 2}, (u_N^{(i')})^{\circ 2} \rangle,$$

where $\cdot^{\circ p}$ denotes the p -th Hadamard product power.

In special cases, the covariance simplifies further. Of particular interest are when the contracting vectors are localized versus delocalized. A straightforward application of (1.5) for $\sigma_\lambda^2 \equiv 1$ yields the following behavior (cf. Example 1.7):

$$\begin{aligned} \|u_N^{(i)}\|_\infty \rightarrow 1 &\implies \mathcal{K}_{i,i'} = \lim_{N \rightarrow \infty} \langle (u_N^{(i)})^{\circ d-2}, (u_N^{(i')})^{\circ d-2} \rangle; \\ \|u_N^{(i)}\|_\infty \rightarrow 0 &\implies \mathcal{K}_{i,i'} = (d-2)! \left(\lim_{N \rightarrow \infty} \langle u_N^{(i)}, u_N^{(i')} \rangle \right)^{d-2}; \\ \text{supp}(u_N^{(i)}) \cap \text{supp}(u_N^{(i')}) = \emptyset &\implies \mathcal{K}_{i,i'}^{(N)} = 0. \end{aligned}$$

Note that the first two cases do not make any assumptions on the shape of the contracting vector $u_N^{(i')}$ used to construct the covarying matrix $\frac{1}{\sqrt{N}} \mathbf{T}_{d,N}[(u_N^{(i')})^{\otimes d-2}]$.

The variance profile $\sigma_\lambda^2 \equiv 1$ does admit an alternative formulation for the covariances $(\mathcal{K}_{i,i'}^{(N)})_{i,i' \in I}$ that can be more tractable computationally. In particular, by considering $(\mathcal{UBP}(d-2), \leq)$ as a poset for the reversed refinement order \leq , we can use the corresponding Möbius function $\text{Möb}_{\mathcal{UBP}}$ to rewrite (1.5) as

$$\mathcal{K}_{i,i'}^{(N)} = \sum_{\rho \in \mathcal{UBP}(d-2)} \left(\sum_{\substack{\pi \in \mathcal{UBP}(d-2) \\ \text{s.t. } \rho \leq \pi}} \text{Möb}_{\mathcal{UBP}}(\rho, \pi) \right) \prod_{B \in \rho} \langle \circ_{l \in B_L} u_N^{(i,l)}, \circ_{r \in B_R} u_N^{(i',r)} \rangle.$$

1.1.2 General contractions

We can also consider contractions by general unit tensors $(\mathbf{U}_{d-2,N}^{(i)})_{i \in I} \subset \mathbb{S}^{N^{d-2}-1}$. For the GOTE variance profile, the new formula for the covariance follows essentially from linearity:

$$\mathcal{K}_{i,i'}^{(N)} = \frac{1}{d(d-1)} \langle \text{Sym}(\mathbf{U}_{d-2,N}^{(i)}), \text{Sym}(\mathbf{U}_{d-2,N}^{(i')}) \rangle.$$

One needs to take care since the contracting tensor $\mathbf{U}_{d-2,N}^{(i)}$ may require an increasing number of basis elements to write as $N \rightarrow \infty$, which is why the result does not follow directly from linearity. The modification of the pictures in Section 2.1 and the corresponding calculations is straightforward but tedious. As before, the upper bound $\mathcal{K}_{i,i}^{(N)} \leq \frac{1}{d(d-1)}$ for the variance is achieved iff $\mathbf{U}_{d-2,N}^{(i)}$ is symmetric. On the other hand, the lower bound from Example 1.8 no longer holds. Indeed, a natural candidate for maximally asymmetric in this context would be an antisymmetric tensor $\mathbf{U}_{d-2,N}^{(i)}$, for which $\text{Sym}(\mathbf{U}_{d-2,N}^{(i)}) = \mathbf{0}_{d-2,N}$ is the zero tensor and $\mathbf{T}_{d,N}[\mathbf{U}_{d-2,N}^{(i)}] \equiv \mathbf{0}_{2,N}$ is the zero matrix.

1.1.3 Independent Wigner tensors

A key aspect of the contracted tensor ensemble is the singular source of randomness coming from the ambient tensor $\mathbf{T}_{d,N}$. One can also consider contractions $(\mathbf{U}_{d-2,N}^{(i,j)})_{i \in I, j \in J} \subset \mathbb{S}^{N^{d-2}-1}$ of independent Wigner tensors $(\mathbf{T}_{d,N}^{(j)})_{j \in J}$, in which case one obtains the expected answer based on the well-established principle of free independence emerging in the large N limit of suitably generic independent random matrices. The augmented covariance matrix $(\mathcal{K}_{i,i',j,j'}^{(N)})_{i,i' \in I, j,j' \in J}$ now takes the form

$$\mathcal{K}_{i,i',j,j'}^{(N)} = \frac{\mathbb{1}\{j=j'\}}{d(d-1)} \langle \text{Sym}(\mathbf{U}_{d-2,N}^{(i)}), \text{Sym}(\mathbf{U}_{d-2,N}^{(i')}) \rangle.$$

1.2 Comparison with existing results

Of course, one can simply view a contracted Wigner tensor as a generalization of a Wigner matrix where one allows for a dependence structure between the entries. Results in this direction naturally depend on the particulars of the dependence structure. To the best of our knowledge, nothing in the existing literature covers the case of a contracted Wigner tensor, even in the single-matrix model. For concreteness, consider a Wigner tensor $\mathbf{T}_{4,N}$ with suitably scaled Rademacher entries under the common contracting vector

$$u_N = \left(0, 0, \frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3(N-4)}}, \dots, \frac{1}{\sqrt{3(N-4)}} \right). \tag{1.6}$$

If one partitions $\mathbf{T}_{4,N}[u_N^{\otimes 2}]$ so that entries across different blocks are independent, one necessarily obtains a block B of size $\#(B) \sim N^2$, which rules out the partitioning scheme in [30]. Moreover, knowing that the entries other than the (3, 3)-th entry of $\mathbf{T}_{4,N}[u_N^{\otimes 2}]$ have each attained their maximal value forces the remaining entry $\mathbf{T}_{4,N}[u_N^{\otimes 2}](3, 3)$ to attain its maximal value as well, which rules out the conditional centeredness condition in [15]. Next, a straightforward calculation shows that

$$\begin{aligned} \text{Cov}(\mathbf{T}_{4,N}[u_N^{\otimes 2}](1, 1), \mathbf{T}_{4,N}[u_N^{\otimes 2}](2, 1)) &= 0; \\ \text{Cov}(\mathbf{T}_{4,N}[u_N^{\otimes 2}](3, 3), \mathbf{T}_{4,N}[u_N^{\otimes 2}](4, 3)) &= \frac{7}{54}, \end{aligned}$$

which rules out the functional representation in [9] and the approximate uncorrelation scheme in [16]. Finally, a similar calculation shows that

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left[\sum_{k \in [N]} \mathbf{T}_{4,N}[u_N^{\otimes 2}](3, k) \mathbf{T}_{4,N}[u_N^{\otimes 2}](4, k) \right] = \frac{1}{18},$$

which rules out the isotropy condition in [8].

2 Background

2.1 Graphs of tensors

Diagrammatic representations for tensors date back to the work of Penrose [27]. Roughly speaking, in Penrose graphical notation, tensors are represented by shapes such as boxes, circles, and triangles; the corresponding indices are notated by lines emanating from the shapes. If a line connects two shapes, then the corresponding index is summed over in both tensors.

In contrast, the RMT convention represents matrices by edges and indices by vertices. One can then keep track of repetitions in the indices by identifying vertices. This bookkeeping allows one to identify leading order terms when computing trace expansions. Such considerations form the basis of a noncommutative probability theory known as *traffic probability* [19, 10, 4]. Of course, one can simply view the contracted tensor $\mathbf{T}_{d,N}[u_N^{(i,1)} \otimes \dots \otimes u_N^{(i,d-2)}]$ as a random matrix and represent it in the usual way; however, the graph then fails to keep track of the dependencies between the entries introduced by the contractions.

We circumvent this issue by amalgamating the two representations. In particular, we work with a bipartite graph with a class of vertices V for the indices and a class of vertices W for the tensors. It will be convenient to further separate the tensors $W = W_1 \sqcup W_2$ into the d -th order tensors W_1 and the contracting vectors W_2 . In our diagrams, we use circles for the indices, diamonds for the tensors, and boxes for the contracting vectors. For example,

$$= \mathbf{T}_{4,N}(k_1, k_2, l_1, l_2) u_N^{(1)}(l_1) u_N^{(2)}(l_2)$$

The symmetry of our tensor $\mathbf{T}_{d,N}$ ensures that this representation is well-defined. We adopt the convention that any circular vertex without a label indicates a summation over all possible indices. For example,


$$= \sum_{l_1 \in [N]} \sum_{l_2 \in [N]} \mathbf{T}_{4,N}(k_1, k_2, l_1, l_2) u_N^{(1)}(l_1) u_N^{(2)}(l_2)$$

$$= \mathbf{T}_{4,N}[u_N^{(1)} \otimes u_N^{(2)}](k_1, k_2)$$

The (unnormalized) trace of a general m -fold product

$$\text{Tr}(\mathbf{T}_{d,N}[u_N^{(i_1,1)} \otimes \dots \otimes u_N^{(i_1,d-2)}] \dots \mathbf{T}_{d,N}[u_N^{(i_m,1)} \otimes \dots \otimes u_N^{(i_m,d-2)}])$$

can then be represented as a decorated cycle



Let $G = (V, W, E)$ be the graph above, where for convenience we suppress the dependence on m . Recall that V is the set of (circular) index vertices and W is the set of (polygonal) tensor vertices. We use \sim to denote adjacency in G . The expected normalized trace associated to the graph G can then be written as

$$\tau[G] := \frac{1}{N} \sum_{\phi: V \rightarrow [N]} \mathbb{E} \left[\prod_{w_1 \in W_1} \mathbf{T}_{d,N}(w_1 | \phi) \right] \prod_{w_2 \in W_2} u_N(w_2 | \phi),$$

where

$$\begin{aligned} \mathbf{T}_{d,N}(w_1 | \phi) &:= \mathbf{T}_{d,N} \left(\prod_{\substack{v \in V: \\ v \sim w_1}} \phi(v) \right); \\ u_N(w_2 | \phi) &:= u_N^{(w_2)}(\phi(v^{w_2})), \end{aligned}$$

and v^{w_2} denotes the unique vertex in V adjacent to w_2 . In other words, the notation $(w | \phi)$ is shorthand for the ϕ -labels of the neighbors of w . By a slight abuse of notation, we assume that the superscript in $u_N^{(w_2)}$ identifies the contracting unit vector $u_N^{(i,j)}$.

Similarly, we define

$$\tau^0[G] := \frac{1}{N} \sum_{\phi: V \hookrightarrow [N]} \mathbb{E} \left[\prod_{w_1 \in W_1} \mathbf{T}_{d,N}(w_1 | \phi) \right] \prod_{w_2 \in W_2} u_N(w_2 | \phi), \tag{2.2}$$

where $\phi : V \hookrightarrow [N]$ denotes an injective map. Note that the functions τ and τ^0 can be extended to arbitrary graphs of tensors in the obvious way. For any such graph, which we again denote by $G = (V, W, E)$, we write $\mathcal{P}(V)$ for the set of partitions of the index

vertices V . For a partition $\pi \in \mathcal{P}(V)$, we define the *quotient graph* G^π as the graph obtained from G by identifying vertices in the same block in π . Formally, $G^\pi = (V^\pi, W, E)$, where

$$V^\pi := V/\sim_\pi = \{B : B \in \pi\} = \{[v]_\pi : v \in V\}.$$

This allows us to formulate the relation

$$\tau[G] = \sum_{\pi \in \mathcal{P}(V)} \tau^0[G^\pi]. \tag{2.3}$$

In the context of matrices, the function τ (resp., τ^0) is known as the *traffic state* (resp., *injective traffic state*) [19].

To use equation (2.3), we need to calculate τ^0 on all possible quotients of G . The advantage of this comes from the fact that the injective labelings $\phi : V^\pi \hookrightarrow [N]$ will allow us to read off the dependence structure of the tensor entries directly from the graph G^π . To recognize the corresponding limit object, we use a characterization of semicircular families implicit in [10] (see also [5, Section 3]). We review this and other free probabilistic notions in the next section.

2.2 Free probability

We content ourselves with a narrow review of the free probability machinery. The interested reader should consult [31, 25, 23] for further details.

Definition 2.1 (Noncommutative probability). *A noncommutative probability space (ncps) is a unital algebra \mathcal{A} over \mathbb{C} paired with a unital linear functional $\varphi : \mathcal{A} \rightarrow \mathbb{C}$. For a family of random variables $\mathbf{a} = (a_i)_{i \in I} \subset \mathcal{A}$, their (joint) distribution is defined as the linear functional*

$$\mu_{\mathbf{a}} : \mathbb{C}\langle \mathbf{x} \rangle \rightarrow \mathbb{C}, \quad P \mapsto \varphi(P(\mathbf{a})),$$

where $\mathbb{C}\langle \mathbf{x} \rangle$ is the free algebra over the indeterminates $\mathbf{x} = (x_i)_{i \in I}$ and $P(\mathbf{a}) \in \mathcal{A}$ is the usual evaluation of noncommutative polynomials. We say that a sequence of families $\mathbf{a}_N = (a_N^{(i)})_{i \in I} \subset (\mathcal{A}_N, \varphi_N)$ converges in distribution if the corresponding sequence of linear functionals $(\mu_{\mathbf{a}_N})_{N \in \mathbb{N}}$ converges pointwise. Note that the limit defines a new ncps $(\mathbb{C}\langle \mathbf{x} \rangle, \lim_{N \rightarrow \infty} \mu_{\mathbf{a}_N})$.

In other words, the distribution in this setting is the information of the noncommutative moments. Just as classical moments have a parameterization in terms of cumulants, noncommutative moments have a parameterization in terms of *free cumulants*.

Definition 2.2 (Free independence). *Let $\mathcal{NC}(n)$ denote the set of noncrossing partitions of $[n]$. We write \leq for the reversed refinement order on $\mathcal{NC}(n)$ and $\text{Möb}_{\mathcal{NC}}$ for the corresponding Möbius function. Note that the minimal element 0_n in $(\mathcal{NC}(n), \leq)$ consists of singletons, whereas the maximal element 1_n is itself a singleton.*

A noncrossing partition $\pi \in \mathcal{NC}(n)$ defines a multilinear functional

$$\varphi_\pi : \mathcal{A}^n \rightarrow \mathbb{C}, \quad (a_1, \dots, a_n) \mapsto \prod_{B \in \pi} \varphi(B)[a_1, \dots, a_n],$$

where a block $B = (i_1 < \dots < i_m) \in \pi$ splits the n -tuple (a_1, \dots, a_n) into a lower order moment

$$\varphi(B)[a_1, \dots, a_n] := \varphi(a_{i_1} \cdots a_{i_m}).$$

The free cumulant $\kappa_\pi : \mathcal{A}^n \rightarrow \mathbb{C}$ is obtained from the Möbius convolution

$$\kappa_\pi[a_1, \dots, a_n] := \sum_{\substack{\sigma \in \mathcal{NC}(n) \\ \text{s.t. } \sigma \leq \pi}} \varphi_\sigma[a_1, \dots, a_n] \text{Möb}_{\mathcal{NC}}(\sigma, \pi).$$

By the Möbius inversion formula, this is equivalent to

$$\varphi_\pi[a_1, \dots, a_n] = \sum_{\substack{\sigma \in \mathcal{NC}(n) \\ \text{s.t. } \sigma \leq \pi}} \kappa_\sigma[a_1, \dots, a_n]. \tag{2.4}$$

The free cumulants inherit both multilinearity and multiplicativity with respect to the blocks. In particular,

$$\kappa_\pi[a_1, \dots, a_n] = \prod_{B \in \pi} \kappa(B)[a_1, \dots, a_n],$$

where $B = (i_1 < \dots < i_m)$ is a block as before and

$$\kappa(B)[a_1, \dots, a_n] := \kappa_{1_m}[a_{i_1}, \dots, a_{i_m}].$$

For notational convenience, we write $\kappa_m = \kappa_{1_m}$.

Subsets $(\mathcal{S}_i)_{i \in I}$ of a ncps (\mathcal{A}, φ) are freely independent if mixed free cumulants on the $(\mathcal{S}_i)_{i \in I}$ vanish: for any a_1, \dots, a_n such that $a_j \in \mathcal{S}_{i(j)}$,

$$\exists i(j) \neq i(k) \implies \kappa_n[a_1, \dots, a_n] = 0.$$

A sequence of families $\mathbf{a}_N \subset (\mathcal{A}_N, \varphi_N)$ is asymptotically free if it converges in distribution to a limit \mathbf{x} that is free in $(\mathbb{C}\langle \mathbf{x} \rangle, \lim_{N \rightarrow \infty} \mu_{\mathbf{a}_N})$.

The noncommutative analogue of the multivariate normal distribution then follows from the noncrossing analogue of Wick's formula.

Definition 2.3 (Semicircular family). Let $(\mathbf{K}_{i,i'})_{i,i' \in I}$ be a real positive semidefinite matrix. A (centered) semicircular family of covariance \mathbf{K} is a family of random variables $s = (s_i)_{i \in I}$ in a ncps (\mathcal{A}, φ) such that

$$\kappa_2(s_i, s_{i'}) = \mathbf{K}_{i,i'}, \quad \forall i, i' \in I,$$

with all other cumulants on s vanishing. By the moment-cumulant relation (2.4), this is equivalent to

$$\varphi(s_{i(1)} \cdots s_{i(n)}) = \sum_{\pi \in \mathcal{NC}_2(n)} \prod_{\{j,k\} \in \pi} \mathbf{K}_{i(j),i(k)}, \tag{2.5}$$

where $\mathcal{NC}_2(n) \subset \mathcal{NC}(n)$ is the subset of noncrossing pair partitions of $[n]$. In particular, the $(s_i)_{i \in I}$ are free iff $(\mathbf{K}_{i,i'})_{i,i' \in I}$ is diagonal, in which case we call $(s_i)_{i \in I}$ a free semicircular family. If $\mathbf{K}_{i,i'} = \mathbb{1}\{i = i'\}$, then we call $(s_i)_{i \in I}$ a standard semicircular family.

Example 2.4. The algebra $L^{\infty-}(\Omega, \mathcal{F}, \mathbb{P}) \otimes \text{Mat}_N(\mathbb{C})$ of random $N \times N$ matrices whose entries have finite moments of all orders equipped with the expected trace $\frac{1}{N} \mathbb{E}[\text{Tr}(\cdot)]$ defines a ncps. A well-known result of Dykema [12, Theorem 2.1], building on earlier work of Voiculescu [32], proves that independent Wigner matrices $(\mathbf{W}_N^{(i)})_{i \in I} \subset (L^{\infty-}(\Omega, \mathcal{F}, \mathbb{P}) \otimes \text{Mat}_N(\mathbb{C}), \frac{1}{N} \mathbb{E}[\text{Tr}(\cdot)])$ converge in distribution to a standard semicircular family.

For random matrices $(\mathbf{A}_N^{(i)})_{i \in I}$, one can also define convergence in distribution almost surely.

Definition 2.5. We say that $(\mathbf{A}_N^{(i)})_{i \in I}$ converges in distribution almost surely if for any finite subset $I_0 \subset I$, the sequence of random functionals

$$\mu_{(\mathbf{A}_N^{(i)})_{i \in I_0}} : \mathbb{C}\langle \mathbf{x} \rangle \times (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow \mathbb{C}, \quad P \mapsto \frac{1}{N} \text{Tr}[P(\mathbf{A}_N^{(i)} : i \in I_0)]$$

converges pointwise almost surely.

In special cases, one can read off the free cumulants from the injective traffic state. To see this, we recall the notion of a *cactus graph*, a connected graph in which every edge belongs to a unique simple cycle. If the only quotients of a cycle that contribute in the injective traffic state are cactus graphs and the contributions are multiplicative with respect to the cycles of the cactus, then one can use a bijection between cactus graph quotients of a cycle of length n and noncrossing partitions of $[n]$ to read off the free cumulants from the cycles of the cactus [10, Section 5] (see also [5, Section 3]). A semicircular family can then be identified by the property that only cactus graphs with cycles of length two contribute (so-called *double trees* [19]). We review this characterization in the next section. The picture for our contracted tensor ensemble is complicated by the additional vertices coming from the contracting vectors. Nevertheless, we show how the double tree criteria can still be adapted to this situation in Section 3.2.

2.3 The double tree criteria for $d = 2$

We briefly review the $d = 2$ case to identify key properties that are sufficient to ensure convergence in distribution to a semicircular family. The techniques will serve as a guide for higher order d , where more sophisticated arguments are required to deal with the technical complications that arise when considering general tensors.

We start by noting that our picture for the trace of a general m -fold product in (2.1) applies equally well to the $d = 2$ case of random matrices $(\mathbf{W}_N^{(i)})_{i \in I}$. Here, we no longer have the outer vertices decorating the cycle since there are $d - 2 = 0$ contractions. We can then simplify our graph G by discarding the diamonds, which were only necessary to keep track of the additional indices in our d -th order tensor, and represent our 2-tensor with a single labeled edge:



Let $H = (\mathcal{V}(H), \mathcal{E}(H))$ be the resulting graph: by construction, H is a simple m -cycle in which each edge represents a matrix in $(\mathbf{W}_N^{(i)})_{i \in I}$. In fact, the picture above shows how we can identify $(\mathcal{V}(H), \mathcal{E}(H)) = (V, W)$ with the vertices of our original graph $G = (V, W, E)$. We use H to formulate two convenient properties that together guarantee convergence to a semicircular family.

Proposition 2.6. *Let $(\mathbf{W}_N^{(i)})_{i \in I}$ be a family of random real symmetric matrices such that*

(P1) *(Supported on double trees)*

$$\lim_{N \rightarrow \infty} \tau^0[G^\pi] = 0$$

whenever H^π is not a double tree;

(P2) *(Multiplicativity) there is a real positive semidefinite matrix $(\mathbf{K}_{i,i'})_{i,i' \in I}$ such that*

$$\lim_{N \rightarrow \infty} \tau^0[G^\pi] = \prod_{\{i,i'\}} \mathbf{K}_{i,i'}$$

whenever H^π is a double tree, where the product runs over the double edges $\{i, i'\}$ in H^π .

Then $(\mathbf{W}_N^{(i)})_{i \in I}$ converges in distribution to a semicircular family $(s_i)_{i \in I}$ of covariance \mathbf{K} .

The proposition follows from a one-to-one correspondence between noncrossing partitions and cactus graph quotients of simple cycles [10, Lemma 5.2] (see also [5,




Figure 3: An example of Lemma 2.7 in action. On the left, the blocks of π are indicated by the coloring of the vertices; the dotted lines connect the pairs in the corresponding Kreweras complement $\mathcal{K}(\pi)$. The vertex identifications produce the quotient H^π on the right, a double tree as guaranteed by the lemma.

Proposition 2.4]). For the convenience of the reader, we outline this correspondence in the special case of double trees. We start by enumerating the vertices $\{v_1, \dots, v_m\}$ of the m -cycle H in counterclockwise order. Similarly, we enumerate the edges $\{e_1, \dots, e_m\}$ so that $v_i \stackrel{e_i}{\sim} v_{i+1}$ with the convention $v_{m+1} = v_1$. The map $[n] \ni i \mapsto v_i \in \mathcal{V}(H)$ defines a one-to-one correspondence between $\mathcal{NC}(n)$ and $\mathcal{NC}(\mathcal{V}(H))$. Similarly, for $[\bar{n}] = \{\bar{1} < \dots < \bar{n}\}$, the map $[\bar{n}] \ni \bar{i} \mapsto e_i \in \mathcal{E}(H)$ defines a one-to-one correspondence between $\mathcal{NC}(\bar{n})$ and $\mathcal{NC}(\mathcal{E}(H))$. By considering the interlacing

$$1 < \bar{1} < \dots < n < \bar{n},$$

we can transport the Kreweras complement [25, Definition 9.21]

$$\text{Kr} : \mathcal{NC}(n) \rightarrow \mathcal{NC}(\bar{n})$$

to a map

$$\mathcal{K} : \mathcal{NC}(\mathcal{V}(H)) \rightarrow \mathcal{NC}(\mathcal{E}(H)).$$

We will need the following lemma, which is a special case of [5, Proposition 2.4]. See Figure 3 for an illustration.

Lemma 2.7 (Double trees and noncrossing pairings). *The graph H^π is a double tree iff $\pi \in \mathcal{NC}(\mathcal{V}(H))$ and $\mathcal{K}(\pi) \in \mathcal{NC}_2(\mathcal{E}(H))$. Moreover, if $\mathcal{K}(\pi) \in \mathcal{NC}_2(\mathcal{E}(H))$, then two edges are paired together iff their incident vertices are identified head-to-tail:*

$$\{e_i, e_{i'}\} \in \mathcal{K}(\pi) \iff [v_i]_\pi = [v_{i'+1}]_\pi \quad \text{and} \quad [v_{i+1}]_\pi = [v_{i'}]_\pi.$$

We can now give a short proof of Proposition 2.6.

Proof of Proposition 2.6. If G is the undecorated m -cycle in (2.1) with $\mathbf{W}_N^{(i_1)}, \dots, \mathbf{W}_N^{(i_m)}$ in counterclockwise order, then

$$\tau[G] = \frac{1}{N} \mathbb{E}[\text{Tr}(\mathbf{W}_N^{(i_1)} \dots \mathbf{W}_N^{(i_m)})].$$

Combining Lemma 2.7 with properties (P1) and (P2), we obtain

$$\begin{aligned} \lim_{N \rightarrow \infty} \tau[G] &= \sum_{\pi \in \mathcal{P}(V)} \lim_{N \rightarrow \infty} \tau^0[G^\pi] \\ &= \sum_{\substack{\pi \in \mathcal{NC}(\mathcal{V}(H)) \text{ s.t.} \\ \sigma = \mathcal{K}(\pi) \in \mathcal{NC}_2(\mathcal{E}(H))}} \lim_{N \rightarrow \infty} \tau^0[G^\pi] \\ &= \sum_{\sigma \in \mathcal{NC}_2(\mathcal{E}(H))} \prod_{\{i, i'\} \in \sigma} \mathbf{K}_{i, i'}, \end{aligned}$$

which is precisely the characterization of a semicircular family of covariance \mathbf{K} given in (2.5). \square

3 Proofs of the main results

3.1 Setup and complications for $d \geq 3$

Let $G^\pi = (V^\pi, W, E)$ be a quotient of the graph G in (2.1). The analogue of (2.2) for the *normalized* contracted tensor ensemble

$$\left(\frac{1}{\sqrt{N}} \mathbf{T}_{d,N} \left[u_N^{(i,1)} \otimes \cdots \otimes u_N^{(i,d-2)} \right] \right)_{i \in I}$$

can be written as

$$\tau^0[G^\pi] = \frac{1}{N^{\frac{m}{2}+1}} \sum_{\phi: V^\pi \hookrightarrow [N]} \mathbb{E} \left[\prod_{w_1 \in W_1} \mathbf{T}_{d,N}(w_1|\phi) \right] \prod_{w_2 \in W_2} u_N(w_2|\phi).$$

Since $|u_N^{(i,j)}(k)| \leq \|u_N^{(i,j)}\|_2 = 1$, our strong moment assumption (1.1) bounds

$$\mathbb{E} \left[\prod_{w_1 \in W_1} \mathbf{T}_{d,N}(w_1|\phi) \right] \prod_{w_2 \in W_2} u_N(w_2|\phi) = O_m(1) \tag{3.1}$$

uniformly in (π, ϕ) . We immediately arrive at the crude asymptotic

$$\tau^0[G^\pi] = O_m(N^{|V^\pi| - \frac{m}{2} - 1}); \tag{3.2}$$

however, in many cases, the independence and centeredness of our tensor entries result in

$$\mathbb{E} \left[\prod_{w_1 \in W_1} \mathbf{T}_{d,N}(w_1|\phi) \right] = 0, \quad \forall \phi: V^\pi \hookrightarrow [N]. \tag{3.3}$$

To understand when this happens, we introduce some notation. For any $w_1 \in W_1$, we define $\mathcal{B}_\pi(w_1)$ to be the set of neighbors of w_1 in G^π with the additional information of edge multiplicity. Formally,

$$\mathcal{B}_\pi(w_1) = \left\{ ([v]_\pi, \#\{e: [v]_\pi \overset{e}{\sim} w_1\}) \right\} \subset V^\pi \times \mathbb{N}.$$

In the special case of $\pi = \{\{v\} : v \in V\}$ and $G^\pi = G$, we simply write $\mathcal{B}(w_1)$. By a slight abuse of notation, we sometimes refer to a vertex $[v]_\pi \in \mathcal{B}_\pi(w_1)$ or a vertex $v \in \mathcal{B}(w_1)$ instead of the tuple. For a map $\phi: V^\pi \hookrightarrow [N]$, the tensor entries $\mathbf{T}_{d,N}(w_1|\phi)$ and $\mathbf{T}_{d,N}(w'_1|\phi)$ are dependent (in fact, identical) iff $\mathcal{B}_\pi(w_1) = \mathcal{B}_\pi(w'_1)$, in which case we say that the diamonds w_1 and w'_1 are *overlaid* in G^π .

Definition 1.3 (i) of our Wigner tensor ensemble implies that (3.3) holds unless

$$\#\mathcal{B}_\pi(w_1) = 1 \text{ or } \exists w'_1 \in W_1 \setminus \{w_1\} : \mathcal{B}_\pi(w_1) = \mathcal{B}_\pi(w'_1), \quad \forall w_1 \in W_1. \tag{3.4}$$

So, we can restrict ourselves to partitions $\pi \in \mathcal{P}(V)$ that satisfy this property. Assuming the suboptimality of $\#\mathcal{B}_\pi(w_1) = 1$ for now, this says that diamond overlays are common. To keep track of these, we define the function

$$\sigma_{(\cdot)} : \mathcal{P}(V) \rightarrow \mathcal{P}(W_1), \quad \pi \mapsto \sigma_\pi,$$

where the partition σ_π satisfies

$$w_1 \sim_{\sigma_\pi} w'_1 \iff \mathcal{B}_\pi(w_1) = \mathcal{B}_\pi(w'_1). \tag{3.5}$$

It will be necessary to classify the index vertices $V = V_{\text{in}} \sqcup V_{\text{out}}$ in G according to their positions in the graph: the *inner* vertices V_{in} are those inside of the cycle (adjacent only to diamonds); the outer vertices V_{out} are those outside of the cycle (adjacent to boxes). We also use an auxiliary graph H^π to compress the information of G^π .




Figure 4: An example of a partition π that overlays every diamond with at least one other diamond. For convenience, we have omitted the boxes corresponding to the contracting vectors. On the left, the blocks of π are indicated by the coloring of the vertices. On the right, the dotted lines indicate the blocks of σ_π in H^π . Note that the quotient graph H^π is the same as the original graph H in this case and that σ_π is a *crossing* pairing.




Figure 5: Another example of a partition π that overlays every diamond with at least one other diamond. Here, the yellow, pink, and purple vertices are identified in a way that resembles the optimal situation from the usual Wigner case. We will see that this is still not enough to overcome the defects in the other parts of the graph.

Definition 3.1 (H and H^π). Let $H = (\mathcal{V}(H), \mathcal{E}(H))$ be the graph with vertices $\mathcal{V}(H) = V_{\text{in}}$ and edges $\mathcal{E}(H) = W_1$. Two vertices $v, v' \in \mathcal{V}(H)$ are connected by an edge $e \in \mathcal{E}(H)$ iff the diamond e is adjacent to v and v' in G . For a partition $\pi \in \mathcal{P}(V)$, we use the shorthand notation H^π for the quotient graph $H^{\pi|_{\mathcal{V}(H)}}$.

As before, H is a simple m -cycle whose vertices are the inner vertices of G and whose edges correspond to the diamonds of G . The construction of H^π does not materially differ from the matrix case of $d = 2$; however, one no longer has a simple characterization of the quotients that produce nonzero terms $\tau^0[G^\pi] \neq 0$. In particular, if $d \geq 3$, then outer vertices can be identified with inner vertices to create a diamond overlay while keeping the inner vertices separate in H^π . Consequently, edges do not even have to be incident in H^π to be in the same block of σ_π . See Figure 4 and Figure 5 for examples.

This already rules out applying our crude asymptotic (3.2), which is sufficient for the analysis in the $d = 2$ case. Roughly speaking, the diamond overlay condition allows for too many vertices in G^π : one can find partitions $\pi \in \mathcal{P}(V)$ such that $\tau^0[G^\pi] \neq 0$ with

$$\#\{\phi : V^\pi \hookrightarrow [N]\} = \Theta(N^{(d-1)\frac{m}{2}+1}).$$

Since $d \geq 3$, this is much larger than the normalization term $N^{\frac{m}{2}+1}$.

The issue comes from the overly generous bound of the summands in (3.1), which can be of much smaller order depending on the vertex labels ϕ . For example, if the contracting vectors are standard basis vectors, then the product $\prod_{w_2 \in W_2} u_N(w_2|\phi)$ will be 0 for most ϕ . The opposite extreme would be if the contracting vectors were completely flat with $N^{-\frac{1}{2}}$ for every entry, in which case the product is $N^{-(d-2)\frac{m}{2}}$ for any ϕ . Treating the general case without knowing the exact shape of the contracting vectors greatly complicates the analysis. Recognizing the importance of the inner-outer interactions, we split our partition $\pi = \pi_{\text{in}} \sqcup \pi_{\text{out}} \sqcup \pi_{\text{mix}}$ into

- π_{in} , the set of blocks consisting only of inner vertices;
- π_{out} , the set of blocks consisting only of outer vertices;
- π_{mix} , the set of remaining blocks, each of which has at least one inner vertex and one outer vertex.

Our goal is then to understand how the number of outer vertices in each block of π relates to the structure of σ_π .

3.2 Convergence in expectation

In the previous section, we saw that a major complication comes from the existence of partitions $\pi \in \mathcal{P}(V)$ such that $\tau^0[G^\pi] \neq 0$ and $\pi_{\text{mix}} \neq \emptyset$. If we assume for the time being that such mixed partitions do not contribute in the limit, then we can prove convergence to a semicircular family in a manner similar to Proposition 2.6. Actually proving that mixed partitions do not contribute in the limit involves a delicate technical argument that we defer to Section 4.1.

Properties (P1) and (P2) in Proposition 2.6 have natural extensions to the general setting of $d \geq 2$. Since we will need to show that the contracted ensemble satisfies these properties, we state them as lemmas.

Lemma 3.2 (Supported on double trees).

$$\tau^0[G^\pi] = O_m(N^{-1})$$

unless H^π is a double tree and $\pi_{\text{mix}} = \emptyset$, where the asymptotic is uniform over all possible contracting vectors $(u_N^{(i,j)})_{i \in I, j \in [d-2]} \subset \mathbb{S}^{N-1}$.

Lemma 3.3 (Multiplicativity). For any partition $\rho \in \mathcal{P}(V_{\text{in}})$ such that H^ρ is a double tree,

$$\sum_{\substack{\pi \in \mathcal{P}(V): \\ \pi \upharpoonright_{V_{\text{in}}} = \rho, \pi_{\text{mix}} = \emptyset}} \tau^0[G^\pi] = \prod_{\{i, i'\}} \mathbf{K}_{i, i'}^{(N)} + O_{m,d}(N^{-1}), \tag{3.6}$$

where the product runs over the double edges $\{i, i'\}$ in the double tree H^ρ and $(\mathbf{K}_{i, i'}^{(N)})_{i, i' \in I}$ is the Gram matrix defined in (1.2). As before, the asymptotic is uniform over all possible contracting vectors $(u_N^{(i,j)})_{i \in I, j \in [d-2]} \subset \mathbb{S}^{N-1}$.

The proof of Lemma 3.2 (resp., Lemma 3.3) is deferred to Section 4.1 (resp., the end of this section). Drawing inspiration from Proposition 2.6, we will now show that Lemmas 3.2 and 3.3 imply that the mixed moments of the contracted tensor ensemble are close to the mixed moments of a semicircular family. We recall the notation

$$P_N(i_1, \dots, i_m) = \left(\frac{1}{\sqrt{N}} \mathbf{T}_{d,N} \left[u_N^{(i_1,1)} \otimes \dots \otimes u_N^{(i_1,d-2)} \right] \right) \dots \left(\frac{1}{\sqrt{N}} \mathbf{T}_{d,N} \left[u_N^{(i_m,1)} \otimes \dots \otimes u_N^{(i_m,d-2)} \right] \right),$$

which allows us to state

Proposition 3.4. For any $i_1, \dots, i_m \in I$,

$$\left| \frac{1}{N} \mathbb{E} \left[\text{Tr} \left(P_N(i_1, \dots, i_m) \right) \right] - \varphi \left(s_{i_1}^{(N)} \cdots s_{i_m}^{(N)} \right) \right| = O_{m,d}(N^{-1}),$$

where $(s_i^{(N)})_{i \in I}$ is a semicircular family of covariance $(\mathbf{K}_{i,i'})_{i,i' \in I}$ and the asymptotic is uniform over all possible contracting vectors $(u_N^{(i,j)})_{i \in I, j \in [d-2]} \subset \mathbb{S}^{N-1}$.

Proof assuming Lemmas 3.2 and 3.3. Recall that the graph G was defined so that

$$\tau[G] = \frac{1}{N} \mathbb{E} \left[\text{Tr} \left(P_N(i_1, \dots, i_m) \right) \right].$$

By Lemma 2.7, if $\rho \in \mathcal{P}(V_{\text{in}}) = \mathcal{P}(\mathcal{V}(H))$, then H^ρ is a double tree iff $\rho \in \mathcal{NC}(V_{\text{in}})$ and $\mathcal{K}(\rho) \in \mathcal{NC}_2(\mathcal{E}(H))$. Combining this with Lemma 3.2, we obtain

$$\tau[G] = \sum_{\pi \in \mathcal{P}(V)} \tau^0[G^\pi] = \sum_{\substack{\rho \in \mathcal{NC}(V_{\text{in}}): \\ \mathcal{K}(\rho) \in \mathcal{NC}_2(\mathcal{E}(H))}} \sum_{\substack{\pi \in \mathcal{P}(V): \\ \pi|_{V_{\text{in}}} = \rho, \pi_{\text{mix}} = \emptyset}} \tau^0[G^\pi] + O_{m,d}(N^{-1}).$$

Applying Lemma 3.3 to the inner sum yields

$$\begin{aligned} \tau[G] &= \sum_{\substack{\rho \in \mathcal{NC}(V_{\text{in}}): \\ \mathcal{K}(\rho) \in \mathcal{NC}_2(\mathcal{E}(H))}} \prod_{\{i,i'\} \in \mathcal{K}(\rho)} \mathbf{K}_{i,i'}^{(N)} + O_{m,d}(N^{-1}) \\ &= \sum_{\sigma \in \mathcal{NC}_2(\mathcal{E}(H))} \prod_{\{i,i'\} \in \sigma} \mathbf{K}_{i,i'}^{(N)} + O_{m,d}(N^{-1}). \end{aligned}$$

Finally, since $\mathbf{K}^{(N)}$ is positive semidefinite, the semicircular family $(s_i^{(N)})_{i \in I}$ is well-defined with

$$\varphi \left(s_{i_1}^{(N)} \cdots s_{i_m}^{(N)} \right) = \sum_{\sigma \in \mathcal{NC}_2(\mathcal{E}(H))} \prod_{\{i,i'\} \in \sigma} \mathbf{K}_{i,i'}^{(N)}$$

by the characterization in (2.5). We conclude that $|\tau[G] - \varphi(s_{i_1}^{(N)} \cdots s_{i_m}^{(N)})| = O_{m,d}(N^{-1})$. \square

Naturally, if the limits $\mathbf{K}_{i,i'} = \lim_{N \rightarrow \infty} \mathbf{K}_{i,i'}^{(N)}$ exist for every $i, i' \in I$, then Proposition 3.4 implies that the contracted tensor ensemble converges in distribution to a semicircular family of covariance $(\mathbf{K}_{i,i'})_{i,i' \in I}$. We conclude this section by proving Lemma 3.3.

Proof of Lemma 3.3. Let $\rho \in \mathcal{P}(V_{\text{in}})$ be such that H^ρ is a double tree. We divide the proof into three steps. First, we identify the partitions $\pi \in \mathcal{P}(V)$ that contribute to the sum in (3.6) and use this to simplify the formula for each of the corresponding $\tau^0[G^\pi]$. Secondly, we show that the values of $\phi|_{V_{\text{in}}}$ can be ignored in these formulas. Finally, we prove that our simplified formulation of the sum in (3.6) can be factored, up to a $O_{m,d}(N^{-1})$ correction, as a product of sums over uniform block permutations, which in turn can be rewritten using the advertised inner product in (1.2).

Step 1: Simplifying the formula for $\tau^0[G^\pi]$

By Lemma 2.7, we know that H^ρ is a double tree iff $\rho \in \mathcal{NC}(V_{\text{in}})$ and $\mathcal{K}(\rho) \in \mathcal{NC}_2(\mathcal{E}(H)) = \mathcal{NC}_2(W_1)$, where $\mathcal{K}(\rho)$ is the pair partition whose pairs are precisely the double edges in H^ρ . The appearance of uniform block permutations (Definition 1.10) in our formulas is a result of the following relabeling of the vertices in G according to $\mathcal{K}(\rho)$ as suggested by Figure 6.




Figure 6: An example of the graph G in (2.1) for $m = 2$. The dotted line down the middle is not part of the graph. We add it to motivate the appearance of uniform block permutations in our formula for the covariance. See the proof of Lemma 3.3 for the precise details.

For each pair of diamonds $D \in \mathcal{K}(\rho)$, we designate one as the *left* diamond $w_1^{D_L}$ and the other as the *right* diamond $w_1^{D_R}$ so that $D = \{w_1^{D_L}, w_1^{D_R}\}$. The outer vertices V_{out} in G can then be partitioned according to their neighboring diamond:

$$V_{\text{out}} = \bigsqcup_{D \in \mathcal{K}(\rho)} [d - 2]_D = \bigsqcup_{D \in \mathcal{K}(\rho)} \left([d - 2]_{D_L} \sqcup [d - 2]_{D_R} \right),$$

where

$$\begin{aligned} [d - 2]_{D_L} &= \{v \in V_{\text{out}} : v \sim w_1^{D_L}\}; \\ [d - 2]_{D_R} &= \{v \in V_{\text{out}} : v \sim w_1^{D_R}\}. \end{aligned}$$

Now, suppose that $\pi \in \mathcal{P}(V)$ satisfies $\pi \upharpoonright_{V_{\text{in}}} = \rho$ and $\pi_{\text{mix}} = \emptyset$. We can then factor $\pi = \pi_{\text{in}} \sqcup \pi_{\text{out}}$ into $\pi_{\text{in}} = \rho$ and $\pi_{\text{out}} \in \mathcal{P}(V_{\text{out}})$. Since H^ρ is a double tree, every diamond in G^π is adjacent to at least two index vertices $[v]_\pi \neq [v']_\pi$. In that case, the centeredness of the off-diagonal tensor entries in Definition 1.3 (i) implies that $\tau^0[G^\pi] = 0$ unless every diamond is overlaid with at least one other diamond in G^π . Since $\pi_{\text{mix}} = \emptyset$ and H^ρ is a double tree, diamonds $w_1, w'_1 \in W_1$ are overlaid in G^π only if $\{w_1, w'_1\} \in \mathcal{K}(\rho)$. It follows that $\sigma_\pi = \mathcal{K}(\rho)$, where we recall that $\sigma_{(\cdot)}$ is the function that takes a partition $\pi \in \mathcal{P}(V)$ to the corresponding partition of the diamonds $\sigma_\pi \in \mathcal{P}(W_1)$ satisfying (3.5). We can then restrict our sum

$$\sum_{\substack{\pi \in \mathcal{P}(V): \\ \pi \upharpoonright_{V_{\text{in}}} = \rho, \pi_{\text{mix}} = \emptyset}} \tau^0[G^\pi] = \sum_{\pi \in \mathcal{P}_\rho} \tau^0[G^\pi]$$

to the class of partitions

$$\mathcal{P}_\rho = \{\pi \in \mathcal{P}(V) : \pi_{\text{in}} = \rho, \pi_{\text{mix}} = \emptyset, \sigma_\pi = \mathcal{K}(\rho)\}.$$

If $\pi \in \mathcal{P}_\rho$, then any diamond in G^π will have at least three neighbors: two neighbors come from the fact that H^ρ is a double tree, and at least one outer vertex neighbor survives from G as a result of the non-mixing $\pi_{\text{mix}} = \emptyset$. Since $\sigma_\pi = \mathcal{K}(\rho)$, the independence assumption together with the variance profile in Definition 1.3 (ii) of our tensor

ensemble imply that for any $\phi : V^\pi \hookrightarrow [N]$,

$$\begin{aligned} \mathbb{E} \left[\prod_{w_1 \in W_1} \mathbf{T}_{d,N}(w_1|\phi) \right] &= \prod_{D \in \mathcal{K}(\rho)} \mathbb{E} \left[\mathbf{T}_{d,N}(w_1^{D_L}|\phi) \mathbf{T}_{d,N}(w_1^{D_R}|\phi) \right] \\ &= \prod_{D \in \mathcal{K}(\rho)} \left(b_1^{D,\phi}, \dots, b_N^{D,\phi} \right)^{-1}, \end{aligned} \tag{3.7}$$

where $b_l^{D,\phi}$ is the multiplicity of $l \in [N]$ as an index in $\mathbf{T}_{d,N}(w_1^{D_L}|\phi)$ (equivalently, in $\mathbf{T}_{d,N}(w_1^{D_R}|\phi)$). Note that while $(b_l^{D,\phi})_{l \in [N]}$ depends on ϕ , the multiset of multiplicities $\{b_1^{D,\phi}, \dots, b_N^{D,\phi}\}$ only depends on the block structure of π restricted to $[d-2]_D$. To keep track of this, we define the partition $\pi_{\text{out}}^D := \pi_{\text{out}} \upharpoonright_{[d-2]_D}$. Since $w_1^{D_L}$ and $w_1^{D_R}$ are overlaid, we know that $\#(B \cap [d-2]_{D_L}) = \#(B \cap [d-2]_{D_R})$ for every $B \in \pi_{\text{out}}$. In particular, this implies that π_{out}^D is a uniform block permutation of $[d-2]_D$. This will prove useful later, but for now we simply note that

$$\prod_{D \in \mathcal{K}(\rho)} \left(b_1^{D,\phi}, \dots, b_N^{D,\phi} \right)^{-1} = \prod_{D \in \mathcal{K}(\rho)} \frac{1}{d!} \prod_{B \in \pi_{\text{out}}^D} \left(\frac{\#(B)}{2} \right)! =: \nu_\pi$$

only depends the choice of $\pi \in \mathcal{P}_\rho$ and not the actual vertex labels $\phi : V^\pi \hookrightarrow [N]$. Combining this with (3.7), we can simplify the formula for $\tau^0[G^\pi]$ to

$$\begin{aligned} \tau^0[G^\pi] &= \frac{1}{N^{\frac{m}{2}+1}} \sum_{\phi: \pi \hookrightarrow [N]} \mathbb{E} \left[\prod_{w_1 \in W_1} \mathbf{T}_{d,N}(w_1|\phi) \right] \prod_{w_2 \in W_2} u_N(w_2|\phi) \\ &= \frac{\nu_\pi}{N^{\frac{m}{2}+1}} \sum_{\phi: \pi \hookrightarrow [N]} \prod_{w_2 \in W_2} u_N(w_2|\phi). \end{aligned} \tag{3.8}$$

Step 2: Discarding the inner labels $\phi|_{V_{\text{in}}}$

Note that the terms $\prod_{w_2 \in W_2} u_N(w_2|\phi)$ appearing in the sum above do not depend on the values of $\phi|_{V_{\text{in}}}$. Thus, for any $\pi \in \mathcal{P}_\rho$,

$$\begin{aligned} &\frac{1}{N^{\frac{m}{2}+1}} \sum_{\phi: \pi \hookrightarrow [N]} \prod_{w_2 \in W_2} u_N(w_2|\phi) \\ &= \frac{1}{N^{\frac{m}{2}+1}} \sum_{\phi: \pi_{\text{out}} \hookrightarrow [N]} \sum_{\substack{\phi': \pi_{\text{in}} \hookrightarrow [N] \text{ s.t.} \\ \phi([N]) \cap \phi'([N]) = \emptyset}} \prod_{w_2 \in W_2} u_N(w_2|\phi) \\ &= \frac{(N - \#(\pi_{\text{out}}))^{\frac{m}{2}+1}}{N^{\frac{m}{2}+1}} \sum_{\phi: \pi_{\text{out}} \hookrightarrow [N]} \prod_{w_2 \in W_2} u_N(w_2|\phi), \end{aligned}$$

where

$$(N - \#(\pi_{\text{out}}))^{\frac{m}{2}+1} = (N - \#(\pi_{\text{out}}))(N - \#(\pi_{\text{out}}) - 1) \cdots (N - \#(\pi_{\text{out}}) - \frac{m}{2})$$

is the falling factorial and we have used the fact that $\#(\pi_{\text{in}}) = \frac{m}{2} + 1$ since H^ρ is a double tree.

At the same time, since $\sigma_\pi = \mathcal{K}(\rho)$ and $\pi_{\text{mix}} = \emptyset$, we know that every block $B \in \pi_{\text{out}}$ is adjacent to at least two boxes in G^π . Combining the elementary bound $|u_N^{(i,j)}(k)| \leq$

$\|u_N^{(i,j)}\|_2 = 1$ with the Cauchy-Schwarz inequality, we see that

$$\begin{aligned} \left| \sum_{\phi: \pi_{\text{out}} \hookrightarrow [N]} \prod_{w_2 \in W_2} u_N(w_2|\phi) \right| &\leq \sum_{\phi: \pi_{\text{out}} \rightarrow [N]} \prod_{w_2 \in W_2} |u_N(w_2|\phi)| \\ &= \prod_{B \in \pi_{\text{out}}} \left(\sum_{k \in [N]} \prod_{\substack{w_2 \in W_2: \\ w_2 \sim_{G^\pi} B}} |u_N^{(w_2)}(k)| \right) \leq 1. \end{aligned}$$

Since $(N - \#\pi_{\text{out}})^{\frac{m}{2}+1} = N^{\frac{m}{2}+1} + O_{m,d}(N^{\frac{m}{2}})$, this allows us to further simplify (3.8) to

$$\tau^0[G^\pi] = \nu_\pi \sum_{\phi: \pi_{\text{out}} \hookrightarrow [N]} \prod_{w_2 \in W_2} u_N(w_2|\phi) + O_{m,d}(N^{-1}).$$

Step 3: Factoring and reassembling

Putting everything together, we obtain

$$\begin{aligned} \sum_{\substack{\pi \in \mathcal{P}(V): \\ \pi|_{V_{\text{in}}} = \rho, \pi_{\text{mix}} = \emptyset}} \tau^0[G^\pi] &= \sum_{\pi \in \mathcal{P}_\rho} \tau^0[G^\pi] \\ &= \sum_{\pi \in \mathcal{P}_\rho} \nu_\pi \sum_{\phi: \pi_{\text{out}} \hookrightarrow [N]} \prod_{w_2 \in W_2} u_N(w_2|\phi) + O_{m,d}(N^{-1}) \\ &= \sum_{\pi \in \mathcal{P}_\rho} \sum_{\phi: \pi_{\text{out}} \hookrightarrow [N]} \prod_{D \in \mathcal{K}(\rho)} \frac{1}{d!} \prod_{B \in \pi_{\text{out}}^D} \left(\frac{\#(B)}{2} \right)! \prod_{v \in B} u_N^{(w_2^v)}(\phi(B)) + O_{m,d}(N^{-1}), \end{aligned} \tag{3.9}$$

where w_2^v denotes the unique box in W_2 that is adjacent to the outer vertex $v \in V_{\text{out}}$. Using the fact that $\pi_{\text{out}}^D \in \mathcal{UBP}([d-2]_D)$, we can factor the non-asymptotic term in (3.9) into

$$\prod_{D \in \mathcal{K}(\rho)} \left(\frac{1}{d!} \sum_{\pi \in \mathcal{UBP}([d-2]_D)} \sum_{\phi: \pi \hookrightarrow [N]} \prod_{B \in \pi} \left(\frac{\#(B)}{2} \right)! \prod_{v \in B \cap [d-2]_{D_L}} u_N^{(w_2^v)}(\phi(B)) \prod_{v' \in B \cap [d-2]_{D_R}} u_N^{(w_2^{v'})}(\phi(B)) \right).$$

It then suffices to show that

$$\begin{aligned} &\frac{1}{d!} \sum_{\pi \in \mathcal{UBP}([d-2]_D)} \sum_{\phi: \pi \hookrightarrow [N]} \prod_{B \in \pi} \left(\frac{\#(B)}{2} \right)! \prod_{v \in B \cap [d-2]_{D_L}} u_N^{(w_2^v)}(\phi(B)) \prod_{v' \in B \cap [d-2]_{D_R}} u_N^{(w_2^{v'})}(\phi(B)) \\ &= \frac{1}{d(d-1)} \left\langle u_N^{(i,1)} \odot \dots \odot u_N^{(i,d-2)}, u_N^{(i',1)} \odot \dots \odot u_N^{(i',d-2)} \right\rangle, \end{aligned}$$

where the vectors $(u_N^{(i,j)})_{j \in [d-2]}$ (resp., $(u_N^{(i',j)})_{j \in [d-2]}$) correspond to the outer vertices $[d-2]_{D_L} = \{v_j\}_{j \in [d-2]}$ (resp., $[d-2]_{D_R} = \{v'_j\}_{j \in [d-2]}$).

To see this, we introduce the set $\mathcal{UBP}_2([d-2]_D)$ of uniform block permutations with blocks exclusively of size two, which can be thought of as the set of permutations $\mathcal{UBP}_2([d-2]_D) \cong \mathfrak{S}_{d-2}$. Note that for any $\pi \in \mathcal{UBP}([d-2]_D)$,

$$\#\{\sigma \in \mathcal{UBP}_2([d-2]_D) : \sigma \leq \pi\} = \prod_{B \in \pi} \left(\frac{\#(B)}{2} \right)!,$$

where \leq denotes the reversed refinement order. This allows us to reassemble the injective labels $\phi : \pi \hookrightarrow [N]$ for $\pi \in \mathcal{UBP}([d-2]_D)$ into general labels $\phi : \sigma \rightarrow [N]$ for

$\sigma \in \mathcal{UBP}_2([d-2]_D)$. Indeed,

$$\begin{aligned} & \frac{1}{d!} \sum_{\pi \in \mathcal{UBP}([d-2]_D)} \sum_{\phi: \pi \leftrightarrow [N]} \prod_{B \in \pi} \left(\frac{\#(B)}{2}\right)! \prod_{v \in B \cap [d-2]_{D_L}} u_N^{(w_2^v)}(\phi(B)) \prod_{v' \in B \cap [d-2]_{D_R}} u_N^{(w_2^{v'})}(\phi(B)) \\ &= \frac{1}{d!} \sum_{\pi \in \mathcal{UBP}([d-2]_D)} \sum_{\substack{\sigma \in \mathcal{UBP}_2([d-2]_D) \\ \text{s.t. } \sigma \leq \pi}} \sum_{\phi: \pi \leftrightarrow [N]} \prod_{B \in \pi} \prod_{v \in B \cap [d-2]_{D_L}} u_N^{(w_2^v)}(\phi(B)) \prod_{v' \in B \cap [d-2]_{D_R}} u_N^{(w_2^{v'})}(\phi(B)) \\ &= \frac{1}{d!} \sum_{\sigma \in \mathcal{UBP}_2([d-2]_D)} \sum_{\phi: \sigma \rightarrow [N]} \prod_{B \in \sigma} \prod_{v \in B \cap [d-2]_{D_L}} u_N^{(w_2^v)}(\phi(B)) \prod_{v' \in B \cap [d-2]_{D_R}} u_N^{(w_2^{v'})}(\phi(B)) \\ &= \frac{1}{d!} \sum_{\sigma \in \mathfrak{S}_{d-2}} \prod_{j=1}^{d-2} \langle u_N^{(i,j)}, u_N^{(i',\sigma(j))} \rangle \\ &= \frac{1}{d!} \text{per} \left[\left(\langle u_N^{(i,j)}, u_N^{(i',k)} \rangle \right)_{j,k=1}^{d-2} \right] \\ &= \frac{1}{d(d-1)} \left\langle u_N^{(i,1)} \odot \dots \odot u_N^{(i,d-2)}, u_N^{(i',1)} \odot \dots \odot u_N^{(i',d-2)} \right\rangle, \end{aligned}$$

where we have used the identification $\mathcal{UBP}_2([d-2]_D) \cong \mathfrak{S}_{d-2}$ and the permanental representation of the symmetrized inner product [22, Theorem 2.2]. \square

3.3 Concentration and almost sure convergence

To prove concentration, we first need to define the random version of $\tau[G]$. For the graph G in (2.1), let

$$\text{Tr}(G) = \frac{1}{N^{\frac{m}{2}}} \sum_{\phi: V \rightarrow [N]} \prod_{w_1 \in W_1} \mathbf{T}_{d,N}(w_1 | \phi) \prod_{w_2 \in W_2} u_N(w_2 | \phi)$$

so that $\tau[G] = \frac{1}{N} \mathbb{E}[\text{Tr}(G)]$. Proposition 3.4 already provides the asymptotics of $\frac{1}{N} \mathbb{E}[\text{Tr}(G)]$. We proceed to an analysis of the higher moments of the centered random variable $\text{Tr}(G) - \mathbb{E}[\text{Tr}(G)]$. We prove the following result in Section 4.2 using an extension of Lemma 3.2 to the disjoint union of multiple copies of G .

Proposition 3.5 (Concentration). *If G is the graph in (2.1), then*

$$\mathbb{E} \left[\left(\frac{1}{N} \text{Tr}(G) - \mathbb{E} \left[\frac{1}{N} \text{Tr}(G) \right] \right)^{2M} \right] = O_{m,d,M}(N^{-2M}). \tag{3.10}$$

Propositions 3.4 and 3.5 allow us to prove our main results: Theorems 1.5 and 1.6.

Proof of Theorems 1.5 and 1.6. Theorem 1.5 follows from Theorem 1.6 via a standard application of the Borel-Cantelli lemma. To prove Theorem 1.6, let $i_1, \dots, i_m \in I$ and recall that $\text{Tr}(G) = \text{Tr}(P_N(i_1, \dots, i_m))$ by construction. Proposition 3.5 implies

$$\varepsilon^{2M} \mathbb{P} \left[\left| \frac{1}{N} \text{Tr}(G) - \mathbb{E} \left[\frac{1}{N} \text{Tr}(G) \right] \right| > \varepsilon \right] = O_{m,d,M}(N^{-2M}).$$

At the same time, we know that $\mathbb{E} \left[\frac{1}{N} \text{Tr}(G) \right] = \phi(s_{i_1}^{(N)} \dots s_{i_m}^{(N)}) + O_{m,d}(N^{-1})$ from Proposition 3.4. Using the triangle inequality, we conclude that

$$\varepsilon^{2M} \mathbb{P} \left[\left| \frac{1}{N} \text{Tr}(G) - \phi(s_{i_1}^{(N)} \dots s_{i_m}^{(N)}) \right| > \varepsilon' \right] = O_{m,d,M}(N^{-2M})$$

for any $\varepsilon' > \varepsilon$, which proves (1.3). \square

4 Proofs of the technical results

4.1 Proof of Lemma 3.2

Recall that we must prove

$$\tau^0[G^\pi] = O_m(N^{-1})$$

unless H^π is a double tree and $\pi_{\text{mix}} = \emptyset$. By our earlier discussion, we can restrict to partitions $\pi \in \mathcal{P}(V)$ that satisfy (3.4).

Note that a diamond overlay $\mathcal{B}_\pi(w_1) = \mathcal{B}_\pi(w'_1)$ does not imply that an inner vertex neighbor $v \in \mathcal{B}(w_1) \cap V_{\text{in}}$ of w_1 gets identified with at least one other inner vertex neighbor $v' \in \mathcal{B}(w'_1) \cap V_{\text{in}}$ of w'_1 ; however, without such an identification, the equality of the neighborhoods can only hold if v is identified with an outer vertex neighbor $v'' \in \mathcal{B}(w'_1) \cap V_{\text{out}}$. See Figure 4 and Figure 5 for examples. This simple observation will be used frequently in the sequel, so we record it below.

Lemma 4.1. *If $\mathcal{B}_\pi(w_1) = \mathcal{B}_\pi(w'_1)$ and $v \in \mathcal{B}(w_1) \cap V_{\text{in}}$ satisfies*

$$[v]_\pi \neq [v']_\pi, \quad \forall v' \in \mathcal{B}(w'_1) \cap V_{\text{in}},$$

then

$$\exists v'' \in \mathcal{B}(w'_1) \cap V_{\text{out}} : [v]_\pi = [v'']_\pi.$$

In particular, $\pi_{\text{mix}} \neq \emptyset$.

Condition (3.4) forces every block $B = [v]_\pi \in \pi_{\text{out}}$ to have size $\#([v]_\pi) \geq 2$. The degrees of freedom in choosing $\phi|_{\pi_{\text{out}}}$ are then negated by the terms coming from the neighboring boxes. To see this, note that

$$\begin{aligned} & \sum_{\phi: V^\pi \rightarrow [N]} \mathbb{E} \left[\prod_{w_1 \in W_1} \mathbf{T}_{d,N}(w_1 | \phi) \right] \prod_{w_2 \in W_2} u_N(w_2 | \phi) \\ &= O_m \left(N^{\#(\pi_{\text{in}})} \sum_{\phi: \pi_{\text{mix}} \rightarrow [N]} \prod_{[v]_\pi \in \pi_{\text{mix}}} \prod_{\substack{w_2 \in W_2: \\ w_2 \sim [v]_\pi}} \left| u_N^{(w_2)}(\phi([v]_\pi)) \right| \right) \\ & \times O_m \left(\sum_{\phi': \pi_{\text{out}} \rightarrow [N]} \prod_{[v]_\pi \in \pi_{\text{out}}} \prod_{\substack{w_2 \in W_2: \\ w_2 \sim [v]_\pi}} \left| u_N^{(w_2)}(\phi'([v]_\pi)) \right| \right), \end{aligned}$$

where we can control the last term

$$\sum_{\phi': \pi_{\text{out}} \rightarrow [N]} \prod_{[v]_\pi \in \pi_{\text{out}}} \prod_{\substack{w_2 \in W_2: \\ w_2 \sim [v]_\pi}} \left| u_N^{(w_2)}(\phi'([v]_\pi)) \right| = \prod_{[v]_\pi \in \pi_{\text{out}}} \left(\sum_{k \in [N]} \prod_{\substack{w_2 \in W_2: \\ w_2 \sim [v]_\pi}} |u_N^{(w_2)}(k)| \right) \leq 1$$

by the Cauchy-Schwarz inequality and the fact that $|u_N^{(i,j)}(k)| \leq \|u_N^{(i,j)}\|_2 = 1$. So, the π_{out} vertices do not actually contribute to the exponent in the crude asymptotic (3.2) and we have the refinement

$$\tau^0[G^\pi] = O_m(N^{\#(\pi_{\text{in}}) + \#(\pi_{\text{mix}}) - \frac{m}{2} - 1}).$$

In fact, we can apply the same argument to $\pi_{\text{mix}} = \pi_{\text{mix}}^{(1)} \sqcup \pi_{\text{mix}}^{(\geq 2)}$, where $\pi_{\text{mix}}^{(1)}$ is the set of mixed blocks such that each block has exactly one outer vertex and $\pi_{\text{mix}}^{(\geq 2)}$ is the set of mixed blocks such that each block has at least two outer vertices. This allows us to

further reduce

$$\begin{aligned}
 & O_m \left(N^{\#(\pi_{\text{in}})} \sum_{\phi: \pi_{\text{mix}} \rightarrow [N]} \prod_{[v]_{\pi} \in \pi_{\text{mix}}} \prod_{\substack{w_2 \in W_2: \\ w_2 \sim [v]_{\pi}}} \left| u_N^{(w_2)}(\phi([v]_{\pi})) \right| \right) \\
 &= O_m \left(N^{\#(\pi_{\text{in}})} \sum_{\phi: \pi_{\text{mix}}^{(1)} \rightarrow [N]} \prod_{[v]_{\pi} \in \pi_{\text{mix}}^{(1)}} \prod_{\substack{w_2 \in W_2: \\ w_2 \sim [v]_{\pi}}} \left| u_N^{(w_2)}(\phi([v]_{\pi})) \right| \right) \\
 &\quad \times O_m \left(\sum_{\phi': \pi_{\text{mix}}^{(\geq 2)} \rightarrow [N]} \prod_{[v]_{\pi} \in \pi_{\text{mix}}^{(\geq 2)}} \prod_{\substack{w_2 \in W_2: \\ w_2 \sim [v]_{\pi}}} \left| u_N^{(w_2)}(\phi'([v]_{\pi})) \right| \right) \\
 &= O_m \left(N^{\#(\pi_{\text{in}})} \prod_{[v]_{\pi} \in \pi_{\text{mix}}^{(1)}} \prod_{\substack{w_2 \in W_2: \\ w_2 \sim [v]_{\pi}}} \langle |u_N^{(w_2)}|, \hat{1}_N \rangle \right) = O_m \left(N^{\#(\pi_{\text{in}}) + \frac{\#(\pi_{\text{mix}}^{(1)})}{2}} \right),
 \end{aligned}$$

where $|u_N^{(w_2)}|$ is the entrywise absolute value of the contracting vector $u_N^{(w_2)}$ and $\hat{1}_N \in \mathbb{R}^N$ is the all-ones vector. The final form of our asymptotic is then

$$\tau^0[G^{\pi}] = O_m \left(N^{\#(\pi_{\text{in}}) + \frac{\#(\pi_{\text{mix}}^{(1)})}{2} - \frac{m}{2} - 1} \right),$$

which is controlled by

Proposition 4.2. *For any partition π satisfying (3.4),*

$$\#(\pi_{\text{in}}) + \frac{\#(\pi_{\text{mix}}^{(1)})}{2} \leq \frac{m}{2} + 1 \tag{4.1}$$

with equality iff m is even, $\pi_{\text{mix}} = \emptyset$, and H^{π} is a double tree. In all other cases,

$$\#(\pi_{\text{in}}) + \frac{\#(\pi_{\text{mix}}^{(1)})}{2} \leq \frac{m}{2}.$$

Proof. First, we introduce some auxiliary sets:

$$\begin{aligned}
 \pi_{\text{mix}}^{(1,1)} &= \{B \in \pi : \#(B \cap V_{\text{out}}) = 1, \#(B \cap V_{\text{in}}) = 1\}; \\
 \pi_{\text{mix}}^{(1,\geq 2)} &= \{B \in \pi : \#(B \cap V_{\text{out}}) = 1, \#(B \cap V_{\text{in}}) \geq 2\}; \\
 \pi_{\text{in}}^{(1)} &= \{B \in \pi : \#(B \cap V_{\text{out}}) = 0, \#(B \cap V_{\text{in}}) = 1\}; \\
 \pi_{\text{in}}^{(\geq 2)} &= \{B \in \pi : \#(B \cap V_{\text{out}}) = 0, \#(B \cap V_{\text{in}}) \geq 2\};
 \end{aligned}$$

and

$$\begin{aligned}
 V_{\pi}^{(1,1)} &= \cup_{B \in \pi_{\text{mix}}^{(1,1)}} B \cap V_{\text{in}}; \\
 V_{\pi}^{(1,\geq 2)} &= \cup_{B \in \pi_{\text{mix}}^{(1,\geq 2)}} B \cap V_{\text{in}}; \\
 V_{\pi}^{(0,1)} &= \cup_{B \in \pi_{\text{in}}^{(1)}} B \cap V_{\text{in}}; \\
 V_{\pi}^{(0,\geq 2)} &= \cup_{B \in \pi_{\text{in}}^{(\geq 2)}} B \cap V_{\text{in}}.
 \end{aligned}$$

For a diamond $w_1 \in W_1$, we use the notation $\mathcal{O}(w_1) = \{v \in V_{\text{out}} : v \sim w_1\}$ for its set of outer vertex neighbors in G .

Note that we only need to prove the only if portion of the iff statement, the converse following from the classical $d = 2$ case. We prove Proposition 4.2 by induction on m with the base cases $m = 1, 2$. If $m = 1$, then (3.4) implies that $\#(\pi) = 1$. In particular, $\#(\pi_{\text{in}}) = 0$ and $\#(\pi_{\text{mix}}^{(1)}) \leq 1$, whence $\#(\pi_{\text{in}}) + \frac{\#(\pi_{\text{mix}}^{(1)})}{2} \leq \frac{1}{2}$.

If $m = 2$, then $\#(\pi_{\text{mix}}^{(1)}) = 0$. Indeed, an internal vertex $v \in B \in \pi_{\text{mix}}$ is necessarily adjacent to the two diamonds w_1 and w_2 . Condition (3.4) imposes the equality

$$\#(\{e : [v]_\pi \overset{e}{\sim} w_1\}) = \#(\{e : [v]_\pi \overset{e}{\sim} w_2\}),$$

which implies that $\#(B \cap V_{\text{out}})$ is even and not equal to 1. Now, if $\#(\pi_{\text{in}}) \leq 1$, then we are done. Otherwise, $\#(\pi_{\text{in}}) = 2$, in which case $\#(\pi_{\text{in}}) + \frac{\#(\pi_{\text{mix}}^{(1)})}{2} = 2$. Since $m = \#(\pi_{\text{in}}) = 2$, we have no inner vertices left to mix $\pi_{\text{mix}} = \emptyset$ and H^π is the double tree consisting of two vertices with two edges between them.

Now suppose that $m \geq 3$. Note that $\pi_{\text{mix}}^{(\geq 2)}$ and π_{out} do not play a role in the statement, so we can assume that

$$\#(\pi_{\text{mix}}^{(\geq 2)} \cup \pi_{\text{out}}) \leq 1 \tag{4.2}$$

by merging every such block. The resulting partition clearly still satisfies (3.4). If $\#(\pi_{\text{in}}^{(1)}) \leq \#(\pi_{\text{mix}}^{(1, \geq 2)})$, then we are done. To see this, note that every block $B \in \pi_{\text{mix}}^{(1, \geq 2)} \cup \pi_{\text{in}}^{(\geq 2)}$ satisfies $\#(B) \geq 2$, whence

$$\begin{aligned} \#(\pi_{\text{in}}) + \frac{\#(\pi_{\text{mix}}^{(1)})}{2} &= \#(\pi_{\text{in}}^{(1)}) + \frac{\#(\pi_{\text{mix}}^{(1, \geq 2)})}{2} + \#(\pi_{\text{in}}^{(\geq 2)}) + \frac{\#(\pi_{\text{mix}}^{(1,1)})}{2} \\ &\leq \frac{\#(\pi_{\text{in}}^{(1)})}{2} + \#(\pi_{\text{mix}}^{(1, \geq 2)}) + \#(\pi_{\text{in}}^{(\geq 2)}) + \frac{\#(\pi_{\text{mix}}^{(1,1)})}{2} \\ &\leq \frac{\#(V_\pi^{(0,1)})}{2} + \frac{\#(V_\pi^{(1, \geq 2)})}{2} + \frac{\#(V_\pi^{(0, \geq 2)})}{2} + \frac{\#(V_\pi^{(1,1)})}{2} \\ &\leq \frac{\#(V_{\text{in}})}{2} = \frac{m}{2}. \end{aligned} \tag{4.3}$$

Henceforth, we assume that

$$\#(\pi_{\text{in}}^{(1)}) > \#(\pi_{\text{mix}}^{(1, \geq 2)}). \tag{4.4}$$

We will need the following result.

Lemma 4.3. *For any block $B \in \pi_{\text{mix}}^{(1, \geq 2)}$, there is at most one block $B' \in \pi_{\text{in}}^{(1)}$ connected to B in H^π with a simple edge.*

Proof of Lemma 4.3. Let $B \sim_e B'$ be as above. Since $B' \in \pi_{\text{in}}^{(1)}$, there is exactly one other edge $e' \neq e$ incident to B' in H^π . By assumption, e' is not incident to B . Condition (3.4) then forces $\{e, e'\} \in \sigma_\pi$, which can only be achieved if $B \cap \mathcal{O}(e') \neq \emptyset$. We conclude that e' is responsible for the lone outer vertex in $B \in \pi_{\text{mix}}^{(1, \geq 2)}$.

The requirement $e \sim_{\sigma_\pi} e'$ precludes the possibility of e' being incident to another block $\tilde{B}' \neq B'$ in $\pi_{\text{in}}^{(1)}$. So, if $\tilde{B}' \in \pi_{\text{in}}^{(1)}$ is also connected to B in H^π with a simple edge \tilde{e} and \tilde{e}' is the only other edge incident to \tilde{B}' in H^π , it must be that $\tilde{e}' \neq e'$. We can then repeat the argument above and arrive at the contradiction that \tilde{e}' is responsible for another outer vertex $B \cap \mathcal{O}(\tilde{e}') \neq \emptyset$ in $B \in \pi_{\text{mix}}^{(1, \geq 2)}$. \square

Assumption (4.4) and Lemma 4.3 guarantee the existence of a block $B_0 \in \pi_{\text{in}}^{(1)}$ that is not connected to any element of $\pi_{\text{mix}}^{(1, \geq 2)}$ with a simple edge. Let e_1, e_2 be the two edges incident to B_0 . We write v_1 and v_2 for the vertices incident to e_1 and e_2 in H respectively with $B_1 = [v_1]_\pi$ and $B_2 = [v_2]_\pi$ the corresponding blocks. Since $B_0 \in \pi_{\text{in}}^{(1)}$, there are no other edges incident to B_0 and (3.4) imposes the constraint $\{e_1, e_2\} \in \sigma_\pi$.

We claim that $B_1 \notin \pi_{\text{mix}}^{(1,1)}$. Otherwise, $B_1 \in \pi_{\text{mix}}^{(1,1)}$ and $B_1 \neq B_2$ since $v_1 \neq v_2$ are distinct inner vertices (recall that $m \geq 3$). The constraint $e_1 \sim_{\sigma_\pi} e_2$ can then only be fulfilled if $\#(B_1 \cap \mathcal{O}(e_2)) = 1$. In that case, the other edge $e_3 \neq e_2$ incident to v_1 in

H ($e_3 \neq e_2$ since $v_1 \neq v_2$) cannot satisfy (3.4): $\#(\mathcal{B}_\pi(e_3)) \neq 1$ since $B_1 \in \pi_{\text{mix}}^{(1,1)}$, and $\mathcal{B}_\pi(e_3) \neq \mathcal{B}_\pi(e_1), \mathcal{B}_\pi(e_2)$ since $B_0 \in \pi_{\text{in}}^{(1)}$.

We then have the following trichotomy:

1. $B_1 \in \pi_{\text{in}}$, in which case $B_1 = B_2$ since $e_1 \sim_{\sigma_\pi} e_2$;
2. $B_1 \in \pi_{\text{mix}}^{(1, \geq 2)}$, in which case $B_1 = B_2$ because B_0 is not connected to any element of $\pi_{\text{mix}}^{(1, \geq 2)}$ by a simple edge by assumption (furthermore, the same parity argument as before shows that $B_1 \cap \mathcal{O}(e_1) = B_1 \cap \mathcal{O}(e_2) = \emptyset$);
3. $B_1 \in \pi_{\text{mix}}^{(\geq 2)}$, in which case $B_2 \notin \pi_{\text{in}} \cup \pi_{\text{mix}}^{(1, \geq 2)}$: otherwise, interchanging B_1 and B_2 , the analysis above would imply that $B_1 = B_2 \in \pi_{\text{in}} \cup \pi_{\text{mix}}^{(1, \geq 2)}$. Thus, it must be that $B_2 \in \pi_{\text{mix}}^{(\geq 2)}$ as well, whence $B_1 = B_2$ by our earlier reduction (4.2).

In any case, $B_1 = B_2$. We construct a smaller graph \tilde{G} from G by removing: the lone vertex in B_0 , the adjacent diamonds e_1, e_2 , the outer vertices $\mathcal{O}(e_1) \cup \mathcal{O}(e_2)$, and every incident edge. We also merge v_1 and v_2 . In other words, we have simply pinched off a small portion of the decorated cycle (2.1), bringing us to the case of $m - 2 \geq 1$.

By a slight abuse of notation, we write $\tilde{\pi} = \pi|_{\tilde{G}}$ for the natural partition of \tilde{G} induced by π (we merged v_1 and v_2 in constructing \tilde{G} , but $[v_1]_\pi = B_1 = B_2 = [v_2]_\pi$ anyway). Note that the partition $\tilde{\pi}$ still satisfies (3.4) since the other diamonds are unaffected (recall that $\{e_1, e_2\} \in \sigma_\pi$). We can then apply the induction hypothesis, which tells us that

$$\#(\tilde{\pi}_{\text{in}}) + \frac{\#(\tilde{\pi}_{\text{mix}}^{(1)})}{2} \leq \frac{m}{2} \tag{4.5}$$

with equality iff $m - 2$ is even, $\tilde{\pi}_{\text{mix}} = \emptyset$, and $\tilde{H}^{\tilde{\pi}}$ is a double tree.

To see what this says about our original partition π , consider a block $B \in \pi_{\text{mix}}^{(1)}$. Since $e_1 \sim_{\sigma_\pi} e_2$, we know that $B \cap \mathcal{O}(e_1) = B \cap \mathcal{O}(e_2) = \emptyset$ by the now familiar parity argument. So, after removing $\mathcal{O}(e_1) \cup \mathcal{O}(e_2)$, the block B becomes a block $\tilde{B} \in \tilde{\pi}_{\text{mix}}^{(1)}$, leading to the inequality

$$\#(\pi_{\text{mix}}^{(1)}) \leq \#(\tilde{\pi}_{\text{mix}}^{(1)}). \tag{4.6}$$

At the same time, removing $\mathcal{O}(e_1) \cup \mathcal{O}(e_2)$ creates a block in $\tilde{\pi}_{\text{in}}$ for each block $B \in \pi_{\text{mix}}$ such that $B \cap V_{\text{out}} \subset \mathcal{O}(e_1) \cup \mathcal{O}(e_2)$. We also lose a block in π_{in} from the removal of $B_0 \in \pi_{\text{in}}^{(1)}$, and so

$$\#(\tilde{\pi}_{\text{in}}) = \#(\pi_{\text{in}}) - 1 + \#(\{B \in \pi_{\text{mix}} : B \cap V_{\text{out}} \subset \mathcal{O}(e_1) \cup \mathcal{O}(e_2)\}).$$

In any case, $\#(\pi_{\text{in}}) \leq \#(\tilde{\pi}_{\text{in}}) + 1$. Combining this with (4.5) and (4.6), we obtain

$$\#(\pi_{\text{in}}) + \frac{\#(\pi_{\text{mix}}^{(1)})}{2} \leq \frac{m}{2} + 1$$

with $\#(\pi_{\text{in}}) + \frac{\#(\pi_{\text{mix}}^{(1)})}{2} \leq \frac{m}{2}$ unless $\tilde{\pi}_{\text{mix}} = \emptyset$, $\tilde{H}^{\tilde{\pi}}$ is a double tree, and $\#(\{B \in \pi_{\text{mix}} : B \cap V_{\text{out}} \subset \mathcal{O}(e_1) \cup \mathcal{O}(e_2)\}) = 0$.

Now assume that the latter conditions hold. Since $\tilde{\pi}_{\text{mix}} = \emptyset$, it follows from (4.6) that $\pi_{\text{mix}}^{(1)} = \emptyset$. Similarly, $\tilde{\pi}_{\text{mix}} = \emptyset$ and $\#(\{B \in \pi_{\text{mix}} : B \cap V_{\text{out}} \subset \mathcal{O}(e_1) \cup \mathcal{O}(e_2)\}) = 0$ imply that $\pi_{\text{mix}}^{(\geq 2)} = \emptyset$, and so $\pi_{\text{mix}} = \emptyset$. Finally, H^π is obtained from $\tilde{H}^{\tilde{\pi}}$ by adding the vertex B_0 and the double edge $\{e_1, e_2\}$ between B_0 and $B_1 = B_2$: if $\tilde{H}^{\tilde{\pi}}$ is a double tree, then so too is H^π . \square

4.2 Proof of Proposition 3.5

The idea of the proof is to realize (3.10) as a sum over vertex labelings of the disjoint union of $2M$ copies of the graph of tensors G in (2.1) satisfying certain properties. We will need to adapt the ideas in the previous section in conjunction with those in [3, Section 3.2] to bound the number of admissible vertex labelings in this extended picture. The gain in the exponent compared to the single graph case of Section 4.1 comes from the fact that one now needs an overlay of diamonds across different copies of G to prevent the vanishing of the expectation due to the centering: as we will see, such identifications are always suboptimal.

We start with some notation. Let $G_k = (V_k, W_1^{(k)} \sqcup W_2^{(k)}, E_k)$ be a graph of tensors (2.1) of length m_k . The statement of Proposition 3.5 assumes that $m_k \equiv m$, but we will prove a more general result that is amenable to induction. For a vertex labeling $\phi_k : V_k \rightarrow [N]$, we define the centered random variable

$$Q_{\phi_k} = \left(\prod_{w_1 \in W_1^{(k)}} \mathbf{T}_{d,N}(w_1 | \phi_k) - \mathbb{E} \left[\prod_{w_1 \in W_1^{(k)}} \mathbf{T}_{d,N}(w_1 | \phi_k) \right] \right) \prod_{w_2 \in W_2^{(k)}} u_N^{(w_2)}(w_2 | \phi_k). \quad (4.7)$$

This allows us to write the expectation of a product of unnormalized traces as

$$\mathbb{E} \left[\prod_{k=1}^K (\text{Tr}(G_k) - \mathbb{E}[\text{Tr}(G_k)]) \right] = N^{-\frac{\sum_{k=1}^K m_k}{2}} \sum_{\substack{(\phi_1, \dots, \phi_K) \\ \text{s.t. } \phi_k : V_k \rightarrow [N]}} \mathbb{E} \left[\prod_{k \in [K]} Q_{\phi_k} \right]. \quad (4.8)$$

We say that vertex labels $\phi_k, \phi_{k'}$ are *matched* if there exist vertices $(w_1, w'_1) \in W_1^{(k)} \times W_1^{(k')}$ such that

$$\phi_k(\mathcal{B}(w_1)) = \phi_{k'}(\mathcal{B}(w'_1)),$$

where the equality is in the sense of multisets. In words, this says that there must be two diamonds such that the ϕ_k -labeled neighborhood of the first matches the $\phi_{k'}$ -labeled neighborhood of the second. We say that a coordinate $\phi_{k'}$ in a K -tuple (ϕ_1, \dots, ϕ_K) is *unmatched* if it is not matched to any other coordinate ϕ_k with $k \neq k'$, in which case

$$\mathbb{E} \left[\prod_{k \in [K]} Q_{\phi_k} \right] = \mathbb{E} \left[\prod_{\substack{k \in [K]: \\ k \neq k'}} Q_{\phi_k} \right] \mathbb{E}[Q_{\phi_{k'}}] = 0$$

by the independence of our tensor entries and the centeredness of $Q_{\phi_{k'}}$. We may then restrict the sum in (4.8) to K -tuples with no unmatched coordinates.

To translate this condition to the setting of partitions and injective labelings, we will need some additional notation. Let $\mathcal{G} = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ be the disjoint union $\sqcup_{k=1}^K G_k$ of the graphs G_k . In particular,

$$\mathcal{V} = \sqcup_{k=1}^K V_k; \quad (4.9)$$

$$\mathcal{W}_1 = \sqcup_{k=1}^K W_1^{(k)}; \quad (4.10)$$

$$\mathcal{W}_2 = \sqcup_{k=1}^K W_2^{(k)}; \quad (4.11)$$

$$\mathcal{W} = \mathcal{W}_1 \sqcup \mathcal{W}_2; \quad (4.12)$$

$$\mathcal{E} = \sqcup_{k=1}^K E_k. \quad (4.13)$$

This allows us to rewrite (4.8) as

$$\mathbb{E} \left[\prod_{k=1}^K (\text{Tr}(G_k) - \mathbb{E}[\text{Tr}(G_k)]) \right] = N^{-\frac{\sum_{k=1}^K m_k}{2}} \sum_{\pi \in \mathcal{P}(\mathcal{V})} \sum_{\Phi : \mathcal{V}^\pi \hookrightarrow [N]} \mathbb{E} \left[\prod_{k \in [2M]} Q_{\Phi|_{V_k}} \right],$$

where, by a slight abuse of notation, we have used the fact that a map $\Phi : \mathcal{V}^\pi \hookrightarrow [N]$ defines a vertex labeling $\Phi|_{V_k} : V_k \rightarrow [N]$ for every $k \in [K]$. The condition about matchings from the previous paragraph then allows us to restrict to partitions π such that every copy G_k has a diamond overlay with at least one other copy $G_{k'}$ with $k' \neq k$.

The diamond overlays between the $(G_k)_{k=1}^K$ define an equivalence relation on $[K]$ with $r_\pi \leq \lfloor \frac{K}{2} \rfloor$ equivalence classes A_1, \dots, A_{r_π} . In particular, elements $k, k' \in [K]$ belong to the same equivalence class A_q iff there are subgraphs G_{k_0}, \dots, G_{k_t} in \mathcal{G} with $k_0 = k$ and $k_t = k'$ such that G_{k_s} has a diamond overlay with $G_{k_{s+1}}$ for $s = 0, \dots, t - 1$. This allows us to factor

$$\mathbb{E} \left[\prod_{k \in [K]} Q_{\Phi|_{V_k}} \right] = \prod_{q \in [r_\pi]} \mathbb{E} \left[\prod_{k \in A_q} Q_{\Phi|_{V_k}} \right].$$

We need the obvious analogue of Definition 3.1 for \mathcal{G} , which we denote by $\mathcal{H} = \sqcup_{k=1}^K H_k$ and $\mathcal{H}^\pi = \mathcal{H}^{\pi \uparrow \mathcal{V}(\mathcal{H})}$. We write \mathcal{H}_q^π for the subgraph spanned by the edges of $(H_k)_{k \in A_q}$ in \mathcal{H}^π . We decompose the vertex set of \mathcal{H}_q^π as before:

$$\mathcal{V}(\mathcal{H}_q^\pi) = \pi_{\text{in},q} \sqcup \pi_{\text{mix},q}^{(1)} \sqcup \pi_{\text{mix},q}^{(\geq 2)}.$$

Recall that the contracting vectors are deterministic: this allowed us to factor out their contribution in the formulation (4.7) of Q_ϕ . So, we can repeat the analysis from Section 4.1 to conclude that

$$\sum_{\Phi: \mathcal{V}^\pi \hookrightarrow [N]} \prod_{q \in [r_\pi]} \mathbb{E} \left[\prod_{k \in A_q} Q_{\Phi|_{V_k}} \right] = O_{m,M} \left(N^{\sum_{q \in [r_\pi]} \#(\pi_{\text{in},q}) + \frac{\#(\pi_{\text{mix},q}^{(1)})}{2}} \right).$$

So, we will be done if we can prove that

$$\#(\pi_{\text{in},q}) + \frac{\#(\pi_{\text{mix},q}^{(1)})}{2} \leq \frac{\sum_{k \in A_q} m_k}{2}. \tag{4.14}$$

Note that (4.14) is localized to an equivalence class A_q . So, without loss of generality, we may assume that there is only one equivalence class $A_q = [K]$. Thus, we need to prove

Proposition 4.4. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{W}, \mathcal{E})$ be the disjoint union as in (4.9)–(4.13) with $K \geq 2$. If $\pi \in \mathcal{P}(\mathcal{V})$ satisfies*

$$\#(\mathcal{B}_\pi(w_1)) = 1 \text{ or } \exists w'_1 \in \mathcal{W}_1 \setminus \{w_1\} : \mathcal{B}_\pi(w_1) = \mathcal{B}_\pi(w'_1), \quad \forall w_1 \in \mathcal{W}_1; \tag{4.15}$$

$$\exists k' \neq k, (w_1, w'_1) \in W_1^{(k)} \times W_1^{(k')} : \mathcal{B}_\pi(w_1) = \mathcal{B}_\pi(w'_1), \quad \forall k \in [K], \tag{4.16}$$

then

$$\#(\pi_{\text{in}}) + \frac{\#(\pi_{\text{mix}}^{(1)})}{2} \leq \frac{\sum_{k=1}^K m_k}{2}.$$

Proof. Before getting started, we note that (4.15) is the analogue of (3.4) in this context: here, the overlaying diamond w'_1 might come from a different graph of tensors $G_k \neq G_{k'}$. The additional constraint on the partition (4.16) is the mandatory crossed overlay necessary to survive the centeredness of the $Q_{\Phi|_{V_k}}$ as discussed earlier.

We adapt the proof of Proposition 4.2 and proceed by induction on $m = \sum_{k=1}^K m_k$ with the base cases $m = 2, 3$. The case of $m = 2$ corresponds to $K = 2$ and $m_1 = m_2 = 1$, meaning G_1 and G_2 each consist of a single diamond with a double edge to a single inner vertex and simple edges to $d - 2$ outer vertices. The crossed overlay (4.16) bounds $\#(\pi_{\text{in}}) \leq 1$ with $\#(\pi_{\text{in}}) = 1$ only if $\#(\pi_{\text{mix}}^{(1)}) = 0$. If $\#(\pi_{\text{in}}) = 0$, then the fact that each G_i has a double edge to its lone inner vertex means that each mixed block must have at least two outer vertices, and so $\#(\pi_{\text{mix}}^{(1)}) = 0$.




Figure 7: An example of the double tree obstruction for \mathcal{H}^π . For convenience, we have omitted the outer vertices. Even if one assumes $\pi_{\text{mix}} = \emptyset$, one cannot identify the remaining vertices in such a way that will produce a double tree for \mathcal{H}^π .

There are two possibilities for the case of $m = 3$. First, it could be that $K = 3$ and $m_1 = m_2 = m_3 = 1$. This is treated almost identically to the $m = 2$ case. Second, it could be that $K = 2$, $m_1 = 1$, and $m_2 = 2$, meaning G_1 is as before and G_2 is as in Figure 6. Let $B \in \pi$ be the block containing the inner vertex in G_1 . If $\#(B \cap \mathcal{V}_{\text{in}}) = 3$, then $\#(\pi_{\text{in}}) + \#(\pi_{\text{mix}}^{(1)}) \leq 1$ and we are done. If $\#(B \cap \mathcal{V}_{\text{in}}) = 2$, then (4.16) imposes the constraint $B \in \pi_{\text{mix}}$ because of the double edge in G_1 . But then $\#(\pi_{\text{in}}) \leq 1$ since there is only one unaccounted inner vertex left in \mathcal{G} and again we are done. Finally, if $\#(B \cap \mathcal{V}_{\text{in}}) = 1$, then (4.16) imposes the constraint $B \in \pi_{\text{mix}}^{(\geq 2)}$ (again, because of the double edge in G_1). In that case, $\#(\pi_{\text{in}}) \leq 2$ and we will be done if we can rule out equality. Since $\#(B \cap \mathcal{V}_{\text{in}}) = 1$, the crossed overlay in (4.16) keeps the inner vertices of the two overlaid diamonds in separate blocks. This forces the inner vertices of each diamond to be merged with outer vertices, and so $\#(\pi_{\text{in}}) = 0$.

Now suppose that $m \geq 4$. From here, the proof largely resembles Proposition 4.2 and we simply outline the argument. As before, we can assume (4.2) by merging every such block and noting that the resulting partition still satisfies (4.15)–(4.16). Similarly, we can assume (4.4) since otherwise we are done using the same line of reasoning as in (4.3). Lemma 4.3 still holds for \mathcal{H}^π with the same proof: together with (4.4), this guarantees the existence of a block $B_0 \in \pi_{\text{in}}^{(1)}$ that is not connected to any element of $\pi_{\text{mix}}^{(1, \geq 2)}$ with a simple edge as before.

Assume that the lone vertex $v_0 \in B_0 \in \pi_{\text{in}}^{(1)}$ comes from H_{k_0} . Here, we need to be careful: in the proof of Proposition 4.2, we knew that $m_{k_0} \geq 3$ because of the induction on m_{k_0} ; now, we are inducting on $m = \sum_{k=1}^K m_k$, and so we cannot immediately assume $m_{k_0} \geq 3$. We can rule out $m_{k_0} = 1, 2$ as follows. If $m_{k_0} = 1$, then there is a unique inner vertex in G_{k_0} , necessarily v_0 , and the crossed overlay in (4.16) forces $B_0 \in \pi_{\text{in}}^{(\geq 2)} \cup \pi_{\text{mix}}$, a contradiction. Similarly, if $m_{k_0} = 2$, then G_{k_0} is as in Figure 6 and the fact that $B_0 \in \pi_{\text{in}}^{(1)}$ precludes a crossed overlay (4.16) with any other G_k , again a contradiction.

We can now repeat the argument leading to the trichotomy in Proposition 4.2, pinch off the edges incident to v_0 in H_{k_0} , and reduce to the case of $\tilde{\pi}$ and $m - 2 \geq 2$, where the induction hypothesis tells us that

$$\#(\tilde{\pi}_{\text{in}}) + \frac{\#(\tilde{\pi}_{\text{mix}}^{(1)})}{2} \leq \frac{m - 2}{2}.$$

As before, we have the inequalities $\#(\pi_{\text{mix}}^{(1)}) \leq \#(\tilde{\pi}_{\text{mix}}^{(1)})$ and $\#(\pi_{\text{in}}) \leq \#(\tilde{\pi}_{\text{in}}) + 1$, which allow us to conclude. \square

Remark 4.5. The suboptimality of the crossed overlay (4.16) and the subsequent gain in (4.14) versus (4.1) comes from the fact that the condition for equality in Proposition 4.2 cannot be achieved in the setting of Proposition 4.4 for \mathcal{H}^π . See Figure 7 for an illustration.

Setting $K = 2M$, $m_k \equiv m$, and $G_k \equiv G$ in Proposition 4.4, it follows that

$$\mathbb{E}\left[\left(\mathrm{Tr}(G) - \mathbb{E}[\mathrm{Tr}(G)]\right)^{2M}\right] = O_{m,d,M}(1),$$

as was to be shown.

References

- [1] G. Akemann, J. Baik, and P. Di Francesco (eds.), *The Oxford handbook of random matrix theory*, Oxford University Press, Oxford, 2011. MR2920518
- [2] G. W. Anderson, A. Guionnet, and O. Zeitouni, *An introduction to random matrices*, Cambridge Studies in Advanced Mathematics, vol. 118, Cambridge University Press, Cambridge, 2010. MR2760897
- [3] B. Au, *Traffic distributions of random band matrices*, Electron. J. Probab. **23** (2018), paper no. 77, 48 pp. MR3858905
- [4] B. Au, G. Cébron, A. Dahlqvist, F. Gabriel, and C. Male, *Freeness over the diagonal for large random matrices*, Ann. Probab. **49** (2021), no. 1, 157–179. MR4203335
- [5] B. Au and C. Male, *Rigid structures in the universal enveloping traffic space*, Preprint. arXiv:2011.05472v1.
- [6] Z. Bai and J. W. Silverstein, *Spectral analysis of large dimensional random matrices*, second ed., Springer Series in Statistics, Springer, New York, 2010. MR2567175
- [7] J. Baik, P. Deift, and T. Suidan, *Combinatorics and random matrix theory*, Graduate Studies in Mathematics, vol. 172, American Mathematical Society, Providence, RI, 2016. MR3468920
- [8] A. Bandeira, M. Boedihardjo, and R. van Handel, *Matrix concentration inequalities and free probability*, Preprint. arXiv:2108.06312v1. MR4635836
- [9] M. Banna, F. Merlevède, and M. Peligrad, *On the limiting spectral distribution for a large class of symmetric random matrices with correlated entries*, Stochastic Process. Appl. **125** (2015), no. 7, 2700–2726. MR3332852
- [10] G. Cébron, A. Dahlqvist, and C. Male, *Traffic distributions and independence II: universal constructions for traffic spaces*, Preprint. arXiv:1601.00168v2.
- [11] J. H. de M. Goulart, R. Couillet, and P. Comon, *A random matrix perspective on random tensors*, Preprint. arXiv:2108.00774v1. MR4577703
- [12] K. Dykema, *On certain free product factors via an extended matrix model*, J. Funct. Anal. **112** (1993), no. 1, 31–60. MR1207936
- [13] L. Erdős and H.-T. Yau, *A dynamical approach to random matrix theory*, Courant Lecture Notes in Mathematics, vol. 28, Courant Institute of Mathematical Sciences, New York; American Mathematical Society, Providence, RI, 2017. MR3699468
- [14] W. Feller, *An introduction to probability theory and its applications. Vol. II*, second ed., John Wiley & Sons, Inc., New York-London-Sydney, 1971. MR0270403
- [15] F. Götze, A. A. Naumov, and A. N. Tikhomirov, *Limit theorems for two classes of random matrices with dependent entries*, Theory Probab. Appl. **59** (2015), no. 1, 23–39. MR3416062
- [16] W. Hochstättler, W. Kirsch, and S. Warzel, *Semicircle law for a matrix ensemble with dependent entries*, J. Theoret. Probab. **29** (2016), no. 3, 1047–1068. MR3540489
- [17] L.-H. Lim, *Singular values and eigenvalues of tensors: a variational approach*, in: 1st IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing, 2005, pp. 129–132.
- [18] L.-H. Lim, *Tensors in computations*, Acta Numer. **30** (2021), 555–764. MR4298222

- [19] C. Male, *Traffic distributions and independence: permutation invariant random matrices and the three notions of independence*, Mem. Amer. Math. Soc. **267** (2020), no. 1300, v+88. MR4197072
- [20] M. Marcus, *The Hadamard theorem for permanents*, Proc. Amer. Math. Soc. **15** (1964), 967–973. MR168585
- [21] M. L. Mehta, *Random matrices*, third ed., Pure and Applied Mathematics (Amsterdam), vol. 142, Elsevier/Academic Press, Amsterdam, 2004. MR2129906
- [22] H. Minc, *Permanents*, Encyclopedia of Mathematics and its Applications, vol. 6, Addison-Wesley Publishing Co., Reading, Mass., 1978, With a foreword by Marvin Marcus. MR504978
- [23] J. A. Mingo and R. Speicher, *Free probability and random matrices*, Fields Institute Monographs, vol. 35, Springer, New York; Fields Institute for Research in Mathematical Sciences, Toronto, ON, 2017. MR3585560
- [24] A. Nica and R. Speicher, *Commutators of free random variables*, Duke Math. J. **92** (1998), no. 3, 553–592. MR1620518
- [25] A. Nica and R. Speicher, *Lectures on the combinatorics of free probability*, London Mathematical Society Lecture Note Series, vol. 335, Cambridge University Press, Cambridge, 2006. MR2266879
- [26] OEIS Foundation Inc., *Entry a023998 in the On-Line Encyclopedia of Integer Sequences*, 2023, <https://oeis.org/A023998>. MR4278491
- [27] R. Penrose, *Applications of negative dimensional tensors*, in: Combinatorial Mathematics and its Applications (Proc. Conf., Oxford, 1969), Academic Press, London, 1971, pp. 221–244. MR0281657
- [28] L. Qi, *Eigenvalues of a real supersymmetric tensor*, J. Symbolic Comput. **40** (2005), no. 6, 1302–1324. MR2178089
- [29] L. Qi and Z. Luo, *Tensor analysis: spectral theory and special tensors*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2017. MR3660696
- [30] J. H. Schenker and H. Schulz-Baldes, *Semicircle law and freeness for random matrices with symmetries or correlations*, Math. Res. Lett. **12** (2005), no. 4, 531–542. MR2155229
- [31] D. V. Voiculescu, K. J. Dykema, and A. Nica, *Free random variables*, CRM Monograph Series, vol. 1, American Mathematical Society, Providence, RI, 1992, A noncommutative probability approach to free products with applications to random matrices, operator algebras and harmonic analysis on free groups. MR1217253
- [32] D. Voiculescu, *Limit laws for random matrices and free products*, Invent. Math. **104** (1991), no. 1, 201–220. MR1094052
- [33] E. P. Wigner, *Characteristic vectors of bordered matrices with infinite dimensions*, Ann. of Math. (2) **62** (1955), 548–564. MR0077805
- [34] J. Wishart, *The generalised product moment distribution in samples from a normal multivariate population*, Biometrika **20A** (1928), no. 1-2, 32–52. MR0050223

Acknowledgments. The authors would like to thank the anonymous referees for their detailed feedback. In particular, the authors are grateful to the anonymous referee who suggested a shorter proof of Proposition 4.2 but insisted on remaining anonymous.